Scrutinizer™

Incident Response System

User Manual

plixer

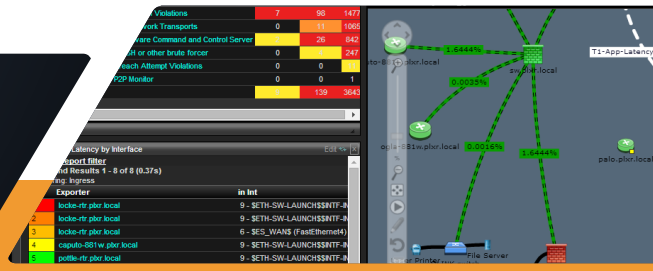

207.324.8805

Table of Contents

Welcome to the Scrutinizer Manual 1
- Overview 1

AdminTab 2

Admin Tab 2
- Overview 2
  - Settings: 2
  - Definitions: 3
  - Security: 4
  - Functional IDs: 5
  - Reports: 7
  - NetFlow Help: 7

Data Aggregation - What is 'Other' Traffic? 8
- Overview 8

Device Details View 10

Report Designer 12
- Overview 12

System Preferences 13
- Overview 13

User Name Reporting 15
- Overview 15
- User Name Reporting 15
- Alarms reporting with User Name 15
- All Exporters Search by User Name 15

User Name Reporting - Active Directory Integration 17
- Upgrade Scrutinizer 17
- Configure a non-admin user to query the Domain Controller Event Logs in Windows 2008 or 2012 17
- Enable Logon/Logoff Audit Policies on the domain controller 21
- Setup IPFIXify on a Windows computer 23
- Example IPFIXify Config 29
- FAQ 31

User Name Reporting - Cisco ISE integration 32
- Enable ERS 32
- Configure Cisco ISE Integration in Scrutinizer 32
  - Other scrut_util options for CiscoISE 33

AlarmsTab 35

Alarms Overview 35
- Bulletin Boards 35

Heat Maps 35
Universal Table Controls 35
Threat Index 35
Alarming on Interface Percent Violation 36
Alarms Configuration 37
Create New Board 37
Flow Analytics Settings 37
Flow Analytics Configuration 37
Policy Manager 37
Exporters in Algorithms 38
Hosts Excluded from Algorithms 38
Edit / Create Policy 39
Overview 39
Notification Profile 41
Overview 41
Threat Index (TI) 43
Threat Index 43
Alarms Views 44
Bulletin Board by Policy 44
Bulletin Board by Violator 44
Notification Queue 45
Orphans 46
Policy Violation Overview 46
Baselining 47
Baselining 47
How Baselining Works 47
Baselining Hosts 47
Enabling/Disabling Baseline Globally 48
Using the Interactive scrut_util.exe utility for Configuring Baselines 48
Adding/Removing default Baselines to an Exporter 48
Reset Baselines to Defaults 50
Monitoring Baseline Processing 50
DashboardTab 53
Dashboard Tab 53
Overview 53
Vitals Dashboard 54
Flow Analytics 56
Flow Analytics 56
FA Navigation 58
FA Configuration 59

Algorithm Activation Strategy 60
Algorithms and Gadgets 62
FA Bulletin Boards 71
Flow Analytics Custom Algorithms 73
Getting Started 73
Creating a Custom Algorithm 73
Adding a Custom Algorithm 73
Testing a Custom Algorithm 74
Configuring a Custom Algorithm 74
Running a Custom Algorithm in Production 74
Creating a Policy for a Custom Algorithm 75
Disabling a Custom Algorithm 75
Deleting a Custom Algorithm 75
Display a list of Custom Algorithms 76
Optimizing FA 77
MapTab 78
Maps Main View 78
Overview 78
Map Connections 79
Overview 79
Map Dependencies 80
Overview 80
Google Maps 81
Overview 81
Map Groups 82
Overview 82
Map Objects 83
Overview 83
Plixer Maps 85
Overview 85
Device Overview 86
Overview 86
Map Settings 87
Overview 87
StatusTab 88
Status Tab 88
Overview 88
CrossCheck 95
Overview 95
CSV Reporting 97

Overview 97
Flow Hopper 98
Flow View 99
Overview 99
Network Traffic Reporting 102
Overview 102
Report Groups 102
Report Types 104
Current Report 108
Current Report (frame contents) 108
Report Options 109
Saved Reports 111
Filters - Include or Exclude Data 111
Thresholds 113
Report Navigation 113
Other Options 114
Report Thresholds 115
Vitals Reporting 118
System 121
Access Denied 121
Backups 122
Overview 122
Database Connection Failure 123
Overview 123
Distributed Collectors 124
Overview 124
Flowalyzer 125
Interactive Scrut_util 126
Overview 126
Language Translations 149
Overview 149
Mailinizer Setup 150
Meta Data Collection 151
Overview 151
How often it is collected 151
What meta data is collected 151
How is the data transferred 151
What Plixer Does with the Data 151
How do I turn it off 152
Multi Tenancy Module 153

Overview 153
NetFlow Knights 154
Search Tool 155
Overview 155
SSL 157
System LEDs 160
Overview 160
Troubleshooting 164
Getting Started Guide 164
Vitals 165
Overview 165
Web Server Port 166
Overview 166
404 Error 167
500 Error 168
Implementation Guides 169
Configuring Amazon Web Services flow streaming 169
Overview 169
Prerequisites 169
Required Information 169
Create an AWS User for Flow Streaming 170
AWS Configuration in Scrutinizer 170
Frequently Asked Questions 171
Troubleshooting AWS 172
Configuring Cisco's FireSIGHT eStreamer Client 173
Overview 173
Prerequisites 173
Required Information 173
Register Scrutinizer with FireSIGHT 174
Configure Scrutinizer's eStreamer Client 177
Wait for Flows 178
Configuring Endace Probe Integration 179
Overview 179
Setting up Endace Packet Capture Integration 179
Accessing Endace Probes in Scrutinizer 179
Configuring a Scrutinizer Collector for Dual/Multi-homing 182
Security Considerations 182
Virtual Routing and Forwarding (VRF) Mode 182
Configuration Examples 182
Creating a Network Map 185

Overview 185
Creating a Plixer Map 185
Creating a Google Map 187
Creating Thresholds and Notifications 188
Overview 188
Setting the Global Threshold 188
Applying a Notification to the Global Threshold 188
How to Create a Notification that is sent via email 188
Creating the Notification Profile to use 189
Adding a Threshold that sends an email notification 189
Elasticsearch / Kibana (ELK) Integration 191
Overview 191
What is ELK? 191
Integration Prerequisites 191
Kibana Searches from Scrutinizer Reports and Alarms 191
Scrutinizer Reporting from within Kibana 192
How to Configure ELK Integration with Scrutinizer 193
Scrutinizer for Splunk Application 195
Overview 195
Splunk Searches from Scrutinizer Reports and Alarms 195
Scrutinizer Reporting from within Splunk 195
How to Configure Splunk Integration with Scrutinizer 197
Using the Scrutinizer Reporting API 199
Overview 199
Getting Started 199
Examples 199
Reference Guide 203
Using the Scrutinizer User API 214
Overview 214
Getting Started 214
Examples 214
Reference Guide 216
Index 229

Welcome to the Scrutinizer Manual

Overview

Welcome to the on-line manual. Click Here for help troubleshooting. You can also get this manual in .pdf format.

You will find the latest updates to the manual online. There are also online webcasts which give quick overviews (i.e. 2 - 5 minutes each) of specific features.

IMPORTANT: Don't struggle, contact support!

AdminTab

Admin Tab

Overview

The Settings page is primarily left to the administrators.

Settings:

- **Alarm Notifications**: enable additional system alarms
- **ASA ACL Descriptions:** Enter the username and password used to SSH into ASA firewalls to retrieve ACL descriptions (Appliance only).
- **Alarm Settings:** modify settings to optimize syslog and SMTP processing.
- **CrossCheck**: Specify the thresholds for changing color and the syslog threshold that the Fault Index must reach to trigger a syslog.
- **Data History**: Specify how long each flow interval is saved.
    -
    - **Historical 1 Min Avg**: Saves 100% of all flows received. Make sure the server has enough disk space to save significant quantities of the raw flows. The 1 minute intervals consume the most disk space as it is not aggregated and flows are in raw format.
    - **Historical 5 minute - 1 week Avg**: These intervals only save the specified Maximum Conversations after aggregation per interval.
    - **Maximum Conversations**: Used when creating large intervals (e.g. 5 minute) from prior intervals (e.g. 1 minute). All flows are aggregated together per router. The top 1,000 (default) based on bytes are saved.
    - **Auto History Trimming:** This option allows for automatic database trimming when available disk space falls below 10GB*. Check the checkbox to activate this option. An alarm will also be generated to alert you that the database is being trimmed (1 minute and 5 minute conversation database tables) and includes how much 1 minute and 5 minute data currently exists in the database (in hours).

      * The Auto History Trimming threshold is higher for Hardware Appliances. See the System LEDs page for more information.

        - **Note** that in a distributed collector environment, each collector will perform the database trimming independent of the other collectors. Auto History Trimming on/off applies to all of the collectors in the cluster, but the database trimming will only occur on the server(s) that fall below 10GB of available disk space.
- **Email Server**: Necessary for on demand and scheduled emailed reports. Make sure the test is successful.
- **FA Configuration**: add NetFlow exporters to FA Algorithms.
- **Flow Analytics**: configure miscellaneous algorithm settings (Jitter, Latency, etc.)
- **Licensing**: Displays the current licensing level, expiration date(s), and unique Machine ID for this installation. **The Machine ID is required by Plixer Customer Service for generating new license keys.** Once a new key is received, to activate the key, copy and paste the entire key in the License Key textbox.
- **Mapping Configuration:** Customization for both Plixer maps and Google maps (e.g. connections, text boxes, etc.). Learn more about mapping.

- **Proxy Server**: Setup the server to work with a proxy server
- **Syslog Server**: Configure the syslog server, port and priority
- **System Preferences**: Other system configuration options

Definitions:

- **3rd Party Integration**: Create links to 3rd party applications and pass variables in URLs
- **Applications**: Setup and modify applications using ranges of ports and IP addresses. This feature is useful for properly labeling in-house applications.
- **Autonomous Systems**: Display and search Autonomous System Names that are shipped with the software, or imported by user. Use import_custom_asns in Interactive scrut_util.exe to import AS Names.
- **Device Details**: Displays the SNMP details of the devices sending flows. Allows custom device and interface names to be defined which override the defaults. Notice that in and out speeds can be configured.
- **Host Names**: Setup and modify known hosts. Use this option to statically assign host names to IP addresses that will not age out. It can also be used to label subnets in the Subnet report types. There are three Resolve DNS options:
    - **Current** - has been or attempted to be resolved already (will expire in whatever days are set in the serverprefs)
    - **Queued** - ready to be resolved by the resolver. User can set it to queued to force a DNS resolve again on the host.
    - **Never** - a permanent address that was manually added by the user. Users can make names permanent by switching this to never. It's not purged.
- **IP Groups**: Define ranges of IP addresses that belong in a specific group (e.g. Marketing, sales, phones, etc.). Run a report on an interface to see the IP Group reports.
- **Language**: use this interface to update languages or create new translations.
- **Manage Collectors**: Details on the servers collecting flow(s) for this Scrutinizer install. (May be multiple if the Distributed solution is deployed.)
    - Delete: This check box can be used to remove collector(s) from the list.
    - Collector: IP Address of the flow collector.
    - State: Current state of the flow collector – ONLINE or OFFLINE
    - Exporter Count: Number of exporters sending flows to the collector.
    - First Flow Time: Timestamp when flows first received by the collector.
    - Last Flow Time: Timestamp when last flows received by the collector.
    - Flow Rate: Current flows per second per collector.
    - Packet Rate: Current packets per second per collector.
    - MFSN Rate: Missed Flows Sequence Number rate in flows per second.
    - Duplicate Rate: Duplicate flows per second.
- **Manage Exporters**: Details on the devices sending flows. This page provides the following information and configuration options as viewed from left to right on the screen:
    - Delete: Check this checkbox to remove the device from the Status tab device tree. The device will be rediscovered immediately if the collector is still receiving flows from the device. Also, templates and interfaces from devices that stop sending flows are aged out.

- Status LED (click to view the Device Details):
    - Green: this exporter is enabled and up on the collector specified.
    - Red: this exporter is enabled and down on the collector specified.
    - Yellow: no flows have been received for this exporter on the collector specified.
    - Gray: This exporter is disabled on the collector specified.
- Last Activity: timestamp when last flow received for this exporter.
- State: enabled, disabled, backup, unlicensed
- Name: Exporter name, or IP Address if unnamed. Clicking on name/IP Address opens a Manage Exporters modal with options to Name the exporter, the domain for the exporter, set Protocol Exclusions for this exporter, SNMP Credential selection, and also attach Additional Notes to the exporter.
- Exporter IP: IP Address of the exporter. The Manage Exporters configuration modal is also available by clicking on the Exporter IP.
- Collector IP: IP Address of the collector receiving flows for this exporter.
- Credential: SNMP Credential in use by this exporter. Clicking on the SNMP Credential opens the Manage Exporters configuration modal to the SNMP section, allowing editing of the credential.
- Additional Notes: Any notes added to this exporter are visible in this column.

- **Notification Manager:** Configure notifications to be applied to Policies in the Alarms tab
- **Policy Manager:** List all of the Policies that are configured for the Alarms Tab
- **SNMP Credentials**: Configure the SNMP Credentials used on each flow exporter. SNMP v1, v2 and v3 are supported.
- **Type of Service (ToS)**: Configure the ToS and DSCP values displayed in the reports. Be sure to Define the "ToS Family" under System Preferences.
- **Well Known Ports**: define port names. In the Well Known Ports report, the following logic is used:
    - Which port is lower the source port or the destination port
    - If the source port is lower and defined, use this as the well known port
    - else, use the destination port if defined as the wellknown port
    - else, display the lower port as the wellknown port

Security:

- **Auditing:** Displays a report of all the administrative actions users have performed within Scrutinizer.
- **Authentication:** Configure general authentication settings, enable or disable different technologies, allow or deny users from different authentication methods, and set the order in which methods are attempted.
- **Authentication Tokens:** These tokens can be used to automate Scrutinizer application logins with user-specific permissions and applicable expiration dates without having to include user name and passwords in the URL.
- **LDAP Configuration:** Server and connection settings for LDAP integration.
    - LDAP User Authentication Process:

1. In the LDAP Configuration, Administrators provide credentials for an LDAP account with permission to see any users they'd like to permit access to Scrutinizer.
    a. This is the account that will be used to search for and authenticate users when users attempt to log in.
    b. The Searchbase defines the group that will be searched for authorized users. This is a required field.
    c. The scope of users in that searchbase who are allowed to authenticate can be limited in two ways:
        i. By specifying one or more Security Groups in the LDAP Configuration
        ii. By specifying individual user account names in Security > Authentication > LDAP
2. A user attempts to log in.
    a. Scrutinizer authenticates as the administrative account provided, then checks a searchbase specified by the Scrutinizer administrator for any account matching the username provided.
        - Authentication with the sAMAccountName, UserPrincipalName, or uid attribute is supported.
3. If the LDAP server responds with an LDAP_REFERRAL code, Scrutinizer will check the referred server.
4. If the Scrutinizer administrator has specified multiple LDAP servers, it will check them all in series until successfully authenticating or failing.
5. Once the user has successfully authenticated for the first time, Scrutinizer checks for any Security Group they're a member of which also exists in Scrutinizer with the same usergroup name. If it does, they're added to the Scrutinizer usergroup automatically.

- **RADIUS Configuration:** Server and connection settings for RADIUS server integration.
- **TACACS+ Configuration:** Server and connection settings for TACACS+ integration.
- **User Groups**: Specifies what a Group login account can access. Some permissions require further explanation:
    - Device Status: Grants permission to see the status of the device (i.e. Flow exporter). Device icons appear blue in maps if the “Device Group” permission is granted without this permission. Mailinizer devices show up here.
    - Interface Statistics: Grants permission to see the statistics of an interface. Mailinizer does not show up here.
    - Device Groups: Grants permission to see a Group (i.e. map). Devices (i.e. Flow Exporters) appear blue and interfaces black unless permission is granted in “Device Status” and “Interface Statistics”.
- **User Accounts**: Configure login preferences for individual accounts. User Accounts must be a member of one or more User Groups. If no group is selected when a User Account is created, they are placed in the default (e.g. Guest) User Group. Permissions are inherited by all User Groups a User Account is a member of.

Functional IDs:

A functional ID is a generic account used for an IT asset. Specific to Scrutinizer, functional IDs are User IDs providing access to different areas and access levels within Scrutinizer.

The following functional IDs are included by default with the installation of Scrutinizer.

Application	User	Usage	Access Level	Description
Operating System				
	root	Interactive	Privileged	Provides user root access to the operating system. User has unrestricted shell access, SSH and console. Some processes run as root (procmon, scheduler, poller) and some scrut_util processes when run by the scheduler.
	plixer	Interactive	Non-privileged	Provides user access to run all plixer services and/or processes. Primary user for the Scrutinizer interactive prompt.
	pgbouncer	Non-interactive	Non-privileged	Database load balancing
	postgresql or mysql	Non-interactive	Privileged	Database user
	apache	Non-interactive	Privileged	HTTP services
Database				
	root	Interactive	Privileged	Provides full read/write to PostgreSQL/MySQL database for local users.
	scrutremote	Interactive	Non-privileged	Provides communications between remote systems.
	scrutinizer	Non-interactive	Privileged	Local DB connections for Scrutinizer.
Web interface				
	admin	Interactive	Privileged	Full access to management functions for Scrutinizer.

- **Usage**
  - Interactive usage allows a User to use that ID, inheriting all privileges that were granted to that ID.
  - Non-interactive is used internally by the system only and cannot be used by a User.
- **Access level**
  - Privileged access level is defined as an ID with elevated entitlements such as a system administrator or super user.
  - Non-privileged access allows just the access level required for the intended functions of the ID.

Reports:

- **Report Designer:** Is used to create new reports that are not part of the core reporting solution. Click here for instructions on using the Report Designer.
- **Report Folders**: Manage Saved Report Folders found in the Status tab under saved reports. Notice the Membership drop down box:
  - Folders: Select a folder and add or remove reports from it.
  - Reports: Select a report and add or remove folders it can be found in.
- **Scheduled Reports**: Manage Scheduled Reports, delete, etc.

NetFlow Help:

- Activating NetFlow, J-Flow, sFlow, NetStream, IPFIX, etc.

Data Aggregation - What is 'Other' Traffic?

Overview

What is other traffic? Other traffic is traffic that didn't make it into the top 10. Some important details should be understood when trying to comprehend Other traffic.

The collector saves 100% of all data in raw format to the 1 minute conversations tables for each router. Every hour it creates a new 1 minute interval table per router. Every 5 minutes, it creates higher intervals using the smaller intervals. This process is called "roll ups".

When the roll ups occur for 5 Min, 30 Min, 2 Hr, 12 Hr, 1 Day and 1 week, two tables are created:

1. Totals: The total in and out byte counts are saved per interface before the data for the conversations table is calculated. This table allows the reporting front end to display accurate total throughput per interface over time and allows the front end to operate with no dependency on SNMP yet still provide accurate total utilization reporting.
2. Conversations: All flows for the time period (e.g. 5 minutes) are aggregated together based on a tuple. Once all flows are aggregated together, the top 1000 (i.e. default) flows based on byte count are saved. The non top 1000 flows are dropped. Remember: the total tables ensure a record of the total in / out utilization per interface over time.

When a report is run on an individual interface within 1 minute intervals, the totals table isn't needed because the conversations table contains 100% of the data. When a report is run on an individual interface with no filters in 5 minute or higher intervals, both the Conversations and Total tables are used in the report. When reporting, the Total tables are used to display the total in and out utilization of the interface and the top 10 from the Conversations table are subtracted out from the total and added back in color.

IMPORTANT: In some cases, a report that doesn't utilize the Total tables can understate the actual utilization of the interface.

The Total tables are not used when:

- Reporting across All interfaces of a device
- A report is run on multiple interfaces from different devices regardless of filters
- Looking at data from a single template (e.g. using a single or multiple Medianet templates).
- Looking at 1 minute intervals in a report. One minute intervals contain 100% of all data exported for the template as no roll ups have occurred. As a result, no Total tables are created for 1 minute intervals.

The Total tables are only being used when

- Looking at 5 minute intervals and higher
- The "Flow Templates" section of the report filter indicates "Available Templates".
- Looking at a single interface without any additional filters

Remember: Only the top 1000 (default) conversations are saved in the roll ups by default. If the server has the available disk space, try increasing the Maximum Conversations under Admin Tab - Settings - Data History to 10,000 and see if it improves the accuracy. Don't configure it right away for the maximum of 100K rather, ease up the number of conversations saved over a few days. Some reports may render more slowly when the Maximum conversations is increased, this is the result of the tables being larger.

A Word about sFlow

When collecting sFlow, make sure the packet samples and the interface counters are both being exported to the collector. The collector will save the packet samples to the Conversation tables and the Interface counters to Total tables even at 1 minute intervals.

WARNING: If the flow device (e.g. router) is exporting multiple templates for different flows it is exporting, utilization could be over stated if the flows contain the same or nearly the same information. The front end of Scrutinizer will render reports using data from all templates with matching information. Be careful when exporting multiple templates from the same device! If you find this to be the case, use the filters to select a single template.

Device Details View

Using this interface, selected interfaces can be hidden from the reporting GUI. The SNMP community string used to communicate with the device can be altered.

Notice at the top: there is a drop down box with all the flow sending devices. Under the devices is a drop down box to select the SNMP community string/credential for the selected device. Next to the community string is a check box for SNMP Enabled. If SNMP Enabled is checked, the Watcher Service will attempt to poll and update SNMP information for the device. By default, the automatic SNMP discovery occurs once a night. The user can disable the automatic SNMP capability by unchecking "Auto SNMP Update" from the Admin Tab, Settings -> System Preferences.

There are several columns displayed for each interface on the NetFlow capable router/switch. Some of them include:

- Instance
- Custom Description: A custom interface name can be entered.
- ifAlias
- ifName
- ifDescr
- ifSpeed: Custom speeds can be specified both inbound and outbound per interface.
- Direction: tells us if NetFlow is collected INGRESS, EGRESS or BOTH on this interface.

Scrutinizer will attempt to build the drop down boxes based on whether or not the following information is available in this order:

- Instance and Custom Name
- Instance, ifAlias and ifDescr
- Instance, ifDescr and ifName
- Instance and ifDescr
- Instance

This interface relies on devices that support the SNMP standard MIB II. SNMP Enterprise MIBs may require 3rd party software or customized scripts to correlate the enterprise instances to match the MIB II instances.

If SNMP is not available, the collector will look for an interface names option template. Some vendors export an interface names option template using NetFlow or IPFIX. This option template contains the names of the interfaces. In Cisco IOS v 12.4(2)T or greater, the command is:

Router(config)# ip flow-export interface-names

SonicWALL and other vendors export a similar options template.

If the Custom Description is filled in, it will over ride the use of the SNMP descriptions. This is also true when the Custom (Mb) is filled in, they will over ride the use of the SNMP ifSpeed. Enter a 0 in the Custom (Bits) ifSpeed to force the Status tab to display the interface in bits in lieu of % utilization.

If any updates are applied to a router or switch, be sure to go back to the device interface and run an update by clicking on the Update button else, the default evening update will take effect.

Direction: Displays how the flows are collected and reported on for the interface. Values are INGRESS, EGRESS or BOTH and are not updated until the collector is restarted. If Direction is unset '-' this means NetFlow is not exporting for this interface.

If the interface row is white then the interface number and traffic values are inferred from NetFlow exported from another interface. If the interface row is gray then the interface number was discovered via SNMP and there will be no traffic values.

Protocol Exclusions are performed to avoid traffic from being counted twice on a given interface. Generally over reporting is caused by VPNs or tunnel traffic. Exclusions can be made per exporter (e.g. router, switch, etc.) or per interface per exporter. They can also be excluded globally across all exporters. Click on the (-) icon to launch the Protocol Exclusions modal.

VERY IMPORTANT By default, the flow collector nightly SNMP polls the switches and routers it is receiving flows from. This software was engineered to be a passive collection tool with minimal SNMP requirements. The best way to update the SNMP information including the information on the interfaces is to click on the "Update" button. NetFlow v9 option templates can be used in place of SNMP to gather interface names and speeds.

Report Designer

Overview

The Report Designer is used to create new reports that are not part of the core reporting solution. It can be used against any flow template even when byte counts are not available. These new report types only appear on devices that are exporting the necessary elements in templates. The steps to design a new report is as follows:

1. Copy an existing report design or select 'New'.
2. Enter a name for the new report design.
3. Select a device that is exporting the template that is needed for the report.
4. Select an element in the template for the first column.
    a. Specify the column name. It is best to try and keep it short. Often the name entered is similar to the Report Field name.
    b. Specify the width. In most cases use 'Dynamic' as specifying pixels is generally used to cut off long element names.
    c. Specify the treatment.
        - Average: takes the average of the total values divided by the number of matches.
        - Count: Counts the number of entries in consideration of the group by columns.
        - Count Distinct: Counts the number of entries in consideration of the group by columns but, if a matching flow shows up more than once, it is only counted once.
        - Max: Display the maximum value.
        - Min: Display the minimum value.
        - Sum: Adds up the values
        - Group By: Group the matching values together.
    d. Rate vs. Total
        - Rate : trend the data by rate per second. Total is not an option in the drop down box.
        - Total : trend the data by total per interval. Rate is not an option in the drop down box.
        - Rate (default) / Total : trend the data by rate per second. Total is an option in the drop down box.
        - Rate / Total (default) : trend the data by total per interval. Rate is an option in the drop down box.
    e. Stack or Unstacked
        - Stacked : trend the data as a stacked trend. Non Stacked is not an option in the drop down box.
        - Non Stacked : trend the data as an unstacked trend. Stack trend is not an option in the drop down box.
        - Stacked (default) / Non Stacked : trend the data as a stacked trend. Non Stacked is an option in the drop down box.
        - Stacked / Non Stacked (default) : trend the data as an unstacked trend. Stack trend is an option in the drop down box.

The new report will show up in the category Designed Reports when the device template(s) contain the elements necessary for the report.

NOTES:

- The report will not work outside of one minute intervals if rollups are not being performed on the template in a format that is supportive of the report created.
- The order of the columns can be reordered. Grab a row in the table with the mouse and move it up or down then release it.

System Preferences

Overview

System Preferences is accessible via Admin Tab -> Settings -> System Preferences.

Included in System Preferences are global configuration settings for Scrutinizer, including optional settings such as CSV Repository, and threshold settings for color changes in Top Interfaces and maps.

Following is the list of options available in System Preferences.

Always display Totals: In Status reports, this option will force the table to always show totals (even if the graph is showing rate).

Auto SNMP Update: If enabled, the server will rediscover all SNMP enabled devices once per day at 1 AM.

Business Hours End: The end of the business day as an integer. 5pm = 17

Business Hours Start: The start of the business day as an integer. 8am = 8

CSV Repository: The directory to store scheduled CSV files.

Display Others on Top: Report Graphs can display the 'Other' traffic on top of or below the top 10

DNS Cache Retention: Number of days to retain DNS names. 0 means never remember. Valid range is 0 - 365.

DNS Resolve: When checked, the Domain Service will actively resolve IP addresses found in recent flow data and add them to the DNS Cache.

DNS Timeout: The number of seconds to try and resolve DNS names.

Google TLD: Defines the Top Level Domain used by Google Maps and Geocoding. For example .cn, .com, .co.kr, .co.uk, etc. Please reference google's supported domain list HERE.

Host Indexing: When checked, the Domain Service will actively index IP addresses found in recent flow data

Inactive Expiration: The number of hours that must pass before interfaces are flagged as no longer sending flows. 1 hour to 168 hours (1 week) are valid. The default is 24 hours.

Inactivity Threshold: The global threshold will trigger alarms if any interface or device does not transmit flows in X number of minutes.

Language: The default system language. Users can also have their own default language which is updated in Admin -> Security -> User Accounts. The default is English. The manual and technical support are only available in English.

Listener Port: The port(s) that the collector will use to listen for any NetFlow/sFlow Traffic sent to it. Default is 2055, 2056, 4432, 4739, 9995, 9996, and 6343. Separate multiple ports by a comma.

Map Refresh Interval: The time (in minutes) between automatic refreshes of your maps on the Maps Tab.

Max graphed intervals: The maximum number of intervals allowed in a graph. Default = 300

Max Report Processes: Each report run will use this as a maximum number of sub process. This breaks reports up by time or exporters depending on which will be faster.

Max Reports per Interval: The maximum number of reports users are able to schedule for the same minute. The default is 5.

Maximum FlowView Exporters: The maximum number of exporters allowed in a filter for Flow View to be an available option.

Push Data Aggregation: Apply data aggregation when pushing temp tables from collector to reporter. (Only applies to Distributed collector environments)

Report Caching Timeout: The number of minutes to cache a list of available reports.

Share Health Statistics: Share information on your servers health with Plixer Support. This allows technicians to proactively reach out when problems arise. No company traffic information is shared. This feature may require granting access through your firewall.

Theme: This option sets the system theme. NOTE: Your preference in Admin Tab -> Security -> User Accounts will override this setting.

Threshold-Yellow: The value to use for a Yellow Threshold. The default is 51. Used in Top Interfaces and Map Connections.

Threshold-Orange: The value to use for an Orange Threshold. The default is 76. Used in Top Interfaces and Map Connections.

Threshold-Red: The value to use for a Red Threshold. The default is 90. Used in Top Interfaces and Map Connections.

Threshold-Utilization: The global threshold will trigger alarms if the total in or out utilization of any interface exceeds the specified percentage of the total bandwidth. Default is 90%.

ToS Family: Select the family or configuration of Quality of Service or Type of Service your organization has implemented. Options are: RFC 791, RFC 1349, or RFC 2474.

Version Checking: Enable Scrutinizer to connect to the internet and check for updates.

User Name Reporting

Overview

User Name reporting (and other User Name features) requires integration with an authentication system such as a Microsoft Domain Controller. Most authentication systems are supported (e.g. Cisco ISE, Enterasys Mobile IAM).

The following sections of the User Manual provide some step-by-step help in configuring the integration.

- User Name Reporting - Active Directory integration
- User Name Reporting - Cisco ISE Integration

Other devices that require authentication, such as firewalls and wireless LAN controllers, can also provide User Name information to Scrutinizer.

Once the User Name integration is in place, the following features are available.

- User Name Reporting
- Alarms reporting with User Name
- All Exporters Search by User Name

User Name Reporting

User Name reports are available under:

- Top reports category
  - Reports the amount of traffic generated per user
- Device-specific report categories (such as SonicWALL, Palo Alto, or Wireless reports)
  - Reports the amount of traffic generated per user
- Source / Destination > User Name by IP reports
  - These reports display the user name(s) associated with the IP Address during the report time frame.

With the User Name reports, you can monitor network usage per user, including alerting on a user's traffic volume or type of traffic if thresholds have been violated.

Alarms reporting with User Name

Alarms can be associated with the user name of the user that has triggered them, helping to reduce the MTTR (Mean Time to Resolution) for network issues by highlighting who was responsible for the alarm.

All Exporters Search by User Name

If it's a specific user that requires investigation and/or monitoring, finding that users traffic is quick and easy with the Search Tool on the Status page, using either "User as Source" or "User as Destination" as the search field.

User Name Reporting - Active Directory Integration

Upgrade Scrutinizer

First, it's recommended to upgrade to the latest version of Scrutinizer. Please contact Plixer support at 207-324-8805 x4 for assistance.

Configure a non-admin user to query the Domain Controller Event Logs in Windows 2008 or 2012

1. Create a domain user for IPFIXify to use
    a. Add the IPFIXify user to the "Event Log Readers" Built-in group

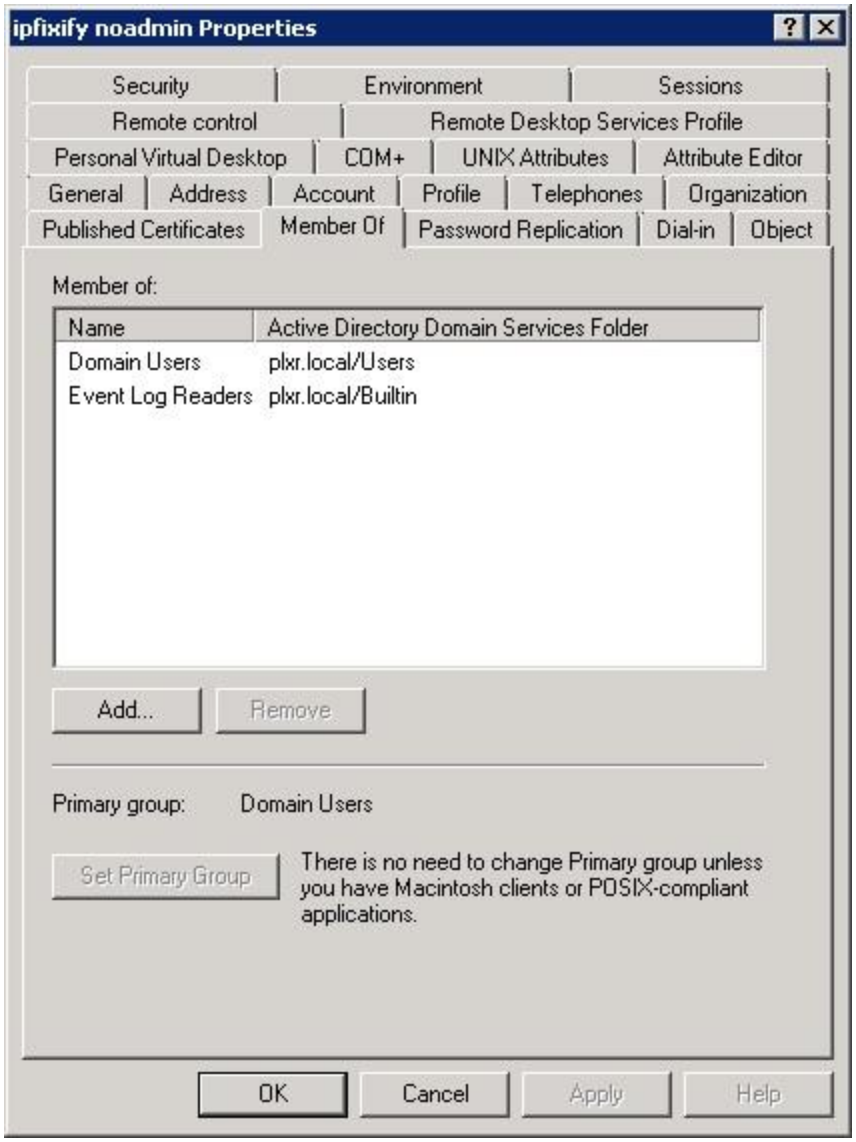

2. Provide WMI Permissions:
    a. Login to the Domain Controller as an administrator
    b. Go to "Start -> Run"
    c. Type "wmimgmt.msc"
    d. Right click on “WMI Control (Local)” and select “Properties”

e. Go to the "Security" tab, click on “Root”, and then select “Security”

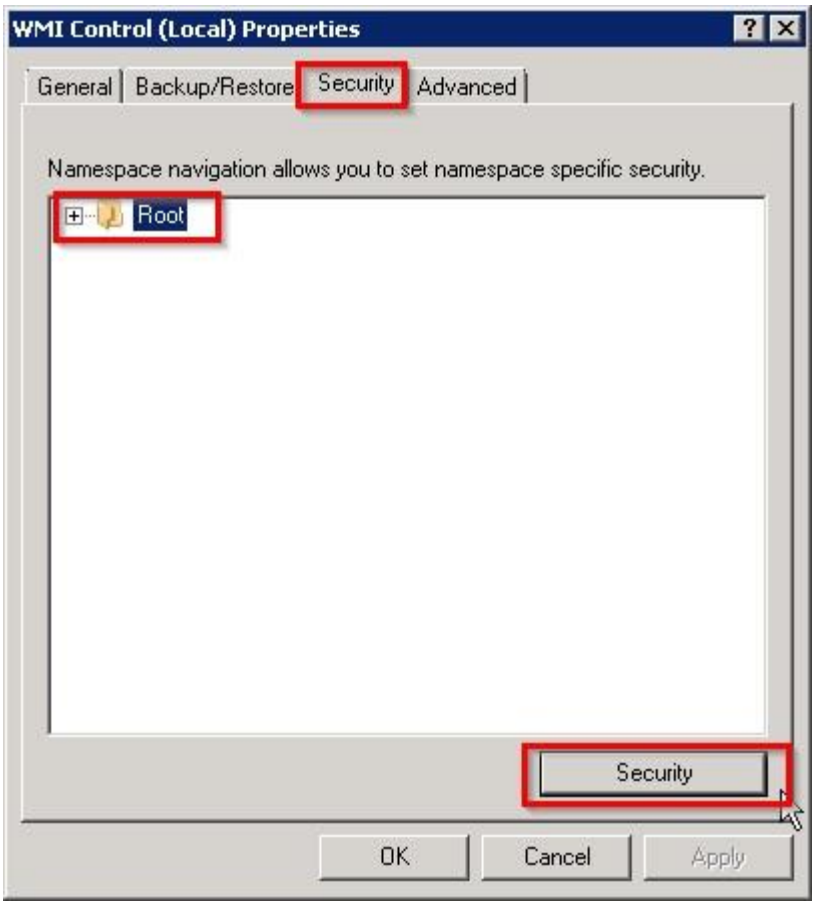

f. In the next popup, select “Advanced”

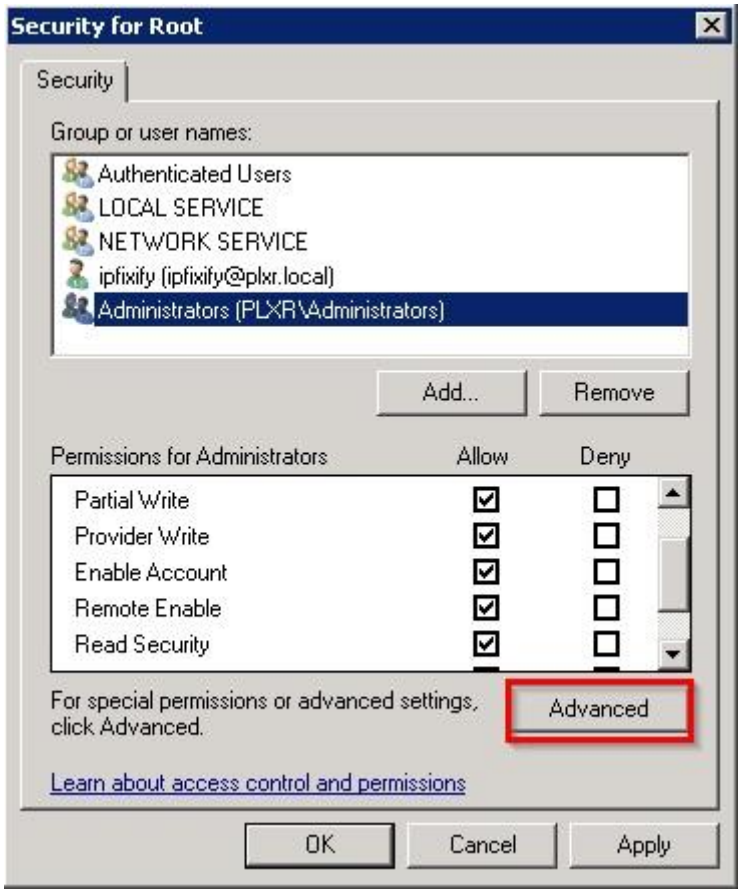

g. Press "Add…" and then enter the ipfixify user.

h. Under the "Apply to:" section, make sure it is configured for "This namespace and subnamespaces".

i. Give the user "Enable Account" and "Remote Enable" Allow privileges. Apply these changes by pressing OK in each of the popup windows.

