

ExtremeAnalytics User Guide Version 8.5



Copyright © 2020 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.
- GTAC Knowledge Get on-demand and tested resolutions from the GTAC Knowledgebase, or create a help case if you need more guidance.
- The Hub 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.
- TERM. 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.

 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. <u>EXPORT REQUIREMENTS</u>. 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. <u>LIMITED WARRANTY AND LIMITATION OF LIABILITY</u>. 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

ExtremeAnalytics® User GuideVersion 8.5	1
	1
Legal Notices	2
Trademarks	2
Contact	2
Extreme Networks® Software License Agreement	4
Table of Contents	11
ExtremeAnalytics™ Help	21
Document Version	21
ExtremeAnalytics Licensing	23
Using Licenses to Establish Flow Rate Capacity	24
Calculating Flow Rate Violations	25
Calculating End-System Violations	25
Enforcing License Violation Limits	25
Clearing License Violations	26
Getting Started with ExtremeAnalytics	27
ExtremeAnalytics Access Requirements	27
ExtremeAnalytics Engine Configuration	27
Enable Flow Collection	27
Enable Jumbo Frames	28
Configuring Enhanced Netflow for Extreme Analytics and Extreme Wireless Controller Version 10.21	29
How to Deploy ExtremeAnalytics in an MSP or MSSP Environment	34
Configuring Extreme Management Center Behind a NAT Router	34

ExtremeAnalytics Application Data Collection	36
Data Collection Overview	36
Collection Targets	37
Collection Statistics	38
Collection Intervals	39
Using Sites to Collect In-Network Traffic	41
Data Collector Types	41
General Usage Collectors	42
Hourly General Usage Collectors	42
High-Rate General Usage Collectors	45
End-System Details Collector	45
Flow Information Sources	46
Enabling ExtremeControl Integration	47
Reports	48
Dashboard Report	48
Browser Reports	49
ExtremeAnalytics Tab Overview	51
Dashboard	51
Browser	51
Application Flows	52
Fingerprints	52
Packet Captures	52
Configuration	52
Reports	52
ExtremeAnalytics Dashboard Overview	54
Insights Dashboard Reports	54

Client/Server Dashboard Reports	55
Applications Browser Dashboard Report	55
Industry Dashboards	55
Enterprise Dashboard	55
Education Dashboard	55
Healthcare Dashboard	55
Venue Dashboard	56
IP Reputation Dashboard	56
Response Time Dashboard	57
Network Service Dashboard	57
Tracked Applications Dashboard	57
ExtremeAnalytics Insights Dashboard	59
Insights	59
Ring Chart	59
Custom Dashboard	60
How to Create an ExtremeAnalytics Insights Custom Dashboard	61
Custom Dashboard	61
Graphs	63
Usage	63
Performance	63
Trending	63
Analytics Events	64
ExtremeAnalytics Response Time Dashboard	64
Overview	65
Application	65
Тор	66

Tracked Applications	66
Filters	66
Network Response Time Graph	67
Application Response Time Graph	67
ExtremeAnalytics Network Service Dashboard	68
Overview	69
Expected Response Time	70
Historical Response Time	71
ExtremeAnalytics Tracked Applications Dashboard	72
Overview	72
Expected Response Time	73
Historical Response Time	75
ExtremeAnalytics Browser Overview	76
Overview	76
Data Aggregation	77
Options	78
Bookmark	81
Save to Report Designer	82
Export to CSV	82
ExtremeAnalytics Application Flows	83
Overview	83
Application Flows Tables	85
Historical Flows	85
Bidirectional Flows	85
Unidirectional Flows	85
Report Features	86

ExtremeAnalytics Bidirectional Flow Table	87
ExtremeAnalytics Unidirectional Flow Table	91
ExtremeAnalytics Historical Flow Table	93
ExtremeAnalytics Fingerprints Overview	97
ExtremeAnalytics Custom Fingerprints	98
Fingerprint Table	98
Menu	98
Column Definitions	99
Delete Custom Fingerprints	102
Deleting a Custom Fingerprint	102
Custom Fingerprint Examples	104
Fingerprints Based on a Flow	104
Fingerprints Based on an Application or Application Group	105
Fingerprints Based on a Destination Address	106
Create Custom Fingerprints Based on Flow	109
Creating Fingerprints Based on a Flow	109
Create Custom Fingerprints Based on Destination Address	111
Creating Fingerprints Based on a Destination Address	111
Create Custom Fingerprints Based on Application or Application Group	113
Creating Fingerprints Based on an Application or Application Group	113
ExtremeAnalytics Packet Captures	115
ExtremeAnalytics Configuration Overview	117
Engines	117
Status	118
Configuration	118
Virtual Sensors	118

Fingerprints	118
Licenses	119
Status	119
Configuration	120
Virtual Sensors	123
Virtual Sensors	123
Virtual Machines	123
ExtremeAnalytics Engine Advanced Configuration	124
Flow Collection Type	126
Collection Privacy Levels	126
Client Aggregation	127
Slow Client Data	127
Max End-Systems in Hourly Details	127
Sensor Log Levels	127
Store Application Site Data	128
ExtremeControl Integration	128
Flow Sources/Application Telemetry Sources	128
Web Credentials	129
Configuration Properties	130
Sensor Modules	130
Auditing	130
Network Settings	131
DNS	131
NTP	132
SSH	132
SNMP	134

Interfaces	134
ExtremeAnalytics Reports	136
Reports	136
ExtremeAnalytics Report Descriptions	138
Report Descriptions	138
Analytics Events	139
Bandwidth for a Client Over Time	139
Interface Top Applications Treemap	139
Sites Using the Most Bandwidth	139
Most Popular Applications	139
Most Used Applications for a Client	140
Most Used Applications for a User Name	140
Network Activity by Site	140
Network Activity by Client	140
Network Activity by Application	140
Slowest Applications by Site	140
Top Applications Group Radar	140
Top Applications Radar	141
Top Applications TreeMap	141
Top Applications for Interface	141
Top Applications for Server	141
Top Clients by Interface	141
Top Interfaces by Application	142
Top N Applications	142
Top N Clients	142
Top N Servers	143

Add and Modify Fingerprints	144
Adding a Fingerprint	144
Modifying a Fingerprint	146
Enabling or Disabling a Fingerprint	148
Deleting a Custom Fingerprint	148
Updating Fingerprints	149
Perform a Fingerprint Update	149
Schedule Fingerprint Updates	151
Add Fingerprints	153
Add a Fingerprint	153
Enable or Disable Fingerprints	156
Enabling or Disabling a Fingerprint	156
Modify Fingerprints	157
Modifying a Fingerprint	157
Update Fingerprints	160
Updating Fingerprints	160
Perform a Fingerprint Update	160
Schedule Fingerprint Updates	162
Custom Fingerprint Examples	164
Fingerprints Based on a Flow	164
Fingerprints Based on an Application or Application Group	165
Fingerprints Based on a Destination Address	166
How to Deploy ExtremeAnalytics in an MSP or MSSP Environment	169
Configuring Extreme Management Center Behind a NAT Router	169
ExtremeAnalytics Virtual Sensor Configuration in Extreme Management	Center 171
Prerequisites	172

Installing the Virtual Sensor Using the Extreme Management Center Server	174
Install Using Extreme Management Center	.175
Configuring the VMware vSphere Module in ExtremeConnect	175
Adding the Virtual Sensor to Extreme Management Center	. 177
Adding the Virtual Sensor in ExtremeAnalytics	. 178
Configuring vCenter Settings for the Virtual Sensor	. 180
Stream Flow Data from ExtremeAnalytics into Splunk	.182
Environment	. 182
Overview	. 182
Part 1 - Making File Level Splunk Modifications	. 183
Part 2 - Creating a New Stream using the Splunk web UI	. 183
Part 3 - Configuring each Analytics Engine to Export IPFIX Data to the Splunk Server	. 185
Appendix	. 185
Files	.185
\$SPLUNK/etc/apps/splunk_app_stream/default/vocabulary/extreme.xml	187
\$SPLUNK/etc/apps/splunk_app_stream/default/streams/netflow (additions)	. 188
Stream Flow Data from ExtremeAnalytics into Elastic Stack	. 191
Environment	. 191
Overview	. 191
Part 1 - Installing and Configuring ElastiFlow and Elastic Stack	. 192
Part 2 - Configuring each Analytics Engine to export IPFIX data to the Elastic Stack server	. 197
Appendix: Files	198
Additions to ipfix.yml in extr_elastiflow_3.4.2.tar.gz	198
Additions to elastiflow.template.json and elastiflow_dynamic.template.json in extr_elastiflow_3.4.2.tar.gz	199

Additions to elastiflow_static.template.json in extr_elastiflow_3.4.2.tar.gz2	20

ExtremeAnalytics™ Help

ExtremeAnalytics provides Layer 7 application visibility on your network. Combining Extreme Management Center, S-Series and/or K-Series devices, and the ExtremeAnalytics engine, this feature integrates application, user, and device data to give you a full understanding of the applications on your network and who's using those applications.

ExtremeAnalytics uses deep packet inspection (DPI) and a rich set of application fingerprinting techniques to provide granular control of private applications (SAP, SOA traffic, Exchange, SQL, etc.), public cloud applications (Salesforce, Google, Email, YouTube, P2P, file sharing, etc.), as well as social media applications (Facebook, Twitter, etc.), guaranteeing a quality user experience for business critical applications.

The combination of patent pending application flow sampling, flow statistics collection in hardware, and custom flow-based network processors (CoreFlow2) is at the heart of ExtremeAnalytics. ExtremeAnalytics samples up to 32 packets of every new flow. The CoreFlow2 ASIC identifies those new flows and sends the packets to the ExtremeAnalytics engine where it is combined with the non-sampled NetFlow traffic for the remainder of the flow, allowing the ExtremeAnalytics engine to process traffic at unprecedented scale. The ExtremeAnalytics engine determines the application, aggregates the data, adds additional context, and then sends it to Extreme Management Center for visibility.

In addition, you can collect application telemetry (SFlow) flows on ExtremeXOS devices. This allows you to use your ExtremeXOS devices as flow collectors, providing a more complete picture of the applications access on your network and the performance of those applications.

ExtremeAnalytics requires the NMS-ADV license.

Document Version

The following table displays the revision history for the ExtremeAnalytics Help documentation.

Date	Revision Number	Description
7-20	8.5 Revision -00	Extreme Management Center 8.5 release

Date	Revision Number	Description
06-17	8.0 Revision -00	Extreme Management Center 8.0 release
04-16	7.0 Revision -00	Extreme Management Center 7.0 release

ExtremeAnalytics Licensing

Since its inception, ExtremeAnalytics has tracked the flow rates across all configured ExtremeAnalytics engines and enforced restrictions on the data if the aggregate of these rates exceeded the licensed limit too many times.

Beginning in Extreme Management Center version 8.5, ExtremeAnalytics licensing now allows deployment flexibility by also granting end-system flow capacity to an entire deployment, regardless of the number of ExtremeAnalytics engines.

NOTE: For information on base licensing and all Extreme Management Center features, see <u>Requirements</u>.

If a flow rate license is applied and no end-system specific license is applied (for example, 88201), flow rate licensing is used. End-system licensing applies in all other instances.

If you have both an end-system license and a flow rate license, Extreme Management Center converts the total FPM allowed to end-systems at a 25:1 ratio. For example, a 1,000 FPM license would be enforced as a 40 end-system license.

The NMS Standard license (NMS) grants a NAC End-System license limit of 250 end-systems, but does not include ExtremeAnalytics end-systems. Additional ExtremeAnalytics licenses can then be added to extend that capacity, if needed.

The NMS Advanced license (NMS-ADV) grants a NAC End-System license limit of 500 end-systems and 100 ExtremeAnalytics end-systems. Additional ExtremeAnalytics licenses can then be added to extend that capacity, if needed.

The following is a list of ExtremeAnalytics flow rate licenses and their limits:

ExtremeAnalytics License Name	Maximum Flow Rate
NMS-ADV-XXX	3,000
PV-FPM-50K	50,000
PV-FPM-100K	100,000
PV-FPM-500K	500,000
PV-FPM-1M	1,000,000
PV-FPM-3M	3,000,000
NETSIGHTEVAL	3,000,000

The following list shows the NAC end-system licenses and their limits:

ExtremeAnalytics End-System License Name	Maximum End-System Limit
NMS-ADV-XXX	100
NETSIGHTEVAL / Subscription	2,000
	Note: If no flow rate license is applied, nms/netsighteval/subscription uses this end-system limit rather than the flow rate limits for licensing purposes.
88201	1,000
88202	3,000
88203	12,000

ExtremeAnalytics licensing is enforced through Extreme Management Center. The capacity allowed by each license is applied to the entire deployment and is calculated by adding together the usage numbers for each unique end-system or by adding together the flow rates for each ExtremeAnalytics engine in the deployment, depending on which license is being used. The capacity is checked on a continual basis.

NOTE: The **Max End-Systems** field in the ExtremeAnalytics Engine Advanced Configuration panel allows you to limit the end-systems from a single ExtremeAnalytics engine persisted in the Extreme Management Center database. This limits the amount of data stored on one engine.

Using Licenses to Establish Flow Rate Capacity

The following example shows how you can select ExtremeAnalytics licenses to achieve a desired flow rate capacity for your deployment.

The NMS-ADV license provides a basic flow rate capacity of 3,000 fpm (flows per minute). If you add additional ExtremeAnalytics licenses, the capacity is increased by the amount provided by the added licenses, up to the system-wide maximum of 3 million FPM.

For example, if you add the PV-FPM-100K license, then you would have a total flow capacity of 103,000.

3,000 + 100,000 = 103,000 FPM

If you then add the PV-FPM-50K license, you would have a total end-system capacity of 153,000.

3,000 + 100,000 + 50,000 = 153,000 FPM

Calculating Flow Rate Violations

Extreme Management Center uses the following steps to calculate ExtremeAnalytics licensing violations:

- Run a task every 15 minutes to calculate a flow rate. This rate is based on a system-wide average, not a maximum flow rate.
- 2. Compare the flow rate to the FPM limit for the licenses present. For example, the limit for an NMS Advanced license is 3,000 FPM.
- 3. Check once a day to see if the calculated FPM for any 15-minute period exceeds the licensed limit. If there was a violation, treat it as 1 violation day.

Calculating End-System Violations

Extreme Management Center checks the end-system count once a day and compares that count with the licensed limit. If this check determines that the count exceeds the licensed limit, the day is considered in violation.

Enforcing License Violation Limits

Extreme Management Center enforces flow rate and end-system license violations in the same way, and there are two violation levels.

The first enforcement level causes the following actions when the user first moves into the violation state:

- Displays a license violation message in the status bar.
- Generates an ExtremeAnalytics License Violation event.
- Generates an ExtremeAnalytics License Violation alarm.

The second enforcement level causes the following actions when the system is in violation for 7 days or more:

 ExtremeAnalytics replace the application name and group of a random 50% of identified flows with License Violation in the 'live' flows in the Flow Grid.



 Appliances also replace future flow report data including records with License Violation.



Clearing License Violations

Extreme Management Center resets the license violation count only after 7 consecutive days without a violation. When this happens, it triggers an **ExtremeAnalytics License Violation Clear** event and alarm.

Getting Started with ExtremeAnalytics

This topic provides information to help you get started using Extreme Management Center ExtremeAnalytics to view network application data in the **Analytics** tab. It includes information on ExtremeAnalytics access requirements, configuring the ExtremeAnalytics engine, enabling NetFlow flow collection, and configuring network locations.

ExtremeAnalytics Access Requirements

Both the ExtremeAnalytics feature and the **Analytics** tab require the Extreme Management Center Advanced (NMS-ADV) license. Contact your sales representative for information on obtaining an Extreme Management Center Advanced license.

In order to view the **Analytics** tab, you must be a member of an authorization group assigned the Extreme Management Center ExtremeAnalytics Read Access or Read/Write Access capability. The Read Access capability allows the ability to access the **Analytics** tab and view the ExtremeAnalytics reports. The Read/Write capability adds the ability to configure ExtremeAnalytics engines and NetFlow Collecting devices. It also adds the ability to create and modify fingerprints.

ExtremeAnalytics Engine Configuration

The ExtremeAnalytics engine provides the engine to monitor and classify layer 7 application information based on data from CoreFlow switches and reports that information to Extreme Management Center, where it is managed and displayed in the **Analytics** tab.

The ExtremeAnalytics engine must be installed and running on your network. For instructions, see the ExtremeAnalytics Engine Installation Guide.

Following installation, the ExtremeAnalytics engine must be added to Extreme Management Center and enforced via the **Configuration** tab in the **Analytics** tab.

Enable Flow Collection

Because the **Analytics** tab displays reports based on NetFlow or Application Telemetry (sflow) flow data, you must enable your network devices that act as the flow sensors, and

enable flow collection for their device interfaces. You must also configure your flow sensor devices to send their flow information to the ExtremeAnalytics engine. In addition, the device interfaces you enable for flow collection must match the interfaces configured for analysis by the engine.

Enable Jumbo Frames

When configuring a device as an Application Telemetry source for ExtremeAnalytics, jumbo frames must be enabled on the device and any device or virtual machine between the device and the ExtremeAnalytics engine.

For example, to enable jumbo frames on an ExtremeXOS device, enter the following in the device CLI:

enable jumbo-frame ports all

Related Information

Configuration - Analytics

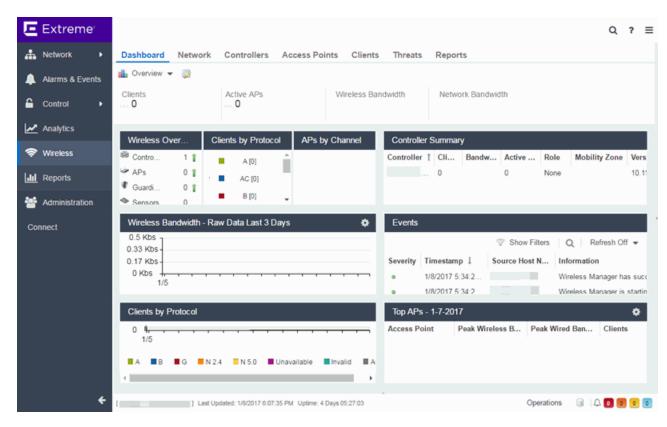
Configuring Enhanced Netflow for Extreme Analytics and Extreme Wireless Controller Version 10.21

When adding a Wireless Controller as a flow source in Extreme Management Center, a mirror port is automatically created. Wireless Controllers on which a firmware version of 10.21 or higher is installed use IPFIX, so the mirror port is unnecessary.

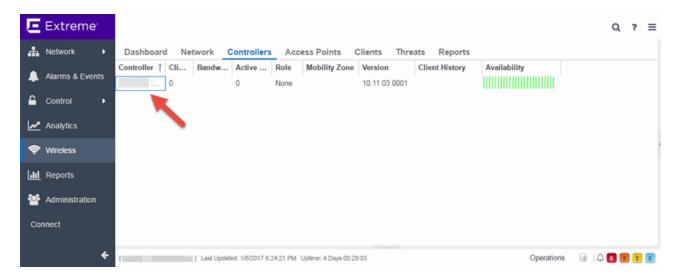
NOTE: Wireless Controllers on which a firmware version lower than 10.21 is installed still require the mirror port be configured.

To remove a mirror port on a Wireless Controller running version 10.21:

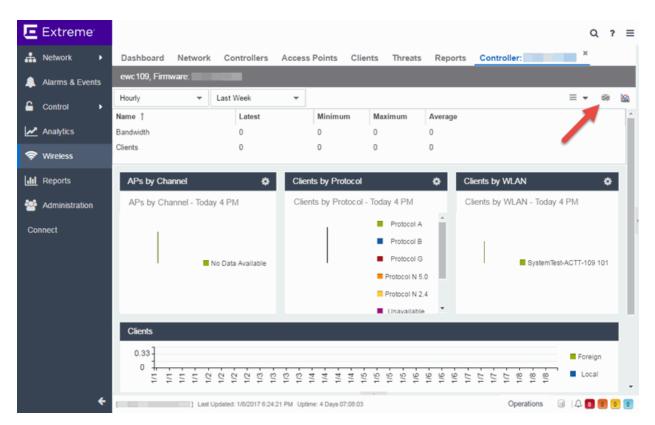
Access the Wireless tab in Extreme Management Center.
 The Wireless tab opens.



Select the Controllers tab.The Controllers tab opens.



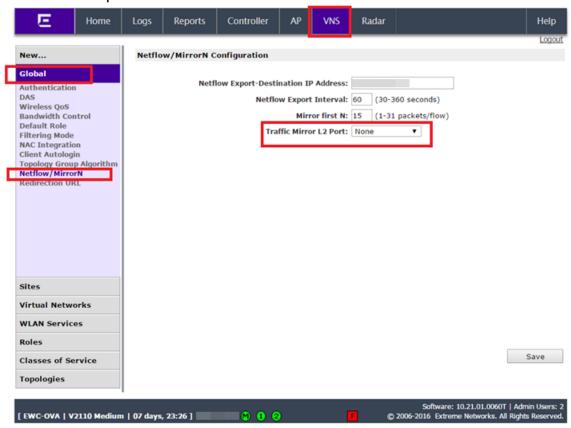
3. Click the **IP address** for the controller, located in the **Controller** column. The Wireless Controller Summary page opens.



4. Click the **WebView** icon (*) at the top right of the Wireless Controller Summary page.

The WebView opens for the controller.

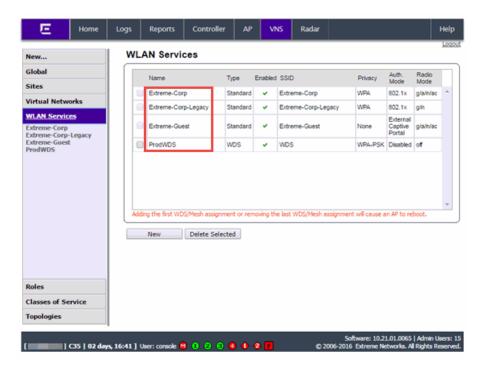
Click the VNS tab.The VNS tab opens.



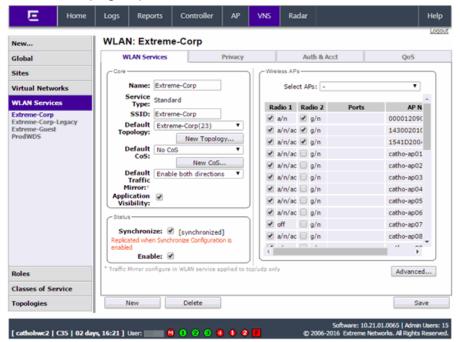
- Select Netflow/MirrorN from the left-panel.
 The Netflow/MirrorN Configuration page opens.
- 7. Select **None** from the **Traffic Mirror L2 Port** drop-down list.
- 8. Click the **Save** button.

NOTE: The Mirror Port in the Wireless Control Flow Sources section of the **Analytics** > **Configuration** > **Configuration** tab is not available once the **Traffic Mirror L2 Port** is disabled.

Select WLAN Services from the left-panel.
 The WLAN Services page opens.

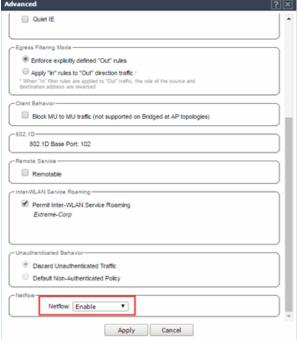


Click a wireless LAN in the table.
 The WLAN page opens for the selected wireless LAN.



Click the Advanced button.
 The Advanced window opens.

12. Scroll to the bottom of the window and ensure the **Netflow** drop-down list is set to **Enable**.



13. Click the Apply button.

The wireless controller is now configured.

NOTE: Rx Packets and Rx Bytes may incorrectly be 0 when flow data is gathered via a wireless controller running version 10.21 or higher. Additionally, application response times and some meta data may be blank. This is a known issue and will be addressed in a future release.

Related Information

For information on related topics:

Wireless

How to Deploy ExtremeAnalytics in an MSP or MSSP Environment

This Help topic presents instructions for deploying ExtremeAnalytics within an MSP (Managed Service Provider) or MSSP (Managed Security Service Provider) environment. It includes the following information:

- Configuring Extreme Management Center Behind a NAT Router
- Defining Interface Services

Configuring Extreme Management Center Behind a NAT Router

If the Extreme Management Center server is located behind a NAT (Network Address Translation) router, use the following steps to add an entry to the nat_config.txt file that defines the real IP address for the Extreme Management Center server. This allows the Extreme Management Center server to convert the NAT IP address received in the ExtremeAnalytics engine response to the real IP address used by the Extreme Management Center server. Not adding the real IP address for the Extreme Management Center server to the nat_config.txt file results in the ExtremeAnalytics engine incorrectly displaying a state of IMPARED (orange) rather than UP (green).

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.

- 1. On the Extreme Management Center server, add the following entry to the <install directory>/appdata/nat_config.txt file.
 <NAT IP address>=<real IP address>
- 2. Save the file.
- 3. If the Extreme Management Center Management server IP address is not configured to use the NAT IP address of the Extreme Management Center server, perform the following steps:
 - a. Enter the following command at the engine CLI:
 /opt/appid/configMgmtIP < IP address>
 Where <IP address> is the NAT IP address of the Extreme Management

Center server.

Press Enter.

b. Restart the appidserver once the new IP address is configured by typing: appidctl restart

Press Enter.

4. On the Extreme Management Center server, add the following text to the <install directory>/appdata/NSJBoss.properties file. In the second to last line, specify the hostname of the Extreme Management Center server.

NOTE: The ExtremeAnalytics engine functions as a client computer independent of the server. Both engines and clients must be able to resolve the hostname you specify.

```
# In order to connect to a NetSight server behind a NAT fi
rewall or a
# NetSight server with multiple interfaces you must define
  these two
# variables on the Extreme Management Center
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 corr
ect IP that
# they want to use as the IP to connect to.
java.rmi.server.hostname=<hostname of NetSight server>
java.rmi.server.useLocalHostname=true
```

- 5. Save the file.
- Add the Extreme Management Center server hostname to your DNS server, if necessary.

NOTE: ExtremeAnalytics engines, remote Extreme Management Center clients, and any ExtremeControl engines must be able to connect to Extreme Management Center using this hostname.

Related Information

For information on related windows:

ExtremeAnalytics Engine Advanced Configuration Panel

ExtremeAnalytics Application Data Collection

The ExtremeAnalytics engine provides an application data collection function that collects and records information about network utilization. It includes:

- General Usage Collection High-level application-centric data, collected hourly and in five-minute intervals.
- Extended Application Collection Detailed data about all end-systems in the network, collected hourly.

Application data collection is based on network flow information. Network utilization for various objects in the network (called targets) is measured, collected, and used to create application data reports in Extreme Management Center.

NOTE: Ensure at least 4GB of swap space is available for flow storage or impaired functionality may occur. Use the free command to verify the amount of available RAM on your Linux system.

This Help topic describes application data collection, including collection targets, statistics, and intervals. It also describes the different collectors used to perform the collection, as well as the sources for flow information.

Data Collection Overview

Application data collection is performed by the ExtremeAnalytics engine. The engine collects flow records from switches in your network. It then augments the collected flow data with detailed application information derived by network packet inspection, resulting in rich analytical data.

For example, if a NetFlow record reports 100 bytes transferred from client Workstation 1 to server Host A, then the collection process would add 100 bytes to the tally for Workstation 1, and 100 bytes to the separate tally for Host A. If the flow is identified as traffic for the Payroll application, then 100 bytes would be added to another tally for Payroll as well. And finally, 100 bytes is added to another tally for the entire network. At the end of a collection interval, the totals for client Workstation 1, server Host A, the Payroll application, and the entire network are written to the database.

Data from network flows is collected in an aggregated form for a period of time (called a collection interval), and then stored in the Extreme Management Center database.

Extreme Management Center uses this data to provide reports that show how your network is being utilized.

