



Extreme Networks Extreme Management Center[®]

Application Analytics User Guide

Copyright © 2016 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/

Support

For product support, including documentation, visit: www.extremenetworks.com/support/

Contact

Extreme Networks, Inc.,
145 Rio Robles
San Jose, CA 95134
Tel: +1 408-579-2800

Toll-free: +1 888-257-3000



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. DEFINITIONS. "Affiliates" means any person, partnership, corporation, limited liability company, or other form of enterprise that directly or indirectly through one or more intermediaries, controls, or is controlled by, or is under common control with the party specified. "Server Application" shall refer to the License Key for software installed on one or more of Your servers. "Client Application" shall refer to the application to access the Server Application. "Licensed Materials" shall collectively refer to the licensed software (including the Server Application and Client Application), Firmware, media embodying the software, and the documentation. "Concurrent User" shall refer to any of Your individual employees who You provide access to the Server Application at any one time. "Firmware" refers to any software program or code imbedded in chips or other media. "Licensed Software" refers to the Software and Firmware collectively.
2. 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. AUDIT RIGHTS. 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. RESTRICTION AGAINST COPYING OR MODIFYING LICENSED

MATERIALS. Except as expressly permitted in this Agreement, You may not copy or otherwise reproduce the Licensed Materials. In no event does the limited copying or reproduction permitted under this Agreement include the right to decompile, disassemble, electronically transfer, or reverse engineer the Licensed Software, or to translate the Licensed Software into another computer language.

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

7. TITLE AND PROPRIETARY RIGHTS

- a. The Licensed Materials are copyrighted works and are the sole and exclusive property of Extreme, any company or a division thereof which Extreme controls or is controlled by, or which may result from the merger or consolidation with Extreme (its "Affiliates"), and/or their suppliers.

This Agreement conveys a limited right to operate the Licensed Materials and shall not be construed to convey title to the Licensed Materials to You. There are no implied rights. You shall not sell, lease, transfer, sublicense, dispose of, or otherwise make available the Licensed Materials or any portion thereof, to any other party.

- b. You further acknowledge that in the event of a breach of this Agreement, Extreme shall suffer severe and irreparable damages for which monetary compensation alone will be inadequate. You therefore agree that in the event of a breach of this Agreement, Extreme shall be entitled to monetary damages and its reasonable attorney's fees and costs in enforcing this Agreement, as well as injunctive relief to restrain such breach, in addition to any other remedies available to Extreme.

8. PROTECTION AND SECURITY. In the performance of this Agreement or in contemplation thereof, You and your employees and agents may have access to private or confidential information owned or controlled by Extreme relating to the Licensed Materials supplied hereunder including, but not limited to, product specifications and schematics, and such information may contain proprietary details and disclosures. All information and data so acquired by You or your employees or agents under this Agreement or in contemplation hereof shall be and shall remain Extreme's exclusive property, and You shall use your best efforts (which in any event shall not be less than the efforts You take to ensure the confidentiality of your own proprietary and other confidential information) to keep, and have your employees and agents keep, any and all such information and data confidential, and shall not copy, publish, or disclose it to others, without Extreme's prior written approval, and shall return such information and data to Extreme at its request. Nothing herein shall limit your use or dissemination of information not actually derived from Extreme or of information which has been or subsequently is made public by Extreme, or a third party having authority to do so.

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

of any unauthorized use thereof. You acknowledge that the Licensed Materials contain valuable confidential information and trade secrets, and that unauthorized use, copying and/or disclosure thereof are harmful to Extreme or its Affiliates and/or its/their software suppliers.

9. MAINTENANCE AND UPDATES. Updates and certain maintenance and support services, if any, shall be provided to You pursuant to the terms of an Extreme Service and Maintenance Agreement, if Extreme and You enter into such an agreement. Except as specifically set forth in such agreement, Extreme shall not be under any obligation to provide Software Updates, modifications, or enhancements, or Software maintenance and support services to You.
10. DEFAULT AND TERMINATION. In the event that You shall fail to keep, observe, or perform any obligation under this Agreement, including a failure to pay any sums due to Extreme, or in the event that you become insolvent or seek protection, voluntarily or involuntarily, under any bankruptcy law, Extreme may, in addition to any other remedies it may have under law, terminate the License and any other agreements between Extreme and You.
 - a. Immediately after any termination of the Agreement or if You have for any reason discontinued use of Software, You shall return to Extreme the original and any copies of the Licensed Materials and remove the Licensed Software from any modular works made pursuant to Section 3, and certify in writing that through your best efforts and to the best of your knowledge the original and all copies of the terminated or discontinued Licensed Materials have been returned to Extreme.
 - b. Sections 1, 7, 8, 10, 11, 12, 13, 14 and 15 shall survive termination of this Agreement for any reason.
11. EXPORT REQUIREMENTS. You are advised that the Software is of United States origin and subject to United States Export Administration Regulations; diversion contrary to United States law and regulation is prohibited. You agree not to directly or indirectly export, import or transmit the Software to any country, end user or for any Use that is prohibited by applicable United States regulation or statute (including but not limited to those countries embargoed from time to time by the United States government); or contrary to the laws or regulations of any other governmental entity that has jurisdiction over such export, import, transmission or Use.
12. UNITED STATES GOVERNMENT RESTRICTED RIGHTS. The Licensed Materials (i) were developed solely at private expense; (ii) contain "restricted computer software" submitted with restricted rights in

accordance with section 52.227-19 (a) through (d) of the Commercial Computer Software-Restricted Rights Clause and its successors, and (iii) in all respects is proprietary data belonging to Extreme and/or its suppliers. For Department of Defense units, the Licensed Materials are considered commercial computer software in accordance with DFARS section 227.7202-3 and its successors, and use, duplication, or disclosure by the U.S. Government is subject to restrictions set forth herein.

13. LIMITED WARRANTY AND LIMITATION OF LIABILITY. The only warranty that Extreme makes to You in connection with this license of the Licensed Materials is that if the media on which the Licensed Software is recorded is defective, it will be replaced without charge, if Extreme in good faith determines that the media and proof of payment of the license fee are returned to Extreme or the dealer from whom it was obtained within ninety (90) days of the date of payment of the license fee.
NEITHER EXTREME NOR ITS AFFILIATES MAKE ANY OTHER WARRANTY OR REPRESENTATION, EXPRESS OR IMPLIED, WITH RESPECT TO THE LICENSED MATERIALS, WHICH ARE LICENSED "AS IS". THE LIMITED WARRANTY AND REMEDY PROVIDED ABOVE ARE EXCLUSIVE AND IN LIEU OF ALL OTHER WARRANTIES, INCLUDING IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE EXPRESSLY DISCLAIMED, AND STATEMENTS OR REPRESENTATIONS MADE BY ANY OTHER PERSON OR FIRM ARE VOID. ONLY TO THE EXTENT SUCH EXCLUSION OF ANY IMPLIED WARRANTY IS NOT PERMITTED BY LAW, THE DURATION OF SUCH IMPLIED WARRANTY IS LIMITED TO THE DURATION OF THE LIMITED WARRANTY SET FORTH ABOVE. YOU ASSUME ALL RISK AS TO THE QUALITY, FUNCTION AND PERFORMANCE OF THE LICENSED MATERIALS. IN NO EVENT WILL EXTREME OR ANY OTHER PARTY WHO HAS BEEN INVOLVED IN THE CREATION, PRODUCTION OR DELIVERY OF THE LICENSED MATERIALS BE LIABLE FOR SPECIAL, DIRECT, INDIRECT, RELIANCE, INCIDENTAL OR CONSEQUENTIAL DAMAGES, INCLUDING LOSS OF DATA OR PROFITS OR FOR INABILITY TO USE THE LICENSED MATERIALS, TO ANY PARTY EVEN IF EXTREME OR SUCH OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES. IN NO EVENT SHALL EXTREME OR SUCH OTHER PARTY'S LIABILITY FOR ANY DAMAGES OR LOSS TO YOU OR ANY OTHER PARTY EXCEED THE LICENSE FEE YOU PAID FOR THE LICENSED MATERIALS.
Some states do not allow limitations on how long an implied warranty lasts and some states do not allow the exclusion or limitation of incidental or consequential damages, so the above limitation and exclusion may not apply

to You. This limited warranty gives You specific legal rights, and You may also have other rights which vary from state to state.

14. JURISDICTION. 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

Legal Notices	1
Trademarks	1
Support	1
Contact	1
Extreme Networks® Software License Agreement	2
Table of Contents	10
Extreme Management Center® Application Analytics™ Help	1
Document Version	1
Application Analytics Licensing	3
Using Licenses to Establish Flow Rate Capacity	4
Using Licenses to Establish Flow Client Capacity	4
Getting Started with Application Analytics	5
Application Analytics Access Requirements	5
Application Analytics Engine Configuration	5
Enable NetFlow Collection	6
Configure Network Locations	6
Application Analytics Application Data Collection	7
Data Collection Overview	7
Collection Targets	8
Collection Statistics	9
Collection Intervals	10
Using Locations to Collect In-Network Traffic	12
Data Collector Types	13
General Usage Collectors	13
Hourly General Usage Collectors	14

High-Rate General Usage Collectors	16
End-System Details Collector	17
Flow Information Sources	17
Enabling Extreme Access Control Integration	18
Reports	19
Dashboard Report	20
Browser Reports	21
Add and Modify Fingerprints	22
Adding a Fingerprint	22
Modifying a Fingerprint	25
Enabling or Disabling a Fingerprint	26
Deleting a Custom Fingerprint	27
Updating Fingerprints	28
Perform a Fingerprint Update	28
Schedule Fingerprint Updates	30
Applications Browser	32
Overview	32
Data Aggregation	33
Options	34
Bookmark the Report	37
Save to Report Designer	38
Export to CSV	38
Custom Fingerprint Examples	40
Fingerprints Based on a Flow	40
Fingerprints Based on an Application or Application Group	41
Fingerprints Based on a Destination Address	42

