
Scrutinizer Documentation

Version 19.5.2

Plixer

October 21, 2024

VERSION 19.5.2

| | | |
|----------|---|------------|
| 1 | Plixer Scrutinizer - Overview | 1 |
| 1.1 | What is Plixer Scrutinizer? | 1 |
| 1.2 | How does Plixer Scrutinizer work? | 1 |
| 1.3 | Main features | 2 |
| 2 | Deployment Guides | 3 |
| 2.1 | Hardware appliance | 3 |
| 2.2 | Virtual appliances | 5 |
| 2.3 | Plixer ML engine | 15 |
| 2.4 | Initial configuration | 26 |
| 3 | Configuration Guides | 31 |
| 3.1 | Alarms and Events | 31 |
| 3.2 | Configuration checklist | 35 |
| 3.3 | Distributed environments | 37 |
| 3.4 | Flow Analytics | 40 |
| 3.5 | Device groups | 49 |
| 3.6 | Plixer ML Engine | 50 |
| 3.7 | Environment sizing | 53 |
| 4 | Use Cases | 71 |
| 4.1 | NetOps Use Cases | 72 |
| 4.2 | SecOps Use Cases | 93 |
| 5 | Features and Functionality | 113 |
| 5.1 | Plixer Scrutinizer web interface | 113 |
| 5.2 | Data aggregation | 292 |
| 5.3 | Machine learning | 294 |
| 6 | Advanced Services | 297 |

| | | |
|----------|------------------------------------|------------|
| 6.1 | Integrations | 297 |
| 6.2 | Interactive CLI | 328 |
| 6.3 | Plixer Scrutinizer APIs | 346 |
| 6.4 | Reverse-path filtering | 383 |
| 6.5 | Streaming to data lakes | 386 |
| 6.6 | Backups | 386 |
| 6.7 | Updates and upgrades | 400 |
| 7 | Additional Resources | 413 |
| 7.1 | Changelog | 413 |
| 7.2 | FAQ | 442 |
| 7.3 | Functional IDs | 447 |
| 7.4 | Localization | 448 |
| 7.5 | Glossary | 448 |
| 7.6 | Third-party attributions | 452 |
| 7.7 | Plixer Technical Support | 475 |

PLIXER SCRUTINIZER - OVERVIEW

1.1 What is Plixer Scrutinizer?

Plixer Scrutinizer is a network monitoring and analysis appliance that collects, interprets, and contextualizes data from every digital exchange and transaction to deliver insightful network intelligence and security reports.

1.2 How does Plixer Scrutinizer work?

Plixer Scrutinizer collects network-related metadata from existing infrastructure, such as switches, firewalls, and packet brokers, and consolidates the information into a single unified database through efficient, dynamic correlation. Combined with a modern, hierarchical architecture, this allows the system to scale up and process millions of flows per second and report the meaningful, actionable data your IT professionals need via an intuitive web interface.

Plixer Scrutinizer is available as a rack-mountable hardware appliance or in virtualized ESX-, Hyper-V-, KVM- or AWS-based packages.

1.3 Main features

Reduces NetOps and SecOps complexity

Gain actionable insights from immense volumes of raw flow data via accessible, context-aware visualizations and reports.

Delivers essential network metadata when you need it

Get accurate, up-to-date statistics covering bandwidth, application, and user utilization that extends from clients through to the cloud with industry-leading reporting protocols.

Continuously monitors traffic for any irregularities

Combine proactive thresholds, alerts, and open RESTful APIs with comprehensive DDoS attack detection to build dynamic and streamlined event response strategies.

Minimizes downtime and loss of revenue

Gain true end-to-end visibility with real-time database updates that quickly identify root causes and reduce time-to-resolution metrics.

Maximizes efficiency with intelligent detection and reporting

Leverage the AI-backed capabilities of the Plixer ML Engine for intelligent threat and anomaly detection.

Advanced integration with other Plixer products

Seamlessly combines Plixer Scrutinizer with other Plixer products to get the exact functionality you need for your environment.

DEPLOYMENT GUIDES

This section covers the installation and initial setup procedure for Plixer Scrutinizer and includes individual guides for each type of deployment.

2.1 Hardware appliance

After removing the Plixer Scrutinizer hardware appliance from the box and verifying that there are no missing accessories (rackmount kit, appliance-locking bezel and keys, power cord), mount the appliance in a standard 19-inch rack or cabinet if desired.

Important: If your box arrives torn, dented, or otherwise damaged, the appliance itself seems damaged, or there are missing parts, *contact Plixer Technical Support* immediately and **do not attempt to install the unit**.

From there, follow these steps to set up the Plixer Scrutinizer hardware appliance:

1. Locate the ports to be used on the rear panel of the appliance (from left to right):
 - iDRAC
 - Serial
 - VGA
 - USB Type-B x 2

- 10GbE SFP x 2 (1 and 2)
- 1GbE RJ45 x 2 (3 and 4)
- Power supply x 2

Hint: To enable high-availability configurations and increase traffic capacity, the Ethernet port pairs are configured for bonding mode 6 (adaptive load balancing). The appliance is also equipped with two power supply units (PSUs) for redundancy.

2. Connect the power cable to one of the power supply sockets and plug the other end to a grounded AC outlet or UPS. To take advantage of the redundant PSUs, ensure that each socket is connected to an independent power source.
3. Depending on the bandwidth requirements of the environment, connect the appliance to the network using either RJ-45 or fiber optic cables. Unused ports may be left uncabled, but connecting both ports of either pair is recommended for high availability.

Note: The Plixer Scrutinizer hardware appliance comes equipped with required SFP+ transceivers to support 10 Gbps transfer rates over the SFP ports. Neither the RJ-45 nor the fiber optics cables are included in the package.

4. [Optional] Connect the iDRAC port to a remote access controller using an RJ-45 cable to enable remote console access for hardware management and monitoring. [Contact Plixer Technical Support](#) for help with configuring alerts for hardware-related events.
5. Using the additional ports provided, connect a monitor and keyboard to use during the appliance's initial configuration.

Hint: The iDRAC Virtual Console can also be used for the initial configuration process.

Once the Plixer Scrutinizer hardware appliance has been set up and cabled, proceed to configuring the appliance by following the steps outlined in the [Initial configuration](#) guide.

Note: Depending on the model of your Plixer Scrutinizer hardware appliance, ports may be positioned differently on the rear panel. Refer to the label on each port for correct cabling.

2.2 Virtual appliances

Plixer Scrutinizer is available in standard virtual appliance packages for VMware ESXi, Microsoft Hyper-V, and KVM environments or as an Amazon Machine Image (AMI) for Amazon Web Services (AWS).

All Plixer Scrutinizer virtual appliance types are available through Plixer or a local reseller, who will assist you with acquiring the evaluation or subscription license key required to activate the product.

System requirements:

Plixer Scrutinizer virtual appliances have the following basic system requirements:

| Component | Minimum (for trial installations) | Recommended (for production environments) |
|-----------|-----------------------------------|---|
| Memory | 16 GB | 64 GB |
| Storage | 100 GB | 1+ TB 15K RAID 0 or 10 configuration |
| Processor | 8 CPU cores , 2.0+ GHz | 12 CPU cores, 2.0+ GHz |

To ensure optimal performance, Plixer Scrutinizer virtual appliances should be provisioned with dedicated rather than shared resources, especially in environments where higher flow rates are expected. Hardware appliances, which are designed to support extremely large exporter counts and flow volumes, are recommended for large-scale enterprise networks.

Note: In clustered virtual environments where the hostname and MAC address of the VM can be changed, assign a static MAC address to the Plixer Scrutinizer NIC to avoid license key issues.

2.2.1 ESXi

This Plixer Scrutinizer virtual appliance deployment guide for ESXi environments is divided into the following subsections:

Deployment guide - ESXi

To deploy the ESXi-based Plixer Scrutinizer virtual appliance, take note of the following additional requirements and proceed with the subsequent setup process:

Additional requirements for ESXi deployments:

- ESXi 6.7 U2+
- VMware vSphere or vCenter
- Plixer Scrutinizer OVF file
- Plixer Scrutinizer VMDK file

Deploying the OVF template

1. After downloading the latest version of the Plixer Scrutinizer virtual appliance, connect to the ESXi host where the appliance will be deployed using VMware vSphere or vCenter.
2. Right-click the host the appliance will be deployed on and select *Deploy OVF Template* from the menu.
3. Select *Local file* and browse to the Plixer Scrutinizer OVF and VMDK files before clicking *Next*.
4. Provide a name for the Plixer Scrutinizer virtual appliance and continue to follow the deployment wizard.
5. When prompted, select the datastore, set the disk format to *Thick Provision* and click *Next*.
6. After selecting the network to be used by the virtual appliance, verify the configuration in the summary and click *Finish* to import the Plixer Scrutinizer virtual appliance. This may take a few moments.
7. Before powering on the Plixer Scrutinizer virtual machine, assign a static MAC address to the NIC for licensing purposes:
 - a. Right-click on the VM and select *Edit Settings...*
 - b. Select the network adapter, set the MAC address to *Manual*, and enter a unique MAC address to assign to the virtual machine NIC.
 - c. While on this page, adjust the other virtual hardware settings to match the recommended specifications outlined in the [virtual appliance deployment overview](#) if necessary.
 - d. Click *OK* to save the current configuration and return to the previous page.
8. Right-click on the Plixer Scrutinizer virtual machine to power it on. Afterwards, click the console preview window and select *Open Remote Console*.

When the new window opens, login to the Plixer Scrutinizer virtual appliance using `plixer:scrutinizer` and proceed with the [basic configuration process](#).

Upgrading the VM hardware version

To upgrade the hardware version of the virtual machine to the latest version of ESXi, follow these steps:

1. With the Plixer Scrutinizer virtual machine powered off, right-click on it in vSphere (or vCenter).
2. Under the **Compatibility** submenu, select *Upgrade VM Compatibility*.
3. When asked, click *Yes* to continue with the virtual machine upgrade.
4. Once the process is complete, power on the virtual machine.

The Plixer Scrutinizer virtual appliance VM will boot up with the latest ESXi hardware version available.

Installing VMware Tools

The Plixer Scrutinizer package includes a script to install VMware Tools, which is not installed by default. The recommended time to install VMware Tools is after the [basic appliance configuration process](#).

Installing VMware Tools will allow you to properly shutdown the Plixer Scrutinizer virtual machine when needed.

To run the script to install VMware Tools, follow these steps:

1. Log in to the appliance as the **plixer** user using the configured password.
2. Launch the interactive `scrut_util` by entering the following at the prompt:

```
[plixer@scrutinizer ~]$ /home/plixer/scrutinizer/bin/scrut_util
```

3. At the Plixer Scrutinizer interactive prompt, enter the following command:

```
SCRUTINIZER> enable vmwaretools
```

4. After the operation completes, type `exit` or `quit` to terminate the interactive prompt.

Once VMware Tools has been enabled on the appliance, select *Shut Down Guest* under the **Power** submenu to shut down the Plixer Scrutinizer virtual machine. Do **NOT** select *Power Off* as this may result in database corruption.

Expanding database size

Depending on the volume of NetFlow data that will be forwarded to the Plixer Scrutinizer virtual appliance, it may be necessary to allocate additional storage space for its database.

This process is divided into several tasks:

Adding a hard drive to the Plixer Scrutinizer virtual machine

1. Power off the Plixer Scrutinizer VM by either logging in and issuing the `sudo shutdown -h now` command or using the [VMware Tools power menu](#).
2. Right-click on the virtual machine and select *Edit Settings...*
3. Click *Add New Device* and select *Hard Disk* from the dropdown.
4. Expand the *New Hard disk* settings, and select the type of Disk Provisioning and adjust the disk capacity before clicking *OK*.
5. Right-click the Plixer Scrutinizer virtual machine and power it on.

Once the hard drive has been added, it will need to be set up for use by the Plixer Scrutinizer virtual appliance.

Configuring Plixer Scrutinizer to use the new drive

1. Log in to the Plixer Scrutinizer virtual appliance as the `plixer` user.
2. Launch the interactive `scrut_util` by entering the following at the prompt:

```
[plixer@scrutinizer ~]$ scrut_util
```

3. At the `SCRUTINIZER>` prompt, enter `show diskpace` to view the current size of the database mounted on `/var/db` and then use `show partitions` to view the available disks.
4. Still at the `SCRUTINIZER>` prompt, issue `set partitions <new_partition>` to make the added hard drive available to the Plixer Scrutinizer virtual appliance.
5. When prompted, select whether or not you have a backup of your data.

Wait for the operation to complete automatically before proceeding to the next task.

Verifying the new filesystem size

To confirm that the new hard drive has been successfully added, enter the `show diskpace` command again at the `SCRUTINIZER>` prompt. The new size of the database should reflect the additional space added with the new hard disk.

2.2.2 Hyper-V

This Plixer Scrutinizer virtual appliance deployment guide for Hyper-V environments is divided into the following subsections:

Deployment guide - Hyper-V

To deploy the Plixer Scrutinizer virtual appliance for Hyper-V environments, take note of the following additional requirements and proceed with the subsequent setup process:

Additional requirements for Hyper-V deployments:

- Generation 2 Hyper-V VM
- Hyper-V 2012
- Hyper-V Manager
- Latest Plixer Scrutinizer virtual appliance for Hyper-V

Deploying the Hyper-V virtual appliance

1. After downloading the latest version of the Plixer Scrutinizer virtual appliance, unzip the file on your Hyper-V server.
2. Open Hyper-V Manager, right-click the virtual machine, and select *Import Virtual Machine...*
3. Browse to the location of the Scrutinizer_Hyper-V folder.
4. Select the Scrutinizer_Hyper-V virtual machine and click *Next*.
5. Use the radio buttons to select the import operation type and click *Next*.
6. Verify the settings in the summary and click *Finish* to import the virtual machine.
7. Right-click on the Plixer Scrutinizer virtual machine and select *Settings...*
8. In the **Settings** menu, set the *Startup RAM:* to 16 GB (if not already set).
9. Select a network adapter and assign it to the appropriate virtual switch.
10. Expand the network adapter settings, select *Advanced Features*, and set the MAC address to *Static*.
11. Enter a unique MAC address and click *OK*.
12. After starting the virtual machine, right-click on it and select *Connect*.

From there, login to the Plixer Scrutinizer virtual appliance using `plixer:scrutinizer` and proceed with the [basic configuration process](#).

Expanding database size

Depending on the volume of NetFlow data that will be forwarded to the Plixer Scrutinizer virtual appliance, it may be necessary to allocate additional storage space for its database.

To add a hard drive to the Plixer Scrutinizer virtual machine, follow these steps:

1. Power off the Plixer Scrutinizer VM by logging in and issuing the `sudo shutdown -h now` command.
2. In Hyper-V manager, right-click on the Plixer Scrutinizer virtual machine and select *Settings*.
3. Under the *IDE Controller* settings, select *Hard Drive* and click *Add*.
4. Under *Virtual hard disk*:, click *New* to start the **New Virtual Hard Disk** wizard.
5. When asked to choose the disk format, select VHDX to allow for expansion past 2 TB.
6. Continue to follow the wizard and provide the requested details.
7. Review the settings in the summary and click *Finish* to complete the operation.
8. Power on the virtual machine and follow the [steps to configure the hard drive for use by the Plixer Scrutinizer virtual appliance](#) under the ESXi subsection of the deployment guides.

When done, verify that the new hard drive has been successfully added by entering `show diskspace` at the `SCRUTINIZER>` prompt and confirming the new size of the database.

2.2.3 KVM

To deploy the KVM-based Plixer Scrutinizer virtual appliance, take note of the following additional requirements and proceed with the subsequent setup process:

Additional requirements for KVM deployments:

- KVM 16
- Latest Plixer Scrutinizer virtual appliance for KVM

Deploying the KVM virtual appliance

1. Create a directory for your install by entering:

```
mkdir kvm/Scrut_VM_Guide/
```

2. Download the latest version of the Plixer Scrutinizer virtual appliance for KVM by entering `wget https://files.plixer.com/Scrutinizer_KVM_Image.tar.gz`.

```
wget https://files.plixer.com/Scrutinizer_KVM_Image.tar.gz
```

Hint: If the above URL does not work, [contact Plixer Technical Support](#).

3. Unzip the file on your KVM server to the new directory:

```
sudo tar xvzf Scrutinizer_KVM_Image.tar.gz
```

4. Run the script to install the Plixer Scrutinizer virtual appliance:

```
sudo ./install-kvm-scrut.shared
```

5. Wait for the confirmation that the virtual machine has been created from the image.

Once Plixer Scrutinizer virtual machine has been created, log in using the command `virsh console Scrutinizer` with the credentials `plixer:scrutinizer` and proceed with the [basic configuration process](#).

2.2.4 AWS (AMI)

This Plixer Scrutinizer virtual appliance deployment guide for AWS is divided into the following subsections:

Deployment guide - AWS

To deploy Plixer Scrutinizer under AWS, start by going to the AWS Marketplace product page to subscribe to the service.

After subscribing, proceed with the following steps to continue the deployment process:

1. Navigate to the EC2 Dashboard in the AWS console and click *Launch Instance*.
2. In the navigation pane, select *My AMIs* and choose the Plixer Scrutinizer AMI from the list.
3. For flow rates up to 10,000 flows per second, select *C5.2xlarge* as the instance type. If expected rates exceed 10,000, contact [Plixer Technical Support](#) for assistance.
4. Navigate to the **Configure Instance Details** page, set *Shutdown Behavior* to *Stop* and enable *Termination Protection*.
5. Select the network and subnet to assign to the instance using the dropdowns and assign the IP addresses to the AMI.

Hint: Because AWS does not allow an active instance's primary private IP address to be released, it is recommended to deploy the AMI with two NICs and use the secondary NIC as the collection interface.

6. When prompted to add storage, leave the root volume size (*/dev/xvda/*) at the default 100 GB capacity. To add additional storage, allocate additional disks after the instance is running by following [these instructions](#).
7. (Optional) Add tags to categorize your AWS resources (e.g., by purpose, owner, or environment) to streamline resource management.
8. When prompted to configure the security group for the instance, either create a new security group to define firewall rules for the Plixer Scrutinizer instance or assign it to an existing group with the necessary firewall rules.
9. Verify the instance launch details via the **Review Launch** before clicking *Launch* to assign a key pair to the instance and complete the process.

Important: Do not lose the SSH key that you will be asked to generate as part of the deployment process, because it is the only way to access the server via SSH.

Once the instance has been launched, SSH to the server as `ec2-user` using the following command and use `sudo su -` to elevate to root:


```
ssh -i /path/to/key/key.pem ec2-user@scrutinizerIPAddress
```

Accessing the web interface

To access the Plixer Scrutinizer AMI web interface, use a supported browser to navigate to its primary private or public IP address.

After accepting the user agreement, log in as the `admin` user with the AMI instance ID as the password.

Changing passwords

After the Plixer Scrutinizer AMI instance has been deployed, follow these steps to change the web interface login credentials:

1. Navigate to the **Admin** tab in the Plixer Scrutinizer web interface and select the `admin` user.
2. Enter a new password and click *Save*.
3. (Optional) To enforce more complex password requirements, navigate to **Admin > Settings > System Preferences** and tick the *Enforce Password Complexity* checkbox.

When *Enforce Password Complexity* is enabled new passwords will need to consist of at least eight characters, including one capital letter, one number, and one special character.

Expanding database size for AWS

Depending on the volume of NetFlow data that will be forwarded to the Plixer Scrutinizer virtual appliance, it may be necessary to allocate additional storage space for its database.

To add storage to the Plixer Scrutinizer AMI, follow these steps:

1. Get the correct availability zone for the instance by checking the *Availability Zone* column of the instance page.
2. Navigate to the **Volumes** page and click *Create Volume*.
3. On the **Create Volume** page, create the new volume with the desired size and in the correct availability zone.

4. After the new volume has been created, right-click on it and select *Attach Volume*.
5. Start typing in the instance name, select it from the dropdown, and click *Attach*.
6. Once the new volume has been attached, follow the [steps to configure it for use by the Plixer Scrutinizer virtual appliance](#) under the ESXi subsection of the deployment guides.

When done, verify that the new hard drive has been successfully added by entering `show diskspace` at the `SCRUTINIZER>` prompt and confirming the new size of the database.

Hint: When adding more than one new drive to a single AWS instance, the `set partitions` command will need to be run for each additional drive.

Adding instance resources

As resource needs change, it may become necessary to resize the Plixer Scrutinizer AMI instance and either change its instance type to one that is compatible with the new configuration or migrate the application to a new instance.

To resize the Plixer Scrutinizer instance size and/or make the necessary changes to the instance type, follow these steps:

1. SSH to the instance and stop all services by entering the following command at the `SCRUTINIZER>` prompt:

```
SCRUTINIZER> services all stop
```

2. Power off the operating system by issuing the `shutdown -h now` command and open the AWS EC2 console.
3. In the navigation pane, click *Instances* and select the Plixer Scrutinizer AMI instance.
4. If the instance has an associated Elastic IP address, take note of it and the instance ID shown in the details pane.
5. Under *Actions*, select *Instance State* and then *Stop*. Confirm the action and wait for the instance to stop.

Note: Once the instance is stopped, the *Elastic IP*, *Public DNS (IPv4)*, *Private DNS*, and *Private IPs* fields will be blank to indicate that the old values are no longer associated with the instance.

6. With the instance still selected, choose *Actions* and then *Instance Settings*.
7. Click *Change Instance Type* and select the desired instance type. If it is not available, the instance type is not compatible with the instance configuration (e.g., because of virtualization type).
8. When done, select the instance again and choose *Actions* > *Instance State* > *Start* to restart the instance. It may take several minutes for the instance to start up again after confirming the action.
9. Once the instance is running, confirm that the *Public DNS (IPv4)*, *Private DNS*, and *Private IPs* fields are populated with the new values assigned.
10. If the instance had an associated Elastic IP address, reassociate it by clicking *Elastic IPs* in the navigation pane and selecting the Elastic IP address that was previously associated with the instance.
11. With the correct Elastic IP address selected, click *Actions* and then *Associate Address*. After that, select the previous instance ID and click *Associate* to complete the operation.

Once the instance has been reconfigured, SSH to the instance and issue the `set tuning` command at the `SCRUTINIZER>` prompt to retune the appliance.

2.2.5 Optimizing datastores

Because of how NetFlow works, large Plixer Scrutinizer deployments can be negatively impacted by insufficient disk I/O throughput. For optimal performance, Plixer Scrutinizer virtual appliances should be deployed on a dedicated 15K RPM RAID 10 datastore whose storage capacity meets your history setting requirements (1.8 TB recommended).

In environments with extremely high flow volumes, it is recommended to use the Plixer Scrutinizer hardware appliance for its dedicated resources and higher collection rates.

2.3 Plixer ML engine

Deploying a machine learning engine means integrating machine learning models into various environments where they can be used to make predictions and decisions in real-time.

Deployment options vary and can include cloud-based platforms such as AWS and Azure, which offer scalability and ease of use, as well as local and on-premise solutions like vSphere, providing greater control and security.

2.3.1 Authentication token

Before anything else, a Plexier Scrutinizer auth token must be manually created. To do this, follow these steps:

1. On the Plexier Scrutinizer web interface, navigate to **Admin > Users & Groups > Authentication Tokens**.
2. Click the + button to add a new authentication token.
3. Set the expiration date of the authentication token to **One Year**.
4. Set the user to **admin**.
5. Click **Generate Token**. Take note of the generated token. You'll need this when running the `setup.sh` later in the setup process.

2.3.2 SSH credentials

During setup, the Plexier ML engine needs to be able to SSH into your Plexier Scrutinizer primary host. Take note of your primary IP address, as well as your `plexier` user's SSH password.

If you require a private key for SSH, make sure that the public key is configured on your Plexier Scrutinizer at `/home/plexier/.ssh/authorized_keys`. Make the private key available on the machine that you are deploying the ML engine from. Take note of the path to the private key file, you'll be asked for it during `setup.sh`.

2.3.3 Using an On-Premises Plexier Scrutinizer

Plexier Enterprise in AWS supports pairing with an On-Premises Plexier Scrutinizer, but a few network factors have to be considered.

VPC Site-to-Site connection

A VPC Site-to-Site connection can be used to enable communication between Plexier Enterprise in AWS and an On-Premises Plexier Scrutinizer. To configure a Site-to-Site connection between a VPC and a customer's environment, follow the instructions here.

Direct Access via Public IP

An alternative to using a Site-to-Site connection is allowing external access to the On-Premises Plexier Scrutinizer using a public IP address.

In this case, the Plexier Scrutinizer instance in your environment will be exposed to external access via the internet. Ports 5432, 22, and 443 will be exposed. This public IP is what you will specify when running `01_aws_infrastructure.sh` or `setup.sh`.

Take note of the internet gateway IP, because this will be manually added to the Plexier Scrutinizer `pg_hba.conf` file to allow access to Postgres. After adding the IP address to `pg_hba.conf`, run the command `psql -c "SELECT pg_reload_conf()"` on Plexier Scrutinizer to reload the configuration.

2.3.4 Setting environment variable on Plixer Scrutinizer

Important: This step applies only to version 19.4 of the Plixer ML engine. Future releases of the Plixer ML engine eliminates the need to set up the environment variable.

Before pairing the ML engine with Plixer Scrutinizer, the environment variable `NO_ML_HEARTBEAT` must be set. To do this, add the line `export NO_ML_HEARTBEAT=1` to the `local_env` file.

```
$ cat ~/scrutinizer/env/local_env
### Add custom environment variable configuration directory that
### should not be overwritten on upgrade here.
### other values may be present here...
export OTHER_VARIABLE=1

export NO_ML_HEARTBEAT=1
```

To make sure that the variable is set:

```
$ echo $NO_ML_HEARTBEAT
1
```

Once the variable is set, follow the corresponding guide below to pair your ML engine with your Plixer Scrutinizer.

AWS (EKS)

The Plixer ML Engine deployment utilizes a separate host VM to deploy, monitor, and control the deployment. All guide content and steps are managed from that host VM, which is distinct from the actual deployment nodes and resources. It is recommended to deploy the Plixer ML VM image as a deployment host because it installs all the required prerequisites.

The following are the software prerequisites required:

- awscli
- Docker
- Terraform

Once all required software is installed, proceed with setting up the following:

AWS secrets

To connect to AWS you will need to provide a secret key pair from an IAM user (Secret Access Key ID and Secret Access Key). You will be asked for these in `01_aws_infrastructure.sh`.

VPC and subnet names

You will need to know which VPC and subnets (2 are required) you want to deploy to. You will be asked for these during `01_aws_infrastructure.sh`, or you can configure them directly in `aws.tfvars`.

Creating a new VPC

`01_aws_infrastructure.sh` will ask you for new VPC details, or you can set `new_vpc_cidr`, `new_vpc_public_cidr`, and `new_vpc_private_cidr` manually in `aws.tfvars`.

Deploying to an existing VPC

If deploying to an existing VPC and existing subnets, the VPC must meet the following requirements:

- A DHCP option set applied to the VPC with **domain-name-servers: AmazonProvidedDNS**
- Two private subnets within the the VPC, in two different AZs
- If no NAT gateway is attached to the private subnets (i.e. subnets have no internet access), then set `airgap_install` to `TRUE`

For more information, see the article on Amazon EKS VPC and subnet requirements and considerations.

Setting up Plixer Kubernetes in AWS

Run the following steps on a cloned VM template that will act as the *deploy host* to manage the rest of the infrastructure.

1. Log in as the `plixer` user, and then accept the EULA and set up networking.
2. Navigate to the `/home/plixer/common/kubernetes` directory.
3. Run `01_aws_infrastructure.sh`, which for a new install should walk you through configuring required values.

Note:

- If you are reconfiguring an existing setup, run `01_aws_infrastructure.sh --configure`, or you can also manually configure the required values in `aws.tfvars`.
 - If you get an error complaining that the `tfstate s3` bucket already exists and that the cluster name that you have set isn't being used by another user, you can run `01_aws_infrastructure.sh --recover-tfstate`, and then continue with step 3 again.
-

4. Run `kubect1 get nodes` to make sure that the Kubernetes infrastructure has been deployed successfully.

Setting up the ML engine in AWS

After the Kubernetes infrastructure has been deployed and is now running successfully, you may proceed with deploying the ML engine in AWS.

1. Run the following command: `/home/plixer/ml/setup.sh`
2. Enter the path to your Plixer Scrutinizer private [SSH key](#) or SSH password.
3. Enter your Plixer Scrutinizer [authentication token](#).
4. Run `kubectl get pods`. The pods will have the **Running** status which means that it has been deployed successfully.

Note: If you encounter the **Error: Provider configuration not present** when deploying in AWS, simply run `terraform -chdir=aws state rm module.eks.data.http.wait_for_cluster` and re-run `01_aws_infrastructure.sh`.

aws.tfvars field reference table

| Field name | Description |
|----------------------------|--|
| creator | REQUIRED: This is the name of the person creating these AWS resources, used as a tag in AWS to track utilization. |
| cost_center | REQUIRED: This is the cost center to use for these AWS resources, used as a tag in AWS to track utilization. |
| aws_certificate | REQUIRED: This is the name of an existing SSH certificate configured in your AWS environment. You can see a list of these in your AWS Console by navigating to EC2 > Network > Security > Key Pairs . |
| instance_type | REQUIRED: This is the AWS instance type to create for EKS worker nodes (i.e. t2.large). |
| fargate | REQUIRED: Use fargate instead of EKS nodes for applicable workloads. Setting the value to TRUE will allow using a smaller instance_type. |
| aws_region | REQUIRED: The AWS region to deploy infrastructure in. |
| public_load_balancer | REQUIRED: Whether or not to deploy a public load balancer for external access. |
| airgap_install | OPTIONAL: If this is an airgapped install (i.e. the vpc_private_subnets don't have a route to a NAT gateway), then set this to TRUE. |
| create_ec2_endpoint | OPTIONAL: If airgap_install = TRUE, this bool controls whether or not to create an EC2 endpoint in the VPC. |
| create_s3_endpoint = true | OPTIONAL: If airgap_install = true, this bool controls whether or not to create an S3 endpoint in the VPC. |
| create_ecr_endpoint = true | OPTIONAL: If airgap_install = true, this bool controls whether or not to create an ECR endpoint in the VPC. |
| create_ssm_endpoint = true | OPTIONAL: If airgap_install = true, this bool controls whether or not to create an SSM endpoint in the VPC. |
| new_vpc_cidr | OPTIONAL: If you want to create a new VPC, then specify the IP address range in this field. |
| new_vpc_public_cidr | OPTIONAL: If you want to create a new VPC, then specify the IP address range for the public subnet in the new VPC. |
| new_vpc_private_cidr | OPTIONAL: If you want to create a new VPC, then specify the IP address range for the private subnet in the new VPC. |
| azs | OPTIONAL: Availability zones corresponding to the subnets you want created in new_vpc_public_cidr and new_vpc_private_cidr. |
| vpc_name | OPTIONAL: Existing vpc_name to create the EKS resources in. |
| vpc_private_subnets | OPTIONAL: List of private subnet names to create the EKS resources in. |
| vpc_public_subnets | OPTIONAL: List of public subnet names to create the EKS resources in. |

Azure (AKS)

The Plixer ML Engine deployment utilizes a separate host VM to deploy, monitor, and control the deployment. All guide content and steps are managed from that host VM, which is distinct from the actual deployment nodes and resources. It is recommended to deploy the Plixer ML VM image as a deployment host because it installs all the required software prerequisites.

The following are the software prerequisites required:

- `azurecli`
- Docker
- Terraform

Once all required software is installed, proceed with setting up the following:

Configuring Azure CLI

1. Download and install the Azure CLI.

Note: Specific steps will depend on your OS. You may refer to the install guide [here](#).

2. Run `az login` to set up the client. This will prompt you to log in to your Azure account.

Important: The user account being used to deploy **MUST** be assigned to the “owner” role in Azure. This is because AKS needs to assign a role for the AKS cluster user, which only owners can perform.

3. Make sure you have a `.azure` directory at your current user’s home directory.

Setting up Plixer Kubernetes in Azure

Run the following steps on a cloned VM template that will act as the *deploy host* to manage the rest of the infrastructure.

1. Log in as the `plixer` user, and then accept the EULA and set up networking.
2. Navigate to the `/home/plixer/common/kubernetes` directory.
3. Modify `azure.tfvars`. You may refer to this [table](#) for the field values in `azure.tfvars`.
4. Run `01_azure_infrastructure.sh`.
5. Run `kubect1 get nodes` to make sure that the Kubernetes infrastructure has been deployed successfully.

Setting up the ML engine in Azure

After the Kubernetes infrastructure has been deployed and is now running successfully, you may proceed with deploying the ML engine in Azure.

1. Run the following command: `/home/plixer/ml/setup.sh`
2. If deploying to a VPC with internet access you will be asked for docker credentials.
3. Enter the path to your Plixer Scrutinizer private [SSH key](#) or SSH password.
4. Enter your Plixer Scrutinizer [authentication token](#).
5. Run `kubectl get pods`. The pods will have the **Running** status which means that it has been deployed successfully.

azure.tfvars field reference table

| Field name | Description |
|---------------------|---|
| cluster_name | REQUIRED: The name to associate with and identify this deployment. |
| vm_type | REQUIRED: This is the Azure VM instance type to create for AKS worker nodes. |
| location | REQUIRED: This is the location to create the AKS worker nodes in (e.g. East US 2). |
| resource_group_name | OPTIONAL: Name of existing resource group to use when deploying assets. If empty, resource group named <code>\${var.cluster_name}-resource-group</code> will be created. <code>resource_group_name</code> must also be in the specified location field. |
| vnet_name | OPTIONAL: Name of existing VNET to deploy AKS in. |
| vnet_subnet_name | OPTIONAL: Name of existing subnet within <code>vnet_name</code> to deploy AKS in. Each subnet can only contain one aks cluster. |
| vnet_address_space | OPTIONAL: If <code>vnet_name</code> is not specified, then use this address space when creating the new VNET to place AKS in. By default, value is set to 172.18.0.0/16. |
| new_subnet_cidr | OPTIONAL (required if <code>vnet_subnet_name</code> is not specified): If <code>vnet_subnet_name</code> is not specified, then use this address space when creating the new VNET subnet to place AKS in. Value must be within the address space of the specified VNET. Default value is set to 172.18.1.0/24. |
| public_node_ip | OPTIONAL: Whether or not to assign public IPs to AKS nodes. By default, value is set to FALSE . |
| service_cidr | OPTIONAL: Service CIDR space for internal k8s services. Must not conflict with the address space of the VNET being deployed to. By default, value is set to 172.19.1.0/24. |
| dns_service_ip | OPTIONAL: Service IP to assign to the k8s internal DNS service. Must be within the address space specified by <code>service_cidr</code> . By default, value is set to 172.19.1.5. |

Local (Single Node)

After the Plixer ML Engine completes its first boot sequence, login with the user `plixer` and password `plixer`. Accept the EULA, and then configure the network settings.

Once the initial configuration is complete, the appliance will restart, and then you can proceed with deploying a standalone (single node) host. To do this, make sure that the following prerequisites are met first:

Setting up Plixer Kubernetes locally

1. Navigate to the `/home/plixer/common/kubernetes` directory.
2. Run `./01_local_infrastructure.sh`.
3. Run `kubectl get nodes` to make sure that the Kubernetes infrastructure has been deployed successfully. It should return like the following:

```
$ kubectl get nodes
```

| NAME | STATUS | ROLES | AGE | VERSION |
|------|--------|---------------------------|-----|----------------|
| | Ready | control-plane,etcd,master | 39d | v1.27.1+rke2r1 |

Make sure that all of the nodes have a running or ready status. If not, wait for a few minutes, and then check again.

Setting up the ML engine locally

After the Kubernetes infrastructure has been deployed and is now running successfully, you may proceed with deploying the ML engine locally.

1. Navigate to `/home/plixer/ml`.
2. Run `setup.sh`.
 - For offline (no internet access) installation, run: `/home/plixer/ml/setup.sh --offline`
 - For online installation, run: `/home/plixer/ml/setup.sh`.
3. Select if your Plixer Scrutinizer instance needs a private key for SSH.
 - If Yes, enter the path to your Plixer Scrutinizer private [SSH key](#) or SSH password.
 - If No, proceed to the next step.
4. Select whether your Plixer Scrutinizer is a distributed cluster or not.
 - If Yes, enter the IP address of your Plixer Scrutinizer reporter.
 - If No, enter the IP address of your Plixer Scrutinizer host.
5. Enter your Plixer Scrutinizer [authentication token](#).
6. Run `kubectl get pods`. The pods will have the **Running** status which means that it has been deployed successfully.

vSphere (Multi-Node) Cluster

The Plixer ML Engine deployment utilizes a separate host VM to deploy, monitor, and control the deployment. All guide content and steps are managed from that host VM, which is distinct from the actual deployment nodes and resources. It is recommended to deploy the Plixer ML VM image as a deployment host because it installs all the required software prerequisites.

The following are the software prerequisites required:

- Ansible
- Docker
- Terraform

Important: Put the Plixer ML VM template in vSphere. Set **vm_template** in `vsphere.tfvars` to the path of the VM template in vSphere.

Once all required software is installed, proceed with setting up the following:

Setting up Plixer Kubernetes in vSphere

Important: Make sure that the `vsphere_user` defined in `vsphere.tfvars` has permissions in vSphere to create VMs and resource groups. You will be prompted for the `vsphere_user` password when running `01_vsphere_infrastructure.sh`.

1. Log in as the `plixer` user, and then accept the EULA and set up networking.
2. Navigate to the `/home/plixer/common/kubernetes` directory.
3. Modify `vsphere.tfvars`. You may refer to this [table](#) for the field values in `vsphere.tfvars`.
4. Run `01_vsphere_infrastructure.sh`.
5. Run `kubectl get nodes` to make sure that the Kubernetes infrastructure has been deployed successfully.

Setting up the ML engine in vSphere

After the Kubernetes infrastructure has been deployed and is now running successfully, you may proceed with deploying the ML engine in vSphere.

1. Run `setup.sh`.
 - For offline (no internet access) installation, run: `/home/plixer/ml/setup.sh --offline`
 - For online installation, run: `/home/plixer/ml/setup.sh`.
2. Enter the path to your Plixer Scrutinizer private [SSH key](#) or SSH password.
3. Enter your Plixer Scrutinizer [authentication token](#).
4. Run `kubect1 get pods`. The pods will have the **Running** status which means that it has been deployed successfully.

vsphere.tfvars field reference

| Field name | Description |
|-----------------------|---|
| create_hosts | Whether or not to create vSphere hosts. If FALSE, then the IPs in <code>vm_master_ips</code> should correspond to the VMs created using the VM template. |
| vm_master_ips | List of IPs to assign to Kubernetes nodes. This must be 1 or 3 hosts (can't be an even number of IPs). |
| vm_haproxy_ip | The virtual IP address to assign to a VM running HAProxy. |
| vsphere_vcenter_ip | The IP address of the vCenter host to deploy on. |
| vsphere_username | vSphere user to connect with. |
| vsphere_datacenter | Datacenter in vSphere to deploy assets into. |
| vsphere_host | Host within the specified datacenter to deploy assets into. |
| vsphere_resource_pool | Resource pool to create for the VMs. |
| vm_folder | Folder name in vSphere to create the VMs in. |
| vm_datastore | The datastore name used to store the files of the VMs. |
| vm_network | The vSphere network name used by the VMs. |
| vm_gateway | The network gateway used by the VMs. |
| vm_dns | The DNS server used by the VMs. |
| vm_domain | The domain name used by the VMs. |
| vm_template | The vSphere template that the VM is based on. |
| vsphere_unverified | Set to TRUE to bypass the vSphere host certificate verification. |
| offline_install | If set to TRUE, then it will be assumed that the template being used to create the VMs already has all assets it needs, and will skip downloading the assets. |
| rke2_airgap | If set to TRUE and <code>offline_install</code> is also TRUE, then the script will attempt to proxy any downloads required for RKE2 Kubernetes setup through the host that <code>./01_vsphere_infrastructure.sh</code> is running on. |

2.4 Initial configuration

Once the Plixer Scrutinizer hardware or virtual appliance has been deployed and powered on, the next steps will involve configuring the appliance to prepare the system for use:

2.4.1 First login

After the Plixer Scrutinizer appliance completes its first boot sequence and a user logs in with the credentials `plixer:scrutinizer`, it will perform a quick setup before rebooting itself.

After the reboot, log in again to start the initial configuration script:

1. Provide the following information when prompted by the script:
 - Static IP address
 - Netmask
 - Gateway
 - FQDN
 - DNS IP address
 - NTP server IP address
2. Continue through the succeeding dialogs and enter any additional information requested.
3. At the end of the script, press *Enter* and wait for the server to reboot again to apply the settings.

After the final appliance reboot, point any supported browser to `http://[configuredIPaddress]` and log in with the default `admin:admin` credentials to access the Plixer Scrutinizer web interface, where the rest of the initial configuration steps will be performed.

2.4.2 Adding a license

To fully activate Plixer Scrutinizer, a valid and active evaluation or subscription license must first be added.

To add a new license:

1. Navigate to the **Admin > Plixer > Plixer Scrutinizer Licensing** in the Plixer Scrutinizer web interface.
2. Contact [Plixer Technical Support](#) and provide the engineer with the Machine ID displayed on the page.
3. After receiving the license key, paste it into the provided text field and click *Save*.

Once Plixer Scrutinizer has been successfully activated, the **Licensing** page will display the following details for the saved license:

- Product type
- Status
- Days left
- Customer ID
- Machine ID
- Server count (number of licensed, including reporters and Collectors)
- Reporter count (number of licensed primary and secondary reporters)
- Exporter count (number of licensed Exporters)
- Deployed servers
- Enabled Exporters

2.4.3 Changing the default admin password

To change the default password for the Plixer Scrutinizer web interface **admin** account, navigate to **Admin > Security > Users** and select the *admin* user from the list and enter the new password (must be entered twice) under the *Password* tab of the **Edit User** menu.

Hint: Settings for all web interface user accounts can be configured from the same page and menu.

2.4.4 Configuring SSL

As part of the initial configuration script/wizard for the Plixer Scrutinizer appliance, a self-signed SSL certificate will be created using default values. SSL support will also be enabled by default.

These settings can be later modified as needed.

Installing a CA-signed SSL certificate

As long as the system is set to use the self-signed SSL certificate created during the initial configuration process, browsers will return an untrusted certificate warning, which users must override to access the web interface.

To correct this behavior, an SSL certificate that has been signed by an internal or commercial Certificate Authority (CA) will need to be installed:

1. Forward the `/etc/pki/tls/private/ca.csr` file to the CA for signing and ask that they return it as base 64 encoded rather than DER encoded.

Important: The Plixer Scrutinizer AMI comes with its own self-signed certificate. This can be replaced with a new certificate by running the `scrut_util` command `set ssl on` as described above.

2. After receiving the CA-signed SSL certificate, stop the Apache service using the following `scrut_util` command:

```
SCRUTINIZER> services httpd stop
```

Note: For additional details on the Plixer Scrutinizer interactive CLI utility `scrut_util`, including usage instructions and command information, see the [Interactive CLI](#) section of this documentation.

3. Rename the new certificate to `ca.crt` and replace the previous certificate file located in `etc/pki/tls/certs`.
4. Start the Apache service again:

```
SCRUTINIZER> services httpd start
```


To verify that the web interface is using the correct SSL certificate, use a browser to navigate to the login page using the FQDN specified in the CA-signed certificate. The browser should no longer return an untrusted certificate warning and the padlock icon in the address bar should be locked instead of open.

Enabling/disabling SSL

If needed, SSL support can later be disabled (and later re-enabled) using the interactive CLI utility `scrut_util` command:

```
SCRUTINIZER> set ssl [off | on]
```

Important: Running `set ssl on` at a later time will prompt the user to enter new certificate details and overwrite the previous certificate.

2.4.5 Metadata sharing

By default, the Plixer Scrutinizer appliance is configured to collect certain types of metadata related to the general health, overall performance, or other targeted metrics, which will be used to improve Plixer Scrutinizer and the Plixer Network Detection and Response platform.

This metadata includes:

- **Vendor** - Indicates whether the product is branded under Plixer or an OEM vendor
- **smtpfromEmail** - Administrator email address configured under server preferences
- **smtp** - Name of email server used to identify the subscribing company configured under server preferences
- **installedVersion** - Current Plixer Scrutinizer version installed
- **License details** - License details, including level and status
- **Active Exporter count** - Number of Exporters that have sent flows in the past 24 hours
- **Server metrics** - Flows received, packets received, and flows dropped per second
- **Flow template details** - Templates and element names captured from Exporters currently sending flows to the system

- **Flow Analytics Exporter statistics** - Number of Exporters currently configured for each *Flow Analytics algorithm*.
- **Flow Analytics history:** - List of violator IP addresses from external sources; used to determine if outside patterns exist in data aggregated between globally installed Plixer Scrutinizer servers

How collected data is transferred

The collected metadata is first encrypted and then sent to ph.plixer.com (ports 443, 80, and 25 are used, depending on availability).

How Plixer uses the data

Plixer uses the collected metadata for support purposes only.

From the data, the *Plixer Technical Support team* is able to determine whether customers may require a support call to help address issues or upgrade their system.

Turning off metadata sharing

To disable this feature, navigate to **Admin > Settings > System Preferences** in the Plixer Scrutinizer web interface and untick the *Share Health Statistics* checkbox.

When done, click *Save* to save the new setting.

CONFIGURATION GUIDES

This section contains detailed guides for configuring Plexer Scrutinizer and tailoring its functions to user and organizational needs.

3.1 Alarms and Events

Plexer Scrutinizer uses various technologies to recognize patterns in system activity and network traffic that may be of interest to NetOps and SecOps teams. These patterns are then reported as Events via the [*Alarm Monitor views*](#).

Combined, the Alarm Monitor interface and Plexer Scrutinizer's library of Alarm Policies allow for a highly configurable and comprehensive reporting interface that offers deep observability into an organization's network.

3.1.1 Life cycle and global settings

Plexer Scrutinizer automatically manages Alarm and Event data based on the following life cycle:

1. Plexer Scrutinizer continuously monitors its environment for observations of system activity or network traffic that match preconfigured criteria.
2. Observations are aggregated and reported/managed as an Event based on the Alarm Policy associated with the identified criteria.

3. The details of the Event are reviewed under the corresponding Alarm Policy via the Alarm Monitor interface.
4. After investigation and/or resolution, the Event is flagged as *acknowledged* by a user to clear it from all Alarm Monitor views.

Event data remains accessible for further review following the configured retention settings.

Global retention settings

The following global settings in the [Admin > Settings > Data History](#) tray can be used to change how Alarm Event data is managed:

| | |
|--------------------------------|--|
| Alarm Retention Days | Sets the maximum number of days Alarm and Event data is retained before being deleted from the system |
| Alarm Retention Size | Sets the maximum amount of disk space that can be used for Alarm and Event data storage |
| Auto-Acknowledge Alarms | Sets the number of days before Events are automatically tagged as <i>Acknowledged</i> (Can also be configured as a Notification Profile action) |

Note: The Alarm retention settings control automatic data deletion for both acknowledged and un-acknowledged Events.

3.1.2 Alarm Policy settings

Individual Alarm Policy settings allow granular customization of what, when, and how Alarms and Events are reported.

The following settings can be accessed from the [Admin > Alarm Monitor > Alarm Policies](#) view:

Status

Sets the Alarm Policy to one of three states:

| Setting | Generates Events | Alarm Monitor | Stored in Database | Notifications by Profile(s) |
|-----------------|------------------|---------------|--------------------|-----------------------------|
| Active | Yes | Yes | Yes | Yes |
| Store | Yes | No | Yes | Yes |
| Inactive | No | No | No | No |

Weight

Assigns each Event generated by the Alarm Policy a numerical weight that is used to calculate the severity reported in the Alarm Monitor views

Event timeout

Sets the number of seconds the system will wait when aggregating observations meeting the same criteria as a single Event

Hint: Setting nonessential Alarm Policies to *Store* or *Inactive* can filter out Events that do not require visibility. This can reduce the number of Alarms being reported (and stored) in the Alarm Monitor views.

3.1.3 Alarm notifications

Alarms/Events in Plixer Scrutinizer can also be configured to trigger one or more notification actions when they are generated/observed.

Notification Profiles

Notification actions are assigned to individual Alarm Policies by way of *Notification Profiles*, each of which can be configured with one or more *actions*.

Note: An Alarm Policy can only be assigned one Notification Profile at a time.

Hint: Notification Profiles can be used in conjunction with the *Store* Alarm Policy status to acknowledge, forward, and/or store the details of an Event without them being reported in the Alarm Monitor views.

3.1.4 Flow Analytics

Plixer Scrutinizer uses a collection of Flow Analytics (FA) algorithms to monitor collected flow data for specific traffic patterns and/or behavior typically associated with threats to a network.

Because FA algorithms rely on associated Alarm Policies for reporting, the *initial configuration and regular tuning of FA-based functions* are integral to optimizing Alarms and Events.

For additional information, see the *Flow Analytics configuration guide*.

3.1.5 Optimizing Alarms and Events

When correctly configured, the Plixer Scrutinizer Alarm Monitor is capable of reporting information that is accurate, relevant, and uniquely tailored to the organization or team using it.

To achieve this, the following configuration steps related to Alarms and Events should be completed as part of deploying Plixer Scrutinizer.

1. Navigate to the **Admin > Settings > Data History** tray and adjust the *Alarm Retention Days*, *Alarm Retention Size*, and *Auto-Acknowledge Alarms* values as needed.
2. In the **Admin > Settings > Alarm Notifications** tray, verify that the Alarm Notifications options are correctly configured.
3. Go to the *Admin > Alarm Monitor > Notification Profiles* page and create *Notification Profiles* to enable additional notification channels.
4. Go to the *Admin > Alarm Monitor > Alarm Policies* page and:
 - Set the status of any Alarm Policies that are unnecessary or irrelevant to the environment to *Inactive* (must be done as a bulk action after selecting at least one Policy).
 - Set the status of Alarm Policies whose Events should be monitored but not reported in the Alarm Monitor to *Store* (must be done as bulk action after selecting at least one Policy).
 - Assign the appropriate Notification Profiles to any Alarm Policies that require them.

Note: The *Timeout* and *Weight* values of an Alarm Policy can be adjusted at a later time, after evaluating reporting behavior for Events under it.

5. Follow the *Flow Analytics configuration guide* to correctly set up FA-based functions and features.
6. Follow the *Plixer ML Engine configuration guide* to correctly set up machine-learning-based functions and features.

After the initial setup has been completed, it is highly recommended to continue to evaluate Alarm and Event reporting behavior and make further adjustments to the various elements' configurations as necessary.

3.2 Configuration checklist

The following checklist outlines the recommended order of configuration steps to fully set up a Plixer Scrutinizer deployment:

Note: Click on a checklist item for additional information and detailed instructions related to that configuration step.

| Configuration step | Function/Benefit |
|-----------------------------|--|
| Deploy instance | Deploy the Plixer Scrutinizer hardware/physical or virtual appliance in your environment. |
| Appliance setup | Run the setup questionnaire from the appliance terminal to configure an IP address, DNS hostname, NTP server, and HTTPS certificate. |
| Send flows | Configure exporters/network devices to send flows to Plixer Scrutinizer (or a Plixer Replicator, if applicable) |
| Email server | Configure an SMTP server to enable email notifications for Alarms/events and on-demand/scheduled email reports. |
| SNMP credentials | Configure SNMP credentials to enable importing of exporter names, interface names, and interface speeds. |
| Users | Create additional accounts/logins to customize settings and preferences for individual users. |
| Usergroups | Create usergroups to manage permissions and access levels for different users. |
| Defined applications | Define rules for applications that are unique to your network to enhance reporting, filtering, and other functions. |
| IP Groups | Assign resources specific to your organization to IP Groups, which are used in reporting, filtering, and managing analytics inclusions/exclusions. |
| External authentication | Improve your security posture and simplify user management by leveraging AD, LDAP, SSO, Radius, and/or TACACS for user authentication. |
| Data history settings | Modify historical data retention settings to support your organization's archival and forensics needs. |
| Security Groups | Populate the default Security Groups (Firewalls, Core Exporters, Edge Exporters, Defender Probes) to automatically add devices to recommended Flow Analytics algorithms. |
| Exclusion IP Groups | Verify that the DNS Servers, Public WiFi, Network Scanners, and SNMP Pollers IP Groups are populated to automatically add recommended exclusions to Flow Analytics algorithms. |
| Flow Analytics inclusions | Define additional inclusions and exclusions (including custom Security Groups and IP Groups) for individual Flow Analytics algorithms. |
| Device/Mapping Groups | Organize exporters into groups to quickly find data sources, enable group report filters, and generate customizable network maps. |
| Dashboards | Create/customize one or more dashboards to consolidate frequently used views and drive workflows through the Plixer One Platform. |
| Saved reports | Create saved reports to quickly re-run the same report configuration with a single click. |
| Report thresholds | Add thresholds to saved reports to proactively watch for specific traffic/behaviors and trigger Alarms (and Notification Profiles) when the specified conditions are met. |
| Scheduled email reports | Set up scheduled email reports to automatically re-run and send important reports as emails to any inbox. |
| Notification Profiles | Create Notification Profiles that can be assigned to Alarm Policies to automatically send emails, forward details to your SIEM, or run custom scripts to do absolutely anything. |
| Deploy the Plixer ML Engine | Deploy and set up the Plixer ML Engine to enable advanced features, including network behavior modeling and anomaly detection, in Plixer One Platform Security. |
| ML inclusions | Review, modify, or add inclusions/sources for ingestion and modeling by the Plixer ML Engine. |

3.3 Distributed environments

Multiple Plixer Scrutinizer appliances/servers can be configured as a distributed environment with a central, primary Reporter and one or more remote Collectors.

Distributed environments are capable of ingesting significantly higher flow volumes from a greater number of Exporters. All admin, management, and reporting functions are handled from the primary Reporter.

3.3.1 Distributed cluster setup

Distributed clusters can include any combination of hardware and/or virtual appliances, regardless of physical location.

To set up a distributed cluster, follow these steps:

1. Deploy the Plixer Scrutinizer hardware or virtual appliances following the [deployment guides](#) in this documentation.

Important: To avoid potential bottlenecks in distributed configurations that include hardware appliances, 10 Gb networking is strongly recommended. If the appliances are geographically dispersed, 10G is only required if the WAN link can support it.

2. Start an SSH session as the `plixer` user with the appliance that should act as the primary Reporter.
3. Launch the [scrut_util](#) interactive CLI by running:

```
/home/plixer/scrutinizer/bin/scrut_util
```

4. At the `SCRUTINIZER>` prompt, designate the current appliance as the primary Reporter by entering:

```
SCRUTINIZER> set selfreporter
```

5. From the primary Reporter, register each additional appliance as a remote Collector:

```
SCRUTINIZER> set registercollector [collector_appliance_ip] <secondary>
```

Hint: The optional `secondary` flag is used to promote the remote Collector to *secondary Reporter status*.

6. After registering all remote Collectors, use the `exit` command to exit the `scrut_util` interactive CLI.

Once the Plixer Scrutinizer distributed cluster has been set up, Exporters can be configured to send flows to any of the remote Collectors.

To access the web interface for the cluster, connect using the IP address of the primary Reporter.

Ports used

If appliances in a distributed cluster are unable to communicate with each other, it may be necessary to whitelist the connections between the remote Collectors and the primary Reporter.

The following network ports are used in communications between appliances in a distributed environment:

| Collector(s) -> Reporter (UDP) | Collector(s) <-> Reporter (TCP) |
|--------------------------------|------------------------------------|
| 514 | 22 80 (or 443) 6432 and 5432 |

Important: To learn more about licensing options for distributed environments or for additional assistance, contact [Plixer Technical Support](#).

3.3.2 High availability

Plixer Scrutinizer distributed clusters support high availability (HA) configurations that include secondary Reporters and/or backup Collectors for redundancy.

Note: Contact [Plixer Technical Support](#) to learn more about HA licensing options.

Secondary Reporters

In distributed deployments, a remote Collector can be registered as a secondary Reporter, which can be used to access the system if the primary Reporter becomes unavailable.

To register a remote Collector as a secondary Reporter, enter the following `scrut_util` command from the primary Reporter.

```
SCRUTINIZER> set registercollector [collector_appliance_ip] secondary
```

After a Collector has been registered as a secondary Reporter, its IP address can be used to access a read-only version of the Plixer Scrutinizer web interface at any time. An updated backup of the primary Reporter's configuration metadata will also be maintained on that Collector.

If the primary Reporter has become permanently unavailable, the secondary Reporter should be promoted using the `set selfreporter scrut_util` command, as outlined in the [distributed environment setup guide](#). This will lift the read-only status and restore full web interface functionality.

Note: When promoting a secondary Reporter to primary status, a new license key is required to complete the process.

Backup Collectors

Distributed clusters can be configured to use backup Collectors to enable high availability for flow collection functions.

To use a remote Collector **Y** as a backup for remote Collector **A**

1. Configure all Exporters sending flows to **A** to also send flows to **Y**.
2. In the web interface, navigate to **Admin > Resources > Manage Exporters** and verify that the selected Exporters are correctly sending flows to both Collectors.
3. From the **Manage Exporters** view, set the status of the duplicated Exporters sending flows to **Y** to *Backup*.

If remote Collector **A** becomes unavailable, the Exporters that were previously set to *Backup* on remote Collector **Y** must be set to *Enabled* to allow for continuous flow collection and reporting. Once **A** is online again, the status of the Exporters should be reverted to *Backup*.

Hint: When managing a large number of Exporters, filter the list to view only relevant Exporters and use the checkboxes to set them to *Backup* or *Enabled* as a bulk action.

HA with Plixer Replicator

[Plixer Replicator](#) can simplify the process of setting up backup Collectors by replicating flows data and forwarding it to multiple destination Collectors.

View the Plixer Replicator online documentation or contact [Plixer Technical Support](#) to learn more.

3.4 Flow Analytics

Plixer Scrutinizer includes a library of Flow Analytics (FA) algorithms, which are applied to all incoming flow data. This allows the system to provide additional traffic-based insights and report activity typically associated with threats to a network.

This section outlines the recommended procedure(s) for the initial configuration of Flow Analytics and includes additional guides and references to assist with the optimization of related functions and features.

3.4.1 Configuring Flow Analytics

To enable FA-based functions, several configuration steps must be completed after Plixer Scrutinizer has been deployed and set up.

This process helps ensure that Plixer Scrutinizer is fully adapted to an organization's NDR requirements.

Enabling/disabling algorithms

Because Plixer Scrutinizer is designed to support the full spectrum of enterprise network types, it may include FA algorithms that do not apply to a specific environment. This will be based on the devices and elements present on the network, the types of flow data available, and/or organizational IT policies.

As part of optimizing the system's monitoring and reporting functions, all unnecessary FA algorithms should be disabled. This includes algorithms that:

- Only benefit devices or elements that are not present on the network
- Require flow data that is not being sent by devices on the network
- Target traffic or patterns that are made irrelevant by the organization's IT policies

The [Admin > Alarm Monitor > Flow Analytics Configuration](#) page lists *all FA algorithms* and shows whether they are currently enabled (default) or disabled.

Hint: Hover over the information (i) icon next to an algorithm to view additional information about it.

Disabling FA algorithms

To disable an algorithm, click on it to open the configuration tray and use the toggle. The algorithm can also be re-enabled this way at any time.

Multiple algorithms can also be disabled or enabled as a bulk action when one or more algorithms are selected.

Adding Exporters

Plixer Scrutinizer selectively applies Flow Analytics to incoming flow data, based on the Exporters defined for each algorithm.

To activate the system's FA-based functions, Exporters must first be added to the enabled algorithms.

Security Groups

Plixer Scrutinizer Security Groups are user-defined groups of Exporters to which the same set of FA algorithms are applied. Security Groups allow the Exporter lists for all FA algorithms to be fully populated without the need to manually configure individual algorithms. Exporters can be added to Security Groups via the [Admin > Alarm Monitor > Security Groups](#) page.

Hint: The default *Firewalls*, *Core Exporters*, *Edge Exporters*, and *Defender Probes* Security Groups are configured with FA algorithms based on the recommended Exporter assignments.

If Flow Analytics is being configured for the first time, Exporters should be added to the *Core Exporters* and *Edge Exporters* a few at a time. This will limit the volume of Alarms that may need to be checked when [testing Flow Analytics settings](#) via the Alarm Monitor.

The **Security Groups** view also allows new groups to be added and the settings for existing groups to be modified.

Adding Exporters individually

For more granular control over Exporter-to-algorithm assignment, Exporters can also be added to FA algorithms via the configuration tray of the [Admin > Alarm Monitor > Flow Analytics Configuration](#) page.

Hint: Exporters can also be added to multiple algorithms as a bulk action when one or more algorithms are selected.

Because Alarm-generating algorithms will only be triggered when the target is an internal address, public IP addresses must be defined as part of an IP Group for them to be considered part of the protected network. For internal-to-internal and internal-to-external monitoring, core routers should be added to the relevant algorithms. For monitoring public assets, the edge routers of the relevant IP Groups should be added to the algorithms.

Defining exclusions

To avoid unnecessary Alarms and excessive processing load on the system, certain devices or traffic should be excluded from monitoring by specific FA algorithms.

Plixer Scrutinizer's factory configuration includes four *IP Groups* that are defined as exclusions under the appropriate algorithms:

- DNS servers
- Public WiFi
- Network Scanners
- SNMP Pollers

These IP Groups should be populated with the correct Exporters to optimize Flow Analytics monitoring and reporting.

Adding exclusions to an FA algorithm

FA algorithms can also be configured with additional exclusions beyond those defined under the above-mentioned IP Groups. This is done via the algorithm's configuration tray from the [Admin > Alarm Monitor > Flow Analytics Configuration](#) page.

Exclusions can be defined by IP address, IP range, subnet, domain (via reverse DNS), or IP Group.

Hint: The default IP Group exclusions for an algorithm are also displayed under the *Exclusions* section of the configuration tray.

Additional options

FA-based functions and features in Plixer Scrutinizer can be further tuned and customized using these additional options.

Global and algorithm settings

The following global and algorithm settings can be used to modify the behavior of Flow Analytics in Plixer Scrutinizer:

Hint: Global FA settings can be changed in the **Admin > Settings > Flow Analytics** tray.

| Setting | Scope | Description |
|------------------------------------|-----------|--|
| Auto-Enable Defender | Global | When checked, allows Flow-Pro Defender to be automatically enabled for supported algorithms |
| Jitter by Interface | Global | Sets the variation in packet delay due to queueing, contention, and/or serialization (Default: 80 ms); Also used for record highlighting in <i>Status</i> reports |
| Latency | Global | Sets the latency value used for record highlighting in <i>Status</i> reports (Default: 75 ms) |
| Share Violations | Global | When checked, allows the system to share details of cyber attacks coming from Internet IP addresses with the Plexier Security Team (May require firewall permissions); This information is used to further improve the global host reputation list. No internal addresses will be shared. |
| Top Algorithm Devices | Global | Controls whether <i>Top X</i> FA algorithms are applied to all Exporters or need to be configured individually |
| Thresholds | Algorithm | Increases or decreases the tolerance of Alarm-generating FA algorithms to the corresponding behavior or traffic; This setting should be adjusted if too many false positives are being reported under an algorithm's associated Alarm |
| 44 | | Policy |
| Algorithm-specific settings | Algorithm | Additional configuration |

Custom reputation lists

The *Host Reputation* FA algorithm is capable of using custom lists in conjunction with Plixer Scrutinizer's default host reputation lists. When a host in any reputation list becomes the target of traffic, the Event is reported under the Host Reputation Alarm Policy.

To import a list of IP addresses as a custom host reputation list, follow these steps:

1. Add the hosts to a file, using one line for each IP address.

Example:

```
10.1.1.1
10.1.1.2
10.1.1.3
```

2. Save the file with a `.import` extension. (e.g., `custom_threats.import`)

Important: The name of the file will be used for artifacts involving the included hosts on the [Alarm Summary page](#).

3. Move the file to the `\scrutinizer\files\threats\` directory.

The file is imported hourly, at the same time that threat lists are updated.

Hint: To manually run the file import operation, use the command `scrut_util --downloadhostreputationlists`.

3.4.2 Reporting options

Each Alarm-generating FA algorithm is associated with one or more Alarm Policies, under which anomalies and other insights are reported via the Plixer Scrutinizer Alarm Monitor. The [settings for these Alarm Policies](#) can also be modified to change the reporting behavior for the individual algorithms.

To learn more about Alarm Policies and the Plixer Scrutinizer Alarm Monitor, see the [Alarms and Events](#) section of this manual.

Notification Profiles

To forward the details of Alarms and Events reported by an FA algorithm to one or more users or external systems, at least one Notification Profile must be created and assigned to the corresponding Alarm Policy.

To learn more about Notification Profiles, see the [Alarm Notifications](#) section.

FA Dashboard Gadgets

Certain Gadgets that can be added to [Plixer Scrutinizer Dashboards](#) rely on one or more FA algorithms for the data they report.

These Gadgets require no further configuration and can be added to any Dashboard as long as the corresponding algorithms have been enabled and correctly configured.

Hint: The **Flow Analytics Summary** Gadget can be used to troubleshoot algorithm configurations. If there are algorithms that are taking longer than 5 minutes to run, check that the correct Exporters have been added.

To learn more about Dashboards and Gadgets, see the [Dashboards](#) topic of this documentation.

3.4.3 Testing and tuning

To ensure that Flow Analytics is properly configured, testing the various definitions, settings, and enabled features is strongly recommended. This can be accomplished by checking what Alarms and Events are being reported in the [Alarm Monitor views](#).

When setting up Flow Analytics for the first time, the following process is recommended:

1. Navigate to **Admin > Definitions > IP Groups** and populate the *DNS Servers*, *Public WiFi*, *Network Scanners*, and *SNMP Pollers* groups to define basic exclusions for FA algorithms.
2. Review the list of FA algorithms in the **Admin > Alarm Monitor > Flow Analytics Configuration** and disable any algorithms that are irrelevant.
3. Define additional exclusions for individual algorithms in their configuration trays as needed.
4. Navigate to **Admin > Alarm Monitor > Security Groups** and add several Exporters each to the *Core Exporters* and *Edge Exporters* Security Groups.

Once the first batch of Exporters has been added, review the Alarm Monitor views to verify that Alarms and Events are being reported correctly. Afterwards, repeat Step 4 of the process and continue checking Alarms and Events until all Exporters have been added to Security Groups.

Note: If there are continuous or unnecessary Alarms or Events being reported, it may also be necessary to define additional exclusions for certain algorithms.

Further tuning

After the initial setup and testing have been completed, the Flow Analytics configuration can be further modified to adapt to changes in the Plexier Scrutinizer environment or fine-tune performance.

Hint: For more efficient reporting and/or analysis, [create one or more Notification Profiles](#) and associate them with the appropriate Alarm Policies.

The [additional options](#) related to Flow Analytics should also be reviewed and edited as necessary.

3.4.4 List of FA algorithms

Plexier Scrutinizer's library of FA algorithms is continuously being updated to maximize reporting accuracy and expand support for varied enterprise NDR scenarios.