To conserve space on your Extreme Management Center server hard drive, your ExtremeAnalytics engines only collect total flow records when the server hard drive drops below 10 GB of free space. If the Extreme Management Center server hard drive drops an additional 1 GB (under 9 GB of free space), your ExtremeAnalytics engines stop collecting all flow data.

NOTE: To change the differential threshold (the additional amount of free space reduction after which all records stop being collected), edit the RM_FREE_SPACE_MINIMUM_ALLOW_SUMMARY_KB value in the NSJBOSS.properties file. The value is set to 1,000,000 KB by default, so ExtremeAnalytics stops collecting all records when free space reaches 10GB - 1,000,000 KB = 9 GB.

Collection Targets

Flow data is collected on objects in your network called targets. Some targets are physical, such as clients and servers, and some are logical, such as applications.

An ExtremeAnalytics engine can track the following target types:

- Client The end-point of a flow that has the client role for that connection.
- Server The end-point of a flow that has the server role for that connection.
- Application An application in ExtremeAnalytics, identified through layer 7 analysis (for example, Facebook).
- Application Group Application categories, such as Cloud Computing or Social Networking.
- Site The client's physical location on the network, based on its IP address. Sites
 are used by ExtremeAnalytics to identify the physical location for the client of an
 application flow.
- Device Family The kind of device determined for a client, such as Windows or iOS.
- Profile An ExtremeControl profile assigned to a client.

In some cases, the engine can also track combinations of targets. For example, it can track the total number of bytes transferred from Workstation 1 for the Payroll application separately from Workstation 2 for Payroll, and from Workstation 1 for Facebook. These target and sub-target pairs provide for Extreme Management Center drill-down reports,

for example, reports to show the top Payroll clients or the top applications for Workstation 1.

This report shows the top 10 applications seen on the network (based on bandwidth) during the last hour.

Applic ations	Application Group	Bytes	Sent Bytes	Received Bytes
nsbuild-linux3	Internal File Downloads	12.65 GB	202.69 MB	12.44 GB
Microsoft SQL Server	Databases	2.91 GB	479.26 MB	2.43 GB
CIFS	Storage	2.17 GB	386.53 MB	1.78 GB
■ WASSP	Protocols	1.99 GB	248.98 MB	1.74 GB
¥ Web	Web Applications	1.75 GB	73.74 MB	1.67 GB
Extreme Networks	Corporate Website	1.54 GB	131.88 MB	1.41 GB
№ SSH	∨PN and Security	1.34 GB	217.46 MB	1.12 GB
Outlook Office365	Mail	1.29 GB	466.91 MB	826.56 MB

Collection Statistics

Collection statistics are quantitative data that can be collected for a target. This includes statistics directly reported in NetFlow records, such as bytes transferred, as well as information that can be derived indirectly, such as the number of unique clients seen using an application.

An ExtremeAnalytics engine can track the following statistics:

- Bytes The number of bytes transferred in both directions, between the client and the server. Also known as bandwidth. You can track sent and received bytes as well as total bytes.
- Flows The number of NetFlow records sent by the switch to report the traffic between the client and the server. You can track inbound and outbound flows as well as total flows.
- Clients The number of unique clients associated with the target.
- Applications The number of unique applications associated with the target.
- Network Response Time The average amount of time to create a connection.
- Application Response Time The average amount of time for a server to respond to a request.

Applications (Application Response Time) - Average: 9.61s - 7/17 11 AM - hourly Applic ations Application Group Application Response Time Network Response Time lighthousesites Web Applications Web Applications Acxiom Advertising Google App Engine Web Content Services Parse Push Real Time and Cloud... Web Applications apigee c analc lima Web Applications Web Applications alternativeto Web Applications healthline Web Applications

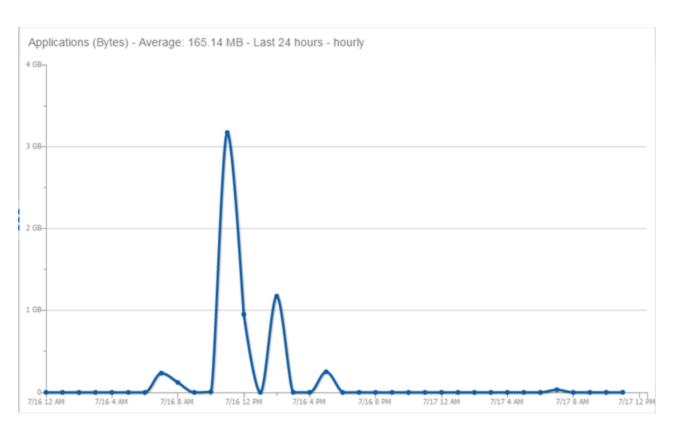
This report shows the average application response times for the top 10 applications during the last hour.

Collection Intervals

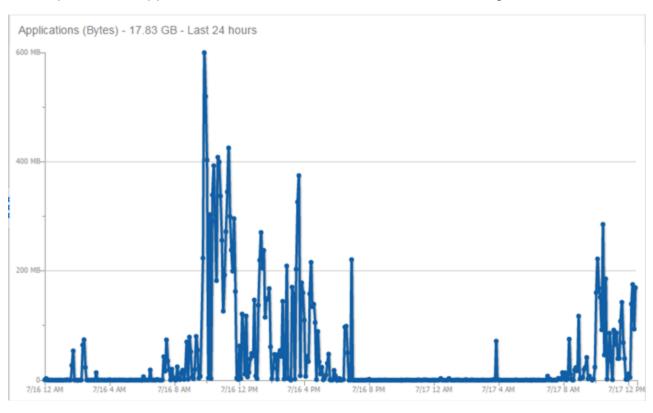
The ExtremeAnalytics engine collects and aggregates flow data for a period of time called an interval. At the end of the interval, the engine writes the totals to the Extreme Management Center database and a new interval begins, with new totals collected starting at zero.

Some statistics are collected and written to the database on an hourly interval. Other statistics are collected at a high-rate interval of every five minutes, providing for a more detailed picture of how traffic changes over time.

This report shows application bandwidth over 24 hours based on an hourly interval.



This report shows application bandwidth over 24 hours based on a high-rate interval.



All statistics can be collected over multiple intervals and averaged. When viewing report data, it is important to know the interval used for any average that is displayed.

Certain statistics, such as bytes and flows, can be collected over multiple intervals to provide a total over time, while other statistics, such as client count, cannot. To illustrate, the number of bytes seen in two hours would be the total of the number of bytes seen in each hour. However, the number of unique clients seen in two hours would not be the total of the number of unique clients seen in each hour, as some clients were probably seen in both hours.

Using Sites to Collect In-Network Traffic

While flow data collection can aggregate data for all flow traffic that is visible, it may be more useful to aggregate data for *in-network* flows only. These are flows used by clients that are located in your internal network. By collecting data for only in-network flows, the overhead of aggregating data over an interval can be reduced.

You can define your internal network by configuring sites. A site is a set of IP masks that defines a well-known portion of your internal network. You can use the World site to identify your entire internal network. If you have already reserved certain IP address ranges for certain physical sites on your network, you can create multiple sites that correspond to these reserved IP ranges. Multiple sites can be created to identify different buildings, sites, or geographical areas of your network. Any IP that matches any site is considered to be in-network. If you define multiple sites, you will be able to analyze data broken down by site.

Data Collector Types

There are two kinds of data collectors used in ExtremeAnalytics.

- General Usage Collectors These are hourly and high-rate collectors that record the top targets during an interval. Many types of targets and target-pairs are supported.
- End-System Details Collector This is an hourly collector that attempts to capture
 and record data for all in-network clients and servers that it detects. All traffic
 collected is tagged with site, profile, device family, and other attributes.

Data from these collectors is stored separately in the database. The collector data used in a report depends on the nature of the report. Higher-level information, such as top applications during an hour, will be based on general usage collector data, since it is relatively inexpensive to access. End-system details data might be used when data for a

specific client or server is needed, or when the information requested is highly specific, for example, top applications used by Android devices in the London site.

General Usage Collectors

General usage collectors collect data about all instances of a target for the interval, and then record only the most significant targets (typically, the 100 most significant targets).

When the top targets are calculated for a collection interval, several different statistics can be used as a basis for choosing the most significant entries. For example, collectors can record the top applications based on bytes, and also record the top applications based on number of clients. For each type of target collected, there are different sets of bases used.

General usage collectors operate at both hourly and high-rate intervals. They can collect data from all flows or from in-network flows only.

Hourly General Usage Collectors

The following table describes the hourly data collected by the general usage collectors.

	Sub-		
Target	Target	Bases	Traffic Used
Total			In-Network
			Flows/
			All Flows
Application		Bytes	In-Network
		Received Bytes	Flows
		Transmitted Bytes	
		Flows	
		Receive Flows	
		Transmit Flows	
		Clients	
		Network Response	
		Time	
		Application Response	
		Time	
Application	Client	Bytes	In-Network
			Flows

Target	Sub- Target	Bases	Traffic Used
Application Group		Bytes Flows Clients	In-Network Flows
Client		Bytes Received Bytes Transmitted Bytes Flows Receive Flows Transmit Flows Applications Network Response Time Application Response Time	All Flows
Device Family		Bytes Flows Clients	In-Network Flows
Site		Bytes Flows Clients Network Response Time Application Response Time	In-Network Flows
Profile		Bytes Received Bytes Transmitted Bytes Flows Receive Flows Transmit Flows Network Response Time Application Response Time	In-Network Flows

	Sub-		
Target	Target	Bases	Traffic Used
Threat		Bytes Flows Application Response Time Network Response Time Received Bytes Sent Bytes Inbound Flows Outbound Flows	In-Network Flows
Threat	Threat End- System Pair	Bytes Flows Application Response Time Network Response Time Received Bytes Sent Bytes Inbound Flows Outbound Flows	In-Network Flows
Server		Bytes Received Bytes Transmitted Bytes Flows Receive Flows Transmit Flows Network Response Time Application Response Time	All Flows
Application	Device Family	Bytes Flows Clients	In-Network Flows
Application	Profile	Bytes Flows Clients	In-Network Flows

High-Rate General Usage Collectors

The following table describes the high-rate data collected by the general usage collectors.

Target	Sub-Target	Bases	Traffic Used
Total			In-Network Flows/ All Flows
Application		Bytes Flows Clients	In-Network Flows
Application Group		Bytes Flows Clients	In-Network Flows
Device Family		Bytes Flows Clients	In-Network Flows
Site		Bytes Flows Clients	In-Network Flows
Profile		Bytes Flows Clients	In-Network Flows

End-System Details Collector

The end-system details collector tracks client/application target pairs.

Unlike general usage collectors, this collector attempts to record data for all in-network clients and servers it sees during the hour. For each client or server, it records data for up to 10 applications, plus an "other" category to capture the remaining traffic. Information such as location, device family, and profile are also recorded for each end-system.

The large number of targets recorded each hour and the amount of detail recorded for each one, can result in a large volume of data being stored in the database. In order to prevent disk space from being over-utilized, there is a total limit of 50,000 clients which can be recorded each hour across all ExtremeAnalytics engines. There is also a 25,000 client limit per engine for most license types. However, if you have an NMS-ADV license

without any ExtremeAnalytics license, the per-hour total limit is 100 clients across all ExtremeAnalytics engines.

Flow Information Sources

The ExtremeAnalytics engine uses NetFlow or SFlow records from the switches and wireless controllers in your network as a source for flow data. Information such as IP addresses, ports, and bytes transferred comes from this flow data source.

This data is augmented with additional layer 7 application information produced by the ExtremeAnalytics engine through deep packet inspection. Information such as application name and network response time comes from this source.

There is additional information that can be obtained from sources other than NetFlow/SFlow records and deep packet inspection.

NOTE: Most of these sources rely on ExtremeControl data. If ExtremeControl is part of your network configuration, then ExtremeControl integration can be enabled (see <u>instructions</u> below) to provide access to these sources. Site data is obtained from sites configured in Extreme Management Center.

The following is a list of information that can obtained from different sources:

- Hostname The client or server's hostname can be derived using ExtremeControl.
 ExtremeControl integration must be enabled.
- Site The site for a flow is the site of the client in the flow. Client and server sites
 are derived from the sites configured on the **Network** tab. If a client does not match
 a site, then the site is empty. If a flow has a site, the flow is considered to be innetwork.
- Detailed Site Detailed site information is derived from the switch and port information resolved for the client end-system. ExtremeControl Integration must be enabled.
- Device Family The device family is a general description of the operating system detected in the client, for example, Windows, Linux, or Android. The device family is derived from network packet inspection. The device family can also be provided by ExtremeControl, if ExtremeControl integration is enabled.
- Profile The client's profile is derived from the ExtremeControl profile assigned to the client end-system. ExtremeControl integration must be enabled.

Username – The client's username is derived from network packet inspection. The
username can also be provided by ExtremeControl, if ExtremeControl integration
is enabled.

It is possible that different sources may provide different values for the same information. For example, network packet inspection may provide the device family name of Window 7, whereas ExtremeControl may provide the device family name of Windows.

Enabling ExtremeControl Integration

If your network configuration includes ExtremeControl, ExtremeControl data can be integrated with flow data to provide additional information. ExtremeControl integration is only useful if you are collecting flows for end-systems managed by ExtremeControl.

When ExtremeControl integration is enabled, if a client in a flow matches an end-system in ExtremeControl, then:

- The client hostname in the flow is derived from the end-system.
- The device family in the flow is derived from the end-system.
- The username in the flow is derived from the end-system.
- The profile in the flow is derived from the end-system's ExtremeControl profile.
- The detailed site in the flow is derived from end-system data.

If a server in a flow matches an end-system in ExtremeControl, then:

The server hostname in the flow is derived from the end-system.

To enable ExtremeControl integration on the ExtremeAnalytics engine:

- 1. If the ExtremeControl distributed end-system cache is not enabled on the Extreme Management Center server, you must enable it using the following steps.
 - a. Select **Administration > Options** from the menu bar to open the **Access Control Options** window.
 - b. Click on Advanced Settings.
 - c. In the End-System Mobility section, select the **Enable distributed end- system cache** option.
 - d. Click the **Reload** button to reload the cache configuration on the Extreme Management Center server. Click **OK**.
- 2. Enable ExtremeControl Integration on each ExtremeAnalytics engine where you want to use ExtremeControl data.

- a. Access the **Analytics** tab.
- Expand each ExtremeAnalytics engine and select Advanced Configuration.
 In the right panel under Configuration Options, select the Enable
 ExtremeControl Integration option.
- c. If your ExtremeControl engines are using Communication Channels, you must select the **ExtremeControl Communication Channel** option and enter the channel name. The ExtremeAnalytics engine is only able to access end-systems in its channel.
- d. Click Save.
- e. Enforce your ExtremeAnalytics engines.

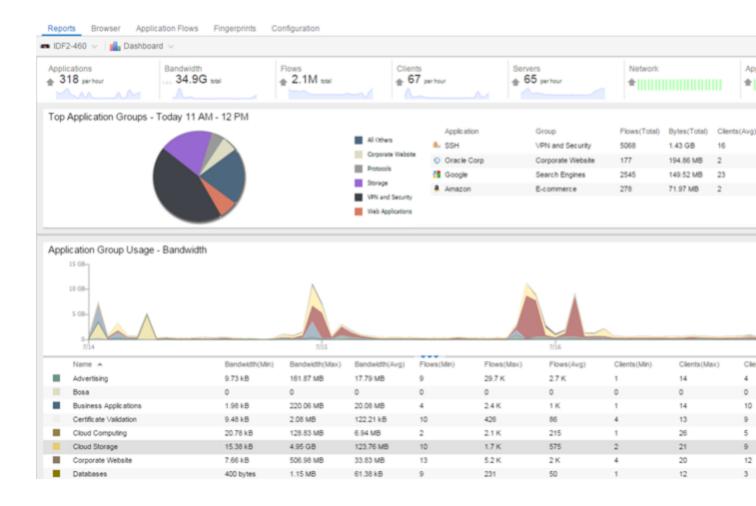
Reports

Data gathered from flow usage collection is the basis of many reports in the Extreme Management Center's **Analytics** tab. Once collection is enabled, these reports begin to exhibit data.

Dashboard Report

The following screen-shot shows the main Dashboard report. It contains data produced by the hourly General Usage collectors, and displays data for a specific hour. Across the top are the hour's totals. Below them are Top Application Groups, as a chart, and Top Applications, as a table, for the same hour. There is also Application Group Usage over the last 3 days, as a chart and as a table.

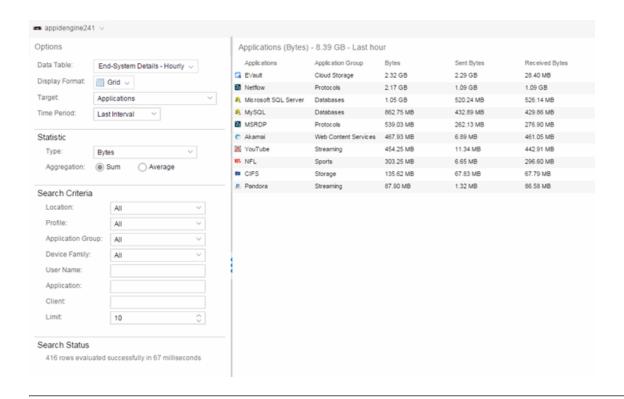
Note that data from different ExtremeAnalytics engines is maintained separately. If you have more than one ExtremeAnalytics engine, you need to select which engine to view, using the engine menu in the top-left corner.



Browser Reports

The Browser provides special reports that lets you select the targets, statistics, and collection interval for your report, as well as define search criteria to further filter report data. Using the Browser, you can create custom queries that provide greater flexibility in defining what data to display and how to display it. When you create a Browser report, you select which type of network activity data to use: end-system details (always hourly), application data hourly, or application data high-rate. For additional information, see Applications Browser.

The following screen-shot shows an example of a Browser report showing application/device family bandwidth usage for the last hour.



Related Information

For information on related ExtremeAnalytics topics:

- Getting Started with ExtremeAnalytics
- Analytics

ExtremeAnalytics Tab Overview

The **ExtremeAnalytics** tab allows you to view and customize its <u>dashboard</u> and <u>browser</u>, as well as ExtremeAnalytics <u>reports</u>, <u>fingerprints</u>, and <u>application flow</u> data. You can also manage and configure your ExtremeAnalytics engines.

Additionally, the **Menu** icon (≡) at the top right of the screen provides links to additional information about your version of Extreme Management Center.

NOTE: ExtremeAnalytics reports and application flow data is not available unless an ExtremeAnalytics engine is configured and you are a member of an authorization group assigned the Extreme Management Center ExtremeAnalytics Read Access or Read/Write Access capability. The Read Access capability allows the ability to access the **Analytics** tab and view the ExtremeAnalytics reports. The Read/Write capability adds the ability to configure ExtremeAnalytics engines and NetFlow Collecting devices. It also adds the ability to create and modify fingerprints.

Viewing ExtremeAnalytics application data requires certain access requirements and prerequisites. Both the ExtremeAnalytics feature and the **Analytics** tab require the Extreme Management Center Advanced (NMS-ADV) license. Contact your sales representative for information on obtaining an Extreme Management Center Advanced license.

Dashboard

The Dashboard tab displays an overview of application usage on your network through a series of graphs. It allows you to view network activity statistics based on client/server, application, industry, IP reputation, and response time for the specified ExtremeAnalytics engine. Many of the reports are links to more detailed pages.

Browser

The Browser tab lets you query information about recent network activity stored in the Extreme Management Center database and display results in various grid and chart report formats. Using the Browser, you can create custom queries based on selected options including a data target, statistic type, and other search criteria.

Application Flows

You can choose from the **View** drop-down list to show you several options in the table on the Application Flows tab, including the latest flows from the specified ExtremeAnalytics engine, the worst network and application response times, classified and unclassified flows, and flows during a specified time frame. The table presents bidirectional flow data (aggregate flows) or unidirectional flow data (base flows).

Fingerprints

A fingerprint is a description of a pattern of network traffic which can be used to identify an application. The **Fingerprints** tab provides detailed information about fingerprints used by ExtremeAnalytics to identify application flows. You can choose to view in-use and customized fingerprint data.

Packet Captures

Use the **Packet Captures** tab to analyze the packets from the flows displayed on the **Application Flows** tab. The packet captures you create are presented in a table, which allows you to view details about the packet capture. Additionally, using this tab you can select a packet capture and view it in a packet analyzer.

Configuration

The **Configuration** tab provides detailed information on the ExtremeAnalytics engines you configure. It also lets you add and enforce your engines, and access engine reports and diagnostics. You must be a member of an authorization group assigned the Extreme Management Center ExtremeAnalytics Read/Write Access capability to view the **Configuration** tab.

Reports

On the **Reports** tab, you can access a selection of reports that provide detailed information on application usage on your network, as well as network activity statistics based on application, user name, client, and site. For many of the reports, you can click on an item in the report to view details or right-click an item to select from other focused reports.

Related Information

- Dashboard Overview
- Browser
- Application Flows
- Fingerprints
- Configuration
- Reports

ExtremeAnalytics Dashboard Overview

Accessible from the **Analytics** tab in Extreme Management Center, the **Dashboard** tab displays an overview of application usage on your network, as well as network activity statistics through a series of real-time reports. The Dashboard is flexible and customizable - you can choose the reports and the design of the page to meet your specific needs. Many of the reports are links to more detailed pages.

The Dashboard includes a drop-down list with links to additional report dashboards:

- Insights
- Client/Server
- Applications Browser
- Industry
- IP Reputation
- Response Time
- Network Service
- Tracked Applications

Several report pages can be launched in the **Reports > Reports Designer** view in Extreme Management Center by clicking the **Launch in Report Designer** icon (**)**.

Insights Dashboard Reports

The Insights dashboard displays graphs with real-time network and application usage and service data, and tools that you can use to customize the dashboard using dragand-drop capabilities.

Five ring charts display real-time Engines, Disk Usage, Flow Rate, Network, and Application usage and service data. The ring charts are links to additional data. The Network and Application charts link to the Network Service and Response Time report dashboards, respectively, which are also accessible from the **Dashboard** drop-down list.

Use the Custom Dashboard to drag and drop only the graphs you want on your dashboard. Each graph is a real-time preview and many are linked to additional detail reports. You can also choose whether the graphs in the Application Group area are organized in columns or rows in the Custom Dashboard area.

Client/Server Dashboard Reports

This dashboard displays reports on clients and servers seen on the network over the last 24 hours. It also displays reports on top clients by bandwidth, flow, or number of applications, and top servers by bandwidth or flow.

Click on the **Info** icon (i) at the top right of the dashboard page to read a description of each report.

Applications Browser Dashboard Report

The Application Browser Dashboard displays bubble maps for top applications by bytes and flows, top profiles by bytes, and top sites by bytes. Hovering over a bubble displays bandwidth use or the number of flows. Use the drop-down menus to change the start date and time for the reports.

Drill-down for more information by clicking on an application bubble to open a new graph of clients, flows, and usage data for that application. In that graph, click on a client link to view application data for that client.

Industry Dashboards

Enterprise Dashboard

The Enterprise Dashboard displays application information specific to the Enterprise network including social applications, storage applications and cloud, business applications and email, and network applications and protocols.

Education Dashboard

The Education Dashboard displays application information specific to the campus network including learning management systems, P2P, streaming, and social applications.

Healthcare Dashboard

The Healthcare Dashboard displays applications used in the healthcare environment including patient care, medical applications, and HIPAA.

Venue Dashboard

The Venue Dashboard displays data grouped according to sports, social media, news and weather applications, as well as software update applications.

IP Reputation Dashboard

This report displays potential threat activity on your network from IP addresses known to be suspicious. IP addresses can be flagged as suspicious for a variety of reasons, including forced IP anonymity through the use of a Tor exit node, being listed as a threat by the Emerging Threats project, or classified as suspicious by internet users. Additionally, each IP address classification has its own recommended course of action, listed below.

- CiArmy Top Attackers The CiArmy reputation feed is a set of IP addresses tied
 to malicious activity defined by a collaborative network security effort backed by
 the Emerging Threats project. Any IP communications to addresses in this list from
 the local network are suspicious and may indicate that the local IP is involved in
 various activities such as command and control communications with the remote
 host. IP addresses classified as CiArmy Top Attackers require further investigation.
- Compromised Hosts Connecting Into the Network IP addresses that match this classification are on a list of IP addresses maintained by the Emerging Threats project. This list consists of a set of IP addresses that appear to have been compromised by malware, individual actors, worms, botnets, or other means. When ExtremeAnalytics detects application flows that match an IP from the Compromised list, this is a likely indicator that systems in the local network are either under attack or have already been compromised (since the communications may be command and control directives emanating from the compromised host).
- Connections to Bad Hosts IP addresses classified as Connections to Bad Hosts are known to function as command and control nodes for various botnets around the Internet. Any flows to or from such IP addresses have a high probability of being associated with botnet command and control traffic.
- Connections to Bad Hosts Based on Port IP addresses flagged in this
 classification are known to function as command and control nodes for botnets
 based on the port number. For example, a botnet command and control node may
 be a legitimate webserver, which is not suspicious. However, if there are flows
 certain botnets are known to use specific ports on a node, these communications
 cause the IP address to be flagged in this classification.
- **DShield Top Attackers** The DShield project is a distributed security analysis effort that collects logs, IDS/IPS events, and other data from volunteers around the

Internet. This data is analyzed by DShield and a list of the top set of IP addresses that appear to be attacking other systems worldwide is provided by DShield. When application flows appear within ExtremeAnalytics that match any of the IP addresses from the DShield top attackers list, it is likely systems in the local network are being actively attacked.

Tor Exit Node, Relay or Router – This reputation feed provides a listing of known
Tor exit nodes, relays, and routers. Tor is a service that provides IP anonymity. It
functions as a distributed set of systems on the Internet and builds sets of "virtual
circuits" through this set of systems on behalf of users that do not want to reveal
their local IP address to destination servers. Typically, Tor is used to mask web
browsing communications, but other services can run over the Tor network.
Matches against this reputation feed indicate Tor usage on the local network.

NOTE: IP addresses that match multiple classifications (e.g. an IP address is listed as both a CiArmy Top Attacker and a DShield Top Attacker) are only classified in the first category in which they match, not in additional categories.

Response Time Dashboard

The Response Time Dashboard displays the response time in milliseconds of application data grouped by different criteria, selected from the drop-down list. The data is displayed as a line graph, which is updated periodically.

Network Service Dashboard

The Network Service Dashboard displays the response time of network services for the top five worst-performing sites as well as the overall average of all sites. The data for each network service at a site is displayed as a bar and line graph, which is updated periodically.

Tracked Applications Dashboard

The Tracked Applications Dashboard displays the response time of the applications you configure in the **Tracked Applications** field on the **Analytics** > **Configuration** > **Configuration** tab. The data for each network service at a site is displayed as a bar and line graph, which is updated periodically. You can choose to organize the graphs in either columns (|||||) or rows (||||).

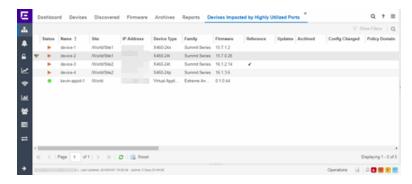
Related Information

• ExtremeAnalytics tab

ExtremeAnalytics Insights Dashboard

Accessible from the **Analytics** tab in Extreme Management Center, the Insights Dashboard displays an overview of application usage on your network, as well as network activity statistics based on client/server, application, industry, IP reputation, and response time.

Use the Insights Dashboard to view graphs that display real-time network and application usage and service data, and tools that you can use to customize your dashboard using drag-and-drop capabilities.



Insights

The Insights Dashboard displays ring charts and a customizable Application Group Dashboard. You can collapse and expand the ring charts and Application Group Dashboard for flexible display capabilities.