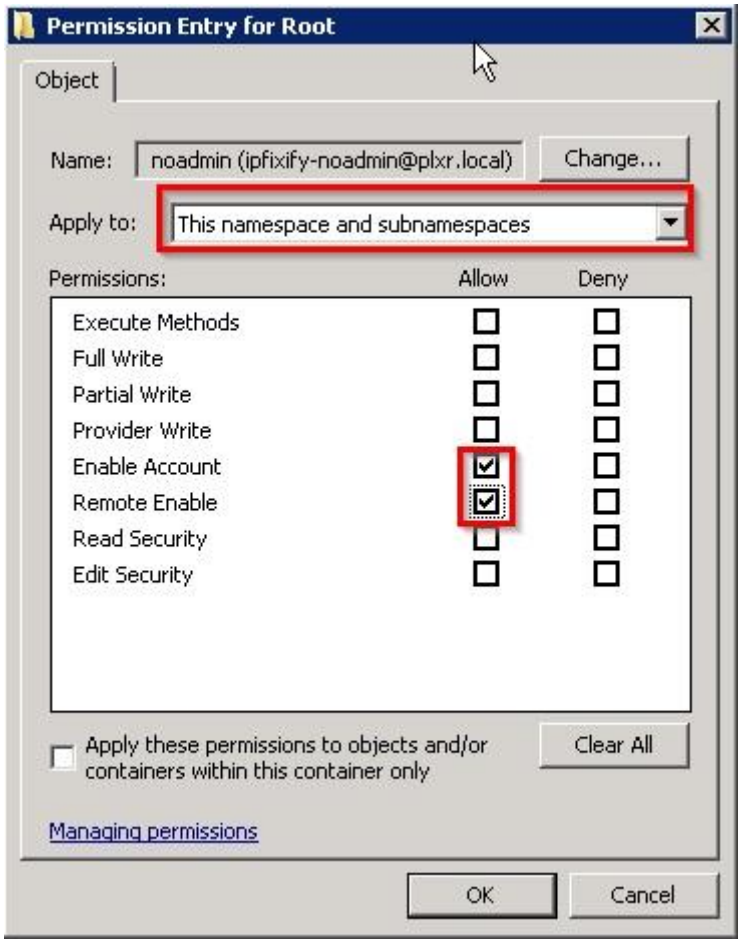

Enable Logon/Logoff Audit Policies on the domain controller

1. Modify the default domain policy for domain controllers and enable the following group policies:
    a. Expand Computer Configuration -> Policies -> Windows Settings -> Security Settings -> Advanced Audit Policy Configuration -> Audit Policies -> Logon/Logoff and then enable success and failure for "Audit Logoff" and "Audit Logon".

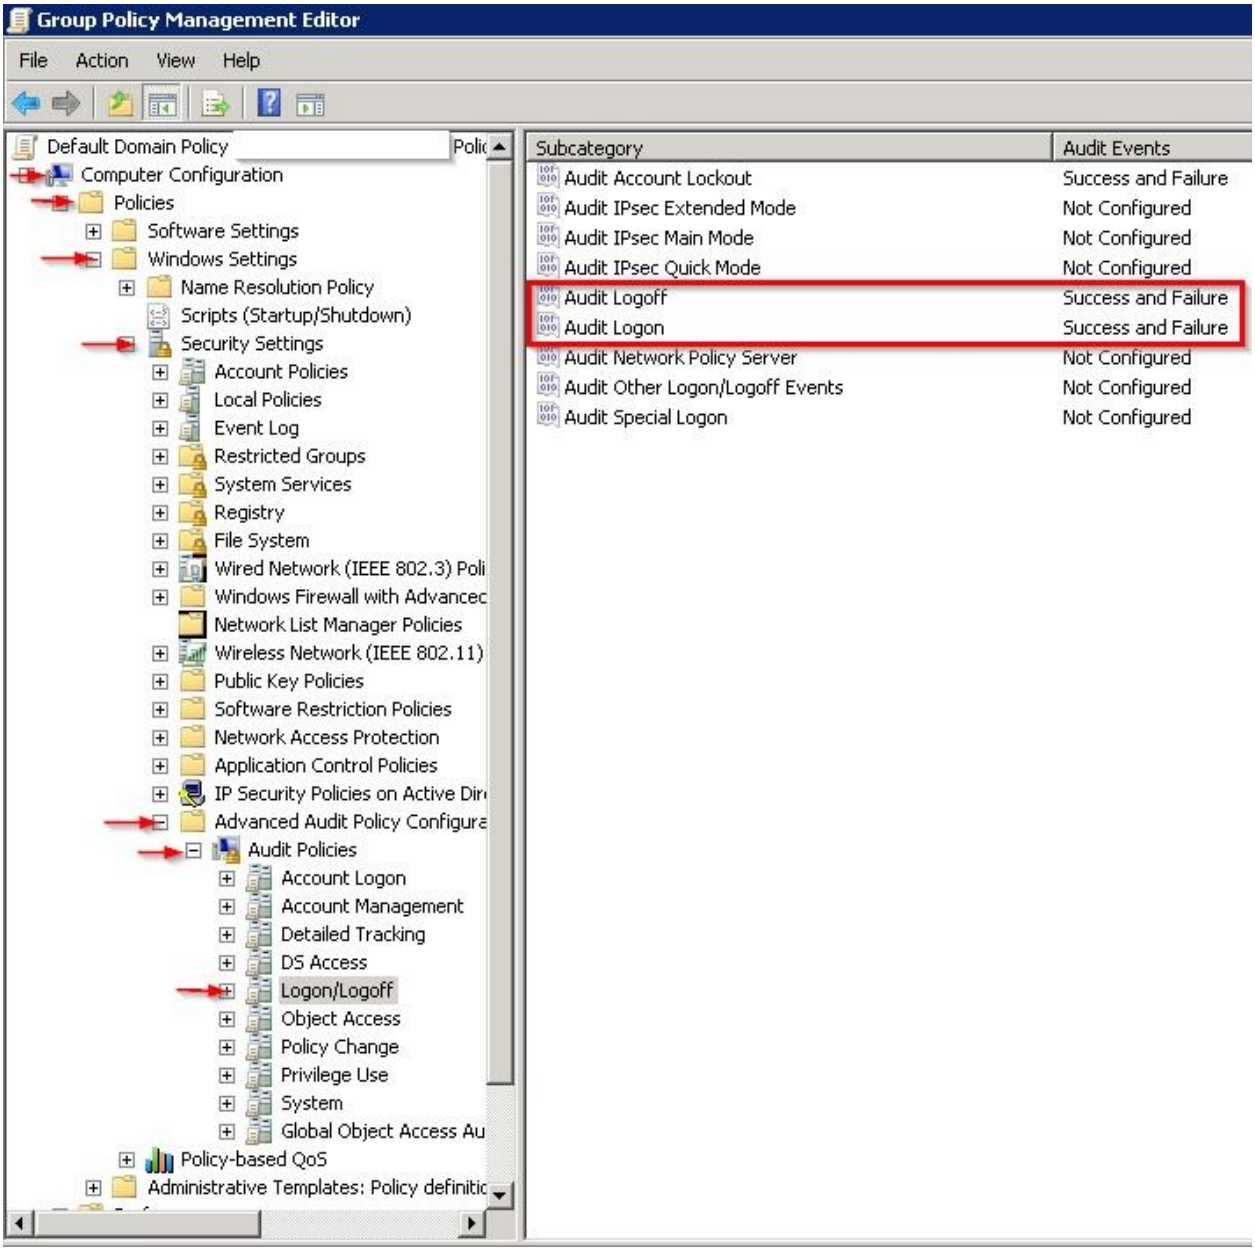

b. The advanced audit policies require that another group policy override setting is enabled under: Computer Configuration -> Policies -> Windows Settings -> Local Policies -> Security Options -> Audit: Force audit policy subcategory settings -> Define this policy setting and set to Enable

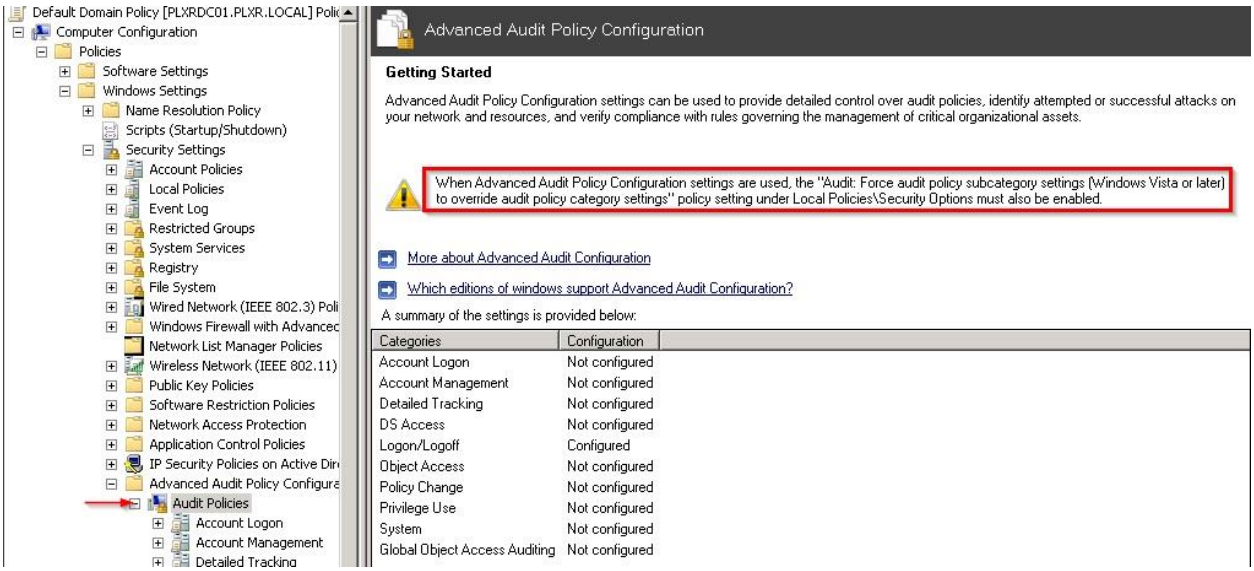

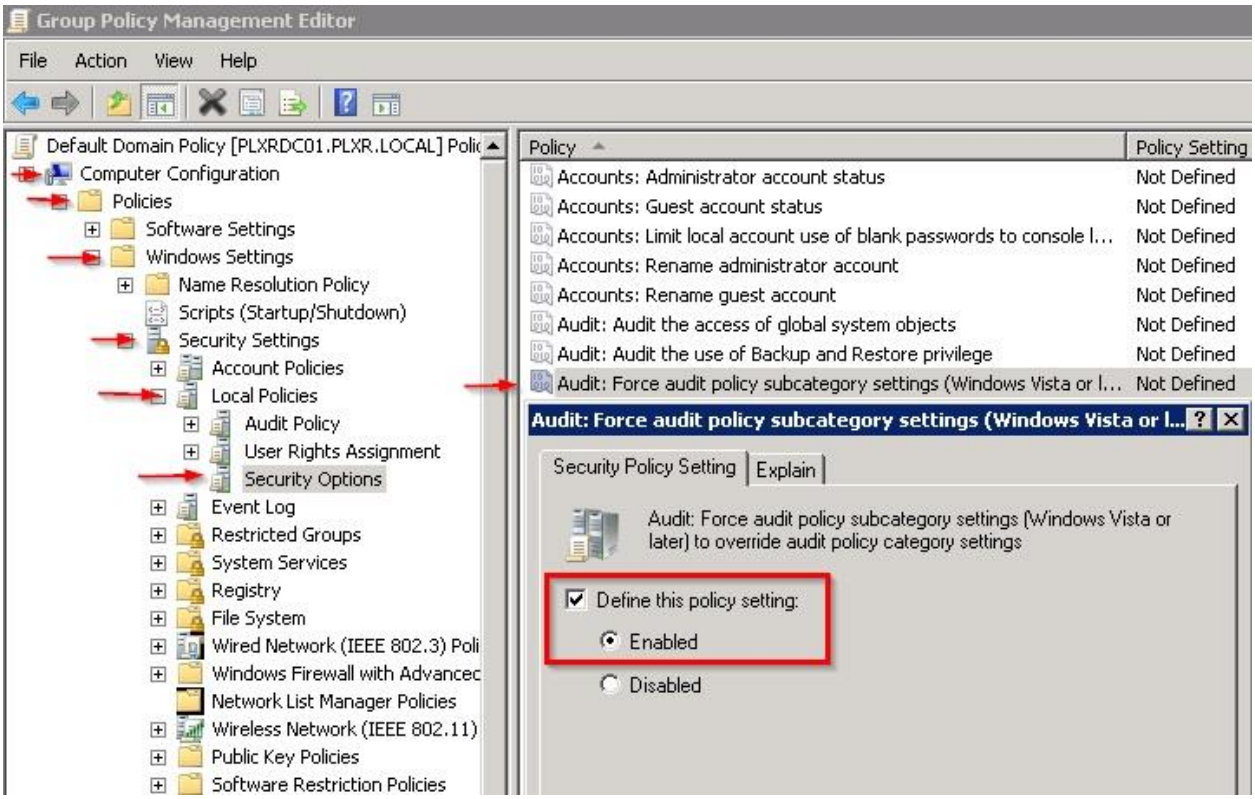

Setup IPFIXify on a Windows computer

1. Move /home/plixer/scrutinizer/files/conf/ipfixify-template.cfg to C:\ipfixify on a Windows computer that will run IPFIXify and also download the Windows IPFIXify executable to this Windows computer.

2. Rename ipfixify-template.cfg to ipfixify.cfg, open the file in a text editor, and then edit the following lines:

    a. collector=NetFlowIP:port
        i. Enter your NetFlow collector's IP and port.
    b. member=DCip
        i. Enter in the IP address of your domain controller. For each additional domain controller, add another member line.
    c. usernamesOnly=yes
        i. If you only want to collect username data, set this value to yes.

3. Configure the IPFIXify user credentials

    a. Open a command prompt and navigate to the directory that contains ipfixify.exe

    b. Run the following command and enter in the ipfixify user and password: ipfixify.exe --credentials ipfixify.cfg

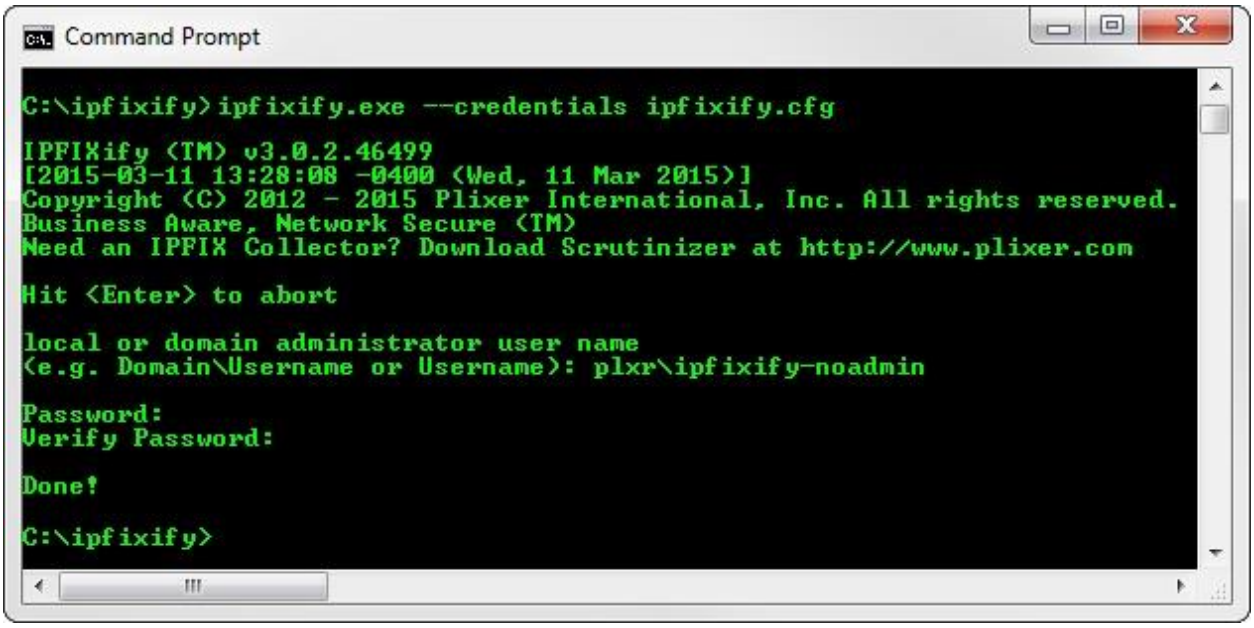

4. Download PSTools.zip and move PsExec.exe to the same directory as ipfixify.exe and ipfixify.cfg.

    a. PSTools.zip download: https://technet.microsoft.com/en-us/sysinternals/bb897553

    b. Before PsExec.exe will function, you must accept the agreement by doing the following:

        i. Hold down shift and then right click on PsExec. In the menu, select "Run as a different user"

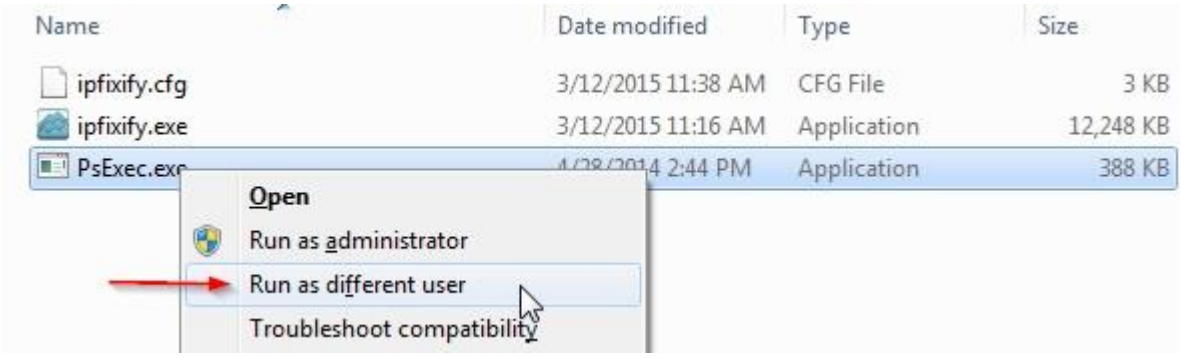

ii. Type in the IPFIXify user and password and press enter.

A. If your user does not have access to the directory that PsExec.exe is in, this will fail and you will need to provide the ipfixify user with access to the directory that PsExec.exe and ipfixify.exe are in.

iii. Agree to the PsExec EULA:

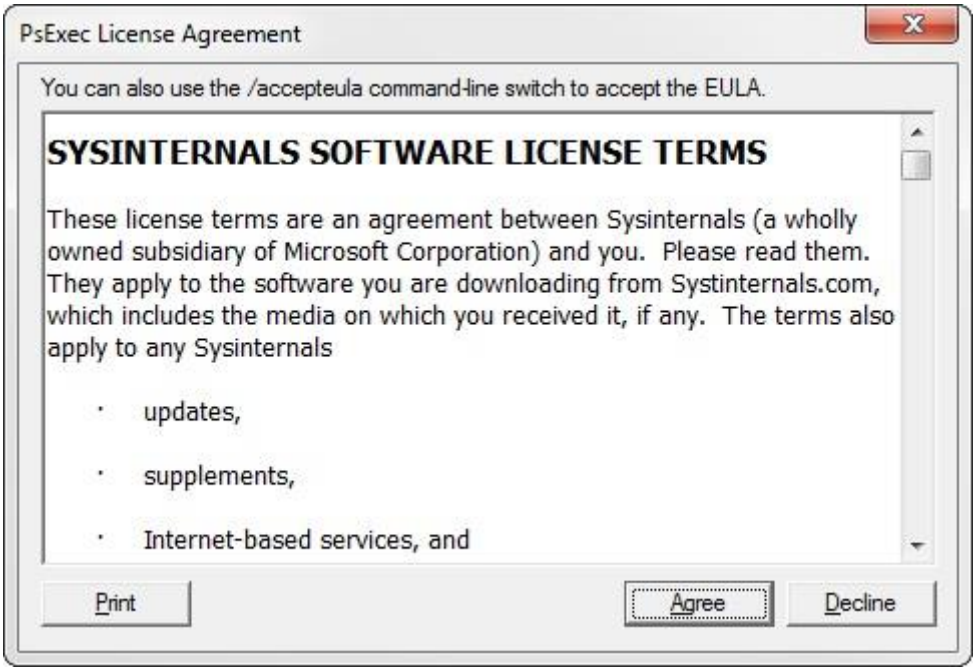

5. From an Administrative command prompt, run the following command to verify that IPFIXify has all the permissions to poll the domain controller (screenshot below):
    a. ipfixify.exe --sysmetrics --config "C:\ipfixify\ipfixify.cfg" –psexec="C:\ipfixify\PsExec.exe" -permtest IPofDC

Administrator: Command Prompt

```
C:\ipfixify>ipfixify.exe --config "C:\ipfixify\ipfixify.cfg" --sysmetrics --psexec="C:\ipfixify\psexec.exe" -permtest 10.1.5.1
--------------------------------------------------------------------------
Current CFG settings
--------------------------------------------------------------------------
$VAR1 = {
  'mode' => 'sysmetrics',
  'vitals' => 'yes',
  'storageAvailability' => 'yes',
  'user' => 'plxr\\ipfixify-noadmin',
  'pollthreads' => '5',
  'pingtimeout' => '2',
  'collector' => [
    '10.1.4.188:4739'
  ],
  'sourceip' => '10.40.0.9',
  'usernamesOnly' => 'yes',
  'psexec' => 'C:\\ipfixify\\psexec.exe',
  'members' => [
    '10.1.5.1',
    '10.1.5.2'
  ],
  'testinterval' => 300,
  'pwd' => 'xxxxxxxxxxxxxxxx',
  'eventlogs' => 'yes',
  'pingthreads' => 1,
  'columns' => undef,
  'pollTimeOut' => '300'
};

--------------------------------------------------------------------------
WMI Connectivity Test
--------------------------------------------------------------------------

Node,LoadPercentage
10.1.5.1,
10.1.5.1,
10.1.5.1,
10.1.5.1,

+ T0    10.1.5.1        PASSED

--------------------------------------------------------------------------
WMI Get Test
--------------------------------------------------------------------------

C:\Windows\SysWOW64\wbem\WMIC.exe /node:10.1.5.1 /user:"plxr\ipfixify-noadmin" /FAILFAST:2 /password:"xxxxxxxxx" OS get FreePhys



Node,FreePhysicalMemory
10.1.5.1,471624


PASSED

--------------------------------------------------------------------------
EventLog Test
--------------------------------------------------------------------------
SYSTEM : ! File exists
SYSTEM : oldest (record: 1), last (record: 28838)
SYSTEM : 28839 record(s)
SECURITY : ! File exists
SECURITY : oldest (record: 17628), last (record: 35698)
SECURITY
$VAR1 = {
  'win2003' => {
    'event_528' => '0',
    'event_538' => '0',
    'event_540' => '0',
    'event_551' => '0'
  },
  'win2008_2012' => {
    'event_4647' => '30',
    'event_4634' => '8720',
    'event_4624' => '9700'
  }
};
SECURITY : 53326 record(s)
APPLICATION : ! File exists
APPLICATION : oldest (record: 28341), last (record: 46763)
APPLICATION : 75104 record(s)

PASSED

--------------------------------------------------------------------------
File Write Test
--------------------------------------------------------------------------

PASSED

--------------------------------------------------------------------------
TEST COMPLETE!
--------------------------------------------------------------------------

C:\ipfixify>
```

6. If all the tests passed, it is time to set up IPFIXify to run as a service. In an administrative command prompt, execute the following command:
    a. ipfixify.exe --install auto --name "Scrutinizer Username Collection" --config "C:\ipfixify\ipfixify.cfg" --sysmetrics --psexec="C:\ipfixify\PsExec.exe"

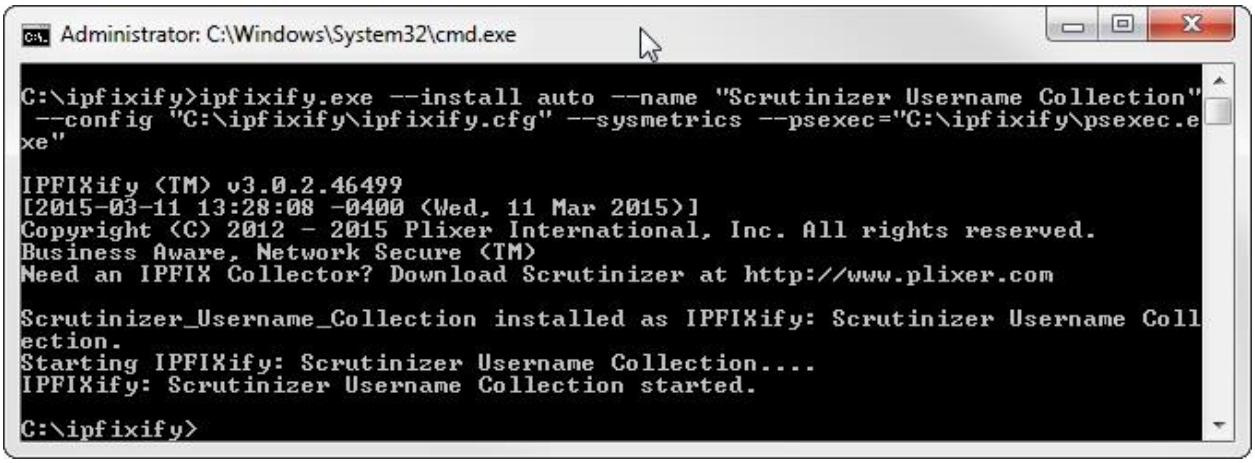

7. Lastly, the IPFIXify service has to be configured to log on as the IPFIXify user.
    a. Go to Start -> Run -> and type "services.msc"
    b. Find the service named "IPFIXIfy: Scrutinizer Username Collection", right click on it and select "Properties"

    c. Click the "Log On" tab, select "This account:", enter in the IPFIXify user and password, and then select "Apply"

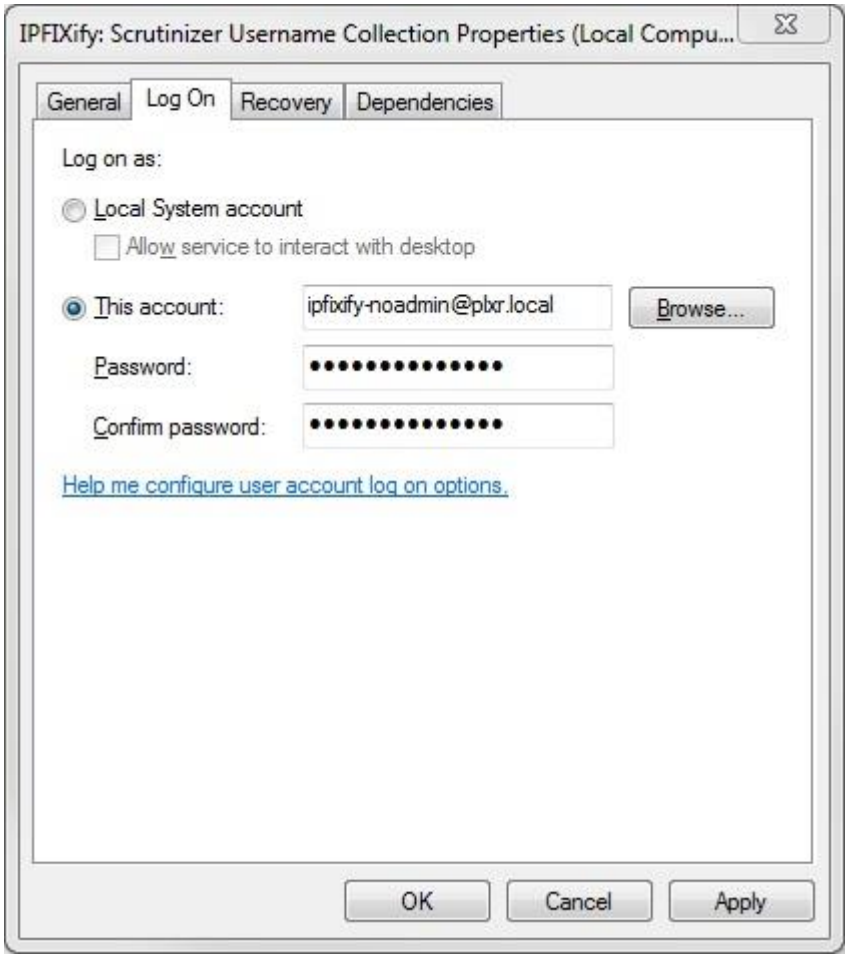

d. When you press OK, if you see a popup that says, "the user has been granted the log on as a service right", it means that the user will not maintain the log on as a service permission across reboots and you will need to provide that permission per this Microsoft document https://technet.microsoft.com/en-us/library/cc794944(v=ws.10).aspx

8. Wait a few minutes and then start checking user names in Scrutinizer.

Example IPFIXify Config

[options]

; The IP Address/Hostname and port of the IPFIX Collector(s) multiple

; collectors can be specified on additional lines

; collector=IP:PORT (e.g. 10.1.4.19:4739)

collector=10.1.4.188:4739

; When accessing remote machines, use the supplied credentials this is

; encoded. So execute the following command to manage it

```
; ipfixify.exe --credentials=<PATH/TO/CFG>
credentials=6e6ff0a30ff3d13d0f9a38a753f52f44283f9a7dfd928511dbaf2f7af1446e57981dc4628c038553
; Number of minutes between ping and WMI test of all members. The default
; is 60 minutes.
testinterval=5
; The number of seconds to try and ping a host during the process of verifying
; a member is reachable. If 0 is used, then the ping test is ignored.
pingtimeout=2
; The number of threads to gather data from the members who responded. If you
; have a small list of members, then this can be a small number (e.g. 1 - 3).
; The more threads you use, the more memory will be consumed by IPFIXify.
pollthreads=5
; If vitals is a true value, then CPU, Memory, and Number of processes running
; data is collected. To disable these statistics, comment out the following
; line.
vitals=yes
; If storageAvailability is a true value, then disk availability is collected.
; To disable these statistics, comment out the following line.
storageAvailability=yes
; If eventlogs is a true value, then System, Security, and Application
; Eventlogs are collected. To disable these statistics, comment out the
; following line.
eventlogs=yes
; usernamesOnly is used in conjunction with the eventlogs option.  If you only
; care about username integration with Scrutinizer this line should be
; un-commented
usernamesOnly=yes
; If processLists is a true value, then running processes  data is collected.
; To disable these statistics, comment out the following line.
;processLists = yes
; If processListCPU is a true value, then CPU per process data is collected.
; To disable these statistics, comment out the following line.
;processListsCPU = yes
; If netstatDetails is a true value, then netstat details are collected.
; To disable these statistics, comment out the following line.
;netstatDetails = yes
; The list below contains the current hosts being polled by the IPFIXify
; Agent. One host or IP Address per line. It is recommended to use IP
; Addresses in case there are DNS issues.
member=10.1.5.1
```

member=10.1.5.2

FAQ

Q. How do I add the 'log on as a service' right to a user account?

A. https://technet.microsoft.com/en-us/library/cc739424(v=ws.10).aspx

User Name Reporting - Cisco ISE integration

User Name Reporting options for Cisco ISE include lists of user names available at Status>Views>Cisco ISE; ability to search across all flows for user names, and also a Cisco ISE option in the Other menu in reports to see which user has generated specific traffic.

Enable ERS

The first part of Cisco ISE integration for User Name Reporting is to enable ERS (External RESTful Services) on the Cisco ISE appliance.

Supported versions of Cisco ISE are ISE 1.2, 1.3, 1.4, 2.0, and 2.1.

1. On the Cisco ISE server, create a new user with the following permissions:
    - ERS Admin
    - ERS User
    - Super Admin
    - System Admin

2. Use the following links for help enabling ERS on:
    - Cisco ISE 1.2
    - Cisco ISE 1.3
    - Cisco ISE 1.4
    - Cisco ISE 2.0
    - Cisco ISE 2.1

3. To test the configuration external to Scrutinizer, use POSTMAN to do a GET Request.
    a. Making a GET Request Using POSTMAN -
    b. Use this URL for the test - https://[ISE_SERVER]/ise/mnt/Session/AuthList/null/null
    c. Tip: When using POSTMAN, first navigate to the server with your browser, tell Chrome it is OK to use a bad certificate, and leave that window open.

Configure Cisco ISE Integration in Scrutinizer

The initial configuration on the Scrutinizer server can be performed from within Interactive scrut_util with just one command.

Login to the Scrutinizer server with administrative permissions and run the following command to open the Interactive scrut_util prompt:

/home/plixer/scrutinizer/bin/scrut_util.exe

Then, at the **SCRUTINIZER>** prompt, enter:

SCRUTINIZER> ciscoise add <ise_ip> <ise_tcp_port> <ise_user>

This command adds a CiscoISE node to the queue to acquire user identity on all active sessions.

The required parameters are the host address <ise_ip>, tcp port <ise_tcp_port>, and user <ise_user> that can access the API (the new user created in step 1 of Enable ERS).

Scrutinizer will prompt the user for the <ise_user> password.

Other scrut_util options for CiscoISE

ciscoise check

Tests polling and outputs the results to the screen for review.

ciscoise kick <ise_id> [<mac_address>] <user_ip>

Kicks the user off the ISE node forcing them to re-authenticate. Minimally the user's IP address is required. Optionally, the <mac_address> can be provided.

ciscoise nodelist

Lists the currently configured CiscoISE nodes.

ciscoise poll

Runs a poll manually and outputs the results to the screen.

ciscoise remove <ise_ip>

Removes a CiscoISE node from Scrutinizer.

The required parameter <ise_ip> is the IP address of the CiscoISE node.

ciscoise update <ise_ip> <ise_tcp_port> <ise_user>

Updates existing configuration settings for a specific CiscoISE node.

The required parameters are the host address <ise_ip>, tcp port <ise_tcp_port>, and user <ise_user> that can access the API.

Scrutinizer will prompt the user for the <ise_user> password.

For further information on CiscoISE scrut_util commands, see the Interactive scrut_util section.

AlarmsTab

Alarms Overview

Bulletin Boards

As messages come in, they are run past the Policy Manager. If the message violates a policy, it can be saved to the history table and may end up being posted to a Bulletin Board. The Bulletin Boards are used to organize alarms into categories. Each policy is associated with a Bulletin Board view.

Heat Maps

A graphical representation of the corresponding Bulletin Board table. Objects appearing in the Heat Map high and to the right are the hosts or policies that often need immediate attention. This is because those objects have the most violators and the most violations combined.

For more detailed information on Bulletin Boards and Heat Maps, go to Alarms Views

Universal Table Controls

For each View, there are some universal controls:

- **Refresh** - Refresh the Bulletin Board for the most up to date information.
- **Show X Entries** - Adjust the number of results shown in the Bulletin Board (10, 25, 50, 100, 200, 300 or 400).
- **Search or Add Filter** - Search the Bulletin Board for specific data using the two dropdown menus in the upper right-hand corner of the view. These dropdowns correspond to the Bulletin Board's specific columns.
- **Delete** *(Orphans View Only)* - The Delete button is tied to the check boxes seen in the first column of the table. Click on check boxes and click the 'Delete' button to clear entries from the Orphan table.
- **Acknowledge** *(Bulletin Board Views Only)* - The Acknowledge button is tied to the check boxes seen in the first column of the table. Click on check boxes and click the 'Acknowledge' button to clear entries from the Bulletin Board.
- **IP / DNS** - Clicking the IP or DNS button will either show IPs or resolve the DNS in the table.
- **Column Headers** - Clicking on column headers in the Bulletin Board will sort the table (ascending or descending) per column clicked.

Threat Index

The Threat Index (TI) is a single value comprised of events with different weights that age out over time. Because any one event could be a false positive, the TI gives the administrator the option of letting the summation events possibly trigger a notification when a configurable threshold is breached.

For example, if a device on the local network reaches out to the Internet to a host with a reputation of being part of a botnet, does that mean it is somehow infected? It could but, probably not. What if the same local PC also receives a few ICMP redirects from the router supporting the subnet. Now can it be discerned that there is an infection that needs to be addressed? Again, probably not but, suspicions are rising.

More Information on Threat Index can be found here

Alarming on Interface Percent Violation

To set this alarm, you need to perform the following steps:

1. Go to Admin >> Settings >> Alarm Notifications
   From here, check "Threshold Violations".
2. Go to Admin >> Settings >> System Preferences
   Make sure "Threshold - Utilization" is set.

If both the "Threshold Violations" and the "Threshold - Utilization" options are set, then alarms will be sent to the Alarms tab once a violation occurs.

In the Alarms Tab, you can view the violations under the Policy, "Scrutinizer: Interface Exceeded Threshold".

From there, you can assign any notification you want to the Policy.

Alarms Configuration

Create New Board

Creates or Removes a Bulletin Board.

Options:

- Create a new Board - Type the name of the board you wish to create. Click “Create”.
- Remove an Existing Board - Choose an existing board from the dropdown menu that you want to remove. Click “Remove”.
- Close - Click to close the modal.

Flow Analytics Settings

Brings you to Admin >> Settings >> Flow Analytics

For more information on Flow Analytics Settings, please go to Flow Analytics.

Flow Analytics Configuration

The overall status of all algorithms and the total runtime and count of violations across all algorithms.

For more information on Flow Analytics Configuration, please go to FA Configuration and Algorithm Activation Strategy.

Policy Manager

Brings you to Admin >> Definitions >> Policy Manager

The Policy Manager lists all of the Policies that are configured for the Alarms Tab.

Policies can match an alarm based on Source Address, Message Match, Syslog Alert Level.

A policy can be configured to do one of three things with an alarm:

- Post it to a Bulletin Board (Alarms posted to a Bulletin Board will also be stored in history).
- Only store in history for reporting.
- Delete the alarm (It is not available in any way).

Policies also determine if a notification should be processed for an alarm by associating alarm messages with a notification profile.

The Policy Manager Table displays:

- **Priority** - The Scrutinizer alarm policy engine compares each alarm against the defined policy list. The order they are checked is based on this priority field.
- **Check Box** - used to select one, multiple or all policies to delete.
- **Name** - Name of the policy.
- **Action** – violations can be posted to a Bulletin Board, stored to history only for future reporting, or deleted.
- **Hits** - The number of times the policy has been violated since counters were last reset.
- **Last Violation** - Date and time of the most recent violation.
- **Notification** - Type of notification.
- **Creation Info** - Date, Time and User that created the policy.

Exporters in Algorithms

A link to add or remove exporters from specific algorithms.

Hosts Excluded from Algorithms

A link to include or exclude specific hosts from an algorithm.

Edit / Create Policy

Overview

The Edit Policy interface is used to create or modify an existing policy. Policies are used to match on messages that will be saved to the history table.

Filters

- Policy Name: Name displayed in the Bulletin Board
- Active: Whether or not the Policy is active
- Message Filter: The text in the body of the message
- Source Address Filter: The host the message came from
- Alert Level Filter: Can be a combination of two fields “facility” and “severity”.
    - Facilities includes: kern, user, mail, daemon, auth, syslog, lpr, news, uucp, cron, authpriv, ftp, unknown, local0...7
    - Severity (Priority) includes: emerg, alert, crit, err, warning, notice, info, debug
- Protocol Filter: Policies can be created for incoming syslogs, SNMP traps or emails
- Excluded Machines: IP addresses that we want excluded from this policy
- IP Range: Hosts that this policy will apply to
- Notes: Information saved with the policy to help administrators remember its useful purpose

Logic

- Match (Default) : Allows for matching on text with Logical And & Or expressions. This is the most common.
- Regex (Advanced) : Requires advanced instruction. A regular expression is a powerful way of specifying a pattern for a complex search. The MySQL database uses Henry Spencer's implementation of regular expressions, which is aimed at conformance with POSIX 1003.2. MySQL uses the extended version to support pattern-matching operations performed with the REGEXP operator in SQL statements. The following does not contain all the details that can be found in Henry Spencer's regex(7) manual page. That manual page is included in MySQL source distributions, in the regex.7 file under the regex directory. In short, a regular expression describes a set of strings. The simplest regular expression is one that has no special characters in it. For example, the regular expression 'hello' matches hello and nothing else. Non-trivial regular expressions use certain special constructs enabling them to match more than one string. For example, the regular expression "hello|word" matches either the string hello or the string word. As a more complex example, the regular expression "B[an]*s" matches any of the strings Bananas, Baaaaas, Bs, and any other string starting with a B, ending with an s. For more references on Regular Expressions, you may visit the following internet pages:
    - Regexp
    - Pattern Matching
    - String Comparison Functions

Select Action

- Bulletin Board: Select and view the foreground and background colors
- History: When the policy is matched, should a message be:
    - Posted to a preferred Bulletin Board and saved to history for later reporting?
    - Stored to history for later reporting but not posted to the Bulletin Board?

    - Deleted immediately with no history on the message
- Save to Policy List:
    - Save to same order in Policy List: Save with the current policy priority (Default)
    - Save to bottom of Policy list: Saves to the bottom of the policy list and will be checked for a match last.
    - Save to top of Policy list: Saves to the top of the policy list and will be checked for a match first.

Notifications
Allows the user to select an action for a policy. See Notification Profiles to add values to the list. Select a Notification Profile or create a new one.

