

ExtremeControl® User Guide



Copyright © 2023 Extreme Networks, Inc. All Rights Reserved.

Legal Notices

Extreme Networks, Inc., on behalf of or through its wholly-owned subsidiary, Enterasys Networks, Inc., reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, please see: www.extremenetworks.com/company/legal/trademarks/

Contact

If you require assistance, contact Extreme Networks using one of the following methods.

- Global Technical Assistance Center (GTAC) for Immediate Support
 - Phone: 1-800-998-2408 (toll-free in U.S. and Canada) or 1-603-952-5000. For the Extreme Networks support phone number in your country, visit: www.extremenetworks.com/support/contact
 - Email: <u>support@extremenetworks.com</u>. To expedite your message, enter the product name or model number in the subject line.
- <u>GTAC Knowledge</u> Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- <u>The Hub</u> A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.
- <u>Support Portal</u> Manage cases, downloads, service contracts, product licensing, and training and certifications.



Extreme Networks Software License Agreement

This Extreme Networks Software License Agreement is an agreement ("Agreement") between You, the end user, and Extreme Networks, Inc. ("Extreme"), on behalf of itself and its Affiliates (as hereinafter defined and including its wholly owned subsidiary, Enterasys Networks, Inc. as well as its other subsidiaries). This Agreement sets forth Your rights and obligations with respect to the Licensed Software and Licensed Materials. BY INSTALLING THE LICENSE KEY FOR THE SOFTWARE ("License Key"), COPYING, OR OTHERWISE USING THE LICENSED SOFTWARE, YOU ARE AGREEING TO BE BOUND BY THE TERMS OF THIS AGREEMENT, WHICH INCLUDES THE LICENSE AND THE LIMITATION OF WARRANTY AND DISCLAIMER OF LIABILITY. IF YOU DO NOT AGREE TO THE TERMS OF THIS AGREEMENT, RETURN THE LICENSE KEY TO EXTREME OR YOUR DEALER, IF ANY, OR DO NOT USE THE LICENSED SOFTWARE AND CONTACT EXTREME OR YOUR DEALER WITHIN TEN (10) DAYS FOLLOWING THE DATE OF RECEIPT FOR A REFUND. IF YOU HAVE ANY QUESTIONS ABOUT THIS AGREEMENT, CONTACT EXTREME, Attn: LegalTeam@extremenetworks.com.

- 1. <u>DEFINITIONS</u>. "Affiliates" means any person, partnership, corporation, limited liability company, or other form of enterprise that directly or indirectly through one or more intermediaries, controls, or is controlled by, or is under common control with the party specified. "Server Application" shall refer to the License Key for software installed on one or more of Your servers. "Client Application" shall refer to the application to access the Server Application. "Licensed Materials" shall collectively refer to the licensed software (including the Server Application and Client Application), Firmware, media embodying the software, and the documentation. "Concurrent User" shall refer to any of Your individual employees who You provide access to the Server Application at any one time. "Firmware" refers to any software program or code imbedded in chips or other media. "Licensed Software" refers to the Software and Firmware collectively.
- 2. <u>TERM</u>. This Agreement is effective from the date on which You install the License Key, use the Licensed Software, or a Concurrent User accesses the Server Application. You may terminate the Agreement at any time by destroying the Licensed Materials, together with all copies, modifications and merged portions in any form. The Agreement and Your license to use the Licensed Materials will also terminate if You fail to comply with any term of condition herein.
- 3. GRANT OF SOFTWARE LICENSE. Extreme will grant You a non-transferable, non-exclusive license to use the machine-readable form of the Licensed Software and the accompanying documentation if You agree to the terms and conditions of this Agreement. You may install and use the Licensed Software as permitted by the license type purchased as described below in License Types. The license type purchased is specified on the invoice issued to You by Extreme or Your dealer, if any. YOU MAY NOT USE, COPY, OR MODIFY THE LICENSED MATERIALS, IN WHOLE OR IN PART, EXCEPT AS EXPRESSLY PROVIDED IN THIS AGREEMENT.

4. LICENSE TYPES.

- Single User, Single Computer. Under the terms of the Single User, Single Computer license, the license granted to You by Extreme when You install the License Key authorizes You to use the Licensed Software on any one, single computer only, or any replacement for that computer, for internal use only. A separate license, under a separate Software License Agreement, is required for any other computer on which You or another individual or employee intend to use the Licensed Software. A separate license under a separate Software License Agreement is also required if You wish to use a Client license (as described below).
- Client. Under the terms of the Client license, the license granted to You by Extreme will authorize You to install the License Key for the Licensed Software on your server and allow the specific number of Concurrent Users shown on the relevant invoice issued to You for each Concurrent User that You order from Extreme or Your dealer, if any, to access the Server Application. A separate license is required for each additional Concurrent User.
- 5. <u>AUDIT RIGHTS</u>. You agree that Extreme may audit Your use of the Licensed Materials for compliance with these terms and Your License Type at any time, upon reasonable notice. In the event that such audit reveals any use of the Licensed Materials by You other than in full compliance with the license granted and the terms of this Agreement, You shall reimburse Extreme for all reasonable expenses related to such audit in addition to any other liabilities You may incur as a result of such non-compliance, including but not limited to additional fees for Concurrent Users over and above those specifically granted to You. From time to time, the Licensed Software will upload information about the Licensed Software and the associated devices to Extreme. This is to verify the Licensed Software is being used with a valid license. By using the Licensed Software, you consent to the transmission of this information. Under no circumstances, however, would Extreme employ any such measure to interfere with your normal and permitted operation of the Products, even in the event of a contractual dispute.
- 6. <u>RESTRICTION AGAINST COPYING OR MODIFYING LICENSED MATERIALS</u>. Except as expressly permitted in this Agreement, You may not copy or otherwise reproduce the Licensed Materials. In no event does the limited copying or reproduction permitted under this Agreement include the right to decompile, disassemble, electronically transfer, or reverse engineer the Licensed Software, or to translate the Licensed Software into another computer language.

The media embodying the Licensed Software may be copied by You, in whole or in part, into printed or machine readable form, in sufficient numbers only for backup or archival purposes, or to replace a worn or defective copy. However, You agree not to have more than two (2) copies of the Licensed Software in whole or in part, including the original media, in your possession for said purposes without Extreme's prior written consent, and in no event shall You operate more copies of the Licensed Software than the specific licenses granted to You. You may not copy or reproduce the documentation. You agree to maintain appropriate records of the location of the original media and all copies of the Licensed Software, in whole or in part, made by You. You may modify the machine-readable form of the Licensed Software for (1) your own internal use or (2) to merge the Licensed Software into other program material to form a modular work for your own use, provided that such work remains modular, but on termination of this Agreement, You are required to completely remove the Licensed Software from any such modular work. Any portion of the Licensed Software included in any such modular work shall be used only on a single computer for internal purposes and shall remain subject to all the terms and conditions of this Agreement. You agree to include any copyright or other proprietary notice set forth on the label of the media embodying the Licensed Software on any copy of the Licensed Software in any form, in whole or in part,

or on any modification of the Licensed Software or any such modular work containing the Licensed Software or any part thereof.

7. TITLE AND PROPRIETARY RIGHTS

- a. The Licensed Materials are copyrighted works and are the sole and exclusive property of Extreme, any company or a division thereof which Extreme controls or is controlled by, or which may result from the merger or consolidation with Extreme (its "Affiliates"), and/or their suppliers. This Agreement conveys a limited right to operate the Licensed Materials and shall not be construed to convey title to the Licensed Materials to You. There are no implied rights. You shall not sell, lease, transfer, sublicense, dispose of, or otherwise make available the Licensed Materials or any portion thereof, to any other party.
- b. You further acknowledge that in the event of a breach of this Agreement, Extreme shall suffer severe and irreparable damages for which monetary compensation alone will be inadequate. You therefore agree that in the event of a breach of this Agreement, Extreme shall be entitled to monetary damages and its reasonable attorney's fees and costs in enforcing this Agreement, as well as injunctive relief to restrain such breach, in addition to any other remedies available to Extreme.
- 8. PROTECTION AND SECURITY. In the performance of this Agreement or in contemplation thereof, You and your employees and agents may have access to private or confidential information owned or controlled by Extreme relating to the Licensed Materials supplied hereunder including, but not limited to, product specifications and schematics, and such information may contain proprietary details and disclosures. All information and data so acquired by You or your employees or agents under this Agreement or in contemplation hereof shall be and shall remain Extreme's exclusive property, and You shall use your best efforts (which in any event shall not be less than the efforts You take to ensure the confidentiality of your own proprietary and other confidential information) to keep, and have your employees and agents keep, any and all such information and data confidential, and shall not copy, publish, or disclose it to others, without Extreme's prior written approval, and shall return such information and data to Extreme at its request. Nothing herein shall limit your use or dissemination of information not actually derived from Extreme or of information which has been or subsequently is made public by Extreme, or a third party having authority to do so.

You agree not to deliver or otherwise make available the Licensed Materials or any part thereof, including without limitation the object or source code (if provided) of the Licensed Software, to any party other than Extreme or its employees, except for purposes specifically related to your use of the Licensed Software on a single computer as expressly provided in this Agreement, without the prior written consent of Extreme. You agree to use your best efforts and take all reasonable steps to safeguard the Licensed Materials to ensure that no unauthorized personnel shall have access thereto and that no unauthorized copy, publication, disclosure, or distribution, in whole or in part, in any form shall be made, and You agree to notify Extreme of any unauthorized use thereof. You acknowledge that the Licensed Materials contain valuable confidential information and trade secrets, and that unauthorized use, copying and/or disclosure thereof are harmful to Extreme or its Affiliates and/or its/their software suppliers.

9. MAINTENANCE AND UPDATES. Updates and certain maintenance and support services, if any, shall be provided to You pursuant to the terms of an Extreme Service and Maintenance Agreement, if Extreme and You enter into such an agreement. Except as specifically set forth in such agreement, Extreme shall not be under any obligation to provide Software Updates, modifications, or enhancements, or Software maintenance and support services to You.

- 10. <u>DEFAULT AND TERMINATION</u>. In the event that You shall fail to keep, observe, or perform any obligation under this Agreement, including a failure to pay any sums due to Extreme, or in the event that you become insolvent or seek protection, voluntarily or involuntarily, under any bankruptcy law, Extreme may, in addition to any other remedies it may have under law, terminate the License and any other agreements between Extreme and You.
 - a. Immediately after any termination of the Agreement or if You have for any reason discontinued use of Software, You shall return to Extreme the original and any copies of the Licensed Materials and remove the Licensed Software from any modular works made pursuant to Section 3, and certify in writing that through your best efforts and to the best of your knowledge the original and all copies of the terminated or discontinued Licensed Materials have been returned to Extreme.
 - b. Sections 1, 7, 8, 10, 11, 12, 13, 14 and 15 shall survive termination of this Agreement for any reason.
- 11. EXPORT REQUIREMENTS. You are advised that the Software is of United States origin and subject to United States Export Administration Regulations; diversion contrary to United States law and regulation is prohibited. You agree not to directly or indirectly export, import or transmit the Software to any country, end user or for any Use that is prohibited by applicable United States regulation or statute (including but not limited to those countries embargoed from time to time by the United States government); or contrary to the laws or regulations of any other governmental entity that has jurisdiction over such export, import, transmission or Use.
- 12. <u>UNITED STATES GOVERNMENT RESTRICTED RIGHTS</u>. The Licensed Materials (i) were developed solely at private expense; (ii) contain "restricted computer software" submitted with restricted rights in accordance with section 52.227-19 (a) through (d) of the Commercial Computer Software-Restricted Rights Clause and its successors, and (iii) in all respects is proprietary data belonging to Extreme and/or its suppliers. For Department of Defense units, the Licensed Materials are considered commercial computer software in accordance with DFARS section 227.7202-3 and its successors, and use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth herein.
- 13. LIMITED WARRANTY AND LIMITATION OF LIABILITY. The only warranty that Extreme makes to You in connection with this license of the Licensed Materials is that if the media on which the Licensed Software is recorded is defective, it will be replaced without charge, if Extreme in good faith determines that the media and proof of payment of the license fee are returned to Extreme or the dealer from whom it was obtained within ninety (90) days of the date of payment of the license fee. NEITHER EXTREME NOR ITS AFFILIATES MAKE ANY OTHER WARRANTY OR REPRESENTATION, EXPRESS OR IMPLIED, WITH RESPECT TO THE LICENSED MATERIALS, WHICH ARE LICENSED "AS IS". THE LIMITED WARRANTY AND REMEDY PROVIDED ABOVE ARE EXCLUSIVE AND IN LIEU OF ALL OTHER WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE EXPRESSLY DISCLAIMED, AND STATEMENTS OR REPRESENTATIONS MADE BY ANY OTHER PERSON OR FIRM ARE VOID. ONLY TO THE EXTENT SUCH EXCLUSION OF ANY IMPLIED WARRANTY IS NOT PERMITTED BY LAW, THE DURATION OF SUCH IMPLIED WARRANTY IS LIMITED TO THE DURATION OF THE LIMITED WARRANTY SET FORTH ABOVE. YOU ASSUME ALL RISK AS TO THE QUALITY. FUNCTION AND PERFORMANCE OF THE LICENSED MATERIALS. IN NO EVENT WILL EXTREME OR ANY OTHER PARTY WHO HAS BEEN INVOLVED IN THE CREATION, PRODUCTION OR DELIVERY OF THE LICENSED MATERIALS BE LIABLE FOR SPECIAL, DIRECT, INDIRECT, RELIANCE, INCIDENTAL OR CONSEQUENTIAL DAMAGES, INCLUDING LOSS OF DATA OR PROFITS OR FOR INABILITY TO USE THE LICENSED MATERIALS, TO ANY PARTY EVEN IF EXTREME OR SUCH OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN

NO EVENT SHALL EXTREME OR SUCH OTHER PARTY'S LIABILITY FOR ANY DAMAGES OR LOSS TO YOU OR ANY OTHER PARTY EXCEED THE LICENSE FEE YOU PAID FOR THE LICENSED MATERIALS. Some states do not allow limitations on how long an implied warranty lasts and some states do not allow the exclusion or limitation of incidental or consequential damages, so the above limitation and exclusion may not apply to You. This limited warranty gives You specific legal rights, and You may also have other rights which vary from state to state.

14. <u>JURISDICTION</u>. The rights and obligations of the parties to this Agreement shall be governed and construed in accordance with the laws and in the State and Federal courts of the State of California, without regard to its rules with respect to choice of law. You waive any objections to the personal jurisdiction and venue of such courts. None of the 1980 United Nations Convention on the Limitation Period in the International Sale of Goods, and the Uniform Computer Information Transactions Act shall apply to this Agreement.

15. GENERAL.

- a. This Agreement is the entire agreement between Extreme and You regarding the Licensed Materials, and all prior agreements, representations, statements, and undertakings, oral or written, are hereby expressly superseded and canceled.
- b. This Agreement may not be changed or amended except in writing signed by both parties hereto.
- c. You represent that You have full right and/or authorization to enter into this Agreement.
- d. This Agreement shall not be assignable by You without the express written consent of Extreme. The rights of Extreme and Your obligations under this Agreement shall inure to the benefit of Extreme's assignees, licensors, and licensees.
- e. Section headings are for convenience only and shall not be considered in the interpretation of this Agreement.
- f. The provisions of the Agreement are severable and if any one or more of the provisions hereof are judicially determined to be illegal or otherwise unenforceable, in whole or in part, the remaining provisions of this Agreement shall nevertheless be binding on and enforceable by and between the parties hereto.
- g. Extreme's waiver of any right shall not constitute waiver of that right in future. This Agreement constitutes the entire understanding between the parties with respect to the subject matter hereof, and all prior agreements, representations, statements and undertakings, oral or written, are hereby expressly superseded and canceled. No purchase order shall supersede this Agreement.
- h. Should You have any questions regarding this Agreement, You may contact Extreme at the address set forth below. Any notice or other communication to be sent to Extreme must be mailed by certified mail to the following address:

Extreme Networks, Inc. 145 Rio Robles San Jose, CA 95134 United States ATTN: General Counsel

Table of Contents

ExtremeControl® User Guide	
Legal Notices	2
Trademarks	2
Contact	2
Extreme Networks® Software License Agreement	3
Table of Contents	8
Getting Started with ExtremeControl	32
Access Requirements	32
Navigating the Control Tab	32
Dashboard	32
Policy	33
Access Control	33
End-Systems	33
Reports	33
Policy	34
Understanding Policy Domains	36
Understanding Roles	37
Role Summary Column	39
Understanding Services	39
Working with Service Groups	40
Understanding Traffic Classification Rules	41
Adding Devices	41
Viewing Port Configuration Information	42
Working with Port Groups	42
Working with VLANS	43
Viewing Classes of Service	43
Saving the Domain	44
Enforcing	44

Enforce Preview	45
Rule Counts Reported by Devices	45
Verifying	46
AP Aware	46
Policy Configuration Considerations	47
General Considerations	47
Authenticating without Policy	47
Terminating Role Override Sessions	48
Port-Level MAC to Role Mappings	49
Import From Device	49
Flood Control	49
C1 Considerations	49
Policy Support	49
Rule Limits	50
N-Series Considerations	50
Role Precedence for the N-Series Platinum	50
C2 and B2 Considerations	50
C3 and B3 Considerations	51
Mixed-Stack C2/C3 and B2/B3 Considerations	51
7100 Considerations	52
ExtremeControl Controller Configuration	53
ExtremeControl Controllers Require Separate Domains	53
Modifying ExtremeControl Controllers Preconfigured Policy	53
Modifying the Downstream Default Policy	53
Configuring LAG on ExtremeControl Controllers	53
Configuring LAG on Layer 3 ExtremeControl Controllers - Upstream Ports	54
Configuring LAG on Layer 3 ExtremeControl Controllers - Downstream Ports	54
Configuring LAG on Layer 2 ExtremeControl Controllers - Upstream Ports	54
Configuring LAG on Layer 2 ExtremeControl Controllers - Downstream Ports	54
Extrama Wirolass Controllar Configuration	5/

Version Supported	54
Policy Rules	55
Supported Rule Types	55
"No Change" Filter Sets	55
Rule Actions	55
Rule Directions	56
Rule Limits	56
Role Default Actions	56
Class of Service	57
Rate Limits	57
Internal VLAN	57
Policy Inheritance	58
Configuring RADIUS Servers	58
Other Considerations	59
ExtremeCloud IQ - Site Engine Policy	60
Policy Tab Overview	60
Details View	60
General	61
Policy Menus	61
Open/Manage Domains Menu	61
Global Domain Settings Menu	62
Tools Menu	63
Policy Enforce Preview	64
Left Panel	64
Right Panel	65
Import from Domain	68
Data Elements to Import	69
Application of Imported Data Elements	71
Import from File	72
Data Elements to Import	73

Global Domain Data	75
Application of Imported Data Elements	75
Assign Devices to Domain	76
Authentication Configuration	79
Device Selection	79
Port Selection	79
Device Configuration	80
Authentication Status	80
Global Authentication Settings	81
MAC Authentication Settings	82
Web Authentication Settings	83
General	83
Guest Networking	84
Web Page Banner	85
Convergence End-Point Settings	86
CEP Role Mappings	86
CEP Detection Tab	87
Port Configuration	89
Authentication Mode	89
Port Mode	89
RFC3580 VLAN Authorization Tab	91
Login Settings	92
Automatic Re-Authentication	94
Authenticated User Counts	95
Convergence End-Point Access	96
Policy Main Window	97
Menu Tabs	97
Dialog Boxes (Messages)	98
lcons	98
Open/Manage Domain Menu Icons	99

Policy Windows	99
Policy Concepts	99
Policy	100
Role	100
What is a Role	100
Default Role	100
Policy Domains	10
Service	10
Rule	102
What is a Rule	102
Disabling Rules	102
Conflict Checking	103
Packet Tagging	103
VLAN to Role Mapping	104
Dynamic Egress	109
Setting Domain GVRP Status	108
Policy VLAN Islands	109
Traffic Mirroring	109
Port Groups	110
User-Defined Port Groups	110
Network Resource Groups	110
Network Resource Topologies	110
Verifying	11
Enforcing	11
Controlling Client Interactions with Locks	112
icy Tab Right-Panel	11
icy Left Panel	113
Roles/Services Tab	113
Roles Tree	114
Service Repository Tree	112

(Class of Service Tab	116
\	/LAN Tab	118
١	Network Resources Configuration	119
	Devices/Port Groups Tab	121
	Devices Tree	121
Sur	nmary (Roles)	123
Ger	neral (Role)	123
	Default Actions	124
S	Services	126
VLA	AN Egress (Role)	126
	Add Egress VLAN Window	127
Мај	ppings (Role)	128
١	MAC to Role Mapping	129
II	P to Role Mapping	130
Т	Fagged Packet VLAN to Role Mapping	130
A	Authentication-Based VLAN to Role Mapping	130
Pre	-configured Domains (Legacy)	130
A	Access Pre-Configured Domains	131
F	Pre-configured Domain Descriptions	131
	Embedded NAC Domain	131
	Generic Services N-Series	132
	Generic Services SecureStack	132
	HealthCare Services	132
	Quickstart	132
	Secure Guest	133
	ShoreTel	133
	VPN Termination Point	133
	Add/Remove Services (Roles)	133
Det	tails View (Service)	134
Service	Repository	138

Local/Global Services	139
Details View (Services)	139
Details View (Service Group)	140
Add/Remove Services (Service Groups)	141
Rule	143
General Area	143
Traffic Description Area	144
Actions Area	145
Create Rule	147
Edit Rule	148
Layer Area	148
Value Area	149
Class of Service Overview	149
Getting Started with Class of Service	150
Class of Service Overview	150
Implementing CoS	151
Configuring CoS	151
Rate Limits	152
Transmit Queues	153
Flood Control	154
Class of Service	154
General	155
Rate Limiting/Rate Shaping	156
Index Numbers	156
General (CoS Components Folder)	158
General (Rate Limits)	158
Details View (Rate Limits Folder)	160
Priority-Based Rate Limits	161
Add/Edit CoS to Rate Limit Mapping	162
Advanced Rate Limiting by Port Type	162

	Configuring Rate Limit Mappings	
	Associating Rate Limits with a Class of Service	164
	Summary (Rate Limit Port Groups Folder)	164
	CoS - Rate Limit Mappings (Rate Limit Port Group)	165
	Ports (Rate Limit Port Group)	167
	Automated Service	169
	Traffic Description Area	170
	Actions Area	17C
	Traffic Classification Rules	172
	Traffic Descriptions	173
	Actions	174
	VLAN Membership (Access Control)	174
	Priority (Class of Service)	174
	Classification Types and their Parameters	175
	Layer 2 Data Link Classification Types	175
	Layer 3 Network Classification Types	176
	Layer 4 Application Transport Classification Types	182
	Layer 7 Application Classification Types	186
	Examples of How Rules are Used	186
	Traffic Containment	186
	Traffic Filtering	187
	Traffic Security	188
	Traffic Prioritization	188
	Ports (Transmit Queue Port Group)	190
	Summary (Transmit Queue Port Groups)	19
	CoS - Transmit Queue Mappings (Transmit Queue Port Group)	192
	Ports (Flood Control Port Groups)	194
FI	ood Control Port Groups	196
	Flood Control Rate Limits (Flood Control Port Groups)	196
	Class of Service Example	197
	Configure the Classes of Service	200

Create the VoIP Core Role	200
Create a VoIP Core Service	200
Create a Rule	200
Creating the VoIP Edge Role	200
Create a VoIP Edge Service	20
Create a Rule	20
Creating the H.323 Call Setup Role	20
Create a H.323 Call Setup Service	20
Create a Rule	20
Apply the Roles to Network Devices	20
ToS/DSCP Value Definition Chart	202
Policy VLAN Tab Overview	202
General	203
Authentication-Based VLAN to Role Mapping	204
Tagged Packet VLAN to Role Mapping	202
Global VLANs	205
Create VLAN	206
Editing an existing VLAN/Class of Service	207
Selection View (Roles)	207
Policy VLAN Islands	208
(VLANs) - VIDs Tab	208
(VLANs) - Role Mappings Tab	209
General	210
Authentication-Based VLAN to Role Mapping	21
Tagged Packet VLAN to Role Mapping	21
Add Devices (VLAN Islands)	212
Island Topology (Policy VLAN Islands)	213
(Island) - VIDs Tab	213
(Island) - Devices Tab	214
Packet Flow Diagram	216

Network Resources Tab Overview	217
Network Resource Group General Tab	217
Network Resource Topology Tab	219
Network Resource Topology Island Domain Wide	219
Details View (Network Resource Topologies Folder)	221
Devices (Devices)	221
User Sessions (Devices)	222
User Sessions Tab	222
Authentication (Device)	227
Authentication Status	227
Current User Counts	228
Global Authentication Settings	229
MAC Authentication Settings	230
Web Authentication Settings	230
General	231
Guest Networking	232
Web Page Banner	233
Convergence End-Point Settings	234
CEP Role Mappings	234
CEP Detection Tab	235
Add/Edit CEP Detection Rule	238
CEP Detection Settings	238
Ports (Authentication)	240
Authentication Mode	241
RFC3580 VLAN Authorization	242
Login Settings	243
MAC	244
802.1X	244
Web Auth	244
Quarantine	245
Auto Tracking	2/15

Automatic Re-Authentication	245
Authenticated User Counts	246
Convergence End-Point Access	247
RADIUS (Device)	248
Authentication Tab	248
RADIUS Authentication Client Settings	248
Authentication RADIUS Server(s) Table	250
Accounting Tab	251
RADIUS Accounting Client Settings	252
Accounting RADIUS Servers Table	253
RADIUS Authentication (Device)	255
RADIUS Authentication Client Settings	255
Authentication RADIUS Server(s) Table	257
RADIUS Authentication (Devices)	259
RADIUS Accounting (Device)	261
RADIUS Accounting Client Settings	262
Accounting RADIUS Servers Table	263
RADIUS Accounting (Devices)	265
Add/Edit RADIUS Server	266
Add RADIUS Accounting Server	269
Ports (Device)	271
Ports (Port Group)	273
Details View (Port Groups)	274
Add/Remove Ports (User-Defined Port Groups)	274
Add/Remove Ports	275
Port Authentication Configuration	278
Authentication Mode	278
Port Mode	278
RFC3580 VLAN Authorization Tab	280
Login Settings	281

4	Jiomatic Re-Authentication	. 285
λι	ıthenticated User Counts	.283
Co	onvergence End-Point Access	285
НС	ow To Use Policy	285
ı	How to Select on Add/Remove Windows	286
	Selecting single items	. 286
	Selecting multiple sequential items	. 286
	Selecting multiple non-sequential items	.286
ı	How to Create and Use Domains	287
	Creating a New Domain	. 287
	Opening a Domain	.288
	Assigning Devices to a Domain	. 288
	Removing Devices From a Domain	289
	Importing a File into a Domain	.289
	Exporting a Domain to a File	. 290
	Importing Data from a Domain	.290
	Saving a Domain	.290
	Renaming a Domain	. 290
	Deleting a Domain	291
	How to Create a Role	291
	Using the Role Tabs	291
	Modifying a Role	292
	Adding Services to Roles	.292
	Removing Services from a Role	.293
	Modifying a Role's Default Class of Service	293
	Modifying a Role's Default Access Control	. 293
	Modifying a Role's Description	. 293
	Modifying a Role's Ports	. 293
	Mapping a Role to an HTTP Redirect Group	. 294
	Deleting a Role	294

How to Assign a Default Role to a Port	294
Assigning and Clearing a Default Role	294
Assigning Default Roles to Ports	294
Clearing Default Roles from Ports	295
How to Create a Quarantine Role	295
Modifying the Quarantine Role	296
Modifying Default Values	296
Adding/Removing Services	296
Setting the Quarantine Role as the Default Role on a Port	296
How to Create a Service	297
Using the Service Tabs	297
Creating an Automated Service	298
Creating a Manual Service	298
Modifying a Service	298
Modifying a Service Description	299
Modifying a Service Name	299
Modifying the Roles for a Service	299
Adding a Service to Roles	299
Modifying the Rules for a Manual Service	300
Modifying an Automated Service	300
Deleting a Service	300
How to Create a Service Group	301
Creating a Service Group	301
Adding Services to a Service Group	301
Removing Services from a Service Group	301
How to Create or Modify a Rule	302
Creating a Rule	302
Disabling/Enabling a Rule	303
Deleting a Rule	303
Harrita Dafia a Data Limita	704

Defining Rate Limits	304
Removing a Rate Limit	305
How to Create a Class of Service	305
Creating a Class of Service	306
Creating Class of Service Port Groups	307
Deleting a Class of Service	308
How to Configure Transmit Queues	308
Transmit Queue Configuration	308
Transmit Queue Rate Shapers	309
How to Define Traffic Descriptions	309
How to Configure Flood Control	310
How to Create Global and Island VLANs	31
Creating a VLAN	312
Editing an Island VLAN ID	312
Deleting a VLAN	312
How to Create a Policy VLAN Island	
Creating a VLAN Island	313
Modifying a VLAN Island	313
Deleting a VLAN Island	312
How to Create a Network Resource	312
How to Add and Delete Devices	315
Adding a Single Device	316
Deleting Devices from the Database	316
How to Create a Port Group	317
Creating a Port Group	317
Adding Ports to a Port Group	317
Removing Ports from a Port Group	317
tremeControl Access Control	318
ExtremeControl Configuration	318
Futures Control Consum Folitan	716

All ExtremeControl Engines	319
ExtremeControl Configuration Considerations	319
ExtremeControl Configuration Tables	320
General Considerations	325
Considerations When Implementing Policy Roles	328
ExtremeWireless Controller Configuration	329
DNS Proxy Functionality for Registration and Remediation	329
Basic Operation	329
Backup DNS Server	330
Troubleshooting	330
Install the Assessment Agent Adapter on a Nessus Server	331
How to Configure Local RADIUS Termination at the ExtremeControl Engine	333
LDAP Authentication	334
User Authentication Considerations	334
Active Directory	334
Other LDAP Servers	335
Local Authentication	336
User Password Considerations	336
Certificate Configuration	336
EAP-TLS Certificate Requirements	336
Deploy ExtremeControl in an MSP or MSSP Environment	338
Configuring ExtremeCloud IQ - Site Engine Behind a NAT Router	338
Defining Interface Services	339
ExtremeControl Concepts	340
Overview of the Access Control Tab	340
ExtremeControl Engines	
Use Scenario	341
ExtremeControl VPN Deployment	343
Access Control Tab Structure	344
ExtramaControl Configuration	711

Rule Components	345
ExtremeControl Profiles	345
AAA Configurations	346
Portal Configurations	346
Access Policies	346
Registration	348
How Registration Works	349
Assessment	350
Assessment Remediation	352
How Remediation Works	353
End-System Zones	353
End-System Zone Use Cases	354
Enforcing	355
Advanced Enforce Options	356
MAC Locking	356
Notifications	357
Access Control	357
Configurations	358
AAA	359
Profiles	359
Captive Portals	359
Notifications	360
Vendor RADIUS Attributes	360
Add Radius Dictionary to ExtremeControl.	360
Global & Engine Settings	360
Configuration Evaluation Wizard	362
User Input	362
Authentication Results Tab	362
Authorization Results Tab	363
ExtremeControl Configuration Rules	366

Accessing ExtremeControl Configuration Rules	366
Viewing Rules in the Table	366
Creating and Editing Rules	368
Add/Edit Rule	370
Add/Edit User to Authentication Mapping	373
AAA Configurations Panel	377
AAA Configurations	379
Accessing the AAA Configuration	379
Basic AAA Configuration	379
Advanced AAA Configuration	381
AAA Configurations Panel	384
Manage LDAP Configurations	386
Add LDAP Configuration	387
Edit LDAP Configuration	393
Manage RADIUS Servers	399
Add/Edit RADIUS Server	401
Authentication Via ExtremeCloud IQ - Site Engine or Captive Portal	402
Configuration	402
Change Server Shared Secret	404
Manage RADIUS Attribute Configurations Window	405
Advanced RADIUS Server Configuration	406
Health Check for UDP	407
Policy Mapping Configuration	408
Column Definitions	409
Add/Edit Policy Mapping	412
Add LDAP Policy Mappings	416
Add Attribute Value to Policy Mapping	416
Edit LDAP Policy Mappings	418
Edit Attribute Value to Policy Mapping Window	418
Access Control Profiles	419

New/Edit ExtremeControl Profile	422
Authorization	423
Assessment	
Edit Assessment Configuration	426
Test Sets	427
Buttons	428
Manage Assessment Servers	428
Manage Assessment Settings	432
Create a Custom Scan for Agent-less Assessment	432
Portal Configuration Overview	436
Accessing the Portal Configuration	436
Default Portal Configuration	436
Network Settings	436
Administration	436
Website Configuration	437
Look and Feel	437
Guest Access and Registration	437
Authenticated Web Access	437
Authenticated Registration	437
Assessment / Remediation	437
External Captive Portal	438
Captive Portal Configuration	439
Accessing the Portal Configuration	439
Default portal configuration	440
Network Settings	440
Administration	440
Website Configuration	440
Look and Feel	440
Guest Web Access	445
Registration Settings	446

I hird-party guest registration	44/
Authenticated Web Access	448
Authentication	449
Redirection	450
Web Access Settings	450
Assessment/Remediation	451
Web Page Settings	451
Remediation Attempt Limits	453
Remediation Links	453
Custom Remediation Actions	453
Portal Web Page URLs	454
External captive portal service	454
External Captive Portal	455
Portal Registration Administration	456
Administration	456
Administration Web Page Settings	457
Manage Notifications	458
Notifications Table Buttons	458
Notifications Table	459
Enable Default Notifications	460
Add/Edit Notification	462
Conditions	465
Actions	466
Result	467
MAC Locking	467
MAC to IP Mappings	468
Access Control Engine Settings	469
Credentials	469
Switch Configuration	470
Admin Web Page Credentials	471

Admin Web Page Authentication	471
EAP-TLS Configuration	472
Network Settings	472
Manage DNS Configuration	474
Manage NTP Configuration	474
Manage SSH Configuration	474
SNMP Configuration	475
Device Type Detection	476
IP Address Resolution	477
Hostname Resolution	482
Username Resolution	483
Reauthentication	484
Miscellaneous	486
Port Link Control	488
NTLM Health Check	488
NetBIOS	489
Kerberos	489
Microsoft NAP	490
Auditing	491
ExtremeControl Engine Groups	491
ExtremeControl Access Control Group Editor	492
Add/Edit Device Type Group	496
End-Systems	498
End-Systems	498
Actions	503
Menu Buttons	504
End-System Events Tab	505
Add/Edit End-System Group	508
End-System Details	
Access Profile Tab	
End-System Tah	514

End-System Events Tab	515
Health Results Tab	516
Health Results	517
Health Result Details	519
Buttons and Paging Toolbar	520
Add/Edit Location Group	522
Create Time Group Window	524
Add/Edit User Group	526
Add/Edit User Group Window	529
Switches	530
Edit Switches in ExtremeControl Engine Group	532
Add Switches to ExtremeControl Engine Group	536
Advanced Switch Settings	539
All Access Control Engines	541
Access Control Engine Enforce Preview	543
Details (ExtremeControl Engine)	543
Details (ExtremeControl Engine Groups)	546
Status	546
Group	546
Engines	547
Access Control Configuration - Default	547
Load Balancing	548
Guest and IoT Configuration	548
Interfaces Window	549
Interface Modes	549
Services	550
DHCP/Kerberos Snooping	552
Captive Portal HTTP Mirroring	552
Tagged VLANs	552
Static Poute Configuration Window	557

How To Use Access Control	554
How to Use Device Type Profiling	555
Device Profiling Use Case	555
How to Configure LDAP for End Users and Hosts via Active Directory	564
How to Change the Assessment Agent Adapter Password	569
How to Set ExtremeControl Options	570
Advanced Settings	570
Assessment Server	571
Data Persistence	571
End-System Event Cache	572
Enforce Warning Settings	573
Setting Features Options	574
Notification Engine Options	574
Policy Defaults	574
Status Polling and Timeout	575
How to Set Up Registration	577
ExtremeControl Gateway Configuration	578
Identifying ExtremeControl Gateway Location	578
Defining the Unregistered Access Policy	578
Creating the Unregistered Access Policy	579
Configuring the Unregistered ExtremeControl Profile	581
Configuring Policy-Based Routing	582
Configuring the Access Control Tab (for ExtremeControl Gateways and Controllers)	584
How to Configure Pre-Registration	586
Configuring Pre-Registration	586
Pre-Registering Guest Users	591
Pre-Registering a Single User	591
Pre-Registering Multiple Users	593
How to Enable RADIUS Accounting	
Considerations for Fixed Switching Devices	
Considerations for ExtremeVOS/Switch Engine Devices	500

Guest and IoT Manager Configuration in ExtremeCloud IQ - Site Engine and A Control (Legacy)	
Connecting GIM to ExtremeControl	600
Configuring the RADIUS Protocol for GIM Authentication	601
Creating and Configuring a GIM Domain	601
Configuring GIM Authentication	603
Local Password Repository	603
LDAP	604
Configuring Multiple Active Directory Domains	606
Requirements	606
Validating Multiple AD Domain Functionality	606
Joining Multiple Active Directory Domains	606
Important Note	607
How to Set Up Access Policies and Policy Mappings	608
Setting Up Your Access Policies	609
How to Configure Credential Delivery for Secure Guest Access	613
Configuration Steps	613
How Secure Guest Access Works	620
How to Configure Verification for Guest Registration	624
Configuration Steps	624
How User Verification Works	627
Configure Sponsorship for Guest Registration	630
How to Implement Facebook Registration	633
Requirements	633
Creating a Facebook Application	634
Portal Configuration	640
How Facebook Registration Works	642
Special Deployment Considerations	642
Wireless Clients	642
Networks using DNS Proxy	642
How to Implement Google Registration	644

	Requirements	644
	Creating a Google Application	645
	Portal Configuration	650
	How Google Registration Works	652
	Special Deployment Considerations	652
	Networks using DNS Proxy	652
Нс	w to Implement Microsoft Registration	654
	Requirements	654
	Creating a Microsoft Application	655
	Portal Configuration	660
	How Microsoft Registration Works	661
	Special Deployment Considerations	661
	Networks using DNS Proxy	662
Нс	w to Implement Yahoo Registration	663
	Requirements	663
	Creating a Yahoo Application	664
	Portal Configuration	666
	How Yahoo Registration Works	668
	Special Deployment Considerations	668
	Networks using DNS Proxy	668
Нс	w to Implement Salesforce Registration	. 670
	Requirements	670
	Creating a Salesforce Application	671
	Portal Configuration	680
	How Salesforce Registration Works	682
	Special Deployment Considerations	682
	Networks using DNS Proxy	682
۸ ـ		CO 4

Getting Started with ExtremeControl

ExtremeCloud IQ - Site Engine's **Control** tab provides end-system and user identity reports and control capabilities, allowing better visibility and control for IT analysts, troubleshooters, and the helpdesk.

Access Requirements

To view the reports in the **Control** tab, you must be a member of an authorization group that has been assigned the appropriate capabilities:

- XIQ-SE OneView > Access OneView
- XIQ-SE OneView > ExtremeControl > Access OneView Identity and Access Reports
- XIQ-SE OneView > ExtremeControl > OneView End-Systems Read Access or Read/Write Access

NOTE: The ExtremeCloud IQ - Site Engine, ExtremeControl, and ExtremeAnalytics Virtual Engine Installation Guide includes an overview of ExtremeCloud IQ - Site Engine, ExtremeControl, and ExtremeAnalytics virtual engine deployment requirements and how to deploy a virtual engine on a VMware® and Hyper-V server.

Navigating the Control Tab

Selecting **Control** in the Menu Bar at the top of ExtremeCloud IQ - Site Engine opens the **Control** tab. The **Control** tab provides access to four sub-tabs:

- <u>Dashboard</u> Displays summary ExtremeCloud IQ Site Engine data including end-system data, system-level information, system events, ExtremeControl engine information, and network health.
- Policy Enables you to create policy profiles, called roles, assigned to the ports in your network.
- Access Control Allows you to configure how end-users connect to your network.
- End-Systems Displays information about end-users connected to your network.
- Reports Provides a variety of system reports that give information about your devices, ports, and network traffic.

Additionally, the **Menu** icon () at the top of the screen provides links to additional information about your version of ExtremeCloud IQ - Site Engine.

Dashboard

Select the **Dashboard** tab to view information about engines and end-systems.

Overview

Provides an overview of end-system connection information. For a description of each report, select the **Info** button in the upper right corner of the view. Enable and disable data display in each chart by selecting the data set in the chart legend. For example, if one segment represents a disproportionately large percentage of the total, mouse over the segment legend to the right of the chart and select it to remove it from the pie chart.

System

Provides system-level information for engines and end-systems. For a description of each report, select the **Info** button in the upper right corner of the view.

Health

Provides reports on end-system assessment and state information. For a description of each report, select the **Info** button in the upper right corner of the view.

Policy

Selecting the **Policy** tab lets you create policies for your network. It allows you to create policies for users and ports, enabling network engineers, information technology administrators, and business managers to work together to create the appropriate network experience for each user in their organization.

Access Control

The Access Control tab lets you manage the end user connection experience and control network access based on a variety of criteria including authentication, user name, MAC address, time of day, and location. The Access Control tab comes with a default ExtremeControl Configuration which is automatically assigned to your ExtremeControl engine. You can use this default configuration as is, or make changes to the default configuration, if desired.

End-Systems

Selecting the **End-Systems** tab displays end-system connection information, and lets you monitor end-system events and view the health results from an end-system's assessment. Double-click on any row in the table to open a browser window that displays End-System Details.

Reports

The **Reports** tab allows you to view information about the end-systems connecting to your network, ExtremeControl authentication information, and the top services and roles based on policy rules. Available reports are accessible via the **Reports** drop-down list at the top of the tab and are grouped into the following reporting areas:

- End-Systems
- Access Control

- Access Control Health
- Policy

Policy


The **Policy** tab, contained in the **Control** tab of ExtremeCloud IQ - Site Engine is a configuration tool that simplifies the creation and enforcement of policies on networks, enabling network engineers, information technology administrators, and business managers to work together to create the appropriate network experience for each user in their organization.

The **Policy** tab enables you to create policy profiles, called roles, which are assigned to the ports in your network. These roles are based on the existing business functions in your company and consist of services that you create, made up of traffic classification rules. Roles provide four key policy features: traffic containment, traffic filtering, traffic security, and traffic prioritization.


Use the following summary to guide you through the basic steps for using the Policy tab.

- 1. Create your Policy Domains (see How to Create and Use Domains.)
- 2. Add your devices to the ExtremeCloud IQ Site Engine Database and assign them to the appropriate domain.
- 3. If desired, group your ports into port groups (see How to Create a Port Group).
- 4. Create services (see How to Create a Service).
- 5. If desired, group services into service groups (see How to Create a Service Group).
- 6. Create roles (see How to Create a Role).
- 7. Write your configuration to your devices (see Enforcing).

The illustration below shows the **Policy** tab relationship hierarchy, with Rules at the base to define specific packet handling behaviors, Roles at the top to identify specific job functions in the organization, and Services in the middle, providing the interface between the two layers.



Using policy configuration tools, you can create multiple roles tailored to your specific needs and set a default policy for some or all of your network devices and ports. These policies can be deployed on multiple devices throughout your switch fabric.



The topic covers the following features:

- Understanding Policy Domains
- <u>Understanding Roles</u>
- Understanding Services
- Working with Service Groups
- Understanding Traffic Classification Rules

- Adding Devices
- Viewing Port Configuration Information
- Working with Port Groups
- Working with VLANs
- Viewing Classes of Service
- Saving the Domain
- Enforcing
- Verifying
- AP Aware

Understanding Policy Domains

The **Policy** tab provides the ability to create multiple policy configurations by enabling you to group your roles and devices into Policy Domains. A Policy Domain contains any number of roles and a set of devices that are uniquely assigned to that particular domain. Policy Domains are centrally managed in the database and shared between **Policy** tab clients.

The first time you launch the **Policy** tab, you are in the Default Policy Domain. You can manage your entire network in the Default Policy Domain, or you can create multiple domains each with a different policy configuration, and assign your network devices to the appropriate domain. The Default Policy Domain is pre-configured with roles and rules. The roles, services, rules, VLAN membership, and class of service in this initial configuration define a suggested implementation of how network traffic can be handled. This is a starting point for a new policy deployment and often needs customization to fully leverage the power of a policy-enabled network.

For more information about domains, see Policy Domains in the Concepts Help topic.

In the Quick Tour, we'll use the Default Policy Domain as a way to explore the basic features and functionality of the **Policy** tab. Later, the Default Policy Domain can be useful as you create your own Policy Domains.

If you have just launched the **Policy** tab for the first time, you are in the Default Policy Domain and you can proceed to the next step, <u>Understanding Roles</u>. If someone else has been using the **Policy** tab before you, use the following steps to create a demonstration domain you can use for the Quick Tour.

NOTE: If someone uses the **Policy** tab before you, ExtremeCloud IQ - Site Engine can prompt to save the previous domain's configuration when you create the new domain. Save the previous domain's configuration if you are going to use that configuration in the future.

To create a policy domain:

- 1. Select Open/Manage Domains > Create Domain. Enter the domain name Demonstration Domain for the new domain and select OK. The new Demonstration Domain opens.
- 2. Select Open/Manage Domains > Assign Devices to Domain. Select the devices to add to the Domain and select OK. The device is added to the left-panel Devices tab.
- 3. Select the left-panel Roles/Services tab. Right-click Roles, Services, or Service Groups and select Create Role, Create Services, or Create Service Groups, respectively to create a role, service, or service group for the domain. For additional information on creating a role, service group, or service, see How to Create a Role, How to Create a Service, or How to Create a Service Group.
- 4. Select the left-panel **Class of Service** tab. Right-click Class of Service and select **Create COS** to create a class of service for the domain. For more information on creating a class of service, see How to Create a Class of Service.
- 5. Select the left-panel **VLANs** tab. Right-click Global VLANs and select **Create VLAN** for the domain. For more information on creating VLANs, see How to Create a VLAN.
- 6. Select the left-panel Network Resources tab. Right-click Network Resources or Global Network Resources (All Domains) and select Create Network Resource to create a network resource for the domain. You can also right-click Network Resource Topologies and select Create Network Resource Topology to create a network resource topology for the domain. For more information on creating a network resource or network resource topology, see How to Create a Network Resource.
- 7. Select **Open/Manage Domains > Save Domain**. The data elements are saved to the new Demonstration Domain.

For more information:

• How to Create and Use Domains

Now that you've created the demonstration domain, we can explore the **Policy** tab in a little more depth.

Understanding Roles

Roles are usually designed to reflect different users in your organization and to provide customized access capabilities based on the role users have in your organization. For example, accounting and engineering personnel have different network access and priority needs and therefore can have different roles.


To view information about existing roles:

- 1. Select the left-panel Roles/Services tab in the Policy tab main window.
- 2. Select the left-panel **Roles** sub-tab in the Roles/Services tab.
- 3. Select a role name to see a description of the role.
- 4. Select the various roles listed in the left panel, and in the right panel you'll see tabs that

display specific information for each role. Select the right-panel tabs to see the information they contain.

A role can be made up of one or more network access services defined in the **Policy** tab. These services determine how network traffic is handled at any network access point configured to use that role. A role can also contain default access control (VLAN) and/or class of service designations applied to traffic not handled specifically by the services contained in the role. A role can contain any number of services or service groups.

To filter through roles easily, select the Show Editable Columns drop down and select iif you wannt to hide or show editable information.



Roles are assigned to users during the authentication process. When a user successfully authenticates, the port is opened, and if a role is assigned to the user, that role is applied to the port. A role can also be directly assigned to a port as a default role for instances when authenticated users are not assigned a role. If an end user on a port is not assigned a role when logging in (authenticating), or if authentication is inactive on a port, then the port uses its default role. However, if a user is assigned a role upon login, then that role overrides any default role on the port.

To create and define a role, right-click Roles and select Create Role.

To create a role:

- 1. In the **Policy** tab left panel, select the **Roles/Services** tab.
- 2. Select the Roles sub-tab.
- 3. Right-click the Roles folder, and select Create Role.
- 4. Enter the role name Office Assistant in the highlighted box and press Ok.

For more information:

- Role
- How to Create a Role

Role Summary Column

The Summary column shows the data for the row iin a condensed form. Hovering over the cell displays the summary data in an expanded, easy too read format. This includes the rule and service usage information, traffic description, action details, automated service relevant network resources, and toplogy information.


Understanding Services

Roles can be made up of one or more network access services. These services determine how network traffic is handled at any network access point configured to use that role. The **Policy** tab enables you to create Local Services (services unique to the current domain) and Global Services (services common to all domains).

Services can be one of two types:

- Manual Service Contain customized classification rules you create.
- Automated Service Associated with a particular set of network resources.

Manual services contain one or more traffic classification rules that define how a network access point handles traffic for a particular network service or application. For example, you might create a Manual service called "Restricted Employee" that contains a classification rule that discards TCP HTTP traffic.



We are creating a Manual service and then adding it to a role. Right now, lets take a look at the services in the domain.

To view information about existing services:

- 1. Select the left-panel Roles/Services tab in the Policy tab main window.
- 2. Expand the **Service Repository** folder and then the **Local Services** folder.


- 3. Expand the **Services** folder to view a list of services.
- 4. Expand a service or two to see the individual classification rules that make up the service.
- 5. Select a service or two in the left-panel to see the right-panel tabs that display specific information for each service. Select the right-panel tabs to see the information they contain.

For more information:

- Service
- How to Create a Service

Working with Service Groups

Services can be grouped together into Service Groups. This enables you to add a set of services to one or more roles.



To view information about existing service groups:

- 1. Select the left-panel Service Repository tab in the Policy tab main window.
- 2. Expand the Service Repository folder and then the Local Services folder. Expand the Service Groups folder.
- 3. Expand the Acceptable Use Policy service group to see its services. These services are also listed under the Services folder.

After you have defined and created your services, you can easily create a Service Group and then add your services to the group.

To create a service group:

- 1. Select the left-panel Roles/Services tab in the Policy tab main window.
- 2. Expand the **Service Repository** folder and then the **Local Services** folder.
- 3. Right-click the Service Groups folder and select Create Service Group.
- 4. Enter the service group name **Trusted User** in the highlighted box and press **Enter**.
- 5. Right-click Service Group, select **Add/Remove Services** and add one or two of the existing Acceptable Use Policy service groups into the Trusted User service group.

For more information:

• How to Create a Service Group

Understanding Traffic Classification Rules

Traffic classification rules enable you to assign access control (VLAN membership) and/or class of service to your network traffic based on the traffic's classification type. Classification types are derived from Layers 2, 3, and 4 of the OSI model and all network traffic can be classified according to specific layer 2/3/4 information contained in each frame.

A traffic classification rule has two main parts:

- Traffic Description Identifies the traffic classification type for the rule.
- Actions Apply access control, class of service, security, and/or accounting behavior to packets matching the rule.

To view existing rules:

- 1. In the left-panel, navigate to the **Service Groups** tab (Roles/Services > Service Repository > Local Services > Service Groups) and expand the **Acceptable Use Policy** service group.
- 2. Expand the **Deny Unsupported Protocol Access** service and select the **Discard AppleTalk** rule.
- 3. Use the **Edit** button to add a description to the service, for example: **AppleTalk not** supported on this network.

For more information:

- Rule
- Traffic Classification Rules
- How to Create or Modify a Rule

Adding Devices

The first step in adding network devices to **Policy** tab, is to add the devices to the Extreme Cloud IQ - Site Engine database. You do this initially, by using the **Discovered** tab on the **Network** tab. This section assumes you have already done this. If you need more information, refer to the **Network** tab Help page.

When you add devices to the ExtremeCloud IQ - Site Engine database, you must assign the devices to a Policy Domain using the **Policy** tab. As soon as the devices are assigned to a domain, they are automatically displayed in the **Policy** tab device tree. Only devices assigned to the domain you are currently viewing are displayed.

To assign devices to a domain:

- 1. In the **Policy** tab main window, right-click **Devices** and select **Assign Devices to Domain**. The Assign Devices to Domain window opens.
 - In the left panel, the Unassigned device tree contains all the devices in the database not assigned to a domain. The right panel displays the devices in the current domain.
- 2. For the Quick Tour, select a couple of devices to add to the domain and select **Add**. Select **OK** to add the devices.

You can also use this window to remove a device from the current domain. This removes the device from the current domain and places it in the Unassigned folder. It does not delete the device from the ExtremeCloud IQ - Site Engine database.

For more information:

- How to Add and Delete Devices
- How to Create and Use Domains

Viewing Port Configuration Information

After importing devices into the **Policy** tab, you can view and configure their ports by selecting a device and displaying its ports in the right-panel **Details View** tab or **Ports** tab.

To view port configuration information:

- 1. Select the left-panel **Devices** tab in the **Policy** tab main window.
- 2. Expand the **Devices** folder and select a device.
- 3. In the right-panel **Ports** tab, expand a **Ports** or **Slot** folder to display ports on the device.
- 4. Right-click a port and select Current Domain > Show Role Details.
- 5. Set Default Role, if necessary.

Working with Port Groups

The **Policy** tab enables you to group ports into User-Defined Port Groups, similar to the way you can group services into service groups. Port groups enable you to configure multiple ports on the same device or on different devices, at the same time. The **Policy** tab also provides you with Pre-Defined Port Groups. Every time one of the Pre-Defined Port Groups is accessed, the **Policy** tab goes to the devices in the current domain and retrieves the ports which fit the pre-defined characteristics of the port group.

To view pre-defined port groups:

- 1. Select the left-panel Port Groups tab in the Policy tab main window.
- 2. Highlight a port group to display information for that port group.

For more information:

• Pre-Defined Port Groups

Working with VLANS

All traffic in a **Policy** tab network is assigned membership in a VLAN. Roles are used to assign VLAN membership to traffic either through the role's default access control or through the role's services which can include traffic classification rules that assign VLAN membership (access control).

When you open a new domain, the Global VLANs folder is prepopulated with the Default VLAN (not to be confused with a default VLAN assigned to a role, although the Default VLAN could be a default VLAN for a role). You can then create additional VLANs and assign them as default access control for a role and/or use them to define traffic classification rules. You can view the roles and services associated with a VLAN by selecting the VLAN in the left-panel. You can also make role and service changes from this window.

Island VLANs are used in Policy VLAN Islands, which enable you to deploy a policy across your network, while restricting user access to only selected local devices. The **Policy** tab enables you to view currently configured Island VLAN information.

To view VLANs:

- 1. From the VLANs tab, expand the Global VLANs folder to see individual VLANs.
- 2. Select the Default VLAN listed and view the VLAN information in the right panel.

For more information:

- How to Create a VLAN
- General Tab (VLAN)
- Policy VLAN Islands

Viewing Classes of Service

The **Policy** tab lets you create a class of service (CoS) that includes one or more of the following components: an 802.1p priority, an IP type of service (ToS) value, rate limits, and transmit queue configuration. You can then assign the class of service as a classification rule action, as part of the definition of an automated service, or as a role default.

To view Classes of Service:

- 1. From the **Policy** tab, select the **Class of Service** tab from the left-hand panel. The Class of Service section expands.
 - Notice that the window is pre-populated with eight static classes of service, each associated with one of the 802.1p priorities (0-7). You can use these classes of service as is, or configure them to include ToS/DSCP, drop precedence, rate limit, and/or transmit queue values. You can also rename them, if desired. In addition, you can also create your own classes of service (user-defined CoS).
- 2. Select the **Class of Service** and all information related to the Class of Service selected is displayed in the right-panel.

For more information:

- Getting Started with Class of Service
- How to Define Rate Limits
- How to Configure Transmit Queues
- How to Create a Class of Service

Saving the Domain

After changing a policy domain, save the domain. This notifies all clients viewing the domain there is a change, which prevents them from saving a domain with an incorrect configuration. The system automatically updates their view with the new configuration.

To save a domain, select **Open/Manage Domains > Save Domain**.

The domain is saved and automatically updates for all clients viewing the domain. To discard unsaved changes you made to a domain, open the **Open/Manage Domains > Open Domain** menu and select the domain in which you are currently working.

For more information:

How to Create and Use Domains

Enforcing

Any time you add, make a change to, or delete a role or any part of it (any of its services and/or rules), the devices in your current domain need to be informed of the change so that your revised policy configuration can take effect. This is accomplished by enforcing — writing your policy configuration to a device or devices. Enforce operations are performed only on the current domain.

To enforce to all devices in the current domain, select **Open/Manage Domains > Enforce Domain**. To enforce to a single device, right-click the device and select **Enforce**.

Enforce Preview

The Enforce preview tool has a very similar setup to the Enforcing Domain tool. To view the enforce preview, select **Open/Manage Domains > Enforce Preview** and select the device to preview from the left dropdown.

Note: If the device has a red exclamation type next to it in the left panel, then it is incompatible with the domain configuration and should be corrected.

Enforcing preview shows you a summary of the stats and info, roles, rules, and services on device. The three preview tabs include:

Device Stats & Info: Shows information on supported role/rule counts, etc.

Roles & Rules: Shows a grid panel with roles and rules that will enforce the device. If supported, it will show a green circle. A yellow circle indicates a rule not being supported, and a red circle denotes a role not being supported. Right-click and select View/Edit which will close enforce preview and bring you to the item you wish to make changes to.

Classes of Service: Shows details of the Class of Service and the related rate limit configuration.

Rule Counts Reported by Devices

Every device has a maximum number of rules that it can follow. Going over the max number of rules on a device will create enforce failures. The max supported rules by rule type are mainly a concern for EXOS/Switch Engine device, which now report the max a type supports via the value returned for etsysPolicyRuleAttributeMaxCreatable for any rule type in that group. For example, reading either instance 1 (macSource(1)) or 2 (macDestination(2)) will return the supported number of layer 2 (MAC) rules. The 4 rule "types" and the rule types () that these include are:

- MAC
 - macSource(1)
 - macDestination(2)
- IPv4
 - ip4Source(12)
 - ip4Destination(13)
 - ipFragment(14)
 - udpSourcePort(15)
 - udpDestinationPort(16)
 - tcpSourcePort(17)
 - tcpDestinationPort(18)
 - ipTtl(20)

- ipTos(21)
- ipType(22),
- IPv6
 - ip6Destination(10)
- L2
- etherType(25)

The total max supported number of rules for EXOS/Switch Engine devices is the sum of these 4 types, NOT the value returned by etsysPolicyRulesMaxEntries (due to that including other things by the FW).

The devices supported number of rules is only read when the device is added to the domain, the firmware is upgraded, or the device is manually refreshed.

For more information:

• Enforcing

Verifying

To determine if the roles currently in effect on your domain devices match the set of roles defined in your current Policy Domain configuration, use the Verify feature.

AP Aware

An AP is assigned "AP Aware," all traffic through this port will not need authentication. This new Role default action is configurable via a new AP Aware setting in the role configurations view. To enable AP Aware:

- 1. Select the left-panel Roles/Services tab in the Policy tab main window.
- 2. Select the left-panel **Roles** sub-tab in the Roles/Services tab.
- 3. Select a role name to see a description of the role.
- 4. Using the scroll bar, scroll to find the AP Aware column.
- 5. Double-click **Disabled**, and in the drop-down, select **Enabled**.

When enforce or verify occurs, the secondary logic runs which inspects all AP Aware enabled roles, and for each role finds all in-use VLANs (rule actions, role default action) and automatically adds them to that role's tagged VLAN egress list if they are not already present. This is then used for the enforce/verify logic, and returned to the client so the domain is updated accordingly.

The domain data can change from doing an enforce/verify, and needs to be saved.

For more information:

Verifying

Policy Configuration Considerations

Review the following configuration considerations when installing and configuring ExtremeCloud IQ - Site Engine's **Policy** tab.

- General Considerations
 - Authenticating without Policy
 - Terminating Role Override Sessions
 - Port-Level MAC to Role Mappings
 - Import From Device
 - Flood Control
- C1 Considerations
 - Policy Support
 - Rule Limits
- N-Series Considerations
 - Role Precedence for the N-Series Platinum
- C2 and B2 Considerations
- C3 and B3 Considerations
- Mixed-Stack C2/C3 and B2/B3 Considerations
- 7100 Considerations
- ExtremeControl Controller Configuration
- Wireless Controller Configuration

General Considerations

Authenticating without Policy

This section discusses how authentication works in a network where end users must authenticate, but there are no roles (policy) for authenticated users defined on the network devices.

The following table shows Authentication Behavior for each device type when the authenticated role is not defined on the device:

Authentication Type	K-Series, S-Series, N-Series Gold and Platinum	E6/E7	E 1	RoamAbout R2 RoamAbout AP3000	C2/B2
802.1X	Successful	Successful	Successful	Successful	Successful
MAC	Successful	Successful	Successful	Successful	Successful
Web-Based	Successful	Successful on firmware version 5.06.x. Failed on older firmware versions.	Successful	Web-Based Auth Not Supported	Successful

The following table shows Authenticated Traffic Behavior for each device type when the authenticated role is not defined on the device:

Authentication Type	N-Series Gold and Platinum 4.11 and earlier	K-Series, S-Series, N-Series 5.01 and later Gold and Platinum	E6/E7	E1	RoamAbout R2 RoamAbout AP3000	C2/B2
802.1X	1	3	2	2	3	2
MAC	1	3	2	2	3	2
Web-Based	1	3	2	2	Web-Based Auth Not Supported	2

- 1 Traffic is forwarded based on the 802.1Q PVID and 802.1p priority for the port, regardless of whether the port has been assigned a default role. Authenticated users display a current role of "None" in the Port Usage tab.
- 2 Traffic is forwarded based on the port's default role and authenticated users will display the default role as their current role in the **Port Usage** tab. If no default role has been assigned to the port, the port's 802.1Q PVID and 802.1p priority are used, and the current role will be "None."
- **3** Traffic is forwarded based on the Invalid Role Action configuration at the device level in the **Policy** tab.

Terminating Role Override Sessions

On Port Usage tabs, you cannot terminate Role Override (IP) or Role Override (MAC) sessions created through the CLI (command line interface).

Port-Level MAC to Role Mappings

Enforcing port-level MAC to Role mappings could potentially remove rules as an intrusion detection response.

Import From Device

If you perform a Verify operation following an Import Policy Configuration from Device, the Verify can fail. This is because the import operation imports only roles and rules from the device, not the complete policy configuration.

Also, if you import from more than one device and the configuration is not the same on each device, Verify fails. This is because the imported configuration will not match the configuration on any one device.

Flood Control

Individual Class of Service granularity is unsupported on fixed switches, so if any CoS is assigned a Flood Control rate, all Class of Service on these devices use that rate.

C1 Considerations

Review the following considerations prior to configuring policy on C1 devices:

Policy Support

Policy support on C1 devices utilizes both a port-level role and a device-level role. In the **Policy** tab, a role is a set of network access services made up of traffic classification rules. It can also contain default Access Control (VLAN) and/or Class of Service settings applied to traffic not handled specifically by the rules contained in the role. Although both the device-level and port-level roles can contain all of these components, only certain portions of each role are used when applied to a port on a C1 device.

On the C1, classification rules are implemented at the device level through a device-level role. The **Policy** tab enables you to set a unique device-level role for each C1 device. The device-level role is a regular role that defines how inbound traffic is handled in terms of classification rules and default Class of Service assignment. In other words, all classification rules are taken from the device-level role, and any rules defined in the port-level role are ignored when applied to a port. The Class of Service setting is also implemented through the device-level role and ignored in the port-level role. However, the default Access Control setting of the device-level role is ignored, and is defined through the port-level role.

Classification rules from the device-level role are only applied to ports which also have a port-level role applied (either statically or dynamically). This enables you to exclude the device-level role from uplink ports and hosts ports, by not applying a port-level role to these ports and not enabling authentication on them.

When a port-level role is applied to a port, it overrides any PVID and Class of Service settings defined on the port through Console or local management. When a device-level role is applied to a port, it also overrides these PVID and Class of Service settings, and overrides any Class of

Service setting defined in the port-level role. It does **not** override any default Access Control setting defined in the port-level role.

In addition, if the port-level role's default Access Control is configured to deny traffic, then **all** inbound traffic will be discarded even if it matches a (forward) classification rule.

Rule Limits

C1 devices limit the number of rules you can create for some classification types. Refer to the C1 information in the ExtremeCloud IQ - Site Engine Release Notes to see which classification types limit the number of rules.

N-Series Considerations

Review the following considerations prior to configuring policy on N-Series devices:

Role Precedence for the N-Series Platinum

The following precedence determines the role (policy) that is being applied on a user/port on a N-Series Platinum device. The precedence used depends on whether the device is configured for multi-user authentication or single user authentication.

Multi-User Authentication:

Devices configured with multi-user authentication use the following precedence when applying a role on a user/port (starting with the highest precedence):

MAC override policy

Authenticated role

MAC-to-Role mapping

IP override policy

IP-to-Role mapping

VLAN-to-Role mapping

Default port role

Single User Authentication:

Devices configured with single user authentication use the following precedence when applying a role on a user/port (starting with the highest precedence):

MAC override policy

MAC-to-Role mapping

IP override policy

IP-to-Role mapping

Authenticated role

VLAN-to-Role mapping

Default port role

C2 and B2 Considerations

Review the following considerations prior to configuring policy on C2 and B2 devices.

- When TCI Overwrite is enabled on a role, C2 and B2 devices support rewriting the 802.1p bit (CoS values) but not the 802.1Q bit (VLAN ID).
- On C2 and B2 gigabit and 10/100 ports, the number of rules per port is restricted. Refer to your C2 and B2 firmware release notes for the maximum number of rules that can be utilized on a port.
- C2 and B2 10/100 ports support two priority-based rate limits (inbound only). When creating a rate limit to be used on C2 and B2 10/100 ports, create the limit with either Low priority to associate the rate limit with priorities 0-3 or High priority to associate the rate limit with priorities 4-7. You can specify both Low and High priorities if you want to associate the rate limit with priorities 0-7.
- C2 and B2 devices do not support setting a default role on a logical port.
- On C2 and B2 devices, it is strongly recommended that you do not enforce rules that assign a Class of Service (CoS) that includes Priority 7. Doing so will interfere with stack communication.
- C2 and B2 devices do not permit a mask for an IP type of service (ToS) rewrite value associated with a class of service (CoS); they will always use ff.
- C2 and B2 devices do not support VLAN ID traffic classification rules. C2 devices (firmware 3.02.xx and newer) and B2 devices (firmware 2.xx.xx) support device-level VLAN to Role mapping. However, VLAN ID traffic classification rules can be configured on C2 devices with firmware versions 3.01.xx or older, using CLI.
- B2 only. Each port on a policy-enabled B2 switch can support up to 100 rules and up to 10 masks. The
 maximum number of unique rules in a single switch or B2 stack is 100, while the maximum number of
 unique masks is 18. These unique rules and masks can be shared across any and all ports in a stack or
 switch.

C3 and B3 Considerations

Review the following considerations prior to configuring policy on C3 and B3 devices.

- B3/C3 devices do not support TCI Overwrite. The B3/C3 does not overwrite 802.1Q VLAN bits, but overwrites the 802.1p Priority bits.
- B3/C3 devices do not support Layer 3 ICMP rules.
- B3/C3 devices support role-based rate limiting. However, on the B3/C3, class of service inbound rate limiting works only on policy roles, not on policy rules.
- C3G and B3 devices have the following additional limitations:
 - Maximum 100 rules per policy role.
 - A system limitation of 768 unique rules.
 - Maximum of 15 roles.
- C3 and B3 devices do not support setting a default role on a logical port.

Mixed-Stack C2/C3 and B2/B3 Considerations

Review the following considerations prior to configuring policy on mixed stacks of C2/C3 and B2/B3 devices.

NOTE: While you can create mixed stacks of C2/C3 devices and mixed stacks of B2/B3 devices, you should not create mixed stacks of C and B devices (e.g. mixed stacks of C2/B2 or C3/B3 devices).

- It is strongly recommended that a C3 device be configured as the controller in a mixed C2/C3 stack.
- It is strongly recommended that a B3 device be configured as the controller in a mixed B2/B3 stack.
- When you have a mixed stack, all devices in the stack have the rule type and Class of Service limitations of a C3 or B3 device, despite the fact that the stack can report itself as a C2 or a B2. The device type that the stack reports is based on what switch is set as the controller.
- Mixed stacks with a B3/C3 controller support role-based rate limiting, however, class of service inbound rate limiting works only on policy roles, not on policy rules.
- A mixed stack containing a C2H or a B2 has the following limitations:
 - A single role limitation of 100 rules and 10 masks.
 - A system limitation of 100 unique rules and 18 unique masks.
 - No support for Layer 2 rules or Layer 3 ICMP type rules.
 - Maximum of 15 roles.
 - No support for rate limiting.
- A mixed stack containing a C2G has the following limitations:
 - A single role limitation of 100 rules and 10 masks.
 - A system limitation of 768 unique rules.
 - No support for Layer 2 rules.
 - Maximum of 15 roles.
 - No support for rate limiting.
- When adding a new device to a mixed stack, the ports should not go active unless the stack supports the policy configuration. When a device has joined the stack, no roles should be enforced that are not supported on all devices. For example:

A C2K is added to an existing C3 stack.

- If the number of masks in the C3 stack's current configuration exceed those permitted by the C2K, its ports cannot go active.
- When the C2K joins the stack, no roles can be enforced that exceed the limitations of any device.

7100 Considerations

- 7100 devices only support fixed IRL index reference mappings for the static CoS. The IRL Index for the CoS needs to match the priority. This is the default configuration for domains, but if it is changed for a static CoS, enforce will fail.
- 7100 devices only support fixed TXQ index reference mappings for the static CoS. The TXQ Index for the CoS needs to match the priority. This is the default configuration for domains, but if it is changed for a static CoS, enforce will fail.

- 7100 devices only support fixed COS transmit queue mappings. The transmit queue specified for a Class of Service must match the 802.1p priority, or enforce will fail.
- TCI Overwrite configuration is not supported on the 7100. It is always enabled, and cannot be turned on or off using the Policy tab.

ExtremeControl Controller Configuration

Review the following considerations prior to configuring policy on ExtremeControl Controller devices.

ExtremeControl Controllers Require Separate Domains

ExtremeControl Controllers must by assigned to their own unique policy domain and cannot be combined with other switch types in a domain.

Modifying ExtremeControl Controllers Preconfigured Policy

ExtremeControl Controllers are shipped with a default policy configuration already configured on the device. To modify this default policy configuration, you must create a domain for the ExtremeControl Controller, assign the ExtremeControl Controller to the domain, then import the policy configuration from the device into the Policy tab (File > Import > Policy Configuration from Device). You can then alter the policy configuration to define the authorization levels for the ExtremeControl process, as appropriate for your environment. If assessment will be enabled in the Extreme Networks ExtremeControl solution, you must add classifications rules to the Quarantine and Assessing policies to permit traffic to be forwarded to the assessment servers deployed on the network. When you have finished modifying the policy configuration, you must enforce it back to the ExtremeControl Controller.

NOTE: If you are using assisted remediation and quarantined end-users will be required to download remediation files via FTP, you will also need to add a rule to the Quarantine policy configuration that opens up ports 49152-65535. If you are concerned with security, you can configure your FTP server to use a smaller range of ports.

Modifying the Downstream Default Policy

Depending on the network configuration or circumstances, it's possible that traffic from the upstream side could be rerouted to the ExtremeControl Controller where it would be authenticated using the upstream source IP address. To avoid this problem, add a Layer 3 IP Address Source rule to the downstream default policy configured on the ExtremeControl Controller, using the upstream IP subnets (or critical servers located in the upstream) and containing the traffic to a VLAN.

Configuring LAG on ExtremeControl Controllers

This section provides instructions for configuring LAG (link aggregation) on your ExtremeControl Controller appliance. The instructions vary depending on whether you are configuring LAG on a Layer 2 or Layer 3 ExtremeControl Controller.

Configuring LAG on Layer 3 ExtremeControl Controllers - Upstream Ports

- 1. Configure LAG on the ExtremeControl Controller PEP (Policy Enforcement Point) using the CLI (Command Line Interface).
- 2. Use the **Policy** tab to assign the appropriate upstream role as the default role on the port. For instructions, see <u>Assigning Default Roles to Ports</u>.

Configuring LAG on Layer 3 ExtremeControl Controllers - Downstream Ports

- 1. Configure LAG on the ExtremeControl Controller PEP (Policy Enforcement Point) using the CLI (Command Line Interface).
- 2. In the **Policy** tab options (Tools > Options), display the Ports panel and uncheck the Hide Logical Ports option.
- 3. Use the **Policy** tab to assign the appropriate downstream role as the default role on the port. For instructions, see Assigning Default Roles to Ports.

Configuring LAG on Layer 2 ExtremeControl Controllers - Upstream Ports

- 1. Configure LAG on the ExtremeControl Controller PEP (Policy Enforcement Point) using the CLI (Command Line Interface).
- 2. In the **Policy** tab options (Tools > Options), display the Ports panel and uncheck the Hide Logical Ports option.
- 3. Use the **Policy** tab to assign the appropriate upstream role as the default role on the port. For instructions, see Assigning Default Roles to Ports.

Configuring LAG on Layer 2 ExtremeControl Controllers - Downstream Ports

- 1. Configure LAG on the ExtremeControl Controller PEP (Policy Enforcement Point) using the CLI (Command Line Interface).
- 2. In the **Policy** tab options (Tools > Options), display the Ports panel and uncheck the Hide Logical Ports option.
- 3. Use the **Policy** tab to assign the appropriate downstream role as the default role on the port. For instructions, see Assigning Default Roles to Ports.
- 4. Use the CLI to set the following command: nodealias maxentries 4096 < lag port>.

ExtremeWireless Controller Configuration

The following sections present information regarding support for the ExtremeWireless Controller in the **Policy** tab. Review the following considerations prior to configuring policy on wireless controller devices.

Version Supported

The Policy tab only supports Wireless Controller version 8.01.03 and higher.

Policy Rules

This section describes wireless controller support for policy rules.

Supported Rule Types

The Wireless Controller supports the following traffic classification rule types:

- Ethertype
- MAC Address Source/Destination/Bilateral
- Priority
- IP Type of Service
- IP Protocol Type¹
- ICMP
- IP Address Source/Destination/Bilateral
- IP Socket Source/Destination/Bilateral
- IP UDP Port Source/Destination/Bilateral
- IP UDP Port Source/Destination/Bilateral Range
- IP TCP Port Source/Destination/Bilateral
- IP TCP Port Source/Destination/Bilateral Range

¹Not all IP Protocols are supported for the wireless controller. Supported IP Protocols for this rule type are: ICMP, TCP, UDP, GRE, ESP, AH.

"No Change" Filter Sets

The wireless controller enables administrators to define policies that do not have any filters of their own, but which instead use the set of filters already assigned to a station by a previously applied policy. This type of policy is said to have a "No Change" set of policy rules. The **Policy** tab does not support policies that have "No change" policy rule sets. Using the ExtremeWireless Assistant, you need to remove any policies containing "No Change" rule sets before the wireless controller can be managed by the **Policy** tab.


Rule Actions

The following list defines the wireless controller support for rule actions:

- Access Control: Permit, Deny, and Contain to VLAN actions are supported.
- Class of Service is supported.
- TCI Overwrite is not supported.
- System Log, Audit Trap, Disable Port, and Traffic Mirror actions are not supported.

Rule Directions

The **Policy** tab rules are applied to incoming data packets based on the source or destination address, whereas the wireless controller applies rules to packets based on In/Out direction. On the wireless controller, "In" means coming from the station into the network and "Out" means going from the network out to the station. The wireless controller applies rules to the destination address of inbound packets and to the source address of outbound packets, as shown in the illustration below.



When you create a rule in the Policy tab that permits traffic to a specific destination, that same rule permits data flow from the destination back to the traffic source. This means that Destination rules in the **Policy** tab map to In/Out rules on the wireless controller. Certain **Policy** tab rule types do not have a Source or Destination designation (such as ICMP); however, these rules still map to In/Out rules on the wireless controller to indicate the filters are applied to traffic in both directions. Unchecking the In or Out flag for non-directional rules via the ExtremeWireless Assistant does not affect the way it is reported to the **Policy** tab. As long as the rule still exists, verify succeeds.

All rules enforced from the **Policy** tab are created as "In" rules, and "Out" rules created on the controller are not reported to the **Policy** tab.

When the egress policy feature is enabled for a VNS, egressing traffic is applied to the defined "In" filters as a "reflected" Out rule (with the source and destination fields reversed) and any explicitly defined "Out" filters created on the controller are ignored. Egress policy can be enabled per VNS by selecting Port Properties for that VNS.

The wireless controller reports to the **Policy** tab any rules created directly on the controller that contain an "In" component. "Out" rules are not reported to the Policy tab. This enables administrators to define and use "Out" rules on the wireless controller in special cases where additional restrictions need to be imposed.

Rule Limits

The wireless controller has a limit of 64 rules per policy role if the policy is enforced at the controller (bridged @ wireless controller or routed topology), and 32 rules per policy role if the policy is enforced at the AP (bridged @ AP).

Role Default Actions

The following list defines the wireless controller support for role default actions:

- Access Control: Permit, Deny, and Contain to VLAN are supported.
- Class of Service: Inbound and outbound rate limits are supported. 802.1p Priority, and ToS/DSCP Marking are supported.
- TCI Overwrite is not supported.
- System Log, Audit Trap, Disable Port, and Traffic Mirror actions are not supported.
- The wireless controller will reject policy configurations that specify a VLAN that does not have an egress port already specified.

Class of Service

The following list defines the wireless controller support for Class of Service (CoS) configuration via the **Policy** tab:

- Inbound and outbound rate limits are supported at the role-level as Class of Service default actions.
- User-based inbound/outbound rate limits are supported for the Default port group for wireless controllers only.
- 802.1p Priority configuration is supported.
- ToS/DSCP Marking is supported.
- TCI Overwrite is not supported.
- Transmit Queue Rate Shaping is not supported.

Rate Limits

The wireless controller supports inbound and outbound rate limits at the role-level as Class of Service (CoS) default actions. There are three states supported for a rate limit:

- Rate limit traffic at the specified rate.
- No Change (the CoS does not specify a rate, and the rate limit is "inherited" from the port's default role or from the global default policy, if one is defined.)

To explicitly prevent traffic from being rate limited for a role, you can map a rate limit with a value of 0 to a CoS, and set that as the default CoS for the role.

Internal VLAN

The wireless controller uses an *internal VLAN* for processing traffic. For controllers with firmware version 8.01.xx, the internal VLAN is set by default to use VID 1 and the static name of "DEFAULT VLAN." For controllers with firmware version 8.11.xx and later, the internal VLAN uses the VID 4094 and the static name of "INTERNAL VLAN."


This internal VLAN cannot be used in your **Policy** tab domain configuration to tag traffic. If the VID for the internal VLAN is used in your domain configuration, the **Policy** tab enforce fails with an error message in the Event Log indicating the internal VID cannot be used.

You can use the Web UI (https:\\<controller IP>:5825 > VNS Config > Topologies > Internal VLAN) to change the internal VLAN to a different value, but your policy domain must not use that new value or the **Policy** tab enforce fails.

NOTE: For controllers with firmware version 8.01.xx. Since using a Default VLAN with a VID of 1 is valid on wired devices, the controller's internal VLAN must be changed to another value to prevent issues with the Policy tab enforcing a configuration that uses this VLAN.

Policy Inheritance

The wireless controller uses the concept of policy inheritance, which specifies that if the authenticated policy's access control (VLAN) or class of service (CoS) is set to "No Change," then the policy inheritance hierarchy is used to determine the VLAN and/or CoS. The policy inheritance hierarchy is as follows:



If the authenticated policy's VLAN and CoS are set to "No Change," then the VLAN and CoS settings for the port's default role is used. If the port's default role does not specify the VLAN and CoS, then the global default policy (specified via the ExtremeWireless Assistant) is used. (In wireless controller terminology, a VNS port's default role is the VNS's default policy.)

It is important to note that the **Policy** tab does not support "No Change" rules (filter set). If any policy's rules (filter set) are set to "No Change," then the **Policy** tab is not able to manage the device until the policy containing the "No Change" configuration is removed.

Configuring RADIUS Servers

When configuring RADIUS authentication and accounting servers, keep in mind the following differences:

• The "Number of Retries" and "Timeout Duration" settings for RADIUS authentication servers are configured on a per-server basis for wireless controller devices. For all other devices, these settings are global to all RADIUS servers, and are specified per device as client defaults.

- The "Update Interval" setting for RADIUS accounting servers is configured on a per-server basis for wireless controller devices. For all other devices, this setting is global to all RADIUS servers, and is specified per device as client defaults.
- For wireless controller devices, the Client Status (Enabled or Disabled) is automatically set to Enabled when a RADIUS server exists and Disabled when it does not. For all other devices, Client Status is configured for each device, enabling you to enable and disable communication between the device and the RADIUS servers.
- If Strict Mode is enabled, up to three RADIUS servers are automatically associated to each WLAN service. If Strict Mode is disabled, RADIUS servers must be manually added to a WLAN service via the ExtremeWireless Assistant.

Other Considerations

- The wireless controller does not support authentication configuration.
- The wireless controller does not support viewing user sessions in the Port Usage tabs.
- The wireless controller must have any VLANs used in a Role's default action already defined on the device and configured with an egress port. If the **Policy** tab enforces a domain configuration to the wireless controller using a VLAN that does not have an egress port specified, enforce fails.

ExtremeCloud IQ - Site Engine Policy

ExtremeCloud IQ - Site Engine **Policy** enables the creation and deployment of role-based policies that dynamically control user access, network security, application prioritization and other parameters. Policy management and role-based administration are keys to effectively enforcing business and IT rules in the network infrastructure.

Contact your sales representative for information on obtaining an ExtremeCloud IQ - Site Engine software license.

Policy Tab Overview

The **Policy** tab simplifies the configuration of policies on networks, and deploys the policies on multiple devices throughout the switch fabric.

With the **Policy** tab, you can create policy profiles, called roles, assigned to the ports in your network. These roles provide four key policy features: traffic containment, traffic filtering, traffic security, and traffic prioritization. When authentication is enabled, users identify themselves to the network and are given customized access capabilities based on the role they serve in the organization.

Using the **Policy** tab configuration tools, you can create multiple roles tailored to your specific needs, and set a default role for all or some of your network devices and ports. Basic **Policy** tab operations include creating, editing, and deleting roles. You can also view role configuration on a per device and per port basis. In addition, the **Policy** tab allows you to verify the roles enforced on your network device match the roles currently configured in the application. The **Policy** tab supports a maximum of 1,000 devices (25,000 ports) and 50 roles per policy domain, and can process a maximum of 250 classification rules with a maximum of 50 classification rules per role.

Details View

Some Details View tabs display a simple list of items for the current selection in the left panel. However, other Details View tabs present more complex tables of information. To access Help topics on those tabs, expand the Details View Tabs folder in the Policy tab Help Table of Contents. The Help topics are named to reflect the item selected in the left-panel tree. For example, the Help topic for the Details View tab with a device selected in the left panel is named Details View Tab (Device).

Most Details View tabs provide the following features:

- Right-click Menus Right-click an item for a menu of options.
- **Column and Table Functions** Details View tab tables include several <u>features and functions</u> that enable you to customize the table data.


General

A **General** tab is available in the right panel of the **Policy** tab main window for many items selected in the left-panel tab. It provides general properties information about the selected item.

Help topics for the right-panel **General** tabs are named to reflect the item selected in the left-panel tree. For example, the Help topic for the **General** tab with a device selected in the left panel is named General Tab (Device). For more complete information on the different **General** tabs, expand the General Tabs section and select the desired tab.

Policy Menus

The drop-down menus on the **Policy** tab provide access to Policy tab functions. The **Open/Manage Domains** menu provides options for the domain currently accessed. The **Global Domain Settings** drop-down list enables you to configure global **Policy** tab settings. Use the **Tools** menu to configure authentication settings and review Policy events.



Open/Manage Domains Menu

The Open/Manage Domains provides the following options for the **Policy** tab:

Open Domain

Provides a list of the available Policy Domains. Selecting a domain opens that domain, allowing you to make changes.

Lock Domain

Lets you lock the current Policy Domain for editing purposes. The **Policy** tab automatically locks the domain when you begin to edit the domain configuration. Other **Policy** tab users are notified that the domain is locked and they are not able to save their own domain changes until the lock is released. For more information, see Controlling Client Interactions with Locks.

Save Domain

Lets you save any changes you made to the current Policy Domain. Only users with the capability to Enforce are able to save the domain.

Enforce Domain

Writes the role and/or any changes you have made to it (rules, services) to all the devices in your current domain. See Enforcing for more information.

Verify Domain

Compares the roles in your current domain to the roles currently enforced on all the devices in the current domain. This is useful for ensuring the roles in your domain are enforced, or, if you use more than one domain, ensuring that the roles in the domain you are currently using matches what is on the devices. See Verifying for more information.

Assign Devices to Domain

Opens the Assign Devices to Domain window where you can assign devices that are in the ExtremeCloud IQ - Site Engine database to the current Policy Domain.

Create Domain

Lets you create and name a new (blank) Policy Domain.

Delete Domain(s)

Opens a window where you can select one or more Policy Domains to delete.

Rename Domain

Lets you rename the current Policy Domain.

Import/Export > Import From Domain

Opens the Import from Domain window where you can import policy configuration data from one Policy Domain into another domain. (This menu option is not available if only one domain exists, as there are no other domains from which to import data.)

Import/Export > Import From File

Opens the Import from File window, which enables you to import policy data from a .pmd file into the current Policy Domain. Be aware that the import overwrites any existing data in the Policy Domain. Any devices in the .pmd file must already exist in the Console database or they won't be imported.

Import/Export > Export to File

Lets you save policy data from the current Policy Domain to a .pmd file or .xml file with the file name and location of your choosing. This file stores all information about roles, services, and rules configured in the current Policy Domain. This allows you to save a Domain configuration prior to making changes so that you can restore the original Domain configuration if required (via Import/Export > Import From File).

Global Domain Settings Menu

The Global Domain Settings Menu provides the following options:

GVRP > Ignore GVRP

To ignore GVRP status on the devices in the current domain, select this menu option and enforce. This means that the **Policy** tab ignores the GVRP configuration on a device during an Enforce operation, allowing you to configure some network devices with GVRP enabled and others with GVRP disabled (using MIB Tools or local management), according to their configuration requirements. Be aware that for devices with GVRP set to disabled, ignoring GVRP configuration during an Enforce may affect connectivity on ports with VLANs that rely on Dynamic Egress.

GVRP > Enable GVRP

To enable GVRP on the devices in the current domain, select this menu option and enforce. If the current domain configuration contains rules that use VLAN containment, Dynamic Egress and GVRP must be enabled on the devices in the domain, or the VLANs must be properly pre-configured on the devices outside of the **Policy** tab.

GVRP > Disable GVRP

If you do not want GVRP enabled on the devices in the current domain, select this menu option and enforce. Be aware that disabling GVRP may affect connectivity through ports with VLANs that rely on Dynamic Egress.

Port Level Role Mappings Enabled

Check this box to enable any port-level Tagged Packet VLAN to role mappings or port-level MAC to role mappings that have been configured and enforced for the current domain. If the box is not checked, all port-level mappings are ignored.

Do Not Use Global Services

Check this box to hide the display of Global Services in the left-panel **Services** tab for this domain. If you use Global Services in some domains but not in others, this option allows you to hide global services in the domains where they are not used so that they won't be inadvertently used or modified.

Role ACL Mode

Select to use ACLs in place of traditional rules on Summit devices. Enabling this feature also facilitates user-specified ordering and support for creating ACL entries that support multi-traffic descriptor matching.

NOTE: Summit devices must have firmware V30.5 or later.

Tools Menu

Authentication Configuration

Opens the Authentication Configuration wizard, where you can configure authentication settings on a device.

RADIUS Configuration

Opens the RADIUS Configuration wizard, where you can configure RADIUS authentication and accounting settings on a device.

Policy Event Log

Opens the **Events** tab filtered to display only Policy events.


Policy Enforce Preview

Use the **Enforce Preview** window in the **Policy** tab to view the information you are writing to your devices, before you actually enforce. Use this window when enforcing to devices that only support certain aspects of policy management. For example, some devices support only the policy features of policy management; some devices support the policy features and classification rules, but do not support VLAN forwarding for certain classification rules; and some devices fully support all policy management features, including policy, classification rules, and VLAN forwarding for all classification rules.

The Enforce Preview window appears in the Policy tab by selecting Open/Manage Domain(s) > Enforce Preview, or selecting the enforce icon in the left panel and selecting Enforce Preview. You can control whether this view automatically appears when you select Enforce with the Show on Enforce checkbox.

What you see in the window depends on whether you are enforcing to all devices or to a subset of devices. The title bar indicates the devices to which the enforce applies. After viewing the information in this window, you can either select **Close** to back out and make changes, or **Enforce** to go ahead with the enforce.

You can view device support for specific roles, services, and rules on the **Roles & Rules** tab. Refer to the ExtremeCloud IQ - Site Engine Firmware Support matrix for complete information on device support for Policy features, and VLAN and Priority classification rules.



Show on Enforce

When this checkbox is checked, the **Enforce Preview** window appears any time you enforce, before the actual enforcement takes place.

Left Panel

The left panel of the **Enforce Preview** window displays folders for different device types. Expand the folders to see your network devices and device groups organized according to device type.

The warning icon (**(!**) alerts you that ExtremeCloud IQ - Site Engine is not writing a staged change to this device type (e.g. rules not supported on a device).

Show all device types

Select the checkbox in the left panel to display all device types in the left panel. When the checkbox is not selected, only the devices you are changing by enforcing are displayed.

Select a specific device type to display the information ExtremeCloud IQ - Site Engine is writing to those devices when you enforce in the right panel.

Right Panel

The right panel provides information about whether certain policy management features are supported and/or enabled for the device type selected in the left panel.

- Additional Warnings If there are additional problems detected with the enforce, you will be directed to see the Event Log for details.
- GVRP Shows whether GVRP is Enabled, Disabled, or Ignored. You can change GVRP status for the domain via the Edit menu.
- Dynamic Egress Shows whether Dynamic Egress is Supported or Not Supported.

Device Stats & Info Tab

Displays the devices for the device type selected in the left panel and provides information about each device. If the number of roles in the domain exceeds the supported number of roles on a device, the enforce fails.

• # of Roles Supported - The maximum number of roles supported by the device.

NOTE: OnExtremeXOS/Switch Engine devices, the maximum number of rules supported is the sum of the maximum L2, MAC and IPv4 rule types reported by the device. In ACL Rule mode, the maximum number of rules supported is reported by the device. Each type (L2, MAC and IPv4) is allocated from the same shared pool of slices for ACLs.

• Domain Role Count Supported - This column says "No" if the number of roles in the domain exceeds the supported number of roles on the device. A "Yes" in this column indicates that the number of roles on the device is equal to or less than the maximum number of supported roles.

Role Statistics - Lists information about each role:

- Number of Rules The number of traffic classification rules the role includes.
- Number of Unique Masks The number of masks defined for the rules included in the role.

There are six tabs that provide specific information about the Roles, Classification Rules, VLANs, Classes of Service, and Mappings that will be enforced. The information displayed depends on the device type you've selected in the left panel, and whether you have the Show All or the Show Errors and Warnings Only radio button selected. In addition, select a role in the Roles tab to filter the information for just that role.

Roles Tab

Incomplete - Lists any roles with unsupported classification rules. These roles will be written to the devices, but without the unsupported rules.

Complete - Lists any roles which do *not* include unsupported classification rules. These roles will be written to the devices as defined.

NOTE: Select a Role to display only those classification rules and VLANs associated with the selected role.

Classification Rules Tab

Excluded - Lists any unsupported classification rules that have been applied to a role. These rules will not be included when the associated roles are written to the devices.

Included - Lists any supported classification rules that have been applied to a role. These rules will be included when the associated roles are written to the devices.

NOTE: On N-Series Platinum devices, range classification rules are achieved through applying subnet masks to values. As such, in order to achieve a user-specified range, the device may need multiple rules with subnets applied to encompass that range. So, although the user created only one rule with a range, this list may show multiple instances of that rule with the name of the rule followed by the portion of the over-all range it applies to.

VLAN Tab

Excluded - Lists any VLANs associated with unsupported classification rules, or VLANs that are not supported by the device. These VLANs will not be written to the devices.

Included - Lists any VLANs associated with supported classification rules and VLANs associated with roles. These will be written to the devices.

Classes of Service Tab

Class of Service Mode - Lists the Class of Service mode that will be written to the devices. Classes of Service Subtab - Lists the classes of service that will be written to the devices:

- Class of Service The name of the class of service.
- 802.1p Priority The priority associated with the class of service.
- ToS Value The IP Type of Service value associated with this class of service, if any.
- Drop Prec The drop precedence associated with this class of service, if any.
- TxQueue Index The transmit queue index associated with the class of service.
- IRL Index The role-based inbound rate limit index associated with the class of service.
- ORL Index The role-based outbound rate limit index associated with the class of service.

For more information, see Getting Started with Class of Service and How to Create a Class of Service.

Inbound/Outbound Role-Based Rate Limit Mappings Subtabs - Lists the rate limit mappings that will be written to the devices:

- Device The device where the rate limit mapping will be in effect.
- IRL/ORL Port Grp The name of the port group that contains the rate limit mapping.
- IRL/ORL Index The logical inbound rate limit (IRL) or outbound rate limit (ORL) index number. This index number is specified in a class of service and dictates the rate limiting behavior for incoming packets.
- Rate Limit The actual rate limit that the IRL/ORL index is mapped to.
- IRL/ORL Port Type The type of ports included in the port group. Port type is based on the number of rate limits the ports support (for example, 8-rate limit ports and 32-rate limit ports).
- Information Information about mapping support.

Transmit Queue/Rate Shaper Mappings Subtab - Lists the transmit queue rate shaper mappings that will be written to the devices:

- Device The device where the transmit queue rate shaper mapping will be in effect.
- TxQ Port Grp The name of the port group that contains the transmit queue rate shaper mapping.
- TxQ Index The logical transmit queue rate shaper index number. This index number is specified in a class of service and dictates the transmit queue and rate shaper behavior for incoming packets.
- Physical Transmit Queue / Rate Shaper The actual transmit queue rate shaper that the index is mapped to.
- TxQ Port Type The type of ports included in the port group. Port type is based on the number of transmit queues the ports support (for example, 4-transmit queue ports and 16-transmit queue ports).
- Information Information about mapping support.

Mappings Tab

WARNING: Enforcing port-level MAC to Role mappings could potentially remove rules created as an intrusion detection response.

MAC to Role Mapping - Lists the device-level and port-level mappings that will be written to the devices:

- Device/Port Level indicates whether the mapping is a device-level mapping (all devices) or a port-level mapping (IP address and port description). Port-level mappings on frozen ports will be enforced.
- MAC Address the MAC address mapped to the role. Masking a MAC address is only supported on N-Series Platinum devices.
- Mask the mask associated with the MAC address.
- Role the role mapped to the MAC address.

IP to Role Mapping - Lists the device-level mappings that will be written to the devices:

- IP Address the IP address mapped to the role.
- Mask the mask associated with each IP address. Masking an IP address is only supported on N-Series Gold and Platinum devices.
- Role the role mapped to the IP address.

Tagged Packet VLAN to Role Mapping - Lists the device-level and port-level mappings that will be written to the devices:

- Device/Port Level indicates whether the mapping is a device-level mapping (all devices) or a
 port-level mapping (IP address and port description). Port-level mappings on frozen ports will be
 enforced.
- VLAN the VLAN mapped to the role.
- Role the role mapped to the VLAN.

Authentication Based VLAN (RFC 3580) to Role Mapping - Lists the mappings that will be written to the devices:

- VLAN the VLAN mapped to the role.
- Role the role mapped to the VLAN.

Event Log Button

Opens the **Events** tab filtered to display events with an **Event Type** of **Policy**.


Enforce Button

Enforces the roles, classification rules and VLANs in the current data file to the devices, based on the level of support available on the devices as indicated in the **Enforce Preview** window.



Import from Domain

This window lets you import policy configuration data from one Policy Domain into another domain. To access the Import from Domain window, select **Open/Manage Domain > Import/Export > Import From Domain**. (This menu option is not available if only one domain exists, as there are no other domains from which to import data.)



Domain

Use the drop-down list to select the domain whose data you want to import.

Data Elements to Import

In this section, you can choose the specific data elements you want to import. Select **Select All** to select all the data import options.

Roles

Select this option to import roles, including the role's name, description, default VLAN (access control), and default class of service. If a role's services already exist in the current domain, or if you are importing them at the same time as the role, the services are associated with the role. Otherwise, the services are not imported.

Services & Rules (Local)

Select this option to import Local services (services that are unique to a specific domain) and their associated classification rules. When you import rules from another domain, the Policy tab checks for rule conflicts (see Conflict Checking for more information).

Service Groups

Select this option to import service group names. If a service group's services already exist in the current domain, or if you are importing them at the same time as the service group, the services will be associated with the group. Otherwise, the services will not be imported.

Devices

Select this option to import devices. Any devices in the .pmd file must already exist in the ExtremeCloud IQ - Site Engine database or they won't be imported. (See How to Add and Delete Devices for more information on using Console to add devices to the ExtremeCloud IQ - Site Engine database.) Devices that are imported are automatically assigned to the current domain and are displayed in the Policy tab Network Elements tree. If the devices being imported were already assigned to another domain, then those devices are reassigned to the current domain. Any devices that are not imported are listed in an Event Log message along with their device type and firmware version.

Port Groups (User-Defined)

Select this option to import user-defined port groups. If you are importing a port group's ports at the same time as the port group, the ports will be associated with the port group. Otherwise, the ports are not imported.

Class of Service

Select this option to import classes of service, role-based rate limit port groups, and transmit queue port groups. For the purposes of importing, a class of service is defined as the class of service name, i.e., priority is not a factor in determining uniqueness. After a class of service is imported, its associated roles, services, and rules are updated. When you import class of service data, the relationship between a class of service and its priority is retained; however, rate limiting characteristics of the priorities are not imported. If you also elect to import rate limits, the rate limits are imported first, then the classes of service are imported. You can then redefine the class of service priorities with some or all of the imported rate limits, if desired. Although ToS characteristics are not used to determine the uniqueness of a class of service for importing, if ToS is a part of a class of service, it is imported as an attribute of the class of service. See append, update and overwrite for information on how those specific actions affect the import of classes of service.

Adv CoS Config

Select this option to import the class of service configuration (basic or advanced) for the domain (whether the Advanced Class of Service Configuration option is selected).

Rate Limits

Select this option to import rate limits. For the purposes of importing, a rate limit is defined as [rate + direction] when determining uniqueness. Any other duplicates on the list are not changed. Because rate limits cannot include conflicting priority values, if a priority is already being utilized by an existing rate limit, it will not be imported. If you also elect to import classes of service, the rate limits are imported first, then the classes of service are imported. See append and update for information on how those specific actions affect the import of rate limits.

NOTE: ZTP+ functionality requires an ExtremeXOS/Switch Engine device on which version 21.1 is installed.

NOTE: Only those network elements that are recognized by the existing domain can be imported as exclusions. Others are ignored.

VLANs

Select this option to import VLANs.

Policy VLAN Islands

If applicable, Policy VLAN Islands and Island VLANs are imported via the Devices and VLANs options.

- If the Devices option is selected and the Policy VLAN Islands feature is enabled in the current domain as well as the imported domain, the Policy VLAN Islands will be imported. The Policy VLAN Island Base ID and Offset settings from the imported data will be used and those in the current domain will be lost.
- If the VLANs option is selected and the Policy VLAN Islands feature is enabled in the current domain as well as the imported domain, the Island VLANs are imported and are added to any existing Policy VLAN Islands.

Whenever Policy VLAN Islands are imported, all the island VLANs are recalculated and the island ranges may change. It is possible to import more islands and VLANs than can be configured. If this is the case, an error appears in the Event Log, asking that the Base ID and Offset settings be changed.

Network Resources

Select this option to import network resource groups. After a Network Resource is imported, the associated services are updated. If a network resource group no longer exists after an import, the service with which it was associated is changed to a manual service on the Automated Service tab for the service.

Port-Level Role Mapping Status

Select this option to import the Port-Level Role Mappings Enabled status for the domain, as specified in the Edit menu.

GVRP Status

Select this option to import the GVRP status for the domain (as specified in the Edit menu).

Do Not Use Global Services Status

Select this option to import the Do Not Use Global Services status for the domain, as specified in the Edit menu.

Domain Mode

Select this option to import the domain mode (active or passive) as specified in the Edit menu.

Application of Imported Data Elements

In this section, you can choose how you want the data elements selected above to update your current domain.

Append domain data to existing elements

Select this option to import only new data elements into your current domain. If any of the selected data elements already exist in your current domain, they will not be changed.

Rate Limits: A rate limit will not be appended if: 1) The Rate, Direction, and 802.1P Priority are already defined. 2) The Priority list is empty.

CoS: A class of service will not be appended if: 1) The name is the same as an existing class of service. 2) The class of service names are different but the rate limits for the imported class of service do not match the existing rate limit settings.

Update existing data with elements from domain

Select this option to 1) replace the selected data elements that exist in your current domain with the imported data elements, and 2) import the selected data elements that don't exist in your current

domain.

Rate Limits: A rate limit will not be updated if the rate limit and direction do not match.

CoS: A class of service will not be updated if: 1) The name does not match an existing class of service. 2) The class of service name matches but the rate limits for the imported class of service do not match the existing rate limit settings.

Overwrite existing elements

Select this option to replace the selected data elements that exist in your current domain with the imported data elements.

CoS: A class of service will not be overwritten if the rate limits for the imported class of service do not match the existing rate limit settings.

NOTE: If you decide that you want to return to the previous configuration (that the import updated), you can perform a File > Read Policy Domain operation to restore the configuration, as long as you have not saved the data you imported.

Select All Button

Selects all of the data elements.


Import Button

Imports the selected data and closes the window.



Import from File

This window lets you import policy data from a .pmd file into a Policy Domain. To access the window, select Open/Manage Domains > Import/Export > Import From File.



Policy Manager Data (PMD) File

Enter the name and path for the data file (.pmd) you want to import, or navigate to the file by selecting the **Select File** button.

Data Elements to Import

In this section, you can choose the specific data elements you want to import. Select **Select All** to select all the data import options.

Roles

Select this option to import roles, including the role's name, description, default VLAN (access control), and default class of service. If a role's services already exist in the current domain, or if you are importing them at the same time as the role, the services will be associated with the role. Otherwise, the services are not imported.

Services & Rules (Local)

Select this option to import Local services (services that are unique to a specific domain) and their associated classification rules. When you import rules from another domain, the **Policy** tab checks for rule conflicts (see Conflict Checking for more information).

Service Groups

Select this option to import service group names. If a service group's services already exist in the current domain, or if you are importing them at the same time as the service group, the services are associated with the group. Otherwise, the services are not imported.

Devices

Select this option to import devices. Any devices in the .pmd file must already exist in the ExtremeCloud IQ - Site Engine database or they won't be imported. (See How to Add and Delete Devices for more information on using Console to add devices to the ExtremeCloud IQ - Site Engine database.) Devices that are imported are automatically assigned to the current domain and are displayed in the Policy tab Network Elements tree. If the devices being imported were already assigned to another domain, then those devices are reassigned to the current domain. Any devices that are not imported are listed in an Event Log message along with their device type and firmware version.

Port Groups (User-Defined)

Select this option to import user-defined port groups. If you are importing a port group's ports at the same time as the port group, the ports are associated with the port group. Otherwise, the ports are not imported.

Class of Service

Select this option to import classes of service, role-based rate limit port groups, and transmit queue port groups. For the purposes of importing, a class of service is defined as the class of service name, i.e., priority is not a factor in determining uniqueness. After a class of service is imported, its associated roles, services, and rules are updated. When you import class of service data, the relationship between a class of service and its priority is retained; however, rate limiting characteristics of the priorities are not imported. If you also elect to import rate limits, the rate limits are imported first, then the classes of service are imported. You can then redefine the class of service priorities with some or all of the imported rate limits, if desired. Although ToS characteristics are not used to determine the uniqueness of a class of service for importing, if ToS is a part of a class of service, it is imported as an attribute of the class of service. See append, update and overwrite for information on how those specific actions affect the import of classes of service.

Adv CoS Config

Select this option to import the class of service configuration (basic or advanced) for the domain (whether the Advanced Class of Service Configuration option is selected).

Rate Limits

Select this option to import rate limits. For the purposes of importing, a rate limit is defined as [rate + direction] when determining uniqueness. Any other duplicates on the list are not changed. Because rate limits cannot include conflicting priority values, if a priority is already being utilized by an existing rate limit, it will not be imported. If you also elect to import classes of service, the rate limits are imported first, then the classes of service are imported. See append and update for information on how those specific actions affect the import of rate limits.

Note: Only those network elements that are recognized by the existing domain can be imported as exclusions. Others will be ignored.

VLANs

Select this option to import VLANs.

Policy VLAN Islands

If applicable, Policy VLAN Islands and Island VLANs are imported via the Devices and VLANs options.

- If the Devices option is selected and the Policy VLAN Islands feature is enabled in the current domain as well as the imported domain, the Policy VLAN Islands will be imported. The Policy VLAN Island Base ID and Offset settings from the imported data will be used and those in the current domain will be lost.
- If the VLANs option is selected and the Policy VLAN Islands feature is enabled in the current domain as well as the imported domain, the Island VLANs are imported and are added to any existing Policy VLAN Islands.

Whenever Policy VLAN Islands are imported, all the island VLANs are recalculated and the island ranges may change. It is possible to import more islands and VLANs than can be configured. If this is the case, an error appears in the Event Log, asking that the Base ID and Offset settings be changed.

Network Resources

Select this option to import network resource groups. After a Network Resource is imported, the associated services are updated. If a network resource group no longer exists after an import, the service with which it was associated is changed to a manual service on the Automated Service tab for the service.

Port-Level Role Mapping Status

Select this option to import the Port-Level Role Mappings Enabled status for the domain.

GVRP Status

Select this option to import the GVRP status for the domain.

Do Not Use Global Services Status

Select this option to import the Do Not Use Global Services status for the domain.

Domain Mode

Select this option to import the domain mode (active or passive) as specified in the Edit menu.

Global Domain Data

Use this option only if you want to append, update, or overwrite the globally defined services and rules in your current domain with the global domain data stored in the .pmd file you are importing. This option will modify or remove any existing global data and will affect all domains. If overwrite is selected, all current global data will be removed and replaced with the global configuration in the file, or nothing if there is no configuration defined.

Global Services & Rules

Select this option to import Global services (services that are common to all domains) and their associated classification rules. When you import rules from another domain, the Policy tab checks for rule conflicts (see Conflict Checking for more information).

Application of Imported Data Elements

In this section, you can choose how you want the data elements selected above to update your current domain.

Append domain data to existing elements

Select this option to import only new data elements into your current domain. If any of the selected data elements already exist in your current domain, they will not be changed.

Rate Limits: A rate limit will not be appended if: 1) The Rate, Direction, and 802.1P Priority are already defined. 2) The Priority list is empty.

Cos: A class of service will not be appended if: 1) The name is the same as an existing class of service. 2) The class of service names are different but the rate limits for the imported class of service do not match the existing rate limit settings.

Update existing data with elements from domain

Select this option to 1) replace the selected data elements that exist in your current domain with the imported data elements, and 2) import the selected data elements that don't exist in your current domain.

Rate Limits: A rate limit will not be updated if the rate limit and direction do not match.

Cos: A class of service will not be updated if: 1) The name does not match an existing class of service. 2) The class of service name matches but the rate limits for the imported class of service do not match the existing rate limit settings.

Overwrite existing elements

Select this option to replace the selected data elements that exist in your current domain with the imported data elements.

CoS: A class of service will not be overwritten if the rate limits for the imported class of service do not match the existing rate limit settings.

NOTE: If you decide that you want to return to the previous configuration (that the import updated), you can perform a File > Read Policy Domain operation to restore the configuration, as long as you have not saved the data you imported.

Select All Button

Selects all of the data elements.

Import Button

Imports the selected data and closes the window.