Ring Chart

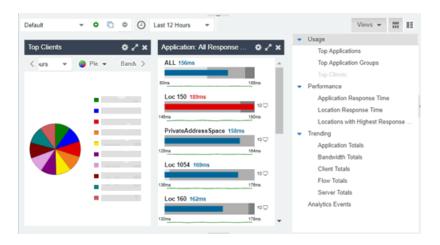
Six ring charts display real-time <u>Engines</u>, <u>Virtual Sensors</u>, <u>Disk Usage</u>, <u>Flow Rate</u>, <u>Network</u>, and <u>Application</u> usage and service data:



- Engines The number at the center of the ring chart indicates how many engines are represented by the chart. The colors in the graph indicate the states of the configured engines. Hover over a ring color to display a tooltip with the status of that engine. Click the graph to display overview and status details.
- Virtual Sensors The number at the center of the ring chart indicates how many
 virtual sensors are represented by the chart. The colors in the graph indicate the
 states of the configured virtual sensors. Hover over a ring color to display a tooltip
 with the status of that virtual sensor. Click the graph to display overview and status
 details. Click the graph to open the <u>Virtual Sensors tab</u>.
- Disk Usage The number at the center of the ring chart indicates the percentage
 of Disk Usage. The colors in the graph display the percentage of disk usage being
 used. Hover over the ring color to display a tooltip with usage percentage and units
 of space details.
 - Click the graph to open the **Configuration** tab, where you can configure the information displayed in the Insights Dashboard.
- Flow Rate The number at the center of the ring chart indicates the flow rate
 percentage. The colors in the graph indicate the flow rates for the different engines
 being used. Hover over a ring color to display a tooltip with status, percentage and
 rate details for each engine. Click the graph to open the Licenses tab.
- Network Response The colors in the graph indicate the network response time
 for the application/site. Hover over a ring color to display a tooltip with status
 details and the number of networks at that status. Click a color in the graph to open
 the Network Service dashboard, which displays network service details.
- Application Response The colors in the graph indicate the application response
 time for the application/site. Hover over a ring color to display a tooltip with
 response time details and the number of applications within the expected response
 time range. Click a color in the graph to open the Response Time dashboard,
 which displays network and application response time charts and details.

Custom Dashboard

The Custom Dashboard is a customizable space for viewing graphs that you select from the **Views** drop-down list. The buttons at the top right of the Applications Group dashboard (allow you to save and copy your dashboard.



Related Information

- Analytics Tab
- How to Use the Application Group Dashboard

How to Create an ExtremeAnalytics Insights Custom Dashboard

The **Dashboard** > **Insights** view in Extreme Management Center provides you with graphs that display real-time network and application usage and service data, and tools that you can use to customize the dashboard using drag-and-drop capabilities.

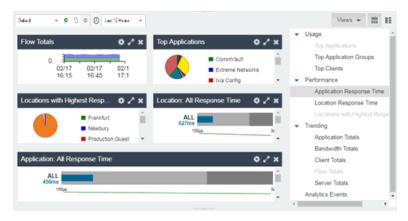
The Custom dashboard in the **Insights** view allows you the flexibility to create your graphs and reports preview area. Choose the graphs you want to view, the data that is displayed in the graphs, and how the graphs are displayed. You can collapse and expand the Custom dashboard for flexible display capabilities.

Custom Dashboard

The Custom dashboard is a flexible space for viewing graphs that you select from the <u>Views</u> drop-down list. The buttons at the top right of the Custom dashboard (allow you to save and copy your dashboard.

 Click the Create button (•) to create a new, empty dashboard where you can drag and drop the graphs you select from the Views drop-down list. The Create

- Dashboard window opens to allow you to name your new dashboard. Click the **Save** button to save the dashboard, which is available to all users in your network.
- Click the **Copy** button () to copy the current dashboard, which you can customize by adding new graphs from the Views drop-down list. Once you have edited the copied dashboard, click the **Save** button to save the new dashboard, which is available to all users in your network.



- Click the Views drop-down list at the far right and select from the graphs in the drop-down list.
- 2. Drag and drop the graph(s) to the open area to the left. Once in place, the link will display as a real-time preview of the graph.
- 3. Choose the orientation of your dashboard by clicking either the row (■) or column (■) button.
- 4. Hover over the data to display a tooltip with usage data.
- 5. Click the **Gear** button (☑) in each graph to further modify your Application Group graphs data:
 - Top Choose the number of top applications, application groups or clients (depending on the graph) to be displayed in the graph
 - Range Adjust the time frame of the data depicted in the graph by choosing from the drop-down list. The Custom Time option allows you to choose any start time, and the Custom Range option allows you to choose any start and end times.
 - Graph style Select from pie, word cloud, tree map or bubble map graph styles in the drop-down list.

 Data – Select from Bandwidth, Flows, Clients data types from the drop-down list.

Graphs

These graphs are available to be added as real-time previews to your Custom dashboard:

Usage

Top Applications – Displays usage data for the top applications. Click any color in the graph to display an encrypted web detail page for that application.

Top Application Groups – Displays usage data for the top application groups. Click any color in the graph to display an encrypted web detail page for that group.

Top Clients – Displays usage data for the top clients. Click any color in the graph to display application and application group detail page for that client.

Performance

Application Response Time – Displays response times for all applications. You can also create response time reports for individual applications and sites that you define.

Site Response Time – Displays response times for all sites. You can also create response time reports for individual applications and sites that you define.

Sites with Highest Response Time – Displays sites with the highest response times. Click any color in the graph to display network and application response time reports for that site.

Trending

Application Totals – Displays the total number of applications based on the date, time, and duration you choose. Click any color in the graph to display an application detail page.

Bandwidth Totals – Displays bandwidths for the application in a line graph based on the date, time and duration you choose. Click any data point in the graph to display a Top Applications by Bandwidth detail page.

Client Totals – Displays the total number of clients for the application based on the date, time, and duration you choose. Click any data point in the graph to display a Top Clients detail page.

Flow Totals – Displays the total number of outbound and inbound flows for the application based on the date, time, and duration you choose. Click any data point in the graph to display a Top Applications by Flows detail page.

Server Totals – Displays the total number of servers for the application based on the date, time, and duration you choose. Click any data point in the graph to display a Top Servers detail page.

Analytics Events

Analytics Events - This report displays the event log filtered to show only the events related to ExtremeAnalytics

Related Information

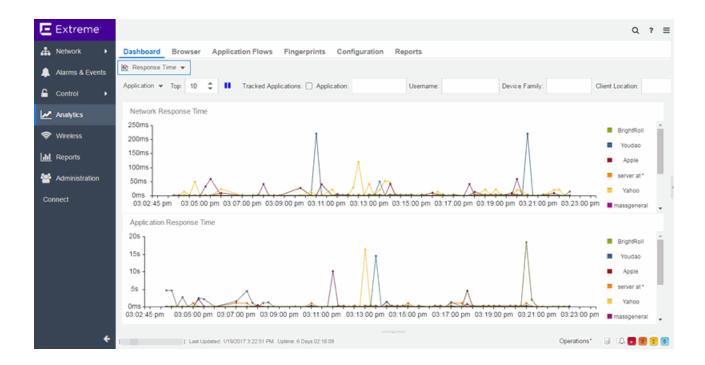
- ExtremeAnalytics tab
- ExtremeAnalytics Insights Dashboard

ExtremeAnalytics Response Time Dashboard

The Response Time Dashboard displays the network and application response time data for the slowest targets on your network based on response time for the last 20 minutes. Use the graph to view response time data for a variety of filters, including application, device family, and username.

Additionally, you can use the dashboard to select the number of targets for which the response time is displayed and you can filter the information based on certain criteria and view flow data specific to the data you select.

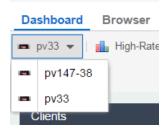
To access the Response Time Dashboard, open the **Analytics > Dashboard** tab and select **Response Time** in the dashboard drop-down list.



Overview

The Response Time Dashboard contains two graphs, one displays the <u>network</u> <u>response time</u> and the other displays the <u>application response time</u>. Data is updated every 15 seconds and displays data over the last 20 minutes.

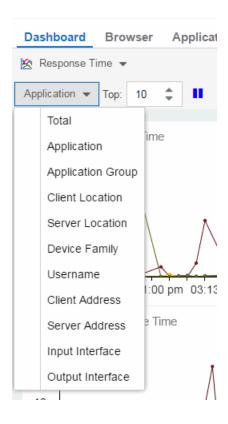
If you have multiple ExtremeAnalytics engines, use the **Engine** drop-down list to select an engine to use as the source for the report data.



Use the toolbar at the top of the window to display data based on criteria you select and updates the two graphs.

Application

Use the **Application** drop-down list to group the data in the Response Time Dashboard by the following criteria:



Top

Use the **Top** field to limit the results in the graphs to display only the top results based on the number you enter.

For example, you can configure the graphs to display the top 3 slowest applications by response time.

Tracked Applications

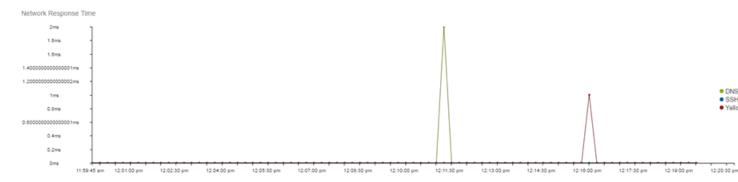
Click the **Tracked Applications** box to add response time results for tracked applications to the Network Response Time and Application Response Time graphs.

Filters

You can also use the filter options at the top of the window to search for specific criteria. Using these fields limits the data to Tracked Applications, Application, Username, Device Family, Client Site, and Server Site. Entering a value in one of these fields filters the results displayed in the graphs below. Clear the data by clicking the Clear (\otimes) button to the right of the filter options.

Network Response Time Graph

The Network Response Time graph displays the response time (in milliseconds) the TCP request took to complete for the Top N slowest Targets. The data in this graph depends on the criteria you select in the toolbar at the top of the window and can be filtered to match specific criteria. Extreme Management Center displays data collected by the ExtremeAnalytics engine over the previous 20 minutes updated every 15 seconds. Use the **Pause** button in the toolbar to stop the graph from updating. Clicking the **Unpause** button resumes the updates and refreshes the graph with the most up-to-date data.



Hover over a point in the graph to see a pop-up with details about that application at that moment in time.

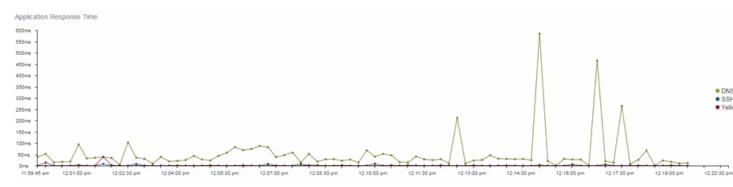
Clicking on a point opens a flow data table for that Target at that time at the bottom of the window, limited to match any <u>filters</u> you applied. Right-click a row in the flow to see additional options for working with that flow. Flows without an identified source are labeled with the device's IP Address.

Click the **Arrow** button () at the top of the flow data table to collapse the table and click the **Arrow** button () on the collapsed table to expand the table again.

Application Response Time Graph

The Application Response Time graph displays the response time (in milliseconds) the application request took to complete for the Top N slowest Targets. The data in this graph depends on the criteria you select in the toolbar at the top of the window and can be <u>filtered</u> to match specific criteria. Extreme Management Center displays data collected by the ExtremeAnalytics engine over the previous 20 minutes updated every 15 seconds. Use the **Pause** button in the toolbar to stop the graph from updating.

Clicking the **Unpause** button resumes the updates and refreshes the graph with the most up-to-date data.



Hover over a point in the graph to see a pop-up with details about that application at that moment in time.

Clicking on a point opens a flow data table for that Target at that time at the bottom of the window, limited to match any <u>filters</u> you applied. Right-click a row in the flow to see additional options for working with that flow.

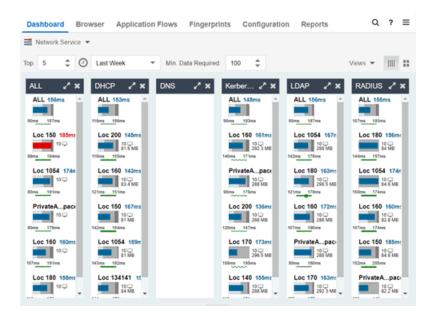
Click the **Arrow** button () at the top of the flow data table to collapse the table and click the **Arrow** button () on the collapsed table to expand the table again.

Related Information

ExtremeAnalytics

ExtremeAnalytics Network Service Dashboard

To access the Network Service Dashboard, open the **Analytics** > **Dashboard** tab and select **Network Service** in the dashboard drop-down list.



Overview

The Network Service Dashboard contains two graphs for each network service: the Expected Response Time bar graph displays the average response time over the selected time period and the Historical Response Time line graph displays the individual response times over that period for each site.

Select the number of sites displayed in each column in the **Top** field.

Use the **Time Period** drop-down list to display the date and time range for which data is displayed. Selecting **Custom** displays additional fields allowing you to indicate a **Start Date** and time and an **End Date** and time.

Use the **Minimum Required Response Time Dashboard Data Points** to configure the minimum amount of data Extreme Management Center requires before displaying a given application or site pair. The data below this threshold is not reliable and may set off a false alarm, however, you can adjust how much data is required based on the individual needs of your network.

The Network Service Dashboard displays the performance (in response time) of your network services. Each column in the dashboard represents a service:

- ALL
- DHCP
- DNS
- Kerberos

- LDAP
- RADIUS

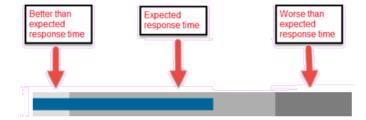
The top graphs for each service displays the average response time of all of the sites for that service, while the following rows indicate the top worst performing sites for that service.

You can display or hide any of the application columns using the **Views** drop-down list. You can also click the **X** at the top of a column to hide the column from the dashboard. Click the **Single Row** icon () to display all columns in a single row, or click the **Double Row** icon () to display the columns in two rows.

The worst performing sites are defined as those whose response time is the slowest when compared to the expected response time observed over the selected time period. For example, a site with an average RADIUS authentication response time of 40 ms over the past seven days that displayed a slowest response time of 50 ms would rank as a better performing site than a site with an average RADIUS authentication response time of 5 ms over the same period that displayed a slowest response time of 30 ms.

Expected Response Time

The Expected Response Time bar graph displays the range of response times, the most recently measured response time, and the expected response time for a network service a specific site during the date range you configure in the Date Range drop-down list. The value displayed on the far right of the graph is the slowest response time observed during the selected time period. The vertical green bar indicates the most recently observed response time for the network service.



Hover over the Expected Response Time graph to display a pop-up with the response time for the network service as well as the date and time the measurement occurred. The Expected Response Time bar graphs also display the client count, represented by a number and a monitor icon (10), and a client byte count observed as of the most recent measured minute. The client count is the number of clients using the service at the site. The client byte count indicates the amount of storage being utilized by clients. The data

used for the client count, the client byte count, and the reported response time are from the same recently observed minute.

NOTE: Client counts and client byte counts are not provided for the bar graphs that display the average response time of all the sites for that service.

Extreme Management Center uses a standard deviation of the values gathered as response times to determine the expected response time for a network service at a site. In the bar graph, the medium gray color indicates a response time that falls within the "expected" range. A response time in the light gray range is better than expected, while a response time in the dark gray is worse than expected.

When a response time is determined to be worse than expected, the site name and the response time indicator turn red to flag the service.

Clicking the Expected Response Time bar graph opens the Response Time dashboard (which is also accessible from the **Analytics > Dashboard** tab) filtered to display the network service. If you click the network service for a particular site, the Response Time dashboard also filters to that site.

Historical Response Time

The Historical Response Time line graph shows all of the response times observed for the network service at a site.



Hovering over a point in the graph causes a dot on the line graph to appear, indicating the point in the response time at which you are looking. Additionally, a pop-up with the date, time, and response time appears for that point.

This is the data set from which Extreme Management Center creates the Expected Response Time graph. The wider the expected response time range in the Expected Response Time graph (indicated by the medium gray color), the greater the variance in the values in this graph.

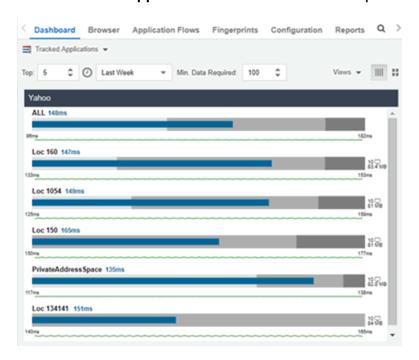
Related Information

ExtremeAnalytics tab

ExtremeAnalytics Tracked Applications Dashboard

The Tracked Application dashboard displays the performance (in response time) of your network for applications you configure in the **Tracked Applications** field on the **Analytics** > **Configuration** > **Configuration** tab.

To access the Tracked Application dashboard, open the **Analytics > Dashboard** tab and select **Tracked Applications** in the dashboard drop-down list.



Overview

The Tracked Applications dashboard contains two graphs for each application, one displays the average response time over the selected time period and the other displays the individual response times over that period for each site. Data is updated every minute and can be manually refreshed by clicking the **Refresh** button (2).

Select the number of sites displayed in each column in the **Top** field. The Tracked Applications dashboard can display up to 25 sites.

Use the **Time Period** drop-down list to display the date and time range for which data is displayed. Selecting **Custom** displays additional fields allowing you to indicate a **Start Date** and time and an **End Date** and time.

Use the Minimum Required Response Time Dashboard Data Points to configure the minimum amount of data Extreme Management Center requires before displaying a given application or site pair. The data below this threshold is not reliable and may set off a false alarm, however, you can adjust how much data is required based on the individual needs of your network.

Each column in the dashboard represents an application. The top row displays the average response time of all of the sites for that application, while the following rows indicate the top worst performing sites for that application.

You can display or hide any of the application columns using the **Views** drop-down list. You can also click the **X** at the top of a column to hide the column from the dashboard. Click the **Single Row** icon () to display all columns in a single row, or click the **Double Row** icon () to display the columns in two rows.

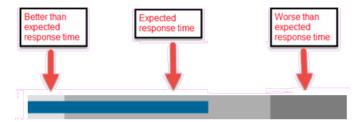
Click the **Maximize** icon (**M**) to expand a single application column.

The worst performing sites are defined as those whose response time is the slowest when compared to the expected response time observed over the selected time period. For example, a site with an average Microsoft Office 365 authentication response time of 40 ms over the past seven days that displayed a slowest response time of 50 ms would rank as a better performing site than a site with an average Microsoft Office 365 authentication response time of 5 ms over the same period that displayed a slowest response time of 30 ms.

Expected Response Time

The Expected Response Time bar graph displays the range of response times, the most recently measured response time, and the expected response time for an application a specific site during the date range you configure in the Date Range drop-down list. The value displayed on the far right of the graph is the slowest response time observed during the selected time period. The vertical blue or red bar indicates the most recently observed response time for the application.

NOTE: The values in this graph are an average of all response times observed every minute.



Hover over the Expected Response Time graph to display a pop-up with the most recent response time for the application as well as the date and time the measurement occurred. The Expected Response Time bar graphs also display the client count, represented by a number and a monitor icon (10), and a client byte count observed as of the most recent measured minute. The client count is the number of clients using the service at the site. The client byte count indicates the amount of storage being utilized by clients. The data used for the client count, the client byte count, and the reported application response time are from the same recently observed minute.

NOTE: Client counts and client byte counts are not provided for the bar graphs that display the average application response time of all the sites for that service.

Extreme Management Center uses the standard deviation of the values gathered as response times to determine the expected response time for an application at a site. In the bar graph, the medium gray color indicates a response time that falls within the "expected" range. This range is the average value of all observed response times plus or minus two standard deviations, or about 95 percent of all response time values. A response time in the light gray range is better than expected, while a response time in the dark gray is worse than expected.

When a response time is determined to be worse than expected, the site name and the response time indicator turn red to flag the application.



Clicking the Expected Response Time bar graph opens the Response Time dashboard filtered to display the application. If you click the application for a particular site, the Response Time dashboard also filters to that site.

Historical Response Time

The Historical Response Time line graph shows all of the response times observed for the application at a site.

NOTE: The values in this graph are an average of all response times observed every hour.

Hovering over a point in the graph causes a dot on the line graph to appear, indicating the point in the response time at which you are looking. Additionally, a pop-up with the date, time, and response time appears for that point.

This is the data set from which Extreme Management Center creates the Expected Response Time graph. The wider the expected response time range in the Expected Response Time graph (indicated by the medium gray color), the greater the variance in the values in this graph.

Related Information

ExtremeAnalytics tab

ExtremeAnalytics Browser Overview

The **Browse**r tab lets you query information about recent network activity stored in the Extreme Management Center database and display results in various grid and chart report formats. Using the Browser, you can create custom queries that provide greater flexibility in defining what data to display and how to display it. You can access the Browser from the Extreme Management Center **Analytics** tab.

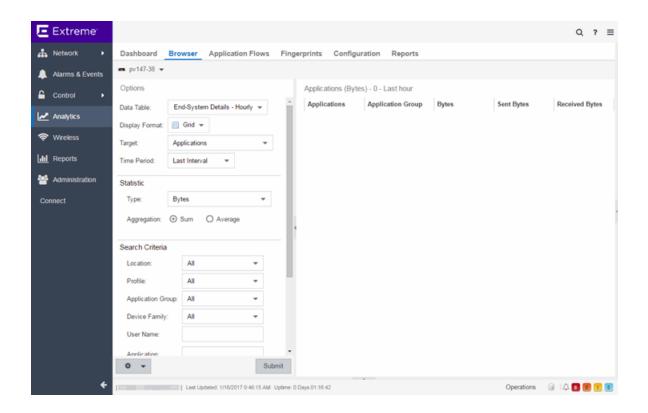
Overview

The Browser allows you to generate reports in several different formats using data based on selected options including a data target, statistic type, start time, and other search criteria.

For example, you can display application response time for the last hour or the last three days. You can view the results as a grid or a chart. You can filter the results to display data for a specific application or site.

If you have multiple ExtremeAnalytics engines, use the **Engine** drop-down list to select an engine to use as the source for the report data. Then, select the desired options on the left side of the Browser view and click **Submit**. The report is displayed on the right side of the view. Click on an item in the report to view details or right-click an item to select from other focused reports.

After you have generated a report, use the **Gear** menu (• •) (at the bottom left of the options panel) to (save) save it to the Report Designer to use as a custom component, () bookmark the report, or () export it as a CSV file.



Data Aggregation

Network data displayed in a report is aggregated from your network by the ExtremeAnalytics engine and sent to Extreme Management Center. The data gathering process begins with the ExtremeAnalytics engine, which monitors network activity on the switch or controller you configure using a traffic mirror and NetFlow or application telemetry. The traffic mirror gathers the first (N) packets of a flow to determine the application in use, while NetFlow (a flow-based data collection protocol) provides information about the amount of data sent and received for the application. The engine holds this information in its cache and transmits the aggregated data to Extreme Management Center every five minutes to update the High-Rate data table information and every hour to update the hourly data table information. Creating a report in the Applications Browser displays the information sent from the ExtremeAnalytics engine to Extreme Management Center based on the criteria you select.

NOTE: Information held in the ExtremeAnalytics engine's cache is not saved. Restarting the ExtremeAnalytics engine before the data in the memory cache is sent to Extreme Management Center results in the loss of that information.

Options

Following are definitions of the different options available when creating your custom query.

Data Table

Select which type of network activity data to query. The correct data table to use depends on the nature of the report.

- End-System Details Hourly End-system data collected every hour. Used when data for a specific client or server is needed, or when the information requested is highly specific, for example top applications used by Android devices in the London site.
- Application Data Hourly Application data collected every hour. Used for higher level information, such as top applications during an hour.
- Application Data High-Rate Application data collected at a higher rate (every five minutes). Used for a more detailed picture of how traffic changes over time.
- Application Telemetry Hourly Application Telemetry flow data collected every hour.

Display Format

Select the display format for the report: Grid, Chart Over Time, Word Cloud, Tree Map, or Bubble Map.

Target

Network traffic information is collected on objects in your network called targets. Some targets are physical, such as clients and servers, and some are logical, such as applications. Select the type of target that you want information about. Available targets vary depending on the selected data table. If you want information on a specific target, specify that target in the Search Criteria options.

- Applications An application in ExtremeAnalytics is identified through layer 7
 analysis of network traffic. For example, an application can be identified as
 Facebook.
- Application/Client Information about applications used by clients, or about clients using an application.
- Application/Device Family Information about applications used by device families, or about device families using an application.
- Application/Interface Information about the applications used by interfaces.

- Application/Profile Information about applications used by profiles, or about profiles using an application.
- Application/Server Information about applications accessed on a particular server, or about severs using an application.
- Application Groups Application categories, such as Cloud Computing or Social Networking, which are implied by the application.
- **Device Family** The kind of device determined for a client, such as Windows or iOS. Device information is only available for some network traffic.
- Interface/Applications Information about interfaces used by applications.
- Application-Interface Pair/Client Displays the applications and interfaces used by clients.
- Interface/Client Information about the interfaces used by clients.
- Sites <u>Sites</u> are used by ExtremeAnalytics to identify the physical location for the client of an application flow. A site is a set of IP address ranges that identify a portion of your network. Multiple sites can be created to identify different buildings, sites, or geographical areas of your network.
- Profiles A profile assigned to a client. Profile information is only collected under certain circumstances.
- Threat Displays a list of the threat classifications that occurred during the Time Period you select.
- Threat/Threat End-System Pair Displays a list of the threat classifications broken down by the IP addresses of the end-systems involved in the flow (the trusted and untrusted hosts) that occurred during the Time Period you select.
- Clients The end-point of a flow which has the client role for that connection.
- Servers The end-point of a flow which has the server role for that connection.
- **Total** The total values for all detected traffic for the interval used by the data table (hourly or high-rate).

Statistic

Statistics are quantitative data that can be collected for the selected target. Available statistics vary depending on the selected target. Select the desired statistic for the report:

- Bytes The number of bytes transferred in both directions, between the client and the server. Also known as bandwidth.
- **Flows** The number of NetFlow records sent by the switch to report the traffic between the client and the server.

- Application Response Time The average amount of time for a server to respond to a request.
- Network Response Time The average amount of time to create a connection.
- Received Bytes The number of bytes received by clients. This may be an
 estimated number of bytes if you are using an Application Telemetry flow.
- Sent Bytes The number of bytes sent by clients. This may be an estimated number of bytes if you are using an Application Telemetry flow.
- Inbound Flows The number of NetFlow records sent by the switch to report the server-to-client traffic. This is a rough indication of the duration of client connections.
- Outbound Flows The number of NetFlow records sent by the switch to report the client-to-server traffic. This is a rough indication of the duration of client connections.
- Clients The number of unique clients that have been seen associated with the target.
- Servers The number of unique servers that have been seen associated with the target.
- Application Count The number of unique applications seen for the selected target.

For byte, flow, and application count statistics, if you select a time range that is larger that the interval, specify whether you want the data aggregated as a summation of all the values for that statistic or as an average of all the values for that statistic.

Start Time

Select the start time (duration) for the report: Last Interval, Today, Yesterday, Last 24 Hours, Last 3 Days, or Last Week. You can also specify a custom start time and end time for the report. The Last Interval is the most recent recorded data covering a time period determined by the selected Data Table.

Search Criteria

Defining search criteria allows you to further filter the report data. Available criteria will vary depending on the selected data table and target. If you select either of the Application Data tables, you can only filter based on the selected target. For example, if you select **Sites** as your target, you can only filter on defined sites. If you select the End-System Details data table, you can filter on additional criteria. For example, if you select **Sites** as your target, you can filter on defined sites as well as flows for iOS devices.

You can enter a partial term in the text field or use the SQL wildcard "%" (as a substitute for multiple characters) or "_" (as a substitute for a single character) for multiple matches. For example, for the Device Family name, you could enter "iPhone %" to match iPhone 3, 4, and 5.

NOTE: Values entered in the text fields that contain multiple, non-alphanumeric characters may cause issues with the returned results. If this happens, alternate values should be used.