- Trigger
    - Threshold Trigger: This is used to notify when the amount of events exceeds the threshold. Remember it could take 10 minutes or greater than 10 months until the threshold is reached.
    - Rate Trigger: This is used to prevent notification for an event until it happens X times in Y minutes.
    - Device Specific: This is checked off when the events coming in must be from the same host in order to trigger the threshold violation alarm.
- Process Notification for:
    - First Violation: Notify once for the threshold violation and don't repeat unless the message is cleared from the bulletin board.
    - Each Violation: Notify every time the threshold is breached.

Notification Profile

Overview

Notification Profiles can be created once and applied to multiple policies or to objects in the maps that are being polled. Notification methods include Email, Logfile, Syslog, SNMPTrap, Script and Auto Acknowledge. Enter the necessary data and select additional details from the "Available Variables for Message" list to ensure the desired information is included in the alert. Specify the number of minutes to wait on triggering the selected action in "Minutes Trigger". (0 is the default.)

If multiple notification alerts are added to the same Notification Profile, the order of notification can be re-ordered by entering lower or higher numbers to the left of each notification and clicking the "Re-order Alerts" button.

Alerts:

Email – send an email alert

Enter the email address the alert is destined for in the “To” field

Note: An email server must be configured in Scrutinizer for these alerts to function. If the email server has not yet been configured in Scrutinizer, click the “Configure” button to set that up.

Logfile – add alert message to a file

Enter log file name with the absolute file path of:

/home/plixer/scrutinizer/files/logs/{logfile_name.txt}

Note: Log files must be placed at this location.

Syslog – send syslog alert to Host address.

Required fields are:

Host: Target server address

UDP Port: Target server port (default 514)

Priority

Facility

SnmpTrap – send snmptrap alert to Host address.

Required fields are:

Host: Target server

Community String

UDP Port

Enterprise OID

Generic ID

Specific ID

Binding OID

From Host

Script – trigger action defined in Script.

Required fields are:

Script: /home/plixer/scrutinizer/files/{alert_script.sh}

Note: Script must be placed in this folder and absolute path must be included in the Script field.

Command-line Arguments: Variables to include in the script from the Available Variables list below.

Auto Acknowledge – automatically acknowledge policy alarms

Policy To Acknowledge: select target policy from dropdown list

Available Variables for Message:

%m	Message
%v	Violator Address
%h	Host
%p	Protocol
%pol	Policy Violated
%notes	Policy Notes
%id	Alarm ID

Threat Index (TI)

Threat Index

The Threat Index (TI) is a single value comprised of events with different weights that age out over time. Because any one event could be a false positive, the TI gives the administrator the option of letting the summation events possibly trigger a notification when a configurable threshold is breached.

For example, if a device on the local network reaches out to the Internet to a host with a reputation of being part of a botnet, does that mean it is somehow infected? It could but, probably not. What if the same local PC also receives a few ICMP redirects from the router supporting the subnet. Now can it be discerned that there is an infection that needs to be addressed? Again, probably not but, suspicions are rising.

In the world of threat detection, reacting to any single odd behavior generally leads to tail chasing because in data networking, normal communications often lead to an occasional odd connection.

The idea behind the threat index is that they rise for an individual host each time it participates in a behavior that is suspicious. Depending on the type of behavior (e.g. scanning the network) the event may increase the TI by a higher value than others (e.g. receiving an ICMP redirect). If the Threat Index of a host hits a threshold (e.g. 100), a notification can be triggered. Keep in mind that the index is a moving value because individual events age out over time. For this reason, an IP address must reach the Threat Index threshold within a configurable window of say 14 days because the same events that increased the counter are also aging out and as a result, the individual TI will go up and down over time.

Alarms Views

Bulletin Board by Policy

In the Bulletin Board by Policy view, the alarms are grouped by policy violated.

The Heat Map in the Bulletin Board by Policy view displays the policies (E.g. threat algorithms) that are violated. Y axis = count, X axis = unique hosts.

The Bulletin Board by Policy Table displays:

- **Policy** - Policies are used to match messages that will be saved to the history table. Click on a Policy name to see all of the messages that violated the Policy from all hosts.
- **Board Name** - Policy categories.
- **Violations** – The number of times a policy has been violated. Click on an individual host to see messages from that host for the desired policy.
- **Events** - The number of events that occurred violating the policy. With Flow Analytics alarm aggregation, 1 violation may consist of multiple events.
- **TI (Threat Index)*** - *Table is sorted by TI by default. The threat index is a function a policies threat multiplier. The higher the TI, the greater the chance these policy violations are a security threat. TI = violations x threat multiplier. Click here for more information on the threat index.
- **HI (Host Index)** – The number of unique secondary IPs associated with a policy. Some algorithms have two IPs associated with the violation. For example, Network transports. If two hosts are seen using an unsanctioned transport, the source becomes the violator and the destination becomes the host. If there is 1 violator and an HI of 6 you know that a single host was communicating with 6 other hosts.
- **Violators** – The number of unique IPs that violated this policy.
- **First Event** - Date and Time of the first violation
- **Last Event** - Date and Time of the last (most recent) violation
- **Last Notification** - Notification methods include Email, Logfile, Syslog, SNMP Trap, Script and Auto Acknowledge

The alarms are grouped by Policy Violated.

Bulletin Board by Violator

In the Bulletin Board by Violator view, the alarms are grouped by violating IP address. The Heat Map in the Bulletin Board by Policy view displays the hosts that are violating policies. Y axis = count, X axis = unique policies.

The Bulletin Board by Violator Table displays:

- **Country / Group** – IF an IP is a public address we determine the IPs country. If it is not a public address we check to see if it is in a defined IP Group.
- **Users** – User is determined based on violator address. The lookup requires eventlog collection be configured. See Username Reporting for details.
- **Violator Address** - the IP and/or DNS associated with the violator
- **Violations** - Number of messages in violation.
- **Events** - The number of violating events generated per violator. With Flow Analytics alarm aggregation, 1 violation may consist of multiple events.

- **TI (Threat Index)*** - *Table is sorted by TI by default. The threat index is a function a policies threat multiplier. The higher the TI, the greater the chance these policy violations are a security threat. TI = violations x threat multiplier summed across all polices violated by that host. Click here for more information on the threat index.
- **HI (Host Index)** - The number of unique secondary IPs associated with a policy. Some algorithms have two IPs associated with the violation. For example, Network transports. If two hosts are seen using an unsanctioned transport, the source becomes the violator and the destination becomes the host. If there is 1 violator and an HI of 6 you know that a single host was communicating with 6 other hosts.
- **Policies Violated** - Name of Policy/Policies that were violated
- **First Event** - Date and Time of the first violation
- **Last Event** - Date and Time of the last (most recent) violation

The alarms are grouped by violating IP address.

Notification Queue

The Notification Queue lists the last 24 hours of notifications that were sent or currently in queue and waiting for execution.

The Notification Queue Table displays:

- **Violator Address** - the IP and/or DNS associated with the violator
- **Policy** - The associated Policy.
- **Notification** - The name of the notification sent.
- **Alert Type** - The type of notification sent (see Edit Policy for available options)
- **Status** – Whether the notification has been sent. If it is set to finished it has been processed. If it is set to available it is waiting to be processed.
- **Notes** – Additional details if available
- **Time Stamp** - Date and Time of the notification.

Once a notification is added you can specify whether it should be triggered on by rate or threshold.

- **Rate:** X alarms within Y minutes need to be seen to trigger a notification
- **Threshold:** Once there are X violations for this alarm on a BB, the notification will be sent. Acknowledging off the BB resets this.

There is also an option to decide whether a notification should be first or each:

- First means once the threshold is breached and the notification is sent, another notification will not be sent until the alarms are acknowledged off the BB.
- Each means a notification will be triggered each time the rate or threshold is met.

Device Specific determines whether the notification thresholds are for all policy violators or are handled per violating address.

For example: with device specific selected, IP address 1.1.1.1 and IP address 2.2.2.2 would each need to breach the threshold for a notification to be sent. Without device specific set, the combined alarms from those IPs would count against the threshold.

Orphans

The Orphans view lists messages that did not violate policies. From this view, new policies can be created to organize alarms.

The Orphan Table displays:

- **Time Stamp** - Date and Time of the notification.
- **Source Address** - the IP and/or DNS associated with the message source.
- **Violator Address** - the IP and/or DNS associated with the violator.
- **Log Level** - The severity and facility of the original syslog
- **Create Policy** - Click here to attach a policy to the orphaned message.
- **Message** - The orphaned message itself.

Policy Violation Overview

Lists the threats detected by Flow Analytics. Displays the Policies and the corresponding violations in the last 5 minutes, last hour, all in history.

The Policy Violation Table displays:

- **Policy Name** - The associated policy name.
- **Last 5 Min** - Number of violations in the last 5 minutes.
- **Last Hour** - Number of violations in the last hour.
- **All** - Number of total violations for the associated policy.
- **Totals** - Total numbers of the 3 previous columns across all violated policies.

Baselining

Baselining

Baselining is analyzing the performance of the network by comparing current performance to historical data (also known as the "baseline"). For example, when measuring the current traffic from an interface, alerts will be sent if any traffic exceeds the baseline. Scrutinizer can be configured to baseline any NetFlow or IPFIX element, and alert whenever that baseline exceeds a higher than normal traffic pattern.

By default, there are several baselines enabled and configured to collect data whenever a new exporter is added to Scrutinizer.

- ingressInterface and octetDeltaCount
- ingressInterface and packetDeltaCount
- egressInterface and octetDeltaCount
- egressInterface and packetDeltaCount
- applicationTag (NBAR) and octetDeltaCount
- applicationTag (NBAR) and packetDeltaCount

Using the Interactive scrut_util.exe command, baselining behavior can be modified.

How Baselining Works

Once an hour, Scrutinizer will analyze historical data and tally network performance of the configured elements for baselining. Exporters that do not have the specified baseline elements in any of the templates exported are skipped. Once the data for recent data is calculated, it is compared against the historical baseline. Alerts will be sent if one of the following occur:

- **(BLACK)** There is no historical baseline
- **(RED)** recent data exceeds the maximum value of any historical baseline value
- **(ORANGE)** recent data exceeds two standard deviations over the average baseline
- **(YELLOW)** recent data exceeds one standard deviation over the average baseline

No alert is generated if the recent data exceeds the average baseline but is less than one standard deviation over the average baseline. Baselining will collect data for one week (by default) before alerts are generated.

Lastly, a weighted average (by default 100% of the value) is applied to the most recent baselining data and stored to compare against the next baseline collection cycle. The weighted average can be modified to allow new data to have a more or less significant role when comparing future baseline values.

Baselining Hosts

Scrutinizer has the ability to baseline hosts that are specified in IP Groups (e.g. 192.168.1.0/24). More than one IP range can be specified. For more information, see the section in the manual on how to configure IP Groups.

Enabling/Disabling Baseline Globally

Using the Flow Analytics Configuration Manager (Admin Tab -> Settings -> FA Configuration), Baselining can be globally enabled or disabled.

NOTE: Baselining is disabled by default.

Using the Interactive scrut_util.exe utility for Configuring Baselines

To enable the interactive scrut_util.exe utility, run:

/home/plixer/scrutinizer/bin/scrut_util.exe

Which will open this Scrutinizer prompt:

SCRUTINIZER>

For more information on using this utility, reference the Interactive Scrut_util section.

Adding/Removing default Baselines to an Exporter

Using the following Interactive scrut_util.exe command, default baselines can be added or removed from a specified exporter.

**** Warning **** These commands will alter the behavior of Scrutinizer baseline functionality. Please use with caution.

Adding

enable baseline <exporter_ip> default

Removing

disable baseline <exporter_ip>

Adding Custom Baselines

Every install is unique. Before deploying additional baselines, contact Plixer Technical Support to assist in planning, sizing, and deploying baselines in Scrutinizer.

Using the following Interactive scrut_util.exe command, custom (manual) baselines can be added to a specified exporter.

**** Warning **** This command will alter the behavior of Scrutinizer baseline functionality. Please use with caution.

enable baseline <exporter_ip> manual <pri_element[,sec_element]> <element> <AVG|COUNT|MIN|MAX|STD|SUM> <dailyhr|busday|sameday>

This 'enable baseline' command enables custom baselines (manual) based on elements from NetFlow and IPFIX templates.

Baselining has several parameters available to customize the specific baseline data to collect with the 'manual' option.

- Replace <exporter_ip> with the exporter to collect the specified baseline data.
- <pri_element> and (optionally <sec_element>) specify which IPFIX elements will be part of this baseline. This is how the data is grouped and calculated.
- <element> must be an IPFIX element with a numeric value that the <pri_element> and/or <sec_element> will be baselined on.

  For example: To collect data on source addresses and bytes, the <pri_element> would be sourceipv4address and the <element> would be octetdeltacount.

  Another example: To collect data off multiple elements such as source address and applicationtag based on packets; the <pri_element> would be sourceipv4address, the <sec_element> would be applicationtag, and the <element> would be packetdeltacount.

  To find the numeric values for IPFIX elements in Scrutinizer, go to Status > System > Templates. Drill in on the Plixer ID to see all details regarding the elements in that template.
- <AVG|COUNT|MIN|MAX|STD|SUM> are options that are used to calculate how to measure the <element>

  AVG = Average

  COUNT = Flow Count

  MIN = Minimum Value

  MAX = Maximum value

  STD = Standard Deviation

  SUM = Sum

- Lastly, there are several ways baselines can be compared:

  dailyhr = same time frame (e.g. 1a - 2a) for each day

  busday = same time frame (e.g. 1a - 2a) for each business day, skipping weekends

  sameday = same day and time each week (e.g. 1a - 2a Mondays)

When baselining ip addresses, IP Groups must be configured with the ranges and subnets of addresses to include in the baseline. This protects the user from baselining addresses that may talk once and therefore eventually alarm as a false positive.

Machine learning takes 7 days by default.

Reset Baselines to Defaults

Using the following Interactive scrut_util.exe command, custom (manual) baselines can be reset to the default baselines for each exporter. Historical data will not be deleted. However, it will expire off based on Scrutinizer's historical settings.

clean baseline

**** Warning **** This command will purge data from Scrutinizer. Please use with caution.

Monitoring Baseline Processing

Use the following Interactive scrut_util.exe command to get the Task ID for the baseline task (to be used in the 'check task' command below)

show task baseline

Displays information regarding the baseline task.

Example:

```
SCRUTINIZER> show task baseline

+------------------+------------+---------------------+---------------------+
| TASK_NAME     | TASK_ID  | EXECUTABLE        | ARGUMENTS         |
+------------------+------------+---------------------+---------------------+
```

```
| baseline          | 224         | scrut_util.exe       | ["-baseline"]         |
+------------------+------------+----------------------+----------------------+
1 Result(s) Found

Done (0.002744 seconds)
```

The 'check task' command provides information on the baseline task processing.

check task <id>

Checks the execution times and error codes for the baseline task. The task id is available by using the show task baseline command.

Example:

```
SCRUTINIZER> check task 224

+----------------------------+---------------+--------------------+--------------------------------
| START_TIME                | RUN_TIME    | RETURN_VAL       | INFO_STR
+----------------------------+---------------+--------------------+--------------------------------
| 2016-01-25 07:10:00      | 5             | 88                 | baseline : scrut_util.exe -
| 2016-01-25 06:10:00      | 5             | 88                 | baseline : scrut_util.exe -
| 2016-01-25 05:10:00      | 5             | 88                 | baseline : scrut_util.exe -
| 2016-01-25 04:10:00      | 5             | 88                 | baseline : scrut_util.exe -
....
```

- RUN_TIME indicates how long the baselining task has run. It should complete in a few seconds.
- RETURN_VAL indicates if any error code was returned. It should be 0. Anything else, something is wrong and Plixer Technical Support should be contacted for assistance.

Baseline Reports

To monitor the progress of the baselining function, baseline reports are available on the collector's exporter. These reports are accessible by following these steps:

1. In the Status tab, under Device Explorer, filter on 'Scrutinizer'
2. Click on 'Scrutinizer'
3. Then click on 'Reports'

4. Then click on ‘Baselines’
5. This will provide a list of available auto-created Baseline reports on your server.
    - Select a report to view
    - These reports will show if Baselining is working as configured.

DashboardTab

Dashboard Tab

Overview

The Dashboard is used to display the favorite gadgets of the user. Dashboards can be customized per login account. Click on the **+** sign in the upper right to expand the panel to add and remove gadgets. Close the panel by clicking on the red x icon in the upper right.

Multiple sub tabs are allowed in Dashboards. Click on "Add a Tab" and then click twice on the newly added tab to rename it. Notice the down arrow next to the selected tab. If clicked, a menu appears:

- Layout Options: select the tab to have 1-5 columns or free form.
- Make Default: make the tab the default tab for Dashboards
- Delete this tab

Each small window is called a gadget. Notice that the header of each gadget can be grabbed with the mouse which allows the window to be moved around Dashboards. Click on Edit and enter the amount of minutes the gadget is to refresh with the latest information.

You can set Dashboards to be the default page when you login. Visit the user settings in the Admin tab to modify the default.

A custom gadget can be added to Dashboards. Click on the + sign in the upper right corner of Dashboards. A box appears, click on the icon to the right of the top drop down menu. Enter the details (e.g. Name and URL and click Save. The gadget will be added to the current Dashboard. Video training on Setting Up Dashboards is also available.

Vitals Dashboard

The Vitals dashboard is created by default in the Admin user's dashboard during a new install. This dashboard provides vital information on how well the servers are handling the NetFlow and sFlow volume and other server metrics. Vitals information is reported for all servers in a Distributed Environment.

The following dashboard gadgets are available in this dashboard:

- **CPU Utilization:** Average CPU utilization for the Scrutinizer server(s).
- **Memory Utilization:** This gadget displays how much memory is available after what is consumed by all programs on the computer is deducted from Total Memory. It is not specific to NetFlow being captured.
    - NOTE: The flow collector will continue to grab memory depending on the size of the memory bucket it requires to save data and it will not shrink unless the machine is rebooted. This is not a memory leak.
- **Storage Available:** The Storage report displays the amount of disk storage space that is available. After an initial period of a few weeks/months, this should stabilize providing that the volume of NetFlow stays about the same.
- **Flow Metric by Exporter:** The following metrics are provided per exporter:
    - MFSN: Missed Flow Sequence Numbers. NOTE: Sometimes MFSN will show up as 10m or 400m. To get the dropped flows per second, divide the value by 1000ms. A value of 400m is .4 of a second. 1 / .4 = 2.5 second. A flow is dropped every 2.5 seconds or 120 (i.e. 300 seconds/2.5) dropped flows in the 5 minute interval displayed in the trend.
    - Packets: Average Packets per second
    - Flows: Average Flows per second: This is a measure of the number of conversations being observed.
        - NOTE: There can be as many as 30 flows per NetFlow v5 packet (i.e. UDP datagram) and up to 24 flows per NetFlow v9 datagram. With sFlow, as many as 1 sample (i.e. flow) or greater than 10 samples can be sent per datagram.
- **Flow Metric by Listening Port:** The above metrics are also available per listening port. The flow collector can listen on multiple ports simultaneously. The defaults are 2055, 2056, 4432, 4739, 9995, 9996 and 6343, however, more can be added at Admin->Settings->System Preferences->Listener Port.
- **Database Statistics:** Provides the following database metrics:
    - Connections by Bytes: Excessive connections can result in reduced performance. NOTE: other applications using the same MySQL database will cause this number to increase.
    - Read Req: The number of requests to read a key block from the cache. A high number requested means the server is busy.
    - Write Req: The number of requests to write a key block to the cache. A high number of requests means the server is busy.
    - Cache Free: The total amount of memory available to query caching. Contact support if you find that your query cache is presently under 1MB. For more information, please reference the MySQL Documentation.

- Queries: Tracks the number of queries made to MySQL. More queries indicates a heavier load to the MySQL server. Generally there will be spikes at intervals of 5 minutes, 30 minutes, 2 hours, 12 hours, etc. This indicates the rolling up of statistics done by the stored procedures. This Vitals report is important to watch if the NetFlow collector is sharing the MySQL server with other applications.
- Threads: Threads are useful to help pass data back and forth between Scrutinizer and the database engine. The MySQL Server currently manages whether or not to utilize the configured amount of threads. For more information, please reference the MySQL Documentation.
- Buffers Used: Key Buffers Used - indicates how much of the allocated key buffers are being utilized. If this report begins to consistently hit 100%, it indicates that there is not enough memory allocated. Scrutinizer will compensate by utilizing swap on the disk. This can cause additional delay retrieving data due to increased disk I/O. On larger implementations, this can cause performance to degrade quickly.
- Users can adjust the amount of memory allocated to the key buffers by modifying the \scrutinizer\mysql\my.ini file and adjusting the key buffer size setting. A general rule of thumb is to allocate as much RAM to the key buffer as you can, up to a maximum of 25% of system RAM (e.g. 1GB on a 4GB system). This is about the ideal setting for systems that read heavily from keys. If you allocate too much memory, you risk seeing further degradation of performance because the system has to use virtual memory for the key buffer.

- **Syslog Received and Processed:** Syslog activity for the servers is provided in this gadget.

Custom dashboard gadgets can be created for any of the other Vitals Reports that are listed in the Vitals Reporting section. The Vitals Dashboard can also be copied to another user, or recreated by selecting the desired gadgets from the gadget panel.

Flow Analytics

Flow Analytics

Overview

Flow Analytics brings the following additional features to Scrutinizer:

- **Functions as a Network Behavior Analysis system** by constantly monitoring all flows for behaviors that could be compromising the health of the network (networks scans, illegal applications, P2P, etc.). It interrogates every flow from every host from selected flow exporting devices for suspicious patterns and anomalies. All flows across selected flow sending devices are monitored at all times.
- **Numerous** additional reports that provide more detailed information on the flows received.
- **DNS resolution** can be enabled to occur automatically to support Domain reporting. To enable DNS resolution, and to control the how long the names are retained in a local cache, visit the Admin tab -> Settings -> System Preferences. Note that this feature places additional load on the system, so monitor the CPU use before and after enabling to ensure proper performance.
- **Performs threshold watches** for saved reports. FA can monitor for nearly any combination of flow characteristics and export a syslog if a match or a high/low threshold is reached.

The FA algorithms that are focused on security events have been rewritten in 12.1 with the objectives of providing alarms that are focused on providing more actionable information while significantly reducing false positives.

A significant change in the way alarms are generated is that Scrutinizer 12.1 has introduced the concept of "internal IPs" – those IP addresses, or assets, that comprise your network, and "external IPs", which is the public Internet and other IP address spaces that are not under your administration. Alarm messages now identify both the source and destination of suspect activity as "external to internal", "internal to external", or "internal to internal", as well as providing additional details that are specific to the alarm type.

Internal Addresses are defined as any addresses entered as your IP Groups plus the following list of IP Address blocks. These addresses are private, non-routable addresses per RFC 1918 and link-local addresses per RFC 3927:

10.0.0.0/8

172.16.0.0/12

192.168.0.0/16

169.254.0.0/16

All other addresses are treated as external IPs.

There are three new algorithms introduced in 12.1, including TCPOddFlags, TCPScan and UDPScan. These are detailed below. Two algorithms (nefarious and unfinished flows) have been deprecated as their functionality has been wrapped into other algorithms.

New FA Features in 15.5

Plixer continues to enhance Flow Analytics to provide more actionable information with fewer false positives. The emphasis for this release was to improve the information display in the alarm manager and to include new capabilities based on FlowPro Defender. Changes for this release of Scrutinizer and FlowPro Defender include:

- Aggregation of repeating alarms into a single concise view
- Changes to the Alarm Manager detail display

  - New columns
  - Aggregated alarm display
  - FlowView displays for the specific alarm

- New alarms focused on integrating detection of suspicious DNS traffic
- Changes to existing alarms

Aggregated Alarms

This version of Scrutinizer supports a new alarm concept called an "aggregated alarm". In previous versions, Scrutinizer would issue a new alarm every five minutes for an event that might last for several hours. Aggregated alarms will combine these events into a single alarm, displaying the original alarm time, the most recent alarm time, and the number of times that the alarm has triggered while it was active. Aggregated alarms will continue to collect matching alarm events until the "Aggregated Alarm Timeout" time has expired with no new alarms. The "Aggregated Alarm Timeout" defaults to two hours (120 minutes). This value is controlled on the Admin / Settings / Flow Analytics page. A value of zero will disable the alarm aggregation.

Alarm Manager – Detailed Alarm Display
Several new columns are displayed in the alarm manager detail display to streamline alarm response and management, as well as to support aggregated alarms. The following columns are new or changed. Also review the "Aggregated Alarm" discussion below:

- **Users**: a new column that displays the user(s) associated with the violator address while the alarm is active.
- **Alarm Time:** This is the time the alarm time, or the time the alarm was first issued in the case of aggregated alarms. This was previously titled "Timestamp".
- **Recent Activity:** A new column that displays the most recent time an alarm was observed for an aggregated alarm. This will display "N/A" for a single alarm incident.
- **Duration:** A new column that displays the time an aggregated alarm has been active. This is the difference between the Recent Activity time and the Alarm Time. This will display "N/A" for a single alarm incident.
- **Events:** A new column that displays the number of individual five minute periods an aggregated alarm has been active without a break in activity longer than the "Aggregated Alarm Timeout".

FlowView Displays for Alarms

The Alarm FlowView display has been simplified and tailored for the specific alarm selected. Clicking the dropdown box beside the alarm message allows the selection of the alarm FlowView. The columns in this selection are used by the algorithm that triggered the alarm, and the first columns displayed are typically the connection 5-tuple (source address, destination address, source port, destination port, protocol). This is a concise view of the flows involved for the alarm and avoids horizontal scrolling.

FA Navigation

The navigation for FA is via gadgets in the Dashboard tab. The primary gadget "Flow Analytics Configuration" should be added to Dashboard.

- At the top, it displays the overall time to run all algorithms and the total count of violations across all algorithms.
- **Name**: This is the name of the algorithm that is checking for abnormal behaviors.
- **Time**: This is the amount of time the algorithm takes to run across all selected routers/switches.
- **Count**: This is the number of violations found the last time the algorithm ran. Click on the trend to view graphs for longer time periods.
- **Time exceeded**: Algorithms that exceed the configured run time will be cancelled.

FA Configuration

Setting Up Flow Analytics (FA):

FA algorithms run sequentially. By default, they do not run against any NetFlow exporters until the NetFlow exporters are added to the selected algorithms.

To add exporters to an algorithm, visit ***Admin > Settings > FA Configuration***:

- Click on the algorithm name in the Flow Analytics Configuration table to configure, change settings and thresholds, and apply exporters and exclusions to the algorithm
- Most of this configuration can also be accomplished by clicking the appropriate link for the algorithm in the Flow Analytics Configuration table

IMPORTANT NOTES:

- Add only a few routers to a few algorithms initially and start off slowly. Pay attention to the Vitals of the server. After 15-30 minutes add few more routers to selected algorithms and slowly ramp up the FA deployment.
- FA has only 300 seconds (i.e. 5 minutes) to finish all enabled algorithms. If it can't finish in 300 seconds, it will stop where it is and start over. All algorithms must finish within 5 minutes as the process repeats every 5 minutes. Optimize performance by paying attention to the Time each algorithm takes to run as well as the overall time shown at the very top of the Flow Analytics Configuration gadget.

Algorithm Activation Strategy

Algorithm Name	Internal / Core Routers	Edge Routers	Public Facing IP Addresses Defined In IP Groups
BotNet Detection	FlowPro Defender	FlowPro Defender	N/A
Breach Attempt Detection	Yes	Yes	Yes
DDoS Detection	No	Yes	Yes
Denied Flows Firewall	Yes	No	No
DNS Command and Control	FlowPro Defender	FlowPro Defender	N/A
DNS Data Link	FlowPro Defender	FlowPro Defender	N/A
DNS Hits	Yes	Yes	Yes
Domain Reputation	FlowPro Defender	FlowPro Defender	N/A
DRDoS Detection	No	Yes	Yes
FIN Scan	Yes	Yes	Yes
Host Reputation	No	Yes	Yes
ICMP Destination Unreachable	Yes	No	No
ICMP Port Unreachable	Yes	No	No
IP Address Violations	Yes	Yes	Yes
Malware Behavior Detection	FlowPro Defender	FlowPro Defender	N/A
Multicast Violations	Yes	Yes	Yes
Null Scan	Yes	Yes	Yes
Odd TCP Flags Scan	Yes	Yes	Yes
P2P Detection	Yes	Yes	Yes
RST/ACK Detection	Yes	Yes	Yes
SYN Scan	Yes	Yes	Yes
TCP Scan	Yes	Yes	Yes
UDP Scan	Yes	Yes	Yes
XMAS Scan	Yes	Yes	Yes

Algorithms for Public Facing IP Addresses: These addresses should be defined as an IP Group, which will cause these addresses be treated as part of your protected network. Algorithms such as DDoS will not trigger an alarm unless the target of the DDoS is an internal address (defined within an IP Group).

If your primary concern is internal – internal and internal – external monitoring, then enable algorithms on the core routers and ensure that any routable IP Addresses you want monitored as part of your internal network are defined within an IP Group. Monitoring your internal – internal and internal to external traffic is highly recommended for identification of traffic patterns that may indicate a compromised asset, and to assist with incident response.

If your primary concern is monitoring your public facing assets, ensure that all public-facing IP Addresses are contained within an IP Group, and you would prefer to use your edge routers for most algorithms.

Algorithms and Gadgets

FA Algorithms may or may not include gadgets. Some algorithms are enabled by default. Others need to have selected flow exporting devices added to them. A few algorithms need to have thresholds configured or modified from the defaults.

FA Gadgets that can be added to Dashboard:

- **Flow Analytics Configuration**: The overall status of all algorithms and the total runtime and count of violations across all algorithms. Algorithms can be ordered alphabetically or by order of execution. LEDs in this gadget are as follows (refresh the gadget in the upper right corner):
    - Name - The name of each algorithm executed by Flow Analytics. Expand Flow Analytics Overall Status. If the status of "Disable All" is toggled, be sure to click Save at the bottom on the gadget. Some algorithms look for network threats and others simply list the Top x across multiple flow exporters. Some algorithms such as Top Network Transports do both.
    - LED Colors:
        - Yellow - incomplete run (time limit caused the algorithm not to run during the last cycle)
        - Lite Green - successfully completed on the last run
        - Gray - disabled
        - Trend - actively executing the algorithm
        - Dark Green - successfully completed on the current run
    - Time - the time given to each algorithm to run. Some algorithms need more time to run depending on the number of flow exporting devices included and the number of flows exported by each device.
    - Count - the number of violations the last time the algorithm ran.
- **Network Volume**: The scale of the traffic traversing through the core network. It lists the volume of unique traffic on the network for the last 5 minute Vs. last 30 hours. Only include a few core routers/switches.
- **Flow Analytics Devices**: Select the flow exporting devices that each algorithm will run against. Enter text and click 'Filter' to find specific devices. Click the 'Clear' button to remove the filter and display all devices. Some algorithms are run against all tables created by flow exporting devices while others are only run against one or two tables (e.g. routers). The "Select Flow Devices" gadget 'scrut_fa_devices.cgi' can be added to Dashboard however it is not necessary because it is best utilized as a popup.
- **Flow Analytics Exclusions:** Exclude hosts from selected algorithms to help prevent false positives. Some hosts will constantly violate the threshold of certain algorithms. This interface helps prevent false positive alarms by allowing selected hosts (i.e. IP addresses) to be excluded from violating one or more algorithms. The "Exclude Hosts" gadget 'scrut_fa_exclusions.cgi' is not necessary in Dashboard as it is bested utilized as a popup.
- **Flow Reports Thresholds:** Saved reports that are given a threshold to compare against every five minutes show up in this gadget.
- **Threats Overview**: Gives Network Administrators an idea on the frequency that each Flow Analytics algorithm is being violated. The colors indicate the frequency within each time interval: Last 5 min, Last Hour and All.
- **Top Applications**: Top Applications on the network. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.

- **Top Conversations**: Top Conversations across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Countries**: Top Countries across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Rev 2nd lvl Domains**: Top reverse 2nd level domains across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Flows**: Top Flow sending end systems across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Hosts**: Top Hosts sending data across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Network Transports**: Top Transport Layer Protocols across selected flow exporting devices. Alarms trigger for protocols that appear that haven't been approved. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Subnet Traffic**: Top IP Subnets across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Well Known Ports**: Top ports across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.

NOTE: Some include algorithms that should only run against core routers/switches. Watch the *Flow Analytics Overall Status* gadgets for algorithms that need more time to run.

FA Algorithms that don't include Gadgets:

Be sure to exclude certain hosts from select algorithms to avoid false positives. This can easily be done from the alarms tab as well by clicking on the host. The interface will prompt for the exclude confirmation.

- **Indicator Correlation Event:** This algorithm escalates multiple Indicator of Compromise (IOC) and security events for a single host to a new alarm on the security event BB. While a single IOC may be indicative of malware, it is much more likely to be a real security concern if there are multiple indicators. By default, This algorithm correlates multiple IOCs along with any events posted to the Security Event BB and issues a alarm for any host that has three or more entries in the IOC and Security Event bulletin boards. Each of the contributing algorithms will be listed in the alarm message.

  By default, three different algorithms are required, but you may adjust this in the threshold setting for Indicator Correlation.

- **Breach Attempt Detection:** This algorithm is examining flow behaviors that may indicate a brute force password attack on an internal IP address. This is accomplished by examining the flow, byte, and packet counts being exchanged in short-duration completed flows between one source and one destination, with specific behaviors observed for common attack vectors such as SSH, LDAP and RDP. If the number of flows that match these characteristics exceeds the alarm threshold, an alarm will be raised. The default flow count threshold is 100. Either IP address can be excluded from triggering this alarm. This algorithm is enabled by default across all flow exporting devices that are exporting the necessary fields.
- **DDoS Detection:** Identifies a Distributed Denial of Service attack targeted at your protected network space. DDoS attacks are often launched by a BotNet, and "reflection attacks" are becoming more common. Scrutinizer may identify attacks against your network as "reflection attacks" if they meet the criteria.

  There are four settings that are used to adjust the sensitivity of the DDoS detection algorithm:

- DDoS Packet Deviation (default: 10) and DDoS Bytes Deviation (default: 10) - These settings control how similar the flows associated with the attack must be. The standard deviation of the byte count and packet counts associated with the flows must be less than this setting.
- DDoS Flows (default: 4) controls the minimum number flows used to identify attacking hosts. The sensitivity of the DDoS attack can be reduced by increasing this setting to six or higher.
- DDoS Unique Hosts (default: 200) controls the threshold for the minimum number of hosts that have sent flows that match the other characteristics required to trigger the alarm.

- **Denied Flows Firewall:** Triggers an alarm for internal IP addresses sending to external IP addresses that cause greater than the threshold of denied flows. The default threshold is set to 5 denied flows. Either the source or destination IP address can be excluded from triggering this alarm.
- **DRDoS Detection:** Identifies a Distributed Reflection Denial of Service attack targeted at your protected network space. DRDoS attacks are often launched by a BotNet, and "reflection attacks" are becoming more common. Scrutinizer may identify attacks against your network as "reflection attacks" if they meet the following criteria.

  Scrutinizer detects the following ten Distributed Reflection Denial of Service (DRDoS) Attacks -

  - DNS
  - NTP
  - SNMP
  - SSDP
  - Chargen
  - NetBIOS Name Server
  - RPC Portmap
  - Sentinel
  - Quote of the Day
  - Trivial File Transfer Protocol

  You also have the option to enable or disable a specific reflection attack via Admin / Settings / Flow Analytics Configuration / DRDoS Detection / Settings.

- **DNS Hits:** Triggers an alarm when a host initiates an excessive number of DNS queries. This identifies hosts that perform an inordinate number DNS lookups. To do this, set the flow threshold to a large value that reflects normal behavior on your network. The default threshold is 2500 DNS flows in five minutes. Either the source or destination IP address can be excluded from triggering this alarm.
- **FIN Scan**: Alerts when a FIN scan is detected. FIN scans are often used as reconnaissance prior to an attack. They are considered to be a "stealthy scan" as they may be able to pass through firewalls, allowing an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.
- **Host Reputation**: This algorithm maintains a current list of active Tor nodes that you should monitor. Some malware families use Tor for Command and Control communications. White-list your users who are authorized to use Tor and regard other uses as suspicious. This algorithm will also monitor any IP address lists that you provide as a custom list as described in the “Custom List” section that follows.
- **Custom List:** This algorithm also supports the use of custom lists where you can add additional reputation lists to the system. A custom list of IP addresses can be imported into the Host

Reputation Algorithm. To do this, Host Reputation needs to be enabled and this list will be compared with traffic on the devices selected under "Configured Flow Analytics".

- **To Enable Host Reputation:**
  - Go to Admin > Settings > FA Configuration.
  - Expand the "Host Reputation Monitor".
  - Make sure Disabled is not checked and exporters are being included.
- **To Create Your Custom List:**
  - The IP Addresses need to be in a file with a single address on each line.
  - The name of the file will be the ***Threat Category Name*** and the ***Alarm Policy Name***.
  - The file must have a .import file extension; for example: "CustomThreatList.import".
  - The file must be placed in the /home/plixer/scrutinizer/files/threats/ directory.
  - Once an hour, this file will be imported into Scrutinizer and used for the next hour of processing.
  - To force a new file import to become active immediately, run:
    - /home/plixer/scrutinizer/bin/scrut_util.exe
    - SCRUTINIZER> download hostreputationlists
  - After the import, the Alarms Policy can be modified to change the threat_multiplier from the default of 0.

- **ICMP Destination Unreachable**: This alarm is generated when a large number of ICMP destination unreachable messages have been sent to the suspect IP address. This may happen as a result of scanning activity, misconfiguration, or network errors. ICMP Destination Unreachable is a message that comes back from a destination host or the destination host gateway to indicate that the destination is unreachable for one reason or another. The default threshold is 100 destination unreachable messages. Either the source or destination IP address can be excluded from triggering this alarm.
- **ICMP Port Unreachable**: This alarm is generated when a large number of ICMP port unreachable messages have been sent to the suspect IP address. This may happen as a result of scanning activity, misconfiguration, or network errors. ICMP Port Unreachable is a message that comes back from the destination host gateway to indicate that the destination port is unavailable for the transport protocol. The default threshold is 100 port unreachable messages. Either the source or destination IP address can be excluded from triggering this alarm.
- **IP Address Violation**: This algorithm compares the traffic from included exporters against a list of allowed subnets. If both source and destination addresses are outside of an allowed subnet, an alarm will be triggered. A common use of this algorithm is to identify unknown or unauthorized internal network addresses that are communicating with the public Internet.
- **Medianet Jitter Violations:** This algorithm compares the jitter values as reported by the Medianet flows to the threshold defined by the user in the Settings section of this algorithm. The default threshold is 80ms.
- **Multicast Violations**: Any multicast traffic that exceeds the threshold that isn't excluded will violate this algorithm. The default threshold is 1,000,000 and the minimum that can be set is 100,000.
- **NULL Scan**: Alerts when a NULL scan is detected. NULL scans are a TCP scan with all TCP Flags cleared to zero. This scan is often used as reconnaissance prior to an attack. They are considered to be a "stealthy scan" as they may be able to pass through firewalls, allowing an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan your internal network,

such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.

- **Odd TCP Flags Scan:** Alerts when a scan is detected using unusual TCP Flag combinations. These types of scans may allow an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.
- **P2P Detection**: Peer to Peer (P2P) traffic such as BitTorrent are identified by this algorithm. The default threshold is a P2P session involving over 100 external hosts, which will detect most P2P applications. However, there are several P2P applications that are stealthier, so you may want to experiment with lower thresholds or periodically lower the threshold to about 20 to determine if other "low and slow" P2P traffic is on your network.
- **RST/ACK Detection**: Alerts when a large number of TCP flows containing only RST and ACK flags have been detected being sent to a single destination. These flows indicate that a connection attempt was made on the host sending the RST/ACK flow, and was rejected. This algorithm may detect other scan types used by an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.
- **SYN Scan**: Alerts when a SYN scan is detected. SYN scans are a TCP scan with the TCP SYN Flag set. This scan is often used as reconnaissance prior to an attack as it is fast and somewhat stealthy. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.
- **TCP Scan:** Alerts when a possible TCP scan is detected from an exporter that does not provide TCP Flag information. These types of scans may allow an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.
- **UDP Scan:** Alerts when a possible UDP scan is detected. These types of scans may allow an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.
    - NOTE: if your policy allows P2P traffic on your network, then you will likely want to exclude the allowed host(s) or disable this alarm as it will often detect P2P control traffic as a UDP Scan violation.
- **XMAS Scan:** Alerts when a XMAS scan is detected. XMAS scans are a TCP scan with the FIN, PSH, and URG TCP Flags set. This scan is often used as reconnaissance prior to an attack. They are considered to be a "stealthy scan" as they may be able to pass through firewalls, allowing an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.