The following table lists all available FA algorithms, along with their functions and recommended applications:

| Algorithm | Function |
|-----------------------------------|--|
| Bogon Traffic | Generates an Alarm if traffic to or from an unallocated public IP space is detected. |
| BotNet Detection | Generates an Alarm when the number of failed unique DNS name lookups exceeds a threshold. |
| Breach Attempt Detection | Generates an Alarm when behavior that may indicate a brute force password attack is detected. |
| DDoS Detection | Generates an Alarm when a Distributed Denial of Service (DDoS) attack is detected. |
| Denied Flows Firewall | Generates an Alarm when the number of denied flows from an internal to an external IP address exceeds a threshold. |
| DNS Command and Control Detection | Monitors DNS TXT communications at the network perimeter and generates an Alarm when a threshold is exceeded. |
| DNS Data Leak Detection | Monitors DNS lookup messages that may contain encoded data and generates an Alarm when a threshold is exceeded. |
| DNS Hits | Generates an Alarm when a host initiates an excessive number of DNS queries. |
| DNS Server Detection | Monitors packet exchanges between clients and servers and generates an Alarm when a threshold is exceeded. |

| Algorithm | Function |
|------------------------------|--|
| Domain Reputation | Monitors the network for traffic from new domains and generates an Alarm when |
| DRDoS Detection | Generates an Alarm when a Distributed Reflection Denial of Service attack tar |
| FIN Scan | Generates an Alarm when a FIN scan is detected |
| Flow Reports Thresholds | Monitors the network for behavior exceeding any <i>thresholds configured in save</i> |
| Host Indexing | Maintains an index of hosts seen on the network that includes additional detail |
| Host Reputation | Maintains a list of active, non-whitelisted Tor nodes for monitoring |
| Host Watchlist | Monitors IP addresses to identify hosts violating a user-defined blacklist |
| ICMP Destination Unreachable | Generates an Alarm when a large number of <i>ICMP Destination Unreachable</i> m |
| ICMP Port Unreachable | Generates an Alarm when a large number of <i>ICMP Port Unreachable</i> message |
| Incident Correlation | Escalates and consolidates multiple Indicator of Compromise (IOC) Events for |
| IP Address Violators | Generates an Alarm when a flow containing a non-authorized IP address as the |
| JA3 Fingerprinting | Checks TLS handshake data against a list of known signatures and generates an |
| Large Ping | Generates an Alarm when an unusually large ICMP Echo Request (ping) is obs |
| Medianet Jitter Violations | Generates an Alarm when jitter values reported by a Medianet flow exceed the |
| Multicast Violations | Generates an Alarm when multicast traffic volume exceeds the configured thre |
| NetFlow Domain Reputation | Generates an Alarm when a DNS lookup from a blacklisted IP is reported via |
| Network Transports | Monitors traffic by transport layer protocol across all flows from configured Ex |
| NULL Scan | Generates an Alarm when a NULL scan is detected |
| Odd TCP Flags Scan | Generates an Alarm when a scan using unusual TCP flag combinations is detec |
| P2P Detection | Monitors flows for P2P traffic and generates an Alarm when a session whose n |
| Packet Flood | Generates an Alarm when a packet flood is detected |
| Persistent Flow Risk | Generates an Alarm when a persistent flow is detected |
| Persistent Flow Risk - ASA | Generates an Alarm when a persistent flow matching a specified 5-tuple is dete |
| Ping Flood | Generates an Alarm when a ping flood is detected |
| Ping Scan | Generates an Alarm when a host suspected of performing a ping scan is observ |
| Protocol Misdirection | Generates an Alarm when traffic not matching the port being used is detected |
| Reverse SSH Shell | Generates an Alarm when potential reverse SSH tunnels to external destination |
| RST/ACK Detection | Generates an Alarm when the system observes a large number of TCP flows co |
| Slow Port Scan | Generates an Alarm when the system observes a large number of ports on the s |
| Source Equals Destination | Generates an Alarm when traffic with the same host and destination is observe |
| SYN Scan | Generates an Alarm when a SYN scan is detected |
| TCP Scan | Generates an Alarm when a potential TCP scan is detected from an Exporter th |
| Top Applications | Monitors application traffic across all flows from configured Exporters |
| Top Autonomous Systems | Monitors traffic to and from autonomous systems across all flows from configu |
| Top Countries | Monitors traffic by country across all flows from configured Exporters |
| Top Hosts | Monitors traffic by host across all flows from configured Exporters |
| Top IP Groups | Monitors traffic by IP Group across all flows from configured Exporters |
| UDP Scan | Generates an Alarm when a potential UDP scan is detected |
| Worm Attack | Generates an Alarm when behavior indicating a potential worm is observed fro |

| Algorithm | Function |
|------------------|--|
| Worm Propagation | Generates an Alarm when successful worm replication across hosts is detected |
| XMAS Scan | Generates an Alarm when a XMAS scan is detected |

3.5 Device groups

Plixer Scrutinizer is able to partition its environment into smaller sets of devices and other objects using several different grouping schemes. These groups are used in various monitoring, reporting, and configuration functions and, when correctly set up, enable more fluid and efficient workflows.

IP Groups

IP Groups are user-defined groups of devices that share certain characteristics, such as device type, ownership/department, and geolocation, for the purpose of monitoring and/or reporting. They are defined as lists of IP addresses, IP address ranges, and/or full subnets and can be configured from the [Admin > Definitions > IP Groups](#) page.

Note: The Plixer Scrutinizer factory configuration includes default IP Groups that are used by certain functions (e.g., [Flow Analytics exclusions](#)) and should be populated as part of tailoring the system to the environment.

Mapping Groups

Mapping Groups are user-defined groups of devices that are used to generate Network Maps and allow for customizable visualization of network topology up to the interface level (not including end devices).

Network Maps can be created and viewed from the [Monitor > Network Maps](#) page of the Plixer Scrutinizer web interface, while additional configuration and management options for Mapping Groups and Objects can be accessed via their respective pages under the [Admin > Settings](#).

Security Groups

Security Groups are user-defined Exporter groups that streamline the process of [enabling the appropriate Flow Analytics algorithms for different Exporter types](#). They can be managed from the [Admin > Alarm Monitor > Security Groups](#) page.

After a Security Group has been created/populated, it can be added to one or more FA algorithms from the [Admin > Alarm Monitor > Alarm Policies](#) page. This will enable the algorithm(s) for all Exporters in the group. Alternatively, algorithms can be assigned to Security Groups from the Security Groups admin page instead.

3.6 Plixer ML Engine

The Plixer ML Engine enables advanced anomaly and threat detection in Plixer Scrutinizer by applying multiple AI-, ML-, and deep-learning-based techniques.

Once *deployed* and configured, the engine ingests flow data through Plixer Scrutinizer and uses it to build behavior models that represent typical, legitimate activity. Using these models as baselines, anomalous behavior and other insights can then be reported as *Alarms and Events*.

Note: To learn more about Plixer ML Engine licensing options, contact [Plixer Technical Support](#).

This configuration guide covers the settings and steps involved in further tuning the Plixer ML Engine. Additional background and recommendations related to its functions and features are also included.

3.6.1 Managing inclusions

The Plixer ML Engine models network behavior using flow data from hosts, subnets, and/or Exporters that have been defined as *inclusions* or *sources*. By default, the 20 hosts that best match the currently enabled *dimensions* are automatically added as inclusions.

To further adapt network models to an organization's unique requirements, inclusions can be added, removed, or reconfigured from the [Admin > Alarm Monitor > Manage ML Inclusions](#) page of the Plixer Scrutinizer web interface.

Configuring inclusions

When defining an inclusion for the Plixer ML Engine, the following settings must be configured:

- Network address of the host, subnet, or Exporter to be added
- *Sensitivity* value (low, medium, or high), which controls how much observed behavior must deviate from expected traffic/activity patterns for it to be considered anomalous or suspicious; Lowering the sensitivity will result in even small deviations from learned activity patterns being reported as Alarms but also increases the risk of false positives.
- The *Malware Detections* setting enables or disables the use of pre-trained classification models to recognize and report network activity associated with common malware classes.

- The *Enabled* toggle is used to enable or disable the host, subnet, or Exporter as an inclusion. Inclusions can be added in the disabled state and enabled at a later time.

Hint: The sensitivity setting for an inclusion can be left at its default value and later changed following a 7-day period of observation (recommended).

After the Plixer ML Engine has been deployed, it is highly recommended to review the **Manage ML Inclusions** page to verify that the pre-added inclusions are the best suited for modeling typical or atypical activity in the current environment. Inclusions should be added or removed as necessary.

3.6.2 Managing dimensions

To ensure that only relevant network traffic data is used when modeling network behavior, the Plixer ML Engine only monitors communications/protocols that have been defined as *dimensions*. Once deployed, the engine defaults to a list of factory-configured dimensions that have been selected to suit most common enterprise environments.

To further adapt network models to an organization's unique requirements, dimensions can be added, removed, or reconfigured from the [Admin > Alarm Monitor > Manage ML Dimensions](#) page of the Plixer Scrutinizer web interface.

Configuring Dimensions

When defining a dimension for the Plixer ML Engine, the following settings must be configured:

- Inclusion type the dimension applies to (hosts/subnets or Exporters)
- Template field to use for grouping (sourceipaddress or destinationipaddress, host dimensions only)
- Aggregation method to use (octetdeltacount or packetdeltacount)
- Port to monitor for the dimension
- Range of communications to monitor (*internal only* or *all*)
- The *Enabled* toggle is used to enable or disable monitoring of the dimension. Dimensions can be added in a disabled state and enabled at a later time.

After deploying the engine, it is highly recommended to review the **Manage ML Dimensions** page to verify that the default list includes the dimensions that are best suited for modeling typical or atypical activity in the current environment. Dimensions should be added or removed as necessary.

3.6.3 About the Plixer ML Engine

Unlike conventional security solutions that alert users only after indications of a breach are discovered, the Plixer ML Engine is designed to actively monitor network activity and alert its users to potential threats in real time.

The Plixer ML Engine relies on several key functions to enable intelligent, multi-layered threat detection in Plixer Scrutinizer:

Comprehensive network behavior modeling

The Plixer ML Engine is capable of ingesting large volumes flow data through Plixer Scrutinizer, as defined by the configured *inclusions* and *dimensions*. This data can then be used to model network behavior at any given time.

Over time, the engine is able to identify typical activity patterns in these models and recognize deviations, such as data accumulation/exfiltration, tunneling, and lateral movement, that may indicate an attack on the network.

Highly configurable ML modeling

To support a wider range of enterprise network scenarios, the Plixer ML Engine supports user-defined inclusions and dimensions. This allows for behavior modeling based on the most relevant hosts and traffic in a given environment.

When building its behavior models, the engine takes into account all characteristics that can make an environment unique, such as legitimate traffic patterns, host/group importance, and seasonality.

ML-based malware detection

Plixer ML Engine uses pre-trained classification models to recognize generic network activity patterns that are associated with common classes of malware, including command and control, remote access trojans, and exploit kits.

When enabled, this provides another layer of protection that further reduces risk and mean time to resolution (MTTR) when threats are detected.

Continuous observation and learning

As it continues to ingest flow data, the Plixer ML Engine updates its behavior models based on a schedule that defines weekdays, weeknights, and weekends.

This allows the engine to not only account for changes in legitimate activity patterns but also recognize more sophisticated threats that attempt to disguise their behavior as normal activity.

Tuning recommendations

The Plixer ML Engine ships with a factory configuration that will allow it to function in common environments out of the box. However, its detection and reporting functions can be further optimized by tailoring its configuration to the environment it has been deployed to.

To tune the engine to report more accurate and relevant Alarms and Events, the following steps are recommended after its deployment:

- Review the **Admin > Alarm Monitor > Manage ML Inclusions** page to verify that the automatically selected hosts are the best suited for modeling the overall state of the environment.
- Review the **Admin > Alarm Monitor > Manage ML Dimensions** page to verify that the default list of dimensions sufficiently covers the types of traffic expected in the environment.
- When adding new inclusions or dimensions, leave the sensitivity setting at its default value and closely monitor all Alarms/events being reported by the engine for a period of at least 7 days. If too many anomalies/deviations are being reported, the sensitivity can be increased to improve accuracy.

3.7 Environment sizing

To ensure consistently optimal performance and continuous availability, Plixer Scrutinizer must be provisioned based on the functions and/or features required by its users.

This section outlines the recommended procedures for calculating the appropriate resource allotments for Plixer Scrutinizer deployments.

Note: Certain steps in these guides require access to the Plixer Scrutinizer web interface. For more accurate results, complete the *initial setup wizard* beforehand.

3.7.1 CPU/RAM

Follow the steps described in this section to calculate the total number of CPU cores and amount of RAM that should be allocated to a Plexier Scrutinizer deployment.

Note: For additional guidelines related to distributed clusters, see [this section](#).

1. Use the recommendations in the table below as starting CPU core count and RAM values. These allocations cover Plexier Scrutinizer's core functions (flow collection, reporting, basic [Alarm Policies](#)) for the expected flow rates and Exporter counts indicated.

CPU cores and RAM based on flow rate and exporter count

| | | Exporters | | | | | | | |
|-------------------------|------|--|--|--|--|--|--|--|--|
| | | 5 | 25 | 50 | 100 | 200 | 300 | 400 | 500 |
| Flows/s | 5k | 8 CPU cores 16 GB RAM | 8 CPU cores 16 GB RAM | 10 CPU cores 20 GB RAM | 14 CPU cores 28 GB RAM | 20 CPU cores 39 GB RAM | 26 CPU cores 52 GB RAM | 32 CPU cores 67 GB RAM | 38 CPU cores 82 GB RAM |
| | 10k | 8 CPU cores 16 GB RAM | 8 CPU cores 16 GB RAM | 12 CPU cores 24 GB RAM | 18 CPU cores 36 GB RAM | 25 CPU cores 50 GB RAM | 32 CPU cores 65 GB RAM | 38 CPU cores 81 GB RAM | 43 CPU cores 97 GB RAM |
| | 20k | 16 CPU cores 32 GB RAM | 16 CPU cores 32 GB RAM | 16 CPU cores 32 GB RAM | 24 CPU cores 48 GB RAM | 32 CPU cores 64 GB RAM | 38 CPU cores 80 GB RAM | 43 CPU cores 96 GB RAM | 48 CPU cores 112 GB RAM |
| | 50k | 32 CPU cores 64 GB RAM | 32 CPU cores 64 GB RAM | 32 CPU cores 64 GB RAM | 32 CPU cores 64 GB RAM | 39 CPU cores 80 GB RAM | 44 CPU cores 96 GB RAM | 48 CPU cores 112 GB RAM | 52 CPU cores 128 GB RAM |
| | 75k | 46 CPU cores 96 GB RAM | 46 CPU cores 96 GB RAM | 46 CPU cores 96 GB RAM | 46 CPU cores 96 GB RAM | 46 CPU cores 96 GB RAM | 49 CPU cores 112 GB RAM | 52 CPU cores 128 GB RAM | 55 CPU cores 144 GB RAM |
| | 100k | | | | | | | | |
| 3.7. Environment sizing | | 52 CPU cores 128 GB RAM | 52 CPU cores 128 GB RAM | 52 CPU cores 128 GB RAM | 52 CPU cores 128 GB RAM | 52 CPU cores 128 GB RAM | 52 CPU cores 128 GB RAM | 55 CPU cores 144 GB RAM | 58 CPU cores 160 GB RAM |

2. Following the table below, compute for the total expected CPU and RAM usage for all feature sets that will be enabled.

| Feature | CPU (cores) | RAM (GB) | FA Algorithms |
|---|-------------|----------|---|
| Streaming (to a Plixer ML Engine or external data lake) | 1 | 0.4 | N/A |
| Data Transfer | 1 | 0.4 | <ul style="list-style-type: none"> • <i>Network Transports</i> • <i>Protocol Misdirection</i> • <i>P2P Detection</i> • <i>Reverse SSH Shell</i> |
| Worm Activity | 0.5 | 0.2 | <ul style="list-style-type: none"> • <i>Lateral Movement Attempt</i> • <i>Lateral Movement</i> |
| FlowPro | 1.25 | 0.5 | <ul style="list-style-type: none"> • <i>BotNet Detection</i> • <i>DNS Data Leak Detection</i> • <i>DNS Command and Control Detection</i> • <i>DNS Server Detection</i> • <i>JA3 Fingerprinting</i> |
| Scanning | 4.25 | 1.7 | <ul style="list-style-type: none"> • <i>ICMP Destination Unreachable</i> • <i>Breach Attempt Detection</i> • <i>Medianet Jitter Violations</i> • <i>Denied Flows Firewall</i> • <i>TCP Scan</i> • <i>Bogon Traffic</i> • <i>Large Ping</i> |
| 3.7. Environment sizing | | | <ul style="list-style-type: none"> • <i>Slow Port Scan</i> • <i>ICMP Port Unreachable</i> • <i>FIN Scan</i> • <i>NULL Scan</i> |

Hint: Each FA algorithm reports detections using one or more Alarm Policies, which are also enabled/disabled as part of the feature set. Policy-to-algorithm associations can be viewed in the [Admin > Alarm Monitor > Alarm Policies](#) view.

Note: The CPU and RAM allocations per feature are recommended for deployments with up to 500 exporters and a total flow rate of 150,000 flows/s.

3. Combine the values obtained from steps 1 and 2, and apply any necessary adjustments to the CPU and RAM allocations for the Plixer Scrutinizer appliance.
4. In the web interface, navigate to [Admin > Resources > System Performance](#) and verify that the correct CPU core count and RAM amount are displayed for the collector.
5. After confirming that CPU and RAM allocations have been correctly applied, go to [Admin > Resources > Feature Resources](#) and enable/disable features according to the selections made for step 2.

Once Plixer Scrutinizer is *fully configured and running*, CPU and RAM utilization can be monitored from the **Admin > Resources > System Performance** page using the *CPU Utilization* and *Available Memory* graphs. These graphs should be reviewed regularly (in addition to after resources are initially allocated), so that any necessary adjustments can be made.

Important: After making any adjustments to the Plixer Scrutinizer's resource allocations, [launch *scrut_util*](#) as the root user and run the [set tuning](#) command to re-tune the appliance.

Alarm Policies under the *System* category are also used to report events related to resource utilization (e.g. collection paused/resumed, feature set paused/resumed, etc.)

Additional factors

In addition to the considerations mentioned above, there are other factors that can impact performance in Plixer Scrutinizer, such as the number/complexity of [Notification Profiles](#) in use, the number of [report thresholds](#) configured, and the number of [scheduled email reports](#) that have been set up. It is recommended to regularly review the **Admin > Resources > System Performance** page to ensure that resource utilization remains within acceptable values.

3.7.2 Storage

The [Admin > Resources > System Performance](#) page of the web interface summarizes disk utilization for individual collectors in a Plixer Scrutinizer environment. A more detailed view that shows actual and expected storage use for historical flow data can also be accessed by drilling into a specific collector.

This section discusses the main factors that influence a Plixer Scrutinizer collector's disk use and provides instructions for anticipating additional storage needs.

Data retention

Plixer Scrutinizer's data history settings (under [Admin > Settings](#) in the web interface) can be used to adjust how long Plixer Scrutinizer stores *aggregated flow data*, alarm/event details, and other information. With the default settings, a collector provisioned with the minimum 100 GB of storage can store up to 30 days of NetFlow V5 data for a maximum of 25 flow-exporting devices with a combined flow rate of 1,500 flows/s.

Note: Plixer Scrutinizer's functions are highly I/O intensive, and there are many factors that can impact the system's disk-based performance, such as the size/complexity of flows being received and flow cardinality. To ensure optimal performance, 15k HDDs or SSDs in a RAID 10 are recommended.

To determine the expected amount of disk space that will be used by a collector for a distinct data retention configuration, *at the current exporter count and total flow rate*, follow these steps:

1. In the web interface, make the necessary edits to the data retention settings in the Admin > Settings > Data History tray and click **Apply**.
2. Navigate to [Admin > Resources > System Performance](#) and click on the collector address in the **Active Collectors** table to open the **Disk Utilization** view.
3. Review the **HD Utilization** graph and **Utilization per Interval** summary table to compare the current disk space used for each historical flow data interval against the predicted/expected utilization based on the configured data history settings (e.g., **X** hours of 1-minute averages, **Y** days of 12-hour averages, etc.).
4. If necessary, make adjustments to the configured data history settings and/or the collector's storage allocation.

Note: The *predicted* values in the **Disk Utilization** view will automatically be updated to reflect any changes applied to Plixer Scrutinizer's data history settings.

Plixer Scrutinizer automatically trims older historical flow data when available disk space falls below the *Minimum Percent Free Disk Space Before Trimming* value configured in the data history settings. This behavior can be disabled by unticking the *Auto History Trimming* checkbox, but flow collection and other functions may be paused when available storage runs low. The amount of storage for the collector can also be increased to retain older records.

Hint: The **Data Retention** graph shows the number of days of historical flow data currently saved compared against the total number of days that will be retained based on the current data history settings.

Host indexing

When enabling *host and host to host indexing*, it may become necessary to allocate additional disk space, *CPU cores, and RAM* to Plixer Scrutinizer collectors.

Host to host indexing can have a significant impact on disk utilization, because the database will include records for two types of host pairs:

- Continuously active pairs, for whom records will not expire
- Ephemeral unique pairs, for whom records will expire but are also replaced at approximately the same rate

Disk space calculations

To approximate the amount of additional disk space that will be used by the host to host index:

1. *Create/run a new a Host to Host pair report* and add all Exporters that were defined as inclusions for the *Host Indexing* FA algorithm.
2. Set the time window to cover a period of at least 24 hours.
3. When the *output of the report* is displayed, click the gear button to open the Options tray and select *Global*.
4. In the secondary tray, select the *5m* option from the **Data Source** dropdown and click *Apply* before returning to the main view.
5. Note the total result count, which will be roughly equivalent to the number of active pairs.

6. Return to the **Options > Global** tray and switch to the *Im* data source option.
7. Subtract the previous result count from the updated total result count to determine the number of ephemeral pairs.

After obtaining the active pair and ephemeral pair counts, the following formula can be used to calculate additional disk space requirements for host to host indexing:

$$(\text{Active pair count} + \text{Ephemeral pair count}) * \text{Exporter count} * 200 \text{ B}$$

where **Exporter count** corresponds to the total number of Exporters/inclusions defined for the *Host Indexing* algorithm.

Utilization alerts

If the combined disk space used by the host and host pair databases reaches 100% of the *Host Index Max Disk Space* setting of the *Host Indexing* algorithm, host and host to host indexing will be suspended until storage becomes available again.

The following Alarm Policies are used to alert users to high disk utilization by host indexing:

| | |
|---|--|
| Host Index Disk Space Warning | Triggered when the disk space used by host indexing functions reaches/exceeds 75% of the specified <i>Host Index Max Disk Space</i> |
| Host Index Disk Space Error | Triggered when host indexing functions are suspended because the <i>Host Index Max Disk Space</i> has been reached |
| Host Index Disk Availability Error | Triggered when host indexing functions are suspended because disk utilization for the volume the host and host pair databases are stored on has reached/exceeded 90% |

3.7.3 Distributed cluster sizing

Distributed configurations consisting of one primary reporting server and multiple remote collectors allow Plexier Scrutinizer to scale beyond the single-appliance ceiling of 500 exporters with a total flow rate of 150,000 flows/s.

This section contains resource allocation guidelines and recommendations for individual appliances in a distributed cluster.

Remote collectors

In a distributed environment, resource allocation for each remote collector should follow the same guidelines/recommendations as that of a single Plexier Scrutinizer appliance:

1. Use the expected flow rate and exporter count for the collector to determine recommended *CPU and RAM allocations for core functions*.
2. Calculate the total additional CPU cores and RAM required to support the *features* that will be enabled for the collector and exporters associated with it.
3. Provision the collector with the minimum 100 GB of disk space and the total CPU and RAM obtained from the first two steps.

After the collector has been *registered as part of the cluster* and is receiving flows, continue to monitor resource utilization via *Admin > Resources > System Performance* page and make adjustments when necessary.

Primary reporter

CPU and RAM requirements for the primary reporter in a distributed environment are primarily based on the number of remote collectors in the cluster:

| | Minimum | Recommended |
|------------------|------------------------------------|------------------------------------|
| CPU cores | 2x the number of remote collectors | 4x the number of remote collectors |
| RAM | 2 GB for every remote collector | 4 GB for every remote collector |

Note: Depending on the scale of the network, the primary reporter may be subject to additional load due to the volume of Alarms/events being forwarded by the collectors.

3.7.4 Plexier ML Engine

Deployed as part of Plexier One Network and Plexier One Security, the Plexier ML Engine is a supplementary appliance that provides advanced anomaly and threat detection through Plexier Scrutinizer.

The following subsections contain sizing guidelines for local and cloud-based Plexier ML Engine deployments:

Hint: Sizing recommendations for the Plexier ML Engine are based on flow rates and asset counts. An “asset” is either an exporter interface (Plexier One Network) or a host (Plexier One Security).

Local deployments

The following table shows the recommended resource allocations for a local Plixer ML Engine install:

CPU cores, RAM, and disk space based on expected flow rate and asset count

| | | Number of assets | | | | | | | | | |
|---------|-----|--------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|--------------------------|
| | | 150 | 300 | 450 | 600 | 750 | 900 | 1050 | 1200 | 1450 | 1700 |
| Flows/s | 10k | 8 CPU cores | 12 CPU cores | 16 CPU cores | 20 CPU cores | 24 CPU cores | 28 CPU cores | 32 CPU cores | 36 CPU cores | 40 CPU cores | 44 CPU cores |
| | | 40 GB RAM | 80 GB RAM | 112 GB RAM | 136 GB RAM | 160 GB RAM | 184 GB RAM | 208 GB RAM | 232 GB RAM | 256 GB RAM | 256 GB RAM |
| | | 0.2 TB disk | 0.4 TB disk | 0.6 TB disk | 0.8 disk | 1.0 TB disk | 1.2 TB disk | 1.4 TB disk | 1.6 TB disk | 1.8 TB disk | 2.0 TB disk |
| | | | | | | | | | | | |
| | 20k | 12 CPU cores | 14 CPU cores | 18 CPU cores | 22 CPU cores | 26 CPU cores | 30 CPU cores | 34 CPU cores | 38 CPU cores | 42 CPU cores | 46 CPU cores |
| | | 80 GB RAM | 112 GB RAM | 136 GB RAM | 160 GB RAM | 184 GB RAM | 208 GB RAM | 232 GB RAM | 244 GB RAM | 256 GB RAM | 288 GB RAM |
| | | 0.4 TB disk | 0.6 TB disk | 0.8 TB disk | 1.0 TB disk | 1.2 TB disk | 1.4 TB disk | 1.6 TB disk | 1.8 TB disk | 2.0 TB disk | 2.2 TB disk |
| | 30k | 16 CPU cores | 18 CPU cores | 20 CPU cores | 24 CPU cores | 28 CPU cores | 32 CPU cores | 36 CPU cores | 40 CPU cores | 44 CPU cores | 48 CPU cores |
| | | 112 GB RAM | 136 GB RAM | 160 GB RAM | 184 GB RAM | 208 GB RAM | 232 GB RAM | 244 GB RAM | 256 GB RAM | 288 GB RAM | 320 GB RAM |
| | | 0.6 TB disk | 0.8 TB disk | 1.0 TB disk | 1.2 TB disk | 1.4 TB disk | 1.6 TB disk | 1.8 TB disk | 2.0 TB disk | 2.2 TB disk | 2.4 TB disk |
| | 40k | 20 CPU cores | 22 CPU cores | 24 CPU cores | 26 CPU cores | 30 CPU cores | 34 CPU cores | 38 CPU cores | 42 CPU cores | 46 CPU cores | 50 CPU cores |
| | | 136 GB RAM | 160 GB RAM | 184 GB RAM | 208 GB RAM | 232 GB RAM | 244 GB RAM | 256 GB RAM | 288 GB RAM | 320 GB RAM | 352 GB RAM |
| 64 | | RAM 0.8 disk | RAM 1.0 TB disk | RAM 1.2 TB disk | RAM 1.4 TB disk | RAM 1.6 TB disk | RAM 1.8 TB disk | RAM 2.0 TB disk | RAM 2.2 TB disk | RAM 2.4 TB disk | RAM 2.6 TB disk |

AWS deployments

When deploying the Plixer ML Engine as an AWS AMI, use the following table to determine the appropriate instance type and amount of storage:

Instance type and Elastic Block Storage (EBS) size based on flow rate and asset count

| | | Number of assets | | | | | | | | | |
|---------|-----|----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|-----------------------------------|
| | | 150 | 300 | 450 | 600 | 750 | 900 | 1050 | 1200 | 1450 | 1700 |
| Flows/s | 10k | r5a.2xlarge 0.2 TB disk | r5a.4xlarge 0.4 TB disk | r5a.4xlarge 0.6 TB disk | r5a.8xlarge 0.8 TB disk | r5a.8xlarge 1.0 TB disk | r5a.8xlarge 1.2 TB disk | r5a.8xlarge 1.4 TB disk | r5a.12xlarge 1.6 TB disk | r5a.12xlarge 1.8 TB disk | r5a.12xlarge 2.0 TB disk |
| | 20k | r5a.4xlarge 0.4 TB disk | r5a.4xlarge 0.6 TB disk | r5a.8xlarge 0.8 TB disk | r5a.8xlarge 1.0 TB disk | r5a.8xlarge 1.2 TB disk | r5a.8xlarge 1.4 TB disk | r5a.12xlarge 1.6 TB disk | r5a.12xlarge 1.8 TB disk | r5a.12xlarge 2.0 TB disk | r5a.12xlarge 2.2 TB disk |
| | 30k | r5a.4xlarge 0.6 TB disk | r5a.8xlarge 0.8 TB disk | r5a.8xlarge 1.0 TB disk | r5a.8xlarge 1.2 TB disk | r5a.8xlarge 1.4 TB disk | r5a.8xlarge 1.6 TB disk | r5a.12xlarge 1.8 TB disk | r5a.12xlarge 2.0 TB disk | r5a.12xlarge 2.2 TB disk | r5a.12xlarge 2.4 TB disk |
| | 40k | r5a.8xlarge 0.8 TB disk | r5a.8xlarge 1.0 TB disk | r5a.8xlarge 1.2 TB disk | r5a.8xlarge 1.4 TB disk | r5a.8xlarge 1.6 TB disk | r5a.12xlarge 1.8 TB disk | r5a.12xlarge 2.0 TB disk | r5a.12xlarge 2.2 TB disk | r5a.12xlarge 2.4 TB disk | r5a.16xlarge 2.6 TB disk |
| | 50k | r5a.8xlarge 1.0 TB disk | r5a.8xlarge 1.2 TB disk | r5a.8xlarge 1.4 TB disk | r5a.8xlarge 1.6 TB disk | r5a.12xlarge 1.8 TB disk | r5a.12xlarge 2.0 TB disk | r5a.12xlarge 2.2 TB disk | r5a.12xlarge 2.4 TB disk | r5a.12xlarge 2.6 TB disk | r5a.16xlarge 2.8 TB disk |
| | 60k | r5a.8xlarge 1.2 TB disk | r5a.8xlarge 1.4 TB disk | r5a.8xlarge 1.6 TB disk | r5a.12xlarge 1.8 TB disk | r5a.12xlarge 2.0 TB disk | r5a.12xlarge 2.2 TB disk | r5a.12xlarge 2.4 TB disk | r5a.12xlarge 2.6 TB disk | r5a.16xlarge 2.8 TB disk | r5a.16xlarge 3.0 TB disk |
| 66 | | 3. Configuration Guides | | | | | | | | | |
| | 70k | r5a.8xlarge 1.4 TB disk | r5a.12xlarge 1.6 TB disk | r5a.12xlarge 1.8 TB disk | r5a.12xlarge 2.0 TB disk | r5a.12xlarge 2.2 TB disk | r5a.12xlarge 2.4 TB disk | r5a.12xlarge 2.6 TB disk | r5a.16xlarge 2.8 TB disk | r5a.16xlarge 3.0 TB disk | r5a.16xlarge 3.2 TB disk |

Azure deployments

When deploying the Plixer ML Engine as an Azure VM image, use the following table to determine the appropriate VM size and amount of storage:

VM and Azure Disk Storage (ADS) sizes based on flow rate and asset count

| | | Number of assets | | | | | | | | | | |
|---------|-----|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|---------|
| | | 150 | 300 | 450 | 600 | 750 | 900 | 1050 | 1200 | 1450 | 1700 | |
| Flows/s | 10k | Standard 0.2 TB disk | Standard 0.4 TB disk | Standard 0.6 TB disk | Standard 0.8 TB disk | Standard 1.0 TB disk | Standard 1.2 TB disk | Standard 1.4 TB disk | Standard 1.6 TB disk | Standard 1.8 TB disk | Standard 2.0 TB disk | SE48_v5 |
| | 20k | Standard 0.4 TB disk | Standard 0.6 TB disk | Standard 0.8 TB disk | Standard 1.0 TB disk | Standard 1.2 TB disk | Standard 1.4 TB disk | Standard 1.6 TB disk | Standard 1.8 TB disk | Standard 2.0 TB disk | Standard 2.2 TB disk | SE48_v5 |
| | 30k | Standard 0.6 TB disk | Standard 0.8 TB disk | Standard 1.0 TB disk | Standard 1.2 TB disk | Standard 1.4 TB disk | Standard 1.6 TB disk | Standard 1.8 TB disk | Standard 2.0 TB disk | Standard 2.2 TB disk | Standard 2.4 TB disk | SE48_v5 |
| | 40k | Standard 0.8 TB disk | Standard 1.0 TB disk | Standard 1.2 TB disk | Standard 1.4 TB disk | Standard 1.6 TB disk | Standard 1.8 TB disk | Standard 2.0 TB disk | Standard 2.2 TB disk | Standard 2.4 TB disk | Standard 2.6 TB disk | SE64_v5 |
| | 50k | Standard 1.0 TB disk | Standard 1.2 TB disk | Standard 1.4 TB disk | Standard 1.6 TB disk | Standard 1.8 TB disk | Standard 2.0 TB disk | Standard 2.2 TB disk | Standard 2.4 TB disk | Standard 2.6 TB disk | Standard 2.8 TB disk | SE64_v5 |
| | 60k | Standard 1.2 TB disk | Standard 1.4 TB disk | Standard 1.6 TB disk | Standard 1.8 TB disk | Standard 2.0 TB disk | Standard 2.2 TB disk | Standard 2.4 TB disk | Standard 2.6 TB disk | Standard 2.8 TB disk | Standard 3.0 TB disk | SE64_v5 |
| 68 | 70k | 3. Configuration Guides | | | | | | | | | | |
| | | Standard 1.4 TB disk | Standard 1.6 TB disk | Standard 1.8 TB disk | Standard 2.0 TB disk | Standard 2.2 TB disk | Standard 2.4 TB disk | Standard 2.6 TB disk | Standard 2.8 TB disk | Standard 3.0 TB disk | Standard 3.2 TB disk | SE64_v5 |

Note: To learn more about Plixer ML Engine licensing options and deployment procedures, contact [Plixer Technical Support](#).

USE CASES

This section details how Plixer Scrutinizer's various functions and features can be applied in a wide range of network and security use cases.

For ease of navigation, these guides are divided into separate subsections for NetOps and SecOps:

NetOps use cases

- Maintain deep visibility via various configurable views
- Monitor health/performance in real time
- Run fully customized reports to investigate issues

SecOps Use Cases

- Get alerted to malware and other threats based on multiple detection methods
- Access historical data to search for traffic indicating malicious activity
- Enhance incident response and threat-hunting workflows with customized views/reports

4.1 NetOps Use Cases

Select a use case to learn more:

| Customer Need | Use Case | Workflows |
|---|--|---|
| Aggregate flow data by any dimension to inspect any host or traffic on the network | <i>Customizable observation points and reporting</i> | – |
| Streamline information sharing and enhance multi-role workflows | <i>Team collaboration</i> | <i>Sharing information via Collections</i> |
| Monitor network health/performance in real time and quickly identify root causes | <i>Investigating network congestion</i> | <i>Monitoring for congestion issues</i> <i>Troubleshooting poor call quality</i> |
| Proactively monitor specified network traffic from any email inbox | <i>Scheduled email reporting</i> | <i>Automating weekly reports</i> |
| Create and customize network maps to visualize what matters to your team | <i>Network mapping and visualization</i> | <i>Mapping your network</i> |
| Maintain multiple customizable dashboards to support unique roles and workflows | <i>NOC dashboards and forensics</i> | <i>Multi-tenancy dashboards</i> |
| Continuously monitor network health/performance and extract additional traffic insights | <i>Network performance monitoring (NPM)</i> | <i>Monitoring for congestion issues</i> |
| Monitor how data circuits are used over time to plan future needs and optimize costs | <i>Capacity planning</i> | <i>Forecasting and meeting business needs</i> |
| Bridge visibility between cloud and on-prem resources without deploying probes | <i>Cloud visibility and detection</i> | – |

4.1.1 Customizable observation points and reporting

With the Plixer One Platform (Core, Network, or Security), users can use Plixer Scrutinizer to configure/run their own *purpose-built reports*. These reports are fully customizable and can be used to visualize network performance, identify problem points, and investigate root causes of network issues. Reports can also be continuously refined to filter, drill down, and/or pivot as part of monitoring or investigative activities.

Overview

Reports in Plixer Scrutinizer aggregate data from *one or more devices/sources based on the dimensions defined in the base report Type*. To further adapt a report to more specific monitoring and investigative needs, there are a range of settings that can be modified.

Configuring reports

In addition to the base type and data sources, Reports use the following settings when they are run:

- Time period covered (either *last X* or custom date/time ranges)
- Graph/visualization type
- Filters

Each report can have multiple filters in any combination of *filter types (device/interface, domain, host addresses, etc.)* defined as either inclusions or exclusions. Additionally, filters can be configured so that they include only source hosts, destination hosts, or both.

Report settings, including the base type and devices/sources, can be set/changed in the report creation wizard or when *refining the output* after the report is run.

Hint: In the *report output view*, table elements can be dragged to *Include* and *Exclude* drop zones to re-define the report's inclusions and exclusions. Additionally, clicking on a dimension element opens a tray that allows the user to pivot to any other report type available for that element.

Additional options

After a report is created/run, it can be *saved and/or exported in several ways* to enhance a wide range workflows.

With Plixer One Network, saved reports can also be used to generate forecasts for *capacity planning* and to enable *more efficient collaboration* between team members.

4.1.2 Team collaboration

To support the growing scale and complexity of enterprise environments, the Plixer One Platform (Core, Network, or Security) includes multiple functions that enable greater efficiency in collaborative processes and workflows:

- Save *custom Reports* and allow other members to access/re-run them at any time
- *Email Reports* directly to concerned parties or export them for use in external systems
- Compile Alarm details and/or Reports into *Collections* for review/investigation by multiple team members
- Assign one or more notification actions (including email alerts) to Alarm Policies through customizable *Notification Profiles*.

Overview

Plixer Scrutinizer includes multiple features and functions that are designed to streamline the sharing of network and incident information between members and teams.

Saved Reports

Plixer Scrutinizer Reports function as a customizable network visibility interface, where you can continuously *filter, drill down, and pivot to different Report types* to monitor specific network elements or *identify problem points*. Once a Report configuration is saved, other users can be given access (through *Usergroups*) to re-run it or add it to their *Dashboards*.

Hint: A saved Report can also be used to set up a *Scheduled Email Report* to automatically run and email the Report to any number of users at regular intervals.

Report/notification emails

Once an *email server* has been configured, Plixer Scrutinizer can be set to send alerts and reports directly to user inboxes:

- *On-demand Email Reports* after any Report is run
- *Scheduled Email Reports* at user-specified intervals

- Alarm/Event email notifications, which are triggered via *Notification Profiles* assigned to *Alarm Policies*

Hint: Both email report types also include a link to run the Report in the Plixer Scrutinizer web interface. PDF and/or CSV copies of the Report may also be attached.

Collections

Collections are compilations of Alarm/Event data or Reports that are assigned to specified users for review, analysis, or resolution. In addition, Collections can be viewed by other users, who are able to add annotations directly to the Collection item's details and/or engage in discussions via threaded notes/comments.

Hint: While reviewing a Collection, a user can click on individual items to quickly jump to more detailed views.

Important: Collections are part of the **Plixer One Network** solution. Contact *Plixer Technical Support* for more information.

Workflows

The following workflow(s) show how the Plixer One Platform can drive more efficient collaborative workflows through various functions:

Sharing information via Collections

The network team discovers suspicious traffic and wants to share the information with an independently operating security team. Instead of exporting the information and sending it via email, they create a Collection containing the relevant Reports and/or Alarm data that can be accessed by other Plixer Scrutinizer users at any time.

Workflow

1. From the Alarm Monitor view or Report page associated with the suspicious traffic, create a new Collection from the [Manage Collections](#) submenu (star button) and set it as the active Collection.

Tip: To use an existing Collection instead, click the star button and select the Collection from the menu.

2. To add an item after the active Collection has been set, click the star button from any relevant Alarm/Event information view or Report and click it a second time, after it has been replaced with a **+**.
3. Repeat the previous step to add additional items to the Collection.

All Plixer Scrutinizer users can access existing Collections via the [Investigate > Collections](#) page of the web interface. When [inspecting a Collection](#), users can add notes to the individual items or for the Collection itself.

Hint: The default view of the Collections page displays all Collections that have been assigned to the current user. To see other Collections, switch to the *Other Collections* tab of the page.

Collections that are no longer relevant can be deleted by selecting them from the main Collections page and clicking the *Delete* button.

4.1.3 Investigating network congestion

In almost any modern enterprise environment, identifying the who, what, where, when, and why behind congestion issues requires tools that go beyond inundating network teams with large volumes of raw data.

Through Plixer Scrutinizer, the Plixer One Platform (Core, Network, or Security) enables multiple approaches to dealing with network congestion issues:

- Drill down into network device/host activity to identify root causes for congestion by applying one or more filters and pivoting between different [Report Types](#).
- Monitor network devices and/or interfaces for congestion in the [Top Interfaces view](#).
- See real-time rates and utilization between devices and other objects in [Network Maps](#) by adding *Connections* with custom color-coded thresholds.
- Get high utilization alerts via the Plixer Scrutinizer [Alarm Monitor](#) by adding user-defined thresholds to Reports.

Overview

Teams can leverage the following Plixer Scrutinizer features/functions to proactively watch for network congestion, collect insights into the root cause(s), and respond efficiently.

Reports

Reports aggregate data from *any number of user-specified devices and dimensions* and can show sources of congestion and bandwidth consumption:

- Identify “Top Talkers” on the network using *Source* and *Destination* Reports.
- View peak and 95th percentile in *Traffic Volume* Reports.
- Check for latency and packet loss with Plixer FlowPro APM *Application Retransmission* Reports.
- *Apply any number of filters* for subnets, applications, usernames and then pivot directly to another Report Type to narrow down your results.

Report Thresholds

Custom Thresholds can be added to *saved Reports* to monitor for congestion and trigger Alarm Monitor *alerts* when those thresholds are reached. With a Report Threshold configured, the Report can be re-run to monitor for min/max bandwidth utilization and mitigate regression after congestion sources are identified.

Hint: If a *Notification Profile* is assigned to the *Report Threshold Violation Alarm Policy*, the threshold can be used to trigger *notification actions*, such as email alerts and CEF notifications for external tools.

Top Interfaces view

The Top Interfaces view (**Explore > Exporters** in the web interface) can be used to monitor all device interfaces, from the most saturated down to the least utilized. This allows network teams to identify which ones are most affected by congestion at a glance. The view can also be used to inspect highwater marks that indicate peak saturation over a period of time.

Hint: The **Explore > Exporters** page can be set to show either *By Interfaces* or *By Exporters* as the default in your user preferences menu.

Map Connections

After a Network Map is populated with *devices and other objects*, it can be further customized with Connections representing activity between devices, objects, and/or interfaces. *Connections* can also be individually configured with utilization thresholds that change the color they're displayed in, giving teams a bird's eye view of potential congestion issues in real time.

Hint: Click on devices or interfaces in a Network Map to quickly jump to the Top Interfaces view filtered on the object.

Workflows

The following workflows show how multiple Plixer One Platform functions can help network teams mitigate, and/or investigate network congestion issues.

Monitoring for congestion issues

Scenario

A user calls in reporting that everything on the network is taking an excessive amount of time to load, indicating network congestion.

Workflow

- Navigate to Explore -> Exporters -> Interfaces
- Identify instantly if any interfaces are congested
- Open a “Conversations” Report to see the top source and destinations of bandwidth
- We may find that a host on the network is performing write intensive backups during the day and eating up all available bandwidth.

Tip: If *Host Indexing* is turned on, you can look up a user's IP and see all network devices that saw that address.

Note: Plixer Scrutinizer records *highwater marks* that represent the peak utilization (based on the configured interface speed) for each interface.

Troubleshooting poor call quality

Scenario

The sales teams reports that outbound calls have been of poor quality recently. Jitter happening sporadically on the call, making it difficult to conduct business efficiently.

Workflow

- Navigate to Reports -> Run Report -> Select Report Types
- Under the *Flowpro APM Reports* category, select a report like 'Host to Host Jitter All by SSRC'
- Open the report and note the report columns such as Source Jitter and Packet Loss
- We may find that we can measure the Jitter and packet loss and see what the RTP payload type was. Perhaps the subnet traffic is not using Class based QOS and voice traffic isn't being prioritized.

Note: Plixer FlowPro is part of the Plixer One platform. To learn more, see the section on [FlowPro integration](#).

4.1.4 Scheduled email reporting

With the Plixer One Platform (Core, Network, or Security), NetOps teams can use Plixer Scrutinizer Reports as a proactive monitoring tool for any type of network meta data by setting up [scheduled email reports](#).

Overview

A Scheduled Email Report is a saved Report that has been set to run at specified intervals using the exact same configuration (*graph, filters, etc.*). Each time the Report is run, its output is automatically emailed to one or more recipients.

Note: Scheduled Email Reports are different from *on-demand Report emails*, which must be sent manually after a Report is run.

All Email Reports contain a direct link to the primary Report and may also include PDF/CSV copies of the Report. One or more additional Reports can also be run and sent in the email.

Setting up a Scheduled Email Report

A Scheduled Email Report can be set up after re-running a saved Report (or after *creating and saving a new Report*).

From there, click the *Export Report/share* button, select *Schedule Report* in the tray, and configure the following:

- A name for the Scheduled Email Report configuration (used for configuration management and as the subject line of the email)
- One or more recipient addresses (comma-separated)
- Frequency and time (minute on the hour) to run and send the Report
- (Optional) PDF and/or CSV format attachments (all included Reports)
- (Optional) Additional Reports to run and include in the email

Once set up, the Report(s) will be run/sent at the specified intervals until the Scheduled Email Report configuration is disabled or deleted.

Hint: To inspect, edit, or disable Scheduled Email Report configurations, navigate to **Admin > Reports > Scheduled Email Reports**.

Workflows

The following workflow(s) show how the Plixer One Platform is able to continuously monitor specific network traffic through scheduled email reports:

Automating weekly reports

Important: To set up scheduled email reports, an email server must first be configured via the [Admin > Integrations](#) page.

Scenario

Management wants to see summarized data concerning the network emailed on a weekly basis.

Workflow

First off, identify the details that are most critical to report on. Some examples are: top applications, top used ports, destination countries, etc.. Regardless of the report types required, the same steps are used to add reports to your scheduled report.

1. Select *Reports* -> *Run Report* -> *Select Report Type* to start a [report](#).
2. Choose *Destination Reports* -> *Countries with AS* and select the appropriate network devices to include in the report.
3. Change the range of the report to *Last Seven Days* to show the entire weeks network data.
4. Save and give this report a name.
5. Export the report as a gadget from the Options tray.

Repeat the same steps for the other reports, making sure the time range is *Last Seven Days*

- Pair Reports -> Conversations Apps
- Top -> Protocols
- Top -> Well Known Ports

Now that the reports that will be sent weekly have been created, they can now be assigned to a [scheduled report](#).

Assign the frequency to 'Weekly' and set time to the day of the week and time to see this email come through, "Friday 5:00pm". Options include adding PDF and CSV attachments along with the email.

Be sure to select the reports that were created for this scheduled email and add them to the include list. After a scheduled report configuration has been set up, it can be viewed or edited from **Admin > Reports > Scheduled Email Reports**.

4.1.5 Network mapping and visualization

With the Plixer One Platform (Core, Network, or Security), network teams can leverage Plixer Scrutinizer's integrated network mapping functions to create and customize maps that are based on user-defined device groups. These maps are continuously updated in real time, allowing them to function as both a high-level view of network health and a starting point for investigating connectivity issues.

Overview

When *creating a new map* in Plixer Scrutinizer, users can select between *Spatial Maps* to fully customize the device layout or *Geographical Maps* for location-based arrangement.

After a network map is initially generated, it can be further *customized/configured* at any time. Existing network maps can be viewed from the **Monitor > Network Maps** page or as *Dashboard Gadgets*.

Spatial Maps

Using the following configuration options, Spatial Maps can be used to design fully customized topologies to meet different visualization requirements:

- Position *map objects* against *custom backgrounds* to recreate office layouts, wiring closet connections, and more.
- Add *custom objects* to represent non-Exporters, such as external hosts
- Define *connections* between objects (devices, interfaces, and/or custom objects) to indicate static links, display interface utilization, or run a saved Report using the connected objects
- Add custom utilization thresholds to connections to show overall network health and potential congestion issues
- Nest mapping groups within each other and create multi-layered maps to support network segment planning and monitoring
- Tailor maps to specific team role or workflow needs and manage access via *dashboards* and usergroups.

Hint: Bulk management functions for mapping objects and groups can be accessed via the Mapping Objects and Mapping Groups pages under **Admin > Settings** in the web interface.

Geographical Maps

Object positions in Geographical Maps are determined by their longitudinal and latitudinal coordinates. Both manual coordinate entry and address lookups via Google Maps are supported.

Hint: Objects can be assigned unique coordinates/addresses for every map/group they are assigned to.

Geographical Maps support similar configuration/customization options as Spatial Maps (except for object positioning and custom backgrounds) and can be used to enhance many of the same workflows. They are also ideal for monitoring the health and performance of geographically segmented networks.

Workflows

The following workflow(s) are examples of workflow enhancements enabled by Plixer Scrutinizer's live network maps in the Plixer One Platform:

Mapping your network

To streamline NOC workflows in their growing environment, the team decides that they need a visual representation of the network and critical applications.

Workflow

To set up the new map, navigate to **Monitor > Network Maps** and [create a new Spatial Map](#):

1. Use a name that matches the coverage of the map (e.g., the entire network).
2. Assign all applicable devices (routers, firewalls, switches) as map objects.
3. Link devices as necessary by creating Connections. Connections can be static lines, interface representations, or saved Reports.

Hint: When a saved Report is used as a Connection, it will represent the traffic aggregated by the Report. This can be anything from a layer 7 application (e.g., YouTube) to firewall events from a Cisco ASA. In the latter case, the Connection will typically be grayed out (inactive), and can serve to quickly alert the network team when it becomes active.

If the network topography changes at a later time, [the map can be updated](#) to reflect the changes.

For larger networks, such as those that span multiple locations, it may be ideal to create smaller maps representing individual network segments and nest them under a larger map as objects. This will create a “global” map with a hub-and-spoke layout.

4.1.6 NOC dashboards and forensics

As ubiquitous as dashboards have become in network operations center (NOC) workflows, many tools remain limited by the lack of customization options for data sources, gadgets, and auxiliary features.

Plixer Scrutinizer Dashboards—part of the Plixer One Platform (Core, Network, or Security)—can be customized to support and enhance any number of unique user roles and/or workflows.

Overview

Plixer Scrutinizer users are able to create any number of uniquely configured dashboards to support and enhance their individual workflows.

Dashboard management

When [creating a new dashboard](#), users can choose between starting with a copy of an existing dashboard or populating a “blank” dashboard with their own selection of gadgets.

Existing dashboards also have the following additional management/configuration options:

| | |
|--------------------------------|---|
| <i>Set as default</i> | Selects the dashboard as the default for the current user |
| <i>Set as read-only</i> | Locks dashboard settings and gadgets until toggled off |
| <i>Modify user access</i> | Shows or hides the dashboard for individual users |
| <i>Modify usergroup access</i> | Shows or hides the dashboard for usergroups |

Hint: To change the layout and gadgets for existing dashboards, switch to [edit mode](#) while dashboard is active.

Custom Gadgets

To complement the preconfigured gadgets bundled with Plixer Scrutinizer, [network maps](#) and [reports](#) can also be added to dashboards as gadgets. This allows users to view/access frequently used maps and reports directly from their preferred dashboard(s) instead of navigating to the corresponding sections of the web interface.

Existing network maps or reports (must be [exported](#) first) can be added when setting up a new dashboard or while in dashboard edit mode, provided the current user has access via their usergroup.

External gadgets

External gadgets are another type of custom gadget that allow Plixer Scrutinizer users to embed valuable data from third-party sites (via URL) in their dashboards and further extend visibility.

Hint: External and Report-based gadgets can be configured with custom refresh intervals to always display the data that is most relevant to users.

Workflows

The following workflow(s) show how the Plixer One Platform is able to enable and enhance UI-driven workflows with Plixer Scrutinizer Dashboards:

Multi-tenancy dashboards

As part of a multi-tenant environment, the operator wants to provide each customer with a dashboard for their network.

Workflow

Assuming two groups (A and B), each group should have exclusive logins so that only content relating to their group is accessible to their users.

This workflow assumes that each of these groups consists of a location with three network devices sending netflow data:

- Firewall
- Core Router
- Switch

The dashboard should contain a single top conversations report for the group's network and be accessible to all users under that group/location.

1. *Create a dashboard* for each group (e.g., Dashboard A and Dashboard B). This will allow you to export the appropriate reports to them after they have been created.

2. *Create Group A's report:*

- a. Start by adding devices and select the IP addresses of Firewall A, Core Router A, and Switch A.
- b. Select *Conversations App* (under the *Recommended* category as the Report Type).
- c. Change the time window/range of the report to *Last 24 hours*.
- d. After running the report, save it under a name associated with Group A (e.g., Top Conversations A)
- e. Click the share button and select *Add to Dashboard*.
- f. In the secondary tray, select Dashboard A from the *Dashboard Tab* dropdown and choose what content to show in the gadget (graph, table, or both).

Note: If a different name is entered in the *Report Name* field, a new, separate Report will be saved. The new name will also be used as gadget label.

3. Repeat the previous steps using Firewall B, Core Router B, and Switch B, and export the report to Dashboard B.
4. Set up the Report Folders for each group:
 - a. Navigate to **Admin > Classic Admin > Reports > Report Folders**.
 - b. Click the *New Folder* button and enter a name for Group A's folder (e.g., Report Folder A).
 - c. Add the report that was created for Group A to the folder by selecting it and clicking the *<- Add* button.
 - d. Repeat the steps to create the folder for Group B and add their report to it.
5. Create a map for each group's network:
 - a. Navigate to **Monitor > Network Maps** and *create a new Spatial Map* for Group A (e.g. Map A).
 - b. Assign Firewall A, Core Router A, and Switch A as map objects.
 - c. Link the devices as necessary using *Connections*.

Hint: The Report previously created for Group A (or any other saved Report) can be used to create a Connection representing that traffic type between devices. These Reports can also be added to dashboards for up-to-the minute display of the traffic covered.

- d. Repeat the steps to create the map for Group B.
6. Set up the usergroups:
 - a. Navigate to **Admin > Users and Groups > Usergroups** and click the + button to create a new group.
 - b. In the tray, enter a name (e.g., Group A Users) and select *Guest* as the starting template from the dropdown.
 - c. After the usergroup has been created, locate it in the main table and click the links under the columns to make the following changes:
 - Devices: Select only Firewall A, Core Router A, and Switch A.
 - Interfaces: Select only interfaces that should be visible to Group A (all interfaces associated with their devices, in most cases)
 - Reports: Select all Reports and Report Folders created for Group A.
 - Dashboard Gadgets: Select only gadgets (based on saved Report names) that were created for Group A.
 - d. Repeat the steps to set up the usergroup for Group B.
 7. Navigate to **Admin > Users and Groups > User Accounts** and click the + button to create login credentials for one or more users for each group. Use the dropdown in the tray to add each user to the appropriate usergroup.

Hint: Users obtained from LDAP or another identity provider can also be added to usergroups.

After everything has been set up, users from each group will only have access to the devices/interfaces, reports, and dashboards/gadgets belonging to their group.

4.1.7 Network performance monitoring (NPM)

Without true visibility into traffic patterns and trends, additional provisioning may seem like the only way to keep up with a network's growth.

With Plixer One Network, network teams can access detailed information related to application performance and performance costs, in addition to being able to examine end-to-end network conversation details through Plixer Scrutinizer's *reporting and filtering functions*. Users can also leverage the Plixer ML Engine to *forecast any future network traffic/behavior*.

Overview

Plixer One Network includes multiple functions/components that can enhance a network team's ability to monitor and manage network performance down to the application level.

Reports

In Plixer Scrutinizer, *reports* can help network teams understand the root causes of traffic saturation on a network's top interfaces. When used in conjunction with *Alarms for interface threshold violations*, they can get alerted to saturated circuits and will have the means to uncover what that traffic consists of.

APM

Plixer One Network provides application performance monitoring functions that are designed to support teams in ensuring consistently optimal experiences for their users:

- Measure application round-trip time (RTT)
- Monitor latency for Layer 7 applications, clients, servers, and VoIP communication
- Diagnose issues using SSRC, ToS, jitter, retransmission rates, and other packet metrics

Forecasts

By combining the capabilities of Plixer Scrutinizer with the Plixer ML Engine, Plixer One Network can provide users with forecasts of future network activity to support capacity planning initiatives. These forecasts can help network teams visualize trends of network growth and predict behavior based on the patterns exhibited by past activity.

Once a report has been configured with the correct *settings and filters*, it can be used to *generate a forecast* that predicts the state of the same traffic into the future.

Data history

Plixer Scrutinizer can be tuned to keep historical data for as long as needed through its data retention settings.

Because raw alarms come in off the wire and are stored each minute, the data stored for that interval offers the most granular historical information. To make more efficient use of disk space, however, Plixer Scrutinizer automatically aggregates that data and rolls it up into 5m averages for up to 2-hour intervals. This allows for historical data to be kept for a longer period of time.

To learn more about how Plixer Scrutinizer aggregates historical data, see [this section](#) of this documentation.

Important: APM-specific reports and forecasting are only available with Plixer One Network. Contact [Plixer Technical Support](#) to learn more.

Workflows

The following workflow(s) show how the functions and features included in Plixer One Network can help teams monitor network and application performance in their environment:

4.1.8 Capacity planning

Through Plixer Scrutinizer, Plixer One Network can leverage the capabilities of the Plixer ML Engine to generate forecasts of future network activity for capacity planning:

- Apply machine learning techniques to create dynamic baselines for network behavior
- Extend any Plixer Scrutinizer Report into the future to forecast trends and predict changes to network activity
- Use AI-/ML-driven data analysis to predict VPN trends, proactively plan capacity, and align investments with business needs
- Gain visibility into encrypted VPN tunnels to detect threats
- Gain visibility into address pool utilization and trend its usage
- Associate users, devices, and applications with the consumption of bandwidth

Overview

Plixer Scrutinizer includes multiple tools and functions that can enhance a network team's capacity planning capabilities.

Traffic/behavior baselining

Using collected flow data, the *Plixer ML Engine* is able to create dynamic machine learning models of baseline network behavior.

Plixer One Network can use these models to deliver additional capacity planning insights in two ways:

- Alarms for behavioral deviations that exceed a certain threshold (based on the configured *sensitivity*) using the *Plixer Network Intelligence Anomaly* policy
- Activity/deviation monitoring via **Behavior** tab when drilling into individual hosts from the *Explore > Entities > Hosts* view.

Reports

Plixer Scrutinizer's customizable reports are designed to help teams get to the bottom of any inquiry.

For capacity planning, they can be used to investigate traffic saturation on top interfaces and help determine whether additional provisioning will be required.

Forecasts

Forecasting is a Plixer One Network feature that allows users to create *forecasts* of future network activity.

A forecast can be *generated from any saved report* and will comprise projections for the traffic included by the *report configuration* (e.g., devices, filters, etc.). This gives teams the ability to define the exact network activity to be forecasted as part of capacity planning.

HD utilization projections

On the *Admin > Resources > System Performance* page, clicking on a Collector opens a view showing predicted HD utilization based on the current data retention settings. These projections can be used to ensure that sufficient disk space is always available to meet historical data storage needs.

Workflows

The following workflow(s) show how teams can leverage Plixer One Network functions to enhance their capacity planning capabilities:

Forecasting and meeting business needs

The network team is asked to predict how long an organization's current infrastructure will continue to support their business needs. To visualize trends in network growth, they create report configurations for various aspects of the environment and use them to create *ML-driven forecasts* in Plixer Scrutinizer.

Workflow

Because Plixer Scrutinizer Forecasts are based on Reports, the environment's current capabilities should be split up into separate capacities, such as:

- WAN usage
- VPN traffic
- Subnet-to-subnet patterns
- BGP traffic
- Core router saturation
- Critical application latency

From there, one or more report configurations should be created and saved for each capacity. These reports can then be used to *generate forecasts* that will show emerging utilization trends. At the same time, any latency problems discovered may also indicate potential capacity issues that need to be addressed, depending on their frequency and degree of deviation from the baseline.

4.1.9 Cloud visibility and detection

The Plixer One Platform (Core, Network, or Security) enables seamless visibility across on-prem and cloud-based resources in cloud or hybrid environments through cloud provider log ingestion in Plixer Scrutinizer.

Overview

After the corresponding cloud storage container is set up to receive log data from an Amazon VPC or Azure NSG, Plixer Scrutinizer can be configured to ingest the information via the container. Containers that have been set up as flow data sources in Plixer Scrutinizer are treated as [Exporters](#) and support the same functions and configuration options as typical flow-exporting devices (e.g., [Flow Analytics Security Groups](#), [Plixer ML Engine inclusion](#), and [Reports](#)).

Amazon VPC flow logs

To enable Amazon VPC flow log ingestion in Plixer Scrutinizer, the VPC must first be set to send log data to an Amazon S3 bucket with the [correct configuration](#). Afterwards, the bucket should be added to Plixer Scrutinizer from the [Admin > Integrations > Flow Log Ingestion page](#) in the web interface.

The following [additional Report types](#) can be run when one or more S3 buckets are [selected as data sources](#) for a Report:

- Action
- Action with Interface
- Action with Interface and Dst
- Action with Interface and Src
- Availability Zones
- Dst Service
- Interface
- Pair Interface
- Pair Interface Action
- Src Service
- Src Service-Dst Service
- Traffic Path
- VPCs

Hint: To view only report types that apply to Amazon VPC flow logs, use the *Amazon AWS* category when selecting a report type.

Azure NSG flow logs

Setting up Azure NSG flow log ingestion in Plixer Scrutinizer requires an Azure Blob Storage container that is *correctly configured* and receiving log data from the NSG. This container should be added to Plixer Scrutinizer from the [Admin > Integrations > Flow Log Ingestion page](#) in the web interface.

When one or more Azure blob containers are *selected as data sources* for a Report, the following *additional Report types* become available:

- Flow Decisions
- Flow Decisions Count
- Flow States
- Flow States Count
- NSG All Details
- Resource IDs

Hint: To view only report types that apply to Azure NSG flow logs, use the *Azure* category when selecting a report type.

4.2 SecOps Use Cases

Select a use case to learn more:

| Customer Need | Use Case | Workflows |
|---|------------------------------------|---|
| Continuously monitor critical services for anomalous usage | <i>Service behavior monitoring</i> | <i>Detecting anomalies and deviations</i> |
| Monitor network activity to identify malware-infected hosts | <i>General malware detection</i> | <i>Alerting on malware activity</i> |
| Drill into numerous data points to examine device behavior and pinpoint Indicators of Attack (IoAs) | <i>Threat hunting</i> | <i>Using host index to identify malicious IPs</i> <i>Reviewing Alarm Monitor alerts for suspicious hosts</i> <i>Investigating off-hour network activity</i> <i>Identifying exfiltration outside business hours</i> |
| Monitor network activity to detect lateral movement behavior | <i>Lateral movement detection</i> | <i>Investigating lateral movement alerts</i> <i>Uncovering data exfiltration</i> |
| Enhance incident response procedures through added visibility and UI-driven workflows | <i>Incident response</i> | <i>Responding to Alarm Monitor security alerts</i> <i>Scrutinizing an infected host</i> |

4.2.1 Service behavior monitoring

Plixer One Security addresses the limitations of traditional security technologies by applying AI and ML techniques to provide early, generic detections for activity associated with advanced persistent threats (APTs).

These detections rely on behaviors rather than signatures and give security teams an additional layer of defense against attempts to use common services to infiltrate, infect, and exploit network resources.

Overview

Plixer One Security’s approach to *anomaly detection* relies on the Plixer ML Engine to turn the flow data collected by Plixer Scrutinizer into behavioral models that represent typical host activity. All incoming flow data can then be compared against these baseline models to proactively scan for potentially malicious activity and alert security teams in real time.

Configuring anomaly detection

The Plixer ML Engine’s anomaly detection functions can be adapted to any type of environment through its *configuration*:

| | |
|--------------------|---|
| <i>Dimensions</i> | Services/applications (protocol and port) whose behavior is modeled and monitored for anomaly detection |
| <i>Inclusions</i> | Hosts (by Exporter or subnet) being monitored for anomalous behavior |
| <i>Sensitivity</i> | The tolerance for deviations from baseline service behavior for hosts associated with the inclusion |

Defining dimensions and inclusions for the engine isolates traffic information to reduce the amount of “noise” and maximize the accuracy of detections. Organizations are also able to tune detections to their unique processes and workflows by adjusting the sensitivity for individual inclusions.

Hint: *Low* sensitivity is generally recommended for critical subnets (e.g., finance, HR, etc.) where all irregularities should be reported, while a *High* can be used for hosts whose security requirements are less strict.

Investigating anomaly detections

Once anomalous behavior is reported via an Alarm, the appropriate response can be determined using a combination of Plixer Scrutinizer workflows, including:

- *Drilling down into the Alarm* (e.g., *Plixer Security Intelligence*, *Lateral Movement Behavior*, etc.) and checking the timeline to determine whether the detection is an isolated observation or an ongoing Event
- *Inspecting event Artifacts* to see which hosts were involved and drilling into them to gain further insights from *Plixer Endpoint Analytics*

- Reviewing activity via the **Behavior** tab when drilling into hosts from the [Explore > Entities > Hosts](#) view.
- Running *Source* and *Destination Reports* on the hosts to check for traffic between them and external IP addresses

Hint: After running an initial Report, it can be [refined](#) directly from the output view to enable further investigation.

Workflows

The following workflow(s) show how Alarms related to anomalous service behavior are used to investigate potential cyber attacks:

Detecting anomalies and deviations

Continuously monitor traffic anomalies or traffic deviations that exceed set thresholds using dynamic ML-modeled baselines.

Workflow

Machine learning allows Scrutinizer to alert users to anomalous traffic utilization patterns typically associated with security incidents.

Note: This workflow requires the Plixer ML Engine for predictive modeling. Contact [Plixer Technical Support](#) to learn more about licensing options.

All incoming flow data can be compared against these baseline models to proactively scan for potentially malicious activity and report discoveries in real time.

From there, the next steps should be to [set up reports](#) and using them to [generate forecasts](#).

Identifying which areas of the network (devices and interfaces) have the majority of traffic:

- What types of traffic would you expect to see – VoIP, HTTP, SQL?

- Business Application traffic like Salesforce, AWS, Azure etc.
- DNS requests to dedicated DNS servers on the network

Now consider traffic that may be anomalous:

- Does Remote Desktop Protocol make sense on this network, is there a business usecase for RDP?
- Should there be SSH traffic to critical hosts?

Based on the above considerations, create/run one or more reports to isolate traffic data for services, hosts, or device groups that are most likely to be involved in malicious activity. Once saved, these reports can then be used to forecast expected traffic patterns and highlight deviations (e.g., an anomalous ICMP data trend in outbound WAN usage for edge devices) that can be analyzed to identify threats.

Next steps would be to *customize alerts* for this behavior or other traffic deviations that exceed *user-defined thresholds* configured for the report(s).

Tip: Plixer Scrutinizer's *Alarm Policies* can be assigned custom *Notification Profiles*. To add one or more *notification actions* for all report thresholds, create a Notification Profile and assign it to the *Report Threshold Violation* policy.

4.2.2 General malware detection

Because all malicious activity leaves footprints in network traffic, the visibility provided by traffic data can be an invaluable asset against modern malware.