- Site Select a site to match or select World. If a site has been added to a map, you
 will also see a selection for that map. If you select custom, you can enter a partial
 site name or use the SQL wildcard characters to match one or more sites.
- Profile Select an ExtremeControl profile to match or select All. If you select custom, you can enter a partial profile name or use the SQL wildcard characters to match one or more profiles. Profile information is only collected under certain circumstances.
- Application Group Select an application group to match or select All. If you
 select custom, you can enter a partial application group name or use the SQL
 wildcard characters to match one or more groups.
- Device Family Select the operating system family to match or select All. If you
 select custom, you can enter a partial device family name or use the SQL wildcard
 characters to match one or more families. Device information is only available for
 some network traffic.
- **User Name** Enter a client's username to match. Username information is only available for some network traffic.
- Application Enter an application name to match.
- Client Enter a client's IP address or hostname to match.
- Engine Select the ExtremeAnalytics engine for which you are generating the report.
- Limit Select the number of results to return, for example, 10 clients.

Display Options

If you have selected Chart Over Time as your report display format, you can select whether to display the data as a line or an area, and also select the color to use in the chart.

Bookmark

After you have generated a report, click the Gear menu (• •) in the lower left corner to save the options you have currently set. A new window opens for the current report with

a link that can be bookmarked in your browser. You can then use the bookmark whenever you want the same search options.

Save to Report Designer

Click the Gear menu (• •) in the lower left corner to access the Save to Report Designer window. This window lets you save the currently defined report to use as a custom component in the Report Designer. The custom component uses the target, statistic, and start time currently defined in the Browser.

Enter a name for the custom component and select any search criteria that you want displayed in the component panel. The search criteria is displayed as fields in the component panel, providing a custom interface that lets you further refine report data. If no search criteria are selected, the saved component only uses the target, statistic, and start time definitions when requesting data, creating a view-only report.



Export to CSV

Click the Gear menu (• •) in the lower left corner and click () to export the report data as a CSV file. The currently defined report opens in a spreadsheet, which can then be saved.

Related Information

ExtremeAnalytics tab

ExtremeAnalytics Application Flows

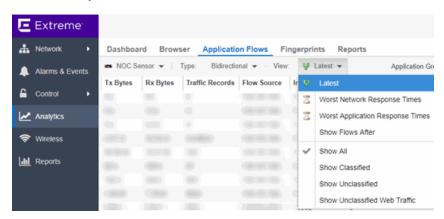
The **Application Flows** tab displays tables that present <u>Historical</u>, <u>Bidirectional</u> or <u>Unidirectional</u> client, server, and application flow data. To access the **Applications Flows** tab, open **Analytics > Application Flows**.

This Help topic provides information on the following topics:

- Overview
- Application Flows Tables
- Report Features

Overview

The **Application Flows** tab includes several functions that allow you to filter and customize your table data.



Appliance Engine

If your network uses multiple ExtremeAnalytics engines, use the **Engine** menu to select an engine to use as the source for the flow data.

Type

Use the **Type** menu to select whether to display <u>Historical</u>, <u>Bidirectional</u> (aggregate flows) or <u>Unidirectional</u> (base flows) flow data.

View

By default, the table displays the latest flows collected. Use the **View** menu to select different display options. The available options vary depending the flow type (bidirectional or unidirectional) selected.

- Latest Displays the latest flows collected by the specified engine.
- Worst Network Response Times Sorts the flows based on the worst TCP response time and displays the flows with the worst time at the top of the chart.
- Worst Application Response Times Sorts the flows based on the worst application response time and displays the flows with the worst time at the top of the chart.
- Show Flows After Allows you to select a start date and time for the flows displayed.
- Show All Show all flows.
- Show Classified Show only flows classified by an application fingerprint.
- Show Unclassified Show only flows not classified by an application fingerprint.
- Show Unclassified Web Traffic Show only web traffic that has not been classified by an application fingerprint.

Application Group

Use the **Application Group** menu to filter the table by application group.

Search

Use the **Search** field at the top right of the table to filter specific flow information. For example, searching on "snmp" or "10.20.30.131/24" filters the table so only flow data related to SNMP or the given subnet is displayed. You can enter one or more filters simultaneously, separated by semicolons. Individual components of a filter is separated by commas. For complete instructions on how to use the Flow Search, rest your cursor on the **Search** field and read the tooltip (select on the "more" link in the tooltip). Press the **Reset** button at the bottom left of the window to clear the Search results and refresh the table.

You can also use the **Search** field to search for a specific application, user name, or IP address from your filtered results:

- 1. Select a user name or IP address from the filtered search results to launch PortView, which provides a detailed topology context for the user.
- Enter meta= before the term for which you are searching includes all variations of that search term in the result set. For example, entering meta=extreme returns extremenetworks.com, www.extremenetworks.com, extreme.boston.com, and any other flows that include the word "extreme".

- 3. Right-click on a flow to access a menu of options including the ability to:
 - Add a new custom fingerprint based on the flow selected in the table.
 - Show all fingerprints associated with the application in the selected flow.
 - Create a UDP or TCP rule using the IP port.
 - Search Extreme Management Center maps for the selected flow client.
 - Open a Flow Details report for the selected flow (bidirectional flows only).
 - Access a variety of reports for the flow.

Refresh

Use the **Refresh** drop-down list at the top right of the window to specify an interval (in seconds) at which the flows data automatically refreshes. To stop auto refresh, select the **Refresh Off** option.

Application Flows Tables

The columns included in the Application Flows tables vary, depending on the type of data flow you select (Historical, Bidirectional and Unidirectional). Additionally, right-click and select **Start Packet Capture** to save a packet capture of the flow on the **Packet Captures** tab.

Historical Flows

The Historical table displays short-term flow data storage you can use to determine trends in your network.

Bidirectional Flows

The Bidirectional table displays bidirectional flow data stored in memory. It provides aggregated flow data for a given client, server, server port, application, and protocol. All matching flows are aggregated to show the flow count, total duration, amount of data transmitted, and additional information. The bidirectional report presents flow data for real-time troubleshooting purposes, and is not designed for historical long-term flow collection. A checkmark (*) in the table denotes a tracked application or a tracked site.

Unidirectional Flows

The Unidirectional table displays unidirectional flow data stored in memory. It provides the raw non-aggregated flow data received from the flow sensors on the network. It

presents flow data for real-time troubleshooting purposes, and is not designed for historical long-term flow collection. A checkmark (•) in the table denotes a tracked application or a tracked site.

Report Features

The Application Flows tables include several report features and functions that allow you to drill down for more detailed application, site, response time, mapping and policy functions. The report features vary, depending on the type of data flow you select (Historical, Bidirectional and Unidirectional).

Interactive Tables

Manipulate table data in several ways to customize the view for your own needs:

- Select on the column headings to perform an ascending or descending sort on the column data.
- **Hide or display different columns** by selecting on a column heading dropdown arrow and selecting the column options from the menu.
- Filter data in each column by selecting on a column heading drop-down arrow and using the Filters option on the menu.

The sort and filter functionality for these two tables behaves differently than for other Extreme Management Center tables. In these tables, Max Rows are considered for display, and then sorting and filtering is applied to these rows. In other tables, sorting and filtering is applied to the entire table, and then Max Rows of the result is displayed. For example, if the Max Rows value is set to 50 and you create a filter for a specific IP address, only those 50 rows will be filtered for the IP, not all the flows maintained in memory on the server.

CSV Export 📓

The <u>CSV Export button</u> allows you to save report data to a CSV file and to provide report data in table form.

Bookmark Bookmark

Use the <u>Bookmark button</u> to save the search, sort, and filtering options you have currently set. It opens a new window for the current report with a link that can be bookmarked in your browser. You can then use the bookmark whenever you want the same search, sort, and filtering options.

Max Rows

By default, the top 100 entries are displayed in the table. However, you can change this value using the Max Rows field at the bottom of the view.

Reset B Reset

The <u>Reset button</u> enables you to clear the search fields and all filters, and to refresh the table.

Aggregate / Base Flows

Aggregate Flows (bidirectional table) and Base Flows (unidirectional table) data uses an X number of days, hh:mm:ss format and includes Current Load and Peak Load calculations in flows per second.

Related Information

ExtremeAnalytics tab

ExtremeAnalytics Bidirectional Flow Table

This table on the **Application Flows** tab displays bidirectional flow data that is stored in memory. Use it to view aggregated flow data for a given client, server, server port, application, and protocol. All matching flows are aggregated to show the flow count, total duration, amount of data transmitted, and additional information. The bidirectional report presents flow data for real-time troubleshooting purposes, and is not designed for historical long-term flow collection. A checkmark () in the table denotes a tracked application or a tracked site.

By default, the top 100 entries are displayed in the table. However, you can change this value using the Max Rows field at the bottom of the view.

Text at the bottom of the table shows:

- The CSV Export icon <a> - allows you to save report data to a CSV file and to provide report data in table form
- Aggregate Flows data uses an X number of days, hh:mm:ss format and includes
 Current Load and Peak Load calculations in flows per second

Following are definitions for the table columns:

Flow Summary

Rest the cursor over the first column in the table and click the ▼ arrow to open the Flow Summary window. Flow summary information can include response times, Uniform Resource Identifier, and header data for the flow. In the Flow Summary window, use the Menu icon ≡ to access additional functionality, such as the ability to modify the application fingerprint or create a policy rule.

Flows

The number of base flows included in the aggregate flow. Click on a link in the Flows column to open a **Flow Details** tab that displays the individual flows that contributed to the aggregate flow.

Client Address

The IP address or hostname of the system where the flow originated. Click on the Client address link to open a **PortView** for the client (if it is in the database) or a **PortView** for the switch configured as the NetFlow sensor.

Server Address

The IP address or hostname of the server handling the flow.

Server Port

Either the TCP or UDP port on the server handling the flow.

Application

The name of the application as identified by the ExtremeAnalytics engine using the Fingerprint database.

Application Group

The flow application group to which the application belongs.

Application Info

Additional information about the flow provided by the ExtremeAnalytics engine. Hover over the flow and a table of the information displays.

Type

The content type of a flow, such as sound, video, or text. Click on the **Type** icon to open the flow's URI.

Network Response

The response time (in milliseconds) that it took for the TCP request to complete.

Application Response

The response time (in milliseconds) that it took the application request to complete.

Site

The name of the site that matches the client's IP address.

Detailed Site

The client's switch IP and switch port (wired), or controller IP, AP, and SSID (wireless).

Device Family

The operating system family for the client end-system.

User

The username used when the client system connected.

Profile

The Extreme Management Center profile assigned to the client end-system.

Threat

Indicates if the flow contains potential threat activity from IP addresses known to be suspicious. IP addresses can be flagged as suspicious for a variety of reasons, including forced IP anonymity through the use of a Tor exit node, being listed as a threat by the Emerging Threats project, or classified as suspicious by internet users.

Protocol

The connection type protocol used by the flow.

Last Seen Time

The last time a unidirectional (base) flow was aggregated into this bidirectional flow.

Duration

The duration of a bidirectional (aggregate) flow is the sum of the durations of the unidirectional (base) flows that make up the bidirectional flow. The duration of a bidirectional flow may be greater than or less than the period of time indicated by the **First Seen** and **Last Seen Time**. This is because there may be times during that time period when no flow is active or when several flows are active at the same time.

NOTE: Bidirectional flows may be greater than the period of time between the First Seen and Last Seen Time columns because they display the sum of all flow records for a client and a server on a server port. For a flow that lasts for 60 seconds, there are two flow records (a client to server flow and a server to client flow), so the total duration may exceed 60 seconds. Multiple simultaneous connections from the client to the same server port (e.g. multiple browser windows open to a web-based email client) can also increase the duration.

Rate

The average bandwidth for the flow based on the total flow duration. Because bandwidth calculations are based on the total duration (not on the **First Seen** and **Last Seen Time**), they represent the average throughput for each flow considered separately, not as an aggregate.

Tx Packets

The number of packets transmitted for this flow. For flows collected via Application Telemetry, this number may be estimated.

Rx Packets

The number of packets received for this flow. For flows collected via Application Telemetry, this number may be estimated.

Tx Bytes

The number of bytes transmitted for this flow. For flows collected via Application Telemetry, this number may be estimated.

Rx Bytes

The number of bytes received for this flow. For flows collected via Application Telemetry, this number may be estimated.

Traffic Records

The number of records received in each flow.

Flow Source

The IP address of the NetFlow source switch, Application Telemetry source switch, or wireless controller sending the NetFlow data to the NetFlow collector.

Input Interface

The interface receiving the flow on the NetFlow sensor.

Output Interface

The interface transmitting the flow on the NetFlow sensor.

Client TOS

The DSCP (Diffserv Codepoint) value for the client to server flow. The TOS/DSCP value is used to configure quality of service for network traffic.

Server TOS

The DSCP (Diffserv Codepoint) value for the server to client flow. The TOS/DSCP value is used to configure quality of service for network traffic.

TTL

The TTL (IP Time to Live) value of the flow. 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. When the value hits zero, the packet is dropped.

Related Information

- ExtremeAnalyticstab
- Application Flows tab

ExtremeAnalytics Unidirectional Flow Table

This table on the **Application Flows** tab displays unidirectional flow data stored in memory. It displays the raw, non-aggregated flow data received from the flow sensors on the network. It presents flow data for real-time troubleshooting purposes, and is not designed for historical long-term flow collection. A checkmark (**) in the table denotes a tracked application or a tracked site.

Hover over an application in the table to display switch data, which is an accumulation of multiple switches into single flow record, as well as the path that flow has taken.

By default, the top 100 entries are displayed in the table. However, you can change this value using the Max Rows field at the bottom of the view.

Text at the bottom of the table shows Base Flows, using X number of days, hh:mm:ss format, and including Current Load and Peak Load calculations in flows per second.

Following are definitions for the table columns:

Flow Summary

Rest the cursor over the first column in the table and click the ▼ arrow to open the Flow Summary window for a specific flow. Flow summary information can include response times, Uniform Resource Identifier, and header data for the flow. In the Flow Summary window, use the Gear menu ≡ to access additional functionality such as the ability to modify the application fingerprint or create a policy rule.

Client/Server Flows

Identifies whether the flow is a Client Flow or a Server Flow. The client/server direction of a flow is calculated by the ExtremeAnalytics engine. Hover over the icon to see a tooltip with more information.

Source Address

The IP address or hostname of the system where the flow originated. Click on the Source address link to open a **PortView** for the client or server (if it is in the database) or a **PortView** for the switch configured as the NetFlow sensor.

Source Port

Either the TCP or UDP port on the client/server handling the flow.

Destination Address

The IP address or hostname of the system that received the flow.

Destination Port

Either the TCP or UDP port on the system that received the flow.

Application

The name of the application as identified by the ExtremeAnalytics engine using the Fingerprint database.

Application Group

The flow application group to which the application belongs.

Application Info

Additional information about the flow provided by the ExtremeAnalytics engine.

Type

The content type of a flow, such as sound, video, or text. Click on the **Type** icon to open the flow's URI.

Network Response

The response time (in milliseconds) that it took for the TCP request to complete.

Application Response

The response time (in milliseconds) that it took the application request to complete.

Site

The site where the flow originated.

Detailed Site

The client's switch IP and switch port (wired), or controller IP, AP, and SSID (wireless).

Device Family

The operating system family for the client end-system.

User

The username used when the client system connected.

Profile

The ExtremeControl profile assigned to the client end-system.

Protocol

The connection type protocol used by the flow.

Last Seen Time

The last time the flow was seen.

Duration

The amount of time that the flow was active.

Rate

The average bandwidth for the flow based on the flow duration.

Packets

The number of packets in this flow. For flows collected via Application Telemetry, this number may be estimated.

Bytes

The number of bytes in this flow. For flows collected via Application Telemetry, this number may be estimated.

NetFlow Records

The number of NetFlow records for this flow.

Flow Source

The IP address of the NetFlow source switch, Application Telemetry source switch, or wireless controller sending the Flow data to the Flow collector.

Input Interface

The interface receiving the flow on the Flow sensor.

Output Interface

The interface transmitting the flow on the Flow sensor.

TOS

The DSCP (Diffserv Codepoint) value for the flow. The TOS/DSCP value is used to configure quality of service for network traffic.

TTL

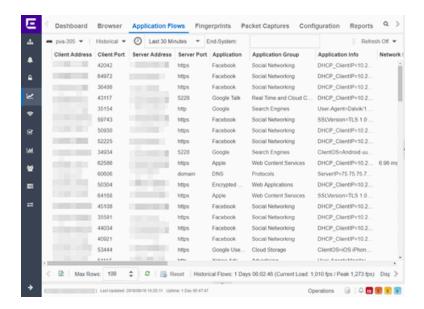
The TTL (IP Time to Live) value of the flow. The TTL field indicates the maximum number of router hops the packet can make before being discarded. The TTL field is set by the packet sender and reduced by every router on the route to its destination. When the value hits zero, the packet is dropped.

Related Information

- ExtremeAnalytics tab
- Application Flows tab

ExtremeAnalytics Historical Flow Table

The Historical Flow table in the **Application Flows** tab displays historical flow data that you can use to analyze trends in your network. The Historical Flow table is not designed for long-term flow collection.



Hover over an application in the table to display switch data, which is an accumulation of multiple switches into single flow record, as well as the path that flow has taken.

By default, the top 100 entries are displayed in the table. However, you can change this value using the Max Rows field at the bottom of the view.

Text at the bottom of the table shows Base Flows, using X number of days, hh:mm:ss format, and including Current Load and Peak Load calculations in flows per second.

Following are definitions for the table columns:

Client Address

The IP address or hostname of the system where the flow originated. Click on the Client address link to open a **PortView** for the client (if it is in the database) or a **PortView** for the switch configured as the NetFlow sensor.

Client Port

Either the TCP or UDP port on the client handling the flow.

Server Address

The IP address or hostname of the server handling the flow.

Server Port

Either the TCP or UDP port on the server handling the flow.

Application

The name of the application as identified by the ExtremeAnalytics engine using the Fingerprint database.

Application Group

The flow application group to which the application belongs.

Application Info

Additional information about the flow provided by the ExtremeAnalytics engine.

Network Response

The response time (in milliseconds) that it took for the TCP request to complete.

Application Response

The response time (in milliseconds) that it took the application request to complete.

Start Time

The start time.

Duration

The amount of time that the flow was active.

Last Seen

The last time the flow was seen.

Client Site

The name of the site that matches the client's IP address.

Server Site

The site of the server.

Detailed Site

Detailed site information

Protocol

The connection type protocol used by the flow.

Rate

The average bandwidth for the flow based on the flow duration.

Client Bytes

The number of bytes in this flow. For flows collected via Application Telemetry, this number may be estimated.

Client Packets

The number of packets in this flow. For flows collected via Application Telemetry, this number may be estimated.

Server Bytes

The number of bytes in this flow. For flows collected via Application Telemetry, this number may be estimated.

Server Packets

The number of packets in this flow. For flows collected via Application Telemetry, this number may be estimated.

Analytics Appliance

The engine to which the server is assigned.

Related Information

- ExtremeAnalytics
- Application Flows

ExtremeAnalytics Fingerprints Overview

The **Fingerprints** tab provides detailed information about fingerprints used by ExtremeAnalytics to identify application flows. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. They can be created based on flow, application or application group, or a destination address. For applications such as Facebook and Google, multiple fingerprints are included to capture the different ways these applications can be used.

Fingerprints are created and stored on the Extreme Management Center server. When a fingerprint is changed or enabled, a flag is raised on the ExtremeAnalyticsengine to show it needs enforcing. Access the Browser from the Extreme Management Center **Analytics** tab.

There are two types of fingerprints: system fingerprints and custom fingerprints.

System fingerprints are provided by Extreme Management Center. They cannot be deleted; however, they can be modified or disabled. When a system fingerprint is modified, it results in a new custom fingerprint that overrides the original system fingerprint.

Custom fingerprints are either new user-defined fingerprints or modifications of system fingerprints. Custom fingerprints can be deleted. If a custom fingerprint was overriding a system fingerprint, then deleting the custom fingerprint will reload the original system fingerprint.

Related Information

ExtremeAnalytics tab

ExtremeAnalytics Custom Fingerprints

Custom fingerprints are either new user-defined fingerprints or modifications of system fingerprints. Custom fingerprints can be deleted. If a custom fingerprint was overriding a system fingerprint, then deleting the custom fingerprint will reload the original system fingerprint.

The <u>Fingerprints</u> view is divided into a left-panel tree and a table with six <u>columns</u>. The left-panel tree displays all the <u>application groups</u> and the fingerprints assigned to that group. The table on the right displays detailed information for each fingerprint. You can filter the information displayed in the table by selecting a single application group or fingerprint in the left-panel.

Fingerprint Table

The Fingerprint table displays detailed fingerprint information. Above the table, in the top left corner, is a Menu icon ≡, where you can access various system and fingerprint actions.

If you have multiple ExtremeAnalytics engines, an **Engine** menu is available that allows you to select an engine to use as the source for the fingerprint <u>Matches</u> data.

Use the **In Use** checkbox to filter the table to only show fingerprints that have had a match for the selected engine. Use the **Customized** checkbox to filter the table to display only custom fingerprints.

Menu

Use the **Menu icon** ≡ to access the following system and fingerprint actions. (You must have a fingerprint selected to enable the **Fingerprint** menu options.) Most of the options are also available by right-clicking on a fingerprint.

- Create Fingerprint Add a new fingerprint.
- Modify Fingerprint Change a fingerprint's description.
- Reset Fingerprint Counters Reset the Matches counters.
- Delete Custom Fingerprint Delete custom fingerprints, which can be identified by a ✓ in the Custom column.
- Fingerprint Definition View the XML definition for a fingerprint.

Column Definitions

Following are definitions for the table columns. All columns are sortable in ascending and descending order and can be filtered by text or numeric values.

Application

Name of the application this fingerprint detects. Click on an **Application** link to view client, flow, and usage information for that specific application.

Fingerprint

Name of the fingerprint.

Confidence

Reliability of this fingerprint. Higher confidence fingerprints override lower confidence fingerprints when determining a match for a traffic flow. The values are from 1 to 100, with 100 being absolutely reliable.

Custom

A check mark \checkmark indicates the fingerprint is a custom (user-defined) fingerprint. It is custom if it is a new fingerprint that has been added, a system fingerprint that has been modified, or a system fingerprint that has been disabled.

Application Group

The group this fingerprint's application belongs to. Application groups organize fingerprints into different types of applications such as Web applications or Business applications. You can sort the **Application Flows** view by application group, making it easier to view data for a specific type of flow. An application may only belong to one application group.

Matches

The total number of times a traffic flow has matched this fingerprint for the selected engine. A match is an occurrence of the ExtremeAnalytics engine making a final determination that a flow matches a fingerprint after all refinements are completed. The corresponding flow in the opposite direction, if there is one, is also matched. See Notes below.

NOTES:

- Matches are stored and displayed per engine. If you have multiple engines, use the Engine menu to select an engine to use as the source for the Hits and Matches data.
- If a flow generates hits on multiple fingerprints, and one fingerprint has a higher confidence than another fingerprint, a hit is counted for each fingerprint, but a match is only recorded for the final, highest confidence fingerprint.
- If you need to reset the Matches counters, use the Reset Fingerprint Counters option from the Menu icon (≡).

Type

The fingerprint type refers to how the fingerprint determines a match.

- FlexFire These fingerprints execute specific matching algorithms encoded into the engine. Disabling the fingerprint disables the specific code that implements the fingerprint.
- PCRE These fingerprints search using Perl Compatible Regular Expressions (PCRE).
- Port-based These fingerprints search for traffic on a specific port (typically, server-only ports). These are very low-confidence fingerprints and are generally just used for wider coverage.
- Web-App Rule These fingerprints search for a specific hostname in the URI of web requests.
- SSL Name These fingerprints search for values in the SSL common name.
- Http Host These fingerprints search for values in the HTTP hostname.
- Decoder These fingerprints extract protocol metadata from a flow that is provided when we generate a match on that flow.
- General Any fingerprint that isn't included in one of the other types. Typically, these fingerprints search for a straight pattern, or for a specific port and/or IP address with custom fingerprints (excluding custom Web-App Rule fingerprints).

Enabled

A \checkmark indicates the fingerprint is enabled. When a fingerprint is enabled, it will be used to identify applications. When it is disabled, it will be ignored.

Last Modified

Date that the fingerprint was last modified.

Created

Date that the fingerprint was created.

Description

Description of the fingerprint.

Related Information

• ExtremeAnalytics tab

Delete Custom Fingerprints

ExtremeAnalytics uses fingerprints to identify to which application a network traffic flow belongs. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. Extreme Management Center provides thousands of system fingerprints with the ExtremeAnalytics feature. In addition, you can modify these fingerprints and create new custom fingerprints.

Deleting a Custom Fingerprint

Delete a custom fingerprint from the **Fingerprints** tab. A custom fingerprint is either a new user-defined fingerprint, a modification of a system fingerprint, or a disabled fingerprint. (Custom fingerprints can be identified by a ✓ in the Custom column.)

When you delete a custom fingerprint, it is removed entirely. If you delete a custom fingerprint overriding a system fingerprint, the original system fingerprint is reloaded. System fingerprints that have not been modified cannot be deleted, however, they can be disabled.

Use these steps to delete a custom fingerprint:

- 1. Select the **Analytics** tab in Extreme Management Center and then select the Fingerprints view
- 2. Right-click on the desired custom fingerprint in the Fingerprints table and select **Delete Custom Fingerprint**. The Delete Fingerprint window opens.



- 3. You can delete only the selected fingerprint or select the option to delete all custom fingerprints that match the application name of the selected fingerprint.
- 4. Click **OK**. If a custom fingerprint overrides a system fingerprint, then deleting the custom fingerprint reloads the original system fingerprint.
- 5. Enforce to push the change to your engines.

Related Information

• ExtremeAnalytics tab

Custom Fingerprint Examples

The ExtremeAnalytics feature uses fingerprints to identify to which application a network traffic flow belongs. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. Extreme Management Center provides thousands of system fingerprints with the ExtremeAnalytics feature. In addition, you can create new custom fingerprints.

For additional information, see Getting Started with ExtremeAnalytics.

This Help topic provides examples of three different types of custom fingerprints you can create:

- · Fingerprints Based on a Flow
- Fingerprints Based on an Application or Application Group
- Fingerprints Based on a Destination Address

For additional information, see Add and Modify Fingerprints.

Fingerprints Based on a Flow

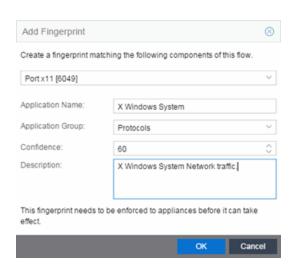
This example demonstrates how to create a custom fingerprint based on X Window System network traffic.

In the Extreme Management Center Flows table (with the Show Unclassified View selected) you notice several flows that had an X Window System source port 6049. Since these flows are not currently identified with a fingerprint, you can create a fingerprint for those flows based on the port that x11 traffic normally runs over.

Use the following steps to create the fingerprint.

- 1. Select the **Analytics** tab.
- 2. Select the **Application Flows** tab.
- 3. In the table, select the **Show Unclassified View**.
- 4. Right-click on a flow with the x11 Source Port and select Fingerprints > Add Fingerprint.

5. The Add Fingerprint window opens.



- 6. Use the drop-down list to select matching **Portx11** [6049].
- 7. Set the Application Name to X Window System.
- 8. Set the Application Group to Protocols.
- 9. Set the **Confidence** level to **60** (the default). A fingerprint with a confidence higher than 60 can supersede this fingerprint, if it also matches the flow.
- 10. Click **OK** to create the fingerprint.
- 11. Enforce to push the new fingerprint to your engines.

Fingerprints Based on an Application or Application Group

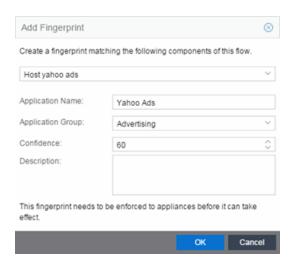
This example demonstrates how to create a fingerprint for some unclassified web traffic.