Note: By default, all of the scan algorithms are looking for "internal to internal" and "internal to external" scanning activity. You now have the option to control which scanning directions the different algorithms look for, and including the addition of "external to internal", which would normally be used to monitor public

facing IP addresses listed in an IP Group. Within each of the scanning algorithms, the settings screen provides a directional control option.

FA Algorithms that Require FlowPro Defender

- **BotNet Detection**

  (Formerly named NXDomain detection)

  This alarm is generated when a large number of unique DNS name lookups have failed. When a DNSlookup fails, a reply commonly known as NXDOMAIN is returned. By monitoring the number of NXDOMAINs detected as well as the DNS name looked up, behavior normally associated with a class of malware that uses Domain Generation Algorithms (DGAs) can be detected.

  The default threshold is 100 unique DNS lookup failures (NXDOMAIN) messages in five minutes. Either the source or destination IP address can be excluded from triggering this alarm.

- **DNS Command and Control**

  This algorithm monitors the use of DNS TXT messages traversing the network perimeter as detected by FlowPro Defender. DNS TXT messages provide a means of sending information into and out of your protected network over DNS, even when you have blocked use of an external DNS server. This technique is used by malware as a method of controlling compromised assets within your network and to extract information back out. Additionally, some legitimate companies also use this method to communicate as a means to "phone home" from their applications to the developer site.

  The algorithm will detect inbound, outbound, and bidirectional communications using DNS TXT messages. Thresholds may be set based either on the number of DNS TXT messages or the number of bytes observed in the DNS TXT messages within a five minute period. The default setting is for any detected traffic to alarm, and alarm aggregation defaults to 120 minutes.

  To suppress alarms from authorized applications in your network, you may add the domain generating the alarm message to the "trusted.domains" list on FlowPro Defender. See the discussion on "trusted.domains" list below.

- **DNS Data Leak**

  This algorithm monitors the practice of encoding information into a DNS lookup message that has no intention of returning a valid IP address or making an actual connection to a remote device. When this happens, your local DNS server will fail to find the DNS name in it's cache, and will pass the name out of your network to where it will eventually reach the authoritative server for the domain. At that point, the owner of the authoritative server can decode the information embedded in the name, and may respond with a "no existing domain" response, or return a non-routable address.

  FlowPro Defender uses proprietary detection algorithms to identify suspicious DNS names that may contain encoded data, and passes this information to Scrutinizer where it is processed by the DNS Data Leak algorithm. Thresholds may be set based either on the number of suspicious DNS names or the number of bytes observed in the suspicious DNS name within a five minute period. The default setting is for any detected traffic to alarm, and alarm aggregation defaults to 120 minutes.

- **Domain Reputation**

  Plixer is introducing domain reputation with this release. Domain reputation provides much more accurate alarming with a dramatic decrease in the number of false positive alarms as compared to IP based Host Reputation. The domain list is provided by Plixer and is updated each hour and currently contains over 400,000 known bad domains.

  FlowPro Defender performs the actual monitoring, and when it detects a domain with poor reputation, it passes the information to Scrutinizer for additional processing. The default setting is for any detected traffic to alarm, and alarm aggregation defaults to disabled so that all DNS lookups observed will result in a unique alarm.

To suppress alarms from authorized applications in your network, you may add the domain generating the alarm message to the "Trusted Domain" list on FlowPro Defender. See the discussion on FlowPro Defender for additional details.

- **Malware Behavior Detection**

  This specific alarm is correlating IP address lookups (i.e. what is my IP address) activity which is commonly performed by malware shortly after the initial compromise with the detection of the BotNet alarm or with a Domain Reputation alert. In other words, this algorithm looks for the following correlation:

  - IP address lookup combined with a Domain Reputation trigger
  - IP address lookup combined with a BotNet trigger

  When either of the two events is detected, this algorithm is triggers as this behavior is a very strong indicator of a compromised asset.

- **Malware Domain Communications**

  This algorithm combines the Domain Reputation algorithm with communications detected going to the IP address that was resolved. Scrutinizer and Defender have detected the following sequence of events:

  1. Defender contains a list, updated every 10 minutes, of over 700,000 known malware domains created by forensic analysis of the actual malware software. These are very high confidence domains.
  2. Defender monitors all of the DNS resolution requests, and generates an IOC alert on detection of a match with a malware domain and saves the resolved "Malware IP Address". This only rates an "IOC" as a browser may "pre-fetch", or resolve address, for all of the links on a web page to make browsing faster. However, as yet, no connection to the malware site has been made.
  3. Scrutinizer then examines all flows for any communications with the "Malware IP Address" resolved by Defender. On detection of any flows to or from that address, a connection to the malware site has been established, and a Malware Domain Communication alert is issued.

  Note: For this algorithm to work, the user must turn ON host indexing. This setting is available in Admin / Settings / System Preferences

Correlation Algorithms

These algorithms demonstrate Plixer's cyber threat correlation capability. Correlation of multiple network behaviors over a long time period provides detection systems with more information allowing for a higher accuracy with fewer false positive alarms.

There are currently three Correlation Algorithms available in Flow Analytics:

- Indicator Correlation Event
- Malware Behavior Detection
- Malware Domain Communications

Trusted Domain List

A "trusted domain list", often called a whitelist, is preconfigured on FlowPro Defender to suppress alarms involving specific domains. The default whitelist contains five entries that you can add to or remove as best fits your environment.

- mcafee.com
- sophos.com
- sophosxl.net

- webcfs03.com
- apple.com

mcafee.com suppresses DNS Data Leak alarms from McAfee AntiVirus software. McAfee encodes information from the anti-virus clients on your network into very long and complex DNS names and captures this information at their DNS server. This is exactly the type of behavior that the DNS Data Leak algorithm is looking for as this technique is also used by some forms of malware.

sophos.com and **sophosxl.net** are related to the Sophos Anti-virus software, and it uses multiple techniques to get information in and out of your network using DNS. In addition to using the same technique as McAfee to send information back to their servers, they also use DNS TXT messages to send information back into the clients on your network. Use of DNS TXT messages to exchange information with an external host is also used by some malware families, and the DNS Command and Control algorithm will alarm on this type of activity. This will prevent Sophos from generating either DNS Data Leak or DNS Command and Control alarms.

webcfs03.com belongs to SonicWALL, and will also generate DNS Data Leak alarms.

apple.com uses DNS TXT messages to apparently exchange settings with their NTP server. This will alarm as a DNS Command and Control alarm.

You may have other authorized software within your networks that use DNS to bypass your firewall for data communications. If so, you will want to add the domain(s) involved to the Trusted Domain list as described below. Once configured for your environment, any other traffic using DNS to communicate will be worth additional investigation.

To Modify the Trusted Domain List:

1) Log on to the FlowPro Defender
2) Enter: "edit trusteddomains"
3) Modify the file contents as desired
4) enter control-x, and select "Y" to save your changes
5) press enter to accept the file name.
6) quit

Untrusted Domain Lists

FlowPro Defender supports both the use of a domain reputation list that is downloaded from Plixer, as well as allowing you to create or enter your own list.

Plixer Domain Reputation List

FlowPro Defender downloads a list of domains from Plixer once each hour. These are domains that have been determined to be "bad domains" with a high probability, and this list is used in the "Domain Reputation" and "Malware Behavior Detection" algorithms. Use of this list can be controlled by the FlowPro Defender:

1) Log on to the FlowPro Defender
2) Enter: "edit plixer.ini"
3) To enable the list (default is enabled), set the value enableDomainReputationList=1
4) or, to disable the list, set the value enableDomainReputationList=0
5) enter control-x, and select "Y" to save your changes
6) quit

User Defined Domain Lists

You may augment the Plixer Domain Reputation list and create one or more domain lists that contain any domains that you want to alarm on. The rules for the domain lists are:

1) The DNS name must contain at least 2 labels, which is often called a second level domain, or 2LD for short (for example, google.com) and no more than 3 labels (maps.google.com), or a 3LD.
2) The labels must contain between 1 and 63 characters, as is required to be a legitimate domain name.
3) Entries that do not match these requirements will be ignored.

To create your own list of domains to detect domainReputation violations:

1) Log on to the FlowPro Defender
2) Enter: "edit my_domain_list_name" NOTE: do NOT enter a file extension. This will be automatically assigned.
3) Modify the file contents as desired
4) enter control-x, and select "Y" to save your changes
5) press enter to accept the file name.
6) quit

To Enable a Domain List

1) Log on to the FlowPro Defender
2) show domainlists
3) Enter: "enable domain_list_name"
4) quit

To Disable a Domain List

1) Log on to the FlowPro Defender
2) show domainlists
3) Enter: "disable domain_list_name"
4) quit

IMPORTANT NOTE: Hosts can easily be excluded from certain algorithms by clicking on the IP address in the Alarms Tab. This will popup the Exclude Hosts table where the IP address can then be excluded from other algorithms as well.

FA Bulletin Boards

FA algorithms are posted to one of three Bulletin boards, depending on their importance and likelihood that a host has been compromised. The different bulletin boards are:

- Policy Events
- Indicators of Compromise (IOC)
- Security Events

Policy Events

Algorithms that post to this bulletin board are those that detect network traffic that generally have no direct security implications, but may violate your network policy. The following algorithms are posted to the Policy Events BB by default:

- Excessive Jitter
- Multicast Violations
- IP Address Violation
- P2P Detection

Indicators of Compromise (IOC)

The algorithms that post to the IOC bulletin board are those that indicate possible malware activity, but with insufficient confidence to initiate a security event alarm. This is a bulletin board that you may want to periodically review to look for recent changes that may be of concern. If multiple IOCs are associated with a single host, these may generate an "Indicator Correlation Event" alarm that is posted to the Security Event BB. This is discussed in more detail below. The following algorithms are posted to the IOC Events BB by default:

- Breach Attempt Detection
- DNS Hits
- ICMP Port Unreachable
- ICMP Destination Unreachable
- Denied Flows
- Domain Reputation
- FIN Scan
- NULL Scan
- Odd TCP Flags Scan
- RST/ACK Scan
- SYN Scan
- TCP Scan
- UDP Scan

- XMAS Scan

Security Events

The algorithms that post to the Security Events bulletin board are high confidence detections that a host is compromised. You should monitor and investigate any events posted to this bulletin board. The following algorithms are posted to the Security Events BB by default:

- BotNet Detection
- DDoS Detection
- DRDoS Detection
- DNS Command and Control Detection
- DNS Data Leak
- Indicator Correlation Event
- Malware Behavior Detection
- Malware Domain Communications
- Host Reputation (Tor, Blackhole, Malware C&C Server, and user defined)

Note: You may edit the policy for any algorithm to change the Bulletin Board from the default settings.

Flow Analytics Custom Algorithms

Developers and administrators can now create their own algorithms that Flow Analytics will execute routinely whenever the pre-packaged algorithms are executed.

The level of complexity in a custom algorithm will vary based on the task(s) that are being performed. An in depth knowledge of Perl and Scrutinizer's database structure are required.

Getting Started

There are important files and directories that are involved with creating and managing custom algorithms.

- scrutinizer/files/algorithms is the directory where all custom algorithms are placed
- scrutinizer/files/algorithms/example.pm is an example of a commented Perl module to help individuals understand the work flow of an algorithm
- Interactive scrut_util is used to add and remove custom algorithms from the Flow Analytics Engine, and also display a list of custom algorithms
- scrutinizer/bin/fa_cli.exe is used to test custom algorithms

A Perl installation on the Scrutinizer server is not required to write algorithms.

Creating a Custom Algorithm

Start by making a copy of example.pm and pasting it into the scrutinizer/files/algorithms directory. Rename the file to something that is unique to help identify the custom algorithm. For purposes of an example, this documentation will reference the custom algorithm as customAlgo1.pm.

Modify customAlgo1.pm to perform the actions intended for this algorithm. You will also have to change the package name at the top of the example script to match the file name of your new algorithm module. When complete, save the file. To test the algorithm, it first must be added to the Flow Analytics engine.

It is up to the algorithm to send out alerts. There is an example of how this is accomplished in the example.pm file.

Adding a Custom Algorithm

Before a custom algorithm can be tested or executed on a regular basis, it must first be added to the Flow Analytics Engine. This is accomplished by running the following command from the Interactive scrut_util prompt.

1. To open the Interactive scrut_util prompt, enter:

   /home/plixer/scrutinizer/bin/scrut_util.exe

2. At the **SCRUTINIZER>** prompt, enter:

   SCRUTINIZER> enable custom_algorithm customAlgo1 My_Custom_Algorithm

   NOTE: Do not include the file extension for 'customAlgo1' in the command and the custom algorithm name (My_Custom_Algorithm) cannot contain spaces.

3. To verify that the custom algorithm has been added, launch the Flow Analytics Configuration Manager (Admin Tab -> Settings -> Flow Analytics Configuration).

Testing a Custom Algorithm

During the testing phase, it is highly recommended that debug is added to ensure everything is working as intended. For production environments, comment out any debug code.

To test a custom algorithm, execute the following command:

```
/home/plixer/scrutinizer/bin/fa_cli.exe --debug 1
```

Configuring a Custom Algorithm

At any time during production use or testing, custom algorithm settings can be managed from the User Interface by launching the Flow Analytics Configuration Manager (Admin Tab -> Settings -> Flow Analytics Configuration).

Settings available to custom algorithms are:

- Threshold
- Exporters to include
- Hosts to exclude
- Enable/Disable Alerts
- Enable/Disable Syslogs
- Enable/Disable Algorithm

Running a Custom Algorithm in Production

When algorithms are run routinely in production, the Flow Analytics Engine manages the execution and exporting of violation count and time taken to execute the algorithm. Flow Analytics will terminate any algorithm that is taking too long to execute. The Flow Analytics Configuration Manager will indicate if any algorithm, including custom algorithms, were terminated before completion.

Creating a Policy for a Custom Algorithm

Once the custom algorithm is violating and sending alerts that show up in the Scrutinizer Orphan View, a policy can be created to match violations and send alerts using notifications. (See Alarm section for more details).

Follow the steps below to create policies for custom algorithms (e.g. customAlgo1):

1. Navigate to Alarms > Configuration > Policy Manager
2. Click New Policy
3. Give the policy a Name
4. Under the message filter, type in *FA:userCustom_customAlgo1*
5. Set any other option for this policy (for more details, reference the Policy Manager section of the documentation)

NOTE: Remember to replace customAlgo1 with the name of the custom algorithm.

The next time a violation occurs, the policy will trigger and perform whatever actions were configured in the policy.

Disabling a Custom Algorithm

Custom algorithms can be disabled from the Flow Analytics Configuration Manager (Admin Tab -> Settings -> Flow Analytics Configuration) by checking the 'Disable' checkbox for the Custom Algorithm.

Deleting a Custom Algorithm

Custom Flow Analytic Algorithms can be deleted from the Flow Analytics Engine by executing the following Interactive scrut_util command.

This command deletes a custom_algorithm at the system level.

1. To open the Interactive scrut_util prompt, enter:

   /home/plixer/scrutinizer/bin/scrut_util.exe

2. At the **SCRUTINIZER>** prompt, enter:

   SCRUTINIZER> delete custom_algorithm custom_algo1

3. To verify that the Algorithm has been deleted, launch the Flow Analytics Configuration Manager (Admin Tab -> Settings -> Flow Analytics Configuration).

Display a list of Custom Algorithms

To display a list of custom algorithms available and whether they are enabled, execute the following command from the Interactive scrut_util interface:

SCRUTINIZER> show custom_algorithms

Optimizing FA

Flow Analytics can be optimized in several different ways:

1. Modify the number of flow exporting devices included in the algorithm
2. Disable selected Algorithms
3. Utilize a second or third copy of Scrutinizer with FA.
4. Contact your vendor to learn about the minimum hardware requirements.

MapTab

Maps Main View

Overview

The Mapping options are primarily utilized by administrators to display all or a portion of the network topology. Right mouse button on the background of a Plixer map (i.e. not Google map) and toggle between Edit mode and View mode to rearrange the icons. Be sure to save changes before leaving the Edit mode. Don't forget to save the position of the icons.

Maps can be created by clicking on Create Maps / Device Groups at the top of the tree menu in either the Maps or the Status tabs, or by clicking the New Group button in Admin > Settings > Maps & Device Groups. Map configuration can be completed in the Maps & Device Groups page, or by right-clicking in the background of an existing map and selecting Map Settings from the menu.

A default map per login account can be selected under Admin -> Security -> User Accounts.

There are several types of images used in the maps:

- Objects: come in four formats.
- Backgrounds: can be added to the maps.
- Links Status: can be read about here.

The device icon color is based on the Fault Index (FI) value in CrossCheck and the corresponding color thresholds. The link color between devices can be based on utilization between the devices. Click on the link to bring up a flow report for the connection.

IMPORTANT: Click on the map and then use the right mouse button to bring up a menu of options.

Map Connections

Overview

Connections between objects:

- A Connection between any two objects can be created using this interface.
- Selecting a **From** Device which is sending flows will cause the **Interface** drop-down box to fill in with the corresponding flow interfaces available.
- Selecting a Group or Icon **From** object results in an empty **Interface** drop-down box. Check off "Display all interfaces in this group" to fill in the **Interface** drop-down box with all interfaces from devices in the group. You can also select "Connect with black line" to connect to the **To** Object.
- Click the **Connect** button and the connection will be displayed in the window below.
- **IMPORTANT**: When creating connections for a Google map, a device name might be followed by (Needs GPS coordinates - Go to Objects Tab). Devices in a Google Map Group will not appear until they are given GPS coordinates or an address using the Objects tab.

Links Status comes in 3 formats:

- Flow link: are links representing flow capable interfaces.
    - Link colors can be green, yellow, orange or red and are based on settings configured in the Admin Tab -> Settings -> System Preferences.
    - Links are blue if there is no bandwidth statement for the interface.
    - Links are dashed gray if flows are not received within the last five minutes from the interface. Click on a link to bring up the current flow information.
- Black line: is a static link between two devices. It is not clickable and doesn't provide a status.
- Saved Reports: connections between objects can be made with existing Saved Reports. The threshold limits for the link color change will be set per Saved Report connection.

Additional Notes on Links:

- Label: displays the percent utilization or the bits received in the last average interval. Set the "average interval" under Admin tab ->Settings -> System Preferences "Status/Link Average". The default is 5 minutes.
- ALT tag: over the Label displays the full interface description.
- Arrow: on the link reflects highest utilization direction.
- Clicking: on the link will bring up the default user preference report on the link for the last few minutes (5 minutes by default) in one minute intervals. Outbound or Inbound traffic is displayed depending on the direction of the arrow when clicked.

Map Dependencies

Overview

Dependencies are configured for a few reasons.

- Up Dependencies (i.e. these devices are 'Up' Dependent on: this device) are configured to prevent excessive notifications behind a single down device. For example, if the poller has to go through Router 1 to reach Routers 2-50 then, routers 2-50 (i.e. These Devices) should be "Up dependent on" Router 1. The result, if Router 1 stops responding to the poller, the administrator will be notified for Router 1 and not Routers 2-50. With Up Dependencies applied, it would be assumed by the poller that routers 2-50 are unreachable and therefore, polling to routers 2-50 would stop until Router 1 responds to a poll.
- Up Dependencies will stop the polling on all of the devices that are Up dependent on a down device. This logic helps maintain optimal availability reports on the devices behind the down device.
- Down Dependencies (i.e. these devices are 'Down' Dependent on: this device) are configured to avoid polling unless a specified device stops responding to a poll. For example, Router B should not be polled unless Router A stops responding to a poll. This configuration would mean that Router B is 'Down' Dependent on Router A. This scenario is desirable in a high availability network configuration.
- Down Dependencies will start the polling on a device that is Down dependent on a down device. For example, if Router A goes down, Polling on Router B would commence until Router A responds to a poll.

Configuring dependencies is a best practice for scaling the network polling process while managing the volume of notifications sent when an outage occurs. If possible, they should be configured all the way back to the poller. This practice limits, the risk of excessive notifications and aids the administrator in isolating the exact location of an incident.

Google Maps

Overview

Google maps can be used to display the topology and connectivity of the network. Visit the connections window to link objects together or to link to groups.

Map Groups

Overview

Groups are the foundation of all maps. Creating a new group creates a map. There are two types of maps:

- Plixer: these maps are entirely local to the Scrutinizer server, do not require any internet access.
- Google: Useful for displaying network devices geographically.

Highlights:

- Flow devices can be added to more then one group/map.
- Flow devices added to groups are removed from Ungrouped.
- Membership: use this to add devices and objects to the group.
- Pass up map status: use this to pass the status of any down devices in a lower map up to parent map.
- Permissions can be set on Group visibility

Map Objects

Overview

Objects are placed in groups. Each Group is a map. Generally Objects on the map represent flow exporting Devices however, polled devices can be added as well. Objects have several properties:

- Label: a read only field determined by the collector
- Poll Using: IP Address, Hostname or disable
- Notification: Specify how you want the alert on the status on the object/device to be sent out.
- Primary Status: This is the background color of icons throughout Scrutinizer. The default primary color of an icon is "Flow". That is an indication of whether we are still receiving flows from an exporter. To change primary status, edit an object under Mapping Configuration and change "Primary Status".
- Secondary Status: This is the colored square that is superimposed on an icon. The secondary or sub icon color is based on CrossCheck status for a device. If the primary status is CrossCheck then there won't be a secondary status.
- Icon image: shape of the icon
- Dependencies: are used to determine how and when the device is polled.
- Membership: Specify the groups / maps the object is a member of.

Modify an Objects Membership to place it in another group/map.

Objects come in four formats:

- **Devices**: are imported from CrossCheck or can be manually added here. These objects change color based on the Fault Index and the threshold settings in CrossCheck. To remove a device that was imported from CrossCheck, the 3rd Party Method must be disabled or the device must be removed from the 3rd Party Method script or removed from the 3rd party application else it will continue to be re imported after deletion.
  - Label: If the device is imported from CrossCheck, this value is imported.
  - Poll Using: Select, IP Address, Hostname or Disable Polling.
  - Notification: Select a Notification Profile which will be triggered when the CrossCheck Fault Index threshold is breached.
  - Icon: The default icon type and size can be modified.
- **Groups**: represent other maps and the status of devices in those maps. They are clickable and bring up the appropriate map.
- **Symbols**: represent devices you want on your maps that don't display a status. They can be assigned labels and made clickable to launch other applications and/or web pages.
- **Text Boxes**: can be placed on maps and generally contain text. Shapes, colors and size can all be defined. As well as the Label and a clickable link. Text boxes are for Plixer maps only. They cannot be placed on Google maps.

NOTE: to modify the Google address of an Object, select a map the Object is on and then edit the Object. Since the same Object can be in multiple maps with different addresses, the map must be selected first. The 'Address' listed is generally the mailing address of the location of the Object. Google uses this 'Address' to locate the GPS coordinates. The actual GPS coordinates can also be manually edited.

Adding custom Device icons:

- **Object Icons**: Save graphic icons to the ~/scrutinizer/html/images/maps directory with the naming convention of <name>_object.gif. Make sure the background of the image is transparent or it may not look very good on the map.
- **Device "Status" Icons**: Save device icons to the ~/scrutinizer/html/images/maps directory with the naming convention of <name>_red.gif and <name>_green.gif. You must have two icons, one for up status (green) and one for down status (red). Make sure the background of the images are transparent or it may not look very good on the map.

Plixer Maps

Overview

The maps can be viewed by clicking on the View this Map link in the tree menu.

Backgrounds can be added to the maps by selecting them in the drop down box. The size of a background image dictates the length and width in pixels of the map. Gray scale background images are ideal as colorful maps distract from the messages displayed by the software.

Save background images to the ~/scrutinizer/html/images/maps/backgrounds directory.

Images can be in .gif, .jpg, or .png format.

NOTE: Maps with background images autoscale to the size of the background image.

Device Overview

Overview

When clicking on a device within a map, will launch the Device Overview. This window will display:

- The SNMP information
- Integration with 3rd party applications
- Applications associated with the device as determined by CrossCheck
- The three busiest interfaces
- Response Time and Availability Trends if the device is being polled
- Any outstanding alarms on the device

Map Settings

Overview

The Map Settings are used to set defaults for all maps:

- Google Maps:
    - Zoom Level: set when using the option "Save Zoom & Position" in a Google map. By default, Google maps auto scale to fit all icons on the map. This option overrides Auto with a favorite position on the map. To undo the Save Level, select 'Auto' and click 'Save'.
- Plixer Maps: Map Settings are available in Admin > Settings > Map & Device Groups by clicking on a map name, or by right-clicking in the background of a map view and selecting Map Settings option.

StatusTab

Status Tab

Overview

The Status tab is one of the most popular views for gaining quick access to the status of all the NetFlow capable devices and interfaces on the network. Choose from one of several views. This tab and a corresponding view can be made the user's default under Admin tab > Security > User Accounts > User Preferences.

The submenu of the Status tab includes:

- **Views**
- **Trends**
- **Search**
- **System**
- **Vendor Specific**

Views

The main window of the Status tab is where the views are displayed.

Top left in the main window is an Options icon (just to the left of the view header). Clicking the Options icon will open a selection of options for the current view. Option selections vary per view, however, most views include:

- **Close when options change** checkbox: When checked, the Options slideout will close after every change/selection made. It is checked by default. Unchecking this option will leave the Options slideout open until you either click the Options icon, or click outside of the Options slideout.
- **Show entries**: select number of entries to display in list
- **Refresh** button: refreshes the data in the main window
- **IP/DNS** toggle button: toggles between IP Addresses and resolved host names.

In the top right of the main window (when in a view), are 3 icons just to the left of the Help (?) icon:

- **@**: Email the view
- **PDF**: Generate a PDF version of the view
- **CSV**: Export the view to a CSV file

Available Views are:

- **CrossCheck**: Provides a list of devices imported from 3rd party applications and a link to:
    - **CrossCheck Summary**: Summary report based on data from the CrossCheck List
- **Device Status**: This view displays the top x exporting devices. All columns are both filterable and sortable.

  Columns included are:

- Device: clicking on the Device icon opens the Device Overview screen for that device. Clicking on the device IP address / host name will open the Report Menu for that device. Selecting a report from here will run the report for ALL interfaces of the device resulting in the inbound traffic matching the outbound traffic. For this reason, this report is displayed inbound by default. The default timeframe for the report is Last 24 hours in 30 minute intervals.
- Last Flow Received
- Template Count: number of flow templates received for this exporter. Clicking an entry in this column opens the Templates list for that device.
- Packets/sec: clicking an entry in this column opens the Vitals > Flow Metrics/Exporter report
- Flows/sec: clicking an entry in this column opens the Vitals > Flow Metrics/Exporter report

- **Interfaces**: This view displays the top x interfaces on the network where flows are being received and is the default Status tab view upon installation. All columns are sortable, and filterable fields include: Device, Flow Version, Group, and Interface.

  The Options slideout also includes Units selection (Bits, Bytes, Packets, or Percent), and Refresh rate for this view in addition to the standard options.

  Columns included are:

  - Device: The flow exporters are the routers, switches and other devices that the collector is receiving flows from. Green means the collector is currently receiving flows from the device. Red means the device has stopped sending flows to the collector for the interface.

    Further information about the device can be obtained from the following actions:

    - Mouse-over on the numerical rank will display an Alt tag including when last flows were received for the device.
    - Mouse_over the Device icon will show an Alt tag with Fault Summary data for the device.
    - Clicking on the Device icon will open the Device Overview screen.
    - Clicking on the version of flows received (E.g. N9, N5, I10) opens a Report menu for the device which includes a Flow Stats report for the device.
  - Interface: clicking on the Interface will open the Report menu for that interface. Selecting a report from here will run an Inbound/Outbound (Bidirectional) report for the last 24 hours in 30 minute intervals.
  - Inbound
  - Outbound
    - The Inbound and Outbound columns represent utilization over the last 5 minutes. Clicking on them will prompt the user to run a report for the last 5 minutes in 1 minute intervals.
    - Clicking on the down arrow to the left of the Inbound/Outbound columns launches a menu including:
      - Reset high watermark(s). Mouse over an Inbound/Outbound value to see the High Water Mark information in the Alt tag.
      - Device Details
      - Device Overview

- **Service Levels**: Response times and availability reports based on polling.
- **Usernames**: This view displays any username information collected from exporters such as Cisco ASA, SonicWALL firewalls, or authentication servers such as Active Directory, RADIUS, etc.

All columns are filterable and sortable and include:

- ipaddress: IP Address the username was seen on
- domain: Domain the username was seen on
- username: User name
- datasource: Source of the username
- First Seen: timestamp when first seen
- Last Seen: timestamp when last seen

Trends

This submenu includes the following All Device reports.

- **Applications Defined**
- **Flow Volume**
- **Hosts (Destination)**
- **Hosts (Source)**
- **Interfaces**
- **Protocols**
- **Run Report Wizard**: Used to create reports with filters

Search

- **Exporters**: This search provides flexible searching for a specific host address, username, wireless host or SSID across some, or all, flow exporters for a specified timeframe.
- **Hosts**: This search is a quick search of all flow exporters for a given host IP Address. Search results returned include a listing of all of the flow exporters that the IP Address has been seen on. Selecting an exporter from the list provides reporting for the last hour of flow collection.
    - This search requires that Host Indexing be turned on in Admin>Settings>System Preferences.

System

- **Available Reports**: provides a list of all Report Types available, including:
    - Action: menu of actions available in this view
        - View: Report Template Details
        - View: Templates
    - Group Name: the Report group they are a member of
    - Report Name: the Report name as displayed in the Report menus
    - reportTypeLang (API ID): the report type ID used by the Reporting API
    - Template Count: How many flow templates have been received that meet each report's requirements
- **Flow Report Thresholds**: quick view of all flow reports that have had thresholds assigned to them including violation counts and timestamp of last violation.

- **Templates**: a listing of all flow templates received from incoming flow records and what device(s) the templates were received from.
- **Vitals**: quick link to the Vitals dashboard

Vendor Specific

The Vendor Specific reports include all flow exporters that meet the criteria for the report. The default interval for these reports is 5 minutes. One minute intervals are available by drilling in on an entry.

- Cisco **ISE**: This option provides a list of all usernames with host addresses, MAC addresses, and domains that were gathered from Cisco Identity Services Engine (ISE)
- Cisco **Interface Compression**: This lists all Cisco interfaces exporting compression information ordered by highest octet compression percentages.
- Cisco **Jitter by Interface**: This lists all Cisco interfaces exporting VoIP jitter information and ordered by most jitter.
- Cisco **Jitter by Interface (EzPM)**: This lists all Cisco interfaces exporting VoIP jitter information via EzPM templates, and ordered by most jitter.
- Cisco **Latency by Interface**: This lists all Cisco interfaces exporting TCP latency information and ordered by most latency.
- Cisco: **Carrier Grade NAT**: This option allows you to search for NAT'd addresses by IP address, port, and timeframe.
- **DPI Applications**: This provides a top applications report gathered from the deep packet inspection application elements exported in the flow data.
- **Endace Probes**: This option allows you quick access to the data collected by the Endace device.

The left frame includes the Device Explorer, Current Report and Saved Reports tabs and is explained below.

Device Explorer

- **Create Maps / Device Groups**: Used to create groups / maps of devices that are currently in 'Ungrouped'. A device can be a member of multiple groups.
- **Groups**
    - **Ungrouped**: By default all flow exporting devices are placed in Ungrouped until they are moved into one or more user created groups.
    - **Group Name**: Create a new Group by selecting "Create Maps / Device Groups"
        - View This Map: Displays the map for the devices in the group.
        - Modify Membership: Modify which flow exporting devices belong to the group.
        - Group CrossCheck: View CrossCheck for the devices in this group.
        - Service Levels: View the Service Level Report for the flow exporting devices in this group.
        - Show Interfaces: Show all active interfaces for the flow exporting devices in this group. The interface list will display in the main window of this screen.

- Exporters: Only devices that are exporting flows show up in the left column. The color of the icon represents the selected primary status for the object. The sub icon represents the fault index value for the device in CrossCheck.
    - Change Address: only appears if the device / object has been given a GPS address. Clicking on this will bring up the Map -> Objects tab for the device where the address or GPS information can be modified.
    - Reports: This option will open a report menu for the exporter. From this menu, select a report to display flow data. Selecting a report from here will run the report for ALL interfaces of the device resulting in the inbound traffic matching the outbound traffic. For this reason, this report is displayed inbound by default. The default timeframe for this report is Last 24 hours in 30 minute intervals.
    - Interfaces: Displays a list of interfaces for the device within the Device Explorer window. Selecting an interface (or the All Interfaces) from this list will open a report menu for that interface. Reports are displayed with the last 24 hour timeframe. ALL Interfaces reports will default to Inbound as described above, selecting a single interface will report on both Inbound and Outbound.
    - Device Overview: Provides the overall status of the device by leveraging data from CrossCheck, the poller and the alarms.
    - Show Interfaces: Shows all active interfaces for the device. This interface list will display in the main window of this screen.
    - Other Options
        - Alarms: Displays the outstanding alarms for the device.
        - Device Details: launches the Device Details view which lists SNMP details about the device including the interface speeds.
        - Flow Templates (Advanced): displays the templates currently being received from the device.
        - Device Overview: Provides the overall status of the device by leveraging data from CrossCheck, the poller and the alarms.
        - HTTP: Launch a web browser to the device.
        - FTP: Launch a FTP session to the device.
        - SSH: Launch a Secure Shell connection to the device.
    - Third Party icons: If any Third Party integration is enabled at Admin > Definitions > 3rd Party Integration, their icon(s) will display here for each exporter.

Current Report

This tab opens when you select a report from the Report menu and includes the following report icons:

- **Clear**
- **Save**
- **Save As**
- **Schedule an emailed report**
- **Place this report in a Dashboard**
- **Print**
- **Export report to a CSV file**

- **Export report to a PDF file**
- **Email this report**

Also included in this view are:

- **Report**: textbox to enter Report name
- **Templates Used**: includes count of templates used by the report and expands to show which templates are used
- **Devices Used**: includes count of devices used by the report and expands to show which devices are used
- **Report Details**: clicking this button opens the Report Details modal with the following tabs:
    - Collector Details
    - Exporter Details
    - Filters: view/edit/remove existing and add new filters to the report
    - Threshold: view/edit/remove existing thresholds or add a threshold to the report
- **Filters**: this button provides the same ability to view/edit/remove existing filters and add new filters to the report as the Filters tab in the Report Details modal described above.
- **Threshold**: this button provides the same ability to view/edit/remove existing thresholds or add a threshold to the report as the Threshold tab in the Report Details modal described above.

Saved Reports

This tab lists any reports that you saved and provides report folder management tools with "Add Folder | Manage Folders"

- **Add Folder:** when you select 'Add Folder' a text box will open to enter a report folder name to use for organizing your saved reports.
- **Manage Folders:** when you select 'Manage Folders', a new browser tab will open to Admin > Reports > Report Folders. From here you can do bulk folder/saved report management, moving several reports into/out of a report folder. You can also create new and delete report folders from here.
- **Saved reports list:** Following the list of report folders (if any) will be the list of any reports that have been saved. Each saved report has two icons:
    - Trash can: to delete the saved report. Deleting the report will also delete any dashboard gadgets or scheduled reports associated with this saved report.
    - Magnifying glass: hovering over this icon will open an Alt tag providing the parameters that the report was saved with, such as who created the report, the date range of the report and other information defining this report. Also included at the top of the Alt tag is the Report ID, which is required for some advanced functions.

Report Folder management is also available from within the Saved Reports tab by dragging and dropping the reports into or out of the desired folders.

Reports can be viewed by clicking on the report name, and can be renamed once the report is in view mode by editing the report name and clicking the Save icon.

The dynamic filter just below the Saved Reports header allows you to easily find reports within the report list or folders.

CrossCheck

Overview

The CrossCheck and the Service Level Reports located in the Status tab provide very important roles in Scrutinizer's architecture. CrossCheck provides the overall status of a device across multiple applications including 3rd parties. The Service Level Report (SLR) provides availability and response time reports using data collected by the poller.

CrossCheck Logic

Each 3rd Party Method in CrossCheck queries the related application (e.g. Flow, Poller, Denika, WhatsUp, etc.) for the devices in its database and finds or adds the device to the CrossCheck list. Duplicate IP addresses are removed. Each 3rd Party Method applies a value to the device. The Fault Index (FI) is the total across all 3rd Party Methods. All 3rd Party Methods are open source and can be modified.

After each 3rd Party Method completes, CrossCheck (I.e. mapping.cgi) performs the following routine:

- Using the data from all 3rd Party Methods, it updates the FI for each device
- It queries the above list for any device that has violated the configurable FI threshold
- For each device over the FI threshold, it sends a syslog to the alarm server which will violate the "Exceeded CrossCheck Fault Index" policy
- Any hosts that are up are removed from the xcheck_notifications table (see below)

The "Exceeded CrossCheck Fault Index" (Policy) in the Alarms tab, performs the following when a CrossCheck syslog comes in:

- The syslog comes into alarm process and violates the "Exceeded CrossCheck Fault Index" policy which triggers the notification profile "XC Notification" (XC)
- XC looks to see if an entry exists in the table called xcheck_notifications
    - If an entry exists for the device (i.e. IP address), this means a notification already went out. The time stamp is then updated.
    - Else notification hasn't already gone out. The IP address of the device is inserted into the xcheck_notifications table. The configured notification profile in mapping for the host is then executed.
    - Else, if the configured notification profile in mapping for the host isn't configured: nothing happens

- **CrossCheck List**: This utility lists all applications monitoring or receiving data from devices on the network and displays an overall Fault Index (FI) across all 3rd Party Methods (e.g. Denika, *WhatsUp Gold, **Solarwinds Orion, etc.). It provides an inventory of devices being monitored by various tools (i.e. 3rd Party Methods). The Fault Index value is used to provide the status of devices throughout Scrutinizer including the maps.
    - Click on the headings to sort.
    - The number to display (e.g. 25) followed by pagination.
    - The query time frame is the last 1 minute by default.
    - Thresholds: This sets the threshold at which the sub icon on a device changes color. Color thresholds are based on percentage of the Fault Index threshold.
    - 3rd Party Methods: Data collected from other applications that will impact the Fault Index (FI) for a device.
        - Flow: Method used to determine if a device is sending flows.
        - Poller: Method used to determine if the poller can ping the device.
        - Denika: Method used to see if the device is being SNMP polled by Denika.
        - WhatsUp: Method used to see if a device is in WhatsUp and the corresponding status.
        - Contact your vendor to create new methods (e.g SNMPc, InterMapper, etc.)

    - Fault Index (FI): The FI provides the overall status of a device across all 3rd Party Methods
    - Overview: Provides a small window with the number of devices discovered in each 3rd Party Method as well as the overall total FI for the devices in the application.

- **CrossCheck Summary**: This is a summary report for the CrossCheck List.
    - The bar chart displays the number of unique hosts found across all 3rd Party Methods as well as the Fault Index.
    - Define Networks: Specify a subnet to group IP addresses found across all 3rd Party Methods.
    - Thresholds: This sets the threshold at which the sub icon on a device changes color. Color thresholds are based on percentage of the Fault Index threshold.

- **Service Level Report**: This is a dual purpose report listing device availability and response time. In the Service Level Report:
    - Click on the headings to sort.
    - The number to display (e.g. 25) followed by pagination.
    - Click on the Poller to run a trend across all devices.
    - Click on the Device to run an availability or response time report on the selected device.
    - Click on the Round Trip Time column value or the Availability percent value to run a report.

NOTE: To remove a device that was imported from CrossCheck, it must be removed from the 3rd Party Method script or removed from the 3rd party application (e.g. Denika(TM), WhatsUp Gold(TM), Solarwinds Orion(TM), etc.).

CSV Reporting

Overview

When checked off, a CSV of the report contents will be sent with the emailed report. This module is often used for service provider usage based billing.

Flow Hopper

Flow Hopper provides end to end visibility into the path a flow took through the network on a router hop by hop basis. Because multiple paths exist between devices, leveraging traceroute or routed topology information may not provide the exact path taken by an end to end flow. Flow Hopper displays the correct path at the time of the flow.

This connection solution requires that most if not all of the flow exporting devices in the path be exporting NetFlow v5 or more recent to the collector. **This feature requires nexthop routing information as well as read only SNMPv2 or v3 access to the router.**

If Flow Hopper determines that an asymmetric flow path exists (i.e. a different route is taken on the return path), the GUI will draw out the connection accordingly. Admins can click on each router or layer 3 switch in the path and view all details exported in the flow template. Changes in element values (e.g. DSCP, TTL, octets, etc.) between ingress and egress metered flows are highlighted.

Flow View

Overview

The Flow View is used to view flows from the received templates at the very basic level with little formatting. The Flow View interface retrieves all of the flows that match the values requested in the URL from the database for the time frame specified.