By ingesting large volumes of network information through Plixer Scrutinizer, Plixer One Security can provide general malware detections and extract additional value from the same flow data.

Overview

The Plexier ML Engine uses *classification* - a machine learning technique that relies on models that have been trained on labeled data - to predict whether a host's behavior is indicative of common classes of malware, including command and control, banking trojans, exploit kits, etc. Each prediction is returned in the form of a percentage, which represents the degree to which the observed traffic patterns match those it has learned to be associated with malware. If that percentage exceeds a preset detection threshold, a high-severity Event is generated under the corresponding *Alarm Policy* in the Plexier Scrutinizer Alarm Monitor.

Enabling malware classification

To optimize resource utilization, malware detection is configured at the ML inclusion level, enabling or disabling classification for all hosts associated with the inclusion. The *Malware Detections* setting can be accessed from the *Manage ML Inclusions* page, where it can be toggled on or off in the inclusion configuration tray.

Investigating malware detections

Once a detection is reported as an Alarm, the appropriate response can be determined using a combination of Plexier Scrutinizer workflows, including:

Note: General ML-driven malware detections are reported under the *ML Engine malware alert* Alarm Policy. A separate *Malware Command and Conquer Activity Detected* policy is used for detections via Flow Analytics.

- *Drilling down into the Alarm* and checking the timeline to determine whether the detection is an isolated observation or an ongoing Event
- *Inspecting Event Artifacts* to see which hosts were involved and drilling into them to gain further insights from *Plexier Endpoint Analytics*
- Running *Source* and *Destination Reports* on the hosts to check for traffic between them and external IP addresses

Hint: After running an initial Report, it can be *refined* directly from the output view to enable further investigation.

Workflows

The following workflow(s) are examples of Plixer One Security's malware detections being used as starting points for investigating suspicious network activity:

Alerting on malware activity

Get alerted to any host demonstrating malware activity and send notification to security team

Workflow

Becoming aware of suspicious activity

Plixer Scrutinizer and the Plixer ML Engine can be used together to help assess possible malware activity on your network.

The ML algorithms used for *malware classification* trigger alerts within Scrutinizer's Alarm Policies for traffic/activity that deviates from dynamic ML-modeled baselines.

Note: This workflow relies on the Plixer ML Engine to report classification-based detections. Additional host analysis and risk assessment functions are enabled through Plixer Endpoint Analytics.

Tip: Plixer Scrutinizer and Plixer FlowPro also use STIX/TAXII and other threat intelligence feeds to identify activity associated with common classes of malware and ransomware.

Responding to potential malware

Review the *Admin -> Alarm Monitor -> Alarm Policies* page and search for the *ML Engine malware alert* policy. Using a custom *Notification Profile*, this policy can be configured to trigger an email to one or more addresses. This can be used to alert security team members whenever there are malware detections that should be reviewed.

Hint: Other automated *notification actions* can also be defined under the same Notification Profile.

From the alarm monitor view with the UI, you could dive into the Alarm Policy and investigate the host with details on top applications and conversations. Plixer Scrutinizer reporting can generate host-to-host reports to show the full extent of the host's communications with other IPs on the network. Any outbound traffic with remote hosts should be investigated by navigating to the **Reports** tab/section of the web interface and running *destination reports*.

Additionally, Plixer Endpoint Analytics may be able to provide MAC details for the host and report its own risk assessment based on internal algorithms, MS Defender, and Tenable.

4.2.3 Threat hunting

Plixer One Security can enhance any team's threat-hunting capabilities by providing them with centralized access to rich, contextualized data accounting for every host and conversation in a network.

Through Plixer Scrutinizer, Plixer One Security is also able to provide real-time alerts for generic malware and other anomalous traffic/activity, drive efficient workflows with its purpose-built UI, and integrate multiple threat intelligence functions. This gives teams the ideal starting point for their threat-hunting operations.

Overview

Plixer Scrutinizer plays two integral roles as part of a security team's threat-hunting program:

1. Collects traffic and host data for the entire environment (including *assets in the cloud*), storing hundreds of thousands of data points for investigations
2. Provides centralized access to all available data through various contextual views and reporting functions

This allows SecOps teams to efficiently search through and analyze device-level behavior and host conversations to search for suspicious activity and potential threats. Historical data can also readily be accessed to hunt for indicators of attack (IoA).

Visibility and workflow enhancements

Security teams using Plixer One Security can leverage the following functions and features to hunt for threats:

Alarm Monitor

The *Alarm Monitor* provides real-time alerts for anomalous behavior and other network activity violating Plixer Scrutinizer *Alarm Policies*. It functions as both a monitoring view for suspicious traffic and an interface for drilling into *activity timelines and individual Event artifacts, and more*.

Customized reports

To further investigate Alarms/Events, users are able to *run reports* that can be *tailored to their exact visibility requirements*. These reports can also be used to *drill deeper into specific data elements* to identify infected hosts or malicious activity.

Configurable detection mechanisms

Configuration options for *Flow Analytics algorithms* and the *Plixer ML Engine* allow users to tailor Plixer Scrutinizer's monitoring and detection functions to their specific requirements. This ensures that detections are always relevant and can greatly reduce investigation/response times for security teams.

Note: Plixer One Security includes additional detection techniques and mechanisms for security events.

Host indexing

With the *Host Indexing FA algorithm* enabled, a user is able to look up any IP address, find out whether or not the host has been seen on their network, and explore all activity associated with it. From the search results, the user can pivot directly to any applicable report and further investigate anomalous traffic originating from or targeting the host.

See also:

For additional details on incident response workflows with Plixer Scrutinizer, see [this](#) use case.

Workflows

The following workflows are sample scenarios where the functions/features bundled with Plixer Scrutinizer are used in threat-hunting activities:

Using host index to identify malicious IPs

Host indexing allows users to quickly look up IP addresses seen on the network, making it ideal for monitoring hosts that have exhibited anomalous or suspicious behavior.

Workflow

To search the host index for malicious IP addresses:

1. Navigate to **Explore > Search** in the web interface.
2. In the *Host Index* subtab, use the dropdown to switch to *Multiple* search mode.
3. Paste in the comma-separated list of IoC (Indicators of Compromise) IP addresses into the field.

4. Review the traffic direction, byte counts, and first/last seen details for each host and, if necessary:
 - Click on the hostname/IP to view additional traffic and alarm information associated with the host.
 - Run a report filtered on the host by clicking the data source and selecting a report from the tray.

Hint: If further investigation is required, continue to [refine the report configuration](#) as needed.

See also:

To learn more about configuring and refining reports, see [this use case](#).

Reviewing Alarm Monitor for suspicious hosts

The Plixer Scrutinizer Alarm Monitor provides users with real-time alerts to both performance issues and security threats and allows them to drill into event details by policy violation or by host.

Workflow

To inspect activity for suspicious hosts using the Alarm Monitor:

1. Navigate to **Monitor > Alarm Monitor** in the web interface.
2. Switch to the **Hosts** subtab and add a filter to show only *Critical* severity violations.
3. Use the dropdown to switch to the *Event Connections* view to look for hosts involved in multiple events.
4. Drill into events or run reports filtered on potential threats as needed.

See also:

To learn more about configuring and refining reports, see [this use case](#).

Investigating off-hour network activity

Plixer Scrutinizer's monitoring and reporting functions can isolate traffic outside business hours and alert teams to potentially malicious activity taking place during an organization's off-hours.

Workflow

To proactively hunt for threats that remain dormant during business hours, security teams can leverage the following report filter options:

- Add a filter that excludes business hours. A report threshold can also be configured, so that any activity exceeding the specified value(s) can be tracked via the Alarm Monitor.
- Define the period of time outside business hours as the report's time window/range.
- Set the report's time window to *Last 24 hours* and compare traffic data during and outside business hours.

Hint: After Plixer Scrutinizer has been deployed, default business hours can be set in the **Admin > Settings > Reporting** tray. These hours can be changed when configuring a business hours report filter.

Important: The Plixer ML Engine uses separate baseline models for network behavior during and outside of business hours. The default 8 am to 5 pm setting can be changed in the **Admin > Settings > Reporting** tray.

Identifying exfiltration outside business hours

Plixer Scrutinizer is able to isolate network activity outside of business hours, allowing teams to quickly identify data exfiltration attempts and other malicious activity taking place outside business hours.

Workflow

Data exfiltration can be identified proactively within Scrutinizer by identifying and reviewing traffic leaving your network. The *Explore -> Exporters -> By Interface View* is a great place to start, as traffic is displayed as inbound/outbound columns.

By default this is sorted so that your most congested interface is displayed at the top. This may be worth reviewing as large amounts of traffic leaving the network may be exfiltration. Even more likely, exfiltration happens in a “low and slow” attack approach where only small amounts of traffic leave the network periodically – avoiding causing spikes in traffic that may cause alarms.

Because inspecting individual interfaces one at a time is inefficient, *Plixer Scrutinizer Reports* can be used to narrow down the scope of information to be reviewed. This allows for a more streamlined approach to proactively searching for unwanted/suspicious traffic.

The following example uses the *Destination Countries with AS* report type:

1. Select *Reports -> Run Report -> Select Report Type* to start an adhoc report.
2. Choose *Destination Reports -> Countries with AS*, add the appropriate device(s), and run the report.

The report is likely to show multiple rows of autonomous systems and the corresponding country they are associated with.

Note: Class A, B, and C addresses are always classified as *Uncategorized* and will often include internal network addresses. In this scenario, these are likely associated with responses to internal destinations through outbound interfaces.

3. Help narrow your search by excluding traffic that you expect to see. What remains may be of use in identifying traffic leaving the network to a destination that is unintended.

When you have a subset of data that is more manageable, e.g., countries your organization does not do business with, you can begin to pivot to other report types. Changing the time frame or “zooming out” can also reveal possible threats in the form of suspicious traffic patterns.

4. Within your report, with same filters, set the timeframe to *Last Seven Days*.

Is there a ping every hour beaconing out? Same packet size of data leaving the network following a pattern?

At this point, your report likely has one or more country, AS, or host filters. Switching to another report type or using extended report options like host reputation or geo IP lookups can lead to additional insights.

Tip: *Run a report* against a core router that is likely to see a majority of your traffic. Alternatively, select 'All Devices' to identify top network conversations across the entire network.

4.2.4 Lateral movement detection

Because indications of a cyber attack are not limited to traffic originating from external hosts, security teams require tools that can monitor internal network activity for potential threats, such as lateral movement.

Plixer One Security employs multiple detection techniques to alert to behavior that may indicate lateral movement through their network by malicious actors.

Overview

Through Plixer Scrutinizer, Plixer One Security combines deep network observability with multiple approaches to lateral movement detection to deliver meaningful alerts that enhance both proactive and reactive workflows.

As it continuously monitors and collects flow data from its environment, Plixer Scrutinizer uses the [Alarm Monitor view](#) to alert users to activity that matches potentially problematic or malicious patterns, including those associated with lateral movement techniques. The Alarm Monitor, [Network Maps](#) and [Dashboards](#) views allow users to pivot to [Reports](#) and launch deeper investigations into typical indicators of lateral movement.

Hint: The [Monitor > Alarm Monitor > ATT&CK](#) tab classifies Alarms using the MITRE ATT&CK framework and can be used to quickly filter for alerts related to lateral movement.

The following [Alarm Policies](#) are used to provide alerts specifically for potential lateral movement and based on different detection approaches/criteria:

Lateral Movement

Lateral Movement Alarms are Flow Analytics detections that are triggered by traffic/activity that is indicative of techniques used to exploit remote services. [Events](#) under this Alarm Policy report the following details for the detection:

- Exporters/devices
- Violating hosts
- Target hosts

Lateral Movement Attempt

Lateral Movement Attempt Alarms are Flow Analytics detections that are triggered by traffic/activity that is indicative of a worm attack on a specific port on a target host. Events under this Alarm Policy report the following details for the detection:

- Type of worm
- Destination/target port
- Violating hosts
- Target hosts

Lateral Movement Behavior

Lateral Movement Behavior Alarms are machine learning detections that are triggered when the behavior of a [monitored host](#) deviates from baseline activity patterns in a way that is indicative of lateral movement. Events under this Alarm Policy report hosts that are communicating with an unusually large number of machines (based on behavior learned by the Plexier ML Engine) as violators.

Hint: The threshold at which irregular traffic/behavior associated with a host is reported as a detection can be adjusted by changing the sensitivity for the [ML Inclusion](#) it belongs to.

Workflows

The following workflows show how lateral movement detections in Plexier Scrutinizer can be used to investigate and respond to potential threats:

Investigating lateral movement alerts

Plixer Scrutinizer uses multiple lateral movement detection techniques, each of which corresponds to a separate Alarm Policy. This provides security teams with additional context on which to base their response strategies.

Workflow

After receiving a lateral movement alert in Plixer Scrutinizer (either directly from or via SIEM), investigate the event:

1. Navigate to **Monitor > Alarm Monitor** in the web interface and search for *Lateral Movement* (FA), *Lateral Movement Attempt* (FA), or *Lateral Movement Behavior* (ML) violations.
2. Click on an Alarm Policy to open the summary view and review the activity timeline and hosts involved.
3. Drill into an event artifact to view a summary of details for a violation associated with a specific host.
4. To further investigate the activity of the host, click on the icon next to its IP address or hostname, and select an automatically filtered report to run.

Hint: For additional context and/or details related to how and why the host was compromised, review all alarms leading up to the lateral movement violation.

Uncovering data exfiltration

While proactively reviewing outbound traffic, the security team discovers activity that indicates a potential attempt to exfiltrate data.

Workflow

After discovering unusually high outbound utilization in the **Explore > Exporters > By Interface** view, run a Report to narrow down the scope of traffic that needs to be reviewed (e.g., *Destination Countries with AS*):

1. [Run a new report](#) for the Exporters/devices exhibiting suspicious behavior, and select *Countries with AS* (under the *Destination Reports* category) as the report type. This will output a list of autonomous systems, along with the countries each one is associated with.

Note: Class A, B, and C addresses are always classified as *Uncategorized* and will often include internal network addresses. In this scenario, these are likely associated with responses to internal destinations through outbound interfaces.

2. Narrow down the scope of the report by dragging rows associated with expected traffic to the *Exclude* drop zone to the left and clicking *Apply* in the Filters tray.
3. After the report has been re-run with the additional exclusions, review the list for traffic bound for unusual destinations.
4. Once a more manageable subset of data (e.g., countries your organization does not transact with) has been achieved, refine the report to gain more insight:
 - “Zoom out” to look for activity patterns by changing the time frame covered by the report.
 - Inspect activity associated with the host, country, or autonomous system by clicking on it and pivoting to a different report type from the tray.
 - Leverage additional tools (under the *Other Options* category in the tray) to obtain additional information.

For further investigation, continue to modify the settings of the report to gain visibility into hosts, traffic, etc. that remain suspicious.

4.2.5 Incident response

Plixer One Security combines Plixer Scrutinizer’s deep, environment-wide visibility and intuitive UI-driven workflows with advanced detection techniques for security events to enhance a team’s ability to respond to threats.

Overview

Plixer Scrutinizer’s “single-pane-of-glass” feature set is designed around providing maximum network observability via synergistic web interface functions and views that streamline monitoring and investigative activities.

Full visibility supporting incident response and other security processes

As part of an incident response plan, Plixer Scrutinizer ensures that SecOps teams have access to all the traffic and device information they need for investigation and remediation:

- Get comprehensive, contextualized details for intrusion detection system (IDS) and intrusion prevention system (IPS) events
- Access full network traffic forensics to watch for and investigate security information management (SIM) events
- View full IP to MAC address mapping history for all connected devices and endpoints
- See real-time and historical endpoint context and location
- Assess endpoint risk through layer 2 historical location tracing
- Glean additional insights from detection details via *MITRE ATT&CK*, *STIX/TAXII*, and *other integrations*

Web interface functions that promote more efficient response strategies and procedures

Plixer Scrutinizer enables more efficient general security and incident response workflows through multiple functions/features, including:

- Highly configurable UI views (*Alarm Monitor*, *dashboards*, *network maps*, etc.)
- *Customizable data aggregation* from any observation point(s) on the network
- Detections and alerts driven by *AI/ML* and *Flow Analytics*
- *Customizable notification options* for Alarm/Event details
- Deep visibility for both on-prem devices and *assets in the cloud*
- *Collaborative features* that promote sharing investigation results/insights between members and/or teams

Workflows

The following workflows show how the additional visibility and workflow enhancements enabled by Plixer Scrutinizer can be leveraged by SecOps teams for monitoring and incident response:

Responding to Alarm Monitor security alerts

Plixer Scrutinizer leverages a range of technologies to alert users to anomalous and potentially malicious network activity through its library of Alarm Policies. Once policy violations are reported via the Alarm Monitor views, security teams can drill into individual event details to evaluate whether further investigation is necessary.

Workflow

To investigate an Alarm Policy (e.g., *Data Exfiltration*, *Data Accumulation*, etc.) violation (e.g. data e) reported in the Alarm Monitor:

1. Click on the Alarm Policy to open the summary view.
2. Review the activity timeline and hosts involved.
3. If further investigation is warranted, drill into individual event artifacts for more details.
4. Click the icon next to an IP address or hostname to run an automatically filtered report and examine additional activity/hosts associated with the event.

Hint: For additional context and/or details related to how and why the host was compromised, review all alarms leading up to the policy violation.

Scrutinizing an infected host

After a user is infected with a virus, the security team must identify what other hosts on the network may have communicated with the infected host.

Workflow

After the infected host is discovered/reported, the following steps can be used to identify other hosts it has interacted with:

Note: This workflow relies on usernames acquired from a network device (router, firewall, etc.) or through enabled integrations (e.g., Active Directory LDAP). If usernames are not available, host IP addresses can be used as identifiers instead.

1. Under **Explore > Exporters > Entities > Usernames**, search for the infected host/username and click on it. A new view will open.
2. Review the Alarms/Events associated with the host, which may include the following violations:
 - *P2P* and *Lateral Movement* (infected host may be attempting to extend access further into the network)
 - *TCP, UCP, XMAS Port Scan* (infected host may be pinging the network for reconnaissance)
3. [Create/run a report](#) with the username applied as a filter to identify all activity where the infected host was either the source or the destination of traffic. Ensure that the time range includes a period before the infection was reported or discovered.

Hint: When viewing information associated with a username, click the graph icon to run a report with the username applied as a filter. The filter will be retained even when pivoting to other report types.

4. Review the output or pivot to different report types for insight related to who, what, when, where, why, and how the infected host communicated on the network:
 - Protocols the host was seen using
 - Countries the host communicated with
 - Firewall events (through vendor-specific report types, e.g., ACL rules, NAT translations, etc.)
 - Destination FQDN reports
 - Activity associated with the host before and after the infection (for additional insight into the techniques used in the initial attack)
5. If the [Host Indexing FA algorithm](#) is enabled, navigate to **Explore > Search** to look up historical data associated with the IP address of the infected host. This information may provide additional insight based on typical communication patterns and reduce mean time to know (MTTK) during the investigation.

Note: If the *Use Host Index* option under **Admin > Settings > Reporting** is enabled, *Group* and *All Device* reports will use the host index to limit the scope of exporters checked when a host filter is applied.

FEATURES AND FUNCTIONALITY

This section contains information on Plixer Scrutinizer's main functions and includes guides for their configuration and use.

5.1 Plixer Scrutinizer web interface

The Plixer Scrutinizer web interface is the primary means of accessing the system's functions and features. It can be accessed by pointing any supported browser to https://scrutinizer_ip/ui/.

Hint: To access the Plixer Scrutinizer Classic UI, use the URL https://scrutinizer_ip/oldui/. The preferred UI can also be set from within the web interface under the user menu.

This section introduces the Plixer Scrutinizer's web interface, explains the basic concepts behind its services, and outlines the main workflows within the UI.

5.1.1 UI overview

Plixer Scrutinizer is equipped with a extensive feature set that allows it to transform raw network flow data into timely, accurate, and fully contextualized intelligence for modern NetOps and SecOps teams.

The Plixer Scrutinizer web interface acts as the system's primary console and is divided into four main sections/tabs corresponding to essential NetOps and SecOps workflows, plus an administration tab for environment configuration and management.

| <i>Monitor</i> | <i>Explore</i> | <i>Investigate</i> | <i>Reports</i> |
|--|--|---|--|
| <ul style="list-style-type: none">• Use customizable Alarm Policies to receive alerts when problematic or dangerous behavior is discovered on the network• Create custom Dashboards using ready-to-use Gadgets that display vital activity summaries and visualizations• Visualize and monitor activity between connected devices with user-defined Network Maps | <ul style="list-style-type: none">• Drill down into flow-generating devices to examine activity, resource usage, and Events generated• Inspect behavior, interactions, and Events generated by individual entities• Look up specific host and host pairs in the system's Host Index to inspect details or verify | <ul style="list-style-type: none">• Define Collections of one or more Alarms, Events, and/or Reports and assign them to analysts for investigation• View available Forecasts to identify resource usage trends and identify future needs | <ul style="list-style-type: none">• Create/run custom or preconfigured network activity Reports that can be saved and used to generate ML-based Forecasts• View/re-run and manage saved Reports |

The functions and workflows under each UI tab are explained in further detail in the succeeding sections of this documentation.

Hint: Click on the the *Help* (?) button in the web interface's navigation bar to open the Plexier Scrutinizer this documentation at any time.

5.1.2 Monitor

The **Monitor** section/tab of the web interface provides access to various high-level NOC/SOC summary views that are actively updated by the system.

The section is divided into three views. To select a view, hover over the **Monitor** item in the menu bar, and then select one of the following:

Alarm Monitor

The **Alarm Monitor** page is Plixer Scrutinizer's main interface for reporting Alarms and Events. The page is divided into three subtabs to support multiple avenues for investigating system activity and anomalous network traffic.

For additional background and recommended configuration steps related to Alarm Monitor functions, see the [Alarms and Events configuration guide](#).

Policies

The **Monitor > Policies** tab lists all Alarm Policies with observations during the specified date/time range. It is the default view of the Plixer Scrutinizer Alarm Monitor.

Hint: Because the date/time range setting uses observation timestamps as its filter, the list will include Alarm Policies with Events that started within the selected time frame, even if the most recent observation was outside of it. To learn more about Alarms, Events, and observations, see [this topic](#) under the [Alarms and Events configuration guide](#).

By default, the **Policies** tab displays the *Event Cards* visualization, which can be used to filter the list by Event Severity, in addition to the list of Alarm Policies.

The following visualizations/shortcuts are also available via the *View* dropdown:

- Event Cards
- Policy Cards
- Host Cards

- Event Timeline
- Policy Timeline
- Host Timeline
- Policy Connections
- Tactic Connections
- Category Connections
- Event Connections

Alarm Policy list

The main table of the **Policies** tab also includes the following details for each Alarm Policy:

- Distribution of individual Events reported under the Policy, based on Severity
- Total number of violators that have triggered Events under the Policy
- Total number of targets in Events reported under the Policy
- Timestamps of the original and most recent Events linked to the Policy
- Policy category

Hint: Additional details/columns can be toggled on for each Alarm Policy via the *Available Columns* button.

Additional actions/options

- The table/list can be sorted using any of the displayed details by clicking on the corresponding column header.
- [Acknowledge](#) one or more Alarms by ticking their checkboxes and clicking *Acknowledge Selected Events*
- Click the shortcut next to an Alarm Policy's *Violators* or *Targets* to open a tray listing all hosts that have been violators or targets under the Policy.
- Add an Alarm (and all Events under it) to the current active [Collection](#) by clicking the star button and selecting *Add to Collection*.

Hint: To set a different Collection as active or create a new Collection, select *Manage Collections*.

- To change an Alarm Policy's settings, open the three-dot menu and select *Edit Settings*.

Alarm Summary

Clicking an Alarm Policy in the main table opens a summary page that consists of a graph or chart and a list of Events/artifacts in table format.

The following visualizations can be selected from the dropdown:

- **Events Scatter Plot** - Shows a visual distribution of the Events and Observations.
- **Events Timeline** (default) - Shows the original and most recent Event timestamps, as well as the number of times the Event was triggered within the given period
- **Entities** - Shows a summary of the Top Violators, Top IP Groups, and Top Targets

Event List

Each Artifact in the Event List links to a summary tray containing all relevant information for the Event (severity, hosts, etc.).

Note: For additional information on Severity, see the [Alarm Policy settings](#) section.

Hovering over the graph icon in the Events List table displays the following options:

- **Explore Event Traffic** - Generates a host-to-host Report for the selected Artifact
- **Export Targets** - Exports the Artifact's list of target hosts as a CSV file
- **Export Violators** - Exports the Artifact's list of violating hosts as a CSV file

Note: Depending on the Alarm Policy, certain options may be absent from the three-dot menu. *Explore Event Traffic (Host to Host Report)* requires the flow data that triggered the Event to be available. *Export Targets* and *Export Violators* require the corresponding host type to be part of the Policy's criteria.

Hosts

The **Hosts** tab of the Alarm Monitor page displays all hosts that are involved in the Events within the designated time period.

The **Event Cards** header visualization is displayed by default in the **Hosts** tab. To switch to a different visualization, click the **View** menu, and then select one of the following:

- Event Cards
- Policy Cards
- Host Cards
- Event Timeline
- Policy Timeline
- Host Timeline

The **Hosts** tab includes a table that shows the following details for each host:

- Distribution of individual Events targeting the host based on their Severity
- Total number of times the host was a target in an Event
- Total number of times the host was a violator in an Event
- Total number of Alarm Policies violated by the host
- Timestamps of the original and most recent Events targeting the host

The following actions are also available in the **Hosts** tab:

- **Hide/Show Columns** - Click the **Hide/Show Columns** icon beside the view mode menu to select which columns to hide or display in the table.
- **Filter Hosts by Severity** - Click one of the color-coded Severity sparkline cards to display only the hosts with the selected Severity.
- **Sort Hosts by Severity** - Click the **Severity** column header in the table to sort the hosts based on their severity.
- **Sort Hosts by Risk** - Click the **Risk** column header in the table to sort the hosts based on their risk level.
- **Add Host to a Collection** - Hover over the star icon and then click the plus button to add the Host and all the Alarms and Events associated with to the active Collection. To switch to a different active collection, hover over the star icon, click **Manage Collections**, and then select a different Collection. For more information on Collections, see the [Investigate](#) section.
- **View more host details** - Hover over the three-dot icon and then select any of the following options: Go to Host View, View Information, View Endpoint, Filter on Host, or Run Report.

Alarm Summary

Each host address in the table links to the Alarm Summary page of the host. The Alarm Summary page has the following views:

- **Alarms** - Displays all Alarm Policies associated with the Events targeting the host as well as other relevant details. This is the default view mode when accessing the Alarm Summary page of the host.
- **Traffic** - Displays the Activity Timeline, the Source IP Groups, top applications, and the Destination IP Groups.

To view information related to the host, click the tray button in the Information section in the **Alarms** view. The tray button opens a quick-access tray that displays the DNS Name, IP Address, and other relevant information. Clicking the **Learn More** button navigates to the **Traffic** view in the **Hosts** tab.

Note: If Plixer Endpoint Analytics (integration) is enabled for the host, an Endpoint section is displayed below the Information section in the **Alarms** view. Clicking the tray button in the Endpoint section opens a quick-access tray that displays the Mapped IP Address, MAC Address, Current Location, System Location, and other information related to the host. Clicking the **Investigate Endpoint** button opens a new tab to the Endpoint Analytics web interface.

In the Alarm Summary page, the **Integrations** menu is available beside the filters icon. The **Integrations** menu displays the following options:

- **GEO IP** - Opens a new window and displays the geographic location information of the associated alarm
- **Alarms** - Navigates back to the **Policies** tab
- **Talos Reputation Center** - Opens a new window displaying the *Reputation Lookup* results of the associated alarm

ATT&CK

The **ATT&CK** tab of the **Alarm Monitor** page uses the MITRE ATT&CK framework to classify malicious Events. Alarms are reported in an Event timeline and sorted into separate lists by MITRE ATT&CK Tactics and Techniques.

Events that align with ATT&CK Tactics and Techniques will be in a category that represents the Tactic, Technique, and Sub-Technique. The ATT&CK view breaks those Plixer Scrutinizer categories down further into individual Techniques and provides a visual representation of how those Techniques were observed over time.

In the ATT&CK view, the Tactics are listed horizontally along the top of the chart, with the Techniques listed vertically below the associated Tactic.

Event timeline actions

- Mouseover - Shows the timestamps for the initial and most recent Events in the block, their severity, and the number of times the Event was observed during the timeframe
- Click - Pulls out a quick-access tray containing links to the Policies and hosts associated with the Alarm as well as MITRE ATT&CK information (tactic, technique, and sub-technique) for the Event(s)

Event category list actions

- Mouseover - Shows basic MITRE ATT&CK information (tactics, techniques, and sub-techniques) for the Event(s) as well as the number of times the Event was observed
- Click - Filters all Alarm Monitor views to show only information for the selected Event block

Hint: Event blocks in both the timeline and the category lists are color-coded based on their severity.

© 2022 The MITRE Corporation. This work is reproduced and distributed with the permission of The MITRE Corporation.

Applying filters

By default, the different Alarm Monitor views are set to show all un-acknowledged Alarms, sorted by severity. This can be changed by applying one or more filters.

Note: Any filters configured will be applied to all Alarm Monitor views until they are reset using the *Clear Filters* button.

Filter categories

Clicking the *Filters* button opens the Filtering Options tray, which is divided into the following categories:

- Policy
- Severity
- Risk
- Hosts
- Violators
- Targets
- Alarm Policy category

To apply a filter, expand the category it falls under and select on or more criteria. Multiple criteria from different categories can be applied at the same time.

Hint: To also display Events that have already been acknowledged and hidden, use the *Show Acknowledged Events* toggle.

Time period covered

To show Alarms generated within a specific time period, click the calendar icon and use the dropdown to select from the following preset ranges:

- Last 5 minutes

- Last 15 minutes
- Last hour
- Today
- Last 24 hours
- Yesterday
- Last 7 days
- Last week
- Last 30 days
- Last month

Custom *From* and *To* ranges can also be applied by selecting *Custom* from the range dropdown.

Exporting filtered tables

When exporting Alarm and Event details using either selection in the Options tray, all active filters (including the date and time period) are automatically applied.

Hint: The *Export CSV (All)* option exports all Alarm and Event data, regardless of the current *Show Entries* setting in the Options tray.

Acknowledging events

Once an event has been investigated and/or resolved, it should be acknowledged to clear it from all Alarm Monitor views. This reduces the volume of active Alarms and/or Events at any given time and can further streamline investigative processes.

Acknowledging events is part of Plixer Scrutinizer's recommended investigation and resolution workflow.

Hint: To show/hide acknowledged event in the Alarm Monitor views, open the filter options tray and toggle the *Show Acknowledged Events* option on/off.

Acknowledging can be done by alarm policy or by event.

Acknowledging by policy

From the main view of the **Policies** tab, acknowledging an alarm policy automatically flags all events generated under the Policy as acknowledged.

To acknowledge by alarm policy:

1. While on the **Policies** tab of the Alarm Monitor view, select the Policy by ticking its checkbox.
2. If acknowledging more than one Policy, verify that the correct Policies have been selected.
3. Click *Acknowledge Selected Events*.

Note: The *Acknowledge Selected Events* button is only available when at least one Policy checkbox is ticked.

Once acknowledged, the alarm policy and all events associated with it will be hidden from all Alarm Monitor views.

Acknowledging by event

Acknowledging can also be used to clear only events that match the same criteria. This allows other events under the same policy (as well as the alarm policy itself) to be retained in Alarm Monitor views.

Events are acknowledged from the summary view of the **Policies** tab as follows:

1. Scroll down to the *Event List* section of the page.
2. Select the artifact linked to the criteria/events to be acknowledged by ticking its checkbox.
3. If selecting more than one artifact, verify that the correct checkboxes have been ticked.
4. Click *Acknowledge Selected Events*.

Note: The *Acknowledge Selected Events* button is only available when at least one policy checkbox is ticked.

Once acknowledged, the event(s) will be hidden from all Alarm Monitor views.

Dashboards

Plixer Scrutinizer **Dashboards** are fully customizable views that provide convenient, “at-a-glance” access to multiple network metadata summaries in a single page. Multiple unique dashboards can be created/saved to enhance different workflows (using *gadgets*) and/or define user roles and responsibilities.

This section discusses the features and functions accessed via the **Monitor > Dashboards** tab/section of the web interface and includes further details related to the creation, customization, and management of dashboards.

Dashboard gadgets

Each Plixer Scrutinizer dashboard can be tailored to a specific task, workflow, or user through its gadget configuration. All gadgets automatically refresh to display the most up-to-date information and can also be clicked to access more detailed views.

Hint: Gadgets can be manually updated outside of their automatic refresh times (displayed in the gadget header) by clicking the refresh button.

Gadgets are divided into several types that can be added to dashboards in any combination:

Core gadgets

Plixer Scrutinizer ships with a core library of general-purpose gadgets that can be added when [creating](#) or [editing](#) a dashboard. These include gadgets for monitoring system health and performance, in addition to those for tracking important network information.

Report gadgets

Any Plixer Scrutinizer report can be added to dashboards after it has been [exported as a gadget](#) from the [output/results view](#). This enables the creation of dashboards that are uniquely customized to monitor any aspect of network performance or behavior.

After a report has been exported, it will be added to the list of available gadgets when [creating](#) or [editing](#) a dashboard.

To learn more about creating and configuring reports, see the [Reports](#) section.

Network maps as gadgets

After a spatial or geographical map is created, it is automatically made available as a dashboard gadget. If the network map is reconfigured at a later time, the gadget will also be updated to reflect any changes made.

When [creating](#) or [editing a dashboard](#), all existing network maps will be included in the list of gadgets that can be added.

To learn more about creating and configuring network maps, see the [network maps](#) section of this documentation.

External gadgets

External gadgets can be imported via their URL and added to any dashboard:

1. While in dashboard edit mode, click *Add Gadget*.
2. Click **New Gadget**.
3. Enter the following details for the gadget:
 - Name/label for the gadget
 - Gadget URL
 - Refresh interval for the gadget (in minutes)

Once a gadget has been imported, it becomes available for use in other dashboards.

Note: URLs for external gadgets must include the `http(s)://` prefix to avoid a 404 error. Additionally, certain gadgets may not load if you specify HTTP content when Plixer Scrutinizer is using HTTPS.

Hint: Access to gadgets can also managed via [usergroup permissions](#).

Feature-based gadgets

Certain gadgets bundled with Plixer Scrutinizer provide additional visibility when specific features are enabled/configured. These include gadgets that complement optional integrations, such as Plixer FlowPro, or leverage additional flow data forwarded by specific devices.

Creating a new dashboard

Creating multiple dashboards allow users to switch between different unique views to segregate monitoring requirements and workflows.

To create/add a new dashboard, follow these steps after selecting *Dashboards* from the **Monitor** submenu:

1. Click the *New Dashboard* button and select the *Create* tab in the popup.

Hint: Select the *Copy* tab to create a copy of an existing dashboard to use immediately or modify.

2. Enter a unique name for the dashboard.
3. (Optional) To set the dashboard as the default for the current user, tick the *Default Dashboard* checkbox.
4. (Optional) To lock the dashboard (cannot be edited), tick the *Read Only* checkbox.
5. Use the arrow buttons to add gadgets to the new dashboard from the list of those available (gadgets can be added or removed later as needed).
6. When done, click *Save* to create the new dashboard with the selected gadgets.

Once the dashboard has been created, it will replace the current view and can be accessed from the dashboards dropdown at any time.

Hint: When selecting dashboards from the dropdown, click the *Click For Menu* button to switch to a tile view showing all available dashboards.

Editing a dashboard

Settings and gadgets for existing dashboards can be modified as long as it has not been set to read only.

To enter edit mode for the current dashboard, click the *Edit Dashboard* button.

Once in edit mode, the following options are available:

- **Basic settings** - Click the pencil icon to change the dashboard name and toggle its default and read only settings. The dashboard can also be deleted from this menu if it is not locked or set as the default.
- **Dashboard Options** - Click the meter icon to create a new dashboard, and copy/delete the current dashboard. User and user group access options can also be accessed from this menu.
- **Add Gadget** - Click *Add Gadget* to configure the dashboard with additional gadgets. Selecting gadgets from the list will automatically add them to the dashboard, while clicking *New Gadget* will bring up a form to import a gadget via a URL. Clicking the checkbox next to a gadget enables the *Delete Selected* button to delete the selected gadget/s.
- **Move Dashboard Gadget** - Mouse over a gadget and click and hold the directional cursor to move the gadget to a new position within the dashboard.
- **Remove Gadgets** - Mouse over a gadget and click the - icon to remove it from the current dashboard.

After making the necessary modifications to the dashboard, click *Exit Edit* to save the current configuration and return to the standard view.

Dashboard management

Dashboard management functions can be accessed either via the *New Dashboard* button or by going into edit mode for the current dashboard and clicking the meter (Dashboard Options) button.

Managing user/usergroup access

To modify user or usergroup dashboard access, follow these steps:

1. Click either the **User access to dashboards** or **Usergroup access to dashboards** button in the **Dashboard Options**.
2. Select a user or usergroup from the dropdown.

3. Use the arrow buttons to move dashboards between the *Visible* (accessible) list and the *Available* (all Dashboards) list.
4. Click *Save* to save the current settings and return to the previous dashboard.

Once the access settings are saved, the selected user or user group is displayed in the dashboards dropdown.

Deleting dashboards

If it has not been set as the default or read only, the current dashboard can be deleted from the basic settings menu (pencil icon).

To delete multiple dashboards from the web interface, follow these steps:

1. Click the *Edit Dashboard* button, and then click **Dashboard Options**.
2. Click either the **User access to dashboards** or **Usergroup access to dashboards** button.
3. Select the dashboards to delete from either list. The dashboards should not be any user's default or set to read only.
4. Click the *Delete* button to permanently delete all selected dashboards and return to the previous page.

Note: The *Delete* button will only be available if it is possible to delete the selected dashboards (not any user's default or set to read only).

Deleted dashboards will no longer be accessible by any user or usergroup and removed from the system entirely.

Network maps

The **Monitor > Network Maps** page is used to view and manage network topology visualizations or *Network Maps*, which are automatically generated from user-defined *Mapping Groups*.

Important: Mapping Groups and IP Groups are separate Plixer Scrutinizer *group types* and have different roles/functions.

This section contains guides and background information related to the use and management of Network Maps and Mapping Groups.

Maps and objects

Network Maps can be generated as one of two types:

- **Spatial Maps** allow *Map Objects* to be positioned and/or grouped in any layout. Custom connections, objects, and labels can also be added and overlaid against uploaded backgrounds (e.g., wiring cabinet/diagram) to create detailed environment representations.
- **Geographical Maps** have Map Objects automatically laid out according to longitudinal and latitudinal coordinates, which can be entered manually or acquired through a Google Maps GPS lookup of the entered address. This type of map can help quickly identify devices with issues, even when there are multiple topologies and/or segments dispersed across different physical locations.

Important: For Geographical Maps to function correctly, a Google Maps browser API key must first be set up under **Admin > Settings > Mapping Groups > Global Settings** (Classic UI only).

When *creating a new map*, a wizard walks the user through the process of *defining the Mapping Group* and adding *Connections*.

Existing maps can be viewed and edited from the main Network Maps page and/or added to *Dashboards*.

Map Objects

To create a new Network Map, the Mapping Group it will be based on must first be populated with Map Objects of the following types (in any combination):

- Devices/Exporters
- Other Mapping Groups
- Custom Map Objects (Spatial Maps only)

In Spatial Maps, Map Objects can be dragged and dropped anywhere on the background to create the desired layout. Objects in Geographical Maps will automatically be repositioned after an address or coordinates have been associated with them.

Important: After a map is first generated, the default zoom may show Map Objects positioned closely together/on top of each other. Object positions will be updated and saved after they have been moved (either manually or based on GPS location).

When clicked, Map Objects also function as shortcuts to the **Explore > Exporters** page with a device filter applied (Spatial Maps only).

Custom Objects

To allow for greater detail and accuracy in Spatial Maps, custom objects may be added to the corresponding Mapping Groups.

Note: Custom objects can be created when *editing a Network Map* or via the Admin > Settings > Mapping Objects page. When adding Map Objects to a *new map*, only existing custom objects can be added to the Mapping Group.

Custom objects can be displayed as an icon (similar to regular Map Objects) or text and configured with the following properties:

| Icon object properties | |
|------------------------|--|
| Icon | Icon used to represent the object in the Network Map view Note: Additional icons in the format of <code><name>_object.gif</code> can be uploaded to the <code>~/scrutinizer/html/images/maps</code> directory |
| IP address | IP address to associate with the object (if applicable) |
| Label | Label to display for the object in the Network Map view |
| Additional notes | Additional notes to associate with the object |
| Link | (Full) URL to open when the object is clicked in the Network Map view |

Hint: The default map icons for devices/Exporters can also be replaced when modifying an object's properties. Additional icons in the format of `<name>_red.gif` and `<name>_green.gif` (both required to display the *up* and *down* device statuses) should be saved to the `~/scrutinizer/html/images/maps` directory. For best results, images with transparent backgrounds should be used.

| Text object properties | |
|------------------------|--|
| Label | Label to display for the object in the Network Map view |
| Shape | Background shape for the text object |
| Type | <p>Only applies when <i>Square</i> is selected as the shape</p> <p><i>Text</i>: Selected background color is used as a highlight for the label</p> <p><i>Background</i>: Label is set against the background shape and color with the specified dimensions</p> |
| Dimensions | Dimensions of the background shape in px (height, radius, etc.) |
| Link | (Full) URL to open when the object is clicked in the Network Map view |
| Color | Color of the background shape |

Once added to a Network Map, custom objects function in the same way as regular Map Objects and can be repositioned or [reconfigured](#) as needed.

Connections

Once a Network Map has been populated, Connections can be added to represent links between Map Objects.

Note: Connections can be added to Network Maps via the [map creation wizard](#), while in [Map Edit mode](#), or from the mapping configuration trays under [Admin > Settings](#).

Connections can be one of three types:

| | |
|---------------------|---|
| Interface | <p>Capable of showing the following color-based states for flow-capable interfaces on the source device:</p> <ul style="list-style-type: none">• <i>Green/yellow/orange/red</i>: Utilization based on thresholds set under Admin > Settings > System Preferences• <i>Blue</i>: No bandwidth statement available• <i>Dashed gray</i>: No flows received from the device in the last 5 minutes <p>Additional flow-related information can be viewed by clicking on or hovering over the Connection.</p> |
| Line | Static link with a specified label and color |
| Saved Report | <p>Functions as a shortcut to run a saved Report for the connected Map Objects and can be configured with custom utilization thresholds</p> <p>Note: A Report must be created/saved before it can be associated with the Connection.</p> |

Connections are unique to the Network Map they were added to and cannot be re-used in other maps/-groups.

Additional map settings

After a Network Map has been created, the following configuration options can be accessed by clicking *Map Settings* while in [Map Edit mode](#):

| | |
|---------------------------|--|
| Auto-add devices | Automatically adds devices with similar resolved hostnames based on the entered regular expressions (RegEx) |
| Pass status | When enabled, allows the status of a map/group to be reflected in its icon when it is used as a object in other Network Maps |
| Truncate labels on | Shortens Map Object labels by omitting the entered string |
| Background | Replaces the default map background with the uploaded image |

Creating a new map

To create a new Network Map, navigate to **Monitor > Network Maps** in the Plexier Scrutinizer web interface and follow these steps:

1. Click *New Map* and select whether to create a Spatial Map or a Geographical Map.
2. Enter a name for the map and click *Next*.

Note: Clicking *Next* automatically saves the Network Map, even if no other steps are completed in the wizard. Further changes to the map configuration can be made in [Map Edit mode](#) or via the Mapping Groups and Mapping Objects pages under **Admin > Settings**.

3. Select the Map Objects to include in the map and use the arrow buttons to add them to the Mapping Group. To filter the list of available Map Objects, use the dropdown and/or search field.
4. Verify that the correct Map Objects have been added and then click *Next*.
5. Use the dropdowns to define the details of the Connection:
 - a. *From/To*: Select the objects to link.
 - b. *Type*: Select the *type of Connection* to create.
 - c. *Color*: Select a color for the Connection.
 - d. *Label*: Enter a label for the Connection (optional).
6. Click *Add Connection* to save the Connection. Repeat the steps to configure additional Connections.
7. Click *Save* to save the current Network Map configuration.

Once created, Network Maps can be viewed at any time from the main **Network Maps** page.

Hint: Additional Map Objects and Connections can also be [added at a later time](#).

Viewing network maps

Existing Network Maps can be accessed from the **Monitor > Network Maps** page (where they can also be *edited*) or viewed in *Dashboards*.

Network Maps page

To open an existing Network Map, select the map from the tree/explorer view of the **Network Maps** page.

Note: If a default map has not been set, the main maps page will display all available maps in tiled or list format.

Maps/groups that have been included in other maps can also be accessed by expanding their parent groups. The default *Ungrouped* group can be used to view all devices/Exporters that have not been assigned to other maps.

To apply a filter to the map/group tree, enter a string to match in the search field and click the search button.

Hint: The icons in the Network Map tree indicate whether the map is a Spatial Map or a Geographical Map. Click the gear icon to access additional options and actions for that map or object, including running Reports and viewing additional details.

The following settings/options can be accessed from the map view pane:

| | |
|-----------------------|---|
| Update map | Force a map refresh with the latest available flow data |
| Toggle details | Switch between IP addresses/DNS hostnames and bitrate/utilization |
| Adjust zoom | Change the zoom level (in, out, reset, stretch) of the map |
| Edit map | Switch to <i>Map Edit mode</i> to modify map membership and/or settings |

Network Maps in Dashboards

To allow for more efficient monitoring workflows, Network Maps can also be added to Plexier Scrutinizer Dashboards, where they can be viewed alongside other *Dashboard Gadgets*.

While *creating* or *editing* a Dashboard, use the dropdown in the Gadget selector to filter for the *Maps* category and select the Network Maps to add. Each Dashboard can include multiple maps.

Reconfiguring maps

Network Map configuration and management functions can be accessed from the main **Monitor > Network Maps** page. Alternatively, the Mapping Groups and Mapping Objects pages, which can be found under **Admin > Settings** in the web interface, can also be used.

Map Edit mode

Settings and membership for an existing Spatial Map can be further modified by entering Map Edit mode while viewing the map. New Map Objects and Connections can also be added while this is toggled on.

Hint: While in Map Edit mode, right-clicking on an object brings up shortcuts for commonly used map editing functions.

In addition, Map Edit mode provides access to the following layout tools:

- Manual single- or multi-object repositioning
- Automatic object alignment
- Object layering (*bring to front, send to back, etc.*)

After making changes, click the *Save Map* button to save the current layout and settings.

Note: To access map and object editing functions for Geographical Maps, left-click on an object or hover over the map in the tree/explorer view and click the gear icon. Objects in Geographical Maps cannot be manually repositioned.

Mapping Group/Object admin pages

Network Maps and/or Map Objects can also be configured via the Mapping Groups and Mapping Objects pages in the [Admin > Settings](#) section of the web interface.

These map management pages may be preferred when setting up or modifying multiple maps or objects without need to view the actual maps.

Important: Spatial Map object layouts can only be modified using Map Edit mode.

5.1.3 Explore

The **Explore** section/tab of the web interface is Plixer Scrutinizer's main inspection and management console for individual devices.

This section contains guides and additional information covering the content and functionality of each of the main **Explore** page's three tabs:

Exporters

The **Explore > Exporters** tab can be used to inspect traffic and Alarm information for individual Exporters. Exporters can be viewed either *By Interface* (default) or *By Exporter*.

Clicking Exporter hostnames/addresses in either view opens a summary page containing activity timelines, Alarm history, and other information associated with the Exporter.

By interface view

In the *By Interface* view, the table lists inbound and outbound activity for each interface, in addition to the Exporter it belongs to.

Hint: The *Inbound* and *Outbound* columns will display utilization percentage for any interfaces whose speeds have been defined. Otherwise, actual rates (in b/s) will be shown instead.

Clicking an interface or the + icon opens the interface details tray with the following options:

- **Reports** - Opens a secondary tray from where available *Reports* for the device can be run
- **Information** - Shows additional details about the interface
- **Exporter** - Opens the interface's Exporter summary page
- **Reset Highwater Inbound** - Resets the highwater mark details for inbound traffic
- **Reset Highwater Outbound** - Resets the highwater mark details for outbound traffic
- **Reset Highwater Both** - Resets the highwater mark details for both inbound and outbound traffic

Clicking the interface to expand the details tray also opens a secondary tray from which *Reports* can be run.

Hint: Selecting one or more instances using the checkboxes opens a bulk actions tray containing options to reset inbound and/or outbound highwater mark details.

By Exporter view

In the *By Exporter* view, Exporter hostnames/addresses are displayed alongside the following details:

- Number of device groups (*SpatialMaps or GeoMaps*) the Exporter belongs to
- Number of interfaces
- Packets per second over the last 12 hours
- Flows per second over the last 12 hours
- Timestamp of the most recent flow received from the Exporter

Clicking the + icon opens the Exporter details tray with the following options:

- **Reports** - Opens a secondary tray from where available *Reports* for the device can be run
- **Information** - Shows additional details about the device
- **Interfaces** - Lists all interfaces associated with the device
- **Integrations** - Lists any third-party integrations using the device
- **Tags** - Lists any custom tags added to the device and allows the user to add new tags
- **Mapping** - Allows the user to view and/or edit the device's mapping properties
- **Admin** - Contains links to device admin pages

Hint: Selecting one or more Exporters using the checkboxes opens a bulk actions tray containing options to run *Reports*, add tags, and modify grouping/mapping properties. In the interface view, the quick access trays will include the option to reset the inbound and/or outbound highwater values.

Tree menu functions

The tree menu/view of the **Explore > Exporters** tab lists all configured device groups/maps.

Expanding a device group in the tree view will show the following options:

- **Actions** - Opens an actions tray with options to run [Reports](#), display the group's network map, and filter the main table by the group's Exporters or interfaces
- **Modify** - Opens a configuration tray where the device group's properties, including Exporters or connections, can be modified
- **Exporters** - Shows the number of Exporters assigned to the group and applies the filter to the main table when clicked
- **Interfaces** - Shows the number of interfaces linked to the Exporters assigned to the group and applies the filter to the main table when clicked

Additional page options

- Filtering options can be accessed by clicking the *Filters* button.
- General page options, such as the number of entries shown and export actions, can be accessed by clicking the *Options* (gear) button.

Entities

The **Entities** tab shows all individual entities on the network—both user-defined and discovered—and sorts them into separate subtabs to streamline searching and inspection. Entities also link directly to a detailed summary pages, which contain activity timelines, statistics, and shortcuts to run [Reports](#).

The following subsections explain the contents of each subtab as well as any additional functions under it:

- [Usernames](#)
- [Applications Defined](#)
- [Hosts](#)
- [Autonomous Systems](#)
- [IP Groups](#)
- [Countries](#)
- [Protocols](#)

Hint: The gear button opens a quick-access tray with global settings for the **Entities** tab and an options to export the data as a CSV or PDF file and refresh the page. The filter button opens a separate tray with filtering options for the entity lists.

Username

The **Username** subtab displays all username-host pairs along with the following details:

- Data source
- Machine name (if available)
- First seen timestamp
- Last seen timestamp

Clicking on either a host address or username will open a summary page for that entity.

Applications Defined

The **Applications Defined** subtab shows all defined applications communicating on the network along with the following details:

- Port number
- Number of Exporters using the application
- Number of flows sent by the application
- Packet rate
- Bit rate

A link to the admin page for adding/defining applications can also be accessed from the **Applications Defined** subtab.

Hosts

The **Hosts** subtab shows all discovered hosts and can be set to list only sources, only destinations, or source-destination pairs.

The following details are also included in the table, regardless of the selected view:

- Number of Exporters
- Number of flows
- Packet rate
- Bit rate

By default, the host summary page is opened on the tab corresponding to the specified role (e.g., destination tab when clicking on the host as a destination), but the other tabs remain accessible from the same page.

Autonomous Systems

The **Autonomous Systems** subtab lists all autonomous systems (ASs) sending or receiving packets through the network. The list can be set to show only sources, only destinations, or source-destination pairs.

The table also includes the following details for each AS listed:

- Number of Exporters
- Number of flows
- Packet rate
- Bit rate

By default, the host summary page is opened on the tab corresponding to the specified role (e.g., destination tab when clicking on the AS as a destination).

IP Groups

The **IP Groups** subtab lists all IP Groups configured under the Plixer Scrutinizer environment and can be set to show only sources, only destinations, or source-destination pairs.

The table also includes the following details for each IP Group listed:

- Number of Exporters
- Number of flows
- Packet rate
- Bit rate

By default, the IP Group summary page is opened on the tab corresponding to the specified role (e.g., destination tab when clicking on the IP Group as a destination).

Countries

The **Countries** subtab shows all countries discovered on the network and can be set to display only sources, only destinations, or source-destination pairs.

The table also includes the following details for each country listed:

- Number of Exporters
- Number of flows
- Packet rate
- Bit rate

Note: The *UNKNOWN* listing and all details provided for it includes all country entities that the system was unable to identify.

By default, the country summary page is opened on the tab corresponding to the specified role (e.g., destination tab when clicking on the country as a destination).

Protocols

The **Protocols** subtab shows all communication protocols being used by devices on the network.

The table also includes the following details for each protocol listed:

- Number of Exporters
- Number of flows
- Packet rate
- Bit rate

Links to the admin pages for whitelisting protocols and defining exclusions can also be accessed from the **Protocols** subtab.

Search

The **Explore > Search** view allows users to search the Plixer Scrutinizer host index to quickly verify whether or not a host has been seen on the network.

Searches can be performed for either individual hosts or pairs (host to host). Simultaneous lookups for multiple hosts or pairs are also supported.

Important: To be able to search for hosts and host pairs, the corresponding indexing feature must be *enabled*.

- Traffic direction (inbound, outbound, A > B, B < A, bidirectional)
- First and last seen timestamps
- Exporter/source of collected data
- Bytes in and out
- Packets in and out
- Flows in and out

Hint: To show fewer details in search results, click the the table button and untick the checkboxes for the columns to be hidden.

In the search results, drilling into a host will display a summary of its activity on the network. Clicking on a data source opens a tray that allows the user to quickly pivot to any supported report type.

Enabling host indexing

When host indexing is enabled, Plixer Scrutinizer will store records for all hosts that pass traffic on the network. Records for host pairs can also be stored (and searched through) by enabling host to host indexing as indicated below.

To enable host indexing:

1. Navigate to [Admin > Alarm Monitor > Flow Analytics Configuration](#).
2. Open the configuration tray for the *Host Indexing* algorithm.
3. Add sources/inclusions for the algorithm either individually or using [Security Groups](#).

Hint: Recommended inclusions for host indexing are internal/core routers, edge routers, and public IP addresses that have been assigned to [IP Groups](#).

4. If there are sources (IP addresses/ranges, domains (by reverse DNS), IP Groups, etc.) that should not be indexed, add them as exclusions.
5. Expand the *Settings* secondary tray to configure the following:
 - Days of Host Index Data Retention
 - Host Index Database
 - Host Indexing Domain Socket
 - Host Index Max Disk Space
 - Host Index Sync Interval Minutes
 - Host to Host Database
 - Window Limit

Important: Entering a database path in the *Host to Host Database* field enables host pair indexing. To disable the feature, leave it blank.

6. Use the toggle to enable the algorithm and close the tray.

Once the algorithm has been configured and enabled, users can use the **Explore > Search** view to search the host or host pair (if enabled) index.

Hint: If the *Use Host Index* option ([Admin > Settings > Reporting](#)) is enabled, only Exporters that a host has been seen on will be searched when data is aggregated for a report. This can significantly reduce the time it takes to run reports.

Resource requirements

When enabling host and/or host to host indexing, it may become necessary to allot *additional CPU cores and RAM* to the Plixer Scrutinizer deployment.

Furthermore, the host to host index will include records for two types of host pairs:

- Continuously active pairs, whose records will not expire
- Ephemeral unique pairs, whose records expire but are replaced at approximately the same rate

To ensure enough storage is available for the host pair database, it is highly recommended to expand the disk space allocated to Plixer Scrutinizer.

Disk space calculation

To approximate the amount of additional disk space needed, start by using a report to calculate the number of host pairs:

1. *Create/run a new a Host to Host pair report* and add all Exporters that were defined as inclusions for the *Host Indexing* FA algorithm.
2. Set the time window to cover a period of at least 24 hours.
3. When the *output of the report* is displayed, click the gear button to open the Options tray and select *Global*.
4. In the secondary tray, select the *5m* option from the **Data Source** dropdown and click *Apply* before returning to the main view.
5. Note the total result count, which will be roughly equivalent to the number of active pairs.
6. Return to the **Options > Global** tray and switch to the *1m* data source option.
7. Subtract the previous result count from the updated total result count to determine the number of ephemeral pairs.

After obtaining the active pair and ephemeral pair counts, the following formula can be used to calculate additional disk space requirements for host to host indexing:

$$(\text{Active pair count} + \text{Ephemeral pair count}) * \text{Exporter count} * 200 \text{ B}$$

where **Exporter count** corresponds to the total number of Exporters/inclusions defined for the *Host Indexing* algorithm.

Utilization alerts

If the combined disk space used by the host and host pair databases reaches 100% of the *Host Index Max Disk Space* setting of the *Host Indexing* algorithm, host and host to host indexing will be suspended until storage becomes available again.

The following Alarm Policies are used to alert users to high disk utilization by host indexing:

| | |
|---|--|
| Host Index Disk Space Warning | Triggered when the disk space used by host indexing functions reaches/exceeds 75% of the specified <i>Host Index Max Disk Space</i> |
| Host Index Disk Space Error | Triggered when host indexing functions are suspended because the <i>Host Index Max Disk Space</i> has been reached |
| Host Index Disk Availability Error | Triggered when host indexing functions are suspended because disk utilization for the volume the host and host pair databases are stored on has reached/exceeded 90% |

Host indexing functions will automatically restart once sufficient storage is available, either due to record expiry or because disk space has been added.

5.1.4 Investigate

The **Investigate** section/tab of the web interface is split into the **Collections** and **Forecasts** pages, which contain specialized packages of information designed for further review and analysis.

This section introduces the two information package types and contains instructions and additional details related to their respective pages/views.

Collections

Collections are bundles of one or more Events and/or Reports that have been compiled and assigned to a specific user for further review and analysis. They can also be added to, annotated, and reassigned allowing multiple users to share workloads and collaborate in investigations.

Important: Collections require a Plixer One Network or Plixer One Security license. Contact [Plixer Technical Support](#) to learn more about licensing options.

In the web interface, Collection-related functions are accessed via the following elements:

Collections page

The **Collections** page of the **Investigate** section lists all existing Collections and is split into two tabs: **Assigned to Me** (current user) and **Other Collections**.

Hint: Click the filter button to view available filtering options for the list.

Along with each Collection's name, the table also shows the following details:

- User who created the Collection
- Date and time the Collection was added
- Date and time the Collection was assigned
- User to whom the Collection is currently assigned
- Number of Alarms, Events and/or Reports that have been added to the Collection

From the main **Collections** page, the following actions are available:

- **Viewing Collections** - Click on a Collection's name to open its [summary page](#).
- **Deleting Collections** - Select one or more Collections to delete by ticking and click the *Delete* button.
- **Reassigning Collections** - Click the username under a Collection's **Assigned User** column to assign it to a different user
- **Setting the active Collection** - Use the radio buttons to set/change the active Collection. For additional information, see the subsection on [managing Collections](#).

Inspecting Collections

A Collection's summary page lists all Alarms, Events, and Reports added to the Collection as links that allow the user to drill down into each item. Annotation can be added to the summary page in threaded view using the **Notes** card.

In addition, the table also lists the following details for each item:

- Type of item
- Additional details, such as the number of individual Events, hosts involved, or Report Type (click + to expand)
- Date item was added to the Collection
- User who added the item
- Any notes related to the Alarm, Event, or Report added by users

Hint: When adding notes to a Report item in a Collection, the text field will be pre-populated with basic information about the Report.

To remove items from the Collection, select one or more items using their checkboxes and click the *Delete* button.

Collection management

The **Manage Collections** submenu can be accessed from the different **Alarm Monitor** views or after running a Report by clicking on the star button.

Creating a new Collection

To create a new Collection, click the *Add New Collection (+)* button in the submenu. Enter a unique name for the Collection and select a user to assign the Collection to.

Note: The name and user fields must both be filled to create a new Collection.

Afterwards, click the + button to save the Collection. Once created, the Collection will be added to the list in the primary menu.

Setting the active Collection

To set/change the current active Collection, open the Collections menu and select it in the list. Only one Collection can be set as active at a time.

The active Collection can also be set from the [main Collections page](#).

Adding Alarms, Events, or Reports to a Collection

To add the current Alarm, Event, or Report to the active Collection, click on the star button to open the Collections menu and then click the button a second time (after it turns into a + button).

Hint: To remove an item from a Collection, click the star button once and click it a second time, after it turns into a - button.

Forecasts

Plixer Scrutinizer Forecasts are AI-/machine-learning-generated projections that provide insight into future network activity and utilization. Each forecast is based on a report, whose configuration determines the data elements to be extrapolated.

Important: Forecasts require an active Plixer One Network license. Contact [Plixer Technical Support](#) to learn more about licensing options.

This section covers the **Investigate > Forecasts** tab/section of the web interface and includes further details on generating, viewing/interpreting, and managing forecasts.

Generating Forecasts

After running a Report from the [Run Report page](#), the user will have the option to generate a Forecast based on the Report's data.

Hint: Forecasts are available for all report types.

The following details from the Report will be used to generate the Forecast:

- Hosts
- Data points
- Time period covered
- Filters applied

At the Report output page, verify that the settings and the data covered by the Report match the Forecast requirements, before clicking the *Forecast* button and entering a name for the new Forecast.

Note: For more detailed information on Reports, see the [Reports](#) section of this documentation.

After the Forecast has been successfully created, the main **Forecasts** page will automatically be opened.

Viewing Forecasts

All available Forecasts are accessible from the **Investigate > Forecasts** page. Forecasts that are marked *Complete* under the **Status** column of the page are ready to view.

The summary page for a Forecast is divided into two sections:

Forecast timeline

The timeline is a visualization of both the data pulled from the base Report (solid lines) and the projected values (broken lines) up to the horizon of the Forecast. The graph will also show the potential variance as a region that can be highlighted by mousing over the corresponding plotted line.

Hint: When mousing over any point of a plotted line will also open a tooltip with additional details about the projections for that particular entity, pair, or group.

The timeline can be viewed as either a line or step graph.

Detail summary

Along with the timeline, the page also includes an accompanying table that summarizes the most notable details of the Forecast:

- Color-based legend for information elements of interest covered by the data
- Links to summary pages for additional objects (sources, destinations, applications, etc.) related to the primary entity/pair/group of interest
- Maximum expected/forecasted value for the measured statistic
- Date and time when the measured statistic is expected to reach the maximum forecasted value
- Upper bound in the variance of the maximum forecasted value

To re-run the Report the Forecast is based on, click the *View Report* button.

Forecast management

The **Forecasts** page shows the following details for each Forecast listed:

- ID number
- Name of the base Report
- Name of user who created the Report
- Current status of the forecast (*Initializing -> Starting -> Data Retrieval -> Processing -> Strategy Selection -> Learning -> Prediction -> Complete*)

Note: In some cases, it may take up to several minutes for the Forecasting task to progress from *Initializing* to *Complete*.

- Date the Forecast became ready to view

The following actions can be also be performed from the **Forecasts** page:

- **Updating Forecasts** - Click the refresh button to update a Forecast to use the latest data from reports with dynamic settings (under the **Custom** dropdown when configuring Report settings) for the time period covered.

Hint: Forecasts based on reports with a specific time window can also be refreshed but will return the same projects as before. To generate a new Forecast with adjusted date and time settings, run a new report of the same type and create a Forecast with it.

- **Viewing/re-running Reports** - Click a Report name to view the Report whose data was used to generate the Forecast*
- **Deleting Forecasts** - Use the checkboxes to select one or more Forecasts and click the *Delete* button to remove them from the system permanently.

5.1.5 Reports

The **Reports** section/tab is used to create, run/view, and manage reports. Auxiliary functions, such as exporting reports, creating forecasts, and adding reports to Collections, can also be accessed from this section of the web interface.

Note: Forecasting is enabled by the Plexier ML Engine under the Plexier One Network solution. Contact [Plexier Technical Support](#) to learn more.

This section contains further information on reports, as well as guides/instructions for further leveraging report-related functions and features.

How reports work

Plexier Scrutinizer Reports automatically aggregate network metadata from any number of observation points based on a specified report type/template.

Reports in Plexier Scrutinizer are network metadata aggregations based on user-configured parameters and options. This allows reports to be tailored to support any network or security *use case* and *refined* to meet more precise visibility requirements.

Report settings

The output of a report in Plexier Scrutinizer is controlled by the following settings:

Network devices

When a report is run, Plexier Scrutinizer aggregates data collected from one or more user-specified network devices or interfaces. These function as the user's "observation points" and determine the scope of the data to be included in the report.

Report types

The base type of a report determines how network metadata from the selected observation points is aggregated (i.e., by **X**).

When *creating a new report configuration* or *refining report results*, report types can be displayed by category for ease of navigation. These categories include:

| | |
|-------------------------------------|---|
| Core | Client Server Counts Destination FQDN Firewall Events Source Top Vitals Volume Summary |
| Integration-/vendor-specific | AWS Azure FlowPro Defender FlowPro APM Palo Alto Networks |

To further simplify report type selection, types are also classified under the supplementary categories *Recommended*, *Favorites*, and *Designed Reports*.

Hint: Available report types vary based on the devices included in the Plixer Scrutinizer environment and those that have been selected as observation points for a report. Additional report types may also become available when certain integrations are configured/enabled for the Plixer Scrutinizer environment.

Time range/window

By default, reports are configured to aggregate data from the past 24 hours. However, this can be changed to a different *last X* window (e.g., last 5 minutes, last week, etc.) or a custom date and time range.

Hint: When a *Last X* time window is selected, clicking the up or down arrow will automatically shift the date/time period covered backward or forward.

Additional filters

The scope of the flow data aggregated by a report can be further limited or expanded through the use of [additional filters](#), which can be defined both during the configuration of a new report and after any report has been run.

Graph type

The output of a report includes a graph or chart plotting the top ten aggregations. The user is able to switch between the available graphs or diagrams (based on the report type) to display a visualization that best suits their current objective.

Plixer Scrutinizer Reports can be set to use any of the following graph/chart types for visualization:

- Line
- Stacked line
- Stacked bar
- Step
- Stacked step
- Pie
- Matrix
- Connection
- Sankey
- Donut

Note: Graph options vary based on the report type.

Custom Reports

To learn more about creating custom Reports, see [this subsection on the Report Designer](#).

Report results/output

The results/output view of a report is divided into two main sections: the *graph* and a paginated table where the complete data set can be reviewed.

Hint: The graph can be hidden by selecting *Hidden* when *creating a new report configuration* or using the *Hide* setting under Graphs in the Options tray.

After a report is run, the results can be continuously refined by *modifying its settings* from the output view. *Report management tools* and other auxiliary report functions can also be accessed from this page.

Additional options

The *Options* tray (gear button) can be used to access the following submenus related to the report:

| | |
|------------------|--|
| <i>Global</i> | General display settings for the current report; can also be used to manually select the summary table/bucket to pull data from for the report |
| <i>Graph</i> | Show or hide the graph in the main output view (<i>Hidden</i> can also be selected from the main view dropdown to hide the graph) |
| <i>Table</i> | Show or hide the peak and/or 95th percentile columns in the results table |
| <i>Threshold</i> | Configure an alarm-generating <i>threshold</i> based on the column the report is currently sorted by |
| <i>Details</i> | View the report's JSON output or additional details about the Exporters or Collectors used |

Hint: Use the *Copy to clipboard* button to copy a report's JSON output for reporting APIs.