How to Deploy Application Analytics in an MSP or MSSP Environment	45
Configuring Extreme Management Center Behind a NAT Router	45
Network Locations	47
Managing Locations	48
Adding Locations	48
Editing Locations	49
Removing Locations	50
Importing Locations	50
Exporting Locations	51
Searching Locations	51
Analytics	52
Dashboard	53
Graph Descriptions	53
Overview	54
Client/Server Dashboard Reports	54
Applications Browser Dashboard Report	55
High-Rate Application Collector Dashboard Report	55
Industry Dashboards	55
IP Reputation Dashboard	55
Application Map	57
Response Time Dashboard	57
Network Service Dashboard	57
Browser	58
Application Flows	58
Bidirectional Flows	59
Unidirectional Flows	63

Report Features	65
Fingerprints	67
Fingerprint Table	68
Gear Menu	68
Column Definitions	68
Configuration	71
Adding an Engine	71
Enforcing an Engine	71
Engine Administrative Options and Reports	72
Overview	72
Application Analytics Engines	73
Application Analytics System	73
Reports	74
Report Descriptions	75
Bandwidth for a Client Over Time	75
Locations Using the Most Bandwidth	75
Most Popular Applications	75
Most Used Applications for a Client	75
Most Used Applications for a User Name	75
Network Activity by Location	75
Network Activity for a Client	76
Network Activity for an Application	76
Slowest Applications by Location	76
Top Applications Group Radar	76
Top Applications Radar	76
Top Applications TreeMap	76

Top N Clients	77
Top N Applications	77
Top N Servers	77
Application Analytics Engine Advanced Configuration	79
Collection Privacy Levels	81
Client Aggregation	81
Slow Client Data	82
Max End-Systems in Hourly Details	82
Sensor Log Levels	82
Access Control Integration	83
Wireless Controller Flow Sources	83
Web Credentials	84
Configuration Properties	84
Sensor Modules	84
Network Settings	85
DNS	85
NTP	86
SSH	87
SNMP	88

Extreme Management Center®

Application Analytics™ Help

Application Analytics provides Layer 7 application visibility on your network. Combining Extreme Management Center, S-Series and/or K-Series devices, and the Application Analytics 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.

Application Analytics 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 Application Analytics. Application Analytics samples up to 32 packets of every new flow. The CoreFlow2 ASIC identifies those new flows and sends the packets to the Application Analytics engine where it is combined with the non-sampled NetFlow traffic for the remainder of the flow, allowing the Application Analytics engine to process traffic at unprecedented scale. The Application Analytics engine determines the application, aggregates the data, adds additional context, and then sends it to Management Center for visibility.

Application Analytics requires the NMS-ADV license.

Document Version

The following table displays the revision history for the Application Analytics Help documentation.

Date	Revision Number	Description
06-16	7.0 Revision -00	Extreme Management Center 7.0 release
07-15	6.3 Revision -00	NetSight 6.3 release
01-15	6.2 Revision -00	NetSight 6.2 release
06-14	6.1 Revision -00	NetSight 6.1 release
02-14	6.0 Revision -00	NetSight 6.0 release

PN: 9034976-01

Application Analytics Licensing

Application Analytics licensing allows deployment flexibility by granting flow rate and flow client capacity to an entire deployment, regardless of the number of Application Analytics engines. The NMS Advanced license (NMS-ADV) grants a basic flow rate and flow client capacity. Additional Application Analytics licenses can then be added to extend that capacity, if needed.

The table below shows the different Application Analytics licenses, and the flow rate capacity and flow client capacity granted by that license. The flow rate capacity is the number of flows per minute (FPM) that can be processed across all Application Analytics engines with a maximum of 3,000,000 FPM. The flow client capacity is the total number of application flow clients that can be reported system-wide with a maximum of 50,000 clients.

License Name	Flow Rate Capacity	Flow Client Capacity
NMS-ADV-XXX	3,000 FPM	100 clients
PV-FPM-50K	50,000 FPM	50,000 clients
PV-FPM-100K	100,000 FPM	50,000 clients
PV-FPM-500K	500,000 FPM	50,000 clients
PV-FPM-1M	1,000,000 FPM	50,000 clients
PV-FPM-3M	3,000,000 FPM	50,000 clients

Application Analytics 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 individual Application Analytics engine in the deployment. The capacity is checked on a continual basis. If the flow rate capacity is exceeded, an event is logged in the Application Analytics event log and a notification is displayed at the bottom of the Management Center screen. If the flow client capacity is exceeded, the client data for the clients beyond the capacity is not persisted.

NOTE: The **Max End-Systems** field in the Application Analytics Engine [Advanced Configuration](#) panel allows you to limit the clients from a single Application Analytics engine persisted in the Management Center database. This ensures that the 50,000 client limit is not collected from one engine.

Using Licenses to Establish Flow Rate Capacity

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

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

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

$$3,000 + 100,000 = 103,000 \text{ FPM}$$

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

$$3,000 + 100,000 + 50,000 = 153,000 \text{ FPM}$$

Using Licenses to Establish Flow Client Capacity

The NMS-ADV license provides a basic flow client capacity of 100. If you add additional Application Analytics licenses, the flow client capacity increases to 50,000, but does not increase beyond the system-wide maximum of 50,000 clients.

For example, if you add the PV-FPM-100K license, then you would have a total flow client capacity of 50,000. If you then add the PV-FPM-50K license, the flow client capacity still remains at 50,000.

Getting Started with Application Analytics

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

Application Analytics Access Requirements

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

In order to view the **Analytics** tab, you must be a member of an authorization group that has been assigned the Management Center Application Analytics Read Access or Read/Write Access capability. The Read Access capability allows the ability to access the **Analytics** tab and view the Application Analytics reports. The Read/Write capability adds the ability to configure Application Analytics engines and NetFlow Collecting devices. It also adds the ability to create and modify fingerprints. For additional information, see [How to Configure User Access to Extreme Management Center Applications](#).

Application Analytics Engine Configuration

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

The Application Analytics engine must be installed and running on your network. For instructions, see the [Application Analytics Engine Installation Guide](#).

Following installation, the Application Analytics engine must be added to Management Center and enforced via the Configuration view in the **Analytics** tab. For additional information, see [Configuration](#).

Enable NetFlow Collection

Because the **Analytics** tab displays reports based on NetFlow data, you must enable NetFlow for your network devices that act as the NetFlow sensors, and enable flow collection for their device interfaces. For additional information, see [How to Enable Flow Collection](#). You must also configure your NetFlow sensor devices to send their NetFlow information to the Application Analytics engine. In addition, the device interfaces you enable for flow collection must match the interfaces that are configured for analysis by the engine.

Configure Network Locations

In order to take full advantage of the reporting features in Application Analytics, it is recommended that you configure network locations. Defining network locations will provide additional client flow data in your Application Analytics reports as well as increase your options when specifying report search criteria.

A network location is a set of IP address ranges that identify a portion of your network. You can create a single network location that identifies which IP address ranges belong to the resources in your network or you can create multiple locations to identify different buildings, sites, or geographical areas of your network. Application Analytics uses the defined network locations to identify the portion of the network where the application flow client resides.

For additional information, see [Network Locations](#).

Related Information

- [Configuration - Analytics](#)
- [Network Locations](#)

Application Analytics Application Data Collection

The Application Analytics 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 Application Analytics engine. The engine collects NetFlow 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 Management Center database. Management Center uses this data to provide reports that show how your network is being utilized.

To conserve space on your Management Center server hard drive, your Application Analytics engines only collect total flow records when the server hard drive drops below 10 GB of free space. If the Management Center server hard drive drops an additional 1 GB (under 9 GB of free space), your Application Analytics 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 Application Analytics stops collecting all records when free space reaches $10\text{GB} - 1,000,000\text{ KB} = 9\text{ GB}$.

Collection Targets


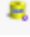






Flow data is collected on objects in your network called targets. Some targets are physical, such as clients and servers, and some are logical, such as applications.

An Application Analytics 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 Application Analytics, identified through layer 7 analysis (for example, Facebook).
- Application Group — Application categories, such as Cloud Computing or Social Networking.
- Location — The client's physical location on the network, based on its IP address. Network locations are used by Application Analytics to identify the physical location for the client of an application flow. For additional information, see [Using Locations to Collect In-Network Traffic](#).
- Device Family — The kind of device determined for a client, such as Windows or iOS.
- Profile — An Extreme Access Control 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 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.

Applications (Bytes) - 42.97 GB - Last hour				
Applications	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	VPN 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

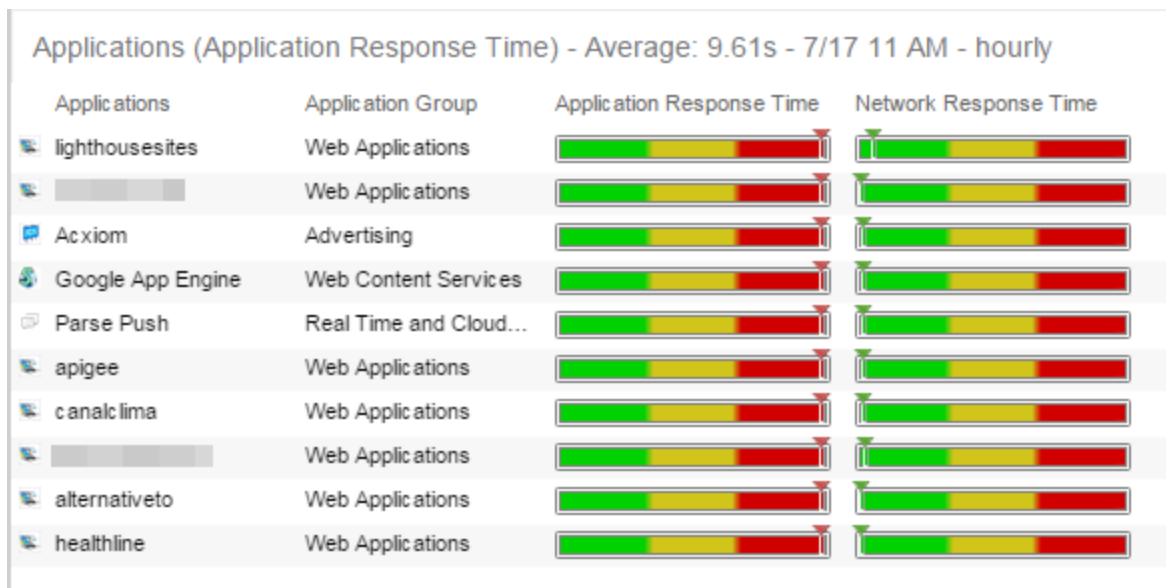
Collection statistics are quantitative data that can be collected for a target. This includes statistics directly reported in NetFlow records, such as bytes transferred, as well as information that can be derived indirectly, such as the number of unique clients seen using an application.