Notice:

- Filters are passed to Flow View when drilling in. Click on the icon at the top of the page to display the contents of the filters passed.
- Use the filters drop down box to find data in specific columns.
    - The sourceOrDestination option is not a column. It is used to filter flows to or from a specific host.
- Click on the column headings to sort.
- Click on the icon in the top left to view which filters have been passed to the Flow View.

Flexible NetFlow and NetFlow Security Event Logs

Flow View is used to view non-traditional NetFlow as generated by Flexible NetFlow, NetFlow Security Event Logs, etc.

The collector can save any type of NetFlow v1, v5, v6 and v9 data inclusive of Flexible NetFlow, NetFlow Security Event Logs (NSEL), IPFIX, NetStream, jFlow, AppFlow and others. once a template is received. Flow View provides access to view any and all flows received by the collector. Some of the columns that may appear in the exports are below.

Flow View Field Names

When looking at data in Flow View some data columns are Plixer specific:

- 'flowDirection' Tells the reporting interface is the follow was collected ingress or egress on the router or switch interface. When direction is not exported, 'ingres*' is displayed which means direction was not exported with the flow and that ingress collection is assumed for the flow. NetFlow v5 does not export the direction bit.
- 'intervalTime' This is the time the collector received the flow.
- 'applicationId' This is the application as determined by settings under Admin tab -> Definitions - Application Groups
- 'commonPort' How the collector determines which port is the application port (also known as WellKnownPort). For example, take a flow with a source port of 5678 and a destination port of 1234. The collector will look at both ports (5678, 1234) and perform the following logic:
    - Which port is lower: 1234! Is it labeled (e.g. http)?
    - Yes: save it as the common port (1234) else,
    - Is 5678 labeled?
    - Yes: save it as the common port (5678) else,
    - Save 1234 as the common port
    - Note: if both were labeled, it would have gone with the lower port.

Fields mapping more or less to IPFIX fields

These field names are overloaded and don't map to any one IPFIX field. IPFIX might send 'sourceIPv4Address' or 'sourceIPv6Address', the column is always named 'sourceIPAddress'. The 'sourceIPAddress' column can store either IPv4 or IPv6.

- 'ipNextHopIPAddress' /* v4 or v6 */
- 'sourceIPAddress' /* v4 or v6 */
- 'destinationIPAddress' /* v4 or v6 */
- 'sourceIPPrefixLength' /* v4 or v6 */
- 'destinationIPPrefixLength' /* v4 or v6 */
- 'ingress_octetDeltaCount'
- 'ingress_packetDeltaCount'
- 'egress_octetDeltaCount'
- 'egress_packetDeltaCount'
- 'snmp_interface' /* (in|e)gress */

NOTE: /* v4 or v6 */ columns are used for both IPv4 and IPv6 formats.

Field Names in both Cisco and IPFIX (The collector uses the name from IPFIX)

The field names below exist only in Cisco docs. Except for the NBAR fields which only exist in Cisco's docs. Notice that the field names are fairly descriptive.

The IPFIX field names and descriptions can be found here. The Cisco fields and descriptions can be found here and here:

- SAMPLING_INTERVAL
- SAMPLING_ALGORITHM
- ENGINE_TYPE
- ENGINE_ID
- FLOW_SAMPLER_ID
- FLOW_SAMPLER_MODE
- FLOW_SAMPLER_RANDOM_INTERVAL
- SAMPLER_NAME
- FORWARDING_STATUS
- NBAR_APPLICATION_DESCRIPTION
- NBAR_APPLICATION_ID
- NBAR_APPLICATION_NAME
- NBAR_SUB_APPLICATION_ID
- NF_F_XLATE_SRC_ADDR_IPV4
- NF_F_XLATE_DST_ADDR_IPV4
- NF_F_SLATE_SRC_PORT
- NF_F_XLATE_DST_PORT
- NF_F_FW_EVENT
- NF_F_FW_EXT_EVENT
- NF_F_INGRESS_ACL_ID
- NF_F_EGRESS_ACL_ID
- NF_F_USERNAME

NOTE: The field names beginning with 'NBAR' were made up by plixer.

Archiving & Rollups

The collector will perform rollups (i.e. 5 min, 30 min, 2 hour, 12 hour, 1 day and 1 week) on any template with a column of octetDeltaCount. Please contact support if you need the rollups to occur on an alternate field. Visit the Admin Tab >> Settings >> Data History to configure how long to save the data.

Network Traffic Reporting

Overview

Reporting direction can be done inbound, outbound or bidirectional. The NetFlow and sFlow report combinations are dynamic and only appear if the data in the templates is available to run them. For example, the Pair Reports >> Flows report will not appear if ToS is not in the template(s) exported by the exporting device.

You can see which reports are available for your flow exports by going to Status > System > Available Reports. The Available Reports view allows you to view, sort, and filter report lists by Group Name, Report Name, and Template Count.

Report Groups

There are three types of Report Groups:

- **Commonly Used**
- **Vendor Specific**
- **System/Management**

The **Commonly Used** Report groups are most often used and also common to many device types and manufacturers. Included are the following.

- **Counts**
- **Destination Reports**
- **DPI Reports**
- **Firewall Events**
- **FQDN Reports**
- **IPFIXify**
- **NAT**
- **Pair Reports**
- **Queue Drops**
- **Source Reports**
- **Top Reports**
- **Volume Reports**
- **Wireless Reports**

Vendor Specific Report groups provide some very advanced reporting based on the unique flow elements exported in the flow templates. If Scrutinizer detects certain elements in the templates, the report group(s) will be listed in the report menu. Following is a list of many of these report groups.

- **AppFlow –** Citrix reporting
- **Astaro**

- **Barracuda**
- **Cisco**
    - **Cisco AnyConnect**
    - **Cisco AVC –** Application Visibility and Control
    - **Cisco CTS –** Cisco TrustSec
    - **Cisco FirePOWER**
    - **Cisco FW -** Firewall
    - **Cisco HSL –** High Speed Logging
    - **Cisco SLT –** Smart Logging and Telemetry
    - **Cisco VQM –** Video Quality Monitoring
- **Exinda**
- **FlowPro**
    - **FlowPro APM Reports**
    - **FlowPro Defender Reports**
- **Gigamon**
- **Ixia Reports**
- **Juniper**
- **Palo Alto Networks**
- **Procera Reports**
- **Riverbed**
- **Saisei**
- **SonicWALL Reports**
- **Stormshield**
- **VMware**
    - **VMware DFW -** Distributed Firewall
    - **VMware VDS -** vSphere Distributed Switch
- **Ziften**

System/Management Report groups provide reporting information regarding different aspects of your Scrutinizer server system, and other miscellaneous network monitoring reporting, such as the Report Group list below indicates.

- **Alarms**
- **Baselines**
- **Chassis**
- **CrossCheck**
- **Honeynet**
- **Mailinizer**
- **Overlay Network –** VxLAN reporting
- **Replicator**

- **SQLify Reports**
- **Syslog**
- **Tests/QA**
- **Threats**
- **Vitals –** Scrutinizer server health reporting

Report Types

Following are some examples of the different report types available, how they are created, and their value to you. This list is a small sample of the wide range of report types included within Scrutinizer. There are hundreds of report types. If you have questions about any not listed here, please contact Plixer support directly.

- **Pair Reports >>**
    - **AS to AS by Tag:** shows the source Autonomous Systems traffic to destination Autonomous Systems. This report is based on AS information included in the flow template and is not available with sFlow, NetFlow v1, or if the router isn't configured properly.
    - **AS to AS by IP:** The autonomous system is determined by comparing the IP addresses to a table in the database. This product includes GeoLite data created by MaxMind, available from http://www.maxmind.com.
    - **Connections by Bytes:** communication between hosts based on similar source and destination ports as well as inbound and outbound interfaces. This reports displays more entries than the Conversations or Host to Host report, but not as many as Flow View.
    - **Connections by Flows:** communication between hosts showing the source and destination IP addresses, ports and the protocol.
    - **Conversations App:** communication between hosts based on an Application Group or the Well-known Port. The random port is ignored.
    - **Conversations WKP:** communication between hosts based on the Well-known Port. The Application Group is ignored. The random port is ignored
    - **Country to Country:** communication between countries. The list of country mappings to IP addresses is updated as necessary in each software update.
    - **Customer VLAN to VLAN:** top dot1qCustomerVlanId and postDot1qCustomerVlanId as defined by the IETF.
    - **Host to Host:** communication between hosts where the application and ports are ignored.
    - **MAC to MAC:** top MAC addresses talking to other MAC addresses.
    - **Rev 2nd lvl Domain pairs:** Reverse 2nd level Domain pairs reports on communication between domains. Domains are determined dynamically as flows are received by doing a DNS reverse lookup. This auto resolution of names can be enabled by checking: Admin > Settings > System Preferences > DNS Resolve. Uncheck "DNS Resolve" to disable this setting if witnessing performance issues. If looking at URLs by domain, keep in mind the issues revolving around DNS as explained in the FlowPro APM URLs report section with the hostmonster.com example.
    - **Subnet to Subnet:** top networks talking to other networks based on IP addresses and subnet masks. This report can be very useful in MPLS environments.
    - **VLAN to VLAN:** top dot1qVlanId talking to postDot11qVlanId or vlandId to postVlanId as defined by the IETF.

- **Source / Destination Reports >>** Reports can be run inbound, outbound, or both, for either source or destination.
  - **Autonomous System by Tag:** either source or destination. This report is based on AS information included in the flow template and is not available with sFlow, NetFlow v1, or if the router isn't configured properly. The Autonomous System report requires that the router be participating in BGP and that the following command be executed on the router:

    router(config)# ip flow-export version 5 [peer-as | origin-as]
  - **Autonomous Systems by IP:** The autonomous system is determined by comparing the IP addresses to a table in the database. This product includes GeoLite data created by MaxMind, available from http://www.maxmind.com.
  - **Countries:** displays the countries sending or receiving the most traffic. The list of country mappings to IP addresses is updated as necessary in each software update.
  - **Customer VLAN:** top dot1qCustomerVlanId as defined by the IETF.
  - **Rev 2nd lvl Domain:** Reverse 2nd level domain displays the domains sending or receiving the most traffic. Domains are determined dynamically as flows are received by doing a DNS reverse lookup. This auto resolution of names can be enabled by checking: Admin > Settings > System Preferences > DNS Resolve. Uncheck "DNS Resolve" to disable this setting if witnessing performance issues. If looking at URLs by domain, keep in mind the issues revolving around DNS as explained in the FlowPro APM URLs report section with the hostmonster.com example.
  - **Hosts:** top hosts inbound or outbound as the Source or Destination.
  - **Host Flows:** displays hosts with the most flows for the time period as well as the total flows per host.
  - **MAC:** top MAC addresses.
  - **Subnets:** top networks based on IP addresses and subnet masks. This report can be very useful in MPLS environments.
  - **User Name by IP:** Displaying the user names that have logged an IP address onto the network requires integration with the local authentication system.
  - **VLAN:** top dot1qVlanId or vlanId as defined by the IETF.

- **Top Reports >>**
  - **Applications Defined:** can be defined as ranges of ports which include IP addresses. Visit the Admin tab to create or modify them.
  - **Next Hop**: Only displayed if the template contains the next hop router of the flow.
  - **Protocols:** top transport layer protocols (e.g. TCP, UDP, ICMP, IGMP, etc.)
  - **Type of Service:** which includes DSCP. The values displayed (i.e. RFC 791, 1349 or 2474) can be set under the Admin tab -> Settings -> System Preferences. NOTE: DSCP is the default which only uses 6 of the 8 bits in the ToS field. The 2 remaining bits are reserved and sometimes used for Explicit Congestion Notification (ECN). This will be indicated when it appears in this report.
  - **Well Known Ports:** top applications based on well-known port. The logic behind this report is explained in the following question:
    - Question: If I have tcp/5678 as the socket on one side of the flow, and tcp/1234 on the other... How do I tell Scrutinizer that I'm interested in the fact that it's tcp/5678 and I don't really care about tcp/1234?
    - Answer: In order to answer this question, it might be best that I explain Scrutinizer's logic by digressing on how Scrutinizer deals with the issue. The collector will look at both ports (5678, 1234) and perform the following logic:
      - Which port is lower: 1234! Is it labeled (e.g. http)?

- Yes: save it as the common port (1234) else,
- Is 5678 labeled?
- Yes: save it as the common port (5678) else,
- Save 1234 as the common port

Note: if both were labeled, it would have gone with the lower port.

- **Volume Reports >>**
  - **Flow Volume:** displays the volume of flows over time frames.
  - **Host Count:** trends the average volume of hosts per second per interval. The table can also list the total hosts per interval. This report is available for either Source (src) or Destination (dst).
  - **Pair Volume:** this report trends the volume of address pairs per second per interval. The table can also display the total address pairs per interval.
  - **Traffic Volume:** in and out total utilization.

- **Cisco FW:** Access Control List (ACL) reports are available for Cisco Firewalls if the required elements are exported in the flow templates.

  When trying to leverage the ACL data, keep the following in mind. The ACL will show up in the report as something like 0xa231c4d3 - 0xb34531ad - 0x00000000. This 12-byte raw ACL ID must be divided into its three constituent parts, as follows:
  - The first four bytes are the ACL Name ID.
  - The second four bytes are the ACL Entry ID (ACE)/Object-Group ID.
  - The final four bytes are the Extended ACL Entry ID.
  - To see them in the ASA, type the following on the console:

    use asa# show access-list

  NOTE: The reports can also display the actual ACL names. Visit Admin Tab > Settings > ASA ACL Descriptions and enter the login credentials. The system will then periodically download the ACL names from all the ASAs it is receiving flows from.

- **Cisco Application Visibility and Control (Cisco AVC)** is an umbrella for several unique exports. Reporting for Cisco AVC includes:
  - **Performance Routing (PfR)**: requires that the NetFlow exporting device be exporting Performance Routing (PfR) information in the templates.
    - **Active**: Reports on Jitter, Jitter out of policy, MOS, MOS out of policy, RTT, RTT out of policy, out of policy
    - **Passive**: Reports on out of policy, RTT and RTT out of policy
  - **NBAR & NBAR2 Applications (NBAR)**: Network Based Application Recognition (NBAR) is used to export layer 7 application details (E.g. Skype, Facebook.com, BitTorrent, Webex, etc.).
  - **Performance Agent (PfA)**: Used to export application performance analytics, traffic statistics, and WAN optimization metrics in a NetFlow Version 9 template-based format, making it easier for management applications to consume. Details include application latency, HTTP Host, WAAS details, etc.
  - **Performance Monitoring (PfM)**: Used to export voice and video performance analytics on UDP and TCP traffic. Details can include, but are not limited to, jitter, packet loss and round trip time.

- **Easy Performance Monitoring (EzPM)**: Similar to Performance Monitoring, but based on the EzPM flow template.

- **FlowPro APM Reports** (all reports are NOT listed below) >>
  - **Jitter**: Available Jitter reports for the FlowPro APM are:
    - **Host Jitter**
    - **Host Jitter By SSRC (Src or Dst)**
    - **Host to Host Jitter By SSRC/Codec**
    - **Host to Host Jitter By SSRC/ToS**
  - **Latency**: Latency of hosts, subnets, and applications broken out by:
    - **Client**: Latency determined by the length of time between the Client TCP flag SYN and the SYN, ACK reply from the server.
    - **Server**: Latency determined by the length of time between the Server TCP flag SYN, ACK and the ACK reply from the client.
    - **Appl**: Application Latency - the time needed by an application to react to a client request. For TCP connections, application latency is computed on the first packet after three-way-handshake, for UDP connections on the first client -> server and server -> client packet.
  - **Top Applications**: Top Applications reporting using layer 7 application data provided in the flow template.
  - **Top URLs**: the top URLs seen by the FlowPro APM. Keep in mind that due to DNS and reverse DNS resolutions, domains in URLs do not always match the IP addresses mapped to the domains. This is because a single IP address leveraged by a hosting site such as hostmonster.com can represent thousands of domains which are all represented by a single IP address (e.g. 74.220.207.152). If you go searching for a site that someone visited (e.g. www.411japan.com) don't expect it to show up in a domain report; expect to see hostmonster.com in its place.

- **FlowPro Defender Reports** (all reports are NOT listed below) >>
  - **ARTT**: Application Round Trip Time is derived from the IPFIX element reverseflowdeltamilliseconds. This value is measured by observing the amount of time two hosts take to communicate with one another. This is not necessarily an indication of network latency as some communications do not elicit an immediate response. For example, a mail server could initiate a TCP connection by sending a SYN and then 2 retransmits due to the lack of a SYN/ACK response. If the client eventually (e.g. after 40 seconds) sends an RST/ACK, that was how long that conversation took for the round trip. The ARTT measurement will be at least 40 seconds even if a ping response would have triggered a response in under 10 milliseconds. The same ARTT metric is also exported for UDP connections. For example, if a phone initiates a VoIP SIP connection and the PBX takes 56 milliseconds to respond, the ARTT will measure at least 56 milliseconds.

    Available ARTT reports are:
    - **App Connections ARTT**
    - **Source ARTT**
    - **Source to Destination ARTT**
    - **VLAN and Source IP ARTT**
    - **VLAN ARTT**

- **Flow View:** Shows the raw flows (aggregated if viewing above 1 minute intervals) for the time frame being viewed. The Flow View report displays all of the raw flows that match the inherited filter. There is no flow aggregation on similar fields as there is in most other reports. All flows are displayed as received by the router/switch/server.

Current Report

The Current Report frame is displayed in the left hand pane when selecting an interface or after selecting the Run Report Wizard from the Trends menu in the Status tab.

The graph and table data for the flow reports will display in the main section of the screen.

You can change the timeframe for the report either by selecting a date range from the slideout Options menu defined in the Report Options section below, or by clicking in the graph and dragging over the desired timeframe.

- **Colors**: In the table below the graph, the top 10 or more entries are displayed. Only the Top 10 are in color. Entries 11 and up are rolled into the color gray. Notice the 'Other' entry at the bottom of the table. This is the total non Top 10 traffic. The 'Total' represents all traffic (i.e. Top 10 and Other traffic added together). These same colors are used in the graph to represent the Top 10 table entries.

  Note: The color selections can be changed in Admin > Security > User Accounts > select a user > Preferences > Rank Colors

WARNING: If the flow device (e.g. router) is exporting multiple templates for different flows it is exporting, utilization could be over stated if the flows contain the same or nearly the same information. The front end of Scrutinizer will render reports using data from all templates with matching information. Be careful when exporting multiple templates from the same device! If you find this to be the case, use the filters to select a single template.

No Data Found

The “No Data Found” message in a report indicates that historical data is not available for the time period requested and could be for either of the following reasons:

1. Historical data settings are too low for the time frame requested. To increase the historical data retention settings, go to Admin tab -> Settings -> Data History.
2. Flows are not (or have not) being received from the exporter(s) during the time frame requested.

Current Report (frame contents)

At the top of the Current Report frame is a row of icons providing the following actions available for the report.

- **Clear** (trashcan) is used to remove all items in the "Current Filter"
- **Save** (diskette) is used to save a collection of report filters and parameters to create a Saved Report.
- **Save As** (double diskette) is used to make a copy of a current Saved Report with a new name, leaving the original report intact.
- **Schedule** (clock) is used to schedule a saved report.
- **Dashboards** (grid) is used to place a saved report in a selected Dashboards sub tab.

- **Print** (printer) is used to print the current report listed in the filter.
- **CSV** (CSV) is used to export the data in the current filter in CSV format.
- **PDF** (PDF) download a pdf file containing the current report.
- **Email** (@) is used to email the report displayed using the current filter. Separate multiple destination email addresses with a comma or semi colon.

Next in the Current Report frame are these additional reporting options.

- **Report**: Enter a name if the report and filter(s) are to be saved for future reference.
- **Templates Used**: Includes count of templates used by the report and expands to show which templates are used
- **Devices Used**: Includes count of devices used by the report and expands to show which devices are used
- **Report Details**: Clicking this button opens the Report Details modal with the following tabs:
    - Collector Details
    - Exporter Details
    - Filters: view/edit/remove existing and add new filters to the report
    - Threshold: view/edit/remove existing thresholds or add a threshold to the report
- **Filters**: This button has the same function as the Filters tab in the Report Details modal above. Any field in a flow template can be included or excluded in a filter. See the section below titled "Filters - Include or Exclude Data" for more details on filters.
- **Threshold**: This button has the same function as the Threshold tab in the Report Details modal above.

Report Options

When navigating in a report, anytime a filter is changed or added to the report, the filters list in the left frame is updated.

These other reporting options are included in the slideout options panel above the report graph:

- **Change Report Type button**: Report types are displayed based on the data available in the templates selected.
- **Direction**: Inbound, Outbound and Bidirectional. In Bidirectional mode, the outbound is displayed on the bottom of the trend. The reporting engine will try to use ingress flows to display inbound traffic however, if ingress flows are not available, it will try to use egress flows if available. The same logic holds true when displaying outbound traffic. The reporting engine will try to use egress flows however if none are available, it will use ingress flows. Switching the configuration on the router from exporting ingress to egress flows or vice versa will not be recognized by the reporting engine until after the top of hour.
- **Rate / Total**: Select Rate to display Rate per second or Total for total amount per interval (e.g. 1 min, 5 min, 30 min, 2 hr, etc.). Some reports (E.g. Cisco Perf Monitor) default to Total. When the report is changed to display 'Rate', this value will not change automatically and will have to be changed back to Total manually. The opposite is also true.
- **Interval: Auto, 1m, 5m, 30m, 2hr, 12hr, 1d, 1w.** This is the time interval in the trend. Generally the default is taken as the database has been optimized for this setting. Selecting 1m (i.e. 1 minute intervals) for a 5 hour time frame can take a significant amount of time to render

depending on the volume of flows coming from the device. Expect results that vary between devices. Excessive intervals in a trend will not be allowed by the interface.

- **Number of Rows: 10, 25, 50, 100, … 10000** this is the top number of results to be displayed in the table below the trend. The default can be set under Admin tab -> Security - User Preferences.
- **Show Host Names**: toggle between displaying IP addresses or DNS Host names in the table data.
- **Show Raw Values**: Formatted/Raw displays the data in certain columns either formatted (5.364 Mb/s) or raw value (5364239).
- **Bits / Bytes / % Util**: can be used when available to change the type of data used for the trend/table. This option does not apply to all report types. Percent utilization (% Util) is not available unless the interface speed is picked up via SNMP. Interface speed can also be entered manually via the Device Details View or as a report filter. When multiple interfaces are included in a report, the calculated interface speed with be the SUM of all interface speeds. Inbound is calculated separately from outbound. The summed port speed is used for percent calculations. All interfaces are required to have a defined speed for percentage reports. If 'Percent' is selected in the drop down box, it represents the overall percent of the entire interface. The preceding percent column that can't be changed represents the percent of the overall bandwidth consumed.
- **Show Peak**: if Yes is selected, a Peak column is added to the report. Peak values are the highest data point in the graph in the same interval the graph is reporting in. The only exception to that is in the Traffic Volume report where the peak values are based on the 1 minute intervals taken from the Totals tables.
- **Show 95th**: if Yes is selected, a 95th (percentile) column is added to the report. The 95th percentile is a mathematical calculation used to indicate typical bandwidth utilization. The top 5% data points in the graph are dropped, making the "95th" data point now the top bandwidth usage point. For example, in a graph with 100 data points, the 5 highest values are removed, and the next highest becomes the 95th percentile.
- **Show Interfaces**: adds an 'in Int' and an 'out Int' column to the report, showing inbound and outbound interfaces for the flow data reported.

Graph Options

- **Graph Type: Line, Step, Bar, Pie, Matrix** is the type of graphical presentation to be displayed. Try clicking and dragging on the line chart to zoom in on time frames. All graphing options are not available for all Report Types. For example, the Matrix graph will only work with reports that have a source and destination field, such as reports in the Pairs report group.
- **Stacked/Unstacked**: Select Stacked to display the total amount. Select Unstacked to display the top 10 individually. Some reports (E.g. Cisco PfM reports) default to Unstacked. When the report is changed to a report normally displayed as Stacked, this value will not change automatically and will have to be changed to Stacked manually.
- **Show Others**: Set this option to 'No' to hide the gray 'other' traffic in the trend or pie chart. Other traffic is discussed in depth in the section on Data Aggregation. This option is often used in sFlow reports. Other traffic:
    - In the trending graph it is the non Top 10 traffic and shows up as gray in color
    - In the table below the graph, the Other value at the bottom of a report table is the total traffic, minus the sum of the line items displayed. Notice as you click through the pagination, the total Other traffic increases.
    - Some report types will have this option set to 'No' by default. When changing to another report, it should be manually changed to 'Yes'.
    - **NOTE**: In a standard interface trend (e.g. Top Protocols) with no filters other than the interface, the graph is first built using data from the totals tables and then the data from the Top 10 in the related conversations table is subtracted from the total and then added back individually to display the colors for each of the Top 10. These two tables are

discussed in further detail in the section below on Filters. As you click on the pagination at the bottom of the table, you can view all of the data that makes up the 11th color: gray.

Date / Time Options Timezone: server timezone is displayed here

- **Reporting & Timezones:**
    - Flow timestamps are stored in epoch format, which is time zone agnostic. When a report is loaded, Scrutinizer uses the browser's time zone setting to format the epoch timestamps into a human-readable date format. Individual users can change their time zone setting in the Admin > Security > [User] view. A setting of "Automatic" will default to the browser's configured time zone.
- **Range**: a drop down box to select a time frame.
- **Report Start / Report End**: the actual date text can be altered or the arrows to the left and right of the displayed time can be clicked to shift the time period displayed. Avoid saving a report with a 'Custom' time frame as each time the report is run, it will execute with the exact same start and end time. If the data necessary for the custom time frame report has been deleted, the report will display with a "no data available" message. Suggested save times include "Last 5 minutes" or "Last 24 hours".
- **Apply Dates**: click this button after making any date / timeframe changes to have the changes take effect.
- **Business Hours**: when checked, only flows reported within the Business Hours defined in Admin > Settings > System Preferences will display.

Saved Reports

Navigate to the bottom of the left frame and click on Saved Reports to view the list of saved reports.

Saved Reports are basically custom reports which display selected data on 1 or more interfaces across potentially several devices.

When you click on a Saved Report to open the report view, you are returned to the Current Report view and have all the features and functions of that view.

Refer to the Saved Reports section in the Status Tab page for more information on the Saved Reports view.

Filters - Include or Exclude Data

It is often necessary to filter on the flow data to narrow in on desired traffic. For this reason, data in a report can be included or excluded. Clicking on the Filters 'Add' button in the left pane of the screen will popup the 'Add New Filter' modal.

1. First option is to select the type of filter. Included in this list are:
    - General filter names are commonly used filters with familiar names. They allow certain boolean expressions for example, host to host, domain to domain, subnet range or Application Defined (i.e. defined range of ports and IP addresses). These filters are not always in the actual NetFlow or sFlow export rather, they are derived via portions or combinations of fields.
        - Not all devices (i.e. switches and routers) include TCP flags or nexthop in their NetFlow exports. If a field is not included in the NetFlow export for a device, it will not be part of the filter list for that device.
    - Advanced filter lists all of the fields that are collectively in all of the templates being used in a report. For example, if the device is exporting MAC addresses in only 1 of two templates being used in a report, MAC address will appear.

- Calculated Column Filter lists any calculated columns available in the current report, ie. sum_octetdeltacount, sum_packetdeltacount.
- Port Speed is a special case filter that allows you to set a port speed for a report.
- Sample Multiplier is another special case filter that allows you to set a multiplier value for sampled flows to recalculate to full flow values.
- Wildcard Mask is another special case filter that allows you to add a custom mask to a network to filter on networks "like" the search criteria.

  For example:

  Network: 10.0.11.3

  Mask: 0.255.128.240

  Results:

  10.1.11.51

  10.30.11.3

  10.27.11.3

  10.26.11.35

  10.26.11.3

  10.26.11.19

2. After selecting a filter type/name, other type specific options will appear. If the filter type has a predefined list of items, a dropdown list will appear to select from, otherwise a textbox will be displayed for entering the filter data. If Source or Destination are applicable, another dropdown selector will appear for selecting Source, Destination, or Both. If it is a calculated column, a dropdown selector of numerical comparisons will appear.
3. The next option is to select whether this will be an Include or an Exclude filter. Include filters will only display flow data where the filter criteria is equal. Exclude filters will display everything except the filtering criteria.
4. When all options are completed, the Add Filter button will appear, allowing the new filter to be added to the existing filters. After adding the new filter, the Update Report button displays and clicking that button is the last step to adding a new filter.
   - Report filters can also be added by simply dragging an Item in the table portion of the report and dropping that item in either the Include Filter (green) or Exclude Filter (red) boxes that display.
   - New or existing filters can be edited at any time by clicking on the edit link for the appropriate filter. After editing is completed, click the Save button in the filter, then click they Apply button at the top of the filter list.

IMPORTANT NOTES:

- Two types of historical tables are maintained via Data Aggregation for each NetFlow exporting device.
  - Totals - this contains the actual amount of total traffic in and out an interface for each interval before flows are rolled up into the conversations table. This table must be maintained as the Conversations table only contains the top 1,000 by default. This can be increased in Admin > Settings > Data History > Flow Maximum Conversations. **If filters are used, this table is no longer part of the report.** A report with only a single interface filter (i.e. selected interface) will use this table so that total utilization is accurate over time.

  - Conversations - this contains the actual flows that are rolled up over time. This can be read about in the Data History section. **If filters are used, these are the only tables used in the report.**

- Including and excluding data using the same filter field twice creates a logical 'OR' relationship (e.g. display all traffic if it includes 10.1.1.1 OR 10.1.1.2)
- Including and excluding data using different filter fields creates a logical 'AND' relationship (e.g. display all 10.1.1.1 traffic AND that uses port 80)
- When adding an 'IP Host' to an 'IP Range' or an 'IP Host' to a 'Subnet' filter, the 'AND' and 'OR' rules still apply. For example, if an IP Range filter of 10.1.1.1 - 10.1.1.255 is added and then an IP Host filter of 65.65.65.65 is added, the flows must match both filters. Obviously, this is impossible and the report will fail with no data.
- When using Source or Destination or Both with IP Host, IP Range or Subnet, keep the following in mind:
  1. If the IP Host filter of 'Source' A (e.g. 10.1.1.4) is applied, then there may be data for inbound, but most likely not outbound. This is because what comes in as the Source, typically doesn't go out the same interface as the Source. The same holds true with Destination addresses.
  2. If the IP Host filter of 'Source' A (e.g. 10.1.1.4) is applied and then a second filter of 'Destination' B (e.g. 10.1.1.5) is applied then only flows where the Source is A and the Destination is B will appear. Although this is adding the same filter 'IP Host' twice, the AND logic applies because host A is the source and host B is the destination and thus are different filter types. Note again that data for inbound may appear, but most likely there won't be any outbound or vice versa. This is because what comes in as the Source, typically doesn't go out the same interface as the Source. The opposite case applies when data appears for outbound using this type of filter.
  3. If the filter "Src or Dst" or 'Both" is applied to an IP Host filter then all flows to or from A will appear and traffic both inbound and outbound will likely display data from A. If a second filter is added as "Src or Dst is B", then traffic again will appear from both hosts in both directions. However, all flows must involve A or B as the Source or Destination.
- The Interface filter is the first option that must be exercised prior to any other filter
  - When mixing NetFlow and sFlow interfaces in a report, NetFlow data will usually dominate. This is due to NetFlow's 100% accuracy with IP traffic where sFlow is sampled traffic.
  - Although sFlow samples packets, it can send interface counters that are 100% accurate. However, the totals tables used for total in / out traffic per interface are not referenced when mixing sFlow with NetFlow interfaces in reports. This leads to understating the 'Other' traffic in reports.
  - When reporting on the 'ALL' Interfaces option for a device, inbound should equal outbound in the trends. What goes in ALL interfaces generally goes out ALL interfaces.

Thresholds

Any report, with any combination of filters, can be turned into a traffic monitoring policy by adding a Threshold to the report. See the Report Thresholds page for more information.

Report Navigation

Clicking on table data in a report will present a report menu of available report types for the device/interfaces used in the current report. When selecting a report in this way, the table data selected will automatically be added as a filter to the new report generated.

If the selected table data is an IP Address, a menu option called “Other Options” will be listed as described below.

Other Options

In Other Options, the IP address selected is passed in the URL to the application. Default menu options are:

- **Report** – Report suspicious behavior
- **Search**
- **Alarms**
- **HTTP**
- **FTP**
- **SSH**
- **Lookup** – Whois Lookup
- **GEO IP** – Geographical lookup
- **Cisco IronPort** – Geographical and detailed IP address information
- You can add new applications by editing the applications.cfg file in the /home/plixer/scrutinizer/html/files/ directory. The format for applications.cfg is: (title),(link),(desc) -- one per line. The description is optional. For example:
    - FTP, ftp://%i, this will launch an ftp session to the IP address
    - Google, http://www.google.com/search?q=%i, this will launch a google search on the IP address.

  Updates to the languages.english table also need to be made for the buttons to show up:

    - You need to create lang entries in the Languages.english table for the buttons to show up. The following is an example for the 'WMI Usernames" script.
        - INSERT INTO languages.english (id, string) VALUES('WMIUsers', 'Current Users'),('WMIUsersDescr', 'Use WMI to identify users currently logged into the address above.');
        - WMIUsers is the lang key for the button name.
        - WMIUsersDescr is the lang key for the description.
        - Then in applications.cfg you have to reference these lang keys:
            - open applications.cfg and add the following line without quotes: "WMIUsers, /cgi-bin/currentUsers.cgi?addr=%i, WMIUsersDescr"
            - **Remember**: in the applications.cfg it’s “button lang key, URL of the file executing the parameters, description lang key”

Report Thresholds

Overview

Any report, with any combination of filters, can be turned into a traffic monitoring policy by adding a Threshold to the report.

The Threshold option is just below the Filters option in the left hand pane of the Report view. Instructions for adding thresholds to reports are detailed below.

To add a threshold to a report:

1. Save a report
    a. Thresholds can only be added to Saved reports.
    b. Enter a report name in the Report: textbox in the left hand pane of the Report view, then click the Save icon above the report name.
        - If you don't save the report first, you will be prompted to enter a report name and save it first thing in the Threshold modal.
2. Click the Add button to the right of Threshold in the left hand pane.
    a. The Report Details modal opens to the Threshold tab with the following text: Trigger alert if rate/total value per table's Total/Per row for inbound/outbound traffic in 5 minute interval.
    b. Selectable options within this modal are:
        - **Rate/Total** – This is taken from the saved report parameter and determines if the threshold is based on the rate of the value selected, or the total amount of the value.
        - **Total/Per row** – This radio button selectable in the Threshold modal indicates whether to threshold against the Total report value or each line/row entry's value (Per row).
        - **Inbound/Outbound** – This variable is also determined by the saved report parameter, whether the selected flow direction is Inbound or Outbound. This is the flow direction that the threshold will be monitoring. If the saved report's flow direction is Bidirectional, the threshold will monitor Inbound traffic.
    c. Thresholds are monitored every 5 minutes, based on the last 5 minute interval.
    d. Threshold comparison options are:
        - **Greater than or equal to** (>=)

        or

        - **Less than or equal to** (<=)
    d. The threshold value is entered in the textbox after the word "than".
    e. The unit of measurement is from the saved report unit setting and can be either bits, bytes, percent, or omitted for counter fields. If bits, bytes, or counter fields, an additional selection for unit quantity is presented:
        - **-** : Integer value of bits/bytes, or counters.
        - **K** : Kilobits/bytes, counter value
        - **M** : Megabits/bytes, counter value

- **G** : Gigabits/bytes, counter value

f. If Per row is selected above, an additional field is presented, Policy Threat Multiplier.

- This value is used to multiply the policy violation counts for this threshold to give its weighted value reported as the Threat Index in the Policy Violations listings. Default is 1.

3. After completing entry of the fields listed above, click the Save Threshold button.

- To exit the Threshold modal without saving, click the Close button.

4. The Select Notification Profile modal displays next.

a. If Notification Profiles have been configured, select the appropriate one from the dropdown selector.

b. To configure new Notification Profiles, click Manage Notifications

i. A new browser window opens to the Notification Manager page

ii. After creating new Notification Profile(s), to assign the profile to the report threshold, click on 'edit' to the right of Threshold, then click Save Threshold, and the Select Notification Profile modal will be displayed again.

c. After selecting the Notification Profile (or leaving the threshold modal without selecting a Notification Profile) click on:

- **Save** – Saves the threshold with the changes made up to this point
- **Close** – Exits without saving the Notification Profile selection
- **Save & Edit Policy** – Saves the threshold settings made so far and opens the Edit Policy modal to edit this threshold policy.

Notes:

- The threshold setting unit of measurement is determined by the report settings, either percent, bits, or bytes. If the report is set to report by bits or bytes, then there is an additional option of K, M, or G for total bits/bytes.

- Thresholds can also be set on other counters such as round trip time, packet loss, jitter, flow count, etc. The K, M, and G option is also available when thresholding against these other counter fields.

- It is good practice to view the FlowView report to get an idea of what the raw data looks like before setting a threshold.

- After saving the threshold, the modal will go to Select Notification Profile. You can select a profile from the dropdown, or you can click Manage Notifications to create one. Or you can select Save and finish without adding a notification to the threshold. An alarm will still be generated to Scrutinizer when the threshold is violated even without a notification included in the threshold configuration.

- Adding a threshold requires that the report be saved first. If you haven't yet saved the report, the first option to be presented in the modal is to name and save the report.

- This thresholding option executes every 5 minutes to compare the volume of data to the threshold setting.

Vitals Reporting

Overview

The Vitals reports provide insight on the health of the Scrutinizer servers (e.g. CPU, Memory usage, Hard drive space available, Flow Metrics, etc.). Vitals information is reported for all servers in a Distributed Environment.

Vitals reports can provide valuable insight into the servers' performance. As with any other flow report type, thresholds can be set on any of the Vitals reports, providing the ability to alert on threshold violations (ie. low disk space, high cpu utilization, etc.)

These reports are accessible at Status->Device Explorer->Scrutinizer server (127.0.0.1)->Reports->Vitals. (A Vitals Dashboard is also created by default for the Admin user and includes many of the reports listed below.)

- **% CPU per Process**: This report displays CPU percentage consumed per process on the server.
- **CPU**: Average CPU utilization for the Scrutinizer server(s).
- **CrossCheck Runtime**: Monitors runtimes for CrossCheck methods (processes).
- **Database**: Provides the following database metrics:
    - **Connections by Bytes:** Excessive connections can result in reduced performance. NOTE: other applications using the same MySQL database will cause this number to increase.
    - **Read Req:** The number of requests to read a key block from the cache. A high number requested means the server is busy.
    - **Write Req:** The number of requests to write a key block to the cache. A high number of requests means the server is busy.
    - **Cache Free:** The total amount of memory available to query caching. Contact support if you find that your query cache is presently under 1MB. For more information, please reference the MySQL Documentation.
    - **Queries:** Tracks the number of queries made to MySQL. More queries indicates a heavier load to the MySQL server. Generally there will be spikes at intervals of 5 minutes, 30 minutes, 2 hours, 12 hours, etc. This indicates the rolling up of statistics done by the stored procedures. This Vitals report is important to watch if the NetFlow collector is sharing the MySQL server with other applications.
    - **Threads:** Threads are useful to help pass data back and forth between Scrutinizer and the database engine. The MySQL Server currently manages whether or not to utilize the configured amount of threads. For more information, please reference the MySQL Documentation.

- **Buffers Used:** Key Buffers Used - indicates how much of the allocated key buffers are being utilized. If this report begins to consistently hit 100%, it indicates that there is not enough memory allocated. Scrutinizer will compensate by utilizing swap on the disk. This can cause additional delay retrieving data due to increased disk I/O. On larger implementations, this can cause performance to degrade quickly. Users can adjust the amount of memory allocated to the key buffers by modifying the \scrutinizer\mysql\my.ini file and adjusting the key buffer size setting. A general rule of thumb is to allocate as much RAM to the key buffer as you can, up to a maximum of 25% of system RAM (e.g. 1GB on a 4GB system). This is about the ideal setting for systems that read heavily from keys. If you allocate too much memory, you risk seeing further degradation of performance because the system has to use virtual memory for the key buffer.