Assign Devices to Domain

This window lets you assign devices in the ExtremeCloud IQ - Site Engine database to a Policy Domain or move devices from one domain to another. A Policy Domain contains any number of roles and a set of devices uniquely assigned to that particular domain. A device can exist in only one Policy Domain. For more information on domains, see How to Create and Use Domains.

Initially, you must add your devices to the ExtremeCloud IQ - Site Engine database. When your devices are in the database, use this window to assign the devices to a Policy Domain. As soon as the devices are assigned to a domain, they display automatically in the **Policy** tab **Devices** tab. Only devices that support policy are displayed in the **Devices** tab.

To access this window, open the domain to which you want to assign devices, and select Open/Manage Domains > Assign Devices to Domain.



Devices

The Devices list displays all the unassigned devices in the database (including devices that do not support policy) but are not assigned to a domain. The panel also displays any other domains and the devices assigned to that domain. Use the navigation trees to select a single domain or All Other Domains.

Current Domain

The Current Domain list displays the current domain and the devices assigned to that domain. To add a device to the current domain, select the device in the left panel and select the right arrow. You can also select and add multiple devices. To remove a device from the current domain, select the device and select the left arrow. This removes the device from the current domain and places it back in the device

tree as either unassigned or as a member of the domain it came from. To remove all devices, select the double left arrow.

Device Domain Membership

This section is only displayed when more than one domain exists. It lists the domain assignment for whatever device or device group you have selected in the Devices panel. This is particularly useful when you have selected All Other Domains from the drop-down list in the Devices panel, as it allows you to quickly see the domain assignment for each device.

Right Arrow Button

Adds the devices selected in the Devices list to the Current Domain list.

Remove Button

Removes the devices selected in the Current Domain list from the current domain and places it back in the Devices list as either unassigned or as a member of the domain from which it came.

NOTE: Removing a device from a domain does not delete the device from the ExtremeCloud IQ - Site Engine database. To delete a device from the database, right-click on the device in the **Network** tab, and select **Device > Delete Device** from the menu. When a device is deleted from the database, it is automatically removed from the **Network** and **Policy** tabs.

Double Left Arrow Button

Removes all the devices from the current domain.


OK Button

Assigns the selected devices to the current domain and displays the devices in the **Policy** tab's **Devices** tab. Only devices that support policy are assigned to the domain and displayed in the **Devices** tab.

Authentication Configuration

The **Authentication Configuration** wizard enables you to configure and change the authentication settings on your devices. Authentication must be configured and enabled on a device in order for individual port authentication settings to take effect (see How to Configure Ports).

To access this tab, select Authentication Configuration from the Tools drop-down list.



Device Selection

Use the **Device Selection** tab to select the devices on which you are configuring authentication settings.

Select a device from the available devices list in the left of the tab and select the right arrow icon to move the device to the selected devices list. Select **Next>** to proceed to the next tab.

Port Selection

Use the **Port Selection** tab to select the ports on which you are configuring authentication settings.

Select a port from the Available Ports list at the top of the tab and select **Add Ports** to move the port to the Selected Devices list. Select **Next>** to proceed to the next tab.

Device Configuration

The **Device Configuration** tab allows you to configure authentication for a device. Use the **Port Configuration** tab to configure authentication settings for individual ports on the device. You can also use the drop-down list at the top of the tab to load device and port configuration settings from a template or import a template from the ExtremeCloud IQ - Site Engine server into ExtremeCloud IQ - Site Engine.

Import Template

Select to open a window from which you can select a device and port configuration template saved on the ExtremeCloud IQ - Site Engine server.

Rename/Delete Template

Select rename or delete a device and port configuration template saved on the ExtremeCloud IQ - Site Engine server.

Save Device & Port Config Settings To Template


Select to save the settings you define on the **Device Configuration** and **Port Configuration** tabs to a template you can load for other devices.

Load Device & Port Config Settings From Template

Select to load a previously saved template of settings you previously defined on the **Device Configuration** and **Port Configuration** tabs.

Authentication Status

Use this section to select the authentication mode and types used on the device.



Use the fields on the left side of this section to select the appropriate single- or multi-user authentication types. Only options supported by the selected device are available for selection. Some devices support multiple authentication types and multiple users (Multi-User Authentication) per port, while others are restricted to only one or two authentication types and single users per port. Refer to the Firmware Support matrix for information on the authentication types supported by each device type.

WARNING: Switching Authentication Types, or changing the Authentication Status from Enabled to Disabled, logs off any currently authenticated users.

Auth Type Precedence (High->Low)

This displays the order in which the authentication types are attempted on the device, with the authentication type on the left having the highest precedence (attempted first). You can edit the precedence order by selecting the field. In the Edit Precedence window, select the authentication type you want to position, and use the **Up** and **Down** buttons to arrange the types in the desired order of precedence.

WARNING: Leave the default precedence, if possible. Changing the Quarantine precedence to be lower than any other type or changing the Auto Track precedence to be higher than any other type may cause problems.

Re-Auth Timeout Action

This setting defines the action for sessions that need to be re-authenticated if the RADIUS server re-authentication request times out. Select the **Terminate** option to terminate the session or the **None** option to allow the current session to continue without disruption.

Maximum Number of Users

This setting applies to devices with Multi-User as their configured authentication type. The maximum number of users that can be actively authenticated or have authentications in progress at one time on this device. You can specify the maximum number of users per port on the port's Port Properties Authentication Configuration tab.

RFC3580 VLAN Authorization


This allows you to enable and disable RFC 3580 VLAN Authorization for the selected device. RFC 3580 VLAN Authorization must be enabled on devices in networks where the RADIUS server is configured to return a VLAN ID when a user authenticates.

When RFC 3580 VLAN Authorization is enabled:

- devices that do **not** support policy tag packets with the VLAN ID.
- devices that support both policy and Authentication-Based VLAN to Role Mapping classify packets according to the role to which the VLAN ID maps.

Global Authentication Settings

This section lets you set session timeout and session idle timeout values for each authentication type.



Session Timeout

This setting represents the maximum number of seconds an authenticated session may last before automatic termination of the session. A value of zero indicates that no session timeout applies. This value may be superseded by a session timeout value provided by the authenticating server. For example, if a session is authenticated by a RADIUS server, that server may send a session timeout value in its authentication response.


NOTE: Non-zero values are rounded to the nearest non-zero multiple of 10 by the device.

Session Idle Timeout

This displays the maximum number of consecutive seconds an authenticated session may be idle before ExtremeCloud IQ - Site Engine automatically terminates the session. A value of zero indicates that no idle timeout applies. This value may be superseded by an idle timeout value provided by the authenticating server. For example, if a session is authenticated by a RADIUS server, that server may send an idle timeout value in its authentication response.

MAC Authentication Settings

This section enables you to set up the MAC password for MAC authentication. In order for MAC authentication to work, you must also configure the RADIUS server with the MAC password as well as the MAC addresses which are allowed to authenticate.



Set Password/Mask

Select this checkbox to set a password and mask for MAC authentication.

MAC User Password

The password passed to the RADIUS server for MAC authentication.

MAC Mask

You can select a mask to provide a way to authenticate end-systems based on a portion of their MAC address. For example, you could specify a mask that would base authentication on the manufacturers ID portion of the MAC address. The MAC Mask is passed to the RADIUS server for authentication after the primary attempt to authenticate using the full MAC address fails.

MAC Address Delimiter

The character used between octets in a MAC address:

- Hyphen A hyphen is used as a delimiter in the MAC address (e.g. xx-xx-xx-xx-xx).


Web Authentication Settings

For users of web-based authentication, this tab lets you specify web authentication parameters using three sections:

- General
- Guest Networking
- Web Login

General

The General section lets you specify the URL of the authentication web page and the IP address of the system where it resides. It also lets you enable certain web authentication features, such as Enhanced Login Mode, on devices that support those features.



Enhanced Login Mode

Enabling the Enhanced Login Mode causes the authentication web page to be displayed regardless of whether the URL or IP address entered into the browser by the end user is the designated Web Authentication URL or IP address. This option is grayed out if the device does not support the mode.

Enhanced Mode Redirect Time(s)

This setting applies for devices with <u>Enhanced Login Mode</u> enabled. It specifies the amount of time (in seconds) before the end-user is redirected from the authentication web page to their requested URL.

An end-system using DHCP requires time to transition from the temporary IP address issued by the authentication process to the official IP address issued by the network. **Enhanced Mode Redirect Time**

specifies the amount of time allowed for the end-system to complete this process and begin using its official IP address.

For example, if an end-user (in **Enhanced Login Mode** and a **Redirect Time** of **30 seconds**) enters the URL of "http://ExtremeNetworks.com", the user is presented the authentication web page. When the user successfully authenticates into the network, the user sees a login success page that displays "Welcome to the Network. Completing network connections. You will be redirected to http://ExtremeNetworks.com in approximately 30 seconds."

WINS/DNS Spoofing

This setting allows you to enable and disable WINS/DNS spoofing for the selected device. Spoofing allows the end-user to resolve the Web Authentication URL name to the IP address using WINS/DNS. The default is Disabled. This option is grayed out if not supported by the device.

Logo Display Status

Specifies whether the Extreme Networks logo is displayed or hidden on the authentication web page window. This option is grayed out if not supported by the device.

Authentication Protocol

This setting is the authentication protocol being used (PAP or CHAP). PAP (Password Authentication Protocol) provides an automated way for a PPP (Point-to Point Protocol) server to request the identity of user, and confirm it via a password. CHAP (Challenge Handshake Authentication Protocol), the more secure of the two protocols, provides a similar function, except that the confirmation is accomplished using a challenge and response authentication dialog.

Web Authentication URL

This is the URL for your authentication web page. Users wishing to receive network services access the web page from a browser using this URL. The http:// is supplied. Alphabetical characters, numerical characters and dashes are allowed as part of the URL, but dots are not. The URL needs to be mapped to the Web Authentication IP address in DNS or in the hosts file of each client. It must be resolvable via DNS/WINS, either on the device or at corporate, assuming the Web Authentication mapping has been set up on the corporate DNS/WINS service. This option is grayed out if not supported by the device.


Web Authentication IP Address

This is the IP address of your authentication web page server. If you have specified a Web Authentication URL, the IP address needs to be mapped to the URL in DNS or in the host file of each client.

Guest Networking

The **Guest Networking** section lets you configure guest networking, a feature that allows any user to access the network and obtain a guest policy without having to know a username or password. The user accesses the authentication web page, where the username and password fields are automatically filled in, allowing them to log access as a guest. If the user does not want to log in as a guest, they can type in their valid username and password to log in.

NOTE: Guest networking is designed for networks using web-based authentication, with <u>port mode</u> set to Active/Discard.



Guest Networking Status

Use the drop-down list to specify guest networking status:

- **Disable** Guest networking is unavailable.
- Local Auth Guest Networking is enabled. The user accesses the authentication web page
 where the username field is automatically filled in with the specified <u>Guest Name</u>. When the user
 submits the web page using this guest name, the default policy of that port becomes the active
 policy. The port mode must be set to Active/Discard mode.
- RADIUS Auth Guest Networking is enabled. The user accesses the authentication web page, where the username field is automatically filled in with the specified <u>Guest Name</u>, and the password field is masked out with asterisks. When the user submits the web page using these credentials, the value of the <u>Guest Password</u> is used for authentication. Following successful authentication from the RADIUS server, the port applies the policy (role) returned from the RADIUS server. The port mode must be set to Active/Discard mode.

Guest Name


The username that Guest Networking uses to authenticate users. The guest name is displayed automatically on the authentication web page. If the user does not want to log in as a guest, they can type in their valid username to override the guest username.

Guest Password

The password that Guest Networking uses to authenticate users when RADIUS Auth is selected.

Web Page Banner

The Web Page Banner section allows you to customize the banner end users see at the top of the authentication web page and set a Redirect Time, if applicable.



Web Page Banner

Use this area to create a banner end users see at the top of the authentication web page. For example, you might include your company name and information on what to do if the user has questions or problems. Because this banner also appears in messages that occur during successful login and failed authentication, as well as on the "Radius Busy" screen, it is not appropriate to include "Welcome to [Your Company]" in the banner.

The **Default** button allows you to reset the banner to default text provided in a text file (pwa_banner.txt). Initially, the default banner text is the Extreme Networks contact information. However, you can customize the text for your network by editing the pwa_banner.txt file, located in the top level of the Policy Manager install directory. Then, when you select the Default button, the new text will be displayed in the Web Page Banner area.

Convergence End-Point Settings


This section provides a way to identify Convergence End-Points (IP phones) connecting to the device, and apply a role to the end-point based on the type of end-point detected. The CEP Detection section lets you create detection rules for identifying the end-points, and the CEP Role Mappings section lets you map a role to each CEP product type.

In addition to configuring CEP on the device, you must also enable CEP protocols on each port using the CEP Access section in the Port Authentication Tab. After you have configured CEP on the device and each port, you can monitor CEP usage on the Port Usage Tab (Port) or Port Usage Tab (Device).

CEP Role Mappings

This section lets you select the CEP product types supported on the device, and map a role for each type. Then, when a convergence end-point (such as an IP phone) connects to the network,

the device identifies the type of end-point (using CEP detection rules) and applies the assigned role.



CEP Type

Lists the CEP types supported by the device.

Role

Lists the role mapped to each CEP Type.

Add

Select a CEP Type and select the **Add** button to open the Add Role Mapping window, where you can select a role for the selected **CEP Type**. Your selections are added to the CEP Role Mappings list.

Remove

Select the CEP Type and select Remove to remove the CEP Type in the CEP Role Mappings list.

CEP Detection Tab


Use this section to create CEP detection rules used to determine if a connecting end-system is a CEP device and the type of CEP device. This allows ExtremeCloud IQ - Site Engine to assign the appropriate role to the port based on the type of CEP device detected.

NOTE: CEP detection rules apply only to Siemens, H.323, and SIP (Session Initiation Protocol) phone detection. Cisco detection uses CiscoDP as its detection method.

CEP detection rules are based on two detection methods:

- TCP/UDP Port Number detection Many CEP vendors use specific TCP/UDP port numbers for call setup on their IP phones. You can create detection rules that identify CEP devices based on specific TCP/UDP port numbers. By default, Siemens Hi-Path phones are detected on TCP/UDP port 4060.
- IP Address detection H.323 phones use a reserved IP multicast address and UDP port number for call setup. You can create detection rules to detect an IP phone based on its IP address in combination with an IP address mask. By default, H.323 phones are detected using the multicast address 224.0.1.41 and the TCP/UDP ports 1718, 1719, and 1720. SIP phones are detected using the multicast address 224.0.1.75

and the TCP/UDP port 5060. H.323 and SIP phones are also detected using only their respective multicast addresses without the TCP/UDP ports.



Priority

The rule priority with one (1) being the highest priority. The rule with the highest priority is used first, so it is recommended the highest priority be given to the predominate protocol in the network to provide for greater efficiency.

Address

If the rule is based on IP address detection, this field displays the IP address that incoming packets matched against. By default, H.323 uses 224.0.1.41 as its IP address, SIP uses 224.0.1.75 as its IP address, and Siemens has no IP address configured.

Address Mask

If the rule is based on IP address detection, this field displays the IP address mask against which incoming packets are matched.

End Point Type

Specifies the end-point type assigned (H.323, Siemens, or SIP) if incoming packets match this rule.

Protocol

If the rule is based on TCP/UDP port detection, this field displays the protocol type used for matching, using a port range defined with the Port Low and Port High values:

- UDP + TCP Match the port number for both UDP and TCP frames.
- TCP Match the port number only for TCP frames.
- UDP Match the port number only for UDP frames.

Port Low

The low end of the port range defined for detection on UDP and/or TCP ports.

Port High

The high end of the port range defined for detection on UDP and/or TCP ports.

Add

Opens the Add/Edit CEP Detection Rule window where you can create CEP detection rules.

Remove

To remove a CEP detection rule, select the entry and select Remove.

Edit

To edit a CEP detection rule, select the rule and select **Edit**. The Add/Edit CEP Detection Rule window opens where you edit the rule's parameters. You can also double-click an entry in the table to open the edit window.

Port Configuration


The Port Configuration tab allows you to configure authentication for the ports of a device.

The **Authentication Configuration** tab has six sections:

- Authentication Mode
- RFC3580 VLAN Authorization
- Login Settings
- Automatic Re-Authentication
- Authenticated User Counts
- CEP Access

Authentication Mode

This section displays general authentication and port mode information about the port.



Port Mode

Port mode defines whether or not a user is required to authenticate on a port, and how unauthenticated traffic will be handled. It is a combination of Authentication Behavior (whether or not authentication is enabled on the port), and Unauthenticated Behavior (whether unauthenticated traffic will be assigned to the port's default role or discarded).

- Authentication Behavior -- Defines whether or not end users are required to authenticate on the port (device).
 - Active -- Normal authentication procedures are implemented. End users are required to authenticate.
 - Inactive -- Authentication of end users is not required.
- Unauthenticated Behavior -- Defines how the traffic of unauthenticated end users will be handled on the port.
 - **Default Role** -- If the end user is unauthenticated, the port will implement its default role. If there is no default role, there will be no role on the port.
 - **Discard** -- If the end user is unauthenticated, no traffic is allowed on the port.

These two settings can be combined to create four possible port modes.

- Inactive/Discard Mode: In this mode, authentication is inactive for the port. All traffic from users connected to the port is discarded. This effectively turns the port off. This port mode is not available for Single User MAC Authentication.
- Inactive/Default Role Mode: In this mode, authentication is inactive for the port. All users connecting to this port will use the default role, if one has been assigned to the port, in combination with any existing static classifications. If there is no default role assigned to the port, the port uses only the static classification rules which exist. If there are no static rules, the port uses the PVID and default class of service for the port. This is the default port mode for ports.
- Active/Discard Mode: In this mode, authentication is active for the port and end users are required to
 authenticate. All traffic from unauthenticated users connected to the port is discarded. The
 Unauthenticated Behavior varies depending on the type of authentication configured on the device.

Single User Web-based Authentication: If authentication is successful, the port is assigned the end user's role as its current role. If unsuccessful, all traffic is discarded. A default role has no meaning on this Active/Discard port, since all unauthenticated traffic is discarded.

Single User 802.1X and 802.1X+MAC Authentication: If authentication is successful, the port is assigned the end user's role as its current role. If unsuccessful, all traffic is discarded. This mode requires that there be **no** default role assigned to the port.

Single User MAC Authentication: This port mode is not available for Single User MAC Authentication.

Multi-User 802.1X and MAC Authentication: If authentication is successful, the port is assigned the end user's role as its current role. If unsuccessful, all traffic is discarded. A default role has no meaning on this Active/Discard port, since all unauthenticated traffic is discarded.

Multi-User Web-based Authentication: This port mode is not available for Multi-User Web-based Authentication.

Advantages of Active/Discard mode: This mode is highly secure, since the end user receives no network

services at all until authentication is successful.

Disadvantages of Active/Discard mode: The unauthenticated end user is unable to connect to any network services, such as the Domain Controller (if using a Microsoft operating system), DHCP services, DNS services, or the Web proxy. In single user web-based authentication, the device spoofs WINS/DNS services (if the functionality is enabled) in order to allow the user to communicate with it for authentication.

Active/Default Role Mode - In this mode, authentication is active for the port and end users are required
to authenticate. If authentication is successful, the port is assigned the end user's role as its current role.
All unauthenticated users connected to the port will use the default role, if one has been assigned to the
port, in combination with any existing static classifications. If there is no default role assigned to the
port, the port uses only the static classification rules which exist. If there are no static rules, the port uses
the PVID and default class of service for the port. For Single User 802.1X and 802.1X+MAC
Authentication, this mode requires that a default role be assigned to the port.

Advantages of Active/Default Role mode: In this mode, a default role is applied to the port to allow unauthenticated end users access to basic services such as the DHCP Server, Domain Services, WINS, and the Web proxy. When the end user is authenticated, that user's role is applied to the port, providing a customized set of services allowed by his or her role. Active/Default Role mode is an alternative to Active/Discard mode, which is limiting in that there are no network services available at all until the end user is authenticated.

Disadvantages of Active/Default Role mode: This mode is less secure than Active/Discard, in that the user receives some network access prior to authentication.

RFC3580 VLAN Authorization Tab

This tab lets you enable or disable RFC 3580 VLAN Authorization on the port and specify an egress state. RFC 3580 VLAN Authorization must be enabled in networks where the RADIUS server has been configured to return a VLAN ID when a user authenticates.

When RFC 3580 VLAN Authorization is enabled:

- ports on devices that do not support policy tag packets with the VLAN ID.
- ports on devices that do support policy and also support Authentication-Based VLAN to Role Mapping classify packets according to the role to which the VLAN ID maps.

You can also enable and disable VLAN Authorization at the device level using the device **Authentication** tab. If the device does not support RFC 3580, this tab is grayed out.



VLAN Authorization Status

Allows you to enable and disable RFC 3580 VLAN Authorization for the selected port. This option is grayed out if not supported by the device.

VLAN Authorization Admin Egress

Allows you to modify the VLAN egress list for the VLAN ID returned by the RADIUS server when a user authenticates on the port:

- None No modification to the VLAN egress list will be made.
- Tagged The port will be added to the list with the egress state set to Tagged (frames will be forwarded as tagged).
- Untagged The port will be added to the list with the egress state set to Untagged (frames will be forwarded as untagged).
- Dynamic The port will use information returned in the RADIUS response to modify the VLAN egress list. This value is supported only if the device supports a mechanism through which the egress state may be returned in the RADIUS response.


The current egress settings for the port are displayed in the VLAN Oper Egress column in the **User Sessions** tab. These options are grayed out if not supported by the device.

Apply Button

Saves any change you made to the VLAN Authorization settings.

Login Settings

This tab displays the current login settings for the port and allows you to change the settings if desired. The options available depend on what type(s) of authentication are enabled on the device.



Number of Attempts Before Timeout

Number of times a user can attempt to log in before authentication fails and login attempts are not allowed. For web-based authentication, valid values are 1-2147483647, zero is not allowed, and the default is 2. For 802.1X and MAC authentication, this value is permanently set to 1.

Hold Time (seconds)

Amount of time (in seconds) authentication will remain timed out after the specified Number of Attempts Before Timeout has been reached. Valid values are 0-65535. The default is 60. (Hold Time is also known as Quiet Period in web-based and MAC authentication.)

Authentication Request Period

For 802.1X authentication, how often (in seconds) the device queries the port to see if there is a new user on it. If a user is found, the device then attempts to authenticate the user. Valid values are 1-65535. The default is 30.

User Timeout

For 802.1X authentication, the amount of time (in seconds) the device waits for an answer when querying the port for the existence of a user. Valid values are 1-300. The default is 30.

Authentication Server Timeout

For 802.1X authentication, if a user is found on the port, the amount of time (in seconds) the device waits for a response from the authentication server before timing out. Valid values are 1-300. The default is 30.

Port Handshake Requests Before Failure

For 802.1X authentication, the number of times the device tries to finalize the authentication process

with the user before the authentication request is considered invalid and authentication fails. Valid values are 1-10. The default is 2.

Quarantine Session Timeout (sec)

For Quarantine authentication, the maximum number of seconds an authenticated session may last before automatic termination of the session. A value of zero indicates that no session timeout will be applied.

Quarantine Session Idle Timeout (sec)

For Quarantine authentication, the maximum number of consecutive seconds an authenticated session may be idle before automatic termination of the session. A value of zero indicates that the device level setting is used.

Auto Tracking Session Timeout (sec)

For Auto Tracking sessions, the maximum number of seconds a session may last before automatic termination of the session. A value of zero indicates that the device level setting is used.

Auto Tracking Session Idle Timeout (sec)

For Auto Tracking sessions, the maximum number of consecutive seconds a session may be idle before automatic termination of the session. A value of zero indicates that the device level setting is used.

Apply Button

Applies the Login Settings changes to the port.

Automatic Re-Authentication

This tab is grayed out if only web-based authentication is enabled on the device. For 802.1X and MAC authentication, the Automatic Re-Authentication tab lets you set up the periodic automatic re-authentication of logged-in users on this port. Without disrupting the user's session, the device repeats the authentication process using the most recently obtained user login information to see if the same user is still logged in. Authenticated logged-in users are not required to log in again for re-authentication, as this occurs "behind the scenes."



802.1X Re-auth Status

If **Active** is selected, the re-authentication feature is enabled for 802.1X authentication. If **Inactive** is selected, the re-authentication feature is disabled.

802.1X Re-auth Frequency (sec)

How often (in seconds) the device checks the port to re-authenticate the logged-in user via 802.1X authentication. Valid values are 1-2147483647. The default is 3600.

MAC Re-auth Status

If **Active** is selected, the re-authentication feature is enabled for MAC authentication. If **Inactive** is selected, the re-authentication feature is disabled.

MAC Re-auth Frequency (sec)

How often (in seconds) the device checks the port to re-authenticate the logged in user via MAC authentication. Valid values are 1-2147483647. The default is 3600.

Authenticated User Counts

This tab provides authenticated user-count information for devices with Multi-User as their configured authentication type. See the device Authentication tab for information on setting the device authentication type.



Current Number of Users

The current number of users actively authenticated or have authentications in progress on this interface. If **Multi-User** authentication is disabled, this number is **0**. Any unauthenticated traffic on the port is not included in this count.

Number of Users Allowed (up to 2048)

The number of users that can be actively authenticated or have authentications in progress at one time on this interface. If you set this value below the current number of users, end-user sessions exceeding that number are terminated.

NOTE: B2/C2 Devices. If you are configuring a single user and an IP phone per port, set this value to 2.

Number of MAC Users Allowed (up to 2048)

The number of users that can be actively authenticated via MAC authentication, or have MAC authentications in progress at one time on this interface. The number of MAC users allowed cannot exceed the number of users allowed. If you set this value below the current number of users, end user sessions exceeding that number are terminated. If MAC is not selected as a **Multi-User** authentication type on the device Authentication tab, this field will be grayed out.

Number of Quarantine Users Allowed (up to 2048)

The number of users that can be actively authenticated via Quarantine authentication, or have Quarantine authentications in progress at one time on this interface. The number of Quarantine users allowed cannot exceed the number of users allowed. If you set this value

below the current number of users, end user sessions exceeding that number are terminated. If Quarantine Auth is not enabled on the device Authentication tab, this field will be grayed out.

Number of Auto Tracking Users Allowed (up to 2048)


The number of Auto Tracking users that can be actively authenticated or have authentications in progress at one time on this interface. The number of Auto Tracking users allowed cannot exceed the number of users allowed. If you set this value below the current number of users, end user sessions exceeding that number will be terminated. If Auto Tracking is not enabled on the device Authentication tab, this field is grayed out.

Convergence End-Point Access

This tab lists all the CEP (Convergence End-Point) protocols supported by the device on which the port resides, and lets you enable or disable them for that port. For devices that do not support CEP, the tab is blank.

NOTE: Port Mode Authentication Behavior must be set to Active (on the <u>General sub-tab</u>) for authentication to be allowed using these CEP Protocols.

Enable CEP protocols for multiple ports using the Port Configuration Wizard. In addition to enabling protocols on the port, you must also configure CEP for the device on which the port resides. Configure CEP for a single device using the device Authentication tab (CEP sub-tab) or for multiple devices using the Device Configuration Wizard.



CEP Access

Lists all the CEP protocols supported by the device on which the port resides. Use the checkboxes to enable or disable CEP protocols on this port. If the device does not support the CEP feature, this area is blank.

Enable All Button

Selects all the checkboxes and enables all the CEP protocols for this port.

Disable All Button

Deselects all the checkboxes and disables all the CEP protocols for this port.

Apply Button

Applies CEP access changes to the port.



Policy Main Window


The **Control** > **Policy** tab main window is the central point for all **Policy** tab tasks. It is divided into a left panel and a right panel. The tabs in the left panel display hierarchical trees that represent the roles, services, network elements, devices and port groups involved in managing policies for your network. There are five left-panel tabs: Roles/Services, Class of Service, VLANs, Network Resources, and Devices. The tabbed pages in the right panel display detailed information about the item selected in the left panel.

Menu Tabs

The Menu tabs on the Policy tab provide access to Policy tab functions. The Open/Manage Domains menu provides options for the domain currently accessed. The Global Domain Settings drop-down list enables you to configure global Policy tab settings. Use the Tools menu to configure authentication settings and review Policy events.

Information on Policy tab features:

- Dialog Boxes (Messages)
- <u>lcons</u>
- Left Panel




Dialog Boxes (Messages)

In the course of using the **Policy** tab, message dialog boxes appear confirming certain tasks are complete, or warning of the consequences of performing a certain action.

Icons

The icons used in the **Policy** tab and their meanings are as follows:



Icon	Definition	Icon	Definition
柔	VLAN or Network Resource Island		Island VLAN
1	Warning	<u>A</u>	CoS (Class of Service)
A CONTRACTOR OF THE CONTRACTOR	802.1p Priority	4	IP Type of Service Value
The second	CoS Port Group		Rate Limit
411-	Transmit Queue	A	Network Resource Topology

Open/Manage Domain Menu Icons

The following icons appear in the Open/Manage Domains drop-down list:



Reminds you the current Policy Domain is locked for editing purposes. You can lock and unlock the domain from the Lock tool bar button.



Reminds you that you've made changes, and you need to save the data to the Policy Domain. Selecting this icon initiates the save operation. Only users with the capability to Enforce are able to save the domain.



Reminds you that you've made changes to roles that you need to enforce. Selecting this icon initiates the enforce operation.

Policy Windows

The Windows Help section contains Help topics describing Policy tab windows and their field definitions.

Policy Concepts

This topic explains concepts used in the **Policy** tab.

Information on:

- Policy
- Role
 - What is a Role
 - Default Role
- Policy Domains
- Service

- Rule
 - What is a Rule
 - Disabling Rules
 - Conflict Checking
- Packet Tagging
- VLAN to Role Mapping
- Dynamic Egress
 - Setting Domain GVRP Status
- Policy VLAN Islands
- Traffic Mirroring
- Port Groups
- Network Resource Groups
 - Network Resource Topologies
- Verifying
- Enforcing
- Controlling Client Interactions with Locks

Policy

In the **Policy** tab, network access policies are called Roles. See Role, below, for a description.

Role

What is a Role

A role is a set of network access services that can be applied at various access points in a policy-enabled network. A port takes on a user's role when the user authenticates. Roles are usually named for a type of user such as Student or Engineering. Often, role names match the naming conventions that already exist in the organization. A role can contain any number of services in the **Policy** tab.

A role can also contain default access control (VLAN) and/or class of service (priority) characteristics that will be applied to traffic not identified specifically by the set of access services contained in the role. The set of services included in a role, along with any access control or class of service defaults, determine how all network traffic will be handled at any network access point configured to use that role.

Default Role

After you have created a role, assign it as the default role for a port (see <u>Assigning Default Roles to Ports</u>).

Policy Domains

The **Policy** tab provides the ability to create multiple policy configurations by allowing you to group your roles and devices into Policy Domains. A Policy Domain contains any number of roles and a set of devices that are uniquely assigned to that particular domain. Policy Domains are centrally managed in the database and shared between the **Policy** tab clients.

In the **Policy** tab, you work in one current domain at a time. Each domain is identified by a unique name. The Domain menu lets you easily switch from one domain to another. There is no limit to the number of domains you can create, however, a device can exist in only one Policy Domain.

The first time you launch the **Policy** tab, you are in the Default Policy Domain. You can manage your entire network in the Default Policy Domain, or you can create multiple domains each with a different policy configuration, and assign your network devices to the appropriate domain. The roles, services, rules, VLAN membership, and class of service in this initial configuration define a suggested implementation of how network traffic can be handled. This is a starting point for a new policy deployment and often needs customization to fully leverage the power of a policy-enabled network.

The **Policy** tab ships with a set of domain configurations that provide ready-made workflows for common policy scenarios. Each domain configuration contains all the elements (roles, services, rules, VLAN membership, class of service) that define how network traffic is handled for each scenario. These domains are listed in the Open/Manage Domain menu.

You can import the data elements from one domain into another domain. You can also import a domain saved as a policy Database file (.pmd file) or data from a Database file into a domain, and you can export a domain or data from a domain to a .pmd file, (one file per domain) for backup and troubleshooting purposes. Verify and Enforce operations are performed only on the current domain.

In order for your network devices to be displayed on the left-panel **Devices** tab, they must be assigned to a Policy Domain. Initially, you must add your devices to the ExtremeCloud IQ - Site Engine database. After devices have been added to the ExtremeCloud IQ - Site Engine database, you can assign the devices to a Policy Domain using the **Policy** tab. As soon as a device is assigned to a domain, it is automatically displayed on the left-panel **Devices** tab. Only devices that support policy are displayed in the **Policy** tab.

The **Policy** tab automatically locks the current Policy Domain when you begin to edit the domain configuration. Other users are notified that the domain is locked and they are not be able to save their own domain changes until the lock is released. For more information, see <u>Controlling Client Interactions with Locks</u>. After a Policy Domain has been changed, you must save the domain to notify all clients viewing that domain of the change and automatically update their view with the new configuration.

Service

Services are sets of <u>rules</u> that define how network traffic for a particular network service or application should be handled by a network access device. A service might consist of only one

rule governing, for example, email priority, or it might consist of a complex set of rules combining class of service, filtering, rate limiting, and access control (VLAN) assignment. The **Policy** tab allows you to create Local Services (services that are unique to the current domain) and Global Services (services that are common to all domains). Global Services let you easily create and manage services shared between all your domains. A service can be included in any number of roles.

As an example, you might create a service called High Priority Internet Web Access that contains priority classification rules for traffic directed toward each of your organization's Internet proxy servers. This service would likely contain one traffic classification rule for each of your Internet proxy servers.

Services can be one of two types: Manual Service or Automated Service.

- Manual Service This service consists of one or more traffic classification rules you create based on your requirements. Manual services are good for applying customized sets of rules to roles.
- Automated Service This service automatically creates a rule with a specified action (class of service and/or access control), for each device in a particular network resource group. You create a network resource group using a list of IP addresses or an IP subnet, and then associate the group with the Automated service (see How to Create a Network Resource Group for more information). Automated rule types include Layer 3 IP Address and IP Socket rules, and Layer 4 IP UDP Port and IP TCP Port rules.

Services provide a common language that network engineers, information technology administrators, and business managers understand. See How to Create a Service for more information.

Rule

What is a Rule

Policy rules define one element of how traffic for a particular network service or application is handled by a network access device. For example, you might create a rule that assigns a certain priority to all email traffic, by adding an 802.1p, ToS, or DiffServ value to all SMTP traffic. A policy rule can be included in any number of services and you can select the types of devices to which the rule applies. You create rules by right-clicking a Service in the Service Repository tab and selecting Create Rule

See Traffic Classification Rules for a detailed explanation of rules.

Disabling Rules

You can elect to disable a rule during or after its creation. If you disable a rule, it is temporarily unavailable for use by the current service, but it can still be copied to other services and enabled, or re-enabled at another time for the current service. Disabling a rule is a way to temporarily remove a rule from your service without having to delete and recreate it. You disable rules by right-clicking a Service in the Service Repository tab and selecting Disable Rule.

Conflict Checking

As you create your Policy view services and rules, you can define conflicting rules. A conflict exists when two rules in the same service or role define different actions for the same traffic description. For example, two rules might have the same traffic description, but forward traffic to different VLANs, or have different priorities. ExtremeCloud IQ - Site Engine ensures that conflicting rules do not coexist in the same role or service by checking rule traffic descriptions and action values, providing a message if conflicts are found, and writing the conflict information to the Event Log. If a rule is <u>disabled</u>, conflicts between that rule and others are ignored.

The one exception to this conflict checking behavior, is when the conflicting rules coexist in the same role, but one rule exists in a Local service and the other exists in a Global service. In this case, the rule defined in the Local service takes precedence over the rule defined in the Global service because the Local service is specific to the current domain. Consider the following example:

In the North Campus domain you have a Local service "A" that assigns an Ethertype IP rule to the Red VLAN. The "A" service is assigned to the Student Role. In addition, a Global service "B" exists that assigns Ethertype IP rules to the Blue VLAN. The "B" service is also assigned to the Student Role. In this case, the Local service takes precedence over the Global service in the North Campus domain. Note that the precedence pertains to the rule's actions: class of service (priority) and access control (VLAN). For example, if a rule in a Local service and a rule in a Global service both have the same traffic description, and the Local rule's actions apply CoS Priority 1 and no access control (no VLAN), while the Global rule's actions apply CoS Priority 2 and VLAN Blue(2), then the rule will be enforced using CoS Priority 1 and VLAN Blue(2). In addition, if either the Local or Global service has the Accounting or Security actions enabled, then they will be enforced to the devices.

Packet Tagging

Packet tagging in a Policy view environment occurs as follows:

Tagged packets and ingress filtering are processed first. Then, VLAN ID and priority are determined.

- VLAN ID: If the packet matches an active VLAN classification rule on the ingress port, the VID (VLAN ID) specified in the matching VLAN classification rule is assigned. Otherwise, if there is an active role on the ingress port and it specifies a default VLAN, the default VID from the active role on the ingress port is assigned. If there is no active role and no classification rule matches, the 802.1Q PVID for the ingress port is assigned.
- Priority: If the packet matches an active priority classification rule on the ingress port, the priority
 specified in the matching priority classification rule is assigned. Otherwise, if there is an active role on
 the ingress port and it specifies a default priority, the default priority from the active role on the ingress
 port is assigned. If there is no active role and no classification rule matches, the 802.1Q_PPRI for the
 ingress port is assigned.

The set of classification rules active on a port includes statically created rules that specify the ingress port on their port list, as well as any rules established as a result of a role being applied on that port. If the port has no active role and thus no default access control (VLAN) or class of service (priority), untagged packets that do not match any classification rules are assigned a VLAN and priority from the 802.1Q and 802.1p defaults for the ingress port.

For a graphical illustration of the packet tagging process in a Policy view scenario, see the Packet Flow Diagram. The packet passes through the decision-making process illustrated in the graphic twice — one time for VLAN tagging and one time for priority tagging.

VLAN to Role Mapping

VLAN to Role mapping lets you assign a role to an end user based on a VLAN ID. There are two kinds of VLAN to Role Mapping: Authentication-Based and Tagged Packet.

Authentication-Based VLAN to Role Mapping (RFC 3580) — Provides a way to assign a role to a user during the authentication process, based on a VLAN Attribute. An end user connects to a policy-enabled device that supports 802.1X authentication using a RADIUS Server. During the authentication process, the RADIUS server returns a VLAN ID in its RADIUS VLAN Tunnel Attribute. The device uses the Authentication-Based VLAN to Role mapping list to determine what role to assign to the end user, based on the VLAN Tunnel Attribute. Authentication-Based VLAN to Role mappings are only configured at the device level (for all devices).

NOTE: When configuring Authentication-Based VLAN to role mapping, you must enable RFC3580 VLAN Authorization on the device via the device Authentication tab. In addition, VLAN IDs must be configured on the RADIUS server for each user authorized to access the network. If a user does not have a configured VLAN ID, the default role (if there is one) or the 802.1Q PVID for the ingress port is assigned. For more information on configuring VLAN ID attributes on the RADIUS server, refer to your device firmware documentation, RFC 3580, and your RADIUS server documentation.

• Tagged Packet VLAN to Role Mapping - Provides a way to let policy-enabled devices assign a role to network traffic, based on a VLAN ID. When a device receives network traffic that has been tagged with a VLAN ID (tagged packet) it uses the Tagged Packet VLAN to Role mapping list to determine what role to assign the traffic based on the VLAN ID. Tagged Packet VLAN to Role mapping can be configured at the device level (all devices) and at the port level (for an individual port on a device). A VLAN can only be mapped to one role at the device level, but the same VLAN can be mapped to a different role at the port level. A mapping does not have to exist at the device level to be created at the port level, and port-level mappings will override any device-level mappings.

NOTE: TCI Overwrite Requirement

- -- Tagged Packet VLAN to Role Mapping will apply the Role definition to incoming packets using a mapped VLAN. This definition will apply a COS and determine if the packet is discarded or permitted, and if TCI Overwrite is enabled will re-specify the VLAN ID defined by the Rule / Role Default. If TCI Overwrite is disabled, the packet will egress (if permitted by the Rule Hit) with the original VLAN ID it ingressed with.
- -- If supported by the device, you can enable TCI Overwrite for an individual role in the role's **General** tab. The stackable devices support rewriting the CoS values but not the VLAN ID.

To configure VLAN to Role Mapping in the Policy view, use the role's **Mappings** tab and/or the VLAN's **General** tab.

Dynamic Egress

In the **VLANs** tab, you can enable Dynamic Egress for a VLAN by selecting the **Dynamic Egress** checkbox when you select a VLAN.

When Dynamic Egress is enabled for a VLAN, any time a device tags a packet with that VLAN ID, the ingress port is automatically added to the VLAN's egress list, enabling the reply packet to be forwarded back to the source. This means you do not need to add the ingress port to the VLAN's egress list manually. (See Example 1, below.)

Dynamic Egress affects only the egress lists for the source and destination ingress ports. However, GVRP (GARP VLAN Registration Protocol) automatically adds the interswitch ingress ports to the egress lists of VLANs. (See Example 2, below.) You can enable GVRP for the domain by selecting the Global Domain Settings > GVRP > Enable menu option.

NOTE: If you do not want GVRP enabled on your network, you can disable it by selecting the Global Domain Settings > GVRP > Disable menu option. If necessary, you can then manually configure the interswitch ports to do what GVRP does automatically, using local management to set up your interswitch links as Q trunks. The trunk ports will be automatically added to the egress lists of all the VLANs at the time of trunk configuration. For more information on using GVRP in the Policy view, see the section on Setting Domain GVRP Status below.


When you disable Dynamic Egress for a VLAN, the VLAN effectively becomes a discard VLAN. Since the destination port is not added to the egress list of the VLAN, the device discards the traffic. If you want a VLAN to act as a discard VLAN, disable Dynamic Egress for that VLAN. (See Example 3, below.)

If an endstation is talking to a "silent" endstation which does not send responses, like a printer, you need to add the silent endstation's ingress port to the VLAN's egress list manually using local management. Dynamic Egress and GVRP take care of adding the other ingress ports to the VLAN's egress list. (See Example 4, below.)

CAUTION: If no packets are tagged with the applicable VLAN on a port within five minutes, Dynamic Egress list entries time out. The result is that ExtremeCloud IQ - Site Engine indicates that the endstation is "silent" if the VLAN has not been used within that time period. For example, if there is a "telnet" rule and two users (A and B) are on ports whose role includes a service containing the "telnet" rule, if User B has not utilized the "telnet" rule within the five minute time frame, User A is not able to telnet to User B. For this reason, the best application of Dynamic Egress is for containing undirected traffic on "chatty" clients which utilize, for example, IPX, NetBIOS, AppleTalk, and/or broadcast/multicast protocols such as routing protocols.


Example 1: Dynamic Egress Enabled

In this example, Dynamic Egress is enabled for VLAN 5. When source endstation A is tagged with VLAN 5, Dynamic Egress places A's ingress port (1) on VLAN 5's egress list. When destination endstation B's traffic is tagged with VLAN 5, Dynamic Egress places B's ingress port (2) on VLAN 5's egress list. The device can then forward traffic to both endstations.



Example 2: Dynamic Egress + GVRP


In this example, Dynamic Egress is enabled for VLAN 5, and the destination endstation, B, is on a different device from the source endstation, A. When A is tagged with VLAN 5, Dynamic Egress places A's ingress port (1) on VLAN 5's egress list. GVRP then places interswitch ingress ports (2) and (3) on VLAN 5's egress list. When B's traffic is tagged with VLAN 5, Dynamic Egress places B's ingress port (4) on VLAN 5's egress list. GVRP then places interswitch ingress ports (5) and (6) on VLAN 5's egress list. The devices can then forward traffic to both endstations.



G = GVRP places interswitchingress ports on VLAN 5's egress list


Example 3: Dynamic Egress Disabled

In this example, Dynamic Egress is disabled. When source endstation A is tagged with VLAN 5, A's ingress port is not placed on VLAN 5's egress list. GVRP places interswitch ingress ports (1) and (2) on VLAN 5's egress list. When B's traffic is tagged with VLAN 5, B's ingress port is not placed on VLAN5's egress list. GVRP places interswitch ingress ports (3) and (4) on VLAN 5's egress list. But VLAN 5 traffic for both A and B is discarded, because VLAN 5 is not aware of the ingress ports for A and B.



Example 4: Silent Endstation

In this example, Dynamic Egress is enabled for VLAN 5, but the destination endstation, B, is a "silent" endpoint, like a printer. Endstation B does not send responses, so the Administrator must place B's ingress port on VLAN 5's egress list manually (1). When A is tagged with VLAN 5, Dynamic Egress places A's ingress port (2) on VLAN 5's egress list. GVRP then places interswitch ingress ports (3) and (4), then (5) and (6) on VLAN 5's egress list. Endstation A is then able to communicate with the printer.



- M = Administrator manually places
 B's ingress port on VLAN 5's egress
 list
- D = Dynamic Egress places A's ingress port on VLAN 5's egress list
- G = GVRP places interswitch ingress ports on VLAN 5's egress list

Setting Domain GVRP Status

The Policy view allows you to set the domain GVRP (GARP VLAN Registration Protocol) status via the Edit menu. There are three GVRP status options. To set the GVRP status for all the devices in the current domain, select a status and then enforce.

- **Ignore** When this option is selected, ExtremeCloud IQ Site Engine ignores the GVRP configuration on a device during an Enforce operation. This allows you to configure some network switches with GVRP enabled and others with GVRP disabled, according to their configuration requirements.
- **Enable** When this option is selected, GVRP is enabled for the devices in the current domain.
- **Disable** Select this option if you do not want GVRP enabled on the devices in the current domain. Disabling GVRP can affect connectivity through ports with VLANs that rely on Dynamic Egress. If GVRP is disabled, rules using VLAN containment do not work properly unless the VLANs have been preconfigured on the devices outside of ExtremeCloud IQ Site Engine.

The following table shows how domain GVRP status affects device-level and port-level GVRP status when an Enforce operation is performed.

Domain GVRP Status	Device Set on Enforce	
Domain GVRP status is set to Ignore .	No GVRP status is written to devices on Enforce.	
Domain GVRP status is set to Enable and the device-level GVRP is enabled.	No GVRP status is written to the device on Enforce.	
Domain GVRP status is set to Enable and the device-level GVRP is disabled.	Device-level GVRP status and port-level GVRP status is set to enabled on Enforce.	
Domain GVRP status is set to Disable and the device-level GVRP is disabled.	No GVRP status is written to the device on Enforce.	

Domain GVRP Status	Device Set on Enforce
Domain GVRP status is set to Disable and the device-level GVRP is enabled.	Device level GVRP status is set to disabled and no change is made to the port-level GVRP status on Enforce.

Policy VLAN Islands

The Policy view offers you the ability to set up Policy VLAN Islands which enable you to deploy a policy across your network, while restricting user access to only selected local devices. For example, if you want to have a guest VLAN but you do not want the guests in one facility to be able to communicate with guests in another facility, you can set up a VLAN island containing only selected devices in each facility, with access controlled by island VLANs.

- Global VLAN Global VLANs are written to all selected devices with the same VID. They are referenced in the format <VID[name]>.
- Island VLAN An Island VLAN is a conceptual VLAN and does not have an actual VID. The VID is assigned automatically based on the island it belongs to.

NOTE: The Policy view provides management of Global VLAN settings, but does not provide management of Island VLANs beyond setting the appropriate VIDs in the Role defaults and Rule access control actions. Also, you must manage separately other related settings in the qBridgeMib such as name, and dynamic egress values.

See How to Create a Policy VLAN Island for more information.

Traffic Mirroring

The Policy view provides policy-based traffic mirroring functionality that allows network administrators to monitor traffic received at a particular port on the network, by defining a class of traffic that will be duplicated (mirrored) to another port on that same device where the traffic can then be analyzed. Traffic mirroring can be configured for a rule (based on a traffic classification) or as a role default action. Only incoming traffic can be mirrored using policy-based traffic mirroring, and the traffic mirroring configuration takes precedence over regular port-based mirroring.

Traffic mirroring uses existing the Policy view port groups (created using the Port Groups tab) to specify the ports where the mirrored traffic will be sent for monitoring and analysis. When an end user connects to the device where the specified ports exist, and is assigned the role that has traffic mirroring configured, then there is a traffic mirror set up for the port the end user connected to. However, if the end user is assigned a role that does not have traffic mirroring configured, or if the end user connects to a device that doesn't have any ports in the specified port groups, then no traffic mirror will exist.

Examples of how traffic mirroring might be used include:

- Mirroring the traffic from suspicious users based on their MAC or IP address.
- Monitoring VoIP calls by IP address or port range.

- Mirroring traffic to optimized IDS systems, for example one system for all HTTP traffic (to look for suspicious websites) or one system for all emails (to look for spam).
- Mirroring traffic to ExtremeAnalytics appliances for use in ExtremeCloud IQ Site Engine application identification reports and analysis.

For information on configuring traffic mirroring, see the Role tab and the Rule General tab.

Port Groups

ExtremeCloud IQ - Site Engine allows ports to be combined into groups, similar to the way services can be combined into service groups. Port groups enable you to configure multiple ports on the same device or on different devices simultaneously, or to retrieve port information from them. You can view port groups on the left-panel **Port Groups** tab.

The Policy view provides you with several commonly used port groups for your convenience, called Pre-Defined Port Groups. You can also create your own port groups, called <u>User-Defined Port Groups</u>.

User-Defined Port Groups

The Policy view also enables you to create your own port groups and select individual ports to add to the group.

Network Resource Groups

Network Resource Groups provide a quick and easy way to define traffic classification rules for groups of network resources such as routers, VoIP (Voice over IP) gateways, and servers. The default Policy domain configuration contains examples of network resource groups that you might want to create, such as Internet Proxy Servers and SAP Servers. Use the Network Resource Configuration window to view and define your network resource groups. See How to Create a Network Resource for more information.

After a network resource group has been defined, you can associate it with an Automated service (see How to Create a Service for more information). The Automated service automatically creates a rule with a specified action (class of service and/or access control), for each resource in the network resource group. Automated rule types include Layer 2 MAC Address rules, Layer 3 IP Address and IP Socket rules, and Layer 4 IP UDP Port and IP TCP Port rules.

Network Resource Topologies

Network Resource Topologies are used to divide the devices in a domain into groups called islands. Each network resource group specifies a topology and can then define a unique resource list for each island within that topology, allowing user access to resources on the network based on the physical location at which they authenticate.

For example, you could create a topology called "Campus Printers" that could be used to restrict printer access to only the printers in the building where the end user is physically located. This topology might define islands such as "Library," "Admissions Office," or "Science

Building." Each island would include the network devices for that location. Then, in the Network Resource Group that specifies this topology, there would be resource lists that define the printers for each of those islands.

In addition to defining topologies based on physical location (such as geographic region, corporate offices, or campus buildings) a topology could also be used to define resources based on the departments within a company (such as Sales, IT, or Human Resources).

When you create a topology, it contains a Default Island that includes all the devices in your domain. You can then create additional islands and distribute your devices between the different islands according to your needs. Each device in a domain must belong to one island in each topology. You can set any island as the Default island for new devices that are added to the domain.

Verifying

The Verify feature lets you verify that the roles in your current domain have been enforced. Verify operations are performed only on the current domain. The Verify operation compares the roles currently in effect (<u>enforced</u>) on your domain devices with the roles defined in the current Policy Domain.

NOTE: If you perform a Verify operation following an Import Policy Configuration from Device, the Verify can fail. This is because the import operation imports only roles and rules from the device, not the complete policy configuration. Also, when you import device-specific rules, these rules are converted to a Rule Type of "All Devices," and this will cause Verify to fail. If you want the rules to be device-specific, you will have to change their Rule Type via the Rule General tab after the import and prior to Enforce.

You can verify using the Open/Manage Domain > Verify Domain menu option, both of which verify the information on all the devices in the current domain. You can also selectively verify on individual devices or device groups in the domain by right-clicking the device or group in the left panel or in the right-panel Details View tab for the Devices folder or Device Group folder, and choosing **Verify** from the menu.

After verifying, you see a window that reports any discrepancies. The title bar of the window lets you know if the verify was done on all devices in the domain, or a subset of devices. From this window, you can select **Enforce Domain** to open the Enforce Preview window, where you can view the effects <u>enforcing</u> the current role set would have, prior to actually enforcing. You can also view the full results of the Verify operation in the event log, which displays any discrepancies and statistics of the operation itself.

Enforcing

In the **Policy** tab, enforcing means writing role information to a device or devices. Enforce operations are performed only on the current domain. Any time you add, make a change to, or delete a role or any part of it (any of its services and/or rules), the devices in your current domain need to be informed of the change, otherwise the role will not take effect. To determine

if the roles currently in effect on your domain devices match the set of roles you have defined in your current Policy Domain configuration, use the Verify feature.

NOTE: Setting up Profiles and Credentials for Enforce. All SNMP operations that are performed from the Policy view client use the SNMP credentials of the logged-in user. For example, when devices are identified, the credentials associated with the user's group are used to communicate with the devices. However, the Enforce operation occurs on the server and uses the Netsight Administrator profile to communicate with devices. Because of this, the Netsight Administrator profile must have write privileges on the devices that users can enforce.

When an Enforce is initiated, the Policy Domain is locked to prevent other clients from enforcing at the same time. Different Policy Domains can be enforced at the same time, but if another user attempts to enforce the same domain at the same time, that user will be notified that the domain is already locked.

To enforce, select the Open/Manage Domains > Enforce Domain menu option. You can also selectively enforce on individual devices by right-clicking the device in the **Devices** tab left panel or in the right-panel **Devices** tab and choosing **Enforce** from the menu. Only users that have been assigned the Enforce capability are allowed to perform an Enforce.

Controlling Client Interactions with Locks

Because the Policy view uses a Client/Server architecture, it is important to maintain a proper sequence of client interactions to ensure a consistent view of Policy Domains among all clients. To do this, the Policy view uses Server Locks to manage user interactions. When a user begins editing a Policy Domain (for example by assigning devices or adding a role), a lock is acquired for that domain at the server. That lock is not released until the same user saves the domain data. This guarantees a consistent view of that domain for all clients. Users are given the option of revoking locks held by other users. This protects against the possibility that users forget they have locked a domain and keep that lock for an extended period of time.

A domain is locked automatically when a user begins to edit the domain data or a user can lock/unlock a domain by selecting the Lock toolbar button. When a domain is locked, the title bar states that the policy data is being edited and specifies the user who has locked the domain. Other Policy view clients are notified that the domain is locked and they will not be able to save their own domain changes until the lock is released.

Here are some important things to remember about locks:

- Locks operate on individual Policy Domains. When a user edits a domain, a lock is acquired for that domain and it remains locked until the same user saves the domain data or the lock is revoked by another user. You cannot save a domain that is locked by another user.
- During Enforce, a lock is acquired on the domain which is being enforced. This ensures a consistent view of the domain while it is being used by the server.
- When devices are being assigned to a Policy Domain, multiple domains can be locked concurrently. This will happen if devices from one domain are being reassigned to another domain. In this case, locks for both domains are acquired.

• When a lock is revoked, the last domain save "wins." While consistency is always maintained by the server, the order of domain saves cannot be guaranteed when locks are revoked, and consequently work done by one user can be lost.

You can view server locks for all clients via the Options > Server Information tab.

◬

Policy Tab Right-Panel

The **Policy** tab main window is divided into two panels: a left panel and a right panel. The Right-Panel Tabs Help section contains Help topics describing the tabs and their field definitions.

The right panel displays different tabs and information depending on the item selected in the left-panel tree. Help topics for right-panel tabs are named in a manner to reflect this. For example, the help topic named Details View Tab (Device Group), provides information on the right-panel **Details View** tab when a device group is selected in the left-panel tree.

Policy Left Panel

The left panel of the **Policy** tab contains tabs that display hierarchical trees representing the roles, services, classes of service, VLANs, network resources, devices, and port groups involved in managing policies for your network. What you select in the left panel determines what is displayed in the right panel. When you first open the Policy tab, the Roles tab is displayed in the left panel, by default.

Features of the left panel include:

- Expanding and collapsing items in the hierarchy: Double-click the item or its icon, or select the turner to the left of the icon.
- *Right-click menus:* Right-click a folder or other item in the left panel, and a menu of the options you can perform on your selection appears.

Information on the left-panel tabs:


- Roles/Services Tab
- Class of Service Tab
- VLAN Tab
- Network Resources Configuration
- Devices/Port Groups Tab

Roles/Services Tab

This tab displays the Roles and Service Repository trees.

Roles Tree

The Roles tree lists the roles defined for the current domain. A role is a set of network access services that can be applied at various access points in a policy-enabled network.



Roles Folder


This folder contains the roles defined for the current domain. See How to Create a Role for more information.

Role

Individual roles are listed by name. Select a role in the left panel, and view information about that role in the right-panel tabs. Only Quarantine roles are displayed with a red icon .

Service Repository Tree

The Service Repository tree displays your Local and Global services and service groups. Services are sets of rules that define how network traffic for a particular network service or application is handled by a network access device. Local Services are services unique to the current domain. Global Services are services common to all domains. The tab also displays your network resource groups.



Local Services Folder

Local Services are services unique to the current domain. This folder contains the local service groups and services defined for the current domain. For more information, see How to Create a Service Group.

Global Services Folder

Global Services are services that are common across all domains. This folder contains the global service groups and services shared by all domains. For more information, see How to Create a Service Group.

Service Groups Folder

The **Policy** tab lets you create categories (service groups) into which you can group services. This folder contains the defined service groups. For more information, see How to Create a Service Group.

Service Group 🧆

Individual service groups are listed by name. Expand the service group to see the services and service groups included in that group.

Services Folder

This folder contains the automated and manual services that have been defined. For more information, see How to Create a Service.

Automated Service 👶

Individual Automated services are listed under the Services Folder or within a service group in the Service Groups folder.

Manual Service 👶

Individual Manual services are listed under the Services Folder. Expand the service to see the rules associated with it.


Rule 🥯

Individual rules are listed by name. If the rule is disabled, the rule icon displays a red X $^{\bigcirc}$. If the rule is device-specific, the rule icon displays a small switch $^{\bigcirc}$.

Class of Service Tab

The left panel Class of Service tab displays your Classes of Service defined for the current domain.

Classes of Service prioritize traffic with an 802.1p priority, and optionally an IP type of service (ToS/DSCP) value, rate limits, and transmit queue configuration. You can then assign the class of service as a classification rule action, as part of the definition of an Automated service, or as a role default. For more information, see Getting Started with Class of Service.



Classes of Service Folder

When you first access the **Policy** tab, the left-panel Classes of Service tab is pre-populated with eight classes of service, each associated with one of the 802.1p priorities (0-7). These are static classes of service and cannot be deleted. You can use these classes of service as is, or configure them to include ToS/DSCP, rate limit, and/or transmit queue values. You can also rename them, if desired. In addition, you can also create your own classes of service. After you have created and defined your classes of service, they are then available when you make a class of service selection for a rule action (Rule tab), a role default (General tab), or an automated service (General tab).

Class of Service[△]

Select a Class of Service in the left panel, and view information about that service in the right-panel tabs. For more information, see How to Create a Class of Service.

CoS Components Folder

This folder contains subfolders of the possible components of a class of service (Rate Limits, Inbound Rate Limit Port Groups, Outbound Rate Limit Port Groups, and Transmit Queue Port Groups).

Rate Limits Folder

This folder contains the currently defined rate limits, listed in the order of precedence. For more information, see How to Define Rate Limits.

Inbound Rate Limit Port Groups

This folders contains the currently defined inbound rate limit port groups. Select a port group in the left panel and view information about that group in the right-panel tabs. For more information, see Creating Class of Service Port Groups.

Outbound Rate Limit Port Groups

These folders contain the currently defined outbound rate limit port groups. Select a port group in the left panel and view information about that group in the right-panel tabs. For more information, see Creating Class of Service Port Groups.

Transmit Queue Port Groups Folder

This folder contains the currently defined transmit queue port groups and the transmit queues defined for each group. For more information, see How to Configure Transmit Queues.

VLAN Tab

The left panel VLAN tab displays the Global VLANs for the current domain. If you have enabled Policy VLAN Islands, it also displays your Island VLANs and Policy VLAN Islands.



Global VLANs Folder

This folder contains your currently defined global VLANs for this domain.

VLAN 🕮

The VLAN icon indicates the access control for the VLAN-- if it is a Discard VLAN, the icon displays a red X . Otherwise, it is a Contain VLAN.

Island VLANs Folder

This folder appears only when the Policy VLAN Islands feature is enabled, and contains your currently defined Island VLANs for this domain.

Policy VLAN Islands Folder


This folder appears only when the Policy VLAN Islands feature is enabled, and contains your currently defined VLAN islands and the devices that belong to them. When you enable Policy VLAN Islands, this folder is pre-populated with a Default Island containing all the devices in the domain.

VLAN Island X

Select a VLAN island to see the devices associated with it listed in the right-panel Details View tab. The Default Island is created by the Policy tab when you enable Policy VLAN Islands, and it cannot be deleted.

Network Resources Configuration

The **Network Resources** left-panel tab displays the network resources and network resource topologies for the current domain.



Network Resources Folder

This folder contains any network resource groups you have created. For more information, see How to Create a Network Resource.

Network Resource

Individual network resource groups are listed by name. Select a resource in the left panel, and view information about that resource in the right-panel tabs.

Global Network Resources Folder

Global Network Resources are network resources that are common across all domains. For more information, see How to Create a Network Resource.

Network Resource Topologies Folder

This folder contains the network resource topologies currently defined for this domain.

Network Resource Topology 📤

A network resource topology can be used to divide the devices in a domain into groups called islands. You can then define a unique network resource list for each island within that topology, allowing user access to resources on the network based on the physical location at which they authenticate. If you are not using custom topologies to group your devices, you will use the Domain Wide topology, which contains just one island for all your domain devices.

Topology Island 🏝


A topology island is a group of devices that have a unique network resource list, allowing you to set up network resource access based on the location where end users authenticate.

Devices/Port Groups Tab

This tab displays the Devices and Port Groups trees.

Devices Tree

The Devices tree displays the devices assigned to the current domain, organized into groups.



Devices

This tab contains all the devices assigned to the current domain. For information on adding devices to the domain, see How to Add and Delete Devices.

ExtremeControl > Policy supports Per-User ACLs (PU-ACL) from third-party vendors passed via RADIUS authentication requests. During a policy enforce, the roles and associated rules are translated into ACLs and pushes them to the appropriate Access Control Engines. You can manage ACL rules onExtremeXOS/Switch Engine devices on which version 30.5 or later is installed. By using ACLs, the access control entries (ACEs) can be ordered by the administrator, allowing for more flexibility in the configuration and better utilization of hardware resources on the device.

The Control > Policy > Devices/Port Groups > Devices tab includes ACL Rule Usage and Rule Hit Count details.


Port Groups

This tab contains the Pre-Defined and User-Defined Port Groups for the current domain. The **Policy** tab allows ports to be combined into groups, similar to the way devices are combined into device groups.

Port groups enable you to configure multiple ports on the same device or on different devices simultaneously, or to retrieve port information from them. For more information, see How to Create a Port Group.

Summary (Roles)

This tab provides a summary view of the domain's roles. To access this tab, select the **Roles** left-panel tab in the Roles/Services tab. Right-click a role to add/remove services, rename the role, or delete the role.




General (Role)

The role **General** tab lets you assign default actions for a role applied to traffic not identified specifically by the set of access services contained in the role. You can also use this tab to enable TCI Overwrite functionality for the role, and enter or edit the description of the role.

The Services section displays a list of the services and service groups associated with the selected role, and provides buttons for adding and removing services, creating a new service, viewing and editing a service or service group, and showing conflicting rules.

To access this tab, select a role in the left panel's **Roles** tab, then select the **General** tab in the right panel. Any additions or changes you make to this tab must be enforced in order to take effect.



Name

Name of the selected role.

Description

Use the **Edit** button to open a window where you can enter or modify a description of the role.

TCI Overwrite

Enable or disable TCI Overwrite functionality for the role. Enabling TCI Overwrite enables the VLAN (access control) and class of service characteristics defined in this role or any of its rules to overwrite the VLAN or class of service (CoS) tag in a received packet if that packet has already been tagged with VLAN or CoS information. If TCI Overwrite is not enabled, tagged packets will egress using the TCI data they already contain. You can also enable TCI Overwrite on a per-rule basis in the Rule Tab.

Default Actions

Default actions for a role are applied to traffic not identified specifically by the set of access services contained in the role.

Access Control

Use the drop-down list to choose a default access control (VLAN) for the role. You can select:

- None No default access control specified.
- Permit Traffic Enables traffic to be forwarded with the port's assigned VID.
- Deny Traffic Traffic will be automatically discarded.
- Contain To VLAN This option contains traffic to the VLAN specified. Use the drop-down list to
 the right to select the desired VLAN. You can also define the Service ID to extend the VLAN
 address space. The Service ID is the implementation of ExtremeCloud IQ Site Engine for the ISID (also called Network Service Identifier = NSI), which increases the number of available
 VLANs.

Class of Service

Use the drop-down list to choose a default class of service (priority) for the role, create a new class of service, or select None if no class of service is desired. The drop-down list displays all of the classes of service for the current domain and also enables you to edit a class of service using the Edit button ...

System Log

When this option is enabled, a syslog message is generated as long as no matching rules specify that sending a syslog message is prohibited (that is, the rule's system log action is set to "Prohibited" on the Rule Tab). When the option is disabled, the system log setting is ignored.

Audit Trap

When this option is enabled, an audit trap is generated as long no matching rules specify that sending an audit trap is prohibited (that is, the rule's audit trap action is set to "Prohibited" on the Rule Tab). When the option is disabled, the audit trap setting is ignored.

Disable Port

When this option is enabled, the port is disabled as long no matching rules specify that disabling the port is prohibited (that is, the rule's disable port action is set to "Prohibited" on the Rule Tab). Ports that have been disabled due to this option are displayed in the device Role/Rule tab. When the option is disabled, the disable port setting is ignored.

Traffic Mirror

Use the drop-down list to specify port groups where mirrored traffic is sent for monitoring and analysis. Select View/Modify Port Groups to open the Port Groups tab where you can define user-defined port groups for selection.

To the right of the drop-down list is an option to mirror only the first (N) packets of a flow. This option is intended for use when mirroring traffic to an ExtremeAnalytics engine. The ExtremeAnalytics engine only needs the initial packets of a flow to properly identify the traffic, and setting this option will reduce network traffic overhead for the switch and engine. By default this number is set to 10, but can be changed by selecting the Edit button. Note that the value you set is used by all mirror actions in use in the current domain.

Services

Name

Lists the names of the services and service groups (local and global) associated with the selected role.

Also Used By Roles

List the other roles using this service. If the service is a global service, the domain name is also displayed if the role is in a different domain.

Add/Remove Services Button

Opens the role Add/Remove Services window, where you can add and remove services and service groups to and from any of the existing roles.

Show Details Button

Select a service or service group in the table and select this button to open the left-panel Services tab. The appropriate service or service group will be selected and you can access its right-panel tabs.

Show Conflicting Rules Button


If the rules in a Global service conflict with the rules in a Local service, the Name column will display a message indicating that the global rules will be overridden by the local rules. Select the **Show Conflicting Rules** button to open a window that displays the rule conflicts and shows specifically which rules will be used and which will be overridden. For more information, see Conflict Checking.



VLAN Egress (Role)

The role VLAN Egress tab displays the list of VLANs on the selected role's egress list, and allows you to add and remove VLANs and set their Egress Forwarding State. Ports that the selected role is active on forwards traffic belonging to the listed VLANs according to the specified forwarding state. Both the role's egress list and the VLAN egress list are checked for egress information. If the lists have duplications, the Forbid Forwarding state takes precedence.

To access this tab, select a role in the left panel's **Roles/Services** tab and select the **VLAN Egress** tab in the right panel. Any changes made on this tab need to be enforced.



VID

The VLAN ID.

Name

The VLAN Name.

Egress Forwarding State

Ports on which the selected role is active forward traffic belonging to this VLAN according to the egress forwarding state: Tagged (frames are forwarded as tagged), Untagged (frames are forwarded as untagged), or Forbid Forwarding (frames are not forwarded; they are discarded).

Add

Opens the Add Egress VLAN Window, where you can choose a VLAN for the role's egress list and specify the egress forwarding state.


Remove

Select a VLAN and select **Remove** to remove the VLAN from the list.

◬

Add Egress VLAN Window

The Add Egress VLAN window appears when you select the **Add** button in the role's VLAN Egress tab. It allows you to add a VLAN to the Role's Egress list and specify the egress forwarding state.



VLAN

This is a drop-down list of the available VLANs.

Forwarding State

Select the desired forwarding state: Tagged (frames are forwarded as tagged), Untagged (frames are forwarded as untagged), or Forbidden (frames are not forwarded; they are discarded).

Mappings (Role)


This tab lets you view and configure four different mapping lists for the selected role:

- MAC to Role Mapping Lets you assign the role to an end user based on the user's MAC address.
- IP to Role Mapping Lets you assign the role to an end user based on the user's IP address.
- Tagged Packet VLAN to Role Mapping Lets you assign the role to network traffic based on the traffic's VLAN ID.
- Authentication-Based VLAN to Role Mapping Lets you assign the role to an end user during the authentication process, based on a VLAN Attribute.

To access this tab, select a role in the left-panel **Roles** tab and select the **Mappings** tab in the right panel. Any additions or changes you make to this tab must be enforced in order to take effect.

NOTE: TCI Overwrite Requirement

- -- Tagged Packet VLAN to Role Mapping applies the Role definition to incoming packets using a mapped VLAN. This definition applies a CoS and determine if the packet is discarded or permitted, and if TCI Overwrite is enabled re-specifies the VLAN ID defined by the Rule / Role Default. If TCI Overwrite is disabled, the packet egresses (if permitted by the Rule Hit) with the original VLAN ID with which it ingressed.
- -- If supported by the device, you can enable TCI Overwrite for an individual role in the role's General tab. The stackable devices support rewriting the CoS values but not the VLAN ID.



Primary Stackable Tagged VLAN Mapping

Use this column to select the device-level VLAN to role mapping used for C2/C3/C5 and B2/B3/B5 devices (C2 firmware version 03.02.xx and higher/B2 firmware version 02.00.16 and higher), and D2, A4, and G3 devices (G3 firmware version 6.03.xx and higher). These devices only support one device-level VLAN to role mapping. If you do not make a selection, there will be no device-level mapping for these devices. Use the Mappings tab in the Enforce Preview window to quickly see which VLAN to role mapping is selected for these devices.

Type

This column indicates the type of mapping: <u>MAC to Role</u>, <u>IP to Role</u>, <u>Tagged Packet VLAN to Role</u>, and Authentication based VLAN to Role.

Value

The MAC addresses, IP addresses, or VLAN mapped to this role.

Src/Dst

Specifies whether the MAC address is a source or destination address.

Device/Port Level

This column indicates whether the mapping is a device-level mapping (all devices) or a port-level mapping (IP address and port description).

Add Button

Opens the Add Role Mapping window, where you can add a new Role mapping by entering the Mapping Type, Value, and Direction.

Remove Button

Remove the selected mapping from the list by selecting **Remove**.

MAC to Role Mapping

MAC to Role mapping provides a way to assign a role to an end station based on its MAC address. This enables you to create a specific role for a group of end stations (such as IP phones), and assign it to them based on their MAC address. When the end stations connect to

the network, the policy-enabled device identifies the source MAC address and applies the mapped role.

IP to Role Mapping

IP to Role mapping provides a way to assign a role to an end station based on its IP address. For example, in networks that haven't deployed authentication, this would enable you to map an individual IP address such as an administrator's laptop, to a specific role. When the end station connects to the network, the policy-enabled device identifies the IP address and applies the mapped role.

Tagged Packet VLAN to Role Mapping

Tagged Packet VLAN to Role mapping provides a way to let policy-enabled devices assign a role to network traffic, based on a VLAN ID. When a device receives network traffic that has been tagged with a VLAN ID (tagged packet) it uses the Tagged Packet VLAN to Role mapping list to determine what role to assign the traffic based on the VLAN ID. For more information, see VLAN to Role Mapping in the Concepts Help topic.

Authentication-Based VLAN to Role Mapping

Authentication-Based VLAN to Role mapping provides a way to assign a role to a user during the authentication process, based on a VLAN Attribute. An end user connects to a policy-enabled device that supports 802.1X authentication using a RADIUS Server. During the authentication process, the RADIUS server returns a VLAN ID in its RADIUS VLAN Tunnel Attribute. The device uses the Authentication-Based VLAN to Role mapping list to determine what role to assign to the end user, based on the VLAN Tunnel Attribute. Use this table to view and configure the VLANs that will map to the selected role. For more information, see VLAN to Role Mapping in the Concepts Help topic.



Pre-configured Domains (Legacy)

To help you quickly achieve the best policy configuration for your network, the **Policy** tab provides pre-configured domains that include roles, services, and rules designed for specific network scenarios. You can use these pre-configured domains as templates, customizing them for your own network requirements.

When you first access the **Policy** tab, it opens to the Default Domain. This domain can be deployed "as-is" for most networks, and provides a complete set of roles, services, and rules, as well as multiple switch platform support.


The **Policy** tab also provides additional pre-configured domains tailored to more specific network scenarios. These domains are named according to the policy configuration they

provide, for example, HealthCare Services and Secure Guest. Below is a brief description of each domain along with some suggestions on how to use the domain in your network environment.

Access Pre-Configured Domains

Access the pre-configured domains from the **Open/Manage Domain** drop-down list. At the top of the menu, the **Open Domain** menu displays all available domains from which you can select. To open a domain, select it in the menu. The domain opens in the **Policy** tab and you can look at the various roles, services, and rules that have been pre-configured in the domain.

As you look at a domain, use the extensive tool tips for the roles and services (as shown below) to view specific information on how to customize the domain to meet your requirements.



Pre-configured Domain Descriptions

The following sections describe the pre-configured domains available from the Domain menu.

Embedded NAC Domain

This domain can be used to configure the policy used by the Embedded ExtremeControl engine. By default it will let traffic through unrestricted as you monitor your network.

Generic Services N-Series

This domain is designed to help networks that use Enterasys N-Series devices to increase security in their existing infrastructure. The roles defined in this domain leverage the capabilities supported on the N-Series. They are based on the best-practice of "least privilege" where all incoming traffic is denied access, and permit rules are used to permit only specific traffic onto the network. The rules also provide appropriate traffic classification.

Start with the roles, services, and rules defined in this domain for your N-Series devices and then expand and customize the domain to meet your own day-to-day business requirements.

Generic Services SecureStack

This domain is designed to help networks that use Enterasys SecureStack devices to increase security in their existing infrastructure. The roles defined in this domain leverage the capabilities supported on the SecureStack products, but will also work on N-Series devices. They are based on the best-practice of "least privilege" where all incoming traffic is denied access, and permit rules are used to permit only specific traffic onto the network. The rules also provide appropriate traffic classification.

Start with the roles, services, and rules defined in this domain for your Securestack devices and then expand and customize the domain to meet your own day-to-day business requirements.

HealthCare Services

The Healthcare Services domain provides a template of roles and services that can be utilized in healthcare industry networks. Roles correspond to the different business roles in health care settings, such as Physician, Nurse, Patient, IT, Hospital Administration, Management, and Guest. Services support a wide range of hospital departments, such as Cardiology, Emergency, Pediatrics, and Payroll/Benefits.

Quickstart

The Quickstart domain gets you up and running quickly with a set of roles, services, and rules that will increase security on your network. Most of the defined roles permit access to the network with certain rules designed to deny or prioritize applications, protocols, and communication traffic on the network. The services are bare minimum examples, and it is suggested that you modify or add roles, services, and rules to meet your day-to-day business requirements. HOW IS THIS DIFFERENT FROM THE DEFAULT DOMAIN?

Note: Before enforcing the policy configuration, set Class of Service mode for the device (select the device in the Policy Manager Network Elements tab, then select the General tab) to "Role-Based Rate Limits / Transmit Queue Configuration". The default Class of Service mode can be specified in the Tools->Options view, and multiple devices can have their Class of Service mode changed using the Device Configuration Wizard in the Tools menu. THIS NOTE IS IN THE TAB DESCRIPTION, IS IT NEEDED?

Secure Guest

Secure Guest is a collection of sample services that you can use to increase security on edge ports where guest users connect.

There is one Secure Guest Access role that enables the end user basic guest services based on the principle of "least privilege" and will permit end users access to HTTP, HTTPS, and PPTP services. Apply this role to an Enterasys policy capable switch port.

The services are bare minimum examples, and it is suggested that you modify or add roles, services, and rules to meet your own business requirements.

ShoreTel

The ShoreTel domain provides a template for traffic prioritization of VoIP traffic on ShoreTel IP Phones that operate with the Media Gateway Control Protocol (MGCP) protocol. Class of Service is configured to provide higher priority to VoIP data, signaling and call control protocol, while lower priority is assigned to other required ShoreTel traffic such as DHCP and TFP.

The defined ShoreTel_IP_Phone role is based on the best practices methodology of "least privilege" where all incoming traffic is denied access, and permit rules are used to permit only specific traffic onto the network. Start by using the template for your N-Series devices, then add custom roles, services, and rules to meet your own network requirements.

VPN Termination Point

VPN Termination Point is a collection of Site-to-Site and Client-to-Site Roles that you can use to enable a VPN Concentrator to initiate, respond-to, and communicate to other VPN termination end points.

Add/Remove Services (Roles)


Add and remove services and service groups from roles using the Add/Remove Services window.

To access the Add/Remove Services window, you must have a role selected in the left-panel **Roles** tab. Select the **Add/Remove** button in the Services section of the Role window.

If you add a service to a role and any or all of the following conditions exist, you are in effect adding an "empty" service, and a warning message displays when you select **OK**:

- No traffic description exists for one or more of the classification rules.
- No access control or class of service has been defined for one or more of the classification rules.
- All of the classification rules are disabled.

When you add a service to a role which already has services associated with it, the **Policy** tab checks for rule conflicts. See Conflict Checking for more information.



All Services & Service Groups

This field displays all the services (local and global) and service groups in the current domain. Select the service groups or services you want to add to the role.

Selected Services & Service Groups

This field displays all the services currently defined for the selected role. Select the services you want to remove from the role.

Right Arrow

Select the **Right Arrow** to add the services or service groups selected in the All Services & Service Groups column to the Selected Services & Service Groups field.

Left Arrow

Select the Left Arrow to remove the services selected in the Selected Services & Service Groups field.

Double Left Arrow

Select the Double Left Arrow to remove all the services in the Selected Services & Service Groups field.

Details View (Service)


This tab displays information about the rules contained in a Manual service or an Automated service. To display this tab:

- 1. Select a service in the left-panel's Roles/Services > Service Repository tab.
- 2. Open either the Local Services tab or Global Services tab, depending on the type of service.
- 3. Select a service from within the **Services** left-panel tab.

The **Details View** tab opens in the right panel. Right-click a rule in the table to see a menu of available options.

NOTE: Rules included in services are read in the order in which they are listed in ExtremeCloud IQ - Site Engine. To configure rules for ExtremeCloud IQ Controller (formally called XCC or XCA) devices, ensure ExtremeCloud IQ - Site Engine lists the rules in the correct order or the service may not execute the correct rule. To reorder rules in the same service, use drag-and-drop capabilities to move from one group to another.

For Manual services, you can double-click on any of the table columns opens the rule's **General** tab.



Name

Name of the rule. For rules contained in an Automated service, this column gives detailed information about the rule including the associated Network Resource (NR), if multiple resource groups are specified. You can rename a rule by right-clicking the rule and selecting **Rename**.

Rule Status

Indicates whether the rule is currently available for use by this service (Enabled), or not (Disabled), as set in the General tab for the rule. If the rule is disabled, the rule icon displays a red X . You can enable or disable a rule by right-clicking and selecting **Enable Rule** or **Disable Rule**, respectively.

Rule Type

Indicates the device types to which the rule applies. (See Create Classification Rule Window for more information.)

Traf Desc Type

Traffic classification type for the rule. (See Classification Types and their Parameters for more information.)

Traf Desc Value

Values associated with the traffic classification type for the rule. (See Classification Types and their Parameters for more information.) Double-clicking on this column opens the Edit Rule window, where you can edit the parameters or values for the rule's classification type.

Access Control

VLAN action associated with the rule. Double-clicking on this column allows you change the setting. You can permit traffic to be forwarded, deny traffic altogether, or select a VLAN to contain traffic. Select **None** to disable access control for this rule.

CoS

Class of service action associated with the rule. Double-clicking on this column allows you change the setting.

System Log

Displays whether the syslog functionality (a syslog message is generated when the rule is used) is enabled, disabled, or prohibited for the rule. Double-clicking on this column allows you change the setting.

- **Enabled** If this option is enabled, a syslog message is generated when the rule is used. This option must be enabled if you are configuring Policy Rule Hit Reporting on your devices.
- **Disabled** If this option is disabled and this rule is hit, it does not generate a Syslog message, but lower-precedence rules and the role default actions may still specify a syslog message be sent for this data packet if there is a match.
- **Prohibited** If this rule is hit, no syslog message is generated for this data packet, even when a lower-precedence rule or the role default actions has the System Log action set to enabled.

Audit Trap

Displays whether the audit trap functionality (an audit trap is generated when the rule is used) is enabled, disabled, or prohibited for the rule. Double-clicking on this column allows you change the setting.

- Enabled If this option is enabled, an audit trap is generated when the rule is used.
- **Disabled** If this option is disabled and this rule is hit, it does not generate an audit trap, but lower-precedence rules and the role default actions may still specify generating an audit trap for this data packet if there is a match.
- **Prohibited** If this rule is hit, no audit trap is generated for this data packet, even when a lower-precedence rule or the role default actions has the Audit Trap action set to enabled.

Disable Port

Displays whether the disable port functionality (ports reported as using this rule will be disabled) is enabled, disabled, or prohibited for the rule. Double-clicking on this column allows you change the setting.

- Enabled If this option is enabled, any port reported as using this rule are disabled.
- **Disabled** If this option is disabled and this rule is hit, it does not disable the port, but lower-precedence rules and the role default actions may still specify disabling the port for this data packet if there is a match.

• **Prohibited** - If this rule is hit, the port is not disabled, even when a lower-precedence rule or the role default actions has the Disable Port action set to enabled.

Traffic Mirror

Displays whether the traffic mirror functionality is enabled, disabled, or prohibited for the rule. Double-clicking on this column allows you change the setting.

- Select port group(s) Use the drop-down list to specify the port groups where mirrored traffic will be sent for monitoring and analysis.
- **Disabled** If this option is disabled and this rule is hit, traffic mirroring will not take place, but lower-precedence rules and the role default actions may still specify traffic mirroring for this data packet if there is a match.
- **Prohibited** If this rule is hit, traffic mirroring is disabled, even when a lower-precedence rule or the role default actions has the Traffic Mirror action specified.

TCI Overwrite

Displays whether TCI Overwrite is enabled, disabled, or prohibited for the rule. Double-clicking on this column allows you change the setting.

- Enabled Enabling TCI Overwrite allows the VLAN (access control) and class of service characteristics defined in this rule to overwrite the VLAN or class of service (CoS) tag in a received packet, if that packet has already been tagged with VLAN or CoS information.
- **Disabled** If this option is disabled the TCI Overwrite option is ignored, but lower-precedence rules and the role default actions may still specify TCI Overwrite for the data packet if there is a match.
- **Prohibited** Do not set TCI Overwrite for this data packet, even when a lower-precedence rule or the role default actions has the TCI Overwrite option set to enabled.

Quarantine Role

Displays whether a Quarantine role is enabled, disabled, or prohibited for the rule. Double-clicking on this column allows you change the setting.

- Select Role Use the drop-down list to select the role that you want to assign as a Quarantine role.
- **Disabled** If this option is disabled and this rule is hit, a Quarantine role will not be assigned, but lower-precedence rules may still specify a Quarantine role for this data packet if there is a match.
- **Prohibited** If this rule is hit, a Quarantine role will not be assigned, even when a lower-precedence rule has a Quarantine role action specified.

_

Details View (Service)

Service Repository

Selecting Service Repository in the Roles/Services navigation panel in the left panel opens the Service Repository panel.

Double-click Local Services to display the service groups and services associated with the current domain or Global Services (All Domains) to display the service groups and services available to all domains.



Name

Displays the Local or Global service groups and services.

Local/Global Services

Selecting Local Services or Global Services (All Domains) in the Roles/Services > Service Repository navigation panel in the left panel opens the Local Services or Global Services (All Domains) panel, respectively.

Double-click Service Groups to display the services that are part of a service group or Services to view services not contained within a service group.



Name




Double-click one of the options to display the Service Groups or Services.

Details View (Services)

This tab lists the Automated and Manual services you create in the **Policy** tab. To display the tab, expand the **Local Services** or **Global Services** left-panel tab in the **Roles/Services > Service Repository** tab, and select the **Services** tab. To see a menu of options available for a service, right-click the service.

For information on the differences between automated or manual services, and local or global services, see the Policy tab Concepts Help topic's section on Services.



Name

Name of the service.

Number of Rules

Number of rules associated with the service.

Included in Roles Directly (Indirectly)

Number of roles in which the service is included.

Parent Service Group

The service group in which the service is included.



Details View (Service Group)

This tab lists information about the services or service groups contained in a **Local** or **Global** service group. To display this tab, select a service group in the left-panel **Roles/Services > Service Repository** tab.



Name

The name of the service or service group.

Number of Rules

The number of rules included in the service or service group.

Included in Roles Directly (Indirectly)

The number of roles where the service or service group exists directly in the role's Services list (as viewed on the role's **General** tab). If a service group also exists indirectly in other roles as part of another service group, that number of roles is displayed in parenthesis. In the example above, the service group called "Authenticated Supplemental Privileges" displays "1 (1)" in this column, showing that it is associated directly with one role (exists in that role's services list) and is also part of a service group associated with one other role.

Parent Service Group(s)

Displays all the "parent" service groups to which the service or service group belongs. This gives you an idea of the service group hierarchy without having to expand the left-panel tree.


• How to Create a Service



Add/Remove Services (Service Groups)

You can add and remove services from service groups using the Add/Remove Services window.

To access the Add/Remove Services window, either select the **Service Groups** tab in the **Local Services** or **Global Services** left-panel tab, right-click on a service group in the right panel and select **Add/Remove Services**. You can also right-click on a service group in the **Service Groups** left-panel tab and select **Add/Remove Services** from the menu.



All Services & Service Groups

This list displays all the local or global services and service groups in the current domain, depending whether you launched the window with a local or global service group selected. Select the services you want to add to the service group.

Selected Services & Service Groups

This list displays all the services currently defined for the selected service group. Select the services you want to remove from the service group.

Right Arrow Button

Select the **Right Arrow** button to add the services selected in the All Services & Service Groups list to the Selected Services & Service Groups list.

Left Arrow Button

Select the **Left Arrow** button to remove the services selected in the Selected Services & Service Groups list


Double Left Arrow Button

Select the **Double Left Arrow** button to remove all the services from the Selected Services & Service Groups list.

Rule

The rule **General** tab displays general information about the rule selected for a Service in the left-panel **Roles/Services > Service Repository > Local or Global Services** tab and enables you to change it. In addition, you can view and change the Traffic Description and Actions associated with the rule. Traffic Description identifies the type of traffic to which the rule pertains. Actions apply class of service, access control, and/or accounting and security behavior to packets matching the rule.

Any additions or changes you make to this tab must be enforced in order to take effect. If you modify an enabled rule's actions, the Policy tab checks for conflicts with other rules in the services and roles with which the newly modified rule is associated. See Conflict Checking for more information.



General Area

Service Name

Displays the name of the rule.

Description

Use the Edit button to open a window where you can enter or modify a description of the rule.

Rule Status

Lets you disable the rule, or enable it if it's already disabled. If the rule is disabled, it is unavailable for use by the current service, but can still be copied to other services and enabled, or re-enabled at another time for the current service. Disabling a rule is an alternative to deleting and recreating it. The rule icon in the left panel displays a red X if the rule is disabled.

Rule Type

Use the drop-down list to select the types of devices to which you wish this rule to apply when enforced. The recommended selection is All Devices, unless there is a specific need for a device-specific rule. If this need arises, the Rule Type feature enables services to be customized to contain rules specific to a device's type when support for a traffic description and/or action is not be available on all managed devices.

For device-specific rules, only those traffic descriptions supported on the device are available when you define the rule's traffic description on this tab. For All Devices rules, all traffic descriptions are available; however, you must be aware that you cannot enforce the rule to a device on which it is not supported.

TCI Overwrite

Specify the TCI Overwrite functionality for the rule:

- Enabled Enabling TCI Overwrite enables the VLAN (access control) and class of service characteristics defined in this rule to overwrite the VLAN or class of service (CoS) tag in a received packet, if that packet has already been tagged with VLAN or CoS information.
- **Disabled** If this option is disabled the TCI Overwrite option is ignored, but lower-precedence rules and the role default actions can still specify TCI Overwrite for the data packet if there is a match.
- **Prohibited** Do not set TCI Overwrite for this data packet, even when a lower-precedence rule or the role default actions has the TCI Overwrite option set to enabled.

Traffic Description Area

The Traffic Description area enables you to view and change the traffic description associated with a rule. The Traffic Description identifies the traffic classification type for the rule. Rules enable you to assign access control (VLAN membership) and/or class of service to network traffic depending on the traffic's classification type.

Type

Displays the Classification Type selected for the rule.

Value

Displays the values/parameters selected for the rule's Classification Type. See Classification Types and their Parameters for parameter information.

Remove Button

Removes the traffic description from the rule.

Edit Button

If a Traffic Description Type has been defined for the rule, selecting Edit opens the Edit Rule window, where you can edit the parameters or values for the rule's classification type.

Actions Area

The Actions area enables you to view and change the actions associated with a rule. Actions apply access control, class of service, security, and/or accounting behavior to packets matching the rule.

Access Control

Use this drop-down list to select the appropriate access control for the rule. You can permit traffic to be forwarded, deny traffic altogether, or contain traffic to a VLAN. Select **None** to disable access control for this rule.

- **Permit Traffic** enables traffic to be forwarded with the port's assigned VID.
- **Deny Traffic** traffic will be automatically discarded.
- Contain to VLAN contains traffic to a specific VLAN. Use the drop-down list to select the desired VLAN.

Class of Service

Use the drop-down list to select a class of service to associate with the rule. The Policy tab lets you define classes of service that each include an 802.1p priority, and optionally an IP type of service (ToS/DSCP) value, rate limits, and transmit queue configuration. You can then assign a class of service as a classification rule action. See Getting Started with Class of Service and How to Create a Class of Service for more information. Select **None** to disable class of service for this rule.

When rule accounting is enabled on a device, each rule keeps a list of the ports on which it has been used. Use the following three options to specify certain rule usage actions to take place when a "rule hit" is reported.

System Log

Specify System Log functionality for the rule.

- Enabled If this option is enabled, a syslog message is generated when the rule is used. This option must be enabled if you are configuring Policy Rule Hit Reporting on your devices.
- **Disabled** If this option is disabled and this rule is hit, it does not generate a Syslog message, but lower-precedence rules and the role default actions can still specify a syslog message be sent for this data packet if there is a match.
- **Prohibited** If this rule is hit, no syslog message is generated for this data packet, even when a lower-precedence rule or the role default actions has the System Log action set to enabled.

Audit Trap

Specify Audit Trap functionality for the rule:

• Enabled — If this option is enabled, an audit trap is generated when the rule is used.

- **Disabled** If this option is disabled and this rule is hit, it does not generate an audit trap, but lower-precedence rules and the role default actions can still specify generating an audit trap for this data packet if there is a match.
- **Prohibited** If this rule is hit, no audit trap is generated for this data packet, even when a lower-precedence rule or the role default actions has the Audit Trap action set to enabled.

Disable Port

Specify Disable Port functionality for the rule:

- Enabled If this option is enabled, any port reported as using this rule will be disabled. Ports that have been disabled due to this option are displayed in the device Role/Rule tab.
- **Disabled** If this option is disabled and this rule is hit, it does not disable the port, but lower-precedence rules and the role default actions can still specify disabling the port for this data packet if there is a match.
- **Prohibited** If this rule is hit, the port is not disabled, even when a lower-precedence rule or the role default actions has the Disable Port action set to enabled.

Traffic Mirror

Specify traffic mirroring functionality for the rule:

- Select port group(s) Use the drop-down list to specify the port groups where mirrored traffic will be sent for monitoring and analysis. Select View/Modify Port Groups to open the Port Groups tab where you can define user-defined port groups for selection.

 To the right of the drop-down list is an option to mirror only the first (N) packets of a flow. This option is intended for use when mirroring traffic to an ExtremeAnalytics engine. The ExtremeAnalytics engine only needs the initial packets of a flow to properly identify the traffic, and setting this option will reduce network traffic overhead for the switch and engine. By default this number is set to 10, but can be changed by selecting the Edit button. Note that the value you set is used by all mirror actions in use in the current domain.
- **Disabled** If this option is disabled and this rule is hit, traffic mirroring will not take place, but lower-precedence rules and the role default actions can still specify traffic mirroring for this data packet if there is a match.
- **Prohibited** If this rule is hit, traffic mirroring is disabled, even when a lower-precedence rule or the role default actions has the Traffic Mirror action specified.

Quarantine Role


Specify the Quarantine Role functionality for the rule:

- Select Role Use the drop-down list to select the role that you want to assign as a Quarantine role. Specifying a role as a Quarantine role turns the role's icon red, denoting its restrictive nature.
- **Disabled** If this option is disabled and this rule is hit, a Quarantine role will not be assigned, but lower-precedence rules can still specify a Quarantine role for this data packet if there is a match.
- **Prohibited** If this rule is hit, a Quarantine role will not be assigned, even when a lower-precedence rule has a Quarantine role action specified.

146 of 685

Create Rule

This window displays when you right-click a service group or the **Services** tab in the left-panel and select **Create Rule**. If you use this window, traffic descriptions and actions can be added to the rule afterwards (see Using the Rule Tabs). In order for a rule to be applied to devices, you must enforce.



Name

Enter a name for the rule.

Type

Select the types of devices to which you wish this rule to apply when enforced. See Rule Type for more information on the consequences of your choice.

OK

Select **OK** to create the rule and close the **Create Rule** window.

Apply

Select Apply to create the rule and remain in the Create Rule window.

Cancel

Select **Cancel** to close the **Create Rule** window without saving your changes.


Edit Rule

The Edit Rule window allows you to change the traffic description associated with a rule. The Traffic Description, which includes the traffic classification layer, traffic classification type, and traffic value, was entered when the rule was created (see How to Create or Modify a Rule).

To display the Edit Rule window, select the rule in the left panel's **Services** tab. In the Traffic Description section, select **Edit** to bring up the Edit Rule window.

If you modify an enabled rule's traffic descriptions, the **Policy** tab checks for conflicts with other rules in the services and roles with which the newly modified rule is associated. See Conflict Checking for more information.

The contents of the Edit Rule window varies according to the selected rule and traffic description.



Layer Area

Traffic Classification Layer

The OSI model classification layer (or All Layers) currently associated with the rule. Each layer has multiple classification types from which you can select. If you change the layer, the Type and Value sections in the window change, and you must make new selections in those sections. See Classification Types and their Parameters for information.

Traffic Classification Type

The traffic classification type currently associated with the rule. Each classification type consists of certain parameters and/or values. If you change the type, the Value section of the window changes, and you must make new selections in that section. See Classification Types and their Parameters for information.

Value Area


This area displays the values currently selected for the traffic classification type, and allows you to change those values. Each traffic classification type requires certain parameters and/or values. See Classification Types and their Parameters for parameter information.

Class of Service Overview

Use this tab to view the Class of Service (CoS) configuration for the current domain. To access this window, select the **Class of Service** left-panel tab from the **Policy** tab.

This window displays the eight pre-populated static classes of service, each associated with one of the 802.1p priorities (0-7). Use these predefined classes of service or create your own classes of service.

Expanding this tab in the left panel allows you to select individual classes of service in the right panel, which opens them in the Class of Service tab, where you can edit the configuration for the selected CoS.



Name

The name of the class of service.

Index

The index number automatically assigned to the class of service.

Priority

The 802.1p priority associated with the class of service. The priority for the eight static classes of service provided by the Policy tab (Priority 0-7), cannot be disabled or changed.

ToS

The IP type of service value associated with this class of service, if any. See IP Type of Service for more information.

Drop Precedence

The drop precedence associated with this class of service. Double-click in the column to select a Drop Precedence value: Low, Medium, or High.

Getting Started with Class of Service

This Help topic provides an overview of **Policy** tab's class of service (CoS) functionality, including information about defining rate limits and configuring transmit queues.

After you have read this topic, look at an example of how a network administrator might use CoS to configure VoIP traffic with appropriate priority, ToS, queue treatment, and flood control by selecting the link: Class of Service Example.

This guide includes the following information:

- Class of Service Overview
- Rate Limits
- Transmit Queues
- Flood Control


Class of Service Overview

Class of Service (CoS) provides the ability to give certain network traffic preferential treatment over other traffic. It classifies traffic into categories such as high, medium, and low, where high-priority traffic gets the best service while low-priority traffic is "drop eligible."

Class of Service helps you manage the bandwidth requirements of a given network flow with the available port resources on your network devices. (In a CoS context, a flow is a stream of packets classified with the same class of service as the packets transit the interface). Using CoS, you can:

- Assign different priority levels to different packet flows.
- Mark or re-mark the packet priority at port ingress with a Type of Service (ToS).
- Sort flows by transit queue. Higher priority queues get preferential access to bandwidth during packet forwarding.
- Limit the amount of bandwidth available to a given flow by either dropping (rate limiting) or buffering (rate shaping) packets in excess of configured limits.

The following figure shows how you can manage network bandwidth requirements by assigning different classes of service to different types of network traffic.



The ICMP protocol, used for error messaging, has a low bandwidth requirement, with a high tolerance for delay and jitter, and is appropriate for a low priority setting. HTTP and FTP protocols, used respectively for browser-generated and file transfer traffic, have a medium to high bandwidth requirement, with a medium to high tolerance for delay and jitter, and are appropriate for a medium priority level. Voice (VoIP), used for voice calls, has a low bandwidth requirement, but is very sensitive to delay and jitter and is appropriate for a high priority level.

Implementing CoS

CoS determines how a given network flow is assigned bandwidth as it transits your network devices. As a preliminary step to using CoS, it is important that you understand the characteristics of the flows on your network and associate these flows with your policy roles. In this sense, CoS is the third step in a three step process:

- 1. Understand your network flows using NetFlow.
- 2. Associate your network flows with a **Policy** tab role.
- 3. Configure your classes of service and associate them with the rules contained in your roles.

Configuring CoS

The **Policy** tab lets you configure multiple classes of service that include one or more of the following components:

- 802.1p priority
- IP type of service (ToS) value
- drop precedence

- inbound and outbound rate limits
- outbound rate shaper per transmit queue.
- flood control rate limits

After you have created and defined your classes of service, they are then available when you make a class of service selection for a rule action (**Rule** tab), a role default (**General** tab), or an automated service (**Automated Service** tab).

To view and configure CoS, open the **Class of Service Overview** tab from the **Policy** tab. It is prepopulated with eight static classes of service, each associated with one of the 802.1p priorities (0-7). You can use these classes of service as is, or configure them to include ToS, drop precedence, rate limit, and/or transmit queue values. In addition, you can also create your own classes of service (user-defined CoS).

Rate Limits


Rate limits are one component of a **Policy** tab class of service. They control the transmit rate at which traffic enters and exits ports in your network. All traffic mapped to a Class of Service on a given port share the bandwidth specified by the rate limit.

For instructions on how to configure rate limits, see How to Define Rate Limits.

Rate limits are tied directly to roles and rules, and are written to a device when the role/rule is enforced. When rate limits are implemented, all traffic on the port that matches the rule with the associated rate limit cannot exceed the configured limit. If the rate exceeds the configured limit, frames are dropped until the rate falls below the limit.

The rate limit remains on the port only as long as the role using the rate limit is active on the port either as the authenticated role or as the port's default role.

The following figure shows how bursty traffic is clipped above the assigned threshold when rate limiting is applied.



The CoS can be configured to perform one or all of the following actions when a rate limit is exceeded:

- Generate System Log on Rate Violation a syslog message is generated when the rate limit is first exceeded.
- Generate Audit Trap on Rate Violation an audit trap is generated when the rate limit is first exceeded.
- Disable Port on Rate Violation the port is disabled when the rate limit is first exceeded.

The **Policy** tab class of service also provides the ability to create rate limit port groups. Port groups let you specify different rate limits within the same class of service. For example, you might create a port group for edge ports and a port group for core ports, and assign two different rate limits. For more information on rate limit port groups, see Creating Class of Service Port Groups.

Transmit Queues

Transmit queue configuration is defined within a class of service and associated with a specific role via a rule action or as a role default. It is implemented based on the role assigned to a port. All traffic received on a port and matching a rule with the associated class of service is forwarded using the defined transmit queue configuration.


For instructions on how to configure transmit queues, see How to Configure Transmit Queues.

There are three components to transmit queue configuration:

- Transmit Queue Configuration enables you to set the transmit queue associated with the class of service.
- Transmit Queue Rate Shapers let you pace the rate at which traffic is transmitted out of that transmit queue.
- Bandwidth Configuration enables you to specify how the traffic in each transmit queue is serviced as it egresses the port.

The transmit queue configuration remains on the port only as long as the role using the configuration is active on the port either as the authenticated role or as the port's default role.

The following figure shows how bursty traffic is smoothed out when it goes above the assigned threshold when rate shaping is applied.



Rate shaping retains excess packets in a queue and then schedules these packets for later transmission over time. Therefore, the packet output rate is smoothed and bursts in transmission are not propagated as seen with rate limiting.

Rate shaping can be used for the following reasons:

- to control bandwidth
- to offer differing levels of service
- to avoid traffic congestion on other network links by removing the bursty property of traffic that can lead to discarded packets

The **Policy** tab class of service also provides the ability to create transmit queue shaper port groups that enable you to isolate certain kinds of sensitive network traffic so that you can vary the bandwidth of the shape for that single queue. For more information on transmit queue port groups, see Creating Class of Service Port Groups.

Flood Control

Flood control provides rate limiting capabilities to individual Class of Service to permit certain types of flooded traffic to be dropped. When enabled, incoming traffic is monitored over one second intervals. Traffic is identified using the following configuration types:

- unknown unicast
- broadcast
- multicast

A traffic control rate sets the acceptable flow for each type, specified in packets per second. If, during a one second interval, the incoming traffic of a configured type reaches the traffic control rate on the port, the traffic is dropped until the interval ends. Packets are then permitted to flow again until the limit is reached.

By default, Flood Control is disabled for each CoS. Similar to CoS Port Groups, a different configuration can be assigned for each group. Since Flood Control is shared across all CoS, when Flood Control is enabled on at least one CoS, those rates apply to all ports that have Flood Control enabled.

For instructions on how to configure flood controls, see How to Configure Flood Control.




Class of Service

This tab lets you view and configure the components of a class of service (CoS). See below for a description of each section. For more information, see How to Create a Class of Service.

Once you have created and defined a class of service, you can then apply it as a classification rule action, as part of the definition of an automated service, or as a role default. For more information, see Getting Started with Class of Service.

To access this tab, select the **Class of Service** left-panel tab on the **Policy** tab. Select a class of service in the tree, and the information for the selected class of service displays in the right panel.



General

Name

Name of the selected class of service.

Description

Use the **Edit** button to open a window where you can add or modify a description for the class of service.

Transmit Queue

This field displays the transmit queue associated with the class of service for each port type. Use the **Edit** button to display a menu where you can select a new transmit queue, if desired.

802.1p Priority

This drop-down list lets you select the 802.1p priority associated with the class of service, if desired. This field is grayed out for the eight static classes of service provided by the Policy tab (Priority 0-7), because the 802.1p priority cannot be disabled or changed.

ToS

Some IP rules enable a ToS value to be written to the ToS field in the IP header of incoming packets. Select the **Edit** button to open the Edit ToS window, where you can enter a ToS value. The value must be an 8-bit hexadecimal number between 0 and FF (see IP Type of Service for more information).

Drop Precedence

The Drop Precedence option is used in conjunction with the Flex-Edge feature available on K-Series and S-Series (Release 7.11 or higher) devices. Flex-Edge provides the unique capability to prioritize traffic in the MAC chip as it enters the switch. When the Class of Service is assigned to a policy role, and that role is applied to a port via a MAC source address mapping or the port default role, the drop precedence dictates the internal priority (within the MAC chip) used for packets received on the port. If congestion occurs, packets with a high drop precedence are discarded first. Therefore, if a packet is important, it should have a low drop precedence. Refer to the K-Series or S-Series Configuration Guide for more information on the Flex-Edge feature and drop precedence.

Rate Limiting/Rate Shaping

This section displays the inbound/outbound rate limits (IRL/ORL) and the outbound transmit queue (TxQ) rate shapers that are configured for the Default port groups associated with the class of service. If you have created additional port groups, the information displays for those groups as well.

With port rate limits, all traffic assigned to this class of service on a given port shares bandwidth specified by the rate limit. Rate shaping paces the rate at which traffic is transmitted out of the transmit queue. You can add or change a rate limit or a rate shaper by double-clicking on the area below a port group name.

If you have ExtremeWireless Controllers (Release 8.01.xx or higher) on your network, you also see the IRL and ORL user rate limits associated with the class of service. User rate limits specify the bandwidth given to each individual user on a port. Currently, user rate limits are only available on wireless controllers.

For more information, see Advanced Rate Limiting by Port Type and How to Configure Transmit Queues.

Index Numbers

At the bottom of the tab there is a section for configuring the rate limit and transmit queue index numbers associated with this class of service. These index numbers are used to map the class of service to the actual rate limits and transmit queue configuration on the device.

Typically, each class of service uses a different index number. The Policy tab automatically assigns these index numbers when you configure a class of services' rate limits and transmit

queue shapers. An index number of "-1" indicates that no mappings are associated with the class of service.

All CoS using the same index will use the same rate limit and rate shaping assignments, and thus all traffic using those CoS will share the bandwidth.

IRL/ORL Index (Inbound/Outbound Rate Limits Index)

The inbound/outbound port rate limit index associated with the class of service. Index numbers map logical rate limit indexes to the actual physical rate limits you have created in the Policy tab. Select the button to open the Rate Limits selection view window, and select an index for the CoS. For convenience, existing index to rate limit mappings are displayed; if one of the existing indexes is selected, the displayed mappings will apply for this CoS. (Selecting an index highlights all the mappings configured for that index number within the selection view.)

TxQ Index (Transmit Queue Index)

The transmit queue index associated with the class of service. Index numbers map logical transmit queue indexes on the ports to the actual physical transmit queues you have configured in the **Policy** tab. If you have selected an 802.1p priority for this class of service, a default transmit queue index is automatically specified based on the selected priority. You can use the default index or change it according to your own transmit queue configuration. Select the button to open the Transmit Queues selection view window, which lists all the possible transmit queues, organized by index number for each existing port type and group. Selecting an index automatically includes all the transmit queues configured for that index number.

IUB/OUB Index (Inbound/Outbound User-Based Rates Index)

If you have ExtremeWireless Controllers (Release 8.01.xx or higher) on your network, you also see the inbound/outbound user rate limits associated with the class of service. User rate limits specify the bandwidth given to each individual user on a port. Currently, user rate limits are only available for these wireless controllers. Select the button to open the Rate Limits selection view window, and select an index for the CoS. For convenience, existing index to rate limit mappings are displayed; if one of the existing indexes is selected, the displayed mappings apply for this CoS. (Selecting an index highlights all the mappings configured for that index number within the selection view.)

Flood Ctrl Port Groups


CoS-based flood control is a form of rate limiting that prevents configured ports from being disrupted by a traffic storm, by rate limiting specific types of packets through those ports. When flood control is enabled on a port, incoming traffic is monitored over one second intervals. During an interval, the incoming traffic rate for each configured traffic type (unknown-unicast, broadcast, or multicast) is compared with the configured traffic flood control rate, specified in packets per second. If, during a one second interval, the incoming traffic of a configured type reaches the traffic flood control rate configured on the port, CoS-based flood control drops the traffic until the interval ends. Packets are then permitted to flow again until the limit is again reached.

NOTE: By default, Flood Control is not managed by the **Policy** tab. To manage flood control configuration on devices in a domain, it can be enabled via the Domain Managed CoS Components drop-down list by selecting All CoS Components or by selecting Flood Control.