An Application Analytics 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.

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




Collection Intervals


The Application Analytics 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 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 Locations 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 Application Analytics locations. A location is a set of IP masks that defines a well-known portion of

your internal network. You can define a single location that identifies your entire internal network. If you have already reserved certain IP address ranges for certain physical locations on your network, you can create multiple network locations that correspond to these reserved IP ranges. Multiple locations can be created to identify different buildings, sites, or geographical areas of your network. Any IP that matches any location is considered to be in-network. If you define multiple locations, you will be able to analyze data broken down by location.

For additional information, see [Network Locations](#).

Data Collector Types

There are two kinds of data collectors used in Application Analytics.

- **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 location, 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 location.

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.

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

Target	Sub-Target	Bases	Traffic Used
Device Family		Bytes Flows Clients	In-Network Flows
Location		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
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

Target	Sub-Target	Bases	Traffic Used
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
Location		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 Application Analytics 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 Application Analytics license, the per-hour total limit is 100 clients across all Application Analytics engines.

Flow Information Sources

The Application Analytics engine uses NetFlow 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 Application Analytics 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 records and deep packet inspection.

NOTE: Most of these sources rely on Access Control data. If Access Control is part of your network configuration, then Access Control integration can be enabled (see [instructions](#) below) to provide access to these sources. Location data is obtained from network locations configured in Application Analytics. For additional information, see [Network Locations](#).

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

- Hostname — The client or server's hostname can be derived using Access Control. Access Control integration must be enabled.
- Location — The location for a flow is the location of the client in the flow. Client and server locations are derived from the network locations configured in Application Analytics. If a client does not match a location, then the location is empty. If a flow has a location, the flow is considered to be in-network. For additional information, see [Using Locations to Collect In-Network Traffic](#).
- Detailed Location — Detailed location information is derived from the switch and port information resolved for the client end-system. Access Control 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 Access Control, if Access Control integration is enabled.
- Profile — The client's profile is derived from the Access Control profile assigned to the client end-system. Access Control integration must be enabled.
- Username — The client's username is derived from network packet inspection. The username can also be provided by Access Control, if Access Control 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 Access Control may provide the device family name of Windows.

Enabling Extreme Access Control Integration

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

When Access Control integration is enabled, if a client in a flow matches an end-system in Access Control, 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 Access Control profile.
- The detailed location in the flow is derived from end-system data.

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

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

To enable Access Control integration on the Application Analytics engine:

1. If the Access Control distributed end-system cache is not enabled on the 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 Management Center server. Click **OK**.
2. Enable Access Control Integration on each Application Analytics engine where you want to use Access Control data.
 - a. Access the **Analytics** tab.
 - b. Expand each Application Analytics engine and select **Advanced Configuration**. In the right panel under **Configuration Options**, select the **Enable Extreme Access Control Integration** option.
 - c. If your Access Control engines are using **Communication Channels**, you must select the **Access Control Communication Channel** option and enter the channel name. The Application Analytics engine is only able to access end-systems in its channel.
 - d. Click **Save**.
 - e. Enforce your Application Analytics engines.

Reports

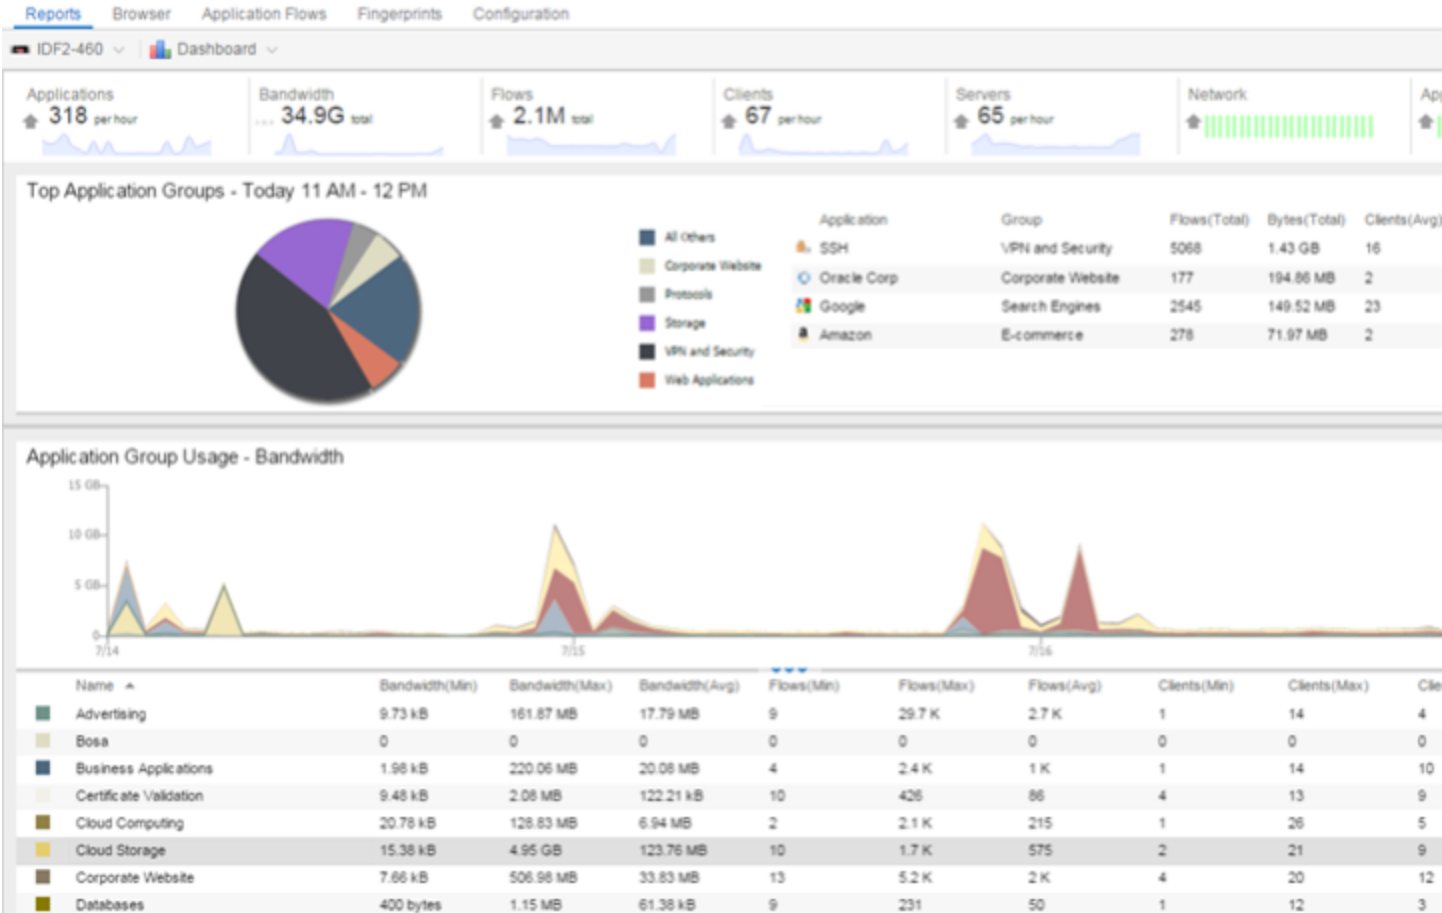
Data gathered from flow usage collection is the basis of many reports in the Management Center's **Analytics** tab. Once collection is enabled, these reports begin to exhibit data.

For additional information, see [Analytics](#).

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 Application Analytics engines is maintained separately. If you have more than one Application Analytics 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.

The screenshot shows a web interface for a network monitoring tool. On the left, there is a sidebar with 'Options' and 'Search Criteria' sections. The 'Options' section includes: Data Table (End-System Details - Hourly), Display Format (Grid), Target (Applications), Time Period (Last Interval), Statistic Type (Bytes), and Aggregation (Sum selected). The 'Search Criteria' section includes dropdowns for Location, Profile, Application Group, and Device Family, all set to 'All'. Below this is a 'Search Status' section indicating '416 rows evaluated successfully in 67 milliseconds'. The main content area is titled 'Applications (Bytes) - 8.39 GB - Last hour' and contains a table with the following data:

Applications	Application Group	Bytes	Sent Bytes	Received Bytes
EVault	Cloud Storage	2.32 GB	2.29 GB	28.40 MB
Netflow	Protocols	2.17 GB	1.09 GB	1.09 GB
Microsoft SQL Server	Databases	1.05 GB	520.24 MB	526.14 MB
MySQL	Databases	862.75 MB	432.89 MB	429.86 MB
MSRDP	Protocols	539.03 MB	262.13 MB	276.90 MB
Akamai	Web Content Services	467.93 MB	6.89 MB	461.05 MB
YouTube	Streaming	454.25 MB	11.34 MB	442.91 MB
NFL	Sports	303.25 MB	6.65 MB	296.60 MB
CIFS	Storage	135.62 MB	67.83 MB	67.79 MB
Pandora	Streaming	87.90 MB	1.32 MB	86.58 MB

Related Information

For information on related Application Analytics topics:

- [Getting Started with Application Analytics](#)
- [Analytics](#)
- [Network Locations](#)

Add and Modify Fingerprints

Application Analytics 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 Application Analytics feature. In addition, you can modify these fingerprints and create new custom fingerprints.

For additional information, see [Getting Started with Application Analytics](#).

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 Management Center Application Analytics Read/Write Access capability. For additional information, see [How to Configure User Access to Extreme Management Center Applications](#).

Adding a Fingerprint

Use the following steps to add a new custom fingerprint based on an existing flow in the Applications Flows view. For additional information, see [Custom Fingerprint Examples](#).