- **Distributed Heartbeat** and **Distributed Synchronization**: These two Vitals reports provide further insight into internal communications in a Distributed environment.
- **FA Counts** and **FA Times**: Flow Analytics Counts and Flow Analytics Times provide metrics on the processing of Flow Analytics Algorithms. FA Times is useful in managing FA algorithms not coming to successful completion.
- **Flow Metrics/Exporter** and **Flow Metrics/Port**: Metrics are provided by exporter and also by listening port for:
  - **MFSN**: Missed Flow Sequence Numbers.
    - **NOTE:** Sometimes MFSN will show up as 10m or 400m. To get the dropped flows per second, divide the value by 1000ms. A value of 400m is .4 of a second. 1 / .4 = 2.5 second. A flow is dropped every 2.5 seconds or 120 (i.e. 300 seconds/2.5) dropped flows in the 5 minute interval displayed in the trend.
  - **Packets**: Average Packets per second.
  - **Flows**: Average Flows per second: This is a measure of the number of conversations being observed.
    - **NOTE:** There can be as many as 30 flows per NetFlow v5 packet (i.e. UDP datagram) and up to 24 flows per NetFlow v9 datagram. With sFlow, as many as 1 sample (i.e. flow) or greater than 10 samples can be sent per datagram.
- **Memory**: The Memory Vitals report displays how much memory is available after what is consumed by all programs on the computer is deducted from Total Memory. It is not specific to NetFlow being captured.
  - **NOTE**: The flow collector will continue to grab memory depending on the size of the memory bucket it requires to save data and it will not shrink unless the machine is rebooted. *This is not a memory leak*.

- **Report Request Time, Report Type Data Time,** and **Report Type Query Time**: These Vitals reports provide reporting performance metrics.
- **Storage**: The Storage report displays the amount of disk storage space that is available. After an initial period of a few weeks/months, this should stabilize providing that the volume of NetFlow stays about the same.
- **Syslogs**: The following metrics are available with the Syslogs report:
    - **Syslogs Received**: The average number of syslogs received per second.
    - **Syslogs Processed**: The average number of syslogs processed per second.
- **Task Runtime**: Task Runtime displays runtimes per Scrutinizer automated tasks such as nightly history expiration, vitals data collection, etc.
- **Totals/Rollups Times**: Time durations for totals, rollups, and data inserts in the MySQL database per flow template per exporter.

System

Access Denied

The Administrator has denied your user account from accessing this tool.

Backups

Overview

This page explains the Alternate Copy and Backup Method for Large Scrutinizer Databases.

Backing up the data

The historical and configuration data for Scrutinizer is stored in the plixer and reporting databases. Users are not recommended to backup these databases directly. The recommended solution is to dump the databases to backup files and optionally run the backup on those files.

Copying the data

Use the following commands to dump the data to files in a backup directory.

```
pg_dump plixer | gzip > plixer.gz
pg_dump reporting | gzip > reporting.gz
```

Choosing when to run the backup

There is no one right time to do a backup, but it is best to try and pick a time when Scrutinizer is least busy on your system. In addition to when the network is busy a few times to keep in mind are:

- 2h are written to starting at the top of even hours
- 12h are written to starting after 12 and midnight
- 1d are written to starting after midnight
- 1w are written after midnight on Saturday

Suggested times

- Weekly : 3am Sunday
- Daily : 3 am Daily

Database Connection Failure

Overview

The system is having trouble connecting to the database. Please contact support directly for more information about your query.

Distributed Collectors

Overview

Scrutinizer supports a distributed architecture where numerous servers can collect flows with a central reporting server for reporting on flows across all of the collectors. Secondary reporting servers can also be configured for redundancy/fail-over purposes.

- The distributed Scrutinizer system allows the collection system to ingest several million flows per second across several servers - regardless of physical location.
- In addition to increased flow volume, the distributed architecture allows for solution to receive flows from tens of thousands of flow exporting devices.
- The reporting server provides a central interface for managing all exporters, flows, interfaces, and alarms across the distributed cluster.
- The distributed architecture solution requires additional configuration. Contact Plixer if interested in this solution.

Flowalyzer

Flowalyzer is a free NetFlow and sFlow Tool Kit (TM) for testing flow technologies. This utility has several tabs each with a unique function:

- **Listener:** The Listener is used to determine whether or not NetFlow or sFlow is being received and what ports they are being received on. Any application currently listening on the same port as Flowalyzer will cause conflicts. Conflicting applications must stop using the ports Flowalyzer is trying to listen on.
- **Generator:** The Generator is used to send NetFlow packets with specified flows. Up to 30 flows can be sent with NetFlow version 5 or up to 24 flows with NetFlow version 9 and IPFIX. If a range of bytes and packets is specified, Flowalyzer will randomly assign these values between the ranges specified in each flow packet sent out. The Flow Time range can also be randomized. Reference the NetFlow v5 format or read RFC 3954 to learn more about NetFlow v9. You can also read the Charter on IPFIX.
- **Configuration:** The Configurator is used to configure NetFlow v5 and v9 on Cisco routers or NetFlow v9 on Enterasys hardware (switches and routers).
- **Communicator:** The Communicator is used to ping destination hosts with specified ports using ICMP, TCP or UDP. It can also perform trace routes.
- **Poller:** Polls Devices found in Scrutinizer or devices can be manually added. The availability and response time for each device polled is sent off to the flow collector in IPFIX datagrams.
- **Trender:** Polls devices for SNMP counters and trends the values returned.

The Flowalyzer page can provide additional information on this free NetFlow and sFlow testing utility. This utility is fully compatible with most NetFlow and sFlow reporting tools.

Interactive Scrut_util

Overview

The interactive scrut_util.exe utility provides access to numerous server maintenance utilities, including password changes, third party integration processes, many routines to access information required for support, and more.

To enable the interactive scrut_util.exe utility, run:

~/scrutinizer/bin/scrut_util.exe

Which will open this Scrutinizer prompt:

SCRUTINIZER>

To close the interactive prompt, type 'exit':

SCRUTINIZER> exit

Exiting...

[root@Scrutinizer ~]#

Help Function

The list of available commands can be displayed from within the Scrutinizer utility shell by entering:

SCRUTINIZER> help

For help with specific commands (for example, the “show” command) enter:

SCRUTINIZER> help show

For help with specific extended commands (for example, the “show groups” command) enter:

SCRUTINIZER> help show groups

Commands

Following are the available top level commands.

check
ciscoise
clean
collect
counteract
delete
disable
download
enable
endace
expire
export
import
moloch
optimize
repair
services
set
show
snoop
system
upload
version

For each top level command, there are several extended commands as listed below.

check

Runs a test or check against the command provided.

Command	Description
check activeif	Checks for active flows by looking at active interface details and lists the last

	timestamp and number of interfaces that received flows.
check data_last_written	Checks the activity of collected flow data written to the database.
check database <db_name> <db_pass>	Checks the specified database for errors.
check dist_info	Checks and displays distributed information of the Scrutinizer Server.
check hdtest	Tests the performance of the hard drive. This is a good way to determine if the hardware is adequate for Scrutinizer's current flow volume.
check heartbeat <database|api>	Checks heartbeat functions to make sure Scrutinizer is internally communicating properly.
check history_index	Checks history_index for 1m interval table activity.
check history_index_empty_tables	List tables with zero rows from history_index. Please stop the collector prior to running this command. This command will *not* delete entries reported. To do so, use delete instead of check.
check history_index_orphans	Checks entries from history_index for which a table does not actually exist. This should never happen, but occasionally when things go wrong we need something like this to make cleanup easier. This command will *not* delete entries reported. To do so, use delete instead of check.
check history_table_orphans	List tables with no history_index entry. Please stop the collector prior to running this command. This command will *not* delete entries reported. To do so, use delete instead of check.
check interfaces [all|cisco|hauwei|sonicwall] [host_ip]	Tries alternative methods to retrieve interface descriptions. For Cisco and SonicWALL that means using NetFlow data. For Huawei, that means using SNMP and referencing their vendor specific MIBs.
check license	Checks and displays license details from the Scrutinizer Server.
check machine_id	Checks and displays the current machine_id of the Scrutinizer Server.
check machine_id_list	Checks and displays the current, possible, and historical Machine IDs of the Scrutinizer Server.

check objects	Verifies xcheck_hosts all have a corresponding row in objects.
check password rootdb	Checks the database root password to make sure its the same password represented in the plixer.ini.
check rollcall	Analyze rollcall and the state of roll ups per time bucket. This is used to confirm the activity of rollups on this Scrutinizer Collector.
check rollups	Lists rollups and their current state. This is used to confirm the activity of rollups on this Scrutinizer Collector.
check route <ip>	Checks device specified to determine if Scrutinizer can access its routing data.
check serverpref <serverpref>	Checks and displays the current value for the specified serverpref.
check simplercv <udp_port>	Runs a simple test to see if udp traffic is seen on the udp port provided. This command is useful to determine if flows are received at the top of the stack (i.e. tcpdump -> collector).
check snmp	Attempts to get SysObjectID for all devices. If SNMP connected successfully, it will return the credential object. Otherwise, it will return the error message.
check ssl	Checks and lists the current settings configured for SSL parameters. Use the set ssl command to modify settings or enable/disable SSL.
check stats_exporters	Lists statistical details related to time and exporter activity.
check task <id>	Checks the execution times and error codes for the specified task <id>. A list of tasks is available by using the show task command.
check thresholdpolicies	Checks and verifies that all threshold policies are correct and functioning properly.
check version	Checks to see if a newer version of Scrutinizer is available.

ciscoise

Manage CiscoISE Node Integration with Scrutinizer.

Command	Description
ciscoise add <ise_ip> <ise_tcp_port> <ise_user>	Adds a CiscoISE node to the queue to acquire user identity on all active sessions. The required parameters are the host address <ise_ip>, tcp port <ise_tcp_port>, and user <ise_user> that can access the API are required. Scrutinizer will prompt the user for the <ise_user> password.
ciscoise check	Tests polling and outputs the results to the screen for review. It's a good way to verify Scrutinizer is collecting user identity information properly.
ciscoise kick <ise_id> <mac_address> <user_ip>	Kicks the user off the ISE node forcing them to re-authenticate. Minimally the users IP address is required. Optionally, the <mac_address> can be provided.
ciscoise nodelist	Lists the currently configured CiscoISE nodes.
ciscoise poll	Runs a poll manually and outputs the results to the screen. When integration is enabled, polling automatically performed routinely. To diagnose issues, run 'ciscoise check' or 'ciscoise test'
ciscoise remove <ise_ip>	Removes a CiscoISE node from Scrutinizer. The required parameter <ise_ip> is the IP address of the CiscoISE node.
ciscoise test	Tests polling and outputs the results to the screen for review. It's a good way to verify Scrutinizer is collecting user identity information properly.
ciscoise update <ise_ip> <ise_tcp_port> <ise_user>	Updates existing configuration settings for a specific CiscoISE node. The required parameters are the host address <ise_ip>, tcp port <ise_tcp_port>, and user <ise_user> that can access the API. Scrutinizer will prompt for the <ise_user> password.

clean

Executes housekeeping tasks that are scheduled to run at various times during Scrutinizer's normal operations.

**** Warning **** These commands will purge data from Scrutinizer. Please use with caution.

Command	Description
clean all	Executes several housekeeping tasks that are scheduled to run at various times during Scrutinizer's normal operations.
clean baseline	Resets all configured baselines to the default baselines for each exporter. Historical data will not be deleted. However, it will expire based on Scrutinizer's historical settings.
clean database	Cleans out temporary database entries manually. This command is executed automatically every 30 minutes by Scrutinizer's task scheduler.
clean ifinfo	Clears entries in the ifinfo db table that do not have an entry in the activeif db table.
clean pcap [<pcapfile>]	Removes all or, if specified, a specific pcapfile from the Scrutinizer server. To see a list of pcap files, execute show pcaplist
clean tmp	Removes any temporary files created by the graphing engine. Executing this will perform an on-demand clean up. By default, it is scheduled to be executed by Scrutinizer routinely.

collect

Manually collect data that is useful for Scrutinizer.

Command	Description
collect asa_acl	Manually collects ASA ACL information from CiscoASA Devices. This task is scheduled and routinely executed as part of normal operations.
collect baseline	Manually collects baseline data and check for alarms. This task is scheduled and routinely executed as part of normal operations. This command is automatically scheduled to run routinely when baselining is enabled.
collect dbsize	Collects database size information.
collect elk <elk_ip>	Manually collect data from Scrutinizer and send it to the configured ELK server. Reference the Elasticsearch / Kibana (ELK) Integration guide for more detailed information on the ELK integration.

collect optionsummary	Manually process flow option data collected by Scrutinizer. This information is routinely processed automatically.
collect pcap <in_sec> [<host>]	Collects a packet capture of the interfaces on the Scrutinizer Server. Requires a timeout (in seconds) and an optional host name in IP format to further filter the capture.
collect snmp	Manually collects SNMP data that is used during Scrutinizers operations. This process is automatically scheduled by Scrutinizer to run regularly.
collect splunk <splunk_ip>	Manually collect data from Scrutinizer and send it over to the configured Splunk server. Reference the Scrutinizer for Splunk Application integration guide for more information on the Scrutinizer for Splunk integration.
collect supportfiles	Collects various log files and server configuration data used by Plixer support to troubleshoot server issues.
collect topology	Collects various types of data from devices and Scrutinizer to help Scrutinizer understand the topology layout of the network.
collect useridentity	Manually process user identity data collected by Scrutinizer. This information is routinely processed automatically.

counteract

Third-party integration support for ForeScout CounterACT servers.

Command	Description
counteract <on|off> <counteract_ip[:port]>	Enables or disables support to ForeScout CounterACT servers. Required parameters are <on|off> and the host name and optional tcp port.

delete

This operation deletes database tables and/or database table entries.

**** Warning **** These commands will purge data from Scrutinizer. Please use with caution.

Command	Description
delete custom_algorithm <identifier>	Deletes a custom_algorithm at the system level. For more information, reference the Flow Analytics Custom Algorithms section. **** Warning **** This command will alter the behavior of Scrutinizer functionality. Please use with caution.
delete history_index_empty_tables	Deletes tables with zero rows from history_index. Please stop the collector, if running, prior to executing this command.
delete history_index_orphans	Deletes entries from history_index for which a table does not actually exist. This should never happen, but occasionally when things go wrong we need something like this to make cleanup easier.
delete history_table_orphans	Deletes tables with no history_index entries. Please stop the collector, if running, prior to executing this command.
delete orphans	**** Warning **** This command will delete all known orphan alarm events.

disable

Disables functionality used by Scrutinizer or incorporatedas part of customized development.

Command	Description
disable baseline <exporter_ip>	Disables all baselines for the specified <exporter_ip>. The historical data will not be deleted. However, it will expire based on Scrutinizer's historical data settings. **** Warning **** This command will alter the behavior of Scrutinizer baseline functionality. Please use with caution.
disable elk http://<ip:port>	Disables ELK (Elasticsearch, Logstash, and Kibana) flows from Scrutinizer to the URL specified. Reference the Elasticsearch / Kibana (ELK) Integration guide for more detailed information on the ELK integration. **** Warning **** This command will alter the behavior of Scrutinizer functionality. Please use with caution.

disable splunk http://<ip:port>	Disables splunk flows from Scrutinizer to the URL specified. Reference the Scrutinizer for Splunk Application integration guide for more information on the Scrutinizer for Splunk integration. **** Warning **** This command will alter the behavior of Scrutinizer functionality. Please use with caution.
disable vmwaretools	Disables vmwaretools for a Virtual Appliance running on VMware. **** Warning **** This command will alter the behavior of Scrutinizer functionality. Please use with caution.

download

Downloads various files and utilities useful to Scrutinizer's operations.

Command	Description
download hostreputationlists	Download the latest Flow Analytics Host Reputations Lists manually. This is also automatically updated.
download installer	Download the Scrutinizer installer to perform upgrades.

enable

Enables functionality used by Scrutinizer or incorporated as part of customized development.

**** Warning **** These commands will alter the behavior of Scrutinizer functionality. Please use with caution.

Command	Description
enable baseline <exporter_ip> default enable baseline <exporter_ip> manual <pri_element[,sec_element]>	Enables default or custom baselines (manual) based on elements from NetFlow and IPFIX templates. Baselining has several parameters available to customize the

<element> <AVG|COUNT|MIN|MAX|STD|SUM> <dailyhr|busday|sameday>	specific baseline data to collect with the 'manual' option. Please reference the Baselining Overview section for detailed configuration instructions.
enable custom_algorithm <identifier> "<algoname>"	Reference the Flow Analytics Custom Algorithms section for specific information on how to configure custom algorithms and create alarm policies.
enable elk http://<ip:port>	Enables ELK (Elasticsearch, Logstash, and Kibana) flows from Scrutinizer to the URL specified. Additional steps are required in Kibana to get the default filters, visualizations, and dashboard. Reference the Elasticsearch / Kibana (ELK) Integration guide for more detailed information on the ELK integration.
enable splunk http://<ip:port> <syslog port>	Enables splunk flows from Scrutinizer to the URL specified. The Scrutinizer for Splunk App is required on the Splunk Server. Reference the Scrutinizer for Splunk Application integration guide for more information.
enable vmwaretools	Enables vmwaretools for a Virtual Appliance running on VMware. Running enable vmwaretools also upgrades an existing install of the vmware agent.

endace

Third-partyintegration support for Endace Probes.

Command	Description
endace add <host_ip> <port> <endace_user> <endace_pass> endace remove <host_ip> endace update <host_ip> <port> <endace_user> <endace_pass>	Manages integration with Endace Probes. For more information on this integration, reference the Configuring Endace Probe Integration guide.

expire

Purges data history older then the number of days defined by Scrutinizer's history settings.

**** Warning **** These commands will purge data from Scrutinizer. Please use with caution.

Command	Description
expire alarms	Expires alarm history from the threatsoverview and fa_transports_violations tables as specified in the Data History Flow Historical 1 Min Avg preference.
expire bulletinboards	Purges alarm bulletin board events older then the number of days defined by Scrutinizer's history settings.
expire dnscache	Purges DNS cache older then the number of days defined by Scrutinizer's history settings.
expire history [trim]	Expires flow data as defined by Scrutinizer's history settings. If the optional 'trim' mode is passed, Scrutinizer will trim older data to make more space on your hard disk.
expire ifinfo	Purges old and outdated interface information.
expire inactiveflows	Expires interfaces from the interface view that have stopped sending flows. Entries are expired based on the number of hours specified in the Scrutinizer Server Preferences. (Admin -> Settings -> Server Preferences)
expire orphans	Purges alarm orphan events older then the number of days defined by Scrutinizer's history settings.

export

Run various export commands to dump data out of Scrutinizer for outside use.

Command	Description
export langtemplate <lang_name>	The <lang_name> parameter is required. If the language exists, then it will create a CSV file that shows the english and <lang_name> keys. If the language does not exist, a blank template will be created. The language file will reside at /home/plixer/scrutinizer/files/pop_languages_<lang_name>_template.csv

export peaks_csv <file> <interval> <dir> <date_range> [<group_id>]	Exports a CSV file listing interfaces and peak values based on criteria specified. Valid options for <interval> are specified as raw minutes (1, 5, 30, 120, 720, 1440, 10080). The valid <ranges> are Last24hours, LastFifteenMinutes, LastFiveMinutes, LastFortyfiveMinutes, LastFullHour, LastHour, LastMonth, LastSevenDays, LastTenMinutes, LastThirtyDays, LastThirtyMinutes, LastThreeDays, LastTwentyMinutes, LastWeek, LastYear, ThisMonth, ThisWeek, ThisYear, Today, or Yesterday. Directory must exist as a sub-directory of Scrutinizer's home directory. If specifying /home/plixer/scrutinizer/temp, then use temp as your directory. <group_id> is optional. To see a list of group_ids use show groups.

import

Run various import commands to bring external sources of data into Scrutinizer.

Command	Description
import aclfile	Imports ACL information from a file. The file must reside at /home/plixer/scrutinizer/files/acl_file.txt. The format is a direct output of SHOW ACCESS-LIST directly on the exporter.
import applications <path/file> [reset]	Import application rules from a CSV file. It is recommended to use this file for your applications import csv file: /home/plixer/scrutinizer/files/application_import.csv A reset option can be passed which will remove all application rules before the bulk import. Expected format is one named application and one application rule per line. Supported rule types are subnet, single ip, ip range, wildcard, port, and child rules. Valid application rule syntax for PostgreSQL installs are: 'subnet rule',10.0.0.0/8 'single ip rule',10.1.1.1 'range rule',10.0.0.1-10.0.0.42

	'wildcard rule',10.0.0.1/0.255.255.0 'parent/child rule','my subnet' 'ports and protocols',0-65535/256 Valid application rule syntax for MySQL installs are: 'subnet rule',10.0.0.0/8 'wildcard rule',10.0.0.1/0.255.255.0 'parent/child rule','my subnet' 'ports and protocols',0-65535/256 Applications must have at least one port rule and one rule of another type. Applications not defined this way will be imported, but may not be tagged properly in flow data. The format for a ports rule is: 'flowclass name',port-port/protocol id 'port-port' can be a range of ports (0-65535), or a single port (443-443) Frequently used protocol ids are: TCP - 6 UDP - 17 All protocols - 256 PostgreSQL installs support up to 100,000 individual application rules. MySQL installs support up to 500 individual application rules.
import asns <path/file> [<delimiter>]	Imports custom asn definitions from a csv file. The <path/file> is a required field. The path should be specified from after the /home/plixer/scrutinizer/ directory. The <delimiter> is an optional parameter and defaults to " " (i.e. space). The csv file name must be all lower-case and requires these elements, in this order: AS Number,AS Name,AS Description,IP Network(s) The fields are comma delimited, whereas the <delimiter> parameter applies specifically to the IP Network(s) element. A comma cannot be used for the IP Network(s) delimiter. Example File:

	213,my_list,what a great autonomous system,10.0.0.0/8 192.168.0.0/16 214,your_list,meh its an okay system,11.0.0.0/8 Example Command: SCRUTINIZER> import asns files/custom_asn.import
import csv_to_gps <csv_file> <group_id> [<col_json>]	Uploads latitude and longitude locations of devices from a csv file and imports them into a Google map. The csv file must be located in the '/home/plixer/scrutinizer' directory. If the csv file is in '/home/plixer/scrutinizer/files/', you will enter 'files/[name_of_file]' as the file name. The default csv file format is 'ip,latitude,longitude'. If your csv file format is different, you can specify that layout as your last command parameter. EXAMPLE "ip,lng,lat" EXAMPLE CSV FILE: 10.169.1.3,37.7749,122.4194 192.168.6.1,40.7128,74.0059 The group_id can be determined by running show groups
import csv_to_membership <csv_file> <group_type> [<col_json>]	Imports group definitions from a csv file. The csv file must be located in the '/home/plixer/scrutinizer' directory. If the csv file is in '/home/plixer/scrutinizer/files/', you will enter 'files/[name_of_file]' as the file name. The <grouptype> field refers to the map type that will be created if the group in the csv file does not already exist and can be either 'plixer' or 'google'. The default csv file format is "ipaddr,group". If your csv file format is different, you can specify that layout as your last command parameter. EXAMPLE "group,ipaddr" EXAMPLE CSV FILE: 10.169.1.3,Routers 192.168.6.1,Firewalls
import hostfile	Imports a custom hosts.txt file that contains a list of IP Addresses and hostnames. The file format is: IPv4orIPv6Address HostName Optional Description

	Example: 10.1.1.4 my.scrutinizer.rocks The Best Software in my company The file must be located at /home/plixer/scrutinizer/files/hosts.txt. **** Warning **** This command will alter database tables in Scrutinizer. Please use with caution.
import ipgroups <path/file> [reset]	Import ipgroup rules from a CSV file. It is recommended to use this file for your ipgroups import csv file: /home/plixer/scrutinizer/files/ipgroup_import.csv A reset option can be passed which will remove all ipgroup rules before the bulk import. Each line of the file is an individual ipgroup with the name of the group as the first field and the rules of the group separated by a space in the second field. Supported rule types are subnet, single ip, ip range, wildcard and child rules. Any child groups must already exist in Scrutinizer or be declared in the import file BEFORE it can be used as a rule in another group. Valid ipgroup rule syntax for PostgreSQL installs are: 'subnet rule',10.0.0.0/8 'single ip rule',10.1.1.1 'range rule',10.0.0.1-10.0.0.42 'wildcard rule',10.0.0.1/0.255.255.0 'parent/child rule','my subnet' Valid ipgroup rule syntax for MySQL installs are: 'subnet rule',10.0.0.0/8 'wildcard rule',10.0.0.1/0.255.255.0 'parent/child rule','my subnet' If a named ipgroup within the import file already exists in Scrutinizer, that group's existing rules will be overwritten by those declared in the file. PostgreSQL installs support up to 100,000 individual ipgroup rules. MySQL installs support up to 500 individual ipgroup rules.

moloch

Third-partyintegration support for Moloch Probes.

Command	Description
moloch <on|off> <moloch_ip> <moloch_port>	Manages integration with Moloch Probes. The <moloch_port> is optional.

optimize

Run various optimization tasks.

Command	Description
optimize common	Optimizes tables that are commonly inserted and deleted. This action keeps things neat and clean for the database. This command is routinely executed as part of normal operations. **** Warning **** This command will alter database tables in Scrutinizer. Please use with caution.
optimize database <db_name> <db_pass>	Optimizes the tables in the database specified.

repair

Runvarious database check and repair commands.

Command	Description
repair business_hour_saved_reports	Saved reports prior to 15.5 that were saved with business hours will require a manual check and repair. This command converts older saved reports with business hours specified to the newer format.

repair database <db_name> <db_pass>	Repairs errors for the database specified.
repair history_tables	Fixes history tables that have the wrong col type for octetdeltacount. It may be updated in the future to address other issues. **** Warning **** This command will alter database tables in Scrutinizer. Please use with caution.
repair policy_priority_order	With some professional services and automated policy creation, some policy IDs have been known to get out of whack (or duplicated). This function fixes that. **** Warning **** This command will alter data in Scrutinizer. Please use with caution.,
repair range_starts	Fixes history tables that may not have a start time that helps identify the range of data within the individual history tables. **NOTE:** This command may take a long time to complete. Only execute under the direction of technical support. **** Warning **** This command will alter database tables in Scrutinizer. Please use with caution.

services

Manages the Scrutinizer services.

Command	Description
services <service|all> <action>	Starts, stops, or restarts the specified service (or all services). **** Warning **** This command will alter Scrutinizer's operations. Please use with caution.

set

Modifies certain behaviors on how Scrutinizer authenticates and performs operations.

Command	Description

set dns	Modify system file to manage list of dns servers. This command will remove any preconfigured dns servers. Use show dns to see what is currently configured. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.
set hostinfo <ip_address> <fqhn>	Set the local machine name to the fully qualified host name provided. Ensures that /etc/hosts is configured to resolve between the given <fqhn> and <ip_address>. ** **Warning** ** This command will alter how Scrutinizer behaves. Please use with caution.
set httpd <port>	Change the web port of non-ssl installs for the Scrutinizer WebUI. Use set ssl to change the SSL port. ** **Warning** ** This command will alter how Scrutinizer and/or users access data. Please use with caution.
set myaddress	Change the IP Address of the current Scrutinizer Server. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.
set ntp	Modify system file to manage list of ntp servers. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.
set partitions <partition_name>	Expand the operating system disk space for hardware and virtual appliances. **NOTE:** Make a backup before using this command. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.
set password rootdb	Modify the database's root password. ** **Warning** ** This command will alter how Scrutinizer and/or users access data. Please use with caution.
set password scrutdb	Modify the database's scrutinizer password. ** **Warning** ** This command will alter how Scrutinizer and/or users access data. Please use with caution.

set password webui <user>	Modify the webui password for the specified user. ** **Warning** ** This command will alter how Scrutinizer and/or users access data. Please use with caution.
set registercollector	Manually register this collector for both stand-alone and distributed use. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.
set reportmenu	Manually recreate the report menu. **NOTE:** The report menu is automatically maintained based on the flows received.
set salt <salt>	Setting a salt value will allow users to mask certain machine characteristics from any license key generated. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.
set selfregister [reset]	Manually registers this Scrutinizer Server to identify itself for both stand-alone or distributed functionality. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.
set selfreporter	Promotes this Scrutinizer Server to a reporter. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.
set sshcollectorkeys	Generate a new SSH key pair, and distribute it to all active, registered machines. Any previous SSH key pairs will be overwritten unconditionally, making this suitable for resynchronizing SSH access should problems arise. This enables future functionality to perform upgrades and other maintenance operations en masse. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.
set ssl <on|off>	Enable or disable SSL support in Scrutinizer. It only works with the local Apache server bundled with Scrutinizer. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution. Please reference the System/SSL section for detailed configuration instructions.

set timezone <timezone>	Set the server's time zone. To see a list of time zones, run show tzlist ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.
set voip <on|off>	Toggles the predefinition of VoIP port ranges on or off. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.
set yum_proxy <host> <port> <user>	Used to set up yum proxy setting in the yum configuration file. This command will remove any previously configured proxy servers. All fields are required. Once all fields are entered on the command line, you will be prompted for the users password. To see what proxy servers are currently configured, use show yum_proxy ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.

show

Shows various details about the Scrutinizer Server.

Command	**Description**
show alarms [filter]	Displays a list of alarms ordered by timestamp, descending.
show custom_algorithms	Displays a list of custom algorithms available and whether they are enabled. For information on managing custom algorithms, reference the Flow Analytics Custom Algorithm section.
show diskspace	Displays details about available storage.
show dns	Displays a list of DNS servers currently used to resolve hostnames. Use the set dns command to change the list of DNS servers.

show exporters [filter]	Displays a list of exporters that are currently sending data to Scrutinizer based on the supplied filter (if any).
show extalarms [filter]	Displays a list of alarms with extended json data ordered by timestamp, descending.
show groups	Displays a list of groups currently configured on this Scrutinizer server.
show interfaces [filter]	Displays a list of interfaces that are currently sending data to Scrutinizer based on the supplied filter (if any).
show ipaddresses	Displays the current ip address(es) on this Scrutinizer server.
show metering [filter]	Displays a list based on the supplied filter (if any) of matching exporter IPs and how each interface is metered (i.e. ingress and/or egress).
show ntp	Displays a list of NTP servers currently used to sync time.
show partitions	Displays a list of partitions on the current Scrutinizer Appliance. This command is only available for Hardware and Virtual Appliances. Use show diskspace if you're looking for diskspace per volume (or partition).
show pcaplist	List what current pcap files have been created and their sizes. pcaps can be removed using the clean pcap command.
show serverpref [filter]	Displays serverprefs and their current values. The filter parameter is optional to narrow the serverprefs to match the string provided.
show task [name]	Displays a list of tasks currently configured in Scrutinizer. The name parameter is optional to narrow the task names to match the string provided.
show timezone	Displays the current timezone of this Scrutinizer Server. Use set timezone command to modify the timezone.
show tzlist [filter]	Displays the list of timezones.
show unknowncolumns	List info elements from exporters that are unknown to Scrutinizer. Don't fret! Just give this list to Plixer and we will add support for you!

show yum_proxy	Displays the currently configured yum proxy settings. To change these settings, use set yum_proxy

snoop

Listens at the interface level for traffic from the specified interface or ip address.

Command	Description
snoop interfaces <interface_name>	Listens at the interface level for traffic from the specified interface.
snoop ipaddresses <ip_address>	Listens at the interface level for traffic from the specified ip address.

system

Scrutinizer system level functions.

Command	Description
system <update|restart|shutdown>	Performs system level functions such as rebooting, shutting down, or applying OS level patches. ** **Warning** ** This command will alter Scrutinizer's operations. Please use with caution.

upload

Uploads files for troubleshooting purposes.

Command	Description
upload pcap <capturefile>	Uploads the specified packet capture collected by the collect pcap command. To see a list of captures on this server, execute show pcaplist

upload supportfiles	Uploads files for troubleshooting purposes.

version

Displays Scrutinizer version.

Command	Description
version	Shows version information about Scrutinizer.

Language Translations

Overview

This software can be translated to another language. To translate or localize Scrutinizer to another language, navigate as follows:

1. Admin Tab -> Definitions -> Language

2. Select a language and make updates. Notice the pagination at the bottom, there are well over 1000 translations.

3. languages are saved as ~scrutinizer/files/localize_LANGUAGENAMEHERE.xls

4. Contact support and they will create a file that can be imported into Scrutinizer to support your language.

Mailinizer Setup

Below are the steps to set up this IPFIXify (TM) Plug-in

1. If you haven't done so, please download the latest version of the IPFIXify Agent from www.plixer.com and extract those files to a directory.

2. Extract the files contained within this plug-in to the directory where the IPFIXify Agent resides.

3. Modify the msexchange_2010.cfg file and change the collector directive at the top to the COLLECTOR_IP:PORT of your IPFIX Collector. If you are unsure of which port to use, port 2055 is the most common port.

4. To verify that the data is being exported properly, run the following command from the IPFIXify directory:

ipfixify.exe --file="path/to/MSGTRK*" --config="path/to/msexchange_2010.cfg" --verbose

Please specify the absolute path to both the msexchange_2010.cfg and

MSGTRK*.log files.

Example: ipfixify.exe --config "C:\Program Files (x86)\IPFIXify\msexchange_2010.cfg" --file "C:\Program Files\Microsoft\Exchange Server\V14\TransportRoles\Logs\MessageTracking\MSGTRK*"

5. IPFIXify should (but not required) start as a service if you wish to automatically collect data whenever the system restarts.

ipfixify.exe --install auto --name="SVCNAME" --config="path/tocfg" --file="path/tofile"

Replace SVCNAME with a short descriptive name and specify the absolute path to your msexchange_2010.cfg and MSGTRK* files. All four parameters are required.

For more information on IPFIXify options, execute ipfixify.exe without any parameters.

Meta Data Collection

Overview

By default, Plixer gathers meta data that determines the general health, overall performance, and targeted metrics to help Plixer improve the Scrutinizer Platform.

How often it is collected

Meta data is collected once a week at a randomly scheduled time during typical peak business hours (M-F 9a to 4p). The actual scheduled time is determined at the time Scrutinizer is installed.

What meta data is collected

- **Vendor** - identifies if the Scrutinizer installer is branded as Plixer or an OEM vendor.
- **smtpfromEmail** - The email address configured as the administrator in the server preferences.
- **smtp** - the email server name which is used to identify the customer's company as configured in the server preferences.
- **installedVersion** - the current version of Scrutinizer installed.
- **License Details** - Identifies license key details such as license level and license state (e.g. valid or expired).
- **How many exporters are actively sending flows?** - A count of exporters that have sent flows in the last 24 hours.
- **Server Metrics** - How many flows per second received, packets per second received, and flows drops per second at the time the data is gathered.
- **Flow Template Details** - All templates and element names exported from the exporters Scrutinizer is currently collecting data.
- **Flow Analytics Exporter Statistics** - A count of exporters configured for each algorithm.
- **Flow Analytics History** - A list of violating IP addresses from external sources. It is used to determine if outside patterns exists from the aggregated data between globally installed Scrutinizer Servers.

How is the data transferred

The data is encrypted. It is sent to ph.plixer.com. Port 443, 80, and 25 are used depending on port availability.

What Plixer Does with the Data

Plixer uses the collected data for support purposes only. From the data we can learn which customers could benefit from a support call to upgrade a system or to fix issues.

How do I turn it off

By default, this functionality is enabled. To turn off this functionality, go to the Admin Tab, select the Settings Menu, and click System Preferences.

Uncheck the option "Share Health Statistics", scroll to the bottom, and click save.

Multi Tenancy Module

Overview

The Multi Tenancy module provides the following features:

- Apply permissions to User Groups per Interface
- Apply permissions to User Groups per Device

The Multi Tenancy module is useful to companies who need to give customers a unique login and restrict what they see to specific devices and or interfaces.

NetFlow Knights

The Help tab is a link to https://www.netflowknights.com. This site is the on-line support community and is used to bring subjects of interest directly to customers and evaluators. This is done in several ways, some of which include:

- A frequently updated blog
- The on-line support forum

Posting a comment or question requires membership. Click here to join.

Search Tool

Overview

The Search Tool is launched by navigating to Status tab ->Search. This tool provides the means to search through all of the flows stored in the database for specific flows.

There are two search options available:

- Exporters search
- Hosts search

NOTE: Only the 1 minute interval tables contain 100% of all flows collected. To make sure you are querying 1 minute interval data, limit the search to under 1 hour of time. Visit the Admin tab -> Settings -> Data History to increase the "Maximum Conversations" saved per interval to increase the volume of flows saved per interval. Be aware that this may require more hard disk space. Visit the Dashboard tab -> Vitals (or Vitals Reporting in the Status tab) to view how much hard drive space is being consumed.

Exporters search

The Exporters search allows you to search on the following fields:

- Source Host
- Destination Host
- Source or Destination Host
- Client
- Server
- User as Source
- User as Destination
- Wireless Host
- Wireless SSID

The User as Source and User as Destination search fields allow you to search by Username.

Other search options for the Exporters search include searching:

- Either All exporting devices or a specific exporter
- Selecting the time range for the search.
    - The time range can be either a predefined time range, such as Last 5 minutes, Last Ten Minutes, etc., or a custom timeframe.

If flows meet the search criteria for the Exporters search, a Host to Host report will return the results of the search.

Hosts search

The Hosts search is used to perform a quick search for hosts in flows from devices that are sending flows to Scrutinizer.

Simply enter the host IP address in the search textbox and click the Search button. If the host is found as either Source or Destination in any flows stored in the database, Scrutinizer will return a list including:

- Device (exporter's IP address)
- First Seen
- Last Seen
- Flow Count

Clicking on an IP address in the Device list will open a Report menu. The report selected will report on the last hour of flows received by the host selected.

The Hosts search requires that Host Indexing in Admin>Settings>System Preferences is enabled.

SSL

Configuring SSL in Scrutinizer

Enabling and disabling SSL support in Scrutinizer is done within the Interactive scrut_util.exe shell. It only works with the local Apache Server bundled with Scrutinizer.

To open the Interactive scrut_util.exe, use the following command:

/home/plixer/scrutinizer/bin/scrut_util.exe

The Scrutinizer prompt will then display:

SCRUTINIZER>

Enabling SSL

To enable SSL, at the Scrutinizer prompt, enter:

SCRUTINIZER> set ssl on

**** Warning **** This command will alter Scrutinizer's operations. Please use with caution.

Scrutinizer will then prompt for the following mandatory fields:

Enter the secure tcp port to be used. ex: 443

Enter the two-letter abbreviation for your country. ex: US

Enter the state/province of your organization. ex: Maine

Enter the city of your organization. ex: Kennebunk

Enter the name of your organization. ex: Plixer

Enter the contact email address. ex: name@company.com

Enter the server name or IP of the Scrutinizer server.
ex: 1.2.3.4 or scrutinizer.company.com

Enter the key encryption size. [2048|4096] ex: 2048

Name Field	Explanation
Country Name	The two-letter ISO abbreviation for your country *example: US = United States*
State / Province	The state/province where your organization is located. Can not be abbreviated. *example: Maine*
City / Locality	The city where your organization is located. *example: Kennebunk*
Organization	The exact legal name of your organization. Do not abbreviate. *example: Plixer*
Email Address	The email address for the CA (who to contact) *example: someone@your.domain*
Common Name	URL that you wish to attach the certificate. *example: 10.1.1.10 or scrutinizer.company.com*
Key Size	2048, 4096 *example: 2048*

Creating a signed certificate

1. Enable SSL by running **ssl on** as described above.
2. Send the /etc/pki/tls/private/ca.csr file to the Certificate Authority (CA) and ask them to sign it and return it as base 64 encoded and not DER encoded.
3. When you receive the signed SSL cert, stop the apache service within Interactive scrut_util:

   SCRUTINIZER> services httpd stop

4. Replace the active SSL Cert with the new one and rename the file to /etc/pki/tls/certs/ca.crt
5. Then start the apache service.

   SCRUTINIZER> services httpd start

Disabling SSL

To disable SSL, at the Scrutinizer prompt, enter:

SCRUTINIZER> set ssl off

System LEDs

Overview

The System LEDs are shown in the upper right hand corner providing the status on server health and other various critical operations of the flow collection and reporting architecture.

Preceding the LEDs is the letter **P** indicating that this is a Primary reporting server. In a single server installation of Scrutinizer, it will always be **P**. In a distributed server environment, this letter can alternatively be **S** for the Secondary reporting server.

The LED's color and icon indicate the status of the LED:

- Operational – Green (checkmark)
- Degraded –Orange (exclamation point)
- Critical – Red (X)

The LED status is based on the most critical level of severity of any item within the detail of the LED modal.

Also, all columns within each LED modal are both sortable and searchable.

P [x] [ ] [ ]

Scrutinizer Server Health: [*status*]

This LED provides vital server statistics information such as CPU utilization and memory and disk storage availability.

The following information is available per Server, with the color change occurring at the threshold levels defined.

All of the entries in the Server Health LED modal link to trend reports except for the Free DB % data entry.