In the Extreme Management Center Application Flows table (with the Show Unclassified Web Traffic View selected) you noticed several flows for the "yahoo ads" application that are part of the Web Applications group. You want to create a fingerprint that provides an application and application group specifically for this traffic, instead of letting it default to the Web Applications group. The new fingerprint categorizes "yahoo ads" flows into the Yahoo Ads Id application and the Advertising application group.

Use the following steps to create the fingerprint.

- 1. Select the Analytics tab in Extreme Management Center.
- 2. Select the **Application Flows** tab.

- 3. In the table, select the **Show Unclassified Web Traffic View**.
- Right-click on a flow with the yahoo ads application and select Fingerprints > Add Fingerprint.
- The Add Fingerprint window opens.



- 6. Use the drop-down list to select matching the "yahoo ads" host.
- 7. Set the Application Name to Yahoo Ads.
- 8. Set the Application Group to Advertising.
- 9. Set the **Confidence** level to **60** (the default). A fingerprint with a confidence higher than 60 can supersede this fingerprint, if it also matches the flow.
- 10. Click **OK** to create the fingerprint.
- 11. Enforce to push the new fingerprint to your engines.

Fingerprints Based on a Destination Address

In both of the previous examples, you created a new custom fingerprint to cover a case where no appropriate fingerprint existed. You may also want to create a new fingerprint for traffic flows already identified as one application, but should be categorized as something else.

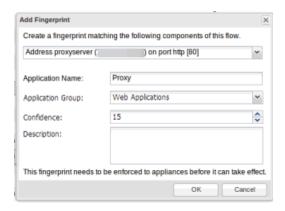
For example, let's say you have a Git repository on your network. Git repositories (a source code management system used in software development) are frequently accessed via SSH on port 22 (the standard TCP port assigned for SSH traffic). In this case, the SSH traffic flows is identified using the system SSH port-based fingerprint.

But what if you would like to more closely monitor who is accessing the Git repository? If you know you are running the Git server on a certain system (10.20.117.102 port 22, for our example), you can create a custom fingerprint to identify the Git traffic flows.

The fingerprint is based on one of the SSH flows using the IP address/port of the Git server and have a higher confidence than the system port-based fingerprint. The higher confidence fingerprint will override the lower confidence fingerprint when determining a match for the traffic flow.

Use the following steps to create the fingerprint.

- 1. Select the **Analytics** tab in Extreme Management Center.
- 2. Select the **Application Flows** tab.
- 3. In the table, right-click on an SSH port-based flow with the Git server destination address and select **Fingerprints > Add Fingerprint**.
- 4. The Add Fingerprint window opens.



- 5. Use the drop-down list to select matching the Git server IP address and port.
- 6. Set the **Application Name** to **Git**.
- 7. Select an Application Group that makes the most sense for your network. It might be Web Collaboration, Databases, Business Applications, or Storage. You can also create a new Application Group using the Create Custom Application Group option available from the gear menu in the Fingerprint Details tab. (You would need to do this before you create the custom fingerprint.)
- 8. Set the **Confidence** level to **60**, which is a higher confidence than the current fingerprint which is set at 10.
- 9. Click **OK** to create the fingerprint.
- 10. Enforce to push the new fingerprint to your engines.

Related Information

For information on related ExtremeAnalytics topics:

- Analytics
- Add and Modify Fingerprints

Create Custom Fingerprints Based on Flow

The ExtremeAnalytics feature uses fingerprints to identify to which application a network traffic flow belongs. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. Extreme Management Center provides thousands of system fingerprints with the ExtremeAnalytics feature. In addition, you can create new custom fingerprints.

Creating Fingerprints Based on a Flow

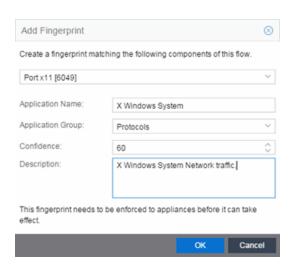
This example demonstrates how to create a custom fingerprint based on X Window System network traffic.

In the Extreme Management Center Flows table (with the Show Unclassified View selected) you notice several flows that had an X Window System source port 6049. Since these flows are not currently identified with a fingerprint, you can create a fingerprint for those flows based on the port that x11 traffic normally runs over.

Use the following steps to create the fingerprint.

- 1. Select the **Analytics** tab.
- 2. Select the Application Flows tab.
- 3. In the table, select the Show Unclassified View.
- 4. Right-click on a flow with the **x11 Source Port** and select **Fingerprints > Add Fingerprint**.

5. The Add Fingerprint window opens.



- 6. Use the drop-down list to select matching Portx11 [6049].
- 7. Set the Application Name to X Window System.
- 8. Set the Application Group to Protocols.
- 9. Set the **Confidence** level to **60** (the default). A fingerprint with a confidence higher than 60 can supersede this fingerprint, if it also matches the flow.
- 10. Click **OK** to create the fingerprint.
- 11. Enforce to push the new fingerprint to your engines.

Related Information

ExtremeAnalytics tab

Create Custom Fingerprints Based on Destination Address

The ExtremeAnalytics feature uses fingerprints to identify to which application a network traffic flow belongs. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. Extreme Management Center provides thousands of system fingerprints with the ExtremeAnalytics feature. In addition, you can create new custom fingerprints.

Creating Fingerprints Based on a Destination Address

Often, you will create a new custom fingerprint to cover a case where no appropriate fingerprint existed. However, you may also want to create a new fingerprint for traffic flows already identified as one application, but should be categorized as something else.

For example, let's say you have a Git repository on your network. Git repositories (a source code management system used in software development) are frequently accessed via SSH on port 22 (the standard TCP port assigned for SSH traffic). In this case, the SSH traffic flows is identified using the system SSH port-based fingerprint.

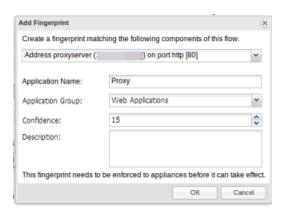
But what if you would like to more closely monitor who is accessing the Git repository? If you know you are running the Git server on a certain system (10.20.117.102 port 22, for our example), you can create a custom fingerprint to identify the Git traffic flows.

The fingerprint is based on one of the SSH flows using the IP address/port of the Git server and have a higher confidence than the system port-based fingerprint. The higher confidence fingerprint will override the lower confidence fingerprint when determining a match for the traffic flow.

Use the following steps to create the fingerprint.

- Select the Analytics tab in Extreme Management Center.
- 2. Select the **Application Flows** tab.
- 3. In the table, right-click on an SSH port-based flow with the Git server destination address and select **Fingerprints > Add Fingerprint**.

4. The Add Fingerprint window opens.



- 5. Use the drop-down list to select matching the Git server IP address and port.
- 6. Set the **Application Name** to **Git**.
- 7. Select an Application Group that makes the most sense for your network. It might be Web Collaboration, Databases, Business Applications, or Storage. You can also create a new Application Group using the Create Custom Application Group option available from the gear menu in the Fingerprint Details tab. (You would need to do this before you create the custom fingerprint.)
- 8. Set the **Confidence** level to **60**, which is a higher confidence than the current fingerprint which is set at 10.
- 9. Click **OK** to create the fingerprint.
- 10. Enforce to push the new fingerprint to your engines.

Related Information

ExtremeAnalytics tab

Create Custom Fingerprints Based on Application or Application Group

The ExtremeAnalytics feature uses fingerprints to identify to which application a network traffic flow belongs. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. Extreme Management Center provides thousands of system fingerprints with the ExtremeAnalytics feature. In addition, you can create new custom fingerprints.

Creating Fingerprints Based on an Application or Application Group

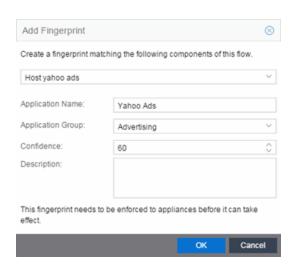
This example demonstrates how to create a fingerprint for some unclassified web traffic.

In the Extreme Management Center Application Flows table (with the Show Unclassified Web Traffic View selected), several flows for the "yahoo ads" application are part of the Web Applications group. The following instructions will allow you to create a fingerprint that provides an application and application group specifically for this traffic, instead of letting it default to the Web Applications group. The new fingerprint categorizes "yahoo ads" flows into the Yahoo Ads Id application and the Advertising application group.

Use the following steps to create the fingerprint.

- 1. Select the **Analytics** tab in Extreme Management Center.
- 2. Select the **Application Flows** tab.
- 3. In the table, select the **Show Unclassified Web Traffic View**.
- 4. Right-click on a flow with the yahoo ads application and select **Fingerprints > Add Fingerprint**.

5. The Add Fingerprint window opens.



- 6. Use the drop-down list to select matching the "yahoo ads" host.
- 7. Set the Application Name to Yahoo Ads.
- 8. Set the Application Group to Advertising.
- 9. Set the **Confidence** level to **60** (the default). A fingerprint with a confidence higher than 60 can supersede this fingerprint, if it also matches the flow.
- 10. Click **OK** to create the fingerprint.
- 11. Enforce to push the new fingerprint to your engines.

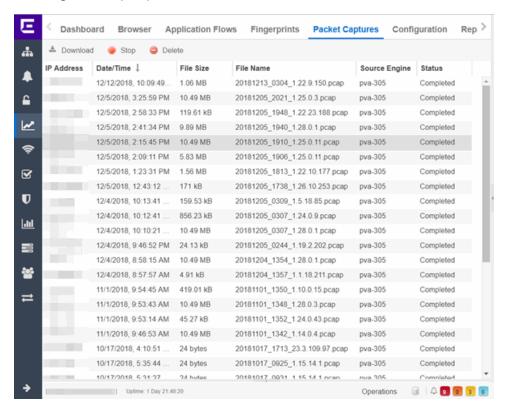
Related Information

ExtremeAnalytics tab

ExtremeAnalytics Packet Captures

Packet Captures (pcaps) consist of data included in flows collected by the ExtremeAnalytics engine you can use to analyze and assess the activity and traffic flow to and from IP addresses accessing devices in your network.

The **Packet Captures** tab displays a table with detailed information about pcaps you create from flows on the **Application Flows** tab. Use the **Download** button to download a pcap file you select in the table you can analyze using a packet analyzer. Use the **Stop** button to stop a currently running packet capture. Use the **Delete** button to remove selected pcap(s) from the table. You can also select **Download** and **Delete** by right-clicking on the pcap file in the table.



IP Address

The IP address of the client accessing the packet.

Date/Time

The date and time of the packet capture.

File Size

The file size of the captured packet, in bytes.

File Name

The file name and path of the captured packet.

Source Engine

The ExtremeAnalytics engine that captured the packet.

Related Information

• ExtremeAnalytics tab

ExtremeAnalytics Configuration Overview

Use the **Configuration** tab to view detailed information on the ExtremeAnalytics engines you configure. You can also use the tab to add and enforce your engines, and access engine reports and diagnostics. You must be a member of an authorization group assigned the Extreme Management Center ExtremeAnalytics Read/Write Access capability to view the **Configuration** tab.

Use the left panel in the Configuration view to access various engine administrative options and reports. This Help topic provides information on the following operations available in the left panel:

- Engines
 - Status
 - Configuration
- Virtual Sensors
- Overview
- Fingerprints
- Licenses
- Status
- Configuration

Engines

View engine status information, configure web credentials, and configure advanced options for an individual engine.

Add

Adds a new ExtremeAnalytics engine to Extreme Management Center.

Delete

Delete the selected engine.

Enforce

Enforce the selected engine.

Enforce All

Enforces all of the ExtremeAnalytics engines added to Extreme Management Center.

Poll

Poll the selected engine.

Restart Collector

Restarts the ExtremeAnalytics engine's collector process.

Status

Select an engine and expand the menu to select Status, where you can view engine status including flow collector, application sensor, CPU and memory, flow sources, and diagnostic information.

Configuration

Select an engine and expand the menu to select **Configuration**, where you can configure advanced options for the selected ExtremeAnalyticsengine:

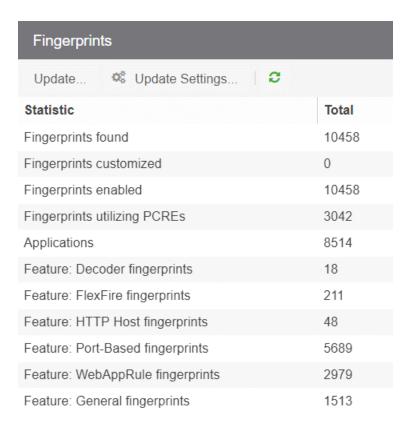
- Set privacy levels.
- · Add and enforce Engines.
- Enable ExtremeControl Integration.
- Add advanced configuration properties.
- Enable sensor modules and sensor module logging.
- Add or remove devices as Application Telemetry flow sources.

Virtual Sensors

Select to display the <u>Virtual Sensors tab</u>, which lists all of the available Virtual Machines on your network as well as those <u>configured as Virtual Sensors</u>.

Fingerprints

View data for the application fingerprints in use.



Use the **Menu** icon (≡) to access the following system fingerprint actions:

Update

Perform a manual one-time update of the fingerprint database.

Update Settings

Schedule fingerprint updates to be performed automatically on a daily or weekly basis.

Licenses

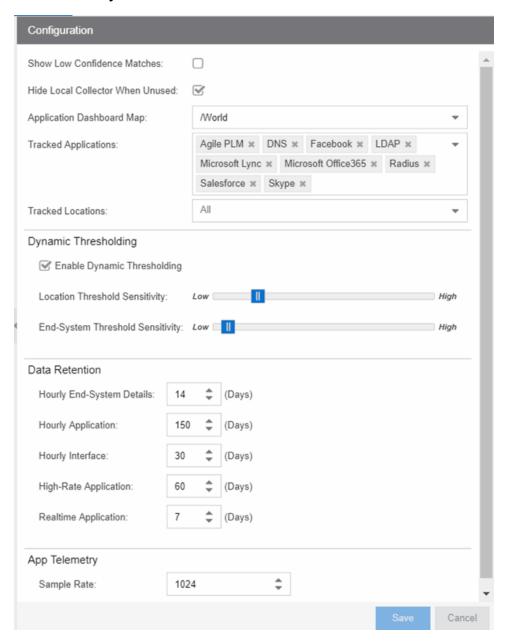
The Licenses window displays data for each license listed, including <u>flow and end-system</u> capacity totals. Select the **Add** button () to increase licensing capacity.

Status

View a collection of ExtremeAnalytics system statistics, including Disk Usage and Approximate Row Counts, as well as Device Families and Profiles.

Configuration

Use the Configuration window to configure the application information displayed in the **ExtremeAnalytics** tab.



Show Low Confidence Matches

Check the box to display flows for which Extreme Management Center has low confidence.

Hide Local Collector When Unused

Check the box to hide local collector information when not in use.

Application Dashboard Map

Select from the drop-down list the map from which to draw application dashboard data.

Tracked Applications

Select the applications to track in the **ExtremeAnalytics** tab.

Tracked Sites

Select the sites to track on the **ExtremeAnalytics** tab. Tracked sites are also indicated on the **Network > Devices > Sites > Endpoint Locations** tab in the right panel.

Dynamic Thresholding

Use the Dynamic Thresholding section of the window to indicate whether to enable dynamic threshold functionality for the Network Service and Tracked Applications dashboards.

When this functionality is enabled, the expected range for application response time is calculated based on past observed response times and a dynamic threshold is assigned. An alarm occurs when any network service or tracked application has a response time measured above its dynamic threshold.

The sliders allow you to adjust the sensitivity of the dynamic threshold by increasing or decreasing the size of the expected response time range. Selecting a lower sensitivity means more time is required for an alarm to occur. Alarms are displayed on the **Alarms & Events > Alarms** tab.

Data Retention

Use this section of the tab to configure the amount of time Extreme Management Center saves flow data.

App Telemetry

This section allows you to configure the default sample rate ExtremeAnalytics uses when configuring ExtremeXOS devices on which Application Telemetry is enabled.

Related Information

- ExtremeAnalytics tab
- Advanced Configuration View

- Add or Enforce Engines in Configuration View
- Add or Remove Devices as Application Telemetry Sources

Virtual Sensors

The Virtual Sensors tab displays all of the available Virtual Machines on your network as well as those <u>configured as Virtual Sensors</u>. The ExtremeAnalytics Virtual Sensor is a virtual machine that monitors application flows across virtual environments.

Virtual Sensors

Use the Virtual Sensors section at the top of the tab to view all of the virtual sensors installed on your network. Click the **Install** button to open the **Install Virtual Sensor on Hypervisor Host** window from which you can add a new virtual sensor. Select a distributed virtual switch from the Virtual Machines section of the window to add it as a virtual sensor. Select a virtual sensor and click **Uninstall** to remove an installed virtual sensor.

Virtual Machines

Use the Virtual Machines section at the bottom of the tab to view all of the virtual machines (VMs) discovered by Extreme Management Center on your network.

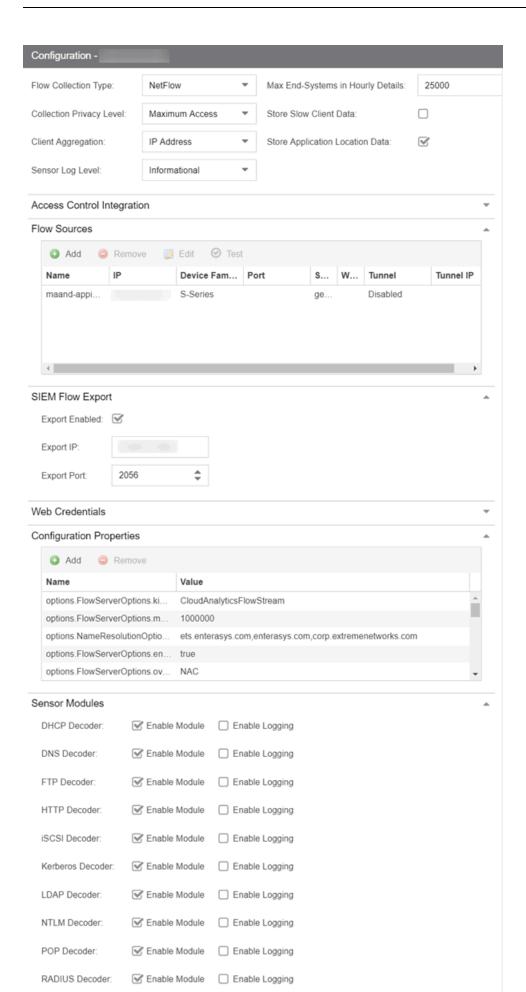
Related Information

- ExtremeAnalytics Virtual Sensor Configuration in Extreme Management Center
- Virtual Sensor Installation Guide

ExtremeAnalytics Engine Advanced Configuration

Use the **Advanced Configuration** panel to configure advanced options for the selected ExtremeAnalytics engine. To access this panel, select the **Configuration** view in the **Analytics** tab in the Extreme Management Center. In the left-panel tree, expand an engine and select **Configuration**.

If you make any changes in this window, be sure to click **Save** and then enforce the engine.



Flow Collection Type

Select from **NetFlow**, **Application Telemetry** or **Both** from the drop-down list to choose the flow data you are using. Selecting **SFlow** changes the Flow Sources table to an Application Telemetry Sources table. These tables allow you to select the devices that collect application telemetry data flow by entering the name, IP and device family for that device.

NOTE: Changing an ExtremeAnalytics engine to Application Telemetry or Both does not update the configuration on the engine. Run the dnetconfig process on the engine to change the configuration.

Collection Privacy Levels

Collection privacy level settings restrict the amount of identifying information collected by the ExtremeAnalytics engine and displayed in the Application Information column of the Application Flows report. (Access this report from the **Analytics** tab. In the Application Flows report, hover over the Application Information column to view the collected information.) This information is also displayed in the Flow Summary window.

Increasing the privacy level protects the end user's identifying information from being viewed by IT staff with access to the Application Flows report. The default privacy level allows maximum access to the information. Increasing the privacy level allows you to restrict the information that is collected and displayed.

There are three privacy levels. For all three levels, passwords are **not** collected or displayed.

- Maximum Access The ExtremeAnalytics engine collects both identifying information and sensitive information. The information displays in the Application Information column.
- Medium Privacy The ExtremeAnalytics engine collects identifying information, but not sensitive information. Identifying information displays in the Application Information column.
- Maximum Privacy The ExtremeAnalytics engine does not collect identifying information or sensitive information. Information does not display in the Application Information column.

Identifying information is data that identifies the end user, such as a username. The ExtremeAnalytics engine collects identifying information when the privacy level is set to Maximum Access or Medium Privacy.

Sensitive information is data an end user may not want to share, such as the caller ID or contact information from an end user's SIP voice call. The ExtremeAnalytics engine collects sensitive information when the privacy level is set to Maximum Access.

Client Aggregation

Use this field to determine how client information is aggregated by the ExtremeAnalytics engine, either by IP Address or MAC Address.

Slow Client Data

Select **Enabled** in the drop-down list to collect additional information about clients with poor response times by the ExtremeAnalytics engine.

Max End-Systems in Hourly Details

Enter the maximum number of client end-systems stored in the Extreme Management Center database for the ExtremeAnalytics engine. This ensures your client limit is not collected from one engine. Once the value set in this field is met, additional end-system data is not collected from the engine.

Sensor Log Levels

The ExtremeAnalytics sensor runs on the ExtremeAnalytics engine and inspects network traffic to identify applications and other information. The sensor log file records diagnostic information about sensor operations, which is useful for troubleshooting engine issues.

In the **Configuration** view, you can enable different levels of logging for the selected engine. Each logging level is inclusive of the levels above it. The five levels are:

- Informational
- Debug
- Verbose Debug

- Trace
- All

The sensor log level should be set to **Informational** unless you are troubleshooting an engine issue. When troubleshooting an issue, Extreme Networks Support may ask you to change the logging level to provide additional information.

To view the log file directly, log into the engine and navigate to the file /opt/appid/logs/appid.log.

You can also use the engine administration web page to view the sensor log. Access the web page using the following URL: https://<EngineIP or hostname>:8443/Admin. The default user name and password is "admin/Extreme@pp." Once you have accessed the web page, navigate to the Log Files/Sensor Log page.

Store Application Site Data

Select this checkbox to allow the high-rate collector to store client count, flow count, bytes, received bytes, sent bytes, application response time, and network response time by application for a site. If this checkbox is selected, the **Application/Site** option is available as a **Target** in the Applications Browser when **Data Table** is **Application Data - High-rate** to display the information.

ExtremeControl Integration

If your network configuration includes ExtremeControl, ExtremeControl data can be integrated with flow data to provide additional information. ExtremeControl integration is only useful if you are collecting flows for end-systems managed by ExtremeControl. For additional information, see Enabling ExtremeControl Integration.

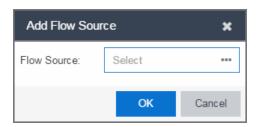
- To enable ExtremeControl Integration for the engine, select the Enable ExtremeControl Integration checkbox.
- If your ExtremeControl engines are using Communication Channels, select the ExtremeControl Communication Channel option and enter the channel name.
 The ExtremeAnalytics engine is only able to access end-systems in its channel.

Flow Sources/Application Telemetry Sources

This section displays the devices set up as flow sources in ExtremeAnalytics.

To add a device as a flow source:

Click the Add button.
 The Add Flow Source window opens.



2. Click the **Ellipsis** icon and select a device from the Select Flow Source Device window.

NOTES: When selecting a controller, only Wireless Controllers that support ExtremeAnalytics and have available L2 ports are listed.

Selecting a Wireless Controller set up as part of a controller pair automatically selects the paired Controller.

To use an ExtremeSwitching X440-G2 switch as an Application Telemetry source for ExtremeAnalytics, install firmware version 22.4.1.4-patch2-5 or higher..

- 3. Select an available L2 port for mirroring in the Mirror Port drop-down list.
- 4. Select a mirror port for the Paired Controller, if necessary.
- 5. Select the appropriate WLANs, if necessary.
- 6. Click OK.
- 7. Verify the L2 ports selected for mirroring are monitored by ExtremeAnalytics.

The configuration is complete.

NOTE: Packets routed through GRE tunnels on an ExtremeAnalytics virtual engine are incorrectly reported as dropped in the ifconfig output on the GRE interface. This scenario was tested and while these packets are reported as dropped, they are being inspected by the ExtremeAnalytics engine.

To remove a device as a flow source, select the device in the Flow Sources section of the window and click the **Remove** button.

Web Credentials

Enter a new **Username** and **Password** for web service requests between the Extreme Management Center server and the ExtremeAnalytics engine. Click the **Show Password**

checkbox to display the Password field unencrypted.

NOTE: By default, the **Username** and **Password** are **admin** and **Extreme@pp**, respectively.

Configuration Properties

Use this section to add properties that provide a solution for a specific problem or task. These properties are supplied directly by Extreme Networks Support. Contact Extreme Networks Technical Support for guidance on using this section.

Sensor Modules

The ExtremeAnalytics sensor uses sensor modules to analyze different types of network traffic. For example, the HTTP decoder decodes HTTP traffic to acquire data needed to match fingerprints against that traffic.

In most cases, it is best to leave the decoders and detectors enabled. For better sensor performance, you can disable decoders for traffic rarely seen on the network; however, doing so prevents some fingerprints from triggering.

You can enable logging for any of the decoders and detectors for debugging purposes. As logging can impact disk space and performance, turn it on only for troubleshooting purposes. Do not enable logging during normal operation.

Auditing

Use this section to enable auditing of users connected to the ExtremeAnalytics engine CLI via SSH.

Enable Auditing

Selecting the **Enable Auditing** option enables the **Auditing Rules** field, where you can configure Extreme Management Center to store all commands entered by a user connected to the ExtremeAnalytics engine CLI via SSH in the engine's local syslog file.

Auditing Rules

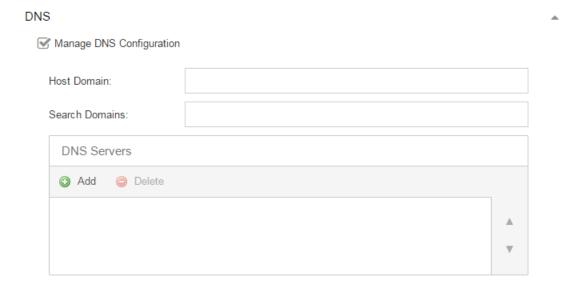
Remove the # symbol from the beginning of a command line to enable the command and store user commands entered using the ExtremeAnalytics engine CLI.

Network Settings

Use the Network Settings section of the window to configure the network settings on an ExtremeAnalytics engine. Selecting a checkbox opens a new section from which you can configure the options for the setting. Click the **Save** button and the bottom of the panel to save your changes.

DNS

Select the **Manage DNS Configuration** checkbox to open the DNS Servers area. This allows you to enter a search domain or add or remove search domains and DNS server IP addresses.



Search Domains

A list of search domains used by the ExtremeAnalytics 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, ExtremeAnalytics appends the search domains in an attempt to resolve the name: server1.domain1 server1.domain2, and so on.

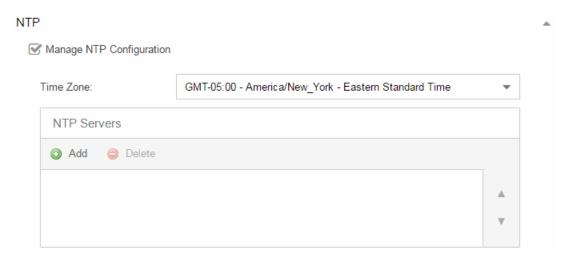
DNS Servers

A list of DNS servers the ExtremeAnalytics engine sends DNS lookups to for name resolution. The list is used by both hostname resolution and by the DNS proxy. Click the **Add** button to open a blank box in which you can enter an IP address. Select an IP address in the table and click the **Delete** button to remove an

IP address. You can enter multiple servers for redundancy. Use the **Up** and **Down** arrows to list the servers in the order they should be used.