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

Flows	Client Address	Server Address	Server Port	Application	Application Group	Application
296	btownsen-ws1		https	Force	Cloud Computing	IssuerIdAtC
30017	10.6.24.41		ldap	LDAP	Protocols	SwitchType
2	rdow-pc		https	Outlook Office365	Mail	IssuerIdAtC
95	dduggan-ubuntu		http	Imgur	Social Networking	URI=/lumb

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.

Add Fingerprint

Create a fingerprint matching the following components of this flow.

Port snmp [161]

Application Name:

Application Group:

Confidence: 60

Description:

This fingerprint needs to be enforced to appliances before it can take effect.

OK **Cancel**

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 menu 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.
10. Click **Save**. The new fingerprint is created on the Management Center server.
11. Enforce to push the new fingerprint to your engines.

TIP: You can also create a custom fingerprint from the [Fingerprints view](#). From the Gear menu, 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.

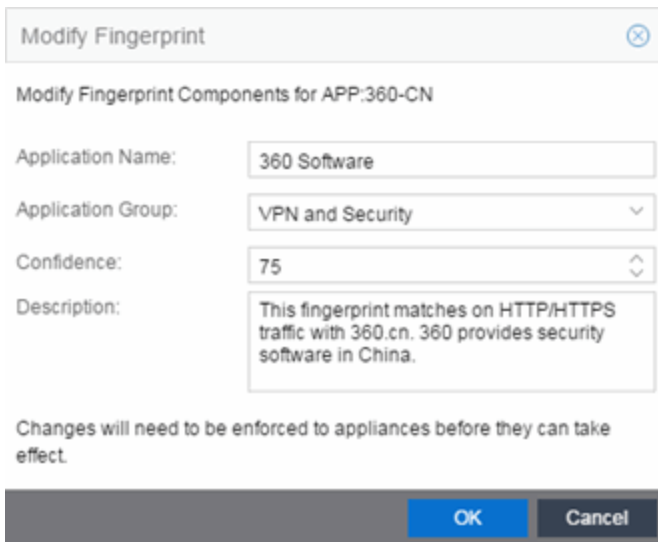
Modifying a Fingerprint

Modify a fingerprint's application name, application group, confidence level, and description from the [Fingerprints view](#).

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



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.

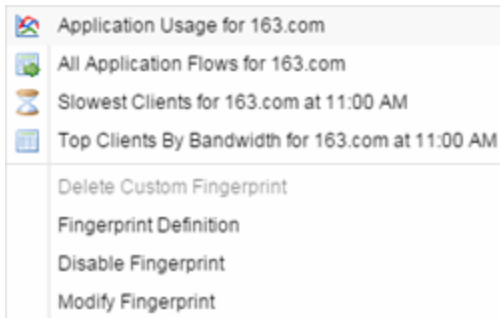
4. Click **OK**.

5. Enforce to push the change to your engines.

Enabling or Disabling a Fingerprint

Enable or disable a fingerprint from the [Fingerprints view](#). When a fingerprint is enabled, it is used to identify applications. When it is disabled, it is ignored.

1. Select the **Analytics** tab and then select the Fingerprints view.
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 view](#). 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 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.
5. 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 view](#). Custom fingerprints are not overwritten when an update is performed.

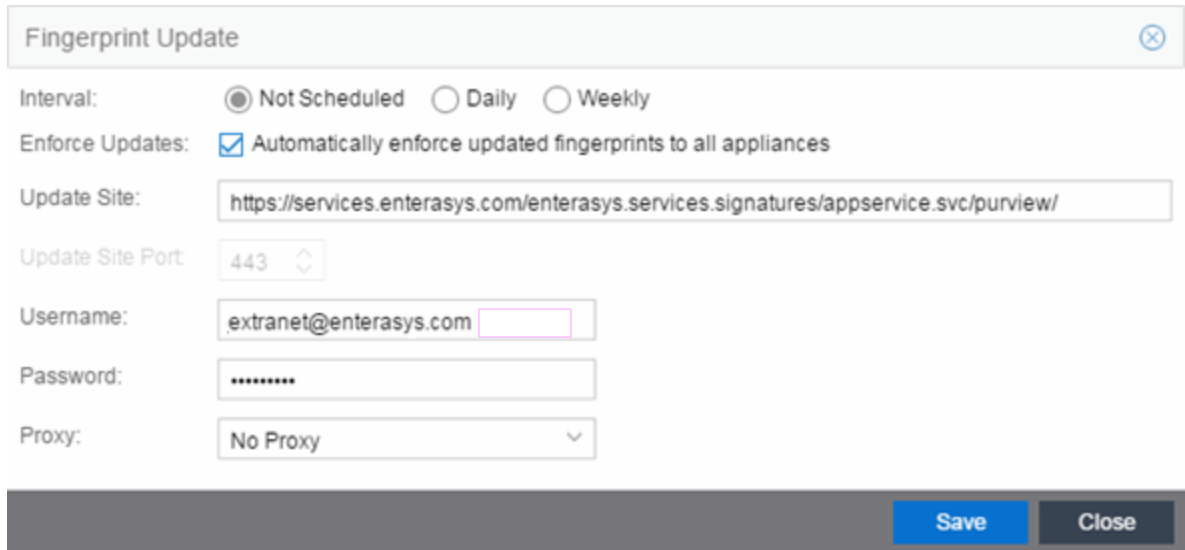
When a fingerprint update is performed, the fingerprint update server is checked for newer fingerprints than what is available on the Management Center server. If there are newer fingerprints, they are downloaded, and the fingerprint definitions are updated with any new fingerprint definition files. 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 Management Center and then select the **Configurations** view.
2. In the left-panel tree, expand the System folder and select **Fingerprints**.
3. Click the gear menu  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.



The screenshot shows a configuration window titled "Fingerprint Update" with a close button in the top right corner. The window contains the following fields and controls:

- Interval:** Three radio buttons: "Not Scheduled" (selected), "Daily", and "Weekly".
- Enforce Updates:** A checked checkbox labeled "Automatically enforce updated fingerprints to all appliances".
- Update Site:** A text input field containing the URL "https://services.enterasys.com/enterasys.services.signatures/appservice.svc/purview/".
- Update Site Port:** A spinner control set to "443".
- Username:** A text input field containing "extranet@enterasys.com" and a small empty input field to its right.
- Password:** A text input field filled with "*****".
- Proxy:** A dropdown menu currently set to "No Proxy".

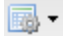
At the bottom right of the window, there are two buttons: "Save" (in a blue box) and "Close" (in a grey box).

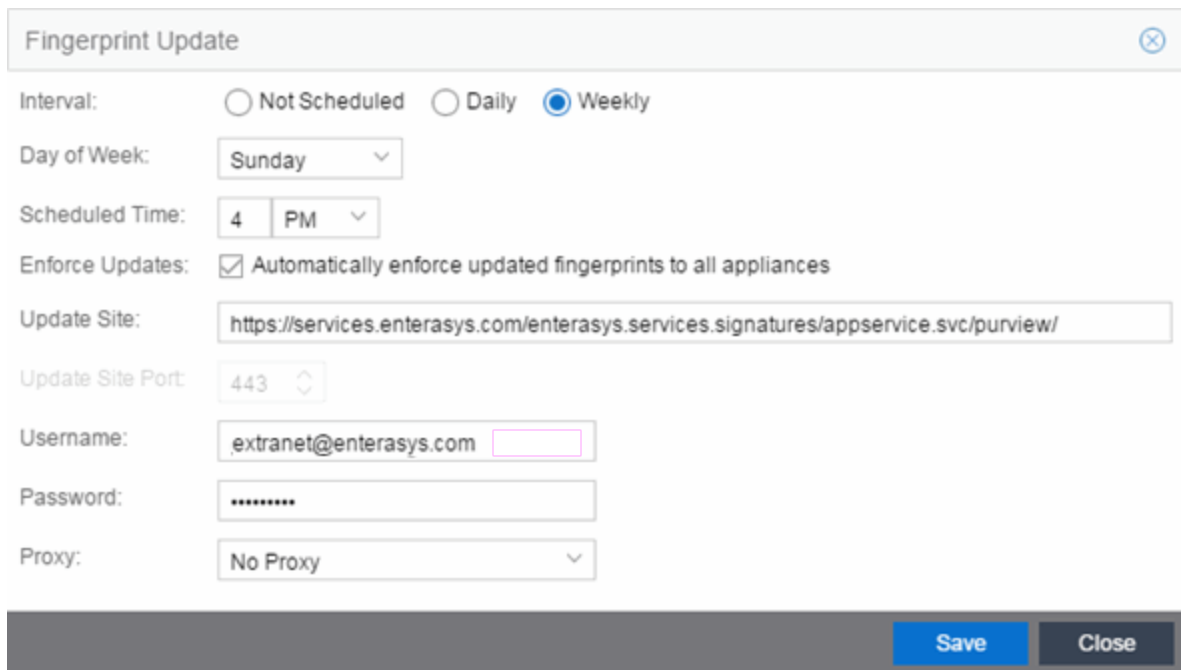
- a. Leave the **Interval** selection as **Not Scheduled**.
 - b. Select the **Enforce Updates** checkbox to automatically update fingerprints on all engines. Not selecting this checkbox requires you to update each engine manually.
 - c. The **Update Site** field displays the default path to the official 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 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 gear menu  and select Fingerprint Update Settings. The Fingerprint Update window opens.



Fingerprint Update

Interval: Not Scheduled Daily Weekly

Day of Week: Sunday

Scheduled Time: 4 PM

Enforce Updates: Automatically enforce updated fingerprints to all appliances

Update Site: <https://services.enterasys.com/enterasys.services.signatures/appservice.svc/purview/>

Update Site Port: 443

Username: extranet@enterasys.com

Password:

Proxy: No Proxy

Save Close

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.

7. 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.
 10. 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 Application Analytics topics:

- [Analytics](#)
- [Custom Fingerprint Examples](#)

Applications Browser

The Applications Browser 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 Management Center [Analytics tab](#).


Viewing Application Analytics application data requires certain prerequisites. For additional information, see [Getting Started with Application Analytics](#).

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 location.