General (CoS Components Folder)

This tab lists the elements that comprise a class of service. It displays when you select the **CoS Components** tab in the **Class of Service** left-panel tab of the **Policy** tab.

See Getting Started with Class of Service for more information about these components.



Inbound Rate Limits

Select this checkbox to enable the **Inbound Rate Limit Port Groups tab** in the **CoS Components** left-panel tab.

Inbound Rate Limits

Select this checkbox to enable the **Outbound Rate Limit Port Groups tab** in the **CoS Components** left-panel tab.

Transmit Queue/TxQ Shapers

Select this checkbox to enable the **Transmit Queue Port Groups tab** in the **CoS Components** left-panel tab.

Flood Control Port Groups

Select this checkbox to enable the Flood Control Port Groups tab in the CoS Components left-panel tab.


General (Rate Limits)

This tab allows you to create and define a rate limit. Rate limits are components of a class of service and are used to control the transmit rate at which traffic enters and exits ports in your network.

To access this window, open the **Control** tab, select the **Policy** tab > **Class of Service** left-panel tab > **CoS Components** left-panel tab > **Rate Limits** tab. Select an existing rate limit to view or

modify a rate limit or right-click the **Rate Limits** left-panel tab and select the **Create Rate Limit** option to create a new rate limit.

To create the rate limit, fill out the window and select **OK** (to create a single rate limit) or **Apply** (to create more rate limits). After you create the rate limit, the General tab for the new rate limit displays, where you can configure additional rate limit parameters.



Name

Specify the name of the rate limit.

Rate Limit

Select the **Edit** button to specify the highest transmission rate at which traffic can enter or exit a port before packets are rate limited:

- % A percentage of the total bandwidth available (not available for priority-based rate limits)
- PPS Packets per second (not available for priority-based rate limits)
- Kb/s Kilobits per second
- Mb/s Megabits per second
- Gb/s Gigabits per second

Actions

Select the action(s) you would like this rate limit to use:


- System Log a syslog message is generated when the rate limit is first exceeded.
- Audit Trap an audit trap is generated when the rate limit is first exceeded.
- Disable Port the port is disabled when the rate limit is first exceeded.

NOTE: N-Series Gold devices do not support rate limit notification.

Details View (Rate Limits Folder)

This tab lists information on any rate limits that have been defined in the **Policy** tab.

To access this tab, select the **Class of Service > CoS Components > Rate Limits** left-panel tab. See How to Define Rate Limits for more information.



Name

Name of the rate limit.

Syslog

Specifies whether a syslog message will be generated when the rate limit is first exceeded.

Audit Trap

Specifies whether an audit trap will be generated when the rate limit is first exceeded.

Disable Port

Specifies whether the port will be disabled when the rate limit is first exceeded.

Priority-Based Rate Limits

Priority-based rate limits are used primarily by legacy devices. They are rate limits that are associated with one or more of the eight 802.1p priorities (0-7). When the associated priority is selected for a class of service, the rate limit becomes part of that class of service.

These rate limits are written directly to each port (unless the port is specified in the rate limit's exclusion list), and are implemented based on the 802.1p priority assigned to a data packet appearing on that port. While priority-based rate limits are not tied directly to roles or rules, they are displayed with the associated priority when you select a class of service while creating a rule, automated service, or role.

When priority-based rate limiting is implemented, the combined rate of all traffic on the port that matches the priorities associated with the rate limit cannot exceed the configured limit. If the rate exceeds the configured limit, frames are dropped until the rate falls below the limit.

When a rate limit is associated with a priority, that priority includes rate limiting wherever and however it is used, until the rate limit is deleted from ExtremeCloud IQ - Site Engine. Also, when a priority-based rate limit is applied to a port, it remains on the port even if the role that originally used the rate limit is no longer associated with the port. For example, if an untagged packet arrives on a port where there is no role or default priority, but the port's 802.1p priority includes a rate limit, that traffic is rate limited. As another example, if the priority of a tagged packet matches a priority-based rate limit on a port, the traffic is rate limited.

To configure a priority-based rate limit, you need to specify the following components:


- Rate Limit The highest transmission rate at which traffic can enter or exit a port.
- *Direction* The direction to which the limit applies (inbound or outbound traffic). In order to control traffic inbound and outbound on the same port, two rate limits must be configured (one inbound and one outbound). Inbound rate limiting takes place after a frame is classified into one of the eight priorities. Outbound rate limiting takes place just before a frame is queued for transmission. A single frame can pass through inbound and outbound rate limits depending on the path it takes through the device and the rate limiting configuration on the device.
- Priority The 802.1p priority or priorities with which the rate limit is associated.
- Precedence The order in which the rate limit is written to supported devices. ExtremeCloud IQ Site Engine allows you to define as many rate limits as you wish; however, the number written to a device is restricted by the number of rate limits supported by the device. Each port on the device can utilize any or all of the defined rate limits up to the number of rate limits it supports.
- Exclusion The devices/ports you wish to be excluded from the rate limit. For example, rate limiting is most often used for edge devices; therefore, you might want to exclude a device group or port group containing non-edge devices or ports.

Add/Edit CoS to Rate Limit Mapping

This window lets you configure the rate limit mappings for a rate limit port group. Rate limit mappings map a logical rate limit index to an actual physical rate limit you have created in ExtremeCloud IQ - Site Engine.

For reference, the CoS IRL/ORL Index table (at the bottom of the window) displays classes of service that already have an IRL/ORL index specified, so that you can see which classes of service are affected by mapping an index to a rate limit.

To access this window, open the select the Add/Edit button on the CoS - Rate Limit Mappings tab (Control tab > Policy tab > Class of Service left-panel tab > CoS Components left-panel tab and select a port group in either the Inbound Rate Limit Port Groups or Outbound Rate Limit Port Groups left-panel tab, depending on the type of rate limit.



IRL/ORL Index

Specify the IRL (Inbound Rate Limit) or ORL (Outbound Rate Limit) Index you are mapping.

Rate Limit

Use the drop-down list to select a rate limit to map to the index. Rate limits are listed by the rate limit name followed by the precedence. For information on how to create a rate limit, see How to Define Rate Limits. Select **None** to remove an existing mapping for the specified port types.

Port Types

These options allow you to create a mapping for all port types, or create a mapping just for specific port types.

Advanced Rate Limiting by Port Type

The **Policy** tab class of service feature provides the ability to create rate limit port groups that let you group together ports with similar rate limiting requirements. For instructions on creating a port group, see Creating Class of Service Port Groups.

This Help topic provides information about an advanced port group feature that lets you specify different rate limits for the different port types contained in a port group: 8-rate limit, 32-rate limit, 64-rate limit, and 100-rate limit port types.

After you have created your port groups, you can use the CoS to rate limit mappings tab to configure rate limit index mappings for each group. These mappings map a logical rate limit index to an actual physical rate limit created in the Policy tab. For each class of service, you can select one mapping index that gives you the desired physical rate limit for each port group (see the Index Numbers section of the CoS General tab for more information on CoS Index Numbers).

The **Policy** tab supports a maximum of 100 logical rate limit indexes and each rate limit port group lets you map all 100 indexes. For 8-rate limit, 32-rate limit, and 64-rate limit ports, this means that the number of logical indexes might be greater than the actual number of rate limits the port supports. The port group can map 100 logical rate limit indexes, but they can only be mapped to a maximum of 8, 32, or 64 different physical rate limits on those ports.

For example, you want to have 25 rate limits for 25 different CoS. You need to define the behavior for the 8-rate port type, since when you get to the 9th rate, you would have no more resources available for the remaining rates (9-25). You would either need to share some of the same resources, or not rate limit with the remaining rates.

The maximum supported indexes for a device is based on the largest number of rates supported for that device. On devices supporting a maximum of 8 rate limits, indexes 0-7 are supported. On devices supporting a maximum of 32 rate limits, indexes 0-31 are supported. On devices supporting 64 rate limits, IRL indexes 0-63 are supported. If a rate limit port group maps indexes greater than the supported value, they are ignored during Enforce (indicated in the Class of Service > Rate Limit Mappings tables of Enforce Preview)

Instructions on:

- Configuring Rate Limit Mappings
- Associating Rate Limits with a Class of Service

Configuring Rate Limit Mappings

Use the following instructions configure rate limit mappings for a port group.

- 1. Open the Class of Service > CoS Components left-panel tab.
- 2. Select either the Inbound Rate Limit Port Groups or Outbound Rate Limit Port Groups left-panel tab.
- 3. Select the right-panel CoS Rate Limit Mappings tab.
- 4. Select Add/Edit to open the Add/Edit CoS to Rate Limit Mappings window.
- 5. In the window, specify the IRL (Inbound Rate Limit) or ORL (Outbound Rate Limit) Index you are mapping.
- 6. Use the drop-down list to select a rate limit to map to the index.

- 7. The port type options enable you to create a mapping for all port types at one time, or create a mapping just for specific port types.
- 8. Select the **OK** button to map all your indexes and close the window. The Mappings tab displays your index to rate limit mapping configuration.

Associating Rate Limits with a Class of Service


After you have configured the rate limit mappings for a port group, you can associate a rate limit mapping index with a class of service.

- 1. Open the Class of Service left-panel tab.
- 2. Select the CoS in the left-panel tree. (If you have not created the class of service, see How to Create a Class of Service.)
- 3. At the bottom of the **Class of Service** tab in the right panel, select the **Edit** button next to the IRL or ORL index that you want to configure. The Edit Index window opens.
- 4. This window lists all the currently mapped rate limits, organized by index number for each existing port type and group. Selecting one index number automatically includes all the rate limits configured for that index number. To configure new mappings for the CoS, you can first select an index that is not currently mapped, then create the mappings as described in Configuring Rate Limit Mappings above. Select **OK**.
- 5. After you have selected the mapping index, the table below displays the actual rate limits used by each rate limit port group for that class of service.
- 6. Select Open/Manage Domains > Save Domain.

Summary (Rate Limit Port Groups Folder)

This tab lists the name of all the inbound or outbound rate limit port groups (depending on the left-panel tab you select). Rate limit mappings map a logical rate limit index (IRL/ORL Index) to an actual physical rate limit. You can configure a port group's mappings on the port group Mappings tab.

To access this tab, open the Class of Service > CoS Components left-panel tab, then, select either the Inbound Rate Limit Port Groups left-panel tab or the Outbound Rate Limit Port Groups tab. The Summary tab displays in the right panel.



Name

The name of the port group.

CoS - Rate Limit Mappings (Rate Limit Port Group)

This tab lets you view and configure the rate limit mappings for a rate limit port group. Rate limit mappings map a logical rate limit index used by classes of service to an actual physical rate limit you create in ExtremeCloud IQ - Site Engine.

Each port group has its own set of index mappings. ExtremeCloud IQ - Site Engine automatically assigns these index numbers when you configure a class of services' rate limits and transmit queue shapers.

The rate limit mappings tab allows you to do two things:


- Map the index to a different rate for different port groups (edge ports versus inter-switch links). See Creating Class of Service Port Groups.
- Map the index to a different rate limit for each port type (8-rate limit, 32-rate limit, 64-rate limit, and 100-rate limit) in a port group. See Advanced Rate Limiting by Port Type.

To access this tab:

- 1. Open the **Control** tab.
- 2. Open the **Policy** tab.
- 3. Open the Class of Service > CoS Components left-panel tab.
- 4. Select either the **Inbound Rate Limit Port Groups** or **Outbound Rate Limit Port Groups** left-panel tab, depending on whether the rate limit is inbound or outbound.
- 5. Select a existing port group in the left panel to open it in the Rate Limit Port Group tab.

NOTE: Create a new port group by right-clicking the **Inbound Rate Limit Port Groups** or **Outbound**Rate Limit Port Groups left-panel tab, selecting Create Port Group, entering a Name for the port group, and selecting **OK**.

6. Select the CoS - Rate Limit Mappings tab in the right panel.



IRL/ORL Index

The logical inbound rate limit (IRL) or outbound rate limit (ORL) index number. This index number is specified in a class of service and dictates the rate limiting behavior for incoming or outgoing packets. For each rate limit port group, use this tab to map the index number to an actual rate limit.

Rate Limit

The actual rate limit to which the IRL/ORL index is mapped.

IRL/ORL Port Type(s)

The type of ports included in the port group. Port type is based on the number of rate limits the ports support (for example, 8-rate limit ports and 32-rate limit ports).

IRL/ORL Index Used By CoS

The classes of service using this IRL/ORL index.

Add/Edit Button

Opens the Add/Edit CoS to Rate Limit Mappings window where you can add or edit rate limit mappings for the rate limit port group

Remove Button

Removes the mapping(s) selected in the table.

Ports (Rate Limit Port Group)

The rate limit port group **Ports** tab lets you view all the ports in the selected port group, as well as add and remove ports to and from the group. It provides information about each port, and lets you view and edit port information (via the port's **General** tab).


To access this tab:

- 1. Open the **Control** tab.
- 2. Open the Policy tab.
- 3. Open the Class of Service > CoS Components left-panel tab.
- 4. Select either the **Inbound Rate Limit Port Groups** or **Outbound Rate Limit Port Groups** left-panel tab, depending on whether the rate limit is inbound or outbound.
- 5. Select a existing port group in the left panel to open it in the Rate Limit Port Group tab.

NOTE: Create a new port group by right-clicking the **Inbound Rate Limit Port Groups** or **Outbound**Rate Limit Port Groups left-panel tab, selecting Create Port Group, entering a Name for the port group, and selecting **OK**.

6. Select the **Ports** tab in the right panel.

Create a new port group by right-clicking the **Inbound Rate Limit Port Groups** or **Outbound Rate Limit Port Groups** left-panel tab, selecting **Create Port Group**, entering a **Name** for the port group, and selecting **OK**.



Name

Name of the port, constructed of the name or IP address of the device and either the port index number or the port interface name.

Rate/Queue Port Type

The number of rate limits the port supports.

Default Role

The Default Role assigned to the port.

Alias

Shows the alias (ifAlias) for the interface, if one is assigned.

Stats

Shows statistics collected for a port, enabled via the Flow Collection & Interface setting in the PortView.

Port Type

Type of port. Possible values include: Access, Interswitch Backplane, Backplane, Interswitch, and Logical.

Neighbor

The port's neighbor port.

Port Speed

Speed of the port. Possible values include: 10/100, speed in megabits per second (for example, 800.0 Mbps), Unknown (displayed for logical ports).

Description

A description of the port.


Add/Remove Ports Button

Opens the <u>Add/Remove Ports window</u>, where you can add and remove ports to and from the port group. When you create new port groups, you add ports from the Default group into your newly defined port groups.



Automated Service

Selecting an Automated Service opens the **Automated Service** tab which enables you to define settings for the service. For more information on services, see How to Create a Service.



Service Name

Name of the selected service.

Description

Use the Edit button to open a window where you can enter or modify a description of the service.

TCI Overwrite

Specify the TCI Overwrite functionality for the service:

- Enabled Enabling TCI Overwrite enables the VLAN (access control) and class of service
 characteristics defined in this service to overwrite the VLAN or class of service (CoS) tag in a
 received packet, if that packet has already been tagged with VLAN or CoS information.
- **Disabled** If this option is disabled the TCI Overwrite option is ignored, but lower-precedence rules and the role default actions can still specify TCI Overwrite for the data packet if there is a match.
- **Prohibited** Do not set TCI Overwrite for this data packet, even when a lower-precedence rule or the role default actions has the TCI Overwrite option set to enabled.

Traffic Description Area

Use this area to provide the specifications for an automated service. Specify the network resource type, the network resources for the service, and the rule type. Some rule types require that you enter certain parameters and/or values. This section is not displayed for a Manual service.

Type

Select the **Edit** button to select the type of rule you want to create for the network resources. Some rule types require you enter certain parameters and/or values. See Classification Types and their Parameters for parameter information. Select and/or enter the required parameters.

Network Resource Type

Select the network resource type (Layer 2 MAC or Layer 3 IP). This will determine the list of network resources available for selection for this service.

Network Resources

Use the drop-down list to select the network resources to associate with the automated service. Use the configuration menu button to the right of the list to add a network resource or view and edit your network resources. For more information, see How to Create a Network Resource.

Actions Area

Use this area to define the access control and/or a class of service for the Automated service rule. This section is not displayed for a Manual service.

Access Control

Use this drop-down list to select the appropriate access control for the rule. You can permit traffic to be forwarded, deny traffic altogether, or contain traffic to a VLAN. Select **None** to disable access control for this rule.

- Permit Traffic enables traffic to be forwarded with the port's assigned VID.
- Deny Traffic traffic will be automatically discarded.
- Contain to VLAN contains traffic to a specific VLAN. Use the drop-down list to select the desired VLAN. Use the Contain to VLAN drop-down list to select a VLAN.

Class of Service

Use the drop-down list to select a class of service to associate with the service. The Policy tab lets you define classes of service that each include an 802.1p priority, and optionally an IP type of service (ToS/DSCP) value, rate limits, and transmit queue configuration. You can then assign a class of service as a classification rule action. See Getting Started with Class of Service and How to Create a Class of Service for more information. Select **None** to disable class of service for this rule. Use the configuration menu button to the right of the drop-down list to add or edit a Class of Service.

When rule accounting is enabled on a device, each rule keeps a list of the ports on which it has been used. The next three options enable you to specify certain rule usage actions to take place when a "rule hit" is reported.

System Log

Specify System Log functionality for the rule:

- **Enabled** If this option is enabled, a syslog message is generated when the rule is used. This option must be enabled if you are configuring Policy Rule Hit Reporting on your devices.
- **Disabled** If this option is disabled and this rule is hit, it does not generate a Syslog message, but lower-precedence rules and the role default actions can still specify a syslog message be sent for this data packet if there is a match.
- **Prohibited** If this rule is hit, no syslog message is generated for this data packet, even when a lower-precedence rule or the role default actions has the System Log action set to enabled.

Audit Trap

Specify Audit Trap functionality for the rule:

- Enabled If this option is enabled, an audit trap is generated when the rule is used.
- **Disabled** If this option is disabled and this rule is hit, it does not generate an audit trap, but lower-precedence rules and the role default actions can still specify generating an audit trap for this data packet if there is a match.
- **Prohibited** If this rule is hit, no audit trap is generated for this data packet, even when a lower-precedence rule or the role default actions has the Audit Trap action set to enabled.

Disable Port

Specify Disable Port functionality for the rule:

- **Enabled** If this option is enabled, any port reported as using this rule is disabled. Ports that have been disabled due to this option are displayed in the device Role/Rule tab.
- **Disabled** If this option is disabled and this rule is hit, it does not disable the port, but lower-precedence rules and the role default actions can still specify disabling the port for this data packet if there is a match.
- **Prohibited** If this rule is hit, the port is not disabled, even when a lower-precedence rule or the role default actions has the Disable Port action set to enabled.

Traffic Mirror

Specify traffic mirroring functionality for the rule:

- Select port group(s) Use the drop-down list to select the port groups where mirrored traffic will be sent for monitoring and analysis. Use the configuration menu button to the right of the drop-down list and select View/Modify Port Groups to open the Port Groups tab where you can define user-defined port groups for selection.
- **Disabled** If this option is disabled and this rule is hit, traffic mirroring will not take place, but lower-precedence rules and the role default actions can still specify traffic mirroring for this data packet if there is a match.
- **Prohibited** If this rule is hit, traffic mirroring is disabled, even when a lower-precedence rule or the role default actions has the Traffic Mirror action specified.

Quarantine Role

Specify Quarantine role functionality for the rule:

- Enabled If this option is enabled, any role reported as using this rule is quarantined.
- **Disabled** If this option is disabled and this rule is hit, it does not quarantine the role, but lower-precedence rules and the role default actions can still specify quarantining the role for this data packet if there is a match.
- **Prohibited** If this rule is hit, the role is not quarantined, even when a lower-precedence rule or the role default actions has the Quarantine Role action set to enabled.

Traffic Classification Rules

Traffic Classification rules allow you to assign VLAN membership and/or class of service to your network traffic based on the traffic's classification type. Classification types are derived from Layers 2, 3, 4, and 7 of the OSI model, and all network traffic can be classified according to specific layer 2/3/4/7 information contained in each frame. In the **Policy** tab, rules are used to provide four key policy features: traffic containment, traffic filtering, traffic security, and traffic prioritization. Examples of how to design rules for each of these features are given below.

A Traffic Classification rule has two main parts: Traffic Description and Actions. The Traffic Description identifies the traffic classification type for the rule. The Actions specify whether traffic matching that classification type will be assigned VLAN membership, class of service, or both. When a frame arrives on a port, the switch checks to see if the frame's classification type matches the type specified in a rule. If it does, then the actions defined in that rule will apply to the frame.

In the **Policy** tab, rules are created and then grouped together into Services, which are then used to define roles. A role is assigned to each port either through end user authentication or as the port's default role. This means that there can be multiple rules active on a port. When a frame is received on a port, if the frame's classification type matches more than one rule, classification precedence rules are used to determine which rule to use.

NOTE: Rules included in services are read in the order in which they are listed in ExtremeCloud IQ - Site Engine. To configure rules for ExtremeCloud IQ Controller (formally called XCC or XCA) devices, ensure ExtremeCloud IQ - Site Engine lists the rules in the correct order or the service may not execute the correct rule. To reorder rules in the same service, use drag-and-drop capabilities to move from one group to another.

The following information is discussed in this file:

- Traffic Descriptions
- Actions
 - VLAN Membership
 - Priority (Class of Service)
- Classification Types and their Parameters
 - Layer 2 Data Link Classification Types
 - Layer 3 Network Classification Types
 - Layer 4 Application Transport Classification Types
 - Layer 7 Application Classification Type
- Examples of How Rules are Used
 - Traffic Containment
 - Traffic Filtering
 - Traffic Security
 - Traffic Prioritization

Traffic Descriptions

When you create a Traffic Classification rule in the **Policy** tab, you must define the rule's traffic description. The traffic description identifies the traffic classification type for that rule. You must select a classification type, and then select or enter certain parameters or values for each type.

Classification types are grouped according to Layers 2, 3, 4, and 7 of the OSI model and there are multiple classification types for each layer.

OSI Model	
Layer 7 - Application	
Layer 6 - Presentation	
Layer 5 - Session	
Layer 4 - Transport	
Layer 3 - Network	
Layer 2 - Data Link	
Layer 1 - Physical	

Specific Layer 2/3/4/7 information contained in each frame is used to identify the frame's classification type. Each layer uses different information to classify frames.

- Layer 2 Data Link -- classifies frames based on an exact match of the MAC address or specific protocol type of each frame.
- Layer 3 Network -- classifies IP or IPX frames based on specific information contained within the Layer 3 header
- Layer 4 Transport -- classifies IP frames based on specific Layer 4 TCP or UDP port numbers contained in the header.
- Layer 7 Application -- classifies frames based on specific Layer 7 application types.

For a complete description of Layer 2, 3, 4, and 7 classifications, refer to <u>Classification Types and Their Parameters</u>.

Actions

When you create a Traffic Classification rule in the **Policy** tab, you must define the actions the rule performs. When a frame arrives on a port, the switch checks to see if the frame's classification type matches the type specified in a rule. If it does, then the actions defined in that rule will apply to the frame. Actions specify whether the frame will be assigned VLAN membership (access control) and/or priority (class of service).

VLAN Membership (Access Control)

In your network domains, you can create VLANs (Virtual Local Area Networks) that allow endsystems connected to separate ports to send and receive traffic as though they were all connected to the same network segment. Using traffic classification rules, you can classify a frame based on the frame's classification type to have membership in a specific VLAN, providing important traffic containment, filtering, and security for your network.

For example, a network administrator could use rules to separate end user traffic into VLANs according to protocol, subnet, or application. Rules could also be used to group geographically separate end-systems into job-specific workgroups.

Priority (Class of Service)

Traffic Classification rules allow you to assign a transmission priority to frames received on a port based on the frame's classification type. For example, a network administrator could use rules to assign priority to one network application over another.

Priority is a value between 0 and 7 assigned to each frame as it is received on a port, with 7 being the highest priority. Frames assigned a higher priority will be transmitted before frames with a lower priority. Each of the priorities is mapped into a specific transmit queue by the switch or router. The insertion of the priority value (0-7) allows all 802.1Q devices in the network to make intelligent forwarding decisions based on its own level of support for prioritization.

The **Policy** tab enables you to utilize priority by creating classes of service that each include an 802.1p priority, and optionally an IP type of service (ToS/DSCP) value, rate limits, and transmit queue configuration. You can then assign the class of service as a classification rule action, as part of the definition of an automated service, or as a role default. See Getting Started with Class of Service for more information.

Classification Types and their Parameters

When you define a rule's traffic description, you select a classification type, and then select or enter certain parameters or values for each type. Classification types are grouped according to Layers 2, 3, 4, or 7 of the OSI model.

Layer 2 -- Data Link Classification Types

Layer 2 classification types allow you to define classification rules based on an exact match of the MAC address or specific protocol type of each frame.

MAC Address Source, MAC Address Destination, MAC Address Bilateral

These classification types are based on an exact match of the source, destination, or bilateral (either source or destination) MAC address contained in an Ethernet frame. Enter a valid MAC address or select Select to open a window where you can select a MAC address read from your network devices. You can specify a mask, however masking a MAC address is not supported on legacy devices.

Ethertype

This classification type is based on the specific protocol type of each frame defined in the two-byte Ethertype field. Select an Ethertype from the list of well-known values, or select **Other** and manually enter a single value in hexadecimal form. You can enter a range of values, however range rules are not supported on legacy devices or N-Series Gold.

Well-known Ethertypes	Values
IP	0x080x0
ARP	0x0806
Reverse ARP	0x8035
Novell IPX 1	0x8137
Novell IPX 2	0x8138
Banyan	0x0bad
AppleTalk	0x809b
AppleTalk ARP	0x80f3
IPv6	0x86dd
Decnet Phase 4	0x6003

DSAP/SSAP

This classification type is based on the specific protocol type of each frame defined in the DSAP and SSAP fields. Select a protocol from the list of well-known values, or select **Other** and manually enter a

custom two-byte value in hexadecimal format (OxFFFF). The LSB of the DSAP address specifies Individual(0) or Group(1), while the LSB of the SSAP address specifies Command(0) or Response(1). For the SNAP frame type, you may enter Advanced DSAP/SSAP configurations. The advanced fields are not supported on legacy devices and are ignored.

Well-known DSAP/SSAP Types	Values
IP	0x0606
IPX	0xe0e0
NetBIOS	0xf0f0
Banyan Vines	Oxbcbc
SNA	0x0404
SNAP	OxAAAA
Other	a two-byte value

VLAN ID

This classification type is based on an exact match of the VLAN tag contained within a frame. Select a VLAN ID (VID) from the list of VLANs defined in the Policy tab. If you select **Other**, you must enter a single VID or specify a range of VIDs in decimal form. Range rules are not supported on legacy devices.

Priority

This classification type is based on an exact match of the Priority tag contained within a frame. Select a Priority value 0 - 7 from the list of well-known values, or select **Other** and enter a value in decimal form.

Layer 3 -- Network Classification Types

Layer 3 Network classification types allow you to define classification rules based on specific information contained within the Layer 3 header of an IP or IPX frame.

IP Time to Live (TTL)

This classification type is based on an exact match of the TTL field contained in the IP header of a frame. The TTL field indicates the maximum number of router hops the packet can make before being discarded. The TTL field is set by the packet sender and reduced by every router on the route to its destination. If the TTL field reaches zero before the packet arrives at its destination, then the packet is discarded. IP Time to Live rules are only supported on K-Series and S-Series devices.

IPX Network Source, IPX Network Destination, IPX Network Bilateral

These classification types are based on specific information contained within the Layer 3 header of an IPX frame. It is a four-byte user-defined value that represents the IPX source, destination, or bilateral (either source or destination) network number. This value must be a valid IPX network address in hexadecimal form. You can enter a range of values, however range rules are not supported on legacy devices or N-Series Gold.

IPX Socket Source, IPX Socket Destination, IPX Socket Bilateral

These classification types are based on specific information contained within the Layer 3 header of an IPX frame. It is a two-byte, user-defined value that represents the IPX source, destination, or bilateral (either source or destination) socket numbers. This value is used by higher layer protocols to target

specific applications running among hosts. Select an IPX Socket type from the list of well-known values, or select **Other** and manually enter the value in decimal form. You can enter a range of values, however range rules are not supported on legacy devices or N-Series Gold.

Well-known IPX Socket Types	Values
NCP	1105
SAP	1106
RIP	1107
NetBIOS	1109
Diagnostics	1110
NSLP	36865
IPX Wan	56868
Other	0-65535

IPX Class of Service

This classification type is based on specific information contained within the Layer 3 header of an IPX frame. This is a one-byte field used for transmission control (hop count) by IPX routers. Enter a valid IPX Class of Service in decimal form, 0-255. You can enter a range of values, however range rules are not supported on legacy devices or N-Series Gold.

IPX Packet Type

This classification type is based on specific information contained within the Layer 3 header of an IPX frame. Select an IPX Packet type from the list of well-known values or select **Other** and manually enter the value in decimal form. You can enter a range of values, however range rules are not supported on legacy devices or N-Series Gold.

Well-known IPX Packet Types	Values
Hello/SAP	0
RIP	1
Echo Packet	2
Error Packet	3
NetWare 386	4
SeqPackProt	5
NetWare 286	17
Other	0-31

IPv6 Address Source, IPv6 Address Destination, IPv6 Address Bilateral

These classification types are based on an exact match of the source, destination, or bilateral (either source or destination) IPv6 address information contained within the IPv6 header of each frame. Enter a valid IPv6 address and optional mask ("/n") in the Value field.

IPv6 Socket Source, IPv6 Socket Destination, IPv6 Socket Bilateral

These classification types are based on an exact match of a specific source, destination, or bilateral (either source or destination) IPv6 address and a UDP/TCP port number (type) contained within the IPv6 header of each frame. Enter an IPv6 address in the Value field. Then, select a UDP/TCP type from the list of well-known values, or select **Other** and manually enter the value in form. (UDP/TCP port numbers are defined in RFC 1700.) If you select **Other**, you can enter a range of values.

Well-known UDP/TCP Types	Values
FTP Data	20
FTP	21
SSH	22
Telnet	23
SMTP	25
TACACS	49
DNS	53
BootP Server	67
BootP Client	68
TFTP	69
Finger	79
HTTP	80
POP3	110
Portmapper	111
NNTP	119
NTP	123
NetBIOS Name Service	137
NetBIOS Datagram Service	138
NetBIOS Session Service	139
IMAP2/IMAP4	143
SNMP	161
IMAP3	220
LDAP	389
HTTPS	443
R-Exec	512
R-Login	513
· · · · · · · · · · · · · · · · · · ·	

Well-known UDP/TCP Types	Values
R-Shell	514
LPR	515
RIP	520
SOCKS	1080
Citrix ICA	1494
RADIUS	1812
RADIUS Accounting	1813
NFS	2049
X11 (Range Start)	6000
X11 (Range End)	6063
Other	0-65535

IPv6 Flow Label

These classification types are based on the exact match of the value in the 20-bit Flow Label field in the IPv6 header. This field is used to identify packets belonging to particular traffic flow that needs special traffic handling. Enter a flow label value and sights mask.

IP Address Source, IP Address Destination, IP Address Bilateral

These classification types are based on an exact match of the source, destination, or bilateral (either source or destination) IP address information contained within the IP header of each frame. Enter a valid IP address and optional mask ("/n") in the Value field.

IP Socket Source, IP Socket Destination, IP Socket Bilateral

These classification types are based on an exact match of a specific source, destination, or bilateral (either source or destination) IP address and a UDP/TCP port number (type) contained within the IP header of each frame. Enter an IP address in the Value field. Then, select a UDP/TCP type from the list of well-known values, or select **Other** and manually enter the value in decimal form. (UDP/TCP port numbers are defined in RFC 1700.) If you select **Other**, you can enter a range of values, however range rules are not supported on legacy devices or N-Series Gold.

Well-known UDP/TCP Types	Values
FTP Data	20
FTP	21
SSH	22
Telnet	23
SMTP	25
TACACS	49

Well-known UDP/TCP Types	Values
DNS	53
BootP Server	67
BootP Client	68
TFTP	69
Finger	79
HTTP	80
POP3	110
Portmapper	111
NNTP	119
NTP	123
NetBIOS Name Service	137
NetBIOS Datagram Service	138
NetBIOS Session Service	139
IMAP2/IMAP4	143
SNMP	161
IMAP3	220
LDAP	389
HTTPS	443
R-Exec	512
R-Login	513
R-Shell	514
LPR	515
RIP	520
SOCKS	1080
Citrix ICA	1494
RADIUS	1812
RADIUS Accounting	1813
NFS	2049
X11 (Range Start)	6000
X11 (Range End)	6063
Other	0-65535

IP Fragment

This classification type is based on Layer 4 information in fragmented frames. IP supports frame fragmentation, where large frames are divided into smaller fragments and sent wrapped in the original Layer 3 (IP) header. When a frame is fragmented, information that is Layer 4 and above is only present in the first fragment. For example, the first fragment may be classified to Layer 4, while subsequent fragments will be classified only to Layer 3. The product line does not support Layer 4 classification for IP frames that have been fragmented, as the Layer 4 information is not present in these frames. Using the IP Fragment classification rule, any frame which is a fragment of a larger frame, is classified according to the information in the original frame. If the first fragment is classified to Layer 4, subsequent fragments will also be classified to Layer 4.

ICMP and ICMPv6

These classification types are based on an exact match of the ICMP (Internet Control Message Protocol) message contained in the ICMP tag within a frame. Select an ICMP well-known value type from the list of well-known values (some well-known value types also let you select a code), or select **Other** and manually enter the value in hexadecimal form. The format of the value is 0xXXYY, where "XX" is the ICMP type, and "YY" is the associated code, if applicable. You can enter a range of values, however range rules are not supported on legacy devices or N-Series Gold.

IP Type of Service

This classification type is based on an exact match of the one-byte ToS/DSCP field contained in the IP header of a frame. The ToS (Type of Service) or DSCP (Diffserve Codepoint) value is defined by an 8-bit hexadecimal number between 0 and FF. Enter a value or select Select to open a window where you can generate a hex value.

Type of Service can be used by applications to indicate priority and Quality of Service for each frame. The level of service is determined by a set of service parameters which provide a three way trade-off between low-delay, high-reliability, and high-throughput. The use of service parameters may increase the cost of service. In many networks, better performance for one of these parameters is coupled with worse performance on another. Except for very unusual cases, at most, two of the parameters should be set.

For a ToS value, the 8-bit hexadecimal number breaks down as follows:

Bits 0-2: Precedence

Bit 3: O=Normal Delay, 1=Low Delay

Bit 4: O=Normal Throughput, 1=High Throughput

Bit 5: 0=Normal Reliability, 1=High Reliability

Bits 6-7: Explicit Congestion Notification

The precedence bits (bits 0-2) break down as follows:

111 - Network Control

110 - Internetwork Control

101 - CRITIC/ECP

100 - Flash Override

011 - Flash

010 - Immediate

001 - Priority

000 - Routine

The Network Control precedence designation is intended to be used within a network only. The actual use and control of that designation is up to each network. The Internetwork Control designation is intended for use by gateway originators only.

For a DSCP value, the value represents codepoints for two Differentiated Services (DS) Per-Hop-Behavior (PHB) groups called Expedited Forwarding (EF) and Assured Forwarding (AF). For more information on these PHB groups, refer to RFC 2597 and RFC 2598.

IP Protocol Type

This classification type is based on the specific protocol type defined in a field contained in the IP header of each frame. Select a protocol from the list of well-known values, or select **Other** and manually enter the value in decimal form. You can enter a range of values, however range rules are not supported on legacy devices or N-Series Gold.

Well-known IP Protocol Types	Values
ICMP	1
IGMP	2
TCP	6
EGP	8
UDP	17
IPv6 (encapsulated in IPv4 packets)	41
RSVP	46
GRE	47
ESP	50
AH	51
ICMPv6	58
EIGRP	88
OSPF	89
PIM	103
VRRP	112
L2TP	115
Other	0-255

Layer 4 -- Application Transport Classification Types

Layer 4 IP classification types allow you to define classification rules based on specific Layer 4 TCP or UDP port numbers contained in the header of an IP frame. You can specify a specific port number or a range of port numbers.

Note: Certain devices do not support Layer 4 classification for IP frames that have been fragmented, as the Layer 4 information is not present in these frames. If a device has an FDDI HSIM installed, Layer 4 classification will not be supported for any frames larger than 1500 bytes. Frames larger than 1500 bytes are fragmented internally in the switch. When creating classification rules based on specific Layer 4 information, using the <u>IP Fragment</u> classification rule will allow fragmented frames to be classified according to the Layer 4 information contained in the original frame.

IP UDP Port Source, IP UDP Port Destination, IP UDP Port Bilateral

These classification types are based on specific Layer 4 UDP port numbers contained within the header of an IP frame. Select a UDP type from the list of well-known values, or select **Other** and manually enter the value in decimal form. (UDP port numbers are defined in RFC 1700.) You can enter a range of values, however range rules are not supported on legacy devices or N-Series Gold. Enter a valid IPv4 or IPv6 address and optional mask ("/n"), if desired. The IP address is an optional field and does not have to be specified. It is only valid for non-range port values.

Well-known UDP Types	Values
FTP Data	20
FTP	21
SSH	22
Telnet	23
SMTP	25
TACACS	49
DNS	53
BootP Server	67
BootP Client	68
TFTP	69
Finger	79
НТТР	80
POP3	110
Portmapper	111
NNTP	119
NTP	123
NetBIOS Name Service	137
NetBIOS Datagram Service	138
NetBIOS Session Service	139
IMAP2/IMAP4	143

Well-known UDP Types	Values
SNMP	161
IMAP3	220
LDAP	389
HTTPS	443
R-Exec	512
R-Login	513
R-Shell	514
LPR	515
RIP	520
SOCKS	1080
Citrix ICA	1494
RADIUS	1812
RADIUS Accounting	1813
NFS	2049
X11 (Range Start)	6000
X11 (Range End)	6063
Other	0-65535

IP TCP Port Source, IP TCP Port Destination, IP TCP Port Bilateral

These classification types are based on specific Layer 4 TCP port numbers contained within the header of an IP frame. Select a TCP type from the list of well-known values, or select **Other** and manually enter the value in decimal form. (TCP port numbers are defined in RFC 1700.) You can enter a range of values, however range rules are not supported on legacy devices or N-Series Gold. Enter a valid IPv4 or IPv6 address and optional mask ("/n"), if desired. The IP address is an optional field and does not have to be specified. It is only valid for non-range port values.

Well-known TCP Types	Values
FTP Data	20
FTP	21
SSH	22
Telnet	23
SMTP	25
TACACS	49
DNS	53

Well-known TCP Types	Values
BootP Server	67
BootP Client	68
TFTP	69
Finger	79
HTTP	80
POP3	110
Portmapper	111
NNTP	119
NTP	123
NetBIOS Name Service	137
NetBIOS Datagram Service	138
NetBIOS Session Service	139
IMAP2/IMAP4	143
SNMP	161
IMAP3	220
LDAP	389
HTTPS	443
R-Exec	512
R-Login	513
R-Shell	514
LPR	515
RIP	520
SOCKS	1080
Citrix ICA	1494
RADIUS	1812
RADIUS Accounting	1813
NFS	2049
X11 (Range Start)	6000
X11 (Range End)	6063
Other	0-65535

IP UDP Port Source Range, IP UDP Port Destination Range, IP UDP Port Bilateral Range

These classification types are based on Layer 4 UDP port numbers contained within the header of an IP frame. When you select this type, you enter a range of UDP port numbers that the port number in the

header will be matched against. Enter the start and end range values in decimal form. UDP port numbers are defined in RFC 1700.

IP TCP Port Source Range, IP TCP Port Destination Range, IP TCP Port Bilateral Range

These classification types are based on Layer 4 TCP port numbers contained within the header of an IP frame. When you select this type, you enter a range of TCP port numbers that the port number in the header will be matched against. Enter the start and end range values in decimal form. TCP port numbers are defined in RFC 1700.

Layer 7 -- Application Classification Types

Layer 7 IP classification types allow you to define classification rules based on specific Layer 7 application types.

Application

This rule type allows management of traffic for a specific application type, for example Apple traffic (Bonjour) using mDNS-SD. The following application types are supported:


- LLMNR (Link Local Multicast Name Resolution) Query/Response
 This protocol is based on the Domain Name System (DNS) packet format. It allows hosts to perform name resolution for hosts on the same local link.
- SSDP (Simple Service Discovery Protocol) Query/Response
 SSDP is a Universal Plug-and-Play (UPnP) based protocol. SSDP uses the NOTIFY and MSEARCH
 HTTP methods to discover and advertise services on the network.
- mDNS-SD (Multicast Domain Name System Service Discovery) Query/Response
 DNS-SD is a service discovery protocol that utilizes the Domain Name System. Multicast DNS is a
 protocol that is mostly compatible with normal DNS but uses link local multicast addressing,
 allowing for zero configuration networking (zeroconf) functionality.

Examples of How Rules are Used

Traffic Classification rules are used to provide four key policy features: Traffic Containment, Traffic Filtering, Traffic Security, and Traffic Priority.

Traffic Containment

Using classification rules, network administrators can group together users of a given protocol, subnet, or application, and control where their traffic can logically go on the network.



The figure above shows a configuration where the network administrator wants to separate end-user traffic into VLANs based on the assigned IP subnet of each department. This can easily be accomplished by creating two Layer 3 classification rules based on the IP subnet range of the respective departments.


Rule 1 - Engineering, which uses the 132.181.28.x subnet, will be assigned to the Red VLAN.

Rule 2 - Sales, which uses the 132.181.29.x subnet, will be assigned to the Blue VLAN.

Based on these two Layer 3 classification rules, the traffic from the Engineering VLAN will be isolated from the Sales VLAN. Since these rules are based on Layer 3 information, an Engineering user could enter the network from a connection in the Sales department, and that user would still be contained in the Engineering VLAN.

Traffic Filtering

Classification rules can also be used to filter out (discard) specific unwanted traffic. Filter criteria can include things such as broadcast routing protocols, specific IP addresses, or even applications such as HTTP or SMTP.



The figure above shows a common configuration in which a routed backbone is using both RIP and OSPF for its routing protocols. The network administrator does not want the multicast

OSPF and broadcast RIP frames propagated to the end stations. The network is designed so that only end users are attached to the E7 devices.

To implement filtering in this scenario, a Layer 3 rule and a Layer 4 rule will be created.


Rule 1 (Layer 3) - Any frame received with an IP Protocol Type of 89 (OSPF) will be discarded.

Rule 2 (Layer 4) - Any frame received with a Bilateral UDP port number of 520 (RIP) will be discarded.

Based on this configuration, all RIP and OSPF frames will be filtered from the end users.

Traffic Security

Traffic Security uses the same concepts as <u>Traffic Filtering</u>. Imagine a scenario where network access is provided to a group of unknown users. There have been problems with these unknown users "hacking" into the router and altering the configuration. A simple classification rule can be put in place that will prevent these types of occurrences.



In the figure above, the network components include a router and an E7 device. In this configuration end-users connect to the ports of the E7 device.


Since the end-users would never need to communicate directly to the router using the router's IP address, a Layer 3 IP classification rule will be used.

Rule - Any frames received by the switch with a destination IP address of the router (129.168.1.2) will be discarded.

The end result is that any frames from a user trying to "hack" into the router will be discarded before ever reaching the router.

Traffic Prioritization

Classification rules can be used to specify that certain network applications receive the highest transmission priority. For example, a network administrator wants to assign priority to three network applications, SAP R/3, web traffic, and email, in that order.



To accomplish the prioritization goals in this example, there are two main steps required: creating the classification rules, and then configuring the priority-to-transmit queue mapping for the switch, if needed.

First, create one Layer 3 and two Layer 4 classification rules.

Rule 1, Layer 3 (SAP R/3) - All frames to or from the IP address of the SAP R/3 server will be tagged with a priority indicator of 7 (highest).

Rule 2, Layer 4 (Web) - All frames with a TCP port number of 80 (HTTP) will be tagged with a priority indicator of 5.

Rule 3, Layer 4 (email) - All frames with a TCP port number of 25 (SMTP) will be tagged with a priority indicator of 3.

Note: An IP address classification was selected for Rule 1 because it has been observed that SAP R/3 dynamically negotiates the TCP/UDP port used, so the port number selections vary from session to session. If this was not the case, a Layer 4 UDP classification could be used.

Then, configure the priority-to-transmit queue mappings. Each switch has default priority-to-transmit queue mappings. You can use these defaults or change the mappings using local management. In addition, the **Policy** tab provides the ability to configure transmit queues as part of the Role-Based Rate Limits and Transmit Queue Configuration class of service mode. This functionality is available only on certain devices such as the S-Series and N-Series Gold and Platinum devices (refer to the ExtremeCloud IQ - Site Engine Firmware Support matrix for specific device/firmware rate limit support).

Based on the default priority-to-traffic queue mapping for an E7 device, the priorities assigned above will work out so that each frame classification type will be mapped to the desired traffic queue. This means that no user configuration of the priority-to-transmit queue mapping would be required.

With the classification rules described above, the network traffic would be prioritized as shown in the table below:

Application	Classification Type	Desired Priority	Priority Value	E7 Traffic Queue
SAPR/3	Bilateral IP	High	7	3
Web	TCP Port Number	Medium	5	2

Email TCP Port Number Low		
Lindii icr foit number Low	3 1	

Ports (Transmit Queue Port Group)


The **Ports** tab lets you view all the ports in the selected transmit queue port group, as well as add and remove ports to and from the group. It provides information about each port, and lets you view and edit port information.

To access this tab:

- 1. Open the **Control** tab.
- 2. Open the Policy tab.
- 3. Open the Class of Service > CoS Components left-panel tab.
- 4. Select either the Transmit Queue Port Groups left-panel tab.
- 5. Select a existing port group in the left panel to open it in the Transmit Queue Port Group tab.

NOTE: Create a new port group by right-clicking the **Transmit Queue Port Groups** left-panel tab, selecting **Create Port Group**, entering a **Name** for the port group, and selecting **OK**.

6. Select the **Ports** tab in the right panel.



Name

Name of the port, constructed of the name or IP address of the device and either the port index number or the port interface name.

Rate/Queue Port Type

The number of rate limits the port supports.

Default Role

The **Default Role** assigned to the port.

Alias

Shows the alias (ifAlias) for the interface, if one is assigned.

Stats

Shows statistics collected for a port, enabled via the Flow Collection & Interface setting in the PortView.

Port Type

Type of port. Possible values include: Access, Interswitch Backplane, Backplane, Interswitch, and Logical.

Neighbor

The port's neighbor port.

Port Speed

Speed of the port. Possible values include: 10/100, speed in megabits per second (for example, 800.0 Mbps), Unknown (displayed for logical ports).

Description

A description of the port.

Add/Remove Ports Button


Opens the <u>Add/Remove Ports window</u>, where you can add and remove ports to and from the port group. When you create new port groups, you add ports from the Default group into your newly defined port groups.



Summary (Transmit Queue Port Groups)

This tab displays the transmit queue port groups. Transmit queue mapping maps a logical transmit queue index (used by a class of service) to an actual physical transmit queue you have configured in the **Policy** tab. You can configure transmit queue mappings for a port group using the CoS - Transmit Queue Mappings tab.

To access this tab, open the **Class of Service > CoS Components** tab. Then, select the select the **Transmit Queue Port Groups** tab in the left panel. The Summary tab displays in the right panel.



Name

The name of the transmit queue port group.

CoS - Transmit Queue Mappings (Transmit Queue Port Group)

This tab lets you view and configure the transmit queue mappings for a port group. Transmit queue mappings map a logical rate limit index used by classes of service to an actual physical rate limit you have created in ExtremeCloud IQ - Site Engine.

Each port group has its own set of index mappings. ExtremeCloud IQ - Site Engine automatically assigns these index numbers when you configure a class of services' rate limits and transmit queue shapers.

The **Transmit Queue Mappings** tab allows you to do two things:


- Map the index to a different rate for different port groups (edge ports versus inter-switch links). See Creating Class of Service Port Groups
- Map the index to a different rate limit for each port type (8-rate limit, 32-rate limit, 64-rate limit, and 100-rate limit) in a port group. See Advanced Rate Limiting by Port Type.

To access this tab:

- 1. Open the **Control** tab.
- 2. Open the Policy tab.
- 3. Open the Class of Service > CoS Components left-panel tab.
- 4. Select either the **Transmit Queue Port Groups** left-panel tab.
- 5. Select a existing port group in the left panel to open it in the Transmit Queue Port Group tab.

NOTE: Create a new port group by right-clicking the **Transmit Queue Port Groups** left-panel tab, selecting **Create Port Group**, entering a **Name** for the port group, and selecting **OK**.

6. Select the CoS - Transmit Queue Mappings tab in the right panel.



TXQ Index

The logical transmit queue index. This index number is specified in a class of service and dictates the queue and shaping behavior for incoming packets.

Transmit Queue

Displays the physical transmit queue used to map to each transmit queue index. To change this value, select the **Edit Index Mapping** button to open the Edit Transmit Queue Mapping window and select a value in the **Transmit Queue** drop-down list.

Rate Shaper

The transmit queue's associated rate shaper. To change this value, select the **Select Rate Shaper** button to open the Select Transmit Queue Rate Shaper window and select a value in the **Rate Limit** field.

TXQ Port Type

The Port Type is based on the number of transmit queues the port supports: 4 transmit queues or 16 transmit queues.


TXQ Index Used By CoS

The Class of Service using this TXQ index.

Ports (Flood Control Port Groups)

The Flood Control Port Group Ports tab provides a table of information about the ports in the selected port group. It also includes buttons that enable you to retrieve the latest information about the ports and to add and remove ports. To access this tab, select a port group in the left-panel Flood Control Port Groups tab, then select the Ports tab in the right panel.

NOTE: The **Ports** tab is only available when a Flood Control port group is selected, and when advanced mode is enabled on the CoS Components tab.



Name

Name of the port, constructed of the name or IP address of the device and either the port index number or the port interface name.

Rate/Queue Port Type

Shows the selected port type rate/queue.

Default Role

Shows the default role for the port. See <u>Default Role</u> in the Concepts topic for information on default roles. For additional information, see <u>Port Mode</u>.

Alias

Shows the alias (ifAlias) for the interface, if one is assigned.

Stats

Shows that statistics are being collected for a port, enabled via the **PortView**.

Port Type

Type of port. Possible values include: Access, Interswitch Backplane, Backplane, Interswitch, and Logical.

Neighbor

Port to which the port is connected.

Port Speed

Speed of the port. Possible values include: 10/100, speed in megabits per second (for example, 800.0 Mbps), Unknown (displayed for logical ports).

Description

A description of the port.

Add/Remove Button

Selecting a port in the table and selecting this button opens the <u>Add/Remove Ports window</u>, which enables you to add and remove ports to and from the port group. This option is available for user-defined port groups only.

Flood Control Port Groups

This panel lists port groups on which you can configure flood control. Each port group supports rate limits for three separate configured traffic types (Unicast, Multicast, and Broadcast).

To access this tab, open Class of Service > CoS Components left panel of the Policy tab and select Flood Control Port Groups.



Name

The name of the port group.

Flood Control Rate Limits (Flood Control Port Groups)


This tab allows you to set individual flood control rates for each traffic type (Unicast, Multicast, and Broadcast).

Choices include:

- None
- Rate limits created in the **Rate Limit** tab. For additional information, see Create Rate Limit/Shaper.

As flood control is enabled/disabled for a Class of Service, when enabled, each column displays a rate limit, or **None**, if no rate has been defined for that portion of flood control.

To access this tab, open the Class of Service > CoS Components left-panel tab. Then, select the Flood Control checkbox from the General tab in the left-panel to display the Flood Control Port Groups tab in the left panel. Expand the Flood Control Port Groups tab, and select a flood control port group in the tree. The Flood Control Port Groups tab is displayed in the right panel.



Unicast Unknown

Select a rate, create a new rate, or edit an existing flood control rate limit for Unicast traffic.

Multicast

Select a rate, create a new rate, or edit an existing flood control rate limit for Multicast traffic.

Broadcast

Select a rate, create a new rate, or edit an existing flood control rate limit for Broadcast traffic.

Class of Service Example

This Help topic provides an example of how class of service (CoS) can be configured on a network to manage bandwidth requirements of network traffic. Before you look at this example, read Getting Started with Class of Service.


In this example, an organization's network administrator needs to assure that VoIP traffic, both originating in and transiting a network of edge switches and a core router, is configured with appropriate priority, ToS, and queue treatment. We also rate limit the VoIP traffic at the edge to 1 Mb/s to guard against DOS attacks, VoIP traffic into the core at 25 Mb/s, and H.323 call setup at 5 PPS. Data traffic retains the default configuration.


This example assumes CEP authentication using H.323 for VoIP. For networks that do not authenticate VoIP end point with CEP H.323 authentication, the VoIP policy needs to be adjusted accordingly. For instance, SIP uses UDP port 5060, not the TCP port 1720.

To simplify the discussion of the configuration process, this example is limited to the VoIP configuration context. The following table provides a set of sample values for priority, inbound rate limit (IRL), and transmit queue across a number of real world traffic types. This table can be used as an aid in thinking about how you might want to apply CoS across your network. Note that Scavenger class is traffic that should be treated as less than best effort: external web traffic, for instance.

			IF	IRL -			Transm	it Queu	е	
CoS Name	CoS Index	Priority	IKL		Queue #		Shaping		Bandwidth	
	IIIGEA	Edge Core Edge	Edge	Core	Edge	Core	Edge	Core		
Scavenger (Static)	0	0	15 Mb/s		0	0	10%		5%	5%
Best Effort (Static)	1	1								
Bulk Data (Static)	2	2			1	1	80%		45%	45%
Critical Data (Static)	3	3								
Network Control (Static)	4	4	40 PPS	1 Mb/s	2	2	1 Mb/s		25%	25%
Network Mgmt (Static)	5	5	2 Mb/s			2	1 1410/2		25%	25%
RTP/Voice/Video (Static)	6	6	1 Mb/s	OF Mb /a	3	3			25%	25%
High Priority (Static)	7	7	1 1410/5	25 Mb/s	3	3			25%	23%
VoIP Call Setup	8	7	5 F	PPS	3	3			25%	25%

The following figure displays the network setup for this example configuration, with the desired Profile/CoS summary for each network device. Each device is configured with VoIP and Data VLANs. Each VoIP VLAN contains four 1-gigabit interfaces for each device.





Edge and Core port groups in the RTP/Voice/Video (Static) CoS provide for the difference in rate limiting needs between the end user and aggregation devices. A VoIP Call Setup CoS provides rate limiting for the setup aspect of the VoIP call.

The Edge, Core, and H.323 Call Setup roles are configured with TCI Overwrite, default CoS 5 (best default priority for voice and video), and default access control that contains traffic to the appropriate VLAN.

Use the Policy tab to configure the policy roles and related services using the following instructions. For more information, see How to Create a Class of Service and How to Define Rate Limits.

Configure the Classes of Service

Use the Class of Service tab to configure the static RTP/Voice/Video CoS with the appropriate edge and core rate limits, and create a new CoS for the call setup rate limits.

- 1. For the static RTP/Voice/Video CoS (CoS Index 6):
 - a. Set the ToS to B8.
 - b. Create two new Inbound RL port groups called Edge and Core.
 - c. Set the Edge port group rate limit to 1 Mb/s and the Core port group rate limit to 25 Mb/s. (You can create these rate limits first.)
 - d. Add the appropriate ports to each port group.
- 2. Create a new class of service and name it VoIP Call Setup (CoS Index 8).
 - a. Set the rate limit to 5 PPS for all port groups. (You can create this rate limit first.)
 - b. Set the ToS to B8.

Create the VoIP Core Role

For the core router, create a policy role for VoIP Core. VoIP Core policy deals with packets transiting the core network using VoIP VLAN 22.

- 1. Name the role VoIPCore-VLAN22.
- 2. Enable TCI overwrite so that ToS is rewritten for this role.
- 3. Set the default access control action to Contain to VLAN 22.
- 4. Set default Class of Service to CoS Index 5.

Create a VoIP Core Service

- 1. Name the service VoIPCore.
- 2. Add the service to the VoIPCore-VLAN22 role.

Create a Rule

- 1. Create a Layer 2 traffic classification rule for VLAN ID 22 within the VoIPCore service.
- 2. Assign the static RTP/Voice/Video CoS (CoS Index 6) as the Class of Service action for the rule.

Creating the VoIP Edge Role

For the edge switches, create a policy role for VoIP Edge. VoIP Edge policy deals with packets transiting the edge network using VoIP VLAN 12.

- 1. Name the role VoIPEdge-VLAN12.
- 2. Enable TCI overwrite so that ToS is rewritten for this role.
- 3. Set the default access control action to Contain to VLAN 12.
- 4. Set default Class of Service to CoS Index 5.

Create a VoIP Edge Service

- 1. Name the service VoIPEdge.
- 2. Add the service to the VoIPEdge-VLAN12 role.

Create a Rule

- 1. Create a Layer 2 traffic classification rule for VLAN ID 12 within the VoIPEdge service.
- 2. Assign the static RTP/Voice/Video CoS (CoS Index 6) as the Class of Service action for the rule.

Creating the H.323 Call Setup Role

The H.323 Call Setup role deals with the call setup traffic for VoIP H.323 authenticated users directly attached to the switch using link ge.1.10.

- 1. Name the role H323CallSetup.
- 2. Enable TCI overwrite so that ToS is rewritten for this policy.
- 3. Set default Class of Service to CoS Index 5.

Create a H.323 Call Setup Service

- 1. Name the service H323CallSetup.
- 2. Add the service to the H323CallSetup role.

Create a Rule

Create a Layer 4 traffic classification rule as follows:

- 1. Traffic Classification Type: IP TCP Port Destination
- 2. Enter in Single Value field: 1720 (TCP Port ID).
- 3. For IP TCP Port Destination value: 10.0.0.1 with a mask of 255.255.255.255.
- 4. Assign the new VoIP Call Setup CoS (CoS Index 8) as the Class of Service action for the rule.

Apply the Roles to Network Devices

After you have created your roles, you must apply them to the network devices as follows:

Core Router

Apply the VoIPCore-VLAN22 role to ports ge.1.2-5.

Edge Switch

Apply the VoIPEdge-VLAN12 role to ports ge.1.10-13.

Apply the H323CallSetup role to port ge.1.10

ToS/DSCP Value Definition Chart

Use this chart to compare ToS and DSCP values.

ToS (Dec)	ToS (Hex)	ToS (Binary)	ToS Preceden ce (Binary)	ToS Preceden ce (Decimal)	ToS Precedence Name	ToS Delay Flag	ToS Through put Flag	ToS Reliabili ty Flag	DSCP (Binary)	DSCP (Hex)	DSCP (Decim	DSCP Class
0	0x00	0000000	000	0	Routine	0	0	0	000000	0x00	0	none
32	0x20	00100000	001	1	Priority	0	0	0	001000	0x08	8	cs1
40	0x28	00101000	001	1	Priority	0	1	0	001010	0x0A	10	af11
48	0x30	00110000	001	1	Priority	1	0	0	001100	0x0C	12	af12
56	0x38	00111000	001	1	Priority	1	1	0	001110	0x0E	14	af13
64	0x40	01000000	010	2	Immediate	0	0	0	010000	0x10	16	cs2
72	0x48	01001000	010	2	Immediate	0	1	0	010010	0x12	18	af21
80	0x50	01010000	010	2	Immediate	1	0	0	010100	0x14	20	af22
88	0x58	01011000	010	2	Immediate	1	1	0	010110	0x16	22	af23
96	0x60	01100000	011	3	Flash	0	0	0	011000	0x18	24	cs3
104	0x68	01101000	011	3	Flash	0	1	0	011010	0x1A	26	af31
112	0x70	01110000	011	3	Flash	1	0	0	011100	0x1C	28	af32
120	0x78	01111000	011	3	Flash	1	1	0	011110	0x1E	30	af33
128	0x80	10000000	100	4	FlashOverrid e	0	0	0	100000	0x20	32	cs4
136	0x88	10001000	100	4	FlashOverrid e	0	1	0	100010	0x22	34	af41
144	0x90	10010000	100	4	FlashOverrid e	1	0	0	100100	0x24	36	af42
152	0x98	10011000	100	4	FlashOverrid e	1	1	0	100110	0x26	38	af43
160	0xA0	10100000	101	5	Critical	0	0	0	101000	0x28	40	cs5
184	0xB8	10111000	101	5	Critical	1	1	0	101110	0x2E	46	ef
192	0xC0	11000000	110	6	InterNetwork Control	0	0	0	110000	0x30	48	cs6
224	0xE0	11100000	111	7	Network Control	0	0	0	111000	0x38	56	cs7

Policy VLAN Tab Overview

The VLAN tab displays information about the VLAN selected in the left panel and lets you configure certain VLAN parameters. If you are using VLAN to Role mapping in your network, you can also use this tab to map the VLAN to a specific role. If you make a change on this tab, you need to enforce it.

Global VLAN: 1[DEFAULT VLAN] Name: DEFAULT VLAN VID: Dynamic Egress Always write VLAN to device(s) Authentication Based VLAN (RFC3580) to Role Mapping Mapped to Role: None Tagged Packet VLAN to Role Mapping NOTE: To forward traffic with the VLAN ID & CoS specified by the mapped Role, TCI Overwrite must be enabled. Select... Device Level Mapping: None Primary C5/B5/A4/C3/B3/G3/C2/B2/D2 mapping Port Level Mappings: Port Role

To view this tab, select **Control > Policy > VLANs** and select a VLAN from the drop down.

General

This area provides general information about the VLAN and enables you to configure the VLAN.

Name

Name of the VLAN selected in the left panel.

VID

Unique number assigned to the VLAN, also called VID (for VLAN ID). This ID was either assigned by an administrator or assigned automatically by the system when the VLAN was created. The value can be anywhere between 1 and 4094, with VID 1 being reserved for the DEFAULT VLAN (a name for a particular VLAN, not to be confused with a role's assigned default VLAN).

Dynamic Egress

Dynamically add all ports which use this VLAN to this VLAN's egress list. Dynamic Egress is enabled by default in Policy Manager. Leave disabled for discard VLANs. See Dynamic Egress for more information.

Always write VLAN to device(s)

If the box is checked, the VLAN is written to the device whether the VLAN is being used in a rule or role, or not. If it is not checked, the VLAN is not written to the device even though it is being used in a rule or role. Enabling this option is a way of ensuring that the device is aware of a VLAN that is being used for something other than policy configuration, and it enables you to configure that VLAN for Dynamic Egress. If the Default VLAN (VID=1) is selected in the left panel, this option is checked and cannot be edited, as the default VLAN is always on the device.

NOTE: On wireless devices (for example, ExtremeWireless and ExtremeCloud Appliance), the VLAN is always written to the device if it is being used in a rule or role, regardless whether this checkbox is checked or not.

Authentication-Based VLAN to Role Mapping

Authentication-Based VLAN to Role Mapping provides a way to assign a role to a user during the authentication process, based on a VLAN Attribute. (For more information, see VLAN to Role Mapping in the Concepts help topic.) This area displays what role (if any) the VLAN is mapped to (at the device-level) and lets you configure a mapping, if desired.

Mapped to Role

The role to which the VLAN is mapped. To select a role, select **Select**, select the **Assign RFC3580 VLAN -** > **Role Mapping** radio button, choose a role in the drop-down list, and select **OK**.

Select

Opens the role Selection View, where you can choose a role to associate with the VLAN.

Tagged Packet VLAN to Role Mapping

Tagged Packet VLAN to Role Mapping provides a way to let policy-enabled devices assign a role to network traffic, based on a VLAN ID. (For more information, see VLAN to Role Mapping in the Concepts help topic.) This area displays what role (if any) the VLAN is mapped to at both the device-level and port-level, and lets you configure mappings, if desired.

NOTE: TCI Overwrite Requirement

Tagged Packet VLAN to Role Mapping will apply the Role definition to incoming packets using a mapped VLAN. This definition will apply a CoS and determine if the packet is discarded or permitted, and if TCI Overwrite is enabled will re-specify the VLAN ID defined by the Rule / Role Default. If TCI Overwrite is disabled, the packet will egress (if permitted by the Rule Hit) with the original VLAN ID it ingressed with.

If supported by the device, you can enable TCI Overwrite for an individual role in the role's General tab. The stackable devices support rewriting the CoS values but not the VLAN ID.

Device Level Mapping

The role the VLAN is mapped to at the device level (all devices). To select a role, select **Select**, choose a role, and select **OK**.

Select

Opens the role Selection View, where you can choose a role to associate with the VLAN at the device level.

Primary C2/B2/D2/C3/B3/G3/C5/B5/A4 mapping

Use this checkbox to specify that this VLAN to role mapping will be the primary mapping for C2/C3/C5 and B2/B3/B5 devices (C2 firmware version 03.02.xx and higher/B2 firmware version 02.00.16 and higher), and D2, A4, and G3 devices (G3 firmware version 6.03.xx and higher). These devices only support one device-level VLAN to role mapping. If you do not make this selection, there will be no device-level mapping for these devices.

Port Level Mappings

This table lists any port-level Tagged Packet VLAN to Role Mappings configured for this VLAN. Port-level mappings override any device-level mapping.

NOTE: This functionality is not yet enabled.



Global VLANs

This tab displays when you select the **Global VLANs** tab in the **VLANs** left-panel tab. It displays a table of information about the existing VLANs.

Right-clicking the **Global VLANs** tab allows you to create a new VLAN by selecting the **Create VLAN** option, while selecting **Reload VLANs** updates the list of VLANs with the latest information.

If you right-click a VLAN in the left-panel tab or in the right-panel table, you have the option to rename and delete the selected VLAN.

Global VLANs				
				Q
Name	VID	Dynamic Egress	Always Write to Device(s)	
DEFAULT VLAN	1	Enabled	Enabled	
VOIP	2		Disabled	
Edge Edge	3		Disabled	
STCOP	4		Disabled	
IMPDEV VLAN- 5	5		Disabled	
IT Staff ∀lan	6		Disabled	
a 7	7		Disabled	
abc abc	8		Disabled	
■ Management Vlan	9		Disabled	
10.20.89.0/32 - 10.20.89.2	10		Disabled	

Name

Name of the VLAN.

VID

Unique number assigned to the VLAN, also called VID (for VLAN ID). For Global VLANs, this ID was either assigned by an administrator or assigned automatically by the system when the VLAN is created. The value can be anywhere between 1 and 4094, with VID 1 being reserved for the DEFAULT VLAN (a name for a particular VLAN, not to be confused with a role's assigned default VLAN).

Dynamic Egress

Indicates whether the Dynamic Egress feature is on (**Enabled**) or off (**Disabled**) for the VLAN. The default is **Enabled**; therefore, this column displays **Enabled** unless a user has turned it off for a particular VLAN.


Always Write to Device(s)

If enabled, the VLAN is written to the device whether or not it is being used in a rule or role.



Create VLAN

This window displays when you right-click the **Global VLANs** left-panel tab and select **Create VLAN**. See How to Create a VLAN, How to Create a Policy VLAN Island, and Roles for additional information.



Name

The name for the VLAN you want to create. VLAN names can be up to 32 characters in length, including spaces. Do not create a VLAN name that uses any letters with diacritical marks. Diacritical marked letters are not supported by SNMP. VLAN names are case sensitive. For example, "Sales" and "sales" would be considered two different VLAN names. You can have multiple VLANs with the same name but with different VLAN IDs in the Policy tab.

VID

Unique numerical identifier for the VLAN, also known as VLAN ID. Can be a value between 1 and 4094, with VID1 being reserved for the DEFAULT VLAN (a name for a particular VLAN, not to be confused with a default VLAN you assign to a role). To select the next VID in sequence, select **Next Available VID**.

Next Available VID Button

Enters the next unassigned VID in the VLAN ID field.


Editing an existing VLAN/Class of Service

OK Button

Creates the VLAN.

Selection View (Roles)

The Roles Selection View displays when you are selecting a role for VLAN to role mapping. It also lets you clear the current VLAN to role mapping. To access this view, select the desired VLAN in the VLANs > Global VLANs left-panel tab, then select the **Select** button in the VLAN to Role Mapping section on the VLAN tab.



Clear RFC3580 VLAN -> Role Mapping

Select this option to clear the current role selection.

Assign RFC3580 -> Role Mapping

Select this option to assign a new role and make a selection from the list of available roles.

Policy VLAN Islands


This tab displays a table of the Island VLANs being used in the Policy VLAN Island, and the names created on the devices in the island. To display this tab, select **Control > Policy > VLANs > Policy VLANs Islands**.

The VLANs Tab provides two sub-tabs:

- (VLAN) VIDs Tab
- (VLAN) Role Mappings Tab

(VLANs) - VIDs Tab

This tab provides information on VIDs assigned to specific islands. When an island is selected, the VIDs tab shows all VIDs for the defined PVI VLANs used for that island.



VLANs

Name of all defined VLANs. Select a VLAN to see the policy VLAN islands in the VLAN Settings section of the window and the VIDs with which that island is associated.

Create

Opens the Create VLAN window from which you can create a PVI VLAN. Unlike global VLANs, PVI VLANs are not created by the Policy tab during enforce. It is left to the user to configure these on the device(s) externally. The Policy tab only associates the appropriate VIDs to the rules during enforce.

Island Name

Shows the names of all VLAN Islands for the PVI VLAN selected in the VLANs section of the window.

Island VLAN ID


Shows the VID used for this PVI VLAN in this Island.

Edit Island VLAN ID

Selecting an island in the table and selecting this button opens the Edit Island VLAN ID window, where you can change the VID for the Island VLAN.

(VLANs) - Role Mappings Tab

This tab displays the role mappings for the Policy VLAN Island.



General

This area provides general information about the VLAN and allows you to configure the VLAN.

Name

Name of the VLAN selected in the left panel.

VID

Unique number assigned to the VLAN, also called VID (for VLAN ID). This ID was either assigned by an administrator or assigned automatically by the system when the VLAN was created. The value can be anywhere between 1 and 4094, with VID 1 being reserved for the DEFAULT VLAN (a name for a particular VLAN, not to be confused with a role's assigned default VLAN).

Dynamic Egress

Dynamically add all ports which use this VLAN to this VLAN's egress list. Dynamic Egress is enabled by default in Policy Manager. Leave disabled for discard VLANs. See Dynamic Egress for more information.

Always write VLAN to device(s)

If the box is checked, the VLAN is written to the device whether the VLAN is being used in a rule or role, or not. If it is not checked, the VLAN is not written to the device even though it is being used in a rule or role. Enabling this option is a way of ensuring that the device is aware of a VLAN that is being used for something other than policy configuration, and it allows you to configure that VLAN for Dynamic Egress. If the Default VLAN (VID=1) is selected in the left panel, this option is checked and cannot be edited, as the default VLAN is always on the device.

NOTE: On wireless devices (for example, ExtremeWireless and ExtremeCloud Appliance), the VLAN is always written to the device if it is being used in a rule or role, regardless whether this checkbox is checked or not.

Authentication-Based VLAN to Role Mapping

Authentication-Based VLAN to Role Mapping provides a way to assign a role to a user during the authentication process, based on a VLAN Attribute. (For more information, see VLAN to Role Mapping in the Concepts help topic.) This area displays what role (if any) the VLAN is mapped to (at the device-level) and lets you configure a mapping, if desired.

Mapped to Role

The role to which the VLAN is mapped. To select a role, select **Select**, select the **Assign RFC3580 VLAN** - **> Role Mapping** radio button, choose a role in the drop-down list, and select **OK**.

Select

Opens the role Selection View, where you can choose a role to associate with the VLAN.

Tagged Packet VLAN to Role Mapping

Tagged Packet VLAN to Role Mapping provides a way to let policy-enabled devices assign a role to network traffic, based on a VLAN ID. (For more information, see VLAN to Role Mapping in the Concepts help topic.) This area displays what role (if any) the VLAN is mapped to at both the device-level and port-level, and lets you configure mappings, if desired.

NOTE: TCI Overwrite Requirement

Tagged Packet VLAN to Role Mapping will apply the Role definition to incoming packets using a mapped VLAN. This definition will apply a CoS and determine if the packet is discarded or permitted, and if TCI Overwrite is enabled will re-specify the VLAN ID defined by the Rule / Role Default. If TCI Overwrite is disabled, the packet will egress (if permitted by the Rule Hit) with the original VLAN ID it ingressed with.

If supported by the device, you can enable TCI Overwrite for an individual role in the role's General tab. The stackable devices support rewriting the CoS values but not the VLAN ID.

Device Level Mapping

The role the VLAN is mapped to at the device level (all devices). To select a role, select **Select**, choose a role, and select **OK**.

Select

Primary C2/B2/D2/C3/B3/G3/C5/B5/A4 mapping

Use this checkbox to specify that this VLAN to role mapping will be the primary mapping for C2/C3/C5 and B2/B3/B5 devices (C2 firmware version 03.02.xx and higher/B2 firmware version 02.00.16 and higher), and D2, A4, and G3 devices (G3 firmware version 6.03.xx and higher). These devices only support one device-level VLAN to role mapping. If you do not make this selection, there will be no device-level mapping for these devices.

Port Level Mappings

This table lists any port-level Tagged Packet VLAN to Role Mappings configured for this VLAN. Port-level mappings override any device-level mapping.

NOTE: This functionality is not yet enabled.



Add Devices (VLAN Islands)


This window enables you to add devices to VLAN islands.

To access the window:

- 1. Select the VLANs > Policy VLAN Islands tab in the left panel.
- 2. Select the **Island Topology** tab in the Policy VLAN Islands right panel.
- 3. Select the Default Island Devices tab in the Island Settings section of the window.
- 4. Select the **Add Devices** button.

Devices contained in an island are assigned a VID for each Island VLAN unique to the island, allowing roles and rules which use the Island VLANs to isolate users to that island. A device must always belong to an island, and shares a common VID assignment for the Island VLANs with all other devices contained in that island.

To add a device to an island, select the Island to which the device is to be added in the **Destination** drop-down list, select the device in the Devices section, and select **Add**. You can also select and add multiple devices.



Destination

Select the VLAN Island to which the device is to be added.

Devices Section

Expand the Island folder from which the VLAN Island is being selected to add the device or devices.

Add Button

Adds the device(s) selected in the Devices panel to the island selected in the Islands panel.

Island Topology (Policy VLAN Islands)


This tab displays a table of information about the Policy VLAN Islands, which shows the VIDs used in the selected island for all defined PVI VLANs. To access this tab, select the Policy VLAN Islands node in the tree of the Access Control Configuration view, and select the Island Topology tab on the right panel.

The Island Topology tab provides two sub-tabs:

- (Island) VIDs Tab
- (Island) Devices Tab

(Island) - VIDs Tab

This tab provides information on VIDs assigned to specific islands. When an island is selected, the VIDs tab shows all VIDs for the defined PVI VLANs that will be used for that island.



Islands

Name of all defined PVI islands. Select an island to see the VIDs and devices associated with that Island.of the VLAN island in which the Island VLAN is being used.

VLAN Name

Shows the defined PVI VLANs in the Domain. Unlike global VLANs, PVI VLANs are not created by the Policy tab during enforce. It is left to the user to configure these on the device(s) externally. The Policy tab only associates the appropriate VIDs to the rules during enforce.

Island VLAN ID

Shows the VID used for this PVI VLAN in this Island.

Edit Island VLAN ID


Selecting an island in the table and selecting this button opens the Edit Island VLAN ID window, where you can change the VID for the Island VLAN.

Create


Opens the Create VLAN Island dialog. For more information, see Creating a VLAN Island.

(Island) - Devices Tab

This tab displays the devices that are part of a Policy VLAN Island. To see a menu of options for a device in the table, right-click the device.



Policy VLAN Islands (PVI) allow Roles and Rules using VLAN containment Access Control to vary the VID across the network based on the Island where a user connects to the network. This can allow the network to isolate resources, for instance putting traffic from visitors in a "Guest" PVI VLAN that uses a different VID for each campus of a company. Below, select an Island to see the specific VID used for each defined PVI VLAN in that island as well as the devices assigned to that island.



Create

Opens the Create VLAN Island dialog. For more information, see Creating a VLAN Island.

Name


The device's IP address.

Add Devices

Opens a separate dialog to add devices to specific Islands. For more information, see Add/Remove Devices window.


۵

Packet Flow Diagram



Network Resources Tab Overview

The **Network Resources** tab displays a table of information about all the network resources in the current domain. To access this tab, select the **Network Resources > Network Resources** left-panel tab on the **Policy** tab. The Details View is displayed in the right panel. Right-click a network resource to rename or delete it. See How to Create a Network Resource for more information on topologies and islands.



Name

Name of the network resource group.

Resource Count

The number of addresses added to the network resource.

Type

The network resource type:

- Layer 2 MAC Define a group of network resource MAC addresses.
- Layer 3 IP Define a group of network resource IP addresses.

Topology


The network resource topology for this group.

Network Resource Group General Tab

This tab lets you configure a network resource group, which is a group of network resource devices associated with an Automated service. You configure the group by selecting a network

resource type (MAC or IP) and typology, and then creating a list of MAC or IP addresses for the resources that are part of the group. When a network resource group is defined, you can associate it with the desired Automated service (see How to Create a Service for more information).

To access this tab, select a network resource group in the **Network Resources** left-panel tab of the **Policy** tab.



Name

Name of the network resource group selected in the left panel.

Description

Use the **Edit** button to open a window where you can add or modify a description for the network resource group.

Type

Select the network resource type:

- Layer 2 MAC Define a group of network resource MAC addresses.
- Layer 3 IP Define a group of network resource IP addresses.

Topology


Use this drop-down list to select a network resource topology for this group. Use the configuration menu button on the right to add a new topology or edit an existing topology.

Network Resource Address List

Lists the addresses included in the selected network resource. Use the address field (IPv4 or IPv6, depending on the selected type) and select the **Add** button to add a new resource to the list.

Network Resource Topology Tab

This tab displays when you select a Network Resource Topology in the left panel of the **Network Resources** tab. It displays a list of the islands defined for the topology and the number of devices assigned to each island. See How to Create a Network Resource for more information on topologies and islands.



Name


Name of the topology island.

Device Count

The number of devices included in that island.

Network Resource Topology Island Domain Wide

The **Domain Wide** tab displays a table of information about all the devices in an island within the network resource topology selected in the left panel. To access this tab, select a network resource island in a network resource topology on the **Network Resources > Network Resource Topologies** left-panel tab on the **Policy** tab. The Domain Wide view is displayed in the right panel. To see a menu of options available for a device, right-click the device.



Name

Name of the device, or its IP address if it does not have a display name.

Device Type

Indicates the type of device. Certain devices can be listed as "Authentication Only" (supports 802.1X and RFC 3580 only; does not support Policy).

CoS Mode

Shows whether the Class of Service mode has been enabled or disabled on the device.

Firmware Version


Shows the current firmware revision for this device.

Add Devices Button

Select the Add Devices button to add devices to the network resource topology.

Details View (Network Resource Topologies Folder)

This tab displays when you select Network Resources > Network Resource Topologies in the left panel of the **Policy** tab. It displays a table of information about the network resource topologies configured in the current domain. See How to Create a Network Resource for more information on topologies.



Name

Name of the network resource topology.

Net Resc Count


The number of network resource groups using this topology.

Network Resources Using

The names of the network resource groups using this topology.

Devices (Devices)

The **Devices** tab displays a table of information about all the devices in the current domain. To access this tab, select the **Devices/Port Groups > Devices** left-panel tab on the **Policy** tab. The Details View is displayed in the right panel. To see a menu of options available for a device, right-click the device.



Name

Name of the device, or its IP address if it does not have a display name.

Device Type

Indicates the type of device. Certain devices can be listed as "Authentication Only" (supports 802.1X and RFC 3580 only; does not support Policy).

CoS Mode

Indicates whether Class of Service is enabled or disabled on the device.

Firmware Version

Shows the current firmware revision for this device.

User Sessions (Devices)

The device **User Sessions** panel displays information related to end user login sessions for a device.

This tab can be accessed in a variety of ways:

- 1. Select a device in the left-panel **Devices** tab, then select the **User Sessions** tab in the right panel.
- 2. Select the My Network navigation tree in the left panel, select a device in the Devices list, and right-click the device or open the tools menu and select **View** > **User Sessions**.
- 3. Open the **Control** > **Policy** tab, select **Devices** in the left panel, and select the **User Sessions** tab in the right panel.

User Sessions Tab

This tab displays information about each login session for the ports on the device, including the current values being collected for a session still in progress, or the final values for the last valid session when there is no session currently active.

Checking the **Show Only Active Sessions** checkbox displays only your active sessions. Deselect the checkbox to display all entries. Active sessions applied to traffic are listed in blue text. Active sessions not being applied are listed in green text.

Some devices support multiple authentication sessions simultaneously per interface. This enables a single user to authenticate via 802.1X, Web-Based, MAC, and CEP all at the same time. However, only one authentication type per interface can be *applied* at a single time. The multiuser authentication type precedence (configured on the device Authentication tab) determines which type is applied. The applied session is the one that provides the role and traffic classification information. The remaining non-applied sessions will only be used if the currently applied session is terminated. For example, if a user authenticates on a port that has multi-user authentication enabled (802.1X, Web-Based, and MAC) the active/applied session will be displayed in blue text and the other two sessions will be in green text. Another example would be if the user authenticates using the MAC authentication type but MAC authentication is disabled on the port, the session would be listed in green text. For devices that do not support multi-authentication, by definition the active session is also applied.

NOTE: Devices configured for multi-user authentication always list *only* active sessions even if the **Show**Only Active Session checkbox is deselected.

Session entries are collected up to the maximum permitted. When the maximum is reached, the oldest session entries are replaced with newer ones. The exception to this is the RoamAbout R2, where older session data is not kept.

For devices that support one authenticated user per port, only one user/current role per port appears in the table. For devices that support multiple authenticated users per port, all users authenticated on its ports are listed in the table, along with the roles under which they are authenticated.

Session Status

The status of the device.

Switch IP

The IP address or name of the device.

Switch Port

A description of the port.

Switch Alias

The alias (ifAlias) for the interface, is one is assigned.

Type

The authentication type of this login session: Web-Based, 802.1X, MAC, CEP, Quarantine, Auto Tracking, or Role Override. If Role Override is displayed, it signifies that a rule has been applied to the port, overriding the user's current role with a different role.

• Role Override (MAC) signifies that a MAC address rule has been applied to the port, overriding the Default role or any authenticated role assigned to the end user.

• Role Override (IP) signifies that an IP address rule has been applied to the port, overriding the Default role or any authenticated role assigned to an end user authenticated with Single User 802.1X. An IP Address rule will **not** override the authenticated role for any authentication type other than Single User 802.1X.

MAC Address

The MAC address of the remote user of this login session.

IP Address

For web-based authentication sessions, this column displays the IP address of the remote user of this login session.

Hostname

The hostname of the remote user of this login session. To determine the hostname, the **Policy** tab takes the IP address (when available) and uses the hostname cache on the ExtremeCloud IQ - Site Engine server. The hostname cache must be explicitly enabled by selecting the **Enable Name Resolution** checkbox in the Administration > Options > tab (by default, this option is disabled).

Role

The role under which the user authenticated on the port. If the user authenticated via RFC 3580 VLAN Authorization, this column displays the role the VLAN is mapped to (configured through Authentication-based VLAN to Role Mapping). If VLAN to Role mapping has not been configured, the port's Default role is displayed (if there is one); otherwise, the column displays "N/A."

Default VID Source

When traffic received on a port doesn't match any rules, it is assigned the default VLAN ID. This column indicates the source for the default VLAN ID:

- Policy Default Access Control The role assigned to the session defines the default VLAN ID via its Default Access Control.
- PVID If the role assigned to the session has no Default Access Control specified, then the 802.1Q PVID for the port is assigned to the traffic.

Default VID

Displays the VLAN ID that comes from the source listed in the Default VLAN ID Source column: Permit (4095), Deny (VLAN ID #), or Contain (VLAN ID #).

RFC3580 VID

If the user authenticated via RFC 3580 VLAN Authorization, this is the VLAN ID that was returned from the RADIUS server. A VLAN ID value of 0 indicates that no VLAN was assigned. If VLAN authentication is not supported on the device, this column will display "N/A."

VLAN Oper Egress

The modification that will be made to the VLAN egress list for the VLAN ID returned by the RADIUS server, if the user authenticated via RFC 3580 VLAN Authorization.

- None No modification to the VLAN egress list will be made.
- Tagged The port will be added to the list with the egress state set to Tagged (frames will be forwarded as tagged).

- Untagged The port will be added to the list with the egress state set to Untagged (frames will be forwarded as untagged).
- Dynamic The port will use information returned in the RADIUS response to modify the VLAN egress list.

If VLAN authentication is not supported on the device, this column will display "N/A."

Start Time

The time and date when the login session started.

Duration

The duration of the user's login session, in the format D + HH:MM:SS.

Auth Status

The authentication status of the login session. Possible values are:

- Authentication Successful
- Authentication Failed
- Authentication in Progress
- Authentication Server Timeout
- Authentication Terminated

Terminate Cause

The reason the login session terminated. For web-based authentication, the possible values are:

- Administratively Terminated
- Authorization Revoked
- Link Down
- Not Applicable
- Port Disabled
- Unknown Termination Cause
- User Logged Out

For 802.1X authentication, the possible values are:

- Authorization Revoked
- Client Restarted
- Link Down (or Lost Carrier)
- Not Applicable
- Port Disabled
- Port Reinitialized
- Reauthentication Failed

- Unknown Termination Cause
- User Logged Out


Authentication Server

The RADIUS server that authenticated the session.

Authentication (Device)

The device **Authentication** tab enables you to configure and change the authentication settings on the selected device. Authentication must be configured and enabled on the device in order for individual port authentication settings to take effect (see How to Configure Ports).

To access this tab, select a device in the left panel under Devices > Devices, then select the **Authentication** tab in the right panel.



Apply


Select this button to save any changes you made to the **Authentication** tab.

Refresh

Select this button to update the tab with your changes.

Authentication Status

Use this section to select the authentication mode and types used on the device.



Use the fields on the left side of this section to select the appropriate single- or multi-user authentication types. Only options supported by the selected device are available for selection. Some devices support multiple authentication types and multiple users (Multi-User Authentication) per port, while others are restricted to only one or two authentication types and single users per port. Refer to the Firmware Support matrix for information on the authentication types supported by each device type.

WARNING: Switching Authentication Types, or changing the Authentication Status from Enabled to Disabled, logs off any currently authenticated users.

Auth Type Precedence (High->Low)

This displays the order in which the authentication types are attempted on the device, with the authentication type on the left having the highest precedence (attempted first). You can edit the precedence order by selecting the field. In the Edit Precedence window, select the authentication type you want to position, and use the **Up** and **Down** buttons to arrange the types in the desired order of precedence.

WARNING: Leave the default precedence, if possible. Changing the Quarantine precedence to be lower than any other type or changing the Auto Track precedence to be higher than any other type can cause problems.

Re-Auth Timeout Action

This setting defines the action for sessions that need to be re-authenticated if the RADIUS server re-authentication request times out. Select the **Terminate** option to terminate the session or the **None** option to enable the current session to continue without disruption.

Maximum Number of Users

This setting applies to devices with Multi-User as their configured authentication type. The maximum number of users that can be actively authenticated or have authentications in progress at one time on this device. You can specify the maximum number of users per port on the port's Port Properties Authentication Configuration tab.

RFC3580 VLAN Authorization


This enables you to enable and disable RFC 3580 VLAN Authorization for the selected device. RFC 3580 VLAN Authorization must be enabled on devices in networks where the RADIUS server is configured to return a VLAN ID when a user authenticates.

When RFC 3580 VLAN Authorization is enabled:

- devices that do **not** support policy tag packets with the VLAN ID.
- devices that support both policy and Authentication-Based VLAN to Role Mapping classify packets according to the role to which the VLAN ID maps.

Current User Counts

This section enables you to specify the maximum number of users on the device and per authentication type.




Current Number of Users

For devices with Multi-User as their configured authentication type. The current number of users that are actively authenticated or have authentications in progress, or that the device is keeping authentication termination information for. Any unauthenticated traffic on the port is not included in this count

NOTE: On E1 and E6/E7 devices, if both 802.1X and MAC authentication are enabled, it is possible for the device to receive a start or response 802.1X packet while a MAC authentication is in progress. If this happens, the device immediately terminates the MAC authentication, and the 802.1X authentication proceeds to completion. Regardless of the success of the 802.1X login attempt, no new MAC authentication logins can occur on the port until 1) the link is toggled; 2) the user executes an 802.1X logout; or 3) the 802.1X session is terminated administratively.

Global Authentication Settings

This section lets you set session timeout and session idle timeout values for each authentication type.



Session Timeout

This setting represents the maximum number of seconds an authenticated session can last before automatic termination of the session. A value of zero indicates that no session timeout applies. This value can be superseded by a session timeout value provided by the authenticating server. For example,

if a session is authenticated by a RADIUS server, that server can send a session timeout value in its authentication response.

NOTE: Non-zero values are rounded to the nearest non-zero multiple of 10 by the device.

Session Idle Timeout

This displays the maximum number of consecutive seconds an authenticated session can be idle before ExtremeCloud IQ - Site Engine automatically terminates the session. A value of zero indicates that no idle timeout applies. This value can be superseded by an idle timeout value provided by the authenticating server. For example, if a session is authenticated by a RADIUS server, that server can send an idle timeout value in its authentication response.

MAC Authentication Settings

This section enables you to set up the MAC password for MAC authentication. In order for MAC authentication to work, you must also configure the RADIUS server with the MAC password as well as the MAC addresses which are permitted to authenticate.



Set Password/Mask

Select this checkbox to set a password and mask for MAC authentication.

MAC User Password

The password passed to the RADIUS server for MAC authentication.

MAC Mask

You can select a mask to provide a way to authenticate end-systems based on a portion of their MAC address. For example, you could specify a mask that would base authentication on the manufacturers ID portion of the MAC address. The MAC Mask is passed to the RADIUS server for authentication after the primary attempt to authenticate using the full MAC address fails.

MAC Address Delimiter

The character used between octets in a MAC address:

- **Hyphen** A hyphen is used as a delimiter in the MAC address (e.g. xx-xx-xx-xx-xx).


Web Authentication Settings

For users of web-based authentication, this tab lets you specify web authentication parameters using three sections:

- General
- Guest Networking
- Web Login

General

The General section lets you specify the URL of the authentication web page and the IP address of the system where it resides. It also lets you enable certain web authentication features, such as Enhanced Login Mode, on devices that support those features.



Enhanced Login Mode

Enabling the Enhanced Login Mode causes the authentication web page to be displayed regardless of whether the URL or IP address entered into the browser by the end user is the designated Web Authentication URL or IP address. This option is grayed out if the device does not support the mode.

Enhanced Mode Redirect Time(s)

This setting applies for devices with <u>Enhanced Login Mode</u> enabled. It specifies the amount of time (in seconds) before the end-user is redirected from the authentication web page to their requested URL.

An end-system using DHCP requires time to transition from the temporary IP address issued by the authentication process to the official IP address issued by the network. **Enhanced Mode Redirect Time** specifies the amount of time permitted for the end-system to complete this process and begin using its official IP address.

For example, if an end-user (in **Enhanced Login Mode** and a **Redirect Time** of **30 seconds**) enters the URL of "http://ExtremeNetworks.com", the user is presented the authentication web page. When the

user successfully authenticates into the network, the user sees a login success page that displays "Welcome to the Network. Completing network connections. You will be redirected to http://ExtremeNetworks.com in approximately 30 seconds."

WINS/DNS Spoofing

This setting enables you to enable and disable WINS/DNS spoofing for the selected device. Spoofing enables the end-user to resolve the Web Authentication URL name to the IP address using WINS/DNS. The default is Disabled. This option is grayed out if not supported by the device.

Logo Display Status

Specifies whether the Extreme Networks logo is displayed or hidden on the authentication web page window. This option is grayed out if not supported by the device.

Authentication Protocol

This setting is the authentication protocol being used (PAP or CHAP). PAP (Password Authentication Protocol) provides an automated way for a PPP (Point-to Point Protocol) server to request the identity of user, and confirm it via a password. CHAP (Challenge Handshake Authentication Protocol), the more secure of the two protocols, provides a similar function, except that the confirmation is accomplished using a challenge and response authentication dialog.

Web Authentication URL

This is the URL for your authentication web page. Users wishing to receive network services access the web page from a browser using this URL. The http:// is supplied. Alphabetical characters, numerical characters and dashes are permitted as part of the URL, but dots are not. The URL needs to be mapped to the Web Authentication IP address in DNS or in the hosts file of each client. It must be resolvable via DNS/WINS, either on the device or at corporate, assuming the Web Authentication mapping has been set up on the corporate DNS/WINS service. This option is grayed out if not supported by the device.


Web Authentication IP Address

This is the IP address of your authentication web page server. If you have specified a Web Authentication URL, the IP address needs to be mapped to the URL in DNS or in the host file of each client.

Guest Networking

The **Guest Networking** section lets you configure guest networking, a feature that enables any user to access the network and obtain a guest policy without having to know a username or password. The user accesses the authentication web page, where the username and password fields are automatically filled in, enabling them to log access as a guest. If the user does not want to log in as a guest, they can type in their valid username and password to log in.

NOTE: Guest networking is designed for networks using web-based authentication, with <u>port mode</u> set to Active/Discard.



Guest Networking Status

Use the drop-down list to specify guest networking status:

- **Disable** Guest networking is unavailable.
- Local Auth Guest Networking is enabled. The user accesses the authentication web page
 where the username field is automatically filled in with the specified <u>Guest Name</u>. When the user
 submits the web page using this guest name, the default policy of that port becomes the active
 policy. The port mode must be set to Active/Discard mode.
- RADIUS Auth Guest Networking is enabled. The user accesses the authentication web page, where the username field is automatically filled in with the specified <u>Guest Name</u>, and the password field is masked out with asterisks. When the user submits the web page using these credentials, the value of the <u>Guest Password</u> is used for authentication. Following successful authentication from the RADIUS server, the port applies the policy (role) returned from the RADIUS server. The port mode must be set to Active/Discard mode.

Guest Name


The username that Guest Networking uses to authenticate users. The guest name is displayed automatically on the authentication web page. If the user does not want to log in as a guest, they can type in their valid username to override the guest username.

Guest Password

The password that Guest Networking uses to authenticate users when RADIUS Auth is selected.

Web Page Banner

The Web Page Banner section enables you to customize the banner end users see at the top of the authentication web page and set a Redirect Time, if applicable.



Web Page Banner

Use this area to create a banner end users see at the top of the authentication web page. For example, you might include your company name and information on what to do if the user has questions or problems. Because this banner also appears in messages that occur during successful login and failed authentication, as well as on the "Radius Busy" screen, it is not appropriate to include "Welcome to [Your Company]" in the banner.

The **Default** button enables you to reset the banner to default text provided in a text file (pwa_banner.txt). Initially, the default banner text is the Extreme Networks contact information. However, you can customize the text for your network by editing the pwa_banner.txt file, located in the top level of the Policy Manager install directory. Then, when you select the Default button, the new text will be displayed in the Web Page Banner area.

Convergence End-Point Settings


This section provides a way to identify Convergence End-Points (IP phones) connecting to the device, and apply a role to the end-point based on the type of end-point detected. The CEP Detection section lets you create detection rules for identifying the end-points, and the CEP Role Mappings section lets you map a role to each CEP product type.

In addition to configuring CEP on the device, you must also enable CEP protocols on each port using the CEP Access section in the Port Authentication Tab. After you have configured CEP on the device and each port, you can monitor CEP usage on the Port Usage Tab (Port) or Port Usage Tab (Device).

CEP Role Mappings

This section lets you select the CEP product types supported on the device, and map a role for each type. Then, when a convergence end-point (such as an IP phone) connects to the network,

the device identifies the type of end-point (using CEP detection rules) and applies the assigned role.



CEP Type

Lists the CEP types supported by the device.

Role

Lists the role mapped to each CEP Type.

Add

Select a CEP Type and select the **Add** button to open the Add Role Mapping window, where you can select a role for the selected **CEP Type**. Your selections are added to the CEP Role Mappings list.

Remove

Select the CEP Type and select Remove to remove the CEP Type in the CEP Role Mappings list.

CEP Detection Tab


Use this section to create CEP detection rules used to determine if a connecting end-system is a CEP device and the type of CEP device. This enables ExtremeCloud IQ - Site Engine to assign the appropriate role to the port based on the type of CEP device detected.

NOTE: CEP detection rules apply only to Siemens, H.323, and SIP (Session Initiation Protocol) phone detection. Cisco detection uses CiscoDP as its detection method.

CEP detection rules are based on two detection methods:

- TCP/UDP Port Number detection Many CEP vendors use specific TCP/UDP port numbers for call setup on their IP phones. You can create detection rules that identify CEP devices based on specific TCP/UDP port numbers. By default, Siemens Hi-Path phones are detected on TCP/UDP port 4060.
- IP Address detection H.323 phones use a reserved IP multicast address and UDP port number for call setup. You can create detection rules to detect an IP phone based on its IP address in combination with an IP address mask. By default, H.323 phones are detected using the multicast address 224.0.1.41 and the TCP/UDP ports 1718, 1719, and 1720. SIP phones are detected using the multicast address 224.0.1.75

and the TCP/UDP port 5060. H.323 and SIP phones are also detected using only their respective multicast addresses without the TCP/UDP ports.



Priority

The rule priority with one (1) being the highest priority. The rule with the highest priority is used first, so it is recommended the highest priority be given to the predominate protocol in the network to provide for greater efficiency.

Address

If the rule is based on IP address detection, this field displays the IP address that incoming packets matched against. By default, H.323 uses 224.0.1.41 as its IP address, SIP uses 224.0.1.75 as its IP address, and Siemens has no IP address configured.

Address Mask

If the rule is based on IP address detection, this field displays the IP address mask against which incoming packets are matched.

End Point Type

Specifies the end-point type assigned (H.323, Siemens, or SIP) if incoming packets match this rule.

Protocol

If the rule is based on TCP/UDP port detection, this field displays the protocol type used for matching, using a port range defined with the Port Low and Port High values:

- UDP + TCP Match the port number for both UDP and TCP frames.
- TCP Match the port number only for TCP frames.
- UDP Match the port number only for UDP frames.

Port Low

The low end of the port range defined for detection on UDP and/or TCP ports.

Port High

The high end of the port range defined for detection on UDP and/or TCP ports.

Add

Opens the Add/Edit CEP Detection Rule window where you can create CEP detection rules.

Remove

To remove a CEP detection rule, select the entry and select Remove.

Edit

To edit a CEP detection rule, select the rule and select **Edit**. The Add/Edit CEP Detection Rule window opens where you edit the rule's parameters. You can also double-click an entry in the table to open the edit window.

Add/Edit CEP Detection Rule

Use this window to add or edit CEP detection rules that are used to determine if a connecting end-system is a CEP device, and what type of CEP device it is. This allows Policy Manager to assign the appropriate role to the port based on the type of CEP device detected. Access the window from the CEP Detection sub-tab in the right-panel Device Authentication tab.

NOTE: CEP detection rules apply only to Siemens, H.323, and SIP (Session Initiation Protocol) phone detection. Cisco detection uses CiscoDP as its detection method.

CEP detection rules are based on two detection methods:

- TCP/UDP Port Number detection Many CEP vendors use specific TCP/UDP port numbers for call setup on their IP phones. You can create detection rules that identify CEP devices based on specific TCP/UDP port numbers. By default, Siemens Hi-Path phones are detected on TCP/UDP port 4060.
- IP Address detection H.323 phones use a reserved IP multicast address and UDP port number for call setup. You can create detection rules detect an IP phone based on its IP address in combination with an IP address mask. By default, H.323 phones are detected using the multicast address 224.0.1.41 and the TCP/UDP ports 1718, 1719, and 1720. SIP phones are detected using the multicast address 224.0.1.75 and the TCP/UDP port 5060. H.323 and SIP phones are also detected using only their respective multicast addresses without the TCP/UDP ports.



CEP Detection Settings

Priority

Enter the rule priority with one (1) being the highest priority. The rule with the highest priority is used

first, so it is recommended the highest priority be given to the predominate protocol in the network to provide for greater efficiency.

IP Address

If the rule is based on IP address detection, enter the IP address against which incoming packets are matched. By default, H.323 uses 224.0.1.41 as its IP address, SIP uses 224.0.1.75 as its IP address, and Siemens has no IP address configured.

Address Mask

If the rule is based on IP address detection, enter the IP address mask against which incoming packets are matched.

End Point Type

Select the endpoint type (H.323, Siemens, or SIP) assigned to incoming packets that match this rule.

Protocol

If the rule is based on TCP/UDP port detection, select the UDP and/or TCP checkbox and define a port range with Port Low and Port High values:

- UDP and TCP Match the port number for both UDP and TCP frames.
- TCP Match the port number only for TCP frames.
- UDP Match the port number only for UDP frames.

Port Low

Define the low end of the port range for detection on UDP and/or TCP ports.

Port High


Define the high end of the port range for detection on UDP and/or TCP ports.

_

Ports (Authentication)

The **Ports (Authentication)** tab allows you to configure and change the authentication settings for a port. Authentication must be configured and enabled on the device in order for individual port authentication settings to take effect. Only those areas of the tab that relate to the authentication type configured on the device are available for editing.

To access the **Ports (Authentication)** tab, select a device in the left-panel **Devices > Devices** tab, then select **Authentication > Ports** in the right panel.



Select a port in the top section to display and configure the authentication settings for that port in the bottom of the window.

Select the Apply button at the top of the window to save changes to this tab.


The Authentication Configuration tab has six sections:

- Authentication Mode
- RFC3580 VLAN Authorization
- Login Settings

- Automatic Re-Authentication
- Authenticated User Counts
- Convergence End-Point Access

Authentication Mode

This tab displays general authentication and port mode information about the port.



This area displays the current port mode for the port, and allows you to change the settings if desired. Port mode defines whether or not a user is required to authenticate on a port, and how unauthenticated traffic is handled. It is a combination of Authentication Behavior (whether or not authentication is enabled on the port), and Unauthenticated Behavior (whether unauthenticated traffic is assigned to the port's default role or discarded). See Port Mode for a complete description of each port mode.

In addition, this section provides checkboxes that allow you to disable a specific authentication type at the port level.

Port Mode (Auth/Unauth Behavior)

Select an option to specify whether or not authentication is enabled on the port. (See <u>Port Mode</u> for more information.)

NOTE: Authentication Behavior must be set to **Active** for authentication to be allowed using CEP Protocols.

Disable 802.1X Auth

Select this checkbox to disable 802.1X authentication at the port level. If the device is only configured with 802.1X authentication, selecting this checkbox results in the port Authentication Behavior being set to **Inactive**.

NOTE: For Single User 802.1X+MAC authentication with Active/Default Role as the selected port mode: Disabling 802.1X authentication also disables MAC authentication on the port. An end user connecting to the port is not able to authenticate via 802.1X or MAC. The port behaves as if Inactive/Default Role is the selected port mode.

Disable Web-Based Auth

Select this checkbox to disable web-based authentication at the port level. If the device is only

configured with web-based authentication, selecting this checkbox results in the port Authentication Behavior being set to **Inactive**.

NOTE: For Multi-User Web-Based authentication with Active/Discard as the selected port mode: This checkbox is automatically selected because multi-user web-based authentication does not support the Active/Discard port mode.

Disable MAC Auth

Select this checkbox to disable MAC authentication at the port level. If the device is only configured with MAC authentication, selecting this checkbox results in the port Authentication Behavior being set to **Inactive**.

Disable Quarantine Auth

Select this checkbox to disable Quarantine authentication at the port level. If the device is only configured with Quarantine authentication, selecting this checkbox results in the port Authentication Behavior being set to **Inactive**.

Disable Auto Tracking Auth


Select this checkbox to disable MAC authentication at the port level. If the device is only configured with Auto Tracking authentication, selecting this checkbox results in the port Authentication Behavior being set to **Inactive**.

RFC3580 VLAN Authorization

This section lets you enable or disable RFC 3580 VLAN Authorization on the port and specify an egress state. RFC 3580 VLAN Authorization must be enabled in networks where the RADIUS server has been configured to return a VLAN ID when a user authenticates. When RFC 3580 VLAN Authorization is enabled:

- ports on devices that do not support policy, will tag packets with the VLAN ID.
- ports on devices that do support policy and also support Authentication-Based VLAN to Role Mapping, will classify packets according to the role that the VLAN ID maps to.

You can also enable and disable VLAN Authorization at the device level using the device Authentication tab. If the device does not support RFC 3580, this tab will be grayed out.



VLAN Authorization Status

Allows you to enable and disable RFC 3580 VLAN Authorization for the selected port. This option is grayed out if not supported by the device.


VLAN Authorization Admin Egress

Allows you to modify the VLAN egress list for the VLAN ID returned by the RADIUS server when a user authenticates on the port:

- None No modification to the VLAN egress list is made.
- Tagged The port is added to the list with the egress state set to Tagged (frames are forwarded as tagged).
- Untagged The port is added to the list with the egress state set to Untagged (frames are forwarded as untagged).
- Dynamic The port uses information returned in the RADIUS response to modify the VLAN egress list. This value is supported only if the device supports a mechanism through which the egress state may be returned in the RADIUS response.

Login Settings

This tab displays the current login settings for the port and allows you to change the settings if desired. The options available depend on what type(s) of authentication are enabled on the device.



MAC

Hold Time (sec)

Amount of time (in seconds) authentication remains timed out after the user fails to login. Valid values are 0-65535. The default is 60. (Hold Time is also known as Quiet Period in web-based and MAC authentication.)

802.1X

Hold Time (sec)

Amount of time (in seconds) authentication remains timed out after the user fails to login. Valid values are 0-65535. The default is 60.

Auth request period (sec)

For 802.1X authentication, how often (in seconds) the device queries the port to see if there is a new user on it. If a user is found, the device then attempts to authenticate the user. Valid values are 1-65535. The default is 30.

User timeout (sec)

For 802.1X authentication, the amount of time (in seconds) the device waits for an answer when querying the port for the existence of a user. Valid values are 1-300. The default is 30.

Auth server timeout (sec)

For 802.1X authentication, if a user is found on the port, the amount of time (in seconds) the device waits for a response from the authentication server before timing out. Valid values are 1-300. The default is 30.

Handshake requests before failure

For 802.1X authentication, the number of times the device tries to finalize the authentication process with the user, before the authentication request is considered invalid and authentication fails. Valid values are 1-10. The default is 2.

Web Auth

Max Requests

Number of times a user can attempt to log in before authentication fails and login attempts are not allowed. For web-based authentication, valid values are 1-2147483647, zero is not allowed, and the default is 2.

Hold Time (sec)

Amount of time (in seconds) authentication remains timed out after the specified **Max Requests** is reached. Valid values are 0-65535. The default is 60.

Quarantine

Session Timeout (sec)

For Quarantine authentication, the maximum number of seconds an authenticated session may last before automatic termination of the session. A value of zero indicates that no session timeout applies.

Session Idle Timeout (sec)

For Quarantine authentication, the maximum number of consecutive seconds an authenticated session may be idle before automatic termination of the session. A value of zero indicates that the device level setting is used.

Auto Tracking

Session Timeout (sec)

For Auto Tracking sessions, the maximum number of seconds a session may last before automatic termination of the session. A value of zero indicates that the device level setting is used.

Session Idle Timeout (sec)

For Auto Tracking sessions, the maximum number of consecutive seconds a session may be idle before automatic termination of the session. A value of zero indicates that the device level setting is used.

Automatic Re-Authentication

This tab is grayed-out if only web-based authentication is enabled on the device. For 802.1X and MAC authentication, the Automatic Re-Authentication tab lets you set up the periodic automatic re-authentication of logged-in users on this port. Without disrupting the user's session, the device repeats the authentication process using the most recently obtained user login information, to see if the same user is still logged in. Authenticated logged-in users are not required to log in again for re-authentication, as this occurs "behind the scenes."



802.1X Re-auth Status

If **Enabled** is selected, the re-authentication feature is enabled. If **Disabled** is selected, the re-authentication feature is disabled.

802.1X Re-auth Frequency (sec)

The length of time (in seconds) the device checks the port to re-authenticate the logged in user. Valid values are 1-2147483647. The default is 3600.