Note: When the global *Data Source* setting is set to auto, Plexier Scrutinizer will automatically use the most suitable *summary table* to pull data from based on the time window of the current report.

Creating/running reports

The **Reports > Run Report** page is the starting point for creating/running new report configurations.

Hint: Access to specific reports and/or report folders, can be defined as part of usergroup permissions from the [Admin > Users & Groups > Usergroups](#) page.

New reports

To create/run a new report from the **Reports > Run Report** page:

1. Select between the two starting points to create a report:

| | |
|---------------------------|--|
| Select Devices | Select one or more devices before choosing a supported report type |
| Select Report Type | Select a report type before choosing one or more eligible devices. |

Note: The wizard automatically displays only supported report types and eligible devices.

2. In the next step, select the type or devices for the report:
 - Report type: Navigate to a category using the dropdown and select from the listed report types.
 - Devices: Check the devices under *Available Devices* and use the arrow buttons to add them to the *Selected Devices* list.
3. Configure the following settings on the following page:
 - [Time Window](#)
 - [Display Type](#)
 - [Additional Filters](#) (optional)
4. Click *Run Report*.

A progress bar is shown as the report is being run. Afterwards, the [report results/output view](#) will be displayed.

Note: Some Reports may take longer than others to run.

Saved reports

By default, the **Reports > Saved Reports** page lists all saved reports and report folders (via the dropdown) available to the current user. These reports can be re-run at any time and used as-is or [reconfigured](#) for other purposes.

To learn more about saved reports, see the section on [report management](#).

Report filters

Plixer Scrutinizer Reports grant full environment observability by aggregating network metadata with any number of user-defined filters applied. This allows reports to be used for both monitoring and investigation.

Basic filters

As part of *creating a new report*, the user is required to configure three *report settings* that function as the main filters:

- Report type
- Devices
- Time window

These settings define how the report should aggregate data (type), which observation points or sources it should use (devices), and the period of time it should cover (window).

Additional filters

Before running a new report and after any report is run, additional filters can be added to tailor the output to the scenario the report is being used in.

The following table lists the additional filters that can be applied to reports:

| Type | Description | Parameter(s) | Option(s) |
|--|---|--|------------------------|
| <i>Applications</i> | Filter results for a selected NBAR application | NBAR application | Restriction |
| <i>Applications defined</i> | Filters results for a selected defined application (based on definitions under Admin > Definitions > Applications) | Defined application | Restriction |
| <i>Autonomous system by tag</i> | Filters results for the selected autonomous system (AS) tags | Autonomous system (by AS number) | Direction, restriction |
| <i>Business hours</i> | Filters results for activity during specified business hours | Start hour, end hour, time zone, days | N/A |
| <i>Calculated column filter</i> | Filters results based on values in one of the report's calculated columns | Filter column, comparison operator and value | N/A |
| <i>Country</i> | Filters results for the selected country | Country | Direction, restriction |
| <i>Device/interface</i> | Filters results for activity associated with the specified devices, interfaces, or mapping groups | Device Interface (if a device is selected) Mapping group (if <i>Group</i> is selected) | N/A |
| <i>Domain</i> | Filters results for the specified domain | Domain | Direction, restriction |
| <i>Flow template</i> | Filters results for the selected template | Flow template | Restriction |
| <i>Host list</i> | Filters results for the specified hosts | Host IP address(es) | Direction, restriction |
| <i>Host to host</i> | Filters results for activity between the specified host pair | Host pair IP addresses | Restriction |
| <i>IP Groups</i> | | IP Group name | Direction, restriction |
| 5.1. Plixer Scrutinizer web interface | Filter interface for the selected IP Group (defined under Admin > Definitions > IP | | 157 |

Direction options: Source, destination, or both *Restriction* options: Include or exclude

Important: The additional filters that can be added to a report vary based on the selected devices/interfaces and report type. More filters may also become available when Plixer Scrutinizer has access to devices from certain vendors or is configured with additional integrations.

Refining a report

After any report is run, one or more of its settings can be modified to further inspect into any included or adjacent data element. This allows the user to create varying levels of visibility from multiple angles and extract deeper insight on the hosts and/or behaviors being investigated.

Editing basic report settings

In the report output view, the report's *basic settings* can be modified as follows:

| Setting | Edit from | Effect |
|--------------------|-----------------------------|--|
| Report Type | Main view (dropdown) | Changes the base Report Type but retains all other applicable settings |
| Time range | Main view (calendar button) | Defines a new period of time to be covered by the Report |
| Graph | Main view (dropdown) | Changes or hides the graph/chart used to model the Report data |

Hint: The **Options** (gear button) tray contains additional settings/options to customize how the report is displayed. *Report thresholds* can also be added from this tray.

Adding/removing filters

As part of refining a report, filters can be added, edited, or removed. By layering the appropriate filters, the user can limit the scope of the report to only hosts and/or traffic relevant to their current investigation.

Note: Devices and interfaces, including those that were initially selected when creating the report, are considered filters and can be edited from the *Filters* tray.

To add a new filter:

1. Click the *Filters* button to open the tray.
2. In the tray, click the + button.
3. Select a *filter type* for the new filter.
4. Configure the additional settings for the filter (varies by filter type)
5. Verify that the settings are correct and click the *Add* button.
6. In the primary tray, click the *Apply* button to re-run the Report with the new filter(s) applied.

Current filters can be modified by clicking the edit (pencil) button and making the necessary changes before clicking the *Save* button. To remove a filter, click the delete (bin) button next to it in the list.

Hint: To avoid having to re-run the Report more than once, add all necessary filters before clicking the *Apply* button.

Inclusion/exclusion dropzones

Data elements in the results table can be dragged into inclusion/exclusion dropzones to the left of the page to automatically add them as filters.

This can be repeated as many times as necessary to set up the appropriate filters before re-running the report.

Note: Dropping an element into the inclusion or exclusion dropzone automatically opens the *Filters* tray, if it was not already open.

Pivoting to different report types

Clickable data elements in the results table open a tray listing all report types (sorted into their respective categories) that can be run using that element. This allows the user to inspect all associated activity using the context best suited for the investigation.

Managing reports

After a report configuration is created and run, it can be saved and later re-run from the **Reports > Saved Reports** page.

Note: A report must be saved before it can be used in certain advanced functions.

Reports can be saved from the results/output view by clicking the save (disk) button.

Report folders

Report folders, which can be used to organize saved reports, can be created and populated from the **Admin > Reports > Report Folders** page.

The default view of the **Reports > Saved Reports** page displays all reports that have not been added to any folders (*Unfolded*). To view the contents of a different folder, select that folder from the dropdown.

Hint: A report can be a member of multiple folders.

User access to specific reports and folders can be controlled from the [Admin > Users & Groups > User-groups](#) page.

Exporting reports

After a report is run, the output can be exported in PDF or CSV format from the *Export* (share button) tray.

Hint: PDF and or CSV copies of a report can also be attached to [email reports](#).

Deleting Reports

A saved report can be deleted from the **Reports > Saved Reports** page by clicking the corresponding delete (bin) button.

Email reports

Once an *email server has been configured*, email reports can be used to provide access to network information via any recipient's inbox.

Hint: Email reports also include a link to view the report in the Plexier Scrutinizer web interface. PDF and/or CSV copies of the report may also be attached.

On-demand email reports

After any report is run, an email report can immediately be sent by selecting *Email Report* in the export tray and entering the following details:

- Sender email address
- Recipient email address (comma-delimited if multiple)
- Subject (optional)
- Message (optional)

After clicking *Send*, a message confirming that the email report has been sent will be displayed.

Scheduled email reports

Scheduled email reports are re-run (using the exact same configuration) and sent to specified recipients at regular intervals. When a *Last X* time window is defined, scheduled reports can be used to continuously monitor any type of network metadata from any email inbox.

Important: A report must be saved before it can be configured as a scheduled report.

Note: Scheduled email reports covering a custom time window (i.e., specific dates and times) will send either the same or no output when they are re-run.

Scheduled reports can be set up by selecting *Schedule Report* in the export tray and entering the following details:

- A name for the scheduled report (used as the email subject line and for scheduled report management)
- Recipient email address (comma-delimited if multiple)
- Frequency and exact minute on the hour that the report should be re-run and emailed

Additional reports may also be added to the scheduled report configuration by selecting them from the list. These reports will also be re-run and displayed in the same email.

Afterwards, click *Save* to save the scheduled report configuration. Scheduled reports will continue to be re-run and sent out as configured until the scheduled report configuration is disabled or deleted.

Hint: Scheduled report configurations can be edited or disabled from the [Admin > Reports > Scheduled Email Reports](#) page.

Report thresholds

Once a report has been saved, it can be used to configure an alarm-generating threshold for the calculated column the report results are currently sorted by. Report thresholds can be used to alert users to behavior or activity that exceeds/falls below specified values.

Hint: In the results table, calculated columns are indicated by the *up* and *down* sorting arrows next to the column header.

Thresholds are configured from the options tray of the report output view as follows:

1. Click *Threshold* to open the threshold settings secondary tray.
2. Select whether the threshold should apply to individual row values or the total value of the current sort column.
3. Select the comparison operator (\geq or \leq) that will trigger the threshold alarm.
4. Enter the desired threshold value and select the measurement prefix to use (kilo-, mega-, or giga-).

Important: The threshold value automatically follows the unit of measurement of the current sort column, e.g., if the report is sorted by packets per second, entering 100 and selecting K from the dropdown will result in a threshold of 100 kilopackets per second.

After a threshold has been configured, violations will be reported via the Alarm Monitor views under the *Report Threshold Violation* policy.

To disable a report threshold, re-run the report and click the delete (X) button in the *Filters* tray of the output view.

Reports gadgets

For more convenient monitoring, reports can also be added to dashboards as gadgets.

Note: For a user to be able to view a report this way, their usergroup must have access to both the report and the dashboard(s) it is added to.

After a report is run, it can be exported as a gadget as follows:

1. Open the export tray and select *Add to Dashboard*.
2. In the *Report Name* field, enter a name for the report gadget.

Note: When exporting a saved report, replacing its name (default) will result in a new report being saved under that name. Unsaved reports will automatically be saved using the name entered.

3. Use the *Dashboard Tab* dropdown to select the dashboard to add the gadget to. If *Don't send to dashboard* is selected, the report gadget will need to be manually added to dashboards at a later time.

Hint: Reports that have been exported directly to a dashboard can still be added to other dashboards.

4. Use the *Display* dropdown to select whether to the gadget should show the report graph only, table only, or both.
5. Click *Save* to save the report gadget.

After a report has been exported as a gadget, it will be included in the list of available gadgets when [creating](#) or [editing](#) a dashboard.

To learn more about customizing dashboards, see the section on [dashboards](#).

Creating forecasts

Using the data aggregated by a report, Plixer Scrutinizer can leverage the capabilities of the Plixer ML Engine to forecast the future states of the report's calculated values.

To create a forecast from a report, click the *Save Forecast* button on the report output view. In the tray, enter a name for the forecast and click *Create*. The forecast can be viewed from the [Investigate > Forecasts](#) page.

Note: After a forecast is saved, the user will automatically be taken to the main **Forecasts** page. Forecasts using certain report types may take additional time before they are ready to be viewed.

See the section on [forecasts](#) to learn more about how forecasts work.

Adding reports to Collections

To add a report to the [active Collection](#), open the *Manage Collections* menu (star button) and then click the button a second time (after it turns into an add (+) button).

Once added to a Collection, a report can be re-run from the [summary page](#) for that Collection.

See the [Collections](#) section to learn more about creating, reviewing, and managing Collections.

5.1.6 Admin

The **Admin** views of the Plixer Scrutinizer web interface are used to access the system's administrative and configuration functions.

For ease of navigation, the different admin pages/views are organized into categories in the **Admin Menu** tray, which can be accessed from any Admin page/view via the three-dot button.

Admin Dashboard

The **Admin Dashboard** provides a visual overview of the functions and performance of the Plexier Scrutinizer environment. It is the default view opened when clicking on the **Admin** text in the web interface header.

The page comprises the following interactive Dashboard Gadgets:

| | |
|-------------------------------|---|
| System Performance | Displays system performance metric in timelines or charts Click on a metric to switch views. Click on the <i>Vitals</i> icon to view server health. |
| Free Disk by Collector | Displays available storage per Collector Click on a storage element to switch views. Click on the <i>Vitals</i> icon to view OS health. |
| Services by Collector | Displays the status of system services per Collector Hover over a chart element to view additional details. Click on the Vitals icon to view Exporter health |
| Activity by User | Shows activity for individual users in a timeline |

Hint: Click the **X** button to close any of the secondary tables.

Virtual LEDs

All Admin pages/views feature three virtual LEDs that can be used to monitor the general health of the Plexier Scrutinizer environment.

Three virtual LEDs are persistent across all Admin pages/views. These can be used to monitor the general health and performance of the Plexier Scrutinizer environment.

The LEDs correspond to the following system components, from left to right:

- Server health
- OS health
- Exporter health

Hint: The virtual LEDs link back to the **Admin Dashboard** with the corresponding vitals table for the component open.

Hover over a virtual LED to view additional details.

Admin Menu tray

The **Admin Menu** tray is main access point for the web interface's Admin views/pages. The tray can be pulled out from any of the various Admin pages/views by clicking on the three-dot button.

Each settings category can be expanded to display all Admin views under it.

Note: Admin views marked with a [-> are still only accessible via the Classic UI of the web interface.

Settings

The **Admin > Settings** category contains global settings to manage the Plexier Scrutinizer system's general behavior.

Hint: In the *Mapping Groups* and *Mapping Objects* management views, bulk actions become available after one or more items are selected.

These settings are further organized into the following subcategories:

| | |
|---------------------------------|--|
| <i>Alarm Notifications</i> | Global Alarm display settings and toggles for <i>Flow Inactivity</i> and <i>Interface Threshold Violation</i> Alarms |
| Data History | Retention settings for historical Alarm and flow data |
| Flow Analytics Settings | <i>Global settings</i> for <i>Flow Analytics</i> |
| Google Maps Proxy Server | Google proxy server settings |
| Login Banner | Custom text for web interface login banner |
| Mapping Groups | Configuration and management options for <i>Network Map</i> device groups |
| Mapping Objects | Configuration and management options for <i>Map Objects</i> |
| Reporting | Global settings for running Reports |
| System Preferences | General Plexer Scrutinizer environment settings |

Hint: *Notification Profiles* can also be assigned to Flow Inactivity and Interface Threshold Violation Alarms via their respective *Alarm Policies*. The interface utilization threshold for violation Alarms can also be adjusted via the **Threshold - Utilization** setting in the *System Preferences* tray.

Note: Assigning a value of 0 to any of the flow history settings under *Data History* does not disable retention of the corresponding data table.

Alarm Notifications

The **Admin > Settings > Alarm Notifications** tray contains the following settings:

| | |
|---------------------------------------|--|
| Alarm Many Crop | Sets the length to which device and host lists in Alarm messages will be truncated |
| Flow Inactivity | Enable/disable Alarms for devices from which flows have not been received for 30 minutes (Reported under the <i>Flow Inactivity</i> Alarm Policy) |
| Hostnames | Enable/disable displaying hostnames for devices and hosts (targets and violators) in Alarm messages |
| Interface Threshold Violations | Enable/disable Alarms when the total utilization (in or out) for any interface exceeds the percentage threshold specified in the Admin > Settings > System Preferences tray (Reported under the <i>Interface Threshold Violation</i> Alarm Policy) |

Important: If *Flow Inactivity* and *Interface Threshold Violations* alerts are disabled in this tray, violations will not be reported/saved, even if the corresponding Alarm Policies are set to [Active or Store](#).

System preferences

The following table lists all options/settings that can be modified via the **System Preferences** tray:

| | |
|-------------------------------------|---|
| Allowed Flow Rate Multiplier | Multiplier applied to the system's maximum supported flow rate to accommodate brief, recoverable traffic spikes. <i>Note: Sustained flow rates exceeding 100% of the rated limit may result in stability issues.</i> |
| Always Display Totals | Toggle on to force tables in Status Reports to display totals, even when the graph shows rate. |
| Auto SNMP Update | Toggle on to enable re-discovery of SNMP devices at 1:00 am every day. |
| CSV Repository | Sets the directory that scheduled CSV files are saved to. |
| Disable File Upload | Toggle on to disable uploading files to the server. |
| Disable Welcome Modal | Toggle on to disable the <i>Welcome to Plixer Scrutinizer</i> modal for new users. |
| DNS Cache Retention | Sets the number of days (0 to 365) to retain DNS names. <i>Note: When the value is set to 0, DNS names are never retained by the system.</i> |
| DNS Timeout | Sets the maximum time (in seconds) allowed for DNS name resolution. |
| Enforce Password Complexity | Toggle on to require at least 8 characters, 1 capital letter, 1 number, and 1 special character for new user passwords. |
| Failed Login Max | Number of failed login attempts allowed before an account is locked (0 = disabled). |
| Failed Login Window | Sets the window of time (in minutes) for failed login attempts; any failed logins outside the window will not count towards the <i>Failed Login Max</i> setting. |
| Inactive Expiration | Sets the number of hours (1 to 168) before an inactive interface is removed from the <i>Top Interfaces</i> view. |
| Inactivity Threshold | Sets the number of hours the last inactive interface values are displayed in the <i>Top Interfaces</i> view. |

continues on next page

Table 1 – continued from previous page

| | |
|--|---|
| Language | Default system language (can also be set per user via user account settings). <i>Note: Documentation and technical support are only available in English.</i> |
| Listener Port | Ports to use to listen for NetFlow or sFlow traffic (separate by comma) |
| Flow Resources Fallback Cooldown Period | Amount of time (in seconds) to wait after <i>Low Resource Fallback</i> settings have been applied before further actions are taken |
| Low Resource Fallback Exporter Chunk Size | Number of Exporters to pause or resume at a time when <i>Low Resource Fallback</i> determines Exporters should be paused or resumed |
| Low Resource Fallback Mode | Actions to be taken by the system to avoid catastrophic failure when there is insufficient CPU and/or RAM |
| Maximum Raw Flow Exporters | Maximum number of Exporters allowed in a filter before the <i>Raw Flows</i> option becomes unavailable |
| Maximum Uploaded File Size in Bytes | Maximum size allowed for uploaded files |
| Minimum Unique Passwords | Number of recent passwords that cannot be reused by users when changing passwords |
| Report Caching Timeout | Amount of time (in minutes) to use when caching a list of available Reports |
| Resolve Hosts at Collection Time | Toggle on to force DNS name resolution for every host seen when flows are collected (only necessary for Flow Analytics domain exclusions and Rev 2nd level domain Reports) <i>*Note: Enabling this feature may result in significant latency at high flow volumes. For assistance, contact Plixer Technical Support.</i> |
| Session Timeout | Amount of time (in minutes) web sessions are allowed to be idle before the user is forcibly logged out (0 = disabled). |
| Theme | Default system theme (can also be set per user via user account settings). |
| Threshold - Yellow | <i>Yellow</i> interface utilization threshold. |
| Threshold - Orange | <i>Orange</i> interface utilization threshold. |
| Threshold - Red | <i>Red</i> interface utilization threshold. |

continues on next page

Table 1 – continued from previous page

| | |
|--------------------------------|---|
| Threshold - Utilization | Total interface utilization (in or out) threshold to trigger an <i>Interface Threshold Violation</i> Alarm. |
| TOS Family | Quality of Service or Type of Service configuration/family used by the organization. |
| Version Checking | Toggle on to allow Plixer Scrutinizer to automatically connect to the Internet and check for updates. |

Definitions

The **Admin > Definitions** category contains management views for the various user-defined elements and groupings used by the Plixer Scrutinizer system.

Hint: In views that include selection checkboxes, bulk actions become available after one or more items are selected.

From the **Definitions** page, the following system elements can be defined or configured:

| | |
|--------------------------------|--|
| Applications | Custom application definitions based on one or more IP addresses and ports |
| Autonomous Systems (AS) | <p>Search for and/or view factory-defined Autonomous Systems (AS)</p> <p>Custom Autonomous System Number definitions can be imported with the <i>import</i> <code>scrut_util</code> command.</p> |
| Host Names | Static hostname-to-IP assignments and subnet labels for use in Reports and other Plixer Scrutinizer functions |
| IP Groups | User-defined subnets and IP address ranges for use in Reports and other Plixer Scrutinizer functions |
| MAC Addresses | Custom labels for collected MAC addresses |
| Protocol Exclusions | Custom rules to exclude certain protocols from collection functions |
| Type of Service | <p>Custom labels for Type of Service (ToS) and Differentiated Services Code Point (DSCP) values in Reports</p> <p>(<i>ToS Family</i> must be defined under <i>System Preferences</i> first)</p> |

Note: Pages under the **Admin > Definitions** tab of the Classic UI are included in this category.

IP Groups

IP Groups are supplementary groupings of network devices or endpoints (e.g., by department, by geographic region, by device type, etc.) that are used in Reports and other Plixer Scrutinizer functions. Each IP Group can comprise IP addresses, IP addresses ranges, and/or subnets as defined by the user.

Hint: Plixer Scrutinizer includes a number of default IP Groups that are used by various core functions (e.g., *Security Groups*). These IP Groups should be populated as part of setting up the system for the first time.

Adding a new IP Group

To add a new IP Group, follow these steps:

1. On the **Admin > Definitions > IP Groups** page, click the (+) button to open the *Add IP Group* tray/form.
2. Enter a name for the group.
3. Select whether the group is internal or external from the dropdown.
4. Click *Save*.
5. In the main view, click the newly created IP Group to open its configuration tray.
6. Expand the *Rules* section of the tray and click the (+) button to add a new rule.
7. In the secondary tray, select the appropriate rule (IP address, subnet, etc.) type from the dropdown.
8. Enter the details required for the rule in the additional fields.
9. Click *Add* to save the rule.

Hint: An IP Group can have multiple rules defining membership.

The settings for an existing IP Group can be edited at any time by clicking on it in the main table view.

Bulk actions

When one or more IP Groups are selected using their checkboxes, the following bulk actions become available:

- Adding new rules to all selected IP Groups
- Deleting all selected IP Groups

Users & Groups

The **Admin > Users & Groups** category contains functions related to user management and access control.

Hint: In views that include selection checkboxes, bulk actions become available after one or more items are selected.

These functions are split across the following views:

| | |
|------------------------------|--|
| <i>Authentication</i> | General authentication options as well as per-user settings |
| <i>Authentication Tokens</i> | Authentication token generation and configuration for individual users |
| <i>LDAP Servers</i> | Server and connection settings for LDAP integration |
| <i>Single Sign-On</i> | SSO provider configuration and management |
| <i>RADIUS Configuration</i> | RADIUS integration settings |
| <i>TACACS+ Configuration</i> | TACACS+ integration settings |
| <i>User Accounts</i> | Web interface user account management |
| <i>Usergroups</i> | Web interface usergroup management |

Note: Pages under the **Admin > Security** tab of the Classic UI are included in this category.

Integrations

The **Admin > Integrations** category consists of the configuration pages for built-in and custom integrations.

Hint: In views that include selection checkboxes, bulk actions become available after one or more items are selected.

The following configuration views can be accessed from this category:

| | |
|------------------------------|---|
| 3rd-Party Integration | Custom settings to pass variables in URLs to 3rd-party applications; Once configured, links to applications will be displayed in device map/trees |
| ASA ACL Descriptions | Credentials and settings for SSH connections to ASA firewalls to retrieve ACL descriptions (appliance only) |
| Email Server | Email server settings (used in Alarm notifications, Report forwarding, and Scheduled Reports) |
| Flow Log Ingestion | Management and configuration views for flow log ingestion from cloud sources - Amazon VPC flow logs - Azure NSG flow logs |
| STIX-TAXII | Settings for STIX-TAXII integration |
| ServiceNow | Settings for bi-directional ServiceNow integration |
| Viptela Settings | Settings for Viptela SD-WAN integration |

Alarm Monitor

The **Admin > Alarm Monitor** settings category covers the following configuration views for Alarm-Monitor-related functions:

Alarm Policies

The main **Alarm Policies** configuration page displays an overview of all current Policy settings in table format. It can be used to manage settings for individual Alarm Policies.

Hint: To apply filters to or export the information in the table, click the corresponding button to see available options.

The table lists the following information for each Alarm Policy:

- State (green: **Active**, blue: **Store**, grey: **Inactive**)
- Source Flow Analytics algorithm
- Category
- Total number of violations
- Number of Exporters being monitored for violations
- Event aggregation timeout
- Weight

Clicking a name opens the configuration tray for that Alarm Policy, where the *state (Active, Store, or Inactive) and other settings* can be configured. *Notification Profiles* can also be assigned to the Policy from this tray.

Notification Profile settings

When Notification Profiles are assigned to an Alarm Policy, their behavior can be further customized using the following options:

| | |
|----------------------------|---|
| Frequency | <p>Specifies how often the actions specified in the Notification Profile are triggered</p> <p><i>Each Observation</i> - Actions are triggered every time observed traffic meets the conditions of the Alarm Policy, regardless of duration.</p> <p><i>Rate</i> - Actions are triggered every Nth Event with the exact same criteria.</p> <p><i>Each Event</i> - Actions are triggered for every Event (aggregated observations based on the Policy's <i>Timeout</i> setting) reported under the Alarm Policy.</p> |
| Notification Filter | <p>Allows Event details (e.g., violators, devices, message contents) to be used as conditions to trigger or bypass notification actions.</p> <p>If no filters are specified, notification actions will be triggered for all observations and/or Events under the Alarm Policy.</p> |

Hint: When setting up notification inclusions or exclusions for observations/Events matching certain criteria, use the Alarm Monitor page to drill down into the Policy -> Event -> Observations views to see which details should be applied as filters.

To add separate notification configurations for different observation/Event criteria, assign multiple Notification Profiles to the Alarm Policy. The same Notification Profile can also be added multiple times with different frequency settings and/or filters.

Important: The frequency setting for each Notification Profile assignment is applied to all notification actions enabled by the configured filters.

For further details on how Alarms work and configuration recommendations, see the [Alarms and Events configuration guide](#).

Flow Analytics Configuration

The **Admin > Flow Analytics** page is used to manage settings for individual FA algorithms. Its main view consists of an overview of all current FA algorithm settings and includes a graph showing frequency by algorithm.

The table lists the following information for each FA algorithm:

- State (green: **Enabled**, grey: **Disabled**)
- Number of Exporters
- Number of defined exclusions
- Number of associated Alarm Policies

For additional information on FA algorithms and configuration recommendations, see the [Flow Analytics configuration guide](#).

Algorithm settings

From the main view, click on an algorithm to open its configuration tray.

The tray is divided into the following sections:

| | |
|-------------------|---|
| Sources | Exporters and Security Groups for which the algorithm has been enabled |
| Exclusions | IP addresses, IP ranges, subnets, domains (by reverse DNS), and IP Groups whose traffic will not be monitored using the algorithm |
| Settings | Additional settings that are exclusive to the current algorithm |

Note: Settings that do not apply to the current algorithm will be excluded from the tray.

Important: Certain features, such as host indexing, top x monitoring, and *Report Threshold* Alarms require the corresponding FA algorithm to be enabled.

An algorithm can also be disabled or re-enabled from its configuration tray.

Bulk actions

When one or more algorithms are selected, the following bulk actions can be accessed via the *Bulk Actions* button:

- Adding sources (Exporters and/or Security Groups)
- Disabling and enabling

Additional page options

- Filtering options can be accessed by clicking the *Filters* button.
- General page options, such as the number of entries shown and export actions, can be accessed by clicking the *Options* (gear) button.

Manage ML Dimensions

The **Admin > Alarm Monitor > Manage ML Dimensions** page is used to manage the communication types used as dimensions/features by the Plexier ML Engine for network behavior modeling.

The main table view of the page includes the following details for each dimension:

| | |
|----------------------|--|
| Status | Current operational status of the dimension (green: <i>Enabled</i> , grey: <i>Disabled</i>) |
| Protocol | Communication protocol monitored |
| Port | Communication port monitored |
| Internal Only | Option to interrogate only internal communications |
| Used For | Type of inclusion/source the dimension is applied to |
| Aggregation | Field used for data aggregation |
| Grouped By | Field used to group observed flow data |
| Created By | ID of dimension creator |
| Last Modified | Date and time the dimension was last modified |

Adding a new ML dimension

To add a new dimension, follow these steps:

1. In the main view, click the (+) button to open the *Add Dimension* tray.
2. Select whether the dimension to be added is for hosts or for Exporters.
3. In the secondary tray, click the + button and fill in the form with the following information:
 - A name for the dimension

Note: Host dimensions are prefixed with CLIENT- while Exporter dimensions are prefixed with NET-.

- Field to use for grouping (can only be changed for host dimensions)
- Aggregation method/field
- Protocol to monitor
- Port to monitor

4. Set the toggles to the desired configuration.
5. Verify that the details and settings entered are correct and then click the *Add* button.

The settings for existing dimensions can be edited at any time by clicking on them to open the configuration tray.

Bulk actions

When one or more dimensions are selected using their checkboxes, the following bulk actions become available:

- Enabling/disabling all selected dimensions
- Deleting all selected dimensions

Additional page options

- Filtering options (by name, protocol, port, aggregation, and grouping field) can be accessed by clicking the *Filters* button.
- General page options, such as the number of entries shown and export actions, can be accessed by clicking the *Options* (gear) button.

Manage ML Inclusions

The **Admin > Alarm Monitor > Manage ML Inclusions** page is used to manage hosts (and subnets) and/or Exporters for the Plixer ML Engine.

The page is divided into two subviews:

- The *By Host* view consists of a table that includes the following details for each host or subnet:

| | |
|----------------------|--|
| Status | Current operational status of the host or subnet as an ML inclusion/-source (green: <i>Enabled</i> , grey: <i>Disabled</i>) |
| CIDR | CIDR number |
| # HOST(s) | Number of hosts included in the subnet |
| Sensitivity | Threshold for classifying observed behavior as anomalous (lower -> less deviation required) |
| Detections | Option to use pre-trained ML algorithms for malware detection (green: <i>Enabled</i> , grey: <i>Disabled</i>) |
| Last Modified | Date and time the host or subnet was last modified |

- The *By Exporter* view table lists all Exporter inclusions, along with their configured sensitivity and *Last Modified* timestamp.

Adding hosts or subnets as inclusions

Hosts and subnets that are configured as ML inclusions/sources are monitored by the Plixer ML Engine (through Plixer Scrutinizer), whose network behavior models are based on their activity and the currently enabled *dimensions*.

To add a new host or subnet as inclusion for the Plixer ML Engine, follow these steps:

1. On the *By Host* view, click the + button to open the *Add ML Host* tray.
2. Enter the network address and select the appropriate mask for the host or subnet to be added.
3. Select between *High*, *Medium*, and *Low* sensitivity from the dropdown.

Note: When setting up inclusions for the Plixer ML Engine for the first time, it is recommended to leave all sensitivity settings at their default values and make adjustments after a period of observation. The *Reset* button can be used to revert the sensitivity setting to its default value.

4. If necessary, use the *Malware Detections* toggle to enable threat detection using pre-trained algorithms for the host or subnet.
5. To immediately enable the host or subnet as an ML inclusion, leave the *Enabled* toggle as is and click the *Apply* button.

The settings for an existing host or subnet inclusion can be edited at any time by clicking it in the main table view.

Adding Exporters as inclusions

To add Exporters instead of hosts or subnets as ML inclusions/sources, follow these steps:

1. Switch to the *By Exporter* view, and click the + button to open the *Add ML Exporter* tray.
2. Select the Exporter to add from the *Network* dropdown.
3. Select between *High*, *Medium*, and *Low* sensitivity from the dropdown.

4. If desired, use the *Malware Detections* toggle to enable threat detection using pre-trained algorithms for the Exporter.
5. To immediately enable the Exporter as an ML inclusion, leave the *Enabled* toggle as is and click the *Apply* button.

The settings for an existing Exporter inclusion can be edited at any time by clicking on it in the main table view.

Deleting inclusions

Inclusions can be completely removed from the system in either view using the *Delete* option in the bulk actions tray.

Note: The button to open the bulk actions tray only becomes available after one or more exclusions are selected using their checkboxes.

Additional page options

- Filtering options can be accessed by clicking the *Filters* button.
- General page options, such as the number of entries shown and export actions, can be accessed by clicking the *Options* (gear) button.

Notification Profiles

Notification profiles allow you to trigger actions when a specified alarm/event is generated. You can assign a notification profile to one or more alarm policies to automatically forward alarm and event data to external systems.

Notification profiles can be assigned to alarm policies from the [Admin > Alarm Monitor > Alarm Policies](#) page. Notification profile and action behavior can be further customized through the *Frequency* or *Notification Filters* settings when assigning them to an alarm policy.

Hint: Notification profiles can also be assigned to [Report Thresholds](#) using the *Report Threshold Violation Alarm Policy*.

This section covers the functions of the **Admin > Alarm Monitor > Notification Profiles** page and provides additional information related to configuration and use of notification profiles.

Creating a notification profile

Notification profiles can be created from the **Admin > Alarm Monitor > Notification Profiles** page as follows:

1. Click the **+** button and enter a name for it in the provided field.
2. Click **Save**.

Once saved, the notification profile will be added to the list and can be further configured.

Hint: The notification profile management page can also be accessed from the alarm policy configuration tray of the **Admin > Alarm Monitor > Alarm Policies** page.

Adding actions to a notification profile

To add actions to a notification profile, follow these steps:

1. Click the name of a notification profile to open the configuration tray.
2. Expand the *Actions* section of the tray and click the **+** button.
3. Use the dropdown to select the type of action to add.
4. Fill in the additional fields (based on the action type) with the required information.
5. Use the *Test* button to verify that the action functions as intended.
6. Click the *Add* button to save the action to the notification profile.
7. Click *Add Action* again to add another action or close the tray to return to the main view.

Hint: A single notification profile can be configured with as many actions as needed.

Batch actions

When selecting one or more notification profiles using their checkboxes, the following batch actions can be performed:

- Add an action to all selected notification profiles
- Delete all selected notification profiles

Note: To trigger notifications for a certain FA algorithm, you must enable the algorithm first, and then add the notification profile to the associated alarm policies.

Action types

Each notification profile can be configured with one or more actions.

When an event/alarm is generated, all actions under the notification profile assigned to it are triggered (unless exceptions/filters are defined).

A notification profile can be configured with one or more of the following action types, in any combination:

- Email event details
- Output event details as a logfile
- Forward event details to a specified host as a CEF notification, SNMP Trap, or syslog
- Run a custom script that uses event details
- Create a ServiceNow ticket
- Automatically acknowledge the event

Note: Notifications will only be processed when the alarm policy status is set to *Active* or *Store*. Notifications will not be processed for alarm policies with *Inactive* status.

Notification variables

When [adding an action to a notification profile](#), you have the option to add variables in the notification message. These variables allow you to customize the content of the notification message and are replaced with actual values when the notification is sent.

%m is the default variable when setting up an action to a notification profile. This means that the notification will show the message that appears in the Alarm Monitor page which includes information specific to the event.

However, you can add one or more of the following variables in the **Message** field, in any combination:

| | | |
|------------------|----------------------------|---|
| %m | Message | Specific message generated by the policy violated |
| %pol | Policy Violated | Indicates the name of the policy that generated the current notification |
| %v | Violator Addresses | IP addresses that violated the policy causing the notification |
| %url | Report Threshold Event URL | The URL of a report for a <i>Saved Report Threshold Violation</i> |
| %h | Host | The host sending the notification, in this case, Plexier Scrutinizer |
| %v_resolved | Resolved Violators | Indicates that the IP addresses that violated the policy are resolved |
| %id | Event ID | The identifier for the logged event that generated the current notification |
| %h_resolved | Host name | Indicates the host name where the resolved violation occurred |
| %viola-tor_users | Violator Usernames | Usernames associated with a violating host |
| %time | Alarm Time | The time that the policy that generated the current notification was violated |
| %p | Protocol | Specifies the name of the violating protocol, if applicable |
| %t | Target Addresses | IP addresses of the target machines |
| %tactic_id | Tactic ID | The ID of the malicious tactic, as identified by MITRE |
| %tactic_name | Tactic name | The name of the malicious tactic, as identified by MITRE |
| %target_users | Target Usernames | Usernames associated with a targeted host |
| %tech-nique_id | Technique ID | The ID of the malicious technique, as identified by MITRE |
| %tech-nique_name | Technique name | The name of the malicious technique, as identified by MITRE |
| %category | Category | The category of the policy violated |

Security Groups

The **Admin > Alarm Monitor > Security Groups** page allows users create and manage Security Groups for Flow Analytics algorithms.

A Security Group can be configured for one or more FA algorithms. When Exporters are assigned to it, those FA algorithms will automatically be enabled for them.

Plexier Scrutinizer ships with four predefined Security Groups that have been configured with the recommended FA algorithms for the corresponding device types.

Adding a new Security Group

To create a new Security Group, click the **+** button. In the tray, enter a unique (recommended) name for the Security Group, and then click the *Save* button.

Adding Exporters to a Security Group

To add one or more Exporters to an existing Security Group, follow these steps:

1. Click the Security Group name to open the configuration tray.
2. Expand the *Active Exporters* section and click the *Add* button.

Note: Exporters can also be removed from the group using the *Delete* icon.

3. In the secondary tray, use the checkboxes to select Exporters.
4. After selecting all Exporters to be added, click the *Add* button.

Once added, the selected Exporters will have the Security Group's FA algorithms enabled for them.

Enabling FA algorithms for a Security Group

To add one or more FA algorithms to an existing Security Group, follow these steps:

1. Click the Security Group name to open the configuration tray.
2. Expand the *Algorithms* section and click the *Add* button.

Note: Algorithms can also be disabled for the group using the *Delete* icon.

3. In the secondary tray, use the checkboxes to select algorithms.
4. After selecting all algorithms to be added, click the *Add* button.

Once added, the selected FA algorithms will be enabled for all Exporters assigned to the Security Group.

Batch actions

When selecting one or more Security Groups using their checkboxes, the following batch actions can be performed:

- Add one or more Exporters to all selected Security Groups
- Enable one or more FA algorithms for all Selected Security Groups
- Delete all selected Security Groups

Reports

The **Admin > Reports** category includes views and tools that extend the capabilities Plixer Scrutinizer's Report functions.

| | |
|--------------------------------|---|
| Auditing Report | Web interface activity logs (including logs for Report-related actions) |
| Flow Report Thresholds | Management functions for configured Report Thresholds |
| Report Designer | Custom Report configuration and management |
| Report Folders | Folder operations for saved Reports |
| Scheduled Email Reports | Scheduled Report email configuration and management |

Plixer

The **Admin > Plixer** page is used to configure and manage integrations with other products in the Plixer ecosystem.

| | |
|-------------------------------------|--|
| Plixer Endpoint Analytics | Configuration/settings for Plixer Endpoint Analytics |
| Plixer FlowPro Licensing | Plixer FlowPro license details (v20.0.0+ only) |
| Plixer Replicator | Configuration/settings for Plixer Replicator |
| Plixer Scrutinizer Licensing | Plixer Scrutinizer license details |

Note: The pages/views in this category can also be accessed from the **Admin > Settings** view of the Classic UI.

Additional licensing may be required to enable certain integration with certain Plixer components. Contact [Plixer Technical Support](#) to learn more.

Resources

The **Admin > Resources** category consists of pages/views that are used to inspect and manage the different elements of the Plixer Scrutinizer environment.

Feature Resources

The **Admin > Resources > Feature Resources** page lists all Plexier Scrutinizer feature sets alongside their expected CPU (core) and RAM usage and other basic details.

Each feature set consists of the FA algorithms used for the specified monitoring function and the Alarm Policies associated with them.

The **Feature Resources** page can be used as both a reference for recommended *resource allocation* and a means of disabling FA algorithms and Alarm Policies by feature set.

Hint: FA algorithms and Alarm Policies can also be enabled/disabled individually from the Admin > Alarm Monitor > *Flow Analytics Configuration* and *Alarm Policies* pages, respectively.

Feature set details

All predefined feature sets are listed in a table alongside the following details for each one:

- Status (green: enabled, grey: disabled)
- Importance

Note: Features with an *Importance* value of 100 will have a “locked” status indicator and cannot be disabled.

- Number of active alarms indicating resource issues for the feature set
- Expected collector CPU (core) usage
- Expected collector RAM usage
- Number of FA algorithms included in the feature set
- Number of Alarm Policies associated with the feature set

Clicking on a feature set opens a tray containing an enable/disable toggle, as well as all FA included algorithms and Alarm Policies.

Resource use graphs

To supplement the summary table, the following graphs (accessed via the dropdown) can be used to visualize the differences in resource usage, FA algorithm counts, and Alarm Policy counts between feature sets:

- FA algorithms
- Expected CPU usage
- Expected RAM usage
- Alarm Policies

Low resource fallback modes

If low resource fallback is enabled, Plexer Scrutinizer will automatically pause certain functions when the expected CPU and RAM utilization for all enabled features at the current exporter count and flow rate exceed the *recommended allocation values*.

Low resource fallback can be enabled under [Admin > Settings > System Preferences](#) and can be set to three modes/states:

- Enabled and pause feature sets before exporters
- Enabled but only pause exporters
- Disabled

The order in which feature sets are paused is based on their *importance* (lowest first). Features with an importance value of 100 will never be paused.

Features and/or exporters will automatically be resumed when the configured CPU core and RAM allocations can support the additional load based on the observed exporter count and flow rate.

Hint: Regularly check the state of the server health (leftmost) virtual LED in the web interface admin views. As long as it remains green, features and/or exporters will not be paused. While in this state, Plexer Scrutinizer will also continuously attempt to resume paused feature sets.

Additional low resource fallback settings

Users can modify the following collector settings (via the **Admin > Settings** page) to customize low resource fallback behavior:

- Cooldown period before pausing the next feature set or group of exporters
- Number of exporters to pause or resume as a group
- Flow rate multiplier/percentage for accommodating brief, recoverable spikes

Manage Collectors

The **Admin > Resources > Manage Collectors** page provides access to management functions for Plexier Scrutinizer Collectors and Plexier ML Engine deployments, which are split into their respective tabs.

Each tab consists of an overview table that lists all Collectors or ML Engines alongside additional details.

Hint: To change the information displayed in either table, click the *Available Columns* button and select the details to show.

Collectors tab

The **Collectors** table includes the following details for each Plexier Scrutinizer server/Collector by default:

| | |
|------------------------|---|
| Rank | Number assigned to the Collector as part of a distributed cluster |
| Collector | IP address or hostname of the Collector |
| Status | Current status of the Collector |
| Exporter count | Number of Exporters sending flows to the Collector |
| First flow time | Timestamp of the first flow received from an Exporter |
| Last flow time | Timestamp of the most recent flow received from an Exporter |
| Flow rate | Average number of flows received per second |
| Packet rate | Average number of packets received per second |
| MFSNs | Average number of MFSNs/missed flows per second |

Collector details can be viewed from the three-dot menu or by clicking the Collector IP address or host-name.

Collectors can be deleted from the *Batch Actions* menu, which becomes accessible when one or more Collectors are selected using their checkboxes.

ML Engines tab

The **ML Engines** table includes the following details for each Plexier ML Engine deployment by default:

| | |
|----------------------|---|
| ES user | Licensed Plexier ML Engine user |
| Hostname | Hostname assigned to the engine (if set) |
| IP address | IP address of the engine |
| Last modified | Timestamp of the most recent configuration update |
| Port | Port used for communication with the engine |

Clicking an ES username opens a configuration tray for the associated Plexier ML Engine, where its settings can be modified. Engine details can be viewed from the three-dot menu.

Additional actions/options

- Collectors or Plixer ML Engine deployments can be deleted from the *Bulk Actions* tray, which becomes accessible when one or more items are selected using the checkboxes.
- Filtering, exporting, and general page option trays can be accessed via the *Filters* and *Options* buttons.

Manage Exporters

The **Admin > Resources > Manage Exporters** page provides access to monitoring and management functions for all Exporters in the Plixer Scrutinizer environment.

The page is divided into an activity timeline for the top 10 Exporters and an overview table listing all Exporters alongside their activity details and settings. The default view includes known Exporters across all Collectors (*By Exporter*), but it can be toggled to display only unique Exporters (*By Collector*) via a dropdown.

Hint: Clicking on the arrows in the table header changes both the table's sorting order and the datapoints displayed in the activity timeline. To change what details are included in the table, click the *Available Columns* button and select the columns to display.

The following table shows what columns can be displayed for each Exporter based on the view mode (*By Exporter* or *By Collector*) selected:

| Column | Description | By Exporter | By Collector |
|--------------------|--|--------------------------------------|--------------|
| Legend | Color assigned to the Exporter when shown in the activity timeline | Available | Available |
| Rank | Exporter rank based on the selected sorting order | Available | Available |
| Status Icon | <p>Green: Exporter is enabled and available</p> <p>Red: Exporter is enabled but unavailable</p> <p>Yellow: No flows from the Exporter have been received</p> <p>Grey: Exporter is disabled</p> | Available | |
| Exporter | IP address or hostname (when view is set to show hostnames) of the Exporter | Available | Available |
| Status | <p><i>Enabled</i>: Flows will be collected, stored, and processed as normal</p> <p><i>Backup</i>: Flows will only be collected for storage</p> <p><i>unresourced Enabled</i>: Temporarily disabled due to low resources</p> <p><i>unresourced Backup</i>: Temporarily disabled due to low resources</p> <p><i>Disabled</i>: Flows will be ignored</p> <p><i>Unlicensed</i>: Flows exceed the Exporter license count and will</p> | Available | Available |
| 192 | be ignored (set by the Collector) | 5. Features and Functionality | |
| MFSN/s | Average number of Mixed Flow Samples | Available | |

Note: In the *By Exporter* view, all values and states displayed for an Exporter are relative to the Collector specified in its **Collector** column.

Hint: The *unresourced* states can be assigned to Exporters that need to be turned off due to limited resources. This allows them to be quickly returned to their previous state when resources become available again.

Configuring Exporter settings

Clicking on an IP address or hostname in the table opens the configuration tray for that Exporter.

The settings tray is divided into the following sections:

| | |
|----------------------------------|--|
| Information | Basic Exporter details |
| Collectors | List of Collectors that have received flows from the Exporter |
| Interfaces | List of interfaces available on the Exporter |
| Protocol exclusions | List of protocols excluded from collection |
| Flow Analytics algorithms | List of FA algorithms applied to flows collected from the Exporter |
| SNMP | SNMP credential (set) used for the Exporter |
| Tags | Additional tags associated with the Exporter |

The tray also includes a dropdown to set the status of the Exporter (*Enabled*, *Disabled*, *Backup*, *unresourced Enabled*, *unresourced Backup*, or *Unlicensed*), as well as toggles to ignore flow durations, MF-SNs, and/or outages.

To make changes to the Exporter's configuration, click the edit (pencil) icon. Certain settings can only be modified via a different view/page (indicated by the link icon).

Note: When in the *By Collector* view mode, only the **Information**, **Collectors**, and **Interfaces** sections, in addition to the status dropdown and ignore toggles, are available in the configuration tray.

Additional actions/options

- Exporter settings can also be modified from the *Bulk Actions* tray, which becomes accessible when when one or more Exporters are selected using the checkboxes.
- Basic details for an Exporter can be viewed from the three-dot menu.
- Filtering, exporting, and general page option trays can be accessed via the *Filters* and *Options* buttons.

Manage FlowPros

The **Admin > Resources > Manage FlowPros** page allows Plixer FlowPro probes to be registered with Plixer Scrutinizer after a valid license has been added via the **Admin > Plixer > Plixer FlowPro Licensing** page (Plixer FlowPro v20.0.0+ only).

Important: After being deployed and configured, Plixer FlowPro probes must be correctly registered to enable communications with Plixer Scrutinizer.

To learn more about Plixer FlowPro licensing options, contact [Plixer Technical Support](#).

Manage Interfaces

The **Admin > Resources > Manage Interfaces** page displays all available instances on active devices/-Exporters.

Hint: To navigate to more detailed Exporter views, drill down into individual devices from the overview table.

The table can be customized to display the following additional details for each instance:

- Custom description (must be set first)
- Interface alias
- Interface name
- Interface description
- Interface speed
- Custom bits in (must be set)
- Custom bits out (must be set)
- Metering (if applicable)

To change the information displayed in the table, click the *Available Columns* button and select the details to show.

Custom instance settings

Clicking an instance number opens a configuration tray, where the following custom settings can be applied:

- Description
- Bits out
- Bits in

Note: These custom settings are only applied through the Collector and will not affect the configuration of the device itself.

The tray also includes a *Hidden* dropdown, which can be used to hide the device from the overview table, and the option to change the SNMP credential(s) used for the device.

Hint: The *Hide* checkbox can also be ticked to hide instances from the table.

Additional actions/options

- Instance settings can also be modified from the *Bulk Actions* tray, which becomes accessible when one or more instances are selected using the checkboxes.
- The three-dot menu contains shortcuts to the [Manage Exporters](#) page filtered on the device and the configuration tray.
- Filtering, exporting, and general page option trays can be accessed via the *Filters* and *Options* buttons.

SNMP Credentials

The **Admin > Resources > SNMP Credentials** page is used to add and manage SNMP credentials that can be associated with multiple devices in the Plexier Scrutinizer environment.

Once defined, SNMP credentials can be assigned to Exporters from the [Manage Exporters](#) page. SNMP v1, v2, and v3 are supported.

Defining new SNMP credentials

To add a new set of SNMP credentials, follow these steps:

1. On the **SNMP Credentials page**, click the *Add* button.
2. Fill in the form with the following information:
 - A *name* to identify the credential(s) by
 - A *description* of the credential(s)
 - The *SNMP credential type/version* (dropdown)
 - The *community* string to send
 - The *port* to use for communication
 - The *timeout* value or number of minutes to wait for a response
 - The number *retries* after a failed request
 - The *backoff* value or number of minutes to wait between retries

Important: If SNMPv3 is selected as the credential type, the additional fields for the username, context, and authentication details (hash function, password, and encryption) must also be filled in.

3. Verify that the information entered is accurate and click the *Save* button to save the credential(s).

Saved credentials can be edited from the **SNMP Credentials** page by clicking on the credential name to open the configuration form.

To delete one or more credentials, tick their checkboxes to select them and click the *Delete* button.

System Performance

The **Admin > Resources > System Performance** page can be used to monitor resource utilization and performance for individual collectors in a Plexier Scrutinizer environment.

The page is divided into a timeline visualization and a summary table listing all active collectors alongside their utilization details.

Utilization timeline

The timeline can be set to display one of the following metrics via the dropdown:

- CPU (%)
- Memory (GB)
- Host Index size (%)
- Alarm database size (%)

Utilization summary table

The **Active Collectors** table lists all active collectors in a Plexier Scrutinizer environment, along with the following details for each:

- Allocated CPU cores
- Allocated RAM
- Total disk space
- Disk space used by the alarm/event database
- CPU utilization peak
- Host index disk space utilization

Collector utilization view

Drilling into a collector from the main **System Performance** page opens a detailed view showing disk space use for *historical flow data*.

This view contains the following information for the collector:

- Current and predicted/expected disk use per interval based on configured data history settings (**HD Utilization** graph and **Utilization per Interval** table)
- Current number of days saved vs. configured number of days to save historical flow data
- Current and predicted disk utilization per roll-up interval bucket
- Current and recommended storage allocation

Hint: [Recommended CPU core and RAM allocation charts](#) can also be accessed from the dropdown on this page.

Hint: The Classic UI Admin page can be accessed via either the icon next to the *Admin* text in the web interface header or the *Classic Admin* link in the tray.

5.1.7 Classic UI

As part of Plixer Scrutinizer 19.0.0, the web interface UI received a major revamp to improve usability and support current and future feature additions.

To maintain continuity, the Classic UI remains accessible either via the URL `https://scrutinizer_ip/oldui/` or by toggling the appropriate preference setting under the user menu.

No EOL date has been announced for the Classic UI.

Dashboards

Overview

Important: The functions and features included in the Classic UI's **Dashboards** tab have been reworked and optimized in more recent releases of Plixer Scrutinizer. They can now be accessed by navigating to **Monitor > Dashboards** in the new UI. To learn more about upgrading to the latest version of Plixer Scrutinizer, see the [Updates and upgrades section](#) of this documentation.

Dashboards are used to create custom views of precisely what the user or group of users wants to see when they log in. Multiple unique dashboards can be created.

- With the right permissions, these dashboards are customizable per login account.
- All dashboards created by any user in a usergroup are available to other users in the same usergroup. The default is read-only access.
- Each dashboard can be manipulated and shared with others.
- The Read-only permission (check box) is used to grant others the ability to manipulate a shared dashboard.

Dashboard administration

In the upper left-hand corner of the dashboard there are three drop down menus.

1. Gear with down arrow:

- If the user has permission, this option can be used to change the dashboard name.
- Set the default dashboard when the Dashboard tab is clicked.
- If the user has permission, the user can make a dashboard **Read-Only** to others whom will be viewing the same dashboard. Leaving unchecked allows them to change the dashboard which includes rearranging as well as adding and removing gadgets.
- The user with ownership of the dashboard is also displayed with the dashboard ID. The dashboard ID can be accessed directly through a URL: `https://<server>/dashboard/id/<dashboard_id>`
- A user wanting to modify a dashboard that doesn't have permission, can copy the dashboard and make changes to the copy. Copying a dashboard requires permission as well.

2. Dashboard name:

- Use this menu to select the desired dashboard to view.
- The default dashboard is displayed at the top of this menu.
- A '*' after the dashboard name indicates that it is read-only.

3. Configuration:

- **Add a New Gadget:** When viewing a dashboard, this option can be used to add additional gadgets. Select the category of gadgets in the drop down box at the top. To add gadgets, click on them.
- **Copy this Dashboard:** Use this option to make a copy of the dashboard which can then be modified by the user. This requires either the “Create New Dashboards” or **Dashboard Admin** permission.
- **Create New Dashboard:** If the user belongs to a usergroup that has permissions, this option can be used to create a new dashboard. This requires either the **Create New Dashboards** or **Dashboard Admin** permission.
- **Remove Dashboard:** Use this option to remove a dashboard from the menu. Both read-only and user created dashboards can be removed and added back to the menu. This is done under **Configuration > User Dashboards**.

Creating a new dashboard

1. Navigate to the **Dashboard Configuration > Create New Dashboard** page.
2. Use the filter on the left to find the desired gadgets. Use the drop down box below the filter to select a category of gadgets.
3. To add gadgets, highlight them in the **Gadgets Available** box and drag them to the **Gadgets Added** box. Use the shift and CTRL keys to select multiple gadgets at once.
4. Uncheck the **Read-only** box if the goal is to give others permission to view AND modify the dashboard. Permission can be granted to give others a read-only view of the dashboard under the **Grant** tabs. Users able to view a read-only dashboard will be able to copy it and manipulate the copy.
5. Give the dashboard a name before saving it.

Note: To add gadgets to a dashboard, one of the following is required: 1) The user must be the creator of the dashboard 2) The creator of the dashboard must have unchecked **Read-Only** in the gear menu or 3) the user must be a **Dashboard Administrator** for the user group.

Creating a new gadget

1. From the *Dashboard name menu** select the dashboard you would like to add a custom gadget to.
2. Navigate to the **Configuration > Add a new gadget** page. Click on the **Add a gadget > New** button.
3. Enter the new gadget name and the Gadget URL.
4. Save the new gadget to a panel. You will now see the new gadget on the dashboard you selected in step 1. It can now be found in the **Custom** gadgets list and added to other dashboards.

Note: External URLs must have an http(s) prefix to avoid a 404 error. Gadgets may not load if you specify HTTP content when Scrutinizer is using HTTPS.

Gadget configuration

There are several configuration options in each gadget or window in the dashboard. Each is represented by an icon, some of which don't appear until the mouse is moved over the window.

- **Timer:** This value decrements to indicate the next refresh of this gadget. Set the refresh frequency by clicking on the gear icon.
- **Gear:** The spinning gear icon can be used to rename the gadget and to set the refresh rate.
- **Refresh:** Press the refresh or recycle icon to force the reload of the contents of the gadget. (Will also happen automatically when the timer runs out.)
- **Move:** Click the four-headed arrow icon, hold, and drag the gadget to a new location in the dashboard.
- **X:** This icon is used to remove the gadget from the dashboard. It can easily be added back later and it will remain in the gadget inventory for use in other dashboards.
- **Resize Arrows:** Located in the lower left and right-hand corners of the window, these icons are used to resize the window.

User and usergroup permissions

- User Dashboards:

Use this option to select the dashboards a user will have visible in their menu of available dashboards.

- Select the user from the drop down box at the top. To see other users, the user must be a member of a usergroup with the **Dashboard Admin** permission.
- Select the dashboards in the “Available” box and move them to the **Visible** box to grant permission.
- Notice the filter on the left. If granting permission to multiple users, administrators generally use the **Usergroup Dashboards** option. This requires the **Dashboard Admin** permission.

- Usergroup Dashboards:

Use this option to select the dashboards a usergroup will have visible in their menu of available dashboards. This feature requires that the User be a member of a User Group with **Dashboard Admin** permission.

- Select the user group from the dropdown box at the top.
- Select the dashboards in the **Available** box and move them to the **Visible** box to grant permission.

Note: Even if a dashboard is added to the menu for a specific user or for all users in a usergroup, individuals can still remove a dashboard from their menu.

Additional permission options can be found under:

- **Admin tab > Security > Users** to set the default dashboard the user will see when first opening the Dashboard tab.
- **Admin tab > Security > Usergroups:**

1. Choose Dashboard Gadgets

- Click the “Dashboard Gadgets” value for the usergroup you want to change
- Uncheck “All Dashboard Gadgets”
- Move the individual gadgets the selected usergroup should be able to view from the “Deny” to “Allow” box

2. Choose Feature Access

- Click the **Configure** link in the "*Features*" column for the usergroup you want to change.
- With **Predefined** selected, add or remove the **Dashboard User** and **Dashboard Administrator** roles.
- With **Advanced** selected, add or remove individual features like **Create Dashboards**.

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, IPFIX and sFlow volume and other server metrics. Vital information is reported for all servers in a Distributed collector Environment.

The following dashboard gadgets are available:

- **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. 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. Other applications using the same database will cause this number to increase.
 - **Read Req:** The number of requests to read a key block from the cache. A high number of requests 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 Plixer Technical Support](#) if the query cache is under 1MB.
 - **Queries:** Tracks the number of queries made to the database. More queries indicates a heavier load to the database 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 database server with other applications.
 - **Threads:** Threads are useful to help pass data back and forth between Scrutinizer and the database engine. The database server currently manages whether or not to utilize the configured amount of threads.
 - **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 database configuration file and adjusting the key buffer size setting. A general rule of thumb is to allocate as much RAM to the key buffer as possible, 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 too much memory is allocated, the risk is seeing further degradation of performance because the system has to use virtual memory for the key buffer. The [check tuning](#) interactive `scrut_util` command can help with recommended system settings.
- **Syslogs 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](#).

Status

Overview

Important: The views/pages included in the Classic UI's **Status** tab, including all the information they provide, have been integrated into the **Monitor**, **Explore**, **Investigate**, and **Reports** tabs of the updated Plixer Scrutinizer UI. To learn more about upgrading to the latest version of Plixer Scrutinizer, see the [Updates and upgrades section](#) of this documentation.

Interfaces

The Top Interfaces is the default view of the Status tab unless it is modified by the user by editing their profile. Be sure to mouse over items on this page before clicking as the tool tip that appears can be very helpful. The columns of this table of interfaces includes:

- **Checkbox:** Check off the interfaces desired to include in a single report and then click the trend icon at the top of this column.
- **Icon color status:** The color is determined by [CrossCheck](#). Mousing over the icon will provide polling details.
- **Flow Version:** 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. Selecting a report from here will run an inbound/outbound (bidirectional) report for the last 24 hours in 30 minute intervals. The user can drill down from there.
- **Arrow down menu:** Clicking this presents a menu:
 - Reset the high watermark(s) in the Inbound/Outbound columns
 - [Interface Details](#)
 - Device Overview
- **Inbound/Outbound:** these 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.

Menus

The Status tab is one of the most popular views for gaining quick access to all the NetFlow capable devices and interfaces that are represented in the flows received. The default view is a list of all flow sending interfaces however, this can be modified under Admin tab > Security > User and then click on a user. Click on Preferences in the modal and find the “Default Status View” and choose from one of several options.

Gear

- Select how many interfaces should be displayed before utilizing the pagination.
- Decide whether the interfaces should be listed by highest percent utilization or by highest bit, byte or packet rate.
- The refresh rate of the top interfaces view.
- Toggle between IP/DNS depending on how the flow devices should be listed.

Top Right icons (mouse over for tool tips) are for:

- Primary Reporting Server: Indicates if the server is a primary server or a collector.
- Scrutinizer Server Health: View the system vitals of the server. Find out where the system needs resources.
- Scrutinizer Software Health: View the status of the system components.
- Exporter Health: View a list of flow exporters. Find out which devices are under performing.
- Magnifying Glass: Search for a specific IP address.
- Down Arrow:
- Scrutinizer Version: The current version the server is running on.
- Check for Updates: Connects to Plixer to see if updates are available.
- Contact Support: Launches a web page to contact Plixer for support.
- Online Help: Launches this manual!

- **Manage Exporters:** Launches > Admin tab > Definitions > Manage Exporters to see what devices are sending flows to the collector(s).
- **Join the beta program:** Fill out a form online to join the beta program.
- **Log Out:** Log out of Scrutinizer

Top Right icons below logout are for:

- **Clock:** Schedule a reoccurring email of the top interfaces view.
- **@:** Email on demand the current interfaces view.
- **PDF:** Create a PDF on the current interfaces view.
- **CSV:** Export a CSV file containing the content of the current interfaces view.

Top menus along the top include:

- **Run Report:** Need to design a custom report? Select from all available elements, operation columns, devices, and time ranges to get the exact data needed.
- **Top:** Not sure what report to run? Select from over a dozen canned reports that will include data from all flow exporters.
- **Search:** Need to find a host or IP address?
 - **Host Index:** Run a report by “Host Index” to quickly determine if the host has ever been on the network. It searches the index rather than the saved flows. This search requires that Host Indexing be turned on in Admin>Settings>System Preferences.
 - **Saved Flows:** Run a search against all “Saved Flows”. This search actually queries the database and can take a bit longer. NOTE: Depending on archive settings, the desired data may have been dropped. This search is more flexible and allows for searching by host address, username, wireless host or SSID across some, or all, flow exporters for a specified timeframe.
- **System:** These are advanced reports used by engineering when trying to understand why something isn’t working. In a future release, they will be moved to the Admin tab.
 - **Available Reports:** Lists the report and the number of templates received that contain the necessary elements.
 - **Flow Report Thresholds:** Lists all the reports that have been saved with a threshold.

- Templates: These display the device and each flow (NetFlow v9/IPFIX) template exported from that device.
- Vitals: These are reports on the system resources from the flow collection and reporting servers.
- Views:
 - *CrossCheck*
 - Device Status: Lists all of the flow sending devices with corresponding details. The color of the icons is determined by CrossCheck.
 - Interfaces: This is the default view of the status tab before a report is run.
 - SLA: Lists all of the flow sending devices which by default are being pinged by the collector. The response from each ping is used to determine the Response times and availability for each device polled.
 - 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.
 - Vendor Specific: lists reports that will work ONLY if the collector is receiving the necessary templates from the flow exporters.

Left hand side menus provide three views:

- Device Explorer: Displays a list of all the Groups of devices. Explained below.
- Current Reports: Displays the current report after a report is run on one or more interfaces. Explained below.
- Saved: Displays all of the saved filters/reports that can be run. Explained below.

Device explorer

Organize devices by moving them into groups.

- New: 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.
- Grouped: A group of devices that typically share one or more attributes.
- View: Displays the map for the devices in the group.
- Reports: Select a report to run against all of the flows collected from all the devices in this group.
- Copy: Make a copy of the group and give it a new name.
- Modify: Modify the membership of the objects in the group.
- CrossCheck: *View CrossCheck* for the devices in this group.
- SLA: 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: 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*. Expand the flow exporter for the menu.
- Reports: Run a report on the flows coming from the device. *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. Click on an interface to run a report. Selecting an interface (or All Interfaces) from this list will open a report menu. Select and run a report for the last 24 hours. ALL Interfaces reports will default to Inbound as described above, selecting a single interface will report on both Inbound and Outbound.

- **Properties:** Modify the properties of the device.
- **Device Overview:** Provides the overall status of the device by leveraging data from CrossCheck, the poller and the alarms.
- **Show Interfaces:** Displays a list of all active interfaces for the device in the main window of the page.
- **Other Options:**
 - **Alarms:** Displays the outstanding alarms for the device.
 - **Interface Details:** launches the [Interface Details](#) view which lists SNMP details about the device including the interface speeds.
 - **Flow Templates (Advanced):** displays the templates (e.g. NetFlow v9, IPFIX, etc.) currently being received from the device.

Current report

This tab opens when a report is selected from the report menu. All of the icons that appear in the top left are explained in [Network Traffic Reporting](#).

Filters can be added to the report by grabbing items in the table and dragging them to the left or by clicking on the “Filters / Details” button.

Saved reports

Saved reports are saved filters or reports which display the selected data on one or more interfaces across potentially several devices. When Saved is clicked the user is returned to the Current Report view and the filter contents are displayed.

This tab lists any reports that were saved and provides folder management utilities:

- **Add Folder:** Select ‘Add Folder’. A text box will open to enter a folder name which is used for organizing saved reports.
- **Manage Folders:** Select ‘Manage Folders’. A new browser tab will open to Admin > Reports > Report Folders. From here, bulk folder/saved report management can be accomplished by moving several reports in and out of a folder. New folders can be created or deleted 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 a tooltip 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 tooltip 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 pulling them out of the desired folders. Reports can be viewed by clicking on the report name. They can also 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 the user to easily find reports within the report list or folders.

Network traffic reporting

Reporting is the interface customers spend the most time in. This page outlines the functionality that can be found in all of the menus of the status tab. If the user is more of a visual learner, training videos are available on the plixer web site.

Templates

Unlike NetFlow v5, NetFlow v9 and IPFIX use templates to dynamically define what is being sent in the flows. Templates are the decoder that is provided by flow exporter. They are used by the flow collector to decipher and ingest the flows.

The reporting options (I.e. menu) available on every flow exporting device is dependent on the values in the template. For example, when clicking on a flow exporting device to launch the report menu, the report “Vendor by MAC” under “Source Reports” will not appear if the MAC address is not exported in the template from the device. If another flow exporting device is selected the user may find that the “Vendor by MAC” report does appear. It all depends on what is being exported in the templates from each device.

This template intelligence becomes critically important when trying to understand why the system is behaving differently with oddly formatted vendor flow exports. For example, some flow exports do not provide an ingress or egress interface. When this is the case, the device will not show up in the interface list of the Status tab. To run reports, the user will have to find the device in the Device Explorer.

The available reports for each device can be observed by navigating to Status > System > Available Reports. The Available Reports view provides the ability to view, sort, and filter report lists by Group Name, Report Name, and Template Count.

Report types

There are hundreds of report types in the database. Most will never appear in the menu because they only appear if the necessary elements are available in the templates exported by the device. When reports are run, they group on the fields displayed. For example, the report Conversation WKP groups on Source IP address, WKP (common port) and Destination IP address. For answers to questions about anything not listed here, please contact Plixer support directly.

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 report is displayed in the main section of the screen to the right of the Current Report frame.

- **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. Greater than 11 entries can be displayed by visiting the gear menu.