	Description	Thresholds		
		Green	Orange	Red
CPU	CPU Utilization percentage	<= 60%	<= 85%	> 85%
Memory	Available memory	>= 8GB	>= 4GB	< 4GB

Free Memory	Available Free memory	>= 2GB	>= 1GB	< 1GB
Free Disk DB*	Available Free Disk space for database.	>= 10GB	>= 5GB	< 5GB
Free Disk DB (Hardware Appliance)	Available Free Disk space for database on Hardware Appliances.	>= 10% of available disk space	>= 25GB	< 5GB
Free DB %	Percent of disk storage that is still available for database.	N/A	N/A	N/A
DB Latency	Server to server database latency (in milliseconds)	< 250ms	>= 250ms	N/A
API Latency	Server to server reporting latency (in milliseconds)	< 1000ms	>= 1000ms	N/A
Clock Drift	Server to server clock difference (in seconds)	0s	<> 0s	N/A

*** Free Disk DB –**

If disk space drops below 10GB and Auto History Trimming is selected in Admin>Settings>Data History, Scrutinizer will automatically start trimming historical data until space available is greater than 10GB again.

If disk space drops below 1.5GB, the collector will stop saving flows. In the event that the collector stops, a utility can be run that expires historical data to free up space. You will need to go to Admin>Settings>Data History, and adjust the current retention settings.

On the server, access the interactive scrut_util prompt with the following command:

/home/plixer/scrutinizer/bin/scrut_util.exe

At the scrut_util prompt, run:

SCRUTINIZER> expire history

When the above command runs, it looks at the settings in the master data history configuration then purges historical data based on the current time. Then run the following command to restart the Plixer Flow Collector service and you will begin receiving and processing flows again.

SCRUTINIZER> services plixer_flow_collector start

PLEASE NOTE: Disk space thresholds are higher for Hardware Appliances

Hardware Appliances are intended for much higher flow volumes, therefore disk space is consumed at a much higher rate.

For that reason, the threshold for when Auto History Trimming will automatically start trimming data on Hardware Appliances is initially set to when disk space drops below 10% of available disk space. The threshold setting for when the collector stops saving flows is when disk space drops below 5 GB.

P [ ] [x] [ ]
Scrutinizer Software Health: [*status*]

The following information is available per Server from this LED:

- Role – primary/secondary/collector (secondary and collector used in distributed server environments)
- API – Reporting (web) interface, Up/Green means the web server is running. Down/Red means the web server is down.
- Database – Up/Green means the database is running. Down/Red means the database is down.
- Collector- Up/Green means the collector service is running. Down/Red means the collector service is down.
- Alarms – Up/Green means the Alarms/syslog service is running. Down/Red means that the Alarms service is down.
- Version – current running Scrutinizer version
- Flow Rate – flows per second
- MFSNs – Missed Flow Sequence Numbers (see Scrutinizer Exporter Health LED for more information on MFSNs)

P [ ] [ ] [x]
Scrutinizer Exporter Health: [*status*]

This LED reports data collected by the Plixer flow collector service. The flow collector receives flow data from network devices, processes it, and stores it in the appropriate database tables. The collector is also responsible for rolling raw 1 minute data into 5 minutes, 30 minutes, 2 hours, 12 hours, 1 day, and 1 week intervals. Currently the collector service supports NetFlow v1, v5, v6, v7, v8, v9, IPFIX and sFlow v2, v4 & v5 as well as jFlow, cflowd, NetStream and others.

- To run the collector from the command prompt (i.e. CLI) type:
  -
  - /home/plixer/scrutinizer/bin/scrut_collector.exe

NOTE: the output is for internal use only.

Information available per Exporter:

- Collector – IP Address of the server collecting the flows
- Exporter – IP Address/name of exporter
- Flows – flow collection rate (green=flows/orange=none)
- Packets – packet rate (green=packets/orange=none)
- MFSN – This led turns orange if Missed Flow Sequence Numbers exceed 300.
  If only one or a few of the flow sending devices report high MFSNs, it is likely the network or the flow exporting device that is dropping or skipping flows. If all devices report high MFSNs, it is likely to be the collector that is dropping flows.

  To improve performance, make sure the server hardware meets the minimum requirements. Visit the Vitals dashboard for trending details of the server.
- Max Flow Duration –

  This LED turns orange if the collector is receiving flows with a total flow duration beyond 60 seconds. Make sure these Cisco or similar commands have been entered on the flow exporting device (e.g. routers or switches):

  ```
  ip flow-cache timeout active 1
  ip flow-cache timeout inactive 15
  ```

  Learn more about the ip flow-cache timeout commands.

- Templates – template count
- Last Flow – timestamp of last flow received

Troubleshooting

Getting Started Guide

Contact Support: For assistance setting up the server or the collector or for navigation techniques.

How to enable NetFlow or sFlow on various hardware.

System LEDs: Familiarize yourself with these. They should all be green.

FAQ: This page lists many common questions we have received over the years.

Webvideos: These are short 2-5 minute videos that offer good general help with different areas of the software.

Vitals

Overview

View vital information on how well the server is handling the NetFlow and sFlow collection volume. More details can be found at Vitals Reporting and Vitals Dashboard.

Web Server Port

Overview

This software runs on the Apache Web Server.

Use the interactive scrut_util.exe utility to change the web server port Scrutinizer is running on.

1. Log on to the Scrutinizer server with administrative permissions
2. Open the Interactive scrut_util prompt with the following command:

   /home/plixer/scrutinizer/bin/scrut_util.exe
3. And run:

   SCRUTINIZER> set httpd <port>

Use this command to change the web port of non-ssl installs of Scrutinizer. Do not edit the httpd.conf manually or certain functionality will not work properly.

**** NOTE: This needs to be run on all servers in the cluster.

SSL Support

To change the port for SSL, run:

SCRUTINIZER> set ssl on

Then follow the instructions provided.

404 Error

We are sorry, the page you requested cannot be found. Please contact support directly for more information about your query.

500 Error

Internal Server Error: Please contact support.

Implementation Guides

Configuring Amazon Web Services flow streaming

Overview

The integration between Amazon Web Services (AWS) and Scrutinizer provides insight into network traffic destined for AWS such as:

- AWS connection latency
- Top AWS users
- Top AWS applications
- Overall traffic load of AWS hosted applications

Configuring Amazon Web Services using Amazon Kinesis, a secure, reliable transport used by Amazon to stream logging data, to send flow streams to Scrutinizer provides the following AWS specific flow reports in Scrutinizer.

- Action
- Action with Interface
- Action with Interface and Dst
- Action with Interface and Src
- Interface
- Pair Interface
- Pair Interface Action

Prerequisites

A minimum of Scrutinizer v16.10 is required for the configuration to successfully complete and for the AWS reports to be available.

Required Information

Before we start the AWS configuration in Scrutinizer, we need to collect some information.

1. Your AWS Account ID
    a. In the AWS interface, click on your name in the upper right hand corner.
    b. Select 'My Account'.
    c. Your Account ID is under 'Account Settings'.
2. What region you will be monitoring
    a. Click the AWS console home link (orange box in upper left hand corner).
    b. Your region will be in the url. Example: us-west-2
3. What VPC you will be monitoring
    a. Go to AWS > VPC > Your VPCs
    b. Make note of the VPC ID of the VPC you want to collect flow data for.

Create an AWS User for Flow Streaming

The next step is to create an AWS user for flow streaming.

1. Go to AWS > Security & Identity > Identity and access management
2. Users > Create User
3. Enter a user name. Example: 'ScrutinizerConsumer'
4. Click 'Create'
5. Click 'Download Credentials' (to be used in 'AWS Configuration in Scrutinizer' below)
6. Edit that user: Permissions > Attach Policy and add:
    - AmazonEC2ReadOnlyAccess
    - CloudWatchFullAccess
    - AmazonDynamoDBFullAccess
    - AmazonKinesisFullAccess
    - AdministratorAccess
        - This policy can be removed from this user after the configuration is complete.

AWS Configuration in Scrutinizer

Next we'll add the AWS configuration settings in the Scrutinizer user interface.

1. Add the following ScrutinizerConfig credentials in Admin > Settings > AWS Configuration
    a. AWS access key ID

        This is included in the credentials downloaded for the flow streaming AWS user just created.

    b. AWS Region Name

        Gathered in the Required Information step.

    c. AWS secret access key

        This is included in the credentials downloaded for the flow streaming AWS user just created.

    d. AWS Stream Name

        Defaults to 'ScrutinizerStream' and is generated in step 5 (aws_config.exe).

    e. Flow Collector IP

        IP Address of the server where step 2 (--enable aws) will be run.

    f. Flow Collector Port

Default 4739

2. Run /home/plixer/scrutinizer/bin/scrut_util.exe –enable aws
    a. Please note: The first time this is run on a server that has been upgraded from an earlier version of Scrutinizer, it will take a few minutes to install all the required libraries.
3. Change to the bin directory:
    a. cd /home/plixer/scrutinizer/bin/

Note: Step 4 is for versions 16.10 and 16.11 only. The aws_config.exe file is included with the install and upgrade of versions above 16.11.

4. Download aws_config.exe to the /home/plixer/scrutinizer/bin/ directory.
    a. wget http://files.plixer.com/support/aws_config.exe
    b. chmod 750 aws_config.exe
5. Run ./aws_config.exe --configure --region XXX --account_id YYY --vpc_id ZZZ
    a. region, account_id, and vpc_id are all from the information gathered in the Required Information section.
    b. This will create the AWS Stream Name 'ScrutinizerStream'.
6. Restart the AWS service
    a. Run: service plixer_aws restart
        - It can take 10 minutes before you see the exporter.
        - Note: If you do not see the exporter after 10 minutes, verify that the exporter is not disabled in Admin > Definitions > Manage Exporters.
7. Now you can detach the AdministratorAccess policy from your AWS user.

Frequently Asked Questions

Q: Why do I see "Unable to load AWS credentials" installing modules?

A: Amazon's package tests assume AWS connectivity is already in place. These can be ignored.

Q: Why are there gaps in the data in 1m and 5m intervals?

A: Amazon flow logs are updated every 10 minutes.

Q: It isn't working, how can I see what is going on?

A: Turn on debugging and review the log files (see the 'Troubleshooting AWS' section below).

Q: All commands ran successfully but I don't see the exporter?

A: If you do not see the exporter after 10 minutes, verify that the exporter is not disabled in Admin > Definitions > Manage Exporters.

Troubleshooting AWS

Overview

AWS integration involves a number of moving parts.

- On the Amazon side of things:
  - Cloudwatch logs are streamed to Kinesis for consumption.
- On the Scrutinizer side:
  - The plixer_aws daemon takes settings from 'Admin > Settings > AWS Configuration' and connects to the Kinesis stream. That data is then sent to Scrutinizer as IPFIX.

Enabling

If upgrading from a version prior to 16.10, required libraries will be installed the first time you run './scrut_util.exe --enable aws'. That will also install the plixer_aws service.

Running

When the service is started all necessary configuration files are rebuilt. That means that changes to settings require a restart of the service.

```
service plixer_aws restart
```

A plixer_aws.log file is created under /home/plixer/scrutinizer/files/logs/ that will show the service start time.

Debugging

1. Stop the plixer_aws service.
2. Go to /home/plixer/scrutinizer/bin/
3. Run this command: './scrut_util.exe --enable aws --debug'
4. Go to /home/plixer/scrutinizer/files/
5. Run this command: './kinesis.sh'
6. Collect plixer_aws_debug.log and plixer_aws.log files from /home/plixer/scrutinizer/files/logs

Configuring Cisco's FireSIGHT eStreamer Client

Overview

Cisco FireSIGHT Management Center centrally manages network security and operational functions for Cisco ASA with FirePOWER Services and Cisco FirePOWER network security appliances.

Configuring the FireSIGHT eStreamer Client to send flows to Scrutinizer will make the following flow reports available:

- App Internet HTTP Host
- Application E-Zone & Sub Type
- Application I-Zone & Sub Type
- Firewall List
- Ingress and Egress Zones
- User App HTTP Host
- User App HTTP URL
- User Application
- Web App & CoS
- Web App Event & Rule Details
- Web App and Source IP

Prerequisites

Minimum required versions are:

- Scrutinizer v16.2

- eStreamer 5.4

Required Information

Before we start the eStreamer Client configuration, we need two pieces of information:

- Your Scrutinizer collector's public IP Address
    - It's 10.30.11.5 in the example

- Your FireSIGHT eStreamer IP Address
    - It's 10.1.2.70 in the example

Register Scrutinizer with FireSIGHT

We'll start the Client configuration in the FireSIGHT Defense Center.

1. Log into your FireSIGHT Defense Center
    a. For Firepower v5.4: Navigate to System > Local > Registration

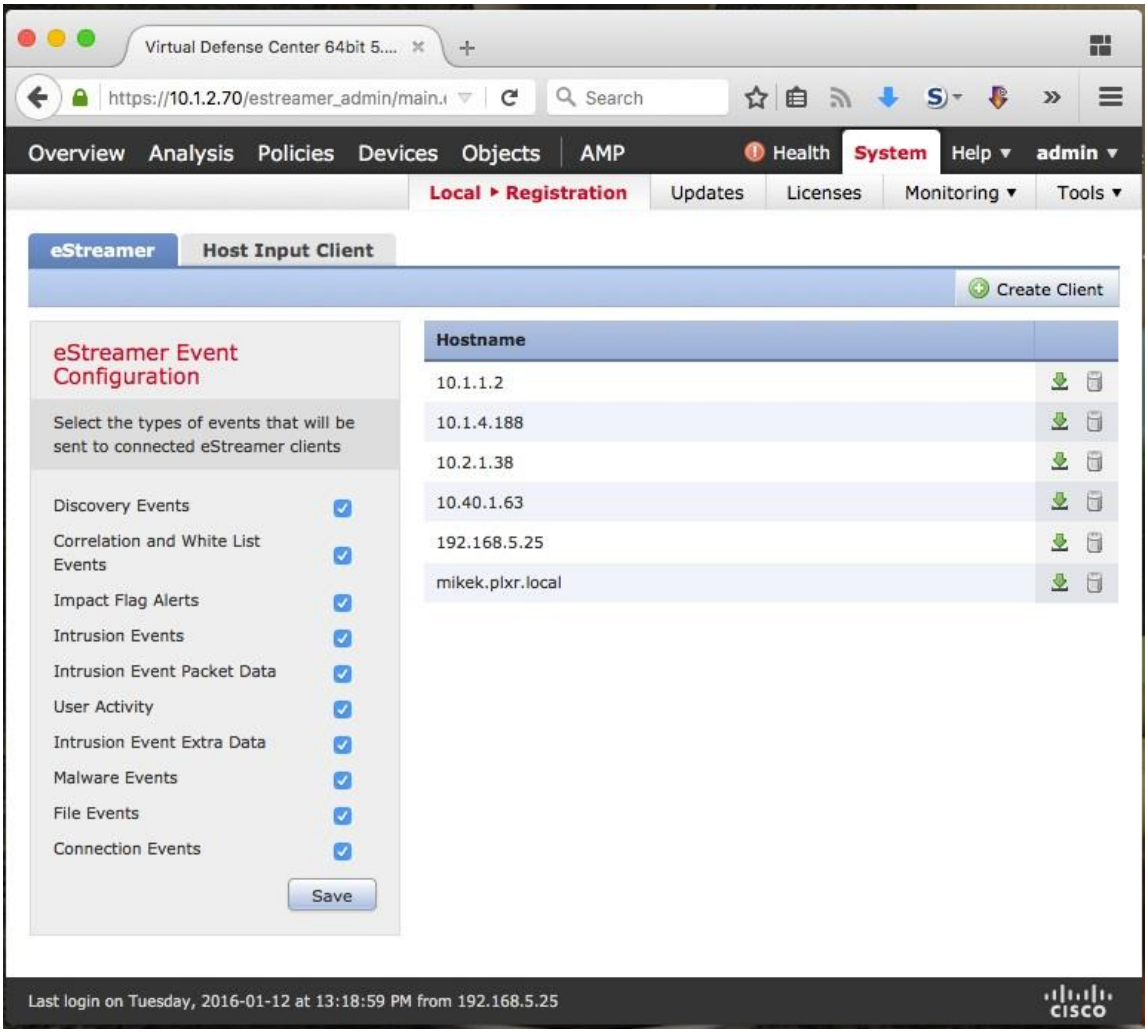

    b. For Firepower v6.x: Navigate to System > Integration > eStreamer

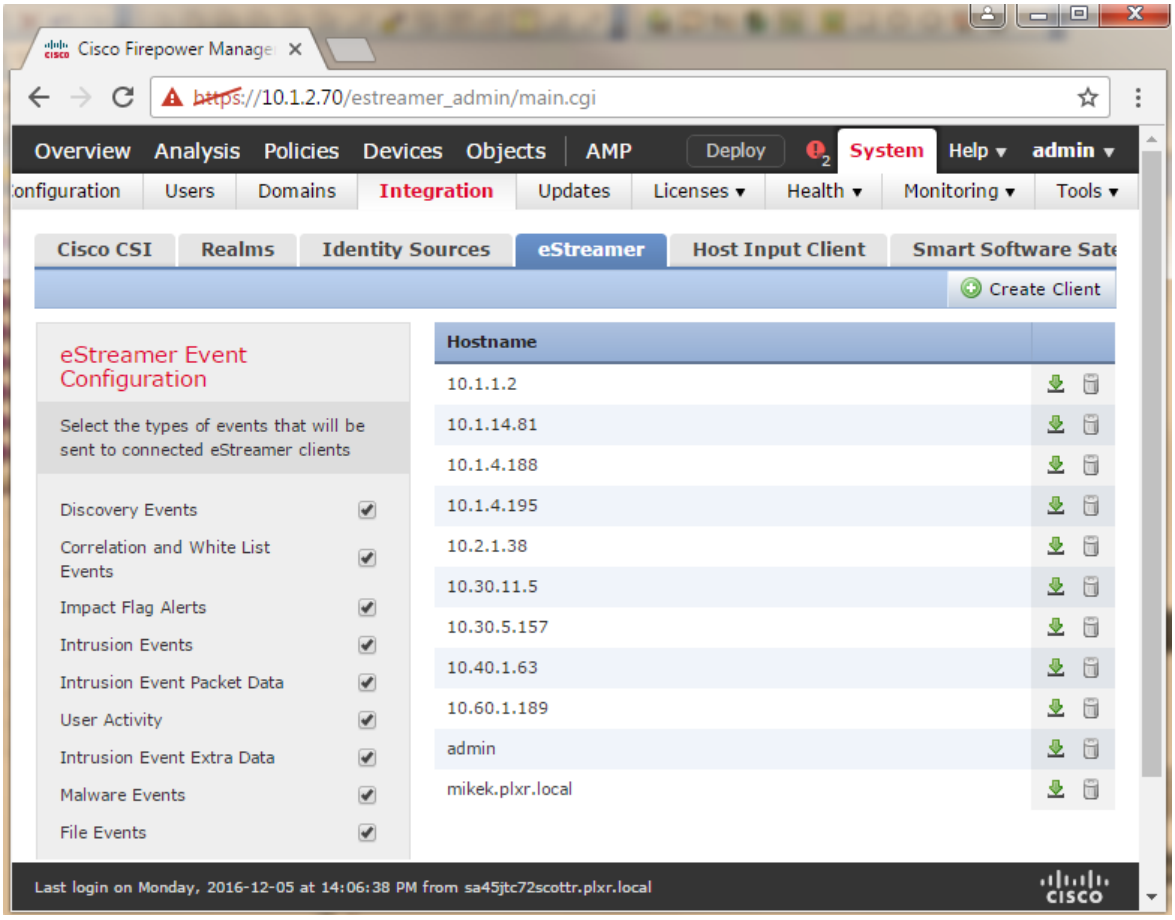

(The remaining steps apply to both versions of Firepower.)

2. Enable all eStreamer Events, and click the "Save" button at the bottom of the list.
    a. Wait for the page to refresh. It may not give any other indication that a change has been made.
3. Click on the "(+) Create Client" button on the right.
4. Enter the Scrutinizer collector's public IP Address
5. Enter a password (optional)
    a. If you do enter a password, remember it. It will be needed in a later step.

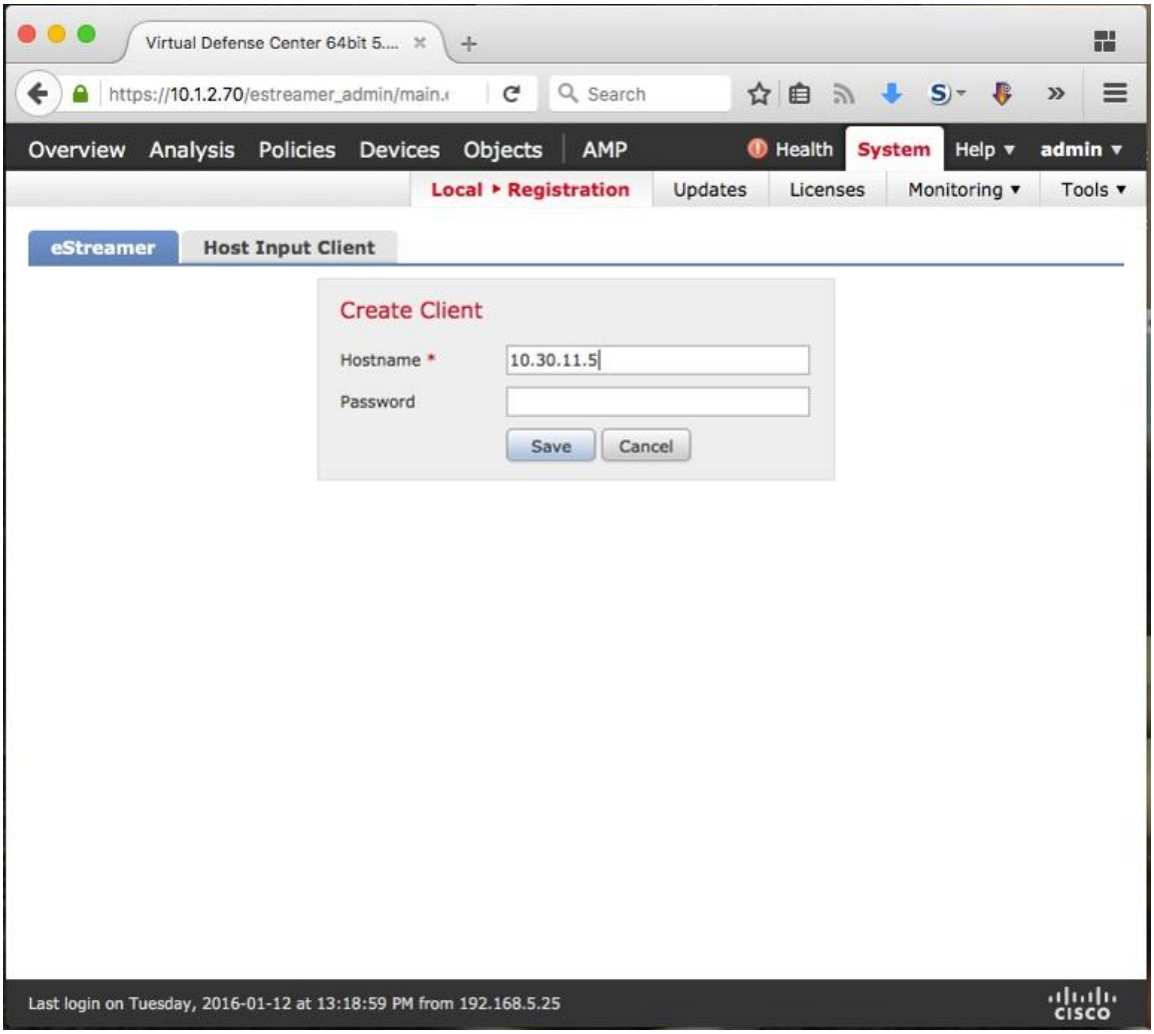

6. Find your newly configured client in the list. and click the download button to the right of your client
    a. Download and save the client certificate

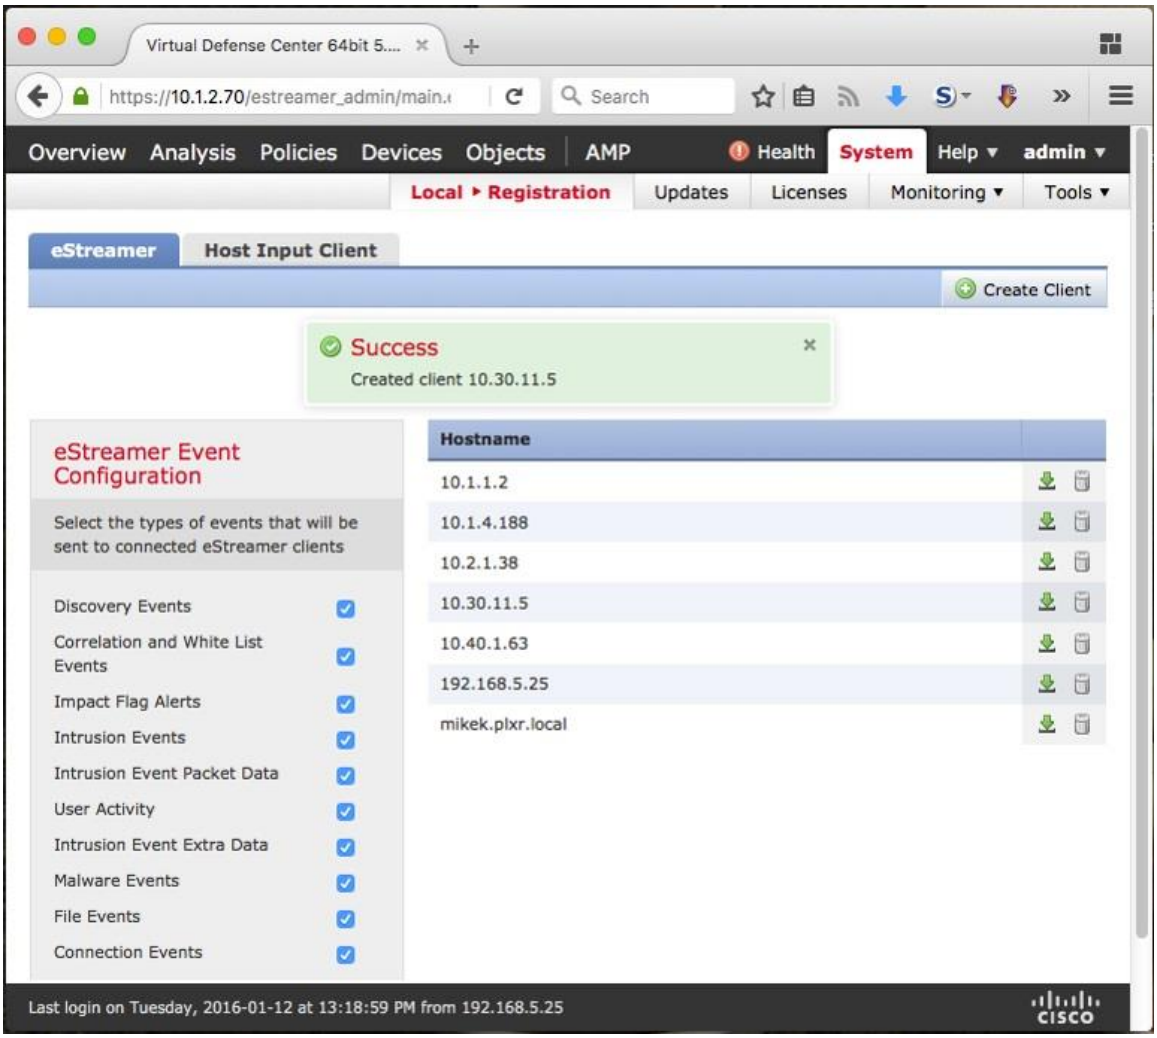

7. License Scrutinizer's eStreamer Client
    a. Upload the client certificate to /home/plixer/scrutinizer/files/ on the Scrutinizer appliance.
    b. example:

```
scp ~/Downloads/10.30.11.5.pkcs12 plixer@10.30.11.5:/home/plixer/scrutinizer/files/
```

Configure Scrutinizer's eStreamer Client

Now we'll move over to the Scrutinizer collector server and configure the client there.

```
; A name I call myself.
[collector me]
        CollectorIp=10.30.11.5
        CollectorPort=2055
        fdi_templates=/home/plixer/scrutinizer/files/fdi_templates/firesight.fdit

[firesight firesight]
        host=10.1.2.70
        port=8302
        pkcs12_file=/home/plixer/scrutinizer/files/10.30.11.5.pkcs12
        pkcs12_password=
        fs_bind_addr=10.30.11.5
        export_to=collector me
~
"/etc/firesight.ini" 13L, 337C written
```

1. Create or edit /etc/firesight.ini with contents like the example above but with details changed to reflect your network
    a. There is also a sample file installed with Scrutinizer in /home/plixer/scrutinizer/files/firesight.ini.sample that you can edit and move to the /etc/ folder.
2. Scrutinizer's eStreamer client will reconfigure itself every time a change is saved to firesight.ini. It may be better to edit a separate file and copy it into place when ready.
3. The eStreamer client will export flows to the collector at CollectorIP and CollectorPort
4. fdi_templates is the path where the export templates are defined. Use the location provided in the example.
5. The eStreamer client will connect to the FireSIGHT at the firesight host and port.
6. pkcs12_file is the location where you uploaded the FireSIGHT certificate.
7. pkcs12_password is the certificate password, or blank if a password wasn't specified.
8. fs_bind_addr is the eStreamer client address registered with FireSIGHT (Scrutinizer collector IP Address). It must be a bindable address that can route to the eStreamer service.
9. export_to tells the eStreamer client which collector or collectors will receive exported flows.

- There can be more than one "collector" and/or "firesight", but they must have different names.
- One "collector" can receive flows from multiple "firesights".
- One "firesight" can export flows to multiple "collectors".

Wait for Flows

You should start seeing flows in Scrutinizer within a minute. Contact Plixer Technical Support if flows do not begin after a minute or so.

Configuring Endace Probe Integration

Overview

Endace captures packets on the network. Searching through what can be thousands of packets or more in the packet capture (pcap) data can be very time consuming and tedious. Using Scrutinizer's flow collection with the Endace probe integration, finding the specific packet capture detail that correlates to the flow data in question is simplified many times over.

With this integration, Scrutinizer allows you to quickly filter down to certain flow data related to the issue you're experiencing, then select Endace Probes from the reporting menus to download just the packets in the pcap related to that specific flow data.

Setting up Endace Packet Capture Integration

In order to configure Scrutinizer to download the packet captures from your Endace probes, you must first enable (add) the probe via the Interactive scrut_util.exe utility on your Scrutinizer appliance.

To access these commands, open the Interactive scrut_util prompt by running:

/home/plixer/scrutinizer/bin/scrut_util.exe

Then, in the **SCRUTINIZER>** prompt, use the following commands to configure your probes.

Add a probe:

SCRUTINIZER> endace add <host_ip> <port> <endace_user> <endace_pass>

Remove a probe:

SCRUTINIZER> endace remove <host_ip>

Change/Update a probe:

SCRUTINIZER> endace update <host_ip> <port> <endace_user> <endace_pass>

Accessing Endace Probes in Scrutinizer

There are three ways to access the probes from within Scrutinizer:

- Vendor Specific Menu
- Violated Alarms

- Reports with IP Addresses

Vendor Specific Menu

From the Status tab > Vendor Specific menu, there is an Endace Probes option. This option allows you to access the Endace device without being in a report. This is especially handy when you already know the IP Addresses, protocol, and ports.

Take the following steps to access Endace packet captures via the Endace Probes option.

1. Select the Status tab
2. Select the Vendor Specific menu
3. Select the Endace Probes menu entry
4. Complete the following fields:
    - Initiator IP
    - Target IP
    - Protocol
    - Initiator and Target Port fields are optional
5. Click Search
6. Download the pcap and open in your favorite packet analyzer.

Violated Alarms

You can investigate more detail on Alarms generated by using the following instructions to access Endace Reports from Violated Alarms.

1. Select the Alarms tab
2. If looking for a specific Alarm type
    a. Select Views>>Bulletin Board by Policy
    b. Select the Policy Violated that you want to get the packet detail for
    c. Or expand the Violators list for that Policy and select the Violator that you want the packet detail for
3. If looking for a specific Violator
    a. Select Views>>Bulletin Board by Violator
    b. Select the Violator Address that you want to get the packet detail for
    c. Or expand the Policies Violated list and select the Policy that you want to the packet detail for
4. In the Bulletin Board Events page
5. Click the dropdown arrow between the Board Name and Message columns
6. Select Endace Probes
    - Any relevant details from the alarm are pre-populated. This is useful because you can get to the actual packets from conversations that triggered alarms.
7. Click Search
8. Download the pcap and open in your favorite packet analyzer.

Reports with IP Addresses

The Reports with IP Addresses option allows you to select the source or destination IP Address (or DNS Name) from a report and get information from your Endace probes.

When you want to investigate a further into a conversation you see in a flow report, this is a great option.

1. Start within a report that includes source and destination IP Addresses
2. Select an IP Address from the report
3. Select Other Options from the dropdown menu
4. Select Endace Probes from the Other Options menu
    - Any relevant details from the conversation are pre-populated. This is useful because you can get to the actual packets from the conversation.
5. Click Search
6. Download the pcap and open in your favorite packet analyzer.

Configuring a Scrutinizer Collector for Dual/Multi-homing

Security Considerations

Scrutinizer is configured with Reverse-Path Filtering disabled by default. This will allow the collector to receive flows from IP addresses it does not know how to route to. This functionality exists in order to allow collectors to receive non-local traffic that may have been forwarded by a proxy or replication appliance and is intended ONLY to be used when Scrutinizer:

1. is in a secure environment
2. has a single interface

Plixer recommends that in multi-interface/multi-homed situations and where strict networking practices are required, best practices defined in RFC 3704 be observed. This is required in order to ensure spoofed/forged packets cannot be used to generate responses out another interface.

It is recommended the following steps be taken:

1. Enable Reverse Path Forwarding –

   Edit the /etc/sysctl.conf and modify the line "net.ipv4.conf.default.rp_filter = 1" to "net.ipv4.conf.default.rp_filter = 0"
2. If you do not want to restart networking after editing this file, simply run the command "sysctl net.ipv4.conf.default.rp_filter = 0" and this should turn it on (editing the file is still required).
3. Make sure your routing tables contain routes to all networks that you will be receiving flows from. Failure to properly configure routing will result in the inability to collect flows from non-local address spaces.

Virtual Routing and Forwarding (VRF) Mode

In certain circumstances you may want to isolate your routing tables from your management network. Common cases for this are either security requirements or your management network may overlap with IP addresses on your collection side interfaces.

In these cases you should create separate routing tables in order to isolate management traffic to the management interface then collection and polling traffic will only impact their respective interfaces.

Configuration Examples

In this example, Scrutinizer has 2 interfaces: eth0 and eth1. Each will have its own routing table. One called "plixer", another called "public".

Configure separate routing tables

In this example you can see that 2 default gateways are configured. This would not be possible under a standard configuration with a single routing table. Each maintains separate isolated IP networks which are not routable to one another on the Scrutinizer box.

1. Add the 2 routing tables to the file named "/etc/iproute2/rt_tables". The 2 lines will result in the file looking like this:

#

```
# reserved values
#
255     local
254     main
253     default
0       unspec
#
# local
#
#1      inr.ruhep
1 public
2 plixer
```

2. Create file /etc/sysconfig/network-scripts/route-eth0 containing the following:

   default via 172.16.2.20 table plixer

3. Create file /etc/sysconfig/network-scripts/route-eth1 containing the following:

   default via 10.1.1.251 table public

4. Interfaces do not need any special configuration. Here is an example based on the above configuration.

```
/etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE="eth0"
BOOTPROTO="none"
HWADDR=""
NM_CONTROLLED="yes"
ONBOOT="yes"
BOOTPROTO="none"
PEERDNS=no
TYPE="Ethernet"
NETMASK=255.255.255.0
IPADDR=172.16.2.7

/etc/sysconfig/network-scripts/ifcfg-eth1
DEVICE="eth1"
BOOTPROTO="none"
HWADDR=""
NM_CONTROLLED="yes"
```

```
ONBOOT="yes"
BOOTPROTO="none"
PEERDNS=no
TYPE="Ethernet"
NETMASK=255.255.0.0
IPADDR=10.1.4.190
GATEWAY=10.1.1.251
```

5. Restart the server in order to make sure networking is reset and operating properly.

6. Also, verify that iptables is configured to accept or deny traffic you want on each interface.

Creating a Network Map

Overview

Network Maps are useful in Scrutinizer to give a quick visual of your overall Network health. They can also easily be added to dashboards for display on a big screen in your Network Operations Center or other common area, so network issues can be quickly detected.

Two types of maps are available: Plixer maps and Google maps.

Plixer Maps allow you to completely design your own Network map by custom arranging the network devices in the map, adding custom background images, and custom objects and text boxes. With Plixer maps, you can create a Network map to reflect exactly how your network is laid out by including an image of your wiring closet, overlaying the flow exporting devices and adding connections between the devices.

You can also include multiple floors, buildings or other locations in one map, and/or have multiple maps with links to maps within maps.

Google Maps can provide a geographical representation of your network. By adding physical global locations to the objects in your maps, Google maps will automatically perform a GPS lookup of longitude and latitude coordinates, then map the objects based on those coordinates.

Google maps are especially handy when you have multiple network locations within a city, state, country, or globally. This type of map not only allows you to see at a glance what network device is having issues, but where in the world it is located.

Creating a Plixer Map

1. To create a Plixer Map, start by clicking on the Maps tab.
2. Select 'Create a Group (Map)'
3. The 'Create a Group (Map)' modal will pop up in the center of the screen.
   - Note: this modal will also pop up by default when you go to the Maps tab if you have not created any maps or groups yet.
4. Enter a Group name
   - This Group name will show in the Status Device Explorer listing for Device Groups and will also be the name of the Map you are creating.
5. Click on 'Create Plixer Group'
6. The Membership selection screen is the first option to display.
   a. Membership defines which devices, objects, or other maps you want to include in your map
   b. Highlight any Non-Members you want in your map (you can use shift and CTRL keys to select multiple)
   c. Click the '<- Add' button at the bottom of the Non-Members list
   d. Clicking the 'Next' button takes us to 'Connections', but let's go to 'Objects' first.
7. Select 'Objects' from the Map menu on the left.
   - Adding new Objects provides the ability to add non-flow objects to your maps, and text objects to increase the descriptions in your maps.

   a. Click the 'New Object' button to the far right.
      - **Symbols** represent devices you want on your maps that don't display a status. They can be assigned labels and made clickable to launch other applications and/or web pages.

- **Text Boxes** can be placed on maps and generally contain text. Shapes, colors, and size can all be defined as well as adding a clickable link. Text boxes are for Plixer maps only, they cannot be placed on Google maps.

b. Enter a Label for your new object and complete the other options as needed also.

- You can also select the icon that will represent this object in the dropdown selector above the Save button.

c. Click Save to save your new object.

8. Now click on 'Connections' in the Map menu

- This is where we connect your objects and devices to each other

a. Select a From and a To device (or object) from the top portion of this modal.

b. If connecting the devices with flow data on an interface, select the Interface (between the From and To options).

   i. The interface list will default to the interfaces available on the From device, but you can check the 'Display all interfaces in this group' to expand the selection.

c. Another option is to Connect with a black line by checking that option on the right.

   i. That will not display any traffic volumes between the two objects, just a static black line.

d. Or you can Connect with a Scrutinizer Report by checking that option on the right.

   i. A 'Report' selection box will display, allowing you to select from any of your Saved Reports to use that data as utilization data on the connection.

      - This is a handy way of graphically displaying traffic volume for multiple Applications, or per user, per IP Groups, or any other groupings.
      - The Default connecting line color is Green. You can set thresholds for color change when those thresholds are met in the Color Change boxes.
      - The values are set in bits and in ascending order from Yellow to Orange to Red, Red being the highest threshold.

e. Once the connection is defined, click the Connect button and the connection will display in the lower half of this modal.

9. You can also add a Background image to your map to dress it up

a. Select Backgrounds from the Map menu

b. Select a background image from the Background dropdown list for your map

c. Or, you can add a custom background image by saving your image as .gif, jpg, swf, or .png to the ~/scrutinizer/html/images/maps/backgrounds directory. Then select that image from the Backgrounds dropdown list.

10. The last step is to arrange your objects and connectors in the map.

a. Click the Map tab

b. Select your map from the list of maps on the left (if you have multiple maps)

c. Click on the Pencil (Edit Map) icon in the vertical toolbar in the map

d. Now that you are in Edit mode, you can drag your objects around the map to arrange them appropriately.

e. To adjust the curve of the connector lines, click on the line until the 'Drag Me' textbox displays, then drag the connector line to the appropriate curve.

f. After all editing is done to the map, click the flashing red Save icon at the top of the map, then click the 'View Map' icon in the toolbar on the left.

11. Your map is complete!

- You can click on the connector lines to bring up reports of the data represented.
- You can click on the devices to get a device overview.
- You can add the map to a Scrutinizer dashboard for quicker access.