If you have multiple Application Analytics engines, use the **Engine** drop-down menu 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 [bookmark the report](#), [save it to the Report Designer](#) to use as a custom component, or [export it as a CSV file](#).

The screenshot displays a network management interface with a navigation menu at the top containing: Network, Alarms and Events, Control, Analytics (highlighted), Wireless, Reports, and Administration. Below the menu is a dashboard with tabs for Dashboard, Browser, Application Flows, Fingerprints, Configuration, and Reports. The main content area is titled "NOC Sensor" and contains several sections:

- Options:** Data Table (End-System Details - Hourly), Display Format (Grid), Target (Applications), Time Period (Last Interval).
- Statistic:** Type (Bytes), Aggregation (Sum selected, Average unselected).
- Search Criteria:** Location (All), Profile (All), Application Group (All), Device Family (All), User Name, Application, Client, Limit (10).
- Applications (Bytes) - 159.14 GB - Last hour:** A table showing application data.

Applications	Application Group	Bytes	Sent Bytes	Received Bytes
ISCSI	Storage	89.21 GB	44.61 GB	44.61 GB
Encrypted Web	Web Applications	15.45 GB	8.51 GB	6.94 GB
CIFS	Storage	11.46 GB	4.24 GB	7.22 GB
Microsoft SQL Server	Databases	8.16 GB	4.12 GB	4.04 GB
Extreme Networks	Corporate Website	7.73 GB	4.68 GB	3.05 GB
YouTube	Streaming	4.96 GB	1.37 GB	3.61 GB
Google	Search Engines	4.46 GB	3.76 GB	694 MB
Outlook Office365	Mail	4.26 GB	597.48 MB	3.67 GB

At the bottom of the interface, there is a status bar showing "Administrator", "Last Updated: 2/4/2016 1:04:15 AM", "Uptime: 5 Days 09:06:18", and "Alarms: 2 35 18 18".

Data Aggregation

Network data displayed in a report is aggregated from your network by the Application Analytics engine and sent to Management Center. The data gathering process begins with the Application Analytics engine, which monitors network activity on the switch or controller you configure using a traffic mirror and NetFlow. 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 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 Application Analytics engine to Management Center based on the criteria you select.

NOTE: Information held in the Application Analytics engine's cache is not saved. Restarting the Application Analytics engine before the data in the memory cache is sent to 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 location.
- **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.

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 Application Analytics 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/Profile** — Information about applications used by profiles, or about profiles 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.
- **Locations** — Network locations are used by Application Analytics to identify the physical location for the client of an application flow. A network location is a set of IP address ranges that identify a portion of your network. Multiple locations can be created to identify different buildings, sites, or geographical areas of your network. For additional information, see [Network Locations](#).
- **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.
- **Sent Bytes** — The number of bytes sent by clients.
- **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 than 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 Locations as your target, you can only filter on defined locations. If you select the End-System Details data table, you can filter on additional criteria. For example, if you select Locations as your target, you can filter on defined locations 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.

- **Location** — Select a network location to match or select All. If a location has been added to a map, you will also see a selection for that map. If you select custom, you can enter a partial location name or use the SQL wildcard characters to match one or more locations. For additional information, see [Network Locations](#).
- **Profile** — Select an Extreme Access Control 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.
- **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 the Report

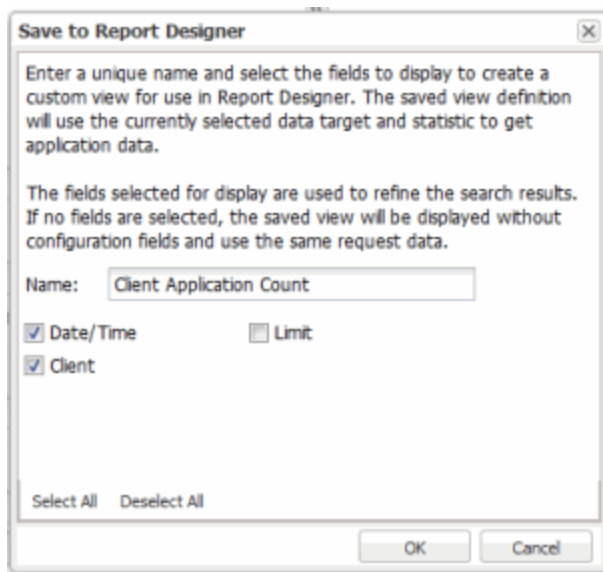
After you have generated a report, click the Gear menu  in the lower left corner to save the options you have currently set. A new window opens for the current

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 to export the report data as a CSV file. The currently defined report opens in a spreadsheet, which can then be saved.

Related Information

For information on related Application Analytics topics:

- [Analytics](#)
- [Getting Started with Application Analytics](#)
- [Network Locations](#)

Custom Fingerprint Examples

The Application Analytics 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 Application Analytics feature. In addition, you can create new custom fingerprints.

For additional information, see [Getting Started with Application Analytics](#).

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

4. The Add Fingerprint window opens.

Add Fingerprint

Create a fingerprint matching the following components of this flow.

Port x11 [6049]

Application Name: X Windows System

Application Group: Protocols

Confidence: 60

Description: X Windows System Network traffic

This fingerprint needs to be enforced to appliances before it can take effect.

OK Cancel

5. Use the drop-down list to select matching Portx11 [6049].
6. Set the **Application Name** to X Window System.
7. Set the **Application Group** to Protocols.
8. 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.
9. Click **OK** to create the fingerprint.
10. 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 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 Management Center.
2. In the Application Flows table, select the **Show Unclassified Web Traffic View**.
3. Right-click on a flow with the yahoo ads application and select **Fingerprints > Add Fingerprint**.
4. The Add Fingerprint window opens.

Add Fingerprint

Create a fingerprint matching the following components of this flow.

Host yahoo ads

Application Name: Yahoo Ads

Application Group: Advertising

Confidence: 60

Description:

This fingerprint needs to be enforced to appliances before it can take effect.

OK Cancel

5. Use the drop-down menu to select matching the "yahoo ads" host.
6. Set the **Application Name** to **Yahoo Ads**.
7. Set the **Application Group** to **Advertising**.
8. 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.
9. Click **OK** to create the fingerprint.
10. 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.

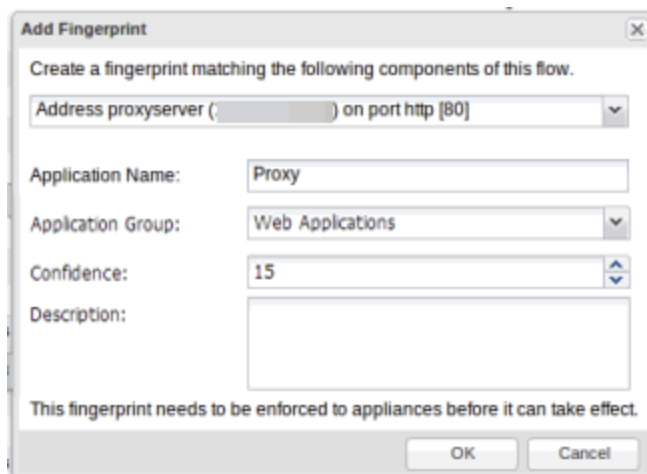
For example, let's say you have a Git repository on your network. Git repositories (a source code management system used in software development) are frequently accessed via SSH on port 22 (the standard TCP port assigned for SSH traffic). In this case, the SSH traffic flows is identified using the system 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 Management Center.
2. In the Application Flows table, right-click on an SSH port-based flow with the Git server destination address and select **Fingerprints > Add Fingerprint**.
3. The Add Fingerprint window opens.



4. Use the drop-down menu to select matching the Git server IP address and port.

5. Set the **Application Name** to **Git**.
 6. 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.)
 7. Set the **Confidence** level to **60**, which is a higher confidence than the current fingerprint which is set at 10.
 8. Click **OK** to create the fingerprint.
 9. Enforce to push the new fingerprint to your engines.
-

Related Information

For information on related Application Analytics topics:

- [Analytics](#)
- [Add and Modify Fingerprints](#)

How to Deploy Application Analytics in an MSP or MSSP Environment

This Help topic presents instructions for deploying Application Analytics 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 Management Center server. This allows the Management Center server to convert the NAT IP address received in the Application Analytics engine response to the real IP address used by the Management Center server. Not adding the real IP address for the Management Center server to the `nat_config.txt` file results in the Application Analytics engine incorrectly displaying a state of **IMPARED** (orange) rather than **UP** (green).

NOTE: The text in the `nat_config.txt` 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 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 Management Center Management server IP address is not configured to use the NAT IP address of the 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 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 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 Management Center server.

NOTE: The Application Analytics 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 firewall 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 correct 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.
6. Add the Management Center server hostname to your DNS server, if necessary.

NOTE: Application Analytics engines, remote Management Center clients, and any Extreme Access Control engines must be able to connect to Management Center using this hostname.

Related Information

For information on related windows:

- [Application Analytics Engine Advanced Configuration Panel](#)

Network Locations

In order to take full advantage of the reporting features in Application Analytics, you must first configure network locations. Defining network locations identify IP ranges for certain end-systems in your network, provides client flow data in your Application Analytics reports, and provides additional options for working with report search criteria. Locations can be imported or exported as a CSV formatted file.

Configuring network locations can be useful if you have already reserved certain IP address ranges for certain physical locations on your network. You can create network locations that correspond to these reserved IP ranges. The network locations are then used to identify the portion of the network where the application flow source resides by matching the client's IP address to the ranges included in each network location. Multiple locations can be created to identify different buildings, sites, or geographical areas of your network. Even if you have no such policies, you can create a single network location that identifies which IP address ranges belong to resources in your network.

A location is defined with a name, description, and one or more IP address ranges specified by an IP address/mask. When a client's IP address matches any IP address/mask in a location, the client is determined to have that location. If the client matches the address/mask of several locations, the location with the most specific mask (the highest CIDR value) is used.

The name of the network location that matches the client's IP address is listed in the Location column of the Application Flows table in the [Analytics tab](#). This allows you to search, sort, and filter flow data according to location. Application Analytics uses this data to provide a summary of the data for locations, which can be viewed as either an hourly or high-rate report in the [Analytics Browser](#).