Tip: 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 overstated 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 this is 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. This could happen for either one of the following reasons:

- Historical data settings are too low for the time frame requested. To increase the historical data retention, go to **Admin tab -> Settings -> Data History**.
- Flows are not being, or have not been, 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 report in CSV format.
- **PDF** (PDF) downloads a pdf file containing the current report.
- **Email** (@) is used to email the report displayed using the current filter(s). 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.
- **Filters / Details:** Button: clicking this opens the Report Details modal with the following tabs:
 - Collector Details: displays the collectors(s) that contained the flow exporters for this report.
 - Exporter Details: details about the exporters that are providing flows for this report.
 - 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.
 - Report JSON (API)

Gear icon

Clicking on the Gear icon will display many more reporting options:

- **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.
- **Data Source:** Auto, 1m, 5m, 30m, 2hr, 12hr, 1d, 1w. This tells the system which tables to take flows from when querying data used in the report. Generally the default is taken as the database has been optimized for this setting. This option allows the system to query several days of 1 minute tables (i.e. non rolled up data) when searching for specific values that may have been dropped in the higher interval data.

Warning: Selecting 1m (i.e. 1 minute tables) for a 24 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 flow exporting devices.

Note: The number of intervals used for granularity is set via the “Target graph interval” setting found under the Admin tab > Settings > reporting.

- **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 [Interface Details View](#) or as a report filter. When multiple interfaces are included in a report, the calculated interface speed will 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.
- **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.
- **Data Mode:** This specifies the source of the data. The two values are Summary or Forensic. Both values at one minute intervals represent 100% of the data with some significant differences:
 - **Summary:** Has been aggregated based on a definable tuple. The default aggregation is on the Well Known Port. This means that the source and destination ports are dropped as is everything else in the flows that isn't needed to run most of the reports. Visit Data Aggregation to learn more about what is kept in Summary tables. As result of this optimization, the table sizes are much smaller which results in faster rendering of reports. This is the default data used to create the higher rollups (E.g. 5 min, 30 min, 2hr, etc. intervals).
 - **Forensic:** This is the raw flows with no aggregation and all of the elements are retained. It is used for vendor specific reports and for a few reports which display the source and destination ports. These tables are not rolled up in SAF mode and therefore, history trends that use the forensic tables will be limited to the length of time that the 1 minute interval data is saved. If however, the server is running in traditional mode, roll ups will occur as summary tables are not created in traditional mode.

How is the 95th percentile calculated?

The data points in the graph are sorted from smallest to largest. Then the number of data points is multiplied by .95 and rounded up to the next whole number. The value in that position is the 95th percentile.

Example :

Data points = [1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,21,22,23,24,25]

25 data points *.95 = 23.75

Round 23.75 up to the next whole number = 24

The value in position 24 is the 95th percentile, which in this example = 24

If a report has less than 21 data points, the largest number is always the 95th percentile. Increase the granularity in the report for increased accuracy.

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.

Note: The system will auto determine the number of intervals or data points in a trend. [Click here](#) to learn how trends determine intervals.

- **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 the pagination is clicked, 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 Summary or Forensic 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 the pagination is clicked at the bottom of the table, all of the data that makes up the 11th color (I.e. gray) comes into view.

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 reporting 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:** This is configured with a filter. See the Business Hours entry in the Filters - Include or Exclude Data section below.

Saved Reports

Refer to the [Saved Reports](#) section in the Status Tab Overview 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 / Details” button in the left pane of the screen will popup a 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 one 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.
- The following special case filters are also available:
- Business Hours filter provides the ability to limit the reporting data between the start and end times, change the reporting timezone, and also select the days of the week for the report. The default Business Hours settings are defined in **Admin > Reports > Settings > Business Hours End and Business Hours Start**. Business hours days of week default to Monday - Friday.
- Port Speed, this filter allows the user to set a port speed for a report.
- Sample Multiplier filter allows the user to set a multiplier value for sampled flows to recalculate to full flow values.
- Wildcard Mask filter allows the user to add a custom mask 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 pre-defined 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 apply 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 on the left.
 - 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 the Apply button at the top of the filter list.

Archived Data: Three types of historical tables are maintained for each NetFlow exporting device.

- Forensic - This was formally the Conversations table. This table contains the actual raw flows.
- Summary - This table contains 100% of the aggregated raw flows with no dropping. By default flows are aggregated based on the WKP (common port). Aggregation can be read about in the [Data History](#) section. **If filters are used, these are the only tables used in the report.**

- Totals - This contains the actual amount of total traffic in and out an interface for each interval before flows are rolled up into the Summary table. This table must be maintained as the 5 minute interval and higher Summary tables only contain the top 1,000 by default for each interval. 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.

Note: Interface utilization reports based on NetFlow or IPFIX flows seldom, if ever, match exactly to the same interface utilization report based on SNMP counters. Remember, it can take 15 or more seconds before a flow is exported. SNMP, on the other hand, is more realtime and the counters include other types of data not reflected in flows (e.g. ethernet broadcasts).

Filter Logic:

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' rule applies. 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.

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 trying to observe traffic between two IP Addresses, use the Host to Host Filter. There is also a filter for subnet to subnet.
4. 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 any value in a row within the table located below the report graphic will present a menu of available report types. Remember, the report options displayed is dependent on the values in the templates coming from the device(s) used in the current report. When selecting a report in this way, the value 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** can pass the IP address selected in the URL to the application. Default menu options are:

- **Report to ISP** - Report suspicious behavior
- **Search**
- **Alarms**
- **Lookup** - Whois Lookup
- **GEO IP** - Geographical lookup
- **Talos Reputation Center** - Leverages the Talos Geographical and detailed IP address information.
- **New applications** can be added by editing the applications.cfg file in the /home/plixer/scrutinizer/-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 new menu option 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 language key for the button name. **WMIUsersDescr** is the language key for the description.

Then, in applications.cfg, add an entry to reference these language keys and associate the URL with them. Add the following line without quotes: .. code-block:: bash

```
“WMIUsers, /cgi-bin/currentUsers.cgi?addr=%i, WMIUsersDescr”
```

Note: The applications.cfg file is located in the /home/plixer/scrutinizer/files/ folder and is used to map the URL of the new menu options to the language keys in the languages database table. (as explained above)

Flow view interface

The Flow view provides 100% access to all the elements that were exported in the raw flows. Some columns or elements are generated by Scrutinizer. The Flow view interface retrieves all of the flows that match the values requested in consideration of the filters applied.

Notice:

- Filters are passed to Flow View when drilling in.
- Use the filters drop down box to find data in specific columns. NOTE: The sourceOrDestination option is not a column.
- Click on the column headings to sort.

IPFIX, NetFlow, sFlow, NSEL, etc

Flow View is used to view flows generated by 100% of all flow technologies. The collector can save any type of NetFlow v1, v5, v6 and v9 data inclusive of IPFIX and other variants including NetFlow Security Event Logs (NSEL), NetStream, jFlow, AppFlow and others. This report provides access to view any and all flows received by the collector given the filters applied. 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 if the flow 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 Well-KnownPort).

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: port 1234
- Is there an entry in the local database for 1234 (e.g. HTTP)
- If Yes: save it as the common port (1234)
- else if: is port 5678 labeled in the local database (e.g. HTTPS)
- If Yes: save it as the common port (5678)
- else save 1234 as the common port (e.g. Unknown)

Note: If both source and destination ports 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 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](#):

Warning: The following names are subject to change depending on the version of firmware running on the hardware.

- 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 at intervals specified under the Admin tab under settings. In order for rollups to occur, the template exported must provide the element: octetDeltaCount. Please [contact Plixer Technical Support](#) to change the rollups to occur on an alternate field. Visit the Admin Tab > Settings > [Data History](#) page to configure how long to save the data.

Report thresholds

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 available by clicking on the “Filters / Details” button located in the left hand frame of the Report view. Instructions for adding thresholds to reports are detailed below. Thresholds are monitored every 5 minutes, based on the last 5 minute interval.

To add a threshold to a report:

1. Save a report. Thresholds can only be added to [Saved reports](#). 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 the report isn't saved first, the interface will prompt the user to enter a report name and save it when they enter the threshold modal.
2. Click the Add button to the right of Threshold in the left hand pane. 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.

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.

Threshold comparison options are:

- **Greater than or equal to (>=)**

or

- **Less than or equal to (<=)**

3. The threshold value is entered in the textbox after the word “than”. 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
4. After completing entry of the fields listed above, click the **Save Threshold** button. To exit the threshold modal without saving, click the **Close** button.
 5. The **Select Notification Profile** modal displays next. If notification profiles have been configured, select the appropriate one from the dropdown selector. To configure new notification profiles, click **Manage Notifications**. A new browser window opens to the Notification Manager page. 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.
 6. 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. Select a profile from the dropdown, or click Manage Notifications to create one. Selecting Save and finish without adding a notification to the threshold is also an option. An alarm will still be generated when the threshold is violated even without a notification included in the threshold configuration.
- Thresholds are checked against whichever column the saved report is ordered by. For example: if the report is ordered by packet rate the threshold is checked against packet rate, if a report is ordered by total bytes the threshold value is checked against total bytes.

Scheduling a report

Prerequisites

- The email server needs to be configured in [Admin > Settings > Email Server](#)
- One or more report(s) need to have been ‘saved’

Schedule reports from the Status tab

- Either create a new saved report or select existing Saved Reports from left pane in Status tab, then select the saved report(s) from the list. Make sure the report is saved with a ‘last’ time frame (E.g. Last Seven Days).
- Click the ‘clock’ icon to Schedule an emailed report. It can be found at the top under Current report. It will launch the Schedule Report modal.
- Schedule Report modal
 - Email Subject: This field is mandatory and is auto filled with the Report name when coming from the Status tab. The subject of the email can be changed here.
 - PDF / CSV: Check these boxes to attach the report in PDF or CSV format.
 - Frequency and Time: This report will kick off on the current day:
 - Hourly: Specify the minute each hour that report(s) will run
 - Daily: Specify the hour, minute, and AM/PM that report(s) will run each day
 - Weekly: Specify the hour, minute, AM/PM, and day of week that report will run each week
 - Monthly: Specify the hour, minute, AM/PM, and day of month that report will run each month
 - Recipients: Enter the email address(es) of recipients here. This field is mandatory and must include at least one recipient’s email address. Multiple email addresses may be separated by commas, semi-colons, or spaces, and may be entered all on one line, or on separate lines.

- **Include/Exclude:** This section shows which reports are in the scheduled report (Included) and which ones are not, but are available to add to this schedule (Excluded). At least one report must be in the Include section. By default, when scheduling from the Status tab, the saved report being viewed will be automatically included. Add more reports to a scheduled report by selecting from the Exclude list and clicking the double left arrows (<<) to move it to the Include list.
- Click 'Save' to add any selections to the Scheduled Report list.
- To monitor and manage the Scheduled Report go to Admin > Reports > Scheduled Reports.

Important: Make sure the report is saved with a 'last' time frame (E.g. Last Seven Days). If the frequency is set to 'Hourly' for example, a report will be emailed every hour which shows the last seven days. Also, in order to avoid excessive processing overhead, try to avoid scheduling multiple reports to run at the same time.

Managing scheduled reports

Scheduled reports can be managed at: [Admin > Reports > Scheduled Email Reports](#). This page will list all existing scheduled reports. Columns in this page include:

- Action
- Edit Schedule – opens the Schedule Report modal allowing changes to any aspect of the scheduled report.
- Send Now – email this report on-demand
- Disable - checkbox
- Email Subject
- Schedule
 - Hourly
 - Daily
 - Weekly
 - Monthly

- Time – scheduled time for report
- Day of Week – scheduled day for report
- Day # - scheduled day of month for report
- Execute Time – the amount of time taken the last time the scheduled report has run
- Last Sent – timestamp for last time the scheduled report has run
- Recipients – email addresses configured to receive this scheduled report

Note: Email Subject and Included Reports do not auto fill when scheduling from the Admin tab.

- The following buttons provide other actions:
 - Delete - deletes any selected Scheduled Reports (leaves the Saved reports intact)
 - Schedule Reports – opens the Schedule Report modal, allowing for scheduling of one or more Saved Reports.

Best practices in scheduling reports

The Admin > Reports > Settings includes all of the server preferences that affect reporting. The following settings are critical to Scheduled Reports:

- Max Report Processes - Each report that is run will use this as a maximum number of sub processes. It breaks reports up by time or exporters depending on the method that will be faster. The default is 4 and the default memory allocation per process is 1024MB.
- Max Reports per Email - The maximum number of saved reports a user is allowed to include in a scheduled email report. Including too many reports in a single email can result in timeouts. The default is 5.
- Max Reports per Interval - This is the maximum number of reports that users are able to schedule for the same minute. The default is 5.

Note: Here's how to calculate how scheduling reports will affect the server. Four processes are created per report x 1024 MB = 4096 MB per report. The maximum scheduled reports per interval is 5 * 4096MB which is equal to 20,480MB. If the server is configured with 16GB of memory, this feature will not work. To continue either decrease the number of reports per interval or add memory to the server. In addition to the memory used by the scheduled email reports, keep in mind the other tasks that are consuming resources.

When possible, schedule reports at off-times, when other processes are resting. Avoid scheduling reports during heavy daytime processing or during server or database backup times. Daily reports can run anytime during the day or night by saving the report with a timeframe of 'Yesterday', which will always run from 00:00 – 23:59 of the previous day.

Run report options

This feature allows the ability to create custom reports. Options available for selection include data elements (fields), operation columns (packets and bits), devices , and timeframe to run the report on. This feature is useful when field combinations not available in predefined report types are required.

Step 1: select data elements

The first step in creating a custom report is choosing the data elements (fields) to include in the report.

The selection list includes the basic tuple elements, plus any Plexier manufactured fields based on those elements.

By default, the selected list is empty, select one or more from the available section and drag to the selected section. A minimum of one data element is required for the report to run.

Step 2: select operation columns

Click on the Step 2 header line to expand this section.

In this step, the packets (packetdeltacount) and bits (octetdeltacount) elements are chosen and configured for which operation will be applied against them.

By default, both packetdeltacount and octetdeltacount are included. Either can be removed by clicking the 'x' to the right of the element. Additional columns of either of these elements can also be added (to include other operations against them) by clicking 'Add Row' and selecting the element.

A custom report requires at least one operation column.

Operations available are:

Sum

Totals the values, per row and a total for the report

Min

Minimum values per row and per report

Max

Maximum values per row and per report

Average

Averages the values per row and per report

Step 3 (optional): select devices

This selection determines which device(s) the custom report will run against and report the data for. The list of devices is limited to those that are exporting the basic tuple elements as shown in the selection box in Step 1.

By default, all devices are selected. Limiting the selection of devices to report against can be done either by:

- Clicking **Select All** and dragging all of the devices to the available section, then select the devices to report on, and drag back to the selected side. This would be the preferable method if there are a large number of devices in the list. The search box can also assist in the selection process.

Or:

- Selecting the devices to NOT include in the report and drag from the selected section to the available section.

Step 4 (optional): select time range

In Step 4, the timeframe that the report is run for can be changed to any of the predefined timeframes, or set to a custom timeframe. If this is not changed, the report will default to the Last Hour.

Step 5: run report

This step is grayed out until:

- At least one data element from Step 1 is selected
- At least one operation column from Step 2 is included
- At least one device from Step 3 is selected

With the criteria met, click the **Run Report** button to generate the custom report.

Saved flows & host index searches

The Search tool is launched by navigating to **Status > 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:

- 1) Saved Flows search
- 2) Host Index search

Note: Only the 1 minute interval tables contain 100% of all flows collected. To make sure the system is querying 1 minute interval data, limit the search to under 1 hour of time. Visit the [Admin>Settings>Data History](#) page and increase the “Maximum Conversations” saved per interval value to increase the volume of flows saved per interval. Be aware that this will likely require more hard disk space. Before making any changes, visit the [Dashboard tab>Vitals](#) (or [Status>System>Vitals](#)) to view how much hard drive space is being consumed.

The **Saved Flows** search allows a 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

Note: The User as Source and User as Destination search fields allow a search by Username if they are being collected from the authentication servers.

Other search options:

- 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 Saved Flows search, a Host to Host report will return the results of the search.

Host Indexing

The **Host Index** search is used to perform extremely fast searches for hosts. The index is a list of all IP addresses that have been seen in flows either as the source or destination of a flow. Because it is an index, it does not contain the entire flow contents.

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 Host Index search requires that Host Indexing in **Admin -> Settings -> System Preferences** is enabled.

Note: The host index will retain IP addresses for 365 days by default. To make changes, visit **Admin tab -> Settings -> Data History** and modify the **Days of host index data**. Keep in mind that even though the host index has the IP address searched on, the flows used to build the index may have been dropped by the rollup process.

Username reporting

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, LDAP, TACACS+, Radius, etc.). 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 in Scrutinizer.

- user name reporting
- Alarms reporting with user name
- Saved Flows search by user name

User name reports are available under:

- Top reports category;
- Device-specific report categories (such as SonicWALL, Palo Alto, or wireless reports);
- Source / Destination > User Name by IP reports.

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.

Saved Flows 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.

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. Since 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, even if the topology has since changed.

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.

Note: This feature requires next-hop 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 user interface 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.

CrossCheck & Service Level reports

The CrossCheck and Service Level Reports located in the [Status tab](#) provide 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 weight of importance to the device depending on any problems found. The Fault Index (FI) is the total value 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`) updates the FI for each device. It queries the above list for any device that has violated the configurable FI threshold. For each device that exceeded the threshold, CrossCheck 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).

Fault Index

The "Exceeded CrossCheck Fault Index" (Policy) in the [Alarms tab](#), performs the following when a Cross-Check syslog comes in:

- The syslog comes into the 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 if, notification hasn't already gone out. The IP address of the device is inserted into the `xcheck_notifications` table. The configured notification profile for the host is then executed.
- Else, if the configured notification profile in mapping for the host isn't configured: nothing happens.

CrossCheck table

The main CrossCheck table displays an inventory of all hosts being monitored as well as their status in each of the applications that are monitoring the device. The overall Fault Index (FI) on the far right provides the status of devices in the Scrutinizer maps and all of the 3rd party applications monitoring them. Click on the headings to sort. The query time frame is the last 1 minute by default. The buttons are explained below.

- CrossCheck Summary: Explained below.
- Thresholds: This sets the threshold at which the sub icon on a device changes color. Color thresholds are based on a percentage of the Fault Index threshold.
- 3rd Party Methods: Data collected from other applications that will impact the Fault Index (FI) for a device.

CrossCheck action:

- Device Overview
- Edit: Polling and Appearance: This launches the modify object view in the mapping utility.
- Run Report: Flow Report: This runs the flow volume report for the device.
- Host: List the IP address or DNS host name for the device. Mouse over for tool tip.
- Flow: Method used to determine if a device is sending flows.
- Poller: Method used to determine if the poller can ping the device.
- Optional Third Party Methods: Contact your vendor to create new methods (e.g PRTG, Solarwinds, etc.). When active, a new column appears.
- Fault Index (FI): The FI provides the overall status of a device across all 3rd Party Methods.

CrossCheck summary

The CrossCheck summary is an bar chart of the CrossCheck list. It displays the number of unique hosts found across all 3rd Party Methods as well as the Fault Index.

- Refresh: Set the interval. The default is every 5 minutes.
- 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.

- **Define Networks:** Specify a subnet to group IP addresses found 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.

Service Level Report

The Service Level Report 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. PRTG™, WhatsUp Gold™, Solarwinds Orion™, etc.).

Vitals reporting

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 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 Plixer Technical Support* if the query cache is presently under 1MB.
 - **Queries:** Tracks the number of queries made to the database. More queries indicates a heavier load to the database 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 database server with other applications.
 - **Threads:** Threads are useful to help pass data back and forth between Scrutinizer and the database engine. The database server currently manages whether or not to utilize the configured amount of threads.
 - **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 resource strapped implementations, this can cause performance to degrade quickly. Users can adjust the amount of memory allocated to the key buffers by modifying the database configuration file and adjusting the key buffer size setting.

A general rule of thumb is to allocate as much RAM to the key buffer as possible, 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 too much memory is allocated, the risk is seeing further degradation of performance because the system has to use virtual memory for the key buffer. The *check tuning* interactive `scrut_util` command can help with recommended system settings.

- **Distributed Heartbeat** and **Distributed Synchronization**: provide further insight into internal communications in a Distributed environment.
- **FA Counts** and **FA 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** display metrics by exporter and also by listening port for:
 - **MFSN**: Missed Flow Sequence Numbers are generated if the device exporting the flows can't keep up with the traffic, the flow packets are being dropped by something on the network, or the flow collector can't keep up with the rate of flows coming in. 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. 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**: 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. 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** provide reporting performance metrics.
- **Storage**: 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** displays runtimes per Scrutinizer automated tasks such as nightly history expiration, vitals data collection, etc.
- **Totals/Rollups Times** shows time durations for totals, rollups, and data inserts in the database per flow template per exporter.

Alarms

Overview

Important: The functions and features included in the Classic UI's **Alarms** tab have been reworked and optimized in more recent releases of Plixer Scrutinizer. They can now be accessed by navigating to the **Monitor** tab of the new UI. To learn more about upgrading to the latest version of Plixer Scrutinizer, see the [Updates and upgrades section](#) of this documentation.

Bulletin boards

As messages come in, they are processed against the list of policies in the policy manager. If the message violates a policy, it can be saved to the history table and may also 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. There are 4 primary menus in the Alarms tab:

- **Views menu** provides options to view some of the more popular reports available in the Alarms tab.
- **Configuration menu:** provides access to the utilities responsible for most of the functionality in the Alarms tab.
- **Reports Menu** provides reports to determine how well the algorithms are performing over time and how frequently the policies are being triggered.
- **Gear menu** configures global settings for the Alarms tab.
- **Show X Entries:** Adjust the number of results shown in the Bulletin Board (10, 25, 50, 100, 200, 300 or 400).
- **Refresh This View:** Set the auto refresh interval.
- **Make this view the default for my profile** Every time the user visits the Alarm tab, this view will be the default.
- **Refresh Button** Refresh the Bulletin Board for the most up to date information.
- **IP/DNS** Display IP addresses or DNS (Host Names)

Heat maps

A heat map is 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.

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 the suspicions are arising.

Views menu

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. With Flow Analytics alarm aggregation, one violation may consist of multiple events.
- **Events** - The number of events triggered by the algorithm.
- **TI (Threat Index)** - This is the default sort by table. The threat index is a function of a policies violation count and the policies threat multiplier. The higher the TI, the greater the chance these policy violations are a security threat. $TI = violations * threat\ multiplier$. [Click here to learn more about 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 one violator and an HI of six, a single host was communicating with six 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.

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 introduces a few new columns that were not outlined above:

- **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. [Click on a Violator address](#) to see all of the alarm events generated by that address.
- **Other columns** - described above.

Notification queue

The notification queue lists the last 24 hours of notifications that were sent or that 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 Notification Profile 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.
- **Rate or Threshold:** Once a notification is added, 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.
- **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.
- **First or Each** 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.

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

This view lists the threats detected by Flow Analytics. It includes the policies and the corresponding violations that occurred in the specified time frame. 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** - Located at the bottom of the table, it provides the totals for the three previous columns across all violated policies. Learn more about [editing policies](#).

Configuration menu

- **Alarm Notifications** allow checking off the entries that the Scrutinizer administrator would like to trigger events for. Events are posted as policy violations in the Alarms tab.
- **Alarm Settings** optimize how notifications are triggered depending on the unique environment. [Contact Plexier Technical Support](#) for assistance.
- **Create New Board** enables the user to create or delete new bulletin boards.

To modify the bulletin board that a policy posts to, visit [Admin tab > Definitions > Policy Manager](#) and edit the corresponding policy.

Note: Bulletin boards can have permissions assigned to them. More details regarding the permissions can be read about under [Usergroup Permissions](#).

- **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.

- **Flow Analytics Settings** brings the user to **Admin > Settings > Flow Analytics Settings**.
- **IP Groups** allows the user to exclude IP addresses, entire subnets or ranges of IPs, as well as child groups from violating specific algorithms.
- **Notification Manager** sets up notifications which can be triggered by policy violations.
- **Policy manager** brings the user to **Admin tab > Definitions > Policy Manager** which lists all of the policies that can be triggered by events. Events are passed through the policies and matches occur based on content in the Message, Source Address or 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.
- **Checkbox:** 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 Username that created the policy.
- **Syslog Server** contains the settings for the syslog server configuration.

Reports menu

Since all of the violations are saved into the database, reports can be run on them to determine how they are performing. The product ships with the sample reports outlined below.

Options:

- **Policies Violated:** This report trends the algorithms violated by violations over the last 24 hours.
- **Threats:** This report trends the Policies violated by violations over the last 24 hours. Columns include the number of violators and the number of Destinations per Policy.
- **Threat Index:** This report trends the top hosts with the highest threat index over the last 24 hours.

Bulletin board events

The Bulletin board events view provides detailed information of the selected alarm events and is useful for isolating specific events and/or violators of alarm events. The view is accessible by clicking on a policy in the Bulletin boards by policy view or a violator in the Bulletin boards by violator view. Filters can be applied to most columns in this view, and the list of events can also be sorted on those columns. Additional actions are included in the Action menu to use against specific alarm events.

The columns available in this view are:

Action - A dropdown menu of available actions per event is provided in this column. Actions available may include excluding the various ip addresses from the Flow Analytics algorithm, view the raw flows for the alarm, view all alarms for the violator address, and several ip address lookup options (GEO IP, Google, HTTP, etc.)

Checkbox - Check this box to acknowledge specific events, or check the box in the header row to select all events for acknowledgment.

The remaining columns are all both sortable and searchable:

Violator Address - IP address that triggered the alarm event.

Host - IP address that the Violator Address was communicating with to trigger the alarm event.

Users - Displays the user(s) associated with the violator address while the alarm is active.

Alarm Time - The time the alarm occurred, or the time the alarm was first issued in the case of aggregated alarms.

Recent Activity - The most recent time an alarm was observed for an aggregated alarm. This will display “N/A” for a single alarm incident.

Duration - 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 - 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”.

Board Name - The name of the Bulletin Board that this event is posted to.

Message - This column provides the full message text of each alarm event.

Editing policies

The Edit policy interface is used to create a new, or modify an existing, policy. Policies are used to match on events that can be saved to the history table and viewed in the Alarms tab. Algorithms, for example, can create events which trigger a policy.

Note: Some policies are read-only and cannot be edited because they are predefined to support specific algorithms that monitor flows or specific events.

Policy Fields

- **Policy Name:** Name displayed in the Bulletin Board
- **Active:** This is a check box that is used to determine whether or not the Policy should be active.

Filters

- **Message Filter:** The text in the body of the message
- **IP Address Filter:** The host the message came from
- **Alert Level Filter:** Can be a combination of two fields “facility” and “severity”.
 - **Facility 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
- **Exclude IPs:** IP addresses to exclude from this policy
- **Include 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 SQL database uses Henry Spencer's implementation of regular expressions, which is aimed at conformance with POSIX 1003.2. The database 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 some 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, 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 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 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.

- **Threat Multiplier:** Enter the value the Threat Index increases by for each violation.
- **Notifications** allow the user to select an action for a policy. 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.

Creating thresholds and notifications

Thresholds are used to receive notification of:

- potential problems on network devices
- excessive utilization on interfaces
- devices that appear to be down
- violation of algorithms in flow analytics

This guide will demonstrate how to properly set thresholds and set up notifications based on violations.

Setting the global threshold

Scrutinizer relies on the SNMP poller to determine the link speed of an interface. These values are used 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 the link speed(s) received from the SNMP poller, admins may want to increase or decrease the values obtained from polling the device. 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 as needed
3. Click Save

Applying a notification to the global threshold

Once the ifSpeed is set and the global threshold is set, notification can be applied. This notification can be in a number of forms (email, logfile, syslog, snmptrap, script, and auto-acknowledge), and will send an alert when the threshold is breached. To add a notification to the global threshold policy, navigate to:

Admin Tab>Definitions>Alarm Policies

1. Enter 'Interface Threshold Violation' in the search field
2. Click the Search button
3. Then click on the 'Interface Threshold Violation' Policy when it's displayed
4. Next, go to the *New Notification* subtab, use the dropdowns to select and configure a Notification Profile and then click the *Save* button.

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

Notification profiles can be created once and applied to multiple *policies*. 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.

Available notification methods:

- **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
- **ServiceNow - Ticket:** automatically create a ServiceNow ticket with the Event message as the description
- **CEF:** send CEF notification with Event details to a host address

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 **Save** button.

Available variables for message

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

Adding a threshold that sends an email notification

Now that a notification profile has been set up, a report which will trigger an email alert can be configured. Adding a threshold to a report requires that the report first be Saved.

Use the following steps to create a Saved Report.

1. Go to the Status tab to bring up the Top Interfaces view
2. Click on the interface name for the reports available list
3. Select Top Reports > Applications Defined from the report menu. This will launch a report for the last 24 hours.
4. To save this report:
 - a. In the upper left, enter a name in the report text box
 - b. Click the Save icon above the report name
5. Next, in the Saved Report, click the Filters/Details button on the left. Click the Threshold tab in the modal. The threshold will use the parameters already defined in the report (Total vs. Rate, Bits or Bytes, etc.)
6. Enter a threshold for the Total amount of traffic reported, or Per Row, which tells Scrutinizer to look at each line in the report table and match it against the threshold. This is useful for applying thresholds to Users or Applications.
7. After completing the fields in the modal, click the Save Threshold button and another window will open prompting the user to select a Notification Profile.

8. Click the dropdown list that says 'None' and select the profile that was created earlier, then click Save.
9. To create a new Notification profile, click the Manage Notifications button.
10. Notice that the Notification Profile was added to the Threshold.

With the threshold set, any time traffic on the specified interface exceeds the value set, an email alert including the specific violation information will be sent.

For additional details or assistance, contact [Plixer Technical Support](#).

Maps

Overview

Important: The functions and features included in the Classic UI's **Maps** tab have been reworked and optimized in more recent releases of Plixer Scrutinizer. They can now be accessed by navigating to **Monitor** > **Network Maps** in the new UI. To learn more about upgrading to the latest version of Plixer Scrutinizer, see the [Updates and upgrades section](#) of this documentation.

Network maps provide a quick visual of the overall network health. They can be added to dashboards for display on a big screen in the network operations center to help identify issues.

Maps are made up of three major parts:

- [Objects](#)
- [Backgrounds](#)
- [Connections](#)

Types of maps

Plixer Maps are used to completely design a topology by arranging the flow sending exporters and other types of network devices in a desired format. Adding custom background images, custom objects and text boxes is also possible. These maps can reflect exactly how the network is laid out by including an image of the wiring closet as a background and then overlaying the flow exporting devices. Connections that represent utilization between the devices can be added.

The feature allows for multiple maps with links between them. Hierarchies can also be established which allows alerts to roll up to the top map.

Google Maps provide a geographical representation of the network. By adding physical addresses to the objects, Google maps will automatically perform a GPS lookup of longitude and latitude coordinates, then place the devices on the map based on those coordinates.

Google maps come especially handy when multiple network topologies are located within a single city, state or country. This type of map not only allows users to see at a glance what network device is having issues, but also where in the world it is located.

Map settings

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.

Note: Learn how the map configuration process works in just a few minutes by watching this video on YouTube

Groups

Groups are the foundation of all maps. Creating a new group creates a map. Flow sending devices that are not assigned to a map are placed in Ungrouped. 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. More details regarding the permissions can be read about under [Usergroup Permissions](#)

Objects

Objects come in four formats:

- **Devices:** are imported from [CrossCheck](#) or can be manually added. 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 in the 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 Plexier maps only. They cannot be placed on Google maps.

Note: To modify the Google address of an object, select a map the object is in 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. Two icons must be provided: one for up status (green) and a second one for down status (red). Make sure the background of the images are transparent.

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 the alert on the status of the object/device should be sent out.
- **Primary Status:** This determines the background color of icons throughout Scrutinizer. The default primary status of a device 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.

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

Connections

The link status comes in 3 formats:

- **Flow links** are links representing flow capable interfaces.
 - Link colors can be green, yellow, orange or red and are based on settings configured in 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** are connections between objects can be made with existing saved reports. The threshold limits for the link color change are set per saved report connection. The values displayed for a Saved Report connection are based on the inbound value for that report.

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. Another option is to select "Connect with black line" to connect to the **To** Object without using a flow interface for the connection.

- 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](#).

Additional notes:

- **Label** displays the percent utilization or the bits received in the last 5 minutes.
- **Tooltip:** mouse over the Label to display 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.

Creating Plixer maps

When creating a Plixer map, the user is presented with the following options:

- **Settings**
 - **Name:** The name given to the map. It can be changed later.
 - **Pass Status:** If some maps are intended to be submaps, the status can be passed to another ‘Parent’ map that contains an icon representing the lower map. This in effect will cause the icon color status of the Parent to change. The status of an icon is determined by [CrossCheck](#) which considers multiple factors.
 - **Auto-add Devices (RegEx):** This option is used to add similar devices quickly using regular expressions. For example, if a number of IP addresses resolve to host names that all contain the text ‘company.local’, this can be entered here. When Save is clicked, all devices that resolve to a host name containing this text will automatically be added to the map.
 - **Truncate Map Labels on:** Sometimes the icon labels can contain excessive amounts of text. Often times, a portion of the trailing text on each icon can be omitted. Enter the text here that shouldn’t be displayed.

- **Objects**

- **Add/Remove objects:** Use this window to move Available objects from the right side to the Members section on the left. Multiple objects can be selected by holding down the shift or CTRL key. Use the filter on the left to quickly locate objects. The search can be performed by IP address or host name by clicking on the button below the filter.
- **New:** Non-flow sending objects can be added to the maps. IP addresses are optional (E.g. text box).

Form fields for **Object Type > Icon** are:

- **Icon:** Use the arrow keys on the key board to scroll through the different icon options.
- **Label:** This names the object and displays in the maps and groups listings.
- **IP Address:** By default, the optional IP address is polled every 60 seconds.
- **Primary Status:** The Primary Status indicator is the largest colored portion of the icon.
- **Link:** The web site that is launched when the icon is clicked in the map.
- **Additional notes:** Help the user understand what the object represents.

Form fields for **Object Type > Text Box** are:

- **Label:** Name of the object, displays in the maps and groups listings.
- **Shape:** Select Rectangle, Circle, Polygon
- **# of Sides:** (Polygon only) Select from 3 - 10 sides for the polygon shape.
- **Height (px) / Width (px):** (Rectangle only) Define the height and width of the rectangle in pixels.
- **Radius (px): (Circle and Polygon):** Define the size of the shape in pixels.
- **Color:** Click on the box to open the color palette to choose the Text Box color.
- **Type:** Choose from Text or Background text box type.
- **Link:** The web site that is launched when the text box is clicked in the map.
- **Connections**

Connections change color based on the utilization settings found in **Admin tab > Settings > System Preferences** and require that an interface speed was collected from, or defined for, the device. Without an interface speed, the connection will stay blue. The arrow on a connection represents the highest flow direction in the last five minutes.

When a map is in view mode, clicking on a link will launch a report showing the last 5 minutes in the highest utilized direction (I.e. inbound or outbound). Connections using a Saved Report are based on the inbound value for that report.

- **Connections:** Click this button to list all of the configured connections with options to either delete or edit a connection. i
- **Create:** Links between devices and objects can be connected using interfaces, saved filters, or a simple black line.
 1. Select a 'From' device
 2. The Type of connection
 3. Fill out the additional options
 4. Then select the 'To' device or object
 5. Click Save
 6. Continue this process to represent the major connections on the network.

- **Background**

There are multiple options to represent the background of a map:

- Existing Map: select a map image from the dropdown selection list that is provided.
- Set background color: click the color in the square to select from the color pallet.
- Upload: create a custom background by transferring an image file to the Scrutinizer server. You can copy the new images directly to the *home/plixer/scrutinizer/files/map_backgrounds* directory.

Maps with background images autoscale to the size of the image. Very light, grayscale backgrounds are ideal as they allow the status of the icons to be visible. The images can be in .gif, .jpg, or .png format. The image size should be at least 800x600 pixels to allow room for icon positioning. Maps with background images autoscale to the size of the background image.

Important: By default, the maximum file size is 5 MB (5000000). You can adjust the setting as well as disable file uploads via the **Admin>Settings>System Preferences** page. The application will discard values below the minimum file size of 200KB (200000).

Laying out the Plixer map

After a new Plixer map has been created and objects added, the objects will be all clustered in the upper left hand corner. To start arranging the icons, the user must enter Edit Mode. When finished editing the map, the user should return to View Mode. Select a blank area on the map and click the RIGHT mouse button, then in the menu, select “Edit Mode”.

- **Gear Menu:** Use these options to set the refresh rate, to display either the IP Address or hostname on the icons, and to reset the zoom level.
- **Edit Mode:** Right click anywhere in the map and select Edit Mode to enter this mode.

The Edit Mode status is then clearly indicated at the top left of the map. In this mode, the icons can be selected with the mouse and dragged to different areas of the map for custom arrangement. Click the right mouse button and notice that several new options present themselves in the Mapping Menu.

- **Align:** (Applies to a group of selected objects only) Aligns selected objects.
- **Auto Arrange:** Select Auto Arrange to get started with laying out the icons and then drag the icons to a more optimal position.
- **Change Background:** Opens Map modal to Background tab, select background as described above.
- **Create a Connection:** (Available only if right clicking on object) Select Create a Connection to connect two devices. The mouse will have a line connected to it. Click on the destination icon. The same two devices can be connected with multiple links.
- **Dependencies:** Configure Map Dependencies
- **Edit Connections:** Edit existing map connections, or create new from the Map modal.
- **Lasso Objects** (or SHIFT+drag mouse): Used to select multiple objects in the map view. Use the crosshair icon to drag over and select a group of objects.

- **Map Settings:** Opens the Map modal to the Settings tab.
- **Objects:** Opens the Map Modal to the Objects tab.
- **Order:** (Available when an object or a group of objects is selected) Indicates object placement. Options are: Bring to front, Send to back, Raise, and Lower.

Background text objects default to being behind all other object types and connections, but their order can be changed using the Order button.

- **Properties:** (Only available when right clicking on an object.) Opens the Edit Object modal to Properties tab.
- **Remove Object:** (Only available when right clicking on an object.) Removes selected object.
- **Save:** Click Save and then select View Mode when finished editing the map.
- **View Mode:** This selection exits Edit Mode.

When finished editing the map, save and exit **Edit Mode**. The status of the devices will update automatically as configured in the Gear menu.

Creating Google maps

To set up the first Google map, an API key has to be generated. To apply a key, navigate to the **Admin > Maps and Device Groups > Global Settings** and paste it into the ****Google Maps - Browser API Key** box.

Note: Google TLD” defaults to .com and should be changed if the install is located in a country that defaults to a top level domain other than .com. For example, in the U.K. change it to .co.uk

Modifying the Google maps involves launching most of the same options found in a [Plixer map](#). There are a few exceptions such as no RIGHT mouse button menu which is reserved by Google for zooming out of the map.

Click on an icon with the LEFT mouse button to launch the menu with the following options:

- **Device Overview:** Launches the device overview including this information.
 - 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
- **Create a connection** works as outlined in the *Plixer Maps* section.
- **GPS Location:**
 1. Placing a device in a specific location requires entering either a physical address or the GPS coordinates. Simply specifying a city in a country will also work.
 2. After entering an address, click (Resolve GPS) to ensure the address is resolved to the new GPS coordinates.
 3. Click Save.
- **Properties** work as outlined in the *Plixer Maps* section.

Admin

Important: The admin functions and settings discussed in this section can also be accessed from **Admin** menus/views of the new UI. To learn more about upgrading to the latest version of Plixer Scrutinizer, see the *Updates and upgrades section* of this documentation.

Definitions

- **3rd Party Integration:** Create links to 3rd party applications and pass variables in URLs. After enabling 3rd Party Integration links will be available in the Device Explorer on the Maps and Status Tabs.

Warning: Please be aware that Solarwinds includes the User ID and Password in plain text in the URL. Using HTTPS will protect the integrity of the credentials over the network, but they will still be visible in the URL, per process set by Solarwinds.

- **Applications:** This feature is useful for properly labeling in-house applications. Some applications utilize multiple IP addresses and ranges of ports. This utility is used to create a single application name that is made up of multiple IP addresses, numerous ports and protocols.
- **Autonomous Systems:** Display and search Autonomous System Names that are shipped with the software, or imported by the user. Use *import asns* in Interactive scrut_util to import AS Names.
- **Host Names:** Set up 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 related 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.
- **Interface 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 the in and out speeds can be entered to override what was collected with SNMP.
- **IP Groups:** IP Groups are used to group ranges of IP addresses or subnets that belong in a specific group or region (e.g. Marketing, sales, phones, Northeast, etc.). A single IP group can contain multiple ranges and / or subnets. Run a report on an interface to see the IP Group reports.

When adding new IP Groups, at least one rule is required for a valid group to be created. Available IP Group rules are:

- IP address: Enter an IP Address in the text box. To enter multiple IP addresses that are not in a range, click **Add** to add additional IP address rules.
- IP range: Defines a range of IP Addresses. Enter the Start IP address and End IP address in the text boxes.

- IP subnet: Enter the subnet in the IP address text box, and select either a subnet mask or a CIDR from the drop-down lists.
- Wildcard mask: Defines a wildcard mask for IP Addresses. Example: IP Address: 10.0.0.1, Wildcard Mask: 0.255.255.0
- Child group: Include other (child) IP Groups in this parent group. Select a child group from the dropdown selection list of existing IP Groups.
- **Language:** Use this interface to update languages or create new translations.
- **MAC Addresses:** Lists MAC Addresses with labels as collected by the utility. It is scheduled to run nightly.
 - MAC address descriptions are collected from Cisco wireless LAN controllers via SNMP.
 - MAC address descriptions are collected from option templates that contain these two elements: 'stamacaddress' and 'username'.
 - Run the scrut_util *'collect optionssummary'* utility to force immediate collection.
 - Manually enter or edit MAC address information here.
- **Manage Collectors:** Provides details on the servers which are collecting flows for this Scrutinizer install. Multiple collectors will be listed if a distributed solution has been deployed.
 - Delete: This checkbox 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 that are currently sending flows to the collector.
 - First Flow Time: Timestamp when flows first received by the collector.
 - Last Flow Time: Timestamp when the last flows were 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:
 - Action / Down Arrow: Use this menu to make several changes to how the flow exporter is represented in the system.
 - * Edit Additional Notes: Add a few comments about the device that can be seen in the Status and Maps tabs.
 - * Edit Name: Give the device a name if it doesn't resolve to an IP address. If it resolved to a host name, this will overwrite it.
 - * Edit Protocol Exclusions: Used to tell the collector to drop flows on certain ports. This was built because some vendors like Cisco export the same flows twice when VPNs or tunnels have been configured.
 - * Edit SNMP Credential: Define the community string to use when querying the device.
 - * Update SNMP: Poll the device for SNMP details on demand.
 - Checkbox: 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. Note that templates and interfaces from devices that stop sending flows are aged out.
 - Round LED: click to view the [Interface 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.
 - Exporter: 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.
- **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](#). Learn more about [editing policies](#).
- **Protocol Exclusions:** Define protocols to exclude during the collection process per exporter, exporter's interface, or for all exporters and interfaces.

Default protocol exclusions for all devices are:

(any private encryption scheme) (99)
(ENCAP) (98)
(ESP) (50)
(ETHERIP) (97)
(GRE) (47)
(IPIP) (94)

Excluding these protocols prevents possible duplication of flow reporting.

- **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 well known port.
 - Else, display the lower port as the well known port.

Settings

- **Alarm Notifications:** Enable additional system alarms.
- **Alarm Settings:** Modify settings to optimize syslog and SMTP processing.
- **ASA ACL Descriptions:** Enter the username and password used to SSH into ASA firewalls to retrieve ACL descriptions (Appliance only).
- **AWS Configuration:** Set parameters for Amazon Web Services flow streaming configuration here.
- **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.

Note: The default value for the Flow Maximum conversations field is 1,000 and the maximum value is 25,000.

-Auto History Trimming: This option allows for automatic database trimming when available disk space falls below 10% (with a minimum threshold of 10GB). Check the checkbox to activate this option. An alarm will also be generated to send an alert 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).

Read more about topics related to this subject:

- [Data Aggregation](#)
- [System LEDs](#)

Note: 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 10% of available disk space.

- **Email Server:** Necessary for on demand and scheduled emailed reports. Make sure the test is successful.
- **Flow Analytics Configuration:** Used to configure the algorithms and monitor their performance.
- **Flow Analytics Exclusions:** Used to manage the Flow Analytics IP Group and hostname exclusions.
- **Flow Analytics Settings:** Used to modify default settings of Flow Analytics relating to FlowPro Defender, jitter, latency, violations and top algorithms.
- **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. See the System > Licensing page for more information.
- *Mapping Groups:* Add and manage Map Groups.
- *Mapping Objects:* Add and manage Map Objects.
- **Proxy Server:** Setup the server to work with a proxy server.
- **Reporting:** Report settings configuration options.
- **Syslog Server:** The syslog server setting tells Plixer Scrutinizer to forward all internal alarms on to an external syslog server/SIEM. For GDPR compliance, select the **Forward Access Log** option. Enabling this will provide a full accounting of all user actions, reports run, and filters applied in Plixer Scrutinizer. For additional information, visit this article on GDPR compliance.
- **System Preferences:** The list of options are global configuration settings for all of the collectors. The explanation for each feature is to the right of the setting.

Security

- **Auditing Report:** Displays a report of all the administrative actions users have performed within Plixer Scrutinizer.
- **Authentication:** Configure general authentication settings, enable or disable different technologies, permit or deny specified groups or named users the ability to use one or more of the supported authentication methods, and set the order in which methods are attempted.
- **Authentication Tokens:** These tokens can be used to automate Plixer 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, the administrator provides credentials for an LDAP service account with permission to perform searches of the remote LDAP directory.
 - a. This is the account that will be used to search for and authenticate users when they attempt to log in (Administrator DN).
 - b. The searchbase identifies the location in the LDAP tree from where the search begins, typically in an OU below the domain level to reduce the scope of the search and improve search performance. This is a required field.
 - c. The scope of users in that searchbase who are allowed to authenticate can be limited in two ways:
 - By specifying one or more security groups in the LDAP Configuration
 - By specifying individual user account names in Security > Authentication > LDAP
2. To configure LDAP integration to use valid certificates, get a PEM encoded version of the Certificate Authority's Certificate and place it into the `/etc/pki/ca-trust/source/anchors/` directory. Provide the full path to the certificate in the "LDAP Server's CA Certificate File" setting. Set the Certificate Verification to required.

3. A user attempts to log in. The system authenticates as the administrative account provided, then checks a searchbase specified by the Plixer Scrutinizer administrator for any account matching the username provided. Authentication with the sAMAccountName, UserPrincipalName, or uid attribute is supported.
4. If the LDAP server responds with an LDAP_REFERRAL code, Plixer Scrutinizer will check the referred server.
5. If the Plixer Scrutinizer administrator has specified multiple LDAP servers, it will check them all until authentication succeeds or fails.
6. Once the user has successfully authenticated for the first time, Plixer Scrutinizer checks for any security group they're a member of which also exists in Plixer Scrutinizer with the same usergroup name. If it does, they're added to the Plixer Scrutinizer usergroup automatically.

LDAP Configuration Example:

| | |
|---------------------------------|--|
| LDAP Server | Server Name |
| LDAP Port | Server TCP Port |
| Domain | example.plixer.com |
| Administrator Password | ***** |
| Administrator DN | CN=ExampleUserName,OU=OptionalOU,DC=PLIXER,DC=com |
| LDAP Server CA Certificate File | |
| Certificate Verification | None |
| ID Attribute | sAMAccountName |
| Searchbase | OU=Example,DC=PLIXER,DC=com |
| Security Groups Allowed | CN=ExampleGroupName,OU=Securitygroups,OU=Applications,DC=PLIXER,DC=com |
| SSL Protocol | tlsv1_2 |
| Timeout | |

Group syncing

When LDAP is enabled and a local usergroup shares the exact same name with an LDAP security group, Plixer Scrutinizer will automatically keep both groups synced by adding or removing users from the local usergroup as they log in.

Examples:

- If a member of the security group *Analysts* logs in to Plixer Scrutinizer using their LDAP credentials, they will automatically be added to the local *Analysts* usergroup (if they were not a member when they logged in).
- If the user is not a member of the *Analysts* LDAP security group, they will be removed from the local *Analysts* usergroup (if they were a member when they logged in).

Important: This feature requires the names of the local usergroup and the LDAP security group to be an *exact* match, including any capitalization and/or punctuation.

LDAP servers

To provide LDAP server redundancy, simply configure multiple distinct LDAP servers on the **Admin > Security > LDAP Servers** page.

Note: When an LDAP user logs into a Plixer Scrutinizer server configured with multiple LDAP servers, authentication attempts will be made against each server in the order they appear in the LDAP Server list until one is successful, otherwise the user authentication fails.

RADIUS configuration

To configure the RADIUS authentication, navigate to the **Admin > Security > RADIUS Configuration** page and provide the following details:

- **RADIUS Server:** The hostname or IP address of the RADIUS server;
- **RADIUS Timeout:** The connection timeout for RADIUS authentication (in seconds);
- **Shared Secret:** The shared secret for the RADIUS server.

Save the changes and attempt to log in with your RADIUS credentials.

TACACS+ configuration

The TACACS + authentication can be set up via the **Admin > Security > TACACS+ Configuration** page.

- **Pre-shared Key:** the pre-shared key for the TACACS+ server.
- **TACACS+ Port:** the TCP port to use when connecting to the TACACS+ server. The default TACACS+ port is TCP 49.
- **TACACS+ Server:** the hostname or IP address of the TACACS+ server.
- **TACACS+ Timeout:** the connection timeout for TACACS+ authentication (in seconds).

Save the changes and attempt to log in with your TACACS+ credentials.

Single sign-on

Plixer Scrutinizer-Azure ADFS SAML integration

To set up the **Plixer Scrutinizer-Azure ADFS SAML integration**, first create the application in Azure.

1. After logging in as an administrator, navigate to **Azure Active Directory > Enterprise Applications**.
2. Click the **New Application** button.
3. In the **Add an Application** dialog, choose **Non-gallery Application**.
4. Enter “Scrutinizer” or any name you prefer in the form that appears, and click **Add**.
5. Once the application is added, you will be redirected to its Overview page. In the toolbar on the left, click **Single Sign-on**.
6. Another dialog with authentication options will appear. **Disabled** is selected by default. Click **SAML** to continue.
7. A form titled “SAML-based sign-on” will have several sections with an “Edit” button in the upper-right of each.

Basic SAML Configuration

| | |
|------------------------|---|
| Identifier (Entity ID) | https://<scrutinizer_server>/ |
| Reply URL | https://<scrutinizer_server>/fcgi/scrut_fcgi.fcgi?rm=usergroups&action=sso_response |
| Sign on URL | https://<scrutinizer_server>/ |
| Relay State | Leave blank |
| Logout URL | Leave blank |

- **User Attributes and Claims**

- Click the claim for `http://schemas.microsoft.com/ws/2008/06/identity/claims/groups`.
- In the panel that appears, select “Security Groups” for “Which groups associated with the user should be returned in the claim?”
- Change “Source attribute” to “sAMAccountName” (unless your organization uses a different AD naming attribute).

- **SAML Signing Certificate**

- Copy the App Federation Metadata URL value.
- Download the Certificate (Base64) file. This document will assume the filename is “azure.cert”

- **Set up Plexier Scrutinizer**

- Copy the Azure AD Identifier value.

Note: The values and the certificate you copied will be required to complete the Plexier Scrutinizer configuration.

This completes the Azure configuration. You should now assign users or groups to the **Plexier Scrutinizer** application in Azure ADFS so that they can successfully authenticate.

Plexier Scrutinizer configuration

Now that you have the required information from Azure’s configuration, you can set up Plexier Scrutinizer’s authentication. Log into Plexier Scrutinizer as an administrator and follow the steps below.

1. Using your favorite client or command line, copy the azure.cert to the following directory on your Plexier Scrutinizer primary reporter: /home/plexier/scrutinizer/.
2. Navigate to the **Admin > Security > Single Sign-On** page and click **Add Server**.
3. In the modal that appears, enter the following values:

| | |
|--------------------|---|
| Name | Enter any unique identifier you prefer (e.g. “Azure ADFS”) |
| IdP Identifier URL | Enter the “Azure AD Identifier” URL you previously copied |
| Entity ID | Enter in the format of https://<scrutinizer_server>/ |
| Assertion URL | Enter in the format of https://<scrutinizer_server>/fcgi/scrut_fcgi.fcgi?rm=usergroups&action=sso_response |
| Audience Value | Enter in the format of https://<scrutinizer_server>/ |
| Name Attribute | Enter http://schemas.xmlsoap.org/ws/2005/05/identity/claims/name |
| Groups Attribute | Optional. Enter http://schemas.microsoft.com/ws/2008/06/identity/claims/groups. It will send usergroup names if the company’s IdP is set up to provide them. |
| IdP Metadata URL | Enter the “App Federation Metadata URL” link you previously copied |
| IdP Metadata XML | Optional. Either the Metadata URL or Metadata XML needs to be entered. Rather than initiating a connection with the IdP to fetch the Metadata URL each time, provide the Metadata XML by pasting it in this field |
| IdP Certificate | Enter “/home/plexier/scrutinizer/azure.cert” |

4. Click **Save** to save the configuration.

A new row will appear in Plexier Scrutinizer’s Single Sign-On Admin view. Log out of your user account. You will notice the URL ends in **/login** – this is the direct access URL to Plexier Scrutinizer’s local and third-party authentication form.

Note: With SSO configured, accessing the root of your server (e.g. <https://scrutinizer.mycompany.com/>) will automatically redirect to Azure ADFS for authentication. If the user or their group has been assigned access to the “Plixer Scrutinizer” application in Azure ADFS, they will be granted access. If the local Plixer Scrutinizer admin account is needed, or if other authentication methods are configured (e.g. LDAP or RADIUS), the login form can be accessed directly at https://<scrutinizer_server>/login.

Plixer Scrutinizer-Okta SAML integration

To enable single sign-on through Okta in Plixer Scrutinizer, you must first create the application in Okta. Launch the Okta Classic UI to perform the steps below. If you see “Developer Console” in a dropdown at the top of your page, click it to switch to Classic UI.

1. Select **Applications** in the navigation bar.
2. Click the **Add Application** button.
3. In the sidebar, pick the green **Create New App** button.
4. In the modal that appears, set **Platform** to **Web**, tick the **SAML 2.0** radio button, and then click **Create**.

Once a new application is created, you will see page 1 of its **General Settings**:

5. Enter **Scrutinizer** for the App name. Click Next.
6. Use the following format for **Single sign on URL**: https://<scrutinizer_server>/fcgi/scrut_fcgi.fcgi?rm=usergroups&action=sso_response
7. Set **Audience URI** to: https://<scrutinizer_server>/
8. Skip the other options and click Next, and then Finish.

You will be redirected to the **Sign On** settings page for the Plixer Scrutinizer application.

9. Locate the section of the page that says: “Identity Provider metadata is available if this application supports dynamic configuration.” Enter the link in the following format: https://<okta_server>/app/identifier/sso/saml/metadata
10. Click **View Setup Instructions**.

11. Set the **Identity Provider Single Sign-On URL** to: `https://<okta_server>/app/application_id_and_name/identifier/sso/saml`
12. Use this link for the **Identity Provider Issuer**: `http://www.okta.com/identifier`
13. Click the **Download Certificate** button and save your `okta.cert` file. We will need to copy it to the Plixer Scrutinizer server later.

With the Okta configuration complete, you should now assign users or groups to the **Plixer Scrutinizer** application so that they will be able to successfully authenticate.

Plixer Scrutinizer configuration

Now that you have the required information from Okta's configuration, you can set up SSO authentication in Plixer Scrutinizer. Log into Plixer Scrutinizer as an administrator and follow the steps below.

1. Using your favorite client or command line, copy the `okta.cert` you previously saved to the following directory on your Plixer Scrutinizer primary reporter: `/home/plixer/scrutinizer/`
2. Navigate to the **Admin > Security > Single Sign-On** page and click **Add Server**.
3. In the modal that appears, enter the following values:

| | |
|--------------------|--|
| Name | Enter any unique identifier you prefer (e.g. "Okta") |
| IdP Identifier URL | Enter the "Identity Provider Issuer" URL you previously copied |
| Entity ID | Enter in the format of <code>https://<scrutinizer_server>/</code> |
| Assertion URL | Enter in the format of <code>https://<scrutinizer_server>/fcgi/scrut_fcgi.fcgi?rm=usergroups&action=sso_res</code> |
| Audience Value | Enter in the format of <code>https://<scrutinizer_server>/</code> |
| Name Attribute | Enter "nameid" to use the name attribute configured and passed back by Okta |
| IdP Metadata URL | Enter the "Identity Provider metadata" link you previously copied |
| IdP Metadata XML | Leave this blank |
| IdP Certificate | Enter <code>/home/plixer/scrutinizer/okta.cert</code> |

4. Click Save.

There will be a new row in the Single Sign-On view. Log out of your user account. You will notice the URL ends in “/login”. This is the direct access URL to Plixer Scrutinizer’s local and third-party authentication form.

Note: With SSO configured, accessing the root of your server (e.g. <https://scrutinizer.mycompany.com/>) will automatically redirect to Okta for authentication. If the user or their group has been assigned access to the **Plixer Scrutinizer** application in Okta, they will be granted access. If the local Plixer Scrutinizer admin account is needed, or if other authentication methods are configured (e.g. LDAP or RADIUS), the login form can be accessed directly at “https://<scrutinizer_server>/login”

- **User Groups:** Specifies what a group login account can access. More details regarding the permissions can be read about under [Usergroup Permissions](#).
- **Users:** 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 for a user account are inherited from all the user groups it is a member of.
- **User Account Lockout:** If a user has a specified amount of failed logins within a defined period of time, that user’s account will be set to ‘locked’ status and will require a user with administrative permissions to unlock it.

These settings are defined in **Admin > Settings > System Preferences** and include:

- Failed Login Max: the maximum failed logins allowed within the Failed Login Window time
- Failed Login Window: the number of minutes that the Failed Login Max value is matching against

For example, the following settings will lock a user account out after two failed logins within a 5-minute timespan:

$$\text{Failed Login Max} = 2 \text{ Failed Login Window} = 5$$

To unlock the account, an administrative user needs to go to **Admin > Security > Users**, select the username that is locked out, then click on the **Authentication Method** tab in the Edit User modal, and change the **Authentication Method** from ‘locked’ to the appropriate method.

Managing devices and interfaces

You can make changes to the device and interface settings from the **Admin > Definitions > Manage Exporters** page. It includes the following information and configuration options as viewed from left to right on the screen:

- **Action / Down Arrow:** Use this menu to make several changes to how the flow exporter is represented in the system.
- **Edit Additional Notes:** Add a few comments about the device that can be seen in the Status and Maps tabs.
- **Edit Name:** Give the device a name if it doesn't resolve to an IP address. If it resolved to a host name, this will overwrite it.
- **Edit Protocol Exclusions:** Used to tell the collector to drop flows on certain ports. This was built because some vendors like Cisco export the same flows twice when VPNs or tunnels have been configured.
- **Edit SNMP Credential:** Define the community string to use when querying the device.
- **Update SNMP:** Poll the device for SNMP details on demand.
- **Checkbox:** 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. Note that templates and interfaces from devices that stop sending flows are aged out.
- **Round LED:**
 - 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.
- **Exporter:** 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.
- **Status:**

- **Enabled:** Flows from this exporter will be collected, stored, and available for reporting.
- **Backup:** Flows from this exporter will be collected and stored, but will not be included in reporting from this collector.
- **Disabled:** Flows from this exporter will be ignored by the collector.
- **Unlicensed:** Set by the collector. This exporter exceeds the exporter license count and flows from it will be ignored. Users wanting to disable specific exporters should use 'disabled'.
- **Last Activity:** Timestamp when the last flow was received for this exporter.
- **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.

Interface details

Selected interfaces can be hidden from the reporting GUI. The *SNMP community* string used to communicate with the device can be altered.

At the top, there is a drop down box containing all the flow sending devices. Type in this box to filter. After a device is selected, a drop down box to select the SNMP community string/credential will appear. Next to the community string is a checkbox 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:

- **Action:** The drop-down arrow is a menu providing options for:
 - **Manage Exporters:** Launches the *Manage Exporters* interface.
 - **Settings:** Provides a modal to provide a custom description for the device and allows for custom In and Out speeds on the interface to be entered.
 - **Update SNMP:** Attempts to update the details using the SNMP credentials.
- **Hide:** Check off to remove the interface from appearing in the Status tab.

- **Interface:** this is the SNMP instance of the interface. Click on it to run the default report.
- **Custom Description:** A custom interface name can be entered.
- **ifAlias:** Collected via SNMP.
- **ifName:** Collected via SNMP.
- **ifDescr:** Collected via SNMP.
- **ifSpeed:** Collected via SNMP. Use the next two columns to customize the in/out speeds.
- **Custom (Bits) In:** Specify a custom inbound speed to override the default. This does not do an SNMP set on the device. Enter a 0 in the Custom (Bits) ifSpeed to force the Status tab to display the interface in bits in lieu of % utilization.
- **Custom (Bits) Out:** Specify a custom outbound speed to override the default. This does not do an SNMP set on the device.
- **Metering:** Indicates whether NetFlow is collected INGRESS, EGRESS or BOTH on this interface. To determine which flows are being used when reporting on an interface, run a report and click on the “Filters / Details” button and then click on the Exporter Details tab.

Scrutinizer labels flow exporter interface names using the following logic in this order if it is available:

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

This requires SNMP access to the devices that are exporting flows. 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.

SNMP

If any updates are applied to a router or switch, be sure to go back to the device interface and run Update SNMP in the down arrow menu, or wait for the daily evening update to run.

Important: By default, the flow collector performs SNMP polls on a nightly basis on 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.

Reports

- **Report Designer** is used to create new reports that are not part of the core reporting solution.
- **Report Folders** manages 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** is used for editing, disabling, and deleting scheduled reports.

Report settings

The Reporting page is accessible via **Admin Tab -> Settings**. This page includes system configuration options related to Scrutinizer reporting.

Following is the list of options available:

- **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 include all rows:** Checkbox. If checked, all rows will be included in the csv instead of the Top X selected in the report.
- **Display Others on Top:** Report Graphs can display the 'Other' traffic on top of or below the top 10.
- **Display raw MAC addresses in reports:** Checkbox. When checked, MAC addresses will appear in reports in raw format. When unchecked, it will display the first 3 bytes as the manufacturer name.
- **Limit All Device report results:** Only this many results will be returned if set to a non-zero value when running all device reports.
- **Max Aggregations from Data Source:** This value limits the number of intervals used to run a report. Click [here](#) for more detailed information on this configuration option.
- **Max Report Processes:** Each report run will use this as a maximum number of sub processes. This breaks reports up by time or exporters depending on which will be faster.
- **Max Reports per Email:** The maximum number of saved reports a user is allowed to include in a scheduled email report. Including too many reports in a single email can result in timeouts. The default is 5.
- **Max Reports per Interval:** The maximum number of reports users are able to schedule for the same minute. The default is 5.
- **Push Data Aggregation:** Checkbox. Apply data aggregation when pushing temp tables from collector to reporter. (Only applies to Distributed collector environments.)
- **Re-use temp tables:** Checkbox. With this option turned on, reports will use existing temp tables when possible.
- **Target graph intervals:** The maximum number of intervals allowed in a graph. Default = 300

Report designer

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:

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 a template from the device. After selecting a template, click [Open Raw Flows] to verify that the element is contained in the template.
5. Select an element in the template for the first column.
6. Specify the column name. It is best to try and keep it short. Specify the treatment.
 - **Average:** takes the average of the total (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.

Rate vs. Total

- **Rate:** Trend the data by rate per second.. Total will not be an option in the drop-down box after the report is run.
- **Total:** Trend the data by total per interval. Rate will not be an option in the drop-down box after the report is run.

- **Rate (default) / Total:** Trend the data by rate per second. Total is an option in the drop down box after the report is run.
- **Rate / Total (default):** Trend the data by total per interval. Rate is an option in the drop down box after the report is run.

Stacked or Unstacked

- **Stacked:** Trend the data as a stacked trend. Non Stacked is not an option in the drop-down box after the report is run.
- **Non Stacked:** Trend the data as an unstacked trend. Stacked trend is not an option in the drop-down box after the report is run.
- **Stacked (default) / Non Stacked:** Trend the data as a stacked trend. Non Stacked is an option in the drop-down box after the report is run.
- **Stacked / Non Stacked (default):** Trend the data as an unstacked trend. Stacked trend is an option in the drop-down box after the report is run.

The new report will show up in the run report menu in a category named “Designed Reports” when the template(s) from the device 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 columns can be reordered. Grab a row in the table with the mouse and move it up or down, then release it.

Multi-tenant configuration

The Multi-tenancy module provides the following features:

- Access to specific tabs (e.g. Dashboard, Maps, Status, Alarms, Admin)
- Ability to apply permissions to User Groups per flow exporting Interface or per device
- Set permissions to see dashboards and even the ability to manipulate or copy a dashboard
- Access to administrative functions

The Multi-tenancy module is useful to companies who need to give customers a unique login and restrict what they see. Restrictions can be set on specific devices and or interfaces.

Usergroup permissions

Users are assigned to usergroups. Usergroups are granted permissions. Users inherit permissions from all the usergroups they are a member of. This functionality also serves as the basis for the enterprise focused multi-tenancy functionality.

- **New User Groups:** Is used to create a new usergroup that individual users can be assigned to. Give the group a name and apply a template from another Usergroup that has similar permissions to the new user group. After creating an account, find the new usergroup on the left and click it to modify.

[Click here](#) for a special note regarding Scrutinizer usergroups and LDAP security groups.

- **Administrators:** This is the admin account and cannot be deleted. Users can be assigned to this group and inherit all of its permissions.
- **Guest:** This is the default guest account which cannot be deleted. Users can be assigned to this group and will have limited permissions.

Important: Permissions for an individual user account will be inherited from all usergroups it is a member of. To view all the usergroups a user account is a member of, visit **Admin tab > Security > Users** and click on a user account. Then open the **Group Membership** tab.

Members

Select the user accounts that will need to have access to this usergroup. A user can be a member of multiple usergroups and inherit all applicable permissions.

Features

Permissions control features the usergroup should have access to within Scrutinizer. Permissions can restrict product features entirely for a usergroup or specific features can be accessed based on your usergroup membership.

Features include:

- Which tab the members of the usergroup should be able to see,
- Administrative permissions the usergroup should have access to,
- Advanced features like acknowledging alarms, scheduling reports, adding/deleting users etc.

Clicking the **Configure** link in the **Features** column will provide a click and drag modal to adjust usergroup permissions. Inside that modal, on the left will two radio buttons with **Predefined** and **Advanced** labels. The following section describes the difference between the two modes, as you must chose one or the other per group.

Predefined roles vs advanced features

The features modal allows Usergroups to use predefined roles or manually specifying features. A Usergroup must use either the Predefined Feature sets **or** the Advanced features that can be manually configured.

Important: You cannot configure manual permissions for a predefined set.

- **Advanced** - Manually configure all permissions available. Use Advanced to create custom feature sets.
- **Predefined roles** - Feature sets for common persona's like "ReportUser" or "DashboardAdministrator"

| | |
|----------------------------|--|
| Pre-defined role | Underlying permissions |
| Alarm-sAdmin-is-trator | ackBBEvent alarmSettings almDelete LogalotPrefs NotificationManager Policy-Manager |
| Alarm-sUser | alarmsTab |
| Dash-board-Admin-is-trator | dashboardAdmin |
| Dash-board-User | createDashTabs myViewTab |
| Map-sAdmin-is-trator | mappingGroupConfiguration mappingObjectConfiguration |
| Map-sUser | adminTab allLogalotReports mapsTab reportFilters statusTab |
| Re-portin-gAdmin-is-trator | ApplicationGroups asnames deleteReport HostNames protocolExclusions report-Settings tos vptelaSettings wkp |
| Re-porting-PowerUser | reportFolders ReportDesigner saveReport scheduledReports srCreate |
| Re- | runReport |

5.1. Plixer Scrutinizer web interface

291

| | |
|-----------------|---|
| Re-portin-gUser | |
| Sys-temAdmin- | 3rdPartyIntegration auditing auth Authentication authLdapServers awsSettings changeUserPasswords createUsers CrossCheck DataHistory deleteUsers DeviceDetails EmailNotifications fa mgmt link faExclusions feedbackForm Flow- |

- **Device status** is used to grant 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.
- **Interface statistics** grants permission to see the statistics of an interface.
- **Groups** are used to grant 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**.
- **Saved reports** allows to select the saved reports/ filters that the usergroup will need to have access to run.
- **Dashboard gadgets** selects the gadgets that the usergroup will need to be able to add to dashboards.
- **Third-party links** controls the vendor third-party integrations that the usergroup will be able to integrate with.
- **Bulletin boards** manages the Bulletin boards that the usergroup will need to be able to access in the Alarms tab.