MAC Re-auth Status

If **Enabled** is selected, the re-authentication feature is enabled. If **Disabled** is selected, the re-authentication feature is disabled.

MAC Re-auth Frequency (sec)

The length of time (in seconds) the device checks the port to re-authenticate the logged in user. Valid values are 1-2147483647. The default is 3600.

Authenticated User Counts

This section provides authenticated user count information for devices with Multi-User as their configured authentication type. See the device Authentication tab for information on setting the device authentication type.



Current Number of Users

The current number of users actively authenticated or are in the process of authenticating on this interface. If multi-user authentication is disabled, this number is 0 (zero). Any unauthenticated traffic on the port is not included in this count.

Number of Users Allowed

The maximum number of users that can actively authenticate or be in the process of authenticating at one time on this interface. If you set this value below the current number of users, end user sessions exceeding that number are terminated.

NOTE: B2/C2 Devices. If you are configuring a single user and an IP phone per port, set this value to 2.

Number of MAC Users Allowed

The number of users that can actively authenticate via MAC authentication, or be in the process of authenticating via MAC authentication at one time on this interface. The number of MAC users allowed cannot exceed the number of users allowed. If you set this value below the current number of users, end user sessions exceeding that number are terminated. If MAC is not selected as a Multi-User authentication type on the device Authentication tab, this field is grayed out.

Number of Quarantine Users Allowed

The number of users that can be actively authenticated via Quarantine authentication, or have Quarantine authentications in progress at one time on this interface. The number of Quarantine users allowed cannot exceed the number of users allowed. If you set this value below the current number of users, end user sessions exceeding that number are terminated. If Quarantine Auth is not enabled on the device Authentication tab, this field is grayed out.

Number of Auto Tracking Users Allowed

The number of Auto Tracking users that can be actively authenticated or have authentications in progress at one time on this interface. The number of Auto Tracking users allowed cannot exceed the number of users allowed. If you set this value below the current number of users, end user sessions exceeding that number are terminated. If Auto Tracking is not enabled on the device Authentication tab, this field is grayed out.

Convergence End-Point Access

This section lists all the Convergence End-Point (CEP) protocols supported by the device that the port resides on, and lets you enable or disable them for that port. For devices that do not support CEP, the section is blank.



Enable Button

Selects all the checkboxes and enables all the CEP protocols for this port.

Disable All Button

Deselects all the checkboxes and disables all the CEP protocols for this port.

CEP Protocols List


Lists all the CEP protocols supported by the device on which the port resides. Highlight a CEP protocol and select the Enable or Disable button to enable or disable CEP protocols, respectively. If the device does not support the CEP feature, this area is blank.

RADIUS (Device)

The device **RADIUS** tab allows you to configure and enable communication between the selected device (the RADIUS client), a RADIUS server or servers, and ExtremeCloud IQ - Site Engine, for the purposes of authentication and accounting.

RADIUS accounting collects various data and statistics, such as the length of time a user has been logged on, and makes that data available to an administrator. It is used by a device to save accounting data on a RADIUS server. The device sends accounting requests to the server. The server acknowledges these requests, and data is passed to the server via accounting updates. For more information on accounting functionality, refer to your RADIUS server documentation.

To display the device RADIUS tab, select a device in the left-panel **Devices** tab, then select the RADIUS tab in the right panel.



Authentication Tab

Use this tab to view and configure the RADIUS authentication servers with which the device (the RADIUS client) can communicate.

RADIUS Authentication Client Settings

This section lets you enable or disable communication between the selected device (the RADIUS client) and the RADIUS authentication servers, and specify connection attempt information.

Authentication Status

Allows you to enable and disable communication between this device and the RADIUS authentication server(s). If enabled, the device becomes a RADIUS client and communicates with a RADIUS authentication server whenever a user logs on to a port on the device, as long as the port itself is enabled for authentication and the device is set up as a client on the RADIUS authentication server. The default is Disabled. For ExtremeWireless devices, the Client Status is automatically set to Enabled when a RADIUS server exists and Disabled when it does not.

Management Access Auth Status Override

Allows you to override the Authentication Status for users accessing the RADIUS authentication server (s) that have requested management access via the console, Telnet, SSH, or HTTP, etc.

Network Access Auth Status Override

Allows you to override the Authentication Status for users accessing the network via 802.1X, MAC, or Web-Based authentication.

Number of Retries

The number of attempts the device will make in contacting each RADIUS authentication server before giving up and trying the next RADIUS authentication server on the list. Valid values are 1-65535. For ExtremeWireless devices, this value is entered when the RADIUS server is added.

Timeout Duration

The total number of seconds the device will wait for the RADIUS authentication server to respond, before trying again. Valid values are 1-65535. For ExtremeWireless devices, this value is entered when the RADIUS server is added.

Management Access Timeout Duration Override (sec)

The total number of seconds the device waits for the RADIUS authentication server to respond before trying again for users accessing the RADIUS authentication server(s) that have requested management access via the console, Telnet, SSH, or HTTP, etc.

Network Access Timeout Duration Override (sec)

The total number of seconds the device waits for the RADIUS authentication server to respond before trying again for users accessing the network via 802.1X, MAC, or Web-Based authentication.

Response Mode

Select the RADIUS response attribute that the device should use for authentication:

- Filter ID The Filter ID (role) is used. If a VLAN Tunnel Attribute (VTA) is returned, it will be ignored.
- VLAN Tunnel Attribute The VLAN Tunnel Attribute is used and the Authentication-Based VLAN to Role Mappings are applied, if present. If a Filter ID is returned, it will be ignored.
- Filter ID With VLAN Tunnel Attribute Both attributes are applied in the following manner: the role is applied to the user, except that the VLAN Tunnel Attribute replaces the role's Default Access Control VLAN (if present). In this case, the Authentication-Based VLAN to Role mappings are ignored (as the role was explicitly assigned). VLAN classification rules are still applied, as defined by the assigned role.

Retransmit Algorithm

Select the authentication retransmission algorithm for this device to use with your RADIUS servers. Devices that do not support this functionality will have the option grayed out.

- Standard Specifies that the primary RADIUS server should always be used for authentication, if it is
 available. The standard RADIUS authentication algorithm focuses on using RADIUS servers for
 redundancy rather than for scale provisioning. The only time secondary RADIUS servers are used, is
 when the primary server is unreachable due to a network outage or because server capacity is
 exceeded.
- Round-Robin The round-robin RADIUS authentication algorithm spreads RADIUS server usage evenly between available RADIUS servers, allowing the load balancing of a large number of authentications across all RADIUS servers. This allows for a maximum authentication throughput for the number of servers configured. Additionally, if a single server is down, only a portion of the authenticating sessions will be affected by the outage.
- Sticky Round-Robin This algorithm uses round-robin when assigning a RADIUS server to each unique authentication session, but specifies that the same RADIUS server should be used for any given authentication session when a session is initiated. In large-scale ExtremeControl deployments, this algorithm is used for switches that are authenticating more users than an ExtremeControl engine supports. For example, an ExtremeControl deployment might have an S-Series device that supports 9000 users deployed at the distribution level and authenticating users to three ExtremeControl engines that support 3000 users each. In this scenario, the sticky round-robin algorithm allows the S-Series device to spread the load across all three ExtremeControl engines while using the same ExtremeControl engine for all RADIUS transactions for a given session (MAC address).

Apply Button

Applies the changes you made in the RADIUS Authentication Client Settings section.

Authentication RADIUS Server(s) Table

This table lists the RADIUS authentication servers with which the device (the RADIUS client) can communicate. Use the buttons to add or remove servers, and edit server parameters. You can also edit a server's parameters by double-clicking the server entry in the list.

Priority

Order in which the RADIUS authentication server is checked, as compared to the other RADIUS authentication servers listed here. The lower the number, the higher the priority.

RADIUS Server IP

IP address of the RADIUS authentication server.

Client UDP Port

UDP port number (1-65535) on the RADIUS authentication server that the device will send authentication requests to; 1812 is the default port number.

Access Type

The type of authentication access allowed for this RADIUS server:

- Any access the server can authenticate users originating from any access type.
- Management access the server can only authenticate users that have requested management access via the console, Telnet, SSH, or HTTP, etc.

• **Network access** — the server can only authenticate users that are accessing the network via 802.1X, MAC, or Web-Based authentication.

Devices that do not support this feature will display N/A in this column.

Current Sessions

The current number of sessions associated with this server when the device is using the <u>sticky round-robin RADIUS authentication algorithm</u>. This value is not used when other algorithms are being used.

Max Sessions

The maximum number of sticky round-robin authentication sessions allowed on the server when the sticky round-robin RADIUS authentication algorithm is configured for the device. This value is not used when other algorithms are being used. In sticky round-robin, if a MAC address needs to re-authenticate, the request is sent to the same RADIUS server as the initial authentication request, unless the current number of authentication sessions for the server has reached the specified Max Sessions value. When this value is reached, re-authentication requests will instead default to the standard round-robin behavior to determine which RADIUS server to send the request to.

Number of Retries

The number of times the device will resend an authentication request if the RADIUS authentication server does not respond. For ExtremeWireless devices, this value is configured per RADIUS server. For all other devices, this value is global to all RADIUS servers, and is specified per device (Client Default) in the RADIUS Authentication Client Settings section.

Timeout Duration

The amount of time in seconds the device will wait for the RADIUS authentication server to respond to an authentication request. For ExtremeWireless devices, this value is configured per RADIUS server. For all other devices, this value is global to all RADIUS servers, and is specified per device (Client Default) in the RADIUS Authentication Client Settings section.

Management Interface

The IP address and VRName used when the switch is communicating with a configured RADIUS server.

Apply Button

Applies any changes you made in the RADIUS Authentication Server(s) tab.

Add Button

Opens the Add RADIUS Authentication Server window, where you can enter the parameters for a server you want to add to the list. When you select **OK** on this window, the new server is added.

Remove Button


Select a RADIUS authentication server in the list and use this button to remove the server.

Edit Button

Select a RADIUS authentication server in the list and use this button to edit the server's parameters. You can also edit the server parameters by double-clicking the server entry in the list.

Accounting Tab

Use this tab to view and configure the RADIUS accounting servers with which the device (the RADIUS client) can communicate.



RADIUS Accounting Client Settings

This section lets you enable or disable communication between the selected device (the RADIUS client) and the RADIUS accounting servers, and specify the update interval.

Accounting Status

Allows you to enable or disable RADIUS accounting. RADIUS accounting is used by a device to save accounting data on a RADIUS accounting server. If accounting is enabled, an accounting session starts after the user is successfully authenticated by a RADIUS authentication server. The default is Disabled. For ExtremeWireless devices, the status is automatically set to Enabled when a RADIUS server exists and Disabled when it does not. Devices that do not support RADIUS accounting will have this field grayed out.

Management Access Auth Status Override

Allows you to override the Accounting Status for users accessing the RADIUS accounting server(s) that have requested management access via the console, Telnet, SSH, or HTTP, etc.

Network Access Auth Status Override

Allows you to override the Accounting Status for users accessing the network via 802.1X, MAC, or Web-Based authentication.

Per Authentication Type Accounting Status

Allows you to enable/disable RADIUS accounting for individual authentication types. Some authentication types do not have RADIUS accounting enabled by default (when global RADIUS accounting is enabled). Enabling these authentication types will give both ExtremeControl and other RADIUS servers more complete information regarding authentication sessions. These options also allow you to disable accounting messages from certain authentication types, for example, Auto-Tracking, which does not actually authenticate end users. Note that the global <u>Accounting Status</u> option controls accounting on a global basis for all authentication types. Devices that do not support this functionality will have these fields grayed out.

Update Interval (minutes)

Collected accounting data is sent from the device to the RADIUS accounting server via accounting updates. The Accounting Update Interval is the amount of time in minutes between accounting updates. Valid values are 1-65535. It is recommended that the value be greater than 10 minutes, and careful consideration should be given to its impact on network traffic. Devices that do not support RADIUS accounting have this field grayed out (with the exception of an SNMPv1 R2 device, which display accounting values but will not allow you to set them.) For ExtremeWireless devices, this value is entered when the RADIUS server is added.

Management Access Timeout Duration Override (sec)

The total number of seconds the device waits for the RADIUS accounting server to respond before trying again for users accessing the RADIUS accounting server(s) that have requested management access via the console, Telnet, SSH, or HTTP, etc.

Network Access Timeout Duration Override (sec)

The total number of seconds the device waits for the RADIUS accounting server to respond before trying again for users accessing the network via 802.1X, MAC, or Web-Based authentication.

Apply Button

Applies the changes you made in the RADIUS Accounting Client Settings section.

Accounting RADIUS Servers Table

This tab lists the RADIUS accounting servers with which the device (the RADIUS client) can communicate. Use the buttons to add or remove servers, and edit server parameters. You can also edit a server's parameters by double-clicking the server entry in the list.

Priority

Order in which the RADIUS accounting server is checked, as compared to the other RADIUS accounting servers listed here. The lower the number, the higher the priority.

RADIUS Server IP

IP address of the RADIUS accounting server.

Client UDP Port

UDP port number (1-65535) on the RADIUS accounting server that the device will send accounting requests to; 1813 is the default port number. Devices that do not support RADIUS accounting will display N/A in this column (with the exception of an SNMPv1 R2 device, which will display accounting values but will not allow you to set them.)

Access Type

The type of authentication access allowed for this RADIUS server:

- Any access the server can authenticate users originating from any access type.
- Management access the server can only authenticate users that have requested management access via the console. Telnet. SSH. or HTTP. etc.
- **Network access** the server can only authenticate users that are accessing the network via 802.1X, MAC, or Web-Based authentication.

Devices that do not support this feature will display N/A in this column.

Number of Retries

The number of times the device will resend an accounting request if the RADIUS accounting server does not respond. Valid values are 0-20. Devices that do not support RADIUS accounting will display N/A in this column (with the exception of an SNMPv1 R2 device, which display accounting values but does not allow you to set them.)

Timeout Duration

The amount of time in seconds the device will wait for the RADIUS accounting server to respond to an accounting request. Valid values are 2-10 seconds. Devices that do not support RADIUS accounting will display N/A in this column (with the exception of an SNMPv1 R2 device, which display accounting values but does not allow you to set them.)

Update Interval

The amount of time in minutes between accounting updates. For ExtremeWireless devices, this value is configured per RADIUS server. For all other devices, this value is global to all RADIUS servers, and is specified per device (Client Default) in the RADIUS Accounting Client Settings section.

Management Interface

The IP address and VRName used when the switch is communicating with a configured RADIUS server.

Apply Button

Applies any changes you made in the RADIUS Accounting Server(s) tab.

Add Button

Opens the Add RADIUS Accounting Server window, where you can enter the parameters for a server you want to add to the list. When you select **OK** on this window, the new server is added.

Remove Button

Select a RADIUS accounting server in the list and use this button to remove the server.


Edit Button

Select a RADIUS accounting server in the list and use this button to edit the server's parameters. You can also edit the server parameters by double-clicking the server entry in the list.

RADIUS Authentication (Device)

The device RADIUS **Authentication** tab enables you to configure and enable communication between the selected device (the RADIUS client), a RADIUS server or servers, and ExtremeCloud IQ - Site Engine, for the purposes of authentication and accounting (for your SNMPv3 devices that support it).

Use this tab to view and configure the RADIUS authentication servers with which the device (the RADIUS client) can communicate.



RADIUS Authentication Client Settings

This section lets you enable or disable communication between the selected device (the RADIUS client) and the RADIUS authentication servers, and specify connection attempt information.

Authentication Status

Enables you to enable and disable communication between this device and the RADIUS authentication server(s). If enabled, the device becomes a RADIUS client and communicates with a RADIUS authentication server whenever a user logs on to a port on the device, as long as the port itself is enabled for authentication and the device is set up as a client on the RADIUS authentication server. For ExtremeWireless devices, the Client Status is automatically set to **Enabled** when a RADIUS server exists and Disabled when it does not.

Management Access Auth Status Override

Enables you to override the Authentication Status for users accessing the RADIUS authentication server (s) that requested management access via the console, Telnet, SSH, or HTTP, etc.

Network Access Auth Status Override

Enables you to override the Authentication Status for users accessing the network via 802.1X, MAC, or Web-Based authentication.

Number of Retries

The number of attempts the device makes in contacting each RADIUS authentication server before giving up and trying the next RADIUS authentication server on the list. For ExtremeWireless devices, this value is entered when the RADIUS server is added.

Timeout Duration (seconds)

The total number of seconds the device waits for the RADIUS authentication server to respond, before trying again. For ExtremeWireless devices, this value is entered when the RADIUS server is added.

Management Access Timeout Duration Override (sec)

The total number of seconds the device waits for the RADIUS authentication server to respond before trying again for users accessing the RADIUS authentication server(s) that requested management access via the console, Telnet, SSH, or HTTP, etc.

Network Access Timeout Duration Override (sec)

The total number of seconds the device waits for the RADIUS authentication server to respond before trying again for users accessing the network via 802.1X, MAC, or Web-Based authentication.

Response Mode

Select the RADIUS response attribute the device uses for authentication:

- Filter ID (Discard VTA) The Filter ID (role) is used. If a VLAN Tunnel Attribute (VTA) is returned, it is ignored.
- VLAN Tunnel Attribute (Discard Tunnel Attribute) The VLAN Tunnel Attribute is used and the
 Authentication-Based VLAN to Role Mappings are applied, if present. If a Filter ID is returned, it is
 ignored.
- Filter ID With VLAN Tunnel Attribute Both attributes are applied in the following manner: the role is applied to the user, except that the VLAN Tunnel Attribute replaces the role's Default Access Control VLAN (if present). In this case, the Authentication-Based VLAN to Role mappings are ignored (as the role was explicitly assigned). VLAN classification rules are still applied, as defined by the assigned role.

Retransmit Algorithm

Select the authentication retransmission algorithm for this device to use with your RADIUS servers. Devices that do not support this functionality have the option grayed out.

- Standard Specifies that the primary RADIUS server should always be used for authentication, if it is available. The standard RADIUS authentication algorithm focuses on using RADIUS servers for redundancy rather than for scale provisioning. The only time secondary RADIUS servers are used, is when the primary server is unreachable due to a network outage or because server capacity is exceeded.
- Round-Robin The round-robin RADIUS authentication algorithm spreads RADIUS server usage evenly between available RADIUS servers, enabling the load balancing of a large number of authentications across all RADIUS servers. This enables a maximum authentication throughput

for the number of servers configured. Additionally, if a single server is down, only a portion of the authenticating sessions are affected by the outage.

• Sticky Round-Robin — This algorithm uses round-robin when assigning a RADIUS server to each unique authentication session, but specifies that the same RADIUS server is used for any given authentication session when a session is initiated. In large-scale ExtremeControl deployments, this algorithm is used for switches authenticating more users than an ExtremeControl appliance supports. For example, an ExtremeControl deployment might have an S-Series device that supports 9000 users deployed at the distribution level and authenticating users to three ExtremeControl appliances that support 3000 users each. In this scenario, the sticky round-robin algorithm enables the S-Series device to spread the load across all three ExtremeControl appliances while using the same ExtremeControl appliance for all RADIUS transactions for a given session (MAC address).

Apply Button

Applies the changes you made in the RADIUS Authentication Client Settings section.

Authentication RADIUS Server(s) Table

This table lists the RADIUS authentication servers with which the device (the RADIUS client) can communicate. Use the buttons to add or remove servers, and edit server parameters. You can also edit a server's parameters by double-clicking the server entry in the list.

Priority

Order in which the RADIUS authentication server is checked, as compared to the other RADIUS authentication servers listed here. The lower the number, the higher the priority with 1 being the highest priority.

Address

IP address of the RADIUS authentication server.

Client UDP Port

UDP port number (1-65535) on the RADIUS authentication server to which the device sends authentication requests; 1812 is the default port number.

Access Type

The type of authentication access enabled for this RADIUS server:

- Any access the server can authenticate users originating from any access type.
- Management access the server can only authenticate users that requested management access via the console, Telnet, SSH, or HTTP, etc.
- Network access the server can only authenticate users accessing the network via 802.1X, MAC, or Web-Based authentication.

Devices that do not support this feature display N/A in this column.

Current Sessions

The current number of sessions associated with this server when the device is using the <u>sticky roundrobin RADIUS</u> authentication algorithm. This value is not used when other algorithms are being used.

Max Sessions

The maximum number of sticky round-robin authentication sessions permitted on the server when the sticky round-robin RADIUS authentication algorithm is configured for the device. This value is not used when other algorithms are selected. In sticky round-robin, if a MAC address needs to re-authenticate, the request is sent to the same RADIUS server as the initial authentication request, unless the current number of authentication sessions for the server has reached the specified Max Sessions value. When this value is reached, re-authentication requests instead default to the standard round-robin behavior to determine the RADIUS server to which to send the request.

Number of Retries

The number of times the device resends an authentication request if the RADIUS authentication server does not respond. For ExtremeWireless devices, this value is configured per RADIUS server. For all other devices, this value is global to all RADIUS servers, and is specified per device (Client Default) in the RADIUS Authentication Client Settings section.

Timeout Duration (sec)

The amount of time in seconds the device waits for the RADIUS authentication server to respond to an authentication request. For ExtremeWireless devices, this value is configured per RADIUS server. For all other devices, this value is global to all RADIUS servers, and is specified per device (Client Default) in the RADIUS Authentication Client Settings section.

Management Interface

The IP address and VRName used when the switch is communicating with a configured RADIUS server.

Add Button

Opens the Add/Edit RADIUS Authentication Server window, where you can enter the parameters for a server you want to add to the list. When you select **OK** on this window, the new server is added.

Edit Button

Select a RADIUS authentication server in the list and use this button to edit the server's parameters. You can also edit the server parameters by double-clicking the server entry in the list.

Remove Button

Select a RADIUS authentication server in the list and use this button to remove the server.

Apply Button

Applies any changes you made in the RADIUS Authentication Server(s) tab.


- Authentication
- Port Properties Authentication Configuration Tab
- Add RADIUS Authentication Server Window
- Add RADIUS Accounting Server Window

Δ

RADIUS Authentication (Devices)

The RADIUS Authentication tab displays authentication RADIUS server information for all the devices in the current domain. You can configure RADIUS server information for an individual device using the device's RADIUS Tab.

To access this tab, select **Devices/Port Groups>Devices** in the left-panel of the **Policy** tab, then select the **RADIUS Authentication** tab in the right panel.



IP Address

IP address of the device.

Auth Client Status

Informs you whether or not the device is enabled as a RADIUS client. If **Enabled**, the device is a RADIUS client and communicates with a RADIUS authentication server whenever a user logs on to a port on the device, as long as the port itself is enabled for authentication. If Disabled, the device is currently not enabled as a RADIUS client.

Auth Retries

Number of attempts the device (RADIUS client) makes to connect to the RADIUS authentication server before giving up and trying the next RADIUS server on the list.

Auth Timeout Duration

Total number of seconds the device (RADIUS client) waits for the RADIUS authentication server to respond before trying again.

Auth Server Address

The IP addresses of the RADIUS servers the client device attempts to contact.

Auth UDP Port

The UDP port number used to send authentication requests.

RADIUS Response Conflict


Indicates the RADIUS response attribute that the device uses for authentication. You can configure the Response Mode in the RADIUS tab for the device.

RADIUS Accounting (Device)

The device RADIUS **Accounting** tab enables you to configure and enable communication between the selected device (the RADIUS client), a RADIUS server or servers, and ExtremeCloud IQ - Site Engine, for the purposes of accounting (for your SNMPv3 devices that support it).

RADIUS accounting collects various data and statistics, such as the length of time a user has been logged on, and makes that data available to an administrator. It is used by a device to save accounting data on a RADIUS server. Accounting requests are sent from the device to the server. The server acknowledges these requests, and data is passed to the server via accounting updates. For more information on accounting functionality, refer to your RADIUS server documentation.

To display the device RADIUS **Accounting** tab, select a device in the left panel Devices > Devices tree, then select **RADIUS** > **Accounting** in the right panel.



RADIUS Accounting Client Settings

This section lets you enable or disable communication between the selected device (the RADIUS client) and the RADIUS accounting servers, and specify the update interval.

Accounting Status

Enables you to enable or disable RADIUS accounting on SNMPv3 devices that support it. RADIUS accounting is used by a device to save accounting data on a RADIUS accounting server. If accounting is enabled, an accounting session starts after the user is successfully authenticated by a RADIUS authentication server. The default is Disabled. For ExtremeWireless devices, the status is automatically set to Enabled when a RADIUS server exists and Disabled when it does not. Devices that do not support RADIUS accounting have this field grayed out.

Management Access Auth Status Override

Enables you to override the Accounting Status for users accessing the RADIUS accounting server(s) that have requested management access via the console, Telnet, SSH, or HTTP, etc.

Network Access Auth Status Override

Enables you to override the Accounting Status for users accessing the network via 802.1X, MAC, or Web-Based authentication.

Per Authentication Type Accounting Status

Enables you to enable/disable RADIUS accounting for individual authentication types (Quarantine, 802.1X, PWA, MAC, CEP, and Auto Tracking). Some authentication types do not have RADIUS accounting enabled by default (when global RADIUS accounting is enabled). Enabling these authentication types gives both ExtremeControl and other RADIUS servers more complete information regarding authentication sessions. These options also enable you to disable accounting messages from certain authentication types, for example, Auto-Tracking, which does not actually authenticate end users. Note that the global <u>Accounting Status</u> option controls accounting on a global basis for all authentication types. Devices that do not support this functionality have these fields grayed out.

Update Interval (seconds)

Collected accounting data is sent from the device to the RADIUS accounting server via accounting updates. The Accounting Update Interval is the amount of time in seconds between accounting updates. This field is greyed out for devices that do not support RADIUS accounting (with the exception of an SNMPv1 R2 device, which displays accounting values but does not permit you to set them.) For ExtremeWireless devices, this value is entered when the RADIUS server is added.

Management Access Timeout Duration Override (sec)

The total number of seconds the device waits for the RADIUS accounting server to respond before trying again for users accessing the RADIUS accounting server(s) that have requested management access via the console, Telnet, SSH, or HTTP, etc.

Network Access Timeout Duration Override (sec)

The total number of seconds the device waits for the RADIUS accounting server to respond before trying again for users accessing the network via 802.1X, MAC, or Web-Based authentication.

Apply Button

Applies the changes you made in the RADIUS Accounting Client Settings section.

Accounting RADIUS Servers Table

This table lists the RADIUS accounting servers with which the device (the RADIUS client) can communicate. Use the buttons to add or remove servers, and edit server parameters. You can also edit a server's parameters by double-clicking the server entry in the list.

Priority

Order in which the RADIUS accounting server is checked, as compared to the other RADIUS accounting servers listed here. The lower the number, the higher the priority with 1 being the highest priority.

Address

IP address of the RADIUS accounting server.

Client UDP Port

UDP port number (1-65535) on the RADIUS accounting server to which the device sends accounting requests; 1813 is the default port number. Devices that do not support RADIUS accounting display N/A in this column (with the exception of an SNMPv1 R2 device, which displays accounting values, but does not permit you to set them.)

Access Type

The type of authentication access permitted for this RADIUS server:

- Any access the server can authenticate users originating from any access type.
- Management access the server can only authenticate users accessing the network via the console, Telnet, SSH, or HTTP, etc.
- Network access the server can only authenticate users accessing the network via 802.1X, MAC, or Web-Based authentication.

Devices that do not support this feature display N/A in this column.

Number of Retries

The number of times the device resends an accounting request if the RADIUS accounting server does not respond. Valid values are 0-20. Devices that do not support RADIUS accounting display N/A in this column (with the exception of an SNMPv1 R2 device, which displays accounting values, but does not permit you to set them.)

Timeout Duration (sec)

The amount of time in seconds the device waits for the RADIUS accounting server to respond to an accounting request. Valid values are 2-10 seconds. Devices that do not support RADIUS accounting display N/A in this column (with the exception of an SNMPv1 R2 device, which displays accounting values, but does not permit you to set them.)

Update Interval (sec)

The amount of time in seconds between accounting updates. For ExtremeWireless devices, this value is configured per RADIUS server. For all other devices, this value is global to all RADIUS servers, and is specified per device (Client Default) in the RADIUS Accounting Client Settings section.

Management Interface

The IP address and VRName used when the switch is communicating with a configured RADIUS server.

Apply Button

Applies any changes you made in the RADIUS Accounting Server(s) tab.

Add Button

Opens the Add RADIUS Accounting Server window, where you can enter the parameters for a server you want to add to the list. When you select **OK** on this window, the new server is added.

Remove Button

Select a RADIUS accounting server in the list and use this button to remove the server.


Edit Button

Select a RADIUS accounting server in the list and use this button to edit the server's parameters. You can also edit the server parameters by double-clicking the server entry in the list.

RADIUS Accounting (Devices)

The RADIUS Accounting tab displays accounting RADIUS server information for all the devices in the current domain. You can configure RADIUS server information for an individual device using the device's RADIUS Tab.

To access this tab, select **Devices/Port Groups>Devices** in the left-panel of the **Policy** tab, then select the **RADIUS Accounting** tab in the right panel.



IP Address

IP address of the device.

Acct. Client Status

Informs you whether or not RADIUS accounting is enabled on the device (the RADIUS client). RADIUS accounting is supported on certain SNMPv3 devices, and is used by the device to save accounting data on a RADIUS server. If accounting is enabled, an accounting session starts after the user is successfully authenticated by a RADIUS server. Devices that do not support RADIUS accounting display N/A in this column (with the exception of an SNMPv1 R2 device, which displays a status.)

Acct. Update Interval

Collected accounting data is sent from the device (RADIUS client) to the RADIUS server via accounting updates. The Accounting Update Interval is the amount of time in minutes between accounting updates. Devices that do not support RADIUS accounting display N/A in this column (with the exception of an SNMPv1 R2 device, which displays a value.)

Acct Server Address


The IP addresses of the RADIUS servers the client device attempts to contact.

Auth UDP Port

The UDP port number used to send accounting requests.

Add/Edit RADIUS Server

This window lets you add a RADIUS server to ExtremeCloud IQ - Site Engine for the purpose of authentication. Access this window by selecting **Add** in the RADIUS Server(s) Authentication sub-tab in the RADIUS tab for a device.



Authentication Server Type

Select the authentication type used on the RADIUS server.

NOTE: DNS servers (on supported devices) can only be added when there is a valid DNS server configured on the Device which allows the DNS name to resolve to an IP address at the time of configuration.

Authentication Server IP

Enter the IP or IPv6 address, or the hostname of the RADIUS authentication server. Not all devices support IPv6 address types.

Authentication Client UDP Port

Enter the UDP port number (1-65535) the device (RADIUS client) uses to send authentication requests to the RADIUS authentication server; 1812 is the default port number.

Server Shared Secret

A string of characters used to encrypt and decrypt communications between the device (RADIUS client) and the RADIUS authentication server. This string must match the shared secret entered when you added the client device on the RADIUS server. Without the shared secret, the server and client are

unable to communicate, and authentication attempts fail. The shared secret must be at least 6 characters long; 16 characters is recommended. Dashes are allowed in the string, but spaces are not.

NOTES: If you are configuring multiple RADIUS servers, the same server shared secret must be used for each RADIUS server. This is because most devices (RADIUS clients) only support one shared secret. Matrix N-Series devices with firmware version 5.0 or above are an exception to this, as these devices **do** support a unique shared secret for each server.

This Server Shared Secret is not to be confused with the Application Shared Secret that encrypts communication between the RADIUS client and ExtremeCloud IQ - Site Engine, entered in the Application Shared Secret area of the RADIUS tab for a device.

Verify Shared Secret

Re-enter the Server Shared Secret you entered above.

Max Sessions (Sticky Round-Robin)

Specifies the maximum number of sticky round-robin authentication sessions allowed on the server when the sticky round-robin RADIUS authentication algorithm is configured for a device. In sticky round-robin, if a MAC address needs to re-authenticate, the request is sent to the same RADIUS server as the initial authentication request, unless the current number of authentication sessions for the server has reached the specified Max Sessions value. When this value is reached, re-authentication requests will instead default to the standard round-robin behavior to determine which RADIUS server to send the request to. Devices that do not support this functionality will have the option grayed out.

Number of Retries

The number of times the device will resend an authentication request if the RADIUS authentication server does not respond. For ExtremeWireless devices, this value is configured for each server. For all other devices, this value is global to all RADIUS servers, and is specified per device (Client Default) in the RADIUS Authentication Client Settings section of the RADIUS tab.

Timeout Duration

The amount of time in seconds the device will wait for the RADIUS authentication server to respond to an authentication request. For ExtremeWireless devices, this value is configured for each server. For all other devices, this value is global to all RADIUS servers, and is specified per device (Client Default) in the RADIUS Authentication Client Settings section of the RADIUS tab.

Authentication Access Type

Use the drop-down list to select the type of authentication access allowed for this RADIUS server:

- Any access the server can authenticate users originating from any access type.
- Management access the server can only authenticate users that have requested management access via the console, Telnet, SSH, or HTTP, etc.
- **Network access** the server can only authenticate users that are accessing the network via 802.1X, MAC, or Web-Based authentication.

This feature allows you to have one set of servers for authenticating management access requests and a different set for authenticating network access requests. Devices that do not support this feature will have this field grayed out.

Server Priority

Order in which the RADIUS authentication server will be checked, as compared to the other RADIUS authentication servers on the device. The lower the number, the higher the priority.


Management Interface

Select the IP address and VRName to use when the switch is communicating with a configured RADIUS server.

NOTE: ExtremeXOS/Switch Engine devices must define a Management Interface.

Add RADIUS Accounting Server

This window lets you add a RADIUS server to ExtremeCloud IQ - Site Engine for the purpose of RADIUS accounting. Access this window by selecting **Add** in the RADIUS Server(s) Accounting sub-tab in the RADIUS tab for a device.



Accounting Server Type

Select the accounting type used on the RADIUS server.

NOTE: DNS servers (on supported devices) may only be added when there is a valid DNS server configured on the Device which allows the DNS name to resolve to an IP address at the time of configuration.

Accounting Server IP

Enter the IP or IPv6 address, or the hostname of the RADIUS accounting server. Not all devices support IPv6 address types.

Accounting Client UDP Port

Enter the UDP port number (1-65535) the device (RADIUS client) uses to send accounting requests to the RADIUS server; 1813 is the default port number. Devices that do not support RADIUS accounting will have this field grayed out (with the exception of an SNMPv1 R2 device, which will display accounting values but will not allow you to set them.)

Server Shared Secret

A string of characters used to encrypt and decrypt communications between the device (RADIUS client) and the RADIUS accounting server. This string must match the shared secret entered when you added the client device on the RADIUS server. Without the shared secret, the server and client will be unable to communicate. The shared secret must be at least 6 characters long; 16 characters is recommended. Dashes are allowed in the string, but spaces are not.

NOTES: If you are configuring multiple RADIUS servers, the same server shared secret must be used for each RADIUS server. This is because most devices (RADIUS clients) only support one shared secret. Matrix N-Series devices with firmware version 5.0 or above are an exception to this, as these devices **do** support a unique shared secret for each server.

This Server Shared Secret is different than the Application Shared Secret that encrypts communication between the RADIUS client and ExtremeCloud IQ - Site Engine, entered in the Application Shared Secret area of the RADIUS tab for a device.

Verify Shared Secret

Re-enter the Server Shared Secret you entered above.

Number of Retries (0-20)

The number of times the device will resend an accounting request if the RADIUS server does not respond. Valid values are 0-20. Devices that do not support RADIUS accounting will have this field grayed out (with the exception of an SNMPv1 R2 device, which will display accounting values but will not allow you to set them.)

Timeout Duration (2 -10 sec)

The amount of time in seconds the device will wait for the RADIUS server to respond to an accounting request. Valid values are 2-10 seconds. Devices that do not support RADIUS accounting will have this field grayed out (with the exception of an SNMPv1 R2 device, which will display accounting values but will not allow you to set them.)

Update Interval (minutes)

The Accounting Update Interval is the amount of time in minutes between accounting updates. For ExtremeWireless Wireless devices, this value is configured per RADIUS server. For all other devices, this value is global to all RADIUS servers, and is specified per device (Client Default) in the RADIUS Accounting Client Settings section of the RADIUS tab. Devices that do not support RADIUS accounting will have this field grayed out.

Accounting Access Type

Use the drop-down list to select the type of accounting access allowed for this RADIUS server:

- Any access the server can send an accounting request for users originating from any access type.
- Management access the server can only send an accounting request for users that have requested management access via the console, Telnet, SSH, or HTTP, etc.
- Network access the server can only send an accounting request users that are accessing the network via 802.1X, MAC, or Web-Based accounting.

This feature allows you to have one set of servers for accounting management access requests and a

Ports (Device)

different set for accounting network access requests. Devices that do not support this feature have this field grayed out.

Server Priority (1-20)

Order in which the RADIUS accounting server will be checked, as compared to the other RADIUS accounting servers on the device. The lower the number, the higher the priority.

Management Interface


Select the IP address and VRName to use when the switch is communicating with a configured RADIUS server.

NOTE: ExtremeXOS/Switch Engine devices must define a Management Interface.



Ports (Device)

The device **Port Groups** tab displays a table of information about the selected device's ports. To access this tab, select a port group from the left panel's **Devices/Port Groups Port Groups** tab.



Name

Name of the port, constructed of the name or IP address of the device and either the port index number or the port interface name.

Instance

Shows the instance for the port.

Dot1dIndex

The index value assigned to the port interface.

Status

Shows the status (Up, Down, or Unknown) of the port.

Default Role

Displays the default role for the port. To set the default role, select a port, right-click and select Set Default Role. The Roles Selection view appears where you can select the desired default role. See Default Role in the Concepts topic for information on default roles.

NOTE: Setting a default role on an ExtremeWireless Controller port that is not yet a VNS, creates a new VNS on the HWC.

Alias

Shows the alias (ifAlias) for the interface, if one is assigned.

Stats

Displays information about the port, if configured in PortView.

Port Type

Type of port. Possible values include: Access, CDP, CDP FTM 1 Backplane, FTM 1 Backplane, and Logical.

Neighbor

The port to which the port is connected.

Port Speed

Speed of the port. Possible values include: 10/100, speed in megabits per second (for example, 800.0 Mbps), Unknown (displayed for logical ports).

VLANs

The VLANs to which the port is associated.

Description

A description of the port and the device.

Port Type Details

Additional information about the type of port.

Serial Number


The serial number of the device.

Retrieve Button

Retrieves the most recent information about the ports on the device.

Ports (Port Group)

The Ports panel in the Port Groups navigation tree lists the ports in the selected port group. You can also add and remove ports (user-defined port groups only) by right-clicking the Port Group in the left-hand navigation tree. To access this panel, select a port group in the left-panel **Devices/Port Groups** > **Port Groups** navigation tree.



Name

Name of the port, constructed of the name or IP address of the device and either the port index number or the port interface name.

Default Role

See <u>Default Role</u> in the Concepts topic for information on default roles. For additional information, see <u>Port Mode</u>.

Alias

Shows the alias (if Alias) for the interface, if one is assigned.

Port Type

Type of port. Possible values include: Access, Interswitch Backplane, Backplane, Interswitch, and Logical.

Port Speed

Speed of the port. Possible values include: 10/100, speed in megabits per second (for example, 800.0 Mbps), Unknown (displayed for logical ports).

Details View (Port Groups)

This tab displays when you select the **Devices/Port Groups > Port Groups** left-panel tab. It displays a table of information about the existing port groups.



Name

Name of the port group.

Number of Ports


Number of ports in the port group.

Add/Remove Ports (User-Defined Port Groups)

Use the Add/Remove Ports window to add and remove ports from user-defined port groups.

To access this window, select the left-panel Port Groups tab. Expand the User-Defined Port Groups folder and select a port group. From this window you can:

- Select the Add/Remove Ports button in the right-panel Ports tab.
- Right-click a Port Group in the left-panel and select Add/Remove Ports.



Default Group — Port Membership

This list displays all the device groups, devices, and port groups in the current domain. Select the ports you want to add to the port group. You can select individual ports, devices, or groups of ports.

Other Groups — Port Membership

This field displays all the ports currently defined for the port group. Select the port you want to remove from the port group.

Add/Move To Button

Select Add/Move To and select the port group to add the ports selected in the Default Group — Port Membership list to the Other Groups — Port Membership list.

Remove Button

Select **Remove** to remove the ports selected in the **Other Groups — Port Membership** list from the port group.

Remove All Button

Select Remove All to remove all the ports in the Other Groups — Port Membership list.


Add/Remove Ports

In this window, you can add and remove ports to and from port groups. Initially, all ports are grouped into a Default port group. When you create new port groups, you add ports from the Default group into your newly defined port groups using this window.

To access this window, open the **Devices > Port Groups** tab. Then, right-click on the port group to which the ports are being added and select **Add/Remove Ports**. The Add/Remove Ports window opens with the ports in the Default port group displayed in the left panel.

Add ports to the port group by selecting the ports in the left-panel, then selecting the port group in the right panel and selecting **Add To Group**.

NOTE: User based ports are not listed because user based port groups can only be one default.



Devices

This field displays the Devices assigned to the Policy Domain. Ports grouped in the Devices list are not members of the Port Group.

Group Port Membership

This field displays any port groups you have created and their currently defined ports.

Add To Group Button

Adds the ports selected under the Devices list to the port group selected on the right.

Remove Button

Select the ports you want to remove from a port group and select **Remove** to return the ports to the Devices list.

Remove All Button

Select a port group and select **Remove All** to remove all ports from the port group and return them to the Devices list.

Port Authentication Configuration


The **Port Configuration** tab allows you to configure and change the authentication settings for a port. Authentication must be configured and enabled on the device in order for individual port authentication settings to take effect. Only those areas of the tab that relate to the authentication type configured on the device are available for editing.

The Authentication Configuration tab has six sections:

- Authentication Mode
- RFC3580 VLAN Authorization
- Login Settings
- Automatic Re-Authentication
- Authenticated User Counts
- CEP Access

Authentication Mode

This section displays general authentication and port mode information about the port.



Port Mode

Port mode defines whether or not a user is required to authenticate on a port, and how unauthenticated traffic will be handled. It is a combination of Authentication Behavior (whether or not authentication is enabled on the port), and Unauthenticated Behavior (whether unauthenticated traffic will be assigned to the port's <u>default role</u> or discarded).

- Authentication Behavior -- Defines whether or not end users are required to authenticate on the port (device).
 - Active -- Normal authentication procedures are implemented. End users are required to authenticate.

- **Inactive** -- Authentication of end users is not required.
- Unauthenticated Behavior -- Defines how the traffic of unauthenticated end users will be handled on the port.
 - **Default Role** -- If the end user is unauthenticated, the port will implement its default role. If there is no default role, there will be no role on the port.
 - **Discard** -- If the end user is unauthenticated, no traffic is allowed on the port.

These two settings can be combined to create four possible port modes.

- Inactive/Discard Mode: In this mode, authentication is inactive for the port. All traffic from users connected to the port is discarded. This effectively turns the port off. This port mode is not available for Single User MAC Authentication.
- Inactive/Default Role Mode: In this mode, authentication is inactive for the port. All users connecting to this port will use the default role, if one has been assigned to the port, in combination with any existing static classifications. If there is no default role assigned to the port, the port uses only the static classification rules which exist. If there are no static rules, the port uses the PVID and default class of service for the port. This is the default port mode for ports.
- Active/Discard Mode: In this mode, authentication is active for the port and end users are required to authenticate. All traffic from unauthenticated users connected to the port is discarded. The Unauthenticated Behavior varies depending on the type of authentication configured on the device.

Single User Web-based Authentication: If authentication is successful, the port is assigned the end user's role as its current role. If unsuccessful, all traffic is discarded. A default role has no meaning on this Active/Discard port, since all unauthenticated traffic is discarded.

Single User 802.1X and 802.1X+MAC Authentication: If authentication is successful, the port is assigned the end user's role as its current role. If unsuccessful, all traffic is discarded. This mode requires that there be **no** default role assigned to the port.

Single User MAC Authentication: This port mode is not available for Single User MAC Authentication.

Multi-User 802.1X and MAC Authentication: If authentication is successful, the port is assigned the end user's role as its current role. If unsuccessful, all traffic is discarded. A default role has no meaning on this Active/Discard port, since all unauthenticated traffic is discarded.

Multi-User Web-based Authentication: This port mode is not available for Multi-User Web-based Authentication.

Advantages of Active/Discard mode: This mode is highly secure, since the end user receives no network services at all until authentication is successful.

Disadvantages of Active/Discard mode: The unauthenticated end user is unable to connect to any network services, such as the Domain Controller (if using a Microsoft operating system), DHCP services,

DNS services, or the Web proxy. In single user web-based authentication, the device spoofs WINS/DNS services (if the functionality is enabled) in order to allow the user to communicate with it for authentication.

Active/Default Role Mode - In this mode, authentication is active for the port and end users are required
to authenticate. If authentication is successful, the port is assigned the end user's role as its current role.
All unauthenticated users connected to the port will use the default role, if one has been assigned to the
port, in combination with any existing static classifications. If there is no default role assigned to the
port, the port uses only the static classification rules which exist. If there are no static rules, the port uses
the PVID and default class of service for the port. For Single User 802.1X and 802.1X+MAC
Authentication, this mode requires that a default role be assigned to the port.

Advantages of Active/Default Role mode: In this mode, a default role is applied to the port to allow unauthenticated end users access to basic services such as the DHCP Server, Domain Services, WINS, and the Web proxy. When the end user is authenticated, that user's role is applied to the port, providing a customized set of services allowed by his or her role. Active/Default Role mode is an alternative to Active/Discard mode, which is limiting in that there are no network services available at all until the end user is authenticated.

Disadvantages of Active/Default Role mode: This mode is less secure than Active/Discard, in that the user receives some network access prior to authentication.

RFC3580 VLAN Authorization Tab

This tab lets you enable or disable RFC 3580 VLAN Authorization on the port and specify an egress state. RFC 3580 VLAN Authorization must be enabled in networks where the RADIUS server has been configured to return a VLAN ID when a user authenticates.

When RFC 3580 VI AN Authorization is enabled:

- ports on devices that do **not** support policy tag packets with the VLAN ID.
- ports on devices that do support policy and also support Authentication-Based VLAN to Role Mapping classify packets according to the role to which the VLAN ID maps.

You can also enable and disable VLAN Authorization at the device level using the device **Authentication** tab. If the device does not support RFC 3580, this tab is grayed out.



VLAN Authorization Status

Allows you to enable and disable RFC 3580 VLAN Authorization for the selected port. This option is grayed out if not supported by the device.

VLAN Authorization Admin Egress

Allows you to modify the VLAN egress list for the VLAN ID returned by the RADIUS server when a user authenticates on the port:

- None No modification to the VLAN egress list will be made.
- Tagged The port will be added to the list with the egress state set to Tagged (frames will be forwarded as tagged).
- Untagged The port will be added to the list with the egress state set to Untagged (frames will be forwarded as untagged).
- Dynamic The port will use information returned in the RADIUS response to modify the VLAN egress list. This value is supported only if the device supports a mechanism through which the egress state may be returned in the RADIUS response.


The current egress settings for the port are displayed in the VLAN Oper Egress column in the **User Sessions** tab. These options are grayed out if not supported by the device.

Apply Button

Saves any change you made to the VLAN Authorization settings.

Login Settings

This tab displays the current login settings for the port and allows you to change the settings if desired. The options available depend on what type(s) of authentication are enabled on the device.



Number of Attempts Before Timeout

Number of times a user can attempt to log in before authentication fails and login attempts are not allowed. For web-based authentication, valid values are 1-2147483647, zero is not allowed, and the default is 2. For 802.1X and MAC authentication, this value is permanently set to 1.

Hold Time (seconds)

Amount of time (in seconds) authentication will remain timed out after the specified Number of Attempts Before Timeout has been reached. Valid values are 0-65535. The default is 60. (Hold Time is also known as Quiet Period in web-based and MAC authentication.)

Authentication Request Period

For 802.1X authentication, how often (in seconds) the device queries the port to see if there is a new user on it. If a user is found, the device then attempts to authenticate the user. Valid values are 1-65535. The default is 30.

User Timeout

For 802.1X authentication, the amount of time (in seconds) the device waits for an answer when querying the port for the existence of a user. Valid values are 1-300. The default is 30.

Authentication Server Timeout

For 802.1X authentication, if a user is found on the port, the amount of time (in seconds) the device waits for a response from the authentication server before timing out. Valid values are 1-300. The default is 30.

Port Handshake Requests Before Failure

For 802.1X authentication, the number of times the device tries to finalize the authentication process with the user before the authentication request is considered invalid and authentication fails. Valid values are 1-10. The default is 2.

Quarantine Session Timeout (sec)

For Quarantine authentication, the maximum number of seconds an authenticated session may last before automatic termination of the session. A value of zero indicates that no session timeout will be applied.

Quarantine Session Idle Timeout (sec)

For Quarantine authentication, the maximum number of consecutive seconds an authenticated session may be idle before automatic termination of the session. A value of zero indicates that the device level setting is used.

Auto Tracking Session Timeout (sec)

For Auto Tracking sessions, the maximum number of seconds a session may last before automatic termination of the session. A value of zero indicates that the device level setting is used.

Auto Tracking Session Idle Timeout (sec)

For Auto Tracking sessions, the maximum number of consecutive seconds a session may be idle before automatic termination of the session. A value of zero indicates that the device level setting is used.

Apply Button

Applies the Login Settings changes to the port.

Automatic Re-Authentication

This tab is grayed out if only web-based authentication is enabled on the device. For 802.1X and MAC authentication, the Automatic Re-Authentication tab lets you set up the periodic automatic re-authentication of logged-in users on this port. Without disrupting the user's session, the device repeats the authentication process using the most recently obtained user login information to see if the same user is still logged in. Authenticated logged-in users are not required to log in again for re-authentication, as this occurs "behind the scenes."



802.1X Re-auth Status

If **Active** is selected, the re-authentication feature is enabled for 802.1X authentication. If **Inactive** is selected, the re-authentication feature is disabled.

802.1X Re-auth Frequency (sec)

How often (in seconds) the device checks the port to re-authenticate the logged-in user via 802.1X authentication. Valid values are 1-2147483647. The default is 3600.

MAC Re-auth Status

If **Active** is selected, the re-authentication feature is enabled for MAC authentication. If **Inactive** is selected, the re-authentication feature is disabled.

MAC Re-auth Frequency (sec)

How often (in seconds) the device checks the port to re-authenticate the logged in user via MAC authentication. Valid values are 1-2147483647. The default is 3600.

Authenticated User Counts

This tab provides authenticated user-count information for devices with Multi-User as their configured authentication type. See the device Authentication tab for information on setting the device authentication type.



Current Number of Users

The current number of users actively authenticated or have authentications in progress on this interface. If **Multi-User** authentication is disabled, this number is **0**. Any unauthenticated traffic on the port is not included in this count.

Number of Users Allowed (up to 2048)

The number of users that can be actively authenticated or have authentications in progress at one time on this interface. If you set this value below the current number of users, end-user sessions exceeding that number are terminated.

NOTE: B2/C2 Devices. If you are configuring a single user and an IP phone per port, set this value to 2.

Number of MAC Users Allowed (up to 2048)

The number of users that can be actively authenticated via MAC authentication, or have MAC authentications in progress at one time on this interface. The number of MAC users allowed cannot exceed the number of users allowed. If you set this value below the current number of users, end user sessions exceeding that number are terminated. If MAC is not selected as a **Multi-User** authentication type on the device Authentication tab, this field will be grayed out.

Number of Quarantine Users Allowed (up to 2048)

The number of users that can be actively authenticated via Quarantine authentication, or have Quarantine authentications in progress at one time on this interface. The number of Quarantine users allowed cannot exceed the number of users allowed. If you set this value below the current number of users, end user sessions exceeding that number are terminated. If Quarantine Auth is not enabled on the device Authentication tab, this field will be grayed out.

Number of Auto Tracking Users Allowed (up to 2048)


The number of Auto Tracking users that can be actively authenticated or have authentications in progress at one time on this interface. The number of Auto Tracking users allowed cannot exceed the number of users allowed. If you set this value below the current number of users, end user sessions exceeding that number will be terminated. If Auto Tracking is not enabled on the device Authentication tab, this field is grayed out.

Convergence End-Point Access

This tab lists all the CEP (Convergence End-Point) protocols supported by the device on which the port resides, and lets you enable or disable them for that port. For devices that do not support CEP, the tab is blank.

NOTE: Port Mode Authentication Behavior must be set to Active (on the <u>General sub-tab</u>) for authentication to be allowed using these CEP Protocols.

Enable CEP protocols for multiple ports using the Port Configuration Wizard. In addition to enabling protocols on the port, you must also configure CEP for the device on which the port resides. Configure CEP for a single device using the device Authentication tab (CEP sub-tab) or for multiple devices using the Device Configuration Wizard.



CEP Access

Lists all the CEP protocols supported by the device on which the port resides. Use the checkboxes to enable or disable CEP protocols on this port. If the device does not support the CEP feature, this area is blank.

Enable All Button

Selects all the checkboxes and enables all the CEP protocols for this port.

Disable All Button

Deselects all the checkboxes and disables all the CEP protocols for this port.

Apply Button

Applies CEP access changes to the port.



How To Use Policy

The **How To** section contains Help topics that give you instructions for performing tasks in the **Policy** tab.

How to Select on Add/Remove Windows

The **Policy** tab includes several Add/Remove windows in which you can add items from a left panel to a right panel, and remove items from the right panel. The following procedures explain how to make single and multiple selections in the panels and move the selections to the opposite panel.

Instructions on:

- Selecting single items
- Selecting multiple sequential items
- Selecting multiple non-sequential items

Selecting single items

To select one item from the left panel and add it to the right panel, select the item, then select the **Right Arrow** button.

To remove one item from the right panel, select the item, then select the Left Arrow button.

Selecting multiple sequential items

To select a sequence of items in the left panel and add them to the right panel:

- 1. Hold down the **Shift** key and select the first and last (or last and first) items in the sequence.
- 2. Select the **Right Arrow** button.

To remove a sequence of items from the right panel:

- 1. Hold down the **Shift** key and select the first and last (or last and first) items in the sequence.
- 2. Select the **Left Arrow** button.

Selecting multiple non-sequential items

To select multiple non-sequential items in the left panel and add them to the right panel:

- 1. Hold down the **Ctrl** key and select each item you want to add.
- 2. Select the **Right Arrow** button.

To remove multiple non-sequential items from the right panel:

- 1. Hold down the **Ctrl** key and select each item you want to remove.
- 2. Select the **Left Arrow** button.

286 of 685

How to Create and Use Domains

ExtremeCloud IQ - Site Engine provides the ability to create multiple policy configurations by allowing you to group your roles and devices into Policy Domains. A Policy Domain contains any number of roles and a set of devices that are uniquely assigned to that particular domain. For example, a university can have a Dormitory domain with a policy configuration created for students, and an Administration domain with a policy configuration for staff members.

You can create multiple domains and easily switch from one domain to another. You can also export policy domain configuration data to a .pmd file, (one file per domain) for backup and troubleshooting purposes, and you can import data from a .pmd file into a policy domain.

In order for your network devices to be displayed in the **Policy** tab's left-panel **Devices** tab, they must be assigned to a Policy Domain. Initially, you must use a device Discover to add your devices to the ExtremeCloud IQ - Site Engine database. After your devices are in the database, you can assign the devices to a Policy Domain. As soon as the devices are assigned to a domain, they are automatically displayed in the **Policy** tab's left-panel **Devices** tab. Only devices that support policy are displayed.

ExtremeCloud IQ - Site Engine automatically locks the current Policy Domain when you begin to edit the domain configuration. Other users are notified that the domain is locked and they are not be able to save their own domain changes until the lock is released. For more information, see Controlling Client Interactions with Locks. After a modification is made, you must save the domain to notify all clients that are viewing that domain of the change, and automatically update their view with the new configuration.

Instructions on:

- Creating a New Domain
- Opening a Domain
- Assigning Devices to a Domain
- Removing Devices From a Domain
- Importing a File into a Domain
- Exporting a Domain to a File
- Importing Data from a Domain
- Saving a Domain
- Reading a Domain
- Renaming a Domain
- Deleting a Domain

Creating a New Domain

Use these steps to create a new Policy Domain.

- 1. Select Open/Manage Domain > Create Domain.
- 2. Enter the name for the new domain. Select **OK**.
- 3. A new (blank) Domain opens.
- 4. Select the **Global Domain Settings > Do Not Use Global Services** checkbox if you don't want the domain to include and display services common to all domains.
- 5. Proceed with assigning devices to the domain and then configuring the desired policies.

Opening a Domain

In ExtremeCloud IQ - Site Engine, you work in one current domain at a time. To change to a different domain, use the **Open/Manage Domain > Open Domain** menu to select the desired domain. If you have made changes to the current domain, you are prompted to update the database with the current domain configuration prior to opening the new domain.

Assigning Devices to a Domain

Initially, you must perform a device Discover to add a device to the ExtremeCloud IQ - Site Engine database. After your devices have been added to the database, you must assign the devices to a Policy Domain. A device can exist in only one Policy Domain. As soon as the devices are assigned to a domain, they are automatically displayed in the **Policy** tab's left-panel **Devices** tab. Only devices assigned to the Policy Domain you are currently viewing are displayed in the tab.

Use these steps to assign devices to a Policy Domain.

- 1. If necessary, open the domain to which you want to assign devices.
- 2. Select Open/Manage Domain > Assign Devices to Domain. The Assign Devices to Domain window opens.
- 3. Devices in the database but not assigned to a domain are listed in the left-panel Unassigned folder (including devices that do not support policy). The left panel also displays any other domains and the devices assigned to those domains. Use the drop-down list to select a single domain or All Other Domains. If you select All Other Domains, use the bottom panel to view the domain to which each device is assigned.

Note: Select the search icon to <u>search</u> for a device. A search box is available to filter through the visible device tree.

- 4. The right panel displays the current domain and the devices assigned to that domain. To add a device to the current domain, select the device in the left panel and select **Add**. You can also select and add multiple devices.
- 5. To remove a device from the current domain, select the device and select **Remove**. This removes the device from the current domain and places it back in the device tree as either unassigned or as a member of the domain from which it came. It does not delete the device from the ExtremeCloud IQ Site Engine database.
- 6. Select **OK**

7. The selected devices are assigned to the current domain and displayed in the **Policy** tab left-panel **Devices** tab. (Only devices that support policy are assigned to the domain and displayed.)

Removing Devices From a Domain

Removing a device from a domain, removes the device from the **Devices** tab and places it in the Unassigned folder in the Assign Devices to Domain window.

NOTE: Removing a device from a domain does not delete the device from the ExtremeCloud IQ - Site Engine database. To delete a device from the database, right-click on the device in the left-panel **Devices** tab, and select **Delete** from the menu. When a device is deleted from the database, it is automatically removed from ExtremeCloud IQ - Site Engine and the **Devices** tab.

- 1. If necessary, open the domain from which you want to remove devices.
- 2. Select Open/Manage Domain > Assign Devices to Domain. The Assign Devices to Domain window opens.
- 3. The right panel displays the current domain and the devices assigned to that domain. To remove a device from the current domain, select the device from the Current Domain right-panel and select the left arrow. This removes the device from the current domain and places it back in the device tree as either unassigned or as a member of the domain from which it came. It does not delete the device from the ExtremeCloud IQ Site Engine database.
- 4. Select **OK**

Importing a File into a Domain

You can import policy data from a PMD file into a Policy Domain.

- 1. Make sure that the domain you want to import a file into is your current domain.
- 2. Select Open/Manage Domain > Import/Export > Import From File. The Import from File window opens.
- 3. Enter the name and path for the data file (PMD) you want to import, or browse to the file. Selecting Select File, opens a dialog box from which you can select a data file by searching your local drive or a network drive.
- 4. Select the specific data elements you want to import or select **Select All** to select all the data import options. See Data Elements to Import for important information on each element and how they are imported.
- 5. To append, update, or overwrite the global rules with the PMD file you are importing, select the **Global Services & Rules** checkbox.
- 6. Select how you want the imported data applied to your current domain. Select the links below for detailed information on how each specific action affects the import of certain data elements.
 - Append data to existing elements
 - Update existing data with elements from domain
 - Overwrite existing elements

7. Select **OK**. The data elements are imported and see a message regarding import status.

Exporting a Domain to a File

You can export policy data from a Policy Domain to a PMD file.

- 1. Select Open/Manage Domain > Import/Export > Export to File.
- 2. Select the **Domain** to save as a PMD file.
- 3. Select Export.
- 4. The Policy Domain is downloaded to the default file download location.

Importing Data from a Domain

You can import policy configuration data from one policy domain into another.

- 1. Ensure your current domain is the domain into which you want to import data.
- 2. Select Open/Manage Domain > Import/Export > Import From Domain. (This menu option is not available if only one domain exists, as there are no other domains from which to import data.) The Import from Domain window opens.
- 3. Use the drop-down list to select the domain whose data you want to import.
- 4. Select the specific data elements you want to import or select **Select All** to select all the data import options. See Data Elements to Import for important information on each element and how they are imported.
- 5. Select how you want the imported data applied to your current domain. Select the links below for detailed information on how each specific action affects the import of certain data elements.
 - Append data to existing elements
 - Update existing data with elements from domain
 - Overwrite existing elements
- 6. Select Import. The data elements are imported and you see a message regarding import status.

Saving a Domain

After a Policy Domain has been changed, you must save the domain to notify all clients using that domain of the change and automatically update their tab with the new configuration. An asterisk (*) is displayed beside the Policy tab title when you have made changes to the domain that need to be saved. You can save a Policy Domain by selecting Open/Manage Domain > Save Domain. To discard unsaved changes you made to a domain, open the Open/Manage Domains > Open Domain menu and select the domain in which you are currently working.

Renaming a Domain

You can rename the current Policy Domain by selecting **Open/Manage Domain > Rename Domain** and entering a new name.

Deleting a Domain

You can delete one or more Policy Domains by selecting **Open/Manage Domain > Delete Domain**.



How to Create a Role

A role is a policy profile consisting of a set of network access services that you can apply at various access points in a policy-enabled network. A port takes on a user's role when the user authenticates.

Creating a role using the role tabs consists of creating a name for the role with the **Create Role** menu option, then defining its characteristics (default class of service, default access control, and/or services) using the role's right-panel tabs. You might also use this method if you are creating a role for which there is default class of service and/or access control, but no services.

If you want to change the characteristics of a role, you can select the role in the left panel and use the right panel to modify it.

Instructions on:

- Using the Role Tabs
- Modifying a Role
- Deleting a Role

Using the Role Tabs

Creating a role using the **Role** tab consists of creating a name for the role, then using the right panel to specify the characteristics of the role (default class of service, default access control, and/or services).

- 1. In the **Policy** tab left panel, select the **Roles/Services > Roles** tab.
- 2. Right-click the **Roles** tab, and select **Create Role**. The Create window opens.
- 3. Type the role name in the highlighted box. The name can be up to 64 characters in length, and special characters are allowed, with the exception of colons (:) and semicolons (;). Duplicate names are not allowed, regardless of case. For example, if you already have a role Faculty and you attempt to name the new role Faculty or faculty, the **Policy** tab creates the role, but with the name New Role, or New Rolen (where n is the sequence number, if there is more than one New Role). You can then rename the new role. Press **Enter** after you've entered the name. (If you don't press **Enter**, the name remains New Role.)
- 4. Select the role in the left panel, and the role opens in the right panel. Use the right panel to add a role description, enable TCI Overwrite, and set the role's default actions (including access control and class of service).

5. In the Services section in the right panel, select the **Add/Remove Services** button to add services to the role. This opens the role Add/Remove Services window.

NOTE: The **Policy** tab checks for rule conflicts when more than one service is added. See Conflict Checking for more information.

- 6. To add a VLAN to the Role's Egress list, select the role and use the VLAN Egress tab in the right panel.
- 7. To configure MAC, IP, and VLAN to role mapping lists for the role, select the role and use the **Mappings** tab in the right panel.
- 8. Now that you have created the role, you can:
 - Assign the role as the default role for a port
 - Modify the role's characteristics
- 9. Enforce to write the new information to the devices.

Modifying a Role

Once you've created a role, you can change its characteristics by selecting the role in the Policy tab's left panel and using the associated tabs in the right panel.

Instructions on:

- Adding Services to Roles
- Modifying a Role's Default Class of Service
- Modifying a Role's Default Access Control
- Modifying a Role's Description
- Modifying a Role's Ports
- Removing Services from Roles

Adding Services to Roles

To add services to roles:

- 1. Select the left panel Roles/Services > Roles tab and expand the Roles tab. Select the role to which you want to add services in the left panel, then select the General tab in the right panel.
- 2. Select Add/Remove Services. This opens the Add/Remove Services window.
- 3. Make sure the role to which you wish to add services is displayed in the Role selection box.
- 4. In the Groups and Services panel, select the services and/or service groups you wish to add to the role, and select the **Right Arrow** button. To remove services, select them in the Selected Services panel and select the **Left Arrow** button.

NOTE: The Policy tab checks for rule conflicts when more than one service is added. See Conflict Checking for more information.

- 5. If you wish, you can select another role, and add or remove services from it.
- 6. Select **OK**.
- 7. Enforce to write the new information to the devices.

Removing Services from a Role

- 1. Select the left panel Roles/Services > Roles tab and expand the Roles folder.
- 2. Select the role from which you want to remove services, then select the **General** tab in the right panel.
- 3. Select Add/Remove Services. This opens the Add/Remove Services window.
- 4. Make sure the role from which you wish to remove services is displayed in the Role selection box.
- 5. In the Selected Services panel, select the services and/or service groups you wish to remove from the role, and select the **Left Arrow** button. To add services, select them in the Groups and Services panel and select the **Right Arrow** button.
- 6. If you wish, you can select another role, and remove services from or add services to it.
- 7. Select OK.
- 8. Enforce to write the new information to the devices.

Modifying a Role's Default Class of Service

Use the role's <u>General tab</u> to change its default class of service settings. Be sure to <u>enforce</u> to write the new information to the devices.

Modifying a Role's Default Access Control

Use the role's <u>General tab</u> to change its default access control. Be sure to <u>enforce</u> to write the new information to the devices.

Modifying a Role's Description

You can edit the description for the role on the role's <u>General tab</u>. Select **OK** to save the change to the database.

Modifying a Role's Ports

You can select a port and choose the default role on the <u>Ports tab</u>. You can also select <u>PortView</u> to open the PortView for the port or make changes to the port settings themselves.

- 1. In the **Policy** tab left panel, select a device in the **Devices** left-panel tab.
- 2. Select the port on which you want to set a default role.
- 3. Right-click the port and select Policy > Set Default Role.
- 4. Select the Assign/Replace Default Role checkbox. The drop-down list is available.

- 5. Select the default role for the port from the drop-down list.
- 6. Select OK.
- 7. Enforce to write the new information to the devices.

Mapping a Role to an HTTP Redirect Group

The HTTP Redirect action allows the role/rule to be mapped to an HTTP Redirect group index. The action widgets contain a menu to edit the group configuration.

Deleting a Role

- 1. In the **Policy** tab left panel, select a device in the **Devices** left-panel tab.
- 2. Select the port on which you want to delete the default role.
- 3. Right-click the port and select Policy > Set Default Role.
- 4. Select the Clear Default Role checkbox.
- 5. Select the default role for the port.
- 6. Select OK.
- 7. Enforce to write the new information to the devices.

How to Assign a Default Role to a Port

In the **Policy** tab, you can specify a default role for the port. To configure ports you use the Set Default Role window.

Assigning and Clearing a Default Role

Configuring a port allows you to set the port mode, establish login settings, set the default role, and enables you to view the current configuration on the port.

- Assigning Default Roles to Ports
- Clearing Default Roles from Ports

Assigning Default Roles to Ports

NOTE: Setting a default role on an ExtremeWireless Controller port that is not yet a VNS, creates a new VNS on the wireless controller.

- 1. Select a device in the left-panel **Devices** tab and expand a slot or ports grouping in the right-panel Details view.
- 2. Right-click the desired port and select **Policy > Set Default Role** from the menu. The Set Default Role window opens.

- 3. Select Assign/Replace Default Role and select a role in the drop-down list.
- 4. Select OK.

Clearing Default Roles from Ports

You can clear the default role from a single port, or from multiple ports.

- 1. Select a device in the left-panel **Devices** tab and expand a slot or ports grouping in the right-panel Details view.
- 2. Right-click the desired port and select **Policy > Set Default Role** from the menu. The Set Default Role window opens.
- 3. Select Clear Default Role.
- 4. Select **OK**

NOTE: If you are replacing the current default role with another one, you don't need to clear the current default role. Selecting the new default role and selecting **OK** clears the previous default role automatically.



How to Create a Quarantine Role

The Quarantine role is a highly restrictive role used to isolate users and restrict network access.

The Quarantine role is used in conjunction with the Extreme Networks Intrusion Prevention System (IPS) to create an automatic response to threats detected on the network. After the Quarantine role has been enforced to the network and the Extreme Networks IPS is properly configured, this role can be automatically set as the default role on any port where a threat has been detected. Normally, roles are applied to ports via authentication.

You can also set the Quarantine role as a port's default role if, for example, you have modified the role to provide some limited access and you want to use it as a "guest" role.

The **Policy** tab default domain includes the Quarantine role. However, if you add a new domain, you need to create the Quarantine role. For information on how to create a role, see <u>How to</u> <u>Create a Role</u>.

After you have created the role, you can modify the role's default class of service and access control settings, and make changes to the role's services and rules using the right-panel tabs, just like any other role. If you make any changes to the Quarantine role, keep in mind that the role can be used by other applications and should remain highly restrictive in nature.

Instructions on:

• Modifying the Quarantine Role: Use the right-panel tabs to modify the Quarantine role's default values and add or remove services.

• <u>Setting the Quarantine Role as the Default Role on a Port:</u> Use the right-panel General tab or the Port Configuration wizard to set the Quarantine role as a default role on a port.

Modifying the Quarantine Role

When you've created a Quarantine role, you can change its characteristics by selecting the role in the **Policy** tab's left panel and using the associated tabs in the right panel.

NOTE: You cannot rename the Quarantine role.

Modifying Default Values

Use the <u>General tab</u> to change the Quarantine role's default class of service and default access control settings, and to add or edit a description.

- 1. Select the Quarantine Role in the left-panel **Roles** tab.
- 2. In the right-panel **General** tab, select the desired default class of service and default access control settings.
- 3. If desired, add or edit the role's description.
- 4. Be sure to perform an Enforce to write the new Quarantine role to the devices.

Adding/Removing Services

Use the General tab to add or remove services to the Quarantine role.

- 1. Select the Quarantine Role in the left-panel Roles tab.
- 2. In the right-panel General tab, select **Add/Remove Services**. This opens the <u>Add/Remove Services</u> window.
- 3. Make sure the Quarantine role is displayed in the Role selection box.
- 4. Select the service or service group in the All Services & Service Groups and select the **Right Arrow** button to add them to the Selected Services & Service Groups list. To remove services, select them in the Selected Services & Service Groups list and select the **Left Arrow** button. To remove all services, select the **Double Left Arrow** button.

NOTE: The **Policy** tab checks for rule conflicts when more than one service is added. See <u>Conflict</u> Checking for more information.

- 5. Select OK.
- 6. Be sure to perform an Enforce to write the new Quarantine role to the devices.

Setting the Quarantine Role as the Default Role on a Port

There can be circumstances when you would like to use the **Policy** tab to assign the Quarantine role as the default role on one or more ports. For example, if you have modified the Quarantine role to provide limited access, you can use it as the default role for guest users on your network.

The Quarantine role is assigned as a default role just like any other role. Refer to <u>Assigning</u> Default Roles to Ports for instructions.

How to Create a Service

Services are sets of rules that define how network traffic for a particular network service or application should be handled by a network access device. A service might consist of only one rule governing, for example, email priority, or it might consist of a complex set of rules combining class of service, filtering, rate limiting, and access control (VLAN) assignment. ExtremeCloud IQ - Site Engine policy allows you to create Local Services (services unique to the current domain) and Global Services (services common to all domains). Global Services let you easily create and manage services shared between all your domains.

Services can be one of two types: Manual Service or Automated Service.

- Manual Service This service consists of one or more traffic classification rules you create based on your requirements. Manual services are good for applying customized sets of rules to roles.
- Automated Service This service automatically creates a rule with a specified action (class of service and/or access control), for each device in a particular network resource group or groups. You create a network resource group using a list of MAC or IP addresses, and then associate the group with the Automated service (see How to Create a Network Resource for more information). Automated rule types include Layer 2 MAC Address rules, Layer 3 IP Address and IP Socket rules, and Layer 4 IP UDP Port and IP TCP Port rules.

To create a service using the service tabs, right-click the Services tab and select **Create Service**. If you are creating a Manual service, you can then use the Create Rule menu option and the tabs for the rule to define the rules for the service. You can also use the service tabs and rule tabs to modify an existing service and its rules.

Once you've created a service, you can apply it to any number of roles in the **Policy** tab. A role may utilize both Manual and Automated services.

Instructions on:

- Using the Service Tabs
- Modifying a Service
- Deleting a Service

Using the Service Tabs

The following steps depend on whether you are creating a <u>Manual</u> or an <u>Automated</u> service. For an Automated service, you create the service, select the newly created service, and define the class of service and/or access control for the service in the right-panel. For a Manual service, you create the service and then use the Create Rule menu option and the tabs for the rule to define the rules for the service.

Creating an Automated Service

- 1. In the left panel, select the **Service Repository** tab.
- 2. Expand either the **Local Services** tab or the **Global Services** tab depending on whether you want the service to be local (unique to the current domain) or global (shared between all your domains).
- 3. Right-click on the **Services** tab and select **Create Automated Service**. A New Service item is created in the left panel in a highlighted box.
- 4. Type the service name in the Create window. The service name is case-sensitive; therefore, ExtremeCloud IQ Site Engine policy sees Engineer and engineer as two different service names. Select **OK**. If you don't do this, the name remains New Service. The right-panel displays the service you created.
- 5. Define the rule's traffic description and actions, and enter a description of the service, if desired. For information on configuring the fields on this tab, see the Automated Service window Help topic.
- 6. Enforce to write the new information to your devices.

Creating a Manual Service

- 1. In the left panel, select the **Service Repository** tab.
- 2. Expand either the **Local Services** tab or the **Global Services** tab depending on whether you want the service to be local (unique to the current domain) or global (shared between all your domains).
- 3. Right-click on the **Services** tab and select **Create Service**. A New Service item is created in the left panel in a highlighted box.
- 4. Type the service name in the Create window. The service name is case-sensitive; therefore, the Policy view sees Engineer and engineer as two different service names. Select **OK**. If you don't do this, the name remains New Service. The service is created.
- 5. Define rules for the service. For more information, see Using the Rule General Tab.

NOTE: When you add more than one rule to a service, ExtremeCloud IQ - Site Engine checks for conflicts with other rules in the service. See Conflict Checking for more information.

6. Enforce to write the new information to your devices.

Modifying a Service

Once you've created a service, you can change its characteristics by selecting the service or its rules in the left-panel **Services** tab and using the menu options or associated right-panel tabs.

- Modifying a Service Description
- Modifying a Service Name
- Modifying the Roles for a Service
- Modifying the Rules for a Manual Service
- Modifying an Automated Service

Modifying a Service Description

You can edit the description for the service by selecting it and selecting the **Edit** button beside the **Description** field in the right-panel. Enter a description in the Edit Description window and select **Save** to save the change to the database.

Modifying a Service Name

- 1. In the left panel, select the **Service Repository** tab.
- 2. Expand the **Local** or **Global Services** tab and then the **Services** tab, and select the service you want to modify.

NOTE: If the service is a member of a service group and it's more convenient, you can find the service under the service group in the Service Groups folder. Any change you make to the name there are also reflected in the **Services** tab.

- 3. Right-click the service whose name you want to change, and select **Rename**.
- 4. Type the new name in the Rename window.
- 5. Select **OK** to save the change to the database.

Modifying the Roles for a Service

You can see all the roles associated with a particular service in the Role/Service Usage window.

- 1. In the left-panel Roles tab, select the Role to which you are adding or removing a service.
- 2. Select the Add/Remove button in the Services section of the window to open the Add/Remove Services window.
 - Add a service by selecting it from the All Services & Service Groups column and moving it to the Selected Services & Service Groups column by selecting the right arrow.
 - Remove a service by selecting it from the Selected Services & Service Groups column and moving it to the All Services & Service Groups column by selecting the left arrow.
- 3. Select **OK** to save the changes.
- 4. Enforce to write the new information to your devices.

Adding a Service to Roles

A newly created service can be added to multiple roles using the Add to Role(s) menu.

- 1. In the left panel, select the Roles/Services drop-down list.
- 2. Right-click the service or service group(s) and select Add to Role(s).
- 3. Select one of more Roles to add to the selected Service/Service Group(s) to.
- 4. Select **OK** to save the changes.

Modifying the Rules for a Manual Service

1. Select the left-panel **Services** tab and locate the service you want to modify.

NOTE: If the service is a member of a service group and it's more convenient, you can find the service under the service group in the **Service Groups** tab. Any change you make to the rule there will also be reflected in the **Services** tab.

- 2. Select the service to display its rules.
- 3. Select the rule you want to change, then use the right-panel tabs to make your changes.
- 4. Enforce to write the new information to your devices.

Modifying an Automated Service

1. Open the left-panel Services tab.

NOTE: If the service is a member of a service group and it's more convenient, you can find the service under the service group in the **Service Groups** tab. Any change you make to the service there are also reflected in the **Services** tab.

- 2. Select the service you want to modify. The Automated Service window opens in the right panel.
- 3. Modify the characteristics of the Automated service as required.
- 4. Enforce to write the new information to your devices.

Deleting a Service

Deleting a service removes the service and its rules. If copies of the rules exist for other services, those copies are not affected by the deletion. However, deleting the service removes it from any service groups and roles with which it was associated, so be sure the service is not needed before you delete it. Deleting a Global service deletes the service from all your domains.

- 1. Select the left-panel Roles/Services > Service Repository tab.
- 2. Expand the **Services** tab in either the **Local Services** or **Global Services** tab, depending on the type of service you are deleting.

NOTE: If the service is a member of a service group and it's more convenient, you can find the service under the service group in the **Service Groups** tab. Any change you make to the service there are also reflected in the **Services** tab.

- 3. Right-click the service you want to delete, and select **Delete**.
- 4. Select **Yes** to confirm, then **OK** to clear the confirmation message.
- 5. Enforce to write the change to your devices.

How to Create a Service Group

ExtremeCloud IQ - Site Engine Policy lets you create service groups into which you can group Local and Global services. A service group can contain any number of services, as well as other service groups. A service can be a part of more than one group.

Instructions on:

- Creating a Service Group
- Adding Services to a Service Group
- Removing Services from a Service Group

Creating a Service Group

- 1. In ExtremeCloud IQ Site Engine, select the **Control** tab.
- 2. Open the **Policy** tab and select **Roles/Services** > **Service Repository** left-panel tab. Expand the **Local Services** or **Global Services** tab.
- 3. Right-click on the Service Groups folder and select **Create Service Group**. This opens the Create window where you can enter a name for the new service group.
- 4. Type the service group name in the highlighted box and select **OK**. You can now <u>add services</u> to the service group. After a service group has been created at the top level under the Service Groups folder, it can be added to another service group.

Adding Services to a Service Group

A service group can contain any number of services, as well as other service groups. You can add services to a service group by

- 1. Right-click the service group from which you wish to remove services, and select Add/Remove Services.
- 2. In the Add/Remove Services window, select the services or service groups you want to add to the service group, and select the **Right Arrow** button.
- 3. Select **OK**.

Removing Services from a Service Group

Use the following steps to remove a service or service group from a service group. Removing a service from a service group does not delete the service itself. If you want to delete the service itself, see Deleting a Service. Keep in mind that if you change the contents of a service group, ExtremeCloud IQ - Site Engine automatically updates the services list for any role that the service group is associated with, affecting the rules in the role.

1. Right-click the service group from which you wish to remove services, and select Add/Remove Services.

- 2. In the Add/Remove Services window, select the services or service groups you want to remove from the service group, and select the **Left Arrow** button.
- 3. Select **OK**.



How to Create or Modify a Rule

Traffic Classification rules enable you to assign a class of service and/or access control (VLAN membership) to network traffic, depending on the traffic's classification type. Classification types are based on layers 2, 3, and 4 of the OSI model, and traffic is classified according to specific layer 2/3/4 information contained in each frame. For more information, see Traffic Classification Rules.

A rule has two main parts: Traffic Description and Actions. The Traffic Description identifies the type of traffic to which the rule pertains. Actions specify whether that traffic is assigned class of service, access control, or both.

In order to create a rule, you must first create a service with which to associate it.

Instructions on:

- Creating a Rule
- Disabling/Enabling a Rule
- Deleting a Rule

Creating a Rule

When you create a rule using the Rule tab, you first create and name the rule using the **Create Rule** menu option, then define its characteristics in the right panel. You can also use the right panel to modify an exiting rule's characteristics.

- 1. In the Policy tab left panel, select the Roles/Services > Service Repository tab.
- 2. Expand either the **Local** or **Global Services** folder, depending on whether the rule is going to be used locally or by all users.
- 3. Expand either the **Service Groups** or **Services** folder and select the service for which you want to create a rule.
- 4. Right-click the service and select **Create Rule**.
- 5. In the Create Rule window, enter a name for the rule and select the rule type. Select **OK**. The rule is created in the left-panel tree.
- 6. Select the rule to and use the associated right-panel **Rule** tab to define the rule. Refer to the Rule tab Help topic for information on configuring the rule.
- 7. Enforce to write the new information to the devices.

Disabling/Enabling a Rule

In the **Policy** tab, you can disable and enable individual or multiple rules. You can also disable and enable all the rules associated with a service, or all the rules for all the services in a service group. The rule icon in the left panel displays a red X if the rule is disabled.

Disabling a rule is an alternative to deleting and recreating it. If you disable a rule, it is temporarily unavailable for use by the service with which it is associated. However, the rule can be copied to another service and enabled for that service.

Disabling/Enabling an Individual Rule

You can enable or disable a rule on the Rule tab or by right-clicking on the rule in the **Service Repository** tab and selecting **Disable Rule(s)** or **Enable Rule(s)**.

- 1. In the **Policy** tab left panel, select the **Roles/Services > Service Repository** tab.
- 2. Expand either the **Local** or **Global Services** folder, depending on whether the rule is going to be used locally or by all users.
- 3. Expand either the **Service Groups** or **Services** folder and select the service for which you want to create a rule.
- 4. Select the rule you want to disable or enable. The Rule tab opens in the right panel.
- 5. Select **Enable** or **Disable** in the **Rule Status** field. Disabling the rule turns on the red X on the rule icon in the left panel, and re-enabling it turns it off.
- 6. Enforce to write the new information to the devices.

Disabling/Enabling the Rules for a Service or Service Group

If a service is associated with more than one service group, disabling or enabling the rules for the service in one service group will disable/enable the rules for the service in the other service groups of which the service is a part.

- 1. In the **Policy** tab left panel, select the **Roles/Services > Service Repository** tab.
- 2. Expand either the **Local** or **Global Services** folder, depending on whether the rule is used locally or by all users.
- 3. Right-click the service or service group containing the rules you want to disable or enable and select **Disable Rule(s)** or **Enable Rule(s)**.
- 4. Select **Yes** to confirm the change.
- 5. Enforce to write the new information to the devices.

Deleting a Rule

Deleting a rule removes the rule from a service. If the service is also part of a service group, the rule is deleted there as well, so be sure the rule is not needed before you delete it.

- 1. In the Policy tab left panel, select the Roles/Services > Service Repository tab.
- 2. Expand either the **Local** or **Global Services** folder, depending on whether you are deleting a rule used locally or by all users.
- 3. Right-click the rule you want to delete, and select **Delete**.
- 4. Select Yes to confirm, then OK to clear the confirmation message. The rule is deleted wherever it exists.
- 5. Enforce to write the new information to the devices.
- Traffic Classification Rules
- Edit Rule Window
- Rule Tab



How to Define Rate Limits

The **Policy** tab allows you to create and define rate limits as components of a class of service. Rate limits are used to control the transmit rate at which traffic enters and exits ports in your network.

The **Policy** tab uses role-based rate limits that are tied directly to roles and rules, and are written to a device when the role/rule is enforced.

Instructions on:

- Defining Rate Limits
- Removing a Rate Limit

Defining Rate Limits

Rate limits are defined within a class of service and associated with a specific role via a rule action or as a role default. When role-based rate limits are implemented, all traffic on the port that matches the rule with the associated rate limit cannot exceed the configured limit. If the rate exceeds the configured limit, frames are dropped until the rate falls below the limit.

The rate limit remains on the port only as long as the role using the rate limit is active on the port either as the authenticated role or as the port's default role.

- 1. Open the Class of Service > CoS Components left-panel tab on the Policy tab.
- 2. Right-click the Rate Limits left-panel tab and select Create Rate Limit.
- 3. Create a new rate limit using the **Rate Limit** tab.
- 4. Select the desired CoS and in the **Class of Service** left-panel tab. Select the **View/Edit** button for the appropriate rate limit to open the Create Rate Limit/Shaper window.

- 5. Fill out the Create Rate Limit/Shaper window:
 - a. Specify the desired rate limit.
 - b. Select the action you would like performed if the rate limit is exceeded:
 - Generate System Log on Rate Violation a syslog message is generated when the rate limit is first exceeded.
 - Generate Audit Trap on Rate Violation an audit trap is generated when the rate limit is first exceeded.
 - Disable Port on Rate Violation the port is disabled when the rate limit is first exceeded.

NOTE: N-Series Gold devices do not support rate limit notification.

c. Select OK.

The rate limit appears in the CoS Configuration table mapped to the CoS.

Role-based rate limits are written to your devices when you enforce the role that includes them.

Removing a Rate Limit

Rate limits remain on a port only as long as the role using the rate limit is active on the port either as the authenticated role or as the port's default role. To remove a rate limit, you must delete it from the **Policy** tab and then enforce. This removes the rate limit from any roles with it is associated.

- 1. Select the Class of Service > CoS Components > Rate Limits left-panel tab on the Policy tab.
- 2. In the right-panel table, right-click on the rate you want to remove.
- 3. Select **Delete**.
- 4. Enforce.

NOTE: If you simply select **None** from the drop-down list, it un-maps the rate from the class of service but it does not remove the rate limit.



How to Create a Class of Service

The **Policy** tab lets you define classes of service (CoS) that can include one or more of the following components: an 802.1p priority, an IP type of service (ToS) value, drop precedence, rate limits, and transmit queue configuration.

Initially, the Class of Service Configuration window (available from the **Policy** tab **Class of Service** left-menu tab) is pre-populated with eight static classes of service, each associated with one of the 802.1p priorities (0-7). You can use these classes of service as is, or configure them to include ToS, rate limit, and/or transmit queue values. In addition, you can also create your own classes of service.

After you have created and defined your classes of service, they are then available when you make a class of service selection for a rule action (**Rule** tab), a role default (**General** tab), or an automated service (**Automated Service** window).

It is recommended that you read Getting Started with Class of Service before creating your classes of service.

Instructions on:

- Creating a Class of Service
- Creating Class of Service Port Groups
- Deleting a Class of Service

Creating a Class of Service

The basic components for a class of service include an 802.1p priority, an IP type of service (ToS) value, drop precedence, rate limits, and transmit queue configuration.

Use the following instructions to create a new class of service using the Class of Service Configuration window.

- 1. Open ExtremeCloud IQ Site Engine and select Control tab > Policy tab > Class of Service left-menu tab.
- 2. Right-click the **Class of Service** tab tree and select **Create COS** from the menu. The Create window opens.
- 3. Enter the name for the CoS in the **Name** field and select **OK**. The new class of service opens in the right panel.
- 4. Select the **Edit** button to enter a description for the CoS.
- 5. Select the **Edit** button next to the **Transmit Queue** field to open the Edit Transmit Queue window, from which you can select a transmit queue for the class of service. If you would like to select a different transmit queue for each port type, select the **Select Q/Port Type** option. Then, when you select **OK**, a window opens where you can specify a different transmit queue for each port type.
- 6. Select an 802.1p priority from the drop-down list to choose the priority (0-7 with 7 being the highest priority).
- 7. Select the **Edit** button to select the ToS option to associate an IP ToS (Type of Service) value with the class of service, if desired (see IP Type of Service for more information). Enter a value in the **Type of Service (ToS)** field.
- 8. Specify a Drop Precedence, if necessary. The Drop Precedence is used in conjunction with the Flex-Edge feature available on K-Series and S-Series (Release 7.11 or higher) devices. Flex-Edge provides the unique capability to prioritize traffic in the MAC chip as it enters the switch. When the Class of Service is assigned to a policy role, and that role is applied to a port via a MAC source address mapping or the port default role, the drop precedence dictates the internal priority (within the MAC chip) that will be used for packets received on the port. If congestion occurs, packets with a high drop precedence are discarded first. Therefore, if a packet is important, it should have a low drop precedence. Refer to the K-

- Series or S-Series Configuration Guide for more information on the Flex-Edge feature and drop precedence.
- 9. If desired, use the Rate Limiting/Rate Shaping section to select a port inbound, outbound, and transmit queue rate limit to associate with the class of service. Select View/Edit next to the IRL Port Group Mappings or ORL Port Group Mappings to open the CoS Rate Limit Mappings tab of the Rate Limit Port Groups window where you can add, edit, or delete a rate limit. The rate limit you select here applies to all IRL/ORL port groups. Select the View/Edit button next to TXQ Port Group Shapers field to open the CoS Transmit Queue Mappings tab to configure transmit queue mappings.
- 10. If you have ExtremeWireless Controllers on your network, you see an option to select inbound and outbound user rate limits to associate with the class of service. User rate limits specify the bandwidth given to each individual user on a port. Currently, user rate limits are only available for wireless controllers.
- 11. Select Open/Manage Domain > Save Domain. The class of service is created and is listed in the Class of Service tab.

After a class of service has been created, you can double-click in the Class of Service Configuration table to modify its characteristics, if necessary.

Creating Class of Service Port Groups

The **Policy** tab provides the ability to create rate limit port groups that let you group together ports with similar rate limiting requirements. For example, you might want to create a class of service where your edge ports would receive one rate limit while your core ports would receive a different rate limit. With port groups, you can create a single class of service that assigns a different rate limit to each group.

It also provides the ability to create transmit queue shaper port groups that enable you to isolate certain kinds of sensitive network traffic so that you can give it a high transmit queue priority. For example, ports on a router might be grouped together and configured with a specific rate shaping parameter. A transmit queue port group can contain multiple port queue types (for example, 4-queue ports and 16-queue ports) depending on the type of devices on your network.

Initially, all ports are grouped into a Default port group. When you create new port groups, you add ports from the Default group into your newly defined port groups.

The following instructions are for creating new port groups for an existing class of service.

- Open the Class of Service left-panel tab and select the Inbound Rate Limit Port Groups, Outbound Limit Port Groups, or Transmit Queue Port Groups tab, depending on the type of port group you want to create.
- 2. Right-click the tab and select **Create Port Group** to create the desired group type: rate limit (RL) port group or transmit queue (TxQ) shaper port group.

 The Create window opens.
- 3. Enter a name for the port group and select **OK**.

- 4. The new port group displays in the **Class of Service** left-panel tab under the appropriate port group type.
- 5. Right-click on the new port group in the left-panel tab and select Add/Remove Ports.
- 6. The Add/Remove Ports window opens with the ports in the Default port group displayed in the left panel. Add ports to the new port group by selecting the ports in the left-panel, then selecting the port group in the right panel, and selecting Add/Move To. Select OK to save the changes and close the window.
- 7. Select Save Domain in the Open/Manage Domain drop-down list.

Deleting a Class of Service

- 1. Open the Class of Service tab.
- 2. Right-click the class of service you want to remove, and select **Delete**.
- 3. Select **OK** to confirm that you want the class of service removed.
- 4. Select Save Domain in the Open/Manage Domain drop-down list.

How to Configure Transmit Queues

The **Policy** tab allows you to configure transmit queues as a component of a <u>class of service</u> (CoS).

There are two transmit queue configuration capabilities:

- Transmit Queue Configuration Allows you to set the transmit queue associated with the class of service.
- TxQ Shaper Transmit Queue Rate Shapers let you pace the rate at which traffic is transmitted out of a transmit queue.

These two capabilities are configured in the <u>Class of Service tab</u> available from the **Policy** tab.

For more information, see the section on transmit queues in <u>Getting Started with Class of Service</u>.

Instructions on:

- Transmit Queue Configuration
- Transmit Queue Rate Shapers

Transmit Queue Configuration

Transmit queues represent the hardware resources for each port used in scheduling packets for egressing the device. By default, the static classes of service 0-7 map to transmit queues 0-7. The actual transmit queue number can vary depending on the number of queues supported by the port.

The Priority column in the Class of Service Configuration window displays the actual transmit queues associated with the class of service for each port type. Double-click in the column to see a drop-down list where you can select a new transmit queue for all port types, or select a different transmit queue for each individual port type.

TIP: For more detailed information, refer to the tooltip that displays when you hover the cursor over the Queue column.

Transmit Queue Rate Shapers

Rate shapers let you pace the rate at which traffic is transmitted out of a transmit queue. Packets received above the configured rate are buffered rather than dropped. Only when the buffer fills are packets dropped.

The following steps describe how to configure rate shapers in the **Policy** tab:

- 1. In the **Class of Service** left-panel tab, select the class of service where you want to configure the transmit queue.
- 2. Select the **Edit** button beside the **Transmit Queue** field and select the desired Transmit Queue from the drop-down list.
- 3. Select Open/Manage Domain > Save Domain to save the configuration change to the database.

For more information, see the section on transmit queues in <u>Getting Started with Class of</u> Service.

NOTE: A rate shaper is associated to a specific transmit queue, not a CoS. This means that the 1) you should select the queue you want to use for a CoS first, then set the shaper and 2) all CoS using that queue uses the same rate shaper. Associating a rate shaper to a transmit queue is accomplished via the CoS - Transmit Queue Mappings tab. For additional information, see the CoS - Transmit Queue Mappings Tab (Transmit Queue Port Group) Help topic.

How to Define Traffic Descriptions

Traffic Classification rules allow you to assign VLAN membership and/or class of service to network traffic based on the traffic's classification type. Traffic descriptions are the part of a rule that defines this classification type. For more information, see Traffic Classification Rules.

The Edit Rule window accessed via the Traffic Description section of the Rule window is used to define traffic descriptions for new rules.

Use the following steps to create a new rule:

- 1. Open the **Control** tab.
- 2. Select the Policy tab.
- 3. In the Policy tab left panel, select the Roles/Services tab.

- 4. Open the Service Repository tab and open either the **Local** or **Global Services** tab, depending on the location of the rule being edited.
- 5. Open either the **Service Groups** or **Services** tab and select the service for which you want to create a rule.
- 6. From the menu bar, select **Tools > Create Classification Rule**. You can also right-click the service and select the option from the menu.
 - The Rule opens in the right panel.
- 7. Select the **Edit** button in the Traffic Description area. The Edit Rule window opens.
- 8. Enter the information for the Traffic Description rule. For additional information, see Edit Rule window.
- 9. Enforce to write the new information to the devices.



How to Configure Flood Control

Flood Control provides rate limiting capabilities to CoS to enable certain types of flooded traffic to be dropped. The flood control traffic types are:

- · unknown unicast
- multicast
- broadcast

When Flood Control is enabled, incoming traffic is monitored over one second intervals. A traffic control rate sets the acceptable flow for each type, specified in packets per second. If, during a one second interval, the incoming traffic of a configured type reaches the traffic control rate on the port, the traffic is dropped until the interval ends. Packets are then permitted to flow again until the limit is reached.

By default, Flood Control is disabled for each CoS. Similarly to CoS Port Groups, a different configuration can be assigned for each group. Since Flood Control is shared across all CoS, when Flood Control is enabled on at least one CoS, those rates apply to all ports that have Flood Control enabled.

How to Display Flood Control Port Groups on the CoS Components Tab

- 1. Select the **CoS Components** left-panel tab on the **Class of Service** left-panel tab. The **CoS Configuration** tab opens.
- 2. Verify that the **Flood Control** checkbox is selected.

How to Create a Flood Control Port Group

- 1. From the left-panel menu, open the CoS Components tab and select the Flood Control Port Groups tab.
- 2. Right-click the Flood Control Port Groups tab and select Create Port Groups.

3. In the Create window, enter a name for the Flood Control Port Group and select **OK**. A New Flood Control item is added to the CoS Configuration Window.

How to Enable/Disable Flood Control for a CoS

Flood Control Rate Limits are shared across all CoS. When a Flood Control rate has been enabled on at least one CoS, that is the rate specified for all Flood Control enabled CoS.

- 1. Open the Flood Control Port Groups tab (Class of Service > CoS Components tab) and select a Port Group.
- 2. Select a rate from the drop-down list for the desired Flood Control broadcast traffic type Unicast, Multicast, or Broadcast.
- 3. Select an existing rate or create a new one.
- 4. Open a CoS in the Class of Service left-panel tab, and enable Flood Control for the CoS by selecting the Enable in the Flood Ctrl Status drop-down list.

How to Add/Remove Ports to Flood Control Port Groups

- 1. From the Class of Service left-panel tab, select the CoS Components > Flood Control Port Groups tab.
- 2. Right-click a Flood Control Port Group, and select **Add/Remove Ports**.
- 3. Add or remove the ports in the Add/Remove Ports window.
- Getting Started with Class of Service
- Class of Service Configuration Tab
- How to Create a Class of Service
- How to Define Rate Limits
- How to Configure Transmit Queues
- General Tab (Rate Limit)
- General Tab (Class of Service)

How to Create Global and Island VLANs

The **Policy** tab **VLANs** left-panel tab used for access control are displayed in the Access Control Configuration window. If you have enabled the Policy VLAN Islands feature, there are two tabs in the VLANs tab: Global VLANs and Policy VLAN Islands. Otherwise, only the Global VLANs folder is displayed. For more information on Policy VLAN Islands, see How to Create a Policy VLAN Islands.

The **Policy** tab provides you with one Global Default VLAN, available when you first access the **Policy** tab. You can create additional VLANs by selecting the **Create VLAN** option available when you right-click on the **Global VLANs** tab.

Once a VLAN is created, you can use it as follows:

- as the default access control for a role, using the role **General** tab.
- as an access control action for a rule using the **Rule** tab.
- as an access control action for an automated service, using the Automated Service tab.
- in a Policy VLAN Island, if that feature is enabled.

See Create VLAN Window and Roles for additional information.

Instructions on:

- Creating a VLAN
- Editing an Island VLAN ID
- Deleting a VLAN

Creating a VLAN

- 1. Open the **Policy** tab.
- 2. Select the left-panel VLANs > Global VLANs tab.
- 3. Right select the **Global VLANs** tab and select **Create VLAN** from the menu.
- 4. Fill out the Create VLAN Window to your specifications.
- 5. Select **OK** to create the VLAN and close the Create VLAN window.
- 6. Enforce to write the new information to the devices.

Editing an Island VLAN ID

- 1. Open the **Policy** tab.
- 2. Expand the VLANs > Policy VLAN Islands left-panel tab.
- 3. Select the VLANs tab in the right panel.
- 4. Select the VLAN with which the policy VLAN island is associated in the VLANs section of the window.
- 5. Select the Island VLAN in the VLAN Settings section of the window and select Edit Island VID.
- 6. Enter the new VLAN ID and select **OK**.
- 7. Enforce to write the new information to the devices.

Deleting a VLAN

Deleting a VLAN removes it and its associations with any roles and services from the NetSight database and from the devices.

WARNING: The delete operation immediately removes the VLAN(s) from the devices in the **Devices** tab and could result in serious consequences if the VLANs are used outside the scope of the **Policy** tab.

- 1. Open the **Policy** tab and select the **VLANs** left-panel tab.
- 2. Expand the Global VLANs left-panel tab.
- 3. Right-click on the VLAN you wish to delete and select **Delete** from the menu. A confirmation window opens.
- 4. Select Yes to delete the VLAN.
- 5. Enforce to write the new information to the devices.

How to Create a Policy VLAN Island

VLAN islands enable you to set up, for example, a guest VLAN that restricts the guests in one facility from communicating with guests in another facility. See Policy VLAN Islands for more information.

Instructions on:

- Creating a VLAN Island
- Modifying a VLAN Island
- Deleting a VLAN Island

Creating a VLAN Island

You can create a Policy VLAN Island as follows:

Note: VLANs used in VLAN islands must be Island VLANs.

- 1. Open the **Policy** tab and select the **VLANs** left-panel tab.
- 2. In the left-panel VLANs tab, select the Policy VLAN Islands tab.
- 3. In the right-panel, select the **VLANs** Tab and select **Create** in the VLANs section.
- 4. In the Create VLAN window, enter a name for the VLAN. Select OK.
- 5. Select Open/Manage Domains > Save Domain.

Modifying a VLAN Island

Once you've created a VLAN island, you can change its characteristics using the right-panel tabs as follows:

- To change a VLAN island name: Right-click the island in the VLANs section of the VLANs > Policy VLAN Islands and select Rename.
- To change a VLAN island description: Use the island's Island Topology tab.
- To edit an Island VLAN ID: Use the Edit Island VLAN ID button on the island's VLANs tab.
- To change a VLAN Island Configuration (Base ID, Offset, Naming Convention): Use the Policy VLAN

Islands tab Island Topology tab.

• To add or remove devices from a VLAN island: Use the VLAN Islands Add/Remove Devices window.

Deleting a VLAN Island

You cannot delete the Default Island.

- 1. Open the **Policy** tab and select the **VLANs > Policy VLAN Islands** left-panel tab.
- 2. Select the VLAN island you want to delete in the VLANs section of the right panel.
- 3. Right-click the island you want to delete and select **Delete**.
- 4. Select **Yes** to confirm the deletion.

How to Create a Network Resource

Network Resource groups provide a quick and easy way to define traffic classification rules for groups of network resources such as routers, VoIP (Voice over IP) gateways, and servers. You create a network resource group by defining a list of MAC or IP addresses for the resources you want included in the group.

In addition, you can use Network Resource Topologies to define a different resource list for different groups of devices in your domain. This enables you to set up network resource access based on the location where end users authenticate.

After a network resource group has been defined, you can associate it with an Automated service (see How to Create a Service for more information). The Automated service automatically creates a rule with a specified action (class of service and/or access control), for each resource address in the network resource group. Automated rule types include Layer 2 MAC Address rules, Layer 3 IP Address and IP Socket rules, and Layer 4 IP UDP Port and IP TCP Port rules.

You can also create Global Network Resources shared between all your domains and can be used by global automated services. Network Resource Topologies are not available for Global Network Resources.



How to Create a Network Resource

- 1. From the **Policy** tab, select the **Network Resources** left-panel tab.
- 2. Right-click the Network Resources folder and select **Create Network Resource**. A New Network Resource item is created in the left panel in a highlighted box. (If you want to create a Global Network Resource, select the Global Network Resources folder.)
- 3. Type the resource name in the Create window and select **OK**.

- 4. In the right-panel **General** tab, use the **Edit** button to add a description of the network resource, if desired.
- 5. Select the network resource Type:
 - Layer 2 MAC Define a group of network resources using MAC addresses.
 - Layer 3 IP Define a group of network resources using IP addresses.
- 6. Select the appropriate network resource topology. Network Resource Topologies are used to divide the devices in a domain into groups called islands. You can then define a unique resource list for each island within that topology, allowing user access to resources on the network based on the physical location at which they authenticate. If you are not using topologies to group your devices, select the Domain Wide topology, which contains just one island for all your domain devices.
- 7. For each topology island included in the selected topology, a tab is available where you can list the resources for that specific island. Use the address field (MAC or IP, depending on the selected type) and select the **Add** button to add a new resource to the list.

After a network resources group has been created and defined, it can be associated with an Automated service (see How to Create a Service for more information).

How to Create a Network Resource Topology

- 1. From the **Policy** tab, select the **Network Resources** left-panel tab.
- 2. Right-click the **Network Resource Topologies** left-panel tab and select **Create Network Resource Topology**. A New Network Resource Topology item is created in the left panel in a highlighted box.
- 3. Type the topology name in the highlighted box.
- 4. Expand the topology to see the Default Island, which contains all the devices in the domain.
- 5. Right-click on the topology and select **Create Network Resource Island**. Type in the island name in the highlighted box and select **OK**. Use this step to create all the islands for this topology.
- 6. Select an island and select the **Add Devices** button to open the Add Devices to Resource Island window, where you can move devices from the Default Island to the islands you just created. Select **Add**.
- 7. Set any island as the [Default] island for new devices that are added to the domain by right-clicking the island and selecting **Set Default**.

The Network Resource Topology is available for selection when you create your network resources.

How to Add and Delete Devices

The ExtremeCloud IQ - Site Engine database contains all the devices in your network and displays them in the left-panel device tree. The **Network** tab and the **Policy** tab share a common view of the device tree, except that only devices that support policy are displayed in the **Policy** tab tree. Any changes you make to the devices are reflected in both trees.

Initially, perform a device Discover to populate the database. After devices have been added to the ExtremeCloud IQ - Site Engine database, you must assign the devices to a Policy Domain using the **Policy** tab. As soon as the devices are assigned to a domain, they are automatically

displayed in the **Policy** tab device tree. Only devices assigned to the domain you are currently viewing are displayed. For more information, see How to Create and Use Domains.

After you have initially added your devices, you can use the **Policy** tab's Add Device window to add a single device to the database and the current domain.

Instructions on:

- Adding a Single Device
- Deleting Devices from the Database

Adding a Single Device

You can add a single device to the ExtremeCloud IQ - Site Engine database using the **Policy** tab's Add Device window. When you add a device, it is assigned to the current domain and automatically listed in the left-panel device tree. Specify the device's SNMP profile. This information is used by the **Policy** tab to access and manage the device.

- 1. Select the **Devices** tab.
- 2. Select Devices folder, right-click and select **Assign Devices to Domain**. The Add Device window opens.
- 3. Enter the IP address of the device you want to add.
- 4. Use the drop-down list to select one of the SNMP profiles that have been defined for device access. The **Edit** button lets you create a profile if one does not already exist.
- 5. Select the checkbox and enter an SNMP context, if desired.
- 6. Select whether to use the default nickname or select **Specify** to assign a unique nickname to this device.
- 7. To add the device and leave the window open, select **Apply**. To add the device and close the window, select **OK**.

Deleting Devices from the Database

When a device is deleted from the ExtremeCloud IQ - Site Engine database, it is removed from all groups where it is a member in both the **Policy** tab device tree (and any other ExtremeCloud IQ - Site Engine plugin applications).

NOTE: If you want to remove a device from a domain without deleting it from the database, you must use the Assign Devices to Domain window. For more information, see Removing Devices from a Domain.

To delete devices from the ExtremeCloud IQ - Site Engine database:

- 1. Open the **Network** tab, select the device being deleted from the Devices table.
- 2. Right-click the device and select **Device > Delete Device** from the menu. A confirmation message advises that you are deleting the device from the ExtremeCloud IQ Site Engine database.
- 3. Select **Yes** to delete the device.

How to Create a Port Group

The **Policy** tab allows you to group ports into user-defined port groups, similar to the way you can group services into service groups. Port groups enable you to configure multiple ports on the same device or on different devices, simultaneously. A port can be a member of more than one group.

When you create a user-defined port group, you select individual ports to add to the group.

The **Policy** tab also provides you with Pre-Defined Port Groups which are automatically populated according to port characteristics. See Pre-Defined Port Groups for more information.

Instructions on:

- Creating a Port Group
- Adding Ports to a Port Group
- Removing Ports from a Port Group

Creating a Port Group

- 1. In the left panel, select the **Devices > Port Groups** tab.
- 2. Right-click on the Port Groups folder and select Create Port Group. This opens the Create window.
- 3. Enter a Name and select OK.

Adding Ports to a Port Group

You can add ports directly from the port group:

- 1. Select the left-panel **Devices > Port Groups** tab. Expand the User-Defined Port Groups folder and select a port group.
- 2. Right-click the port group and select Add/Remove Ports from the menu.
- 3. In the Add/Remove Ports window, select the ports you want to add to the port group in the Devices list and select **Add to Group** to move the port to the Group Port Membership list.
- 4. Select **OK**

Removing Ports from a Port Group

This procedure applies to user-defined port groups.