You must be a member of an authorization group that has been assigned the Extreme Management Center Application Analytics Read/Write Access capability in order to manage network locations. For additional information, see [Getting Started with Application Analytics](#).

This Help topic provides the following information about managing Management Center network locations:

- [Adding Locations](#)
- [Editing Locations](#)
- [Removing Locations](#)


- [Importing Locations](#)
- [Exporting Locations](#)
- [Searching Locations](#)

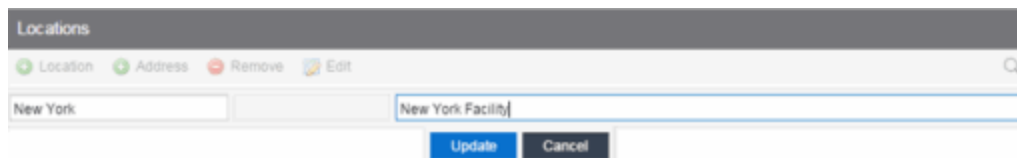
Managing Locations

Locations are created and managed in the Configuration view of the **Analytics** tab.

Adding Locations

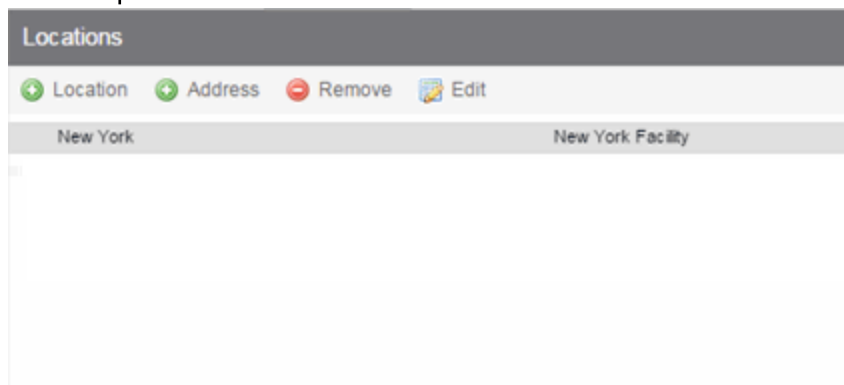
To add a location:

1. Access the **Analytics** tab and select the **Configuration** view.
2. In the left-panel tree, expand System and select **Locations**.
3. In the right-panel Locations view, click the **Add Location** button ( Location) and enter a name for the new location in the first text box. If desired, enter a description for the location in the second text box.





The screenshot shows the 'Locations' configuration interface. At the top, there are buttons for '+ Location', '+ Address', '- Remove', and 'Edit'. Below these are two text input fields. The first field contains the text 'New York' and the second field contains 'New York Facility'. At the bottom of the form are two buttons: 'Update' and 'Cancel'.

4. Click **Update**. The new location is added.




The screenshot shows the 'Locations' configuration interface after the 'Update' button was clicked. The 'New York' location is now listed in the table below the form. The table has two columns: 'Name' and 'Address'. The first row contains 'New York' and 'New York Facility'. Below the table is a large empty text area.

5. Make sure that the location is selected in the list. Click the **Add Address** button ( Address) and enter an IP address/mask in the field in CIDR notation.




6. Click **Update**. The IP address/mask is added to the list under the location. Repeat steps 5 and 6 to add as many IP addresses/masks as necessary. The following image shows the Locations panel with multiple locations defined.



Editing Locations

To edit a location name, description, or address/mask:

1. Access the Management Center **Analytics** tab and select the Configuration view.
2. In the left-panel tree, expand System and select **Locations**.
3. In the right-panel Locations view, select the location name, description, or address/mask that you want to edit.
4. Click the **Edit** button  and make the desired changes to the location name, description, or address/mask. You can also double-click on a name,

description, or address/mask to make the desired changes, instead of using the **Edit** button.

5. Click **Update**.

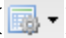
Removing Locations

To remove a location or address/mask:

1. Access the **Analytics** tab and select the Configuration view.
2. In the left-panel tree, expand System and select **Locations**.
3. In the right-panel Locations view, select the location name or address/mask that you want to remove.
4. Click the **Remove** button.

Importing Locations


To import locations from a CSV file:

1. Access the Management Center **Analytics** tab and select the Configuration view.
2. In the left-panel tree, expand System and select **Locations**.
3. In the right-panel Locations view, click the **Gear** button () and select **Import from CSV**.
The Import Locations window appears.
4. Click the **Select File** button and navigate to the folder in which the CSV file is located.
5. Select the CSV file and click **Open**.
6. Click one of the following options in the Import Options section of the window to determine how Management Center handles existing locations:
 - a. Select **Discard all locations and import new ones** to replace all locations currently listed in the Locations view with the locations imported from the CSV file.
 - b. Select **Import locations, overwriting existing locations** to replace existing locations currently listed in the Locations view with locations imported from the CSV file, but leave all other locations in the Locations view unchanged.

- c. Select **Import locations, but do not change existing locations** to add new locations to the Locations view, but prevent existing locations currently listed in the Locations view from being overwritten by locations imported in the CSV file.
7. Click **Import**.
The locations are imported to the Management Center.

Exporting Locations

To export locations and save them as a CSV file:

1. Access the Management Center **Analytics** tab and select the Configuration view.
2. In the left-panel tree, expand System and select **Locations**.
3. In the right-panel Locations view, click the **Gear** button () and select **Export to CSV**.
The CSV file is saved to the web browsers default download location.

Searching Locations

To search for a specific location or address/mask in the Locations view, enter the location name or address/mask in the **Search** field and press **Enter**. The search results are displayed in the view.

Related Information

For information on related Application Analytics topics:

- [Getting Started with Application Analytics](#)
- [Analytics](#)

Analytics

The **Analytics** tab lets you view and customize Application Analytics reports and application flow data, as well as manage and configure your Application Analytics engines. Additionally, the [Menu at the top of the screen](#) provides links to additional information about your version of Extreme Management Center.

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

Viewing Application Analytics application data requires certain prerequisites. For additional information, see [Getting Started with Application Analytics](#).

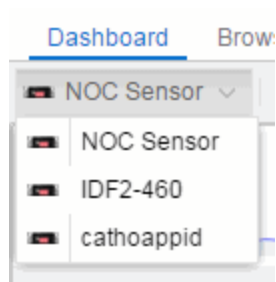
This Help topic provides information on the different reports available from the **Analytics** tab, as well as engine configuration information.

- [Dashboard](#)
 - [Graph Descriptions](#)
- [Browser](#)
- [Application Flows](#)
 - [Bidirectional Flows](#)
 - [Unidirectional Flows](#)
 - [Report Features](#)
- [Fingerprints](#)
 - [Fingerprint Table](#)
- [Configuration](#)
 - [Adding an Engine](#)
 - [Enforcing an Engine](#)
 - [Engine Administrative Options and Reports](#)
- [Reports](#)
 - [Report Descriptions](#)


Dashboard

The **Dashboard** view 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. For many of the graphs, you can click on an item to view details.

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



Then use the **Report** drop-down menu to the right to access the different reports.

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 location. 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, alternate values should be used.

Graph Descriptions


This section provides a description for each of the available graphs.


Overview

The Dashboard report contains an info bar and two summary graphs for top applications and application group usage. The info bar provides a selection of sparkline graphs showing different network statistics for the last 24 hours, with arrows that indicate trends compared to the previous reporting period.

- The **per hour** statistics show the average unique applications, clients, or servers seen per hour for the past 24 hours.
- The **total** statistics show the total bandwidth or flows reported for the past 24 hours.

The Network and Application graphs show the average reported response times for the last 24 hours. Rest your cursor on the graph to see a tooltip showing the response time including the time and date of the data sample. Clicking on a graph for which historical data is available displays the available historical data for the category, with options to change the reporting time period displayed.

The Top Application Groups report displays the top five applications for the last hour. Use the **Gear** button  to change the start time for the report and whether the data is displayed as a pie chart, word cloud, tree map or bubble map. If you change the reporting start time, the data in the Dashboard info bar changes accordingly. Clicking an application link in the table to the right displays the list of clients using that application. Right-click on a client to open various reports, launch PortView, or search Management Center maps.

The Application Group Usage graph provides a longer view of application usage. Use the **Gear** button  to change the start date and time and the number of days of data to display. You can also select whether to display bandwidth, flow, or client data in the graph. Use the arrows at the ends of the graph to quickly change the reporting period displayed. The table below the graph presents the individual bandwidth, flow, and client statistics for each group.

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** button  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 locations 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.

High-Rate Application Collector Dashboard Report

The High-Rate Application Collector Dashboard shows the number of clients, flows and bytes collected during the high-rate collection interval for the time period configured at the top of each section.

Click on the **Info** button  at the top right of the dashboard page to read a description of each report.

Industry Dashboards

- 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.
- The Education Dashboard displays application information specific to the campus network including learning management systems, P2P, streaming, and social applications.
- The Healthcare Dashboard focuses on applications used in the healthcare environment including patient care, medical applications, and HIPAA.
- 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 Application Analytics 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 Application Analytics 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.

Application Map

The Application Map provides a global overview of top application groups by location, displayed in the **Network** tab World map. The application data is displayed in pie charts and is based on application data for Application Analytics locations linked to the **Network** tab map.

For information on configuring the Management Center World map to show application data, see Show Application Data in the Advanced Map Features section of the *Extreme Management Center User Guide*.

NOTE: By default, the Application Map displays the Management Center World map. You can specify a different map to use by changing the Application Dashboard Map option. On the **Configuration** tab, select the **System > Advanced** options in the left-panel and select a new Application Dashboard Map in the right-panel.

Response Time Dashboard

The Response Time Dashboards present the response time in milliseconds of application data grouped by different criteria, selected from the dropdown menu. The data is displayed as a line graph, which is updated periodically. For additional information, see [Response Time Dashboard](#).

Network Service Dashboard

The Network Service Dashboard displays the response time of network services for the top five worst-performing locations as well as the overall average of all locations. The data for each network service at a location is displayed as a bar and line graph, which is updated periodically. For additional information, see [Network Service Dashboard](#).