5.2 Data aggregation

Plixer Scrutinizer's *SAF* (Summary and Forensic) data aggregation method is an optimized system of storing flow data that makes use of summary tables to condense collected information without compromising transparency or accuracy.

How SAF works

With SAF, any incoming flow template with the required data elements is aggregated into a new template definition based on a tuple that includes *commonPort*. The resulting “summarized” template will omit all data elements that prevent aggregation (e.g., source and destination transport ports) but still contain all information required for the vast majority of reporting needs.

Hint: The aggregation logic used to create summary tables can be modified to suit different scenarios. Contact [Plixer Technical Support](#) for assistance.

The data elements retained in the summary tables are but not limited to:

- `intervalTime`

- commonPort
- ingressInterface
- egressInterface
- sourceIpAddress
- destinationIpAddress
- octetDeltaCount
- octetDeltaCount_rev
- packetDeltaCount
- packetDeltaCount_rev
- flowDirection
- applicationId
- protocolIdentifier

Once five 1m summary tables are available, the data averages for the top 1000 (default) conversations are rolled up into 5m tables, and the system continues the rollups to create 30m, 2h, and 12h tables.

Hint: If a Collector's disk capacity will support it, the [Flow Maximum Conversations](#) value under **Admin > Settings > Data History** can be increased, which may improve reporting accuracy. Since this results in larger tables and certain Report types taking longer to run, it is recommended to gradually increase the value over several days.

Note: When *Auto History Trimming* (under **Data History** settings) is enabled, 1m and 5m historical tables are trimmed to maintain the configured *Minimum Percent Free Disk Space before Trimming* value. Automatic trimming is also used to retain a similar level of historical data for all configured exporters.

Benefits of SAF aggregation

Because the summary tables created under SAF aggregation are drastically smaller in size than regular full-template tables, they benefit the Plixer Scrutinizer system in the following ways:

- Reduced disk utilization per table
- Increased historical data capacity
- Improved Report render times
- Faster lookups before drilling into forensic data

While only summary data is rolled up into higher interval tables, Plixer Scrutinizer still retains the original forensic data, which is used by a handful of Reports that require data elements not included in the summary tables. At the same time, the system also maintains a separate totals table for in/out byte counts per interface to allow for accurate utilization reporting without relying on SNMP.

Note: Systems that have been upgraded from versions prior to 18.x may still use the legacy data aggregation method that was the default in their original installs. To check, navigate to **Admin > Settings > Data History** and if the *Rollup Type* is not set to **Summary and Forensic**, contact [Plixer Technical Support](#) for assistance with switching.

Notes on collecting sFlow

When collecting sFlow, packet samples and interface counters should both be forwarded to the Collector. Packet samples will be saved to the raw tables, and interface counters will be saved to the totals tables at 1-minute intervals.

Important: Having an sFlow-exporting device (e.g., switch) that sends multiple templates for different flows may result in overreporting, if the flows contain the same or very similar information. Plixer Scrutinizer's frontend will run Reports using data from all templates that match the information. To avoid this, use filters to specify a single template.

5.3 Machine learning

Through the Plixer ML Engine, Plixer Scrutinizer is able to leverage advanced AI, machine learning, and deep learning technologies to provide real-time anomaly detection and reporting.

Note: To learn more about Plixer ML Engine licensing options, contact [Plixer Technical Support](#).

Once set up, the engine enables the following functions in Plixer Scrutinizer:

5.3.1 Anomaly recognition

As it ingests data through Plexier Scrutinizer, the Plexier ML Engine compiles datasets based on the *hosts* and *dimensions* it has been configured to use. These datasets are then used by the engine to build behavior models that encompass all network activity, including applications and communications to/from external hosts, at a given time.

When a sufficient volume of data has been acquired, the Plexier ML Engine is able to use models that represent typical, legitimate activity patterns as a baseline and recognize deviations that may indicate threats and other anomalies. Deviations that exceed the specified thresholds are then reported as Alarms and Events via the Plexier Scrutinizer web interface.

The Plexier ML Engine's detection and reporting functions can be adapted to any type of enterprise network by defining the *inclusions, dimensions, and sensitivity/threshold values* that best suit an organization's environment.

5.3.2 Malware detection

Because irregular behavior by itself is only indicative of a possible threat and may or may not need remediation, the Plexier ML Engine utilizes additional pre-trained ML models to classify the anomalies it observes through Plexier Scrutinizer and report whether the anomaly actually constitutes malicious activity.

Note: The pre-trained models packaged with the Plexier ML Engine are IP-agnostic and allow Plexier Scrutinizer to alert users to potential threats without needing previously known domain or IP-based signatures.

This classification process is divided into four steps:

1. The engine ingests flow data containing anomalous traffic streamed from Plexier Scrutinizer.
2. The data is preprocessed by the Plexier ML Engine into feature vectors that can be used by the pre-trained ML models.
3. The resulting data is used as the input for the different pre-trained ML models.
4. Each ML model outputs a probability score, which represents the likelihood that the anomaly observed constitutes malicious behavior.

Once probability scores have been obtained, Plexier Scrutinizer compares them to a user-configurable threshold to determine whether or not an Alarm should be generated for the host.

Note: The Plexier ML Engine regularly checks for updates that may include newer versions of the pre-trained ML models it uses.

5.3.3 Continuous learning

To combat the growing sophistication of modern threats, the Plixer ML Engine is also equipped with deep learning capabilities that take advantage of the large quantities of flow data collected by Plixer Scrutinizer to identify complex behavioral patterns and enable advanced features, such as link prediction.

The Plixer ML Engine's deep learning-based threat detection processes can be summarized in the following steps:

1. Flow data collected by Plixer Scrutinizer is forwarded to a datastore module for preprocessing.
2. Once preprocessed, the data is forwarded to the engine, which runs it through a multi-layered neural network designed to discover behavioral patterns in the data.
3. The neural network uses the patterns to learn how devices on the network typically interact with each other.
4. After an anomaly has been detected and classified, the system uses link detection to analyze the device's interactions with other devices on the network.
5. If the deviation from what the Plixer ML Engine has learned as typical behavior exceeds a set threshold, the device involved is added to an endpoint monitoring protocol.

Devices that have been flagged for further monitoring will trigger Alarms under Plixer Scrutinizer's Alarm Monitor, allowing security teams to decide whether immediate action is necessary.

ADVANCED SERVICES

This section introduces Plexer Scrutinizer’s advanced functions and includes configuration guides as well as additional background information related to their use.

6.1 Integrations

Plexer Scrutinizer utilizes standards and protocols that facilitate integration with a wide range of networking tools and services.

This section contains guides for configuring integrations with industry-leading third-party products as well as information and instructions for setting up integrations with other networking solutions.

6.1.1 Plexer Replicator

Plexer Replicator integration enables automatic load balancing across multiple Plexer Scrutinizer servers/-Collectors in a *distributed environment*.

Important: To learn more about Plexer Replicator licensing options, contact *[Plexer Technical Support](#)*.

Once enabled, Plexer Replicator generates a single “seed” Profile and additional Profiles for each Plexer Scrutinizer server in the cluster.

Enabling load balancing

To enable Plixer Replicator integration in the default configuration, follow these steps:

1. Navigate to **Admin > Plixer > Plixer Replicator** (or **Admin > Settings > Plixer Replicator** in the Classic UI) and tick the *Enable* checkbox.
2. Fill in the form with the following details for the Plixer Replicator deployment:
 - admin account password
 - Port used for inbound flows
 - Hostname
 - Name for the seed Profile
 - Port used to send flows to Plixer Scrutinizer
3. Verify that the information is correct and click *Save* to save the settings.
4. Start an SSH session with main Plixer Scrutinizer server as the **plixer** user and run the following command:

```
scrut_util --autoreplicate
```

5. Edit the file `/home/plixer/scrutinizer/files/autoreplicate.conf` as described [here](#).
6. Log into Plixer Replicator web interface and add all Exporters whose flows should be load balanced to the seed Profile.
7. Re-run the `scrut_util --autoreplicate` command.

Note: Running `scrut_util autoreplicate` the first time will create the load balancing configuration file, and re-running it will initiate Exporter flow processing.

Hint: To forward all incoming Exporter flows to the Plixer Scrutinizer distributed cluster, add a Policy for `0.0.0.0/0` to the seed Profile.

Once auto-replication has started, an Alarm containing the configuration details will be generated under the Plexier Scrutinizer Alarm Monitor.

Editing the auto-replication configuration file

The *autoreplicate.conf* file created in the */home/plexier/scrutinizer/files* directory can be edited to control how Plexier Replicator autoreplicates Exporter flows across a distributed Plexier Scrutinizer cluster.

Before making changes to the file, take note of the following additional details:

- The `collector_capacities` section of the file should contain an entry for each Plexier Scrutinizer server/Collector in the cluster and must be updated each time a new Collector is added. This will allow the system to automatically create and manage a Policy for each Collector.

Important: The `collector_capacities` variable also controls the maximum number of Exporters allowed. When that limit is exceeded, Exporters will automatically be removed.

- The `exporters` variable sets the maximum number of Exporters allowed to send flows to a single Collector.
- The `flow_rate` variable controls the maximum number of flows/s a Collector is allowed to receive.
- The `seed_profile` refers to the Plexier Replicator Profile that should contain all Exporters sending flows to the distributed Plexier Scrutinizer cluster. Adding an Exporter to this Profile will include its flows in Plexier Replicator's auto-replication/load balancing operations.

Advanced configurations

Plexier Scrutinizer supports advanced configurations, such as running multiple binaries and pooling, for Plexier Replicator integration.

In addition, the *auto-replication configuration file* supports several optional parameters, which will override any related settings configured via the web interface:

- `replicator_host`
- `replicator_pass`
- `replicator_seed_profile`

- replicator_receive_port
- replicator_send_port

Note: To add an AES256-encrypted password in the configuration file, generate the password using the command `--autoreplicate --encrypt <password>` and paste the output as the value for the `replicator_pass` variable.

To learn more about running multiple binaries and pooling or for assistance with setting up advanced configurations, contact [Plixer Technical Support](#).

For FAQs related to Plixer Replicator integration, see the [FAQs section](#) of this documentation.

6.1.2 Plixer Endpoint Analytics

When Plixer Endpoint Analytics integration is enabled, an additional tab becomes available when inspecting individual hosts (e.g., in the **Monitor** > **Hosts** view)

This tab will show the following details for the endpoint:

- MAC address
- Plixer Endpoint Analytics profile
- OS
- Switch port location
- Risk Profile, etc.

Note: To learn more about Plixer Endpoint Analytics and additional licensing options, contact [Plixer Technical Support](#).

Configuration Guide

After setting up a Plixer Endpoint Analytics account, configure integration in Plixer Scrutinizer as follows:

1. Navigate to **Admin** > **Plixer** > **Endpoint Analytics** and tick the *Enable* checkbox.

2. Enter the IP address or hostname to send API requests to.
3. Enter the password to send with API requests.
4. Enter the port to use for sending API requests.
5. Use the dropdown to select the communication protocol for API requests.
6. Enter the username to send with API requests.
7. Click *Save* to save the entered settings.

Important: Plixer Scrutinizer retains date and time data reported by Plixer Endpoint Analytics, which is based on the time zone of the account used for integration.

Troubleshooting

If there are issues with the integration, try the following steps:

- Check Plixer Scrutinizer logs for errors.
- Verify that the correct credentials were entered during configuration.

For additional assistance, contact [Plixer Technical Support](#).

6.1.3 Plixer FlowPro

The integration of Plixer FlowPro with Plixer Scrutinizer involves setting up Plixer FlowPro as a data source within Plixer Scrutinizer to enable the seamless transfer and analysis of network traffic data.

Configuration Guide

After setting up a Plixer FlowPro, configure integration in Plixer Scrutinizer as follows:

1. Navigate to **Admin > Plixer > FlowPro Licensing**.
2. Paste the Plixer FlowPro license key.
3. Click *Save*.

Note: To learn more about Plixer FlowPro licensing options, contact [Plixer Technical Support](#).

When configured and integrated, Plixer Scrutinizer is able to monitor and provide insights to network performance in real time, diagnose issues, and quickly identify the root causes.

The following Plixer FlowPro APM Reports also become available under the **Reports** section:

| | |
|---------------------------------------|---|
| Application Latency | Application latency reporting measures the delay or time an application takes to send a request and receive a response. It is a critical metric for assessing the responsiveness of applications. |
| Application Latency (old) | This report refers to historical data on application latency, providing insights into how latency has changed. Analyzing historical data can help identify trends and potential issues. |
| Host Jitter | Jitter is the variation in the delay of received packets. Host Jitter measures the irregularity in the packet arrival timing at the destination host. It is crucial for understanding network stability and potential performance issues. |
| Host Jitter By SSRC (Dst) | This report breaks down the host jitter by Synchronization Source (SSRC) at the destination. SSRC is a unique identifier assigned to each synchronization source in a multimedia session. |
| Hosts Latency (Dst) | Measures the latency at the destination host; it provides insights into the delay experienced by packets as they reach their destination. |
| Hosts Latency (Src) | Like Hosts Latency (Dst), this report measures latency at the source host. It helps in understanding the delay introduced by the source system. |
| Host to Host Latency | Host to Host Latency measures the overall latency between two hosts, from source to destination. It considers the complete round trip time for data transfer between the specified hosts. |
| Re-transmission By Application | Indicates the number of times an application has to retransmit data due to packet loss or other network issues. High re-transmission rates may suggest network congestion or unreliable connections. |
| Re-transmission Host to Host | Like Re-transmission By Application, this metric focuses on retransmissions between two hosts. |
| Top Applications | This report provides information on the network's most used or resource-intensive applications. Monitoring top applications helps identify bandwidth consumption and potential performance bottlenecks. |

Troubleshooting

If there are issues with the integration, try the following steps:

- Check Plixer Scrutinizer logs for errors.
- Verify that the correct credentials were entered during configuration.

For additional assistance, contact [Plixer Technical Support](#).

6.1.4 Flow log ingestion

Plixer Scrutinizer can be configured to ingest Amazon VPC and Azure NSG flow logs for traffic monitoring and analysis, as well as to enable additional functionality specific to each platform.

VPC and/or NSG flow log ingestion is enabled by forwarding logs to their respective cloud storage services, which can then be set up as a data source in Plixer Scrutinizer.

Amazon Web Services VPC flow logs

With AWS Virtual Private Cloud (VPC) flow log ingestion enabled, Plixer Scrutinizer is able to report additional insights for network traffic destined for AWS, including top AWS users and applications, as well as traffic load generated by AWS-hosted applications.

The following AWS-specific *report types* become available to run:

| | |
|--------------------------------------|---|
| Action | Aggregation based on the <i>Action</i> (ACCEPT, REJECT, or DROP) associated with the traffic |
| Action with Interface | Aggregation based on the action applied and the interface associated with the flow |
| Action with Interface and Dst | Aggregation based on the action applied, the associated interface, and the traffic's destination |
| Action with Interface and Src | Aggregation based on the action applied, the associated interface, and the traffic's source |
| Availability Zones | Aggregation based on the AWS Availability Zone associated with the traffic |
| Dst Service | Aggregation based on the AWS service the traffic was destined for |
| Interface | Aggregation based on the source or destination interface associated with the traffic |
| Pair Interface | Aggregation based on the source and destination interfaces associated with the traffic |
| Pair Interface Action | Aggregation based on the <i>Action</i> applied and the source and destination interfaces of the traffic |
| Src Service | Aggregation based on the AWS service the traffic originates from |
| Src Service-Dst Service | Aggregation based on AWS services the traffic originated from and was destined for |
| Traffic Path | Aggregation based on the traffic path used by egress traffic to reach its destination |
| VPCs | Aggregation based on the VPC ID associated with the traffic |

Setting up S3 storage

Before setting up AWS VPC flow log ingestion in Plexier Scrutinizer, the Amazon S3 storage bucket(s) that will be used should be configured as follows:

- Set the VPC(s) to be monitored to send flow logs to the bucket. The flow log format *must include* the following fields:
 - log-status
 - vpc-id
 - interface-id
 - flow-direction
- The bucket should be reserved for exclusive use by Plexier Scrutinizer. If the flow logs need to be archived or used for other purposes, send the flow logs to a separate S3 bucket, and then automate the replication/duplication of those logs to the bucket that will be used by Plexier Scrutinizer.
- Versioning should be disabled.

Note:

- When upgrading from older versions of Plexier Scrutinizer, it may be necessary to delete the existing VPC flow log configuration and create a new one that includes the `interface-id` and `flow-direction` fields.
- When creating a VPC flow log, leaving the *Maximum Aggregation Interval* setting at the default 10 minutes will minimize processing load on the Plexier Scrutinizer collector at the cost of longer update times and data spikes. Setting the maximum aggregation interval to 1 minute will result in more granular data but also increase resource utilization.
- After an S3 bucket is first configured for ingestion, Plexier Scrutinizer purges all older flow logs from the bucket before starting to collect and delete the most recent 15 minutes of logs as normal. If any historical data needs to be retained, it should be copied off the bucket before the integration is configured. Manually clearing the bucket of any log data older than 15 minutes will also allow Plexier Scrutinizer to become current more quickly.

Configuring AWS VPC flow log ingestion

To add an S3 bucket as a flow log ingestion source in Plexier Scrutinizer, follow these steps:

1. Navigate to **Admin > Integrations > Flow Log Ingestion** in the web interface.
2. Click the **+** button and select *AWS VPC FlowLogs* in the tray.
3. In the secondary tray, configure the bucket for flow log ingestion as follows:
 - Enter a name identify the bucket/source by.
 - Select the Plixer Scrutinizer server to use as the log downloader for the bucket.
 - Select the Plixer Scrutinizer server to use as the collector for the bucket.
 - Enter the name of the bucket.
 - Select the AWS region where the bucket is hosted from the dropdown.
 - Enter the credentials to use to access the bucket (AWS access key ID and secret access key)
4. Click the **Save** button to add the bucket with the current settings.

Once added, the bucket will be listed in the main Admin > Integrations > Flow Log Ingestion view under the configured name. An exporter associated with the VPC will also be added to the device lists for Plixer Scrutinizer's various functions (*Flow Analytics, network maps, reports*, etc.).

Note:

- The credentials entered must grant full access to the bucket.
- To verify that an AWS VPC flow log source has been successfully added, look for an exporter labeled **vpc-** in the **Explore > Exporters > By Exporters** view or the **Admin > Resources > Manage Exporters** page (after ~1 hour).
- Flow log ingestion processes are divided between the *log downloader* (downloads the flow logs from the bucket) and the *flow collector* (collects and processes the downloaded logs). A different Plixer Scrutinizer server can be used for each role. Multiple downloaders and collectors can also be assigned to a single bucket.
- In distributed environments, avoid using the primary reporter as a log downloader or flow collector.

For assistance with any issues, consult the [troubleshooting guide](#) or contact [Plixer Technical Support](#).

Enabling role-based IAM for AWS deployments

Role-based IAM can be enabled for Plixer Scrutinizer AMI instances by ticking the checkbox in the configuration tray. The role assigned to the EC2 instance should be provisioned with the following permissions:

```
{ "Version": "2012-10-17",
  "Statement": \[
    { "Sid": "VisualEditor0",
      "Effect": "Allow",
      "Action": \[ "s3:GetObject", "s3:DeleteObject" \],
      "Resource": \[ "arn:aws:s3:::<S3BUCKET>/\*" \]
    },
    { "Sid": "VisualEditor1",
      "Effect": "Allow",
      "Action": "s3:\*",
      "Resource": "arn:aws:s3:::<S3_BUCKET_NAME>"
    }
  \]
}
```

Note: Role based authentication is only available when all Log Downloaders are hosted in AWS.

Importing AWS entity descriptions

To allow description reporting and filtering by AWS entity identifiers (`interface-id`, `vpc-id`, etc.) directly in the Plixer Scrutinizer UI, follow these steps:

1. Provision the user or IAM role with the following additional permissions:

```
ec2:DescribeInstances
ec2:DescribeSubnets
ec2:DescribeVpcs
ec2:DescribeNetworkInterfaces
```

2. Start an SSH session with the Plixer Scrutinizer server (or the primary Reporter in distributed deployments), and run the following command via the **scrut_util** interactive CLI:

```
SCRUTINIZER> awssync
AWS entities synced!
```

Once entity descriptions have been synced, AWS entity identifiers will automatically be replaced with their descriptions whenever an AWS-specific Report is run. The *awssync* task will automatically be run every hour thereafter.

Microsoft Azure NSG flow logs

With Azure NSG flow log ingestion enabled, Plixer Scrutinizer can monitor and run reports on IP traffic traversing an NSG.

Once NSG flow data is being received, the following additional *report types* can be run:

| | |
|-----------------------------|---|
| Flow Decisions | Aggregation based on decision (<i>accept</i> or <i>deny</i>) applied to traffic via NSG rules |
| Flow Decisions Count | Flow count aggregation for each NSG traffic decision |
| Flow States | Aggregation based on distinct states reported for individual network flows |
| Flow States Count | Flow count aggregation for each network flow state |
| NSG All Details | Aggregation based on full range of flow details, including the NSG rule and application associated with the traffic |
| Resource IDs | Aggregation based on NSG resource IDs |

Setting up Azure blob storage

Before setting up NSG flow log ingestion in Plixer Scrutinizer, the Azure Storage blob container(s) that will be used should be configured as follows:

- Set the NSG(s) to be monitored to send flow logs to the container. Both version 1 and version 2 flow logs are supported, but the latter format is recommended to enable volume-based reports.
- The container should be reserved for exclusive use by Plixer Scrutinizer. If the flow logs need to be archived or used for other purposes, send the flow logs to a separate blob container, and then automate the replication/duplication of those logs to the container that will be used by Plixer Scrutinizer.
- Versioning should be disabled.

Note: Once a blob container is configured for NSG flow log ingestion, Plixer Scrutinizer will routinely collect the most recent 15 minutes of logs and delete all inactive log files (i.e., not updated in the past ~1 hour). If any historical data needs to be retained, it should be copied off the container before the integration is configured. Manually clearing the container of inactive log files will also allow Plixer Scrutinizer to become current more quickly.

Configuring NSG flow log ingestion

To add an Azure Storage blob container as a flow log ingestion source in Plixer Scrutinizer, follow these steps:

1. Navigate to **Admin > Integrations > Flow Log Ingestion** in the web interface.
2. Click the **+** button, and then select *Azure NSG FlowLogs* in the tray.
3. In the secondary tray, configure the container for flow log ingestion as follows:
 - Enter a name to identify the bucket/source by.
 - Enter the container name (in most cases, `insights-logs-networksecuritygroupflowevent`).
 - Select the collector to assign to the container from the dropdown.
 - Enter the credentials to use to access the container (storage account name and key).
 - Enter the service URL for the container (in most cases, formatted as `https://STORAGE-ACCOUNT-NAME.blob.core.windows.net/`).
4. Click the **Save** button to add the container with the current settings.

Once added, the container will be listed in the main Admin > Integrations > Flow Log Ingestion view under the configured name. An exporter associated with the NSG will also be added to the device lists for Plixer Scrutinizer's various functions (*Flow Analytics*, *network maps*, *reports*, etc.).

Note:

- To verify that an Azure NSG flow log source has been successfully added, look for an exporter whose hostname matches the NSG in the **Explore > Exporters > By Exporters** view or the **Admin > Resources > Manage Exporters** page (after ~1 hour).
 - In distributed environments, avoid using the primary reporter as the collector for the container.
-

For assistance with any issues, consult the [troubleshooting guide](#) or contact [Plixer Technical Support](#).

Troubleshooting

MFSNs and a buildup of log files in an S3 bucket or Azure blob container are indications that the rate of flow and/or log generation exceeds the capacity of the Collector assigned to the flow log source.

The following are potential solutions for an overloaded Collector:

- If the Collector is a VM, allocate additional resources to it.
- If the Collector is ingesting flow logs from only one source (bucket or container), distribute the logs across multiple sources, which can then be assigned to different Collectors.
- If the Collector is ingesting flow logs from multiple sources, reassign sources across multiple Collectors.
- If the Collector license has a flow rate limit, the license may need to be upgraded.

Hint: In distributed deployments, it is recommended to start with a 1:1 pairing of sources and Collectors.

If a VPC or NSG is not listed in the **Admin > Resources > Manage Exporters** view:

- Navigate to **Admin > Integrations**, open the configuration tray for the Collector assigned to and use the *Test* button to verify that the correct details were entered.
- Verify that flow logs are correctly being sent to the bucket or container.
- Check the Collector log file in `/home/plixer/scrutinizer/files/logs/` for errors.
- Check `awss3_log.json` (AWS) or `aznsg_log.json` (Azure) for possible source-side issues.

Note: The **Manage Exporters** view also displays Exporters that have been disabled. Because each VPC or NSG counts as an Exporter, one or sources may be disabled automatically (in last-in/first-out order) if the Exporter count limit of the current license is reached.

6.1.5 Third-party

These guides cover the set-up procedures for Plixer Scrutinizer's built-in third-party integrations, including any additional details related to their configuration and use.

EndaceProbe

With EndaceProbe integration enabled, Plixer Scrutinizer can use specific flow data to generate Endace-Probe-based Reports that automatically download the relevant packets.

To set up EndaceProbe integration, follow these steps:

1. Launch the `scrut_util` prompt by running:

```
/home/plixer/scrutinizer/bin/scrut_util
```

2. At the `SCRUTINIZER>` prompt, use the following commands to configure the probes:

- **Adding an EndaceProbe:**

```
SCRUTINIZER> endace add
```

- **Removing an EndaceProbe:**

```
SCRUTINIZER> endace remove
```

- **Updating an EndaceProbe:**

```
SCRUTINIZER> endace update <host_ip> <port> <endace_user> <endace_  
↩pass>
```

Cisco FireSIGHT eStreamer

Plixer Scrutinizer can be configured to receive flows from a Cisco FireSIGHT system via its Event Streamer (eStreamer) service.

After this integration is enabled, the following Reports will be available in Plixer Scrutinizer:

- 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

Important: The minimum supported eStreamer version is 5.4.

Registering Plixer Scrutinizer with FireSIGHT

Before setting up the integration in Plixer Scrutinizer, the server/Collector must be registered under the FireSIGHT Defense Center:

1. Log into the FireSIGHT Defense Center.

For Firepower v5.4: Navigate to **System > Local > Registration**

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

2. Enable all eStreamer Events and click the Save button.
3. Click on the Create Client (+) button and enter the IP address of the Plixer Scrutinizer Collector.
4. [OPTIONAL] Enter a password.
5. Locate the Plixer Scrutinizer client in the list and click the *Download* button to download the client certificate.
6. Upload the client certificate to the `/home/plixer/scrutinizer/files/` directory on the Plixer Scrutinizer appliance.

Configuring Plixer Scrutinizer as an eStreamer client

After the Plixer Scrutinizer Collector has been registered, it will need to be configured to start receiving FireSIGHT flows:

1. Start an SSH session with the Plixer Scrutinizer Collector.
2. Edit the the `/home/plixer/scrutinizer/files/firesight.ini` file to reflect your Plixer Scrutinizer Collector and FireSIGHT configuration:
 - `CollectorIp` - Plixer Scrutinizer Collector IP address
 - `CollectorPort` - Plixer Scrutinizer receiving port for FireSIGHT flows
 - `fdi_templates` - Path where export templates are defined (default: `/home/plixer/scrutinizer/files/fdi_templates/firesight.fdit`)
 - `host` - FireSIGHT server address
 - `port` - FireSIGHT server outbound port
 - `pkcs12_file` - Location of the FireSIGHT eStreamer client certificate (default: `/home/plixer/scrutinizer/files/<Plixer_Scrutinizer_IP>.pkcs12`)
 - `pkcs12_password` - Password entered during registration process; leave blank if no password was set
 - `fs_bind_addr` - eStreamer client address (Collector IP address)
 - `export_to` - Collector name set at the beginning of the file

Note: Editing the provided `firesight.ini` file is recommended, but a new file can also be created in the specified directory. The Plixer Scrutinizer eStreamer client configuration will automatically be updated whenever the file is modified.

Important: Multiple Collectors and FireSIGHT servers with unique names can be set up within the same `firesight.ini` file. A Collector can be configured to receive flows from more than one source and a FireSight server can send flows to more than one destination.

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 FireSIGHT certificate was updated.
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 (Plixer 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.

Important: There can be more than one collector and/or firesight, but they must have different names. A single collector can receive flows from multiple firesights. A firesight exporter can send flows to multiple collectors.

10. In the /home/plixer/scrutinizer/env/local_env file, change the value for export PLIXER_NO_FIRESEER=1 to 0.
11. Restart the Collector using the command:

```
service plixer_flow_collector restart
```

After the restart, Plixer Scrutinizer should start receiving FireSIGHT flows within 1 minute. For assistance with the configuration process or troubleshooting help, contact [Plixer Technical Support](#).

PRTG

When PRTG integration is enabled, users can view PRTG-based device information when inspecting Exporters in the Plixer Scrutinizer web interface.

To set up PRTG integration in Plixer Scrutinizer, navigate to **Admin > Definitions > 3rd Party Integrations** and follow these steps:

1. Select **PRTG** from the dropdown and untick the *Disabled* checkbox.

2. Fill in the additional fields:

- **Protocol** - Protocol used by the PRTG server
- **Server IP** - PRTG server address
- **Port** - Port used by the PRTG server
- **User** - Username to be used to log in to the PRTG server
- **Password** - Password to be used to log in to the PRTG server

Important: Default values assume the PRTG server is running on HTTPS. If necessary, modify these values to match what is configured under **PRTG Administration Tool > Web Server** on the PRTG server.

3. Click the *Save* button to save the configuration.

Once configured, the option to view PRTG details will be available from the **Integrations** menu when inspecting Exporters.

Important: In the Plixer Scrutinizer Classic UI, PRTG details can be viewed from the Exporter trees under the **Status** tab.

SD-WAN solutions

Plixer Scrutinizer comes with built-in integrations for several leading SD-WAN providers/solutions.

Silver Peak

Plixer Scrutinizer can act as a collector for Silver Peak flow data.

This data can then be used in any combination to generate custom reports using the Plixer Scrutinizer Report Designer tool.

VeloCloud

When enabled, VeloCloud integration in Plixer Scrutinizer makes the following VeloCloud-data-based Reports from the web interface:

- Application Flow Path
- Application Link Policy
- Application Policies
- Application Priority
- Application Route Type
- Application Traffic Type
- Conv Dst Edge
- Dst Edge
- Flow Path
- Interface Jitter
- Interface Latency
- Interface Metrics
- Interface Packet Loss
- Link Utilization
- Packet Loss Conv
- Packet Loss Edge
- Remediation Events
- Traffic Type

Viptela

When enabled, Viptela integration in Plixer Scrutinizer makes the following Reports available when the vManage Exporter is selected when running a Report:

- Carrier Performance
- Transport Performance
- Tunnel Performance
- Application Performance
- Status All Components
- vEdge Health
- SLA Events
- Policies Added
- Policies Removed

Setting up Viptela integration

Viptela integration is enabled and configured via the Plixer Scrutinizer web interface.

1. Navigate to **Admin > Integrations > Viptela Settings**.
2. Tick the checkbox to enable the Viptela integration and fill in the fields with the following information:
 - Viptela vManage NMS IP address or hostname
 - Maximum number of concurrent Viptela API requests that can be processed (default: 10)
 - Maximum number of records that should be returned by each Viptela API request (default: 1000)
 - Password of the user account to be used to connect with Viptela
 - Port number to be used by the Viptela vManage NMS to communicate with Plixer Scrutinizer (default: 8443)
 - Protocol to use for communications between Plixer Scrutinizer and Viptela (default: HTTPS)
 - Username of the user account to be used to connect with Viptela
3. Click the *Save* button to save the configuration.

Important: The user account configured to connect to the Viptela API must have full read access.

If Viptela integration has been correctly configured, the new Reports can be run from the **Reports > Run Report** page of the web interface.

Additional tips

If the configured settings are not working, try the following troubleshooting steps:

- Check the Plexier Scrutinizer Collector log for errors.
- Verify the credentials you entered in Plexier Scrutinizer are correct.
- Use the **Test** button on the **Viptela Settings** page to confirm that Plexier Scrutinizer user can access the Viptela SD-WAN API.

ServiceNow (bi-directional)

Bi-directional ServiceNow integration streamlines troubleshooting ticket creation and management by linking incident reports directly to the relevant data in Plexier Scrutinizer.

When a Collection is flagged for ticketing, ServiceNow generates an incident that links back to more detailed views in Plexier Scrutinizer. Alarm Policies can also be configured to send notifications with optional JSON parameters for automatic incident generation.

Important: ServiceNow integration requires additional licensing to enable. Contact [Plexier Technical Support](#) to learn more.

Configuring ServiceNow integration

To configure a ServiceNow instance to Plexier Scrutinizer, follow these steps:

1. Navigate to **Admin > Integrations > ServiceNow**
2. Click the *Add* button and enter the following details for the ServiceNow instance to be added:
 - Unique name for the instance (used only within Plexier Scrutinizer)
 - Instance URL
 - Username to be used to connect to the ServiceNow instance
 - Password associated with the username

Important: The ServiceNow user registered in Plixer Scrutinizer must be assigned the `sn_incident_write` role.

3. Verify that the details entered are correct and click *Save* to save the ServiceNow instance.

Hint: The *Test* button can be used to confirm that the ServiceNow instance has been correctly configured.

Once ServiceNow integration has been enabled, the ServiceNow instance name will be added as an option when managing Collections or configuring Notification Profiles for Alarm Policies.

SolarWinds

When SolarWinds integration is enabled, users can view SolarWinds-based device statistics when inspecting Exporters in the Plixer Scrutinizer web interface.

To set up SolarWinds integration in Plixer Scrutinizer, navigate to **Admin > Definitions > 3rd Party Integrations** and follow these steps:

1. Select **SolarWinds** from the dropdown and untick the *Disabled* checkbox.
2. Fill in the additional fields:
 - **Server IP** - SolarWinds server IP address
 - **User** - Username to be used to log in to the SolarWinds server
 - **Password** - Password associated with the entered SolarWinds login
 - **API Port** - API port users by the SolarWinds server (default: 17778)
3. Click the *Save* button to save the configuration.

Once configured, the option to view SolarWinds details will be available from the **Integrations** menu when inspecting Exporters.

Note: In the Plixer Scrutinizer Classic UI, PRTG details can be viewed from the Exporter trees under the **Status** tab.

Important: When accessing SolarWinds details from the Plixer Scrutinizer web interface, the username and password are included in the URL used to open the page. The use of HTTPS will protect the integrity of the credentials over the network, but they will still be visible as outlined in this SolarWinds support article.

Plixer Scrutinizer integration in SolarWinds

The SolarWinds Network Performance Monitor supports pivoting from the **Node Details** page to a Report in Plixer Scrutinizer.

Note: This integration was configured for Solarwinds NPM 12.2 and is not guaranteed to work on older installations.

To setup Plixer Scrutinizer integration in SolarWinds, follow these steps:

1. Navigate to **Settings > All Settings**. Under **Node & Group Management**, select *Manage Custom Properties*.
2. Click *Add Custom Property* and select **Nodes** from the dropdown list.
3. Fill out the name and description fields for the property (e.g., Plixer Scrutinizer) and then click *Next*.
4. Assign the property to at least one existing Node by clicking the *Select Nodes* button and using the *Add* arrow to add exporters.
5. Fill in the value box for the added node(s) with the following code:

```
<a href="http://SCRUTINIZER_IP_ADDRESS/search.html?el=${IP_Address}&
↪reportType=conversations">
    </img>
</a>
```

Hint: The above code block opens the *Conversations WKP* Report type as the default, but this can be modified by replacing `conversations` with a different Report name API as the value for `reportType`.

6. When done, click *Submit* to save the configuration.

If the integration has been correctly configured, a custom property widget for all selected nodes/exporters will be added to the **Node Details** page. To run the default Report, click on the Plexier Scrutinizer in the widget.

STIX-TAXII

STIX-TAXII integration allows Plexier Scrutinizer to import comprehensive and up-to-date threat intelligence in the industry-standard Structured Threat Information eXchange (STIX) format via the Trusted Automated eXchange of Indicator Information (TAXII) protocol from external systems and organizations. This greatly enhances Plexier Scrutinizer's already robust IP detection capabilities.

Important: STIX-TAXII integration requires additional licensing to enable. Contact [Plexier Technical Support](#) to learn more.

Importing STIX files via CLI

To have Plexier Scrutinizer automatically import IP/domain watchlists, download the files in STIX format (v1 or v2) and copy them to the `/home/plexier/scrutinizer/files/threats` directory on the appliance. The name of the file will also be used as the category.

Important: Domain watchlists are currently only used in AI-based threat detection algorithms and need not be imported for deployments that do not include the Plexier ML Engine.

Note: Plexier Scrutinizer supports `.stix`, `.stix1`, and `.stixv1` extensions for v1 (XML) and `.stix2` and `.stixv2` extensions for v2 (JSON).

Configuring STIX-TAXII feeds

To configure a new STIX-TAXII feed the Plexier Scrutinizer web interface, follow these steps:

1. Navigate to **Admin > Integrations > STIX-TAXII** and click the *Add* button to create a new feed.
2. Fill in the following fields:
 - Feed name
 - API Root (**not** the Discovery URL)
 - Collection ID
 - Login credentials for the feed
3. Click the *Save* button to save the settings.
4. Use the *Test* button to verify that Plixer Scrutinizer can access the feed with the configured settings.

After the feed has successfully been added, Plixer Scrutinizer will attempt to pull the lists from the TAXII server every time the host reputation list download service runs.

Once imported, STIX-TAXII threat intelligence will be added to Plixer Scrutinizer's (IP only) and the Plixer ML Engine's (IP and domain) reputation algorithms for Alarm and Event reporting under their respective Alarm Policies.

Additional tips

- Import IP watchlists only. All other indicators will be ignored but can cause the import of IP indicators to fail.
- Don't attempt to import IP watchlists that use complex boolean logic to trigger matches.
- The feature will ingest only independent IP indicators. It will ignore more complex ones.

Note: A complicated indicator included with more basic ones will not prevent them from being imported.

Username reporting

Plixer Scrutinizer supports username reporting via Microsoft Active Directory (AD) or Cisco Identity Services Engine (ISE).

To enable and configure username reporting integration, follow the corresponding guide below:

Microsoft Active Directory over LDAP

When Microsoft Active Directory (AD) username reporting is enabled, Plixer Scrutinizer is able to retrieve domains, datasources, and first/last seen details for AD users and report the information in various web interface views and functions.

This integration relies on the Plixer AD Users utility to retrieve username data and forward it to Plixer Scrutinizer as IPFIX flows.

The Plixer AD Users utility reads a Windows event log file, continually parses authentication events, and sends event data to an IPFIX collector (Plixer Scrutinizer) for viewing in the **Explore > Entities > Usernames** table in the web UI. If the AD Users service is stopped, the last sent event record ID is saved to *last_recordID.txt*. If this file exists, only events with record IDs greater than the number in the file will be sent to Plixer Scrutinizer. This feature helps avoid duplicate events being sent to the collector or a lapse in the authentication events processed should the program restart.

Plixer AD Users 2.0.0

Plixer AD Users 2.0.0 allows integration with the Plixer ML engine 19.4.0. When configured, both Active Directory authentication events and sign-in logs from designated Azure storage containers are processed by the ML engine. The result is model generation used to detect anomalies, send alerts, and generate reports for usernames or email (Azure) login data.

Note: Azure login data typically categorizes users by email.

For Azure sign-in logs, there is a configuration setting to pull logs older than the specified number of minutes to account for any missed events back to a certain time. The Plixer ML engine will remove duplicates of both AD and Azure events for model generation.

Configuring the servers

User permissions

By default, the Plixer AD Users installer configures the program to run using a local system account and this is the recommended configuration. However, the program can also be configured to run as a different user.

If not using a local system account, the user who is configured to run the Plixer AD Users service needs to:

- Have administrator privileges
- Have permissions to query domain controller event logs by being added to the event log readers built-in group
- Have *Log on as a service* rights if running as a service

Domain controller audit policies

To allow authentication events to be collected, logon/logoff audit policies on the domain controller must be enabled.

To do this, make the following changes to the domain controller's default policies:

1. Expand **Computer Configuration > Policies > Windows Settings > Security Settings > Advanced Audit Policy Configuration > Audit Policies > Logon/Logoff**.
2. Enable **Success and Failure** for *Audit Logoff* and *Audit Logon*.

The advanced audit policies require that another group policy override setting is enabled. To do this, follow these steps:

1. Expand **Computer Configuration > Policies > Windows Settings > Local Policies > Security Options**.
2. Select **Audit: Force audit policy subcategory settings**.
3. Tick **Define this policy setting**, and then tick **Enable**.

Event forwarding

Running Plixer AD Users directly on the Active Directory server does not require any additional configuration, other than ensuring that the [config file](#) points to `Security.evtx`.

To run Plixer AD Users on a separate event collection server joined to the same domain as the Active Directory server/domain controller, follow these steps:

1. On the Active Directory server(s), run the following command from an elevated-permissions command prompt: `C:\> winrm quickconfig`
2. On the event collection server, run the following command from an elevated-permissions command prompt: `C:\> wecutil qc`
3. Establish a subscription by performing the following on the event collection server:
 - As an Administrator, launch **Event Viewer**, and then click **Subscriptions**.
 - In the **Actions** pane, click **Create Subscription**.
 - Enter a subscription name.
 - Select **Computers**, and then enter your Active Directory server(s).
 - Go to **Destination log > Forwarded Events**, and then select **Keep User Account as Machine Account**.
 - Select **Events**, then select **Security for Event logs**, and then enter the following event IDs to include: 4624,4634,4647,6272-6274,6278,6279.

Plixer ML engine

The Azure storage account name and Azure storage account key will need to be entered during the setup script. Apart from this, no further configuration is required on the Plixer ML engine.

If the user did not configure AD Users to send Azure data during the initial setup process of their ML engine (i.e. responded “no” to the example prompt included below), user can re-run `/home/plixer/ml/setup.sh` on the ML engine, which will prompt the user again if they would like to configure Azure for AD Users.

Re-running `setup.sh` on an ML engine that has already been configured and running does not take the user through the entire configuration process again. It only applies any changes or updates made to the ML engine configuration since the last time it was run (e.g. Azure credentials that were not previously entered).

The following is an example of a Plixer ML engine setup script:

```
Successfully upgraded PXI deployment
Will you configure AD Users to send Azure data to your ML engine?_
↪ (yes/no):
yes
Configuring Azure storage account info...
Enter your Azure storage account name:
<Your Azure storage account name>
Enter your Azure storage account key:
<Your Azure storage account key>
secret/azure-backup created
```

Note: Plixer ML engine 19.4.0 only supports integration with one Azure storage account. If storage account credentials were entered previously (i.e. for configuring an Azure backup or previous sign-in log storage), they should be changed accordingly to the storage account containing Azure sign-in data for AD Users.

Existing Azure storage account credentials that were submitted to the ML engine previously (either for AD Users or as an Azure backup) can be changed with the following commands:

```
kubectl -n pxi delete secret azure-backup --ignore-not-found
kubectl -n pxi create secret generic azure-backup --from-
↪ literal=azure.client.secondary.account=<Your storage account name>
↪ --from-literal=azure.client.secondary.key=<Your storage account_
↪ key>
```


Microsoft Azure

To archive both interactive and non-interactive sign-in logs to a storage account, create a storage account, and then set up archiving for the sign-in logs to that storage account.

To set up a storage account, do the following:

1. Sign in to the Microsoft Azure portal.
2. Under **All Services**, select **Storage Accounts**, and then select **Create**.
3. If necessary, configure your storage account if needed. No AD Users-specific configuration is required here.
4. Take note of the storage account name, and then after configuring, select **Create**.
5. Under **Storage Accounts**, select **Access Keys** under **Security + Networking** to generate and save the value of the key for the storage account.

Setting up the Plixer AD Users utility

Once the domain controller has been correctly configured, set up Plixer AD Users on a Windows computer as follows:

1. Download the Plixer AD Users 2.0.0 product package: `ad-users-2.0.0.zip`
 - This archive includes the `ad-users.exe` executable file, the `ad-users-installer.exe` installer, and the `ad-users.yml` config file.
 - The checksum for the zip file can be found here: `ad-users-2.0.0.zip.sha256`
2. Run `ad-users-installer.exe`, and then go through the installation steps.

Important: Make sure that you select **No** to use recommended system account, and to tick **Open config file** to set the collector value.

3. Run a verification test via command prompt with command-line argument `test`. For example, `ad-users.exe test`.

Editing the config files

ad-users.yml

| Name | Re- quired? | Description | Ex- ample Value | De- fault Value |
|------------------------|----------------|--|-----------------------|-----------------------|
| chunk- ing | Yes | Number of authentication log events to collect, and then send at a time. Set to 0 to send each event as it is parsed. | 1000 | 0 |
| flush_wait_ seconds | Yes | Time in seconds to periodically send any events in the buffer. Set to 0 if you want to use chunking value for sending events instead. | 60 | 0 |

Cisco Identity Services Engine (ISE)

When Cisco Identity Services Engine (ISE) username reporting is enabled, Plixer Scrutinizer is able retrieve username lists, search flows for specific usernames, and run additional reports related to Cisco ISE user traffic.

Important: Username reporting integration in Plixer Scrutinizer supports Cisco ISE versions 1.2, 1.3, 1.4, 2.0, 2.1, and 2.3.

Enabling ERS

Before setting up Cisco ISE username reporting in Plixer Scrutinizer, External RESTful Services (ERS) should first be enabled on the ISE appliance as follows:

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

2. Test the configuration using an external host via a **Postman** GET request using the URL: `https://[ISE_server_address]/ise/mnt/Session/AuthList/null/null`

Hint: When creating the GET request using **Postman**, navigate to the server using a browser and agree to use a bad certificate. Leave that window open.

Visit the Cisco website to learn more about enabling ERS for the supported ISE versions.

Configuring steps in Plixer Scrutinizer

1. SSH into the Plixer Scrutinizer server as the `plixer` user and run `/home/plixer/scrutinizer/bin/scrut_util` to launch the *scrut_util interactive CLI*.
2. At the `SCRUTINIZER>` prompt, enter:

```
SCRUTINIZER> ciscoise add [ISE_IP] [ISE_TCP_port] [ISE_user>]
```

This adds a Cisco ISE node from which username data for active sessions can be retrieved. `ISE_IP` and `ISE_TCP_port` refer to the the ISE server's address and TCP port number and `ISE_user` refers to the user previously created on the same server.

3. When prompted, enter the password for the ISE user.

After all configuration steps have been completed, all functions associated with Cisco ISE username reporting will immediately be enabled.

Note: It may take several minutes before usernames are displayed in the web interface.

scrut_util commands for Cisco ISE

Information about other **scrut_util** commands related to Cisco ISE username reporting can be found [here](#).

6.2 Interactive CLI

Plixer Scrutinizer’s interactive **scrut_util** utility provides users with access to various system-level functions through a command line interface (CLI).

These functions include administrative processes, environment configuration, and maintenance routines, as well as third-party integration management.

6.2.1 Launching the utility

To launch **scrut_util** and access the SCRUTINIZER> prompt, start an SSH session with the appliance as the **plixer** user, and then run:

```
/home/plixer/scrutinizer/bin/scrut_util
```

The SCRUTINIZER prompt indicates that the system is ready to accept commands.

Exiting scrut_util

To leave the SCRUTINIZER prompt, enter the **exit** command:

```
SCRUTINIZER> exit
Exiting...
[root@Scrutinizer ~] #
```

Command help

For details about the available **scrut_util** commands, use the **help** command at the SCRUTINIZER prompt:

| | |
|----------------------------|--|
| help | Displays the full scrut_util command list |
| help [command] | Displays information about a specific command e.g., help show |
| help [command] [directive] | Displays information about extended commands e.g., help show groups |

6.2.2 Command list

The following are the available top level commands:

- *aws*
- *check*
- *ciscoise*
- *clean*
- *collect*
- *convert*
- *delete*
- *disable*
- *enable*
- *endace*
- *expire*
- *export*
- *import*
- *moloch*
- *optimize*
- *remove*
- *repair*
- *rotate*
- *services*
- *set*
- *show*
- *snoop*
- *system*
- *unlock*
- *upload*
- *version*

Note: Each top level command may have several extended commands.

aws

| | |
|----------|--|
| Function | Manages AWS flow log integration with Plixer Scrutinizer |
| Syntax | <ul style="list-style-type: none">• <i>aws sync</i> - Synchronizes IDs and descriptions from AWS |

check

| Function | Returns information on the specified resource, setting, or function |
|----------------------|--|
| Syntax | <ul style="list-style-type: none">• <i>check activeif</i> - Checks for active flows• <i>check collectorclass</i> <class> <subsystem> - Logs the Collector's running state• <i>check data_last_written</i> - Checks the activity of collected flow data written to the database• <i>check database</i> [db_name] [db_pass] - Checks the specified database for errors• <i>check dist_info</i> - Displays information on distributed Plixer Scrutinizer servers• <i>check hdtest</i> - Tests the performance of the hard drive• <i>check heartbeat</i> <database api> - Checks heartbeat functions• <i>check history_index</i> - Checks the table activity every minute• <i>check history_index_empty_tables</i> - Lists empty tables• <i>check history_index_orphans</i> - Checks for tables that do not exist• <i>check history_table_orphans</i> - Lists tables without <i>history_index</i> entries• <i>check</i> interfaces [all cisco hauwei sonicwall] [host_ip] - Uses alternative methods to retrieve interface descriptions• <i>check license</i> - Displays license details of the Plixer Scrutinizer server• <i>check machine_id</i> - Displays the current Machine ID of the Plixer Scrutinizer Server• <i>check machine_id_list</i> - Displays the historical, current, and possible Machine IDs of the Plixer Scrutinizer Server• <i>check objects</i> - Verifies that xcheck_hosts have corresponding rows in objects• <i>check password rootdb</i> - Verifies that the database root password matches the password in plixer.ini• <i>check rollcall</i> - Analyzes rollcall and the state of rollups per time bucket |
| 6.2. Interactive CLI | <ul style="list-style-type: none">• <i>check rollups</i> - Lists rollups and their current state• <i>check route</i> [ip] - Checks the device specified to determine if Plixer Scrutinizer can access its routing data |

ciscoise

| | |
|----------|--|
| Function | Manages CiscoISE Node integration with Plixer Scrutinizer |
| Syntax | <ul style="list-style-type: none">• <i>ciscoise add [ise_ip] [ise_tcp_port] [ise_user]</i> - Adds a CiscoISE Node to the queue to acquire user identity on all active sessions• <i>ciscoise check</i> - Tests polling and displays the results• <i>ciscoise kick [ise_id] [mac_address] [user_ip]</i> - Forcibly logs the specified user off the ISE Node and requires re-authentication• <i>ciscoise nodelist</i> - Lists all currently configured CiscoISE nodes• <i>ciscoise poll</i> - Executes a manual poll and displays the results• <i>ciscoise remove [ise_ip]</i> - Removes a CiscoISE Node from Plixer Scrutinizer• <i>ciscoise test</i> - Tests polling and displays the results• <i>ciscoise update [ise_ip] [ise_tcp_port] [ise_user]</i> - Updates current settings for the specified ISE Node |

clean

| |
|---|
| Warning: These commands will purge data from Plixer Scrutinizer and should be used with caution. |
|---|

| | |
|----------|---|
| Function | Executes Plixer Scrutinizer housekeeping tasks outside of their regularly scheduled run times |
| Syntax | <ul style="list-style-type: none">• <i>clean all</i> - Executes all Plixer Scrutinizer housekeeping processes that are configured to run at scheduled times• <i>clean database</i> - Deletes all temporary database entries• <i>clean ifinfo</i> - Deletes all entries in the <i>ifinfo</i> db table that do not have an entry in the <i>activeif</i> db table• <i>clean old_logs</i> - Deletes old log files that are set to the backup status• <i>clean pcap</i> - Deletes all pcap files from the Plixer Scrutinizer server• <i>clean pcap [pcapfile]</i> - Deletes the specified pcap file from the Plixer Scrutinizer server• <i>clean tmp</i> - Deletes all temporary files created by the graphing engine |

collect

| | |
|----------|--|
| Function | Collects data that can be utilized by Plexer Scrutinizer on demand |
| Syntax | <ul style="list-style-type: none">• <i>collect asa_acl</i> - Collects ASA ACL information from Cisco ASA Devices• <i>collect dbsize</i> - Collects database size information• <i>collect optionssummary</i> - Processes flow option data collected by Plexer Scrutinizer• <i>collect pcap [in_sec] [host]</i> - Collects a packet capture on the interfaces of the Plexer Scrutinizer server• <i>collect snmp</i> - Collects SNMP data that is used during Plexer Scrutinizer's operations• <i>collect supportfiles</i> - Collects various log files and server configuration data used by Plexer Technical Support for troubleshooting• <i>collect topology</i> - Polls SNMP-enabled devices (including non-Exporters) to collect data related to network topology• <i>collect useridentity</i> - Processes user identity data collected by Plexer Scrutinizer |

convert

| |
|--|
| Warning: These commands will alter the database tables in Plexer Scrutinizer and should be used with caution. |
|--|

| | |
|----------|---|
| Function | Converts all encrypted information stored in Plexer Scrutinizer to use AES 256 encryption |
| Syntax | <ul style="list-style-type: none">• <i>converttoaes</i> - Converts all encrypted information stored in Plexer Scrutinizer to use AES 256 encryption |

delete

Warning: These commands will purge data from Plexer Scrutinizer and should be used with caution.

| | |
|----------|--|
| Function | Deletes database tables and/or database table entries |
| Syntax | <ul style="list-style-type: none"> • <i>delete history_index_empty_tables</i> - Deletes empty tables • <i>delete history_index_orphans</i> - Deletes tables that do not exist • <i>delete history_table_orphans</i> - Deletes tables without <i>history_index</i> entries • <i>delete orphans</i> - Deletes all known orphan Alarms and Events |

disable

Warning: These commands will alter Plexer Scrutinizer functionality and should be used with caution.

| | |
|----------|--|
| Function | Disables a specific Plexer Scrutinizer function or service |
| Syntax | <ul style="list-style-type: none"> • <i>disable ipv6</i> - Disables IPv6 in <i>sysctl.conf</i> for all interfaces • <i>disable user [username]</i> - Disables an account with <i>scrut_util</i> access on the Plexer Scrutinizer server • <i>disable unresponsive</i> - Disables pinging for unresponsive Exporters • <i>disable hypervtools</i> - Disables Hyper-V Integration Tools for a virtual appliance running on Hyper-V • <i>disable vmwaretools</i> - Disables <i>vmwaretools</i> for a virtual appliance running on VMware |

enable

Warning: These commands will alter Plexer Scrutinizer functionality and should be used with caution.

| Function | Enables a specific Plexer Scrutinizer function or service |
|----------|--|
| Syntax | <ul style="list-style-type: none">• <i>enable dbpool [pool_port]</i> - Enables database connection pooling for PostgreSQL• <i>enable ipv6</i> - Enables <i>ipv6</i> in <i>sysctl.conf</i> for all interfaces• <i>enable perl_support</i> - Installs additional Perl packages to support custom scripting• <i>enable user [username] [security_level]</i> - Creates an account that has access to <i>scrut_util</i> with one of the following security levels:<ul style="list-style-type: none">– 1 - Commands that stop data collection are disabled.– 2 - Commands remove/disable integrations and stop data collection are disabled.– 3 - Only commands to collect information about Plexer Scrutinizer and the operating system are enabled.• <i>enable hypervtools</i> - Enables Hyper-V Integration Tools for a virtual appliance running on Hyper-V• <i>enable vmwaretools</i> - Enables <i>vmwaretools</i> for a virtual appliance running on VMware |

endace

| | |
|----------|---|
| Function | Manages EndaceProbe integration |
| Syntax | <ul style="list-style-type: none"> • <i>endace add [host_ip] [port] [endace_user] [endace_pass]</i> • <i>endace remove [host_ip]</i> • <i>endace update [host_ip] [port] [endace_user] [endace_pass]</i> <p>For additional details on these commands, see this guide</p> |

expire

Warning: These commands will purge data from Plixer Scrutinizer and should be used with caution.

| | |
|----------|---|
| Function | Deletes expired historical data (based on the configured history retention settings) |
| Syntax | <ul style="list-style-type: none"> • <i>expire alarms</i> - Purges expired Alarm history from the threatoverview and fa_transports_violations tables based on <i>I Min Avg</i> flow history retention setting • <i>expire bulletinboards</i> - Purges expired Alarm bulletin board Events • <i>expire dnscache</i> - Purges expired DNS cache data • <i>expire history <trim></i> - Purges expired flow data; if the trim argument is passed, purges older flow data to free up disk space • <i>expire ifinfo</i> - Purges old and outdated interface information • <i>expire ifinactiveflows</i> - Purges inactive interfaces (based on the <i>Inactive Expiration</i> setting) from interface views • <i>expire orphans</i> - Purges expired orphan Events • <i>expire templates</i> - Purges flow template metadata for templates that have not been observed for 30 days |

export

| | |
|----------|---|
| Function | Exports data from Plexier Scrutinizer for external use |
| Syntax | <ul style="list-style-type: none">• <i>export langtemplate [lang_name]</i> - Exports the definition template for the specified language• <i>export peaks_csv [file] [interval] [dir] [date_range] [group_id]</i> - Exports a CSV file listing interfaces and peak values based on the criteria specified |

import

| | |
|----------|--|
| Function | Runs various import commands to bring external sources of data into Plexier Scrutinizer |
| Syntax | <ul style="list-style-type: none">• <i>import aclfile</i> - Imports ACL information from the specified file• <i>import applications [path/file] <reset></i> - Imports application rules from a CSV file• <i>import asns [path/file] [delimiter]</i> - Imports custom autonomous system number (ASN) definitions from a CSV file• <i>import csv_to_gps [csv_file] [group_name/group_id] <create_new> [file_format]</i> - Imports geographic location information of devices from a CSV file and uploads them to an existing Google map• <i>import csv_to_membership [csv_file] [groupype] [file_format]</i> - Imports group definitions from a CSV file• <i>import hostfile</i> - Imports a custom hosts.txt file that contains a list of IP Addresses and hostnames• <i>import ipgroups [path/file] <reset></i> - Imports <i>ipgroup</i> rules from a CSV file |

moloch

| | |
|----------|--|
| Function | Manages integration with Moloch probes |
| Syntax | <ul style="list-style-type: none">• <i>moloch</i> <on off> [<i>moloch_ip</i>] [<i>moloch_port</i>] |

optimize

Warning: These commands will modify Plexier Scrutinizer database tables and should be used with caution.

| | |
|----------|--|
| Function | Runs various optimization processes |
| Syntax | <ul style="list-style-type: none">• <i>optimize common</i> - Optimizes tables that are commonly inserted and deleted• <i>optimize database</i> [<i>db_name</i>] [<i>db_pass</i>] - Optimizes the tables in the database specified |

remove

Warning: These commands will alter Plexier Scrutinizer's functions and should be used with caution.

| | |
|----------|---|
| Function | Removes a configured setting from the system |
| Syntax | <ul style="list-style-type: none">• <i>remove address ipv6</i> - Removes any configured IPv6 address (requires an IPv4 address to be set first) |

repair

| | |
|----------|--|
| Function | Runs various database checks and repair processes |
| Syntax | <ul style="list-style-type: none">• <i>repair business_hour_saved_reports</i> - Converts older saved reports with business hours specified to the newer format• <i>repair database [db_name] [db_pass]</i> - Repairs errors in the database specified• <i>repair history_tables</i> - Repairs history tables that have the wrong col type for octet-deltacount• <i>repair policy_priority_order</i> - Repairs duplicate policy IDs• <i>repair range_starts</i> - Repairs history tables that may not have a start time to help identify the range of data within them; should only be used when instructed by a Plixer Technical Support engineer |

rotate

| |
|---|
| Warning: These commands will alter Plixer Scrutinizer's functions and should be used with caution. |
|---|

| | |
|----------|--|
| Function | Rotates Plixer Scrutinizer's keys and certificates |
| Syntax | <ul style="list-style-type: none">• <i>rotatekeys</i> - Creates a new encryption key and re-encrypts all encrypted fields in the database• <i>rotatecerts</i> - Creates new database certificates used for authentication |

services

Warning: These commands will alter Plixer Scrutinizer's functions and should be used with caution.

| | |
|----------|---|
| Function | Manages Plixer Scrutinizer services |
| Syntax | <ul style="list-style-type: none">• <i>services</i> <<i>service all</i>> <<i>start stop restart</i>> - Starts, stops, or restarts the specified service or all services |

set

| Function | Modifies certain behaviors related to authentication and general operation |
|----------|--|
| Syntax | <ul style="list-style-type: none"> • <i>set columnmoniker [old_element] [new_element] [element_list]</i> - Renames an information element • <i>set dns</i> - Modifies the system file to manage the list of DNS servers • <i>set hostinfo [ip_address] [fqhn]</i> - Sets the local machine name to the fully qualified host name provided • <i>set httpd [port]</i> - Changes the web port of non-SSL installs for the Plixer Scrutinizer WebUI • <i>set myaddress [ip_address] [netmask] [gateway]</i> - Changes the IPv4 address of the current Plixer Scrutinizer server • <i>set myaddress [ipv6_address/cidr] [gateway]</i> - Changes the IPv6 address of the current Plixer Scrutinizer server • <i>set ntp</i> - Modifies system file to manage the list of NTP servers • <i>set partitions [partition_name] <extend></i> - Expands the operating system disk space for hardware and virtual appliances • <i>set password plixer</i> - Resets the CentOS plixer user's password • <i>set password webui [user]</i> - Modifies the WebUI password for the specified user • <i>set permissions</i> - Resets file and directory permissions • <i>set registercollector [collector_ip] [secondary]</i> - Manually registers a Collector for use in a distributed environment • <i>set reportmenu</i> - Manually recreates the Reports menu • <i>set salt [salt]</i> - Sets a salt value to allow the users to mask certain machine characteristics from any license key generated • <i>set selfregister <reset></i> - Manually registers the current Plixer Scrutinizer server to identify itself for both standalone and distributed functionality |
| 342 | <ul style="list-style-type: none"> • <i>set selfreporter</i> - 6. Advanced Services Plixer Scrutinizer server to the primary Reporter role in a distributed environment • <i>set sshcollectorkeys</i> - Generates a new SSH key pair and distributes it to all active agents |

show

| | |
|----------|--|
| Function | Shows various details about the Plixer Scrutinizer server |
| Syntax | <ul style="list-style-type: none">• <i>show alarms [filter]</i> - Displays a list of Alarms sorted by timestamp (newest first)• <i>show diskspace</i> - Displays available storage• <i>show dns</i> - Displays a list of DNS servers currently used to resolve hostnames• <i>show exporters [filter]</i> - Displays a list of Exporters that are currently sending data to Plixer Scrutinizer based on the supplied filter• <i>show extalarms [filter]</i> - Displays a list of Alarms with extended JSON data sorted by timestamp (newest first)• <i>show groups</i> - Displays a list of device groups currently configured on the Plixer Scrutinizer server• <i>show interfaces [filter]</i> - Displays a list of interfaces that are currently sending data to Plixer Scrutinizer based on the supplied filter• <i>show ipaddresses</i> - Displays the current IP addresses on the Plixer Scrutinizer server• <i>show metering [filter]</i> - Displays a list of matching Exporter IP addresses and how each is metered (i.e. ingress and/or egress)• <i>show ntp</i> - Displays a list of NTP servers currently used to synchronize time• <i>show partitions</i> - Displays a list of partitions on the current Plixer Scrutinizer appliance• <i>show pcaplist</i> - Displays a list of all created pcap files and their sizes• <i>show serverpref [filter]</i> - Displays all serverpref elements matching the supplied filter and their current values• <i>show task [name]</i> - Displays a list of tasks currently configured in Plixer Scrutinizer• <i>show timezone</i> - Displays the current time-zone of the Plixer Scrutinizer server• <i>show tzlist [filter]</i> - Displays a list of time-zones matching the supplied filter• <i>show unknowncolumns</i> - Displays a list of Exporter information• <i>show yum_proxy</i> - Displays the currently configured yum proxy |

snoop

| | |
|----------|---|
| Function | Listens for traffic at the interface level |
| Syntax | <ul style="list-style-type: none">• <i>snoop interfaces [interface_name]</i> - listens for interface traffic from the specified interface• <i>snoop ipaddresses [ip_address]</i> - lists for interface traffic from the specified IP address |

system

| | |
|----------|---|
| Function | Performs system-level functions for Plexier Scrutinizer |
| Syntax | <ul style="list-style-type: none">• <i>system <restart shutdown></i> - Reboots or shuts down the system |

unlock

| | |
|----------|--|
| Function | Unlocks accounts that have exceeded the configured maximum number of failed login attempts |
| Syntax | <ul style="list-style-type: none">• <i>unlock [username] [auth_method]</i> - Unlocks a locked account using the specified authorization protocol |

upload

| Function | Uploads files to assist with troubleshooting issues |
|----------|--|
| Syntax | <ul style="list-style-type: none">• <i>upload pcap [capturefile>]</i> - Uploads the specified capture file collected by the <code>collect pcap</code> command• <i>upload supportfiles</i> - Uploads support files for troubleshooting purposes |

version

| Function | Displays the current version of Plixer Scrutinizer |
|----------|---|
| Syntax | <ul style="list-style-type: none">• <i>version</i> - Shows the system's current version information |

6.3 Plixer Scrutinizer APIs

6.3.1 IP Groups

The IP Groups API functionality is a simple way to add, remove, and edit IP Groups.

Prerequisites

When using the API, the following will be used on all requests:

- **authToken** - The authentication token from Plixer Scrutinizer that allows access to API
- **rm** - The runmode for accessing the API. It is specific to each section of the product. `ipgroups` will be used for each of the following examples

- **action** - The list of available actions will change with each request. Below are the actions available within the `user_api` run mode:

| Action | Description |
|------------------|---|
| saveRule | Create a defined IP Group |
| update | Redefine or modify an existing IP Group |
| loadTreeRootFast | Load condensed list all IP Group names and IDs |
| search | Search for an IP Group by name |
| loadRules | View all rule definitions for an IP Group |
| deleteRule | Remove a rule from an IP Group |
| delete | Delete an IP Group |
| deleteAll | Delete all defined IP Groups from Plexier Scrutinizer |

Rules

IP host

One or multiple IPs can be used to define a rule.

```
[
  {
    "type": "ip",
    "sip": "10.1.1.1"
  }
]
```

```
[
  {
    "type": "ip",
    "sip": "192.168.1.1"
  },
  {
    "type": "ip",
    "sip": "192.168.2.2"
  }
]
```

IP range

Uses a range of IPs instead of multiple IP rules.

```
[
{
  "type": "range",
  "sip": "10.1.1.1",
  "eip": "10.1.1.254"
}
```

Note: The 'range' IP rule type requires a start IP (sip) and end IP (eip) to define the start and end of the range.