- 1. In the left-panel **Devices > Port Groups** tab, right-click the port group from which you wish to remove a port, and select **Add/Remove Ports**.
- 2. In the Add/Remove Ports window, select the ports you want to remove from the port group, and select **Remove**.
- 3. Select OK.

Alternatively, you can right-click a single port under the port group in the left panel or multiple ports in the right-panel Ports tab, and select **Remove Port(s) from Group**.

ExtremeControl Access Control

The Access Control tab provides secure, policy-based management for the ExtremeControl solution. It configures and manages ExtremeControl gateways, provides user to device location mapping services, generates network endpoint audit reports and interfaces with other security management applications.

Contact your sales representative for information on obtaining an ExtremeCloud IQ - Site Engine software license.

The Access Control tab contains three main navigation trees in the left-panel:

- ExtremeControl Configuration
- ExtremeControl Group Editor
- All ExtremeControl Engines

ExtremeControl Configuration

The ExtremeControl Configuration lets you manage the end-user connection experience and control network access based on a variety of criteria including authentication, user name, MAC address, time of day, and location. ExtremeCloud IQ - Site Engine comes with a default ExtremeControl Configuration which is automatically assigned to your ExtremeControl engines. You can use this default configuration as is, or make changes to the default configuration, if desired.

Configure a <u>registration</u> that forces any new end-system connected on the network to provide the user's identity in a web page form before being allowed access to the network. End users are automatically provisioned network access on demand without time-consuming and costly network infrastructure reconfigurations. In addition, IT operations gains visibility into the end-systems and their associated users (for example, guests, students, contractors, and employees) on the network.

Via the ExtremeControl **Configuration**, you can also configure agent-less or agent-based security posture assessment of endpoints. The **Access Control** tab uses assessment servers to assess and audit connecting end-systems and provide details about an end-system's patch levels, running processes, anti-virus definitions, device type, operating system, and other information critical in determining an end-system's security compliance. End-systems that fail assessment can be dynamically quarantined with restrictive network access to prevent security threats from entering the network.

Assisted remediation is a process that informs end users when their end-systems have been

quarantined due to network security policy non-compliance, and allows end users to safely remediate their non-compliant end-systems without assistance from IT operations. After the remediation steps have been successfully performed and the end-system is compliant with network security policy, the appropriate network resources are allocated to the end-system, again without the intervention of IT operations.

ExtremeControl Group Editor

The ExtremeControl Engine Groups tree presents groups of ExtremeControl engines you configure into engine groups. Information for engine groups is organized into four tabs in the right-panel, each showing different information relating to the engine group selected:

- **Details** Displays basic information about the engine group as well as information about how the engines in the group are configured.
- Switches Shows the switches monitored by the gateway engines in the group and allows you to add, delete, and edit the switch configuration.
- End-Systems Displays end-systems monitored by the ExtremeControl engines in the selected engine group.
- ExtremeControl Engines Displays the ExtremeControl engines added to the engine group. Right-clicking an engine in the table displays a menu from which you can configure the engine. You can also preview the changes you are making to an engine when you enforce by selecting Enforce Preview.

All ExtremeControl Engines

The <u>All ExtremeControl Engines</u> tree displays all of your ExtremeControl engines. Selecting an engine displays information in three tabs:

- **Details** Displays basic information about the engine, provides a summary of the interface, and allows you to disable ExtremeControl authentication and assessment.
- End-Systems Displays end-systems monitored by the ExtremeControl engine.
- Switches Shows the switches monitored by the gateway engine and allows you to add, delete, and edit the switch configuration.

ExtremeControl Configuration Considerations

Review the following configuration considerations when installing and configuring ExtremeCloud IQ - Site Engine ExtremeControl.

- ExtremeControl Configuration Tables
- General Considerations
- Considerations When Implementing Policy Roles

- ExtremeWireless Controller Configuration
- DNS Proxy Functionality for Registration and Remediation

ExtremeControl Configuration Tables

The following tables provide valuable information to help guide you through the deployment of Extreme Networks ExtremeControl for your network. The first table displays suggested ExtremeControl configurations to use for different network deployment circumstances (e.g. type of end-systems on the network, network topology, authentication method deployed, etc.). The second table displays details and information for each of the different suggested ExtremeControl configurations. The information in the tables assumes that DHCP is deployed on the network.

Suggested ExtremeControl Configuration for Different Deployments

Policy/VLA N Switch Configuratio n	Number of Devices Allowed to Connect to Authenticatio n-enabled Edge Port	Type of End- Systems	Authenticati on Method Deployed	Switch Suppor t IEEE 802.1X MIB	Switch Support, Session Timeout and Terminatio n Action RADIUS Attributes	Suggested Configuratio n
- Policy Only (without changing of VLANs)	*	*	*	*	*	А
- VLAN only - Policy and VLAN - Policy Only (with changing of VLANs)	Multiple	Microsoft XP S P1 with KB82259 6 installed ¹	802.1X ²	Yes	*	A
- VLAN only - Policy and VLAN - Policy Only (with changing of VLANs)	Multiple	*	802.1X ²	Yes	*	В

Policy/VLA N Switch Configuratio n	Number of Devices Allowed to Connect to Authenticatio n-enabled Edge Port	Type of End- Systems	Authenticati on Method Deployed	Switch Suppor t IEEE 802.1X MIB	Switch Support, Session Timeout and Terminatio n Action RADIUS Attributes	Suggested Configuratio n
- VLAN only - Policy and VLAN - Policy Only (with changing of VLANs)	Multiple	*	802.1X ²	No	Yes	С
- VLAN only - Policy and VLAN - Policy Only (with changing of VLANs)	Multiple	*	802.1X ²	No	No	D
- VLAN only - Policy and VLAN - Policy Only (with changing of VLANs) [for Enterasys switch]	Multiple	*	MAC Authentication	*	*	В
- VLAN only - Policy and VLAN - Policy Only (with changing of VLANs) [for non- Enterasys switch]	Multiple	*	MAC Authentication	*	Yes	С

Policy/VLA N Switch Configuratio n	Number of Devices Allowed to Connect to Authenticatio n-enabled Edge Port	Type of End- Systems	Authenticati on Method Deployed	Switch Suppor t IEEE 802.1X MIB	Switch Support, Session Timeout and Terminatio n Action RADIUS Attributes	Suggested Configuratio n
- VLAN only - Policy and VLAN - Policy Only (with changing of VLANs) [for non- Enterasys switch]	Multiple	*	MAC Authentication	*	No	D
- VLAN only - Policy and VLAN - Policy Only (with changing of VLANs)	Single	Microsoft or MAC OS	*	*	*	E
- VLAN only - Policy and VLAN - Policy Only (with changing of VLANs)	Single	Linux	*	*	*	F
Wireless Device	Multiple	*	*	*	*	G

^{* =} Any value.

N/A = Not applicable.

 $^{^1}$ For more information on this patch, see the following link: http://support.microsoft.com/default.aspx?scid=kb;enus;KB822596

 $^{^2}$ When 802.1X is implemented to authenticate multiple users on a single switch port, the downstream device providing connectivity to the users must support the forwarding of EAP frames. Unintelligent devices such as repeaters and switches with newer firmware releases should forward EAP frames. However, some switches do not forward EAP frames therefore preventing the 802.1X authentication of multiple users on a single port.

ExtremeControl Configuration Details

Configuration	Port Link Control	Assessing Session Timeout	Assessing Policy Configuration	DHCP Server Configuration Considerations	Other Considerations			
А	Disabled	Disabled	*	No	N/A			
	NOTE: This is the simples	NOTE: This is the simplest of configurations.						
В	Disabled NOTES:	Disabled	Initial Scan Only	- Set short lease times (e.g. 1 min) for the unauthenticated, Assessing, and Quarantine VLANs - Normal lease times can be configured for the Accept (Production) VLANs	N/A			
	When an end-system transitions from the unauthenticated, Assessing, or Quarantine VLAN to another VLAN, the end-system will soon renew its IP address via DHCP to automatically re-establish connectivity to the network. When a compliant end-system on the Production VLAN is subsequently quarantined after failing a re-assessment, the end-system's connectivity to the network will be lost until expiration of the DHCP lease for the Accept (Production) VLANs.							
С	Disabled	Enabled	Initial Scan Only	- Set short lease times (e.g. 1 min) for the unauthenticated, Assessing, and Quarantine VLANs - Normal lease times can be configured for the Accept (Production) VLANs	N/A			
	NOTES: When an end-system transitions from the unauthenticated, Assessing, or Quarantine VLAN to another VLAN, the end-system will soon renew its IP address via DHCP to automatically re-establish connectivity to the network. Furthermore, the end-system will continually reauthenticate to the network while it is being scanned. When a compliant end-system on the Production VLAN is subsequently quarantined after failing a re-assessment, the end-system's connectivity to the network will be lost until expiration of the DHCP lease for the Accept (Production) VLANs.							

Configuration	Port Link Control	Assessing Session Timeout	Assessing Policy Configuration	DHCP Server Configuration Considerations	Other Considerations			
D	Disabled	Disabled	Initial Scan Only	- Set short lease times (e.g. 1 min) for the unauthenticated, Assessing, and Quarantine VLANs - Normal lease times can be configured for the Accept (Production) VLANs	Set short reauthentication interval manually on edge switches (e.g. 2 min)			
		NOTE: This is not a very scalable configuration model, and therefore should not be implemented for a network with a large number of end-systems.						
Е	Enabled	Disabled	*	No	N/A			
	NOTE: End-system will b	e reauthenticated a	nd will renew its IP addr	ess via DHCP with link dowr	n/up execution.			
F	Enabled	Disabled	Initial Scan Only	- Set short lease times (e.g. 1 min) for the unauthenticated, Assessing, and Quarantine VLANs - Normal lease times can be configured for the Accept (Production) VLANs	N/A			
	NOTES: End-system will be reauthenticated with link down/up execution and will automatically re-establish network connectivity via DHCP upon lease expiration of the IP address in the unauthenticated, Assessing, and Quarantine VLANs. When a compliant end-system on the Production VLAN is subsequently quarantined after failing a re-assessment, the end-system will be reauthenticated and will renew its IP address via DHCP with link down/up execution.							
G	Disabled	*	*	*	RFC 3576 Reauthentication Enabled			
	NOTES: ExtremeCloud IQ - Site Engine supports RFC 3576 which provides for forced reauthentication (Force Reauth) of end-systems connected to an RFC 3576-capable switch. RFC 3576 defines new RADIUS messaging that enables the ExtremeControl Gateway to send Disconnect or Change of Authorization (CoA) RADIUS messages to the authenticating switch or AP to force reauthentication on a currently authenticated end-system.							

^{* =} Any value.

N/A = Not applicable.

General Considerations

- Gateway RADIUS Attributes to Send Send RFC 3580 Only Feature. This feature (configured in the
 Add/Edit Switches to Identity and Access Appliance Group panel) lets you specify that an
 ExtremeControl Gateway sends a VLAN (instead of a policy) via RFC 3580-defined RADIUS Tunnel
 attributes to the RFC 3580-enabled switches in your network. Keep in mind the following considerations
 when configuring this feature:
 - Send RFC 3580 Only is not supported on Matrix E7 Devices. Matrix E7 devices should not be configured with the "Gateway RADIUS Attributes to Send" parameter set to RFC 3580 Only.
 - Send RFC 3580 Only does not support end-systems with static IP addresses. The Send RFC 3580 Only feature is not-supported for end-systems with static IP addresses. This is because end-systems transitioned between VLANs must be assigned an IP address on the appropriate subnet to maintain IP connectivity to the network, which is facilitated dynamically through DHCP.
 - Send RFC 3580 Only requires a particular DHCP configuration for Active/Default Role port mode. When the Send RFC 3580 Only feature is configured, the Active/ Default Role port mode on network devices requires a particular DHCP configuration. The DHCP lease time for the pool of IP addresses that corresponds to the default role's VLAN must be short (e.g. less than 1 minute) because the Active/Default Role port mode enables end-systems to obtain IP addresses via the DHCP protocol before they are authenticated to a VLAN.
 - Switch management fails with Send RFC 3580 Only and certain Auth Access Types. Switch management via TELNET/WebView fails with the following configuration in the Add/Edit Switches to Identity and Access Appliance Group window:

Auth Access Type = "Management Access" or "Any Access" Gateway RADIUS Attributes to Send = "RFC 3580 Only"

This is because switches check the "mgmt" attribute in the Filter-ID for Telnet management. To avoid this problem, set the Auth Access Type to "Network Access."

- Enable Port Link Control Option. Port link control is required if you are using VLAN only (RFC 3580) switches or if you are using policy with VLANs on policy-enabled switches. When an end-system is transitioned between VLANs with a new VLAN being assigned to a switch port, the end-system is required to obtain a new IP address for the assigned VLAN. To do this, the ExtremeControl Gateway links down the port (using the ifAdmin MIB), waits the configured amount of time, and then links up the port, causing the end-system to make a new DHCP request and get a new IP address.
 - Port Link Control is not supported on authentication-enabled switch ports providing connectivity
 to multiple end-systems. Do not enable port link control for switches authenticating multiple
 users per port. When an ExtremeControl Gateway is configured to return only the VLAN RADIUS
 attribute, the gateway links down the authenticated port to force the end-system to release and
 then renew the DHCP IP address when port link control is enabled. This action interrupts IP
 connectivity of other authenticated end-systems on the port. If the switch is an Enterasys switch,
 protection is automatically provided by reading the number of users currently on the port prior
 to linking down an port.
 - Port Link Control is only supported on Windows XP or later. Port link control is only supported for end-users that are authenticating from end-systems running Windows XP or later. When an

ExtremeControl Gateway is configured to return only the VLAN RADIUS attribute, the gateway links down the authenticated port to force the end-system to release and then renew the DHCP IP address when port link control is enabled. However, other systems such as NT workstations, do not release their DHCP IP address when the port is linked down. To account for this scenario, disable port link control, set the ExtremeControl Profile to "Use Assessment Policy During Initial Assessment Only," and set the DHCP lease time for the IP address pools that correspond to the VLAN(s) associated to the Quarantine and Assessing access policies, as well as the default VLAN associated to the unauthenticated state of the port, to a low value (e.g. 1 minute). This forces an end-system to send DHCP Request messages every 30 seconds while it is unauthenticated, being assessed, and quarantined. Upon passing assessment, the end-system is dynamically assigned an IP address on the production VLAN shortly after assessment is complete, establishing connectivity to the network on the production VLAN.

• ExtremeControl Gateway DHCP Snooping:

- Option 1: Locate the ExtremeControl Gateway on the same subnet as the DHCP server. If the ExtremeControl engine is in the same subnet (relay router interface) as the end-system, it is able to hear ACK responses from the DHCP server, enabling it to have more accurate DHCP entries unless the relay router (or DHCP server) sends unicast ACK responses directly to the end-system. Note: Whether the ACK response is sent using unicast or broadcast is normally determined by how the end-system requests the packet. If the end-system sends out a DHCP discover/request with a unicast bootp flag, then the DHCP server (or relay router) sends the ACK response using unicast. This is typically what happens. Sometimes, the end-system can request the DHCP discover/request with a broadcast bootp flag set. In this case, the end-system gets the ACK response with broadcast, and the ExtremeControl engine hears the ACK response if it is in the same broadcast domain.
 - The benefit of using option 1 over the helper-address implementation described in option 2, is that the helper-address implementation only gets the requests from the end-systems that might not have the correct IP address. When an ExtremeControl Gateway learns a MAC/IP address pair, it sends a message to all other ExtremeControl Gateways, so only one ExtremeControl Gateway needs to live on each subnet with a DHCP server on it, to leverage this technique.
- Option 2: Add the ExtremeControl Gateway IP address as a helper address on default gateway routers. To increase the accuracy of the MAP-to-IP resolution, the ExtremeControl Gateway listens for DHCP traffic on port 67 and saves the MAC/IP address pairs it learns. In order to receive DHCP traffic, the IP address of any ExtremeControl Gateway must be added as a helper address on default gateway routers on the network. Routers permit multiple IP helper address entries, so the ExtremeControl Gateway's IP address can be added along with the actual DHCP server IP addresses. When an ExtremeControl Gateway learns a MAC/IP address pair, it sends a message to all other ExtremeControl Gateways, so only one ExtremeControl Gateway IP address needs to be added.
- Configure RADIUS settings on 3rd-party switches. You must manually configure the RADIUS settings on your third-party switches communicating to the ExtremeControl Gateway. In addition, make sure that the shared secret on the switches matches the shared secret you entered in the Advanced Switch Settings window. This is the shared secret the switches uses to communicate with ExtremeControl

Gateways.

- Configuring Agent-based Assessment Test Sets with Hotfix Checks. When configuring an Agent-based test set to perform multiple hotfix checks, make sure that the Monitoring Interval is set to at least 5 minutes, so that the assessment agent does not take a lot of CPU cycles trying to monitor these settings.
- Supported desktop browsers for end-systems connecting through ExtremeControl. The following browsers are supported for desktop end-systems connecting to the network through Extreme Networks ExtremeControl:
 - Microsoft Edge
 - Mozilla Firefox 34 and later
 - Google Chrome 33.0 and later
- Supported mobile browsers for end-systems connecting through ExtremeControl. The following browsers are supported for mobile end-systems connecting to the network through the Mobile Captive Portal of Extreme Networks ExtremeControl:
 - Microsoft Edge
 - Microsoft Windows 10 Touch Screen Native (Surface Tablet)
 - iOS 9+ Native
 - Android 4.0+ Chrome
 - Android 4.4+ Native
 - Dolphin
 - Opera

A native browser indicates the default, system-installed browser. Although this can be Chrome (Android), this also includes the default, system-controlled browser used for a device's Captive Network Detection. Typically, this is a non-configurable option for Wi
NOTES: Fi Captive Network Detection, but default Android, Microsoft of iOS devices are tested for compatibility with the Mobile Captive Portal.

A mobile device can access the standard (non-mobile) version of the Captive Portal using any desktop-supported browsers available on a mobile device.

- For other browsers, the Mobile Captive Portal requires the browser on the mobile device be compatible with Webkit or Sencha Touch. To confirm compatibility with Webkit or Sencha Touch, open http://<ip_of_engine>/mobile_screen_preview using your mobile web browser. If the browser is compatible, the page displays properly.
- RADIUS Configuration on E1 Devices. The ExtremeControl engine opens an SSH/Telnet session on the E1 device and enable RADIUS by running a script of CLI commands. CLI credentials for the device are obtained from the device profile and must be configured in the Authorization/Device Access tool.

- RADIUS Authentication and Accounting Configuration on ExtremeXOS/Switch Engine Devices.
 ExtremeCloud IQ Site Engine uses CLI access to perform RADIUS configuration operations on ExtremeXOS/Switch Engine devices. CLI credentials for the device are obtained from the device profile and must be configured in the Authorization/Device Access tool.
- RADIUS Accounting Configuration on Fixed Switching Devices. ExtremeControl uses CLI to configure RADIUS accounting on Enterasys fixed switching devices (A-Series, B-Series, C-Series, D-Series, G-Series, and I-Series). CLI credentials for the device are obtained from the device profile and must be configured in the Authorization/Device Access tool. This does not apply to A4, B5, and C5 devices running firmware version 6.81 and higher. Those devices support RADIUS accounting configuration using SNMP. For more information, see How to Enable RADIUS Accounting.

Considerations When Implementing Policy Roles

This section describes the communication that takes place between ExtremeControl engines and end-systems connecting to the network. This communication should be taken into account when defining and deploying policy roles and rules on your network. It is particularly critical because certain policy roles and rules can discard traffic that is necessary for communication between the end-system and the engine. For example, in a Guest policy role, NetBIOS traffic is probably discarded, but doing so could impact the MAC to IP resolution process.

Review the following information and verify that the policy roles and rules deployed on your network will permit the required communication between end-systems and your ExtremeControl engines.

IP resolution via NetBIOS

MAC Resolution via NetBIOS

ExtremeControl engine UDP Port 137 <==> End-System Port 137

Remediation and Registration

ExtremeControl engine (TCP or UDP) Port 80 <==> End-System Port (determined on the client) - HTTP

ExtremeControl engine (TCP or UDP) Port 443 <==> End-System Port (determined on the client) - HTTPS

ExtremeControl Agent Discovery via HTTP

ExtremeControl engine Port TCP 8080 <==> End-System Port (determined on the client) ExtremeControl Agent Heartbeat via HTTPS

ExtremeControl engine Port TCP 8443 <==> End-System Port (determined on the client)

ExtremeControl Agent-less Assessment

All ports determined by the selected test set.

The following software is optional and can be installed with agent-less Assessment: SAMBA add-on enabled

TCP Ports 149 and 195, and UDP Ports 137 and 138.

End-System Reachability Test (Assessment Configurations - does not apply to agent-based assessment)

ICMP Ping Test => ICMP Protocol (1), ICMP Type (8)

TCP Ping Test => Default TCP Ports: 21, 22, 23, 25, 79, 80, 111, 135, 139, 445, 497, 515, 548, 1025, 1028, 1029, 1917, 5000, 6000, 9100

ExtremeWireless Controller Configuration

- The NAS IP address used for the wireless controller should be either the management IP address or an IP address of one of its physical data ports, or all zeros to force ExtremeControl (ExtremeControl) to use the source IP. If a logical IP address is used, then ExtremeControl is unable to reauthenticate end-systems.
- If you have configured Assisted Remediation, you must perform the following steps if your network includes wireless controllers:
 - Enable the "ToS override for ExtremeControl" option configured through Wireless Manager in the Edit WLAN Service > Authentication Mode Configuration > Settings window.
 - If Policy Manager is **not** being used to configure policy on the wireless controller, use Wireless Manager to manually add the following rule to the VNS Quarantine, Assessing, and Unregistered filters to permit HTTP traffic to pass through (IN/OUT) the controller when end-systems are proxied to the Internet during remediation.

 0.0.0.0/0 tcp port 80 (Allow traffic In/Out)
 - If Policy Manager **is** being used to configure policy for the wireless controller, use the Classification Rule Wizard to add an "Allow HTTP" rule to a service currently included in your Quarantine, Assessing, and Unregistered policy roles. The rule would be a traffic classification type "IP TCP Port Destination" with the TCP type set to HTTP (80) and the Access Control set to "Permit Traffic."

DNS Proxy Functionality for Registration and Remediation

ExtremeControl (ExtremeControl) Gateway engines provide DNS proxy functionality for use in networks that are deploying registration and/or remediation, but cannot configure the policy-based routing that is required to redirect network traffic to the web portal. Using DNS proxy, any end-system that needs to be redirected to the remediation and registration web portal has its DNS packets spoofed to direct all web page requests to the ExtremeControl Gateway engine. This enables networks that do not have a router to deploy registration and remediation.

Basic Operation

To set up DNS proxy, the ExtremeControl engine is configured as a secondary DNS server in the DHCP scope, in addition to the primary DNS server on the network. When an end-system is required to register or undergo remediation, access to the primary DNS server is blocked and the end-system sends its DNS requests to the DNS proxy on the ExtremeControl Gateway engine.

The DNS proxy must determine whether to spoof the packet or forward the request to the primary DNS server. If the end-system is unregistered or quarantined, the DNS proxy spoofs the DNS packet and send back a DNS response to the end-system with the ExtremeControl engine IP address. This redirects the end-system traffic to the web portal where the end user can

register or remediate. After the end user has registered or remediated their end-system, their DNS requests are forwarded to the primary DNS server.

For third-party devices, a dynamic ACL is configured to block access to the primary DNS server for end-systems undergoing registration or remediation. This causes the DNS requests to be sent to the DNS proxy. The DNS proxy determines whether spoofing is necessary or not by checking the state of the end-system in the database. If the end-system is unregistered or quarantined, the DNS proxy spoofs the DNS packet.

To permit access to hosts or domains for any protocol other than http, you must add the host or domain to the list of allowed web sites configured in the Network Settings view of the ExtremeControl Edit Portal Configuration window. The DNS proxy uses this list of permitted domains to determine if the end-system is permitted access to the requested domain. This can be useful if you want to enable end-systems to perform specific functions such as anti-virus updates or software updates that run over TCP/UDP ports.

You can also define post authorization assessment behavior using DNS proxy. End-systems in the scan state are granted access according to the assessment settings in your ExtremeControl profile.

- If an assessment policy is **not** defined, the user is permitted access while being scanned.
- If an assessment policy is defined for initial assessment only, the user is permitted access if they passed the last scan. If the first or last scan resulted in quarantine, the user is redirected to the ExtremeControl Gateway.
- If an assessment policy is defined for all assessments, the user is redirected to the ExtremeControl Gateway.

Backup DNS Server

Because the DNS proxy forwards DNS requests to the primary DNS server, it is important to configure a backup DNS server on your network, in case the primary server is down. The DNS proxy polls the primary DNS server every minute. If the primary server is down, a backup DNS server is used. If both servers are down, all DNS requests forwarded by the DNS proxy are dropped.

Troubleshooting

DNS proxy error messages are logged in the /var/log/dnsProxy.log file on the ExtremeControl engine. You can enable diagnostics for DNS proxy by going to the ExtremeControl engine administration web page and enabling the DNS Proxy diagnostic group to provide troubleshooting information. Launch the ExtremeControl engine administration web page by using the following URL: https://<ExtremeControlengineIP>:8443/Admin. The default user name and password for access to this web page is "admin/Extreme@pp." Select the Diagnostics page and then the Server Diagnostics page. View the output in the /var/log/dnsProxy.log file or on the Log Files > Server Log web page.

Install the Assessment Agent Adapter on a Nessus Server

This document provides instructions to install the Extreme Networks Assessment Agent Adapter software on a Nessus Server. The Assessment Agent Adapter is required for communication between the ExtremeControl engine and the Nessus server.

NOTE: As of ExtremeCloud IQ - Site Engine version 23.04.11, only Nessus Version 6 is officially supported.

- Go to the Network Management Suite (NMS) Download web page to download the Assessment Agent Adapter: https://extranet.extremenetworks.com/downloads/Pages/NMS.aspx. Select the version of ExtremeCloud IQ - Site Engine you are using.
- 2. Scroll down to find the Identity and Access Tools section of the web page. The install file is named "Assessment Adapter (for 3rd party assessment integration)". Download the file and copy it to the Nessus server.
- 3. Open a shell and "cd" to the directory where you downloaded the install file.
- 4. Change the permissions on the install file by entering the following command at the shell prompt: chmod 755 EXTRAssessmentServerAgentAdapter x.x.x.bin
- 5. Run the install program by entering the following command at the shell prompt: $./\texttt{EXTRAssessmentServerAgentAdapter} \underline{\texttt{x.x.x.x.bin}}$
- 6. The Introduction screen displays. Press Enter.
- 7. Enter Nessus as the agent type to install. Press **Enter**.
- 8. The Choose Install Folder screen displays, where you can choose the installation folder or directory. Enter an absolute path or press **Enter** to accept the default installation folder /root/AssessmentAgent. The installer requires 100 MB of memory. If the installation folder does not have enough memory, an error displays.
- 9. The Pre-Installation Summary screen displays. This screen shows you the locations you have chosen for the installation process and disk space requirements. Review this information to ensure its accuracy.

 Press Fnter
- 10. The Nessus Server Information screen displays. You must enter information in several fields in this screen.
- 11. Enter the port on which the Nessus daemon is running. The default value is 1241. Press Enter.
- 12. Enter the username you created when you installed the Nessus server. Press **Enter**. If you did not create a user when you installed the Nessus server, from a shell prompt, type: cd /nessus installation directory/sbin followed by
 - nessuscli adduser *username* and follow the prompts to add a user to the application. Press **Enter**.
- 13. Enter the password for the Nessus user. Press **Enter**.

- 14. The SSL Server Information screen displays. Enter the port on which the HTTPS daemon is running. The default port number is 8445. Press **Enter**. The Assessment Agent Adapter begins installing.
- 15. If you are upgrading to a newer version of the Assessment Agent Adapter, you are asked if you want to overwrite several files: launchAS.sh, bin/nessus_cmd, and version.txt. Enter the letter "y" to answer yes and press Enter.
- 16. The Installation Complete screen displays. The installation is complete and the Assessment Agent Adapter has been installed on the server.
- 17. Start the Assessment Agent Adapter as a background process by entering the following command at the shell prompt:

/assessmet agent adapter installation directory/launchAS.sh &

18. Make sure that the Nessus daemon and the Assessment Agent Adapter are started each time the system is started, by adding this command into your rc.local script:

/assessment agent adapter installation directory/launchAS.sh &

- 19. To verify the Assessment Agent Adapter is running on the system, from the shell prompt enter:

 netstat -an | grep port number

 where port number is the port you entered that has the HTTPS daemon running on it. The default value
- for this is 8445. Returned entries containing ESTABLISHED or LISTEN is displayed.

 20. To verify the Nessus application is running on the system, from the shell prompt enter:

ps -eaf | grep nessusd

A return entry similar to: "nessusd: waiting for incoming connections" is displayed. This is an indication that the Nessus process is running correctly on the system.

How to Configure Local RADIUS Termination at the ExtremeControl Engine

This Help topic provides information on how to configure authentication using the ExtremeControl engine RADIUS server to locally terminate 802.1X EAP authentication requests. There are three methods that can be used to do this, depending on the protocol that is used:

- LDAP Authentication Uses a backend Active Directory server or LDAP server, and RADIUS server and client certificates (if required) to authenticate users.
- Local Authentication Uses a local password repository, and RADIUS server and client certificates (if required) to authenticate users.
- RADIUS Certificates only Uses only RADIUS server and client certificates to authenticate users (no password is required).

The following chart lists the protocols that are supported for local RADIUS termination, and shows whether the protocol uses RADIUS certificates and/or passwords to authenticate users. If passwords are required, you can then decide whether to use LDAP or local authentication for password verification. The chart also lists the hash types supported by each protocol for user password encryption. Note that PEAP (TLS) is not supported for local RADIUS termination and is only supported in a proxy RADIUS configuration.

Protocol	RADIUS Certificates Required	Password Required	Supported Password Hash Types
PAP	No	Yes	PKCS5 Reversible, SHA1, NT Hash
CHAP	No	Yes	PKCS5 Reversible
MsCHAP	Yes	Yes	PKCS5 Reversible, NT Hash
PEAP (EAP-MsCHAPv2)	Yes	Yes	PKCS5 Reversible, NT Hash
EAP-TTLS	Yes	Yes	PKCS5 Reversible, SHA1, NT Hash
EAP-TLS	Yes	No	N/A
EAP-MD5	No	Yes	PKCS5 Reversible

Instructions on:

- LDAP Authentication
 - User Authentication Considerations
- Local Authentication
 - User Password Considerations
- Certificate Configuration
 - EAP-TLS Certificate Requirements

LDAP Authentication

LDAP authentication uses a backend Active Directory server or LDAP server defined in your AAA Configuration to authenticate users. Additionally, some protocols also require RADIUS server and client certificates to be used in conjunction with LDAP authentication (see Certificate Configuration).

Before configuring LDAP authentication, read through the User Authentication considerations described below.

User Authentication Considerations

If you are using LDAP authentication, the type of LDAP server you select depends on the protocol you are using. With Active Directory, NAC Manager provides a more feature-rich integration and supports a large number of protocols, while with other LDAP servers such as OpenLDAP, NAC Manager provides a more basic integration with limited protocol support.

Active Directory

Supported Protocols: PAP, MsCHAP, PEAP, EAP-MsCHAPV2, and EAP-TTLS with tunneled PAP.

PAP or EAP-TTLS with tunneled PAP protocols

During the authentication process, the ExtremeControl engine sends an LDAP bind request to the Active Directory domain controller using the password retrieved from the end user's authentication request. Therefore, the LDAP protocol must be permitted between the ExtremeControl engine and the Active Directory domain controller for the authentication process to take place.

MsCHAP, PEAP, and EAP-MsCHAPv2 protocols

These three protocols work with Active Directory (and not other LDAP servers) because they use NT Hash for password encryption, which is the same password hash type used by the Microsoft Active Directory domain controller.

Authentication requests are made by the ExtremeControl engine sending an ntlm_auth request to the Active Directory domain controller. The ExtremeControl engine attempts to join the Active Directory domain using the LDAP configuration and the administrator username and password. In your LDAP configuration, the administrator username used to connect to the LDAP server must be a member of the built-in Domain Administrator group or Account Operators group. (See the Active Directory Permissions section below.)

Additionally, the DNS configuration must be set up so that the ExtremeControl engine can resolve the domain by name. To do this, you should configure the DNS server to be one of the domain controllers for that domain, and verify that the domain name is configured correctly on the ExtremeControl engine. If users authenticate to multiple domains, you must also configure the domains to fully trust each other. Refer to the following Microsoft documentation for information on how to set up domain trusts:

https://technet.microsoft.com/en-us/library/cc740018%28WS.10%29.aspx.

Note: For these protocols to work when the active directory domain server is set to only permit NTLMv2 authentication, your version of Samba must pass a flag during authentication to permit NTLMv1 to work for 802.1x MSCHAPv2 when the AD is set to the highest security setting (NTLMv2 only). On earlier versions, these protocols do not work if the active directory is set to only permit NTLMv2 because these protocols do not use NTLMv2 and the hash passed to NAC Manager is rejected by the active directory server. Permitting only NTLMv2 authentication only works if NAC Manager proxies the 802.1x request to Microsoft IAS/NPS. Microsoft IAS/NPS permits this lower level of authentication because it is in a TLS session, which Microsoft believes makes it as secure as NTLMv2. For more information, see https://technet.microsoft.com/en-us/library/cc772468.aspx

Active Directory Permissions

Active Directory is supported on Windows 2008, Windows 2012, and Windows 2016 systems. ExtremeControl can fail to join Active Directory when accessing as a **Standard Domain User with Descendant Computer Objects** group member.

To enable this functionality, add the following permissions:

- Reset Password
- Validated write to DNS host name
- Validated write to service principal
- Read and write account restrictions
- Read and write DNS host name attributes
- Write servicePrincipalName

Active Directory with User Log On Restrictions

In Active Directory, it is possible to configure an option that restricts a user domain log on to specific computers. This configuration is enforced during the domain log on process.

In an ExtremeControl environment where users authenticate using 802.1X and NAC Manager is configured to proxy RADIUS requests, no additional configuration is required. The 802.1X authentication process completes normally and the determination of whether the user is permitted to log on to the domain from the specific computer is enforced at that time.

In an ExtremeControl environment where NAC Manager is terminating 802.1X authentications locally, NAC Manager performs an NTLM authentication to authenticate the 802.1X session. This process simulates the domain log on process. Therefore, NAC Manager indicates the incoming authentication request for the user is coming from a computer (the ExtremeControl engine) that the user is not permitted to log on to, and the authentication attempt is rejected.

The solution in this scenario is to add the ExtremeControl engines to the list of computers the user is permitted to log on to. This enables the 802.1X authentication process to complete and successfully authenticate the user. The enforcement of whether the user is permitted to log on to the specific computer takes place during the domain log on process.

Other LDAP Servers

Supported Protocols: PEAP, PAP, and EAP-TTLS with tunneled PAP.

During the authentication process, the ExtremeControl engine attempts an LDAP(S) bind with the LDAP server to authenticate the end user's credentials. Ensure that LDAP(S) between the ExtremeControl engine and LDAP server is not blocked by an ACL or firewall.

Local Authentication

Local authentication uses a local password repository defined in your AAA Configuration to authenticate users. Additionally, some protocols also require RADIUS server and client certificates to be used in conjunction with local authentication (see Certificate Configuration). Before configuring local authentication, read through the user password considerations described below.

User Password Considerations

When you add or edit a user in your local password repository, you can specify the password hash type used to encrypt the user's password in the ExtremeCloud IQ - Site Engine and NAC Manager databases. Select from two supported hashing algorithms, depending on the protocol you are using:

- SHA 1 a non-reversible hashing algorithm
 Supported Protocols: PAP and EAP-TTLS with tunneled PAP
- PKCS5 a reversible hashing algorithm
 Supported Protocols: PAP, CHAP, MsCHAP, PEAP, EAP-MsCHAPV2, EAP-TTLS with tunneled PAP, and EAP-MD5

Certificate Configuration

If the protocol you are using requires RADIUS certificates for authentication (see the table above), review the certificate configuration information in this section.

During installation, ExtremeControl generates a unique private key and server certificate for the NAC Manager RADIUS server. This certificate provides basic functionality while you are configuring and testing your NAC Manager deployment. To integrate with the certificate structure you already have on your network, update to a certificate generated by a Certificate Authority that your connecting end-systems are already configured to trust.

In addition, configure the AAA Trusted Certificate Authorities to designate which client certificates can be trusted.

EAP-TLS Certificate Requirements

Server Certificate:

Enhanced Key Usage: Server Authentication (1.3.6.1.5.5.7.3.1) Key Usage:

Digital Signature, Key Encipherment

Client Certificate:

Enhanced Key Usage:

Client Authentication (1.3.6.1.5.5.7.3.2)

Key Usage:

Digital Signature, Key Encipherment

Deploy ExtremeControl in an MSP or MSSP Environment

This topic describes deploying ExtremeControl within an MSP (Managed Service Provider) or MSSP (Managed Security Service Provider) environment,. It includes the following information:

- Configuring ExtremeCloud IQ Site Engine Behind a NAT Router
- Defining Interface Services

Configuring ExtremeCloud IQ - Site Engine Behind a NAT Router

If the ExtremeCloud IQ - Site Engine server is located behind a NAT (Network Address Translation) router, use the following steps to add an entry to the nat_config.text file that defines the real IP address for the ExtremeCloud IQ - Site Engine server. This allows the ExtremeCloud IQ - Site Engine server to convert the NAT IP address received in the ExtremeControl engine response to the real IP address used by the ExtremeCloud IQ - Site Engine server.

NOTE: The text in the nat_config.text file refers to a remote IP address and a local IP address. For this configuration, the NAT IP address is the remote IP address and the real IP address is the local IP address.

- On the ExtremeCloud IQ Site Engine server, add the following entry to the <install directory>/appdata/nat_config.text file.
 <NAT IP address>=<real IP address>
- 2. Save the file.
- 3. Configure your ExtremeControl engines to use the NAT IP address for the IP address of the ExtremeCloud IQ Site Engine server. For information on how to configure or change your engine settings, refer to your ExtremeControl engine Installation Guide.

If you have remote ExtremeCloud IQ - Site Engine clients connecting to the NAT IP address, perform the following additional steps.

1. On the ExtremeCloud IQ - Site Engine server, add the following text to the <install directory>/appdata/NSJBoss.properties file. In the second to last line, specify the hostname of the ExtremeCloud IQ - Site Engine server.

```
# In order to connect to a ExtremeCloud IQ - Site
Engine server behind a NAT firewall or a
# ExtremeCloud IQ - Site
Engine server with multiple interfaces you must define these two
# variables on the ExtremeCloud IQ - Site
```

```
Engine server. The java.rmi.server.hostname
# should be the hostname (not the IP) if multiple IPs are being used
# so that each client can resolve the hostname to the correct IP that
# they want to use as the IP to connect to.
java.rmi.server.hostname=<hostname of ExtremeCloud IQ - Site Engine
server>
java.rmi.server.useLocalHostname=true
```

- 2. Save the file.
- 3. Add the ExtremeCloud IQ Site Engine server hostname to your DNS server.

Defining Interface Services

The advanced interface configuration mode available in ExtremeCloud IQ - Site Engine allows you to define which services are provided by each of the ExtremeControl engine's interfaces. This provides the very granular out-of-band management that is often required in MSP or MSSP environments.

For instructions, see the Interface Configuration Window Help topic.

ExtremeControl Concepts

This Help topic explains some of the concepts you'll need to understand in order to make the most effective use of **Access Control** tab.

Information on:

- Overview of the Access Control Tab
- ExtremeControl Engines
 - Use Scenario
 - ExtremeControl VPN Deployment
- Access Control Tab Structure
 - ExtremeControl Configuration
 - Rule Components
 - ExtremeControl Profiles
 - AAA Configurations
 - Portal Configurations
- Access Policies
- Registration
- Assessment
 - Assessment Remediation
- End-System Zones
- Enforcing
- MAC Locking
- Notifications

Overview of the Access Control Tab

Extreme Networks ExtremeControl is a centralized network access control solution located in the Access Control tab that combines authentication, vulnerability assessment, and location services to authorize network access and determine the appropriate level of service for an end-system. The ExtremeControl solution ensures that only valid users and devices with appropriate security postures at the proper location are granted access to your network. For end-systems which are not compliant with defined security guidelines, the ExtremeControl solution provides assisted remediation, enabling end users to perform self-service repair steps specific to the detected compliance violation.

The Access Control tab is the management component in the Extreme Networks ExtremeControl solution. The Access Control tab and ExtremeControl engines work in conjunction to implement network access control. The Access Control tab provides one centralized interface for configuring the authentication, authorization, assessment, and remediation parameters for your ExtremeControl engines. After these configurations are enforced, the ExtremeControl engines can detect, authenticate, assess, authorize, and remediate end-systems connecting to the network according to those configuration specifications.

ExtremeControl Engines

The ExtremeControl engine is required for all Extreme Networks ExtremeControl deployments. It provides the ability to detect, authenticate, and effect the authorization of end devices attempting to connect to the network. It also integrates with, or connects to, vulnerability assessment services to determine the security posture of end-systems connecting to the network. After authentication and assessment are complete, the ExtremeControl engine effects the authorization of devices on the network by allocating the appropriate network resources to the end-system based on authentication and/or assessment results.

If authentication fails and/or the assessment results indicate a non-compliant end-system, the ExtremeControl engine can either totally deny the end-system access to the network or quarantine the end-system with a highly restrictive set of network resources, depending on its configuration. The ExtremeControl engine also provides the remediation functionality of the ExtremeControl solution by means of the remediation web server that runs on the engine. Remediation informs end users when their end-systems have been quarantined due to network security policy non-compliance, and enables end users to safely remediate their non-compliant end-systems without assistance from IT operations.

Use Scenario

The ExtremeControl Gateway engine provides out-of-band network access control for networks where intelligent wired or wireless edge infrastructure devices are deployed as the authorization point for connecting end-systems. End-systems are detected on the network through their RADIUS authentication interchange. Based on the assessment and authentication results for a connecting device, RADIUS attributes are added/modified during the authentication process to authorize the end-system on the authenticating edge switch. Therefore, the ExtremeControl Gateway can be positioned anywhere in the network topology with the only requirement being that IP connectivity between the authenticating edge switches and the ExtremeControl Gateways is operational.

It is important to note that if the wired edge of the network is non-intelligent (unmanaged switches and hubs) and is not capable of authenticating and authorizing locally connected end-systems, it is possible to augment the network topology to enable implementation of inline ExtremeControl with the ExtremeControl Gateway. This can be accomplished by adding an intelligent edge switch that possesses specialized authentication and authorization features. The Extreme Networks K-, S-, or N-Series switch is capable of authenticating and authorizing

numerous end-systems connected on a single port through its Multi-User Authentication (MUA) functionality, and can be positioned upstream from non-intelligent edge devices to act as the intelligent edge on the network. In this configuration, the K-, S-, or N-Series switch acts as the intelligent edge switch on the network, although not physically located at the access edge.

For end-systems connected to EOS policy-enabled switches, a *policy role* is specified in the **Access Control** tab (policy roles are defined and distributed to those switches by the **Policy** tab) to authorize connecting end-systems with a particular level of network access. For end-systems connected to RFC 3580-compliant switches (Enterasys and third-party), a VLAN is specified in the **Access Control** tab to authorize connecting end-systems with a particular level of network access, facilitated using dynamic VLAN assignment via Tunnel RADIUS attributes.

When a user or device attempts to connect to the network, the end-system is authenticated and assessed according to configurations defined in the Access Control tab. The Access Control tab uses the results of the authentication and assessment to determine if that device meets the requirements for a compliant end-system. If the results of the authentication and security assessment are positive, ExtremeCloud IQ - Site Engine authorizes the end-system with network access by assigning a designated policy role or VLAN on the switch port to which the end-system is connected. If the result of the security assessment is negative, ExtremeCloud IQ - Site Engine restricts network access by assigning the user or device to a Quarantine policy role or VLAN on the switch port until the end-system is remediated and brought into a compliant state. If the result of the authentication is negative, ExtremeCloud IQ - Site Engine can deny all network access for the endpoint as an invalid device or user on the network, setting the switch port to the unauthenticated state.

Depending on the engine model, the ExtremeControl Gateway provides either on-board (integrated) vulnerability assessment server functionality and/or the ability to connect to external assessment services, to determine the security posture of end-systems connecting to the network. (On-board assessment requires a separate license.)

The number of ExtremeControl Gateways you deploy on the network depends on the number of end-systems on the network. The following table displays the number of end-systems supported per ExtremeControl Gateway model. Use this table to help determine the number of gateways to deploy.

Model	Number of End-Systems Supported	Notes
IA-A-20	6000	Configured ExtremeControl Features: Authentication and OS/Device Fingerprinting, but no Registration or Assessment.
	4500	Configured ExtremeControl Features: All features excluding Assessment.
	3000	Configured ExtremeControl Features: All features including Assessment.

	Number of End-Systems	
Model	Supported	Notes
IA-A-300	12000	Configured ExtremeControl Features: Authentication and OS/Device Fingerprinting, but no Registration or Assessment.
	9000	Configured ExtremeControl Features: All features excluding Assessment.
	6000	Configured ExtremeControl Features: All features including Assessment.
IA-V	See Notes	The IA-V is used in conjunction with an ExtremeControl Enterprise license (IA-ES-12K).
NAC-V-20	3000	The NAC-V-20 is a virtual engine and requires an ExtremeControl VM license in the ExtremeCloud IQ - Site Engine Server.
NAC-A-20	3000	
SNS-TAG-ITA	3000	
SNS-TAG-HPA	3000	
SNS-TAG-LPA	2000	

It is important to configure ExtremeControl Gateway redundancy for each switch. This is achieved by configuring two different ExtremeControl Gateway engines as a primary and secondary gateway for each switch. When connection to the primary gateway engine is lost, the secondary gateway is used. Note that this configuration supports redundancy but not load-sharing, as the secondary gateway engine is only used in the event that the primary gateway becomes unreachable. To achieve redundancy with load-sharing for two ExtremeControl Gateways, it is suggested that one half of the switches connecting to the gateways are configured with "ExtremeControl Gateway A" as the primary and "ExtremeControl Gateway B" as the secondary, and the second half are configured with "ExtremeControl Gateway B" as the primary and "ExtremeControl Gateway A" as the secondary. In this way, ExtremeControl Gateways are configured in redundant active-active operation on the network.

ExtremeControl VPN Deployment

Extreme Networks ExtremeControl provides out-of-band support for VPN remote access with specific VPN concentrators (see the Release Notes for a list of supported VPN concentrators). Out-of-band VPN support provides visibility into who and what is accessing the network over VPN. If RADIUS accounting is used, you also have the ability to determine who was on the network at any given time. In the VPN remote access use scenario, the VPN concentrator acts as a termination point for remote access VPN tunnels into the enterprise network. In addition, the Extreme Networks ExtremeControl Gateway engine is deployed to authenticate and authorize connecting end-systems on the network and implement network access control.

The process begins when the user's end-system successfully establishes a VPN tunnel with the VPN concentrator, and the VPN concentrator sends a RADIUS authentication request with the

associated credentials to the ExtremeControl Gateway. The ExtremeControl Gateway proxies the authentication request to a backend authentication server (RADIUS or LDAP) to validate the identity of the end user/device or can authenticate with a local password repository within ExtremeCloud IQ - Site Engine. If authentication fails, the ExtremeControl Gateway can deny the end-system access to the network by sending a RADIUS access reject message to the VPN concentrator.

After the end-system is authenticated, the ExtremeControl Gateway requests an assessment of the end-system, if assessment is configured. After authentication and assessment are complete, the ExtremeControl Gateway allocates the appropriate access control to the end-system based on authentication and/or assessment results. Access control can be implemented using one of two methods. With the first method, access control is applied directly at the VPN concentrator via RADIUS response attributes, if the VPN concentrator supports this. For example, with a Cisco ASA security engine, this can be accomplished by using the filter-ID response attribute to specify the name of a valid ACL.

With the second method, an Extreme Networks K-Series, S-Series, or N-Series device is added between the VPN's internal port and the internal network as a Policy Enforcement Point (PEP). This enables the ExtremeControl Gateway to provide a more granular access control mechanism using IP to Policy Mappings. This method must be used if you are implementing remediation on your network. If the end-system fails assessment, the ExtremeControl Gateway can apply a Quarantine policy on the PEP to quarantine the end-system. When the quarantined end user opens a web browser to any web site, its traffic is dynamically redirected to a Remediation web page that provides steps for the user to execute in order to achieve compliance. After executing the steps, the end user can reattempt network access and start the process again.

Access Control Tab Structure

The Access Control tab components are contained in three major navigation trees.

At the top are the following navigation trees:

- Engine Groups Lists the ExtremeControl engines added to the selected engine group, the endsystems connected to those engines, the switches added to the Gateway engines in the engine group, and general information about the engine group.
- All ExtremeControl Engines Lists all ExtremeControl engines added to ExtremeCloud IQ Site Engine, the end-systems connected to those engines, the switches added to the Gateway engines, and general information about the engine.
- ExtremeControl Configurations Provides options to configure the end-user connection experience and control network access based on a variety of criteria including authentication.

ExtremeControl Configuration

The ExtremeControl Configuration lets you manage the end user connection experience and control network access based on a variety of criteria. The **Access Control** tab comes with a

default ExtremeControl Configuration which is automatically assigned to your ExtremeControl engines. You can use this default configuration as is, or make changes to the default configuration, if desired.

The ExtremeControl Configuration determines what ExtremeControl Profile will be assigned to an end-system connecting to the network. It contains an ordered list of rules that are used by the configuration to assign an ExtremeControl Profile to a connecting end-system based on rule criteria. It also specifies the Default Profile which serves as a "catch-all" profile for any end-system that doesn't match one of the rules. By default, all end-systems match the Default Profile.

When an end-system connects to the network, the rules are evaluated in a top-down fashion, similar to the way an ACL would be evaluated. End-systems that do not match any of the rules are assigned the Default Profile.

Rule Components

The rules defined in an ExtremeControl Configuration provide very granular control over how end-systems are treated as they come onto the network. The following criteria can be used to define the rules used in your ExtremeControl Configuration:

- Authentication Type for example, 802.1X or MAC authentication.
- End-System Groups enables you to group together devices that have similar network access requirements or restrictions. For example, a list of MAC addresses, IP addresses, or hostnames.
- Device Type enables you to group together end-systems based on their device type. The device type can be an operating system family, an operating system, or a hardware type, such as Windows, Windows 7. Debian 3.0. and HP Printers.
- Locations enables you to specify network access requirements or restrictions based on the network location where the end user is connecting. For example, a list of switches, wireless devices, switch ports, or SSIDs.
- Time of Day enables you to specify network access requirements or restrictions based on the day and time when the end user is accessing the network. For example, traditional work hours or weekend work hours
- User Groups enables you to group together end users having similar network access requirements or restrictions. For example, a list of usernames, an LDAP users group, or a RADIUS user group.

For more information, see the Manage Rule Groups window.

ExtremeControl Profiles

ExtremeControl Profiles specify the authorization and assessment requirements for the endsystems connecting to the network. Profiles also specify the security policies applied to endsystems for network authorization, depending on authentication and assessment results.

The Access Control tab comes with ten system-defined ExtremeControl Profiles:

- Administrator
- Allow
- Default
- Guest Access
- Notification
- Pass Through
- Quarantine
- Registration Denied Access
- Secure Guest Access
- Unregistered

If desired, you can edit these profiles or you can define your own profiles to use for your ExtremeControl Configurations. For more information, see the Manage ExtremeControl Profiles window.

AAA Configurations

The AAA Configuration defines the RADIUS servers, LDAP configurations, and Local Password Repository that provide the authentication and authorization services for all end-systems connecting to your ExtremeControl engines. The **Access Control** tab comes with a default Basic AAA Configuration that ships with each ExtremeControl engine. You can use this default configuration as is, or make changes to the default configuration, if desired. For more information, see the Edit Basic AAA Configurations window.

Portal Configurations

If your network is implementing <u>Registration</u> or <u>Assisted Remediation</u>, the Portal Configuration defines the branding and behavior of the website used by the end user during the registration or remediation process. ExtremeControl engines are shipped with a default Portal Configuration. You can use this default configuration as is, or make changes to the default configuration, if desired. For more information, see <u>Portal Configuration</u>.

Access Policies

Access policies define the authorization level that the ExtremeControl assigns to a connecting end-system based on the end-system's authentication and/or assessment results. There are four access policies used in the **Access Control** tab: Accept policy, Quarantine policy, Failsafe policy, and Assessment policy. In your ExtremeControl Profiles, these access policies define a set of network access services that determine exactly how an end-system's traffic is authorized on the network. How access policies are implemented depends on whether your network utilizes ExtremeControl Controller engines and/or ExtremeControl Gateway engines.

For end-systems connected to EOS policy-enabled switches, ExtremeControl Gateway engines inform the switch to assign a policy role to a connecting end-system, as specified by the access

policy. These policy roles must be defined in **Policy** tab and enforced to the EOS policy-enabled switches in your network.

For end-systems connected to RFC 3580-enabled switches, policy roles are associated to a VLAN ID. This enables your ExtremeControl Gateways to send a VLAN ID instead of a policy role to those switches using Tunnel RADIUS attributes.

For ExtremeControl Controller engines, authorization of the end-system is implemented locally on the ExtremeControl Controller engine by assigning a policy role to the end-system, as specified by the access policy. In this scenario, all policy roles must be defined in the ExtremeControl Controller policy configuration.

Here is a description of each the **Access Control** tab access policy, and some guidelines for creating corresponding policy roles in the **Policy** tab.

Accept Policy: The Accept access policy is applied to an end-system when it has been authorized locally by the ExtremeControl Gateway and when an end-system has passed an assessment (if an assessment was required), or if the Accept policy has been configured to replace the Filter-ID information returned in the RADIUS authentication messages. For EOS policy-enabled switches, a corresponding policy role (created in the Policy tab) would allocate the appropriate set of network resources for the end-system depending on their role in the enterprise. For example, you might associate the Accept policy in the Access Control tab to the "Enterprise User" role that is defined in the Policy tab demo.pmd file. For RFC 3580-compliant switches, the Accept access policy can be mapped to the Production VLAN. ExtremeControl Controllers are shipped with a default policy configuration that includes an Enterprise User policy role.

Quarantine Policy: The Quarantine access policy is used to restrict network access to end-systems that have failed assessment. For EOS policy-enabled switches, a corresponding Quarantine policy role (created in the Policy tab) should deny all traffic by default while permitting access to only required network resources such as basic network services (e.g. ARP, DHCP, and DNS) and HTTP to redirect web traffic for Assisted Remediation. For RFC 3580-compliant switches, the Quarantine access policy can be mapped to the Quarantine VLAN. ExtremeControl Controllers are shipped with a default policy configuration that includes a Quarantine policy role.

Failsafe Policy: The Failsafe access policy is applied to an end-system when it is in an Error connection state. An Error state results if the end-system's IP address could not be determined from its MAC address, or if there was an assessment error and an assessment of the end-system could not take place. For EOS policy-enabled switches, a corresponding policy role (created in the Policy tab) allocates a nonrestrictive set of network resources to the connecting end-system so it can continue its connectivity on the network, even though an error occurred in the ExtremeControl Solution operation. For RFC 3580-compliant switches, the Failsafe access policy can be mapped to the Production VLAN. ExtremeControl Controllers are shipped with a default policy configuration that includes a Failsafe policy role.

Assessment Policy: The Assessment access policy can be used to temporarily allocate a set of network resources to end-systems while they are being assessed. For EOS policy-enabled switches, a corresponding policy role (created in the **Policy** tab) should allocate the appropriate

set of network resources needed by the Assessment server to successfully complete its end-system assessment, while restricting the end-system's access to the network.

Typically, the Assessment access policy enables access to basic network services (e.g. ARP, DHCP, and DNS), permits all IP communication to the Assessment servers so the assessment can be successfully completed (using destination IP address "Permit" classification rule), and HTTP to redirect web traffic for Assisted Remediation. For RFC 3580-compliant switches, the Assessment access policy can be mapped to the Quarantine VLAN. ExtremeControl Controllers are shipped with a default policy configuration that includes an Assessing policy role.

It is not mandatory to assign the Assessment policy to a connecting end-system while it is being assessed. The policy role received from the RADIUS server or the Accept policy can be applied to the end-system, enabling the end-system immediate network access while the end-system assessment is occurring in the background. In this case, the policy role or Accept policy (or the associated VLAN for RFC 3580-compliant switches) must be configured to permit access to the appropriate network resources for communication with the Assessment servers.

NOTE: The Assessment server sends an ICMP Echo Request (a "ping") to the end-system before the server begins to test IP connectivity to the end-system. Therefore, the Assessment policy role, the router ACLs, and the end-system's personal firewall must permit this type of communication between end-systems and Assessment servers in order for the assessment to take place. If the Assessment server cannot verify IP connectivity, the Failsafe policy is assigned to the end-system.

For more information, refer to the How to Set Up Access Policies Help topic.

Registration

The Extreme Networks ExtremeControl Solution provides support for Registration, a solution that forces any new end-system connected on the network to provide the user's identity in a web page form before being permitted access to the network, without requiring the intervention of network operations. This means that end users are automatically provisioned network access on demand without time-consuming and costly network infrastructure reconfigurations. In addition, IT operations has visibility into the end-systems and their associated users (e.g. guests, students, contractors, and employees) on the network without requiring the deployment of backend authentication and directory services to manage these users. This binding between user identity and machine is useful for auditing, compliance, accounting, and forensics purposes on the network.

End-system or user groups can be configured to exempt certain devices and users from having to register to the network, based on authentication type, MAC address, or user name. For example, a end-system group for the MAC OUI of the printer vendor for the network can be configured to exempt printers from having to register for network access.

The Registration solution has minimal impact on the end user's experience by initially redirecting guests, contractors, partners, students, or other pre-defined end users to a web page for registering their end-system when it is first connected to the network. After successful

registration, the end-system is permitted access, and possibly assessed for security posture compliance checking, until the registration is administratively revoked.

Registration is supported on ExtremeControl Gateway engines and/or Layer 2 ExtremeControl Controller engines. (Registration is not supported on the Layer 3 Identity and Access Controller engines.) Registration provides flexibility in implementation by offering the following capabilities:

- Determine "valid" end users by prompting each end user for a username with additional information such as full name and email address, or a username and password (for example, email address and student ID number) which can be validated against an existing database on the network.
- Enable end users to register to the network when approved by a "sponsor" who is an internal trusted user to the organization. This is referred to as "Sponsored Registration." With sponsored registration, end users are only permitted to register to the network when approved by a sponsor. Sponsorship can provide the end user with a higher level of access than just guest or web access and enables the sponsor to fine-tune the level of access for individual end users.
- Configure the introductory message for the Registration web page (displayed to end-systems before registering to the network) to state that the end user is agreeing to the Acceptable Use Policy for the network upon registering their device.
- Specify the maximum number of registered MAC addresses per user.
- Control areas on the network where Registration is enabled.
- Provide a web-based administrative interface served over HTTPS where registrations can be viewed, manually added, deleted, and modified by administrators and sponsors without requiring access to the Access Control tab.

The Extreme Networks ExtremeControl Solution utilizes a Registration Web Server installed on the ExtremeControl engine to provide this registration functionality to end-systems. Note that an ExtremeControl engine can implement both assisted remediation and registration concurrently.

There are specific network configuration steps that must be performed when using Registration in your ExtremeControl Solution. In addition, you must configure Registration in the **Access Control** tab.

How Registration Works

Here is a description of how Registration works in the Extreme Networks ExtremeControl (ExtremeControl) Solution:

An unregistered end-system attempts to connect to the network and is assigned the unregistered
access profile without being assessed by the ExtremeControl engine. For example, if connected to a
Layer 2 ExtremeControl Controller, the end-system can be assigned to the "Unregistered" policy as
defined in the ExtremeControl Controller's default policy configuration. If connected to an EOS policyenabled switch, the end-system can be assigned to the "Unregistered" policy as defined in the
ExtremeCloud IQ - Site Engine Policy tab and enforced to the policy-enabled switches. Or, if connected

to an RFC 3580-compliant switch, the end-system can be assigned to the "Unregistered" VLAN.

- The user on the unregistered end-system opens up a web browser to any URL and is redirected to the Registration Web Page served by the ExtremeControl engine.
- The end user registers its end-system on the network by entering information such as username, full name, email, and possibly a password or sponsor's email address into the Registration Web Page, and selecting the "Complete Registration" button.
- The Registration Web Server assigns the end user to an end-system group based on the Registration Behavior configured in the ExtremeControl Configuration.
- The end-system is then automatically re-authenticated to the network by the ExtremeControl engine. Upon re-authentication, the end-system is authenticated, assessed, and authorized as defined by the profile specified in the ExtremeControl Configuration for the newly registered system. If the profile specifies to assess the end-system, an assessment of the end-system takes place at this time.

Assessment


The Extreme Networks ExtremeControl Solution integrates with assessment services to determine the security posture of end-systems connecting to the network. It uses assessment servers to assess and audit connecting end-systems and provide details about an end-system's patch levels, running processes, anti-virus definitions, device type, operating system, and other information critical in determining an end-system's security compliance. End-systems that fail assessment can be dynamically quarantined with restrictive network access to prevent security threats from entering the network.

When an assessment is performed on an end-system, a *Health Result* is generated. For each health result, there can be several *Health Result Details*. A health result detail is a result for an individual test performed during the assessment. Each health result detail is given a score ranging from 1 to 10, and based on this score, the health result is assigned a risk level. The **Access Control** tab uses this risk level to determine whether or not the end-system will be quarantined.

In addition, assessment tests are assigned a *scoring mode* which determines whether the resulting health result detail is applied towards the quarantine decision, or is used only for informational or warning purposes. Informational health result details can be used to gather information about the security risks on your network, while warning health result details enable you to notify end users when they have security risks that should be remediated. Informational or warning health result details have scores, however these health result details do not impact the end-system's overall risk level.

The Access Control tab lets you create multiple assessment configurations that can define different assessment requirements for end-systems. Assessment configurations define the following information:

- What assessment tests to run (determined by the selected test sets).
- What resources to use to run the tests (determined by the selected Assessment Resources).
- How to score assessment results (determined by the selected Risk Level and Scoring Override configurations).



Test sets let you define what type of assessment to execute, what parameters to pass to the assessment server, and which assessment server resources to use. The **Access Control** tab provides three default test sets; one for each type of assessment agent that is either supplied or supported by the **Access Control** tab. You can use these default test sets "as is" or edit them, if desired.

When you define your assessment server resources for a test set, you can specify to balance the assessment load between your all your assessment servers, or, you can specify an assessment server pool. For example, if you have four Nessus assessment servers, you can put server A and server B in server pool 1, and server C and server D in server pool 2. Then, in your test set configuration you can specify which server pool that test set should use.

You can use risk level and scoring override configurations to define how each assessment configuration will interpret an end-system's health results. The risk level configuration determines what risk level is assigned to an end-system (high, medium, or low) based on the end-system's health result details score. The scoring override configuration lets you override the default score and scoring mode assigned to a particular assessment test ID.

After you have defined your assessment configurations, they are available for selection when creating your ExtremeControl Profiles. In addition, the **Access Control** tab provides a default assessment configuration that is already set up with default assessment parameters and is ready to use in your ExtremeControl Profiles.

Before beginning to configure assessment on your network, read through the following information presented in the **Access Control** tab online Help.

How to Set up Assessment - Provides information on the steps that must be performed in the Access
Control tab prior to deploying assessment on your network, including managing your assessment
servers and adding external assessment servers. It also includes basic information on how to use the

default assessment configurations provided by the **Access Control** tab, and enable assessment for your ExtremeControl Configuration.

- ExtremeControl Assessment Phased Deployment Guide This guide describes the phased approach to introducing assessment into your ExtremeControl deployment using Informational, Warning, and Quarantine assessment. The guide also provides information on the Access Control tab tools that can be used to monitor and evaluate assessment results, and diagnose and troubleshoot problems.
- How to Configure Assessment Provides step-by-step instructions for configuring assessment using the
 phased approach described in the ExtremeControl Assessment Phased Deployment Guide. Instructions
 are provided for configuring phased assessment using agent-less or agent-based assessment, or a
 combination of both.
- <u>How to Deploy Agent-Based Assessment</u> If you are deploying agent-based assessment, this Help topic provides the configuration steps specific to deploying agent-based assessment in a Windows and Mac network environment. It includes instructions for configuring agent deployment and provides information about the agent icon and notification messages that appear on the end-user's system. It also includes instructions on performing a managed deployment or installation of the agent.
- How to Set Up Assessment Remediation Because Warning and Quarantine assessment provides endsystem remediation, you must enable remediation for your ExtremeControl Configuration. This Help topic provides the specific steps that must be performed when setting up assisted remediation in your network.

Assessment Remediation

Remediation is a process that informs end users when their end-systems have been quarantined due to network security policy non-compliance, and enables end users to safely remediate their non-compliant end-systems without assistance from IT operations. The process takes place when an end-system connects to the network and assessment is performed. End users whose systems fail assessment are notified that their systems have been quarantined, and are instructed in how to perform self-service remediation specific to the detected compliance violation. After the remediation steps have been successfully performed and the end-system is compliant with network security policy, the appropriate network resources are allocated to the end-system, again without the intervention of IT operations.

The Extreme Networks ExtremeControl Solution implements local Remediation Web Server functionality to provide web notification to end users indicating when their end-systems are quarantined and what remediation steps the end user must take. The Remediation Web Server is installed on the ExtremeControl engine.

There are specific network configuration steps that must be performed when using assisted remediation in your ExtremeControl Solution. In addition, you must configure assisted remediation in the **Access Control** tab. For more information, see How to Set up Assessment Remediation and Portal Configuration Help topics.

How Remediation Works

Here is a description of how assisted remediation works in the Extreme Networks ExtremeControl Solution:

- An end-system connects to the network (where assessment has been configured) and is authorized with the level of network access defined by the Assessment access policy configuration.
- The end-system is assessed by the assessment server for security threats and vulnerabilities.
- When the end-system opens a web browser to any web site, the HTTP traffic is redirected to the ExtremeControl engine and a web page indicating that the end-system is currently being assessed is displayed.
- When the assessment is complete, the assessment server sends the results to the ExtremeControl engine. If the end-system failed assessment, the end-system is authorized with the level of network access defined by the Quarantine access policy configuration.
- When the quarantined end user opens a web browser to any web site, its traffic is dynamically redirected to the ExtremeControl engine.
- The ExtremeControl engine returns a web page formatted with self-service remediation information for the quarantined end-system. This web page indicates the reasons the end-system was quarantined and the remediation steps the end user must take.
- After taking the appropriate remediation steps, the end-user selects a button on the web page and attempts to reconnect to the network. A re-assessment of the end-system is initiated. If the end-system is now compliant with network security policy, the ExtremeControl engine authorizes the end-system with the appropriate access policy. If the end-system is not compliant, the Quarantine access policy is again utilized to restrict the authorization level of the end-system and the process starts again.
- After a specified number of attempts and/or maximum time to remediate have expired, the end user can be redirected to a web page requiring them to contact the helpdesk for further assistance, and a notification is sent to the helpdesk system with information regarding the non-compliant end-system.

End-System Zones

The Access Control tab end-system zones enable you to group end-systems into zones, and then limit an ExtremeCloud IQ - Site Engine user's access to ExtremeCloud IQ - Site Engine end-system information and configuration based on those zones.

End-system zones are configured and managed in the **Access Control** tab, and are enforced for ExtremeCloud IQ - Site Engine end-system information and configuration.

When an end-system authenticates to the network, ExtremeControl rules are used to assign an ExtremeControl profile and an end-system zone to the end-system. This enables you to use a variety of rule components (such as End-System Groups, Location Groups, and User Groups) to determine which zone an end-system should be assigned to.

You can create any number of end-system zones in your network. An end-system can only be assigned to one zone (but does not have to be assigned to a zone). You can view which zone an

end-system is currently assigned to in the end-systems table in the **Access Control** tab in ExtremeCloud IQ - Site Engine.

A user's authorized zones are determined by their ExtremeCloud IQ - Site Engine user group membership. User groups are created and configured in the ExtremeCloud IQ - Site Engine Authorization/Device Access Tool (accessed from the Tool menu), and authorized zones are assigned to each user group in the **Access Control** tab.

In addition to using end-system zones, you can also limit a user's access to ExtremeCloud IQ - Site Engine operations by assigning authorized rule groups. Whenever a user initiates a change to a rule group, such as adding or removing an end-system to or from a group, a check is performed to verify that the user is authorized to change that rule group. Similar to end-system zones, a user's authorized rule groups are determined by their ExtremeCloud IQ - Site Engine user group membership.

A third component that should be taken into consideration is the ability to limit user access to ExtremeCloud IQ - Site Engine using authorization group capabilities. For example, you can assign a user group the ExtremeCloud IQ - Site Engine End-Systems Read Access capability to enable read-only access to ExtremeCloud IQ - Site Engine end-system information, and use end-system zones to limit which end-systems can be viewed. You can assign a user group the ExtremeCloud IQ - Site Engine End-Systems Read/Write Access capability to enable the ability to modify rule groups, and use rule group authorization to limit which rule groups the user can perform these operations on.

Capabilities are assigned to user groups using the Authorization/Device Access Tool. The Netsight Administrator group is always assigned all capabilities.

For more information, see Authorization Group Capabilities.

End-System Zone Use Cases

Here are several network scenarios where using end-system zones could be beneficial.

- A Service Provider with multiple tenants. If a service provider serves multiple tenants and each tenant has a clearly delineated set of switches, user access can be configured to enable each tenant's IT staff to only view the end-systems connecting to their own switches.
- A large enterprise with network administrator groups. In a large enterprise where specific groups of
 network administrators are responsible for specific groups of switches on shared engines, user access
 can be configured so that each administrator can view reports and other information only for their
 switches and end-systems.
- A large business segmented by business function. In a large enterprise where division of control is not closely tied to switches or engines, user access can be configured so that administrators only have the ability to view and manage the appropriate end user groups.

In each of these scenarios, a restricted set of authorization group capabilities must be used to prevent users from viewing and accessing information that does not pertain to their area.

Enforcing

In the Access Control tab, enforcing means writing ExtremeControl configuration information to one or more ExtremeControl engines. Any time you add or make a change to the ExtremeControl Configuration, the engines need to be informed of the change through an enforce, otherwise the changes do not take effect. When an engine needs to be enforced, the Enforce icon displays on that engine in the left-panel tree.

To enforce, use the **Enforce** All button in the **Enforce** menu at the bottom of the left-hand panel which writes the information to all the ExtremeControl engines. You can enforce to an individual engine or engine group by selecting the **Enforce** menu and selecting **Selection**.

TIP: For a preview of ExtremeCloud IQ - Site Engine is enforcing/updating on an individual engine, right-click the engine and choose **Enforce Preview** from the menu. The <u>Access Control Engine Enforce Preview window</u> displays, which indicates the information changing.

The enforce operation is performed in two stages: first an engine configuration audit is performed and then the actual enforce to engines is performed.

The configuration audit takes place automatically after you start the enforce operation. It looks for a wide-range of engine configuration problems including a review of the ExtremeControl Configuration, ExtremeControl Profile, rule configuration, AAA configuration, and portal configuration. The audit results are displayed in the Enforce window, enabling you to view any warning and error information. To see warning or error details, use the + icon in the left column to expand the Details information (as shown below) or select **Show Details** to open the information in a new window.

If you choose to correct any problems at this point, you must close the Audit Results window. When you have made your changes, select the Enforce All button to start the enforce operation and perform a new audit.

From the Enforce window, you can select the **Enforce All** button to enforce all engines, or use the checkboxes in the Select column to select some of the engines to enforce and select the **Enforce** button. In order for the enforce operation to be carried out, none of the selected engines can have an error associated with it. Even if one of the selected engine has passed the audit, it will not be enforced if other selected engines have errors.

If none of the selected engines have errors, but a selected engine has a warning associated with it, you are given the option to acknowledge the warning and proceed with the enforce anyway. When you acknowledge the warning and select **OK**, the enforce is performed.

TIP: If there are warning messages that are regularly displayed during Enforce engine audits, you can use the <u>Enforce Warning Settings</u> to specify that these messages should be ignored and not be displayed.

The Enforce window displays the enforce operation status, as shown below.

Advanced Enforce Options

In the Enforce window, there are two Advanced enforce options available. The two options can be used for the following situations:

- Force Reconfiguration for All Switches This option can be used if the switch RADIUS settings were manually changed via CLI or the Policy tab. Since Identity and Access does not reconfigure the switches every time there is an enforce, selecting this option forces reconfiguration of RADIUS settings on all switches to ensure they are configured correctly.
- Force Reconfiguration for Captive Portal During an enforce, captive portal settings are not enforced unless they have changed. You can use this option to force reconfiguration of the portal to ensure the state of the captive portal processes.

MAC Locking

MAC Locking lets you lock a MAC address to a specific switch or port on a switch so that the end-system can only access the network from that port or switch. If the end-system tries to authenticate on a different port or switch, it is rejected or assigned a specific policy based on an action that you specify when you create the MAC Lock. Access the Add MAC Lock window to set up your MAC Locks.

NOTE: MAC Locking to a specific port on a switch is based on the port interface name (e.g. fe.5.1). If a switch board is moved to a different slot in a chassis, or if a stack reorders itself, this name will change and break the MAC Locking settings.

Here are some examples of ways to use MAC Locking:

- A university might lock end-users on a specific floor in a dormitory to a switch that services that floor.
- A printer, server, or other end-system could be permitted network access only when it is connected to a
 port specified by IT operations. This prevents security issues that could result if the device was moved to
 a different area of the network.
- A company could lock an IP phone to a specific port on a switch. This would enable exact identification of the phone's location in case an emergency (911) call was placed from the phone.

NOTE: For ExtremeControl Controller Engines.


- -- On Layer 3 ExtremeControl Controllers, do not use MAC Locking to lock a MAC address to the Controller PEP IP address **and** a port on the PEP. You can however, lock a MAC address to the PEP IP and **not** the port, which would restrict movement of the MAC address away from the Layer 3 Controller
- -- On Layer 2 ExtremeControl Controllers, a MAC address can be locked to the Controller PEP IP address and port, or just the PEP IP address, but this only controls the movement of the end-system between the downstream ports on the PEP (IP address and port) and not the actual edge of the network.
- -- On Layer 3 ExtremeControl Controllers, there can be cases where the **Access Control** tab cannot determine the MAC address of the connecting end-system (for example, DHCP is disabled and a firewall is enabled on the end-system, or the end-system is connecting through a VPN), and the MAC address for the end-system is displayed as "Unknown." In these cases, the MAC Locking feature is not supported.

Notifications

Notifications provide the ability for the **Identity and Access** tab to notify administrators or helpdesk personnel of important information through email, Trap, or Syslog messages. These notifications help administrators understand what is going on in their system on a real-time basis. For example, the **Access Control** tab could be configured to send a notification when a new end-system is learned on the network, when a MAC lock is violated, or when a new MAC address is registered on the network.

Access Control

Access Control Configuration provides a central location to view the configuration parameters for all aspects of your ExtremeControl system. Access this window by selecting **Control** > **Access Control**. Expand the tab to display the options:




The following tabs are included in the Configuration tree:

- Configurations
- <u>AAA</u>
- <u>Profiles</u>
- Captive Portals
- Notifications
- <u>Vendor RADIUS Attributes</u>
- Global & Engine Settings

Configurations

Expand Configurations to access to the following Access Control system components.



Each engine group uses one Access Control configuration that contains an ordered list of rules used to determine which Access Control profile is assigned to the end-systems connecting to the engines in that group. Access Control configurations include the following components:

Name

The Name by which the Access Control Configuration is known.

Portal

If your network is implementing <u>Registration</u> or <u>Assisted Remediation</u>, use the Portal Configuration to define the branding and behavior of the website used by the end user during the registration or remediation process.

AAA

AAA configurations define the RADIUS and LDAP configurations, and Local Password Repository that provide the authentication and authorization services to your ExtremeControl engines.

AAA

The <u>AAA</u> tab defines the RADIUS and LDAP configurations that provide the authentication and authorization services to your ExtremeControlengines.

Profiles

The <u>Profiles</u> tab displays ExtremeCloud IQ - Site Engine's system-defined ExtremeControl profiles that define the authorization and assessment requirements for the end-systems connecting to the network.

Captive Portals

The <u>Captive Portals</u> tab enables you to define the branding and behavior of the portal website used by the end user, if your network is implementing registration or Assessment/Remediation.

Notifications

The <u>Notifications</u> tab displays all the notifications you create, and enables you to add, edit, and test specific notification rules. Notifications enable you to create alert actions performed when specific events or triggers take place in ExtremeCloud IQ - Site Engine

Vendor RADIUS Attributes

The **Vendor RADIUS** Attributes tab displays all the vendors and a list of known vendor RADIUS dictionary <u>attributes</u> that have been discovered from the managed engines. Select a vendor name in the table to display the vendor attribute details, including Attribute Name, Attribute Data Type, Attribute Type, and Options.

Add Radius Dictionary to ExtremeControl.

1. Upload the custom RADIUS dictionary to all Access Control engines:

/opt/tag/radius/share/freeradius

2. Update the permissions for the file:

chmod 644 /opt/tag/radius/share/freeradius/*

3. Restart the service:

nacctl restart

- Custom radius dictionaries are not part of the backup. The procedure may need to be repeated after the software upgrade.
- A non-compatible radius dictionary can cause the solution to be non-operational.

NOTES:

- Renaming existing VSAs in radius dictionaries can cause the system to be non-operational.
- Duplicating existing VSAs in radius dictionaries can cause the system to be nonoperational.
- Extreme can not guarantee that third party radius dictionary will work.

Global & Engine Settings

The Global & Engine Settings tab provides you access to the following additional tabs:


 MAC Locking - Use this tab to view settings for locked MAC addresses or to lock a MAC address to a specific switch or port on a switch so that the end-system can only access the network from that port or switch.

- MAC to IP Mappings Use this tab to view MAC to IP address mappings for devices with statically assigned IP addresses, and import a file of MAC to IP mappings to the list. You can also Add, Edit, Delete, and Export mappings from this tab.
- Manage End System Zones

The <u>Engine Settings tab</u>, which is accessible when you expand the Global & Engine Settings tab, to view and configure advanced configuration options for ExtremeControlengines. ExtremeCloud IQ - Site Engine includes a default engine settings configuration. You can also define your own settings to use for your ExtremeControlengines.

Configuration Evaluation Wizard

This Configuration Evaluation Wizard is used to test the rules defined in your Access Control Configuration in order to determine what behavior an end-system encounters when it is authenticated on an Access Control engine. To access this window, select Configurations in the left-panel of the **Access Control** tab, select an Access Control Configuration in the main panel, and select the **Evaluate** button in the toolbar.




User Input

Enter the end-system data on which you are evaluating the Access Control configuration in this tab.

Authentication Results Tab

This tab displays the set of RADIUS servers and LDAP servers by which ExtremeControl processes an end-system request.



Authentication Result Details

- Authentication Rule A description of the authentication type and user name expression used for the AAA entry that the ExtremeControlengine uses to authenticate the end-system. For a Basic AAA Configuration, this is always: Authentication: Any, User Pattern"*". Additionally, indicates whether fall-through functionality is enabled for the Configuration.
- Authentication For MAC authentication requests, this field displays whether the request is authenticated locally or proxied to the RADIUS server.
- LDAP Configuration The LDAP configuration used to obtain any LDAP data for the end-system, if applicable.

Authentication RADIUS Server Details


This section lists the IP address, port, shared secret, timeout, and retries listed for all the RADIUS servers used to authenticate the end-system request, if it needs to be proxied.

Detailed Reasons

This section is only applicable for an Advanced AAA Configuration. It lists why a request passed or failed the definition of each AAA entry as well as whether <u>fall-through functionality</u> is enabled.

Authorization Results Tab

This tab displays information detailing the method by which the end-system is authorized, according to the parameters and rules of the selected Access Control Configuration. The results also factor in any RADIUS user attributes you enter on the **User Input** tab when the evaluation is run.



Authorization Result Details

- Authentication Request Displays whether the ExtremeControl engine processes the request, or reject the request based on a MAC Lock or a rule that assigns an Access Control Profile configured to reject the user.
- Rule Name The name of the rule that the end-system passed.
- NAC Profile The Access Control Profile assigned to the end-system by the rule.
- Assessment Configuration The assessment configuration used by the Access Control Profile, if any.
- MAC Lock The MAC Lock assigned to the end-system, if any.

Authorization Policy Details

This section displays the RADIUS response attributes returned for end-systems in specific states. Possible states are Accept, Quarantine, Assessing, and Failsafe. Expand each state to view the RADIUS attributes. These are the RADIUS attributes returned for the switch IP that is listed in the End-System Details section.

Detailed Reasons

This section lists all the rules from the Access Control Configuration that were evaluated during the endsystem authentication. Rules are only evaluated until one of them is passed. Each rule listing can be expanded to view why the end-system passed or failed that rule.

ExtremeControl Configuration Rules

The Rules panel in the **Access Control** tab displays a list of rules used by the ExtremeControl Configuration to assign an ExtremeControl Profile to a connecting end-system based on rule criteria.

This Help topic provides information for accessing and configuring ExtremeControl Configuration Rules:

- Accessing ExtremeControl Configuration Rules
- Viewing Rules in the Table
- Creating and Editing Rules
- Advanced Location-Based Registration and Web Access Allows you to configure different access
 features for end users based on the location of a connecting end-system, as determined by the location
 groups you have defined for your network.


Accessing ExtremeControl Configuration Rules

Use the following steps to view and edit your ExtremeControl Configuration rules.

- 1. Open the **Control** tab in ExtremeCloud IQ Site Engine.
- 2. Select the Access Control tab.
- 3. In the left-panel tree, expand the Access Control Configurations tree.
- 4. Expand an ExtremeControl Configuration and select Rules. The table of your ExtremeControl rules is displayed in the right panel. See below for an explanation of the table columns.
- 5. Use the toolbar buttons at the top of the right-panel to create a new rule or edit existing rules. See below for a description of each button.

Viewing Rules in the Table

The Rules table displays the rule name, whether the rule is enabled, and summary information about the rule. It also shows the ExtremeControl Profile assigned to any end-system that matches the rule and the portal redirection action, if applicable. Rules are listed in order of precedence. End-systems that do not match any of the listed rules are assigned the Default Catchall rule.



TIP: Right-click a rule in the table to access a menu of options including the ability to edit the ExtremeControl profile and any user groups included in the rule.

Expand (11) and Collapse (12) Icons

Select the icon to expand a Rule in the table to display additional conditions, end-system, profile, and portal details. Select the icon to collapse a row and hide the additional Rule details.

Enabled

This column displays whether the rule is enabled by displaying a check mark icon or disabled, with no check mark. Select the **Edit** button to enable or disable the rule. You cannot disable any of the system rules provided by ExtremeCloud IQ - Site Engine.

Rule Name

This column displays the rule name. Double-click on the rule to open the Edit Rule window where you can edit the rule name, if desired. You cannot change the name of the system rules provided by ExtremeCloud IQ - Site Engine.

Conditions

This column displays the criteria an end-system must meet in order to be assigned the rule, including the authentication method and rule groups that the end-system or user must match. Double-click on the rule to open the Edit Rule window where you can edit the rule criteria, if desired. You cannot change the criteria for the system rules provided by ExtremeCloud IQ - Site Engine. Select a rule group name to open a window where you can edit the group's parameters.

User Group

This column, hidden by default, displays the user group you configured. User groups limit an ExtremeCloud IQ - Site Engine user's access based on the LDAP, RADIUS, or Username group to which they are assigned. To edit the **User Group**, select the user group in the **Conditions** column, which opens the **Add/Edit User Group** window.

Zone

This column displays the end-system zone you configured. End-system zones allow you to group end-systems into zones, and then limit an ExtremeCloud IQ - Site Engine user's access to end-system information and configuration based on those zones.

Actions

This column displays the actions the rule takes when an end-system matches the rule's criteria. This includes the profile assigned to the end-system, the network policy, and the portal configuration the end user sees. If you want to edit an action, select the profile, policy, or portal to open a window where you can make the changes.

Add or remove a column by selecting the down arrow at the right of a column header and selecting a checkbox associated with a column from the Columns menu.

Creating and Editing Rules

Use the Rules toolbar buttons to create, edit, and modify the rules in the table. Any changes made in this table are written immediately to the ExtremeCloud IQ - Site Engine database.



Opens the Create Rule window where you can define a new rule to use in the ExtremeControl configuration.

TIP: To add a new rule at a specific location in the table, select the rule that you want the new rule to follow, right-click and select **Add Rule** after Selection. When you create the new rule and select **OK**, it is added after the selected rule. The selected rule must be a custom (user-defined) rule, or it can be the blocked list or Assessment Warning rule.

Fdit... Edit Rule

Opens the Edit Rule window where you can edit the rule criteria for a selected rule.

Copy... Copy Rule

Opens the Copy Rule window where you can copy the rule criteria of an existing rule for a new rule.

Delete Delete Selected Rules


Deletes any rules selected in the table.



Move rules up and down in the list to determine rule precedence.

Configuration

Opens the Configuration drop-down list:




Refresh

Updates the Rules details.

Apply Group Label

Opens the Apply Group Label window where you can add a group label to selected rules to create a new group. When the group label is applied, the new group appears in the Rules window and is collapsible.



Expand Rule Details

Expands all Rules in the table to display additional conditions, end-system, profile, and portal details.

Collapse Rule Details

Collapses all expanded Rules in the table.

Advanced Locations Advanced Locations

Use the Advanced Locations tab to define location-based access configurations.


Add/Edit Rule

Use this window to add a new rule or edit an existing rule in an ExtremeControl configuration. End-systems that match the criteria selected for the rule are assigned the ExtremeControl profile that is specified.

To access this window:

- 1. Open the **Control** tab in ExtremeCloud IQ Site Engine.
- 2. Select the ExtremeControl tab.
- 3. In the left-panel tree, select ExtremeControl Configurations > Default > Rules. A table of rules for the ExtremeControl configuration is displayed in the right panel.
- Select the Add button in the table toolbar to open the Create Rule window.
 or
 Select a rule in the table and select the Edit button in the toolbar to open the Edit Rule window.

The image below shows a rule created to provide a different ExtremeControl profile for authenticated registered users on mobile devices. Descriptions of the different fields and options in the window are provided below.



NOTES: For the following rule criteria:

- -- If you select **Any** then the criteria is ignored during the rule match process.
- -- If you select the Invert checkbox, it is considered a rule match if the end-system does **not** match the selected value.

Name

Enter a name for a new rule or change the name of an existing rule, if desired.

Rule Enabled

Select this checkbox to enable this rule in the ExtremeControl configuration.

Description

Enter a description of the rule.

Group Label

If this rule is part of a group, select the group name from the drop-down list or enter a new group label here.

Authentication Method

Select the authentication method that end-systems must match for this rule.

User Group

Select the user group that the end user must be a member of to match this rule. Select the Edit button to edit the selections available in this drop-down list.

End-System Group

Select the end-system group that the end-system must be a member of to match this rule. Select the **Edit** button to edit the selections available in this drop-down list.

Device Type Group

Select the device type group that the end-system must be a member of to match this rule. Select the **Edit** button to edit the selections available in this drop-down list.

Location Group

Select the network location (switch and interface) that the end-system must originate from to match this rule.

Time Group

Select a time frame that the connection request must match for this rule.

Profile

Select the ExtremeControl profile assigned to any end-system matching this rule from the drop-down list. Select New to add a new profile in the Create New Profile window. Select Manage from the drop-down list to be redirected to the Engine Group > Switches tab and allows you to make additions or edits to the switches in this engine group.

Portal

Select the portal configuration from the drop-down list to any end-system matching this rule. Select New to add a new portal configuration in the Add New Portal Configuration window. Select Manage from the drop-down list to be redirected to the Engine Group > Switches tab and allows you to make additions or edits to the switches in this engine group.


Zone

This field only displays if you have displayed the Zone column in the ExtremeControl Configuration Rules table. Select the end-system zone assigned to any end-system matching this rule. Enter a new zone name if none exists. See End-System Zones for more information.

Select Save to save your changes.

Add/Edit User to Authentication Mapping

This window lets you add or edit the user to authentication mappings that define your Advanced AAA configurations. You can access this window from the **Add** or **Edit** buttons in the AAA Configuration window.



Authentication Type

Select the authentication type that the end-system must match for this mapping. Note that individual types of 802.1X authentication are not available for selection because at this point in the authentication process, the fully qualified 802.1X authentication type cannot be determined. Select **Any** if you don't want to require an authentication match. Select **802.1X (TTLS-INNER-TUNNEL)** or **802.1X (PEAP-INNER-TUNNEL)** to authenticate via another RADIUS server using an inner tunnel to protect the authentication request.

The Management Login authentication type enables you to set up a mapping specifically for authenticating management login requests, when an administrator logs into a switch's CLI via the

console connection, SSH, or Telnet. This enables you to send management requests to a different authentication server than network access requests go to. This authentication type can be used to authenticate users locally, or proxy them to specific RADIUS or LDAP servers. Make sure that the Management Login mapping is listed above the "Any" mapping in the list of mappings in your Advanced AAA Configuration. In addition, you must set the Auth. Access Type to either "Management Access" or "Any Access" in the Add/Edit Switches window for this authentication type.

User/MAC/Host

Select the **Pattern** radio button and enter the username, MAC address, or hostname that the end-system must match for this mapping. Or, select the **Group** radio button and select a user group or end-system group from the drop-down list. If you enter a MAC address, you can use a colon (:) or a dash (-) as an address delimiter, but not a period (.).

Location

Select the location group that the end-system must match for this mapping, or select "Any" if you don't want to require a location match. You can also add a new location group or edit an existing one.

Authentication Method

Select the authentication method that the end-system must match for this mapping: Proxy RADIUS, LDAP Authentication, or Local Authentication.

Primary RADIUS Server — Use the drop-down list to select the primary RADIUS server for this mapping to use. You can also add or edit a RADIUS server, or manage your RADIUS servers.

Secondary RADIUS Server — Use the drop-down list to select the backup RADIUS server for this mapping to use. You can also add or edit a RADIUS server, or manage your RADIUS servers.

3rd - 8th RADIUS Server — Use the drop-down list to select the backup RADIUS server for this mapping to use. You can also add or edit a RADIUS server, or manage your RADIUS servers.

Inject Authentication Attrs — Use the drop-down list to select attributes to inject when proxying authentication requests to the back-end RADIUS servers. You can also add or edit a RADIUS attribute configuration, or manage your RADIUS attribute configurations. Select **ExtremeGuest** when configuring a Captive Portal that redirects users to ExtremeGuest.

Inject Accounting Attrs — Use the drop-down list to select attributes to inject when proxying accounting requests to the back-end RADIUS servers. You can also add or edit a RADIUS attribute configuration, or manage your RADIUS attribute configurations. Select **ExtremeGuest** when configuring a Captive Portal that <u>redirects users to ExtremeGuest</u>.

LDAP Authentication — If you select LDAP Authentication, specify the LDAP configuration for this mapping to use.

Local Authentication — If desired, select the option to configure a password for all authentications that match the mapping. This option could be used with MAC authentication where the password is not the MAC address. For example, you can have MAC (PAP) authentication configured for all your switches, with the exception of MAC (MsCHAP) authentication configured for a wireless controller. For the

wireless controller, you would add a new AAA mapping with the authentication type set to MAC (MsCHAP), the location set to the wireless controller location group, and the authentication method set to Local Authentication with the password for all authentications set to the static password configured on the wireless controller.


LDAP Configuration

Use the drop-down list to select the LDAP configuration for the LDAP servers on your network that you want to use for this mapping. You can also add or edit an LDAP configuration, or manage your LDAP configurations. You must specify an LDAP configuration if you have selected LDAP Authentication as your authentication method. However, you might also specify an LDAP configuration if you use Proxy RADIUS to a Microsoft NPS server that is running on a domain controller. The domain controller is also an LDAP server that can do RADIUS requests and LDAP requests for users on that server.

LDAP Policy Mapping

Select the LDAP Policy Mapping for this mapping from the drop-down list. If you have selected an LDAP configuration, this option enables you to use a different LDAP policy mapping. This is useful if the LDAP configuration uses user attribute values that overlap with another LDAP configuration. For example, in the case of multiple companies where company A's Sales department uses one policy, but company B's Sales department uses a different policy.

Select Manage from the drop-down list to <u>Add</u>, <u>Edit</u>, Copy, or Delete the LDAP policy mappings for the LDAP configuration:



Fall-through if Authentication Failed

Select the checkbox to authenticate against the next AAA authentication rule in the event the authentication configured as the first AAA authentication rule results in authentication failure or the Directory Service is unreachable. The fall-through functionality only occurs for those rules on which the

checkbox is selected and only in the event the first authentication rule fails. When this checkbox is enabled and an authentication rule fails, the Access Control engine continues checking the end-user against the remaining rules until it finds a matching rule. If it does not find a matching rule, authentication continues using the previous authentication response.

AAA Configurations Panel

The AAA Configurations panel provides a list of your AAA configurations and buttons to add, edit, or delete configurations. AAA configurations define the RADIUS and LDAP configurations that provide the authentication and authorization services to your ExtremeControl engines.

Access the ExtremeControl Configurations panel in the **Control > ExtremeControl** tab by expanding the **ExtremeControl Configurations** tree in the left-panel and expanding the AAA Configurations tree. Your configurations are listed within the tree.


The following configurations are available in the left-panel tree:

LDAP Configuration

This panel lets you view and define the LDAP configurations used in ExtremeCloud IQ - Site Engine. Any changes made are written immediately to the ExtremeCloud IQ - Site Engine database. For more information about LDAP Configuration, visit the Manage LDAP Configuration topic.

Local Password Repository

The local password repository specified for this AAA configuration. ExtremeCloud IQ - Site Engine supplies a default repository that can be used to define passwords for administrators and sponsors accessing the Registration administration web page and the sponsor administration web page. The default password is Extreme@pp.





Use these buttons to add, edit, or delete local repository in the table. Select the **Add button** to open the Add User window, where you can define a new user and password for the Repository. Select the **Edit button** to open the Edit User, window where you can edit the selected user entry. Use the **Delete button** to delete the selected user entry. You cannot delete a user that is

referenced by an Administrative Login Configuration (as configured in the Edit Portal Configuration Window > Administration).

The following columns are displayed in the default Local Password Repository table:

Enabled

Displays whether the user is enabled or disabled. If a user is disabled, they are not able to log in. This feature is useful if you want to enable a user only at certain times, such as when they are onsite. You can enable or disable a user by editing the user entry (select the entry and select the **Edit** button).

First Name

The user's first name (for administrative information only).

Last Name

The user's last name (for administrative information only).

Display Name

The display name is used on the voucher for pre-registration in the captive portal.

Username

The user's login user name.

Repository

The name of the Local Password Repository of which the user is a member.

GIM

Indicates if the local repository is used by the Guest and IoT Manager (GIM).

Description

A description of the repository.

RADIUS Servers

This panel lets you view and define the RADIUS servers used in ExtremeCloud IQ - Site Engine. RADIUS servers can be used in ExtremeCloud IQ - Site Engine server authentication configurations and in ExtremeControl AAA configurations. For more information about RADIUS Servers, visit the Manage RADIUS Server topic.

AAA Configurations

The AAA Configuration defines the RADIUS and LDAP configurations that provide the authentication and authorization services to your ExtremeControl engines. A AAA Configuration can be a basic or advanced configuration. Basic AAA Configurations define the authentication and authorization services for all end-systems connecting to your ExtremeControl engines Advanced AAA configurations allow you to define different authentication and authorization services for different end users based on end-system to authentication server mappings.

This Help topic provides the following information for accessing and configuring the AAA Configuration:

- Accessing the AAA Configuration
- Basic AAA Configuration
- Advanced AAA Configuration

NOTE: Users with a AAA configuration using NTLM authentication to a back-end active directory domain whose passwords expire are prompted via windows to change their domain password.


Accessing the AAA Configuration

Use the following steps to edit or change your AAA Configuration.

- 1. Open the **Control** tab in ExtremeCloud IQ Site Engine.
- 2. Select the Access Control tab.
- 3. Select **AAA Configurations** within the left-panel tree. The AAA Configuration is displayed in the right panel.
- 4. Use the fields in the right panel to edit or modify the configuration. See the sections below for a description of each field and option in the panel.
- 5. Select **Save** to save your changes.

Basic AAA Configuration

Basic AAA Configurations define the RADIUS and LDAP configurations for all end-systems connecting to your ExtremeControl engines.



Authenticate Requests Locally

This option lets you specify that MAC authentication requests are handled locally by the ExtremeControl engine. Select this option if all MAC authentication requests are to be authorized, regardless of the MAC authentication password (except MAC (EAP-MD5) which requires a password that is the MAC address). The Accept policy is applied to end-systems that are authorized locally.

Select one or more MAC authentication types:

- MAC (All) includes MAC (PAP), MAC (CHAP), MAC (MsCHAP), and MAC (EAP-MD5) authentication types.
- MAC (PAP) this is the MAC authentication type used by Extreme Networks wired and wireless devices.
- MAC (CHAP)
- MAC (MsCHAP)
- MAC (EAP-MD5) this MAC authentication type requires a password, which must be the MAC address.

Primary/Backup RADIUS Servers

If your ExtremeControl engines are configured to proxy RADIUS requests to a RADIUS server, use these fields to specify the primary and backup RADIUS servers to use. Use the drop-down list to select a RADIUS server, add or edit a RADIUS server, or manage your RADIUS servers.

LDAP Configuration

Use this field to specify the LDAP configuration for the LDAP server on your network that you want to use in this AAA configuration. Use the drop-down list to select an LDAP configuration, add or edit an LDAP configuration, or manage your LDAP configurations.

Local Password Repository

Use this field to specify the local password repository you want for this AAA configuration. ExtremeCloud IQ - Site Engine supplies a default repository to define passwords for administrators and

sponsors accessing the Registration administration web page and the sponsor administration web page. The default password is Extreme@pp. Use the drop-down list to select a repository.

Advanced AAA Configuration


Advanced AAA configurations allow you to define different authentication and authorization services for different end users based on end-system to authentication server mappings. Mappings can be based on:

- authentication type
- username/user group
- MAC address/end-system group
- hostname/hostname group
- location group
- · authentication method
- RADIUS user group
- LDAP user group

NOTE: LDAP User Group is only available with an Authentication Type of Registration.

For example, in a higher education setting, you may want faculty members authenticating to one RADIUS server and students authenticating to another. You can also create mappings specifically for authenticating management login requests, when an administrator logs into a switch's CLI via the console connection, SSH, or Telnet.

Mappings are listed in order of precedence from the top down. If an end-system does not match any of the listed mappings, the RADIUS request is dropped. Because of this, you might want to use the "Any" mapping (created automatically when you add a new advanced AAA configuration) as your last mapping in the list.



Authenticate Requests Locally for

This option lets you specify that MAC authentication requests are handled locally by the ExtremeControl engine. Select this option if all MAC authentication requests are to be authorized, regardless of the MAC authentication password (except MAC (EAP-MD5) which requires a password that is the MAC address). The Accept policy is applied to end-systems authorized locally.

Use the drop-down list to specify a particular type of MAC authentication:

- MAC (All) includes MAC (PAP), MAC (CHAP), and MAC (EAP-MD5) authentication types.
- MAC (PAP) this is the MAC authentication type used by Extreme Networks wired and wireless devices.
- MAC (CHAP)
- MAC (MsCHAP)
- MAC (EAP-MD5) this MAC authentication type requires a password, and the password must be the MAC address.

Local Password Repository

Use this field to specify the local password repository you want for this AAA configuration. ExtremeCloud IQ - Site Engine supplies a default repository that can be used to define passwords for administrators and sponsors accessing the Registration administration web page and the sponsor administration web page. The default password is Extreme@pp. Use the drop-down list to select a repository.

Join AD Domain

Use the drop-down list to explicitly select which LDAP configuration of the Active Directory domain the ExtremeControl engine joins in order to authenticate users to all Active Directory domains configured for that engine or select **Auto Detect** to let the ExtremeControl engine determine the domain. Auto Detect starts at the first entry set to LDAP Authentication in the table and attempt to join that domain. If it cannot join that domain, it goes to the next entry set to LDAP Authentication and attempt to join that domain, and so on until one succeeds.

You can also join multiple Active Directory domains by selecting **All Domains** and configuring multiple authentication rules with an **Authentication Method** of **LDAP Authentication** in the **Advanced AAA Configuration** tab.

NOTE: There are configuration considerations when joining multiple Active Directory Domains.

User to Authentication Mapping Table

This table lists mappings between groups of users and authentication configurations. The table displays the username to match along with the defined configuration parameters for that mapping. Mappings are listed in order of precedence from the top down. If an end-system does not match any of the listed mappings, the RADIUS request is dropped. Because of this, you might want to use an "Any" mapping as your last mapping in the list. Use the Mappings toolbar buttons to perform actions on the mappings.

▲ Up ▼ Down Move Mappings Up/Down

Move mappings up and down in the list to determine mapping precedence. Mappings are listed in order of precedence from the top down.

Add... Add New Mapping

Opens the Add User to Authentication Mapping window where you can define a new mapping.

Edit... Edit Mapping

Opens the Edit User to Authentication Mapping window where you can edit the selected mapping.

Delete Delete Selected Mappings

Deletes any mappings selected in the table.

AAA Configurations Panel

The AAA Configurations panel provides a list of your AAA configurations and buttons to add, edit, or delete configurations. AAA configurations define the RADIUS and LDAP configurations that provide the authentication and authorization services to your ExtremeControl engines.

Access the ExtremeControl Configurations panel in the **Control > ExtremeControl** tab by expanding the **ExtremeControl Configurations** tree in the left-panel and expanding the AAA Configurations tree. Your configurations are listed within the tree.


The following configurations are available in the left-panel tree:

LDAP Configuration

This panel lets you view and define the LDAP configurations used in ExtremeCloud IQ - Site Engine. Any changes made are written immediately to the ExtremeCloud IQ - Site Engine database. For more information about LDAP Configuration, visit the Manage LDAP Configuration topic.

Local Password Repository

The local password repository specified for this AAA configuration. ExtremeCloud IQ - Site Engine supplies a default repository that can be used to define passwords for administrators and sponsors accessing the Registration administration web page and the sponsor administration web page. The default password is Extreme@pp.





Use these buttons to add, edit, or delete local repository in the table. Select the **Add button** to open the Add User window, where you can define a new user and password for the Repository. Select the **Edit button** to open the Edit User, window where you can edit the selected user entry. Use the **Delete button** to delete the selected user entry. You cannot delete a user that is

referenced by an Administrative Login Configuration (as configured in the Edit Portal Configuration Window > Administration).

The following columns are displayed in the default Local Password Repository table:

Enabled

Displays whether the user is enabled or disabled. If a user is disabled, they are not able to log in. This feature is useful if you want to enable a user only at certain times, such as when they are onsite. You can enable or disable a user by editing the user entry (select the entry and select the **Edit** button).

First Name

The user's first name (for administrative information only).

Last Name

The user's last name (for administrative information only).

Display Name

The display name is used on the voucher for pre-registration in the captive portal.

Username

The user's login user name.

Repository

The name of the Local Password Repository of which the user is a member.

GIM

Indicates if the local repository is used by the Guest and IoT Manager (GIM).

Description


A description of the repository.

RADIUS Servers

This panel lets you view and define the RADIUS servers used in ExtremeCloud IQ - Site Engine. RADIUS servers can be used in ExtremeCloud IQ - Site Engine server authentication configurations and in ExtremeControl AAA configurations. For more information about RADIUS Servers, visit the Manage RADIUS Server topic.

Manage LDAP Configurations

This panel lets you view and define the LDAP configurations used in ExtremeCloud IQ - Site Engine. You can access this panel by selecting LDAP Configurations from the left-panel in the ExtremeControl Configurations > AAA Configurations tree or from <u>AAA Configuration</u>, by selecting the drop-down list in the LDAP Configuration field. Any changes made are written immediately to the ExtremeCloud IQ - Site Engine database.



LDAP Configurations Table

The name of the configuration and the LDAP server connection URLs specified for that configuration.

Test Configuration Button

Use this button to run a connection test for the selected configuration. The connection to the LDAP server is tested and a report on connection test results is provided. There is also a user search that lets you search on a user entry value and display the attributes associated with the user.

Add Configuration Button

Opens the Add LDAP Configuration window where you can define a new LDAP configuration.

Edit Configuration Button

Opens the Edit LDAP Configuration window where you can edit the selected LDAP configuration.

Delete Configuration Button


Deletes the selected LDAP configuration(s).

Add LDAP Configuration

Use the Add LDAP Configuration window to configure the LDAP servers on your network. You can access this window from the Control > Access Control tab. Expand the Configuration > Configurations > AAA > LDAP Configurations folder in the right panel and select Add. You can also access this window from the Manage LDAP Configurations tab. Any changes made in this window are written immediately to the ExtremeCloud IQ - Site Engine database.

NOTE:

If you are using LDAPS, your ExtremeCloud IQ - Site Engine/ExtremeControl environment must be configured to accept the new LDAPS server certificate. For information, see Server Certificate Trust Mode in the Secure Communications Help topic.



Configuration Name

Enter a name for the LDAP configuration.

LDAP Connection URLs

Use this table to add, edit, or delete connection URLs for the LDAP server and any backup servers you have configured. (The backup servers are redundant servers containing the same directory information.) Use the Up and Down arrows to arrange the order that the URLs are listed.

The format for the connection URL isldap://host:portwhere host equals hostname or IP address, and the default port is 389. For example, ldap://l0.20.30.40:389. If you are using a secure connection, the format is ldaps://host:portand the default port is 636. For example, ldaps://l0.20.30.40:636. If you are using LDAPS, yourExtremeCloud IQ - Site Engine/ExtremeControl environment must be configured to accept the new LDAPS server certificate. For information, see Server Certificate Trust Mode in the Secure Communications Help topic.

If you are creating an LDAP configuration for Novell eDirectory, be aware that the eDirectory may require that the universal password lookup be done using LDAPS. If you configure the URL for LDAP only, the lookup may fail.

Authentication Settings

Enter the administrator username and password that will be used to connect to the LDAP server to make queries. The credentials only need to provide read access to the LDAP server. The timeout field lets you specify a timeout value in seconds for the LDAP server connection.

Search Settings

For the three fields, enter the root node of the LDAP server. To improve search performance, you can specify a sub tree node to confine the search to a specific section of the directory. The search root format should be a DN (Distinguished Name).

Schema Definition

Provide information that describes how entries are organized in the LDAP server.

Schema Definition fields:

- User Object Class- enter the name of the class used for users.
- User Search Attribute- enter the name of the attribute in the user object class that contains the user's login ID.
- Keep Domain Name for User Lookup- If selected, this option will allow the full username to be used when looking up the user in LDAP. For example, you should select this option when using the User Search Attribute: userPrincipalName.


If the option is not selected, the domain name will be stripped off the username prior to performing the lookup. For example, you should deselect this option when using the User Search Attribute: sAMAccountName. Two examples of the domain name being stripped off would be: user@domain.com -> user

DOMAIN\user -> user

- User Authentication Type- Specify how the user is authenticated. There are 4 options:
 - LDAP Bind- This is the easiest option to configure, but only works with a plain text password. It is useful for authentication from the captive portal but does not work with most 802.1x authentication types.
 - NTLM Authentication- This option is only useful when the backend LDAP server is really a Microsoft Active Directory server. This is an extension to LDAP bind that uses ntlm_auth to verify the NT hash challenge responses from a client in MsCHAP, MsCHAPV2, and PEAP requests. If you want to run a NTLM Health Check, see NTLM Health Check and Advanced for the additional configuration steps.
 - NT Hash Password Lookup- If the LDAP server has the user's password stored as an NT hash that is readable by another system, you can have ExtremeControl read the hash from the LDAP server to verify the hashes within an MsCHAP, MsCHAPV2, and PEAP request.
 - Plain Text Password Lookup
 — If the LDAP server has the user's password stored unencrypted and that attribute is accessible to be read via an LDAP request, then this option reads the user's password from the server at the time of authentication. This option can be used with any authentication type that requires a password.
- User Password Attribute- This is the name of the password used with the NT Hash Password Lookup and Plain Text Password Lookup listed above.
- Host Object Class- enter the name of the class used for hostname.
- Host Search Attribute- enter the name of the attribute in the host object class that contains the hostname.
- Use Fully Qualified Domain Namecheckbox use this checkbox to specify if you want to use the Fully Qualified Domain Name (FQDN) or just hostname without domain.
- OU Object Classes- the names of the classes used for organizational units.