Creating a Google Map

1. From the Map tab, click on the 'Create a Group (Map)' option in the top left.
2. Enter the Group/Map name for your Google Map.
3. Click on the 'Create Google Group' button.
4. Select the devices/objects from the Non-Members list and add them to the Members list in the Membership form.
    - You can also include other Groups/Maps in your maps by selecting them from the Non-Members list
5. Select 'Objects' from the Map menu after you have selected all Members for this map.
6. Select the first object in the list by clicking on its Object Name
    a. The right side of this form is now populated with the objects properties with default Latitude and Longitude coordinates
        i. You can overwrite these coordinates with the location of this object
        ii. Or click 'Lookup coordinates by address' and an Address textbox will display for you to enter the street, city, state/province, and country where this object/device is located.
    b. Click 'Save' and the GPS Lookup will be performed and coordinates saved.
7. Repeat step 6 for each object in the map.
8. You can add connections between the objects using the instructions from Step 8 of 'Creating a Plixer Map' above.
9. Your map is complete!
    - To view it, click on the Maps tab and select your Google map from the list of maps on the left.
    - The map will auto-size to the largest area covered by the latitude/longitude coordinates. You can zoom in by clicking on the map and save that zoom level by clicking 'Save Zoom' from the options in the top right.
    - Google maps can also be added to Scrutinizer dashboards for convenience.

Creating Thresholds and Notifications

Overview

Thresholds are an advanced feature available to Scrutinizer users. When applied correctly, they can be used to receive notification of potential problems on network devices.

This guide will demonstrate how to properly set thresholds and set up notifications based on violations of the thresholds.

Setting the Global Threshold

Scrutinizer uses the SNMP poller to determine the link speed of an interface. The link speed is used in Scrutinizer to calculate interface utilization percentages.

- Link speed is commonly referred to as ifSpeed
- The link speed can also be changed manually per interface

When the interface utilization percentage reaches a specific level, an alarm is triggered to indicate high utilization. The default utilization percentage set in Scrutinizer is 90%.

Depending on your link speed(s), you may want to increase or decrease this level.

To change the Global threshold for utilization, navigate as follows:

Admin Tab>Settings>System Preferences

1. Scroll down to Threshold – Utilization
2. Edit the percentage to your requirements
3. Click Save

Applying a Notification to the Global Threshold

Once the ifSpeed is known and the global threshold set to your specifications, you can apply a notification to this threshold. This notification can be in a number of forms (email, logfile, syslog, snmptrap, script, and auto-acknowledge), and will alert you when the threshold is breached.

To add a notification to the global threshold policy, navigate to:

Admin Tab>Definitions>Policy Manager

1. Enter 'Interface Exceeded' in the search field
2. Click the Search button
3. Then click on the 'Scrutinizer: Interface Exceeded Threshold' Policy when it's displayed
4. Next, click on the Assign button, select the Notification profile you set up (see below to set up a Notification profile), then click Save

How to Create a Notification that is sent via email

Example: I want to run a default report that monitors total bandwidth on a particular interface.

When it exceeds a threshold that I will specify, I want to have it send an email to me.

Creating the Notification Profile to use

Navigate to:

Admin Tab>Definitions>Notification Manager

1. Select New and enter a profile name for this notification
2. Then click Add Alert
    a. You need to specify the email address to send the report to
3. Press Save Alert when you are done.

Adding a Threshold that sends an email notification

Now that you have a notification profile set up, you can specify a report to trigger an email alert.
Adding a threshold to a report requires that you first 'Save' the report.
Use the following steps to create a Saved Report.

1. Go to the Status tab to bring up the Top Interfaces view
2. Then click on the Interface name for the top interface in the list
3. Select Top Reports>>Applications Defined from the report menu
    • You now have the Top Applications report for the last 24 hours for your busiest interface.
4. To save this report:
    a. Enter a report name in the Report text box
    b. Then click the Save icon above the report name
5. Next, in your Saved Report, click the Add button next to Threshold.
    • The threshold will use the parameters already defined in your report (Total vs. Rate, Bits or Bytes, etc.)
6. A modal will pop up that asks you to define your threshold; again, keep in mind this is using the parameters from your report.
7. You can add a threshold for the Total amount of traffic reported, or Per Row, which allows you to look at each line in the report and match it against the threshold.
    • This is useful for applying thresholds to Users or Applications, for example.
8. After completing the fields in the modal, click the Save Threshold button and another window will open asking you to select a Notification Profile.
9. Click the dropdown list that says 'None' and select the profile that you created earlier, then click Save.

10. To create a new Notification profile, click the Manage Notifications button.
11. You can now see that the Notification Profile was added to the Threshold.

With this threshold set, any time traffic on that interface exceeds the value you set, you will get an email alert including the specific violation information.

If you run into any issues or need assistance, please contact us.

Elasticsearch / Kibana (ELK) Integration

Overview

What does ELK integration with Scrutinizer do for you?

- From within Scrutinizer, you can launch searches for IP Addresses in Kibana.
- From Kibana, you can view Scrutinizer Vitals and FA TopN gadget details.

What is ELK?

Elasticsearch - A searching service to look through the stored data collected by Logstash.

Logstash - A means to collect logs and events (like syslogs) and filter them in a specific way to be stored for later analysis.

Kibana - Front-end to present data by creating dashboards and visualizations, similar to Scrutinizer's dashboards and gadgets.

Integration Prerequisites

Kibana 4.2

Scrutinizer v15.12+

NOTE: The following configuration instructions apply to Scrutinizer v16.7 and later. If you have an earlier version of Scrutinizer, it is recommended you contact Plixer directly for assistance.

Kibana Searches from Scrutinizer Reports and Alarms

The following steps will walk you through how to search Kibana's database from Scrutinizer reports.

1. Select a Scrutinizer report that includes IP addresses
    a. Select a host of interest and click on it
    b. Select 'Other Options' from the Reports menu
    c. From the 'Other Options' menu, select 'Kibana (ELK)'
2. The IP address and report timeframe are passed to Kibana's search engine for detailed Kibana reporting.

For more detail (from Kibana) on Scrutinizer Alarms, follow these steps.

1. Go to the Alarms tab in Scrutinizer and select either:
    a. Bulletin Board by Violator
        i. Select the violator that you want more detail on.
    b. Bulletin Board by Policy
        i. Select the policy that you want more detail on.
2. In the Bulletin Board Events view that opens, click on the dropdown arrow to the left of the Message column for the alarm that you have selected.
3. Select 'Kibana (ELK)' from the Available Options menu and the Violator's IP address and timeframe of the violation are passed to Kibana.
    - If the Violator's IP address and an alarm time (not timeframe) are being passed, then 30 minutes before and after the alarm time is searched.

Scrutinizer Reporting from within Kibana

With the Kibana (ELK) integration with Scrutinizer, you can set up dashboards that include:

- Scrutinizer Vitals information
- Flow Analytics TopN Algorithms

The Scrutinizer Vitals dashboard in Kibana can include:

- CPU
- Memory
- Disk Usage
- Flows per collector
- Status per collector

Dashboards created with the TopN Algorithm gadgets from Flow Analytics include:

- Top Applications
- Top Countries
- Top Rev 2nd lvl Domains (Top reverse 2nd level domains)
- Top Flows
- Top Hosts
- Top Jitter
- Top Networks

How to Configure ELK Integration with Scrutinizer

There are two components to the ELK Integration.

1. Preparing Scrutinizer
2. Importing the necessary files into Kibana's UI

Preparing Scrutinizer

NOTE: Flow Analytics must be enabled and collecting statistics for the Top X Algorithms.

1. Log on to the Scrutinizer server with administrative permissions
2. Open the Interactive scrut_util.exe prompt with the following command:

   /home/plixer/scrutinizer/bin/scrut_util.exe
3. And run:

   SCRUTINIZER> enable elk http://<ip:port>

   - <ip:port> is the ELK server's IP address and port

After a few moments, Scrutinizer will begin to send events to ELK.

- To test the data export, from within the scrut_util shell, run:

  collect elk <elk_ip>

- To disable the data export, run:

  disable elk http://<ip:port>

Preparing Kibana

Integrating ELK with Scrutinizer displays details in Kibana that have been collected and processed by Scrutinizer. For more information, visit Plixer's Elasticsearch / Kibana Integration page.

1. After enabling the ELK integration on Scrutinizer, refresh the index on Logstash for Scrutinizer's fields to show up.
   a. In Kibana, go to Indices > Logstash
      i. Click on the Reload field list icon at the center top of the screen.

2. Download the Kibana Integration Plugin from the Elasticsearch / Kibana Integration page.
    a. Extract the files from the scrutinizer-elk.zip file.
3. In Kibana, go to Settings > Objects > Import > Visualizations and navigate to the *elk-scrutinizer-visualizations.json* file extracted in Step 2a above and click Open.
4. Then go to Settings > Objects > Import > Dashboards and navigate to the *elk-scrutinizer-dashboards.json* file extracted in Step 2a above and click Open.
5. The Kibana dashboards and visualizations are now all imported and events have been configured to be coming from Scrutinizer also.
6. From the Visualize tab, scroll to the bottom and you can filter for a specific visualization. Typing Scrutinizer in the filter will show all of the Scrutinizer visualizations.
7. From the Dashboard tab, you can navigate around the various Scrutinizer dashboards imported.

Scrutinizer for Splunk Application

Overview

What does the Scrutinizer for Splunk Application do for you?

- From within Scrutinizer, you can launch searches for IP Addresses in Splunk.
- From Splunk, you can view Scrutinizer Vitals and FA TopN gadget details.

Splunk Searches from Scrutinizer Reports and Alarms

The following steps will walk you through how to search Splunk's database from Scrutinizer reports.

1. Select a Scrutinizer report that includes IP addresses
    a. Select a host of interest and click on it
    b. Select 'Other Options' from the Reports menu
    c. From the 'Other Options' menu, select 'Search Splunk'
2. The IP address and report timeframe are passed to Splunk's search engine for detailed Splunk reporting.

For further details (from Splunk) on Scrutinizer Alarms, follow these steps.

1. Go to the Alarms tab in Scrutinizer and select either:
    a. Bulletin Board by Violator
        i. Select the violator that you want more detail on.
    b. Bulletin Board by Policy
        i. Select the policy that you want more detail on.
2. In the Bulletin Board Events view that opens, click on the dropdown arrow to the left of the Message column for the alarm that you have selected.
3. Select 'Search Splunk' from the Available Options menu and the Violator's IP address and timeframe of the violation are passed to Splunk.
    - If the Violator's IP address and an alarm time (not timeframe) are being passed, then 30 minutes before and after the alarm time is searched.

Scrutinizer Reporting from within Splunk

With the Scrutinizer for Splunk Application, you can set up dashboards that include:

- Scrutinizer Vitals information
- Flow Analytics TopN Algorithms

The Default Dashboard view for the Scrutinizer for Splunk Application is for the Vitals information and includes:

- CPU
- Memory
- Disk Usage
- Flows per collector
- Status per collector

Other Splunk menus can include links to the Scrutinizer tabs:

- Dashboards
- Maps
- Status and Reports
- Alarms
- Admin and Settings

Additionally, you can create panels based on the TopN Algorithm gadgets from Flow Analytics:

- Top Applications
- Top Countries
- Top Rev 2nd lvl Domains (Top reverse 2nd level domains)
- Top Flows
- Top Hosts
- Top Jitter
- Top Networks
- And more

* Clicking on any entities in the graphs or detail in a table report will run a Splunk search for that detail and time range. As mentioned earlier, those searches can also be initiated directly from Scrutinizer reports or alarms.

You can also have a menu consisting of useful links from our websites:

- plixer.com
- Configure NetFlow

- Internet Threat Center
- Latest Blogs
- NetFlow Knights

How to Configure Splunk Integration with Scrutinizer

There are two components to the Splunk Integration.

1. Preparing Scrutinizer
2. Installing the Scrutinizer for Splunk Application on Splunk

Preparing Scrutinizer

NOTE: Flow Analytics must be enabled and collecting statistics for the Top X Algorithms.

***** The following configuration instructions apply to Scrutinizer v16.7 and later. If you have an earlier version of Scrutinizer, it is recommended you contact Plixer directly for assistance. *****

1. Log on to the Scrutinizer server with administrative permissions
2. To enable the Scrutinizer/Splunk integration, from command line:
   a. Run:

   /home/plixer/scrutinizer/bin/scrut_util.exe

   b. At the SCRUTINIZER> prompt, run:

   SCRUTINIZER> enable splunk http://<ip:port> <syslog port>

   <ip:port> is Splunk IP address and port

   <syslog port> is port used to send syslogs to Splunk

After a few moments, Scrutinizer will begin to export data to Splunk.

- To manually collect data from Scrutinizer and send to Splunk, run:

  SCRUTINIZER> collect splunk <splunk_ip>

- To disable the data export, run:

  SCRUTINIZER> disable splunk http://<ip:port>

Installing the Scrutinizer for Splunk Application on Splunk

The Scrutinizer for Splunk App displays details that have been collected and processed by Scrutinizer. For more information, visit Plixer's Splunk Integration page.

1. Download the Splunk Plugin using the link at the bottom of the Splunk Integration page.

2. Log into Splunk
3. Click "Apps > Manage Apps"
4. Click "Install app from file"
5. Click "Choose File" and locate the PlixerScrutinizerForSplunk.spl file within the Splunk Plugin file downloaded in Step 1.
    a. If you are upgrading the plugin, click the "Upgrade App" checkbox below the "Choose File" button.
6. Click the "Upload" button.
7. Follow the onscreen instructions and restart Splunk.
8. Navigate to the "Apps > Manage Apps" menu.
9. Locate the "Scrutinizer for Splunk" app in the list below and click the "View Objects" link associated to the application.
10. Locate the "Default" link and click it.
11. Replace the text "http://ADD_SCRUTINIZER_ADDR_HERE" with the link to your Scrutinizer server. (e.g. https://10.1.1.12:88)
12. Click the "Save" button and then access the Scrutinizer application in the Apps menu.

After a few minutes data will begin to show up in Splunk and the graphs will begin to fill in.

Using the Scrutinizer Reporting API

Overview

The Reporting API (Application Program Interface) is a simple way to access the report data via HTTP or command line. The API only accepts requests from hosts that are part of the authorized IPs whitelist. See the Getting Started > Authentication section below for more details.

Getting Started

Authentication

The Reporting API uses a whitelist of IP Addresses for authentication. This whitelist is defined in $homeDir/files/auth_ips.cfg (one host IP per line). You will need to add the IP of the host you are going to access the API from to this list.

The response for an unauthenticated request is:

```
{"msg":"Request cannot be authenticated.","err":"oauthFailure"}
```

Requesting a Report via HTTP

To request a report via HTTP using the API you need to make a POST or GET request to this base URI: "http://[SCRUTINIZER SERVER]/fcgi/scrut_fcgi.fcgi", passing the following four fields:

- rm=report_api
- action=get
- rpt_json={}
- data_requested={}

rm=report_api is always the same when accessing the Report API.

action tells the API what we want to do. Currently, the only action supported is 'get'.

rpt_json takes the JSON string containing the details for the report (See Reference Guide > rpt_json for more information).

data_requested is used to indicate what data from the report we need. It also takes a JSON string.

These four fields are mandatory.

Requesting a Report via the Command Line (Linux or Windows)

Using the Report API in the command line doesn't require authentication since you need console access to the Scrutinizer server. To request a report:

1. Navigate to $homeDir/cgi-bin/
2. Run 'reporting.cgi -getReport -rpt_json=[JSON STRING WITH REPORT DETAILS] -data_requested=[JSON STRING WITH DATA REQUESTED DETAILS]'

Examples

HTTP

https://SCRUTINIZER_SERVER/fcgi/scrut_fcgi.fcgi?rm=report_api&action=get&rpt_json={"reportTypeLang":"flowview","reportDirections":{"selected":"inbound"},"times":{"dateRange":"LastFiveMinutes","start":"","end":""},"filters":{"sdfDips_0":"in_0A0101FB_ALL"}}&data_requested={"inbound":{"table":{"query_limit"":{"offset":0,"max_num_rows":11000}}}}

Javascript

This example illustrates how to make an AJAX request to the reporting API using Javascript & jQuery.

```
var url = "http://[SCRUTINIZER SERVER]/fcgi/scrut_fcgi.fcgi";
var report_details = {
      "reportTypeLang" : "conversations",
      "reportDirections" : {
      "selected" : "both"
      },
         "times" : {
                  "start" : 1426001640,
                  "end" : 1426088100
         },
         "filters" : {
                  "sdfDips_0" : "in_0A040101_0A040101-0"
         },
         "dataGranularity" : {
                  "selected" : "auto"
         }
};
var data_i_need = {
         "inbound" : {
                  "graph" : "all",
                  "table" : {
                           "query_limit" : {
                                    "offset" : 0,
                                    "max_num_rows" : 1000
                           }
                  }
         },
         "outbound" : {
                  "table" : {
                           "query_limit" : {
                                    "offset" : 0,
```

```
                                "max_num_rows" : 1000
                        }
                }
        }
};
$.ajax({
        type: "POST",
        url: url,
        data: {
                rm : "report_api",
                action : "get",
                rpt_json : JSON.stringify( report_details ),
                data_requested : JSON.stringify( data_i_need )
        },
        success: function( response ) {
                if ( response.err ) {
                        alert(response.msg);
                } else {
                        console.log(response);
                }
                return;
                }
});
```

Perl

There are two example scripts that are included in Scrutinizer in the /scrutinizer/files/api_examples/ directory.

- api_report_example.pl
- api_host_example.pl

api_report_example.pl

This script will run a report and return the results. It consists three subroutines:

- host_search_request: Builds the request JSON
- get_scrutinizer_data: Makes the request to Scrutinizer
- format_report_data: Formats the returned data

api_host_example.pl

This script uses our host index API to do a quick search. This will determine if a host has been seen anywhere on your network. It consists of two subroutines:

- get_scrutinizer_data: Makes a host search request to Scrutinizer
- format_search_data: Formats the returned data

Python

- Report Type: flowView
- Time Frame: LastFiveMinutes
- Exporter/Interface: IP is 10.1.1.251 which converts to 0A0101FB.

  in_0A0101FB_ALL means that we are going to include all interfaces for the exporter 10.1.1.251
- Max Num Rows: 11000

This python script makes the same HTTP request but saves the json to data.txt.

```
import urllib
import urllib2
import json
url = 'http://SCRUTINIZER_SERVER/fcgi/scrut_fcgi.fcgi'
report_details = {
	'reportTypeLang' : 'flowView',
	'reportDirections' : {
		'selected' : 'inbound'
	},
	'times' : {
		'dateRange' : 'LastFiveMinutes'
	},
	'filters' : {
		'sdfDips_0' : 'in_0A0101FB_ALL',
	},
	'dataGranularity' : {
		'selected' : '1m'
	}
}
data_i_need = {
	'inbound' : {
```

```
            'table' : {
                    'query_limit' : {
                            'offset' : 0,
                            'max_num_rows' : 11000
                    }
            }
        },
 }
data = {
        'rm' : 'report_api',
        'action' : 'get',
        'rpt_json' : json.dumps( report_details ),
        'data_requested' : json.dumps( data_i_need )
}
data = urllib.urlencode( data )
req = urllib2.Request( url, data )
response = urllib2.urlopen( req )
report = response.read()
report_obj = json.loads( report )
with open('./data.txt', 'w') as outfile:
  json.dump(report_obj, outfile)
```

Reference Guide

rm=report_api (required HTTP/ static)

This URL field is required when making an HTTP request. The value for this field will always be report_api

action (required HTTP)

Indicates what we want to do. Right now the only action supported is 'get' to get a report.

getReport (required CMD)

This field is required when using the CMD to access the API.

rpt_json

reportTypeLang (required)

Is the lang id of the report type we want to use for the report.

You can see what reports are available in your install by looking at Scrutinizer > Status > Views (sidebar)> Available Reports

reportDirection (required)

It takes an object with a single attribute 'selected'. The values for 'selected' are:

- inbound
- outbound
- both

times (required)

dateRange (required if no start and end)

Indicates the time for which we want to run the report using one of the predefined date ranges. Available dataRange values:

- LastFiveMinutes
- LastTenMinutes
- LastFifteenMinutes
- LastTwentyMinutes
- LastThirtyMinutes
- LastFortyfiveMinutes
- LastHour
- LastFullHour
- LastThreeDays
- LastSevenDays
- LastThirtyDays
- Today
- Yesterday
- Last24Hours
- ThisWeek
- LastWeek
- ThisMonth
- LastMonth
- ThisYear
- LastYear

start

Epoch time that we want to use for the start time of the report.

end

Epoch time that we want to use for the end time of the report.

clientTimezone

If we want to base the report times on the time zone where the client is located, we need pass the time zone using clientTimezone, otherwise the server time zone will be used.

It takes the time zone names as they are defined in the IANA time zone database, such as "Asia/Shanghai", "Asia/Kolkata", "America/New_York".

filters (required)

It takes an object that defines the different filters that can be applied to the report. We have many filter types that can be applied. Explaining each filter type and how filters work is too lengthy for this document, but here are a couple of notes to get you going:

- There has to be at least an sdfDips filter. sdfDips filters are used to indicate which exporter/interface we want to use for the report. sdfDips are defined like this: "in_[IP_OF_EXPORTER_IN_HEX]_[INTERFACE_ID]"
- filters are defined in the filters object like this:

```
{
"[filter_type]_[index]" = "filter_definition",
}
```

- index is used to differentiate filters of the same type and to indicate in which order they should be applied.

dataGranularity

It indicates which data interval (in minutes) we want to use. It takes an object with one key, 'selected'. Available values for 'selected' are:

- auto (lets the API decide which interval to use)
- 1
- 5
- 30
- 120
- 720
- 1440
- 10080

rateTotal

Indicates if we want rates or totals. Totals or rates are not available for all reports. Default depends on the report type.

```
{
"selected" : "rate" //"total
}
```

byInt

Indicates whether we want to include the interface details in the report table. Values are 0 or 1. Only applicable to reports with interface information available. Defaults to 0.

```
{
"selected" : 1
}
```

data_requested

It is used to indicate to the API what part of the report we need. Reports return two data sets, one for rendering the graph and the other for rendering the table. Each one is divided into outbound and inbound.

Return Data

Example

```
{
  "report": {
    "table": {
      "inbound": {
        "totalRowCount": "50",
        "columns": [
          {
            "title": "Rank",
            "label": "",
            "klassth": "rankWidth"
          },
          {
            "colHasBbp": 0,
            "defaultSort": null,
            "klassth": "",
            "klassLabel": "",
            "format": {},
            "elementName": "sourceipaddress",
            "canBeRate": null,
            "label": "Source",
            "title": "sourceipaddress"
          },
          {
            "colHasBbp": 0,
            "defaultSort": null,
            "klassth": "commPortWidth",
            "klassLabel": "",
            "format": {},
            "elementName": "commonport",
            "canBeRate": null,
            "label": "Well Known",
            "title": "commonport"
          },
          {
            "colHasBbp": 0,
            "defaultSort": null,
            "klassth": "",
            "klassLabel": "",
```

```
            "format": {},
            "elementName": "destinationipaddress",
            "canBeRate": null,
            "label": "Destination",
            "title": "destinationipaddress"
          },
          {
            "colHasBbp": 0,
            "defaultSort": null,
            "klassth": "alignRight dataWidth sortable",
            "klassLabel": "",
            "format": {
              "niceLtr": "p",
              "perSecMode": 0,
              "bitOverRide": 0
          },
          "elementName": "sum_packetdeltacount",
          "canBeRate": 1,
          "label": "Packets",
          "title": "sum_packetdeltacount"
        },
        {
          "colHasBbp": 0,
          "defaultSort": null,
          "klassth": "alignRight dataWidth",
          "klassLabel": "",
          "format": {},
          "elementName": "percenttotal",
          "canBeRate": null,
          "label": "Percent",
          "title": "percenttotal"
        },
        {
          "colHasBbp": 1,
          "defaultSort": "DESC",
          "klassth": "alignRight dataWidth sortable hasBbp ",
          "klassLabel": " sortdesc ",
          "format": {
            "niceLtr": "b",
```

```
            "perSecMode": 0,
            "bitOverRide": 0
          },
          "elementName": "sum_octetdeltacount",
          "canBeRate": 1,
          "label": "sum_octetdeltacount",
          "title": "sum_octetdeltacount"
        }
      ],
      "rows": [
        [
          {
            "klasstd": "rank1",
            "title": "Rank",
            "label": "1"
          },
          {
            "rawValue": "10.1.4.105",
            "dataJson": "{\"column\":\"sourceipaddress\"}",
            "klasstd": "alignLeft",
            "klassLabel": "ipDns",
            "title": "10.1.4.105",
            "label": "10.1.4.105"
          },
          {
            "rawValue": "av-emb-config (2050 - TCP)",
            "dataJson": "{\"column\":\"commonport\"}",
            "klasstd": "alignLeft",
            "klassLabel": "",
            "title": "()",
            "label": "av-emb-config (2050 - TCP)"
          },
          {
            "rawValue": "10.1.11.51",
            "dataJson": "{\"column\":\"destinationipaddress\"}",
            "klasstd": "alignLeft",
            "klassLabel": "ipDns", "title": "10.1.11.51",
            "label": "10.1.11.51"
          },
```

```
            {
              "rawValue": "20132651",
              "dataJson": "{\"column\":\"sum_packetdeltacount\"}",
              "klasstd": "alignRight",
              "klassLabel": "",
              "title": "20132651",
              "label": "20.13 Mp"
            },
            {
              "rawValue": null,
              "dataJson": "{\"column\":\"percenttotal\"}",
              "klasstd": "alignRight",
              "klassLabel": "",
              "title": "19.53 %",
              "label": "19.53 %"
            },
            {
              "rawValue": "26274904910",
              "dataJson": "{\"column\":\"sum_octetdeltacount\"}",
              "klasstd": "alignRight",
              "klassLabel": "",
              "title": "26274904910",
              "label": "26.27 Gb"
            }
          ]
        ],
        "footer": [
          [
            {
              "title": "Other",
              "label": "Other"
            },
            {
              "klasstd": "alignLeft",
              "translations": "",
              "elementName": "sourceipaddress",
              "title": "Group by columns are not summarized",
              "label": " - ",
              "menuType": "",
```

```
  "filterType": ""
},
{
  "klasstd": "alignLeft",
  "translations": "",
  "elementName": "commonport",
  "title": "Group by columns are not summarized",
  "label": " - ",
  "menuType": "",
  "filterType": ""
},
{
  "klasstd": "alignLeft",
  "translations": "",
  "elementName": "destinationipaddress",
  "title": "Group by columns are not summarized",
  "label": " - ",
  "menuType": "",
  "filterType": ""
},
{
  "klasstd": "alignRight",
  "translations": "",
  "elementName": "sum_packetdeltacount",
  "title": 83442022,
  "label": "83.44 Mp",
  "menuType": "",
  "filterType": ""
},
{
  "klasstd": "alignRight",
  "translations": "",
  "elementName": "percenttotal",
  "title": "Group by columns are not summarized",
  "label": " - ",
  "menuType": "",
  "filterType": ""
},
{
```

```
      "klasstd": "alignRight",
      "translations": "",
      "elementName": "sum_octetdeltacount",
      "title": 108236252030,
      "label": "108.24 Gb",
      "menuType": "",
      "filterType": ""
    }
  ],
  [
    {
      "title": "(from conv tables)",
      "label": "Total*"
    },
    {
      "klasstd": "alignLeft",
      "translations": "",
      "elementName": "sourceipaddress",
      "title": "Group by columns are not summarized",
      "label": " - ",
      "menuType": "",
      "filterType": ""
    },
    {
      "klasstd": "alignLeft",
      "translations": "",
      "elementName": "commonport",
      "title": "Group by columns are not summarized",
      "label": " - ",
      "menuType": "",
      "filterType": ""
    },
    {
      "klasstd": "alignLeft",
      "translations": "",
      "elementName": "destinationipaddress",
      "title": "Group by columns are not summarized",
      "label": " - ",
      "menuType": "",
```

```
            "filterType": ""
          },
          {
            "klasstd": "alignRight",
            "translations": "",
            "elementName": "sum_packetdeltacount",
            "title": "103574673",
            "label": "103.57 Mp",
            "menuType": "",
            "filterType": ""
          },
          {
            "klasstd": "alignRight",
            "translations": "",
            "elementName": "percenttotal",
            "title": "Group by columns are not summarized",
            "label": " - ",
            "menuType": "",
            "filterType": ""
          },
          {
            "klasstd": "alignRight",
            "translations": "",
            "elementName": "sum_octetdeltacount",
            "title": "134511156940",
            "label": "134.51 Gb",
            "menuType": "",
            "filterType": ""
          }
        ]
      ]
    }
  }
 }
}
```


report (Object)

Is the object that contains the data returned by the API. The data is divided into "table" and "graph".

table (Object)

The table data is divided into inbound or outbound. For each direction you get:

- totalRowCount: integer representing the total number of rows available.
- columns: See columns section below.
- rows: See rows section below.
- footer: See the footer section below.

columns (Collection)

It is an array of objects that represent the definition of each column in the table. Most of the data is used to create the html table in the browser. Here is a list of the most relevant information included in each object:

- elementName: Name of the element that the data in the column is for.
- format: Details about how the data in the column needs to be formatted.
- label: Label used in the table header.

rows (Array)

Array of collection. Each collection in the array represents a row and each object in the collection represents a table cell for that row. Here are some details about the keys for each object representing a cell:

- rawValue: The unformatted value as it is return from the database
- label: formatted value

footer (Array)

The footer[0] represents the "Others" data for the operations columns (columns which value is the product of an arithmetic operation). That is, the amount for the data not included in the rows.

The footer[1] represents the "total". Total for a column is equal to the sum of the rows for that column plus the "Others" value for that column.

Using the Scrutinizer User API

Overview

The User API is a simple way to add, remove, and edit user and usergroup accounts via HTTP or the command line. The API only accepts requests from hosts that are part of the authorized IPs whitelist.

Getting Started

Authentication

First, you need to get the IP of the host from where you are going to access the API in the list of IP that can access it. The User API uses a whitelist of IPs for authentication. The whitelist is defined in $homeDir/files/auth_ips.cfg (one host IP per line).

The response for unauthenticated request is: {"msg":"Request cannot be authenticated.","err":"oauthFailure"}

Accessing the User API via HTTP

To access the API via HTTP you need to make a POST or GET request to this base URI: "http://[SCRUTINIZER SERVER]/fcgi/scrut_fcgi.fcgi" passing the following fields:

- *rm=user_api*
- *action=*
- *json={}*

rm=user_api is always the same every time we want to access the user API. *action* tells the API what we want to do. *json* will contain a JSON object with the details of the action that we want the User API to perform.

These three fields are mandatory.

Examples

Javascript

This example illustrates how to make an AJAX request to the User API using Javascript & jQuery

```
 var url = "http://[SCRUTINIZER
SERVER]/fcgi/scrut_fcgi.fcgi";
 var action_details = {
    "users":[
       {
          "name":"MyAdmin",
          "pass":"MyPass",
          "membership":[1]
       },
```

```
        {
            "name":"MyGuest",
            "pass":"OtherPass",
            "membership":[2]
        }
    ]
};

$.ajax({
    type: "POST",
    url: url,
    data: {
        rm : "user_api",
        action : "createUser",
        json : JSON.stringify(
action_details ),
    },
    success: function( response ) {
        if ( response.err ) {
            alert(response.msg);
        } else {
            console.log(response);
        }
        return;
    }
});
```

Python

```
import urllib
import urllib2
import json
url = 'http://[SCRUTINIZER
SERVER]/fcgi/scrut_fcgi.fcgi'
user_details = {
    'users'  :  [
        {
            'name' : 'MyAdmin',
            'pass' : 'MyPass',
```

```
            'membership' : [1]
        },
        {
            'name' : 'MyGuest',
            'pass' : 'OtherPass',
            'membership' : [2]
        }
    ]
}

data = {
    'rm' : 'user_api',
    'action' : 'createUser',
    'json' : json.dumps(
user_details )
}
data = urllib.urlencode( data )
req = urllib2.Request( url, data )
response = urllib2.urlopen( req )
user_data = response.read()
user_obj = json.loads( user_data )
```

Reference Guide

rm=user_api

This URL field is required when making an HTTP request. This field is always equal to user_api.

action

Indicates what action we want the User API to perform with the JSON we provide. Below is a list of available actions.

createUser

Allows the creation of user accounts within Scrutinizer. They can be added to user groups at the same time.

- **JSON Object Expected**

```
{"users":[
    {
        "name":"MyAdmin",
        "pass":"MyPass",
        "membership":[1]
```

```
    },
    {
        "name":"MyGuest",
        "pass":"OtherPass",
        "membership":[2]
    }
]}
```


* users: An array of all users to be created, so multiple users can be created at once. If only one user is needed, this will be an array of one.

* name: The name of the user to be created.

* pass: The password for the new user account.

* membership: An array of usergroup_ids from the plixer.usergroups table. The user will be added to these groups when the account is created.

- **JSON Object Returned**


```
{"data":[
    {
        "id":3,
        "name":"MyAdmin"
    },{
        "id":4,
        "name":"MyGuest"
    }
]}
```


* data: An array of responses for each user account that Scrutinizer attempted to create.

* id: The new user_id of the user account that was created.

* name: The name of the account created (by design this is identical to the name passed in).

delUser

Allows the deletion of user accounts within Scrutinizer.

- **JSON Object Expected**

```
{"delUsers":[
    2,
    "GuestUser",
    7
]}
```

* delUsers must contain an array of user IDs or usernames of accounts to be deleted.

- **JSON Object Returned**

```
{"data":[
    "Deleting user 2",
    "Deleting user GuestUser",
    "Failed to delete 7 (No such account exists)"
]}
```

* data: An array of responses for each user account that Scrutinizer attempted to delete.

createUsergroup

Allows the creation of usergroups within Scrutinizer. Members can be added at the same time.

- **JSON Object Expected**

```
{"usergroups":[
    {
        "name":"GroupA",
        "template_usergroup":1,
        "users":[1,2]
    },
```

```
        {
            "name":"GroupB",
            "template_usergroup":2,
            "users":["MyUser","MyUser2"]
        }
]}
```


* name: The name of the usergroup to be created.

* users: An array of all users to be added, by user ID or username. If only one user is needed, this will be an array of one. An empty array will create an empty usergroup.

- **JSON Object Returned**


```
{"data":[
{
    "id":5,
    "name":"GroupA",
    "members":[1,2]
},
{
    "name":"GroupB",
    "error":"A usergroup
already exists with that name"
}
]}
```


* data: An array of responses for each usergroup that Scrutinizer attempted to create.

* id: The new usergroups_id of the usergroup that was created.

* name: The name of the usergroup created (by design this is identical to the name passed in).

* members: An array of user IDs or usernames for the members successfully added to the group

* error: Any errors encountered during the creation of a particular usergroup

delUsergroup

Allows the deletion of usergroups within Scrutinizer.

- **JSON Object Expected**

```
{"delUsergroups":
  [
    3,
    "My Usergroup"
  ]
}
```

* delUsergroups must contain an array of usergroup IDs or usergroup names to be deleted.

- **JSON Object Returned**

```
{"data":
  [
    "Usergroup 3 was deleted",
    "Failed to delete My Usergroup (No such group exists)"
  ]
}
```

* data: An array of responses for each usergroup that Scrutinizer attempted to delete.

prefs

Changes user preferences for individual users

- **JSON Object Expected**

{"user_id":5,
"prefs":[
{
"pref":"statusTopn",
"setting":10
},
{
"pref":"language",
"setting":"english"
}
]}

* user_id is required for the user with preferences to change

* prefs contains an array of preferences and new settings

- **JSON Object Returned**

{"data":
"updated":[
"statusTopn updated to 10 for user_id 5",
"language updated to english for user_id 5"
],
"errors":[
"Failed to update user_id 5's preference favoriteIceCream: No such preference exists"
]
}

* data: An array of responses for each preference change updated or attempted

* updated: Messages for any preference successfully changed

* errors: Any errors encountered while changing preferences

changeUsergroupNames

Changes user preferences for individual users

- **JSON Object Expected**

```
{"changeUsergroupNames":
  [
    {
      "id":3,
      "newname":"My Best
Group"
    },
    {
      "oldname":"Sad Group",
      "newname":"Happy
Group"
    },

  ]
}
```

* changeUserGroupNames is an array of changes to be made to usergroups

* id is the usergroups_id of a group you wish to change

* oldname is an alternative to the "id" parameter, and contains the current usergroup name

* newname is the name this usergroup will be changed to have

- **JSON Object Returned**

```
{"data":
  "updated":[
    "Usergroup 3 renamed to My Best Group"
  ],
  "errors":[
    "Failed to rename Sad Group to Happy Group:
Another usergroup already exists with this name."
```

```
    ]
  }
```

* data: An array of responses for each name change attempted

* updated: Messages for any name successfully changed

* errors: Any errors encountered while changing names

changeUsername

Changes usernames

- **JSON Object Expected**

```
  {"changeUsername":
    {
      "<user_id|oldname>":"<N|"myUser">,
      "newname":"OpSCT"
    }
  }
```

* user_id is the user ID of the account to be changed

* oldname is an alternative to user ID, and contains the current name of the user

* newname contains the name to which you wish to change this user

- **JSON Object Returned**

```
  {"data":
    {
      "message": "User myUser successfully renamed to OpSCT"
    }
  }
```

* message will contain either a statement of success or an error explaining why the name change failed

membership

Changes usergroup membership

- **JSON Object Expected**

```
{"membership":
  {
    "add":[
        {
          "user_id":"N",
          "user_name":"USER1",
          "usergroup_id":"X",
          "usergroup_name":"USERGROUP1"
        },
        {
          "user_id":"M",
          "user_name":"USER2",
          "usergroup_id":"Y",
          "usergroup_name":"USERGROUP2"
        }
    ],
    "remove":[
        {
          "user_id":"L",
          "user_name":"USER3",
          "usergroup_id":"Z",
          "usergroup_name":"USERGROUP3"
        }
    ]
  }
}
```

* membership contains two arrays, "add" and "remove," which

have information on each membership change

* user_id is the ID from plixer.users you wish to add or remove

* user_name is an alternative to user_id, and can be the plain text name of the user

* usergroup_id is the ID from plixer.usergroups that the user will be added to or removed from

* usergroup_name is an alternative to usergroup_id, and can be the plain text name of the usergroup

- **JSON Object Returned**

```
{"data":
  "added":[
    "User N added to usergroup X",
    "User M added to usergroup Y",
  ],
  "removed":[
    "User L removed from usergroup Z"
  ]
}
```

* added will contain an array of either statements of success or errors explaining why the membership change failed

* removed will contain an array of either statements of success or errors explaining why the membership change failed

permissions

Changes usergroup permissions

- **JSON Object Expected**

```
{"permissions":
  {
```

```
        "add":[
            {
                "usergroup_id":X,
                "usergroup_name","USERGROUP1",
                "type":<usergroups_permissions.fk>,
                "secCode":?
            }
        ],
        "remove":[
            {
                "usergroup_id":X,
                "usergroup_name","USERGROUP3",
                "type":<usergroups_permissions.fk>,
                "secCode":?
            }
        ]
      }
    }
```

* permissions will contain two arrays, "add" and "remove," for each permission to be changed

* usergroup_id is the ID from plixer.usergroups that the user will be added to or removed from

* usergroup_name is an alternative to usergroup_id, and can be the plain text name of the usergroup

* type can be found in plixer.usergroups_permissions in the 'fk' field. It differentiates the types of permissions (e.g. "int," "dev," etc.)

* secCode another field in plixer.usergroups_permissions and contains the permission. It will be different depending on the "type" (e.g. "1" for interface 1, "0A010107" for a device, etc.)

- **JSON Object Returned**

```
    {"data":
      "updated":[
          "Added permission int-1 to usergroup USERGROUP1",
      ],
      "errors":[
```

```
        "Failed to remove permission dev-0A010107 from usergroup
USERGROUP3: Permission not currently assigned to this usergroup"
      ]
   }
```

* updated will contain an array of either statements of success or errors explaining why the permission change failed

* errors will contain an array of either statements of success or errors explaining why the permission change failed

Index

A

access 121

alarms 35

apache 166

Autonomous Systems 102

C

collector 160

configuration 2

D

database 123

denied 121

dropped 160

E

Executive Summary 53, 88

F

failure 123

Flow Analytics 59

I

interfaces 88

IP addresses 102

M

main page 1

main view 88

manual 1

Mapping 78, 79, 83, 85, 86, 87

Maps 78

Mashup 53

MyView 53

N

NetFlow 102

O

overview 88

P

preferences 2

S

settings 2

sFlow 102

status 88

status tab 88

summary 88

T

trouble shooting 164

W

web server 166