NTP

Select the **Manage NTP Configuration** checkbox to open the NTP (Network Time Protocol) Servers area. NTP configuration is important for protocols such as SNMPv3 and RFC3576 which incorporate playback protection. In addition, having accurate time configured on the ExtremeAnalytics engine is essential for event logging and troubleshooting.



Time Zone

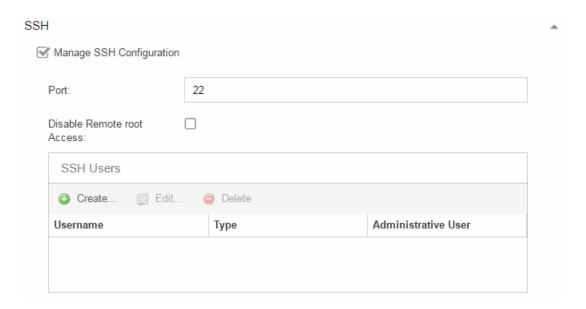
Select the appropriate **Time Zone** from the drop-down list to allow ExtremeAnalytics to manage date/time settings.

NTP Servers

A list of NTP servers. You can enter multiple servers for redundancy. Click the **Add** button to open a blank box in which you can enter an IP address. Select an IP address in the table and click the **Delete** button to remove an IP address. Use the **Up** and **Down** arrows to list the servers in the order they should be used.

SSH

Select the **Manage SSH Configuration** checkbox to open the SSH Users area. SSH configuration provides additional security features for the ExtremeAnalytics engine.



Port

The port field allows you to configure a custom port 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. Make sure that you don't disable root access unless you have configured RADIUS authentication or this disables remote access to the ExtremeAnalytics engine.

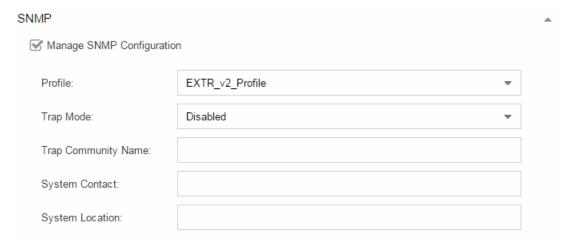
SSH Users

Use the toolbar buttons to create a list of users allowed to log in to the ExtremeAnalytics engine using SSH. Click the **Add** button to open a blank box in which you can enter an IP address. Select an IP address in the table and click the **Delete** button to remove an IP address. 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.

SNMP

Use the SNMP configuration section to deploy SNMP credentials for the ExtremeAnalytics engine. The credentials can include different read/write credentials, for example, use "public" as the read credential and "private" as the write credential. In addition, basic host traps can be enabled from the ExtremeAnalytics engine. Select the Manage SNMP Configuration checkbox and provide the following SSH information.



Profile

Use the drop-down list to select a device access profile to use for the ExtremeAnalytics engine.

Trap Mode

Use the drop-down list to set the trap mode.

Trap Community Name

Enter the trap community name.

Interfaces

Use the Interfaces section to configure the ExtremeAnalytics engine interface Extreme Management Center uses to monitor applications accessed on your network.



Monitor Mode

Use the drop-down list to configure the ExtremeAnalytics engine interface used to monitor network traffic.

Related Information

For information on related ExtremeAnalytics topics:

- Analytics Tab
- Enabling ExtremeControl Integration

ExtremeAnalytics Reports

The **Analytics** tab lets you view and customize ExtremeAnalytics reports and application flow data, as well as manage and configure your ExtremeAnalytics engines.

NOTE: ExtremeAnalytics reports and application flow data is not available unless an ExtremeAnalytics engine is configured and you are a member of an authorization group assigned the Extreme Management Center ExtremeAnalytics Read Access or Read/Write Access capability.

Reports

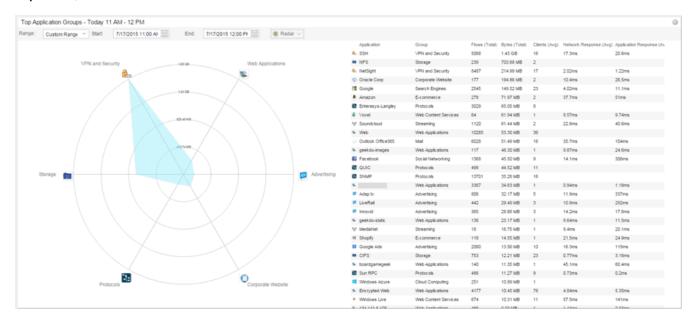
In the **Reports** tab, you can access a selection of reports that provide detailed information on application usage on your network, as well as network activity statistics based on application, user name, client, and site. For many of the reports, you can click on an item in the report to view details or right-click an item to select from other focused reports.

If you have multiple ExtremeAnalytics engines, use the **Engine** drop-down list to select an engine to use as the source for the report data. Then use the Report drop-down list to the right to access the different reports:

- Analytics Events
- Bandwidth for a Client Over Time
- Interface Top Applications Tree Map
- Sites Using the Most Bandwidth
- Most Popular Applications
- Most Used Applications for a Client
- Most Used Applications for a User Name
- Network Activity by Site
- Network Activity by Client
- Network Activity by Application
- Slowest Application by Site
- Top Applications Group Radar
- Top Applications Radar

- Top Applications Tree Map
- Top Clients by Interface
- Top Interfaces by Application
- Top N Applications
- Top N Clients
- Top N Servers

In most of the reports, use the **Gear** button (on the right side of the view) to display a **Start Time** option that allows you to change the length of the reporting period displayed. Depending on the report, you can also change the type and/or format of the data reported, and the number of results to return.



Some of the reports are based on a specific object (target), such as a user name, client, application, or site. In those reports, enter the required information and then click the **Submit** button to generate the report. You can enter a partial value in the text field or use the SQL wildcard "%" (as a substitute for multiple characters) or "_" (as a substitute for a single character) to generate a report with multiple matches.

NOTE: Values entered in the text fields that contain multiple, non-alphanumeric characters may cause issues with the returned results. If this happens, use alternate values.

Related Information

ExtremeAnalytics tab

ExtremeAnalytics Report Descriptions

The **Analytics** tab lets you view and customize ExtremeAnalytics reports and application flow data, as well as manage and configure your ExtremeAnalytics engines.

NOTE: ExtremeAnalytics reports and application flow data is not available unless an ExtremeAnalytics engine is configured and you are a member of an authorization group assigned the Extreme Management Center ExtremeAnalytics Read Access or Read/Write Access capability.

Report Descriptions

In the <u>Reports</u> tab, you can access a selection of reports that provide detailed information on application usage on your network, as well as network activity statistics based on application, user name, client, and site. For many of the reports, you can click on an item in the report to view details or right-click an item to select from other focused reports.

If you have multiple ExtremeAnalytics engines, use the **Engine** drop-down list to select an engine to use as the source for the report data. Then use the Report drop-down list to the right to access the different reports:

- Analytics Events
- Bandwidth for a Client Over Time
- Interface Top Applications
- Sites Using the Most Bandwidth
- Most Popular Applications
- Most Used Applications for a Client
- Most Used Applications for a User Name
- Network Activity by Site
- Network Activity by Client
- Network Activity by Application
- Slowest Application by Site
- Top Applications Group Radar
- Top Applications Radar

- Top Applications Tree Map
- Top Applications for Interface
- Top Applications for Server
- Top Clients by Interface
- Top Interfaces by Application
- Top N Applications
- Top N Clients
- Top N Servers

Analytics Events

This report displays the event log filtered to show only the events related to ExtremeAnalytics.

Bandwidth for a Client Over Time

This report displays the bandwidth used by the specified client, provided as a line chart showing average bytes used over time. Enter a client's IP address or hostname and then click the **Submit** button to generate the report.

Interface Top Applications Treemap

This report displays the top applications for the top switch interfaces (devices) with application telemetry enabled.

NOTE: You need to first enable the application telemetry feature on ExtremeXOS switches from the **Analytics > Configuration** tab.

Sites Using the Most Bandwidth

This report displays the network sites with the highest bandwidth, provided as a bubble map.

Most Popular Applications

This report displays the applications used the most, based on the number of unique client IP addresses associated with them. Click on an application name to open a report showing the top clients for that application. Click on a client from the report to display an End-System Applications Summary for that client

Most Used Applications for a Client

This report displays the applications used the most by the specified client, based on bandwidth. Enter a client's IP address or hostname and then click the **Submit** button to generate the report.

Most Used Applications for a User Name

This report displays the applications used the most by the specified user, based on bandwidth. Enter a client's user name and then click the **Submit** button to generate the report.

Network Activity by Site

This report displays network traffic statistics and application and network response time for each site.

Network Activity by Client

This report displays network traffic statistics for the specified client. Enter a client's IP address or hostname and then click the **Submit** button to generate the report.

Network Activity by Application

This report displays network traffic statistics for the specified application. Enter an application name and then click the **Submit** button to generate the report.

Slowest Applications by Site

This report displays the applications with the highest application response times for the specified site. Select a site from the drop-down list to match or select All and then click the **Submit** button to generate the report. If a site has been added to a map, you also see a selection for that map. If you select custom, you can enter a partial site name or use the SQL wildcard characters to match one or more sites.

Top Applications Group Radar

In the **Top Applications Group Radar** report, the info bar provides an overview of application group usage in a radar format. Use the **Start** calendar to select the start date and time and the format to display.

Top Applications Radar

In the **Top Applications Radar** report, the info bar provides an overview of application usage in a radar format. Use the **Start** calendar to select the start date and time and the format to display.

Top Applications TreeMap

This report displays hierarchical data on application bandwidth usage, grouped by application group and displayed in sets of colored nested rectangles. This design allows you to easily see patterns of bandwidth usage that might otherwise be difficult to spot. Click on an application group to zoom in and view data for that group. Hover over an application cell to view bandwidth for a particular application. Right-click on an application cell to access additional reports for that application.

Use the **Gear** button to change the start date and time to display. Set the scale to Linear to view the data scaled proportionately; set the scale to Log to make smaller rectangles of data more visible. Use the combo box to change how the data is displayed: by bandwidth, client count, or flow count.

Top Applications for Interface

This report displays the top applications for a specified interface (device) with application telemetry enabled (wildcards allowed).

NOTE:

You need to first enable the application telemetry feature on ExtremeXOS switches from the **Analytics > Configuration** tab.

Top Applications for Server

This report displays the top applications for a device configured as a server with application telemetry enabled (wildcards allowed).

NOTE:

You need to first enable the application telemetry feature on ExtremeXOS switches from the **Analytics > Configuration** tab.

Top Clients by Interface

This report displays the top clients for a specified switch interface (device) with application telemetry enabled (wildcards allowed).

NOTE:

You need to first enable the application telemetry feature on ExtremeXOS switches from the **Analytics > Configuration** tab.

Top Interfaces by Application

This report displays the top interfaces (device) for a specified application with application telemetry enabled (wildcards allowed).

NOTE:

You need to first enable the application telemetry feature on ExtremeXOS switches from the **Analytics > Configuration** tab.

Top N Applications

This report displays application information, provided as a bar graph. Use the fields in the menu to configure the information displayed in the report:

- Top N Select the number of clients displayed in the chart.
- Start Select the start date and time.
- # Hours Select the amount of time for which data is displayed from the date and time selected in Start.
- Statistic Select the statistic by which the top clients are listed.
 - Bandwidth
 - Flows
 - Client Count

Top N Clients

This report displays client information, provided as a bar graph. Use the fields in the menu to configure the information displayed in the report:

- Top N Select the number of clients displayed in the chart.
- Start Select the start date and time.
- # Hours Select the amount of time for which data is displayed from the date and time selected in Start.
- Statistic Select the statistic by which the top clients are listed.
 - Bandwidth
 - Flows
 - Number of Applications

Top N Servers

This report displays server information, provided as a bar graph. Use the fields in the menu to configure the information displayed in the report:

- Top N Select the number of clients displayed in the chart.
- Start Select the start date and time.
- # Hours Select the amount of time for which data is displayed from the date and time selected in **Start**.
- Statistic Select the statistic by which the top clients are listed.
 - Bandwidth
 - Flows

Related Information

ExtremeAnalytics tab

Add and Modify Fingerprints

ExtremeAnalytics uses fingerprints to identify to which application a network traffic flow belongs. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. Extreme Management Center provides thousands of system fingerprints with the ExtremeAnalytics feature. In addition, you can modify these fingerprints and create new custom fingerprints.

For additional information, see Getting Started with ExtremeAnalytics.

This Help topic provides the following information:

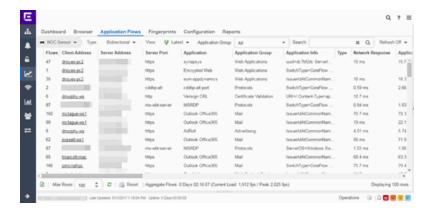
- Adding a Fingerprint
- Modifying a Fingerprint
- Enabling or Disabling a Fingerprint
- Deleting a Custom Fingerprint
- Updating Fingerprints

In order to add and modify fingerprints, you must be a member of an authorization group assigned the Extreme Management Center ExtremeAnalytics Read/Write Access capability.

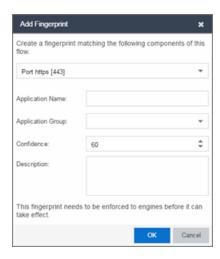
Adding a Fingerprint

Use the following steps to add a new custom fingerprint based on an existing flow in the Applications Flows view.

1. Select the **Analytics** tab and then select the **Application Flows** view.



- 2. Select the flow in the table that you want to base your new custom fingerprint on.
- Right-click on the flow and select the Fingerprints > Add Fingerprint option. The Add Fingerprint window opens.



- 4. Use the drop-down list to select the flow components on which to base the fingerprint. The options vary depending on the fingerprint you initially selected.
 - Port <port number> Creates a fingerprint that identifies traffic either coming from or going to the specified port.
 - Address <IP address> on port <port number> Creates a fingerprint that
 identifies traffic either coming from or going to this IP address on the
 specified port.
 - Address <IP address> with mask on port <port number> Creates a fingerprint that identifies traffic either coming from or going to the specified subnet on the specified port. For example, an IP address of 192.168.0.0 with a mask of 16 would result in all traffic either coming from or going to the 192.168 subnet on the specified port to be identified by the fingerprint.
 - Host <host name> Creates a fingerprint that identifies a specific hostname
 in the URI of web traffic.
 - HTTP Header Creates a fingerprint that identifies traffic containing specified HTTP header information, if HTTP header information is included in the flow's metadata.

Note that there may be two port number or IP address options listed: one for the flow's source port/IP address and one for the flow's destination port/IP address.

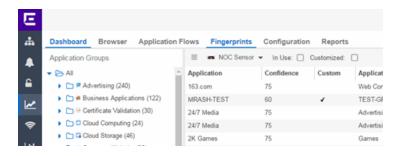
- 5. If you selected an IP address with mask option, you need to specify a subnet of IP addresses. Enter the IP CIDR mask, which is a mask on the flow IP, with 0-32 for IPv4 and 0-128 for IPv6.
- 6. Enter the name of the application for which the fingerprint is defined.
- 7. Use the drop-down list to select the application group to which the application belongs. If none of the existing groups are appropriate, you can enter a new group name and the new group is automatically created.
- 8. Select the fingerprint's confidence level. The confidence level defines the reliability of this fingerprint. Higher confidence fingerprints override lower confidence fingerprints, if multiple fingerprints match a flow. Values are 1-100, with 100 being absolutely reliable.
- 9. Enter a description of the fingerprint, if desired.
- Click Save. The new fingerprint is created on the Extreme Management Center server.
- 11. Enforce to push the new fingerprint to your engines.

TIP: You can also create a custom fingerprint from the **Fingerprints** tab. Click the **Menu** icon and select **Create Fingerprint**. The Add Fingerprint window opens where you can select all the flow components you want for the fingerprint. The new fingerprint is not based on an existing fingerprint and you need to enter values for all required fields such as **IP** or **Hostname**, **Application Name**, and **Application Group**. The new fingerprint must be enforced to engines before it can take effect.

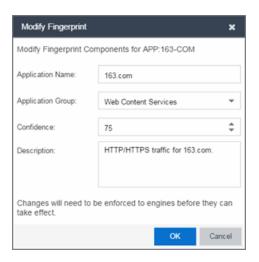
Modifying a Fingerprint

Modify a fingerprint's application name, application group, confidence level, and description from the **Fingerprints** tab.

Select the Analytics > Fingerprints tab.



2. Right-click on the desired fingerprint and select **Modify Fingerprint** from the menu. The Modify Fingerprint window opens.



- 3. Make the desired changes:
 - Application Name The name of the application that the fingerprint detects.
 If you change the application name, you are prompted to select whether to
 change the application name for only the currently selected fingerprint or for
 all fingerprints that have that same application name.

NOTE: If you change both the Application Name and Application Group: If the new Application Name matches an existing name, the application group changes to the new group for all fingerprints with that new name, regardless of whether you choose to change the name for only the selected fingerprint or for all fingerprints with that name.

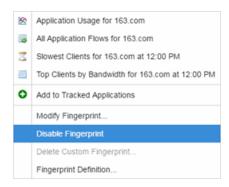
- Application Group Organizes fingerprints into different types of applications such as Web applications or Business applications. You can sort the Application Flows view by application group, making it easier to view the data. If you change the application group for a fingerprint, it changes the group for all fingerprints with that same application name. If none of the existing groups are appropriate, you can create a new group by entering a new group name.
- Confidence Defines the reliability of this fingerprint. Higher confidence fingerprints override lower confidence fingerprints, if multiple fingerprints match a flow. Values are 1-100, with 100 being absolutely reliable. The confidence level only applies to the currently selected fingerprint.

- **Description** A description of the fingerprint. The description only applies to the currently selected fingerprint.
- 4. Click OK.
- 5. Enforce to push the change to your engines.

Enabling or Disabling a Fingerprint

Enable or disable a fingerprint from the **Fingerprints** tab. When a fingerprint is enabled, it is used to identify applications. When it is disabled, it is ignored.

- 1. Select the **Analytics** > **Fingerprints** tab.
- 2. Right-click on the desired fingerprint in the Fingerprints table and select either **Enable Fingerprint** or **Disable Fingerprint**.



3. Enforce to push the change to your engines.

NOTE: If you disable a system fingerprint, it becomes a custom fingerprint. If you then enable the fingerprint, it remains a custom fingerprint. Deleting the custom fingerprint reloads the original system fingerprint.

Deleting a Custom Fingerprint

Delete a custom fingerprint from the **Fingerprints** tab. A custom fingerprint is either a new user-defined fingerprint, a modification of a system fingerprint, or a disabled fingerprint. (Custom fingerprints can be identified by a ✓ in the Custom column.)

When you delete a custom fingerprint, it is removed entirely. If you delete a custom fingerprint overriding a system fingerprint, the original system fingerprint is reloaded. System fingerprints that have not been modified cannot be deleted, however, they can be disabled.

Use these steps to delete a custom fingerprint:

- Select the **Analytics** tab in Extreme Management Center and then select the Fingerprints view
- 2. Right-click on the desired fingerprint in the Fingerprints table and select **Delete Custom Fingerprint**. The Delete Fingerprint window opens.



- 3. You can delete only the selected fingerprint or select the option to delete all custom fingerprints that match the application name of the selected fingerprint.
- 4. Click **OK**. If a custom fingerprint overrides a system fingerprint, then deleting the custom fingerprint reloads the original system fingerprint.
- Enforce to push the change to your engines.

Updating Fingerprints

New and updated fingerprints are provided via a fingerprint update website. Perform a one-time manual update of the fingerprint database or configure a scheduled update to be performed automatically from the **Configuration** tab. Custom fingerprints are not overwritten when an update is performed.

When a fingerprint update is performed, the fingerprint update server is checked for newer fingerprints than what is available on the Extreme Management Center server. If there are newer fingerprints, they are downloaded, and the fingerprint definitions are updated with any new fingerprint definition files. You need to enforce your engines following an update to push the updated fingerprints to the engines.

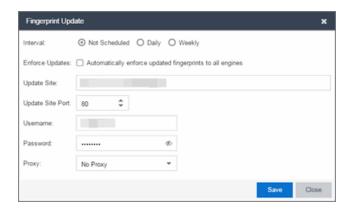
Perform a Fingerprint Update

Perform a manual one-time update of the fingerprint database. To access the update website, you need to create an Extranet account at ExtremeNetworks.com and define a username and password for the account. You need the username and password in order to perform updates.

1. Select the **Analytics** tab in Extreme Management Center and then select the **Configurations** view.

- 2. In the left-panel tree, expand the System folder and select Fingerprints.
- Click the Menu icon (≡) and select Update Fingerprints. If you have already
 configured your Fingerprint Update settings, the update is performed immediately.

If you have not configured your settings, the Fingerprint Update window opens.



- a. Leave the Interval selection as Not Scheduled.
- Select the Enforce Updates checkbox to automatically update fingerprints on all engines. Not selecting this checkbox requires you to update each engine manually.
- c. The **Update Site** field displays the default path to the official fingerprint update site. Typically, this field does change unless for security reasons the system does not have access to the internet and an internal update site must be used.
- d. The **Update Site Port** is the port on the update site to which the update connects. The port cannot be changed unless you are using a custom update site.
- e. Enter the credentials used to access the fingerprint update website. These are the username and password credentials you defined when you created an Extranet account at ExtremeNetworks.com.
- f. If your network is protected by a firewall, you need to configure proxy server settings to use when accessing the website. In the Proxy field, select Use Proxy or Use Proxy with Credentials and enter your proxy server address and port ID. (Consult your network administrator for this information.) If your proxy server requires authentication, enter the proxy username and password credentials. The credentials you add here must match the credentials configured on the proxy server.
- g. Click **Save**. The Fingerprint Update is performed immediately.

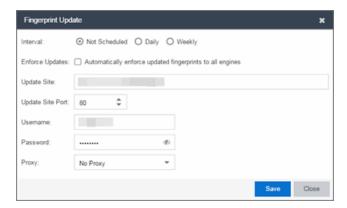
4. If you did not select the **Enforce Updates** checkbox, enforce to push the changes to your engines when the update is complete.

Schedule Fingerprint Updates

You can schedule fingerprint updates performed automatically on a daily or weekly basis.

To access the update website, you need to create an Extranet account at ExtremeNetworks.com and define a username and password for the account. You need the username and password in order to schedule updates.

- 1. Select the **Analytics** tab in Extreme Management Center and then select the Configuration view.
- 2. In the left-panel tree, expand the System folder and select Fingerprints.
- 3. Click on the **Menu** icon (≡) and select Fingerprint Update Settings. The Fingerprint Update window opens.



- 4. Select the update interval which defines how frequently the update is performed: **Daily** or **Weekly**.
- 5. If you have selected **Weekly**, select the day of the week you would like the update performed.
- 6. Enter the scheduled time you would like the update performed.
- Select the Enforce Updates checkbox to automatically update fingerprints on all engines. Not selecting this checkbox requires you to update each engine manually.
- 8. The **Update Site** field displays the default path to the official fingerprint update site. Typically, this field does not change unless for security reasons the system does not have access to the internet and an internal update site must be used.

- 9. The **Update Site Port** is the port on the update site to which the update connects. The port cannot be changed unless you are using a custom update site.
- Enter the credentials used to access the fingerprint update website. These are the
 username and password credentials you defined when you created an Extranet
 account at ExtremeNetworks.com.
- 11. If your network is protected by a firewall, configure proxy server settings to use when accessing the website. In the Proxy field, select Use Proxy or Use Proxy with Credentials and enter your proxy server address and port ID. (Consult your network administrator for this information.) If your proxy server requires authentication, enter the proxy username and password credentials. The credentials you add here must match the credentials configured on the proxy server.
- 12. Click Save.
- 13. If you did not select the **Enforce Updates** checkbox, enforce to push the changes to your engines when the update is complete.

Related Information

For information on related ExtremeAnalytics topics:

- Analytics
- Custom Fingerprint Examples

Add Fingerprints

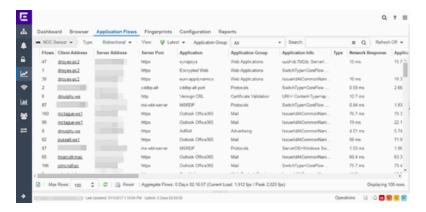
ExtremeAnalytics uses fingerprints to identify to which application a network traffic flow belongs. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. Extreme Management Center provides thousands of system fingerprints with the ExtremeAnalytics feature. In addition, you can modify these fingerprints and create new custom fingerprints.

In order to add and modify fingerprints, you must be a member of an authorization group assigned the Extreme Management Center ExtremeAnalytics Read/Write Access capability.

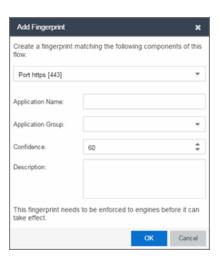
Add a Fingerprint

Use the following steps to add a new custom fingerprint based on an existing flow in the Applications Flows view. You can also create a new fingerprint based on an application or application group, or on a destination address.

1. Click Analytics > Application Flows view.



- 2. Select the flow in the table that you want to base your new custom fingerprint on.
- 3. Right-click on the flow and select the **Fingerprints > Add Fingerprint** option. The Add Fingerprint window opens.



- 4. Use the drop-down list to select the flow components on which to base the fingerprint. The options vary depending on the fingerprint you initially selected.
 - Port <port number> Creates a fingerprint that identifies traffic either coming from or going to the specified port.
 - Address <IP address> on port <port number> Creates a fingerprint that identifies traffic either coming from or going to this IP address on the specified port.
 - Address <IP address> with mask on port <port number> Creates a fingerprint that identifies traffic either coming from or going to the specified subnet on the specified port. For example, an IP address of 192.168.0.0 with a mask of 16 would result in all traffic either coming from or going to the 192.168 subnet on the specified port to be identified by the fingerprint.
 - Host <host name> Creates a fingerprint that identifies a specific hostname
 in the URI of web traffic.
 - HTTP Header Creates a fingerprint that identifies traffic containing specified HTTP header information, if HTTP header information is included in the flow's metadata.

Note that there may be two port number or IP address options listed: one for the flow's source port/IP address and one for the flow's destination port/IP address.

- 5. If you selected an IP address with mask option, you need to specify a subnet of IP addresses. Enter the IP CIDR mask, which is a mask on the flow IP, with 0-32 for IPv4 and 0-128 for IPv6.
- 6. Enter the name of the application for which the fingerprint is defined.

- 7. Use the drop-down list to select the application group to which the application belongs. If none of the existing groups are appropriate, you can enter a new group name and the new group is automatically created.
- 8. Select the fingerprint's confidence level. The confidence level defines the reliability of this fingerprint. Higher confidence fingerprints override lower confidence fingerprints, if multiple fingerprints match a flow. Values are 1-100, with 100 being absolutely reliable.
- 9. Enter a description of the fingerprint, if desired.
- Click Save. The new fingerprint is created on the Extreme Management Center server.
- 11. Enforce to push the new fingerprint to your engines.
- TIP: You can also create a custom fingerprint from the **Fingerprints** tab. Click the **Menu** icon and select **Create Fingerprint**. The Add Fingerprint window opens where you can select all the flow components you want for the fingerprint. The new fingerprint is not based on an existing fingerprint and you need to enter values for all required fields such as **IP** or **Hostname**, **Application Name**, and **Application Group**. The new fingerprint must be enforced to engines before it can take effect.

Related Information

ExtremeAnalytics tab

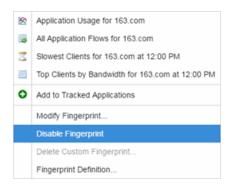
Enable or Disable Fingerprints

ExtremeAnalytics uses fingerprints to identify to which application a network traffic flow belongs. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. Extreme Management Center provides thousands of system fingerprints with the ExtremeAnalytics feature. In addition, you can modify these fingerprints and create new custom fingerprints.