Browser

The Applications Browser lets you query information about recent network activity stored in the 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. For additional information, see [Applications Browser](#).

Application Flows

The Application Flows table presents bidirectional flow data (aggregate flows) or unidirectional flow data (base flows).

If you have multiple Application Analytics engines, use the **Engine** menu to select an engine to use as the source for the flow data. Use the **Type** menu to select whether to display bidirectional or unidirectional flow data.



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 TCP 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.

- Top N — These reports provide aggregated flow data for individual applications, clients, or servers, with results sorted based on bandwidth, number of flows, number of packets, or number of connections.
- Show All — Show all flows.
- Show Classified — Show only flows that have been classified by an application fingerprint.
- Show Unclassified — Show only flows that have not been classified by an application fingerprint.
- Show Unclassified Web Traffic — Show only web traffic that has not been classified by an application fingerprint.

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

Use the **Search** field at the top right of the table to search for a specific application, user name, or IP address. From the filtered search results, click a user name or IP address to launch PortView, which provides a detailed topology context for the user. Entering **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".

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. For additional information, see [Create Policy Rule](#).
- Search 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.

Bidirectional Flows



This table displays bidirectional flow data that is 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.

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 how long the data has been collected, using X number of days, hh:mm:ss format.

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 **Gear** menu  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 Application Analytics 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 Application Analytics 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.

Location

The name of the network location that matches the client's IP address. For additional information, see [Network Locations](#).

Detailed Location

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 Access Control 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 considered suspicious for a variety of reasons. For additional information, see [IP Reputation Dashboard](#).

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.

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.

Rx Packets

The number of packets received for this flow.

Tx Bytes

The number of bytes transmitted for this flow.

Rx Bytes

The number of bytes received for this flow.

NetFlow Records

The number of NetFlow records received in each flow.

Flow Source

The IP address of the NetFlow 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.

Unidirectional Flows



This 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.

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 view shows how long the data has been collected, using X number of days, hh:mm:ss format.

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 Application Analytics engine. Mouse 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 Application Analytics 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 Application Analytics 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.

Location

The network location where the flow originated. For additional information, see [Network Locations](#).

Detailed Location

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 Access Control 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.

Bytes

The number of bytes in this flow.

NetFlow Records

The number of NetFlow records for this flow.

Flow Source

The IP address of the NetFlow 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.

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.

Report Features

The Application Flows table (bidirectional and unidirectional) includes the following features:

Search

The **Search** field can be used 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 (click 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.

Refresh Interval

Use the **Refresh** drop-down menu 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.

Create Policy Rule

Right-click on a flow in the table and select **Create Policy Rule** to open the Create Policy Rule window, which allows you to create a UDP or TCP rule using the IP port. You can also enter a **Rule Name**, if applicable. In the Policy Manager domain that you select, two services are created, each with their own rule: one that is server-based and one that is client-based. For example, for an SNMP flow, the following two rules would be created:

- Client Traffic - To Server Port: snmp[161]
- Server Traffic - From Server Port: snmp[161]

Optionally, the IP address of the flow can be used when creating the rule, which would add the IP address to the rule name, for example:

- Client Traffic - To Server Port: snmp[161](10.20.30.131)
- Server Traffic - From Server Port: snmp[161](10.20.30.131)

These are simplified rules that have no associated action and are not added to any roles. You must use Policy Manager to configure actions for the rules and assign them to the appropriate role.

Interactive Tables

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

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

Bookmark Report

Use the **Bookmark** button 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.

CSV Export

Save report data to a CSV file to provide report data in table form.

Fingerprints

The **Fingerprints** view provides detailed information about fingerprints used by Application Analytics to identify application flows. A fingerprint is a description of a pattern of network traffic which can be used to identify an application. 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 Management Center server. When a fingerprint is changed, a flag is raised on the Application Analytics engine to show it needs enforcing.

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

System fingerprints are provided by 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.

For additional information, see [Add and Modify Fingerprints](#).

The **Fingerprints** view is divided into a left-panel tree and a table. The left-panel tree displays all the application groups 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 is a **Gear** menu , where you can access various system and fingerprint actions. See below for a description of the menu options.

If you have multiple Application Analytics engines, an **Engine** menu is available that allows you to select an engine to use as the source for the fingerprint [Hits](#) and [Matches](#) 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.

Gear Menu

Use the **Gear** menu  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 — For additional information, see [Creating a Fingerprint](#).
- Delete Custom Fingerprint — For additional information, see [Deleting a Custom Fingerprint](#).
- Fingerprint Definition — View the XML definition for a fingerprint.
- Enable/Disable Fingerprint — Enable or disable a fingerprint. When a fingerprint is enabled, it will be used to identify applications. When it is disabled, it will be ignored.
- Modify Fingerprint — Change a fingerprint's description. For additional information, see [Modifying a Fingerprint](#).
- Reset Fingerprint Counters — Reset the Hits and Matches counters.

Column Definitions

Following are definitions for the table columns:

Application Name

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

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 ✓ 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.

Hits

The total number of times a hit has been recorded for this fingerprint for the selected engine. A hit is an occurrence of the Application Analytics engine matching a fingerprint in a flow. It may refine the application detected with other fingerprint hits on the same flow. This column is not displayed by default. You must display this column by clicking on a column heading drop-down arrow and selecting the Hits column option from the menu. See Notes below.

Matches

The total number of times a traffic flow has matched this fingerprint for the selected engine. A match is an occurrence of the Application Analytics 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:**
- Hits and 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.
 - A single hit, applied to one direction, may result in two matches, one in each direction.
 - If you need to reset the Hits and Matches counters, use the **Reset Fingerprint Counters** option from the **Gear** menu .
-

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 ✓ 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.

Description

Description of the fingerprint.

Last Modified

Date that the fingerprint was last modified.

Created


Date that the fingerprint was created.

Configuration

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

Adding an Engine

Use the following steps to add a Application Analytics engine to Management Center.

1. Select the **Analytics** tab in the Management Center and then select the **Configuration** view.
2. Select **Overview** in the left-panel tree.
3. Click on the **Gear** menu  and select **Add Engine**.
4. Enter the IP address of the management interface of the engine and a name for the engine.




The engine is added to Management Center if it does not already exist.

5. Select the SNMPv3 profile to use for the engine.
6. Click **OK**. The engine is added to the engine list.
7. Enforce the engine.


Enforcing an Engine

You need to enforce an engine whenever there are any changes made in Management Center that need to be sent to the Application Analytics engine. This includes changing system settings, changing engine settings, and changing fingerprints.

Use the following steps to enforce a Application Analytics engine.

1. Select the **Analytics** tab and then select the **Configuration** view.
2. Select **Overview** in the left-panel tree to display a list of configured engines. The orange Enforce icon  is displayed above an engine that needs to be enforced.
3. Hover the mouse over the engine that needs to be enforced. Click on the yellow Enforce icon  to the right of the engine to enforce the engine.
4. To enforce all engines, click on the **Gear** menu  and select **Enforce All Engines**.


The orange Enforce icon is also displayed in the **Applications** view status bar along with the number of engines that need to be enforced. Mouse over the icon to see a tooltip that lists the engines that need to be enforced. Click the icon to enforce the engines.



Engine Administrative Options and Reports

Use the left panel in the Configuration view to access various engine administrative options and reports.

Overview

View a list of configured engines and their engine statistics. Access the following options from the **Gear** menu . For some of the options, you must first select an engine in the list.

- Add Engine — Adds a new Application Analytics engine to Management Center.
- Delete Engine — Delete the selected engine.
- Enforce Engine — Enforce the selected engine.
- Poll Engine — Poll the selected engine.
- Restart Collector Process — Restarts the Application Analytics engine's collector process.
- Enforce All Engines — Enforces all of the Application Analytics engines added to Management Center.


Application Analytics Engines

View engine status information, configure web credentials, and configure advanced options for an individual engine. Selecting the engine name opens

- Status — View engine status including flow collector, application sensor, CPU and memory, flow sources, and diagnostic information. Click on **Help Tips** to read a description of the various reports.
- Web Credentials — Configure web credentials for an engine.
- Advanced Configuration — For additional information, see [Application Analytics Engine Advanced Configuration](#).
 - Set privacy levels.
 - Enable Access Control Integration.
 - Add advanced configuration properties.
 - Enable sensor modules and sensor module logging.

Application Analytics System


Configure and manage components of the Application Analytics system.