Advanced

Advanced LDAP Configuration only accessible when User Authentication Type is set to NTLM Authentication. The LDAP configuration information you enter here specifies a user account and domain to the user for the NTLM Health Check. To configure the Health Check tests:



- 1. Configure the interval and timeout for the test. See NTLM Health Check.
- 2. Select NTLM Health Check.
- 3. Enter the **Username**. **Password**. and the **Domain** to use for the health check tests.
- 4. Select OK.

The Access Control Engine expects a positive response from the domain controller for the health check authentication. If timeout happens or a negative response is received, the failover occurs and the Access Control Lost Partial Contact with LDAP Service alarm is generated.

• You should only use the health check in an environment where you have multiple domain controllers deployed.

WARNING:

- The health check should only be enabled if you have experienced this issue.
- The Domain password policy requirement for periodic password check should be disabled for the health check account. The credentials used by the health check should always be working.

Test Button

The connection to the LDAP server is tested and a report on connection test results is provided. There is also a user/host search that lets you search on a user entry or host entry value and display the attributes associated with those values.

Populate Default Values Button

Select from the defaults available from the menu:


- Active Directory: User Defaults- Settings that allow user authentication when ExtremeControl is set to proxy to LDAP and the server is an Active Directory machine.
- Active Directory: Machine Defaults- Settings that allow machine authentication when ExtremeControl is set to proxy to LDAP and the server is an Active Directory machine.
- Open LDAP Defaults -Settings that allow ExtremeControl to verify the user's password via an OpenLDAP server. See the NAC Manager How to Configure PEAP Authentication via OpenLDAP Help topic for information.
- Novell eDirectory Defaults- Settings that allow ExtremeControl to read the universal password from Novell eDirectory. You must configure eDirectory to allow that password to be read. See the NAC Manager <u>How to Configure PEAP Authentication via eDirectory</u> help topic for information.

Edit LDAP Configuration

Use the Edit LDAP Configuration window to configure the LDAP servers on your network. You can access this window from the **Users** tab in the Authorization/Device Access tool, or in NAC Manager from the AAA Configuration window, by selecting an LDAP configuration from the drop-down list in the LDAP Configuration field. Any changes made in this window are written immediately to the ExtremeCloud IQ - Site Engine database.

NOTE:

If you are using LDAPS, your ExtremeCloud IQ - Site Engine/ExtremeControl environment must be configured to accept the new LDAPS server certificate. For information, see Server Certificate
Trust Mode in the Secure Communications Help topic.



Configuration Name

The name for the LDAP configuration you defined.

LDAP Connection URLs

Use this table to add, edit, or delete connection URLs for the LDAP server and any backup servers you have configured. (The backup servers are redundant servers containing the same directory information.) Use the Up and Down arrows to arrange the order that the URLs are listed.

The format for the connection URL is ldap://host:port where host equals hostname or IP address, and the default port is 389. For example, ldap://l0.20.30.40:389. If you are using a secure connection, the format is ldaps://host:port and the default port is 636. For example, ldaps://l0.20.30.40:636. If you are using LDAPS, your ExtremeCloud IQ - Site Engine/ExtremeControl environment must be configured to accept the new LDAPS server certificate. For information, see Server Certificate Trust Mode in the Secure Communications Help topic.

If you are creating an LDAP configuration for Novell eDirectory, be aware that the eDirectory may require that the universal password lookup be done using LDAPS. If you configure the URL for LDAP only, the lookup may fail.

Authentication Settings

Enter the administrator username and password that will be used to connect to the LDAP server to make queries. The credentials only need to provide read access to the LDAP server. The timeout field lets you specify a timeout value in seconds for the LDAP server connection.

Search Settings

For the three fields, enter the root node of the LDAP server. To improve search performance, you can specify a sub tree node to confine the search to a specific section of the directory. The search root format should be a DN (Distinguished Name).

Schema Definition

Provide information that describes how entries are organized in the LDAP server.

Schema Definition fields:

- User Object Class enter the name of the class used for users.
- User Search Attribute enter the name of the attribute in the user object class that contains the user's login ID.
- Keep Domain Name for User Lookup If selected, this option will allow the full username to be used when looking up the user in LDAP. For example, you should select this option when using the User Search Attribute: userPrincipalName.


If the option is not selected, the domain name will be stripped off the username prior to performing the lookup. For example, you should deselect this option when using the User Search Attribute: sAMAccountName. Two examples of the domain name being stripped off would be: user@domain.com -> user

DOMAIN\user -> user

- User Authentication Type Specify how the user is authenticated. There are 4 options:
 - LDAP Bind This is the easiest option to configure, but only works with a plain text password. It is useful for authentication from the captive portal but does not work with most 802.1x authentication types.
 - NTLM Auth This option is only useful when the backend LDAP server is really a Microsoft Active Directory server. This is an extension to LDAP bind that uses ntlm_auth to verify the NT hash challenge responses from a client in MsCHAP, MsCHAPV2, and PEAP requests.
 - NT Hash Password Lookup If the LDAP server has the user's password stored as an NT hash that is readable by another system, you can have ExtremeControl read the hash from the LDAP server to verify the hashes within an MsCHAP, MsCHAPV2, and PEAP request.
 - Plain Text Password Lookup If the LDAP server has the user's password stored unencrypted and that attribute is accessible to be read via an LDAP request, then this option reads the user's password from the server at the time of authentication. This option can be used with any authentication type that requires a password.
- User Password Attribute This is the name of the password used with the NT Hash Password Lookup and Plain Text Password Lookup listed above.
- Host Object Class enter the name of the class used for hostname.
- Host Search Attribute enter the name of the attribute in the host object class that contains the hostname.
- Use Fully Qualified Domain Name checkbox use this checkbox to specify if you want to use the Fully Qualified Domain Name (FQDN) or just hostname without domain.
- OU Object Classes the names of the classes used for organizational units.

Advanced

Advanced LDAP Configuration only accessible when User Authentication Type is set to NTLM Authentication. The LDAP configuration information you enter here specifies a user account and domain to the user for the NTLM Health Check. To configure the Health Check tests:



- 1. Configure the interval and timeout for the test. See NTLM Health Check.
- 2. Select NTLM Health Check.
- 3. Enter the **Username**. **Password**. and the **Domain** to use for the health check tests.
- 4. Select OK.

The Access Control Engine expects a positive response from the domain controller for the health check authentication. If timeout happens or a negative response is received, the failover occurs and the Access Control Lost Partial Contact with LDAP Service alarm is generated.

• You should only use the health check in an environment where you have multiple domain controllers deployed.

WARNING:

- The health check should only be enabled if you have experienced this issue.
- The Domain password policy requirement for periodic password check should be disabled for the health check account. The credentials used by the health check should always be working.

Test

The connection to the LDAP server is tested and a report on connection test results is provided. There is also a user/host search that lets you search on a user entry or host entry value and display the attributes associated with those values.

Populate Default Values


Select from the defaults available from the drop-down list:

- Active Directory: User Defaults Settings that allow user authentication when ExtremeControl is set to proxy to LDAP and the server is an Active Directory machine.
- Active Directory: Machine Defaults Settings that allow machine authentication when ExtremeControl is set to proxy to LDAP and the server is an Active Directory machine.
- OpenLDAP Defaults Settings that allow ExtremeControl to verify the user's password via an OpenLDAP server. See the NAC Manager How to Configure PEAP Authentication via OpenLDAP Help topic for information.
- Novell eDirectory Defaults Settings that allow ExtremeControl to read the universal password from Novell eDirectory. You must configure eDirectory to allow that password to be read. See the NAC Manager <u>How to Configure PEAP Authentication via eDirectory</u> Help topic for information.

Manage RADIUS Servers

This panel lets you view and define the RADIUS servers used in ExtremeCloud IQ - Site Engine. RADIUS servers can be used in ExtremeCloud IQ - Site Engine server authentication configurations and in ExtremeControl AAA configurations.

You can access this panel by selecting RADIUS Servers from the ExtremeControl Configurations > AAA Configurations > RADIUS Servers in the left-panel tree, or from the Configure Device window or AAA Configuration window. Any changes made are written immediately to the ExtremeCloud IQ - Site Engine database.



RADIUS Server IP

The IP address of the RADIUS server.

Auth Port

The UDP port number (1-65535) on the RADIUS server to which the ExtremeCloud IQ - Site Engine server or ExtremeControl engine sends authentication requests; 1812 is the default port number.

NOTE: If you are enforcing to an ExtremeControl engine for an Extreme Management Center version prior to Version 8.5, you must use different ports to configure UDP Auth. and Accounting. UDP will not function if the Auth and Accounting are configured for the same port for previous versions of ExtremeCloud IQ - Site Engine.

The TCP port number (1-65535) on the RADIUS server that the ExtremeCloud IQ - Site Engine server or ExtremeControlengine sends authentication requests to; 1812 is the default port number.

The TLS port number (1-65535) on the RADIUS server that the ExtremeCloud IQ - Site Engine server or ExtremeControlengine sends authentication requests to; 2083 is the default port number.

NOTE: For versions prior to ExtremeCloud IQ - Site Engine Version 8.5, TCP and TLS settings are not supported and cannot be enforced to ExtremeControl engines.

Acct Port

The UDP port number (1-65535) on the RADIUS server to which the ExtremeControl engine sends accounting requests; 1813 is the default port number.

NOTE: If you are enforcing to an ExtremeControl engine for an Extreme Management Center version prior to Version 8.5, you must use different ports to configure UDP Auth. and Accounting. UDP will not function if the Auth and Accounting are configured for the same port for previous versions of ExtremeCloud IQ - Site Engine.

The TCP port number (1-65535) on the RADIUS server that the ExtremeControl engine sends accounting requests to; 1813 is the default port number.

The TLS port number (1-65535) on the RADIUS server that the ExtremeControl engine sends accounting requests to; 2083 is the default port number.

NOTE: For versions prior to ExtremeCloud IQ - Site Engine Version 8.5, TCP and TLS settings are not supported and cannot be enforced to ExtremeControl engines.

Timeout Duration

The amount of time, in seconds, the ExtremeCloud IQ - Site Engine server or ExtremeControl engine waits for the RADIUS server to respond to an authentication or accounting request. Valid values are 2-60 seconds.

Number of Retries

The number of times the ExtremeCloud IQ - Site Engine server or ExtremeControl engine resends an authentication or accounting request if the RADIUS server does not respond. Valid values are 0-20.

Shared Secret

The shared secret used to encrypt and decrypt communication between the ExtremeCloud IQ - Site Engine server or ExtremeControl engine and the RADIUS server. In ExtremeControl, this is also the shared secret used between the switch and the RADIUS server if the ExtremeControl engine is bypassed or if you configured the Management RADIUS Server options when you added the switch.

Show Shared Secrets

When checked, the shared secrets are shown in text. When unchecked, the shared secrets are shown as a string of asterisks.

Used By Button

This button is only available when the panel is launched from ExtremeControl. Opens the RADIUS Server (s) Used By window which shows where the selected servers are in use by AAA configurations.

Add Button

Opens the Add RADIUS Server window where you can define a new RADIUS server.

Edit Button

Opens the Edit RADIUS Server window where you can edit the values for the selected RADIUS server.


Delete Button

Deletes the selected RADIUS server. You cannot delete servers currently in use.

Add/Edit RADIUS Server

Use the Add/Edit RADIUS Server window to configure the RADIUS servers used in your ExtremeCloud IQ - Site Engine applications. RADIUS servers can be used in ExtremeCloud IQ - Site Engine server authentication configurations and in ExtremeControl AAA configurations.

You can access this window from the Manage RADIUS Servers window. Any changes made in this window are written immediately to the ExtremeCloud IQ - Site Engine database.



RADIUS Server IP

The IP address of the RADIUS server.

Response Window

This setting is used by ExtremeControl when proxying a RADIUS request to a backend RADIUS server. ExtremeControl keeps a status on all backend RADIUS servers instead of going to the primary RADIUS server for every request. If a RADIUS server does not respond in the amount of time specified here, that

server is marked as down until it can be verified as being up. See the Health Check section of the Advanced RADIUS Server Configuration window for information on how ExtremeControl determines the health of a RADIUS server.

Authentication Via ExtremeCloud IQ - Site Engine or Captive Portal

Timeout Duration

The amount of time in seconds the ExtremeCloud IQ - Site Engine server or ExtremeControl waits for the RADIUS server to respond to an authentication or accounting request. Valid values are 2-60 seconds. This setting is only used for logging into ExtremeCloud IQ - Site Engine via RADIUS or logging into the ExtremeControl Captive Portal via RADIUS.

NOTE: The ExtremeControl engine times out a RADIUS server if it takes more than "(retries +1) * timeout" or 20 seconds, whichever is greater, for the server to respond. For example, if the number of retries is set to 1 and the timeout duration is set to 2 (the default values), then the engine times out a RADIUS server if it takes longer than 20 seconds to respond, because that is the greater value (20 to 4). If the RADIUS server times out, then ExtremeControl fails over to the backup RADIUS server until it determines that the primary server is back up. At that point, ExtremeControl starts proxying RADIUS requests to the primary server again.

Number of Retries

The number of times the ExtremeCloud IQ - Site Engine server or ExtremeControl engine resends an authentication or accounting request if the RADIUS server does not respond. Valid values are 0-20. This setting is only used for logging into ExtremeCloud IQ - Site Engine via RADIUS or logging into the ExtremeControl Captive Portal via RADIUS.

Configuration

UDP Button

Select the UDP button to configure the UDP port on the RADIUS server to receive authentication and accounting requests.

NOTE: If you are enforcing to an ExtremeControl engine for an Extreme Management Center version prior to Version 8.5, you must use different ports to configure UDP Auth. and Accounting. UDP will not function if the Auth and Accounting are configured for the same port for previous versions of ExtremeCloud IQ - Site Engine.

Auth. Client UDP Port

The UDP port number (1-65535) on the RADIUS server that the ExtremeCloud IQ - Site Engine server or ExtremeControl engine sends authentication requests to; 1812 is the default port number.

Accounting Client UDP Port

The UDP port number (1-65535) on the RADIUS server that the ExtremeControl engine sends accounting requests to; 1813 is the default port number.

TCP Button

Select the TCP button to configure the TCP port on the RADIUS server to receive authentication and accounting requests.

NOTE: For versions prior to ExtremeCloud IQ - Site Engine Version 8.5, TCP settings are not supported and cannot be enforced to ExtremeControl engines.

Auth. Client TCP Port

The TCP port number (1-65535) on the RADIUS server that the ExtremeCloud IQ - Site Engine server or ExtremeControl engine sends authentication requests to; 1812 is the default port number.

Accounting Client TCP Port

The TCP port number (1-65535) on the RADIUS server that the ExtremeControl engine sends accounting requests to; 1813 is the default port number.

RADSec Button

Select the RADSec button to configure the TLS (Transport Layer Security) port on the RADIUS server to receive authentication and accounting requests.

NOTE: For versions prior to ExtremeCloud IQ - Site Engine Version 8.5, TLS settings are not supported and cannot be enforced to ExtremeControl engines.

Auth. Client TLS Port

The TLS port number (1-65535) on the RADIUS server that the ExtremeCloud IQ - Site Engine server or ExtremeControl engine sends authentication requests to; 2083 is the default port number.

Accounting Client TLS Port

The TLS port number (1-65535) on the RADIUS server that the ExtremeControl engine sends accounting requests to; 2083 is the default port number.

Proxy RADIUS Accounting Requests

Select this checkbox to enable the ExtremeControl engine to proxy RADIUS accounting requests to the RADIUS server. This option must be enabled if you are doing RADIUS accounting in an ExtremeControl environment where the primary RADIUS server is being used for redundancy in a single ExtremeControl engine configuration (Basic AAA configuration only).

Change Server Shared Secret

Server Shared Secret

The shared secret is a string of characters used to encrypt and decrypt communication between the ExtremeCloud IQ - Site Engine server or ExtremeControl and the RADIUS server. In ExtremeCloud IQ - Site Engine, this is also the shared secret used between the switch and the RADIUS server if the ExtremeControl engine is bypassed or if you configured the Management RADIUS Server options when you added the switch. The shared secret must be at least 6 characters long; 16 characters is recommended. Dashes are allowed in the string, but spaces are not.

Verify Shared Secret

Re-enter the Server Shared Secret you entered above.

Show Shared Secret


Displays the secret in the Server Shared Secret and Verify Shared Secret fields.

Advanced Button

Use this button to open the Advanced RADIUS Server Configuration window, where you can configure advanced RADIUS settings used by ExtremeControl when proxying access requests to a backend RADIUS server.

Manage RADIUS Attribute Configurations Window

Use this window to view attributes injected when authentication or accounting requests are proxied to a back-end RADIUS server. Attributes you inject provide additional information about the users on your network. You can access the RADIUS Attribute Configurations window from the Add/Edit User To Authentication Mapping window.



Preview With Policy

Presents a preview of the attributes defined for selected attribute configuration.

Name

The names of the available attribute configurations. You cannot edit the name of a configuration.

Add

Select the **Add** button to open the Create New RADIUS Attribute Settings window, which allows you to create a new attribute configuration.

Edit


Select the **Edit** button to open the Edit RADIUS Attribute Settings window, which allows you to edit an existing attribute configuration.

Delete

Select an attribute and select the **Delete** button to remove an existing attribute configuration.

Advanced RADIUS Server Configuration

Use this window to configure advanced RADIUS settings used by ExtremeCloud IQ - Site Engine when proxying authentication requests to a backend RADIUS server. You can access this window by selecting the **Advanced** button at the bottom of the Add/Edit RADIUS Server window.



Username Format

This field is used by ExtremeCloud IQ - Site Engine to determine what format to use for the username when proxying a request to the backend RADIUS server. There are two options:

- Strip Domain Name (default) This option removes a domain name from the username when proxying the request. Select this option unless the backend RADIUS server requires the domain name to be included.
- Keep Domain Name This option keeps any domain names on the username when proxying the request to the backend RADIUS server. If the backend RADIUS server is a Microsoft IAS or NPS server, this option could cause the RADIUS server to time out if a guest comes onto the network with another domain. In that scenario, if the request is proxied to the backend RADIUS server with the domain name, the server does not respond to the request because it is from an unknown domain. Therefore, if you use this option with a Microsoft IAS or NPS server, use an advanced

AAA configuration so that only requests for the desired domain(s) are sent to the backend RADIUS server, and all unknown domains are processed locally so they are rejected.

Require Message-Authenticator

Enable this checkbox if the backend RADIUS server requires a message authenticator to be part of the request. If enabled, ExtremeCloud IQ - Site Engine adds the message authenticator when proxying the request.

Health Check for UDP

ExtremeCloud IQ - Site Engine uses the options in this section to determine how to check the health of a backend RADIUS server, if that server stops responding to requests.

NOTE: For backend RADIUS server options other than UDP (for example, TCP or RADSec), all fields except Revive Interval in the Health Check for UDP are not available.

Use Server-Status Request

When selected, ExtremeCloud IQ - Site Engine attempts to use Server-Status RADIUS packets as defined by RFC 5997, to determine if the backend RADIUS server is up.

Use Access Request

When selected, ExtremeCloud IQ - Site Engine attempts to use an access request message to determine if the RADIUS server is up. The request is made using the username and password specified below. The username and password do not need to be valid, as ExtremeCloud IQ - Site Engine is looking for a response and a reject also works. The username/password fields are provided in case you want to prevent rejects from being logged in the backend RADIUS server.

Check Interval

The interval to wait between checks to see if the RADIUS server is up. This is only applicable if the Server-Status request or Access request methods are used.

Number of Answers to Alive

The number of times the RADIUS server must respond before it is marked as alive. This is only applicable if the Server-Status request or Access request methods are used.

Revive Interval

If Server-Status requests and Access requests are not allowed or supported by the RADIUS server, then ExtremeCloud IQ - Site Engine waits the amount of time specified here before allowing requests to go to a backend RADIUS server, if it stops responding. Only use this if there is no other way to detect the health of the backend RADIUS server.

Policy Mapping Configuration

In your ExtremeControl profiles, each access policy (Accept, Quarantine, Failsafe, and Assessment) is associated to a *policy mapping* that defines exactly how end-system traffic is handled on the network. Each mapping specifies a policy role (created in the **Policy** tab) and/or any additional RADIUS attributes included as part of a RADIUS response to a switch.

The RADIUS attributes required by a switch are specified in the Gateway RADIUS Attributes to Send field configured in the Edit Switch window. The actual switch RADIUS attribute values (Login-LAT-Port, Custom 1, etc.) are defined within each policy mapping configured in this window. Each policy mapping is associated with the access policy selected in your ExtremeControl profiles.

When an end-system authenticates to the network, the ExtremeControl profile is applied and the appropriate RADIUS response attributes are extracted from the mapping based on the switch the authentication request originated from. The attributes are returned to the switch in the RADIUS Access-Accept response.


For more information on configuring policy mappings, see How to Set Up Access Policies and Policy Mappings. For a description of each ExtremeControl access policy, and some guidelines for creating corresponding policy roles in the **Policy** tab, see the section on Access Policies in the Concepts file.

To access this window, select the **Policy Mappings** left-panel option in the **ExtremeControl Configurations** > **Access Control** left-panel menu.


The columns displayed in this window vary depending on whether you are using a Basic or Advanced policy mapping configuration. For a definition of each column, see below.

Basic AAA Configuration

Basic AAA Configurations define the RADIUS and LDAP configurations for all end-systems connecting to your ExtremeControl engines.



Advanced Policy Mapping Configuration



Column Definitions

Name

The policy mapping name.

Policy Role

The policy role assigned to this mapping. All policy roles used in your mappings must be part of your ExtremeControl (ExtremeControl) Controller policy configuration and/or defined in the **Policy** tab and enforced to the policy-enabled switches in your network.

Location

Policy mapping locations permit authentication requests that match the same ExtremeControl rule and corresponding ExtremeControl profile to be authorized to different accept attributes (policy/VLAN/Custom Attribute) based on the location the request originated from. For example, in the Policy Mapping Configuration screenshot above, the Administration policy mapping has five entries, with each entry assigning a different VLAN (for RFC 3580-enabled switches) for authentication requests matching the specified location. Requests originating from the 1st floor South location will be authorized to VLAN 100, and requests originating from the 2nd floor North location (matching the same ExtremeControl rule) is authorized to VLAN 220. Using locations in this manner lets you authorize end-systems to different access criteria using a single ExtremeControl rule, whereas the alternative would be to create multiple location-based ExtremeControl rules each with an ExtremeControl Profile that corresponds with the desired access value.

When policy mapping locations are used in this manner, it is important to include a catch-all policy mapping (the fifth Administration mapping in the example above) that has a location of "any" and sets the access behavior for an authorization originating from any other location. The access behavior could be a policy/VLAN/Custom Attribute that grants some form of restricted access, or denies access altogether. If a catch-all mapping is not included, a warning message appears on enforce indicating that there is no catch-all mapping configured, and authorizations that match the policy but do not originate from a defined location, can result in errors or unpredictable behavior.

VLAN Name

If you have RFC 3580-enabled switches in your network, this column displays the VLAN name assigned to this mapping.

VLAN Egress

If you have RFC 3580-enabled switches in your network, this column displays the VLAN ID assigned to this mapping.

Filter

This value is only displayed in Basic mode if ExtremeWireless Controllers have been added to ExtremeCloud IQ - Site Engine. The Filter column typically maps to the Filter-Id RADIUS attribute. This value applies to ExtremeWireless Controllers and other switches that support the Filter-Id attribute.

Login-LAT-Group

If your network devices require a Login-LAT-Group, it displays here.

Login-LAT-Port

If you have ExtremeWireless Controllers on your network, the Login-LAT-Port is an attribute returned in the default RADIUS response. The Login-LAT-Port value is used by the controller to determine whether the authentication is fully authorized. A value of "1" indicates the authentication is authorized, where a value of "0" indicates that authorization is not complete. The value of "0" is used by the controller to determine that additional authentication is required and is a signal for the controller to engage its external captive portal and use HTTP redirection to force HTTP traffic from the end-system to the defined ExtremeControl engine. This is used in conjunction with the Registration and Assessment features of ExtremeControl.

Management

The authorization attribute returned for successful administrative access authentication requests that originate from network equipment configured to use RADIUS as the authentication mechanism for remote management of switches, routers, VPN concentrators, etc. Examples of management values for EOS devices are: "mgmt=su:", "mgmt=rw:", or "mgmt=ro:". The management attribute determines the level of access the administrator will have when authorized to access the device: superuser, read/write, or read-only.

Custom

Some network devices require additional RADIUS response attributes in order to provide authorization or define additional parameters for the authenticated session. These additional attributes can be defined in the five available Custom option fields.

Attribute List 1-3


The **Attribute List** fields display additional RADIUS response attributes in a single mapping. For example, you can use each field to provide a complete ACL for a different third-party vendor.

Add/Edit Policy Mapping

Use this window to add a new policy mapping or edit an existing policy mapping. A policy mapping specifies a policy role (created on the **Policy** tab) and/or any additional RADIUS attributes included as part of a RADIUS response to a switch (as defined in the Gateway RADIUS Attributes to Send field configured in the Edit Switch window). For additional information about configuring policy mappings, see How to Set Up Access Policies and Policy Mappings.

Access this window by selecting the **Add** or **Edit** toolbar buttons in the Edit Policy Mapping Configuration window.

The fields in this window vary depending on whether you are using a basic or advanced policy mapping configuration. For a definition of each field, see below.



Name

Enter a name for the policy mapping.

Map to Location

Allows you to specify a certain location for the mapping. You should first configure your locations using the Location Group (**Control** tab > **ExtremeControl** > ExtremeControl Configurations > Group Editor > Location Groups) or you can select the **Edit** button to the right of the field to add a location group to the list. For more information on using the Location option in Policy Mappings, see the Edit Policy Mapping Configuration Window Help topic.

Policy Role

Use the drop-down list to select a policy role, or enter a policy role in the field. The drop-down list displays any policy roles you have created and saved in the **Policy** tab and/or all the policy roles contained in the ExtremeControl Controller policy configuration. Roles from all your policy domains are listed; if there are duplicate names, only one is listed. The list is not case sensitive, so "Enterprise User" and "enterprise user" are considered duplicate policy names. All policy roles used in your mappings must be part of your ExtremeControl) Controller policy configuration and/or defined in **Policy** tab and enforced to the EOS policy-enabled switches in your network.

NOTE: Entering a new policy role does **not** create a new role in the **Policy** tab.

VLAN [ID] Name

Use the drop-down list to select the appropriate VLAN associated with the policy. This list displays any VLANs defined in ExtremeCloud IQ - Site Engine. Select the configuration menu button to the right of the field to add a VLAN to the list. VLANs you add remain in the list only as long as they are used in a mapping and they are **not** added to the ExtremeCloud IQ - Site Engine database.

VLAN Egress

Use the drop-down list to select the appropriate VLAN the egress forwarding state: Tagged (frames are forwarded as tagged), Untagged (frames are forwarded as untagged), Same as Ingress (frames are forwarded as specified by the VLAN Ingress), or User Defined (you define how frames are forwarded).

Filter

If your network devices require a custom Filter-Id, enter it here. The Filter column typically maps to the Filter-Id RADIUS attribute. This value applies to ExtremeWireless Controllers and other switches that support the Filter-Id attribute.

Port Profile

For ExtremeXOS/Switch Engine devices on which legacy firmware is installed, this field indicates the profile used by Extreme Policy.

Login-LAT-Group

If your network devices require a Login-LAT-Group, enter it here.

Login-LAT-Port

If you have ExtremeWireless Controllers on your network, the Login-LAT-Port is an attribute returned in the default RADIUS response. The Login-LAT-Port value is used by the controller to determine whether the authentication is fully authorized. A value of "1" indicates the authentication is authorized, where a value of "0" indicates that authorization is not complete. The value of "0" is used by the controller to determine that additional authentication is required and is a signal for the controller to engage its external captive portal and use HTTP redirection to force HTTP traffic from the end-system to the

defined ExtremeControl engine. This is used in conjunction with the Registration and Assessment features of ExtremeControl.

Custom

If your network devices require additional RADIUS response attributes in order to provide authorization or define additional parameters for the authenticated session, you can define them in the five available Custom option fields.

Organization 1-3


Enter additional RADIUS response attributes in a single mapping in the **Organization** fields. For example, you can use each field to provide a complete ACL for a different third-party vendor.

Management

Enter a management attribute used to authenticate requests for administrative access to the selected switches, for example, "mgmt=su:", "mgmt=rw:", or "mgmt=ro:". The management attribute determines the level of access the administrator will have to the switch: superuser, read/write, or read-only. Be sure to include the final colon (":") in the attribute, or the management access will not work.

Add LDAP Policy Mappings

Use the Add LDAP Policy Mapping window to add LDAP Policy authentication mappings that define what policy will be assigned to an end-system, based on LDAP information.



The Add LDAP Policy Mapping window includes the following information:

Name

A unique name used to identify the LDAP Policy Mapping.

Attribute

The LDAP attribute for which the value to policy mappings are defined. This is the attribute that will be queried in the LDAP database to determine which policy to assign to a given end-system.

Value - Policy Table

Lists mappings between the LDAP database attribute value and authentication policies.

Use the buttons in the Add LDAP Policy Mappings window to perform the following functions:

Add

Opens the Add Attribute Value to Policy Mapping window, where you can define a new entry.

Edit


Opens the Edit Attribute Value to Policy Mapping window, where you can edit the selected entry.

Delete

Enables you to delete a selected entry.

Add Attribute Value to Policy Mapping

Use the Add Attribute Value to Policy Mapping window to edit the values and attributes to an end-system.



The Add Attribute Value to Policy Mapping window includes the following information:

Value

The specific value that the Attribute of the LDAP Policy Mapping must match in order to assign a given end-system a policy mapping.

Policy Mapping

The policy mapping, and by extension, the policy, which is assigned to an end-system that matches the Attribute-Value.


Lookup Attributes from Existing LDAP Configurations

Use this field to query an LDAP database of an existing LDAP configuration, which can help you determine what attributes and values to use for a given policy mapping. This field has no impact on the configuration; it is only meant to aid the user in the configuration.

Edit LDAP Policy Mappings

Edit LDAP Policy Mappings

Use the Edit LDAP Policy Mapping window to add or edit LDAP Policy authentication mappings that define what policy will be assigned to an end-system, based on LDAP information.



The Edit LDAP Policy Mapping window includes the following information:

Name

A unique name used to identify the LDAP Policy Mapping.

Attribute

The LDAP attribute for which the value to policy mappings are defined. This is the attribute that will be queried in the LDAP database to determine which policy to assign to a given end-system.

Value - Policy Table

Lists mappings between the LDAP database attribute value and authentication policies.

Use the buttons in the Edit LDAP Policy Mappings window to perform the following functions:

Add

Opens the Add Attribute Value to Policy Mapping window, where you can define a new entry.

Edit


Opens the Edit Attribute Value to Policy Mapping window, where you can edit the selected entry.

Delete

Enables you to delete a selected entry.

Edit Attribute Value to Policy Mapping Window

Use the Edit Attribute Value to Policy Mapping window to edit the values and attributes to an end-system.



The Edit Attribute Value to Policy Mapping window includes the following information:

Value

The specific value that the Attribute of the LDAP Policy Mapping must match in order to assign a given end-system a policy mapping.

Policy Mapping

The policy mapping, and by extension, the policy, which is assigned to an end-system that matches the Attribute-Value.

Lookup Attributes from Existing LDAP Configurations

Use this field to query an LDAP database of an existing LDAP configuration, which can help you determine what attributes and values to use for a given policy mapping. This field has no impact on the configuration; it is only meant to aid the user in the configuration.

Access Control Profiles

ExtremeCloud IQ - Site Engine includes ten system-defined ExtremeControl profiles that define the authorization and assessment requirements for the end-systems connecting to the network.


The system-defined profiles are:

- Administrator
- Allow
- Default
- Guest Access
- Notification
- Pass Through
- Quarantine
- Registration Denied Access

- Secure Guest Access
- Unregistered

Use the Access Control Profiles window to view and edit these profiles, and define new profiles if desired. Any changes made in this window are written immediately to the ExtremeCloud IQ - Site Engine database.

To open the Access Control Profiles window, navigate to the **Access Control** tab and select the ExtremeControl Profiles tab in the left-panel.



The window includes the following buttons and functionality:

Add Button Add...

Use this button to open the New ExtremeControl Profile window, where you can add an ExtremeControl profile.

Edit Button p Edit...

Use this button to open the Edit ExtremeControl Profile window, where you can edit an existing ExtremeControl profile.

Delete Button (a) Delete

Use this button to add an ExtremeControl profile.

The Access Control Profiles table includes the following columns:

Name

The name of the ExtremeControl profile.

Accept Policy

The Accept policy defined for this profile. An Accept policy is applied to an end-system when

- an end-system has been authorized locally by the ExtremeControl engine and has passed an assessment (if assessment in enabled).
- authentication is configured to replace the attributes returned from the RADIUS server with the Accept policy.

NOTES:

- If your Accept policy is "Use User/Host LDAP Policy Mappings," an Accept Policy will be assigned, based on the end-system information in the LDAP database and the <u>LDAP Policy Mappings</u> configured in the Authentication Mapping.
- Authenticated Guest and IoT Management provisioners cannot match a rule associated with an Accept Policy =

 No Policy --. Guest and IoT Management authenticated provisioners must match a rule in Control, mapped to an Accept Policy that is not mapped to "- No Policy --".

Reject Policy

Indicates whether all authentication requests are rejected.

Failsafe Policy

The Failsafe policy defined for this profile. A Failsafe policy is applied to an end-system if the end-system's IP address cannot be determined from its MAC address, or if there has been a scanning error and a scan of the end-system could not take place.

Assessment Configuration

The assessment configuration defined for this profile. The configuration define the assessment requirements for end-systems

Assessment Interval

If assessment is required, this defines the interval between required assessments for an end-system.

Quarantine Policy

The Quarantine policy defined for this profile. A Quarantine policy is applied to an end-system if the end-system fails an assessment.

Assessment Policy

The Assessment policy defined for this profile. An Assessment policy is applied to an end-system while it is being assessed.

Hide Assessment/Remediation Details

Denotes whether the option to hide assessment or remediation information on the Remediation Web Page has been selected.

New/Edit ExtremeControl Profile


ExtremeControl Profiles specify the authorization and assessment requirements for the end-systems connecting to the network. Profiles also specify the security policies that will be applied to end-systems for network authorization, depending on authentication and assessment results.

ExtremeCloud IQ - Site Engine comes with ten system-defined ExtremeControl profiles:

- Administrator
- Allow
- Default
- Guest Access
- Notification
- Pass Through
- Quarantine
- Registration Denied Access
- Secure Guest Access
- Unregistered

You can edit these profiles or you can define your own profiles to use for your ExtremeControl configurations. Use this window to create a new profile, or edit an existing profile. When you create a new profile, it is added to the Manage ExtremeControl Profiles window. When you edit a profile, it changes the profile wherever it is used, so you don't have to do individual edits for each profile.

To create a new profile, select the **Add** button in the Manage ExtremeControl Profiles window. To edit an existing profile, select a profile in the Manage ExtremeControl Profiles window and select the **Edit** button or select it from the left-panel.



Name

Enter a name for a new profile. If you are editing a profile, the name of the profile is displayed and cannot be edited. To change the name of a profile, right-click on the profile name in the ExtremeControl Profiles left-hand panel navigation tree and select **Rename** from the menu.

Reject Authentication Requests

If you select this checkbox, all authentication requests are rejected.

Authorization

Accept Policy

Use the drop-down list to select the Accept policy you want to use in this ExtremeControl profile. An Accept policy is applied to an end-system when:

- an end-system has been authorized locally (MAC authentication) by the ExtremeControl engine and has passed an assessment (if assessment in enabled).
- you have selected the **Replace RADIUS Attributes with Accept Policy** option.

If you select "No Policy," then the ExtremeControl engine does not include a Filter ID or VLAN Tunnel Attribute in the RADIUS attributes returned to the switch, and the default role configured on the port is assigned to the end-system. This option is necessary when configuring single user plus IP phone authentication supported on C2/C3 and B2/B3 devices.

If you select "Use User/Host LDAP Policy Mappings," an Accept Policy will be assigned, based on the end-system information in the LDAP database and the <u>LDAP Policy Mappings</u> configured in the <u>Authentication Mapping</u>.

Replace RADIUS Attributes with Accept Policy

When this option is checked, the attributes returned from the RADIUS server are replaced by the policy designated as the Accept policy. If the RADIUS server does not return a Filter ID or VLAN Tunnel attribute, the Accept policy is inserted. When this option is unchecked, the attributes returned from the RADIUS server are forwarded back "as is" and the Accept Policy would only be used to locally authorize

MAC authentication requests. If the RADIUS server does not return a Filter ID or VLAN Tunnel attribute, no attributes are returned to the switch.

Use Quarantine Policy

Select this checkbox if you want to specify a Quarantine policy. The Quarantine policy is used to restrict network access for end-systems that have failed the assessment. You must have the **Enable Assessment** checkbox selected to activate this checkbox.

If a Quarantine policy is not specified and you have configured RADIUS in your AAA configuration, then the policy from the RADIUS attributes would be applied (unless Replace RADIUS Attributes with Accept Policy has been selected, in which case the Accept policy would be used.) If Authorize Authentication Requests Locally has been selected in your AAA configuration, then the Accept policy would be applied to those end-systems that are authorized locally. This allows an end-system onto the network with its usual network access even though the end-system failed the assessment.

Use Failsafe Policy on Error

Select this checkbox if you want to specify a Failsafe policy to be applied to an end-system when it is in an Error connection state. An Error state results if the end-system's IP address could not be determined from its MAC address, or if there was a scanning error and a scan of the end-system could not take place. A Failsafe policy should allocate a nonrestrictive set of network resources to the connecting end-system so it can continue its work, even though an error occurred in ExtremeControl operation.

If a Failsafe policy is not specified and you have configured RADIUS in your AAA configuration, then the policy from the RADIUS attributes would be applied (unless Replace RADIUS Attributes with Accept Policy has been selected, in which case the Accept policy would be used.) If Authorize Authentication Requests Locally has been selected in your AAA configuration, then the Accept policy would be applied to those end-systems that are authorized locally. This allows end-systems onto the network with their usual network access when an error occurs in ExtremeControl operation.

Assessment

Enable Assessment

Select the **Enable Assessment** checkbox if you want to require that end-systems are scanned by an assessment server.

NOTES: If you require end-systems to be scanned by an assessment server, you need to configure the assessment servers performing the scans. The Manage Assessment Settings window is the main window used to manage and configure assessment servers. To access this window, select **Assessment** from the ExtremeControl Configurations > ExtremeControl Profiles left-hand panel navigation tree.

The ExtremeControl engine restarts when you enforce if **Enable Assessment** is selected the first time in an ExtremeControl profile. The ExtremeControl engine also restarts when you enforce when **Enable Assessment** is deselected for all ExtremeControl profiles.

Assessment Configuration

Use the drop-down list to select the assessment configuration you would like to use in this ExtremeControl Profile. Use the **Edit** button to add a new assessment configuration or edit a configuration, if needed. After you create an assessment configuration, it becomes available for selection in the list.

Assessment Interval

Enter an assessment interval that defines the interval between required assessments:

- Minutes 30 to 120
- Hours 1 to 48
- Days 1 to 31
- Weeks 1 to 52
- None

Hide Assessment Details and Remediation Options from User

If you select this option, the end user does not see assessment or remediation information on the Remediation Web Page. They are informed that they are quarantined, and told to contact the Help Desk for assistance.

Use Assessment Policy

Select this checkbox if you want to specify a certain policy to be applied to an end-system while it is being assessed. Use the drop-down list to select the desired policy.

Select when to apply the policy:

- During Initial Assessment Only Only initial assessments receive the assessment policy. If the end-system is being re-assessed, it remains in its current policy.
- During All Assessments All end-systems being assessed receive the specified assessment policy.

If an assessment policy is not specified and you have configured RADIUS in your AAA configuration, then the policy from the RADIUS attributes are applied (unless "Replace RADIUS Attributes with Accept Policy" is selected, in which case the Accept policy is used.) If "Authorize Authentication Requests Locally" is selected in your AAA configuration, then the Accept policy is applied to those end-systems authorized locally. This allows the end-system immediate network access without having to wait for assessment to be complete.


Edit Assessment Configuration

Use the Assessment Configuration window to view and configure the assessment configurations that define the assessment requirements for end-systems. Assessment configurations define the following information:

- How to score assessment results (determined by the selected Risk Level and Scoring Override configurations).
- What assessment tests to run (determined by the selected test sets).

After you have defined your assessment configurations, they are available for selection when creating your ExtremeControl configurations.

To access this window, select ExtremeControl Configurations > ExtremeControl Profiles > Assessment in the left-hand menu to open the Manage Assessment Settings window. Select an existing configuration and select Edit to open the Edit Assessment Configuration window, or you can select Add to add a new assessment configuration, and then open the Edit Assessment Configuration window.



Scoring Override Configuration

Use the drop-down list to select the scoring override configuration for this assessment configuration. Scoring overrides let you override the scoring mode and test result scores for a particular assessment test. The default scoring override configuration provided by ExtremeCloud IQ - Site Engine specifies no overrides, but can be edited to contain overrides, if desired.

Risk Level Configuration

Use the drop-down list to select the risk level configuration for this assessment configuration. The risk level configuration determines what risk level is assigned to an end-system (high, medium, or low) based on the end-system's health result details score.

Enable Assessment Warning Period

This section allows you to enable assessment warning periods. Warning periods let you specify a grace period and probation period used for assessment warnings.

Grace Period

Specify the number of days the end user has to resolve the warning issues before the endsystem is guarantined.

Probation Period

The number of days after an end user is quarantined that additional warnings results in immediate quarantine. This allows administrators to block repeat offenders by limiting their access to the network. When the probation period has passed, the end user can again receive assessment warnings. Setting the probation period to 0 is the same as having no probation period.

Test Sets

Select one or more test sets to run for this assessment configuration. Test sets define which type of assessment to launch against the end-system, what parameters to pass to the assessment server, and what assessment server resources to use.

For networks that use on-board agent-less assessment, you can <u>create a custom Saint scan</u> and add it to your agent-less test set configuration and use it for your end-system assessment.

If you select multiple agent-based test sets, the first test set you select is called the Controller test set. A Controller test set includes the Agent Configuration settings, the Advanced Settings, and all the specified test cases. Each subsequent agent-based test set that you select for the configuration is a "supporting" test set. For supporting test sets, only the "Application" test cases are used; all other configuration values are ignored. In the list of Test Sets, Controller test sets have a "(Master)" designation after them.

For example, you might want to use multiple agent-based test sets if you are managing multiple networks, and you have a unique agent-based test set for each network as well as secondary test sets for specific application tests that all the networks would use. In the assessment configuration for each network, select the unique test set as the Controller test set and then select any number of secondary test sets to be included in the configuration as well.

If the Controller test set is deselected, then a new controller is automatically selected. To specify a different test set as Controller, deselect all test sets, select the desired Controller test set first, and select the additional supporting test sets.

Name

The name of the test set.

Type

The type of assessment server used in the test set.

Assessment Resources

Specifies the network assessment servers that perform the assessments for the test set:

- Load Balance All The assessment load is balanced across all the servers of the specified type on the network.
- Use Assessment Server Pool As a more granular approach, you can specify an assessment server pool. For example, if you have four agent-less assessment servers, you can put server A and server B in server pool 1, and server C and server D in server pool 2. Then, you can specify which server pool the configuration should use.
- Use Onboard Assessment The onboard assessment server is used to perform the assessments.

Buttons

Use the configuration menu buttons to perform the following functions:

Used By

Opens a window that lists all assessment configurations currently using the selected test sets.

Add

Select to add a new test set.

Edit

Select to edit the selected test set.

Delete

Select to delete the selected test set(s).

Manage


Select to open the <u>Manage Assessment Servers</u> window, where you can view and define the assessment servers that are used in your assessment configurations.

- Manage Assessment Servers
- Manage Assessment Settings

Manage Assessment Servers

Use the Manage Assessment Servers window to view and configure the assessment servers that perform the end-system assessments in your network. After you have configured your assessment servers, they can be added to an assessment server pool and participate in assessment server load-balancing, if desired.

Agent-less Assessment Servers are automatically displayed in this window and cannot be edited or deleted. In order to enable your Agent-less Assessment Servers to participate in assessment server load-balancing and server-pools, you must add them manually to this window.



The following columns are included in the Manage Assessment Servers table:

Name

The name of the assessment server. This is the name that is entered when you add an assessment server. For on-board assessment servers, the name is determined by the name of the ExtremeControl engine. For example, if you create an ExtremeControl engine and name it MyExtremeControl engine, then the on-board assessment server name is listed as MyExtremeControl engine as well.

IP Address

The IP address of the assessment server. This is the IP address entered when you add an assessment server. For on-board assessment servers, the IP address is determined by the address of the ExtremeControl engine. For example, if you create an ExtremeControl engine with an IP address of 10.20.80.8, then the on-board assessment server IP address is listed as 10.20.80.8 as well.

Port

The port number on the assessment server to which the ExtremeControlengine sends assessment requests.

Type

The assessment server type: Agent-less, Nessus, or a FusionAssessmentAgent.

Poolable

A check mark in this column indicates that the assessment server can be part of an assessment server pool. If you have multiple assessment servers on your network, creating assessment server pools enables you to control which assessment server resources are used for each assessment configuration. External assessment servers are "poolable," however, in order to enable your agent-less on-board assessment servers to participate in server-pools, you must add them manually to this window.

Assessment Agent Adapter Version

The version of assessment agent adapter software that is installed on the assessment server.

Scanner Version

The version of scanner software installed on the assessment server. When an upgrade for the software is available, the upgrade icon displays. The Upgrade feature is only available for on-board agent-less assessment servers and enables you to upgrade the scanner software installed on the assessment server. When you select the row, the Upgrade button becomes active and you can select the button to initiate the upgrade.

Status

When the assessment server is operational, then the status is Normal. Otherwise, this column provides status information regarding an upgrade procedure: Downloading, Download failed, Updating..., Update complete, or Update failed.

Used By Button

Opens a window that lists the assessment server pools currently using the selected assessment servers.

Add Button

Opens the Add Assessment Server window, where you can define a new assessment server.

Edit Button


Opens the **Edit Assessment Server** window, where you can edit the settings for the selected assessment server. You cannot edit on-board assessment server settings.

Delete Button

Deletes the selected assessment server. You cannot delete on-board assessment servers or servers that are currently in use.

Check for Updates Button

This button opens the **Check for Assessment Updates**, which lists any assessment software updates available for download. The download operation downloads any updated software but does not perform the actual upgrade to the assessment server. The actual upgrade must be performed using the **Upgrade** button here in this window.



Upgrade Button

This feature is only available for Agent-less Assessment Servers. Use it to upgrade the scanner software installed on the assessment server. When an upgrade is available, the upgrade icon displays in the Scanner Version column. When you select the row, the Upgrade button becomes active and you can select the button to initiate the upgrade.

Upgrades are available through the Web Update feature accessed via Help > Check For Assessment Updates or by selecting the **Updates** button. This check downloads any updated

software, but does not perform the actual upgrade to the assessment server. The actual upgrade must be performed using the **Upgrade** button here in this window.

Perform the Check for Assessment Updates and the Upgrade operation at least every two weeks to ensure that the assessment servers are running the latest scanner software that includes the most up-to-date virus definitions. You can schedule the check for assessment updates using the Assessment Server Web Update option.

- Because the on-board Agent-less Assessment license is subscription-based, the Upgrade
 operation must be performed at least one time a month in order to upgrade the license. If the
 ExtremeCloud IQ Site Engine (ExtremeCloud IQ Site Engine) server is unable to contact
 the upgrade server, contact Extreme Networks Support so that a special license can be
 provided.
- If the ExtremeCloud IQ Site Engine Server does not have internet access (and cannot use the Web Update feature), you can perform an upgrade by copying the upgrade file to the ExtremeCloud IQ Site Engine Server install directory and extracting the file in the ExtremeCloud IQ Site Engine directory (it extracts the entire path from there).
- To perform the upgrade:
 - 1. Select the **Upgrade**
 - 2. Select http://www.extremenetworks.com/netsight-renew/netsight-saint/ to download the multi-file archive.
 - 3. Search for a filename using this naming convention saint_latest.zip.XXX. Use 7zip to unpack the multi-file archive before copying it to the install directory.

Refresh Button


Reloads the latest assessment server information in the table. You can also refresh just the version information by right-clicking on a row in the table and selecting **Refresh Version Info**.

NOTES:

Manage Assessment Settings

The Manage Assessment Settings panel is the main panel used to manage and configure the assessment servers performing the end-system assessments in your network. To access this window, select ExtremeControl Configurations > ExtremeControl Profiles > Assessment from the menu bar.

Assessment configurations define the different assessment requirements for end-systems connecting to your network. When you create an ExtremeControl profile, you select an assessment configuration that defines the assessment requirements for the end-systems using that profile. You can also select the **Used By** button to view a list of all assessment configurations currently being used by ExtremeControl configurations.



Name

The name of the assessment configuration. This is the name that is entered when you add an assessment configuration in the Edit Assessment Configuration window.

Scoring Override Config

The scoring override configuration for this assessment configuration. The scoring override configuration lets you override the default scoring assigned by the assessment server to a particular assessment test ID.

Risk Level Config

The risk level configuration for this assessment configuration. The risk level configuration determines what risk level is assigned to an end-system (high, medium, or low) based on the end-system's health result details score.

Test Sets

The test sets that runs for this assessment configuration. Test sets define which type of assessment to launch against the end-system, what parameters to pass to the assessment server, and what assessment server resources to use.

Create a Custom Scan for Agent-less Assessment

You can create a custom Saint scan for networks that use on-board agent-less assessment.

The custom scan feature is useful if you are already using Saint assessment and want to integrate existing custom scans into ExtremeControl. It also allows you to create a custom scan with assessment criteria that requires only a limited number of port scans and tests.

To create a custom scan, you must connect to the Saint web service and use the Saint web interface to configure the scan. After you have created the scan, you will be able to add it to your agent-less test set configuration and use it for your end-system assessment.

Use the following steps to create a custom scan:

- 1. Connect a monitor and keyboard to your ExtremeControl engine, or connect via SSH.
- 2. From the CLI, "cd" to the directory /opt/nac/saint/saint.

NOTE: On some ExtremeControl engines, the second Saint directory includes a version number. For example, /opt/nac/saint/saint-8.5.11.

3. Start the Saint web service by entering the following command line argument: ./custom_policy_editor.pl -r -h <ip> where <ip> is the IP address of the system that is going to connect to the Saint web service and configure the custom scan (for example, your laptop system).

NOTE: You cannot run custom_policy_editor.pl from any directory. You must "cd" to the directory /opt/nac/saint/saint.

- 4. During the web service start-up, you are asked to create login user names and passwords for two accounts: saint and admin. The accounts are disabled by default, but they become enabled when you provide a password for them. After you complete the start-up by providing the user names and passwords, you are ready to connect to the web service and configure your custom scan.
- 5. From the connecting system, connect to the Saint web service by entering the following URL in a web browser window: http://ip of Extreme Access Control engine>:1414
- 6. Login using the admin user name and password that you created during the web service startup. (The Welcome screen automatically displays the Saint username and password; you need to change it to the admin username and password.)
- 7. Select the Create option in the Custom Scan Level Selection screen after you have logged in.
- 8. Create a new scan by entering a name, choosing a template, and selecting the Add button.
- 9. Configure your custom scan by selecting the Vulnerability Checks, Port Scans, and other desired options in the Custom Scan Setup screen. Select **Save** at the bottom of the web page to save your scan. (You might need to scroll down to see this button).
- 10. The custom scan is created. Close your web browser window.
- 11. Enter the name of the scan in your agent-less test set in ExtremeControl:
 - a. From the Extreme Access Control engine command line, cd to the /opt/nac/saint/saint/config/policy directory to determine the name of the scan.


NOTE: On some ExtremeControl engines, the second Saint directory will include a version number. For example, /opt/nac/saint/saint-8.5.11/config/policy.

b. In the policy directory, there are two files that contain the name of the scan as you entered it in the Saint web interface. For example, if you named the scan "MyCustom," you'll see the following two files in the directory:

```
saint_data_MyCustom.probe
and
saint data MyCustom.conf.
```

In this example, the scan name that you enter into ExtremeControl is saint_data_MyCustom. You can rename the scan if desired, as long as you rename both the .probe and .conf files. If you rename the scan, enter the new name into ExtremeControl.

- c. Select ExtremeControl Configurations > ExtremeControl Profiles > Assessment in the left-hand menu to open the Manage Assessment Settings window.
- d. In the **Assessment Configurations** tab, select any configuration and select **Edit**. The <u>Edit</u> <u>Assessment Configuration</u> window opens. You can also select **Add** to add a new assessment configuration, and then open the Edit Assessment Configuration window.



e. The Test Sets section of the window includes a list of all the test sets available for your assessment configurations. Select the agent-less test set that will be configured to use the custom scan, select the test set you want to configure, and select **Edit**. (Select **Add Agent-less** if you need to create a new test set.)

- f. In the Scanning Level section of the Edit Agent-less Test Set window, select **Custom** from the drop-down list and enter the scan name as determined in step b. Select **OK**.
- g. The agent-less test set with the custom scan can now be used in your assessment configurations.
- How to Set Up Assessments
- Edit Assessment Configuration

Portal Configuration Overview

If your network is implementing <u>registration</u> or <u>assessment / remediation</u>, you define the branding and behavior of the portal website used by the end user during the registration or assessment/remediation process using a Portal Configuration. ExtremeCloud IQ - Site Engine allows you to create two types of portal configurations.

ExtremeControl engines ship with a default Portal Configuration. You can use this default configuration as is, or make changes to the default configuration using this window, if desired.

If you network is using an external captive portal service (for example, ExtremeGuest), use the <u>External</u> configuration type when creating a new portal configuration.

Accessing the Portal Configuration

Use the following steps to access the Portal Configuration:

- 1. Open the Control > Access Control tab.
- 2. In the left-panel, expand **Configuration**.
- 3. Expand Captive Portals.
- 4. Expand a Portal Configuration.

Default Portal Configuration

The following settings relate to the default type portal configuration:

Network Settings

Use this panel to configure common <u>network</u> web page settings that are shared by both the <u>Assessment / Remediation</u> and the Registration portal web pages.

Administration

Use this panel to configure settings for the <u>Registration Administration</u> web page and grant access to the page for administrators and sponsors.

The Registration Administration web page allows Helpdesk and IT administrators to track the status of registered end-systems, as well as add, modify, and delete registered end-systems on the network

Website Configuration

Use this tab to <u>configure</u> the common settings used by the different registration web pages, including selecting guest access, authentication settings, and whether assessment and remediation is supported.

Look and Feel

Use the <u>Look and Feel</u> panel to configure common web page settings shared by both the <u>Assessment / Remediation</u> and the Registration portal web pages.

Guest Access and Registration

<u>Guest Web Access</u> provides a way for you to inform guests that they are connecting to your network and lets you display an Acceptable Use Policy (AUP).

<u>Guest Registration</u> forces any new end-system connecting on the network to provide the user's identity in the registration web page before being allowed access to the network.

Secure Guest Access provides secure network access for wireless guests via 802.1x PEAP by sending a unique username, password, and access instructions for the secure SSID to guests via an email address or mobile phone (via SMS text). Secure Guest Access supports both preregistered guests and guests self-registering through the captive portal. No agent is required.

Authenticated Web Access

<u>Authenticated Web Access</u> provides a way to inform end users that they are connecting to your network and lets you display an Acceptable Use Policy. End users are required to authenticate to the network using the Authenticated Web Access login page. However, end users are only granted one-time network access for a single session, and no permanent end user registration records are stored. Authentication is required each time a user logs into the network, which can be particularly useful for shared computers located in labs and libraries.

Authenticated Registration

<u>Authenticated Registration</u> provides a way for existing corporate end users to access the network on end-systems that don't run 802.1X (such as Linux systems) by requiring them to authenticate to the network using the registration web page. After successful registration, the end-system is permitted access until the registration expires or is administratively revoked.

Assessment / Remediation

Use this panel to configure settings for the <u>Assessment / Remediation</u> portal web page.

External Captive Portal

Use this tab to <u>configure</u> an external captive portal, outside of ExtremeCloud IQ - Site Engine.

Captive Portal Configuration

If your network is implementing <u>registration</u> or <u>assessment/remediation</u>, you define the branding and behavior of the portal website used by the end user during the registration or assessment/remediation process using a Portal Configuration. ExtremeControl engines ship with a default Portal Configuration. You can use this default configuration as is, or make changes to the default configuration using this window, if desired.

ExtremeControlengines ship with a default Portal Configuration. You can use this default configuration or reconfigure it. If you network is using an external captive portal service (for example, ExtremeGuest), use the External Captive Portal configuration type when creating a new portal configuration.

This Help topic provides the following information for accessing and configuring the Portal Configuration:

- Accessing the Portal Configuration
- Default portal configuration
 - Network Settings
 - Administration
- Website Configuration
 - Look and Feel
 - Guest Web Access
 - Authenticated Web Access
 - Assessment/Remediation
- External captive portal service
 - External Captive Portal

Accessing the Portal Configuration

Use the following steps to access the Portal Configuration:

- 1. Open the **Control** > **Access Control** tab.
- 2. In the left-panel, expand Configuration.
- 3. Expand Captive Portals.
- 4. Expand a Portal Configuration.

Default portal configuration

Network Settings

Use this panel to configure common <u>network</u> web page settings that are shared by both the <u>Assessment / Remediation</u> and the Registration portal web pages.

Administration

Use this panel to configure the settings for the <u>Registration Administration</u> web page and grant access to the page for administrators and sponsors.


The Registration Administration web page allows Helpdesk and IT administrators to track the status of registered end-systems, as well as add, modify, and delete registered end-systems on the network.

Website Configuration

Use this tab to <u>configure</u> the common settings used by the different registration web pages, including selecting guest access, authentication settings, and whether assessment and remediation is supported.

Look and Feel

Use the Look and Feel panel to configure common web page settings shared by both the Assessment / Remediation and the Registration portal web pages.



Display Powered by Logo

Select this checkbox to display the Extreme Networks logo at the bottom of all of your portal web pages.

Header

Select the **Edit** button to open a window where you can configure the link for the header image displayed at the top of all portal web pages. By default, the header image is configured as the Extreme Networks logo acting as a link to the Extreme Networks website. Text entered in this window can be formatted in HTML.

Footer

Select the **Edit** button to open a window where you can configure the footer displayed at the bottom of all portal web pages. By default, the footer is configured with generalized information concerning an organization. Change the *example* text in this section to customize the footer to your own organization. Text entered in this window can be formatted in HTML.

Helpdesk Information

Select the **Edit** button to open a window where you can configure the Helpdesk contact information provided to end users in various scenarios during the assessment/remediation and registration process (e.g. an end-system exceeded the maximum number of remediation attempts). By default, this section is configured with generalized Helpdesk information, such as contact URL, email address, and phone number. Change the *example* text to customize the Helpdesk information for your own organization. Text entered in this window can be formatted in HTML. In addition, the entire contents of the Helpdesk Information section are stored in the variable "HELPDESK_INFO". By entering "HELPDESK_INFO" (without the quotation marks) in any section that accepts HTML in the Common Page Settings (or any other settings), all information configured in this section will be displayed in place of "HELPDESK_INFO".

Title

Select the **Edit** button to open a window where you can modify the text that appears in the title bar of the registration and web access page browser tabs. The default page title is "Enterprise Registration."

Welcome Message

Select the **Edit** button to open a window where you can modify the message displayed to users on the menu bar of any registration or web access page. The default welcome message is "Welcome to the Enterprise Network's Registration Center."

User Registration Success

Select the **Edit** button to open a window where you can edit the message displayed to the end user after successfully registering their end-system to the network.

Images

Using the dropdown menus, you can specify the image files used in the portal web pages. All image files used for Assessment/Remediation and Registration portal web pages must be defined in this list. The image files defined here are sent to the ExtremeControl engine along with the web page configuration. Use the **Add** button to select an image file to add to the list. You can select an image in the list and use the **Preview** button to preview the image.

When an image file is defined here, it is available for selection from the configuration drop-down lists (for example, when you configure the <u>Access Granted Image</u>), and may be referenced in the sections supporting HTML. Available drop-down lists include:

• Header Background Image

Select the background image displayed behind the header image at the top of all portal web pages. The drop-down list displays all the images defined in the <u>Images window</u> for your selection. To add a new image, select **Add** to open the Images window.

Header Image

Select the image displayed at the top of all portal web pages. The drop-down list displays all the images defined in the <u>Images window</u> for your selection. To add a new image, select **Add** to open the Images window.

Favorites Icon

Select the image displayed as the Favorites icon in the web browser tabs. The drop-down list displays all the images defined in the <u>Images window</u> for your selection. To add a new image, select **Add** to open the Images window.

· Access Granted Image

Select the image displayed when the end user is granted access to the network either based on compliance with the network security policy or upon successful registration to the network. The drop-down list displays all the images defined in the Images window for your selection. To add a new image, select Add to open the Images window.

Access Denied Image

Select the image you would like displayed when the end user has been denied access to the network. The drop-down selection list displays all the images defined in the Images window. To add a new image, select Manage Images to open the Images window.

Error Image

Select the image displayed when there is a communication error with the ExtremeCloud IQ - Site Engine Server. The drop-down list displays all the images defined in the Images window window for your selection. To add a new image, select Add to open the Images window.

Busy Image

Select the progress bar image displayed to the end user when the web page is busy processing a request. The drop-down list displays all the images defined in the Images window for your selection. To add a new image, select Add to open the Images window.


Colors

Select the **Background** or **Text** color box corresponding to each item listed to choose your colors for the portal web pages:

- Page Define the background color and the color of all primary text on the web pages.
- Header Background Color Define the background color displayed behind the header image.
- Menu Bar Define the background color and text color for the menu bar.
- Menu Bar Highlight Define the background color and text color used for the menu bar highlights in the Administration pages.
- Footer Define the background color and text color for the footer.

- Table Header Define the background color and text color for the table column headers in the Administrative web pages.
- In-Progress Define the background color and text color for task in-progress images.
- Hyperlink Define the color used for hyperlinks on the web pages.
- Hyperlink Highlight Define the color of a hyperlink when it is highlighted.
- Accent Define the color used for accents on various parts of the web pages.

Select **OK** to save the changes.



Style Sheets

Select the **Desktop** or **Mobile** buttons to open the Edit Style Sheet window where you can create a style sheet that adds to or overwrites the formatting styles for the portal, or mobile version of the portal web pages, respectively.

Locales

This field lists the locales (languages) presented as options to the user in the captive portal, in addition to the default locale.

You can also define the default locale (language), displayed to any captive portal user unless the client locale detected from their browser matches one of the defined supplemental locales. The list of available locales includes the current default locale and any supplemental defined locales.

Display Locale Selector

Select this checkbox if you want a locale (language) selector to display as a drop-down list in the menu bar on the captive portal welcome and login pages. This is useful for a shared machine where the users of the machine may speak different languages. (On the mobile captive portal, the selector is displayed as a list of links at the bottom of the welcome screen.)


Guest Web Access

Guest Web Access provides a way for you to inform guests that they are connecting to your network and lets you display an Acceptable Use Policy (AUP).

End users are initially redirected to the captive portal when they first connect to the network. After the user enters the required information on the Guest Web Access login page (typically, their name and email address), they are allowed access on the network according to the assessment and authorization defined in the Guest Access profile.

Guest web access provides a single session, and no permanent end user records are stored. This provides increased network security, and also allows you to minimize the number of registration records stored in the ExtremeCloud IQ - Site Engine database.

Implementing guest web access requires web redirection or DNS proxy.



Introduction Message

Select the **Edit** button to open a window where you can edit the introductory message displayed to end users when gaining web access as guests. It may include an introduction to the network and information stating that the end user is agreeing to the Acceptable Use Policy (AUP) for the network upon

registering their device. A link to the URL that contains the full terms and conditions of the network's AUP can be provided from this introductory message. Note that the URL for this link must be added as an Allowed URL <u>Allowed Web Sites</u> accessed from the <u>Network Settings</u>. By configuring the introductory message with this information, end users can be held accountable for their actions on the network in accordance with the terms and conditions set forth by the network's AUP. This message is shared by Guest Web Access and Guest Registration. Changing it for one access type also changes it for the other

Customize Fields

Select the **Open Editor** button to open <u>Manage Custom Fields</u> where you can manage the fields displayed in the Guest Web Access login page. These settings are shared by Guest Web Access, Guest Registration, and Secure Guest Access. Changing them for one access type also changes them for the others.

Redirection

There are four Redirection options that specify where the end user is redirected following successful access, when the end user is allowed on the network. The option selected here overrides the Redirection option specified on the <u>Network Settings</u>. This setting is shared by Guest Web Access, Guest Registration, and Secure Guest Access. Changing it for one access type also changes it for the others.

- Use Network Settings Redirection Use the Redirection option specified on the Network Settings.
- **Disabled** This option disables redirection. The end user stays on the same web page where they were accepted onto the network.
- To User's Requested URL This option redirects the end user to the web page they originally requested when they connected to the network.
- To URL This option lets you specify the URL for the web page where the end user will be redirected. This would most likely be the home page for the enterprise website, for example, "http://www.ExtremeNetworks.com."

Registration Settings

Verification Method

User verification requires that guest end users registering to the network enter a verification code that is sent to their email address or mobile phone (via SMS text) before gaining network access. This ensures that network administrators have at least one way to contact the end user. For more information and complete instructions, see How to Configure Verification for Guest Registration.

Select from the following **Verification Methods**:

- **Email** The end user must enter an email address in the Guest Web Access login page. The Email Address field must be set to **Required** in the Manage Custom Fields.
- SMS Gateway The end user must enter a mobile phone number in the Guest Web Access login page. The Phone Number field must be set to Required in Manage Custom Fields.

- SMS Gateway or Email The end user must enter a mobile phone number or email address in the Guest Web Access login page. The Phone Number and Email Address fields must be set to Visible in the Manage Custom Fields.
- SMS Text Message The end user must enter a mobile phone number in the Guest Web Access login page. The Phone Number field must be set to Required in the Manage Custom Fields.
- SMS Text or Email The end user must enter either a mobile phone number or email address in the Guest Web Access login page. The Phone Number and Email Address fields must be set to Visible in the Manage Custom Fields.

If you have selected the "SMS Text Message" or the "SMS Text or Email" Verification method: select the Service Providers Edit button (below the verification method) to configure the list of mobile service providers from which end users can select the Registration web page. This setting allows ExtremeControl to correctly format the email address to which to send an email. This email is then received by the service provider and converted to an SMS text which is sent the user. The default configuration provides lists of the major US cellular service providers.

NOTE: Not all cellular service providers provide a way to send SMS text messages via email.

If you have selected the "SMS Gateway" or "SMS Gateway or Email" method: enter the SMS Gateway Email address provided by the SMS Gateway provider.

For all methods: use the Message Strings Edit button (below the verification method) to open the Message Strings Editor and modify the registration verification messages displayed to the user during the verification process. For example, if you have selected Email, you need to modify the "registrationVerificationEmailSentFromAddress" message string to be the appropriate email address for your company.

For all methods: set the Verify Pin Characters and Verify Pin Length options to define the characteristics and length of the verification code that is sent to the guest end user. This setting is shared by Guest Registration and Guest Web Access. Changing it for one access type also changes it for the other.

Third-party guest registration

Select additional registration types to implement guest registrations using a third-party as a way to obtain end user information. The registration portal provides the end user with an option to log into a third-party account in order to complete the registration process.

NOTE: Guest OAuth (for example, Google, Yahoo) may not support native mobile browsers and display a "user agent" error. To access the network, use a standard browser application (e.g. Google Chrome).

For more information, see the appropriate topic:

How to Implement Facebook Registration

How to Implement Google Registration

How to Implement Microsoft Registration

How to Implement Yahoo Registration

How to Implement Salesforce Registration


How to Implement Generic Provider Registration

Authenticated Web Access

Authenticated web access provides a way to inform end users that they are connecting to your network and lets you display an Acceptable Use Policy.

End users are required to authenticate to the network using the Authenticated Web Access login page. However, end users are only granted one-time network access for a single session, and no permanent end user registration records are stored. Authentication is required each time a user logs into the network, which can be particularly useful for shared computers located in labs and libraries.

Implementing authenticated web access requires web redirection or DNS proxy.



These settings are shared by the Authenticated Web Access and Authenticated Registration access types. Changing them for one type also changes them for the other.

Login or Register Message

Select the **Edit** button to open a window where you can edit the message displayed to the end user when they are registering. By default, the message states that the end user is required to register before being allowed on the network.

Introduction Message

Select the **Edit** button to open a window where you can edit the introductory message displayed to the end user when they are registering. By default, the message states that the end user is agreeing to the terms and conditions in the Acceptable Use Policy.

Failed Authentication Message

Select the **Edit** button to open a window where you can edit the message displayed to the end user if the end user fails authentication. By default, this message advises the end user to contact their network administrator for assistance. Note that the default configuration of the message references the "HELPDESK_INFO" variable which represents the <u>Helpdesk Information</u> that is defined in the <u>Look and Feel Settings</u>.

Customize Fields (Shared)

Select the **Open Editor** button to open the <u>Manage Custom Fields</u> where you can manage the fields displayed in the Registration web page.

Authentication

These settings are shared by the Authenticated Web Access and Authenticated Registration access types. Changing them for one type also changes them for the other.

AAA Configuration

This section displays the name of the AAA configuration being used by the Access Control configuration and provides a link to open the AAA Configuration window where you can make changes to the AAA Configuration, if desired. If the portal configuration is shared between multiple ExtremeControl Configurations using different AAA configurations, the different AAA configurations are listed here (maximum of 3), allowing you to open the appropriate AAA configuration.

The section also displays the method(s) utilized for validating the credentials entered during registration (LDAP, RADIUS, and/or a Local Password Repository) as specified in the AAA configuration(s).

- Authentication to End-System Group Select the Change button to open the User Group to
 End-System Group Map window where you can map the LDAP/RADIUS/Local User Group to the
 appropriate end-system group to specify end user access levels. When an end-system group has
 been mapped to a user group, the icon for the end-system group changes to display a key
 indicating that it is no longer available for general use. You can use the Move Up/Move Down
 arrows to set the precedence order for the mappings, allowing you to change the authentication
 order that takes place during the user authenticated registration.
- Local Password Repository If you are using a local repository, authenticated end users are assigned to the Web Authenticated Users group. Select the **Default** button to open a window

where you can edit the Local Password Repository. Multiple links may be listed if there are different repositories associated with different AAA configurations.

Max Failed Logins

Select this checkbox to specify the maximum consecutive number of times an end user can attempt to authenticate on an end-system and fail. You can specify a lockout period that must elapse before the user can attempt to log in again on that end-system.

Redirection

These settings are shared by the Authenticated Web Access and Authenticated Registration access types. Changing them for one type also changes them for the other.

Redirection

There are four Redirection options that specify where the end user is redirected following successful registration, when the end user is allowed on the network. The option selected here overrides the Redirection option specified on the Network Settings.

- Use Network Settings Redirection Use the Redirection option specified on the Network Settings.
- **Disabled** This option disables redirection. The end user stays on the same web page where they were accepted onto the network.
- To User's Requested URL This option redirects the end user to the web page they originally requested when they connected to the network.
- To URL This option lets you specify the URL of the web page to which the end user is redirected. This is typically the home page for the enterprise website, for example, "http://www.ExtremeNetworks.com."

Web Access Settings

Enable Agent-Based Login

If this option is enabled, when the end user connects to the network with an agent installed, the login dialog is displayed in an agent window instead forcing the user to go to the captive portal via a web browser. This allows you to provide authenticated web access without having to set up the captive portal. Agent-based login is useful for shared access end-systems running an agent because it prompts for a login dialog and also provides a logout option. Login credentials are limited to username/password and an Acceptable Use Policy is not displayed.

You can customize the messages in the Agent Login window using the <u>Message Strings Editor</u> available in the <u>Look and Feel</u> settings. Use the **agentLoginMessage** string to change the message. Any changes you make in the Message Strings Editor override the internationalized messages used in the Agent Login window.

NOTE: If you configure both <u>guest registration</u> and authenticated registration for an area on your network, the end user is presented with a choice on the registration web page whether or not to authenticate.

Assessment/Remediation

Assessment/Remediation allows you to configure the settings for the portal web page.

keep

Web Page Settings

Title

Select the **Edit** button to open a window where you can modify the message displayed in the title bar of the Assessment/Remediation web pages. The default page title is "Enterprise Remediation."

Welcome Message

Select the **Edit** button to open a window where you can modify the message displayed in the banner at the top of the Assessment/Remediation web page. The default welcome message is "Welcome to the Enterprise Remediation Center."

Display Violations

Use the checkboxes to select the assessment violation information that displays to the end user:

- None No violations are displayed to the web page. This option might be used for an ExtremeControlengine that is serving web pages to guest users, when you do not want the guest users to attempt to remediate their end-system.
- **Description** Only the description is displayed for violations. This provides the end user with information concerning what violation was found, but no information concerning how it can be fixed. This configuration may be appropriate for scenarios where the user population of the network does not possess technical IT knowledge and is not expected to self-remediate. It provides the Helpdesk personnel with technical information about the violation when the end user places a call to the Helpdesk.
- Solution Only the solution is displayed for violations, allowing the end user to perform self-service remediation without knowing what the violation is. This configuration may be appropriate for scenarios where the user population on the network does not possess technical IT knowledge but is expected to self-remediate.
- Description and Solution Both the description and solution are displayed for violations. This
 provides the end user with information concerning what violation was found and how to fix it.
 Providing complete information concerning the violation gives the end user the best chance of
 self-remediation, however, the technical details of the violation may result in end user confusion.
 Therefore, this configuration may be appropriate for scenarios where the user population of the
 network possesses more technical IT knowledge.

Do Not Allow Rescan

Select this checkbox if you do not want the end-user to have the ability to initiate a rescan of their end-system when quarantined. When selected, the **Reattempt Network Access** button is removed from the Assessment/Remediation web page, and the user is not provided with any way to initiate a rescan ondemand for network access. The end user is forced to contact the Help Desk for assistance. You can edit the "Permanently Removed Message" which, by default, advises the end user to contact the Helpdesk to

obtain access to the network. Note that the default configuration of the "Permanently Removed Message" references the "HELPDESK_INFO" variable which represents the <u>Helpdesk Information</u> that is defined in Look and Feel.

Allow Blacklist Remediation

Select this checkbox if you want end users added to the blocked list to have the ability to remediate their problem and attempt to reconnect to the network. When selected, a "Reattempt Network Access" button is added to the blocked list web page, allowing end users to remove themselves from the blocked list and reauthenticate to the network.

Permanently Removed Message

Select the **Edit** button to open a window where you can modify the message displayed when users can no longer self-remediate and must contact the Help Desk for assistance. Note that the default message references the "HELPDESK_INFO" variable which represents the <u>Helpdesk Information</u> that is defined in <u>Look and Feel</u>.

Custom Agent Install Message

Select the **Edit** button to open a window where you can create a message containing additional agent install information to add to the default text on the Install Agent portal web page.

Access Denied Image

Select the image you want displayed when the end user is quarantined and denied access to the network. The drop-down list displays all the images defined in the Images window for your selection.

Image During Reattempt

Select the image you want displayed when the end-user is reattempting network access after they repair their system. The drop-down list displays all the images defined in the <u>Images window</u> for your selection.

Agent Scan in Progress Image

Select the progress bar image you want displayed while the end-user is being scanned. The drop-down list displays all the images defined in the Images window for your selection.

Redirection

There are four Redirection options that specify where the end-user is redirected following successful remediation, when the end-user is allowed on the network. The option selected here overrides the Redirection option specified in the Network Settings for Remediation only.

- Use Network Settings Redirection Use the Redirection option specified in the Network Settings.
- **Disabled** This option disables redirection. The end-user stays on the same web page where they were accepted onto the network.
- To User's Requested URL This option redirects the end user to the web page they originally requested when they connected to the network.
- To URL This option lets you specify the URL of the web page to which the end-user is redirected. This is typically the home page for the enterprise website, for example, "http://www.ExtremeNetworks.com."

Remediation Attempt Limits

Limit Remediation Attempts

Select this checkbox to limit the maximum number of times an end-user is allowed to initiate a rescan of their end-system after initially being quarantined, in an attempt to remediate their violations. If selected, enter the number of attempts allowed.

Limit Time for Remediation

Select this checkbox to limit the total interval of time an end user is allowed to initiate a rescan of their end-system after initially being quarantined, in an attempt to remediate their violations. If selected, enter the amount of time in minutes.

Remediation Links

This table lists the links displayed on the Assessment/Remediation web page for the end users to use to remediate their end-system violations. There are two default remediation links: Microsoft Support and MAC OS Support. Use this tab to add additional links such as an internal website for patches. Links must contain a valid protocol prefix (http://, https://, ftp://).

Select **Add** to open a window where you can define a new link's name and URL. Select a link and select **Edit** to edit the link's information. Select **Delete** to remove a URL from the table.

Custom Remediation Actions

Use this table to create your own custom remediation action for a particular violation to use in place of the remediation action provided by the assessment server.

Use the following steps to add a custom remediation action:

- 1. Select the **Add** button to open the Add Custom Remediation Action window.
- 2. Enter the Test Case ID for the particular violation being remediated by the custom action. Test Case ID is found in the Health Results Details subtab in the End-Systems tab.
- 3. Add a custom description of the violation (required) and an optional custom solution.
- 4. If you have multiple portal configurations and you want to use this custom remediation action in all of your configurations, select the **Add to All Portal Configurations** option. This option overwrites any existing custom actions defined for the test case ID.
- 5. Select **OK**. Whenever the test case ID is listed as a violation on the web page, the custom violation description and solution you define is displayed instead of the remediation actions provided by the assessment server.

Select the **Define Default Custom Action** checkbox to advise end-users to contact the Helpdesk regarding additional security violations not explicitly listed with custom remediation actions. If this checkbox is selected, only the violations and associated custom remediation actions listed in the table would be presented to the user, along with a message advising them to contact the

Helpdesk for any other security violations not explicitly configured with a custom remediation action. Select the **Edit** button to edit this message.

To copy a custom action to another portal configuration, select the action in the table and select the **Copy To** button. A window opens where you can select the portal configurations where you want to copy the action, and whether you want it to overwrite any existing custom remediation actions already defined for that test case ID.

Portal Web Page URLs

The following table provides a list of URLs for accessing commonly used portal web pages. You can also access these web pages using the **Engine Portal Pages** button at the bottom of the Portal Configuration window.

Web Page	URL
Preview Web Page Allows you to preview the web pages that may be accessed by the end user during the assessment/remediation and registration process.	https:// ExtremeControl enginelP/screen_ preview
Registration Administration Page Lets administrators view registered devices and users, and manually add, delete, and modify users.	https:// ExtremeControl engine IP/administration
Registration Sponsor Page Lets sponsors view registered devices and users, and manually add, delete, and modify users.	https:// ExtremeControl engineIP/sponsor
Pre-Registration Page The pre-registration web page lets selected personnel easily register guest users in advance of an event, and print out a registration voucher that provides the guest user with their appropriate registration credentials.	https:// ExtremeControl engineIP/pre_ registration
Self-Registration Page Allows an authenticated and registered user to self-register additional devices that may not have a web browser (for example, game systems).	https:// ExtremeControl engineIP/self_ registration

External captive portal service

Use the External Captive Portal if you are using an external captive portal service, for example, ExtremeGuest.

External Captive Portal

Use this panel to configure an external captive portal service (for example, ExtremeGuest) based on parameters in ExtremeCloud IQ - Site Engine and your ExtremeControl engine.

To use this configuration, redirect all traffic to your ExtremeControl engine to http://nac_ip_address/redirect_with_info. All traffic that passes this address is redirected again to an external captive portal service of define.

If your service serves captive portals to unregistered users, modify their Policy profiles to allow traffic to this specific domain.

Base URL

Enter the URL or IP address of the service that processes attributes you send.

Shared Secret

Enter a **Shared Secret** to use ExtremeCloud IQ - Site Engine features with the external captive portal. The **Shared Secret** is a string of characters used to encrypt and decrypt communications between ExtremeCloud IQ - Site Engine and the ExtremeControl engine via a REST interface. This string entered in this field must match the shared secret entered for the ExtremeControl engine at the following site https://nac_ip_address/rest/method. The shared secret must be at least 6 characters long; 16 characters is recommended. Dashes are allowed in the string, but spaces are not.

Configuration

Select the **Configuration**, which is a set of attributes, that ExtremeControl sends to the external captive portal:

- ExtremeGuest Select this configuration if you are using ExtremeGuest for your external portal service. This configuration is system-defined and cannot be edited.
- **Custom** Select this configuration to specify the attributes ExtremeControl is sending to your external portal service yourself. When **Custom** is selected, the **Attributes to Send** section of the panel is available.

Attributes to Send

Selecting **Custom** in the **Configuration** drop-down list allows you to enter values in the Attributes to Send section of the window. Use this section to define the information sent to the external captive portal service

• Original URL — The URL requested by the end-system before being redirected to the ExtremeControl engine. Most Extreme switches encode this information automatically when redirection occurs.

NOTE: The redirecting agent can use the "X-Forwarded-Host" header to inject this attribute.

- Client IP The IP address of the redirected end-system.
- Client MAC The MAC address of the redirected end-system. This is only included if the endsystem is present in ExtremeCloud IQ - Site Engine.
- Captive Portal Name The name of the captive portal configured on the ExtremeControl engine.

- NAC IP The IP address of the ExtremeControl engine.
- NAC Group The ExtremeControl Group to which the ExtremeControl engine belongs.
- Switch IP The IP address of the device to which the end-system connected. This is only included if the end-system is present in ExtremeCloud IQ Site Engine.
- Switch Port The port of the device to which the end-system connected. This is only included if the end-system is present in ExtremeCloud IQ Site Engine.
- Switch Port ID The Port ID of the device to which the end-system connected. This is only included if the end-system is present in ExtremeCloud IQ Site Engine.

Preview


Displays a preview of the redirection URL for the external captive portal service to which the endsystem is redirected.

Portal Registration Administration

The Registration Administration web page allows Helpdesk and IT administrators to track the status of registered end-systems, as well as add, modify, and delete registered end-systems on the network.

Administration

Use this panel to configure settings for the Registration Administration web page and grant access to the page for administrators and sponsors.



Administration Web Page Settings

Welcome Message

Select the **Edit** button to open a window where you can modify the message displayed to users when they log into the administration or sponsor portal. The default welcome message is *Registration System Administration*.

Force Administration HTTPS

Select this checkbox to force the administration web page to be served securely over HTTPS (instead of HTTP) to administrators and sponsors on the network. It is recommended this is enabled for additional security.

Session Timeout (Minutes)

This field specifies the length of time an administrator can be inactive on the administration web page before automatically being logged out. The default value is 10 minutes.

Login Failure Image

Select an image to display when the end user fails to correctly log in to the web page. The drop-down selection menu displays all the images defined in the Images window for your selection. To add a new image, access the Look & Feel panel.


Limit Sponsor's View to Own Users

Select this checkbox if you want to limit a sponsor's view to only the users they have sponsored. This option is valid only if you configure LDAP or RADIUS authentication of your sponsors. If you select this checkbox, you must enter the LDAP Email Address Attribute Name or RADIUS Email Address Attribute Name so a sponsor's login name can be matched to their email address, and only the registered users for that sponsor are displayed.

Manage Notifications

Use the **Notifications** tab to review all the notifications you create, and to add, edit, and test specific notification rules. Notifications enable you to create alert actions performed when specific events or triggers take place in ExtremeCloud IQ - Site Engine. Notification actions include sending an email, creating a syslog entry, sending an SNMP trap, and launching a custom program or script.

To access this window, expand **Access Control> Configuration** in the left-panel and select **Notifications**.



Notifications Table Buttons

Use these buttons to add, edit, delete, or test a notification.

Add

Select to open the Add Notification window, where you can define a new notification rule.

Edit

Select to open the Edit Notification window, where you can edit notification rule actions for selected notification(s).

Delete

Select to delete notification(s) you select in the table.

Configuration

Use the configuration menu button to <u>create</u> default SIEM Notifications or <u>change</u> the default SIEM server:

Create Default SIEM Notifications - Creates five default notifications that enable the notification feature to integrate with SIEM (Security Information and Event Manager) by sending syslog messages to your SIEM server. The notifications are based on the following conditions and triggers:

- Any Registration event
- Any Health Result
- End-System events:
 - End-system added
 - End-system moved
 - End-system state changed

The generated syslog messages include the following information:

- IP address
- MAC address
- Username
- Switch IP address
- Switch port
- Hostname
- Operating system
- State
- Extended State
- Reason
- NAC Appliance

Change Default SIEM Server - Use this option to change the default SIEM server IP address used when you generate new default SIEM notifications. The specified default SIEM server only applies to newly generated notifications; manually edit previously generated notifications to change the server.

Notifications Table

The following columns are included in the Notifications Table:

Enabled

The checkbox indicates whether the notification is enabled. When a notification is enabled, the defined action takes place when the trigger occurs and the conditions are met.

Name

The name of the notification.

Type

The notification type defines the source of the event triggering the notification: End-System Group, End-System, User Group, Health Result, or Registration.

Trigger

The trigger determines when a notification action occurs, based on filtering for a specific event.

Action

The actions that take place when a notification is triggered.

Override Content

Specifies whether Override Content is enabled or disabled for the notification.

Notes

A short description of the notification rule. This description is created when a new notification is added.

Enable Default Notifications

ExtremeControl includes four default notifications you can enable and edit. To enable a default notification, perform the following steps:

- 1. Select the notification in the table and select the Edit button to open the Edit Notification window.
- 2. Use the **Edit Email Lists** button and change the default address to an address specific to your network. Default notifications are configured to send an email to this address.
- 3. Configure the **SMTP E-Mail Server** option in the SMTP Email Options to identify the SMTP email server used for outgoing messages generated by the Notification feature.
- 4. Select the **Enable Notification** check box and then select **OK** in the Edit Notification Action window. The default notification is now enabled in the Manage Notifications window.


The following examples show how notifications can be used to alert you of changes or events in your network:

- Send an email to the Helpdesk when an end-system changes location, for example if it moves from a wired connection in a building to a wireless connection outside.
- Send a trap if an end-system fails registration.
- Send a syslog message if an end-system reports a high risk assessment result.
- Send an email if an end-system that is reported as a stolen laptop authenticates on the network.
- Send an email if someone logs into the network after normal work hours.

- Send an email when an end-system is added or removed from an end-system group, such as the blocked list end-system group or other defined end-system group.
- Send an email when a user is added or removed from a user group, such as an Administrator or Help Desk user group.

Add/Edit Notification

The Add/Edit Notification window lets you edit an existing notification or create a new one. In the window, you can enable or disable the notification, specify the notification type and trigger, define the required conditions, and configure the actions that occur when the notification is activated. At the bottom of the window, provide a summary description of the notification's properties.



To create a new notification, select the **Add** button on the **Notifications** tab. To edit a notification, select a notification on the **Notifications** tab and select the **Edit** button.

Enable

Select the checkbox to enable the notification. When a notification is enabled, then the defined action takes place when the trigger occurs and the conditions are met.

Name

Enter a name for the notification.

Notes

Enter notes for the notification that describe the notification action or other notification details. This information is displayed on the **Notifications** tab.

Type

The notification type defines the source of the event that activates the notification. Use the drop-down list to select one of the following notification types:

- End-System
- Captive Portal Registration
- Guest and IoT Manager Provisioning
- End-System Group
- User Group
- · Health Result

Trigger

Triggers allow you to determine when a notification action occurs based on filtering for a specific event. Use the drop-down list to select the event for which you want to filter. The list of triggers changes according to the notification type you have selected. Selecting "Any" or "Any Change" means that no filtering occurs.

- End-System the actions are performed based on:
 - o an end-system being added, deleted, or moved
 - o an end-system state or a state change
 - o an authentication type or device type change
 - o a custom field change
 - whether the end-system is registered
 - an end-system IP address change. An event is generated when an end-system is added with a static IP, the end-system IP changes after IP resolution, or the end-system IP changes due to DHCP rediscover.
 - when an end-system is added to a MAC-based end-system group. Note that a notification
 is not generated if the end-system is already a member of three end-system groups and is
 added to an additional group, unless the option "Remove from Current Group
 Assignments" is enabled when the end-system is added to the group.
 - certain errors occurring
- Captive Portal Registration the actions occur when a registered user or device is added, removed, or updated.
- Guest and IoT Manager Provisioning the actions occur when a user or device is added, removed, or updated via Guest and IoT Manager.

- End-System Group the actions are performed when entries in the group are added or removed. "Any Change" would include added, removed, and modified.
- User Group the actions occur when entries in the group are added or removed. "Any Change" would include added, removed, and modified.
- Health Result the actions occur based on the risk level of a health result.

Conditions

This section lets you define additional conditions that, in addition to the trigger, determines when actions occur. Conditions can be used to limit the scope of events that trigger a notification action. The list of conditions changes according to the notification type you have selected.

Access Control Engines

Filter end-system notifications based on the engines you select here. Only end-systems being managed by the selected engines trigger the notification actions.

Profile

End-System events are filtered based on the ExtremeControl profile assigned to the end-system. Use the drop-down list to select the desired profile.

Device Type Group

Specify a device type group to use as a filter for the End-System, Health Result, and Registration notification types. When the end-system's device type matches the device type group, then the notification actions are performed.

End-System Group

Select an end-system group to use as a filter for the End-System Group notification type. When the end-system is a member of this end-system group, then the notification actions are performed. If you don't select this checkbox and specify a group, then the notification is sent if any end-system group is matched.

Location Group

Specify a location group to use as a filter for the End-System, Health Result, and Registration notification types. When the location where the end-system (the source of the event) connects to the network matches the location group, then the notification actions are performed.

Time Group

Specify a time group to use as a filter for the End-System, Health Result, and Registration notification types. When the day and time that the end-system (the source of the event) connects to the network matches the time group, then the notification actions are performed.

User Group

Select a user group to use as a filter for the User Group notification type. When the end-system is a member of this user group, then the notification actions are performed. If you don't select this checkbox and specify a group, then the notification is sent if any user group is matched.

Guest and IoT Manager Domain

Select the GIM Domain or Domains in which the Trigger must occur for the Actions to be invoked.

Guest and IoT Manager Onboarding Templates

After you select a **Guest and IoT Domain**, select the GIM Onboarding Template or Templates to which the Provisioner performing the event defined in the **Trigger** must be assigned for the **Actions** to be invoked.

Actions

Use the checkboxes to specify the actions you want to take place when a notification is triggered and the conditions are met. You can test a notification by selecting the **Test** button. (A notification must be saved before it can be tested.)

If an action depends on details from the triggered notification, the **Test** button triggers the notification, but the action might not complete successfully.

For example, if the action is to execute a Script or Workflow, selecting the Test button will not successfully complete the action if the script or workflow is using variables from the notification itself because the notification does not contain the details of the variables.

Email

Select this checkbox if you want an email sent when the notification is triggered. Use the drop-down list to select one of your pre-defined email lists. If no lists have been defined, the menu is empty and you can select the **Edit Email Lists** button to define a list.

Syslog to Server(s)

Select this checkbox if you want to create a syslog message when the notification is triggered. Enter the IP address or hostname for each syslog server where the message is sent. Multiple syslog servers can be listed, separated by either a comma or a space.

Trap Server

Select this checkbox if you want to send an SNMP trap when the notification is triggered. Enter the IP address for a trap receiver where the trap is sent. Valid trap receivers are systems running an SNMP Trap Service. From the Credential drop-down list, select the appropriate SNMP credential used when sending the trap to the trap receiver. Credentials are defined in the **Profiles/Credentials** tab in the Authorization/Device Access window (Tools > Authorization/Device Access).

Execute Program

Select this checkbox to specify a custom program or script run on the ExtremeCloud IQ - Site Engine Server when the notification is triggered. In the **Workflow** field, select the workflow from the drop-down list. Select the **Test** button to run the workflow.

Access Control Events Workflow

Select this checkbox if you want an Access Control event workflow run when the notification is triggered. To configure Access Control event workflows, create a workflow on the **Workflows** tab and select **Access Control Events** in the **Menus** drop-down list on the **Menus** tab of the Workflow Details section.

Override Content

Select this checkbox if you want to override the default content contained in the action message. Use the **Edit Content** button to open the **Edit Action Overrides** window, where you can change the defaults

for this specific notification only. Additionally, select the **Show Keywords** button in the **Edit Action Overrides** window to view the **keywords** available for the overrides.

Result

This section summarizes the notification type, trigger, conditions, and specified actions.

MAC Locking

This tab displays the settings for locked MAC address. MAC Locking lets you lock a MAC address to a specific switch or port on a switch so that the end-system can only access the network from that port or switch. If the end-system tries to authenticate on a different switch/port, it is rejected or assigned a specific policy. You can add or edit MAC locks from the End-Systems tab.

NOTE: MAC Locking to a specific port on a switch is based on the port interface name (e.g. fe.5.1). If a switch board is moved to a different slot in a chassis, or if a stack reorders itself, this name changes and breaks the MAC Locking settings.

MAC Address

The locked MAC address.

Switch IP

The IP address of the switch on which the MAC address is locked.

Port

The port on the switch for which the MAC address is locked.

Lock to Switch and Port

Indicates whether the MAC address is locked to a specific port on the switch, and enter the port interface name.

Failed Action

The action ExtremeCloud IQ - Site Engine takes when this MAC address tries to authenticate on a different port and/or switch:

- Reject The authentication request is rejected.
- Use Policy Use the drop-down list to select the policy that you want applied. This policy must exist in the **Policy** tab and be enforced to the switches in your network.

MAC to IP Mappings

Use the MAC to IP Mappings tab to view MAC to IP address mappings for devices with statically assigned IP addresses. You can also import a file of MAC to IP mappings to the list.

The MAC to IP mappings are sent to the ExtremeControl engines in the configuration enforce. The ExtremeControl engines use this table to resolve IP addresses.

MAC Address

The MAC address mapped to the static IP address.

IP Address

The statically assigned IP address.

Description

A description of the mapping; for example, a description of the device with the statically assigned IP address.

Add Button

Opens the Add MAC to IP Mapping window where you can add a new mapping and description to the table.

Edit Button

Opens the Edit MAC to IP Mapping window where you can edit the IP address and description for a mapping.

Delete Button

Deletes the selected MAC to IP mapping.

Import Button

Use the **Import** button to import a file of MAC to IP mappings to the list. In the file, MAC to IP mappings must be listed in CSV format, with one mapping per line. For example:

02.0A.40.0B.01.44, 122.111.45.66, description of mapping

34.34.44.44.48, 122.111.45.48, description of mapping

MAC addresses can be delimited with colons (:), periods (.), or dashes (-), but they display in the table with colons. Lines starting with "#" or "//" are ignored.

Access Control Engine Settings

Engine settings provide advanced configuration options for ExtremeControl engines. ExtremeCloud IQ - Site Engine comes with a default engine settings configuration. If desired, you can edit these default settings or you can define your own settings to use for your ExtremeControl engines.

Launch the **Engine Settings** window by selecting the Control > Access Control tab, expanding the Engines left-panel menu, selecting an ExtremeControl engine, and selecting the **Engine Settings** button. The **Engine Settings** window contains the following tabs available for configuration:

- <u>Credentials</u>
- Network Settings
- <u>Device Type Detection</u>
- IP Address Resolution
- Hostname Resolution
- Username Resolution
- Reauthentication
- Miscellaneous
- Auditing


To access status and diagnostic information for an ExtremeControl engine, launch the ExtremeControl Engine administration web page by right-clicking on the ExtremeControl engine in the left-panel tree and selecting WebView. You can also access the administration web page using the following URL: https://<ExtremeControlEngineIP>:8444/Admin. The default user name and password for access to this web page is "admin/Extreme@pp." The username and password can be changed in the Web Service Credentials field on the

NOTE:

Credentials

Use this tab to configure various parameters for your network engines including switch configuration, web service credentials, and EAP-TLS configuration.

Credentials Tab in the Engine Setting window.



Switch Configuration

Enter the shared secret that switches uses when communicating with ExtremeControl engines.

Shared Secret

A string of alpha-numeric characters used to encrypt and decrypt communications between the switch and the ExtremeControl engine. The shared secret is shown as a string of asterisks. Select the **Eye** icon to reveal the **Shared Secret**.

RADIUS Timeout

The amount of time (in seconds) that a switch waits before re-sending a RADIUS request to the ExtremeControl engine. The default is 15 seconds and the maximum is 60 seconds. Note that the time specified should be long enough to allow the ExtremeControl engine to receive a response from the RADIUS server.

NOTE:

Although this option allows a maximum of 60 seconds, the actual maximum time allowed varies depending on the switch model. If a switch does not support the timeout value specified here, then the value is not set on the switch and an error message displays in the ExtremeControlengine log. Check your switch documentation to verify supported values

RADIUS Timeout Retry Count

The number of times the switch attempts to contact an ExtremeControl engine with a RADIUS request, when an attempted contact fails. The default setting is 3 retries, which means that the switch retries a timed-out request three times, making a total of four attempts to contact the engine.

Use Primary RADIUS Server for Redundancy in Single Engine Configuration

If your ExtremeControl deployment has only one ExtremeControl Gateway engine, this option allows you to configure redundancy by using the primary RADIUS server as a backup when configuring the switches. This option would not apply to ExtremeControl deployments using advanced AAA configurations with more than one set of RADIUS servers, or if you have configured primary and secondary ExtremeControl Gateways.

SNMP Timeout

The amount of time (in seconds) that ExtremeCloud IQ - Site Engine waits before re-trying to contact the ExtremeControl engine. The value for this setting must be between 1 and 60.

NOTE:

When SNMP requests are redirected through the server, all SNMP timeouts are extended by a factor of four (timeout X 4) to allow for the delays incurred by redirecting requests through the server.

Admin Web Page Credentials

ExtremeControl Engine Web Service Credentials

The credentials specified here provide access to the ExtremeControl engine administration web page and the web services interface between the ExtremeCloud IQ - Site Engine server and the ExtremeControl engine. ExtremeCloud IQ - Site Engine provides default credentials that can be changed, if desired. Changes to the credentials are propagated to the ExtremeControl engines on Enforce.

NOTE:

The Web service credentials are used to communicate between ExtremeCloud IQ - Site Engine and ExtremeControl Engine. If you use non-default credentials and add a new ExtremeControl Engine, then you must pre-provision the new ExtremeControl Engine with the new credentials manually by running /opt/nac/configWebCredentials <username> <password>.

Admin Web Page Authentication

By default, the ExtremeControl engine administration web page (https://<ExtremeControlEngineIP>:8444/Admin/) uses the above Web Service Credentials for authentication. However, you can configure the web page to use the AAA Configuration assigned to that engine for authentication as well. This allows you to use LDAP or RADIUS authentication for the web page.

There are three steps for setting up the web page to use LDAP or RADIUS authentication:

- 1. Verify that the ExtremeControl Configuration assigned to the engine has LDAP or RADIUS authentication configured in its AAA Configuration.
- 2. Create a local user account on the ExtremeControl engine that matches the user name of the user logging in. Use the useradd command on the ExtremeControl engine CLI to create the local user account.
- 3. Select the **Use ExtremeControl AAA Configuration for Admin Web Page authentication** option here on the Credentials tab. Select **OK**. Enforce the change to the engine.

The ExtremeControl engine begins using the AAA configuration for the administration web page authentication. Note that it may take the Linux operating system on the ExtremeControl engine up to two minutes to recognize that the new user is valid.


EAP-TLS Configuration

Server Private Key Passphrase

The Server Private Key Passphrase is used to encrypt the private key created during certificate request generation of server certificates for use by ExtremeControl engines during Local EAP-TLS Authentication. The passphrase must be identical for all ExtremeControl engines, and must be configured properly, or Local EAP-TLS Authentication does not operate successfully.

Network Settings

Use this tab to configure the following network services for the ExtremeControl engine: DNS, NTP, SSH, and SNMP.



Manage DNS Configuration

Select the Manage DNS Configuration checkbox and enter a list of search domains and DNS servers.

Search Domains

A list of search domains used by the ExtremeControl engine when doing lookups by hostname. When an attempt to resolve a hostname is made, these domain suffixes are appended to the hostname of the device. For example, if someone does a ping to server1, ExtremeControl appends the search domains in an attempt to resolve the name: server1.domain1 server1.domain2, and so on.

DNS Servers

A list of DNS servers the ExtremeControl engine sends DNS lookups to for name resolution. The list is used by both hostname resolution and by the DNS proxy. You can enter multiple servers for redundancy. Use the Up and Down arrows to list the servers in the order they should be used.

Manage NTP Configuration

NTP (Network Time Protocol) configuration is important for protocols such as SNMPv3 and RFC3576 which incorporate playback protection. In addition, having accurate time configured on the ExtremeControl engine is essential for event logging and troubleshooting. Select the Manage NTP Configuration checkbox, specify the appropriate time zone, and create a list of NTP servers.

Time Zone

Select the appropriate time zone. This allows ExtremeControl to manage all date/time settings.

NTP Servers

A list of NTP servers. You can enter multiple servers for redundancy. Use the Up and Down arrows to list the servers in the order they should be used.

Manage SSH Configuration

SSH configuration provides additional security features for the ExtremeControl engine. Select the Manage SSH Configuration checkbox and provide the following SSH information.

Port

The port field allows you to configure a custom port to be used when launching SSH to the engine. The standard default port number is 22.

Disable Remote root Access

Select this option to disable remote root access via SSH to the engine and force a user to first log in with a real user account and then su to root (or use sudo) to perform an action. When remote root access is allowed, there is no way to determine who is accessing the engine. With remote root access disabled, the /var/log/message file displays users who log in and su to root. The log messages looks like these two examples:

```
sshd[19735]: Accepted password for <username> from 10.20.30.40 port 36777 ssh2
```

```
su[19762]: + pts/2 <username>-root
```

Enabling this option does not disable root access via the console. Do not disable root access unless you have configured RADIUS authentication or this disables remote access to the ExtremeControl engine.

RADIUS Authentication

This option lets you specify a centralized RADIUS server to manage user login credentials for users that are authorized to log into the engine using SSH. Select a primary and backup RADIUS server to use, and use the table below to create a list of authorized RADIUS users.

Authorized Users Table

Use the toolbar buttons to create a list of users allowed to log in to the ExtremeControlengine using SSH. You can add Local and RADIUS users and grant the user Administrative privileges, if appropriate. A user that is granted administrative rights can run sudo commands and commands that only a root user would be able to run. For example, some commands that require administrative rights to run would be:

sudo nacctl restart

sudo reboot

sudo nacdb

If a user is not granted administrative rights, they can log in, view files, and run some commands such a ping and ls.

SNMP Configuration

The SNMP configuration section allows you to deploy SNMP credentials for the ExtremeControl engine. The credentials can include different read/write credentials, for example, the read credential can be "public" and the write credential can be "private". In addition, basic host traps can be enabled from the ExtremeControl engine. Select the Manage SNMP Configuration checkbox and provide the following SNMP information.

Profile

Use the drop-down list to select a device access profile (or multiple profiles) to use for the ExtremeControl engine.

Trap Mode

Set the trap mode.

Trap Community Name

Supply the trap community name.

System Contact


Allows you to specify contact information for the person maintaining the device. Additionally, enter a backslash "\" between contacts to create a device group in a tiered tree structure. For example, to move the device into a device group called "John's Devices" within a device group called "Quality Assurance Testing", enter Quality Assurance Testing\John's Devices in this field.

System Location

The physical location of the device. Additionally, enter a backslash "\" between locations to create a device group in a tiered tree structure. For example, to move the device into a device group called "London" within a device group called "Europe", enter **Europe\London** in this field.

Device Type Detection

The device type detection settings are advanced settings with complex requirements. Before editing these mappings, contact your Extreme Networks representative or Extreme Networks Support for information and assistance.



Device Type Detection

When the device type detection option is selected, ExtremeControl determines the end-system's device type using the selected detection methods below. Device type can be an operating system family, an operating system, or a hardware type, such as a printer or a smartphone. ExtremeControl uses the selected methods in the order configured in the <u>detection source precedence</u>. When this option is deselected, all device type detection functionality on the ExtremeControl engine is disabled.

Use Agent-Based Assessment

This option causes the ExtremeControl engine to query connected agents for the end-system device type. This is the most accurate method of device type detection.

Use Agent-less Assessment

This option allows the ExtremeControl engine to use the results of an agent-less scan to determine the end-system's device type.

Use DHCP Fingerprinting

This option enables passive device type detection by fingerprinting DHCP packets snooped from an

end-system. Select the **Edit** button to <u>change the mapping of the DHCP packet properties</u> to map to a different operating system or physical hardware type.

Utilize Hostname When Using DHCP Fingerprinting

This option allows the end-system hostname to be used to fine-tune device type detection results using DHCP fingerprinting. With certain device types, if DHCP fingerprinting does not result in a unique device type match, the hostname can be used as one possible tie-breaker. For example, with Apple iOS devices, the hostname can be a good indicator of the device type.

Use Captive Portal Data

This option allows the ExtremeControl engine to detect the end-system's device type by using the agent string returned from the end-system's browser. This is the least secure method for device type detection, since it can be faked by the end-system. However, this option should be enabled if you have configured agent-based assessment with the "Allow Agent Unreachable for Unsupported Operating Systems" option enabled, so that the operating system can be detected when the end-system gets the Remediation web page when it is quarantined.

Allow Device Type Overwrite When Device Type Family Changes


This option allows the device type to be changed by a lower precedence detection method, if the device type family has changed. This option is required if you are supporting dual boot systems and have configured agent-based assessment with the "Allow Agent Unreachable for Unsupported Operating Systems" option enabled. For example, let's say Microsoft Windows XP SP3 was detected by an agent running on a dual boot end-system. If the system is rebooted and switched to Red Hat Linux 4.4, and the end-system is quarantined for not running the agent, the device type detection using captive portal data (a lower precedence method) would yield the device type family of Linux instead of Windows. The device type would be updated and would now pass the unsupported operating system test, and be allowed onto the network.

Source Precedence Order

This list specifies the precedence for the source of information used to determine end-system device type, with the highest precedence listed first. Select an item in the list and use the Move Up and Move Down arrows to change its position in the list. Manual Set refers to device type information that has been hard-coded via ExtremeCloud IQ - Site Engine Web Services. Typically, Manual Set (Accurate) has the highest precedence because the exact device type is known and the remaining sources of detection aren't needed, while Manual Set (Best-Guess) has the lowest precedence because it is a best-guess of the device type and should be used only when the other detection methods cannot provide a device type.

IP Address Resolution

The IP Address Resolution tab is used to define how and when ExtremeControl resolves an end-system's MAC address to an IP address for the end-system. These parameters are applicable for ExtremeControl Gateways and L2 ExtremeControl Controllers, but not L3 ExtremeControl Controllers.



Resolve IP Address

Specify when an ExtremeControl engine resolves the IP address for end-systems:

- Always (Default) Resolve the IP address for every end-system that ExtremeControl sees.
- Only for Assessment Resolve the IP address for end-systems that need to be assessed (scanned).

IP Address Resolution Timeout

Enter the maximum time an ExtremeControl engine waits trying to resolve an IP address from an end-system's MAC address before giving up and returning the Error state (MAC to IP Resolution Timed Out) for that end-system.

Allowed Retries on Failure

The number of attempts made to resolve the IP address after the first attempt fails. The default setting is 2 retries, which means that ExtremeControl retries a timed-out request two times, making a total of three attempts to resolve the IP address. Enter the amount of delay time in seconds that ExtremeControl waits before retrying to resolve the IP address.

Delay Between Failures

Enter the amount of time an ExtremeControl engine waits after failing to resolve an IP address before attempting again.