Enabling or Disabling a Fingerprint

Enable or disable a fingerprint from the **Fingerprints** tab. When a fingerprint is enabled, it is used to identify applications. When it is disabled, it is ignored.

- Select the Analytics > Fingerprints tab.
- 2. Right-click on the desired fingerprint in the Fingerprints table and select either **Enable Fingerprint** or **Disable Fingerprint**.



3. Enforce to push the change to your engines.

NOTE: If you disable a system fingerprint, it becomes a custom fingerprint. If you then enable the fingerprint, it remains a custom fingerprint. Deleting the custom fingerprint reloads the original system fingerprint.

Related Information

ExtremeAnalytics tab

Modify Fingerprints

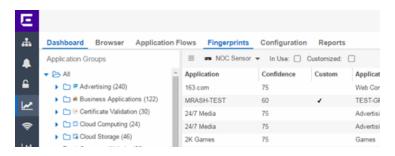
ExtremeAnalytics uses fingerprints to identify to which application a network traffic flow belongs. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. Extreme Management Center provides thousands of system fingerprints with the ExtremeAnalytics feature. In addition, you can modify these fingerprints and create new custom fingerprints.

In order to add and modify fingerprints, you must be a member of an authorization group assigned the Extreme Management Center ExtremeAnalytics Read/Write Access capability.

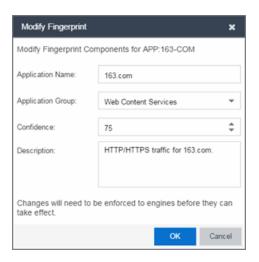
Modifying a Fingerprint

Modify a fingerprint's application name, application group, confidence level, and description from the **Fingerprints** tab.

Select the Analytics > Fingerprints tab.



2. Right-click on the desired fingerprint and select **Modify Fingerprint** from the menu. The Modify Fingerprint window opens.



- 3. Make the desired changes:
 - Application Name The name of the application that the fingerprint detects.
 If you change the application name, you are prompted to select whether to
 change the application name for only the currently selected fingerprint or for
 all fingerprints that have that same application name.

NOTE: If you change both the Application Name and Application Group: If the new Application Name matches an existing name, the application group changes to the new group for all fingerprints with that new name, regardless of whether you choose to change the name for only the selected fingerprint or for all fingerprints with that name.

- Application Group Organizes fingerprints into different types of applications such as Web applications or Business applications. You can sort the Application Flows view by application group, making it easier to view the data. If you change the application group for a fingerprint, it changes the group for all fingerprints with that same application name. If none of the existing groups are appropriate, you can create a new group by entering a new group name.
- Confidence Defines the reliability of this fingerprint. Higher confidence
 fingerprints override lower confidence fingerprints, if multiple fingerprints
 match a flow. Values are 1-100, with 100 being absolutely reliable. The
 confidence level only applies to the currently selected fingerprint.
- Description A description of the fingerprint. The description only applies to the currently selected fingerprint.
- 4. Click OK.
- 5. Enforce to push the change to your engines.

Related Information

• ExtremeAnalytics tab

Update Fingerprints

ExtremeAnalytics uses fingerprints to identify to which application a network traffic flow belongs. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. Extreme Management Center provides thousands of system fingerprints with the ExtremeAnalytics feature. In addition, you can modify these fingerprints and create new custom fingerprints.

In order to add and modify fingerprints, you must be a member of an authorization group assigned the Extreme Management Center ExtremeAnalytics Read/Write Access capability.

Updating Fingerprints

New and updated fingerprints are provided via a fingerprint update website. Perform a one-time manual update of the fingerprint database or configure a scheduled update to be performed automatically from the **Configuration** tab. Custom fingerprints are not overwritten when an update is performed.

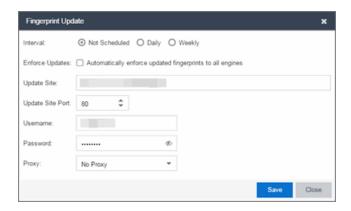
When a fingerprint update is performed, the fingerprint update server is checked for newer fingerprints than what is available on the Extreme Management Center server. If there are newer fingerprints, they are downloaded, and the fingerprint definitions are updated with any new fingerprint definition files. You need to enforce your engines following an update to push the updated fingerprints to the engines.

Perform a Fingerprint Update

Perform a manual one-time update of the fingerprint database. To access the update website, you need to create an Extranet account at ExtremeNetworks.com and define a username and password for the account. You need the username and password in order to perform updates.

- 1. Select the **Analytics** tab in Extreme Management Center and then select the **Configurations** view.
- 2. In the left-panel tree, expand the System folder and select **Fingerprints**.
- Click the Menu icon (≡) and select Update Fingerprints. If you have already configured your Fingerprint Update settings, the update is performed immediately.

If you have not configured your settings, the Fingerprint Update window opens.



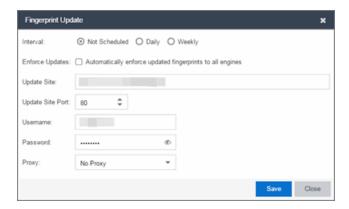
- a. Leave the Interval selection as Not Scheduled.
- Select the Enforce Updates checkbox to automatically update fingerprints on all engines. Not selecting this checkbox requires you to update each engine manually.
- c. The **Update Site** field displays the default path to the official fingerprint update site. Typically, this field does change unless for security reasons the system does not have access to the internet and an internal update site must be used.
- d. The Update Site Port is the port on the update site to which the update connects. The port cannot be changed unless you are using a custom update site.
- e. Enter the credentials used to access the fingerprint update website. These are the username and password credentials you defined when you created an Extranet account at ExtremeNetworks.com.
- f. If your network is protected by a firewall, you need to configure proxy server settings to use when accessing the website. In the Proxy field, select Use Proxy or Use Proxy with Credentials and enter your proxy server address and port ID. (Consult your network administrator for this information.) If your proxy server requires authentication, enter the proxy username and password credentials. The credentials you add here must match the credentials configured on the proxy server.
- g. Click Save. The Fingerprint Update is performed immediately.
- 4. If you did not select the **Enforce Updates** checkbox, enforce to push the changes to your engines when the update is complete.

Schedule Fingerprint Updates

You can schedule fingerprint updates performed automatically on a daily or weekly basis.

To access the update website, you need to create an Extranet account at ExtremeNetworks.com and define a username and password for the account. You need the username and password in order to schedule updates.

- Select the **Analytics** tab in Extreme Management Center and then select the Configuration view.
- 2. In the left-panel tree, expand the System folder and select **Fingerprints**.
- 3. Click on the **Menu** icon (≡) and select Fingerprint Update Settings. The Fingerprint Update window opens.



- 4. Select the update interval which defines how frequently the update is performed: **Daily** or **Weekly**.
- 5. If you have selected **Weekly**, select the day of the week you would like the update performed.
- 6. Enter the scheduled time you would like the update performed.
- Select the Enforce Updates checkbox to automatically update fingerprints on all engines. Not selecting this checkbox requires you to update each engine manually.
- 8. The **Update Site** field displays the default path to the official fingerprint update site. Typically, this field does not change unless for security reasons the system does not have access to the internet and an internal update site must be used.
- 9. The **Update Site Port** is the port on the update site to which the update connects. The port cannot be changed unless you are using a custom update site.

- Enter the credentials used to access the fingerprint update website. These are the
 username and password credentials you defined when you created an Extranet
 account at ExtremeNetworks.com.
- 11. If your network is protected by a firewall, configure proxy server settings to use when accessing the website. In the Proxy field, select Use Proxy or Use Proxy with Credentials and enter your proxy server address and port ID. (Consult your network administrator for this information.) If your proxy server requires authentication, enter the proxy username and password credentials. The credentials you add here must match the credentials configured on the proxy server.
- 12. Click Save.
- 13. If you did not select the **Enforce Updates** checkbox, enforce to push the changes to your engines when the update is complete.

Related Information

ExtremeAnalytics tab

Custom Fingerprint Examples

The ExtremeAnalytics feature uses fingerprints to identify to which application a network traffic flow belongs. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. Extreme Management Center provides thousands of system fingerprints with the ExtremeAnalytics feature. In addition, you can create new custom fingerprints.

For additional information, see Getting Started with ExtremeAnalytics.

This Help topic provides examples of three different types of custom fingerprints you can create:

- Fingerprints Based on a Flow
- Fingerprints Based on an Application or Application Group
- Fingerprints Based on a Destination Address

For additional information, see Add and Modify Fingerprints.

Fingerprints Based on a Flow

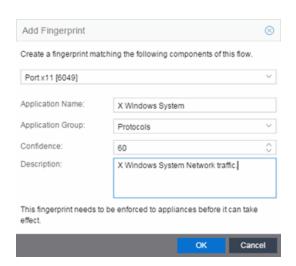
This example demonstrates how to create a custom fingerprint based on X Window System network traffic.

In the Extreme Management Center Flows table (with the Show Unclassified View selected) you notice several flows that had an X Window System source port 6049. Since these flows are not currently identified with a fingerprint, you can create a fingerprint for those flows based on the port that x11 traffic normally runs over.

Use the following steps to create the fingerprint.

- 1. Select the **Analytics** tab.
- 2. Select the **Application Flows** tab.
- 3. In the table, select the **Show Unclassified View**.
- 4. Right-click on a flow with the x11 Source Port and select Fingerprints > Add Fingerprint.

5. The Add Fingerprint window opens.



- 6. Use the drop-down list to select matching **Portx11** [6049].
- 7. Set the Application Name to X Window System.
- 8. Set the Application Group to Protocols.
- 9. Set the **Confidence** level to **60** (the default). A fingerprint with a confidence higher than 60 can supersede this fingerprint, if it also matches the flow.
- 10. Click **OK** to create the fingerprint.
- 11. Enforce to push the new fingerprint to your engines.

Fingerprints Based on an Application or Application Group

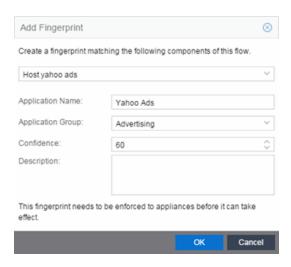
This example demonstrates how to create a fingerprint for some unclassified web traffic.

In the Extreme Management Center Application Flows table (with the Show Unclassified Web Traffic View selected) you noticed several flows for the "yahoo ads" application that are part of the Web Applications group. You want to create a fingerprint that provides an application and application group specifically for this traffic, instead of letting it default to the Web Applications group. The new fingerprint categorizes "yahoo ads" flows into the Yahoo Ads Id application and the Advertising application group.

Use the following steps to create the fingerprint.

- 1. Select the **Analytics** tab in Extreme Management Center.
- 2. Select the Application Flows tab.

- 3. In the table, select the **Show Unclassified Web Traffic View**.
- Right-click on a flow with the yahoo ads application and select Fingerprints > Add Fingerprint.
- The Add Fingerprint window opens.



- 6. Use the drop-down list to select matching the "yahoo ads" host.
- 7. Set the Application Name to Yahoo Ads.
- 8. Set the Application Group to Advertising.
- 9. Set the **Confidence** level to **60** (the default). A fingerprint with a confidence higher than 60 can supersede this fingerprint, if it also matches the flow.
- 10. Click **OK** to create the fingerprint.
- 11. Enforce to push the new fingerprint to your engines.

Fingerprints Based on a Destination Address

In both of the previous examples, you created a new custom fingerprint to cover a case where no appropriate fingerprint existed. You may also want to create a new fingerprint for traffic flows already identified as one application, but should be categorized as something else.

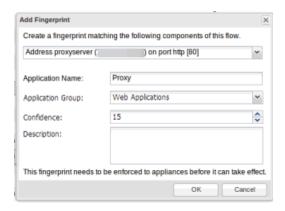
For example, let's say you have a Git repository on your network. Git repositories (a source code management system used in software development) are frequently accessed via SSH on port 22 (the standard TCP port assigned for SSH traffic). In this case, the SSH traffic flows is identified using the system SSH port-based fingerprint.

But what if you would like to more closely monitor who is accessing the Git repository? If you know you are running the Git server on a certain system (10.20.117.102 port 22, for our example), you can create a custom fingerprint to identify the Git traffic flows.

The fingerprint is based on one of the SSH flows using the IP address/port of the Git server and have a higher confidence than the system port-based fingerprint. The higher confidence fingerprint will override the lower confidence fingerprint when determining a match for the traffic flow.

Use the following steps to create the fingerprint.

- 1. Select the **Analytics** tab in Extreme Management Center.
- 2. Select the **Application Flows** tab.
- 3. In the table, right-click on an SSH port-based flow with the Git server destination address and select **Fingerprints > Add Fingerprint**.
- 4. The Add Fingerprint window opens.



- 5. Use the drop-down list to select matching the Git server IP address and port.
- 6. Set the **Application Name** to **Git**.
- 7. Select an Application Group that makes the most sense for your network. It might be Web Collaboration, Databases, Business Applications, or Storage. You can also create a new Application Group using the Create Custom Application Group option available from the gear menu in the Fingerprint Details tab. (You would need to do this before you create the custom fingerprint.)
- 8. Set the **Confidence** level to **60**, which is a higher confidence than the current fingerprint which is set at 10.
- 9. Click **OK** to create the fingerprint.
- 10. Enforce to push the new fingerprint to your engines.

Related Information

For information on related ExtremeAnalytics topics:

- Analytics
- Add and Modify Fingerprints

How to Deploy ExtremeAnalytics in an MSP or MSSP Environment

This Help topic presents instructions for deploying ExtremeAnalytics within an MSP (Managed Service Provider) or MSSP (Managed Security Service Provider) environment. It includes the following information:

- Configuring Extreme Management Center Behind a NAT Router
- Defining Interface Services

Configuring Extreme Management Center Behind a NAT Router

If the Extreme Management Center server is located behind a NAT (Network Address Translation) router, use the following steps to add an entry to the nat_config.txt file that defines the real IP address for the Extreme Management Center server. This allows the Extreme Management Center server to convert the NAT IP address received in the ExtremeAnalytics engine response to the real IP address used by the Extreme Management Center server. Not adding the real IP address for the Extreme Management Center server to the nat_config.txt file results in the ExtremeAnalytics engine incorrectly displaying a state of IMPARED (orange) rather than UP (green).

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.

- 1. On the Extreme Management Center server, add the following entry to the <install directory>/appdata/nat_config.txt file.
 <NAT IP address>=<real IP address>
- 2. Save the file.
- 3. If the Extreme Management Center Management server IP address is not configured to use the NAT IP address of the Extreme Management Center server, perform the following steps:
 - a. Enter the following command at the engine CLI:
 /opt/appid/configMgmtIP < IP address>
 Where <IP address> is the NAT IP address of the Extreme Management

Center server.

Press Enter.

b. Restart the appidserver once the new IP address is configured by typing: appidctl restart

Press Enter.

4. On the Extreme Management Center server, add the following text to the <install directory>/appdata/NSJBoss.properties file. In the second to last line, specify the hostname of the Extreme Management Center server.

NOTE: The ExtremeAnalytics engine functions as a client computer independent of the server. Both engines and clients must be able to resolve the hostname you specify.

```
# In order to connect to a NetSight server behind a NAT fi
rewall or a
# NetSight server with multiple interfaces you must define
  these two
# variables on the Extreme Management Center
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 corr
ect IP that
# they want to use as the IP to connect to.
java.rmi.server.hostname=<hostname of NetSight server>
java.rmi.server.useLocalHostname=true
```

- 5. Save the file.
- Add the Extreme Management Center server hostname to your DNS server, if necessary.

NOTE: ExtremeAnalytics engines, remote Extreme Management Center clients, and any ExtremeControl engines must be able to connect to Extreme Management Center using this hostname.

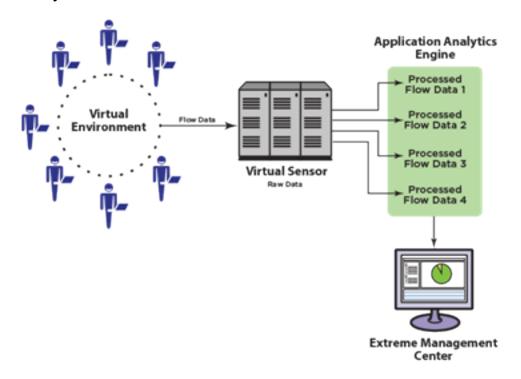
Related Information

For information on related windows:

ExtremeAnalytics Engine Advanced Configuration Panel

ExtremeAnalytics Virtual Sensor Configuration in Extreme Management Center

The ExtremeAnalytics Virtual Sensor is a virtual machine that monitors application flows across virtual environments. The Virtual Sensor sends network traffic information to your ExtremeAnalytics engine for processing. The ExtremeAnalytics engine then sends the processed information to Extreme Management Center, where it is displayed on the **Analytics** tab.



In a typical environment, one Virtual Sensor is deployed on each physical ESX host. Each Virtual Sensor is counted as 1/10 of a device towards your Extreme Management Center device count license.

Use these instructions to add one or more Virtual Sensors to Extreme Management Center. While it is possible to add the Virtual Sensor without using the VMware vSphere ExtremeConnect module in Extreme Management Center, the first method outlined in this topic is highly recommended. Refer to Virtual Sensor Installation Guide for more information about using the Virtual Sensor.

IMPORTANT: Configuration of the virtual environment via Extreme Management Center is limited to distributed virtual switches only.

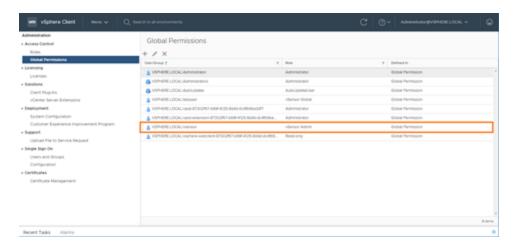
Configuring a Virtual Sensor includes the following steps:

- Prerequisites
- 2. Installing the Virtual Sensor using the Extreme Management Center Server
- 3. Adding the Virtual Sensor in ExtremeAnalytics
- 4. Configuring vCenter Settings for the Virtual Sensor

Prerequisites

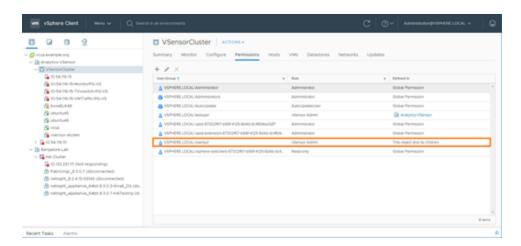
Before attempting to add Virtual Sensors to Extreme Management Center, ensure the following:

- Your vCenter vSphere login is configured with an administrative role in at least the following Global Permissions:
 - Distributed Switch
 - VSPAN operation
 - o Datastore
 - Allocate space
 - Browse datastore
 - Host
 - Local operations
 - Create virtual machine
 - Delete virtual machine
 - Reconfigure virtual machine
 - Network
 - Assign network
 - vApp
 - Import
 - View OVF environment



- Your vCenter vSphere login is configured with an administrative role in at least the following permissions in the cluster the Virtual Sensor is monitoring :
- O Distributed Switch
 - VSPAN operation
 - Datastore
 - Allocate space
 - Browse datastore
 - Remove file
 - Host
 - Local operations
 - · Create virtual machine
 - · Delete virtual machine
 - Reconfigure virtual machine
 - Network
 - Assign network
 - o Tasks
 - Create task
 - Update task
 - o vApp
 - Import
 - View OVF environment

- Virtual machine
 - Change Configuration
 - · Add new disk
 - · Advanced configuration
 - Edit Inventory
 - Create new
 - Remove
 - Interaction
 - Power off
 - Power on



 The <u>VMware Open Virtualization Format Tool</u> (OVFTool) is installed on the Extreme Management Center server.

NOTE: Instructions for this prerequesite are not included. You must be a VMware customer to download the OVF Tool. Please consult the VMware site for download and installation instructions.

Installing the Virtual Sensor Using the Extreme Management Center Server

The Virtual Sensor is installed as an .OVA file using your Extreme Management Center server.

You can install the Virtual Sensor in two different ways:

- Install Using Extreme Management Center To install using ExtremeConnect functionality in Extreme Management Center, use the instructions in this topic.
- Install Using the vSphere Web Client To install using the vSphere web client, refer to the section "Installing Virtual Sensor using vSphere Web Client" in the ExtremeAnalytics Virtual Sensor 1.0.0 Software Installation Guide.

We recommend installing the Virtual Sensor using Extreme Management Center unless you do not have the required permissions from your VMware Administrator. When installing the Virtual Sensor using the vSphere client, some information does not populate on the Analytics > Configuration > Virtual Sensors tab in Extreme Management Center, including the **Physical Host**, **Monitored Switch**, **Port Group**, and **VMs Monitored** fields in the Virtual Sensors table at the top of the tab and the Virtual Machines table at the bottom of the tab.

Install Using Extreme Management Center

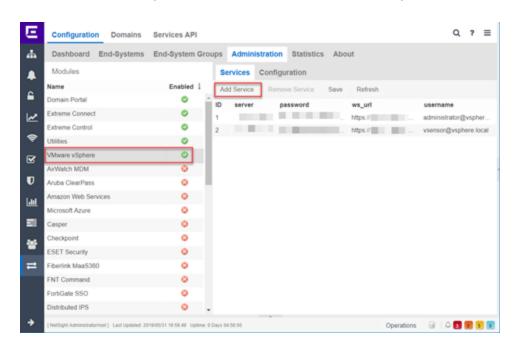
Installing the Virtual Sensor using Extreme Management Center includes two steps:

- Configuring the VMware vSphere Module in ExtremeConnect
- 2. Adding the Virtual Sensor to Extreme Management Center

Configuring the VMware vSphere Module in ExtremeConnect

Use the following steps to configure your vCenter server:

1. Open the **Connect** > **Configuration** > **Administration** tab in Extreme Management Center.



2. Select VMware vSphere in the Modules list in the left-panel.

- 3. Open the **Services** tab in the right-panel.
- 4. Click Add Service.

A new row appears.

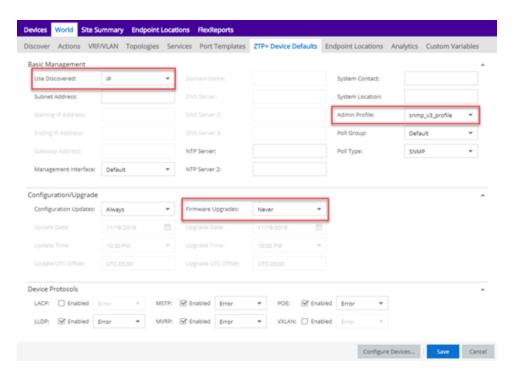
- 5. Enter the information for your vCenter server in the new row:
 - username Username used to access the vCenter server.
 - password Password used to access the vCenter server.
 - server vCenter server IP address.
 - ws_url vCenter server web service URL in the following format: https://<IP
 <p>Address of vCenter Server>:<Non-Standard Port of vCenter Server>/sdk.
 For example, https://10.20.30.40:8008/sdk.
- 6. Click Save.
- 7. Open the **Configuration** tab.
- 8. Select **true** in the Value column of the Module enabled row.
- 9. Configure the other settings according to your preferences for your vCenter server.
- 10. Click Save.

The vCenter server is configured in ExtremeConnect.

Adding the Virtual Sensor to Extreme Management Center

After configuring the vCenter server in ExtremeConnect, download and add the Virtual Sensor .OVA file in Extreme Management Center.

- 1. Access the **Network > Devices > Sites > World** site.
- 2. Open the ZTP+ Device Defaults tab in the right panel.
- 3. Select IP in the Use Discovered drop-down list.
- 4. Select an **Admin Profile** from the drop-down list that uses SNMPv3 and has a CLI credential password that is not empty.
- 5. Select **Never** in the **Firmware Upgrades** drop-down list.



6. Download the Virtual Sensor .OVA file from the Extreme Portal to your local client machine.

You need to access this location in Step 10.

NOTE: Download the **Small** .OVA file if you are using a VS100 license or the **Medium** .OVA file if you are using a VS250 license.

Microsoft Internet Explorer downloads the .OVA file as a .TAR file. Change the file extension back to .OVA once the download is complete. Using Google Chrome avoids this step and downloads the file as an .OVA file.

- 7. Navigate back to Extreme Management Center.
- 8. Open the **Network** > Firmware tab.
- Click Upload.
 The Upload Firmware to Server window displays.
- 10. Add the .OVA file to the **Upload Firmware to Server** window through one of the following methods:
 - Access the location on the local client machine into which you saved the .OVA file in Step 6 and drag and drop the Virtual Sensor .OVA file into the box at the top of the window.
 - Click the box at the top of the window and browse to the location into which you saved the .OVA file in Step 6.
- 11. Use the default transfer type setting in **Directory**.
- 12. Leave the **Subdirectory** field blank.
- 13. Click Upload.

Extreme Management Center uploads the .OVA file.

Once the .OVA file is installed, proceed to <u>Adding the Virtual Sensor in</u> ExtremeAnalytics

Adding the Virtual Sensor in ExtremeAnalytics

Once you add the Virtual Sensor .OVA file to Extreme Management Center, install the Virtual Sensor in ExtremeAnalytics.

- 1. Open the **Analytics** > Configuration tab.
- Click Virtual Sensors in the left-panel.
 The Virtual Sensors tab displays. All of the virtual machines configured in VMware vSphere are listed in the Virtual Machines section of the tab.
- Click Install in the Virtual Sensors section of the tab.
 The Install Virtual Sensor on Hypervisor Host window displays.
- 4. Click the ellipsis button (***) in the **Hypervisor Host** field and select the hypervisor host on which you are installing the Virtual Sensor.
- 5. Click the ellipsis button (***) in the **Datastore** field and select the datastore, if the field does not automatically populate.
- 6. Select the **License Type** for your Virtual Sensor from the available license types listed:

- VS100 One virtual CPU core with a capacity of 1Gbps. Select this license type when installing the Small .OVA file.
- VS250 Two virtual CPU cores with a capacity of 2.5Gbps. Select this license type when installing the Medium .OVA file.

NOTE: The Extreme Management Center evaluation license provides access for up to 10 of each license type until the evaluation expiration date. To purchase additional licenses, contact your Extreme Networks sales representative.

- 7. Select the Analytics Engine to which the Virtual Sensor sends flow data.
- 8. Click the ellipsis button (***) in the **Management Interface Network** field and select the network you use for management,

 The Virtual Sensor must communicate with the Extreme Management Center server and the ExtremeAnalytics engine using this network.
- 9. Click the ellipsis button (***) in the **Monitored Interface Port Group** field and select the port group for the monitoring interface.

The port group must be on the distributed virtual switch you are monitoring.

IMPORTANT: Select a port group that does not include an uplink port.

10. Change the Name of the Virtual Sensor, if desired.

NOTE: By default, the Virtual Sensor is named using the following format: *<IP Address of host><PortGroupName>-*VS.

- Select Use DHCP, if using a DHCP server to dynamically assign an IP address for the Virtual Sensor.
- 12. Enter an **IP Address** and **Default Gateway** for the Virtual Sensor if you are not using DHCP.

NOTE: If you enter an IP address included in a range defined in a site, the DNS Server and NTP Server are automatically populated from that site, if defined.

- 13. Enter the **Root Password** for the Virtual Sensor.
- 14. Enter the **Domain** name (for example, extremenetworks.com).
- 15. Enter the **DNS Server** IP address and **NTP Server** IP address, if applicable.
- 16. Click Install.

NOTE: The installation may take several minutes. The Operations panel shows the progress of the installation. It is normal for the progress to pause at 89 or 90 percent for a short amount of time. During this time, the OVA is deployed and Extreme Management Center is waiting for the new Virtual Sensor to be discovered via ZTP+. If Extreme Management Center is not configured to on-board ZTP+ devices automatically, you need to access to the Network > Discovered tab, and configure the new Virtual Sensor for Extreme Management Center.

The Virtual Sensor is installed and displays in the Virtual Sensors table at the top of the **Analytics > Configuration > Virtual Sensors** tab.