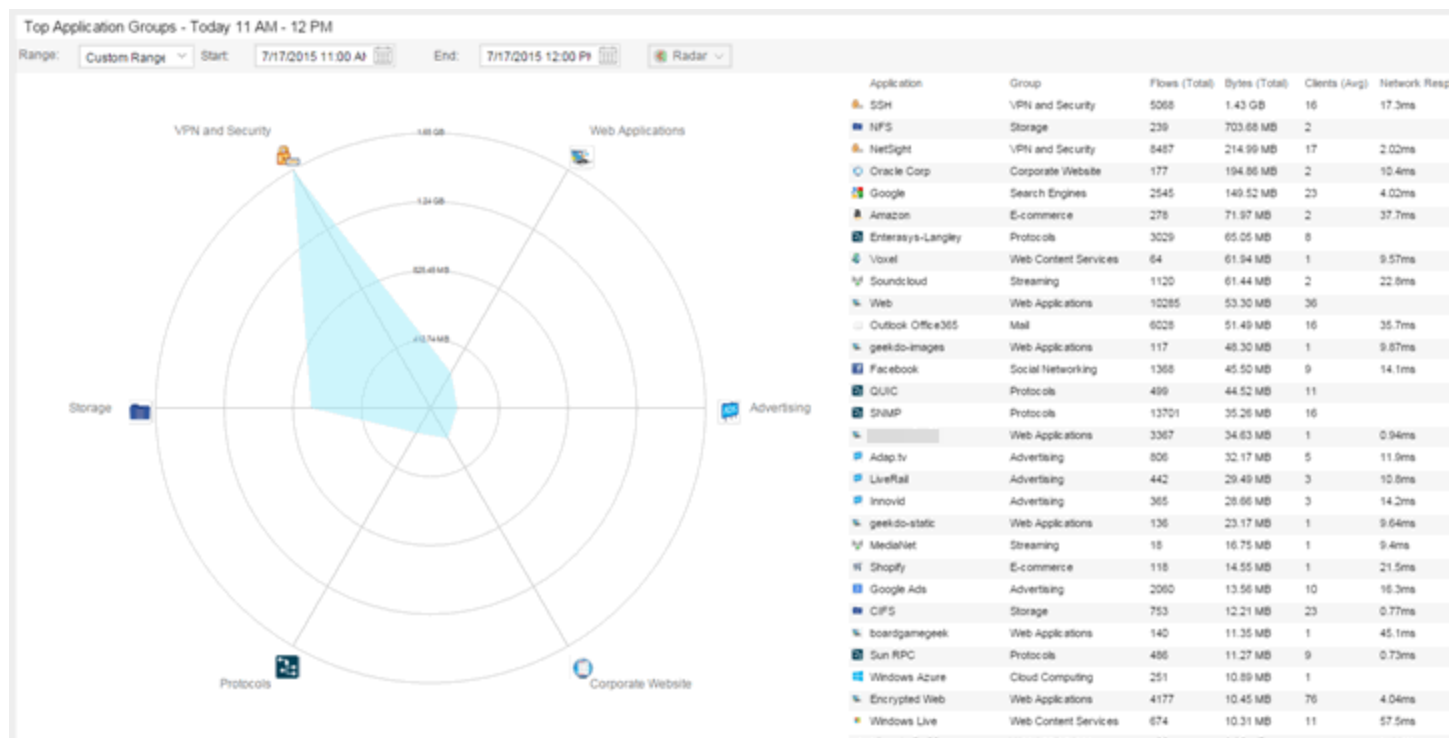
- Locations — Configure and manage network locations. For additional information, see [Network Locations](#).
- Fingerprints — View a summary of the kinds of application fingerprints in use. Use the **Gear** menu  to access the following system fingerprint actions:
 - Update Fingerprints — Perform a manual one-time update of the fingerprint database. For additional information, see [Updating Fingerprints](#).
 - Fingerprint Update Settings — Schedule fingerprint updates to be performed automatically on a daily or weekly basis. For additional information, see [Updating Fingerprints](#).
- Licenses — Add an engine flow rate increase license. Click on Help Tips to read a description of the various sections.
- Advanced — Configure global options for the Application Analytics system.
- Status — View a collection of Application Analytics system statistics.

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 location. 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 Application Analytics engines, use the **Engine** drop-down menu to select an engine to use as the source for the report data. Then use the Report drop-down menu to the right to access the different reports.

In most of the reports, you can 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 location. 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, alternate values should be used.

Report Descriptions

This section provides a description for each of the available reports.

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.

Locations Using the Most Bandwidth

This report displays the network locations 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.

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 Location

This report displays network traffic statistics for each network location.

Network Activity for a 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 for an 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 Location

This report displays the applications with the highest application response times, for the specified location. Select a network location to match or select All and then click the **Submit** button to generate the report. If a location has been added to a map, you also see a selection for that map. If you select custom, you can enter a partial location name or use the SQL wildcard characters to match one or more locations. For additional information, see [Network Locations](#).

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 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:

- **Start** — Select the start date and time.
- **Top N** — Select the number of clients displayed in the chart.
- **# 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 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:

- **Start** — Select the start date and time.
- **Top N** — Select the number of clients displayed in the chart.
- **# 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 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:

- **Start** — Select the start date and time.
 - **Top N** — Select the number of clients displayed in the chart.
 - **# 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

For information on related Application Analytics topics:

- [Getting Started with Application Analytics](#)
- [Add and Modify Fingerprints](#)
- [Custom Fingerprint Examples](#)
- [Network Locations](#)

Application Analytics Engine Advanced Configuration

The **Advanced Configuration** panel lets you configure advanced options for the selected Application Analytics 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.

Configuration - [redacted]

Collection Privacy Level: Max End-Systems in Hourly Details:

Client Aggregation: Sensor Log Level:

Store Slow Client Data:

Access Control Integration ⓘ

Enable Access Control Integration:

Access Control Communication Channel:

Wireless Controller Flow Sources ⓘ

Name	IP	Port	WLANs
------	----	------	-------

SIEM Flow Export ⓘ

Export Enabled:

Export IP:

Export Port:

Web Credentials ⓘ

Username:

Password:

Configuration Properties ⓘ

Name	Value
options.FlowServerOptions.maxFlows	1000000
options.NameResolutionOptions.ho...	ets.enterasys.com,enterasys.com,corp.extremenetworks.com
options.FlowServerOptions.override...	NAC
options.FlowServerOptions.override...	NAC
options.AppldCollectorOptions.max...	2500
options.NameResolutionOptions.ho...	true
SecondNetSightTest_enabled	true

- Sensor Modules** ⓘ
- DHCP Decoder: Enable Module Enable Logging
 - DNS Decoder: Enable Module Enable Logging
 - FTP Decoder: Enable Module Enable Logging
 - HTTP Decoder: Enable Module Enable Logging
 - iSCSI Decoder: Enable Module Enable Logging
 - Kerberos Decoder: Enable Module Enable Logging
 - LDAP Decoder: Enable Module Enable Logging
 - NTLM Decoder: Enable Module Enable Logging
 - POP Decoder: Enable Module Enable Logging
 - RADIUS Decoder: Enable Module Enable Logging
 - SIP Decoder: Enable Module Enable Logging
 - SSL Decoder: Enable Module Enable Logging
 - Carrier Detector: Enable Module Enable Logging
 - OS Detector: Enable Module Enable Logging
 - Reputation Detector: Enable Module

Network Settings ⓘ

DNS ⓘ

Manage DNS Configuration

NTP ⓘ

Collection Privacy Levels

Collection privacy level settings restrict the amount of identifying information that is collected by the Application Analytics 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](#).

This allows you to protect 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 Application Analytics engine collects both identifying information and sensitive information. The information displays in the Application Information column.
- **Medium Privacy** — The Application Analytics engine collects identifying information, but not sensitive information. Identifying information displays in the Application Information column.
- **Maximum Privacy** — The Application Analytics 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 Application Analytics 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 Application Analytics engine collects sensitive information when the privacy level is set to Maximum Access.

Client Aggregation

This field determines how client information is aggregated by the Application Analytics engine, either by **IP Address** or **MAC Address**.

Slow Client Data

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

Max End-Systems in Hourly Details

Enter the maximum number of client end-systems stored in the Management Center database for the Application Analytics 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 Application Analytics sensor runs on the Application Analytics 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.

Access Control Integration

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

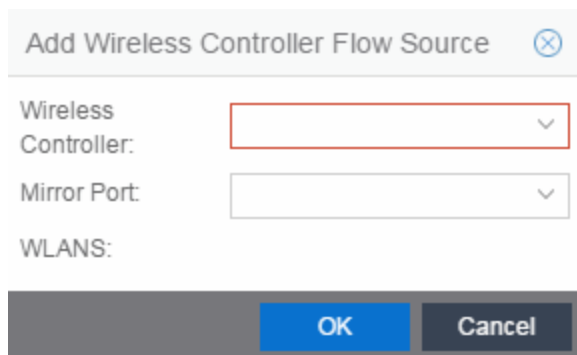
- To enable Access Control Integration for the engine, select the **Enable Access Control Integration** checkbox.
- If your Access Control engines are using Communication Channels, select the **Access Control Communication Channel** option and enter the channel name. The Application Analytics engine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

Wireless Controller Flow Sources

This section displays the Wireless Controllers set up as flow sources in Application Analytics.

To add a Wireless Controller as a flow source:

1. Click the **Add** button.
The Add Wireless Flow Source window opens.



The screenshot shows a dialog box titled "Add Wireless Controller Flow Source". It contains three input fields: "Wireless Controller:" (with a red border and a dropdown arrow), "Mirror Port:" (with a dropdown arrow), and "WLANS:" (with a dropdown arrow). At the bottom, there are "OK" and "Cancel" buttons.

2. Select a **Wireless Controller** from the drop-down menu.

NOTES: Only Wireless Controllers that support Application Analytics and have available L2 ports are listed.
Selecting a Wireless Controller set up as part of a controller pair automatically selects the paired Controller.

3. Select an available L2 port for mirroring in the **Mirror Port** drop-down menu.
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 Application Analytics.

The configuration is complete.

To remove a Wireless Controller as a flow source, select a Controller in the Wireless Controller 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 Management Center server and the Application Analytics engine. Click the **Show Password** check box 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 Application Analytics 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, you should turn it on only for troubleshooting purposes. Do not enable logging during normal operation.

Network Settings

The Network Settings section of the window allows you to configure the network settings on an Application Analytics 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.

DNS

Manage DNS Configuration

Search Domains:

DNS Servers

+ Add - Delete

--

▲
▼

Search Domains

A list of search domains used by the Application Analytics 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, Application Analytics 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 Application Analytics 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 Application Analytics engine is essential for event logging and troubleshooting.

NTP

Manage NTP Configuration

Time Zone: GMT-05:00 - America/Kentucky/Louisville - Eastern Standard Time ▼

NTP Servers

+ Add - Delete

1.1.1.1

▲
▼

Time Zone

Select the appropriate **Time Zone** from the drop-down menu to allow Application Analytics 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 Application Analytics engine.




SSH

Manage SSH Configuration

Port:

Disable Remote root Access:

SSH Users

 Create
 Edit
 Delete

Username	Type	Administrative User

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 Application Analytics engine.

SSH Users

Use the toolbar buttons to create a list of users allowed to log in to the Application Analytics 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

The SNMP configuration section allows you to deploy SNMP credentials for the Application Analytics 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 Application Analytics engine. Select the Manage SNMP Configuration checkbox and provide the following SSH information.

SNMP

Manage SNMP Configuration

Profile:

Trap Mode:

Trap Community Name:

Profile

Use the drop-down menu to select a device access profile to use for the Application Analytics engine.

Trap Mode

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

Trap Community Name

Enter the trap community name.

Related Information

For information on related Application Analytics topics:

- [Analytics](#)
- [Enabling Extreme Access Control Integration](#)