DHCP Resolution Delay Time

The number of seconds an ExtremeControl engine should wait after learning about an end-system before attempting to resolve the end-system's IP address. This delay is used to allow the end-system to negotiate its DHCP IP address. If Port Link Control is enabled, this delay is used after the ExtremeControl engine links down/up the port to force the end-system to request a new IP address on the new VLAN.

NOTE:

If the delay time specified here is less than the amount of time the end-system needs to renew its IP address, then the ExtremeControlengine may resolve the end-system's IP address incorrectly. This is a problem when assessment is enabled and may cause the engine to scan the incorrect IP address. Be sure to take into account the amount of time required for an end-system to get a new IP address when setting the delay time value.

Use DHCP Request IPs

Specify when, if ever, an IP address learned from a DHCP request packet could be used when resolving an end-system's IP address. This option is applicable only for ExtremeControl Gateways, since an inline ExtremeControl Controller should always hear the DHCP response as well.

- Always Always consider the IP address learned from a DHCP request for an end-system's IP, after all more reliable methods have been exhausted.
- Never Never consider an IP address learned from a DHCP request when resolving an endsystem's IP address. In a situation where the ExtremeControl Gateway receives DHCP packets from both the client and server, the gateway uses this IP when these packets are received during the IP resolution process. With subsequent authentications for which there is no additional DHCP exchange, ExtremeControl uses the enabled resolution options to resolve the IP address but does not use any previously learned DHCP information to resolve the IP.
- For Non-VLAN Switches Only (Default) Only consider IP addresses learned from DHCP request
 packets when the NAS switch the end-system was authenticated for does not use VLANs for
 access control. The IP addresses from request packets in a VLAN environment is always
 incorrect, because as an end-system transitions through VLANs, it always requests the IP from
 the previous VLAN.

Rediscover IP on DHCP Request

When this option is selected, ExtremeControl re-runs IP resolution on an authenticated end-system if a DHCP request causes its IP address to change. In this instance, the ExtremeControl policy applies to the new IP address and removed from the old IP address, and assessment scans and port resolution are not performed.

Always Use Fully Trusted DCHP IPs

When this options is selected, the ExtremeControl engine runs a DHCP table lookup to see if DHCP IP address is fully trusted for the end-system. If the address is fully trusted in the table, ExtremeControl resolves the IP address for the end-system without attempting additional resolution processes. If the address is not fully trusted or not found, the ExtremeControl engine attempts to resolve the IP address as normal. When this option is not selected, there is no fast IP resolution using DHCP IP packets.

Use Agent-based Assessment IPs

Specify when, if ever, an IP address reported by a connected agent could be used when resolving an end-system's IP address. This process looks for the end-system's MAC address in the list of MAC addresses from known connected agents. If an agent is connected and heartbeats during the IP Resolution process, then ExtremeControl uses the IP address of that agent.

- Always Always consider the IP address reported by a connected agent for an end-system's IP, after all more reliable methods have been exhausted.
- Never (Default) Never consider an IP address reported by a connected agent when resolving an end-system's IP address.
- For Non-VLAN Switches Only Only consider IP addresses reported by a connected agent when the NAS switch the end-system was authenticated for does not use VLANs for access control.

Use RADIUS Accounting Packets (IPv4 address)

When this option is selected, if the ExtremeControl engine receives a RADIUS accounting packet with a Framed-IP-Address in it, the engine skips IP resolution and uses the IP address in the RADIUS accounting packet.

Use RADIUS Accounting Packets (IPv6 address)

When this option is selected (ExtremeControl engine enforce required), the Framed-IPv6-Address attributes learned from the RADIUS accounting packet are shown in the IPv6 Address column.

NOTE:

IPv6 address resolution by SNMP is not used if enabled, (even when the 'Enable IPv6 Addresses for End-Systems' is enabled in **Options> Access Control> Advanced**).

Clear Duplicate IPs on Switches

Select this option to have an ExtremeControl engine clear out duplicate entries in the node alias and ARP tables of the NAS switch the end-system was authenticated for, if duplicates are found while trying to resolve the IP address of an end-system. The ExtremeControl engine then tries to re-read the IP address from the table to find the most recent entry.

Re-Read Delay

Specify the amount of time in seconds that an ExtremeControl engine waits after clearing duplicate IPs on a switch or a router before re-reading the node alias or ARP tables.

NetBIOS IP Filtering

This option causes the ExtremeControl engine to make NetBIOS requests to a list of IP addresses, if multiple IP addresses are found when trying to resolve the IP address of an end-system. See <u>NetBIOS</u> <u>Timeout</u> and <u>NetBIOS</u> <u>Timeout</u> Retry Count on the Miscellaneous tab.

Enable Router IP Discovery

This option causes ExtremeControl to make requests to an end-system's gateway router ARP table to try to resolve the IP address for an end-system, if the ExtremeControl engine was unable to resolve the IP address by querying the NAS switch. The gateway router for an end-system can be discovered by the relay router field of a DHCP packet or by using the gateway router defined for an IP subnet for the VLAN an end-system is put into by ExtremeControl. See IP Subnets.

Clear Duplicate IPs on Routers

This option causes an ExtremeControl engine to clear out duplicate entries in the ARP tables of an end-system's gateway router, if duplicates are found while trying to resolve the IP address of an end-system. The ExtremeControl engine then tries to re-read the IP address from the table to find the most recent entry. See <u>Clear Duplicate IP Re-Read Delay</u>.

Default Router Profile

The profile used to make SNMP requests to the gateway router for an end-system, if one is not defined for a specific router's interface IP address as part of an IP subnet. Use the **Edit** button to open the Profiles/Credentials tab in the Authorization/Device Access window where you can define authentication credentials and create the profiles that use those credentials.

Default Router SNMP Context

The SNMP context used when making requests to the router, if the credentials used for the router are SNMP v3 and the specific router's interface IP address has not been defined as part of an IP subnet.

IP Subnets

IP subnets are used to assist in IP resolution in the following three scenarios:

- If a switch is using RFC3580 (VLAN enforcement of access control), the process for determining
 an IP address is much more difficult. In this scenario, IP subnets can be defined for each VLAN to
 provide an IP range filter, which can be used to filter the list of IPs discovered on the switch. IP
 subnets also provide a way to specify a gateway router for the VLAN's subnet, which can be used
 for doing SNMP reads on a router if DHCP snooping did not capture the relay router.
- When VRRP or HSRP is used, and you want ExtremeControl to query the router if needed,
 ExtremeControl needs to know the primary/secondary router relationship. This order of
 precedence can be defined in the IP subnet and ensures that ExtremeControl queries the primary
 router first to get the most accurate data. This is needed in a VRRP or HSRP environment,
 because both routers send out a DHCP inform message, and it is most likely that the
 ExtremeControl Gateway gets the secondary router's message last causing it to query the
 incorrect router.
- When DHCP snooping is used, the router SNMP credentials are not the same for all routers. In this scenario, if you want ExtremeControl to query the router for IP resolution, the IP subnets can be used to define the mapping between the relay router IPs and the correct SNMP credentials to use for them.


You can add, edit, or delete IP subnets using the toolbar buttons at the top of the table. There is also a

File Import button that lets you import a file of IP subnets; see the File Import window for the file format that must be used.

The **Global IP subnets** option is used to create a global list of IP ranges used for the purpose of IP Resolution. The IP Resolution process ignores any IP address outside the configured ranges. The checkbox is disabled unless there is at least one subnet configured.

Hostname Resolution

The tab is used to define how and when ExtremeControl resolves an end-system's hostname and an end-system's username. These parameters are engine for ExtremeControl Gateways, L2 ExtremeControl Controllers, and L3 ExtremeControl Controllers.



Hostname Resolution

Use this checkbox to enable or disable hostname resolution for ExtremeControl engines. Hostname resolution is only performed for end-systems for which ExtremeControl has an IP address.

DNS Hostname Resolution

This option allows the use of reverse DNS lookup on the ExtremeControl engine to resolve an end-system's hostname. In order for this option to work, a valid DNS server IP address must have been specified when the ExtremeControl engine was installed. Use the **DNS Timeout** field to specify the amount of time in seconds that the ExtremeControl engine waits after making a reverse DNS lookup prior to giving up and moving on to the next hostname resolution mechanism.

NetBIOS Hostname Resolution

This option allows the ExtremeControl engine to make a NetBIOS request to the end-system to query the end-system for its hostname. See <u>NetBIOS Timeout</u> and <u>NetBIOS Timeout Retry Count</u> on the Miscellaneous tab.

Kerberos Hostname Resolution


This options allows the ExtremeControl engine to do a lookup in the table of data learned from Kerberos snooping, to resolve the end-system's host name.

DHCP Hostname Resolution

This options allows the ExtremeControl engine to do a lookup in the table of data learned from DHCP snooping, to resolve the end-system's host name.

Username Resolution

The tab is used to define how and when ExtremeControl resolves an end-system's hostname and an end-system's username. These parameters are engine for ExtremeControl Gateways, L2 ExtremeControl Controllers, and L3 ExtremeControl Controllers.



Username Resolution

Use this checkbox to enable or disable username resolution, which allows the ExtremeControl engine to try resolve the name of a user currently on an end-system when the username was not part of the authentication request. MAC authentication and L3 ExtremeControl Controller authentication are the two cases where username resolution can currently be used.

Registration Username Resolution

This options causes ExtremeControl to use the username used for authenticated registration or the user's full name for unauthenticated registration in the format: Last Name, First Name.

Kerberos Username Resolution


This options allows the ExtremeControl engine to do a lookup in the table of data learned from Kerberos snooping, to resolve the name of the user currently logged into the end-system.

Ignored Kerberos Usernames

The table is used to define usernames for which Kerberos data is ignored. This is useful when applications running on an end-system use a global user over the Kerberos protocol to pass information for a program. Two known cases of this would be Sophos Anti-Virus software and the IBM Rational ClearCase source control system. You can add, edit, or delete entries using the toolbar buttons at the top of the table.

Reauthentication

This tab is used to define global session-timeout behavior for L2 ExtremeControl Controllers and ExtremeControl Gateways, and how ExtremeControl Gateways reauthenticates end-systems on various NAS switches. This tab is not applicable for L3 ExtremeControl Controllers.



Set Reauthentication Time for Accepted End-Systems to Assessment Interval

This option allows the ExtremeControl engine to set session-timeouts for accepted end-systems, causing the end-system to be reauthenticated the next time a scan needs to be performed. This option is required for networks using 802.1X authentication on wireless switches that do not support the IEEE

802.1X Port Reauthenticate MIB. It is also required for networks using MAC or Web-Based authentication on third-party switches. These switches do not have a mechanism to force re-authentication on end-systems when assessment is complete. This checkbox does not apply for Layer 3 ExtremeControl Controller engines.

Accept Session Timeout

If enabled, this timeout applies to all end-systems that are accepted, but not considered by ExtremeControl to be unregistered end-systems. If both this option and the "Set Reauthentication Time For Accepted End-Systems to Assessment Interval" option are enabled, ExtremeControl uses the lower value. The timeout can be either:

- Enabled For Session-Timeout Switches (Default) The timeout only applies to accepted endsystems authenticated for a NAS switch where ExtremeControl cannot reauthenticate sessions on demand via SNMP or RFC3576.
- Enabled for All Switches The timeout is applied to any accepted end-system (not considered by ExtremeControl to be unregistered) on any switch.

Quarantine Session Timeout

If enabled, this timeout applies to all end-systems quarantined by ExtremeControl. The timeout can be either:

- Enabled For Session-Timeout Switches (Default) The timeout is only applied to quarantined end-systems that were authenticated for a NAS switch where ExtremeControl cannot reauthenticate sessions on demand via SNMP or RFC3576.
- Enabled for All Switches The timeout is applied to any quarantined end-system on any switch.

Unregistered Session Timeout

If enabled, this timeout applies to all end-systems determined to be unregistered by ExtremeControl. The timeout can be either:

- Enabled For Session-Timeout Switches (Default) The timeout only applies to unregistered endsystems authenticated for a NAS switch where ExtremeControl cannot reauthenticate sessions on demand via SNMP or RFC3576.
- Enabled for All Switches The timeout is applied to any unregistered end-system on any switch.

Assessing Session Timeout

If enabled, this timeout applies to all end-systems being scanned by ExtremeControl. The timeout can be either:

- Enabled For Session-Timeout Switches (Default) The timeout applies to end-systems being assessed authenticated for a NAS switch where ExtremeControl cannot reauthenticate sessions on demand via SNMP or RFC3576.
- Enabled for All Switches This option tells ExtremeControl to apply the session timeout for endsystems being assessed on any switch.

Session Deactivate Timeout

This option can be used to provide more up-to-date information about which end-systems are still active on the network. When it is enabled, ExtremeControl checks periodically to determine if an authentication request is received from an end-system within the specified time. If a user is still on the

network, then the user is reauthenticated and a new event is generated stating the user is still active on the network. If the user is no longer on the network, the session is removed on the switch and the end-system is displayed in ExtremeControl with the **Disconnected** state. (Note that when a user leaves the network within the period of time specified, ExtremeControl does not display them as "**Disconnected** until the specified time has passed.) While this option does provide a more up-to-date list of active end-systems, RADIUS accounting should be used to provide real-time connection status. This option is useful when RADIUS accounting is not desired or is not supported on certain network devices.

NOTE:

The timeout process could be off by approximately 60 seconds from the specified time, depending on when ExtremeControl runs the check for authentication requests.

Switch Reauthentication Configuration

This table is used to configure the reauthentication method an ExtremeControl engine uses on a switch. For example, you may want to add support for another wireless switch. In this case, you would add an entry for the new switch by selecting the Add button, entering the sysObjectId of the switch, and setting the Reauthentication Type to either RFC3576 (if the switch supports it) or Session Timeout. This table is also where you can disable port link control for switches by selecting the switch, selecting the Edit button, and setting the Port Link Control option to disabled.

If you've deleted or edited any of the default configurations, the **Restore Defaults** button restores them to their original state and add back any that are missing. Any custom entries you added are retained unless they have the same sysObjectId as a default configuration. Following a restore, you need to save the configurations.

Miscellaneous

Use this tab to configure various parameters for your network engines including port link control, NetBIOS, Kerberos, and Microsoft NAP.

Default					
Credentials	Port Link Control				
Network Settings	Enabling will cause NAC to link down and link up the port for end-systems whenever the end-system changes state. This setting is ignored for NAC Controllers and Extreme equipment with multi-authentication enabled. Intended for use in VLAN environments				
Device Type Detection	to force end-systems to get new IP addresses. Option must be manually disabled for third-party environments with multi-authentication.				
IP Address Resolution	☐ Enable Port Link Control				
Hostname Resolution	Port Down Time:	2	\$		
Username Resolution	Enable For Authentication	✓ MAC	✓ CHAP	▼ 802.1X	
Reauthentication	Type(s):	▼ PAP	 MsCHAP		
Miscellaneous	NTLM Health Check				
Auditing	Interval (in seconds):	60	\$		
	Timeout (in seconds):	5	<u> </u>		
	NetBIOS				
	NetBIOS Timeout (in seconds):	3	\$		
	NetBIOS Timeout Retry Count:	1	‡		
	Kerberos				
	Allow use of data from Kerberos request packets				
	Reauthenticate users when a Kerberos username change is detected				
	Reset Authentication Type on Kerberos login for MAC and IP authentication				
	✓ Kerberos Age Out Time (in ho	ours) 12	-		
	Microsoft NAP				
	Reset Authentication Type for NAP enabled end-systems				
	Override Quarantine Policy for NAP enabled end-systems				
	Proxy NAP attributes to the s	witch			
	Administrative Requests				
	Allow EAP-Message attribute in administrative requests				
				Save Cancel	

Port Link Control

Enable Port Link Control

Use this checkbox to enable or disable port link control. Port link control is required if you are using VLAN only (RFC 3580) switches or if you are using policy with VLANs on EOS policy-enabled switches. When a VLAN is assigned to a switch port, the end-system needs to get a new IP address for the assigned VLAN. To do this, the ExtremeControl engine links down the port, waits the configured amount of time, and then links up the port, causing the end-system to make a new DHCP request and get a new IP address.

Be aware that when multiple devices are connected to a switch port where authentication is enabled (such as an IP phone cascaded with a PC on a single port), port link down disconnects all devices. In this scenario, you may want to disable port link control, set the ExtremeControl profile to "Use Assessment Policy During Initial Assessment Only," and set the DHCP lease time for the IP address pools that correspond to the VLAN(s) associated to the Quarantine and Assessment access policies to a low value (e.g. 1 minute).

This setting is ignored for ExtremeControl Controllers and EOS equipment with multi-authentication enabled. The option must be manually disabled for third-party environments with multi-authentication.

In the **Port Down Time** field, enter the amount of time in seconds that the engine waits before linking up the port. The time must be sufficient to cause the end-system to make the DHCP request.

In the **Enable for Authentication Types** field, you can enable port link control for only specific authentication types, depending on the checkboxes you select. For example, you can disable port link control for 802.1x, but have it enabled for MAC authentication so that a port is only linked down when a MAC authentication session changes VLANs.

NTLM Health Check

Windows provides an authentication protocol called New Technology LAN Manager (Windows NTLM). NTLM is a challenge-response authentication protocol used to authenticate a client to a remote on an Active Directory. NTLM Health Check is an Access Control test that can be enabled in LDAP configurations using NTLM Authentication, and actively used in Access Control AAA configuration rules.

When enabled, the health check runs at regular intervals and test the current domain controller for the specified domain. The test domain is specified in the LDAP configuration. The interval and a timeout can be configured in the Miscellaneous section of ExtremeControl settings (Control > Access Control > Configuration > Global Engine Settings > Engine Settings > Default > Miscellaneous).

Interval (in seconds)

This value tells Access Control how often to run the health check.

Timeout (in seconds)

This value tells Access Control how long to wait for an authentication response from the Active Directory. If the timeout threshold is exceeded, the failover occurs and the Access Control Lost Partial Contact with LDAP Service alarm is generated.

To complete the health check setup, refer to the documentation sections on NTLM
Authentication and Advanced. To configure these additional settings in Access Control, go to Control > Access Control > Configuration > AAA > LDAP Configuration.

NetBIOS

This section controls the timeout and retries that an ExtremeControl engine uses when making NetBIOS requests for IP resolution, MAC resolution, or hostname resolution.

NetBIOS Timeout

The amount of time in seconds that an ExtremeControl engine waits for a response to a NetBIOS request to an end-system, before giving up on that request and retrying.

NetBIOS Timeout Retry Count

The number of times an ExtremeControl engine retries making a NetBIOS request to an end-system, if the end-system does not respond.

Kerberos

Controls how an ExtremeControl engine deals with data it receives from Kerberos snooping.

Allow Use of MAC Resolution for Kerberos Data Processing

When end-systems are behind a router, the ExtremeControl engine uses MAC resolution to resolve an end-system's MAC address from its IP address. This is because when end-systems are behind a router (not in the local network), the Kerberos packets carry the MAC address of the router instead of the end-system. This option allows you to turn off the use of MAC resolution for Kerberos processing, if desired.

Allow Use of Data from Kerberos Request Packets

This option allows the use of data such as username and hostname, from Kerberos request packets. The data in the request packet is provided by the user, and is not guaranteed to be accurate, since it is not authenticated.

Reauthenticate Users on Kerberos Username Change Detected

This option causes the ExtremeControl engine to reauthenticate a user if the username in the Kerberos packet changes.

Reset Authentication Type on Kerberos Login for MAC and IP Authentication

This option is supported for ExtremeControl deployments with inline ExtremeControl Controllers that can capture the end user login. When a user logs in via Kerberos, (for example, a user logs into a Windows domain,) the ExtremeControl Controller resets the authentication type from MAC (for an L2 ExtremeControl Controller) or IP (for an L3 ExtremeControl Controller) to Kerberos. The Kerberos authentication type can then be used by rules to give elevated access to users that have successfully logged into a Windows domain.

Kerberos Age Out Time

This option provides a way to disable the aging out of Kerberos authentication data. This authentication data is used by ExtremeControl to provide elevated access to end-systems. By default, the authentication data is automatically aged out every 12 hours. During that 12-hour period, any time the end-system reauthenticates with ExtremeControl, the user would receive their elevated access privileges. After the 12 hours is exceeded and the authentication data is aged out, the end-system must log in again to get their elevated access. You can use this option to change the age out time or disable the aging altogether. For example, you might want to change the 12 hours to 8 hours, based on a shorter 8-hour workday.

WARNING:

Keep in mind that disabling the age out would create a potential security hole. Elevated access is tied to the end-system, so if it isn't aged out, the elevated access is always available. For example, if a user leaves their laptop and someone logs them out and then logs in as a local user, that person continues to have the elevated access privileges of the original user. Also, a person could spoof someone else's MAC address and receive their elevated access, if the access isn't aged out.

Microsoft NAP

This section provides options related to Microsoft NAP for Windows.

Reset Authentication Type for NAP Enabled End-Systems

When this option is enabled, the ExtremeControl engine resets the authentication type from 802.1x to MS NAP (Microsoft NAP), if the end-system authenticating is NAP-enabled (Windows XP SP2 or higher) and the 802.1x authentication request was proxied to a NAP-enabled server. The MS NAP (Microsoft NAP) authentication type can then be used by rules to assign a different ExtremeControl profile. To configure ExtremeControl to perform as it did in ExtremeControl version 3.1.x, you can create a rule that maps the MS NAP (Microsoft NAP) authentication type to the Pass Through ExtremeControl Profile. With this profile, ExtremeControl does not assess the end-system, and uses the NAP determination of whether or not to guarantine a user.

Override Quarantine Policy for NAP Enabled End-Systems

This option allows ExtremeControl to replace the quarantine policy for NAP-enabled end-systems, using the quarantine policy defined in the profile's Use Quarantine Policy field. Be aware that when this NAP option is enabled, the Use Quarantine Policy checkbox becomes active for all ExtremeControl profiles, even if assessment is disabled. However, you can deselect the checkbox for an individual profile, in which case the policy from the RADIUS attributes is applied.

Proxy NAP Attributes to Switch

This option is disabled by default. When disabled, the following attributes are **not** proxied to the switch if they are present in the response from the backend RADIUS server:


- MS-Machine-Name
- MS-Extended-Quarantine-State

- MS-RNAP-Not-Quarantine-Capable
- MS-Quarantine-State

If the option is enabled, the attributes are proxied to the switch.

Auditing

Use this tab to enable auditing of users connected to the ExtremeControl engine CLI via SSH.



Enable Auditing

Selecting the **Enable Auditing** option enables the **Auditing Rules** field, where you can configure ExtremeCloud IQ - Site Engine to store all commands entered by a user connected to the ExtremeControl engine CLI via SSH in the engine's local syslog file.


Auditing Rules

Remove the # symbol from the beginning of a command line to enable the command and store user commands entered using the ExtremeControl engine CLI.

ExtremeControl Engine Groups

The ExtremeControl Engine Groups panel is displayed in the right panel when you select the ExtremeControl Engine Groups folder in the left panel. (The ExtremeControl Engine Groups folder is only displayed if you have created engine groups.) The tab displays a table of information about the engine groups in the folder.

Use the table options and tools to filter, sort, and customize table settings. You can access the options by selecting the down arrow in the right corner of any column header.



Name

The name of the engine group.

ExtremeControl Configuration

The ExtremeControl Configuration currently selected for this engine group.

Portal Configuration

If your network is implementing Registration or Assisted Remediation, the Portal Configuration that defines the branding and behavior of the website used by the end user during the registration or remediation process.

AAA Configuration

The AAA Configuration used by this engine group.

Policy Mapping

The Default policy mapping can be viewed in the ExtremeControl Configurations tree (under ExtremeControl Profiles) or accessed from the Edit ExtremeControl Profile window.

Engine Settings

The Engine Settings configured for the group. Use the Edit Engine Settings window to specify and configure engine settings.

ExtremeControl Access Control Group Editor

This panel lists the various rule groups used to define the criteria for the rules used in your ExtremeControl configuration. You can use this window to view and edit the defined rule groups and also to add new rule groups for use in your ExtremeControl configuration. Any changes made in this window are written immediately to the ExtremeCloud IQ - Site Engine database.

ExtremeCloud IQ - Site Engine comes with system-defined rule groups. ExtremeCloud IQ - Site Engine also contains system-defined end-system groups that automatically populate. The Assessment Warning end-system group includes end-systems that have assessment warnings and must acknowledge them before being granted access to the network. The blocked list end-system group includes end-systems denied access to the network. The other system-defined groups are populated as the end-systems register through the Registration portal.


Select from the following rule group categories when you create a new rule group:

Category	Group Types	Value Types		
All Groups	All Types	A list of all group types.		
Device Type Groups	Device Type	A list of device types.		
End-System Groups	Hostname	A list of hostnames, which can be an exact match or wild card (for example, *.extremenetworks.com).		
	IP	A list of IP addresses or subnets.		
	LDAP Host Group	A way to group hosts by doing an LDAP lookup on the resolved hostname of the end-system detected on the network, which can be an exact match or wild card.		
	MAC	A list of MAC addresses, MAC OUI, or MAC masks.		
Location Groups	Location	A list of switches, switches and ports, or switches and SSIDs.		
Time Groups	Time of Week	A list of the times of the week when the end user is accessing the network.		
User Groups	LDAP User Group	A list imported from an LDAP Server, organized by Organization Unit (OU), which can be an exact match or wild card.		
	RADIUS User Group	A list of attributes returned by the RADIUS server, which can be an exact match or wild card.		
	Username	A list of usernames, which can be based on an exact match or a wild card.		

To access this window:

- Access the ExtremeCloud IQ Site Engine > ExtremeControl tab.
- Select the Access Control tab.
- Expand the **Group Editor** tab in the left-panel.

The right-panel rule group detail table opens.



The following buttons are included in the rule group detail table:

Add 🕥 Add...

Use this button to add rule groups or to import MAC entries from a file for viewing and assigning to various end-system groups.

Edit p Edit...

Use this button to edit existing rule groups.

Сору 🕞 Сору...

Use this button to copy a selected rule group.

Delete 🔘 Delete

Use this button to delete existing rule groups.

Refresh 🥰

Use this button to reload group entries in the table.

Import

Use this button to import MAC entries into groups.

Reset Reset

Use this button to clear the search field and any filters, and to update the data in the table.

The following columns display in the rule group detail table:

Name

The name of the rule group.

Type

The type selected for the specific rule group; for example, an end-system group could have a type of MAC.

Used By

The name of the Identity and Access configuration using this rule group.

Description

A description of the rule group.

Add/Edit Device Type Group

There are nine system-defined operating system family device type groups that are automatically populated by ExtremeCloud IQ - Site Engine: Android, Apple iOS, Blackberry, Chrome OS, Game Console, Linux, Mac, Windows, and Windows Mobile. You can view these system-defined groups and your other device type groups by expanding the ExtremeControl Configurations > Group Editor > Device Type Groups left-panel tree.

Device type groups are comprised of entries that ExtremeControl uses to determine if an end-system's device type matches the group. Entries can be a specific device type or a wildcard, such as Windows 7 or win*. If an entry does not already contain a wildcard, ExtremeCloud IQ - Site Engine creates a wildcard by adding an asterisk (*) to the beginning and end of the entry. For example, if the entry is **Gentoo**, the match pattern is *Gentoo* allowing a match for any end-system device type that contains Gentoo. This allows you to restrict the match to a very specific value that might include a version number or model number, or expand the match to include all versions and model numbers of a certain operating system or hardware family.

For additional information about how to use device type groups, see How to Use Device Type Profiling.

NOTE: Changes to rule groups do not require an enforce. Changes are automatically synchronized with engines on the next status update. Changes do not affect end-systems until the next authentication and/or assessment occurs.

To access the Add New Group window, select **Add** (Add...) in the Device Type Groups right panel.

The Add New Group window opens.



Name

Enter a new name for the device type group. After a group is created, you cannot edit the name of the group.

Description

Enter a description of the device type group.


Type

To create a new device type group, select **Device Type** from the drop-down list.

Select the Create button to open the Device Type Entry Editor section of the window.

Select the **Select from Existing Types** button () to open the Select Device Types window from which

you can choose a list of predefined entries. Select the **Add** button in the Device Type Entry Editor section of the window to open the Add Entry window.



Use this window to add a new entry by entering a device type or a wildcard, such as Google Pixel or *pixel. Alternately, you can select a type from a list of entries that already appear in existing device type groups from the Select Device Types window. This window can be accessed by selecting the **Select from Existing Types** button. This list allows you to multi-select entries, and each entry appears as a separate row in the table. The list also allows you to select **Unknown** that matches against any device that does not have an operating system name, either due to failed detection or because detection hasn't happened yet.


All entries selected from the list are assigned the same description. If you would like a separate description for each type, you need to add each entry individually.

End-Systems

Use the **End-Systems** tab to view end-system connection information for a single ExtremeControl engine, all ExtremeControl engines, or all the engines in an engine group, depending on what you select in the left-panel tree. You can also monitor end-system events and view the health results from an end-system's assessment.

The **End-Systems** tab is available from the **Control** tab. You can also access the tab by selecting a single ExtremeControl engine, the All Engines folder, or an engine group in the left-panel tree, then selecting the **End-Systems** tab in the right panel. Selecting a single engine or engine group displays only the end-systems accessing the network via the selected engines.

Use the table options and tools to <u>filter, sort, and customize</u> table settings. Access the options by selecting the down arrow in the right corner of any column header.



End-Systems

This table displays the last known connection state for each end-system that has attempted connection.

State

The end-system's connection state:

- Scan The end-system is currently being scanned.
- Accept The end-system is granted access with either the Accept policy or the attributes returned from the RADIUS server.
- Quarantine The end-system is quarantined because the assessment failed.
- Reject The end-system was rejected because the assigned ExtremeControl profile was set to Reject, the MAC Locking test failed, or the RADIUS server was reachable but rejected the authentication request.
- Disconnected All sessions for the end-system are disconnected. This state is only applicable for end-systems connected to switches that have RADIUS accounting enabled.

- Error Indicates one of nine problems:
 - the MAC to IP resolution failed, if assessment is enabled
 - the MAC to IP resolution timed out, if assessment is enabled
 - all RADIUS servers are unreachable
 - the RADIUS request was non-compliant
 - all assessment servers are unavailable
 - the assessment server can't reach the end-system
 - no assessment servers are configured
 - the assessment server is not compatible with the current version of ExtremeControl
 - the username and password configured in the <u>Assessment Server panel</u> of the ExtremeControl options (Administration > Options > ExtremeControl > Assessment Server) are incorrect for the assessment server.

ID

The device identification number.

Last Seen

The last time the end-system was seen by the ExtremeControlengine.

The End-Systems table is sorted by the Last Seen Time by default. Sorting using any other column will automatically pause the table to allow sorting on those columns (except the OUI Vendor and Switch Nickname columns - these columns cannot be sorted). Reverting to a Live view will revert back to the "Last Seen Time"

Note:

IP Address

The end-system's IP address. ExtremeControl engine.

sort, in descending order.

OV MAC Address

The end-system's OV MAC address.

MAC Address

The end-system's MAC address. MAC addresses can be displayed as a full MAC address or with a MAC OUI (Organizational Unique Identifier) prefix.

MAC OUI Vendor

The vendor associated with the MAC OUI.

Host Name

The end-system's hostname.

Device Family

The hardware family or the operating system family for the end-system.

Device Type

The hardware type or the operating system type for the end-system.

User Name

The user name used to connect.

Site

The site of the switch to which the end-system is connected.

Switch IP

The IP address of the switch to which the end-system is connected. If the end-system is connected to an ExtremeControl Controller engine, this is the ExtremeControl Controller PEP (Policy Enforcement Point) IP address.

Switch Nickname

An alternate name for the switch.

NOTE: Configure the nickname on the <u>Device Annotation tab</u> in the <u>Configure Device</u> window.

Switch Port

The port alias (if defined) followed by the switch port number to which the end-system connected. If the end-system is connected to a Layer 2 ExtremeControl Controller engine, this is the ExtremeControl Controller PEP (Policy Enforcement Point) port. However, for Layer 3 ExtremeControl Controller engines, this column is blank.

- If you add or update the port alias on the switch, you must enforce the ExtremeControl engine in order for the new information to be displayed in the End-Systems table.
- If you don't want the port alias displayed, remove the PORT_DESCRIPTION_FORMAT variable from the /opt/nac/server/config/config.properties file. If this variable is removed, only the switch port number is displayed.

Policy

The name of the ExtremeControl policy role assigned to the end-system when it connected to the network.

Authorization

The attributes returned by the RADIUS server for this end-system. If the end-system is connected to a switch that supports multi-authentication, then this column may not reflect the actual active policy for the authenticated user. For Layer 3 ExtremeControl Controller engines, this column displays the policy assigned to the end-system for its authorization.

Risk

The overall risk level assigned to the end-system based on the health result of the scan:

- Red High Risk
- Orange Medium Risk
- Yellow Low Risk
- Green No Risk
- Gray Unknown

Profile

The name of the ExtremeControl profile assigned to the end-system when it connected to the network.

Reason

Provides information about the reason the ExtremeControl profile is assigned to the end-system.

Authentication Type

Identifies the latest <u>authentication method</u> used by the end-system to connect to the network. (For Layer 3 ExtremeControl Controller engines, this column displays "IP.")

State Description

This column provides more details about the end-system state. For example, if the end-system's connection state is Reject, this column might list the RADIUS server (primary or secondary) that rejected the authentication request.

Extended State

Provides the reasons why the end-system is in its particular connection state. It gives you an idea as to why a certain policy was applied to the end-system or why the end-system was rejected.

ExtremeControl Engines/Source IP

The ExtremeControl engine to which the end-system is connecting.

Engine Group

This column is only displayed if you have multiple engine groups. It displays what engine group the ExtremeControl engine was in when the end-system event was generated. For example, if the engine was in Engine Group A when an end-system connected, but then later the engine was moved to Engine Group B, this column would still list Engine Group A for that end-system's entry.

RFC3580 VLAN

For end-systems connected to RFC 3580-enabled switches, this is the RFC3580 VLAN ID assigned to the end-system.

Warning Time

Shows the time for warning. This column is hidden by default.

Last Quarantined

The last time the end-system was quarantined.

Score

The total sum of the scores for all the health details that were included as part of the quarantine decision.

Top Score

The highest score received for a health detail in the health result.

Actual Score

The actual score is what the total score would be if all the health details including those marked Informational and Warning were included in the score.

Switch Port Index

The SNMP index (ifIndex) of the port to which the end-system connected.

Switch Location

The physical location of the switch to which the end-system connected. If the end-system is connected to an ExtremeControl Controller engine, this is the ExtremeControl Controller PEP (Policy Enforcement Point) location.

ELIN

An extended set of data for an end-system based on a MAC address.

Port Info Raw

Displays unformatted information as it is received from the port.

All Authentication Types

This column displays all the authentication methods the end-system has used to authenticate. The authentication types are listed in order of precedence from highest to lowest: Switch Quarantine, 802.1X, CHAP, PAP, Kerberos, MAC, CEP, RADIUS Snooping, Auto Tracking. View details about each authentication session (such as the ExtremeControl profile that was assigned to the end-system for each authentication type) in the End-System Events tab.

Last Scan Result

The last scan result assigned to the end-system: Scan, Accept, Quarantine, Reject, Error. This is the state assigned to the end-system as a result of the last completed scan. This typically matches the end-system State if scanning is currently enabled and has been performed recently.

Last Scanned

The last time an assessment (scan) was performed on the end-system.

First Seen

The first time the end-system was seen by the ExtremeControl engine.

NAP Capable

Indicates whether the end-system is Microsoft NAP (Network Access Protection) capable: Yes or No

Custom

Use this column to add additional information about the end-system. To add or edit custom information, right-click on the table and select **Edit Custom Information**. You can add information for up to four Custom columns. The columns for Custom 2, Custom 3, and Custom 4 are hidden by default. To display these columns, select the down arrow to the right of the table header and select Columns > Column 2, Column 3, or Column 4.

NOTE: Change the name of the Custom columns in the ExtremeControl options.

Registered User

The registered username supplied by the end-user during the registration process.

Registered Email

The registered email address supplied by the end-user during the registration process.

Registered Phone

The registered phone number supplied by the end-user during the registration process.

Sponsor

The registered user's sponsor, if sponsorship is enabled.

Registration

Custom information supplied by the end-user during the registration process.

Registration Description

The device description supplied by the end user during the registration process.

Groups

Displays any end-system and/or user groups to which the end-system belongs.

Group 1-3

Displays the names of up to three end-system and/or user groups to which the end-system belongs.

Zone

Displays the end-system zone to which the end-system is assigned.

Request Attributes

Indicates if RADIUS attributes are requested.

Registration Type

Shows the type of end-system connection (for example, **Transient**).

RADIUS Server IP

The IP address of the RADIUS server to which the end-system authenticated.

Source

Displays the origin of the end-system in the network:

- Access Controlengine An Access Control engine.
- Wireless Manager An ExtremeWireless Controller or AP.
- ExtremeXOS/Switch Engine ID Manager An Extreme switch running ExtremeXOS/Switch Engine with the Identify Manager feature configured to send events to ExtremeCloud IQ Site Engine.
- OneFabric Connect An ExtremeConnect module (e.g. Solutions Architecture and Innovation (SAI) integration)
- One Controller The Extreme SDN Controller.

TLS Client Certificate Expiration

Expiration date of the TLS Client Certificate issued for 802.1x authentication.

TLS Client Certificate Issuer

Name of the issuer of the TLS Client Certificate issued for 802.1x authentication.

TLS Client Certificate Fingerprint

The attributes in an SSL handshake used for identifying the end-system.

Actions

TIP: These actions are also available from the right-click menu off an end-system entry in the table.

Force Reauthentication

Forces the selected end-system to re-authenticate. End-systems authenticated to a VPN device are disconnected from the VPN.

Force Reauth and Scan

Forces the selected end-system to re-authenticate and undergo an assessment (scan). (End-systems authenticated to a VPN device are disconnected from the VPN.) The assessment only takes place if scanning is enabled in the ExtremeControl profile assigned to the end-system.

Add to Group

Lets you add the selected end-system to a specific end-system or user group. If the end-system is a registered device, it can be added to a registration group. After adding an end-system to a group, any rules created that involved that group apply to the end-system as well. Changes to end-system group membership do not require an enforce and are synchronized with engines immediately. Changes do not affect the end-system until the next authentication or assessment occurs.

Lock MAC

Opens the <u>Add MAC Lock window</u> where you can lock the MAC address of the selected end-system to a switch or switch and port.

Show Details

Opens the <u>End-System Details tab</u> where you can view summary information for the end-system selected in the table.

Delete

Deletes the selected end-system entries from the table and also deletes the associated end-system events. You are given the option to delete any custom information, group assignment, MAC locks, and registration and web authentication associated with the end-systems.

The Force Delete of End-System option completely deletes the end-system from ExtremeCloud IQ - Site Engine, regardless of whether the end-system reauthentication is successful when the delete is executed. The option is deselected by default. When deselected, it prevents possible synchronization conditions where the authentication session remains active on the switch even though the end-system has been deleted from ExtremeCloud IQ - Site Engine. These conditions can occur when there are underlying issues that prevent the end-system reauthentication from completing properly.

NOTES: The Delete operation does not remove an end-system from the blocked list group. Blocked list is a special group that requires end-systems to be manually removed using the Edit End-System Group window.

Deleting an end-system from the table also deletes the user's current authentication. If the user is connected to the network at the time of the delete, they are forced to re-authenticate.

Menu Buttons

The menu at the top of the window contains most of the options available via a right-click previously mentioned in the <u>Actions</u> section above, as well as the End-System Events button, described below.


All End-System Events

Opens the <u>End-System Events tab</u> where you can view information about events for all end-systems accessing your network.

End-System Events Tab

This tab displays historical connection information for all end-systems accessing your network. End-system events are stored daily in the database. In addition, the end-system event cache stores in memory the most recent end-system events and displays them here in this tab. This cache allows ExtremeCloud IQ - Site Engine to quickly retrieve and display end-system events without having to search through the database. You can configure parameters for the event cache (such as the number of events to display) using the End-Systems Event Cache options in the ExtremeControl Options view (Administration > Options > ExtremeControl > End-Systems Event Cache).

NOTE: The **End-System Events** tab displays events up to the most recent delete event for the end-system, if one exists. If you want to see events that happened prior to the most recent delete event, use the **Search for Older Events** button.



State

The end-system's connection state:

- Scan The end-system was scanned.
- Accept The end-system was granted access with either the Accept policy or the attributes returned from the RADIUS server.
- Quarantine —The end-system was quarantined because the assessment failed.

- Reject The end-system was rejected because the assigned ExtremeControl profile was set to Reject, the MAC Locking test failed, or the RADIUS server was reachable but rejected the authentication request.
- Disconnected This end-system session was disconnected, however other sessions for the end-system may still be active. For example, the end-system may have a disconnected session with an authentication type of 802.1X, but still have an active MAC authentication session. This state is only applicable for end-systems connected to switches that have RADIUS accounting enabled.
- Error Indicates one of nine problems:
 - the MAC to IP resolution failed
 - the MAC to IP resolution timed out
 - all RADIUS servers are unreachable
 - the RADIUS request was non-compliant
 - all assessment servers are unavailable
 - the assessment server can't reach the end-system
 - no assessment servers are configured
 - the assessment server is not compatible with the current version of ExtremeCloud IQ -Site Engine
 - the username and password configured in the <u>Assessment Server panel</u> of the ExtremeControl options (Administration > Options > ExtremeControl > Assessment Server) are incorrect for the assessment server

Time Stamp

The date and time the end-system connected.

ExtremeControl Engine/Source IP

The IP address of the ExtremeControl engine on which the event occurred.

Profile

The name of the ExtremeControl profile assigned to the end-system when it connected to the network.

IP Address

The end-system's IP address.

MAC Address

The MAC address of the end-system on which the event occurred. MAC addresses can be displayed as a full MAC address or with a MAC OUI (Organizational Unique Identifier) prefix.

User Name

The username used to connect.

Host Name

The end-system's host name.

Device Family

The hardware family or the operating system family for the end-system.

Device Type

The hardware type or the operating system type for the end-system.

State Description

This column provides more details about the end-system state. For example, if the end-system's connection state is Reject, this column might list the RADIUS server (primary or secondary) that rejected the authentication request.

Extended State

Provides additional information about the end-system's connection state.

Reason

Provides additional information about the reasons why the end-system is in its particular connection state. It provides information as to the reason a policy is applied to the end-system or the reason the end-system is rejected.

Authorization

The attributes returned by the RADIUS server. If the end-system is connected to a switch that supports multi-authentication, then this column may not reflect the actual active policy for the authenticated user. For Layer 3 ExtremeControl Controller engines, this column displays the policy assigned to the end-system for its authorization.

Auth Type

Identifies the authentication method used by the end-system to connect to the network. For Layer 3 ExtremeControl Controller engines, this column shows **IP**.

Switch IP

The IP address of the switch to which the end-system connected. If the end-system is connected to an ExtremeControl Controller engine, this is the ExtremeControl Controller PEP (Policy Enforcement Point) IP address.

Switch Nickname

The nickname defined for the switch to which the end-system is connected.

Switch Port

The switch port number to which the end-system is connected. If the end-system is connected to a Layer 2 ExtremeControl Controller engine, this is the ExtremeControl Controller PEP (Policy Enforcement Point) port. However, for Layer 3 ExtremeControl Controller engines this column is blank.

Switch Location

The physical location of the switch to which the end-system is connected. If the end-system is connected to an ExtremeControl Controller engine, this is the ExtremeControl Controller PEP (Policy Enforcement Point) location.

Last Scan Time

Displays the last time ExtremeCloud IQ - Site Engine scanned the end-system on which the event occurred.

Zone

Displays the end-system zone to which the end-system is assigned. For additional information, see <u>End-System Zones</u>.

Registration Type

Shows the type of end-system connection (for example, Transient).

RADIUS Server IP

The IP address of the RADIUS server to which the end-system authenticated.

Event Source

Displays the origin of the end-system in the network:

- Access Control engine An Access Control engine.
- Wireless Manager An ExtremeWireless Controller or AP.
- ExtremeXOS/Switch Engine ID Manager An Extreme switch running ExtremeXOS/Switch
 Engine with the Identify Manager feature configured to send events to ExtremeCloud IQ Site
 Engine.
- OneFabric Connect An ExtremeConnect module (e.g. Solutions Architecture and Innovation (SAI) integration)
- One Controller The Extreme SDN Controller.

Engine Group

This column is only displayed if you have multiple engine groups. It displays what engine group the ExtremeControl engine is in when the end-system event was generated. For example, if the engine began in Engine Group A when an end-system connected, then the engine is moved to Engine Group B, this column still lists Engine Group A for that end-system's entry.

Search for Older Events

This button lets you search for older events stored in the database outside of the end-system events cache. The maximum search parameters for this extended search are configured in the End-System ExtremeControl Options view (Administration > Options > ExtremeControl > End-System Event Cache). The search is ended when any one of the parameters is reached.

- Maximum number of results to return from search
- Maximum time to spend searching for events (in seconds)
- Maximum number of days to go back when searching

For information on related topics:

- Add MAC Lock Window
- End-System Details Tab


Add/Edit End-System Group

Use this window to add a new end-system group or edit an existing end-system group. End-system groups are rule components that enable you to group together devices having similar network access requirements or restrictions. You can access the Add/Edit End-System Group window from the Manage Rule Groups window or from the end-system group field in the Create Rule window.

There are six system-defined end-system groups automatically populated by ExtremeCloud IQ - Site Engine. The first is the Assessment Warning end-system group that includes end-systems that have assessment warnings and must acknowledge them before being granted access to the network. The second is the blocked list end-system group that includes end-systems denied access to the network. The other four system-defined groups are populated as end-systems register through the Registration portal.

You can access the Create Group window by accessing the **Access Control** tab and selecting ExtremeControl Configurations > Group Editor > End-System Groups in the left-panel menu and selecting the **Add** button in the right panel.

NOTE: Changes to rule components do not require an enforce. Changes are automatically synchronized with engines on the next status update. Changes do not affect end-systems until the next authentication and/or assessment occurs.



Name

Enter a new name for the end-system group. You cannot edit the name of a group.

Description

Enter a description of the end-system group. If you are using Data Center Manager (DCM), the end-system group description contains the DCM specific settings as key/value pairs.

Type

Specify whether the end-system group be based on:

- MAC a list of MAC addresses, MAC OUI, or MAC Masks.
- IP a list of IP addresses or subnets.
- Hostname a list of hostnames: exact match or wild card (for example, *.extremenetworks.com).

LDAP Host Group - a way to group hosts by doing an LDAP lookup on the resolved hostname of
the end-system detected on the network. Note for the standard use with Active Directory, the
Engine Settings > Hostname Resolution must be configured to use DNS Hostname Resolution so
ExtremeCloud IQ - Site Engine can resolve the Fully Qualified Domain Name. In the LDAP
configuration, you must also have the "Use Fully Qualified Domain Name" checkbox selected.

Value

The MAC address, IP address, Hostname, or Attribute value of the end-system.

Description

The description of the end-system group.

Mode

For LDAP Host Groups, the mode option lets you specify whether to match any or match all of the LDAP attributes listed below. You can also use "Exists" to just check to see if a host is present in LDAP.

Custom 1

Displays additional information about the end-system. Up to four custom columns can be <u>added</u> to the table. The columns for Custom 2, Custom 3, and Custom 4 are hidden by default. To display these columns, select the down arrow next to the Custom 1 column header and select **Columns > Custom 2**, **Custom 3**, or **Custom 4**.

Add Button () Add...

Select the **Add** button to open the Add Entry window, from which you can add an entry to the table. To add or edit <u>custom</u> information, right-click on the table entry and select Edit Custom Information. You can add information for up to four Custom columns.



Edit Button 📝 Edit...

Select an entry in the Entry Editor section of the window and select the **Edit** button to open the Edit Entry window, from which you can edit an existing entry.

Delete Button 🔵 Delete

Select an entry in the Entry Editor section of the window and select the **Delete** button to delete an existing entry.

Save Button

Select the **Save** button to save the location group.



Use the Multiple MAC OUI Entries button to open a window where you can select MAC OUI vendors.

Filter

Use the <u>Filter functions</u> to filter for a specific entry based on a numeric value or text.

End-System Details

The End-System Details window provides connection state and assessment information for a single end-system. It is launched from the End-Systems View in the **Control** tab, by double-clicking any end-system in the table or selecting an end-system and then selecting **Show Details** from the Tools menu.


The End-System Details window has four tabs. The **Access Profile** tab provides end-system summary information. The End-System tab provides end-system connection state information. The **End-System Event** tab displays end-system event information. The **Health Results** tab displays end-system assessment result information.

This Help topic provides information on the four tabs:

- Access Profile Tab
- End-System Tab
- End-System Events Tab
- Health Results Tab

Access Profile Tab

The Access Profile tab presents a graphical view of end-system and health result information, providing an at-a-glance end-system summary. Select the information in each section to link to more detailed information.



Access Type

Displays the switch IP address, port index, and port that the end-system is connected to. Select to open a PortView for the switch in a new tab.

Top Application Flows

Lists the top five applications and flow counts for the end-system, listed in descending order by flow count. Select to open the Applications Dashboard in a new tab.

Device Family

Displays the end-system's operating system (OS) family (for example: Windows, Linux, Android) and OS name. Use the device family icon to quickly determine the end-system type. Select to open the **End-System** tab where you can view additional end-system details.

Health

Displays health data from the latest scan, including risk level, total score, and last scan time. Use the health icon to quickly determine risk level by color. Select to open the **Health Results** tab where you can view additional health result information and details.

Registration

Displays the end-system's registration state, user name, and sponsor. Select to open the **End-System** tab where you can view additional registration information.

Activity

Displays the last seen and first seen times for the end-system. Select to open the **End-System** tab where you can view additional end-system details.

Location

Displays location summary information, including end-system zone membership, access point information, engine group, and engine IP address. Select to open the **End-System** tab where you can view additional location information.

Physical Device Identity

Displays the end-system's MAC address, IP address, and host name. The device icon displays the end-system's physical device type with a small OS-based icon in the corner. Select to open the **End-System** tab where you can view additional end-system details.

Virtual Device Identity

If the end-system is a virtual machine, this section displays virtual device information, including VM name, ID, Guest Name, and manufacturer. Use the icon to quickly determine the virtual machine's operating system. If the end-system is not a virtual machine, this section is replaced by Custom Data.

Custom Data


Displays any custom information associated with the end-system. Custom information for an end-system is added in the End-Systems tab or End-Systems View. If the end-system is a virtual machine, this section is replaced by Virtual Device Identity.

Access Control

Displays the end-system's user name, authentication type, connection state, policy, and profile. Select to open the **End-System** tab where you can view additional end-system authentication session details.

End-System Tab

This tab presents detailed information on the selected end-system's connection, authentication, and registration. Expand the sections using the arrow buttons to see additional information.




For a definition of various fields, see the column definitions included in the **End-Systems topic**.

Changes to group membership do not require an enforce and will be synchronized with engines immediately. Changes will not affect the end-system until the next authentication or assessment occurs.

End-System Events Tab

The End-System Events tab shows all the events for the selected end-system.




You can manipulate the table data in this window in several ways to customize the view for your own needs:

- Select the column headings to perform an ascending or descending sort on the column data.
- Hide or display different columns by selecting a column heading and selecting the column options from the menu.
- Rearrange columns by dragging a column heading to the desired position.
- Filter the data in each column in the table.

Health Results Tab

The top table in the **Health Results** tab provides summary information on scan results obtained for the selected end-system. The bottom table presents the individual health result details for the scan selected in the top table. Double-click any row in the bottom table to open the Health Result Details window and view a description, solution, and result for the health result. Information is displayed in this tab only if assessment is enabled on the network and there are health results in the database.



Health Results

This table presents health results for all the scans performed on the end-system.

Risk

The overall risk level assigned to the end-system based on the health result of the scan:

- Red High Risk
- Orange Medium Risk
- Yellow Low Risk
- Green No Risk
- Gray Unknown

Start Scan

The date and time the scan started.

MAC Address

The end-system's MAC address.

Reason

The reason the health result was placed into the specified risk level. This is based on the risk level configuration that was used for the assessment, for example, if there was one or more health result detail with a score greater than 7. If the end-system is NAP capable, then this is based on the values returned from NAP.

Summary

A list of all the test cases that were run against the device during assessment. The test case name will be listed, or if that is not available, the test case ID will be listed.

Test Sets

The list of test sets that were run during assessment, for example, Default Nessus, Default Agent-less, and Default Agent-based. Test sets are defined as part of the assessment configuration. If the end-system is NAP capable, then this column displays Microsoft NAP indicating that NAP performed the assessment.

Total Score

The total sum of the scores for all the health details that were included as part of the quarantine decision, followed by the actual score in parenthesis. The actual score is what the total score would be if all the health details were included as part of the quarantine decision. It includes all scores, including those marked Informational and Warning. If the total score and the actual score are the same, only one score is shown.

Top Score

The highest score received for a health detail that was included as part of the quarantine decision. Scores that are marked as Informational or Warning are not considered.

IP Address

The end-system's IP address.

End Scan

The date and time the scan ended.

Server Name

The name of the assessment server. For on-board assessment servers, the name is determined by the name of the ExtremeControl engine. For example, if you create an ExtremeControl engine and name it MyAccessControlengine, then the on-board assessment server name will be listed as MyAccessControlengine as well.

Server IP

The IP address of the assessment server. For on-board assessment servers, the IP address is determined by the address of the ExtremeControl engine. For example, if you create an ExtremeControl engine with an IP address of 10.20.80.8, then the on-board assessment server IP address is listed as 10.20.80.8 as well.

Server Port

The port number on the assessment server to which the ExtremeControl engine sends assessment requests.

Host Unreachable

Displays whether the end-system was unreachable and could not be scanned: Yes or No.

Warning Count

The total number of health result details that are marked as Warnings.

Health Result Details

This table displays the individual health result details for the scan selected in the top table. Double-click any health result detail to open the Health Result Details window that displays a description, solution, and result for the health result.

Risk

The risk level assigned to the problem found on the port:

- Red High (corresponds to a Hole)
- Orange Medium (corresponds to a Warning)
- Yellow Low (corresponds to a Note)
- Black No Result Available

Name

This column lists the name of the test that is reported by the health result detail.

Test Case ID

The unique number assigned to the test case.

Score

The score assigned to the test case. The score is a value between 0.0 and 10.0. In the case of agent-based test cases, the score will be either 0.0 for a passed test, or 10.0 for a failed test, unless specifically overwritten by the scoring override configuration.

Scoring Mode

The scoring mode that was used at the time the test was performed.

- Applied The score returned by this test was included as part of the quarantine decision.
- Informational The score returned by this test was reported, but did not apply toward a quarantine decision.
- Warning The score returned by this test was only used to provide end user assessment warnings via the Notification portal web page.

CVE ID

The CVE (Common Vulnerability and Exposures) ID assigned to the security vulnerability or exposure. For more information on CVE IDs, refer to the following URL: https://cve.mitre.org/.

Description

This column lists information about the health result detail.

Solution

A solution for the problem found in the health result detail.

Port ID

The port on the end-system that the security risk was detected on.

Protocol ID

The well-known number (ID) assigned to the IP Protocol Type.

Value

What this specific test case is testing or checking for on the end-system.

Assessment Type

The type of assessment server used in the test set.

Remediation Success

For agent-based assessment, this column lists the results of remediation attempts: Remediation Successful, Remediation Failed, or Not Applicable.

Type

A "type" is assigned to each security risk found on a port during an assessment, and is used to determine whether to Quarantine an end-system. Types are configurable on the assessment agent. There are three types:

- Hole The port is vulnerable to attack.
- Warning The port may be vulnerable to attack.
- Note There may be a security risk on the port.

Buttons and Paging Toolbar

Add to Group

Lets you add the selected end-system to a specific end-system or user group. After adding an end-system to a group, any rules that have been created that involved that group will now apply to the end-system as well. Changes to end-system group membership do not require an enforce and will be synchronized with engines immediately. Changes will not affect the end-system until the next authentication or assessment occurs.

Force ReAuth

Forces the selected end-system to re-authenticate.

Lock MAC

Opens the Add MAC Lock window where you can lock the MAC address of the selected end-system to a switch or switch and port.

Edit Registration

Opens a window where you can edit the expiration time and maximum registered device count for the end user.

Refresh

Use the refresh button to update the data in the table.

Paging Toolbar

The <u>paging toolbar</u> provides four buttons that let you easily page through the table: first, previous, next, and last page.

Reset

The reset button clears the search field and search results, clears all filters, and refreshes the table.

Bookmark


Use the <u>bookmark button</u> to save the search, sort, and filtering options you have currently set.

Add/Edit Location Group

Use this window to add a new location group or edit an existing location group. Location Groups are rule components that enable you to specify network access requirements or restrictions based on the network location where the end-user is connecting. For example, in an enterprise environment, an engineer logging on to the network from the corporate cafeteria could receive different network access than an engineer logging on from the engineering development area.

You can access the Add/Edit Location Group window by accessing the **ExtremeControl** tab and selecting ExtremeControl Configurations > Group Editor > Location Groups in the left-panel menu and selecting the **Add** button in the right panel.

NOTE: Changes to rule components do not require an enforce. Changes are automatically synchronized with engines on the next status update. Changes do not affect end-systems until the next authentication and/or assessment occurs.



Name

Enter a name for a new location group. You cannot edit the name of a group.


Description

Enter a description of the location group.

Type

Select **Location** to create a Location group.

Select **Create** to display the Entry Editor section of the window. This section varies depending on the **Type** selected.



Switch

The IP address of the switches added to the location.

Port/SSID

The port or port range for a wired switch or the SSIDs for a wireless switch.

AP ID

The access point identifiers for a wireless switch.

Description

The description of the location group.

Add Button (Add...

Select the **Add** button to open the Add Entry window, from which you can add an entry to the Entry Editor section.

Edit Button 📝 Edit...

Select an entry in the Entry Editor section of the window and select the **Edit** button to open the Edit Entry window, from which you can edit an existing entry.

Delete Button Delete

Select an entry in the Entry Editor section of the window and select the **Delete** button to delete an existing entry.

Save Button


Select the **Save** button to save the location group.

Create Time Group Window

Use this window to add a new time group or edit an existing time group. Time groups are rule components that enable you to specify network access requirements or restrictions based on the day and time when the end user is accessing the network. For example, in an enterprise environment, an employee could be assigned different access privileges based on whether they log in during traditional work hours or after hours.

You can access the Add/Edit Time Group window from the Manage Rule Groups window or from the time group field in the Create Rule window.

NOTE: Changes to rule components do not require an enforce. Changes will be automatically synchronized with engines on the next status update. Changes will not affect end-systems until the next authentication and/or assessment occurs.



Name

Enter a name for a new time group. You cannot edit the name of an existing group. If you want to change the name, you must create a new time group with a new name and then delete the old time group.

Description

Enter a description of the time group. This description displays in the Manage Rule Groups window.

Calendar


Use the calendar to select the desired weekly time periods. Select to choose a specific day and time, or select and drag to quickly select a time sequence or series of days. For example, you can select Monday at 8 AM and drag down to select that hour for Monday through Friday. The select and drag feature makes it easy to select an entire week or chunk of time with just one action. Right-click on a selected

square to access menu options that let you select all or clear all squares, and undo the last action. If a square is the first or last in a series, right-click to access the Refine Time Range Start/End options that let you specify hourly increments for the start and end times.

Add/Edit User Group

Use this window to add a new user group or edit an existing user group. User groups are rule components that allow you to group together end users having similar network access requirements or restrictions. You can access the Add/Edit User Group window from the Manage Rule Groups window or from the user group field in the Create Rule window.

NOTE: Changes to rule components do not require an enforce. Changes are automatically synchronized with engines on the next status update. Changes do not affect end-systems until the next authentication and/or assessment occurs.



Name

Enter a name for a new user group. You cannot edit the name of a group.

Description

Enter a description of the user group.

Type

Select **User** to create an end-system group. Specify whether the user group is based on:

- Username a list of usernames which can be based on an exact match or a wild card.
- LDAP User Group a list imported from an LDAP Server, organized by Organization Unit (OU), or a custom attribute lookup for any user or MAC address if they match a AAA configuration entry that assigns the request a valid LDAP Configuration.
- RADIUS User Group a list of attributes returned by the RADIUS server.

Select **Create** to display the Entry Editor section of the window. This section varies depending on the **Type** selected.



Match Mode

For LDAP and RADIUS user groups, the **Match Mode** option lets you select whether to match any or match all of the LDAP or RADIUS User Group entries (attribute names) listed below.

For LDAP User Groups, you can also select **Exists**, as the username can be used to verify this criteria after the initial authentication (i.e., using Registration). The **Exists** mode is not available for RADIUS User Groups because they cannot be verified after an initial registration as the user credentials are not stored on the ExtremeControl engine for re-verification.

Attribute Name

The name of the LDAP or RADIUS Attribute.

Value


The Attribute value of the user group or username.

Add Button (Add...

Select the **Add** button to open the Add Entry window, from which you can add an entry to the Entry Editor section.

Edit Button 🔯 Edit...

Select an entry in the Entry Editor section of the window and select the **Edit** button to open the Edit Entry window, from which you can edit an existing entry.



IMPORTANT Commas are generally used to separate the attribute/value pairs in an entry to ensure they are evaluated separately. Adding a comma can impact how wildcards (*) are handled. To force the entry to be treated as a single value, do not use a comma before a second '='.

> For Example: ou=NacDev, DC=com is evaluated as two separate entries; ou=ou=NacDev, DC=com is evaluated as a single entry.

Delete Button Delete

Select an entry in the Entry Editor section of the window and select the Delete button to delete an existing entry.

Tools

Use the **Tools** menu button to either open a window where you can select a file for importing usernames (if you are creating username entries) or open a window where you can configure an LDAP OU import (if you are creating an LDAP user group).




Use the Filter functions to filter for a specific entry based on a numeric value or text.

Add/Edit User Group Window

Use this ExtremeControl window to add a new user group or edit an existing user group. User groups are rule components that allow you to group together end-users having similar network access requirements or restrictions. You can access the Add/Edit User Group window from the Group Editor or from the user group field in the Add Rule window.

NOTE: Changes to rule components do not require an enforce. Changes automatically synchronize with the engines on the next status update. Changes do not affect end-systems until the next authentication and/or assessment occurs.



Name

Enter a name for a new user group. You cannot edit the name of a group.

Description

Enter a description of the user group.

Type

Specify the criteria on which the user group is based:

- Username a list of usernames which can be based on an exact match or a wild card.
- LDAP User Group a list imported from an LDAP Server, organized by Organization Unit (OU), or a custom attribute lookup for any user or MAC address if they match a AAA configuration entry that assigns the request a valid LDAP Configuration.
- RADIUS User Group a list of attributes returned by the RADIUS server.

Match Mode

For LDAP and RADIUS user groups, the Match Mode option lets you select whether to match any or match all of the LDAP or RADIUS User Group entries (attribute names) listed below.

For LDAP User Groups, you can also select "Exists", since the username can be used to verify this criteria after the initial authentication (i.e., using Registration). The "Exists" mode is not available for RADIUS User Groups because they cannot be verified after an initial registration as the user credentials are not stored on the ExtremeControl engine for re-verification.

Username Entry Editor

Use the buttons to add, edit, or delete entries in the group. Usernames can be an exact match or use wildcards.



Use the Filter functions to filter for a specific entry based on a numeric value or text.


Switches

This tab provides information about the switches assigned to an ExtremeControl Gateway engine or ExtremeControl Engine Group. To access this tab, select a gateway or engine group in the left-panel tree, then select the **Switches** tab in the right panel.

You can right-click on one or more switch for a menu of options.

If you are using the **Policy** tab, you can also right-click on one or more switch and select from the options in the Policy menu.

Use the table options and tools to <u>filter, sort, and customize</u> table settings. You can access the options by selecting the down arrow in the right corner of any column header.



Switch IP Address

The switch's IP address.

Switch Nickname

The nickname assigned to the switch when it is added to the ExtremeCloud IQ - Site Engine database.

Switch Status

The current operational status of the switch, based on the ExtremeCloud IQ - Site Engine device poll. If the device poll did not update the status of a switch, and a Verify RADIUS Configuration operation is

performed on that switch, the switch status in the **Switches** tab can differ from the switch status in the Verify RADIUS Configuration window.

Switch System Name

The assigned name of the device as stored in the device's sysName MIB object.

Primary Gateway

The name and IP address of the switch's primary ExtremeControl Gateway. If load balancing has been configured for the engine group, the ExtremeCloud IQ - Site Engine server determines the primary and secondary gateways at Enforce, and this field displays "Determined by Load Balancer."

Secondary Gateway

The name and IP address of the switch's secondary ExtremeControl Gateway. If load balancing has been configured for the engine group, the ExtremeCloud IQ - Site Engine server determines the primary and secondary gateways at Enforce, and this field displays "Determined by Load Balancer."

Policy/VLAN

The RADIUS attributes included as part of the RADIUS response.

Policy Domain

The Policy Manager domain the switch is assigned to (if any). You can populate this field by right-clicking on a switch and selecting Policy > Verify Domain. This information does not automatically update if there are domain assignment changes. You need to re-select the menu option to update the domain information.

Auth Access Type

The type of authentication access allowed for this switch:

- Any access the switch can authenticate users originating from any access type.
- Management access the switch can only authenticate users that have requested management access via the console, Telnet, SSH, or HTTP, etc.
- Network access the switch can only authenticate users accessing the network via the following authentication types: MAC, PAP, CHAP, and 802.1X. If RADIUS accounting is enabled, then the switch also monitors Auto Tracking, CEP (Convergence End Point), and Switch Quarantine sessions.
- Monitoring RADIUS Accounting the switch monitors Auto Tracking, CEP (Convergence End Point), and Switch Quarantine sessions. ExtremeControl learns about these session via RADIUS accounting. This allows ExtremeControl to be in a listen mode, and to display access control, location information, and identity information for end-systems without enabling authentication on the switch.
- Manual RADIUS Configuration RADIUS configuration was performed manually on the switch using Policy Manager or CLI.

Switch Type

Specifies the switch type: a switch that authenticates layer 2 traffic via RADIUS to an out-of-band ExtremeControl gateway, or a VPN concentrator being used in an ExtremeControl VPN deployment.

Switch Location

The physical location of the switch.

Switch Contact

The person responsible for the switch.

Switch Description

A description of the switch, which can include its manufacturer, model number, and firmware revision number.

Management RADIUS Servers

RADIUS servers used to authenticate requests for administrative access to the switch.

RADIUS Accounting

Displays whether RADIUS accounting is enabled or disabled on the switch. RADIUS accounting can be used to determine the connection state of the end-system sessions on the ExtremeControl engine, providing real-time connection status in ExtremeCloud IQ - Site Engine. RADIUS accounting is also used to monitor switches for Auto Tracking, CEP (Convergence End Point), and Switch Quarantine authentication sessions, when used in conjunction with the Monitoring or Network Access switch authentication access types. For more information, see the Auth. Access Type section of the Add/Edit Switch Window Help topics.

IP Subnet for IP Resolution

Displays the IP subnet that the switch is using as an inclusive list for MAC to IP resolution. Specifying an IP subnet in a static IP network allows for a router to be used for IP resolution in cases where it would not be discovered via DHCP. IP Subnets also contain an IP range which can be used to filter out secondary IP addresses that are not valid for the network.

Policy Enforcement Points

If the switch is a VPN device (see Switch Type column), this column displays the Policy Enforcement Points that are being used to provide authorization for the connecting end-systems.

Add Switch

Opens the Add Switches to ExtremeControl Engine Group window where you can select switches to add to the engine or engine group.

Edit

Select a switch and select this button to open the Edit Switches in ExtremeControl Engine Group window where you can change the switch's primary and secondary ExtremeControl Gateway (Gateway), and also edit other switch attributes, if desired.


Delete

Select a switch and select this button to delete the switch from ExtremeCloud IQ - Site Engine's device database. The switch's primary gateway enforces its own primary RADIUS server as both the primary and secondary RADIUS servers on the switch.

Edit Switches in ExtremeControl Engine Group

Use this window to change a switch's primary and secondary ExtremeControl Gateway, and also edit other switch parameters including the switch's authentication access type and the RADIUS attributes to send, if desired.

You can access this window by selecting an engine or engine group in the left-panel tree. Then, in the right-panel **Switches** tab, select the switches you wish to edit and select the **Edit** button.



Switch Type

Use the drop-down list to change the type of switch:

- Layer 2 Out-Of-Band A switch that will do authentication on layer 2 traffic via RADIUS to an out-of-band ExtremeControl gateway.
- Layer 2 Out-Of-Band Data Center A switch within a data center where virtualization and mobility are a factor. If an end-system changes location but does not move to a different ExtremeControl engine, ExtremeCloud IQ Site Engine removes the end-system authentication from their prior port/switch. This allows VMs that quickly move from one server to another and then back again to still have their location updated in ExtremeCloud IQ Site Engine, because only one authenticated session is allowed per end-system within ExtremeCloud IQ Site Engine.
- Layer 2 RADIUS Only In this mode, ExtremeControl does not require any information from the switch other than the end-system MAC address (from Calling-Station-Id or User-Name). The NAS-Port does not need to be specified. If the switch supports RFC 3576, you can set the Reauthentication Behavior in the Advanced Switch Settings window. IP resolution and reauthentication occasionally do not work in this mode.
- VPN A VPN concentrator being used in an ExtremeControl VPN deployment. In this case, you should specify one or more Policy Enforcement Points below. If you do not specify a Policy Enforcement Point, then ExtremeControl is unable to apply policies to restrict access after the user is granted access.

Primary Gateway

Use the drop-down list to select the primary ExtremeControl Gateway for the selected switches. If load balancing has been configured for the switch, this field is not displayed.

Secondary Gateway

Use the drop-down list to select the secondary ExtremeControl Gateway for the selected switches. If load balancing has been configured for the switch, this field is not displayed.

Auth Access Type

Use the drop-down list to select the type of authentication access allowed for these switches. This feature allows you to have one set of switches for authenticating management access requests and a different set for authenticating network access requests.

WARNING: For ExtremeXOS/Switch Engine devices only. ExtremeControl uses CLI access to perform configuration operations on ExtremeXOS/Switch Engine devices.

- Enabling an Auth type of "Any Access" or "Management Access" can restrict access to the switch after an enforce is performed. For management requests handled through ExtremeControl, make sure that an appropriate administrative access configuration is in place by assigning a profile such as "Administrator ExtremeControl Profile" to grant proper access to users. Also, verify that the current switch CLI credentials for the admin user are defined in the database against which ExtremeControl authenticates management login attempts.
- Switching from an Auth type of "Any Access" or "Management Access" back to
 "Network Access" can restrict access to the switch after an enforce is performed.
 Verify that the current switch CLI credentials for the admin user are defined locally
 on the switch.
- Any Access the switch can authenticate users originating from any access type.
- Management Access the switch can only authenticate users that have requested management access via the console, Telnet, SSH, or HTTP, etc.
- Network Access the switch can only authenticate users accessing the network via the following
 authentication types: MAC, PAP, CHAP, and 802.1X. If RADIUS accounting is enabled, then the
 switch also monitors Auto Tracking, CEP (Convergence End Point), and Switch Quarantine
 sessions. If there are multiple sessions for a single end-system, the session with the highest
 precedence will be displayed to provide the most accurate access control information for the
 user. The ExtremeControl authentication type precedence from highest to lowest is: Switch
 Quarantine, 802.1X, CHAP, PAP, Kerberos, MAC, CEP, RADIUS Snooping, Auto Tracking.
- Monitoring RADIUS Accounting the switch will monitor Auto Tracking, CEP (Convergence End Point), and Switch Quarantine sessions. ExtremeCloud IQ Site Engine learns about these session via RADIUS accounting. This allows ExtremeCloud IQ Site Engine to be in a listen mode, and to display access control, location information, and identity information for end-systems without enabling authentication on the switch. If there are multiple sessions for a single end-system, the session with the highest precedence displays to provide the most accurate access control information for the user. The ExtremeControl authentication type precedence from highest to lowest is: Switch Quarantine, 802.1X, CHAP, PAP, Kerberos, MAC, CEP, RADIUS Snooping, Auto Tracking.

• Manual RADIUS Configuration — ExtremeCloud IQ - Site Engine does not perform any RADIUS configurations on the switch. Select this option if you want to configure the switch manually using the **Policy** tab or CLI.

Virtual Router Name

Select the checkbox to enter the name of the Virtual Router. The default value for this field is **VR-**

WARNING: For ExtremeXOS/Switch Engine devices only. If ExtremeCloud IQ - Site Engine has not detected and populated this field, enter the Virtual Router Name carefully. Incorrectly entering a value in this field causes the RADIUS configuration to fail, which is not reported when enforcing the configuration to the switch.

Gateway RADIUS Attributes to Send

Use the drop-down list to select the RADIUS attributes settings included as part of the RADIUS response from the ExtremeControl engine to the switch.

RADIUS Accounting

Use the drop-down list to enable RADIUS accounting on the switch. RADIUS accounting can be used to determine the connection state of the end-system sessions on the ExtremeControl engine, providing real-time connection status in ExtremeCloud IQ - Site Engine. It also allows ExtremeControl to monitor Auto Tracking, CEP (Convergence End Point), and Quarantine (anti-spoofing) sessions.

Management RADIUS Server

Use the drop-down list to specify RADIUS servers used to authenticate requests for administrative access to the selected switches. Select from the RADIUS servers you have configured in ExtremeCloud IQ - Site Engine, or select **New** or **Manage** to open the Add/Edit RADIUS Server or Manage RADIUS Servers windows.

Network RADIUS Server

This option lets you specify a backup RADIUS server to use for network authentication requests for the selected switches. This allows you to explicitly configure a network RADIUS server to use if there is only one ExtremeControl engine. (This option is only available if a Secondary Gateway is not specified.) Select from the RADIUS servers you have configured in Extreme Control, or select **New** or **Manage** to open the Add/Edit RADIUS Server or Manage RADIUS Servers windows.

Policy Domain

Use this option to assign the switch to a **Policy** tab domain and enforce the domain configuration to the switch. The switch must be an Extreme Networks switch.

NOTE: Selecting -- Do Not Set -- for an ExtremeControl engine on which a Policy Domain is configured does not unassign the Policy Domain. To unassign a Policy Domain, use the **Policy** tab.

Advanced Settings


Select this button to open the Advanced Switch Settings window.

Add Switches to ExtremeControl Engine Group

Use this window to add switches to a gateway engine or engine group. The window allows you to select one or more switches from the device tree, and set the primary and secondary ExtremeControl Gateways for the switches. It also lets you set other parameters including the authentication access type for the switches and the RADIUS attributes to send.

NOTE: If desired, you can set only the primary ExtremeControl Gateway for the switches; ExtremeCloud IQ - Site Engine does not require the secondary ExtremeControl Gateway to be set. If only the primary ExtremeControl Gateway is set, then by default that gateway uses its primary proxy RADIUS server as a secondary direct RADIUS server to the switch. This allows for redundancy without the requirement for a secondary ExtremeControl Gateway. In this scenario, if contact with the ExtremeControl Gateway fails, authentication traffic would bypass the ExtremeControl gateway, but normal authentication would continue in the network, and still provide some security.

You can access this window by selecting an engine or engine group and selecting the **Add Switch** button in the right-panel **Switches** tab.



Device Tree

This area displays the device tree. Expand the tree and select the switches you want to add to the engine or engine group.

Add Device

Opens the Add Device window where you can add a device to the ExtremeCloud IQ - Site Engine database. The device is displayed in the My Network folder in the device tree.

Switch Type

Use the drop-down list to select the type of switch you are adding:

- Layer 2 Out-Of-Band A switch that authenticates on layer 2 traffic via RADIUS to an out-of-band ExtremeControl gateway.
- Layer 2 Out-Of-Band Data Center A switch within a data center where virtualization and mobility are a factor. If an end-system changes location but does not move to a different ExtremeControl engine, ExtremeControl removes the end-system authentication from their prior port/switch. This allows VMs that quickly move from one server to another and then back again to still have their location updated in ExtremeCloud IQ Site Engine, because only one authenticated session is allowed per end-system in ExtremeCloud IQ Site Engine.
- Layer 2 RADIUS Only In this mode, ExtremeCloud IQ Site Engine does not require any information from the switch other than the end-system MAC address (from Calling-Station-Id or User-Name). The NAS-Port does not need to be specified. If the switch supports RFC 3576, you can set the Reauthentication Behavior in the Advanced Switch Settings window. IP resolution and reauthentication might not work in this mode.
- VPN A VPN concentrator being used in an ExtremeControl VPN deployment. In this case, you should specify one or more Policy Enforcement Points below. If you do not specify a Policy Enforcement Point, then ExtremeCloud IQ Site Engine is unable to apply policies to restrict access after the user is granted access.

Primary Gateway

Use the drop-down list to select the primary ExtremeControl Gateway for the selected switches. If load balancing has been configured for the engine group, the ExtremeCloud IQ - Site Engine server determines the primary and secondary gateways at Enforce, and this field displays **Determined by Load Balancer**.

Secondary Gateway

Use the drop-down list to select the secondary ExtremeControl Gateway for the selected switches. If load balancing has been configured for the engine group, the ExtremeCloud IQ - Site Engine server determines the primary and secondary gateways at Enforce, and this field displays **Determined by Load Balancer**.

NOTE: To configure additional redundant ExtremeControl Gateways per switch (up to four), use the Display Counts option in the Display options panel (Administration > Options > ExtremeControl).

Auth. Access Type

Use the drop-down list to select the type of authentication access allowed for these switches. This feature allows you to have one set of switches for authenticating management access requests and a different set for authenticating network access requests.

WARNING: For ExtremeXOS/Switch Engine devices only. ExtremeControl uses CLI access to perform configuration operations on ExtremeXOS/Switch Engine devices.

- Enabling an Auth type of "Any Access" or "Management Access" can restrict access
 to the switch after an enforce is performed. Make sure that an appropriate
 administrative access configuration is in place by assigning a profile such as
 "Administrator ExtremeControl Profile" to grant proper access to users. Also, verify
 that the current switch CLI credentials for the admin user are defined in the
 database that ExtremeCloud IQ Site Engine authenticates management login
 attempts against.
- Switching from an Auth type of "Any Access" or "Management Access" back to
 "Network Access" can restrict access to the switch after an enforce is performed.
 Verify that the current switch CLI credentials for the admin user are defined locally
 on the switch.
- Any Access the switch can authenticate users originating from any access type.
- Management Access the switch can only authenticate users that have requested management access via the console, Telnet, SSH, or HTTP, etc.
- Network Access the switch can only authenticate users that are accessing the network via the
 following authentication types: MAC, PAP, CHAP, and 802.1X. If RADIUS accounting is enabled,
 then the switch also monitors Auto Tracking, CEP (Convergence End Point), and Switch
 Quarantine sessions. If there are multiple sessions for a single end-system, the session with the
 highest precedence displays to provide the most accurate access control information for the
 user. The ExtremeControl authentication type precedence from highest to lowest is: Switch
 Quarantine, 802.1X, CHAP, PAP, Kerberos, MAC, CEP, RADIUS Snooping, Auto Tracking.
- Monitoring RADIUS Accounting the switch monitors Auto Tracking, CEP (Convergence End Point), and Switch Quarantine sessions. ExtremeCloud IQ Site Engine learns about these session via RADIUS accounting. This allows ExtremeCloud IQ Site Engine to be in a listen mode, and to display access control, location information, and identity information for end-systems without enabling authentication on the switch. If there are multiple sessions for a single end-system, the session with the highest precedence displays to provide the most accurate access control information for the user. The ExtremeControl authentication type precedence from highest to lowest is: Switch Quarantine, 802.1X, CHAP, PAP, Kerberos, MAC, CEP, RADIUS Snooping, Auto Tracking.
- Manual RADIUS Configuration ExtremeCloud IQ Site Engine does not perform any RADIUS configurations on the switch. Select this option if you want to configure the switch manually using the Policy tab or CLI.

Virtual Router Name

Enter the name of the Virtual Router. The default value for this field is VR-Default.

WARNING: For ExtremeXOS/Switch Engine devices only. If ExtremeCloud IQ - Site Engine has not detected and populated this field, enter the Virtual Router Name carefully. Incorrectly entering a value in this field causes the RADIUS configuration to fail, which is not reported when enforcing the configuration to the switch.

Gateway RADIUS Attributes to Send

Use the drop-down list to select the RADIUS attributes included as part of the RADIUS response from the ExtremeControl engine to the switch. You can also select **New** or **Manage** from the menu to open the RADIUS Attribute Settings window where you can define, edit, or delete the available attributes.

RADIUS Accounting

Use the drop-down list to enable RADIUS accounting on the switch. RADIUS accounting can be used to determine the connection state of the end-system sessions on the ExtremeControl engine, providing real-time connection status in ExtremeCloud IQ - Site Engine.

Management RADIUS Server 1 and 2

Use the drop-down list to specify RADIUS servers used to authenticate requests for administrative access to the selected switches. Select from the RADIUS servers you have configured in ExtremeCloud IQ - Site Engine, or select New or Manage RADIUS Servers to open the Add/Edit RADIUS Server or Manage RADIUS Servers windows.

Network RADIUS Server

This option lets you specify a backup RADIUS server to use for network authentication requests for the selected switches. This allows you to explicitly configure a network RADIUS server to use if there is only one ExtremeControl engine. (This option is only available if a Secondary Gateway is not specified.) Select from the RADIUS servers you have configured in ExtremeCloud IQ - Site Engine, or select New or Manage RADIUS Servers to open the Add/Edit RADIUS Server or Manage RADIUS Servers windows.

Policy Enforcement Point 1 and 2

Select the Policy Enforcement Points used to provide authorization for the end-systems connecting to the VPN device you are adding. The list is populated from the N-Series, S-Series, and K-Series devices in your Console device tree. If you do not specify a Policy Enforcement Point, then ExtremeControl is unable to apply policies to restrict end user access after the user is granted access.

Policy Domain

Use this option to assign the switch to a policy domain and enforce the domain configuration to the switch. The switch must be an Extreme Networks switch.


Advanced Settings

Select the Advanced Settings button to open the Advanced Switch Settings window.

Advanced Switch Settings

This window allows you to configure settings for switches that require a different configuration than your standard switch settings set in the Engine Settings window.

You can access the window from the <u>Add Switch to ExtremeControl Engine Group window</u> or from the <u>Edit Switches in ExtremeControl Engine Group window</u>.



IP Subnet for IP Resolution

Select the drop-down list to display a list of the IP subnets configured in the Engine Settings window. If you select a subnet, the switch uses it as an inclusive list for MAC to IP resolution. Specifying an IP subnet in a static IP network allows for a router to be used for IP resolution in cases where it would not be discovered via DHCP. IP subnets also contain an IP range which can be used to filter out secondary IP addresses that are not valid for the network.

Shared Secret

A string of alpha-numeric characters used to encrypt and decrypt communications between the switch and the ExtremeControl engine. The shared secret is shown as a string of asterisks. When the Show Password option is selected, the shared secret is shown in text.

Reauthentication Type

Select the reauthentication type for the switch:

- SNMP uses SNMP to trigger reauthentication using various OIDs in different MIBs. The ExtremeControl engine checks a series of proprietary Enterasys MIBs, standardized MIBs, and proprietary third-party MIBs to determine availability, and forces reauthentication using any available SNMP method.
- Session Timeout causes ExtremeControl to return a session timeout and terminate action to the
 end-system via RADIUS response attributes. The use of this mechanism causes the user to be
 automatically reauthenticated at a specified interval by the switch to which they are connected.
 Only use this option for wireless switches that do not have RFC 3576 support or wired switches
 that do not have SNMP support.
- RFC 3576 a method of reauthenticating RADIUS sessions through the use of Disconnect-Request messages as defined by RFC 3576. (For more information, see
 http://www.ietf.org/rfc/rfc3576.txt). RFC 3576 configurations must be customized to work with the specific vendor implementation for each device type. To add, edit, or delete an RFC 3576 configuration, select the Manage RFC 3576 Configurations button.

Enable Port Link Control


Port link control allows the toggle of the operational mode of a port. Select this option to enable port link control for specific switches.

All Access Control Engines

The All ExtremeControl Engines tab is displayed in the right panel when you select the All ExtremeControl Engine tree in the left panel or when you select the ExtremeControl Engines tab when an ExtremeControl Engine Group is selected. The panel displays a table of information about the engines in the folder or group. Right-click an engine for a menu of options.

Use the table options and tools to filter, sort, and customize table settings. You can access the options by selecting the down arrow in the right corner of any column header.

NOTE: The ExtremeControl Engine administration web page allows you to access status and diagnostic information for an ExtremeControl engine. Access the administration web page using the following URL: https://ExtremeControlEnginelP:8444/Admin. The default user name and password for access to this web page is "admin/Extreme@pp."



Name

The name of the ExtremeControl engine (assigned when the engine is created).

IP Address

The ExtremeControl engine's IP address.

Engine Type

The ExtremeControl engine type: ExtremeControl Gateway, ExtremeControl Layer 2 (L2) Controller, or ExtremeControl Layer 3 (L3) Controller.

Primary Count

The number of switches for which the ExtremeControl engine is the primary engine.

Secondary Count

The number of switches for which the ExtremeControl engine is the secondary engine.

Model

The ExtremeControl engine's model number.

Version

The ExtremeControl engine's version number.

CPU Load (0-100%)

The percentage of the engine's CPU currently being used. This value gives you an indication of how busy the engine is and helps you determine if your network needs additional engines, or if you need to change your network configuration so that the load is more evenly distributed among your existing engines.

Memory Used

The amount of memory used by the engine.

Memory Available

The amount of memory available on the engine.

Connected Agents


The number of assessment agents connected to the engine.

Capacity

The engine's current capacity, which is the number of end-systems that have authenticated within the last 24 hours out of the maximum number of authenticating end-systems supported for the engine.

Access Control Engine Enforce Preview

Use this window to preview what you are changing on an ExtremeControl engine by performing an enforce. You can access the Access Control Engine Enforce Preview window by right-clicking an engine in the Engines list on the Access Control tab and selecting **Enforce Preview** from the menu.




The window displays details of the changes you are making on the ExtremeControl engine.

Select the **Export** button to export the results to a text file.

Details (ExtremeControl Engine)

This tab provides information about an ExtremeControl engine's configuration. The information changes depending on the type of engine selected in the left-panel tree.

To access this tab, select an ExtremeControl engine in the left-panel tree, then select the **Details** tab in the right panel.



General Information

This section displays general information about the ExtremeControl engine, including its name, IP address, type (ExtremeControl Gateway or Layer 2/Layer 3 ExtremeControl Controller), the engine version, the IP address of the ExtremeCloud IQ - Site Engine Management server, and the ExtremeControl engine status.

End-System Capacity

This field lists the engine's current capacity, which is the number of end-systems that authenticated within the last 24 hours out of the maximum number of authenticating end-systems supported for the engine.

ExtremeControl Configuration

Displays the ExtremeControl Configuration assigned to the engine. The ExtremeControl Configuration determines the ExtremeControl Profile assigned to an end-system connecting to the network.

Engine Settings

The engine settings configuration being used by your ExtremeControl engine. Engine settings are configurable in the Engine Settings window by selecting the Engine Settings button.

Certificates

Select Manage to update the ExtremeControl certificates in the Manage Certificates window.

Interface Summary

Displays a summary of the current engine interface configuration.

Select **Edit** to open the **Interfaces** window, where you can change the engine Host Name and Gateway.. Select **Static Routes** to open the **Static Routes** window, where you can add or edit the static routes used for advanced routing configuration..

ExtremeControl Bypass Configuration

The ExtremeControl Bypass Configuration feature allows you to bypass ExtremeControl processing of authentication requests from end-systems connecting to the network and also disable the ExtremeControl assessment process. For ExtremeControlauthentication bypass, ExtremeControl either configures the switch to authenticate directly to a RADIUS server to which ExtremeControl is configured to proxy authentication requests, or it disables RADIUS authentication on the switch. This capability is useful for troubleshooting purposes. For example, if there is a problem with an ExtremeControl Configuration, the **Disable** button lets you remotely disable ExtremeControl functionality until the problem is resolved. You can then use the **Enable** button to re-enable ExtremeControl functionality on the engines. When ExtremeControl authentication or assessment is disabled, the ExtremeControl engine name and IP address display in red text in the left-panel tree indicating the engine is in Bypass mode.

For ExtremeControl Gateway engines, when you select the option to disable ExtremeControl authentication processing, if proxy RADIUS servers are configured for authentication in a Basic AAA Configuration, the ExtremeControl Engine configures the switches to send RADIUS packets directly to the primary and secondary RADIUS servers (from the Basic AAA Configuration), instead of talking to the RADIUS proxy through the ExtremeControl gateway. RADIUS authentication is not disabled on the switch, and end users still need to authenticate in order to connect to the network. The switches must be defined in the back-end proxy RADIUS server as RADIUS clients with the same shared secret used by the ExtremeControl Gateway engines. If there are no proxy RADIUS servers configured in a Basic AAA Configuration, or if an Advanced AAA Configuration is used, RADIUS authentication on the switch is disabled when ExtremeControl authentication processing is disabled.

NOTES: If you have disabled ExtremeControl authentication processing and then enforce with new switches, the new switches are configured to send RADIUS packets directly to the primary and secondary RADIUS servers. These switches are reconfigured to talk to the RADIUS proxy when you enable ExtremeControl; a second enforce is not necessary.

Bypass is not an option for switches set to Manual RADIUS Configuration or ExtremeWireless controllers not configured for RADIUS strict mode.

For ExtremeControl Controller engines, when you disable ExtremeControl authentication, then the ExtremeControl Controller does **not** send RADIUS packets directly to the RADIUS servers. Authentication **is** disabled on the ExtremeControl Controller and end-systems do not need to authenticate to the network. Traffic from the end-systems bypass the ExtremeControl Controller and go directly onto the network.


The **Status** fields provide the current status of the ExtremeControl authentication or assessment process. The authentication status field also includes a link to the Verify

RADIUS Configuration on Switches feature. This feature is available for ExtremeControl Gateway engines and Layer 2 ExtremeControl Controllers, and can be used to alert you to any RADIUS configurations that are out of sync and could cause RADIUS authentication problems on the network.

Details (ExtremeControl Engine Groups)

This tab provides information about the ExtremeControl Details being used by your ExtremeControl engines.

To access this tab, select an engine group from within the Engine Group tree in the left-panel tree, then select the **Details** tab in the right panel.



Status

Status

Displays status of engines in the engine group.

Group

Policy Domain

Displays the policy domain for the ExtremeControl engines in the folder. Select the **Edit Policy Domain** button to select a new policy domain for the engine group.

RADIUS Monitor Clients

Displays whether RADIUS Monitor Clients are enabled for the ExtremeControl engines in the folder. RADIUS monitoring tools monitor ExtremeControl engine performance and availability.

Select the **Edit RADIUS Monitor Clients** button to open the **Configure RADIUS Monitor Clients** window, from which you can select a new client, change the monitoring client, and delete a client.

Distributed End-System Cache

Displays whether the Distributed End-System Cache option is enabled for the ExtremeControl engines in the folder

Engines

Engine Settings

The engine settings configuration being used by your ExtremeControl engines. Engine settings are configurable in the Engine Settings window by selecting the Engine Settings button.

Engine Count

The number of engines in the engine group.

Access Control Configuration - Default

Access Control Configuration

The name of the ExtremeControl Configuration being used by your ExtremeControl engines. The ExtremeControl Configuration determines the ExtremeControl Profile assigned to an end-system connecting to the network.

Default Profile

The name of the Default Profile specified in the ExtremeControl Configuration. The Default Profile serves as a "catch-all" profile for any end-system that doesn't match one of the rules listed in the ExtremeControl Configuration.

Registration

Whether a registration/web access feature is enabled or disabled for the ExtremeControl Configuration.

Assessment/Remediation

Whether the assessment/remediation feature is enabled or disabled for the ExtremeControl Configuration.

Portal Configuration

The name of the Portal Configuration specified in the ExtremeControl Configuration. If your network is implementing Registration or Assisted Remediation, the Portal Configuration defines the branding and behavior of the website used by the end user during the registration or remediation process.

AAA Configuration

The name of the AAA Configuration specified in the ExtremeControl Configuration.

Load Balancing

Edit Internal Load Balancing

The Load Balancing panel displays the status of ExtremeXOS/Switch Engine/EOS firmware. By default, ExtremeXOS/Switch Engine/EOS Firmware status will be set to Manual Configuration.


Select the Edit Internal Load Balancing button to open the Internal Load Balancer Window, which includes the following configuration types. The default configuration is Manual Configuration.

- Manual Configuration
- Standard
- Round Robin
- Sticky Round Robin

Edit External Load Balancing

The Load Balancing panel displays the status of External Load Balancer(s). By default, the External Load Balancer status is Disabled.

Select the **Edit External Load Balancing** button to open the External Load Balancer Window, where you can add, edit, delete or reorder Load Balancer IP addresses. If you delete all the load balancer addresses, a confirmation message window displays after the last IP address is deleted:



Select Yes to continue.

Guest and IoT Configuration

Domain

This section allows you to configure a <u>GIM domain</u>, which contains all of the Guest and IoT configuration information. GIM domains are created in ExtremeCloud IQ - Site Engine and the configuration within that domain is configured in GIM.


Local Password Repository

This section allows you to configure a password repository in the **Local Password Repository** for GIM. You can also customize the sponsor retrieval in your GIM Domain and choose a different LDAP configuration and Search Root specifically for sponsor look-ups.

Interfaces Window

Use this ExtremeCloud IQ - Site Engine window to configure the interfaces on an ExtremeControl engine. Interface configuration enables you to separate management traffic from end-system traffic, providing another layer of protection for sensitive data. It also provides the ability to snoop mirrored traffic on other ports.

This window is accessed from the **Control** > **ExtremeControl** tab by selecting an ExtremeControl engine, opening the **Details** tab, and selecting the **Edit** button in the Interface Summary section.



Interface Modes

There are five different modes that can be configured for an interface: Management, Registration & Remediation, Management Only, Registration & Remediation Only, Listening Only, Advanced Configuration, and Off. The mode determines the type of traffic permitted on the interface and the <u>services</u> provided by the interface.

You can configure all the interfaces on an engine; however, you cannot change the management interface and you are only permitted to configure one interface to enable management traffic.

Management, Registration & Remediation - This mode is the in-band management mode where both management traffic and registration, assessment, and remediation traffic use the same interface. In this mode, the engine does not limit traffic to each of the services.

Management Only - In this mode, the engine binds all management services to this interface. This includes:

- traffic to ExtremeCloud IQ Site Engine and other engines (JMS and HTTP)
- all traffic to switches
- all LDAP and RADIUS traffic
- traffic for the following services: SSH daemon, SNMP daemon, and RADIUS server
- traffic for captive portal administration, sponsorship, pre-registration, and screen preview (on ports 80 and 443)
- traffic for WebView pages and ExtremeCloud IQ Site Engine web services (on ports 8080 and 8443)

Registration & Remediation Only - In this mode, the engine binds all registration and remediation services to this interface. All traffic to end-systems is initiated through this interface, including:

- assessment traffic
- NetBIOS for IP and hostname resolution
- traffic for registration pages, remediation pages, and self-registration (on ports 80 and 443)
- all agent communication traffic (on ports 8080 and 8443)

Listen Only - In this mode, the engine enables DHCP and Kerberos snooping to be performed on the interface. No IP address or hostname can be assigned to the interface.


Advanced Configuration - This mode enables you to configure the services that are provided by the selected interface, using the link in the <u>Services</u> field. This is useful for ExtremeControl deployments in MSP or MSSP environments.

Off - The interface is disabled and not used in any way.

Services

The Services field displays the services that are provided by the ExtremeControl engine interface, as determined by the selected interface mode. Each mode provides a different set of services on the interface.

If the mode is set to Advanced Configuration, the services list becomes a link that launches an Edit window where you can select or deselect the services provided by the interface. This granularity is useful for ExtremeControl deployments in MSP or MSSP environments.



NOTE:

Only one interface can have **End-System** enabled when using the OAUTH2 social login. The End-System service is part of the Management, Registration, and Remediation mode, so it can also be enabled in Advanced Configuration.

The following list describes the various services that are provided by the different modes:

- Management The communication to and from the ExtremeCloud IQ Site Engine server. Sub-services include JMS. Web Services. and Syslog.
 - **NOTE:** The Management service cannot be moved from eth0.
- Monitoring Services The services used to monitor or contact an engine. Sub-services include the SSH daemon and SNMP agent.
- **Network Services** The communication to external servers that provide networking services. Subservices include DNS servers and NTP servers.
 - **NOTE:** The Network Services service can only be applied to one interface.
- AAA Servers The communication used by external servers for authentication and authorization. Subservices include RADIUS servers and LDAP servers.
 - **NOTE:** The AAA Servers service can only be applied to one interface.
- **Device** The communication to and from a NAS (switch, router, VPN, or wireless controller). Subservices include SNMP, RADIUS, RFC3576, SSH/Telnet, and TFTP.
- Portal: Management the captive portal registration management services for an engine.
- End-System The communication to and from end-systems. Sub-services include portal registration and remediation, assessment, NetBIOS, and DNS proxy.
- Traffic Snooping DHCP and Kerberos snooping on the interface. This service is listed if the DHCP/Kerberos Snooping option is set to Enabled.

DHCP/Kerberos Snooping

Use the DHCP/Kerberos Snooping option to enable or disable DHCP and Kerberos snooping on the interface. DHCP snooping is used for IP resolution and OS detection. Kerberos snooping is used for user name detection and elevated access.

Captive Portal HTTP Mirroring

This is an advanced option that enables the interface to accept mirrored HTTP traffic which is used to display the captive portal to end users. This option is an alternative to using Policy-Based Routing and DNS Proxy.


Tagged VLANs

If the mirrored traffic includes an 802.1Q VLAN tag, then the list of VLANs to capture must be explicitly stated in this field by entering a comma-separated list of VLAN IDs from 1 to 4094. If the mirrored traffic is not tagged then this field can be left blank.

Static Route Configuration Window

This window displays the static routes used for advanced routing configuration. Use the toolbar buttons to add, edit, or delete a route.

This window is accessed from the **Control** > **ExtremeControl** tab by selecting an ExtremeControl engine, opening the **Details** tab, and selecting the **Static Routes** button in the Interface Summary section.



Interface

The ExtremeControl engine interface used for the static route.

Destination

The IP address used to define the subnet or individual device whose traffic is assigned to the route.

Gateway IP

The IP address of the device where traffic matching the Network value is sent.

Metric

A number used to configure route precedence. The lower the number, the higher the precedence.

How To Use Access Control

The **How To** section contains Help topics that give you instructions for performing tasks in the **Access Control** tab.

How to Use Device Type Profiling

This Help topic describes how to set up device type profiling in your ExtremeControl Configuration using device type rule groups. Device type profiling lets you assign ExtremeControl profiles to end-systems based on operating system family, operating system, or hardware type. This allows you to use the end-system's device type to determine the end user's level of network access control and whether the end-system is scanned. For more information on device type groups, see the Add/Edit Device Type Group Window Help topic.

NOTE: Assessment provides the most accurate determination of device type. If the initial device type determination is not based on assessment results, it can be less reliable. For that reason, device type rule groups should be based on broad families of device types.

Here are some examples of how device type profiling can be used to determine network access:

- When an end user with valid credentials logs in to the network on a registered iPad versus a registered Windows 10 machine, they receive a lower level of network access.
- When an end user registers a Windows machine using its MAC address, another user cannot spoof that MAC address using a Linux system. (Device profiling does not resolve this issue in environments with dual boot machines.)
- If an end user exports a certificate from a corporate PC to an iPad and successfully authenticates with 802.1x, the iPad is not allowed full network access.

Device Profiling Use Case


This section provides high-level instructions for configuring device type profiling for a sample use case. In this scenario, the network administrator has the following network access requirements:

- All Windows registered devices should be assigned the "Default ExtremeControl Profile."
- All Windows 10 registered devices should be assigned the "Windows10 Profile."
- All Linux registered devices should be assigned the "Default ExtremeControl Profile." In addition, a new Linux version called SuperLinux needs to be added to the Linux family device type.
- All HP Printers should be assigned the "HP Printer Profile."


To do this, create four rules in your ExtremeControl configuration that use device type as criteria for matching rules to end-systems authenticating to the network. The following instructions assume that you already created your profiles: Basic Profile, Windows10 Profile, and HP Printer Profile.

Expand the Default left-panel tree (Control > ExtremeControl > ExtremeControl Configurations > Default).

2. Select the Rules left-panel option and select the Add button in the right panel.




3. Create a rule that assigns the Default ExtremeControl Profile to all Registered Guests using Windows devices as shown below.




- 4. Create a rule that assigns the Windows10 Profile to all Windows 10 registered devices. To do this, you need to create a new Windows 10 device type group.
 - a. From the ExtremeControl Configurations left-panel tree, expand the Group Editor tree.


b. Select Device Type Groups and select the **Add** button in the right panel.




c. Create a new device type group with the name Windows 10.




d. Select **Create**. The Device Type Entry Editor displays.




e. Select the **Add** button. The Add Entry window displays.



f. Select the **Select from Existing Types** button and in the Select Device Types window, select Windows 10.




- g. Select the Add Selected button.
- h. Select the **Save & Close** button on the Add New Group window.
- i. You can then create the rule.
- j. Select the ExtremeControl Configurations > Default > Rules left-panel option and select the **Add** button in the right panel.
- k. In the Profile drop-down list, select **New**. The Create New Profile window displays.




I. Enter the name Windows10 in the Name field and select the Create button.

The ExtremeControl Profile window opens.


- m. Select Save.
- n. Configure the rule as shown in the screenshot below.



- o. Select Save.
- 5. Create a rule that assigns the Default ExtremeControl Profile to all Linux registered devices and add the SuperLinux version to the Linux family device type. To do this, you need to create a new Linux device type group that includes SuperLinux.
 - a. Create the My Linux device type group to include the devices in the Linux device type group using the **Select from Existing Types** button in the Add Entry window as discussed in step 4f above.



b. Select the **Add** button and in the Add Entry window, create the **SuperLinux** Device Type as shown below.




- c. Select **Add** to save the SuperLinux device type to the My Linux device type group.
- d. Select the **Save & Close** button on the Add New Group window.
- 6. Create a rule that assigns the HP Printer Profile to all HP printers on the network. To do this, create a new HP Printers device type group.


a. Open the Add New Group window by selecting the **Add** button on the ExtremeControl Configurations > Group Editor > Device Type Groups panel.



- b. Select Create. The Device Type Entry Editor section displays.
- c. Add the HP Printers via the Add Entry window by selecting the Add button as shown below.



- d. Select **Save & Close** to save the HP Printers group.
- e. Select Rules in the left-panel tree (ExtremeControl Configurations > Default > Rules).
- f. Select **Add** in the right-panel to open the Add Rule window.
- g. Select the New option in the Profile drop-down list and create the HP Printer Profile.



h. Create the HP Printers rule using the following criteria.


- i. Select Save.
- 7. Your ExtremeControl Configuration now contains the following rules used to determine network access and assessment requirements based on device type.

How to Configure LDAP for End Users and Hosts via Active Directory


This Help topic provides instructions for creating LDAP configurations in Access Control that provide authentication and authorization for network end users and host machines via Active Directory.

In Access Control, you can create an Advanced AAA configuration that contains one mapping rule for your host machines and two mapping rules for your users. These mappings are the same except for their LDAP configuration. You need to create two LDAP configurations: one for the hosts mapping and one for the users mapping. The LDAP configurations are identical except for the User Search Attribute. When you have completed these instructions, Access Control uses the new AAA configuration to authenticate both end users and host machines via your Active Directory server.


- 1. Select Control > Access Control > Configuration tab.
- 2. In the left-panel tree, select the **AAA** tab to open the AAA Configuration window to the right.




- 3. Select the **Add** button in the AAA Configuration panel create a new AAA Configuration.
- 4. Select LDAP Configuration in the left-panel tree to open the LDAP Configuration window.
- 5. Create an LDAP configuration for use with end users that authenticate to the network using the sample below as a guide. Select **Save**.




6. Open the Add LDAP Configuration window to add another LDAP configuration that will be used for host machines that authenticate to the network using the sample below as a guide. Note that the only difference between the two LDAP configurations is the User Search Attribute. Select **Save**.




- 7. In the left-panel tree, select an AAA Configuration to open the Advanced AAA Configuration window.
- 8. In the Authentication Rules panel of the Advanced AAA Configuration window, select the **Add** button to open the Add User to Authentication Mapping window.
- 9. Create your first mapping rule to capture machine authentications using the sample below as a guide. In the example below, host/*.nac2003.com captures the machine authentications for the NAC2003 active directory domain. Be sure to select the host LDAP Configuration you create. Select OK.



10. Create your second mapping rule to capture end user authentications using the sample below as a guide. In the example below, *@nac2003.com captures all users logging in to the NAC2003 active directory domain when they authenticate with their username in the format <username>@<domain>. Be sure to select the end user LDAP Configuration you create. Select OK.



11. Create your third mapping rule to capture other end user authentications using the sample below as a guide. In the example below, NAC2003* captures all users logging in to the NAC2003 active directory domain when they authenticate with their username in the format <domain>\cusername>. Be sure to select the end user LDAP Configuration you create. Select OK.



12. In the left-panel tree, select an AAA Configuration to open the Advanced AAA Configuration window. Use the **Up** and **Down** buttons to move your new mappings above the "Any" mappings in the list of mappings. Select **Save**.

You can configure your LDAP policy mappings and/or LDAP user groups based on the attributes from either your host or user LDAP configurations.

How to Change the Assessment Agent Adapter Password

This Help topic provides instructions for changing the password on the assessment agent adapter on your network assessment servers, including agent-less, Nessus, or a third-party assessment agent (an assessment agent not supplied or supported by ExtremeCloud IQ - Site Engine). The assessment agent adapter enables communication between the ExtremeControl engine and the assessment servers, and the password is used by the assessment agent adapter to authenticate ExtremeControl engine assessment requests.

This password must match the password specified in the ExtremeControl Options as the Assessment Agent Adapter Credentials (Administration > Options > Identity and Access > Assessment Server). If you change the password on the assessment agent adapter, change assessment agent adapter credentials in the ExtremeControl options as well, or connection between the engine and assessment servers is lost and assessments is not performed.

To change the assessment agent adapter password:

- 1. Go to the install directory for the assessment agent adapter on the assessment server. This can be a Nessus server or the ExtremeControl engine if you are using on-board agent-less assessment. On an ExtremeControl engine, the install directory is /opt/nac/saint.
- 2. Run the **sha1.sh** script (on an ExtremeControl engine, the script is located in **/opt/nac/saint/util**) using the new password as the argument. The script produces a hash string that looks something like:
 - 9ba2db465ff11b0bdfd188f7ee87b10fc3a145dc
- 3. Open the users.properties file (on an ExtremeControl engine, the file is located in /opt/nac/saint/users.properties) and replace the existing hash string with the new one: admin=<new string>
- 4. Restart the assessment agent adapter. On an ExtremeControl engine, the command is aglsctl restart.

How to Set ExtremeControl Options

Use the Options window (Administration > Options) to set options for ExtremeControl. In the Options window, the right-panel view changes depending on what you have selected in the left-panel tree. Expand the ExtremeControl folder in the tree to view all the different options you can set.

Instructions on setting the following ExtremeControl options:

- Advanced Settings
- Assessment Server
- Data Persistence
- End-System Event Cache
- Enforce Warning Settings
- Features
- Notification Engine
- Policy Defaults
- Status Polling and Timeout

Advanced Settings

Use the Advanced Settings panel to configure advanced settings for ExtremeControl. These settings apply to all users on all clients.

- 1. Select Administration > Options in ExtremeCloud IQ Site Engine. The Options window opens.
- 2. In the left-panel tree, expand the ExtremeControl folder and select Advanced Settings.
- 3. Use the **Resource Allocation Capacity** option configure the ExtremeCloud IQ Site Engine resources allocated to end-system and configuration processing services. The greater the number of end-systems and engines in your ExtremeControl deployment, the more resources it requires.
 - Low For low performance shared systems.
 - Low-Medium For medium performance shared systems, or low performance dedicated systems
 - Medium For medium performance shared systems, or medium performance dedicated systems.
 - Medium-High For high performance shared systems, or medium performance dedicated systems.
 - High For high performance dedicated systems.
 - Maximum For extremely high performance dedicated systems.