IP subnet

Uses a subnet and mask instead of multiple IP rules or ranges.

```
[
{
  "type": "network",
  "address": "192.168.0.0",
  "mask": "16"
}
```

Note: A subnet rule uses the 'network' type. A subnet mask is required.

All IPs

Specifies all IPs to be used in an IP Group definition.


```
[
  {
    "type": "ipall",
    "all": 1
  }
]
```

Wildcard

Specifies a rule based on mask of bits that indicates which parts of an IP address are to be used for defining the IP Group hosts.

```
[
  {
    "type": "wildcard",
    "address": "10.0.4.0",
    "mask": "0.255.0.255"
  }
]
```

Note: The example above tags all hosts with the first octet of '10' and the third octet of '4'. Therefore, IPs such as 10.1.4.1, 10.2.4.250, 10.99.4.98, etc., would be included in the defined IP Group as the first and third octets match in the wildcard rule.

Child group

Nests IP Groups to create a hierarchy with child groups' rules being more specific than their parent.

```
[
  {
    "type": "child",
    "child_id": "16900062"
  }
]
```

Important: A child group definition is based on the parent group. Define smaller and more distinct child groups, then create the parent group so that you can add the child group that already exists.

1. Create **UK Datacenter** and **UK Office** groups for their respective subnets/IPs.
2. Create a parent group, **UK**, that will include child groups, *UK Datacenter* and *UK Office*.
3. Create **Germany Datacenter** and **Germany Office** groups for their respective subnets/IPs.
4. Create a parent group, **Germany**, that will include child groups, *Germany Datacenter* and *Germany Office*.
5. Create a parent group **European Offices** that will include child groups *UK* and *Germany*.

The workflow above will create the following hierarchy:

```
* European Offices      10.0.0.0/8
  ** UK                  10.30.0.0/16
    *** UK Datacenter    10.30.10.1/32
    *** UK Office        10.30.20.0/24
  ** Germany             10.40.0.0/16
    *** Germany Datacenter 10.40.10.1/32
    *** Germany Office     10.40.20.0/24
```

Creating an IP Group

When creating IP Groups with the API, use the `saveRule` action. The following are the additional fields that can be used:

- **new_fc** - Provides a name for the IP Group
- **added** - Specifies a JSON array of rules to add to/define the IP Group

Below is an example of how to use the `added` field for tagging a single IP address:

JSON object expected:

```
[
  {
    "type": "ip",
    "address": "10.1.4.66"
  }
]
```

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
→fcgi' \
--form 'authToken={{authToken}}' \
--form 'rm=ipgroups' \
--form 'action=saveRule' \
--form 'new_fc=UK Data Center' \
--form 'added=[
    {
        "type": "ip",
        "address": "10.30.10.1"
    }
]'
```

JSON object returned:

```
{
  "removed": [],
  "updated": [],
  "added": [
    {
      "rule_id": 506588,
      "cid": null,
      "type": "ip",
      "address": "10.30.10.1"
    }
  ],
  "warnings": [],
  "fc_id": 16900006,
  "myrules": "IP Address:10.30.10.1",
  "fc_name": "UK Datacenter",
  "rule_id": 506588,
```

(continues on next page)

(continued from previous page)

```
{  
  "total": 1  
}
```

Updating IP Groups

You can update or remove existing rules or add new ones with one request. There are four optional and additional fields that you can use with the `update` action:

- **name** - If specified, the name of the IP Group will be updated.
- **added** - Specifies a JSON array of rules to add to/define the IP Group.
- **updated** - Specifies a JSON array of rules to modify the IP Group. The `rule_id` field must be defined to change a rule.
- **removed** - Specifies a JSON array of rule IDs to be removed.

Important: When updating IP Groups, the rule type must remain the same. For example, you can not change an IP rule to a subnet rule. The workflow in that case is to remove the old rule type and create a new one.

Note: You can leave any of the [name|added|updated|removed] fields empty.

Below is an example of the `added` field for tagging traffic for a single IP address on a specific port:

```
[  
  {  
    "type": "ip",  
    "address": "10.1.4.66"  
  }  
]
```

Use the `updated` field to change an existing rule:

```
[
  {
    "rule_id": "84",
    "type": "network",
    "address": "10.30.0.0",
    "mask": "16"
  }
]
```

Below is an example of the removed field syntax:

```
[ 81, 82, 83 ]
```

Note: The three examples above result to the following:

- Create two new rules for the IP Group, an IP rule, and a port rule.
- Update the rule with ID 84 to change IPs to match on.
- Remove rules 81, 82, and 83 that already existed in the IP Group definition.

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
→fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken={{authToken}}' \
--form 'rm=ipgroups' \
--form 'action=update' \
--form 'fc_id={{new_ipgroup_fcid}}' \
--form 'name=Renamed Group' \
--form 'added=[
    {
        "type": "ip",
        "address": "10.1.4.66"
    }
]' \
--form 'updated=[
```

(continues on next page)

(continued from previous page)

```
{
    "rule_id": "84",
    "type": "ip",
    "address": "192.1.0.0"
}
]' \
--form 'removed=[114]'
```

Search IP Groups

This feature is useful for searching Plexer Scrutinizer for IP Groups with a partial string or a full name. This feature can be used to search for IP Groups with comparisons such as “like ‘UK Office’” or “notLike ‘Email Server’”. There are four additional fields that you can use with the **search** action:

- **name** - The name of the IP Group (or string) to find in Plexer Scrutinizer.
- **fc_name_comp** - Specifies the comparison criteria for search results. Valid options are [like|notLike].
- **page** - Default is 1, which loads the first page of pagination.
- **maxRows** - The number of results per page returned in the API response.

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.  
↳fcgi' \  
--header 'Content-Type: application/json' \  
--form 'authToken={{authToken}}' \  
--form 'rm=ipgroups' \  
--form 'action=search' \  
--form 'name={{search_term}}' \  
--form 'fc_name_comp={{search_type}}' \  
--form 'page=1' \  
--form 'maxRows=10'
```

Deleting entries from IP Groups

Use this command to remove a single IP from a specific IP Group. The `rule_id` is required and can be obtained by viewing the rules of an IP Group. The update action also has an optional `removed` field that can be used to delete rules from an IP Group.

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
→fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken={{authToken}}' \
--form 'rm=ipgroups' \
--form 'action=deleteRule' \
--form 'rule_id=506588'
```

JSON object returned:

```
{
  "fc_id": 16900006,
  "success": 1,
  "myrules": "",
  "rule_id": "506588",
  "total": 0
}
```

Deleting IP Groups

You can remove more than one IP Group at a time by specifying more IDs in the array. The following is an additional field that is required with the `delete` action:

json

Array of IP Group IDs to be deleted

For example, here's how you define the `json` field to remove a single IP Group:

```
[
  {
    "id": "16900032"
  }
]
```

This is the json field to remove multiple IP Groups:

```
[
  {
    "id": "16900032"
  },
  {
    "id": "16900033"
  }
]
```

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
↳fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken={{authToken}}' \
--form 'rm=ipgroups' \
--form 'action=delete' \
--form 'json=[
    {
      "id": "16900032"
    }
  ]'
```

JSON object returned:

```
{
  "processedCount": 1,
  "removed": [
    "16900006"
  ]
}
```


6.3.2 Reporting

The reporting API is a simple way to access the report data via HTTP or the command line.

Prerequisites

- **authToken** - The authentication token from Plexier Scrutinizer that allows access to API.
- **rm** - The runmode for accessing the API. It is specific to each section of the product. `report_api` will be used for each of the following examples.
- **action** - The list of available actions will change with each request. Below is an action within the `report_api` run mode:

| Field | Description |
|-------|---|
| get | the action used to execute flow reports |

- **rpt_json** - An array of details to specify most of the options found in the gear menu of the UI, where you can specify report type, date range, etc. See the [See Available Parameters](#) section on how to fill out the `rpt_json` field. Every report has an API tab for obtaining JSON for it.
- **data_requested** - Indicates what part of the report is needed. 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. While `rpt_json` tells Plexier Scrutinizer how to prepare the data (timeframe, filters, aggregation, etc.), [data_requested](#) tells Plexier Scrutinizer what data should be returned with the request (inbound, outbound, table, graph, or just a table name).

Available parameters

rpt_json

This field tells Plexier Scrutinizer how to prepare the data: report type, timeframe, filters, aggregation, etc.

Expected JSON object:

```
{
  "reportTypeLang": "conversations",
  "filters": {
    "sdfDips_0": "in_0A190101_ALL"
  },
  "reportDirections": {
    "selected": "inbound"
  },
  "times": {
    "dateRange": "LastFiveMinutes",
    "clientTimezone": "America/New_York"
  },
  "dataMode": {
    "selected": "saf"
  },
  "rateTotal": {
    "selected": "total"
  },
  "dataGranularity": {
    "selected": "auto"
  },
  "bbp": {
    "selected": "bits"
  }
}
```

The example above shows the minimum needed option for each field. The following table is a breakdown of the available options for each of the fields from above:

reportTypeLang

A language keycode that represents a report type

| Report Lang | Report Description |
|-----------------|---|
| conversations | Conversations WKP (the default report) |
| host2host | Host to Host |
| ipGroupGroup | IP Group to IP Group |
| applications | Applications defined |
| country2country | Country to Country |
| ..etc.. | Each report has an API tab, with the report lang, in the UI |

filters

| | | |
|-----------|---|------------------------|
| sdfDips_0 | The exporter and interface to use to generate report with. The syntax for this field is in_[IP Hex]_[interface(s)]. | |
| | Value | Description |
| | in_0A190101_ALL | Uses all interfaces |
| | in_0A190101_0A190101_1 | Uses interface index 1 |

reportDirections

| | | |
|---------------|---|--|
| se- lected | Specifies which direction to query, relative to the interfaces in the filter. If not applicable, use <i>inbound</i> . | |
| | Value | |
| | inbound | |
| | outbound | |

times

Specifies the *dateRange* in the report

| | | |
|-----------------|---|--|
| dateRange | Indicates the timeframe to include in the report | |
| | Value | |
| | LastFiveMinutes | |
| | LastTenMinutes | |
| | LastFifteenMinutes | |
| | LastTwentyMinutes | |
| | LastThirtyMinutes | |
| | LastFortyfiveMinutes | |
| | LastHour | |
| | LastFullHour | |
| | LastThreeDays | |
| | LastSevenDays | |
| | LastThirtyDays | |
| | Today | |
| | Yesterday | |
| | Last24Hours | |
| | ThisWeek | |
| | LastWeek | |
| | ThisMonth | |
| | LastMonth | |
| | ThisYear | |
| | LastYear | |
| | Custom | |
| start | Start date of the data to include in the report. Use <i>Custom</i> in the <i>dateRange</i> field to set the start date. | |
| end | End date of the data to include in the report. Use <i>Custom</i> in the <i>dateRange</i> field to set the end date. | |
| clientTime-zone | Displays dates local to your time zone | |
| | Value | |
| | America/New_York | |
| | America/Los_Angeles | |

dataMode

Saves and rolls up data to condense collected information

| | | |
|----------|--|------------------------------|
| selected | Indicates the timeframe to include in the report | |
| | Value | Description |
| | saf | Default <i>dataMode</i> type |
| | traditional | Used for legacy support |

rateTotal

Specifies whether to display data as a rate or as total traffic

| se- lected | Value | Description |
|---------------|-------|---|
| | rate | Displays row data as a rate. For example, packets per second or bits per second |
| | total | Displays row data as total traffic. For example, total traffic seen within the time-frame |

dataGranularity

Specifies how to retrieve data for the report

| se- lected | Value | Description |
|---------------|-------|---|
| | auto | Lets the API select the data aggregation method |
| | 1m | Displays row data as total traffic. For example, total traffic seen within the time-frame |
| | 5m | Shows data from 1-minute granularity |
| | 30m | Shows data from 5-minute granularity |
| | 2h | Shows data from 2-hour granularity |
| | 12h | Shows data from 12-hour granularity |

bbp

Determines how to display the data

| se- lected | Value | Description |
|---------------|-------|---|
| | bits | Displays data as bits per second or total bits, depending on the <i>rateTotal</i> selection |
| | bytes | Displays data as bytes per second or total bytes, depending on the <i>rateTotal</i> selection |
| per- cent | | Displays data as a percentage, <i>rateTotal</i> selection is disregarded |
| | | |

data_requested

This field tells Plixer Scrutinizer how to prepare the data for graphs, table pagination, etc.

Important: The direction specified in `data_requested` must match the `reportDirections` selected value in `rpt_json` (e.g. `inbound/inbound` or `outbound/outbound`).

Expected JSON object expected:

```
{
  "inbound": {
    "graph": "none",
    "table": {
      "query_limit": {
        "offset": 0,
        "max_num_rows": 10
      }
    }
  }
}
```

The example above shows the minimum needed option for each field.

Running Reports

The following example is an API call to run a default report, over the last 5 minutes, on all interfaces of a device.

Note: In the example below, you must replace `{{scrutinizer_ip_address}}` with your Plexer Scrutinizer's IP address, as well as `{{authToken}}` with an Authentication Token that can be obtained from the UI.

Example API call:

```
curl --location --request POST 'https://{{scrutinizer_ip_address}}/fcgi/scrut_
↳fcgi.fcgi' \
--form 'authToken={{authToken}}' \
--form 'rm=report_api' \
--form 'action=get' \
--form 'rpt_json=
```

(continues on next page)

(continued from previous page)

```

{
  "reportTypeLang": "conversations",
  "filters": {
    "sdfDips_0": "in_0A190101_ALL"
  },
  "reportDirections": {
    "selected": "inbound"
  },
  "times": {
    "dateRange": "LastFiveMinutes",
    "clientTimezone": "America/New_York"
  },
  "dataMode": {
    "selected": "saf"
  },
  "rateTotal": {
    "selected": "total"
  },
  "dataGranularity": {
    "selected": "auto"
  },
  "bbp": {
    "selected": "bits"
  }
}
}' \
--form 'data_requested=
{
  "inbound": {
    "graph": "none",
    "table": {
      "query_limit": {
        "offset": 0,
        "max_num_rows": 10
      }
    }
  }
}
}'

```

The API call above is processed by the reporting engine, and then the server returns a JSON response.

JSON object returned:

```
{
  "report": {
    "request_id": "0xed184820e4b611eab58f1fc02130f7f9",
    "table": {
      "inbound": {
        "totalRowCount": 1,
        "footer": [],
        "columns": [],
        "rows": []
      }
    },
    "time_details": {},
    "exporter_details": {},
    "graph": {}
  }
}
```

The sample response above is condensed to show the typical structure of a JSON response. The following table is a breakdown of the most important fields from the `report` field/key from the above response:

| Field | | Description | |
|-------|-----------------|--|--|
| table | | The table data is divided into inbound or outbound. For each direction the following are provided: | |
| | columns | 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. | |
| | | 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 | Each collection in the array represents a row and each object in the collection represents a table cell for that row. | |
| | | rawValue | The unformatted value as it is returned from the database |
| | | label | The formatted value including bits, bytes, or percent |
| | footer | | |
| | | The footer[0] represents the ‘Others’ data of operations columns (columns in which the value is the product of an arithmetic operation). That is, the amount of the data not included in rows. | |
| | | The footer[1] represents the ‘total’. Total for a column is equal to the sum of the rows of that column plus the ‘Others’ value of that column. | |
| | total-Row-Count | Integer representing the total number of rows available | |
| graph | | | |
| | | all | Returns all available graph types |
| | | pie | Values used in drawing pie chart of table data |
| | | timeseries | Values used in drawing line graph chart of table data |
| | | none | Only returns data for pie, as that is the default |

6.3.3 User management

The user API functionality allows creating user accounts within Plexier Scrutinizer and adding them to user groups at the same time.

Prerequisites

When using the API, pass the following mandatory fields:

- **authToken** - The authentication token from Plexier Scrutinizer that allows access to the API.
- **rm** - The runmode for accessing the API. It is specific to each section of the product. `user_api` will be used for each of the following examples.
- **action** - The list of available actions will change with each request. Below is a list of available actions within the `user_api` run mode:

| Action | Description |
|------------------|--|
| createUser | Allows the creation of user accounts within Plexier Scrutinizer. You can add the accounts to user groups at the same time. |
| delUsers | Allows the deletion of user accounts within Plexier Scrutinizer by ID or name. |
| createUser-group | Creates user groups. You can add members at the same time. |
| delUser-groups | Deletes user groups. |
| member-ship | Changes usergroup membership. When adding or removing a user, use only either <code>user_id</code> or <code>user_name</code> . |
| prefs | Makes changes to the user preferences for individual accounts. It contains an array of preferences and new settings. |
| permissions | Makes changes to a user group's permissions. |
| changeUser-name | Changes the name of a user account. |

Creating users

The `createUser` action allows creating users within the API. It calls for an additional field:

json

An array of users each contains the name, password, and group template that user will be a member of.

Expected JSON object:

```
{
  "users": [
    {
      "name": "MyAdmin",
      "pass": "secretAdminPass",
      "membership": [2]
    },
    {
      "name": "MyGuest",
      "pass": "myGuestPass",
      "membership": [2]
    }
  ]
}
```

| Field | Description |
|------------|---|
| users | An array of all users to be created. Multiple users can be created at once. If only one user is needed, this will be an array of one. |
| name | The name for the new user account. |
| 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. You can add your own usergroups. By default, there are two groups installed. Group 1 is the Administrators group. Group 2 is reserved for Guests. |

JSON object returned:

```
{
  "data": [
    {
      "id": 3,
      "name": "MyAdmin"
    },
    {
      "id": 4,
      "name": "MyGuest"
    }
  ]
}
```

| Field | Description |
|-------|---|
| data | An array of responses for each user account that Plexier 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). |

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.  
↳fcgi' \  
--form 'authToken={{authToken}}' \  
--form 'rm=user_api' \  
--form 'action=createUser' \  
--form 'json=  
{  
  "users": [  
    {  
      "name": "MyAdmin",  
      "pass": "MyPass",  
      "membership": [ 1 ]  
    },  
    {  
      "name": "MyGuest",  
      "pass": "OtherPass",  
      "membership": [ 2 ]  
    }  
  ]  
}'
```

Note: If Plexier Scrutinizer is using a self-signed certificate, add `--insecure` to the header options to tell curl to ignore it.

Deleting users

The `delUser` action allows the deletion of user accounts within Plexier Scrutinizer by ID or name. There is an additional field used with the action:

json

An array of users where each contains the name, password, and group template that user will be a member of.

Expected JSON object:

```
{
  "delUsers": [
    11,
    "MyGuest",
    207
  ]
}
```

| Field | Description |
|--------------|---|
| delUsers | An array of all users to be deleted. You can delete multiple users at once. If only one user is needed, this will be an array of one. |
| id/- name | The user_id of the user to be deleted. Alternatively, the name of the user can be used. |

JSON object returned:

```
{
  "data": [
    "Deleting user id 11 (1 matched)",
    "Deleting user named 'MyGuest' (1 matched)"
    "Deleting user id 207 (0 matched)",
  ]
}
```

| Field | Description |
|-------|--|
| data | An array of responses for each user account that Plixer Scrutinizer attempted to delete. The example includes a failure message when a user did not exist. |

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.  
→fcgi' \  
--form 'authToken={{authToken}}' \  
--form 'rm=user_api' \  
--form 'action=delUsers' \  
--form 'json=
```

(continues on next page)

(continued from previous page)

```
{
  "delUsers": [
    11,
    "MyGuest",
    207
  ]
}
```

Creating user groups

You can use the `createUsergroup` action to create user groups and add members at the same time. It requires an additional field:

json

An array of usergroups. Each entry contains the name of the usergroup, the ID of the usergroup to use as a template, and the id or name of the users to be added to the group.

Expected JSON object:

```
{
  "usergroups": [
    {
      "name": "GroupA",
      "template_usergroup": 1,
      "users": [1, 2]
    },
    {
      "name": "GroupB",
      "template_usergroup": 2,
      "users": ["MyUser", "MyUser2"]
    }
  ]
}
```

| Field | Description |
|---------------------|--|
| usergroups | An array of responses for each usergroup that Plexier Scrutinizer attempted to create. |
| name | The name to be applied to the usergroup. |
| tem-plate_usergroup | The ID of the user group to use as a template for creating a new user group. |
| 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 creates an empty user group. |

JSON object returned:

```
{
  "data": [
    {
      "id": 5,
      "name": "GroupA",
      "members": ["1", "2"]
    },
    {
      "name": "GroupB",
      "error": "A usergroup already exists with that name"
    }
  ]
}
```

| Field | Description |
|---------|---|
| data | An array of responses for each user group that Plexier Scrutinizer attempted to create. |
| id | The new usergroups_id of the user group that was created. |
| name | The name of the user group created (by design this is identical to the name passed in) |
| members | An array of user IDs or user names for the members successfully added to the group |
| error | Any errors encountered during the creation of a particular user group |

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
→fcgi' \
--form 'authToken={{authToken}}' \
--form 'rm=user_api' \
--form 'action=createUsergroup' \
--form 'json=
{
  "usergroups": [
    {
      "name": "GroupA",
      "template_usergroup": 1,
      "users": [1,2]
    },
    {
      "name": "GroupB",
```

(continues on next page)

(continued from previous page)

```
        "template_usergroup":2,  
        "users":["MyUser","MyUser2"]  
      }  
    ]  
  }'  
'
```

Deleting user groups

The delUsergroups action deletes user groups. It has an additional field:

json
An array of usergroups. Each entry contains the name or ID of the usergroup to be deleted.

Expected JSON object:

```
{  
  "delUsergroups": [  
    3,  
    "My Usergroup"  
  ]  
}
```

| Field | Description |
|---------------|---|
| delUsergroups | An array of responses for each user group that Plixer Scrutinizer attempted to delete |
| id/name | The usergroups_id or exact name of the user group to be deleted |

JSON object returned:

```
{  
  "data": [  
    "Deleting usergroup named '3' (1 matched)",  
    "Deleting usergroup id My Usergroup (0 matched)"  
  ]  
}
```


| Field | Description |
|-------|---|
| data | An array of responses for each usergroup that Plixer Scrutinizer attempted to delete. |

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
→fcgi' \
--form 'authToken={{authToken}}' \
--form 'rm=user_api' \
--form 'action=delUsergroups' \
--form 'json=
{
    "delUsergroups": [
        3,
        "My Usergroup"
    ]
}'
```

Modifying group membership

The membership action changes user group membership. When adding or removing a user, use either `user_id` or `user_name`, but not both as shown below. There is an additional field that can be used with the membership action in the user API:

json

Contains two arrays, `add` and `remove`, which have information on each membership change.

Expected JSON object:

```
{
  "membership":
  {
    "add": [
      {
        "user_id": 13,
        "usergroup_id": 2
      },
      {
```

(continues on next page)

(continued from previous page)

```
        "user_name": "USER2",
        "usergroup_name": "USERGROUP2"
      },
    ],
    "remove": [
      {
        "user_name": "USER3",
        "usergroup_id": 4
      }
    ]
  }
}
```

| Field | Description |
|-----------------|---|
| membership | Contains two arrays, add and remove, which have information on each membership change |
| user_id | Required for the user with preferences to change |
| user_name | An alternative to user_id. It can be the plain text name of the user. |
| usergroup_id | The ID from plixer.usergroups that the user will be added to or removed from. |
| user-group_name | An alternative to usergroup_id. It can be the plain text name of the user group. |

JSON object returned:

```
{
  "data": {
    "added": [
      "User 13 added to usergroup 1",
      "User 14 added to usergroup 3",
    ],
    "removed": [
      "User 15 removed from usergroup 4"
    ]
  }
}
```

| Field | Description |
|----------|---|
| data | An array of responses for membership updated |
| added | Contains an array of either statements of success or statements of errors explaining why the membership change failed |
| re-moved | Contains an array of either statements of success or statements of errors explaining why the membership change failed |

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.  
↳fcgi' \  
--form 'authToken={{authToken}}' \  
--form 'rm=user_api' \  
--form 'action=membership' \  
--form 'json=  
{  
  "membership": {  
    "add": [  
      {  
        "user_id": 3,  
        "usergroup_id": 17  
      },  
      {  
        "user_id": "USER2",  
        "usergroup_id": "GROUPB"  
      }  
    ],  
    "remove": []  
  }  
}'
```

Editing user preferences

The `prefs` action makes changes to the user preferences for individual accounts. It contains an array of preferences and new settings. There is an additional field used with the `prefs` action in the user API:

json

The `user_id` and array of `prefs` each contains the pref code and setting value to be modified.

Expected JSON object:

```
{
  "user_id": 11,
  "prefs": [
    {
      "pref": "statusTopn",
      "setting": 10
    },
    {
      "pref": "language",
      "setting": "english"
    }
  ]
}
```

| Field | Description |
|---------|--|
| user_id | Required for the user with preferences to change |
| prefs | An array of user preferences and setting values |
| pref | The Plexier Scrutinizer user preference to edit |
| setting | The value that will be set for the user_id specified |

JSON object returned:

```
{
  "data": {
    "updated": [
      "statusTopn updated to 10 for user_id 11",
      "language updated to english for user_id 11"
    ],
    "errors": []
  }
}
```

| Field | Description |
|---------|---|
| 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 |

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
→fcgi' \
--form 'authToken={{authToken}}' \
--form 'rm=user_api' \
--form 'action=prefs' \
--form 'json=
{
    "user_id":11,
    "prefs":[
        {
            "pref":"statusTopn",
            "setting":10
        },
        {
            "pref":"language",
            "setting":"english"
        }
    ]
}'
```

Changing permissions

The `permission` action makes changes to a user group's permissions. Users inherit permissions from their user group. There is an additional field used with the `permissions` action in the user API:

json

An array of permissions each contains a usergroup identifier (name or ID), the security code, and permission type to be modified.

Expected JSON object:

```
{
  "permissions": {
    "add": [
      {
        "usergroup_name": "Dashboarders",
        "permission_type": "gadget",
        "seccode": "lLabelCPU"
      }
    ]
  }
}
```

(continues on next page)

(continued from previous page)

```

    ],
    "remove": [
      {
        "usergroup_name": "ReadOnlyReporters",
        "permission_type": "plexer",
        "seccode": "allGadgets"
      }
    ]
  }
}

```

| Field | Description | |
|-----------------|--|--|
| permissions | Contains two arrays, add and remove, which have information on each permission change | |
| add/remove | Each contains an array of objects of permissions and user groups | |
| usergroup_id | The ID from plixer.usergroups that the user will be added to or removed from | |
| usergroup_name | An alternative to usergroup_id, and can be the plain text name of the user group | |
| permission_type | Differentiates the types of permissions. Options include: | |
| | Permission | Description |
| | device | A hex representation of the IP address of a device |
| | interface | Hyphenated, a hex representation of the IP address |
| | group | The group ID of a map/device group from plixer |
| | report | The saved_id of a saved report from reporting. |
| | gadget | The gadget_id of a dashboard gadget, from plixer |
| | thirdparty | The id of a third-party link from plixer.third_party |
| | bboard | The bulletin board id from plixer.alm_bulletin_board |
| | plexer | A static string of permission codes we use for devices |
| 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.) | |
| | seccode | Description |
| | 3rdPartyIntegration | Create, edit, and delete third-party integration links |
| | ackBBEvent | Ability to acknowledge events on Alarms tab bulletin board |
| | adminTab | Access the Admin Tab |
| | alarmSettings | Configure alarm notifications |
| | alarmsTab | Access the Alarms tab |
| | allBBoards | View all Alarms bulletin boards |
| | allDevices | The status of all devices and each of their interfaces |
| | allGadgets | Every gadget created on the Dashboards tab, by plixer |
| | allGroups | Access to all maps/device groups created in Scrutinizer |
| | allInterfaces | Report on interfaces for any device |

| | | |
|--|----------------------------|---|
| | allLogalotReports | All Logalot Reports |
| | allReportFolders | Permission to all saved report folders |
| | allReports | Saved reports created by any user |
| | allThirdparty | All configured third-party links will be available |
| | almDelete | Permission to permanently delete alarms |
| | ApplicationGroups | Configure Application Groups |
| | asnames | Configure AS Names |
| | auditing | Access the Auditing report containing logs of |
| | auth | Manage external authentication tokens |
| | Authentication | Manage external authentication types |
| | authLdapServers | Manage LDAP server configuration used for |
| | awsSettings | AWS configuration |
| | changeUserPasswords | The ability to change the passwords of other |
| | createDashTabs | Create new Dashboards |
| | createUsers | The ability to create new local Plexier Scrutinizer |
| | CrossCheck | View and edit CrossCheck configuration, wh |
| | crossCheckView | Access to the CrossCheck methods table view |
| | dashboardAdmin | Manage all dashboards created by any user |
| | DataHistory | Configure settings that control how long Plexier |
| | deleteReport | Ability to delete saved reports regardless of c |
| | deleteUsers | The ability to delete local Plexier Scrutinizer |
| | DeviceDetails | Edit device interface details |
| | EmailNotifications | Configure the mailserver Plexier Scrutinizer v |
| | faExclusions | Configure Flow Analytics exclusions |
| | fa_mgmt_link | Configure Flow Analytics thresholds and sett |
| | feedbackForm | Access the link to send feedback to Plexier |
| | FlowAnalyticsSettings | Global Flow Analytics settings |
| | helpTab | Access the Help tab |
| | HostNames | Edit Host Name information |
| | IPGroups | Configure Plexier Scrutinizer IP Groups |
| | language | Create and edit language localization settings |
| | licensing | Configure Plexier Scrutinizer product licensin |
| | LogalotPrefs | Configure global alarms settings |
| | MACAddresses | Configure device MAC Address information |
| | ManageCollectors | Manage the devices collecting flow data for P |
| | ManageExporters | Manage the devices exporting flow data to P |
| | mappingGroupConfiguration | Create and edit Maps/Groups |
| | mappingObjectConfiguration | Create and edit Mapping Objects |
| | mapsTab | Access the Maps tab |
| | myViewTab | Access the Dashboards tab |

Table

| | |
|---------------------|--|
| NotificationManager | Manage alarm notifications |
| PolicyManager | Manage alarm policies |
| protocolExclusions | Edit which protocols are discarded from flow reports |
| proxySettings | Configure proxy server settings in Plexier Scrutinizer |
| radiusConf | Manage RADIUS server configuration used for authentication |
| ReportDesigner | Design new custom report types |
| reportFilters | Permission to update the filters used in Status Tab |
| reportFolders | Manage saved report folders |
| reportSettings | Reporting engine configuration options |
| runReport | Ability to run flow reports |
| saveReport | Ability to name and save flow reports |
| scheduledReports | Create, edit, and delete scheduled email reports |
| sf_asa_acls | Configure ASA ACL descriptions |
| SNMPCredentials | Manage SNMP credentials used to poll device information |
| srCreate | Schedule a saved report to be emailed on a regular basis |
| sso | Add, Delete, and Edit Identity Provider configuration |
| statusTab | Access the Status Tab |
| syslogNotifications | Syslog server configuration |
| SystemPreferences | Administrative access to global Plexier Scrutinizer settings |
| tacacsConf | Manage TACACS+ server configuration used for authentication |
| tos | Edit TOS Configuration |
| userAccounts | Access to the Users view on the Admin Tab, list users |
| usergroups | Manage Scrutinizer usergroups |
| viewUserIdentity | View identity and access information relevant to users |
| viptelaSettings | Viptela Settings |
| Vitals | View the Plexier Scrutinizer server vitals reports |
| wkp | Edit WKP Configuration |

JSON object returned:

```
{
  "data": {
    "errors": [],
    "updated": [
      "Added gadget permission lLabelCPU to usergroup 26 ",
      "Removed plixer permission allGadgets from usergroup 27 "
    ]
  }
}
```


| Field | Description |
|---------|---|
| data | An array of responses for each permission change updated or attempted |
| updated | Messages for any successful changes to permissions |
| errors | An array of errors explaining why the permission change failed |

Example API call

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.  
→fcgi' \  
--form 'authToken={{authToken}}' \  
--form 'rm=user_api' \  
--form 'action=permissions' \  
--form 'json=  
{  
  "permissions": {  
    "add": [  
      {  
        "usergroup_id": 23,  
        "permission_type": "plexer",  
        "seccode": "statusTab"  
      }  
    ],  
    "remove": []  
  }  
}'
```

Changing usernames

The changeUsername action allows editing the name of a user account. It requires an additional field:

json

An array of user groups each containing the existing name of user, the new user name to be set. Alternatively, the user_id can be used instead of the oldname field.

Expected JSON object:

```
{
  "changeUsername":
  {
    "user_id": "14",
    "newname": "OpSCT"
  }
}
```

| Field | Description |
|---------|--|
| user_id | The user ID of the account to be changed |
| oldname | 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"
  }
}
```

| Field | Description |
|---------|--|
| data | An array of responses for each preference change updated or attempted |
| message | Contains either a statement of success or an error explaining why the name change failed |

Example API call:

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
↳fcgi' \
--form 'authToken={{authToken}}' \
--form 'rm=user_api' \
--form 'action=changeUsername' \
--form 'json=
{
  "changeUsername":
  {
```

(continues on next page)

(continued from previous page)

```
        "oldname": "myUser",  
        "newname": "OpSCT"  
    }  
}'
```

6.4 Reverse-path filtering

When reverse-path filtering is enabled, a Plexier Scrutinizer Collector is able to receive flows from IP addresses that it is unable to route to normally, such as non-local hosts whose traffic data is forwarded by a proxy or replication appliance.

This configuration should only be used when the Plexier Scrutinizer server/Collector is both **in a secure environment** and **using a single interface**.

Important: In multi-interface/multi-homed scenarios and/or where strict networking practices are observed, the recommendations in RFC 3704 should be followed. This ensures that spoofed/forged packets cannot be used to generate responses that are sent out over a different interface.

6.4.1 Enabling reverse-path filtering

To enable reverse-path filtering on a Plexier Scrutinizer Collector, find the following line in `/etc/sysctl.conf`:

```
net.ipv4.conf.default.rp_filter = 1
```

And change its value from 1 to 0.

In addition, the following steps are also recommended:

- To bypass having to restart networking after editing the file, run the command `sysctl net.ipv4.conf.default.rp_filter = 0` to turn reverse-path filtering on.
- Verify that the routing tables include routing data for all networks to be monitored to ensure that flows can be collected from non-local address spaces.

6.4.2 VRF (Virtual Routing and Forwarding) Mode

In some scenarios, such as when there are special security requirements or if the management network IP addresses overlap with collection-side interfaces, routing tables may need to be isolated from the management network.

Separate routing tables can be created to isolate management traffic to the management interface, so collection and polling traffic only impact their respective interfaces.

Sample routing table configuration

This example outlines the steps to configure two separate routing tables called `plexer` and `public` corresponding to interfaces `eth0` and `eth1` on a Plexier Scrutinizer deployment.

1. Add the two routing tables to `/etc/iproute2/rt_tables` after the line `#1 inr.ruhep`:

```
#
# reserved values
#
255 local
254 main
253 default
0 unspec
#
# local
#
#1 inr.ruhep
1 public
2 plexer
```

2. Create the files `route-eth0` and `route-eth1` under `/etc/sysconfig/network-scripts/` containing the following lines to define the default gateway for each table:

`route-eth0`

```
default via 172.16.2.20 table plexer
```

`route-eth1`

```
default via 10.1.1.251 table public
```

3. Add the gateway for each interface in `/etc/sysconfig/network-scripts/ifcfg-eth0` and `ifcfg-eth1` (no other changes are necessary) as follows:

`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
GATEWAY=172.16.2.20
```

`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
```

4. Reboot the server to restart networking.
5. Verify that networking is functioning and confirm that IP tables are configured to accept or deny the correct traffic on each interface.

6.5 Streaming to data lakes

Plixer Scrutinizer supports data streaming to customer data lakes.

For assistance with the configuration process, contact [Plixer Technical Support](#).

6.6 Backups

The Plixer Scrutinizer filesystem includes utilities that automate the process of creating or restoring system backups.

Note:

- These utilities are recommended for most long-term backup scenarios, because they include all database configuration and historical data for a Plixer Scrutinizer instance. Native snapshots may still be used as a short-term recovery option when there is no need to store the data, e.g., when upgrading the instance.
 - For Plixer Scrutinizer instances deployed on AWS, backups should be created and/or restored using native AWS functionality.
-

These utilities allow several types of backup and restore operations to be performed by the user.

6.6.1 Full backups

Full or comprehensive backups are disaster-recovery-grade images of a Plixer Scrutinizer instance and include the following elements of the filesystem:

- Application data and collected NetFlow in the PostgreSQL database
- Host index data in BadgerDB databases
- Plixer Scrutinizer's third-party encryption key - `/etc/plixer.key`
- Apache Server TLS certificate and key

Important:

- The license key (if the instance is a primary reporter) and the TLS certificates and keys generated by Plixer Scrutinizer are **not** backed up and **cannot be restored**.
 - Any files not included in full backups must be manually backed up and restored, including:
 - Custom threat lists created under `/home/plixer/scrutinizer/files/threats`
 - Custom notifications created under `/home/plixer/scrutinizer/files`
 - LDAP authentication certificates
-

Creating full backups

The Plixer Scrutinizer filesystem includes the `backup.sh` utility, which automates the creation of full backups. This script is located under `home/plixer/scrutinizer/files`.

The following instructions cover the process of creating and saving full Plixer Scrutinizer instance backups to a specified remote host:

1. SSH to the Plixer Scrutinizer server to be backed up and start a `tmux` session to prevent timeouts:

```
tmux new -s backup
```

2. Allow others to use FUSE mounts:

```
sudo grep -Eq "^user_allow_other" /etc/fuse.conf || \  
sudo sed -i '$ a user_allow_other' /etc/fuse.conf
```

3. Create the backup directory locally and mount it to an empty directory on the remote host:

```
BACKUPDIR=/mnt/backup  
sudo mkdir -p $BACKUPDIR  
sudo chown plixer:plixer $BACKUPDIR  
sshfs -o allow_other -o reconnect REMOTE_USER@REMOTE_HOST:REMOTE_  
↵ DIRECTORY $BACKUPDIR
```

Important: Verify that the remote directory to be used is empty and there is sufficient storage available, before running the backup script in the next step. For a rough estimate of the backup file size, run the following on the Plixer Scrutinizer instance:

```
BACKUPDIR=/mnt/backup ~plexer/scrutinizer/files/backup.sh ``df -h /var/
↳db | awk '!/^Filesystem/ {print "Space Required: "$3}'`` on the Plexer
↳Scrutinizer instance.
```

4. Run `backup.sh` as the `plexer` user, with the mounted remote directory set as the backup file location:

```
BACKUPDIR=/mnt/backup ~plexer/scrutinizer/files/backup.sh
```

5. Once the script confirms that the backup file has been saved, unmount the remote backup directory:

```
BACKUPDIR=/mnt/backup
fusermount -u $BACKUPDIR
sudo rmdir $BACKUPDIR
```

Full backup files are created as `scrutinizer-VERSION-backup-DATE.tar.gz` at the specified location and owned by the `plexer` user.

Note:

- A *second Plexer Scrutinizer instance* can be used as the remote backup host, provided it has sufficient disk space available and is running the same Plexer Scrutinizer version as the instance to be backed up. However, doing so is only recommended for redundancy.
- If no remote hosts are available, backups *saved locally*, on the same Plexer Scrutinizer instance. However, due to the size of full backup files, this will limit the amount of storage available for system functions and is not recommended.

For further details or assistance with issues, contact *Plexer Technical Support*.

Backing up additional files

When creating a full backup of a Plexer Scrutinizer server, any files not *covered by the script* must be manually backed up and should be stored on an external host/system.

These files should also be manually restored, after running the *restore script*.

Restoring from a full backup

To restore a Plexier Scrutinizer instance from a full backup file, use the `restore.sh` utility located under `home/plexier/scrutinizer/files`.

The script will fully restore *all backed up elements* of a Plexier Scrutinizer instance, provided the following conditions are met:

- A valid full backup file is accessible by the `plexier` user at the specified (`$BACKUPDIR`) remote location.
- The Plexier Scrutinizer instance to be used for the restore has been freshly deployed.
- The version of the backup matches the version of the fresh Plexier Scrutinizer instance to restore *to* (e.g. a 19.3.0 backup can only be restored to a new 19.3.0 instance).

Important:

- A restore completely overwrites the state of the target instance and deletes the source backup file. It is highly recommended to always restore from a **copy** of a backup file.
 - If the restore target is the primary reporter in a distributed cluster, contact [Plexier Technical Support](#) for assistance.
-

The following instructions cover the process of restoring from a backup file on a remote host to a fresh Plexier Scrutinizer deployment:

1. SSH to the target Plexier Scrutinizer server for the restore, and start a `tmux` session to prevent time-outs:

```
tmux new -s restore
```

2. Allow others to use FUSE mounts:

```
sudo grep -Eq "^user_allow_other" /etc/fuse.conf || \  
sudo sed -i '$ a user_allow_other' /etc/fuse.conf
```

3. Create the backup directory locally and mount the remote directory containing the backup file(s):

```
BACKUPDIR=/mnt/backup
sudo mkdir -p $BACKUPDIR
sudo chown plixer:plixer $BACKUPDIR
sshfs -o allow_other -o reconnect REMOTE_USER@REMOTE_HOST:REMOTE_
↳ DIRECTORY $BACKUPDIR
```

4. Run `restore.sh` as the `plixer` user, with the remote directory set as the backup file location:

```
BACKUPDIR=/mnt/backup ~plixer/scrutinizer/files/restore.sh
```

5. When prompted, enter `yes` to select the backup file to use for the restore or `no` to have the script continue searching (if the backup file was not previously specified).

Hint: To specify the file to use for the restore, use `BACKUPDIR=/mnt/backup`
`BACKUP=restore_filename.tar.gz ~plixer/scrutinizer/files/restore.sh` at the previous step instead.

6. Once the script confirms that the restore has been completed, unmount the remote backup directory:

```
BACKUPDIR=/mnt/backup
fusermount -u $BACKUPDIR
sudo rmdir $BACKUPDIR
```

Important: The `restore.sh` utility does not restart Plixer Scrutinizer services after it completes running.

Based on the role of the Plixer Scrutinizer instance, proceed to finalize setup of the restored server:

- If the restored instance is a **standalone** server, run the following to restart all services and register it:

```
scrut_util --services --name all --switch restart
scrut_util scrut_util --set selfregister --reset
```

These commands may take several minutes to complete.

- If the restored instance is a **remote collector** in a distributed cluster, run the following on the primary reporter to register it:

```
scrut_util --set registercollector --ip [IP address of restored  
→instance]
```

- If the restored instance is a primary reporter in a distributed cluster or a standalone server, and its Machine ID is different from that of the backup file, contact [Plixer Technical Support](#) to obtain a new license key.

Alternative backup methods

Because full backup files are extremely large and intended for use in disaster recovery scenarios, saving and storing backup files to remote hosts serving ssh is highly recommended.

In scenarios where this is not possible, the following alternative backup methods can be used:

Backup to a second Plixer Scrutinizer instance

If a separate host is not available to save backups to, a second Plixer Scrutinizer instance can be used for backup file storage instead. The versions of the two instances must match.

Important: Due to how Plixer Scrutinizer is designed to optimize the use of all available disk space, it will likely be necessary to add more storage and/or modify the data retention settings of the second instance. For assistance, contact [Plixer Technical Support](#).

The following instructions cover the additional steps required for creating backups on a second Plixer Scrutinizer instance (using the default location):

1. Set the location/directory to use for backup files:

```
#!/bin/bash  
BACKUPDIR=${BACKUPDIR:='/var/db/big/pgsql/restore'}  
REMOTE=YOUR_REMOTE_SCRUTINIZER_INSTANCE
```

2. Create the backup directory on both instances:

```
sudo mkdir -p $BACKUPDIR
sudo chown plixer:plixer $BACKUPDIR
ssh plixer@$REMOTE "sudo su -c 'mkdir -p $BACKUPDIR && chown_
↳ plixer:plixer $BACKUPDIR'"
```

3. Allow other users to use FUSE mounts:

```
sudo grep -Eq "^user_allow_other" /etc/fuse.conf || \
sudo sed -i '$ a user_allow_other' /etc/fuse.conf
```

4. Mount the remote instance's backup directory on the local instance:

```
sshfs -o allow_other -o reconnect plixer@$REMOTE:$BACKUPDIR $BACKUPDIR
```

5. Run `backup.sh` using the directory mounted from the remote instance as the backup file location:

```
BACKUPDIR=/var/db/big/pgsql/restore ~plixer/scrutinizer/files/backup.sh
```

Important: Before running the backup utility, verify that the remote directory to be used is empty and there is sufficient storage available. For a rough estimate of the backup file size, run the command `df -h /var/db | awk '!/^Filesystem/ {print "Space Required: "$3}'` on the Plixer Scrutinizer instance.

6. After the backup is complete, unmount the remote Plixer Scrutinizer directory:

```
#!/bin/bash
BACKUPDIR=${BACKUPDIR:='/var/db/big/pgsql/restore'}
fusermount -u $BACKUPDIR
```

Local backups

By default, both `backup.sh` and `restore.sh` are set to use `/var/db/big/pgsql/restore` on the local Plixer Scrutinizer filesystem for full backup files. However, in most cases, the backup operation will likely fail unless additional disk space is allocated to or created on the Plixer Scrutinizer instance. Running

the command `df -h /var/db | awk '!/^Filesystem/ {print "Space Required: "$3}'` will provide a rough estimate of the storage required for the backup.

To force a checkpoint, enter `psql plixer -c "CHECKPOINT"` after the script has finished running.

Note:

- In v19.2, the backup file path must be defined in the `backup.sh` and `restore.sh` scripts before they are run.
 - Storing backup files locally will severely limit the storage Plexier Scrutinizer can use for its primary functions. As such, backup files saved to the instance should be transferred to a separate resource as soon as possible.
-

6.6.2 Configuration backups

For more “lightweight” backup and restore operations, the `scrut_conf_dump.sh` and `scrut_conf_restore.sh` scripts (both located in `/home/plixer/scrutinizer/database/utils`) can be used to target only the application/configuration data of a Plexier Scrutinizer instance, including:

- User-added maps
- Dashboards
- IP groups
- Saved reports
- 3rd-party integration settings

Configuration backups do not include any collected flow data.

Note: In distributed environments, the primary reporter regularly syncs application/configuration data to remote collectors. Only the configuration backup of the primary reporter is needed to perform a restore for the cluster.

`scrut_conf_dump.sh` and `scrut_conf_restore.sh` use Postgres’s `pg_dump` and `pg_restore` utils and respect the same set of environment variables:

| Variable | Description | Default |
|------------|---|-------------|
| DUMP | Location of the backup file | ./conf.dump |
| PGHOST | IP address or hostname of the PostgreSQL database | localhost |
| PGUSER | Role/user used to connect to PGHOST | plexer |
| PGDATABASE | The database to access at PGHOST | plexer |

Note: To avoid potential issues, use the [services command](#) to stop the `plexer_flow_collector` service before attempting to restore from a configuration backup.

Backing up configuration data

To create a backup of a Plexier Scrutinizer server's current configuration data, run the following script:

```
~/scrutinizer/database/utils/scrut_conf_dump.sh
```

To dump configuration data to a specified location/file:

```
mkdir /tmp/CONF_BACKUP_DIR
touch /tmp/CONF_BACKUP_DIR/CONF_BACKUP.dump
DUMP=/tmp/CONF_BACKUP_DIR/CONF_BACKUP.dump
~/scrutinizer/database/utils/scrut_conf_dump.sh
```

Restoring configuration data

To restore configuration data to a Plexier Scrutinizer server from a backup file, run the following script:

```
~/scrutinizer/database/utils/scrut_conf_restore.sh
```

To restore configuration data from a specified dump file:

```
PGHOST=42.42.42.42
DUMP=/tmp/CONF_BACKUP_DIR/CONF_BACKUP.dump
~/scrutinizer/database/utils/scrut_conf_restore.sh
```

Important: The `scrut_conf_restore.sh` script is meant to be used for restoring configuration data for the same Plexier Scrutinizer server/appliance. To apply a configuration backup to a different server, follow the steps for [backup migrations](#) in the [migration guides](#).

Additional notes

- `pg_restore` errors typically only cause the restore to fail for the table associated with the error. Other tables should still be restored successfully.
- Errors associated with **duplicate keys** usually indicate a conflict between existing rows in the table and the rows being restored.

```
pg_restore: [archiver (db)] Error from TOC entry 51348; 0 17943 TABLE
↳DATA exporters plixer
pg_restore: [archiver (db)] COPY failed for table "exporters": ERROR:
↳duplicate key value violates unique constraint "exporters_pkey"
DETAIL:  Key (exporter_id)=(\x0a4d4d0a) already exists.
```

The conflicting keys should be removed from the table before attempting to restore again.

- If you are swapping IP addresses, the database keys should be rotated using `scrut_util --pgcerts --verbose`, because the backed up keys will be associated with the old address.

6.6.3 Certificate utilities

Several utilities are bundled with Plexier Scrutinizer to help manage the TLS certificates used by the system.

Note: These scripts rely on Plexier Scrutinizer's default ssh connectivity.

`generate_requests.sh`

This script generates certificate requests from all TLS keys in a distributed Plexier Scrutinizer cluster. It should be run on the cluster's primary reporter as the `plexier` user. Certificate details can be set via the script's variables.

```
~/scrutinizer/files/generate_requests.sh
```

All certificate requests are placed in `/tmp/request`. `/tmp/request/apache_server.csr` is the certificate request for the primary reporter's web server, and requests from the rest of the cluster are organized in subdirectories.

install_certs.sh

This script installs signed TLS certificates across a distributed Plexier Scrutinizer cluster. It should be run on the cluster's primary reporter as the `plexier` user.

```
~/scrutinizer/files/install_certs.sh
```

`.cer` files are expected at `/tmp/signed` and should follow the filename conventions used by `generate_requests.sh`. `/tmp/signed/ca.cer` should be the Certificate Authority's root certificate.

scrut_util --rotatecerts --reset

This [scrut_util command](#) automatically resets and restores database certificates. It can be used if either of the former scripts causes unexpected issues or when DB connection issues are observed.

```
scrut_util --rotatecerts --reset
```

This will regenerate all TLS keys and certificates in a distributed cluster and should restore normal operations, at the expense of any existing signed certificates.

For further information or assistance, contact [Plexier Technical Support](#).

6.6.4 Migration

In addition to the `restore.sh` script for [full backups](#), Plexier Scrutinizer also includes a [scrut_util](#) migration utility that can be used to migrate specific subsets of an existing environment to a new server/deployment.

Important:

- Because the steps outlined here can potentially result in an irrecoverable state for the source and/or destination Plixer Scrutinizer server, it is highly recommended to contact [Plixer Technical Support](#) for assistance.
- The migration utility saves logs to `/var/log/migrate.log`. The Plixer support engineer will require these logs to better assist with any issues/errors encountered during the process.
- To view additional information on the `--migrate` command, run the following:

```
/tmp/scrut_util --help migrate
```

The migration utility supports three types of migration:

- [Configuration migrations](#)
- [Data migrations](#)
- [Backup migrations](#)

Pre-migration requirements

Before any type of migration, the source and destination servers must both be prepared for the process.

Preparing the source server

1. On the source server, create a temporary role for migrations in the Plixer Scrutinizer database:

```
psql -c "CREATE USER migrator WITH SUPERUSER ENCRYPTED PASSWORD  
→ '[PASSWORD]'"
```

2. Allow migration connections from the destination's IP address by running the following on the source server:

```
DESTINATION=[DESTINATION_IP]  
sudo sed -i -e "lihost plixer migrator $DESTINATION/32 md5" /var/db/big/  
→pgsql/data/pg_hba.conf  
psql -c "SELECT pg_reload_conf()"
```

Important: For source servers on 18.20, modify the above `psql` commands to use the database user (i.e., `psql -Uroot plixer -c...`)

Preparing the destination server

The destination Plixer Scrutinizer server must be fully deployed and licensed prior to the migration. It should also already be collecting flows.

On the destination server, update `/etc/migrate.ini` with the details of the source server and the temporary migration role.

Configuration migrations

Configuration migrations are used to transfer application configuration and user preferences on the source Plixer Scrutinizer server to the destination. Configuration migrations are only supported by source servers on **v19.4.0** and above.

To initiate a configuration migration, run the following on the destination server:

```
SOURCE=[SOURCE_IP]
scp $SOURCE:/etc/plixer.key /etc/plixer.key
sudo systemctl stop plixer_flow_collector
scrut_util --migrate config
sudo systemctl start plixer_flow_collector
```

Additional notes

- By default, the distributed environment configuration of the destination server is retained during configuration migrations. To also migrate the distributed environment settings from the source to the destination (and still preserve the latter's IP address and machine ID), add the `--force_dist` option as follows:

```
scrut_util --migrate config --force_dist
```

- While configuration migrations include the Plixer Scrutinizer application configuration, settings that must be applied to the hardware or virtual appliance (e.g., custom listening ports) will not be migrated. These settings must be removed and re-applied in Plixer Scrutinizer after the migration.
- Because the migration utility relies on Plixer Scrutinizer's [configuration backup and restore scripts](#), any additions to the factory SQL codebase are not included in migrations by default. To migrate such modifications, update the `scrut_conf_dump.sh` to include them.

Data migrations

Data migrations are used to transfer all flow data still retained on the source Plixer Scrutinizer server to the destination. Data migrations are supported by source servers on **v18.2.0** and above.

Note: To maintain reporting continuity after the migration, the destination server should be collecting flows from the same exporters as the source.

To initiate a data migration, run the following on the destination server:

```
scrut_util --migrate data
```

Confirm to begin the migration when prompted. The migration may take some time.

Backup migrations

Backup migrations are used to apply application configuration and user preferences to the destination Plixer Scrutinizer server from an existing *configuration backup*. This type of migration is only supported between **a source (of the backup) and destination on the same version**.

Note: Backup migrations should only be used to transfer configuration data to a different destination appliance/server. To restore from a backup file to the same server, use the *restore script* instead.

To initiate a backup migration, run the following on the destination server:

```
scrut_util --migrate restore
```

This will restore/migrate `/tmp/conf.dump` to the destination server.

Restoring from a specified backup file

To use a different backup file, add the `--dump_file` option as follows:

```
scrut_util --migrate restore --dump_file [BACKUP_FILE]
```

Post-migration

After any type of migration, the hba rule and temporary migrator role must be removed from the source Plixer Scrutinizer server.

To do so, run the following:

```
sudo sed -i '/migrator/d' /var/db/big/pgsql/data/pg_hba.conf
psql -c "SELECT pg_reload_conf()"
psql -c "DROP ROLE migrator"
```

6.7 Updates and upgrades

To ensure a consistently feature-rich and secure experience, all supported versions of Plixer Scrutinizer will continuously be updated. When installed, update packages may add new features, improve existing functionality, and/or apply patches for emerging security threats. All update packages will have been applied to Plixer's own QA servers and extensively tested before they are made available.

This section provides details on the different types of update packages that may be released and includes instructions for their installation.

Important: While it is possible to install Plixer Scrutinizer update packages without assistance, it is highly recommended to contact [Plixer Technical Support](#) and allow our engineers to guide you through the process.

6.7.1 Update preparations

Before attempting to install any type of update package, the following procedures should be observed:

1. Verify that the version currently installed can be upgraded to the target version (e.g., v18.20 or v19.x -> v19.4.0).
2. Back up the current install:
 - Virtual appliances: Take a snapshot, ideally with the appliance powered off.

- Hardware appliances: Perform a [full](#) or [configuration](#) backup. For further details, see the [Backups](#) subsection of this documentation or contact [Plixer Technical Support](#).
3. **Hardware appliances only** - Log in to iDRAC and perform a hardware health check. Any hardware issues discovered should be escalated to Dell for resolution. A reboot is also recommended as an additional check for underlying hardware issues.
 4. Confirm that all Plixer Scrutinizer Collectors/servers have access to <https://files.plixer.com>. This check can be performed by downloading the checksum file using the following command:

```
curl -o scrutinizer-checksums.txt -L https://files.plixer.com/plixer-  
repo/scrutinizer/<version_number>/scrutinizer-checksums.txt
```

For Plixer Scrutinizer deployments that do not have Internet access, follow the steps to perform [offline updates](#).

5. Collect the following details and check the [Plixer Scrutinizer sizing guide](#) to confirm that sufficient resources will be available to the system after the upgrade:
 - Flows per second
 - Number of active Exporters
 - CPU (number of cores, clock speeds)
 - Amount of RAM
 - Disk speed and RAID type
 - Flow Analytics algorithms enabled
6. Obtain a valid license key for the upgrade if one has not been acquired.
7. Delete any older versions of `scrutinizer-installer.run` on the Plixer Scrutinizer instance. This will prevent them from being used instead of the correct installer.
8. Enter `crontab -e` and inspect the table for lines containing `* * * * * /home/Plixer/scrutinizer/files/collector_restart.sh`. These should be commented out by adding a `#` at the beginning of the line to prevent scheduled restarts from interfering with the upgrade process.
9. **Distributed cluster upgrades only** - If there are Palo Alto firewalls configured for the cluster, whitelist the connections between the Reporter and the Collectors. This will prevent the firewall from identifying the ~113 SSH connections created during the Collector registration process as a threat. Alternatively, the rate at which the SSH connections are established can be slowed down by adding `sleep 5` to the `/home/plixer/.bashrc` file on each remote Collector.

10. **AWS flow log integration only** - As of version 19.2, Plixer Scrutinizer requires four log fields to be configured for AWS flow log collection: `log-status`, `vpc-id`, `interface-id`, and `flow-direction`. For further details, see the [AWS flow log integration guide](#).

These steps are meant to identify and resolve any underlying issues with the current Plixer Scrutinizer install and help ensure that upgrade will be applied without issue.

Once completed, follow the [appropriate upgrade guide](#) to update Plixer Scrutinizer to the latest version.

Hint: All install logs will be saved to `/var/log/Scrutinizer-Install.log`.

Note: As of v19.1+ Plixer Scrutinizer no longer requires the use of the root OS user, and the `plixer` user is the recommended user for command line access.

6.7.2 Version upgrades

Version upgrades update Plixer Scrutinizer to the latest major or minor release (e.g., 19.4) and include significant improvements over the previous version. These updates may include additional functionality, performance enhancements, and/or QoL improvements, in addition to implementing fixes for certain types of issues.

Upgrading to v19.5.2

Because Plixer Scrutinizer v19.5.2 uses a new operating system (OS), a v19.4.0 install is **required** for the upgrade. If you are running an older version, follow [this guide to upgrade your server/appliance to v19.4.0](#).

For assistance or clarifications, contact [Plixer Technical Support](#).

Note:

- The upgrade will take **at least one hour to complete**.
 - The `plixer` user SSH password will be needed during the upgrade. If necessary, it can be reset when the OS upgrade script is run.
 - If root login is enabled on the Plixer Scrutinizer server, it will be disabled as part of the upgrade.
 - If the Plixer Scrutinizer server is able to access `files.plixer.com`, the `REPO_HOST` variable should be set to `files.plixer.com` for the steps outlined below. For [offline upgrades](#), the IP address of the offline repo should be used instead.
-

Upgrade process

The process of upgrading a v19.4.0 Plixer Scrutinizer server to v19.5.2 involves the following steps:

- [Backing up the v19.4.0 database and server-specific files](#)
- Downloading the operating system upgrade script, `olmigrate.run`, and running it a total of four times (with a reboot between runs)
- Downloading and running the Plixer Scrutinizer v19.5.2 installation script (`scrutinizer-install.run`)
- Verifying that the v19.4.0 data has been successfully migrated after v19.5.2 is installed

Pre-upgrade preparation

- [Hardware appliances] Create a [full backup](#) of the current Plixer Scrutinizer install and store it on an external system/drive.
- [Virtual appliances] Back up the current Plixer Scrutinizer install by taking a VM snapshot.
- Review the [general upgrade preparation guide](#) and complete any steps that apply.
- [Offline upgrades] If the Plixer Scrutinizer server does not have access to `files.plixer.com`, [set up an offline repository for this upgrade](#).

OS migration

Once all preparation steps have been completed, follow these steps to migrate the v19.4.0 appliance to the new operating system:

Important:

- For offline upgrades, `REPO_HOST` should point to the IP address of the [offline repo](#) instead of `files.plixer.com`.
- In distributed clusters, complete the upgrade for all remote collectors before upgrading the primary reporter.
- To verify the current progress of the OS upgrade at any time:

```
cat /etc/motd
```

or check versions between runs (NAME= and VERSION= lines):

```
cat /etc/os-release
```

- If any errors are encountered during the upgrade process, run the following to collect log files:

```
sudo tar -czf /tmp/olmigrate_logs.tar.gz /var/log/olmigration/ /var/log/  
↳leapp/ /var/log/messages /var/log/Scrutinizer-Install.log
```

Afterwards, move `/tmp/olmigrate_logs.tar.gz` off the server before reverting. [Plixer Technical Support](#) will require the logs to better assist you with any issues.

1. SSH to the v19.4.0 server to be upgraded as the `plixer` user.
2. Verify that the current working directory is correct (`plixer`):

```
cd /home/plixer/
```

3. Download the OS upgrade script and its checksum file:

```
REPO_HOST=files.plixer.com  
curl -k -o olmigrate.run https://$REPO_HOST/plixer-repo/scrutinizer/19.5.  
↳2/olmigrate.run  
curl -k -o olmigrate.run.sha256 https://$REPO_HOST/plixer-repo/  
↳scrutinizer/19.5.2/olmigrate.run.sha256
```

4. Validate the integrity of `olmigrate.run`:

```
sha256sum -c olmigrate.run.sha256
```

5. Update permissions for the OS upgrade script:

```
chmod a+x olmigrate.run
```

6. Run the `olmigrate.run` script a total of four times:


```
REPO_HOST=files.plixer.com ./olmigrate.run -- -k
```

Important: Reboots between runs of the OS upgrade script (`olmigrate.run`) can take a long time. Before trying to reconnect to the server, start a PING to the Plexier Scrutinizer IP address and wait for it to become available again. **Do NOT manually reboot the server.**

After the fourth `olmigrate.run` run (there will be no reboot), the OS migration will be complete.

Upgrading to Plexier Scrutinizer 19.5.2

Once the appliance is on the *new OS*, Plexier Scrutinizer can be upgraded to v19.5.2 as follows:

1. Change directories to `/tmp`:

```
cd /tmp/
```

2. Download the Plexier Scrutinizer v19.5.2 installation script and its checksum file:

```
REPO_HOST=files.plixer.com
curl -k -o scrutinizer-install.run https://$REPO_HOST/plixer-repo/
↪scrutinizer/19.5.2/scrutinizer-install.run
curl -k -o scrutinizer-install.run.sha256 https://$REPO_HOST/plixer-repo/
↪scrutinizer/19.5.2/scrutinizer-install.run.sha256
```

3. Validate the integrity of `scrutinizer-install.run`:

```
sha256sum -c scrutinizer-install.run.sha256
```

4. Update permissions for the installation script:

```
chmod a+x scrutinizer-install.run
```

5. Run `scrutinizer-install.run` to begin the upgrade to Plexier Scrutinizer v19.5.2:

```
REPO_HOST=files.plixer.com ./scrutinizer-install.run -- -k
```

6. After the installation script finishes running, reboot the appliance:

```
sudo shutdown -r now
```

7. After the reboot, run the following commands to verify that the system is in working order:

```
scrut_util --check heartbeat --type database  
scrut_util --check heartbeat --type api
```

Important: For distributed environments, the heartbeat checks should only be run on remote collectors *after the primary reporter has been upgraded, and the the cluster has been reestablished.*

If the heartbeat checks are successful, then the Plixer Scrutinizer appliance has been successfully upgraded to v19.5.2.

Offline upgrades to v19.5.2

The following instructions for setting up an offline repo are intended for upgrading to Plixer Scrutinizer v19.5.2 only.

1. *Deploy a new Plixer Scrutinizer v19.4.0 VM* and *assign an IP address to it*.
2. SSH to the VM as the `plixer` user:

```
ssh plixer@SCRUTINIZER_VM_IP
```

3. Create the offline repo directory and assign it the correct permissions:

```
sudo mkdir /var/db/big/offline  
sudo chown plixer:plixer /var/db/big/offline
```

4. Download the offline tar file for 19.5.2 and its checksum file:

```
curl -o /var/db/big/offline/19.5.2_offline.tgz https://files.plixer.com/  
↳plixer-repo/scrutinizer/19.5.2_offline.tgz  
curl -o /var/db/big/offline/19.5.2_offline.tgz.sha256 https://files.  
↳plixer.com/plixer-repo/scrutinizer/19.5.2_offline.tgz.sha256
```

5. Validate the integrity of 19.5.2_offline.tgz:

```
sha256sum -c /var/db/big/offline/19.5.2_offline.tgz.sha256
```

6. Extract the offline tar file:

```
cd /var/db/big/offline  
tar xvf 19.5.2_offline.tgz
```

7. Create a symlink in the html directory to the offline repo:

```
ln -s /var/db/big/offline/plixer-repo /home/plixer/scrutinizer/html/  
↳plixer-repo
```

After the offline repo has been set up, the VM's IP address should be used in place of `files.plixer.com` for `REPO_HOST` in the [upgrade instructions](#).

Upgrading to v19.4.0

To upgrade to Plixer Scrutinizer v19.4.0, a v18.20 or v19.x install is required.

The following instructions cover the upgrade process for both standalone and distributed environments:

1. Perform a [backup](#) of the current Plixer Scrutinizer install as described in the [recommended upgrade preparation steps](#).
2. SSH to the primary Reporter and start a new tmux session.

```
tmux new -s upgrade
```

3. Download the v19.4.0 installer.

```
cd /tmp
curl -o scrutinizer-install.run https://files.plixer.com/plixer-repo/
↪scrutinizer/19.4.0/scrutinizer-install.run
```

4. Download the checksum file and validate the integrity of the `scrutinizer-install.run` file.

```
curl -o scrutinizer-checksums.txt https://files.plixer.com/plixer-repo/
↪scrutinizer/19.4.0/scrutinizer-checksums.txt
cat scrutinizer-checksums.txt
sha256sum scrutinizer-install.run
```

5. Set the correct permissions for the installer.

```
sudo chmod 755 scrutinizer-install.run
```

Important: For AMI cluster upgrades, it is also necessary to run `chmod 755 /home/ec2-user` before starting the installer. This will ensure that the SSH key used for communication between Plixer Scrutinizer Collectors has the correct permissions. Contact [Plixer Technical Support](#) for assistance.

6. Run `scrutinizer-install.run` on the primary Reporter as the `root`, `plixer` (recommended), or `ec2-user` user.

```
./scrutinizer-install.run
```

Note: If the server was previously upgraded from 18.20, you will be asked whether to delete the `data.old` backup that was created during that upgrade. Because a new backup should have been created prior to starting the current upgrade process, `data.old` can safely be deleted.

7. **Distributed cluster upgrades only** - When asked how the installer should log in to remote Collectors in the cluster, enter either **existing** (recommended) or **passwords**.

Important: When upgrading a distributed cluster from v18.20, **passwords** should be used. Prior to this step, the installer will also prompt the user to create a new *plixer control key*, which should be left blank unless encrypted keys are required.

8. After the installer has finished running, the following heartbeat checks should be run to verify that all nodes (including standalone deployments) are able to communicate normally:

```
scrut_util --check heartbeat --type database
scrut_util --check heartbeat --type api
```

If all heartbeat checks are successful, then the systems have been upgraded successfully.

Offline upgrades

To upgrade Plixer Scrutinizer collectors/servers that do not have access to <https://files.plixer.com>, an offline repository can be created on the primary reporter. For assistance, contact [Plixer Technical Support](#).

To set up the offline repository and start the upgrade process, follow these steps:

Important: In the following instructions, VERSION should be replaced with the full target version number (e.g., when upgrading to v19.4.0, replace VERSION with 19.4.0).