NOTE: If the installation process takes longer than five minutes, the Operation panel in Extreme Management Center may incorrectly indicate the installation failed. If the Operations panel indicates a failure, open the Network > Discovered tab and look for the Virtual Sensor. Continue to add the Virtual Sensor via this tab. If the Virtual Sensor is NOT listed on the Discovered tab, the cause may be one of the following:

- Extreme Management Center is not communicating with the management interface of the Virtual Sensor.
- The Virtual Sensor is configured with the Use DHCP option selected (step 11), but the DHCP process did not complete successfully.

Open the **Analytics** > Application Flows tab to ensure the Virtual Sensor is collecting flow data. Any fragmented packets the Virtual Sensor receives out of order are reordered for processing and then forwarded in the correct order. Once flows are collected by the Virtual Sensor, you can create packet captures for those flows.

NOTES: If the Virtual Sensor receives traffic as it is initializing, packet drops (reported as mbuf allocation failure) may be observed.

Configuring vCenter Settings for the Virtual Sensor

To ensure the Virtual Sensor is performing optimally, configure the following vCenter settings for your Virtual Sensor:

- 1. In vCenter, power off the Virtual Sensor virtual machine.
- 2. Right-click the virtual machine and select **Edit Settings**.
- 3. On the **Virtual Hardware** tab, expand **CPU** and allocate the CPU capacity according to the size of the OVA:

- Medium OVA 4,594 MHz
- Small OVA 2,297 MHz

The CPU capacity required (Reservation of CPU capacity) is calculated as the number of vCPUs (Small = 1, Medium = 2) * CPU speed of ESXi (2,297 MHz).

- 4. Click OK.
- 5. Right-click the virtual machine and select **Compatibility** > **Upgrade VM Compatibility**.

The virtual machine is upgraded to the latest supported version.

NOTE: If the Virtual Sensor is already running the latest supported version, the **Upgrade VM Compatibility** option is not available.

Related Information

- Virtual Sensors
- Virtual Sensor Installation Guide

Stream Flow Data from ExtremeAnalytics into Splunk

ExtremeAnalytics includes the ability to stream flow data from an Analytics engine to Splunk. To help you use Splunk with ExtremeAnalytics, we added a Splunk directory to the Extreme Management Center NetSight/appdata/Purview directory.

The Splunk directory contains the following:

- A PDF with instructions describing how to add Extreme's enterprise IPFIX fields into the Splunk vocabulary and adjust the Splunk streaming app to process the Extreme IPFIX format.
- Files that you can copy to the Splunk server to facilitate integration, instead of manually editing the files.

Use the procedures in this section to send Splunk-enriched network flow data using IPFIX.

Environment

- Extreme Management Center 8.2 and later
- Splunk 7.2.6 (single server deployment) and later
- Splunk Stream 7.1.3 and later

Overview

You can configure the Splunk Stream app to process Netflow/IPFIX flow records and add the data into the Splunk data store. Configure this partly by editing text files on the file system, and partly by using the web UI.

The instance of Splunk Stream at any site may already be configured to import one or more flow sources. Because of this, you must take care to merge the needed changes for ExtremeAnalytics with the existing file contents. After you make the file system changes, restart Splunk. Then, define a new "stream" using the user interface. Finally, enable and deploy the IPFIX exporter of ExtremeAnalytics from the Extreme Management Center user interface.

Part 1 - Making File Level Splunk Modifications

- Connect to the Splunk server via SSH.
- 2. Enter cd \$SPLUNK HOME/etc/apps/splunk app stream/local.
- 3. Copy the streamfwd.conf file. (If there is no streamfwd.conf file present, skip this step.)
- 4. Copy Extreme's version of the streamfwd.conf file and paste it into streamfwd.conf. Alternately, merge Extreme's version of streamfwd.conf settings into the existing streamfwd.conf file.
- 5. Enter cd \$SPLUNK_HOME/etc/apps/Splunk_TA_stream/local.
- 6. Copy the streamfwd.conf file. (If there is no streamfwd.conf file present, skip this step.)
- 7. Copy the streamfwd.conf file from the splunk_app_stream/local directory to this directory.
- 8. Enter cd \$SPLUNK_HOME/etc/apps/splunk_app_stream/default/vocabulary.
- 9. Copy the extr.xml file to this directory.
- 10. Enter cd \$SPLUNK_HOME/etc/apps/splunk_app_ stream/default/vocabulary/streams.
- 11. Make a copy of the netflow file.
- 12. Merge the contents of our extr.netflow file to the netflow file.

Part 2 - Creating a New Stream using the Splunk web UI

- 1. Log in to Splunk (by default, the web server is on port 8000).
- 2. Navigate to the Splunk Stream App.
- 3. Select Configure Streams from the Configuration menu.
- 4. Optionally, disable all existing streams if you installed Splunk Stream solely to integrate Analytics flow data.

- 5. Create a new stream.
 - a. Click on **New Stream > Metadata Stream** to start creating a stream for ExtremeAnalytics.
 - b. Select the **Netflow** protocol.
 - c. Type a name and description for your stream.
 - d. If you are an advanced Splunk user, you can choose an appropriate Aggregation method. This can be changed later, as well. We will leave aggregate off in this topic.
 - e. In the **Fields** dialog box, you do not have to deselect any filters. You can deselect any selected filters you do not need later.
 - f. In the **Filters** dialog box, do not deselect any filters. You can deselect any selected filters you do not need later.
 - g. In the **Settings** dialog box, select **Enabled** for Status.
 - h. In the **Groups** dialog box, leave the default, and click **Create Stream**.
 - i. After you click **Create Stream**, Splunk confirms that your stream was created.

Part 3 - Configuring each Analytics Engine to Export IPFIX Data to the Splunk Server

- 1. Log in to Extreme Management Center.
- 2. Navigate to Analytics > Configuration.
- 3. Use the following steps for each engine you want exporting flows to Splunk,
 - a. Click on its Configuration page.
 - b. Expand the IPFIX/Netflow Exporter section and fill out the required fields.
 - c. Ensure that **Export Enabled** is checked.
 - d. Set the **Export IP** to the Splunk server IP address.
 - e. Set the **Export Port** to 2055 unless this has been customized in Splunk.
 - f. Set the **Protocol** to either IPFIX or IPFIX with Padded Strings.
 - g. Select **Save**.

The Metadata contains some protocol-specific data for the analysis of DNS, HTTP, etc. This additional data can double the size of the records.

4. Enforce the changes by expanding Engines in the left-panel of **Analytics** > **Configuration**, selecting the engine in the left-panel menu, and clicking the **Enforce** button. The Splunk Search displays "netflowdata" within a minute or two.

Appendix

Files

\$\$PLUNK/etc/apps/splunk_app_stream/local/streamfwd.conf \$\$PLUNK/etc/apps/\$plunk_TA_stream/local/streamfwd.conf

```
[streamfwd]
port = 8889
ipAddr = 127.0.0.1
```

```
netflowReceiver.0.ip = 10.54.116.45
netflowReceiver.0.port = 2055
netflowReceiver.0.decoder = netflow
#netflowElement.997.enterpriseid = 1916
netflowElement.997.id = 96
netflowElement.997.termid = extr.appName
netflowElement.998.enterpriseid = 1916
netflowElement.998.id = 371
netflowElement.998.termid = extr.userName
netflowElement.999.enterpriseid = 1916
netflowElement.999.id = 372
netflowElement.999.termid = extr.appGroupName
netflowElement.1000.enterpriseid = 1916
netflowElement.1000.id = 1000
netflowElement.1000.termid = extr.srcHostName
netflowElement.1001.enterpriseid = 1916
netflowElement.1001.id = 1001
netflowElement.1001.termid = extr.dstHostName
netflowElement.1002.enterpriseid = 1916
netflowElement.1002.id = 1002
netflowElement.1002.termid = extr.netResponseTime
netflowElement.1003.enterpriseid = 1916
netflowElement.1003.id = 1003
netflowElement.1003.termid = extr.appResponseTime
netflowElement.1004.enterpriseid = 1916
netflowElement.1004.id = 1004
netflowElement.1004.termtype = ipaddress
netflowElement.1004.termid = extr.serverAddress
netflowElement.1005.enterpriseid = 1916
netflowElement.1005.id = 1005
netflowElement.1005.termid = extr.nacProfile
netflowElement.1006.enterpriseid = 1916
netflowElement.1006.id = 1006
netflowElement.1006.termid = extr.detailedLocation
netflowElement.1007.enterpriseid = 1916
netflowElement.1007.id = 1007
netflowElement.1007.termid = extr.oneSidedFlow
```

```
netflowElement.1008.enterpriseid = 1916
netflowElement.1008.id = 1008
netflowElement.1008.termid = extr.clientLocation
netflowElement.1009.enterpriseid = 1916
netflowElement.1009.id = 1009
netflowElement.1009.termid = extr.serverLocation
netflowElement.1010.enterpriseid = 1916
netflowElement.1010.id = 1010
netflowElement.1010.termid = extr.metaData
```

\$\$PLUNK/etc/apps/splunk_app_ stream/default/vocabulary/extreme.xml

```
<?xml version="1.0" encoding="UTF-8"?>
<CmConfig xmlns="http://purl.org/cloudmeter/config" version="7.1.1">
       <Vocabulary id="extr">
         <Locked>true</Locked>
         <Name>ExtremeAnalytics Netflow Protocol Vocabulary
               <Term id="extr.appName">
                       <Type>string</Type>
                       <Comment>Application Name - ExtremeAnalytics//Comment>
               </Term>
               <Term id="extr.userName">
                       <Type>string</Type>
                       <Comment>User Name - ExtremeAnalytics</Comment>
               </Term>
               <Term id="extr.appGroupName">
                       <Type>string</Type>
                       <Comment>Application Group Name - ExtremeAnalytics/Comme
               </Term>
               <Term id="extr.srcHostName">
                       <Type>string</Type>
                       <Comment>Source Host Name - ExtremeAnalytics</Comment>
               </Term>
               <Term id="extr.dstHostName">
                       <Type>string</Type>
                       <Comment>Destination Host Name - ExtremeAnalytics/Commen
               </Term>
               <Term id="extr.netResponseTime">
                       <Type>uint64</Type>
                       <Comment>TCP Response Time - ExtremeAnalytics//Comment>
               </Term>
               <Term id="extr.appResponseTime">
```

<Type>uint64</Type>

```
<Comment>Application Response Time - ExtremeAnalytics/Co
               </Term>
               <Term id="extr.serverAddress">
                      <Type>shortstring</Type>
                      <Comment>Server IP Address - ExtremeAnalytics//Comment>
               </Term>
               <Term id="extr.nacProfile">
                      <Type>string</Type>
                      <Comment>Client NAC Profile - ExtremeAnalytics//Comment>
               </Term>
               <Term id="extr.detailedLocation">
                      <Type>string</Type>
                      <Comment>Client Detailed Location - ExtremeAnalytics/Com
               </Term>
               <Term id="extr.oneSidedFlow">
                      <Type>uint8</Type>
                      <Comment>One Sided Flow Boolean - ExtremeAnalytics/Comme
              </Term>
               <Term id="extr.clientLocation">
                      <Type>string</Type>
                      <Comment>Client Location - ExtremeAnalytics</Comment>
               </Term>
               <Term id="extr.serverLocation">
                      <Type>string</Type>
                      <Comment>Server Location - ExtremeAnalytics
              </Term>
               <Term id="extr.metaData">
                      <Type>string</Type>
                      <Comment>Extra Data - ExtremeAnalytics
               </Term>
       </Vocabulary>
</CmConfig>
$$PLUNK/etc/apps/splunk_app_
stream/default/streams/netflow (additions)
       "aggType": "value",
       "desc": "Application Name - ExtremeAnalytics",
       "enabled": true,
       "name": "app-name",
       "term": "extr.appName"
},
       "aggType": "value",
```

```
"desc": "Application Group Name - ExtremeAnalytics",
       "enabled": true,
       "name": "application-group",
       "term": "extr.appGroupName"
},
       "aggType": "value",
       "desc": "User Name - ExtremeAnalytics",
       "enabled": true,
       "name": "x-user-name",
       "term": "extr.userName"
},
       "aggType": "value",
       "desc": "Source Host Name - ExtremeAnalytics",
       "enabled": true,
       "name": "src host name",
       "term": "extr.srcHostName"
},
       "aggType": "value",
       "desc": "Destination Host Name - ExtremeAnalytics",
       "enabled": true,
       "name": "dst host name",
       "term": "extr.dstHostName"
},
       "aggType": "value",
       "desc": "TCP Response Time - ExtremeAnalytics",
       "enabled": true,
       "name": "tcp-response-time",
       "term": "extr.netResponseTime"
},
{
       "aggType": "value",
       "desc": "Application Response Time - ExtremeAnalytics",
       "enabled": true,
       "name": "app-response-time",
       "term": "extr.appResponseTime"
},
       "aggType": "value",
       "desc": "Server IP Address - ExtremeAnalytics",
       "enabled": true,
       "name": "server-ip-address",
       "term": "extr.serverAddress"
```

```
},
       "aggType": "value",
       "desc": "Client NAC Profile - ExtremeAnalytics",
       "enabled": true,
       "name": "client-nac-profile",
       "term": "extr.nacProfile"
},
{
       "aggType": "value",
       "desc": "Client Detailed Location - ExtremeAnalytics",
       "enabled": true,
       "name": "client-detailed-location",
       "term": "extr.detailedLocation"
},
       "aggType": "value",
       "desc": "One Sided Flow Boolean - ExtremeAnalytics",
       "enabled": true,
       "name": "one-sided-flow",
       "term": "extr.oneSidedFlow"
},
{
       "aggType": "value",
       "desc": "Client Location - ExtremeAnalytics",
       "enabled": true,
       "name": "client-location",
       "term": "extr.clientLocation"
},
       "aggType": "value",
       "desc": "Server Location - ExtremeAnalytics",
       "enabled": true,
       "name": "server-location",
       "term": "extr.serverLocation"
},
{
       "aggType": "value",
       "desc": "Extra Meta Data - ExtremeAnalytics",
       "enabled": true,
       "name": "meta-data",
       "term": "extr.metaData"
}
```

Stream Flow Data from ExtremeAnalytics into Elastic Stack

ExtremeAnalytics includes the ability to stream flow data from an ExtremeAnalytics engine to Elastic Stack (aka ELK stack). To help you use Elastic Stack with ExtremeAnalytics, we added an ELK directory to the Extreme Management Center NetSight/appdata/Purview directory.

The ELK directory contains the following:

- A PDF describing how to add the open-source "Elastiflow" module to an ELK server and how to update this deployment to make Elastiflow aware of Extreme's IPFIX format.
- Files that you can copy to the ELK server to assist with the customization.

Use the procedures in this section to send Extreme Networks-enriched network flow data to Elastic Stack using IPFIX and ElastiFlow.

Environment

- Extreme Management Center 8.2 and later
- Elastic Stack 6.7 (single server deployment) and later
- ElastiFlow 3.4.2 and later (version compatible with Elastic Stack 6.7) running on Ubuntu Server 18.04

Overview

Use ElastiFlow to collect IPFIX flow data and visualize the results using Elastic Stack. ElastiFlow requires a working Elastic Stack and it must be configured by editing text files on the file system and by using the Kibana user interface.

The installation steps assume that IPFIX will only be sent over UDP on port 2055 from Extreme Management Center. After making the file system and UI changes, you must restart the Elastic Stack components. Finally, enable and deploy the IPFIX exporter of ExtremeAnalytics from the Extreme Management Center user interface.

Part 1 - Installing and Configuring ElastiFlow and Elastic Stack

ElastiFlow installation instructions can be found at:

https://github.com/robcowart/elastiflow/blob/master/INSTALL.md

Other helpful installation links are:

- https://www.catapultsystems.com/blogs/install-elastiflow-on-ubuntu-18-04-part-1/
- https://sadsloth.net/post/elastiflow-ubuntu1804/
- 1. Download and install the Oracle Server JRE 8 in one of the following ways:
 - a. Via the Ubuntu bundle by entering the following commands:

```
$ sudo apt install openjdk-8-jre
```

Verify Oracle Server JRE 8 is installed properly by entering the following:

```
$ java -version
openjdk version "1.8.0_212"
OpenJDK Runtime Environment (build 1.8.0_212-8u212-
b03-0ubuntu1.18.04.1-b03)
OpenJDK 64-Bit Server VM (build 25.212-b03, mixed mode)
```

- b. Via the external Java web site by following the instructions at: https://www.oracle.com/technetwork/java/javase/downloads/server-jre8-downloads-2133154.html
- 2. Install Elastic Stack 6.7 (6.7.2 at time of writing) by entering the following commands:

```
wget -q0 - https://artifacts.elastic.co/GPG-KEY-elasticsearch |
sudo apt-key add
echo "deb https://artifacts.elastic.co/packages/6.x/apt stable
main" | sudo tee -a
/etc/apt/sources.list.d/elastic-6.x.list
```

a. Install Elastic Search:

```
sudo apt update; sudo apt -y install elasticsearch
(You may need to run sudo apt-get -f install to update packages before
you can install elasticsearch .)
```

```
sudo /bin/systemctl daemon-reload
sudo /bin/systemctl enable elasticsearch.service
sudo systemctl start elasticsearch.service
```

b. Install Kibana:

```
sudo apt -y install kibana
sudo /bin/systemctl daemon-reload
sudo /bin/systemctl enable kibana.service
sudo systemctl start kibana.service
```

c. Install Logstash:

```
sudo apt -y install logstash
```

3. Configure Elastic Stack components.

a. Configure Elasticsearch:

```
Change #network.host: 192.168.0.1 to network.host: localhost
Save the file.
```

sudo nano /etc/elasticsearch/elasticsearch.yml

b. Configure Kibana:

WARNING: The following changes allow external access to Kibana. You must follow best practices to restrict public access to the system.

```
sudo nano /etc/kibana/kibana.yml
```

```
Change #server.host: "localhost" to server.host: "YOUR_
KIBANA_SERVER_IP".
Save the file.
```

```
sudo systemctl restart kibana
```

The default port for the Kibana's server is tcp port 5601. Create a firewall rule to allow users access to the Kibana server. The rule should be something like:

sudo ufw allow from YOUR_MANAGEMENT_SUBNET to any port 5601 proto tcp

c. Configure Logstash:

Edit JVM setting in /etc/logstash/jvm.options.

sudo nano /etc/logstash/jvm.options

Change -Xms1g to -Xms4g. Change -Xmx1g to -Xmx4g.

Save the file.

Add required Logstash plugins.

sudo /usr/share/logstash/bin/logstash-plugin update
logstash-codec-netflow;
sudo /usr/share/logstash/bin/logstash-plugin update
logstash-input-udp;
sudo /usr/share/logstash/bin/logstash-plugin update
logstash-filter-dns;
sudo /usr/share/logstash/bin/logstash-plugin update
logstash-filter-geoip;
sudo /usr/share/logstash/bin/logstash-plugin update
logstash-filter-translate

4. Download and extract ElastiFlow v3.4.2.tar.gz to /usr/local/src.

wget

https://github.com/robcowart/elastiflow/archive/v3.4.2.tar.gz
sudo tar xvzf v3.4.2.tar.gz -C /usr/local/src

5. Copy logstash configuration.

```
cd /usr/local/src
sudo cp -arv elastiflow-3.4.2/logstash/elastiflow/.
/etc/logstash/elastiflow
```

- 6. Merge Extreme Networks specific IPFIX definitions with ElastiFlow.
 - a. To prepare for this step, copy the extr_elastiflow_3.4.2.tar.gz file to the /etc/logstash directory.

- b. cd /etc/logstash
 sudo tar xvzf extr elastiflow 3.4.2.tar.gz
- 7. Configure logstash pipelines.yml.
 - a. sudo nano /etc/logstash/pipelines.yml
 - b. Add:

```
- pipeline.id: elastiflow
path.config: "/etc/logstash/elastiflow/conf.d/*.conf"
```

NOTE: Be careful about spacing and extra blank lines with the following file.

Make sure there are no blank lines between the main definition and the elastiflow definition.

```
- pipeline.id: main
path.config: "/etc/logstash/conf.d/*.conf"
- pipeline.id: elastiflow
path.config: "/etc/logstash/elastiflow/conf.d/*.conf"
```

Do not add a <CR> at the end of the file. Save the file.

The following configuration example receives Extreme Networks-enriched IPFIX on UDP port 2055 only. We can modify the Logstash configuration and service parameters to limit the Logstash plugins that are loaded.

To prepare for this step, copy the extr_udp_2055_logstash.tar.gz file to the / directory.

```
cd /
sudo tar xvzf extr_udp_2055_logstash.tar.gz
cd /etc/logstash/elastiflow/conf.d
sudo mv 10_input_netflow_ipv4.logstash.conf 10_input_netflow_
ipv4.logstash.conf.disabled;
sudo mv 10_input_sflow_ipv4.logstash.conf 10_input_sflow_
ipv4.logstash.conf.disabled;
sudo mv 20_filter_20_netflow.logstash.conf 20_filter_20_
netflow.logstash.conf.disabled;
sudo mv 20_filter_40_sflow.logstash.conf 20_filter_40_
sflow.logstash.conf.disabled;
```

8. To prevent packet drops:

```
sudo cp -arv /usr/local/src/elastiflow-3.4.2/sysctl.d/87-
elastiflow.conf /etc/sysctl.d/.
```

To apply sysctl changes without restart:

```
sudo sysctl --system
```

- 9. Set up Kibana index patterns.
 - a. Download https://github.com/robcowart/elastiflow/archive/v3.4.2.tar.gz to the computer that you will use to view ElastiFlow.
 - b. Extract the files to a temporary directory.
 - c. In the web browser, go to http://YOUR_KIBANA_SERVER_IP:5601.
 - d. In the Kibana UI, select **Management** on the left side of the screen and then **Saved Objects**.
 - e. Click Import.
 - f. Select elastiflow-3.4.2\kibana\elastiflow.kibana.6.7.x.json.
 - g. Click Import.

10. Start Logstash:

- a. sudo /usr/share/logstash/bin/system-install sudo systemctl daemon-reload sudo systemctl start logstash
- b. Run sudo tail -f /var/log/logstash/logstash-plain.log to see messages from Logstash.

```
It is normal to see info messages such as the following in the log:
```

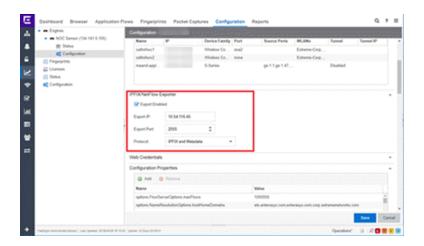
```
[INFO ][logstash.config.source.local.configpathloader] No
config files found in path
{:path=>"/etc/logstash/conf.d/*.conf"}
```

Depending on your system, it can take a few minutes for Logstash to start up. When it starts, you should see a message that says: Successfully started Logstash API endpoint.

11. Go to https://github.com/robcowart/elastiflow/blob/master/INSTALL.md and perform "Recommended Kibana Advanced Settings."

Part 2 - Configuring each Analytics Engine to export IPFIX data to the Elastic Stack server

- 1. Log in to the Extreme Management Center.
- 2. Navigate to Analytics > Configuration.
- 3. Use the following steps for each engine that you want exporting flows to ElastiFlow:
 - a. Click on the engine's Configuration page.
 - b. Expand the IPFIX/Netflow Exporter section.
 - c. Fill out the required fields:
 - · Set Export IP to the Logstash IP address.
 - Set Export Port to 2055.
 - Set Protocol to either IPFIX or IPFIX and Metadata.
 Metadata contains some protocol-specific data for the analysis of DNS, HTTP, etc. This additional data can double the size of the records.



- 4. Enforce the changes.
 You should see Ipfixdata in ElastiFlow within a minute or two.
- 5. Go to Kibana at http://YOUR_KIBANA_SERVER_IP:5601.
- 6. Select Dashboard.
- 7. Click on **ElastiFlow: Overview**. You should see some data.

8. The Logstash log file displays the following message:

```
... [WARN ] [logstash.codecs.netflow ] Can't (yet) decode flowset id xxx from observation domain id xxxx, because no template to decode it with has been received.
```

This message is normal. It goes away after one minute when Logstash receives the IPFIX data template, and this message will stop being added to the log file.

Appendix: Files

Additions to ipfix.yml in extr_elastiflow_3.4.2.tar.gz

```
# Extreme Networks (formerly 'Enterasys')
 1916:
       0:
       - :skip
       371:
       - :string
       - :extr userName
       372:
       - :string
       - :extr_appGroupName
       1000:
       - :string
       - :extr srcHostName
       1001:
       - :string
       - :extr dstHostName
       1002:
       - :uint64
       - :extr netResponseTime
       1003:
       - :uint64
       - :extr appResponseTime
       1004:
       - :ip4 addr
       - :extr serverAddress
       1005:
       - :string
       - :extr nacProfile
       1006:
```

```
- :string
- :extr_detailedLocation
1007:
- :uint8
- :extr_oneSidedFlow
1008:
- :string
- :extr_clientLocation
1009:
- :string
- :extr_serverLocation
1010:
- :string
- :extr_metaData
```

Additions to elastiflow.template.json and elastiflow_dynamic.template.json in extr_elastiflow_3.4.2.tar.gz

```
{
       "ipfix.extr userName": {
       "path match": "ipfix.extr userName",
       "mapping": {
       "type": "keyword"
   }
 },
       "ipfix.extr appGroupName": {
       "path match": "ipfix.extr appGroupName",
       "mapping": {
       "type": "keyword"
 },
       "ipfix.extr srcHostName": {
       "path match": "ipfix.extr srcHostName",
       "mapping": {
      "type": "keyword"
   }
 },
       "ipfix.extr dstHostName": {
       "path match": "ipfix.extr dstHostName",
```

```
"mapping": {
     "type": "keyword"
  }
},
     "ipfix.extr netResponseTime": {
     "path_match": "ipfix.extr_netResponseTime",
     "mapping": {
     "type": "long"
  }
},
     "ipfix.extr_appResponseTime": {
     "path match": "ipfix.extr appResponseTime",
     "mapping": {
     "type": "long"
},
{
     "ipfix.extr serverAddress": {
     "path match": "ipfix.extr_serverAddress",
     "mapping": {
     "type": "ip"
  }
},
{
     "ipfix.extr nacProfile": {
     "path match": "ipfix.extr nacProfile",
     "mapping": {
     "type": "keyword"
  }
},
     "ipfix.extr detailedLocation": {
     "path match": "ipfix.extr detailedLocation",
     "mapping": {
     "type": "keyword"
  }
},
{
```

```
"ipfix.extr oneSidedFlow": {
     "path match": "ipfix.extr oneSidedFlow",
     "mapping": {
     "type": "long"
  }
},
     "ipfix.extr clientLocation": {
     "path match": "ipfix.extr clientLocation",
     "mapping": {
     "type": "keyword"
  }
},
     "ipfix.extr serverLocation": {
     "path match": "ipfix.extr serverLocation",
     "mapping": {
     "type": "keyword"
 }
},
     "ipfix.extr metaData": {
     "path match": "ipfix.extr metaData",
     "mapping": {
     "type": "keyword"
  }
},
```

Additions to elastiflow_static.template.json in extr_elastiflow_3.4.2.tar.gz

```
"extr_userName": {
     "type": "keyword"
},
"extr_appGroupName": {
     "type": "keyword"
},
"extr_srcHostName": {
     "type": "keyword"
},
"extr_srcHostName": {
```

```
"type": "keyword"
"extr_netResponseTime": {
  "type": "long"
"extr appResponseTime": {
 "type": "long"
"extr serverAddress": {
  "type": "ip"
},
"extr_nacProfile": {
  "type": "keyword"
"extr detailedLocation": {
  "type": "keyword"
},
"extr oneSidedFlow": {
  "type": "long"
} ,
"extr_clientLocation": {
 "type": "keyword"
} ,
"extr serverLocation": {
 "type": "keyword"
} ,
"extr_metaData": {
 "type": "keyword"
},
```