- 4. Use the **Hybrid Mode** option to enable Hybrid Mode for Layer 2 Controllers. Hybrid Mode enables a Layer 2 ExtremeControl Controller engine to act as a RADIUS proxy for switches, like an ExtremeControl Gateway engine. Select this option to enable Hybrid Mode for your Layer 2 Controllers at a global level. When the option is selected, the **Configuration** tab for a Layer 2 Controller displays an option to enable Hybrid Mode for that specific controller. Disabling Hybrid Mode at the global level when a controller has switches has a similar effect to deleting a gateway: the switches have the controller removed as a reference.
- 5. Select **Save** or select the **Autosave** checkbox.

Assessment Server

Use the Assessment Server view to provide assessment agent adapter credentials. The options apply to all users on all clients.

The assessment agent adapter credentials are used by the ExtremeControl engine when attempting to connect to network assessment servers, including Extreme Networks Agent-less, Nessus, or a third-party assessment server (an assessment server that is not supplied or supported by ExtremeCloud IQ - Site Engine). The password is used by the assessment agent adapter (installed on the assessment server) to authenticate assessment server requests. ExtremeControl provides a default password you can change, if desired. However, if you change the password here, you need to change the password on the assessment agent adapter as well, or connection between the engine and assessment agent adapter is lost and assessments are not performed. For instructions, see How to Change the Assessment Agent Adapter Password.

- 1. Select **Administration > Options**. The Options window opens.
- 2. In the left-panel tree, expand the ExtremeControl folder and select Assessment Server.
- 3. Specify the assessment agent adapter credentials.
- 4. Select **Save** or select the **Autosave** checkbox.

Data Persistence

Use the <u>Data Persistence view</u> to customize how ExtremeCloud IQ - Site Engine ages-out or deletes end-systems, end-system events, and end-system health results (assessment results) from the tables and charts in the <u>End-Systems tab</u>. These settings apply to all users on all clients.

- 1. Select **Administration > Options**. The Options window opens.
- 2. In the left-panel tree, expand the ExtremeControl folder and select Data Persistence.
- 3. In the **Age End-Systems** section, enter the number of days the Data Persistence Check uses as criteria for aging end-systems. Each day, when the Data Persistence check runs, it searches the database for end-systems ExtremeCloud IQ Site Engine has not received an event for in the number of days specified (90 days by default). It removes those end-systems from the tables in the End-Systems tab.

- 4. If you select the Remove Associated MAC Locks and Occurrences in Groups checkbox, the aging check also removes any MAC locks or group memberships associated with the end-systems being removed. The Remove Associated Registration Data checkbox is selected by default, so the aging check also removes any registration data associated with the end-systems being removed.
- 5. In the **End-System Event Persistence** section, select the checkbox if you want ExtremeCloud IQ Site Engine to store non-critical end-system events, which are events caused by an end-system reauthenticating. End-system events are stored in the database. Each day, when the Data Persistence check runs, it removes end-system events which are older than the number of days specified (90 days by default).
- 6. In the **End-System Information Event** section, select the checkbox if you want ExtremeCloud IQ Site Engine to generate an ExtremeControl event when end-system information is modified.
- 7. In the Health Result Persistence section, specify how many health result (assessment results) summaries and details are saved and displayed in the <u>End-Systems tab</u> for each end-system. By default, the Data Persistence check saves the last 30 health result summaries for each end-system along with detailed information for the last five health result summaries per end-system.
 There are two additional options:
 - You can specify to only save the health result details for quarantined end-systems (with the exception of agent-based health result details, which are always saved for all end-systems).
 - You can specify to save duplicate health result summaries and detail. By default, duplicate health
 results obtained during a single scan interval are **not** saved. For example, if the assessment
 interval is one week, and an end-system is scanned five times during the week with identical
 assessment results each time, the duplicate health results are not saved (with the exception of
 administrative scan requests such as Force Reauth and Scan, which are always saved). This
 reduces the number of health results saved to the database. If you select this option, all duplicate
 results are saved.
- 8. Set the time you would like the Data Persistence Check to be performed each day.
- 9. In the **Transient End-Systems** section, configure the number of days to keep transient end-systems in the database before they are deleted as part of the nightly database cleanup task. The default value is 1 day. A value of 0 disables the deletion of transient end-systems. Transient end-systems are Unregistered end-systems and have not been seen for the specified number of days. End-systems are not deleted if they are part of an End-System group or there are MAC locks associated with them. Select the **Delete Rejected End-Systems** checkbox if you want end-systems in the Rejected state to be deleted as part of the cleanup. You can also delete transient end-systems using the Tools > End-System Operations > Data Persistence option.
- 10. Select **Save** or select the **Autosave** checkbox.

End-System Event Cache

End-system events are stored daily in the database. In addition, the end-system event cache stores in memory the most recent end-system events and displays them in the End-System Events tab. This cache enables ExtremeCloud IQ - Site Engine to quickly retrieve and display end-system events without having to search through the database. Use the End-System Event

<u>Cache view</u> to configure the amount of resources used by the end-system event cache. This setting applies to all users on all clients.

- 1. Select **Administration > Options** in the menu bar. The Options window opens.
- 2. In the left-panel tree, expand the ExtremeControl folder and select End-System Event Cache.
- 3. Specify the parameters to use when searching for older events outside of the cache. (The search is initiated by using the **Search for Older Events** button in the <u>End-System Events tab</u>.) The search is ended when any one of the parameters is reached.
 - Maximum number of days to go back when searching
 - Maximum number of results to return from search
 - Maximum time to spend searching for events
- 4. Specify the number of events to cache. Keep in mind the more events you cache, the faster data is returned, but caching uses more memory.
- 5. The End-System Event Cache also keeps a secondary cache of events by MAC address. This means a particular end-system's events can be more quickly accessed in subsequent requests. Specify the number of MAC addresses kept in the secondary cache. Keep in mind that the more MAC addresses you cache, the more memory used. Also, note the secondary cache can include events not in the main cache, but were retrieved by scanning the database outside the cache boundary.
- 6. Select **Save** or select the **Autosave** checkbox.

Enforce Warning Settings

Use the <u>Enforce Warning Settings view</u> to specify warning messages you don't want displayed during the Enforce engine audit.

When an engine configuration audit is performed during an Enforce operation, warning messages can display in the audit results listed in the Enforce window. If an engine has a warning associated with it, you are given the option to acknowledge the warning and proceed with the enforce anyway.

These settings enable you to select specific warning messages that you do not want to have displayed in the audit results. This enables you to proceed with the Enforce without having to acknowledge the warning message. For example, you can have an ExtremeControl configuration that always results in one of these warning messages. By selecting that warning here, it is ignored in future audit results and you no longer have to acknowledge it before proceeding with the Enforce.

- 1. Select **Administration > Options** in the menu bar. The Options window opens.
- 2. In the left-panel tree, expand the ExtremeControl folder and select Enforce Warnings. The Enforce Warnings view opens.
- 3. Select the checkbox in the Ignore column next to the warning messages you don't want displayed.
- 4. Select **Save** or select the **Autosave** checkbox.

Setting Features Options

Use the <u>Features view</u> to automatically create new Policy mappings and profiles. If you are not using these features, you can disable them to remove sections that pertain only to those features from certain ExtremeCloud IQ - Site Engine windows.

Notification Engine Options

Use the <u>Notification Engine view</u> to define the default content contained in ExtremeControl notification action messages. For example, with an email notification action, you can define the information contained in the email subject line and body. With a syslog or trap notification action, you can specify certain information you want contained in the syslog or trap message. These settings apply to all users.

There are certain "keywords" that you can use in your email, syslog, and trap messages to provide specific information. Following is a list of the most common keywords used. For a complete list of available keywords for ExtremeControl notifications, see the Keywords Help topic.

- \$type the notification type.
- \$trigger the notification trigger.
- \$conditions a list of the conditions specified in the notification action.
- \$ipaddress the IP address of the end-system that is the source of the event.
- \$macaddress the MAC address of the end-system that is the source of the event.
- \$switchIP the IP address of the switch where the end-system connected.
- \$switchPort the port number on the switch where the end-system connected.
- \$username the username provided by the end user upon connection to the network.
- 1. Select Administration > Options. The Options window opens.
- 2. In the left-panel tree, expand the ExtremeControl folder and select Notification Engine. The Notification Engine view opens.
- 3. Use the fields to define the default content contained in notification action messages. For a definition of each field, see the <u>Notification Engine view</u> Help topic.
- 4. In the Advanced section, set parameters for the Action and Event queues processed by the Notification engine.
- 5. Select **Save** or select the **Autosave** checkbox.

Policy Defaults

Use the <u>Policy Defaults view</u> to specify a default policy role for each of the four <u>access policies</u>. These default policy roles display as the first selection in the drop-down lists when you create an

ExtremeControl profile. For example, if you specify an Assessment policy called "New Assessment" as the Policy Default, then "New Assessment" automatically displays as the first selection in the Assessment Policy drop-down list in the New ExtremeControl Profile window.

ExtremeCloud IQ - Site Engine supplies seven policy role names from which you can select. You can add more policies in the <u>Edit Policy Mapping window</u>, where you can also define policy to VLAN associations for RFC 3580-enabled switches. When a policy is added, it becomes available for selection in this view.

- 1. Select **Administration > Options**. The Options window opens.
- 2. In the left-panel tree, expand the ExtremeControl folder and select Policy Defaults.
- 3. Select the desired policies.
 - The Accept policy is applied to an end-system when an end-system has been authorized locally
 by the ExtremeControl Gateway and has passed an assessment (if an assessment was required),
 or the "Replace RADIUS Attributes with Accept Policy" option is used when authenticating the
 end-system.
 - The **Assessment policy** is applied to an end-system while it is being assessed (scanned).
 - The Failsafe policy is applied to an end-system when it is in an Error connection state. An Error state results if the end-system's IP address could not be determined from its MAC address, or if there was a scanning error and an assessment of the end-system could not take place.
 - The Quarantine policy is applied to an end-system if the end-system fails an assessment.
- 4. Select **Save** or select the **Autosave** checkbox.

Status Polling and Timeout

Use the <u>Status Polling and Timeout view</u> to specify polling and timeout options for ExtremeControl engines. These settings apply to all users on all clients.

- 1. Select **Administration > Options**. The Options window opens.
- 2. In the left-panel tree, expand the ExtremeControl folder and select Status Polling and Timeout.
- 3. In the ExtremeControl Appliance Enforce Timeout section, specify the amount of time ExtremeCloud IQ Site Engine waits for an enforce response from the engine before determining the ExtremeControl engine is not responding. During an enforce, an ExtremeControl engine responds every second to report that the enforce operation is either in-progress or complete. Typically, you do not need to increase this timeout value, unless you are experiencing network delays that require a longer timeout value.
- 4. In the ExtremeControl Inactivity Check section, you can enable a check to verify end-system ExtremeControl activity is taking place on the network. If no end-system activity is detected, an ExtremeControl Inactivity event is sent to the ExtremeControl Events view. You can use the Alarms and Events tab to configure custom alarm criteria based on the ExtremeControl Inactivity event to create an alarm, if desired.
- 5. In the **Status Polling** section, select the **Length of Timeout**, which specifies the amount of time ExtremeCloud IQ Site Engine waits when communicating with ExtremeControl engines for status

polling before determining contact failed. If ExtremeCloud IQ - Site Engine does not receive a response from an engine in the defined amount of time, ExtremeCloud IQ - Site Engine considers the engine to be "down" and the engine icon changes from a green up-arrow to a red down-arrow in the left-panel tree. The engine status refers to Messaging connectivity, not SNMP connectivity. This means that if the engine is "down," ExtremeCloud IQ - Site Engine is not able to enforce a new configuration to it.

- 6. Specify the **Polling Interval**, which is the frequency ExtremeCloud IQ Site Engine polls the ExtremeControl engines to determine engine status.
- 7. Select **Save** or select the **Autosave** checkbox.

How to Set Up Registration

The Extreme Networks ExtremeControl Solution provides support for Registration which forces any new end-system connected on the network to provide the user's identity in a web page form before being permitted access to the network. Registration utilizes Registration Web Server functionality installed on an ExtremeControl engine to enable end users to register their end-systems and automatically obtain network access without requiring the intervention of network operations. For more information on Registration and an overview of how it works, see the Registration section of the Concepts help file.

NOTE: For important information on web browser requirements for end-systems connecting through ExtremeCloud IQ - Site Engine, refer to the ExtremeControl Configuration Considerations Help topic.

This Help topic describes the specific steps that must be performed when deploying Registration on your network. The steps vary depending on whether you are using ExtremeControl Gateway engines and/or Layer 2 ExtremeControl Controller engines on your network. (Registration is not supported on the Layer 3 ExtremeControl Controller engines.)

For ExtremeControl Gateway engines you must:

- Identify the location in your network topology for the ExtremeControl Gateway installation.
- Define the access policy for authorizing unregistered end-systems.
- Configure policy-based routing on your network.
- Configure Registration parameters in ExtremeCloud IQ Site Engine.

For Layer 2 ExtremeControl Controller engines you must:

• Configure Registration parameters in ExtremeCloud IQ - Site Engine.

The Registration Web Server is pre-installed on the ExtremeControl engine. For instructions on installing and configuring a ExtremeControl engine, refer to your engine Installation Guide.

NOTE: It is important to add a DNS entry from the Fully Qualified Domain Name (FQDN) of the ExtremeControl engine (both ExtremeControl Gateways and ExtremeControl Controllers) into the DNS servers deployed on the network so that the device running ExtremeCloud IQ - Site Engine is able to resolve queries to these DNS servers. Otherwise, a short delay occurs in returning the Registration web page to end users on the network.

Information and instructions on:

- ExtremeControl Gateway Configuration
 - Identifying ExtremeControl Gateway Location
 - Defining the Unregistered Access Policy

- Configuring Policy-Based Routing
- Configuring ExtremeCloud IQ Site Engine (for ExtremeControl Gateway and ExtremeControl Controller Engines)

ExtremeControl Gateway Configuration

Perform the following steps when you are deploying Registration in a network that utilizes ExtremeControl Gateway engines. These steps are not necessary if you are utilizing only ExtremeControl Controller engines on your network.

Identifying ExtremeControl Gateway Location

Although several ExtremeControl Gateways can be deployed on the entire network depending on the number of connecting end-systems, only one ExtremeControl Gateway is required to serve as the Registration Web Server. The location of this ExtremeControl Gateway is important for the implementation of web redirection for unregistered end-systems on the network. The ExtremeControl Gateway serving as the Registration Web Server must be installed on a network segment directly connected to a router or routers that exist in the forwarding path of HTTP traffic from unregistered end-systems. This is because policy-based routing is configured on this router or routers to redirect the web traffic sourced from unregistered end-systems to this ExtremeControl Gateway. It is important to note that only the ExtremeControl Gateway that you wish to serve as the Registration Web Server needs to be positioned in such a manner. All other ExtremeControl Gateways can be positioned at any location on the network, with the only requirement being that access layer switches are able to communicate to the gateways.

Typically, the ExtremeControl Gateway serving as the Registration Web Server is positioned on a network segment directly connected to the distribution layer routers on the enterprise network, so that any HTTP traffic sourced from unregistered end-systems that are connected to the network's access layer can be redirected to that ExtremeControl Gateway. As an alternative, the ExtremeControl Gateway can be positioned on a network segment directly connected to the router providing connectivity to the Internet or internal web server farm. In this scenario, the HTTP traffic sourced from unregistered end-systems would be redirected to the ExtremeControl Gateway before reaching the Internet or internal web servers.

Defining the Unregistered Access Policy

When you implement Registration, you assign the Unregistered ExtremeControl Profile defined in ExtremeCloud IQ - Site Engine (ExtremeCloud IQ - Site Engine) as the Default Profile for all end-systems connected to the engine group. The Unregistered ExtremeControl Profile specifies that end-systems are **not** assessed for security posture compliance (at this time) and authorizes end-systems on the network with the "Unregistered" access policy. With this configuration, end-systems are first forced to register to the network, and after successful registration, can be assessed for security posture compliance and subsequently quarantined or permitted network access.

Note that an end-system group can be configured to exempt certain devices from having to register to the network, based on authentication type, MAC address, or user name. For example, an end-system group for the MAC OUI of the printer vendor for the network can be configured to exempt printers from having to register for network access.

Creating the Unregistered Access Policy

The Unregistered access policy must permit unregistered end-systems access to ARP, DHCP, DNS, and HTTP; particularly HTTP communication to the ExtremeControl Gateway implementing the Registration Web Server functionality. For a network composed of EOS policy-enabled switches in the access layer, you must create the appropriate network access services and rules for the Unregistered *policy role* in ExtremeCloud IQ - Site Engine's **Control** > **Policy** tab to meet these requirements, and enforce those changes to the policy-enabled switches. For a network composed of RFC 3580-enabled switches, you must ensure appropriate network services are enabled for the VLAN(s) associated to the Unregistered access policy.

For EOS policy-enabled Access Layer Switches

When configuring the Unregistered policy role (using ExtremeCloud IQ - Site Engine's **Policy** tab) for EOS policy-enabled switches, there are two required configurations:

- A rule must be added that permits HTTP traffic (i.e. TCP destination port equaling 80) on the network.
- The rule must specify a class of service action that rewrites the ToS value of the HTTP traffic to a value of 'y'. This value should match the decimal equivalent used in your policy-based routing that is used on the router.

If Assisted Remediation is already deployed with the Quarantine policy role appropriately configured for web redirection on EOS policy-enabled access layer switches, the simplest way to configure the Unregistered policy role in ExtremeCloud IQ - Site Engine is to copy and paste the Quarantine policy role under the **Roles** tab in ExtremeCloud IQ - Site Engine and rename this new policy role "Unregistered".

In addition, the **Policy** tab's Default Policy Domain includes an Unregistered role that is already configured with a service called Redirect Web Services, that includes an "Allow HTTP and Redirect" rule configured with the ExtremeControl Web Redirect Class of Service.

Perform the following steps in ExtremeCloud IQ - Site Engine to configure your Unregistered policy role.

NOTE: The ExtremeCloud IQ - Site Engine Default Policy Domain includes an ExtremeControl Web Redirect Class of Service you can use. Make sure that the ToS rewrite value is set to the appropriate value for your network. If you already created a Class of Service with ToS rewrite functionality for Assisted Remediation, you can use that same Class of Service for Registration and start with step number 3 below.

1. In ExtremeCloud IQ - Site Engine, access the **Administration** > **Options** tab and select Policy Manager in the left-panel.

- 2. In the Default Class of Service Mode section, select Role-Based Rate Limits/Transmit Queue Configuration to enable the Role-based Class of Service mode on your network devices.
- 3. Create a new Class of Service that implements the ToS rewrite functionality:
 - a. Open the Class of Service left-panel (Control > Policy tab > Class of Service).
 - b. Right-click the Class of Service navigation tree and select Create CoS. The Create CoS window opens.
 - c. Enter a name for the class of service (for example, "Web Redirection").
 - d. Select OK.
 - e. Select the **802.1p Priority** checkbox and use the drop-down list to select the **802.1p priority** to associate with the class of service.
 - f. Select the **Edit** button next to the ToS field and enter a value (hex).
 - g. The new Class of Service is automatically saved.
- 4. Create an "Allow HTTP" rule to a service currently included in your Unregistered policy role.
 - a. Right-click a service in the Roles/Services left-panel and select **Create Rule**.
 - b. Enter a name for the rule (for example, "Allow HTTP") and select **All Devices** in the **Rule Type(s)** drop-down list.
 - c. Select OK.
 - d. Select the new rule in the left-panel to display the rule details in the right panel.
 - e. Enter a **Description** for the rule.
 - f. Select **Enabled** in the **Rule Status** drop-down list.
 - g. In the Traffic Description section, select the **Edit** button.

The Edit Traffic Description window displays.

- h. Select Layer 4 Application Transport in the Traffic Classification Layer drop-down list.
- i. Select IP TCP Port Destination in the Traffic Classification Type drop-down list.
- j. Select HTTP (80) in the Well-Known Value drop-down list.
- k. Do not enter an IP address value.
- I. Select OK.
- m. In the Actions section, select **Permit Traffic** in the **Access Control** drop-down list.
- n. Select CoS you created in step 2 ("Web Redirection") in the Class of Service drop-down list.
- o. In the Open/Manage Domain(s) drop-down list at the top of the tab, select Save Domain.
- 5. Enforce these policy configurations to your network devices by selecting **Enforce Preview** in the **Enforce** drop-down list.

The **Enforce Preview** window displays.

6. Verify the information you are enforcing is correct and select the **Enforce** button.

For RFC 3580-compliant Access Layer Switches

A VLAN must be identified to which unregistered end-systems will be assigned upon connecting to the network. You can make this the same VLAN assigned to end-systems when they are being assessed or quarantined. The VLAN must provision network services to an unregistered end-system that permit the end-system to open a web browser; specifically HTTP, DHCP, ARP, and DNS. Furthermore, it is required that IP connectivity between the end-system and the ExtremeControl Gateway implementing the Registration Web Server functionality is operational.

The VLAN to which unregistered end-systems are assigned must be appropriately configured on all access layer switches where end-systems will be registering to the network. Access Control lists can be configured at the default gateway router's interface for the unregistered VLAN to restrict particular types of traffic sourced from end-systems within this VLAN to other areas of the network; withstanding the previously described provisioning requirements for this VLAN.

For Both EOS policy-enabled and RFC 3580-compliant Access Layer Switches

Now that you have defined the Unregistered policy role for EOS policy-enabled switches and/or the VLAN assigned to unregistered end-systems for RFC 3580-compliant switches, you must associate this policy role to the appropriate VLAN on the **Access Control** tab.

- 1. In ExtremeCloud IQ Site Engine, access the Control > Access Control tab.
- 2. Select the **Unregistered NAC Profile** entry in the Configuration > Profiles left-panel menu.
- 3. In the Accept Policy drop-down list, select Manage Policy Mappings.

The Manage Policy Mappings window displays.

4. Select the **Unregistered** policy and select the **Edit** button.

The Edit Policy Mapping window displays.

- 5. Select **Unregistered** in the **Policy Role** drop-down list.
- 6. Select the Save button.
- 7. Select **Close** to close the **Manage Policy Mappings** window.

Your Unregistered access policy is now configured to permit unregistered end-systems the ability to communicate to the ExtremeControl Gateway serving as the Registration Web Server. In the next step, the authentication, authorization, and assessment of unregistered end-systems will be specified.

Configuring the Unregistered ExtremeControl Profile

Now that you have created the Unregistered access policy, you can customize the Unregistered ExtremeControl Profile. The Unregistered NAC Profile is defined by default in ExtremeCloud IQ - Site Engine to specify that an unregistered end-system is **not** assessed for security posture compliance and that it is authorized on the network with the "Unregistered" policy. Therefore, unregistered end-systems are immediately assigned to the "Unregistered" policy when

connected to EOS policy-capable access layer switches without being assessed. The authentication, assessment, and authorization settings of the Unregistered NAC profile can be changed as required by your organization. When you have configured the Unregistered NAC Profile, it can be selected as the default profile for an engine group (as described in a later section) where end-systems will be required to register to the network.

To change the Unregistered NAC Profile, use the following steps.

- 1. In ExtremeCloud IQ Site Engine, access the Control > Access Control tab.
- 2. Expand Configuration > Profiles in the left panel.
- 3. Select the Unregistered NAC Profile in the left panel.
- 4. Select the desired authentication, assessment, and configuration settings.
- 5. Select **Save**.

Configuring Policy-Based Routing

As described above, the ExtremeControl Gateway serving as the Registration Web Server must be located on a network segment directly connected to a router or routers that exist in the transmission path of all traffic from any end-system that is not registered. This is because policy-based routing (PBR) must be configured on the routers to redirect the web traffic sourced from unregistered end-systems to that ExtremeControl Gateway.

If EOS policy-enabled switches are deployed on the network, this is done by configuring policy-based routing to forward all HTTP traffic with a ToS field of 'y' to the next-hop address of the Gateway serving as the Registration Web Server. If RFC 3580-enabled switches are deployed on the network, this is done by configuring policy-based routing to forward all HTTP traffic with the source IP address on the subnet(s)/VLAN(s) associated to the Unregistered access policy, to the next-hop address of the Gateway serving as the Registration Web Server.

In addition, if you are adding multiple ExtremeControl Gateways for redundancy, the network needs to be configured for redundant policy-based routing as well.

For EOS policy-enabled Access Layer Switches

Let's consider an example where the Unregistered access policy is associated to a policy role on EOS policy-enabled switches that uses the "Allow HTTP" classification rule to assign HTTP traffic the "Web Redirection" class of service. This class of service rewrites the ToS field in the HTTP traffic to a value of 0x40 (or 64 base 10), equivalent to a DSCP value of 16. (The DSCP is the value defined in the six most significant bits of the 8-bit ToS field.) Furthermore, the Unregistered access policy is associated to VLANs 10, 20, and 30 on RFC 3580-enabled switches on the network which map to subnets 10.1.10.0/24, 10.1.20.0/24, and 10.1.30.0/24, respectively. The following steps describe how to configure policy-based routing on an N-Series router or Cisco IOS-based router when Registration is deployed for EOS policy-enabled access layer switches.

1. Configure an entry in the access-list 102 to identify HTTP traffic with a DSCP of 16. access-list 102 permit tcp any any eq 80 dscp 16

2. Use a route-map to configure the access-list 102 ACL to redirect HTTP traffic from end-systems to the next-hop IP address of the ExtremeControl Gateway serving as the Registration Web Server, where "xxx.xxx.xxx" is the IP addresses of the Gateway. Note that multiple next hop IP addresses can be specified in the route-map if multiple Gateways are serving as Registration Web Servers. route-map 101

match ip address 102 set next-hop xxx.xxx.xxx

3. Apply the route map for the PBR configuration to the routed interface receiving the HTTP traffic from unregistered end-systems by entering the routed interface configuration prompt and executing the following command.

ip policy route-map 101

For RFC 3580-compliant Access Layer Switches

Let's consider an example where the Unregistered access policy is associated to VLANs 10, 20, and 30 on RFC 3580-enabled switches on the network which map to subnets 10.1.10.0/24, 10.1.20.0/24, and 10.1.30.0/24, respectively. The following steps describe how to configure policy-based routing on an N-Series router or Cisco IOS-based router when Registration is deployed for RFC 3580-compliant access layer switches.

1. Configure an entry in the access-list 102 to identify HTTP traffic sourced from subnets 10.1.10.0/24, 10.1.20.0/24, and 10.1.30.0/24.

access-list 102 permit tcp 10.1.10.0.0.0.0.255 any eq 80 access-list 102 permit tcp 10.1.20.0.0.0.0.255 any eq 80 access-list 102 permit tcp 10.1.30.0.0.0.0.255 any eq 80

2. Use a route-map to configure the access-list 102 ACL to redirect HTTP traffic from end-systems to the next-hop IP address of the ExtremeControl Gateway serving as the Registration Web Server, where "xxxx.xxx.xxx" is the IP addresses of the Gateway. Note that multiple next hop IP addresses can be specified in the route-map if multiple Gateways are serving as Registration Web Servers.

route-map 101 match ip address 102

set next-hop xxx.xxx.xxx.xxx

3. Apply the route map for the PBR configuration to the routed interface receiving the HTTP traffic from unregistered end-systems by entering the routed interface configuration prompt and executing the following command.

ip policy route-map 101

Setting up Redundancy on ExtremeControl Gateways

When adding multiple ExtremeControl Gateways for redundancy, the network needs to be configured for redundant policy-based routing as well. This is performed on the router in which policy-based routing is configured. Use the same commands described in the previous two sections except for the two following changes:

• In step 2, in addition to the single IP address set as the next-hop IP address, enter a list of IP addresses of the redundant Gateways. For example:

• In step 3, when adding the ip policy route-map to the router interface, specify an additional command called "ip policy pinger on". This command attempts to ping the first IP address that is specified in the next-hop to determine its availability. If it is not available, the next IP in the list of next-hops will be pinged and then used, if it is available.

For example:

ip policy route-map 101

ip policy pinger on

With policy-based routing and the Unregistered NAC Profile configured, Registration settings can be specified and then enabled on the network, as described in the next section.

Configuring the Access Control Tab (for ExtremeControl Gateways and Controllers)

Perform the following steps when you are deploying Registration in a network that utilizes ExtremeControl Gateway engines and/or Layer 2 ExtremeControl Controllers. (Registration is not supported on Layer 3 Controller engines.)

Use the Configuration section of the Access Control tab left-panel menu to configure parameters for the Registration web pages served from the ExtremeControl engine. All ExtremeControl engines are initially assigned a default portal configuration. Use this tab to view and edit the default configuration or create new configurations. When you define your portal configuration, enforce the Access Control configuration to your engine(s).

Use the following steps to define your portal configuration and enforce it to the engine:

- 1. In ExtremeCloud IQ Site Engine, access the Control > Access Control tab.
- 2. In the left panel, expand the Configuration section and select Captive Portals.
- 3. Select an existing captive portal and select **Edit** or select **Add** to create a new portal.
- 4. Select the portal configuration settings for your network using the Network Settings, Administration, and Website Configuration tabs, available in the left panel:
 - a. Network Settings view network web page parameters. These parameters are shared by both the Remediation and the Registration web pages. Be aware that if you deploy both the assessment/remediation and registration features, any changes will affect the web pages for both features.
 - b. <u>Administration</u> configure settings for the registration administration web page and grant access to the page for administrators and sponsors.
 - c. <u>Website Configuration</u> configure <u>Guest Settings</u>, <u>Authentication Settings</u>, <u>Survivable</u>
 <u>Registration</u>, and <u>Assessment/Remediation</u>. Additionally, use this to configure the <u>Look & Feel</u> of the website.
- 5. When you have finished making your changes to the portal configuration, select Save.
- 6. Enforce the Access Control configuration to the engine group.

7. To exempt certain end-systems or end users from having to register to the network, you can configure end-system groups based on authentication type, MAC address, or user name. For example, an end-system group for the MAC OUI of the printer vendor for the network can be configured to exempt printers from having to register for network access.

Registration is now enabled for all end-systems connecting to this engine group, with the exception of those end-systems and end users that have been exempted based on group membership.

How to Configure Pre-Registration

This Help topic describes how to configure and use the ExtremeControl pre-registration feature as a part of Secure Guest Access or Authenticated Registration. With pre-registration, guest users can be registered in advance and given a username and password, allowing for a more streamlined and simple registration process when the guest user connects to the network. This can be particularly useful in scenarios where guest users are attending a company presentation, sales seminar, or a training session.

Pre-registration allows IT to delegate control of the network registration process to less technical personnel such as company receptionists, administrative assistants, or training personnel. Using the pre-registration web portal, selected personnel can easily register guest users in advance of an event, and print out a registration voucher that provides the guest user with their appropriate registration credentials. The guest user then follows the instructions on the voucher to connect to the corporate network.

This topic includes information and instructions on:


- Configuring Pre-Registration
- Pre-Registering Guest Users
 - Pre-Registering a Single User
 - Pre-Registering Multiple Usersv

Configuring Pre-Registration


Following are instructions for configuring pre-registration in your portal configuration.

- 1. Open the **Control** > **Access Control** tab.
- 2. Select **Portal Configurations > Website Configuration** in the left-panel navigation tree.
- 3. Select <u>Guest Access</u> or <u>Authenticated Registration</u> (depending on the access type you are configuring).


NOTE: If neither panel is available in the Website Configuration navigation tree, select Website Configuration in the left-panel and select the appropriate configuration.




- 4. Select the **Enable Pre-Registration Portal** checkbox and specify whether personnel are able to register a single user, multiple users, or both single and multiple users.
- 5. Set the Generate Password Characters and Generate Password Length options. ExtremeControl uses these options when generating passwords for guest users to use when connecting to the network. These settings are shared by Authenticated Registration and Secure Guest Access. Changing it for one access type also changes it for the other.
- 6. For Authenticated Registration, select the Network Settings view to configure the connection URL specified on the Guest User Voucher (for example, www.ExtremeNetworks.com). Enter the URL in the Redirection To URL field. For Secure Guest Access, the Guest User Voucher provides instructions for connecting directly to the secure SSID.




- 7. Select Save to save your changes. Enforce your ExtremeControl Configuration to your engines.
- 8. Access the Pre-Registration Portal by entering the following URL in a browser window: https://<ExtremeControlEngineIP>/pre registration



- 9. At the top of the portal web page are instructions for the people performing the pre-registrations. To modify and edit these instructions:
 - a. In the **Control** > **Access Control** tab, select I&A Configurations > Portal in the left-panel navigation tree.
 - b. Select a Portal Configuration and select Website Configuration > Look & Feel to open the Look & Feel panel.



- c. Select the Message Strings Launch Message Strings Editor button. The Message Strings Editor window opens.
- d. Scroll down to the "preregIntroMulti" or "preregIntroSingle" message key and double-click that line. The Modify Localized Entry window opens.



- e. Enter any changes or modifications you wish to make to the instructions, and select **OK** to close the window.
- f. Enforce the changes to your engines.
- g. Refresh the browser window to see the new instructions in the Pre-Registration Portal.
- 10. The following sections provides information on how to pre-register a single user (when you want to pre-register one user at time) or multiple users (when you have a larger group of users to pre-register).


Pre-Registering Guest Users

After you have configured pre-registration, provide the URL for the Pre-Registration Portal (https://<ExtremeControlEngineIP>/pre_registration) to the personnel who are pre-registering guests. This can be network administrators or it can be personnel such as company receptionists, administrative assistants, or training personnel. (These users must be configured with administrative login privileges to access the web page).

The following sections provide steps for pre-registering single or multiple users in the Pre-Registration Portal.

Pre-Registering a Single User

Use the instructions in this section to pre-register a single end user using the Single User panel in the Pre-Registration Portal.



- 1. Enter the information for the guest user you want to pre-register. Fields with a red asterisk are required.
 - User Name Enter the user name for the guest user when connecting to the network. Usernames must be unique and cannot already exist in the local password repository. Usernames are case sensitive. For example, "JSmith" and "ismith" would be considered two different usernames.
 - First Name/Last Name Enter the guest user's first and last name. The name is printed on the voucher along with their registration credentials.
 - Password/Confirm Password Enter and confirm the password for the guest user connecting to the network. Select the Generate Password checkbox if you want ExtremeCloud IQ - Site Engine to automatically generate a password for you.
 - Password Repository When you pre-register the user, their credentials are automatically added to the local password repository specified here. Local Password Repositories are configured in the AAA Configuration window. (You only see this field if you have multiple repositories.)
 - Expires Time Select a registration expiration date from the calendar. The time is automatically set to 0:00:00, which is midnight. You can enter a specific time, if desired.

You can add additional fields to be displayed here using the Manage Custom Fields window accessed from the Customize Fields link in the Edit Portal Configuration window's NOTE: Authenticated Registration view or Secure Guest Access view. However the Pre-Registration web page always displays the First Name and Last Name fields even if they are not selected as visible/required in the Manage Custom Fields window. This is because it is important for the first and last name to be included on the pre-registration voucher printed out.

2. Select the **Pre-Register User** button to register the user. The user is added to the local password repository and added to the Registration Administration web page.

3. A voucher (see <u>example</u> below) is generated that provides registration instructions and the guest user's registration credentials. Print out this voucher to give to the guest user.


IMPORTANT:

The voucher must be printed out immediately, as there is no way to go back and print out a voucher after you leave the web page. If you do not print out the voucher, the voucher needs to be created by hand. In the event that the "Generate Password" option was used, you need to modify the guest user password using the registration administration page or local repository administration.


1. To register another user, you must re-access the Pre-Registration page by using the browser's back button or re-entering the URL.

Pre-Registering Multiple Users

Use the instructions in this section to pre-register multiple end users at one time using the Multiple Users panel in the Pre-Registration Portal. When pre-registering multiple users, create a CSV file to provide all the user credential information in table form. Then, upload the file to ExtremeCloud IQ - Site Engine to perform the pre-registration.



1. Select the CSV Template link to open a template CSV file where you create your list of guest users to pre-register. You can use a CSV template that includes password and password repository fields or not, depending on your network requirements. Do not change any of the column headings in the file.



Following is an explanation of the columns that need to be filled in for each user, depending on the template you selected.

- User Name Enter the username for the guest user connecting to the network. Usernames must be unique and cannot already exist in the local password repository. Usernames are case sensitive. For example, "JSmith" and "jsmith" would be considered two different usernames. (If you do try to pre-register existing usernames along with new usernames, you are notified of the error and given the option to continue registering the new names.)
- Password Enter the password for the guest user connecting to the network. If you want ExtremeCloud IQ - Site Engine to automatically generate end user passwords, leave the password column blank and select the **Generate Passwords** checkbox on the Multiple Users panel.
- Password Repository When you pre-register the user, their credentials are automatically be added to the local password repository specified here. Local Password Repositories are configured in the AAA Configuration window. If you are using the Default repository, you can use the Password Repository drop-down list (in the Multiple Users section) to select Default, and then you don't have to enter the Password Repository for each entry.
- First Name/Last Name Enter the guest user's first and last name. The name is printed on the voucher along with their registration credentials.


NOTE: You can add additional columns to be included in the template using the Manage Custom Fields window accessed from the Customize Fields link in the Edit Portal Configuration window's Authenticated Registration view and Secure Guest Access view, however, the template always displays the First Name and Last Name fields even if they are not selected as visible/required in the Manage Custom Fields window. This is because it is important for the first and last name to be included on the preregistration voucher you print.

- 2. When you have finished entering the guest user information, save and close the file.
- 3. Back in the Multiple Users panel, enter the path and filename for the CSV file by using the **Browse** button to browse to the file on your system.
- 4. If your CSV file includes a Password Repository, use the Password Repository drop-down list to specify whether to use the default repository or the repository specified in the file.
- 5. Select the **Upload** button. Users are added to the local password repository and to the Registration Administration web page.
- 6. Individual vouchers (see an <u>example</u> below) are generated that provide registration instructions and the guest user's registration credentials for each guest user. Print out these vouchers to give to the guest users.

IMPORTANT: Vouchers must be printed out immediately, as there is no way to go back and print out a voucher after you leave the web page. If you do not print out the vouchers, the vouchers have to be created by hand. In the event that the "Generate Password" option is used, you need to modify the guest user passwords using the registration administration page or local repository administration.

7. To register another user, you must re-access the Pre-Registration Portal by using the browser's back button or re-entering the URL.

Sample Guest User Voucher



• Portal Configuration

How to Enable RADIUS Accounting

This Help topic describes how to use RADIUS accounting to provide real-time end-system connection status in ExtremeCloud IQ - Site Engine. RADIUS accounting collects various end-system session data that ExtremeCloud IQ - Site Engine uses to determine connection status for each end-system session. This can be useful for compliance purposes, enabling you to determine both when an end-system session started and when it was terminated.

RADIUS accounting is also used to monitor switches for Auto Tracking, CEP (Convergence End Point), and Switch Quarantine authentication sessions, when used in conjunction with the Monitoring or Network Access switch authentication access types. (For more information, see the Auth. Access Type section of the Add/Edit Switch Window Help topics.)

You must be running ExtremeControl engine version 4.0 or higher to take advantage of RADIUS accounting functionality in ExtremeCloud IQ - Site Engine.

For Extreme Networks stackable and standalone devices (A-Series, B-Series, C-Series, D-Series, G-Series, and I-Series), ExtremeCloud IQ - Site Engine uses a combination of SNMP and CLI (command line interface) to configure RADIUS accounting on the switch. Before enabling RADIUS accounting on these devices, read through <u>Considerations for Fixed Switching Devices</u> below.

NOTES: RADIUS accounting is not supported on the ExtremeControl Controller.

Use the following steps to enable RADIUS accounting:

1. Enable RADIUS accounting on your switches and controllers using the instructions appropriate for your devices.

For Extreme Networks devices or ExtremeWireless Controller devices running firmware version 9.21.x.x or newer:

- a. **If you are editing an existing device:** In the right-panel **Switches** tab, select the devices you want to perform RADIUS accounting and select the **Edit** button. The Edit Switches in ExtremeControl Appliance Group window opens.
 - If you are adding a new device: Select Add in the right-panel Switches tab and the Add Switches to ExtremeControl Appliance Group window opens.
- b. Set the RADIUS Accounting option to **Enabled**. Select **OK**.
- c. Enforce to your engines.

For ExtremeWireless Controller devices running firmware versions older than 9.21.x.x:

a. RADIUS accounting must be enabled manually on the controller using the ExtremeWireless Assistant or the device CLI (command line interface).

b. Be sure to configure the ExtremeControl engine IP address as the IP address of the RADIUS server. Refer to your wireless controller User Guide for instructions on enabling RADIUS accounting via the ExtremeWireless Assistant, or the CLI Reference Guide for the exact CLI command syntax to use.

For third-party switching devices:

- a. RADIUS accounting must be enabled manually on the device using the device CLI (command line interface).
- b. Be sure to configure the ExtremeControl engine IP address as the RADIUS accounting server. Refer to your device documentation for the exact command syntax.
- 2. If you are doing RADIUS accounting in an ExtremeControl environment where the primary RADIUS server is being used for redundancy in a single ExtremeControl engine configuration (Basic AAA configuration only), then enable the Proxy RADIUS Accounting Requests option in the Edit RADIUS Server window.
 - a. In the Edit Basic AAA Configurations window, use the Configuration Menu button in the Primary RADIUS Server field to open the Manage RADIUS Servers window.
 - b. Select the RADIUS Server and select Edit.
 - c. Enable the Proxy RADIUS Accounting Requests option. Select **OK**.
 - d. Enforce to your engine.

With RADIUS accounting enabled, you now see real-time connection status in the ExtremeCloud IQ - Site Engine **End-Systems** tab and Dashboard.

Considerations for Fixed Switching Devices

ExtremeCloud IQ - Site Engine uses a combination of SNMP and CLI (command line interface) to configure RADIUS accounting on Extreme Networks stackable and standalone devices (A-Series, B-Series, C-Series, D-Series, G-Series, and I-Series). Due to a limitation on the SNMP interface, the configuration can be read via SNMP, but must be written to the device via CLI. Before enabling RADIUS accounting on these devices, read through the following considerations.

NOTE: These considerations do not apply to A4, B5, and C5 devices running firmware version 6.81 and higher. Those devices support RADIUS accounting configuration using SNMP.

- The devices must be assigned a Device Access profile that provides Write access and includes CLI credentials for Telnet or SSH. Profiles and CLI credentials are configured using the Authorization/Device Access tool's **Profiles** tab.
- Before you enforce a new RADIUS server configuration to your fixed switching devices, you should verify that your CLI credentials are configured according to the settings in your new configuration. This is because the Enforce process first writes the RADIUS server configuration to the switch using SNMP, and then writes the RADIUS accounting configuration to the switch using Telnet or SSH. If CLI credentials are not configured according to the new RADIUS server configuration, then the RADIUS

accounting configuration are not written to the switches.

For example, by default you can Telnet to a fixed switching device using username=admin (with no password or a blank password). But, if you configure a new RADIUS configuration with an Auth Access Type (or Realm Type)=Any, then change the Device Access for the switches to use the IAS credentials, in order for ExtremeCloud IQ - Site Engine to successfully write the RADIUS accounting information to the switches during Enforce.

Fixed switches only permit one accounting server to be configured. If a primary and secondary ExtremeControl gateway are configured for the switch, only the primary gateway's accounting configuration is written to the switch. If a secondary gateway is configured, a warning is displayed.

Considerations for ExtremeXOS/Switch Engine Devices

ExtremeCloud IQ - Site Engine uses CLI access to perform RADIUS accounting configuration operations on ExtremeXOS/Switch Engine devices. CLI credentials for the device are obtained from the device profile and must be configured in the Authorization/Device Access tool.

Guest and IoT Manager Configuration in ExtremeCloud IQ - Site Engine and Access Control (Legacy)

Guest & IoT Manager (GIM) is an application that allows you to access and manage guest user and end-system (device) activity information. Through ExtremeCloud IQ - Site Engine and ExtremeControl, GIM provides non-IT personnel with the tools to configure limited system access for guest users and/or devices based on authorization constraints you define.

NOTE:

Beginning in ExtremeCloud IQ - Site Engine 23.04.11, GIM performs a version compatibility check as it connects to ExtremeCloud IQ - Site Engine. If you are attempting to connect to an incompatible version of ExtremeCloud IQ - Site Engine, GIM displays an error message.

Non-IT personnel who are designated as provisioners can provide limited access to other guest users for a specified amount of time for specific purposes. For example, your company is conducting product training for customers at one of your offices. You provide the front desk employee at the site provisioner access so he or she can provide participating customers limited guest user access to your system for that day only. Refer to Extreme Control Guest and IoT Manager Configuration for more information about provisioner and guest user access.

Connecting GIM to ExtremeControl

GIM uses a REST API to communicate with ExtremeCloud IQ - Site Engine through an Access Control engine. In order for GIM to access the REST API, it must be authorized to do so by configuring the appropriate GIM capability in the Authorization Group configuration in ExtremeCloud IQ - Site Engine. The REST API allows GIM to store its configuration data in the ExtremeCloud IQ - Site Engine database.

Use the following steps to create an Authorization Group and add users to that Authorization Group:

- 1. Open the **Administration > Users tab** in ExtremeCloud IQ Site Engine.
- 2. Create a new Authorization Group for users with access to the GIM REST API.
- 3. Select Save.
- 4. Create users and add them to the new Authorization Group.
- Select Save.
- 6. Access the Administrator Application of GIM.
- 7. Open the **Administration** > **Access Control Engine** tab in GIM.
- 8. Open the Engine Details tab.

9. Enter the information for the Access Control engine you are using for GIM. For additional information, see Configuring Engine Details on page 49 of the Extreme Control Guest and IoT Manager Configuration document.

NOTE:

Enter the credentials of the user or users added to the GIM REST API Authorization Group in the **Admin Username** and **Admin Password** fields.

Configuring the RADIUS Protocol for GIM Authentication

After adding users to the GIM Authorization Group, enter the IP address and RADIUS shared secret in ExtremeCloud IQ - Site Engine and in GIM to allow the Access Control engine to authenticate provisioners in GIM.

- 1. Open the **Control** > **Access Control** tab in ExtremeCloud IQ Site Engine.
- 2. Expand the Engines folder in the left panel.
- 3. Select the Engine Group through which provisioners are authenticating.
- 4. Open the Guest and IoT Managers tab in the right panel.
- 5. Select Add.

The Add Guest and IoT Manager window opens.

- 6. Enter the GIM IP address.
- 7. Enter a Shared Secret and copy it to a safe location.

NOTE:

The shared secret functions as a password, allowing GIM and the RADIUS server (the Access Control engine) to communicate. Use a strong shared secret difficult for others to guess.

- 8. Access the Administrator Application of GIM.
- 9. Open the **Administration** > **Access Control Engine** tab in GIM.
- 10. Open the RADIUS tab.
- 11. Enter the RADIUS information on the tab. For additional information, see Configuring RADIUS Settings on page 50 of the Extreme Control Guest and IoT Manager Configuration document.


The Access Control engine is configured as the RADIUS server for GIM. Configure each GIM application with an IP Address and Shared Secret in ExtremeCloud IQ - Site Engine.

Creating and Configuring a GIM Domain

A GIM domain contains all of the configuration information. GIM domains are created in ExtremeCloud IQ - Site Engine and the configuration within that domain is configured in GIM.

To create a GIM domain in ExtremeCloud IQ - Site Engine:

- 1. Open the Control > Access Control tab in ExtremeCloud IQ Site Engine.
- 2. Expand the Engines folder in the left panel.
- 3. Select the Engine Group through which provisioners are authenticating.
- 4. Open the **Details** tab in the right panel.
- 5. Select **Edit** in the Guest and IoT Configuration section of the tab. The **Edit Guest and IoT Manager Configuration** window opens.
- 6. Select **New** from the drop-down list to create a new domain. The **Create Guest and IoT Management Domain** window opens.




- 7. Enter the name of your GIM domain and select **New** in the **Local Password Repository** drop-down list to create a new password repository for GIM. The **Create Repository** window opens.
- 8. Enter a name for the local password repository you are using for your GIM provisioners and users.
- 9. Select Create. The Edit Local Password Repository window opens.
- 10. Select Add. The Edit User window opens.
- 11. Enter the information for at least one user.
- 12. Select OK.
- 13. Select the local password repository you created in the **Local Password Repository** drop-down list in the **Create Guest and IoT Management Domain** window.
- 14. To customize the sponsor retrieval in your GIM Domain, select the GIM Sponsor Retrieval Advanced Configuration check box and choose a different LDAP configuration and Search Root specifically for sponsor look-ups.

NOTE:

Not enabling the GIM Sponsor Retrieval Advanced Configuration defaults sponsor retrieval to use an LDAP configuration based solely on the Sponsor Group configured in GIM.

a. The Sponsor Retrieval panel displays.



- b. Select, create, or modify the LDAP Configuration from the drop-down list.
- c. Enter the Search Root.
- d. Optionally, you can select the check box to enable the GIM Sponsor LDAP Group filter, which further filters the search for a sponsor using the GIM-configured Sponsor LDAP Group.

NOTE:

Not selecting the GIM Sponsor LDAP Group filter ignores yours Sponsor Group setting in GIM and uses the LDAP configuration and Search Root you define in the Create Guest and IoT Management Domain window for sponsor look-ups.

15. Select Save.

The templates, users, devices, and other information configured in the GIM application are stored in the GIM domain.

NOTE:

While the domain is stored in ExtremeCloud IQ - Site Engine, the only part of the GIM domain configured in ExtremeCloud IQ - Site Engine is the authentication method used by GIM provisioners and users.

Configuring GIM Authentication

In GIM, the Administrator creates provisioners via the Administration login. Provisioners then provide network access to users or devices using the Provisioner login.

Local Password Repository

When you create a provisioner while logged into GIM as an Administrator, ExtremeCloud IQ - Site Engine saves the provisioner credentials in the default local password repository associated with the GIM Domain.

When you provide network access to users or devices in GIM, those credentials are also saved in the local password repository associated with the GIM Domain.

LDAP

Provisioners can also authenticate via <u>Active Directory</u> associated with an LDAP Configuration in ExtremeCloud IQ - Site Engine. For provisioners for which both LDAP and a local password repository are available as authentication methods, the methods can be independent or work in conjunction with each other (for example, if LDAP authentication fails, ExtremeCloud IQ - Site Engine checks the local password repository for valid credentials).

To configure LDAP as an authentication method:

- 1. Access GIM as the Administrator.
- 2. Open the Onboarding Template tab and select Add.
- 3. Open the **Advanced** tab.
- 4. Enter the Active Directory field against which authentication is verified (for example, cn=gimGroup1, dc=extremenetworks, dc=com). The entire path must match for authentication to be successful.

Some common Active Directory objects used include:

- cn=common name
- dn=distinguished name
- dc=domain controller
- ou=organizational unit
- 5. Access ExtremeCloud IQ Site Engine.
- 6. Open the Control tab.
- 7. Select the Access Control tab.
- 8. Select the **Configuration > Configurations** tab in the left-panel tree.
- 9. Expand the Access Control Configuration associated with the Access Control Engine Group to which GIM is associated.
- 10. Select AAA.
- 11. <u>Configure the LDAP configuration</u> to provide authentication and authorization for network end users and host machines via Active Directory.
- 12. Save the LDAP configuration.
- 13. Expand the Access Control Configuration associated with the Access Control Engine Group to which GIM is associated.
- 14. Select AAA.

- 15. Configure the <u>Authentication Rules table</u> to authenticate via your LDAP configuration, your local repository, or both by adding both to the table. If using both authentication methods, ensure the authentication method you want to take precedence is listed first in the table.
- 16. Select Save.

The Access Control Engine now authenticates GIM users based on the Access Control Configuration.

Via the legacy NAC Manager java

application, ensure Manual Set (Accurate) is listed first in the Device Type Detection Source Precedence Order in the Edit

IMPORTANT: Appliance Settings window on the Device

Type Detection tab. This is the default precedence, and is required for GIM-assigned device types to affect

authentication.

Once GIM is fully connected to ExtremeCloud IQ - Site Engine and Access Control, follow the steps outline in the Extreme Control Guest and IoT Manager Configuration document.

Configuring Multiple Active Directory Domains

You can configure multiple Active Directory (AD) domains to authenticate users that reside on Active Directories that do not have trust between them. Additionally, you can configure multiple authentication rules so that if authentication to one fails, ExtremeControl can automatically attempt to authenticate against a second domain.

Requirements

Prior to configuring multiple AD domains:

- Ensure all AD servers communicate using DNS name.
- Validate multi-domain functionality works for your network.

Validating Multiple AD Domain Functionality

To ensure you can configure multiple AD domains for authentication on your network, ExtremeControl must be able to resolve all Directory service domains correctly. DNS resolution is required for multiple AD domain functionality to work properly. For example, if you are using a third-party DNS server (e.g. Infoblox), ExtremeControl is able to resolve all domains correctly. If one of AD's is acting as a DNS server, configure it (using DNS conditional forwarding) to resolve other Domains.

Additionally, ExtremeControl runs the wbinfo command line tool to check the reachability of AD servers to which it joined. In this multi-join scenario, ExtremeControl runs wbinfo against all joined Directory Services.

Joining Multiple Active Directory Domains

After you verify you can configure multiple Active Directory domains on your network, perform the following to configure the functionality:

- 1. Access the **Advanced AAA Configurations** tab.
- 2. Select All Domains in the Join AD Domain drop-down list.

NOTE: If multiple Active Directory domains are configured, ExtremeControl attempts to join them all.

- 3. Select **Add** in the Authentication Rules section to open the **Add/Edit User to Authentication Mapping** window.
- 4. Configure multiple authentication rules with an **Authentication Method** of **LDAP Authentication** in the Authentication Rules section.

- 5. Select the Fall-through if Authentication Failed checkbox if you want ExtremeControl to attempt to authenticate a user against the next AAA authentication rule in the table if the current authentication rule fails or times out. If this checkbox is not selected and authentication fails, the user is not authenticated and ExtremeCloud IQ Site Engine does not attempt to authenticate using any other rules in the table.
- 6. Select **OK**.
- 7. Select **Save**.

ExtremeControl attempts to join to all Domains you configure in the AAA authentication rules. If ExtremeControl is not able to join to any Domains, then a timer runs and attempts to keeps trying to rejoin. When ExtremeControl joins a particular domain, then a separate health check timer runs to ensure AD server is reachable.

Multiple AD domains are configured and if you enabled fall-through for your rules, ExtremeControl automatically attempts to authenticate against the next rule in the table.

Important Note

If duplicate users exist in multiple Active Directory domains with the same password, the AAA rule(s) with user pattern (for example, Domain*) needs to be configured for the user to match the domain name and use the AAA rule correctly.

For example, a user **administrator** exists in 2 Active Directory domain servers and the following is configured in AAA rule:

- All LDAP Authentication using Domain A.com server fall through enabled
- All LDAP Authentication using Domain B.com server

When administrator joined, the Domain_B domain tries to authenticate the user. The administrator user is successfully authenticated to the Domain_A.com server because the user does exist in Domain_A.com server. To avoid this, configure the AAA rule with user pattern as seen below:

- User matching Domain_A* (or *@domain_a.com) using Domain_A.com server fall through enabled
- User matching Domain_B* (or *@domain_b.com) using Domain_B.com server

How to Set Up Access Policies and Policy Mappings

Access policies define the appropriate level of access to network resources allocated to a connecting end-system based on the end-system's authentication and/or assessment results. There are four access policies defined in an ExtremeControl profile: Accept policy, Quarantine policy, Failsafe policy, and Assessment policy. When an end-system connects to the network, it is assigned one of these access policies, as determined by the ExtremeControl profile assigned to the matching ExtremeControl rule and the end-system state.

In your ExtremeControl profiles, each access policy is associated to a *policy mapping* that defines exactly how an end-system's traffic is handled when the access policy is applied.

A policy mapping specifies the policy role (created in the **Policy** tab) and other RADIUS attributes included as part of a RADIUS response to a switch. The RADIUS attributes required by the switch are defined in the Gateway RADIUS Attributes to Send field configured in the Edit Switch window. Policy mappings are configured in the Edit Policy Mapping Configuration window.

How you set up your access policies depends on whether your network utilizes ExtremeControl Controller engines and/or ExtremeControl Gateway engines. In addition, if your network utilizes ExtremeControl Gateway engines, your setup depends on whether your network contains EOS switches that support Policy, third-party switches that support RFC 3580, or switches that support RADIUS attributes that are defined manually.

For ExtremeControl Controllers:

If your network utilizes ExtremeControl L2/L3 controller engines, the access policies specified in ExtremeControl profiles are mapped to policy roles that are defined in a default policy configuration already configured on the controller. It is recommended that you review this default policy configuration using the **Policy** tab. To do this, you must create a policy domain in the **Policy** tab specifically for the ExtremeControl Controller, assign the ExtremeControl Controller to the domain, then import the policy configuration from the device into **Policy** tab. Review the policy roles and make any rule changes required for your environment. When you have finished modifying the policy configuration, you must enforce it back to the ExtremeControl Controller.

For ExtremeControl Gateway Appliances:

If your network utilizes ExtremeControl Gateway engines, the access policies specified in ExtremeControl profiles are mapped to policy roles that must be created and defined in the **Policy** tab and enforced to the policy-enabled switches in your network. If you have RFC 3580-enabled switches in your network, ExtremeCloud IQ - Site Engine lets you associate your policy roles to a VLAN ID or VLAN Name using the Policy Mappings panel. This allows your ExtremeControl Gateway engines to send the appropriate VLAN attribute instead of a policy role to those switches that are RFC 3580-enabled.

Policy mappings have a Location option that allows different VLAN IDs to be returned for a policy based on the location the authentication request originated from. This is useful in networks that have a VoIP/voice VLAN that is defined on multiple switches, but that VLAN maps to a unique VLAN ID on each switch. (For more information, see the section on Location in the Edit Policy Mapping Configuration Window Help topic.)

NOTE: If you have RFC 3580-enabled switches in your network, be sure to verify that the DHCP Resolution Delay Time option is set correctly in your Appliance Settings (Tools > Manage Advanced Configurations > Global and Appliance Settings). This option specifies the number of seconds an ExtremeControl engine waits after an authentication completes before attempting to resolve the end-system's IP address. When modifying this delay, keep in mind that for RFC 3580 devices, the engine links down/up a port to force the end-system to get a new IP address when ExtremeCloud IQ - Site Engine determines that the VLAN has changed. If the delay time specified is less than the amount of time the end-system needs to renew its IP address, then the ExtremeControl engine can resolve the end-system's IP address incorrectly (to the previously held IP), or additional delay can be introduced as the resolution process attempts to resolve the address based on the configured retry interval. This is a problem when either registration or assessment is enabled: the registration process never completes or takes an unacceptable amount of time to complete, or the ExtremeControl engine could attempt to scan the incorrect IP address. Be sure to take into account the amount of time required for an end-system to get a new IP address when setting the delay time value.

Setting Up Your Access Policies

Before you begin working with the **Access Control** tab, use these steps to define the policy mapping criteria (policy roles, corresponding VLAN IDs, etc.) available for selection for each access policy.

For each ExtremeControl profile, create a worksheet listing the four ExtremeControl policies. For each
access policy, associate a policy role (created in the **Policy** tab), and the policy role's corresponding
VLAN ID, if you are using RFC 3580-enabled switches in your network. For a description of each access
policy, and some guidelines for creating corresponding policy roles, see the section on Access Policies in
the Concepts file.

NOTE: If your network uses ExtremeControl Gateway engines with only RFC 3580-enabled switches, instead of listing policy roles, simply create a list of policy names that correspond to the VLANs you are using in your network. One tip is to use policy names that identify the corresponding VLAN name for ease of selection when you are creating your ExtremeControl profiles.


Here's an example of a worksheet for an ExtremeControl profile that contains both policy-enabled and RFC 3580 switches:

Access Policy	Policy Role	VLAN ID
Accept Policy	Enterprise User	[2] Enterprise User VLAN

Access Policy	Policy Role	VLAN ID
Quarantine Policy	Quarantine	[4] Quarantine VLAN
Failsafe Policy	Failsafe	[5] Failsafe VLAN
Assessment Policy	Assessing - Strict	[6] Assessing - Strict VLAN

- 2. For ExtremeControl Controllers, use the **Policy** tab to verify that the policy configuration contains the required policy roles, and that the configuration has been enforced to the ExtremeControl Controller. See the instructions above.
- 3. For ExtremeControl Gateways, verify each policy role listed on your worksheet is created in ExtremeCloud IQ - Site Engine's **Policy** tab and enforced to the policy-enabled switches in your network. If you have RFC 3580-enabled switches in your network, verify that your VLANs have been created on the switches in your network.
- 4. Define the policy mappings that map each access policy to the appropriate policy role as specified in your worksheet.
 - a. Select a policy mapping configuration from the ExtremeControl Configurations > ExtremeControl Profiles > Policy Mappings left-panel option.
 - b. In your ExtremeControl profile, your policy mappings are available for selection when you define your Accept, Quarantine, Failsafe, or Assessment access policy.






c. Select between a Basic policy mapping and an Advanced policy mapping, depending on your network needs by selecting **Switch to Advanced** or **Switch to**

Basic at the top of the panel. Typically, the Basic policy mapping configuration is used unless your devices require customization or when using locations in your mappings. If Basic Policy Mapping is used, then the **Add** new policy mapping, as well as **Edit** policy mapping, gives the option to show the advanced options.

ExtremeControl provides a list of default policy mappings you can use. Be aware if you use one of the default mappings, you still need to verify that the policy role specified in the mapping is part of your ExtremeControl Controller policy configuration and/or is created and enforced to the policy-enabled switches in your network via the **Policy** tab.

d. To add a new policy mapping, select the **Add** button to open the Add Policy Mapping window.



For the new policy mapping, enter a mapping name and specify a policy role (created in the **Policy** tab) and other required RADIUS attributes included in the RADIUS response to a switch. Select **OK** to add the mapping. Note that the required RADIUS attributes for your switches are defined in the Gateway RADIUS Attributes to Send field configured in the Edit Switch window, as shown below.

e. Select **OK** to close the Edit Policy Mapping Configuration window.

How to Configure Credential Delivery for Secure Guest Access

Secure Guest Access provides secure network access for wireless guests via 802.1x PEAP by sending a unique username, password, and access instructions for the secure SSID to guests via an email address or mobile phone (via SMS text). Use the instructions in this Help topic to configure the method used to send guests their credentials and access instructions for the secure SSID.

Configuration Steps

The Credential Delivery method is configured in your portal configuration. Depending on the method you specify, the appropriate custom fields must be configured for display on the Registration web page, so that end users can enter the required information.


The following table provides a description of each credential delivery method and lists their custom field requirements.

User Verification Method	Description	Custom Field Requirement
Captive Portal	The credential information is displayed on the Registration web page.	There are no Custom Field requirements.
Email	The end user must enter a valid email address on the Registration web page.	The Email Address Custom Field must be set to Required .
SMS Gateway	The SMS Gateway provider must support SMTP API. The SMS Gateway provider converts the email to an SMS text message. The end user must enter a mobile phone number on the Registration web page.	The Phone Number Custom Field must be set to Required .
SMS Gateway or Email	The SMS Gateway provider must support SMTP API. The SMS Gateway provider converts the email to an SMS text message. The end user must enter a mobile phone number or email address on the Registration web page.	The Phone Number and Email Address Custom Fields must be set to Visible.
SMS Text Message	The mobile provider converts the email to an SMS text message. The end user must enter a valid mobile phone number on the Registration web page.	The Phone Number Custom Field must be set to Required .


User Verification Method	Description	Custom Field Requirement
SMS Text or Email	The mobile provider converts the email to an SMS text message. The end user must enter a valid mobile phone number or email address on the Registration web page.	The Phone Number and Email Address Custom Fields must be set to Visible.

Use the following steps to configure credential delivery for Secure Guest Access in your portal configuration.

- 1. In the Access Control tab, access the Portal Configuration. Select Secure Guest Access in the Portal Configuration tree. (If you don't see this selection, select Features in the tree and enable the Secure Guest Access feature.)
- 2. In the Secure Guest Access panel, use the drop-down list to select the desired Credential Delivery Method (refer to the table above).



3. If you selected the SMS Text Message or the SMS Text or Email Credential Delivery method, select the Service Providers Edit button to configure the list of mobile service providers from which end users can select the Registration web page. The Mobile Service Provider List provides a default list of providers that can be edited to include the appropriate service providers for your geographic location.




You can comment out entries by preceding each line with either a # or // to enable temporary editing of the file without removing the text.

The list requires one service provider entry per line, using the following format: <Provider>:phonenumber@<specificdomain>.


When the end user registers, they only see the <Provider> portion in the drop-down list of providers on the Registration web page.

Select **OK** to close the window.

4. If you have selected the SMS Gateway or SMS Gateway or Email method, enter the SMS Gateway Email address provided by the SMS Gateway provider.



5. For all methods, select the Message Strings **Edit** button to open the Message Strings Editor where you can customize the text displayed on the Registration web page and the messages sent to the end user.



You need to modify different message strings sent to the end user, depending on the delivery method or methods you selected. Double-click the message to open a window where you can edit the message text.

NOTE: When customizing message strings for text messaging (SMS Gateway or SMS Text Message) it is best to keep the message length as short as possible (under the maximum 160 characters limit). Some providers break long messages into multiple messages and other providers truncate the message, which could cause important information to be missing from the text message the guest receives.

• **Email** — This method uses the following strings:

 $secure Guest Access Email Msg Body-the\ default\ message\ shouldn't\ need\ to\ be\ changed.$

secureGuestAccessEmailSentFromAddress — you need to change the default message to the appropriate email address for your company.

secureGuestAccessEmailSentFromName — the default message shouldn't need to be changed.

 $secure Guest Access Email Subject-the\ default\ message\ shouldn't\ need\ to\ be\ changed.$

• SMS Gateway — Depending on your SMS Gateway provider and their required format, modify the following message strings using appropriate variables to customize the dynamic data such as phone number.

secureGuestAccessSMSMsgBody


secureGuestAccessSMSSubject

• SMS Text Message — This method uses the following strings. The default messages shouldn't need to be changed.

secureGuestAccessSMSMsgBody

secureGuestAccessSMSSubject


Select **OK** to close the window.



6. Select the Customize Fields **Open Editor** button to open the Manage Custom Fields window.

7. Set the appropriate custom fields to display on the Registration web page, depending on the delivery method you selected (refer to the <u>table</u> above). If you do not set these fields, ExtremeControl automatically sets them for you based on your delivery method.

These settings are shared by Guest Web Access, Guest Registration, and Secure Guest Access. Changing them for one access type also changes them for the others. For more information, see the Manage Custom Fields Window.




- 8. Select **OK** to close the window.
- 9. Back in the Portal Configuration, select **Save** to save your changes.
- 10. Enforce the new portal configuration to your engine(s).


Credential delivery is now configured for your secure guest access.

How Secure Guest Access Works


When a guest attempts to access the network, the Registration web page asks for their email address and/or phone number, and any other required/configured information.



When they select the **Complete Registration** button, they see the following screen that notifies them to check their email or phone for instructions on how to gain access to the network.



They are sent a username, password, and access instructions via an email or a phone text message.





When they connect to the Secure Wireless network, they will enter their username and password in this screen to gain access to the network.



For information on related help topics:

• Portal Configuration

How to Configure Verification for Guest Registration

Guest registration requires end users to enter their name and contact information on a Registration web page in order to gain access to the network. However, in many cases, end users provide false names and contact information because they don't want their personal information to be used for other purposes. In those cases, network administrators do not have a way to contact the user in the event of an Acceptable Use Policy (AUP) violation or in the case of an emergency.

With verification, guest end users registering to the network are required to enter a verification code that is sent to their email address or mobile phone (via SMS text) before gaining network access. This ensures that network administrators have at least one way to contact the end user.

Configuration Steps

The verification feature is supported for both Guest Registration and Guest Web Access, and is configured using the Verification Method options in your portal configuration. Depending on the verification method you specify, the appropriate custom fields must be configured for display on the Registration web page, so that end users can enter the required information.


The following table provides a description of each verification method and lists their custom field requirements.

User Verification Method	Description	Custom Field Requirement
Email	The end user must enter a valid email address on the Registration web page or Guest Web Access login page.	The Email Address Custom Field must be set to Required .
SMS Gateway	The SMS Gateway provider must support SMTP API. The SMS Gateway provider converts the email to an SMS text message. The end user must enter a mobile phone number on the Registration web page or Guest Web Access login page.	The Phone Number Custom Field must be set to Required .
SMS Gateway or Email	The SMS Gateway provider must support SMTP API. The SMS Gateway provider converts the email to an SMS text message. The end user must enter a mobile phone number or email address on the Registration web page or Guest Web Access login page.	The Phone Number and Email Address Custom Fields must be set to Visible.

User Verification Method	Description	Custom Field Requirement
SMS Text Message	The mobile provider converts the email to an SMS test message. The end user must enter a valid mobile phone number on the Registration web page or Guest Web Access login page.	The Phone Number Custom Field must be set to Required .
SMS Text or Email	The mobile provider converts the email to an SMS test message. The end user must enter a valid mobile phone number or email address on the Registration web page or Guest Web Access login page.	The Phone Number and Email Address Custom Fields must be set to Visible.

Use the following steps to configure verification in your portal configuration.

- 1. In ExtremeCloud IQ Site Engine, access the Portal Configuration. Select the Guest Registration or Guest Web Access selection in the Portal tree, depending on what access type your network is using. (If you don't see these selections, select Website Configuration in the tree and enable the appropriate feature.)
- 2. In the Guest Registration or Guest Web Access panel, use the drop-down list to select the desired Verification Method (refer to the <u>table</u> above). The Guest Registration panel is shown below.



3. If you selected the SMS Text Message or the SMS Text or Email User Verification method, select the Service Providers link to configure the list of mobile service providers from which end users can select the Registration web page or Guest Web Access login page. The Mobile Service Provider List provides a default list of providers that can be edited to include the appropriate service providers for your geographic location.

You can comment out entries by preceding each line with either a # or // to enable temporary editing of the file without removing the text.

The list requires one service provider entry per line, using the following format: <Provider>:phonenumber@<specificdomain>.

When the end user registers, they will see only the <Provider> portion in the drop-down list of providers on the Registration web page.

Select **OK** to close the window.

- 4. If you have selected the SMS Gateway or SMS Gateway or Email method, enter the SMS Gateway Email address provided by the SMS Gateway provider.
- 5. For all methods, select the Message Strings link to open the Message Strings Editor where you can customize the text displayed on the Registration web page or Guest Web Access login page, and the messages sent to the end user.

You need to modify different message strings sent to the end user, depending on the verification method or methods you selected. Double-click on the message to open a window where you can edit the message text.

• Email - This method uses the following strings:

registrationVerificationEmailMsgBody - the default message shouldn't need to be changed.

registrationVerificationEmailSentFromAddress - you need to change the default message to the appropriate email address for your company.

registrationVerificationEmailSentFromName - the default message shouldn't need to be changed.

registrationVerificationEmailSubject - the default message shouldn't need to be changed.

• SMS Gateway - Depending on your SMS Gateway provider and their required format, modify the following message strings using appropriate variables to customize the dynamic data such as phone number.

registrationVerificationSMSMsgBody

registrationVerificationSMSSubject

• SMS Text Message - This method uses the following strings. The default messages shouldn't need to be changed.

registration Verification SMSMsgBody

registrationVerificationSMSSubject

Select **OK** to close the window.


- 6. In the Web Page Customizations (Shared) section, select the Customize Fields link to open the Manage Custom Fields window.
- 7. Set the appropriate custom fields to display on the Registration web page or Guest Web Access login page, depending on the verification method you selected (refer to the <u>table</u> above). When you save your portal changes, the correct configuration of the custom fields are verified. These settings are shared by Guest Web Access, Guest Registration, and Secure Guest Access. Changing them for one access type also changes them for the others. For more information, see the Manage Custom Fields Window.

Select **OK** to close the window.


8. Back in the Portal Configuration, select **Save** to save your changes. Close the Portal Configuration window. Enforce the new portal configuration to your engine(s). Verification is now configured for your guest registration.


How User Verification Works

When a guest attempts to access the network, the Registration web page or Guest Web Access login page asks for their email address and/or phone number and mobile service provider, along with their normal contact information.




When they select the **Complete Registration** button, they are sent a verification code via an email or a phone text message.





The web page then prompts them for the code. When they enter the correct code that was generated for them and select the **Complete Registration** button, they are permitted access to the network. The verification code is valid for 15 minutes and cannot be reused after it is validated.




Configure Sponsorship for Guest Registration

This topic describes how to configure sponsorship for Guest Registration and Secure Guest Access. Sponsorship is configured as part of your portal configuration, and is accessed from the Guest Registration and Secure Guest Access views in the Portal section of the Portal Configuration panel.

With sponsored registration, end users are only allowed to register to the network when approved by a "sponsor," an internal trusted user to the organization. Sponsorship can provide the end user with a higher level of access than just guest access and allows the sponsor to fine-tune the level of access for individual end users. The end user registers and declares a sponsor's email address. The sponsor is notified and approves the registration, and can assign an elevated level of access, if desired.

To configure sponsorship:

- 1. Access the Control > Access Control tab.
- 2. In the left-panel tree, expand the **Access Control** Configurations > Portal and select the Guest Registration view or the Secure Guest Access view (depending on the access type you are configuring). The screenshot below shows the Guest Registration view.



- 3. In the Sponsorship section, select the **Sponsorship Mode** required. Additional settings display when you select optional or required sponsorship.
 - None Sponsorship is not required and the end user is assigned to the Registered Guests End-System Group.
 - Optional The end user is assigned to the Registered Guests End-System Group until sponsored. At that time, the sponsor can assign elevated access, if desired.
 - Required The end user has no access until the sponsor approves the registration. The end user is added to the Registration Pending Access end-system group and is presented the sponsorship pending page until approved.
- 4. **Sponsored Registration Introduction** Select the **Edit** button to open a window where you can edit the introductory message displayed to the end user.
- 5. Admin/Sponsor Email Enter the person or group to notify when an end user requests sponsorship, typically the network ExtremeControl administrator, for example "IT@CompanyA.com." This email address is always notified, in addition to the sponsor email address entered by the end user when they register to the network.

- 6. **Sponsor Email Field** Select an option for the sponsor email field on the registration web page.
 - **Do Not Display** The field is not displayed, and the end user is not required to enter a sponsor email address. In this case, only the admin/sponsor email address (defined above) is notified when the end user registers.
 - **Display Predefined Sponsor List** The end user must select a sponsor email from a list of predefined sponsors (defined below). The end user sees a drop-down list of sponsor email addresses and select the appropriate sponsor.
 - User Specifies Any Email as Sponsor The end user can enter any email address as a sponsor's email address.
 - User Must Specify Predefined Sponsor Email The end user must enter an email address that matches one of the predefined sponsors (defined below).
- 7. Predefined Sponsors Enter one or more sponsor email addresses. If you have selected Display Predefined Sponsor List as your Sponsor Email Field option (above), these addresses are presented to the end user as a drop-down list, allowing them to select a sponsor email address. If you have selected User Must Specify Predefined Sponsor Email as your Sponsor Email Field option, then the sponsor email address entered by the end user must match an email address listed here. Email addresses can be separated by semi-colons (;) or commas (,) for example, jdoe@CompanyA.com;rsmith@CompanyA.com. Because commas are accepted separators, they should not be used in actual email addresses.
- 8. In the Portal Configuration window, select **Save** to save your changes. You need to enforce the new portal configuration to your engine(s).

For information on related help topics:

• Portal Configuration

How to Implement Facebook Registration

This Help topic describes the steps for implementing guest registration using Facebook as a way to obtain end user information.

In this scenario, the Guest Registration portal provides the option to register as a guest or log into Facebook in order to complete the registration process. If the end user selects the Facebook option, ExtremeCloud IQ - Site Engine OAuth to securely access the end user's Facebook account, obtain public end user data, and use that data to complete the registration process.

NOTE: Guest OAuth (for example, Google, Yahoo) may not support native mobile browsers and display a "user agent" error. To access the network, use a standard browser application (e.g. Google Chrome).

Guest Registration using Facebook has two main advantages:

- It provides ExtremeCloud IQ Site Engine with a higher level of user information by obtaining information from the end user's Facebook account instead of relying on information entered by the end user.
- It provides an easier registration process for the end user. ExtremeCloud IQ Site Engine retrieves the public information from the end user's Facebook account and uses that information to populate the name and email registration fields.

This topic includes information and instructions on:

- Requirements for Facebook Registration
- Creating a Facebook Application
- Portal Configuration for Facebook
- How Facebook Registration Works
- Special Deployment Considerations
 - Networks using DNS Proxy

Requirements

These are the configuration requirements for Facebook Registration.


- The ExtremeControl engine must have Internet access in order to retrieve user information from Facebook.
- The ExtremeControl Unregistered access policy must provide access to the Facebook site (either enable all SSL or make allowances for Facebook servers).
- A Unique Facebook application must be created on the Facebook Developers page (see instructions below).

• The Portal Configuration must have Facebook Registration enabled and include the Facebook Application ID and Secret (see instructions below).

Creating a Facebook Application


When implementing guest registration using Facebook, you must first create a Facebook application. This generates an Application ID and Application Secret that are required as part of the ExtremeCloud IQ - Site Engine OAuth process. Use the following steps to create a Facebook application.

- 1. Access the Facebook Developers page at https://developers.facebook.com/apps/. If you already have a Developers account you can log in, otherwise you must create a Developers account.
- 2. When logged in, select the Add a New App button.



The Add a New App window opens.


3. Select the **basic setup** link at the bottom of the window.



The Create a New App ID window opens.


4. Enter a **Display Name**, enter a **Contact Email**, and select a **Category** for your app.

The **Display Name** is the name of the app presented to the end-user when they grant ExtremeCloud IQ - Site Engine access to their Facebook information and should clearly indicate what its purpose is, for example, Extreme Networks Guest Registration.




5. Select Create App ID.

The Dashboard panel opens and displays information about the new app including an App ID and an App Secret.




6. Select **Settings** in the left panel.

The Settings panel's **Basic** tab opens.



- 7. Enter in a valid domain name for the ExtremeControl engines in the **App Domains** field. For example, if the ExtremeControl engine to which users are connecting is ExtremeControl engine. AbcCompany.com, enter "abccompany.com" in the **App Domains** field.
- 8. Select Add Platform.

The Select Platform window opens.




9. Select Website.

The Website panel displays on the **Basic** tab.




- 10. Enter the domain name you added in the **App Domains** field in step 7 in the **Site URL** field.
- 11. Select Save Changes.
- 12. Select **Add Product** in the left panel.

The Product Setup panel opens.



13. Select the Facebook Login **Get Started** button.

The Getting Started panel opens.




14. Enter the Valid OAuth redirect URIs. A redirect URI is required to redirect the user back to the engine with an Access Token ExtremeCloud IQ - Site Engine uses to access the user account and retrieve the user data. The Redirection URI should be in the following format:

https://<ExtremeControlengineFQDN>/fb oauth

A Redirection URI must be added for each ExtremeControl engine where end users can register via Facebook.

- 15. Select Save Changes.
- 16. Select **App Review** in the left panel.

The App Review panel opens.



17. Select the No button in the Make < Display Name > public field to change the button to Yes.

A Confirmation window displays.

18. Select Confirm.

The Approved Items section displays a list of default permissions that provide access to end user data. (For more information on setting permissions, see https://developers.facebook.com/docs/facebook-login/permissions#reference.)


Your application is created and ready to use.

You need to add the App ID and App Secret to your portal configuration.

Portal Configuration

The Application ID and Application Secret assigned during the creation of the Facebook application must be provided in the Portal Configuration in order for the entire process to complete properly.

- 1. Open the Control > Access Control tab.
- 2. In the left-panel tree, expand the ExtremeControl Configurations > Portal tree and select Guest Registration.



- 3. In the Customize Fields section, select the **Open Editor** button to open the Manage Custom Fields window where you can change registration portal fields. Facebook registration uses only the First Name, Last Name, and Email Address fields, and the Display Acceptable Use Policy (AUP) option. All other fields only apply to regular guest registration. If the Display AUP option is selected, the captive portal verifies that the AUP has been acknowledged before redirecting the user to Facebook.
- 4. Select the Facebook Registration checkbox.
- 5. Enter the Facebook App ID and Facebook App Secret.
- 6. Select **Save**. Warning messages display stating that Verification Method and Sponsorship are not used for Facebook registration, and that an FDQN is required will be enabled.

7. Enforce the new configuration to your engines.

How Facebook Registration Works

After you have configured Facebook registration using the steps above, this is how the registration process works:

- 1. The end user attempts to access an external Web site. Their HTTP traffic is redirected to the captive portal.
- 2. In the Guest Registration Portal, the end user selects the option to register using Facebook.
- 3. The end user is redirected to the Facebook login. If Acceptable Use Policy option is configured, the captive portal verifies that the AUP has been acknowledged before redirecting the user to Facebook.
- 4. When logged in, the end user is presented with the information that ExtremeCloud IQ Site Engine receives from Facebook.
- 5. The end user grants ExtremeCloud IQ Site Engine access to the Facebook information and is redirected back to the captive portal where they see a "Registration in Progress" message.
- 6. Facebook provides the requested information to ExtremeCloud IQ Site Engine, which uses it to populate the user registration fields.
- 7. The registration process completes and network access is granted.
- 8. The word "Facebook" is added to the user name so you can easily search for Facebook registration via the Registration Administration web page.

Special Deployment Considerations

Read the following deployment consideration prior to configuring Facebook Registration.

Wireless Clients

To provide access to your network via a wireless connection, create an L7 host record for the **Unregistered Role** on your Wireless Controller for facebook.com. This domain is subject to change and can vary based on location.

Networks using DNS Proxy

Facebook Registration for networks redirecting HTTP traffic to the captive portal using DNS Proxy requires additional configuration.

In order for Facebook Registration to work properly with DNS Proxy, **all** domains/URLs necessary to properly load the Facebook web page must be added to the Allowed URLs/Allowed Domains section of the captive portal configuration. Otherwise, the ExtremeControl engine resolves DNS queries for these components to the ExtremeControl engine IP causing the page to not load properly.

As of July 26, 2014, you must add the following domains in order for Facebook registration to work with DNS Proxy. These domains are subject to change and can vary based on location.

Facebook.com fbstatic-a.akamaihd.net fbcdn-profile-a.akamaihd.net fbcdn-photos-c-a.akamaihd.net

• Portal Configuration

How to Implement Google Registration

This Help topic describes the steps for implementing guest registration using Google as a way to obtain end user information.

In this scenario, the Guest Registration portal provides the option to register as a guest or log into Google in order to complete the registration process. If the end user selects the Google option, ExtremeCloud IQ - Site Engine OAuth to securely access the end user's Google account, obtain public end user data, and use that data to complete the registration process.

NOTE: Guest OAuth (for example, Google, Yahoo) may not support native mobile browsers and display a "user agent" error. To access the network, use a standard browser application (e.g. Google Chrome).

Guest Registration using Google has two main advantages:

- It provides ExtremeCloud IQ Site Engine with a higher level of user information by obtaining information from the end user's Google account instead of relying on information entered by the end user.
- It provides an easier registration process for the end user. ExtremeCloud IQ Site Engine retrieves the public information from the end user's Google account and uses that information to populate the name and email registration fields.

This topic includes information and instructions on:

- Requirements for Google Registration
- Creating a Google Application
- Portal Configuration for Google
- How Google Registration Works
- Special Deployment Considerations
 - Networks using DNS Proxy

Requirements

These are the configuration requirements for Google Registration.


- The ExtremeControl engine must have Internet access in order to retrieve user information from Google.
- The ExtremeControl Unregistered access policy must allow access to the Google site (either allow all SSL or make allowances for Google servers).
- The ExtremeControl Unregistered access policy must allow access to HTTPS traffic to the Google OAuth servers.
- A Unique Google application must be created on the Google Developers page (see instructions below).

• The Portal Configuration must have Google Registration enabled and include the Google Application ID and Secret (see instructions below).

Creating a Google Application


When implementing guest registration using Google, you must first create a Google application. This generates an Application ID and Application Secret that are required as part of the ExtremeCloud IQ - Site Engine OAuth process. Use the following steps to create a Google application.

- 1. Access the Google Developers page at https://console.developers.google.com/projectselector/apis/library.
- 2. Log into your existing Developers account or create a new Developers account.
- 3. Select the Create a project button.




The New Project window opens.

- 4. Enter a **Project name** and select **Create**.
- 5. Select the **Credentials** link in the left-panel.




The Credentials panel opens.

6. Select the Create credentials button to open the drop-down list and select OAuth client ID.




The Create client ID panel displays.

7. Select Configure consent screen to open the OAuth consent screen panel.




8. Select your email address, enter your product name, and enter the URL to any of the applicable resources for your company, then select **Save**.

The Create client ID panel opens.



9. Select Web application.

The panel expands to display additional fields.



- 10. Enter a name for the application in the **Name** field. Use a name that clearly indicates what its purpose is, for example, Extreme Networks Guest Registration.
- 11. Enter an Authorized redirect URI in the following format https://<AccessControlengineFQDN>/google_oauth. Google uses the Authorized redirect URI to redirect the user back to the engine with an Access Token.

NOTES: Google OAuth APIs require your engine's FQDN resolves to a top level domain (.com, .net, .edu, .org, .mil, .gov, or .int. You cannot use a domain not classified as top level (e.g. MyGateway.MyCompany.Local) or the engines IP address, which can require you to reclassify your domain and hosts.

Use only lowercase when entering the host and domain suffix (e.g. .com).

- 12. Enter the **Authorized redirect URI** for any additional ExtremeControl engines registering end-users via Google.
- 13. Select Create.

The **OAuth client** window displays, displaying your client ID and secret.




Your application is created and ready to use.

You need to add the client ID and client secret to your portal configuration.

Portal Configuration

The client ID and client secret assigned during the creation of the Google application must be provided in the Portal Configuration in order for the entire process to complete properly.

- 1. Open the **Control** > **Access Control** tab.
- 2. In the left-panel tree, expand the Configuration > Captive Portals > Website Configuration > and select Guest Registration.



- 3. In the Customize Fields section, select the **Open Editor** button to open the Manage Custom Fields window where you can change registration portal fields. Google registration uses only the First Name, Last Name, and Email Address fields, and the Display Acceptable Use Policy (AUP) option. All other fields only apply to regular guest registration. If the Display AUP option is selected, the captive portal verifies that the AUP has been acknowledged before redirecting the user to Google.
- 4. Select the Google Registration checkbox.
- 5. Select Edit.
- 6. Enter the client ID in the Google App ID field and the client secret in the App Secret field.
- 7. Select **Save**. Warning messages display stating that Verification Method and Sponsorship are not used for Google registration, and that an FDQN is required will be enabled.
- 8. Enforce the new configuration to your engines.

How Google Registration Works

After you have configured Google registration using the steps above, this is how the registration process works:

- 1. The end user attempts to access an external Web site. Their HTTP traffic is redirected to the captive portal.
- 2. In the Guest Registration Portal, the end user selects the option to register using Google.
- 3. The end user is redirected to the Google login. If Acceptable Use Policy option is configured, the captive portal verifies that the AUP has been acknowledged before redirecting the user to Google.
- 4. When logged in, the end user is presented with the information that ExtremeCloud IQ Site Engine receives from Google.
- 5. The end user grants ExtremeCloud IQ Site Engine access to the Google information and is redirected back to the captive portal where they see a "Registration in Progress" message.
- 6. Google provides the requested information to ExtremeCloud IQ Site Engine, which uses it to populate the user registration fields.
- 7. The registration process completes and network access is granted.
- 8. The word "Google" is added to the user name so you can easily search for Google registration via the Registration Administration web page.

Special Deployment Considerations

Read the following deployment consideration prior to configuring Google Registration.

To allow traffic to your network via a wireless connection, create an L7 host record for the **Unregistered Role** on your Wireless Controller for accounts.google.com and gstatic.com. These domains are subject to change and can vary based on location.

Networks using DNS Proxy

Google Registration for networks redirecting HTTP traffic to the captive portal using DNS Proxy requires additional configuration.

In order for Google Registration to work properly with DNS Proxy, **all** domains/URLs necessary to properly load the Google web page must be added to the Allowed URLs/Allowed Domains section of the captive portal configuration. Otherwise, the ExtremeControl engine resolves DNS queries for these components to the ExtremeControl engine IP causing the page to not load properly.

As of February 2017, you must add the following domains in order for Google registration to work with DNS Proxy. This domain is subject to change and can vary based on location.

Accounts.google.com

• Portal Configuration

How to Implement Microsoft Registration

This Help topic describes the steps for implementing guest registration using Microsoft as a way to obtain end user information.

In this scenario, the Guest Registration portal provides the option to register as a guest or log into Microsoft in order to complete the registration process. If the end user selects the Microsoft option, ExtremeCloud IQ - Site Engine OAuth to securely access the end user's Microsoft account, obtain public end user data, and use that data to complete the registration process.

NOTE: Guest OAuth (for example, Google, Yahoo) may not support native mobile browsers and display a "user agent" error. To access the network, use a standard browser application (e.g. Google Chrome).

Guest Registration using Microsoft has two main advantages:

- It provides ExtremeCloud IQ Site Engine with a higher level of user information by obtaining information from the end user's Microsoft account instead of relying on information entered by the end user.
- It provides an easier registration process for the end user. ExtremeCloud IQ Site Engine retrieves the public information from the end user's Microsoft account and uses that information to populate the name and email registration fields.

This topic includes information and instructions on:

- Requirements for Microsoft Registration
- Creating a Microsoft Application
- Portal Configuration for Microsoft
- How Microsoft Registration Works
- Special Deployment Considerations
 - Networks using DNS Proxy

Requirements

These are the configuration requirements for Microsoft Registration.


- The ExtremeControl engine must have Internet access in order to retrieve user information from Microsoft.
- The ExtremeControl Unregistered access policy must provide access to the Microsoft site (either enable all SSL or make allowances for Microsoft servers).
- The ExtremeControl Unregistered access policy must provide access to HTTPS traffic to the Microsoft OAuth servers.

- A Unique Microsoft application must be created on the Microsoft Developers page (see instructions below).
- The Portal Configuration must have Microsoft Registration enabled and include the Microsoft Application ID and Secret (see instructions below).


Creating a Microsoft Application

When implementing guest registration using Microsoft, you must first create a Microsoft application. This generates an Application ID and Application Secret that are required as part of the ExtremeCloud IQ - Site Engine OAuth process. Use the following steps to create a Microsoft application.

- 1. Access the Microsoft Developers page at https://apps.dev.microsoft.com/#/appList.
- 2. Log into your existing account or create a new account by selecting the **Sign in** link in the top-right corner of the window.




3. Select the **Add an app** button.



The New Application Registration window opens.

4. Enter a **Name** for the application. Use a name that clearly indicates it's purpose (e.g. Extreme Networks Guest Registration) and select **Create application**.

The Application Registration window opens.




5. Select Add Platforms under Platforms.

The Add Platform window opens.



6. Select Web.

Additional fields display under Platforms enabling you to configure a web platform.




7. Enter a **Redirect URI** in the following format https://<AccessControlengineFQDN>/ms_oauth. Microsoft uses the **Redirect URI** to redirect the user back to the engine with an Access Token.


NOTE: Microsoft applications can only use a limited set of <u>redirect URI values</u>.

8. Select **Add Url** to enter the **Redirect URI** for any additional ExtremeControl engines registering endusers via Microsoft.

9. Copy the Application Id under Properties.




10. Select Generate New Password under Application Secrets.



The New password generated window displays.

11. Copy the application password.

IMPORTANT: Ensure you copy the password accurately. After the window is closed, you cannot access the password again.



12. Select Save.


Your application is created and ready to use.

You need to add the **Application Id** and application password to your portal configuration.

Portal Configuration

The Application Id and application password assigned during the creation of the Microsoft application must be provided in the Portal Configuration in order for the entire process to complete properly.

- 1. Open the **Control** > **Access Control** tab.
- 2. In the left-panel tree, expand the ExtremeControl Configurations > Portal tree and select Guest Registration.



- 3. In the Customize Fields section, select the **Open Editor** button to open the Manage Custom Fields window where you can change registration portal fields. Microsoft registration uses only the First Name, Last Name, and Email Address fields, and the Display Acceptable Use Policy (AUP) option. All other fields only apply to regular guest registration. If the Display AUP option is selected, the captive portal verifies that the AUP has been acknowledged before redirecting the user to Microsoft.
- 4. Select the **Microsoft Registration** checkbox.
- 5. Select **Edit**.
- 6. Enter the Application Id in the **Microsoft App ID** field and the application password in the **Microsoft App Secret** field.
- 7. Select **Save**. Warning messages display stating that Verification Method and Sponsorship are not used for Microsoft registration, and that an FDQN is required and will be enabled.
- 8. Enforce the new configuration to your engines.

How Microsoft Registration Works

After you have configured Microsoft registration using the steps above, this is how the registration process works:

- 1. The end user attempts to access an external Web site. Their HTTP traffic is redirected to the captive portal.
- 2. In the Guest Registration Portal, the end user selects the option to register using Microsoft.
- 3. The end user is redirected to the Microsoft login. If Acceptable Use Policy option is configured, the captive portal verifies that the AUP has been acknowledged before redirecting the user to Microsoft.
- 4. When logged in, the end user is presented with the information that ExtremeCloud IQ Site Engine receives from Microsoft.
- 5. The end user grants ExtremeCloud IQ Site Engine access to the Microsoft information and is redirected back to the captive portal where they see a "Registration in Progress" message.
- 6. Microsoft provides the requested information to ExtremeCloud IQ Site Engine, which uses it to populate the user registration fields.
- 7. The registration process completes and network access is granted.
- 8. The word "Microsoft" is added to the user name so you can easily search for Microsoft registration via the Registration Administration web page.

Special Deployment Considerations

Read the following deployment consideration prior to configuring Microsoft Registration.

To provide access to your network via a wireless connection, create an L7 host record for the **Unregistered Role** on your Wireless Controller for login.live.com and auth.gfx.ms. These domains are subject to change and can vary based on location.

Networks using DNS Proxy

Microsoft Registration for networks redirecting HTTP traffic to the captive portal using DNS Proxy requires additional configuration.

In order for Microsoft Registration to work properly with DNS Proxy, **all** domains/URLs necessary to properly load the Microsoft web page must be added to the Allowed URLs/Allowed Domains section of the captive portal configuration. Otherwise, the ExtremeControl engine resolves DNS queries for these components to the ExtremeControl engine IP causing the page to not load properly.

As of February 2017, you must add the following domains in order for Microsoft registration to work with DNS Proxy. These domains are subject to change and can vary based on location.

Login.live.com

• Portal Configuration

How to Implement Yahoo Registration

This Help topic describes the steps for implementing guest registration using Yahoo as a way to obtain end user information.

In this scenario, the Guest Registration portal provides the option to register as a guest or log into Yahoo in order to complete the registration process. If the end user selects the Yahoo option, ExtremeCloud IQ - Site Engine OpenID to securely access the end user's Yahoo account, obtain public end user data, and use that data to complete the registration process.

NOTE: Guest OAuth (for example, Google, Yahoo) may not support native mobile browsers and display a "user agent" error. To access the network, use a standard browser application (e.g. Google Chrome).

Guest Registration using Yahoo has two main advantages:

- It provides ExtremeCloud IQ Site Engine with a higher level of user information by obtaining information from the end user's Yahoo account instead of relying on information entered by the end user.
- It provides an easier registration process for the end user. ExtremeCloud IQ Site Engine retrieves the public information from the end user's Yahoo account and uses that information to populate the name and email registration fields.

This topic includes information and instructions on:

- Requirements for Yahoo Registration
- Creating a Yahoo Application
- Portal Configuration for Yahoo
- How Yahoo Registration Works
- Special Deployment Considerations
 - Networks using DNS Proxy

Requirements

These are the configuration requirements for Yahoo Registration.


- The ExtremeControl engine must have Internet access in order to retrieve user information from Yahoo.
- The ExtremeControl Unregistered access policy must provide access to the Yahoo site (either enable all SSL or make allowances for Yahoo servers).
- The ExtremeControl Unregistered access policy must provide access to HTTPS traffic to the Yahoo OpenID servers.
- A Unique Yahoo application must be created on the Yahoo Developers page (see instructions below).

• The Portal Configuration must have Yahoo Registration enabled and include the Yahoo Application ID and Secret (see instructions below).

Creating a Yahoo Application

When implementing guest registration using Yahoo, you must first create a Yahoo application. This generates an Application ID and Application Secret that are required as part of the ExtremeCloud IQ - Site Engine OpenID process. Use the following steps to create a Yahoo application.

- 1. Log into your existing account or create a new account.
- 2. Access the Create Application page at https://developer.yahoo.com/apps/create/.




- 3. Enter a name for the application in the **Application Name** field. Use a name that clearly indicates what its purpose is, for example, Extreme Networks Guest Registration.
- 4. Select Web Application for the Application Type.
- 5. Enter an Callback Domain in the following format https://<accessControlengineFQDN>. Yahoo uses the Callback Domain to redirect the user back to the engine with an Access Token.

NOTES: Yahoo OAuth APIs require your engine's FQDN resolves to a top level domain (.com, .net, .edu, .org, .mil, .gov, or .int. You cannot use a domain not classified as top level (e.g. MyGateway.MyCompany.Local) or the engines IP address, which can require you to reclassify your domain and hosts.

Use only lowercase when entering the host and domain suffix (e.g. .com).

6. Select Create App.

The Client ID and Client Secret display at the top of the window.




Your application is created and ready to use.

You need to add the client ID and client secret to your portal configuration.

Portal Configuration

The client ID and client secret assigned during the creation of the Yahoo application must be provided in the Portal Configuration in order for the entire process to complete properly.

- 1. Open the **Control** > **Access Control** tab.
- 2. In the left-panel tree, expand the Configuration > Captive Portals > Website Configuration > and select Guest Registration.



- 3. In the Customize Fields section, select the **Open Editor** button to open the Manage Custom Fields window where you can change registration portal fields. Yahoo registration uses only the First Name, Last Name, and Email Address fields, and the Display Acceptable Use Policy (AUP) option. All other fields only apply to regular guest registration. If the Display AUP option is selected, the captive portal verifies that the AUP has been acknowledged before redirecting the user to Yahoo.
- 4. Select the Yahoo Registration checkbox.
- 5. Select Edit.
- 6. Enter the Client ID in the App ID field and the Client Secret in the App Secret field.
- 7. Select **Save**. Warning messages display stating that Verification Method and Sponsorship are not used for Yahoo registration, and that an FDQN is required will be enabled.
- 8. Enforce the new configuration to your engines.

How Yahoo Registration Works

After you have configured Yahoo registration using the steps above, this is how the registration process works:

- 1. The end user attempts to access an external Web site. Their HTTP traffic is redirected to the captive portal.
- 2. In the Guest Registration Portal, the end user selects the option to register using Yahoo.
- 3. The end user is redirected to the Yahoo login. If Acceptable Use Policy option is configured, the captive portal verifies that the AUP has been acknowledged before redirecting the user to Yahoo.
- 4. When logged in, the end user is presented with the information that ExtremeCloud IQ Site Engine receives from Yahoo.
- 5. The end user grants ExtremeCloud IQ Site Engine access to the Yahoo information and is redirected back to the captive portal where they see a "Registration in Progress" message.
- 6. Yahoo provides the requested information to ExtremeCloud IQ Site Engine, which uses it to populate the user registration fields.
- 7. The registration process completes and network access is granted.
- 8. The word "Yahoo" is added to the user name so you can easily search for Yahoo registration via the Registration Administration web page.

Special Deployment Considerations

Read the following deployment consideration prior to configuring Yahoo Registration.

To provide access to your network via a wireless connection, create an L7 host record for the **Unregistered Role** on your Wireless Controller for login.yahoo.com. This domain is subject to change and can vary based on location.

Networks using DNS Proxy

Yahoo Registration for networks redirecting HTTP traffic to the captive portal using DNS Proxy requires additional configuration.

In order for Yahoo Registration to work properly with DNS Proxy, **all** domains/URLs necessary to properly load the Yahoo web page must be added to the Allowed URLs/Allowed Domains section of the captive portal configuration. Otherwise, the ExtremeControlengine resolves DNS queries for these components to the ExtremeControlengine IP causing the page to not load properly.

As of February 2017, you must add the following domains in order for Yahoo registration to work with DNS Proxy. This domain is subject to change and can vary based on location.

login.yahoo.com

• Portal Configuration

How to Implement Salesforce Registration

This Help topic describes the steps for implementing guest registration using Salesforce as a way to obtain end user information.

In this scenario, the Guest Registration portal provides the option to register as a guest or log into Salesforce in order to complete the registration process. If the end user selects the Salesforce option, ExtremeCloud IQ - Site Engine uses OpenID to securely access the end user's Salesforce account, obtain public end user data, and use that data to complete the registration process.

NOTE: Guest OAuth (for example, Google, Yahoo) may not support native mobile browsers and display a "user agent" error. To access the network, use a standard browser application (e.g. Google Chrome).

Guest Registration using Salesforce has two main advantages:

- It provides ExtremeCloud IQ Site Engine with a higher level of user information by obtaining information from the end user's Salesforce account instead of relying on information entered by the end-user.
- It provides an easier registration process for the end user. ExtremeCloud IQ Site Engine retrieves the public information from the end user's Salesforce account and uses that information to populate the name and email registration fields.

This topic includes information and instructions on:

- Requirements for Salesforce Registration
- Creating a Salesforce Application
- Portal Configuration for Salesforce
- How Salesforce Registration Works
- Special Deployment Considerations
 - Networks using DNS Proxy

Requirements

These are the configuration requirements for Salesforce Registration.


- The ExtremeControl engine must have Internet access in order to retrieve user information from Salesforce.
- The ExtremeControl Unregistered access policy must provide access to the Salesforce site (either enable all SSL or make allowances for Salesforce servers).
- The ExtremeControl Unregistered access policy must provide access to HTTPS traffic to the Salesforce OpenID servers.

- A Unique Salesforce application must be created on the Salesforce Developers page (see instructions below).
- The Portal Configuration must have Salesforce Registration enabled and include the Salesforce Application ID and Secret (see instructions below).

Creating a Salesforce Application


When implementing guest registration using Salesforce, you must first create a Salesforce application. This generates an Application ID and Application Secret that are required as part of the ExtremeCloud IQ - Site Engine OpenID process. Use the following steps to create a Salesforce application.

- 1. Access the Salesforce Developers page at https://developer.salesforce.com/signup.
- 2. Log into your existing Developers account or create a new Developers account.
- 3. Select the My Developer Account button from the profile drop-down list.




The Developer Account login window opens.

4. Select your account.




The **Developer Home** window opens.

5. Select My Settings from the Profile drop-down list.




The My Settings window opens.

6. Select App & Administration Setup.




The App & Administration Setup window opens.

7. Select **Apps** from within the Build > Create menu.



The **Apps** window opens.

8. Select the **New** button in the Connected Apps section.



The New Connected App window opens.

9. Enter a Connected App Name, API Name, Contact Email, and select the Enable OAuth Settings checkbox.


The API (Enable OAuth Settings) section of the window expands to display additional fields.

- 10. Select Enable OAuth Settings.
- 11. Enter a Callback URL in the following format https://<AccessControlengineFQDN>/Salesforce_oauth. Salesforce uses the Authorized redirect URI to redirect the user back to the engine with an Access Token.

NOTES: Salesforce OpenID APIs require your engine's FQDN resolves to a top level domain (.com, .net, .edu, .org, .mil, .gov, or .int. You cannot use a domain not classified as top level (e.g. MyGateway.MyCompany.Local) or the engines IP address, which can require you to reclassify your domain and hosts.

Use only lowercase when entering the host and domain suffix (e.g. .com).


- 12. Select Access your basic information (id, profile, email, address, phone), Full access (full), and Allow access to your unique identifier (openid), then select the Add icon in the Selected OAuth Scopes section of the window to add the scopes to the Selected OAuth Scopes list.
- 13. Select the Require Secret for Web Server Flow, Include ID Token and Include Standard Claims checkboxes.



14. Select Save.

Your application is created and ready to use.

The **New Connected App** window opens.



15. Select Continue.

The Connected App window opens.

16. Select the Click to reveal link in the Consumer Secret field and copy the Consumer Secret and Consumer Key.




You need to add the Consumer Key and Consumer Secret to your portal configuration.

Portal Configuration

The client ID and client secret assigned during the creation of the Salesforce application must be provided in the Portal Configuration in order for the entire process to complete properly.

- 1. Open the **Control** > **Access Control** tab.
- 2. In the left-panel tree, expand the ExtremeControl Configurations > Portal tree and select Guest Registration.



- 3. In the Customize Fields section, select the **Open Editor** button to open the Manage Custom Fields window where you can change registration portal fields. Salesforce registration uses only the First Name, Last Name, and Email Address fields, and the Display Acceptable Use Policy (AUP) option. All other fields only apply to regular guest registration. If the Display AUP option is selected, the captive portal verifies that the AUP has been acknowledged before redirecting the user to Salesforce.
- 4. Select the Salesforce Registration checkbox.
- 5. Enter the Consumer Key in the App ID field and the Consumer Secret in the App Secret field.
- 6. Select **Save**. Warning messages display stating that Verification Method and Sponsorship are not used for Salesforce registration, and that an FDQN is required will be enabled.
- 7. Enforce the new configuration to your engines.

How Salesforce Registration Works

After you have configured Salesforce registration using the steps above, this is how the registration process works:

- 1. The end user attempts to access an external Web site. Their HTTP traffic is redirected to the captive portal.
- 2. In the Guest Registration Portal, the end user selects the option to register using Salesforce.
- 3. The end user is redirected to the Salesforce login. If Acceptable Use Policy option is configured, the captive portal verifies that the AUP has been acknowledged before redirecting the user to Salesforce.
- 4. When logged in, the end user is presented with the information that ExtremeCloud IQ Site Engine receives from Salesforce.
- 5. The end user grants ExtremeCloud IQ Site Engine access to the Salesforce information and is redirected back to the captive portal where they see a "Registration in Progress" message.
- 6. Salesforce provides the requested information to ExtremeCloud IQ Site Engine, which uses it to populate the user registration fields.
- 7. The registration process completes and network access is granted.
- 8. The word "Salesforce" is added to the user name so you can easily search for Salesforce registration via the Registration Administration web page.

Special Deployment Considerations

Read the following deployment consideration prior to configuring Salesforce Registration.

To provide access to your network via a wireless connection, create an L7 host record for the **Unregistered Role** on your Wireless Controller for login.Salesforce.com. This domain is subject to change and can vary based on location.

Networks using DNS Proxy

Salesforce Registration for networks redirecting HTTP traffic to the captive portal using DNS Proxy requires additional configuration.

In order for Salesforce Registration to work properly with DNS Proxy, **all** domains/URLs necessary to properly load the Salesforce web page must be added to the Allowed URLs/Allowed Domains section of the captive portal configuration. Otherwise, the ExtremeControl engine resolves DNS queries for these components to the ExtremeControl engine IP causing the page to not load properly.

As of February 2017, you must add the following domains in order for Salesforce registration to work with DNS Proxy. This domain is subject to change and can vary based on location.

login.Salesforce.com

• Portal Configuration

Add/Edit MAC Lock

Use this window to add a new locked MAC address or edit the settings for an existing locked MAC address. MAC Locking lets you lock a MAC address to a specific switch or port on a switch so that the end-system can only access the network from that port or switch. If the end-system tries to authenticate on a different switch/port, it is rejected or assigned a specific policy. You can add or edit MAC locks from the End-Systems tab.

NOTE: MAC Locking to a specific port on a switch is based on the port interface name (e.g. fe.5.1). If a switch board is moved to a different slot in a chassis, or if a stack reorders itself, this name changes and breaks the MAC Locking settings.



MAC Address

Enter the MAC address that you want to lock.

Switch IP

Enter the IP address of the switch on which you want to lock the MAC address.

Lock to Switch and Port

Select this checkbox if you want to lock the MAC address to a specific port on the switch, and enter the port interface name.

Failed Action

Select the action to take when this MAC address tries to authenticate on a different port and/or switch:

- Reject The authentication request is rejected.
- Use Policy Use the drop-down list to select the policy that you want applied. This policy must exist in the **Policy** tab and be enforced to the switches in your network.