1. Perform a [backup](#) of the current Plixer Scrutinizer install as described in the [recommended upgrade preparation steps](#).
2. Download https://files.plixer.com/plixer-repo/scrutinizer/VERSION_offline.tgz to a computer with Internet access.
3. Download <https://files.plixer.com/plixer-repo/scrutinizer/VERSION-checksum-offline.txt> and validate the checksum of the *.tar* file.
4. Enter `df -h` and confirm that the primary reporter has at least 84 GB of free disk space under `/var/db/big`.
5. Start an SSH session with the primary reporter as the `plixer` user.
6. Create a new directory for the offline installation files and set the correct permissions to give the `plixer` user access to it.

```
sudo mkdir /var/db/big/offline
sudo chown plixer:plixer /var/db/big/offline
```

7. Move the *.tar* file to the appropriate directory on the primary Reporter.

```
scp VERSION_offline.tgz plixer@x.x.x.x:/var/db/big/offline/VERSION_
↳offline.tgz
```

8. Extract the contents of the file to the same directory.

```
sudo tar -zxvf /var/db/big/offline/VERSION_offline.tgz -C /var/db/big/
↳offline
```

Important: Extracting the contents of the *.tar* file will exhaust all available space on the appliance.

9. Create a *symlink* to the offline directory from the html directory, so the files can be served by Apache.

```
sudo ln -s /var/db/big/offline/plixer-repo /home/plixer/scrutinizer/html/
↳plixer-repo
```

10. From the appliance to be upgraded, download the installer from the offline repository.

```
curl -o scrutinizer-install.run -L -k https://127.0.0.1/plixer-repo/
↳scrutinizer/VERSION/scrutinizer-install.run
```

11. Download the checksum file and validate the integrity of the `scrutinizer-install.run` file.

```
curl -o scrutinizer-checksums.txt -k https://127.0.0.1/plixer-repo/
↳scrutinizer/VERSION/scrutinizer-checksums.txt

cat scrutinizer-checksums.txt

sha256sum scrutinizer-install.run
```

12. Set the correct permissions for the installer.

```
sudo chmod 755 scrutinizer-install.run
```

13. Run the `scrutinizer-install.run` file as the `root`, `plixer` (recommended), or `ec2-user` user. `x.x.x.x` should be replaced with the IP address or hostname of the offline repository host.

```
REPO_HOST=local LOCAL_REPO_BASEDIR=/var/db/big/offline ./
↪scrutinizer-install.run -- -k
```

Important: For AMI cluster upgrades, it is also necessary to run `chmod 755 /home/ec2-user` before starting the installer. This will ensure that the SSH key used for communication between Plixer Scrutinizer Collectors has the correct permissions. Contact [Plixer Technical Support](#) for assistance.

Note: If the server was previously upgraded from 18.20, you will be asked whether to delete the `data.old` backup that was created during that upgrade. Because a new backup should have been created prior to starting the current upgrade process, `data.old` can safely be deleted.

14. **Distributed cluster upgrades only** - When asked how the installer should log in to remote collectors in the cluster, enter either **existing** (recommended) or **passwords**.

Important: When upgrading a distributed cluster from v18.20, passwords should be used. Prior to this step, the installer will also prompt the user to create a new *plixer control key*, which should be left blank unless encrypted keys are required.

15. After the installer has finished running, the following heartbeat checks should be run to verify that all nodes (including standalone deployments) are able to communicate normally:

```
scrut_util --check heartbeat --type database
scrut_util --check heartbeat --type api
```

If all heartbeat checks are successful, then the systems have been upgraded successfully.

6.7.3 General and CVE patches

From time to time, customers may be notified that general and/or CVE patches are available for the Plixer Scrutinizer version they are currently running. These patches typically address noncritical system issues and/or improve protections against new security threats.

Note: General and CVE patches do not increment the Plixer Scrutinizer version number.

To apply these updates, follow the [version upgrade instructions](#) to download and run the latest installer for the current Plixer Scrutinizer version. Going through the [standard update preparations](#) is also highly recommended.

When run, the installer will automatically download and apply all available patches.

6.7.4 Verifying vulnerability patches

Some vulnerability scanning and auditing solutions may report vulnerabilities that have already been patched in the most recent update. This is typically the combined result of a backported security patch and the tool only scanning for component version numbers.

If this happens, there are two ways to verify the validity of the vulnerability report:

- Check the package changelog for the CVE identifier/number of the vulnerability (e.g., CVE-2017-3169)
- Download and install the latest OVAL Definitions from oval.cisecurity.org/repository, which will allow any compatible tools to determine the status of vulnerabilities, even when security patches have been backported.

For additional assistance, contact [Plixer Technical Support](#).

ADDITIONAL RESOURCES

This section includes additional resources and materials relevant to the use of Plixer Scrutinizer and this user manual.

7.1 Changelog

KEY: Description (Ticket Number)

Ex. Thresholds based on outbound traffic (1640)

Please reference our End of Life Policy for details regarding the end of life schedule. For more information on Plixer Scrutinizer, please reference the online documentation or visit our website.

7.1.1 Version 19.5.2 - July 2024

Note:

- This release addresses CentOS going EOL. To migrate the OS to Oracle Linux 9, Plixer Scrutinizer must be on version 19.4.0. Please contact [Plixer Technical Support](#) with any questions.
 - AWS instances of Plixer Scrutinizer use Amazon Linux 2 and do not need to be updated to 19.5.2. A later release, which will include new features and bug fixes, will be made available for Plixer Scrutinizer deployments on AWS.
-

Plixer Scrutinizer

New Features

- Proxmox support

Fixes

- Addressed various security issues
- Memory leak (4363)
- Missing AS and country names (4353)
- SNMP polling issue (4364)
- Data migration fails when destination expires history (4364)

Plixer ML Engine (Version 19.4.0 - August 2024)

New Features

- User behavior analytics for O365/Azure AD
- Encrypted Traffic Analytics (ETA)
- New base operating system: Ubuntu
- LDAP rogue service detection
- Support for Hyper-V deployments
- New alert type in Plixer Scrutinizer for Suricata TLS alerts
- Support for IP groups in exclusions

Fixes

- Behavior tab missing after upgrading to 19.3 (745)
- Time zone set as UTC causing training data to not be associated correctly with workdays/weeknights/weekends (750)
- An ML engine can now pair with a multi-collector distributed Plixer Scrutinizer (757)
- Certain utilities not functioning properly with 19.3 (786)
- Improved ML event messages by including additional details (810)

7.1.2 Version 19.5.1 - June 2024

Note:

- This release addresses CentOS going EOL. To migrate the OS to Oracle Linux 9, Plixer Scrutinizer must be on version 19.4.0. Please contact [Plixer Technical Support](#) with any questions.
-

Plixer Scrutinizer

New Features

- Check for supported CPU architecture in olmigrate
- Automatic disabling of root login in olmigrate
- Check for multiple interfaces in olmigrate

Fixes

- Addressed an issue where a recursive directory is created if olmigrate is run more than once for the same upgrade stage

7.1.3 Version 19.5.0 - May 2024

Note:

- This release addresses CentOS going EOL. To migrate the OS to Oracle Linux 9, Plixer Scrutinizer must be on version 19.4.0. Please contact [Plixer Technical Support](#) with any questions.
-

Plixer Scrutinizer

New Features

- Oracle Linux v9.4
- System Migration Utility

Fixes

- Addressed various security issues
- Filtering issue with NSG FlowLogs (4225)
- Slow load times for Admin > Manage Exporters (4080)
- Primary server being down prevented collector services from restarting (4147)

7.1.4 Version 19.4.0 - October 2023

Plixer Scrutinizer

New Features

- AWS Flowlog consumption 35x faster
- AWS Flowlog consumption and processing can be spread across multiple collectors
- Azure flow log ingestion
- Azure NSG Reports
- Security Groups for enabling groups of Exporters in Flow Analytics
- Userpreferences Modifiable Template
- Include Custom Designed Reports in Scrutinizer Configuration Backup
- Support 18.20 -> 19.X offline upgrades where the repo server is the Scrutinizer server
- sFlow vlan/sub-interface report
- Merged target and violator alarm views into consolidated hosts view

- Host entity alarm timeline view
- Endpoint Analytics Risk and details into Alarm Monitor views
- Multiple new Alarm Monitor visualizations
- Connections graph type in reporting
- New Admin interfaces
- Default Flow Analytics Exclusion Groups under IP Groups
- Include port name in DrDoS alarm messages
- Support for FlowPro version 20

Fixes

- Addressed various security issues
- On demand PDF/email/csv use server time zone when they should use user time zone (1069)
- Optimize TCP/UDP FA algorithms (2695)
- Turning SSL off breaks the UI (2728)
- Double quotes in SSL serverprefs (2927)
- Distributed upgrades should have collectors run a curl check for Internet access (2958)
- Exporters Not Deleting with Domain Exclusions (3110)
- Event severity timeline (3329)
- Editing FA host exclusions doesn't update caches (3332)
- Distributed Upgrade Installer handle proxy configuration prompt (3339)
- System Performance View shows red when resources exceed the matrix (3351)
- Implement sFlow version 4 (3357)
- Filter all FA sliding windows by streamexporter (3377)
- set myaddress fails on Hardware appliances (3386)
- CSV export column header shifted by one position for Connection reports (3400)
- Reporting - Source / Destination Port EXCLUDE Port Range - Error: "report failed" (3402)
- Setting timezone can pause alarms (3405)
- Issue with units label for application latency report threshold messages (3471)
- Added paged requests to LDAP authentication to handle large lists of Active Directory Security Groups (3481)
- Fix overstated utilization when sFlow counters are dropped (3485)
- Optimize Explore By Exporters view (3541)
- Clean history table orphans in batches (3555)
- RADIUS shared secret needed to be re-entered after v19.3 upgrade (3591)

- scrut_util ‘set ssl on/off’ requires root - but should not be run as root (3624)
- Escape special characters in interface details (3636)
- Fix a logs-based disk space leak (3667)
- Store AWS interface in the aws_interface element - don’t map to ingressinterface (3687)
- Optimize FlowPro FA algorithms (3695)
- Optimize Packet Flood FA algorithm (3699)
- Optimize Slow Port Scan Algorithm (3700)
- Legacy Baselineing is now EOL (3704)
- Made UDP receive buffers configurable (3711)
- Mixing include and exclude advanced filters could restrict more results than necessary (3757)
- Don’t allow “Host Index Max Disk Space” setting to exceed available disk space (3779)
- Manage Exporters and Manage Collectors were removed from the classic Admin UI (3808)
- Monitor.top_stdout Parsing Errors (3828)
- AWS Upgrade package dependency problem (3862)
- scrut_util check heartbeat database as root user error (3873)
- Move slog directory out from under html (3882)
- No packet or octet values for exporter sending samplingpacketspace of 0 (3903)
- distributed_stats_exporters wasn’t being cleaned out (3931)

Plixer Scrutinizer UI

Fixes

- Reports: Restructure to allow proper placement of app-page-toolbar and tray (1001)
- Dashboards: Too much air in vitals (1054)
- Dashboard Recent Alarms gadget is out of sync with current alarms (1114)
- Alarms: Show DNS & IP information in messages (1252)
- Acknowledged Alarms View Doesn’t auto-refresh (1417)
- Explore Event Traffic links do not respect PlixCal filter (1441)
- Exported CSV from Entities page displays host names with ‘Show Host Names’ deselected. (1463)
- Spatial Map: Modified timestamp gets wrongly updated to all the existing maps (1498)
- Explore>Entities: “dbQueryError” seen in console when applying filters (1574)
- Admin: Default Status, Tab & View (1591)
- Default map defined for user does not open when accessing Network Maps in new UI (1598)
- Change Endpoint “Identity Score” to “Profile Match” (1617)
- Reporting: Phantom selected select box (1712)
- Issue with displaying child groups with a parent group filter (1877)

7.1.5 Version 19.3.2 - September 2023

Plixer Scrutinizer

Fixes

- Addressed various security issues
- AWS Upgrade package dependency problem (3826)

7.1.6 Version 19.3.1 - April 2023

Plixer Scrutinizer

Fixes

- Addressed various security issues
- AWS interface IDs no longer used as observation domain (3568)
- Deleting collector log wouldn't always return diskpace (3667)
- Reduced output to logfile for Feature Resources (3675)
- Optimized query for Explore exporter view (3684)
- Upgrades needed a forced reboot for chromium (3692)
- Update LDAP login to get the *defaultRoute* preference (3697)
- Changed default view for Explore to Top Interfaces (3703)

7.1.7 Version 19.3.0 - December 2022

Plixer Scrutinizer

New Features

- MITRE ATT&CK Visualization
- MITRE ATT&CK details for notification profiles
- Support for using hostname when configuring an ML Engine
- Support for redirecting to a proxy address after Single Sign-On
- LRFM: No audit trail from manual enable/disable
- sFlow: Add support for VLAN tags in sampled Ethernet headers
- sFlow: Support for sampled IPv6 headers
- Ability to pass custom parameters when opening ServiceNow issues

Fixes

- Moloch Integration Link not clickable in the new UI (1035)
- Admin Tab permission is required to logout (1269)
- Report selection stuck open without selection (1372)
- Report Data Source Values Show Twice (1374)
- FA Configuration > DRDoS > Settings is missing details (1391)
- Top Interfaces are duplicated for exporters in multiple device groups (3204)
- Undefined Error when modifying Guest Permissions (3219)
- CSV export of Volume reports shows incorrect rate data when resolution doesn't match datasource (3226)
- Error when filtering alarms by violator (3230)
- Add search.html type route to the new UI (3234)
- S3 Integration: Fix a crash when the database disappears at certain times (3235)
- Adding Show Interface option to a report shows outbound exporter as NA (3263)
- LDAP Authentication Fails due to Primary Key Duplicate Restraints (3281)
- Flow Collection Resumed Message Displays First Message instead of Last Message (3292)
- Host Index searches show 'first_seen' as the date of the host_index import (3334)
- Totals values could be doubled when an interface is metered both ingress and egress (3370)
- Severity card time frames don't match date selector (3434)
- Kafka logging can crash server processes (3437)
- Report links from Host Index would pop up a broken window (3483)
- Host Index cleanup tasks fail if H2H Index is turned off (3498)

Plixer Scrutinizer UI

Fixes

- Entities: Alarms: Events: Incidence correlation resize scrollbar (1336)
- Top Src/Dst Host pivot from an IP Group entity view opens a Username Entity view (1412)
- Setting custom interface speed to 0 to override displaying as percent utilization (1416)
- Dashboard issues: Excessive scroll bars on Windows and report gadget graph legends difficult to read (1602)
- CEF: timestamps for start/end times (3369)
- Support multiple usernames per host in alarms (3372)

7.1.8 Version 19.2.2 - September 2023

Plixer Scrutinizer

Fixes

- Addressed various security issues
- AWS Upgrade package dependency problem (3826)

7.1.9 Version 19.2.0 - May 2022

Plixer Scrutinizer

New Features

- Added option to toggle how device group hierarchy is displayed (153)
- Prioritize exporters that get disabled last in the event that a license overage causes some exporters to be disabled (203)

- Ship Scrutinizer with sysbench and a test script in files (1269)
- Expand CEF message content to include ports and usernames (2001)
- Improve messaging on “Unapproved Transport Protocols” alarm page (2161)
- AWS flowlogs: add support for new version 5 fields (2410)
- Workflow Issue: Unapproved Protocol Policy report pivot should include protocol filter (2426)
- AWS S3 Test Button: test the required permissions (2428)
- Improved alarm policies report link filters (2468)
- Run Report on Packet Flood event does not filter on the traffic that triggered the alert (2499)
- Don't use unencrypted connections for upgrades (port 80) (2607)
- Include shortened report URL in Report Threshold policy (2636)
- Create some new AWS reports for v5 elements (2651):
- Audit log entries for key management/encryption changes (2723)
- Ability to set a key lifetime (2724)
- VPC flow logs now require interface-id and flow-direction. (2817)

Fixes

- Addressed various security issues
- Fixed issue where configuration wouldn't synchronize when all settings are removed (473)
- Admin > Settings > Proxy Server has been renamed 'Google Maps Proxy Server' (941)
- PDFs for large reports show the “painting a Plixer” screen for the report screen shot (1054)
- Device tree hierarchy doesn't carry over to usergroups with explicit device group permissions (1500)
- Restore username details to alarm notifications (1999)
- Distributed data expiry errors without events/trends (2190)
- Deactivate Sliding Windows when FA algos are disabled (2310)
- ACL 'Like' filters don't work for ACL Descriptions (2312)
- DDoS and DRDoS alarms no longer present CSV access to the offender source list (2343)
- AWS S3 Test Button: test from the specified collector (2355)
- Improved Incident Correlation Algorithm (2380)
- Emailed reports from Report Threshold alert sometimes have incomplete report images (2413)
- ipfixify-template filepath updated in manual (2445)
- Unable to Export Report to PDF or Email Report for SSL not using port 443 (2463)
- “Report Direct Link” doesn't work for on-demand emailed reports (2485)
- Run Report option in Report Threshold Violation event list does not use the saved report filters (2491)
- Unable to export saved reports to CSV with space in saved report name (2506)

- Report Threshold Violation Email's URL should load the timeframe of the violation (2539)
- inserter.pm stops polling for SAFs, sampled SAFs, totals if the database is temporarily unavailable (2556)
- Graph and Table show in different timezones (2562)
- Top asn overstates exporter count (2595)
- Proxy server support needed for online upgrades (2608)
- Remove ICMP Ping check from upgrades and pass through variables (2609)
- Enable SSL as the default for offline repo servers (2618)
- SonicWALL IPFIX extension templates not being read correctly in v19.X (2622)
- AWS Flow reports - can't filter on the interface (2630)
- AWS flowlogs temp dir missing after upgrade to 19.1.0 (2670)
- allowed transports aren't sync'd to all collector nodes (2675)
- FA NULL scan Algo doesn't exclude destinations (2681)
- scrut_util -enable ram_spools blows away /etc/fstab (2684)
- Sflow inserting - Extra data after last expected column (2697)
- Latency Value ingesting from Ixia not show up properly on Scrutinizer UI (2709)
- Special case sFlow interface instances missing (2712)
- FA Worm Algos don't exclude hosts (2732)
- Update docs.plixer.com to reflect how syslog alerts are configured (2773)
- events.backfill_summaries() crashing with ddos events (2774)
- FA Breach algo doesn't exclude servers (2805)
- An offline update server with self signed certificates may try http (rather than https) and fail (2812)
- Host Index is now configured in Flow Analytics (2856)
- %m in syslog notifications includes CEF (2870)
- Reparser will not redefine templates without hard restart (2882)
- Running single direction report via the top interfaces view returns 'No Template' (2883)
- Scrutinizer device inactivity threshold is not triggering violations (2890)
- Remove plixer_syslogd from systemctl on upgrade (2892)
- FCGI Timeout settings removed after upgrade (2893)
- Install fails with dependency error on 'device-mapper-multipath' (2905)
- Distributed Upgrade hanging at TASK [Gathering Facts] (2907)
- Disabling an Algorithm does not remove its exporters from plixer.streams_config (2944)
- FA Reverse Shell doesn't exclude source (2952)
- Low spool disk space "FA streaming was disabled" does not disabe FA streaming (2979)
- Event Policy Customization Improvements (2985)
- Events with empty target/violator lists crash the policy view (3010)

Plixer Scrutinizer UI

New Features

- Unapproved Protocol Policy third donut chart now has top hosts using protocol (966)
- Include Time Zone in the report date/time display (1012)
- Monitor -> Network Maps Grid view delete option (1030)
- Better DNS Resolve Setting description (1053)
- Latest alarm message to events table (1199)
- CSV links in Policy entity (1207)

Fixes

- Naming a dashboard “Network” in V19.0.2 renames it to “Subnet” (909)
- History Navigation shows Alarms by ID instead of English Description (924)
- Navigating into alarm monitor sometimes throws an ExpiredRequestID error (975)
- inbound and outbound interface reports from explore device tab do not apply the correct filter (988)
- Regression: Traffic %, Other, and Total displaying for sFlow reports (1004)
- New UI doesn’t use the time zone user preference in reports (1013)
- Time Stamps on Line and Step Stacked 1m data source, 1m resolution overlap (1017)
- Deleting the default collection causes “notExists” error when trying to add to the default collection (1027)
- Host Entity View -Top Alarms bell icon mouseover text does not align with click action. (1029)
- Reports against an exporter with no current flow data does not allow for timeframe changes. (1031)
- New UI | Explore -> Interfaces -> Refresh Rate is not saved (1033)
- Changing Report Options triggers direction back to INBOUND when bidirectional is allowed (1038)
- Clicking the add or remove selected buttons keeps the tooltip on screen (1050)
- Recent Alarms Dashboard gadget shows UTC timestamp for Last Event and Last Notification (1112)
- Explore: Devices not using User Default Unit setting - Shows Percent always (1113)
- Toggling Hostname resolution does not change IPs to hostnames in alarm policy views (1135)
- Device/Interface report filter inconsistent with the Show DNS or IP modes (1216)
- Host to Host Index search doesn’t render a report menu when clicking exporter hyperlinks (1218)
- Alarms Monitor Filtering Option by Violators/Targets returning “noDataAvailable” (1221)
- CSV export of a report loses DNS names (1241)

- PDF export of report only shows 10 lines (1242)
- Peak and 95th Percentile not showing on saved reports (1244)
- Report filters not showing up in the “Additional Filters” drop down (1259)
- Show Others displaying when set to No (1267)

Machine Learning Engine

New Features

- Add ML Engine metrics to Vitals reports (338)
- Support high availability (419)
- Support Zerologon detection (446)
- Support SIGRed detection (447)

7.1.10 Version 19.1.1 - September 2021

Plixer Scrutinizer

New Features

- Automatically shut down non-critical features when systems are overwhelmed (2703)

Fixes

- Addressed various security issues
- “Sizing your environment” guide
- Timeout when migrating large historical host_index tables (2337)
- Upgrades didn’t stop on database upgrade error (2638)
- Full alarm message not getting into ServiceNOW tickets (2640)

- AMI didn't have spools on RAM disk / tuning didn't run on AMI deployment (2646)
- Resizing disks with AWS C5 instances (2696)
- Performance issues with host_index process (2701)
- Inefficiency in building TopN view (2710)
- Max locks wasn't set high enough for some upgrades from v18 (2751)
- Report links from threshold violations had the wrong timeframe (2785)
- Registering a new collector could overwrite meta data on the primary (2788)
- Character encoding issues synchronizing binary data (2794)
- Pulling STIX TAXII threat list (2831)

Plixer Scrutinizer UI

Fixes

- URL too long error from report wizard with large exporter counts (852)
- Line and step graphs wouldn't load after switch from a Traffic Volume report (1032)
- Graph and tables in a report could show different timezones (1059)
- Changing Report Options triggers direction back to INBOUND (1060)
- Flow data with a single direction could break the gear menu (1062)

7.1.11 Version 19.1.0 - May 2021

Plixer Scrutinizer

New Features

- Scrutinizer services not required to run as root (187)
- Client - Server reports (261)
- Encrypt stored keys (516)
- Copy to clipboard button to api json tab (733)
- Option to toggle Show System Policies (786)

- Expanded and reworked Host Index and H2H Search (883)
- Target / Violator views and filtering in Alarm Monitor(898)
- Show Host Names and Show Acknowledged Events for Alarms(948)
- Include collector IP address in all vitals reports for grouping and filtering(1971)
- Refactor Alarms backend for better performance (2053)
- Flexible notification policies based on event criteria (2060)
- Autoreplicate support for multiple replicators (encrypt multiple passwords) (2111)
- Ability to set Alarm policies to inactive or store (2231)
- root login disabled on new deployments (2361)
- Cisco SDWan (Viptela) integration updated to support version 20 (2374)

Fixes

- Addressed various security issues
- Mapping: add checks and errors for duplicate map connections (313)
- Sorting by bytes does not account for units in Entity Views (724)
- New UI reports do not display Host Names (793)
- PDF Export of Summary Reports Top N and Overview failure (805)
- Classic View option from user menu doesn't work (893)
- Fix scrolling issues for Exporter Details list in Report Settings (939)
- Alarms takes too long to load and acknowledge (1586)
- Reverse DNS exclusions for alarms (1798)
- Reparser crash when Linux ARP cache filled (1970):
- Adding a notification profile to a saved report threshold doesn't work (1977):
- Child Groups not enforced for FA exclusion (2030):
- Vitals process crashing with extremely high MFSNs in flow streams (2090):
- Custom URL Dashboard Gadgets not working (2214):
- Valid licenses with Expired PNI/PSI eval's prevent the upgrade from running (2217):
- Stream bloat on heavily loaded systems could cause disk space problems (2235):
- Running out of file descriptors on heavily loaded systems (2250):
- Invalid certificates in distributed upgrades (2273):
- TopN views are not always populated (2279):
- LDAP login takes too long with a very large list of security groups (2300):
- P2P Alarm report link not working (2307):
- Improve handling of truncated sFlow sampled headers (2336):

- Flow collection doesn't resume at the end of a network outage (2346):
- Set webui_timeout not working (2358):
- Scheduled report tasks called wrong binary name after upgrade (2379):
- IP exclusion only checking source IP for RST/ACK and Host Reputation (2382):
- Fix incorrect or missing sFlow interface numbers for instances above 63 (2393):
- AES key not syncing on upgrade affecting SNMP, AWS, and other credentials needed on a collector (2401):
- License Exceeded alarm detail shows no data in Alarm Monitor (2414):
- Addressed CVE-2021-28993 (2457):

7.1.12 Version 19.0.2 - January 2021

Plixer Scrutinizer

Fixes

- Disabling User Does Not Invalidate Session (2075)
- Input validation needed in some forms (2076)
- Session cookie value stored in local storage (2080)
- Postgres log noise from unnecessary scheduled analytics command (2118)
- Distributed upgrade issue coming from 19.0.0 (2198)
- pg_cron memory leak (2202)
- Fresh v19.0.1 OVA does not use the 19.0.1 repository (2205) F

7.1.13 Version 19.0.1 - December 2020

Plixer Scrutinizer

New Features

- DDOS: Support IPv6 (12)

- Add AWS Role Based Authentication for use in AWS (377)
- Allow AWS flowlog polling at 1m frequency (940)
- Enforce password policy on password change and restrict from using last four values (1235)
- Summary Reports added to new UI (1459)
- Add “scrut_util –show datasize” to enumerate DB schemas and their disk usage. (1539)
- Define Allegro IEs (1633)
- Support for new format of VPC flow logs (1890)
- Provide descriptions for AWS entity IDs (1891)
- Add Velocloud 4.0 IEs (tcpRttMs and tcpRetransmits) (1899)
- Document new AWS integration requirements (1992)

Fixes

- Mapping: Show Utilization only works for percent (54)
- Not excluding protocols by default (304)
- Secondary reporters show incorrect clock drift (696)
- Apache HTTP Server 2.4.0 - 2.4.39 Remote Open Redirect Vulnerability in mod_rewrite (739)
- Cannot Filter on S3 Bucket Element aws_account_id in a designed report (765)
- Internal Server Error when emailing PDF report name includes / (1065)
- Unable to Exclude IP address from DDoS algorithm (1316)
- Collector log error sflow buffer overrun at ./protocol/sflow/buffer.hpp line 146 (1480)
- VPC Flow Logs should be cleaned up more aggressively (1482)
- The plixer.idp.login_url field appears to be vestigial (1579)
- Other Options > GeoIP links not working (1592)
- Login banners are not working (1660)
- Interface names with special characters cause errors when triggering thresholds (1728)
- Alarm when disabling algorithms or ML stream (1734)
- Group Labels retain original input on Maps Dashboard Widget (1743)
- Host2host and host index lookups to work in distributed setup (1744)
- pgbouncer wont start after yum update (1796)
- Some reports were unable to display in percent interface view (1797)
- Reparser freezes on error during minutely exporter status updates (1812)
- No drillp-down into Connection on Maps (1813)
- Reparser memory leak in sFlow parser (1817)
- Devices blue after upgrade to version 19 (1840)

- ServiceNow Integration doesn't work when server response is too large (1842)
- Reporting: No Data for Timeframe automatically sends to start report wizard (1879)
- Sliding windows falling behind after upgrade to v19 (1911)
- Fix rollout issue for droppedPacketDeltaCount<unsigned64> (1912)
- Closing the report modal doesn't keep the report open (1917)
- Entity Views: sorting by bytes does not account for units (1918)
- Using LDAP user is authenticated but never added to a group when group list was too long (1920)
- Unable to disable unlicensed FA features (1930)
- Unrecognized key type: AWSLogs/xxxxxxxxxxx/ inc/lib/plixer/scrutinizer/awss3.pm line 547 (1941)
- Awss3.pm:373 – get_flowlogs() encountered an error while processing s3_connection_list: Invalid data Invalid data(unknown) for aws_account_id (1942)
- get_flowlogs() encountered an error while processing s3_connection_list: Invalid data (-) @ 1084 for transform (1945)
- Alarm Report data interval default empty for large time frame events (1946)
- NetFlow v5 sampling crashes postgres (1969)
- Too many open files (1981)
- multicast send failure 22: Invalid argument (1984)
- CEF notifications missing 'Device Version' (1988)
- Set 'ssl_prefer_server_ciphers' by default (1994)
- Missing sflow records after an upgrade (2002)
- Report values as rates in tables are incorrect after drilling in on a graph (2021)
- Distributed: AWS S3 secret failing when assigned to remote collector (2029)
- The application is running a vulnerable version of Apache (2068)
- The application is running a vulnerable version of Perl (2069)
- XSS Vulnerability in old UI mechanism to create groups (2070)
- Local file inclusion (2072)
- Autoreplicate support for multiple replicators (encrypt multiple passwords) (2111)
- Formula injection vulnerability in the ability to create third-party CrossCheck methods (2071)

Plixer Scrutinizer UI

New Features

- Entities: Hosts: Anomaly Chart (652)
- Summary Reports: Filtering (692)

Fixes

- Report filter descriptions don't always fill in (657)
- Dashboards not deleted (685)
- Drilling into Policy from Collection loses consistency vs Monitor View (688)
- Apache httpd: CWE-345: Insufficient verification of data authenticity (693)
- Reporting: Summary reports not stretching on page (744)
- Stop 'topping' the graphs (765)

7.1.14 Version 19.0.0 - August 2020

Important: Custom alarm policies are no longer supported. The Report Threshold Violation policy can be assigned one notification profile only.

New Features

- New workflow-based user interface (9)
- DDOS: Support IPv6 (12)
- Address data encryption in Scrutinizer (370)
- Initial Collections implementation (371)
- magicbus_fdw: Avro serialization (476)
- Advanced threat intelligence feeds (481)
- SNMP Enterprise MIB support for Viptela (717)
- Support for new VeloCloud information elements (727)
- Use tenant_id for db ROLE (740)
- Require a license key for free mode (780)
- Support for content updates (781)
- Streaming support for customer data lakes (782)
- Host to host flow connection search (783)
- Plixer Replicator integration (784)
- Update the Silverpeak IPFIX information elements (874)
- Advanced security algorithms (903)

- STIXV1 IP watchlist import (1006)
- STIXV2 IP watchlist import (1007)
- TAXII 2 feed support for IP indicators (1008)
- Domain reputation checking (1142)
- JA3 fingerprinting support (1144)
- Machine learning for security-specific events (1152)
- Machine learning for network-specific events (1153)
- New licensed features (1215)
- ML forecasting in Scrutinizer (1256)
- ServiceNow integration (1258)
- CEF notification action (1411)

Fixes

- Failed “system updates” report “no updates available” (541)
- scrut_util.exe –collect asa_acl gives error Use of uninitialized value \$debug in concatenation (614)
- Saved Reports Folder changes are not audited (636)
- Insecure Direct Object Reference (749)
- Vitalser Memory Leak (767)
- Define missing Cisco IEs (unknown_9_20000) (820)
- Define the unknown_elements for Viptela IPFIX exports (865)
- scrut_util –collect db_size is timing out (1196)

7.1.15 Version 18.20 - April 2020

New Features

- Optimized sFlow collection (496)
- New VeloCloud information elements (2073)
- Security updates (2154)
- SNMP Enterprise MIB support for Viptela (2164)
- Updated Silverpeak IPFIX information elements (2165)

- CentOS 7 : kernel update (2176)
- PostgreSQL security release 10.12 (2177)
- Change default eval key to 14 days (2190)

Fixes

- sFlow traffic discrepancies (2156)
- Saved report dashboard gadgets always display in totals (2167)
- Reporting issues when 0 byte flows are excluded (2179)
- Fixed issue with totals when both ingress and egress flows are exported (2196)

7.1.16 Version 18.18 - December 2019

New Features

- New VeloCloud reports (1939)
- Set admin password to instance_id for AMIs (2036)
- Add SSO authentication method to the manual (2039)
- Many updates, improvements, and clarifications in documentation (2051)
- New Viptela reports (2124)
- Option template based descriptions for VeloCloud LinkUUID (2133)

Fixes

- Create scheduled reports was also requiring admin tab permission (421)
- Auto refreshing pages would prevent session timeout (1441)
- Resolve timeout for FA reverse DNS exclusions wasn't using setting from admin tab (1405)
- We now exclude 0 byte flows biFlow records for reporting and FA (1536)

- Protocol exclusions were not audited (1756)
- 255 character limitation for 'Security Groups Allowed' when configuring LDAP integration (1816)
- Improved column naming in some VeloCloud reports (1936)
- Resolve a harmless UDP receive buffer error (1985)
- Viptela reports would sometimes not show all vEdge hosts (1992)
- Session timeout based on backend activity, not frontend activity (2030)
- PDF report displays no data when data is present (2040)
- Expand Disk scrut_util commands now support NVME drives (2041)
- If an IdP certificate is not provided, SAMLRequests should be unsigned (2106)
- SSO - Submitting metadata XML via the admin view form incorrectly parses out tags (2107)
- Fixed memory leak in vitalser (2041)

7.1.17 Version 18.16 - September 2019

New Features

- Viptela SD-WAN reports (16)
- Permission configuration on a role basis (270)
- Changed AWS Flow Log collection to use S3 buckets and added support for multiple regions and customer IDs (378)
- VeloCloud SD-WAN reports (550)
- Service Now Notification support (569)
- Appliance self migration from CentOS 6 to CentOS 7 (826)
- Ability to Add/remove/update Defined Applications via the API (891)
- Single-Sign-On support through SAML 2.0 (897)
- Alarm when authentication tokens will expire in 30 days or have expired (937)
- Deleting an exporter doesn't block collection (992)
- Removed device specific status notifications (1099)
- Audit logs can now be expired after a configurable duration (1171)
- FDW option to Database migrator for faster PostgreSQL migrations (1205)
- Flow inactivity alarms are now checked across a distributed cluster and are per exporter rather than per interface (1254)
- Support for Fortinet application names (1425)
- Support Nokia (formerly 'Alcatel-Lucent') IPFIX (1735)
- Support for Gigamon Application Intelligence (1832)

Fixes

- Schedule emails will now use the theme from Admin > Settings > System Preferences (185)
- The ability to use an auth token with any URL (308)
- UTF8 issue with Japanese characters in email alert notifications (636)
- ‘Truncate map labels’ was grabbing an extra character sometimes (700)
- Addressed an issue with flow class sequence numbers with distributed upgrades (753)
- Removed admin restriction on running group level reports (841)
- Clarify several log error messages, and reduce their volume (846)
- Some Scrutinizer custom gadgets break the ability to add any gadget for all users (900)
- AMI: set partitions doesn’t remount pg_stat_tmp as a RAM drive (1066)
- Issue where deleted exporters may not be cleared out of LED stats table (1079)
- Issue where system updates could revert a setting causing “Panic: Can’t find temp dir” errors and the interface failing to load (1082)
- Higher default timeouts for collect asa_acl task (1085)
- Issue with special characters in PRTG integration (1117)
- Warnings when an exporter sends the same multiplier data two different ways as long as what it sends is consistent (1120)
- UNION SELECT errors in migrator (1132)
- Autofilling IP on host search from report tables (1140)
- Scheduled reports last sent time used incorrect (1142)
- SQL GROUP BY ERROR in the collector log (1145)
- Issue with Auto SNMP Update not disabling all SNMP calls (1158)
- PostgreSQL logs using too much disk space (1209)
- Special characters in notification profile breaks threshold’s ‘save & edit policy’ option (1229)
- Added stray columnar file check and alarm policy (1231)
- Monitor association of /var/db/fast and RAM spools (1239)
- Issue with running yum update on AWS EC2 instances (1249)
- Issue with load time of Admin > Host names view (1272)
- Defined application changes now realized on distributed collectors w/o a collector restarts (1297)
- Issue with alarm details and FQDN data for clusters using DB encryption (1314)
- DB disk usage stats did not always expire on distributed installs (1322)
- Collect support files includes the PostgreSQL log (1385)
- Allow snmpSystem details longer than 255 characters (1392)
- Errors from set tuning when two changes require a collector restart (1422)
- Getting Internal Server Error (500) when trying to access Maps > CrossCheck and Service Level Reports (1431)

- Some administrative changes for authentication did not generate audit events (1440)
- Addressed issue with ASA ACL collection when the reporter can not communicate with all firewalls (1447)
- * Issue with LDAP/TACACS usernames being case sensitive (1458)
- LDAP authentication was not failing over to try other servers (1489)
- Backup method documentation on docs.plixer.com (1506)
- Advanced TCP flag filters using strings would generate log noise (1527)
- Improved performance of Persistent Flow Risk algorithm (1536)
- Developer tasks_view hours filter causes Internal Server Error (500) (1542)
- Dashboards with multiple saved report gadgets cause oops errors (1544)
- Reporting across migrated data and new data doesn't use the migrated totals tables (1553)
- Migrated totals tables have the wrong scrut_templateid (1556)
- Peak values being less then the total values in the volume -> traffic volume reports (1588)
- Some English values in foreign language themes were out of date (1599)
- New reparser performance (1632)
- Migration from 16.3 mysql to 18.14 removed dashboard gadget permissions (1663)
- LDAP group checking was using sAMAccountName instead of the value specified in the configuration page (1668)
- Map object icons change colors based on polling availability (1691)
- The default group was not being set correctly for new users (1731)
- Payload size preventing CSV rendering of reports (1733)
- Saved reports belonging to users that no longer exist would not show up in report folders (1789)

NOTE: (1458)*

User accounts are no longer case sensitive when being checked on login. If multiple user accounts existed in Scrutinizer prior to the upgrade which were identical except for case, the excess accounts should be deleted from the interface.

7.1.18 Version 18.14 - May 2019

New Features

- Now including cstore table conversion script in utils (873)
- Improved default work_mem settings (951)

Fixes

- DB process needs priority over other processes when system runs out of memory (640)
- Acknowledging Multiple Pages of an Alarm, acknowledges all alarms (676)
- ‘unhandled multicast message’ in the collector log (714)
- Report Designer not saving added row (778)
- Drilling into Palo Alto User Report generates a blank pop up (780)
- Top Interfaces summarization timing out with high interface count (784)
- Issue when upgrading from version 16.7 (790)
- Issue where exporters sending bad timestamps would freeze spool file processing (793)
- “Save password” error when navigating from group membership (832)
- Large number of DrDOS violations could crash process (849)
- Error when changing exporter status (850)
- Backup exporters count against licensing even if same IP is already active (851)
- Interface thresholds would only violate if there was both inbound and outbound traffic (872)
- IP group detection not working for v6 addresses (894)
- Cleanup logging for sFlow exports from Cumulus Router (895)
- Not all interface names are collected from FireSIGHT (896)
- Issue with business hours ending at midnight (903)
- First time LDAP authentication would fail if local authentication is disabled (904)
- Scheduled reports attaching wrong pdf to email (956)
- Drilling in on an interval from volume reports could display the wrong timeframe (963)
- A slow connection could impact API latency LED for other collectors (971)
- Issue with NTP daemon not starting automatically on some installs (990)
- Updated DRDOS thresholds to be ratios instead of fixed packet counts (1004)
- TACACS authentication would work if disabled but configured (1009)
- Issue with the scale APM outbound jitter was displayed in (1019)
- Reparser could not connect to the DB with a space in the password (1063)
- One exporter not collecting when at maximum license count for exporters (1130)

7.1.19 Version 18.12.14 - January 2019

New Features

- Realtime DDOS and DRDOS detection before data is written to disk (10)
- FQDN reports are back and better performing (87)
- Interface threshold checks are now done once a minute and check one minute of data (105)
- FireSIGHT integration includes username support (111)
- FireSIGHT integration includes interface names (112)
- Group reports now include members of child groups (274)
- “User Accounts” permission to allow restriction of Scrutinizer user account creation (299)
- Added option to disable CrossCheck threshold notifications (447)

Fixes

- Faster report CSV generation (132)
- FireSIGHT integration detects connection loss and attempts to reconnect to FirePOWER (167)
- Top interfaces values were understated for sFlow exporters sending multiple totals flows per minute (177)
- PostgreSQL log rotation (263)
- Rate values for Trend reports are now based on graph interval (267)
- Link Back Host set to the wrong port on a deployed AMI (301)
- Installer no longer displays post install script errors (319)
- Add Audit messages when connections to LDAP servers fail (26415)
- Fixed username filtering when name is based on IPv6 address (26768)
- Faster Defined Application tagging (26874)

7.1.20 Version 18.9 - September 2018

Fixes

- Fixed issue with multiple defined applications on the same IP (26874)
- Improved contrast for some icons in dark themes (26511)
- System user was counting against licensing limits (26536)
- Fixed issue with top N gadgets and exporters only sending egress flows (26550)
- Fixed the Analytics Violation Overview link on the Alarms tab (26557)
- Fixed issue using Gmail to send emails (26579)
- Fixed issue with emailing table views (26587)
- Fixed issue with TopN subnets gadget and SAF aggregation (26600)
- Fixed issue with editing designed reports (26602)
- Backslash in LDAP passwords caused issue on upgrade (26613)
- Fixed issue with map labels in dashboards (26619)
- Multiple subnet filters issue in MySQL (26629)
- Fixed issue with threshold details not being cleared out when switching reports (26632)
- Fixed issue editing designed reports with some manufactured columns in them (26650)
- Fixed issue with interface permissions in mapping (26652)
- Fixed issue with row limiting in CSV files (26655)
- Fixed issue with flow vitals when packets contain multiple flow sets for the same template (26699)
- Reporting: Top 10 rows on any page are now color coded as the graph (26731)
- Postgres installs - improved reporting temp table performance (26735)

7.1.21 Version 18.7 - July 2018

New Features

- Added QRadar Integration (23542)
- Changed dashboard gadget behavior to improve usability and clearly display gadget titles (26194)
- Numerous improvements to the manual (26310)

Fixes

- Flickering issue with report graphs when loading a report (24546)
- Formatting issues in Maps Tab alerts (25156)

- Double tooltip when mousing over report graph (25504)
- Audits from IPv6 hosts are now correctly received and recorded (26042)
- Issues with input parameters for the Users API (26298)
- Optimized rollups (26317)
- Decreased time necessary to run upgrades (26318)
- Links from alarms heatmap were not working (26342)
- Tuning would too aggressively set roller memory (26345)
- Addressed upgrade issue related to DB locking (26350)
- Improved dashboard gadget behavior based on customer feedback (26358)
- Reparser: Fix understatement of NetFlow v9 flow volume in vitals report (26360)
- AWS instances would not upgrade if on Postgres 9.5 (26370)
- Maps couldn't be saved in dashboard gadgets (26371)
- Could not generate PDFs of reports in Japanese (26372)
- Fixed issue with Japanese characters in emailed reports (26373)
- Other Options > Search link not working (26395)
- Peaks in totals tables were 5 minute byte counts rather than 1 minute byte counts (26399)
- Forensic filters were not forcing change to forensic data (26406)
- Fixed filtering on AS number under Admin > Definitions > Autonomous Systems (26431)
- Fixed issue with making dashboards visible to a user group (26451)

* This is the last supported release for the CentOS 6 and MariaDB platforms

7.1.22 Version 18.6 - June 2018

New Features

- Test button for LDAP/RADIUS/TACACS setup (9911)
- Ability to acknowledge alarms with any combination of filters (15154)
- scrut_util command to disable ping for devices that have not responded (16826)
- Manufactured columns can be included in the report designer (17589)
- Full back button support (18291)
- Automatically detect which SNMP credentials to use for exporters (19981)
- Ability to manage interface details via API (20068)
- Ability to filter on a port range (21522)
- All interface reports now account for metering on each interface in the report (21744)

- Host -> AS -> Host reports for additional BGP reporting (21770)
- Major release upgrade to PostgreSQL 9.6 and 10 (22220)
- scrut_util command to enable/disable ipv6 (22773)
- User can be locked out after n failed login attempts (23267)
- Full foreign datastore support in collection and rollups (23478)
- Ability to exclude domain names from flow analytics (23924)
- Ability to edit URLs for custom gadgets (24134)
- Milliseconds now included with formatted timestamps where applicable (24164)
- Columnar store support for AWS Scrutinizers (24297)
- Ability to customize the login page (24452)
- Improved support for configuration of multiple LDAP servers and domains (24600)
- Ability to grant dashboards to other users / groups (24661)
- Default PostgreSQL datastore is columnar. Better disk space utilization and IO performance. (24781)
- Performance improvements for flow class lookups (24948)
- Support IPv4-mapped IPv6 addresses in subnet and ipgroup filters (PostgreSQL) (25077)
- Report IP Group with protocol and defined applications (25216)
- Support for Flowmon probe elements (25289)
- DrDoS detection for memcached and CLDAP attacks (25396)
- Ability to schedule operating system updates (26187)

Fixes

- Flow metrics vitals times now align with ingestion time (12972)
- Ungrouped now visible by non-admin users (22530)
- Tidy up loose ends when deleting exporters. Deleted exporters will stay deleted. (22588)
- Stop showing disabled exporters in the exporters LED (22654)
- Some timezones were duplicated in the selector (24107)
- Latency reports per exporter (24115)
- Addressed issue reporting on multiple interfaces with different metering configured (24659)
- Issue with generating PDF with device group filters (24703)
- Restrict PaloAlto username collection to only internal IPs (24790)
- Donut/Pie Graph not available in Top -> Interfaces report (24875)
- Map interface utilization arrows always pointed in the same direction (24893)
- 'cancel report' button truly cancels backend reporting requests. (24899)
- Device menu in Google maps (24993)

- Cleaned up log noise from Cisco ISE data collection (25027)
- Scheduled reports font issue on AWS (25111)
- Remove memcached external exposure CVE-2017-9951 (25317)
- FlowPro APM jitter report (25323)
- Audit report times now display as clients timezone (25399)
- Addressed CVE-2014-8109 (25419)
- Issue with Queue Drops >> Queue Drops By Hierarchy (25660)

7.1.23 Version 17.11 - November 2017

New Features

- Support for Oracle cloud (24685)

Fixes

- Vitals errors when a user with a long UID is created (24500)
- Save button for filters would go away if field was selected, but not changed (24560)
- Localhost Unlicensed after upgrade to 17.10 (24586)
- Collector appears down after Daylight Savings Time change (24616)
- Potential short gap in rollups after collector restart (24647)

7.2 FAQ

To quickly look up answers to frequently asked questions, select a category:

- [*Plixer Scrutinizer core system*](#)
- [*Plixer Replicator integration*](#)

Important: For additional questions or concerns, contact [*Plixer Technical Support*](#).

7.2.1 Plixer Scrutinizer core system

Q) Can we try Plixer Scrutinizer before paying for a full subscription?

A) To try out Plixer Scrutinizer or any other Plixer product, contact [Plixer Technical Support](#) and ask about an evaluation license.

Q) How do I view the details of our current Plixer Scrutinizer license?

A) To add a new license or view the details of the currently applied Plixer Scrutinizer license, navigate to the **Admin > Plixer > Licensing** page.

Q) What happens when a license key expires?

A) Evaluation keys will cease to function after their expiry date. Plixer Scrutinizer subscription keys include a 60-day grace period where data collection continues on, but access to the data is unavailable until a new key is added. Legacy perpetual licenses will never expire, but deployments they are applied to cannot be upgraded.

Q) Is Plixer Scrutinizer available in other languages?

A) The Plixer Scrutinizer web interface supports localization to other languages via the **Admin > Definitions > Language** page.

Q) What does it mean when we get an unexpected inconsistency error while trying to power on our Plixer Scrutinizer ESXi virtual appliance?

A) An unexpected inconsistency error during the ESXi virtual appliance startup indicates that the server clock is not correctly set, resulting in the disk checks failing. To resolve this issue, set your ESXi host to sync with an NTP server and then redeploy the Plixer Scrutinizer OVF.

Q) How do I stop or start Plixer Scrutinizer services?

A) To stop or start individual services, use the following command:

```
SCRUTINIZER> services <service_name | all> <stop | start>
```

Valid service_name values:

- plixer_flow_collector
- plixer_syslogd

- `httpd`
- `plexer_db`

Q) Why am I unable to log in to the appliance using the default scrutinizer password?

A) If you have yet to change the password for the `plexer` user and are unable to log in with the password `scrutinizer` check if caps lock is turned on. Additionally, if you are using a keyboard with a non-ANSI layout, you will need to enter the password using the same key positions as a standard ANSI keyboard (e.g. `scrutiniyer` on QWERTZ keyboards).

Q) I got locked out of my account after several failed log-in attempts. How can I log back in to the web interface?

A) When an account is locked due to multiple failed logins, there are two methods that an `Admin` user can use to unlock it:

- In the web interface, go to **Admin > Security > Users**. Select the locked username/account, click the **Authentication Method** tab in the Edit User modal, and then change the Authentication Method from 'locked' to the appropriate method.
- Launch the `scrut_util` utility and enter the following command at the `SCRUTINIZER>` prompt:

```
unlock [username] <auth_method>
```

Q) Why is the backup file we used to perform an instance restore missing?

A) In addition to overwriting the Plexier Scrutinizer instance the backup was restored to, the restore script is also set to delete the backup file used to perform the operation. As such, it is best to always create a copy of the backup file before initiating a restore.

Q) Why is the Aggregated Alarm Timeout setting missing for certain FA algorithms?

A) Not all FA algorithms support (or benefit from) Aggregated Alarms. The Aggregated Alarm Timeout setting is only available for algorithms with continuous Alarm Events that can be combined.

Q) What do I do if I forgot to assign a static MAC address to the Plexier Scrutinizer NIC?

A) If this happens when deploying virtual appliance to a distributed cluster, contact [Plexier Technical Support](#). to obtain a new license key.

Q) How do I free up disk space on my Plexier Scrutinizer server?

A) Historical data can be trimmed to free up disk space. You can do either of the following:

- In the web interface, go to **Admin > Settings > Data History**, and then adjust the current retention settings.
- Launch the `scrut_util` utility and enter the following command at the `SCRUTINIZER>` prompt:


```
SCRUTINIZER> expire history
```

7.2.2 Plexier Replicator integration

Q: If a new Collector is added to the distributed Plexier Scrutinizer cluster, will it automatically be included under auto-replication/load balancing?

A: After adding a new Plexier Scrutinizer Collector to the distributed cluster, the configuration file will need to be updated. After that, either run `scrut_util --autoreplicate` again or wait for the next scheduled run, if configured.

Q: What happens when auto-replication is enabled and a Collector goes down? Will Plexier Replicator detect the offline Collector and automatically re-balance flows?

A: If a Plexier Scrutinizer Collector goes offline while auto-replication is enabled, the configuration file should be updated. After that, either run `scrut_util --autoreplicate` again or wait for the next scheduled run (if configured) for the flows to be re-balanced across the available Collectors.

Q: Why do the flows being sent to our distributed cluster seem unbalanced? Shouldn't auto-replicate distribute load evenly across all configured Collectors?

A: Because every Exporter sends flows at a different rate that changes throughout the day, a rate of 200 flows/s is assumed for each Exporter added to a Collector. An Exporter's flows will only be reassigned when a Collector exceeds its defined inbound limits as the real rates are observed.

Q: How does the auto-replicate function assign new Exporters to Collectors in a distributed cluster?

A: Exporters are automatically distributed across the Collectors in the cluster using a “round-robin” system until Collectors reach their defined flow rate limits.

Q: When a Collector becomes over-provisioned, how does the system determine which Collector the Exporter will be moved to?

A: When a Collector reaches or exceeds its flow limits, it will be excluded from the round robin and additional Exporters will be assigned to other available Collectors defined in the `autoreplicate.conf` file.

Q: Our auto-replicate Collector threshold is set to 40,000 flows/s but the system is reporting spikes of more than 40,000 flows/s several times a day. Shouldn't Exporters be reassigned to other Collectors once the threshold is reached?

A: The auto-replicate function uses 24-hour averages to determine whether Exporters need to be reassigned to different Collectors.

Q: How many Plixer Replicators installations does the auto-replication/load balancing functionality support?

A: One Plixer Replicator can be defined via the Plixer Scrutinizer web interface, but the system supports an unlimited number of Plixer Replicator appliances through manual configuration. For additional details, see the section on [advanced configurations](#) for Plixer Replicator or contact [Plixer Technical Support](#).

Q: How many Plixer Replicator seed Profiles and/or unique listening ports does auto-replication/load balancing functionality support?

A: One seed Profile and one unique listening port are supported by each auto-replicate configuration. For additional details, see the section on [advanced configurations](#) for Plixer Replicator or contact [Plixer Technical Support](#).

Q: Why does the `scrut_util --autoreplicate` command need to be run manually?

A: By default, the `scrut_util --autoreplicate` command is not scheduled to run automatically to allow users to run it only when necessary.

Q: Can Exporters be statically assigned to specified Collectors with auto-replication enabled?

A: To define static Exporters for a specific Collector, create a new Profile under Plixer Replicator as normal and configure the Exporters to send flows to the Collector.

Q: Is auto-balancing affected by Exporters that are configured to send flows directly to the Plixer Scrutinizer appliance?

A: When auto-replication is enabled, Exporters that do not send their flows through Plixer Replicator are considered *Rogue Exporters*. These Exporters will still count against the Collector's Exporter count and flow limits.

Q: What happens if all Collectors in a distributed cluster have reached their collection rate thresholds and new Exporters are added to the configuration?

A: If there are no longer any Collectors with available bandwidth for load balancing, Plixer Replicator will stop assigning new Exporters but continue its auto-replication using the most recent viable configuration.

Q: Where does Plixer Scrutinizer log any auto-replicate changes that are made?

A: These changes can be found under `/home/plixer/scrutinizer/files/logs/`, inside an epoch-stamped file that contains a final output after the `autoreplicate` command completes.

Q: Does auto-replication take Missed Flow Sequence Numbers (MFSNs) when load balancing?

A: MFSNs are not taken into account when flows are assigned across a distributed Plixer Scrutinizer cluster for load balancing.

Q: Are old Profiles automatically removed after the configuration changes?

A: Older Profiles are not automatically removed, but it is safe to delete any Profiles that are no longer relevant to the current configuration file.

7.3 Functional IDs

The Plexer Scrutinizer system relies on a number of generic functional accounts/IDs to control access to the environment's different components and their respective functions.

The following table lists all default functional IDs used by a Plexer Scrutinizer installation:

| System Component | Account/ID | Type | Access Level | Function |
|-------------------------|------------|-----------------|----------------|---|
| Operating system | root | Interactive | Privileged | Provides root access to the Plexer Scrutinizer OS, with unrestricted shell, SSH, and console access |
| | plexer | Interactive | Non-privileged | Primary user for the interactive <code>scrut_util</code> CLI utility and provides access to run all Plexer Scrutinizer processes and services |
| | pgbound | Non-interactive | Non-privileged | Used to manage remote database access between nodes, e.g. user/role access, load balancing, etc. |
| | postgres | Non-interactive | Privileged | Used for database operations during deployment |
| | apache | Non-interactive | Privileged | Primary HTTP services user |
| Database | plexer | Interactive | Privileged | Primary database role used by application processes for both local and remote access |
| | postgres | Non-interactive | Privileged | Used for local database access during deployment, upgrades, and scheduled <code>pg_cron</code> tasks |
| Web interface | admin | Interactive | Privileged | Provides full access to web interface management functions |

Types:

- *Interactive* - can be used to grant a user all privileges inherent to the ID
- *Non-interactive* - reserved for internal use by the system and cannot be assigned to users

Access levels:

- *Privileged* - has elevated permissions, such as superuser or system admin access
- *Non-privileged* - granted only the access rights required for the ID's intended function(s)

7.4 Localization

Plixer Scrutinizer supports translation of the web interface for localization purposes.

To add or modify translations of UI elements:

1. Navigate to Admin > Definitions > Language.
2. Select a language from the dropdown menu.
3. Click on a key type to enter or modify the translation for that UI element.
4. Repeat the process to translate additional UI elements.

Language translations are saved as `/home/plixer/scrutinizer/files/localize_languageName.xls`.

7.5 Glossary

This glossary is meant to serve as a reference for terms that are specific to Plixer Scrutinizer, Machine Learning Engine, and general computer networking concepts.

7.5.1 Plixer Scrutinizer terms and concepts

Alarm Policy

Rule sets that define what types of network behavior or activity should be monitored as Events and trigger Alarms

Flow Analytics

A library of field-tested algorithms used to analyze network behavior, detect unexpected activity, and report Events and Alarms

IPFIXify

A software agent that reads text-based logs, syslog messages, Windows EventLogs and various other types of data sources and sends the information in flows using the IPFIX protocol

Plixer ML Engine

Software component providing AI capabilities to allow the ingestion and processing of extremely large volumes of flow data for intelligent anomaly and threat detection

Protocol Exclusions

Defines protocols to exclude during the collection process per Exporter, Exporter interface, and/or all Exporters and interfaces

Reverse-Path Filtering

Allows Collectors to receive non-local traffic that may have been forwarded by a proxy or flow replication solution, such as Plixer Replicator

SAF (Summary and Forensic)

An optimized system of storing flow data that uses summary tables to condense collected information without compromising transparency or accuracy

TI (Threat Index)

A single value comprised of events with different weights that age out over time

7.5.2 Machine Learning Engine terms

Deep learning

A progression of supervised and unsupervised learning to create an artificial neural network that can learn and make intelligent decisions on its own

K-means clustering

An algorithm that groups behaviors into common clusters

Link prediction

A method that detects anomalies and analyzes a device's interactions with other devices rather than just a particular behavior

Supervised learning

The process of training a machine learning algorithm using labeled data sets

Unsupervised learning

The process of training a machine learning algorithm to identify patterns or classifications in untagged data sets

7.5.3 General networking

ACK (Acknowledgment Code)

A unique signal sent by a computer to show that it has successfully transmitted data

API (Application Programming Interface)

A software component that allows applications to share data and functionality

CIDR (Classless Inter-Domain Routing)

An IP addressing method that improves the efficiency of allocating IP addresses

CLI (Command-line Interface)

A text-based interface for applications and operating systems that allows a user to enter commands and receive

DNS (Domain Name System)

A system by which computers and other devices on the Internet or Internet protocol networks are uniquely identified using names matched to their IP addresses

ICMP (Internet Control Message Protocol)

A protocol used for devices within the network to determine possible network issues

IPFIX (Internet Protocol Flow Information Export)

A protocol intended to collect and analyze the of flow data from supported network devices

LDAP (Lightweight Directory Access Protocol)

An open, cross-platform protocol used to access and maintain directory services for assets in an Internet protocol network

MTTR (Mean Time to Resolution)

The the average amount of time between the detection and remediation of a security threat or incident

NDR (Network Detection and Response)

A cybersecurity solution that use machine learning to detect cyber threats and aid remediation

NTP (Network Time Protocol)

A networking protocol used to synchronize device clocks over the Internet

NXDOMAIN (No Existing Domain)

An error message that means that a domain mentioned in the Domain Name System (DNS) query does not exist

RADIUS (Remote Authentication Dial-In User Service)

A client-server AAA (authentication, authorization, accounting) protocol used to manage remote user access to a network

SNMP (Simple Network Management Protocol)

An IP network protocol used to collect data related to state and/or behavior from devices on a network

SSDP (Simple Service Discovery Protocol)

A network protocol used for advertising and discovering network services

SSH (Secure Shell Protocol)

A network communication protocol that allows network services to be used securely over an unsecured network

SYN scan

A port scanning technique that allows for the discovery of the status of a communications port without establishing a full connection

Syslog

A cross-platform network logging protocol used to send and/or receive alerts between different devices on a network

STIX (Structured Threat Information eXchange)

An industry-standard file format for the exchange of threat information between organizations and platforms

TAXII (Trusted Automated eXchange of Indicator Information)

A protocol that allows the transmission of threat information, primarily in STIX format, between systems and organizations

TACACS+ (Terminal Access Controller Access-Control System)

A protocol where the remote access server and the authentication server provide validation for users attempting to access the network

TLS handshake

The process that starts secure communication between a client and a server

TCP (Transmission Control Protocol)

A connection-oriented protocol that enables the bidirectional exchange of messages between devices on the same network

UDP (User Datagram Protocol)

A communication protocol for transmitting messages between applications and programs in a network

Virtual appliance

A pre-configured virtual machine image with pre-installed software that is meant to serve a specific function

VPC (Virtual Private Cloud)

A secure and private cloud hosted in a public cloud

VRF (Virtual Routing and Forwarding)

A technology that separates routing tables to isolate management traffic to the management interface

7.6 Third-party attributions

Certain open source or other third-party software components are integrated and/or redistributed with Plixer Scrutinizer software and Plixer machine learning software. The licenses are reproduced here in accordance with their licensing terms, these terms only apply to the libraries themselves, not Plixer Scrutinizer software and/or Plixer machine learning software.

Copies of the following licenses can be found in the licenses directory at `/home/plixer/scrutinizer/files/licenses/`.

7.6.1 Plixer Scrutinizer

Apache 2.0 License

Apache Giraph

<http://giraph.apache.org/>

Copyright (c) 2011-2016, The Apache Software Foundation

Apache Kafka

<http://kafka.apache.org/>

Copyright (c) 2016 The Apache Software Foundation

Bean Validation

<http://beanvalidation.org/>

Copyright (c) 2007-2013 Red Hat, Inc.

code-prettify

<https://github.com/google/code-prettify>

Copyright (c) 2006 Google Inc.

cstore_fdw

https://github.com/citusdata/cstore_fdw

Copyright (c) 2016 - 2017 Citus Data, Inc.

Explorer Canvas

<https://github.com/arv/ExplorerCanvas>

Copyright (c) 2006 Google Inc.

fonts

<http://code.google.com/p/fonts>

Copyright (c) 2009 Google Inc.

Guava

<https://github.com/google/guava>

Copyright (c) Google, Inc.

Kafka

[hogan.js](#)

<https://github.com/twitter/hogan.js>

Copyright (c) 2011 Twitter, Inc.

Jackson JSON Processor

<https://github.com/FasterXML/jackson>

Copyright (c) Jackson Project

Javassist

<https://github.com/jboss-javassist/javassist>

Copyright (c) 1999-2013 Shigeru Chiba. All Rights Reserved.

Javax Inject

<http://code.google.com/p/atinject>

Copyright (c) 2010-2015 Oracle and/or its affiliates

Jetty

<https://github.com/eclipse/jetty.project>

Copyright (c) 2008-2016 Mort Bay Consulting Pty. Ltd., Copyright (c) 1996 Aki Yoshida, modified April 2001 by Iris Van den Broeke, Daniel Deville.

Keyczar

<http://code.google.com/p/keyczar/>

Copyright (c) 2008 Google Inc.

Log4j

<http://logging.apache.org/log4j/>

Copyright (c) 2007 The Apache Software Foundation

LZ4 Java

<https://github.com/jpountz/lz4-java>

Copyright (c) 2001-2004 Unicode, Inc

RocksDB

<http://rocksdb.org/>

deflate 1.2.8 Copyright (c) 1995-2013 Jean-loup Gailly and Mark Adler, inflate 1.2.8 Copyright (c) 1995-2013 Mark Adler

Snappy for Java

<https://github.com/xerial/snappy-java>

Copyright (c) 2011 Taro L. Saito

WenQuanYi Micro Hei fonts

<https://github.com/anthonyfok/fonts-wqy-microhei>

Copyright (c) 2005-2010 WenQuanYi Board of Trustees

ZkClient

<https://github.com/sgroschupf/zkclient>

Copyright (c) 2009 Stefan Groschupf

ZooKeeper

<https://zookeeper.apache.org>

Copyright (c) 2009-2014 The Apache Software Foundation

Artistic 1.0 License

business-isbn

<https://github.com/briandfoy/business-isbn/>

Copyright (c) 2001-2013, Brian D Foy

Common-Sense

<http://search.cpan.org/~mlehmann/common-sense/>

Terms of Perl - No Copyright Author - Marc Lehmann

Compress-Raw-Zlib

<http://search.cpan.org/~pmqs/Compress-Raw-Zlib/>

Copyright (c) 2005-2009 Paul Marquess.

Compress-Zlib

<http://search.cpan.org/~pmqs/IO-Compress-2.066/lib/Compress/Zlib.pm>

Copyright (c) 1995-2009 Paul Marquess.

crypt-ssleay

<https://github.com/gisle/crypt-ssleay/>

Copyright (c) 2006-2007 David Landgren, Copyright (c) 1999-2003 Joshua Chamas, Copyright (c) 1998 Gisle Aas, Copyright (c) 2010-2012 A. Sinan Unur

DBD-mysql

<http://search.cpan.org/dist/DBD-mysql/>

Large Portions Copyright (c) 2004-2013 Patrick Galbraith, 2004-2006 Alexey Stroganov, 2003-2005 Rudolf Lippman, 1997-2003 Jochen Wiedmann, with code portions Copyright (c) 1994-1997, their original authors

Digest-MD5

<http://search.cpan.org/dist/Digest-MD5/>

Copyright (c) 1995-1996 Neil Winton., Copyright (c) 1990-1992 RSA Data Security, Inc., Copyright (c) 1998-2003 Gisle Aas

Encode-Locale

<http://search.cpan.org/dist/Encode-Locale/>

Copyright (c) 2010 Gisle Aas

ExtUtils-MakeMaker

<http://search.cpan.org/~bingos/ExtUtils-MakeMaker/>

Terms of Perl - No Copyright

extutils-parsexs

<https://github.com/dagolden/extutils-parsexs/>

Copyright (c) 2002-2009 by Ken Williams, David Golden and other contributors

HTML::Template::Pro

<http://search.cpan.org/~viy/HTML-Template-Pro-0.9510/>

Copyright (c) 2005-2009 by I. Yu. Vlasenko., copyright (c) 2000-2002 Sam Tregar

HTML-Parser

<http://search.cpan.org/dist/HTML-Parser/>

Copyright (c) 1995-2009 Gisle Aas, Copyright (c) 1999-2000 Michael A. Chase.

HTML-Tagset

<http://search.cpan.org/~petdance/HTML-Tagset/>

Copyright (c) 1995-2000 Gisle Aas., Copyright (c) 2000-2005 Sean M. Burke., Copyright (c) 2005-2008 Andy Lester

HTTP::Cookies

<http://search.cpan.org/~oalders/HTTP-Cookies-6.04/lib/HTTP/Cookies.pm>

Copyright (c) 1997-2002 Gisle Aas, Copyright (c) 2002 Johnny Lee

HTTP::Daemon

<http://search.cpan.org/~gaas/HTTP-Daemon-6.01/lib/HTTP/Daemon.pm>

Copyright (c) 1996-2003 Gisle Aas

HTTP::Date

<http://search.cpan.org/~gaas/HTTP-Date-6.02/lib/HTTP/Date.pm>

Copyright (c) 1995-1999 Gisle Aas

HTTP::Negotiate

<http://search.cpan.org/~gaas/HTTP-Negotiate-6.01/lib/HTTP/Negotiate.pm>

Copyright (c) 1996, 2001 Gisle Aas.

http-message

<https://github.com/php-fig/http-message>

Copyright 1995-2008 Gisle Aas.

IO-Compress

<http://search.cpan.org/dist/IO-Compress/>

Copyright (c) 2005-2009 Paul Marquess.

IO-HTML

<http://search.cpan.org/~cjm/IO-HTML-1.001/lib/IO/HTML.pm>

Copyright (c) 2012-2013 Christopher J. Madsen

IO-Socket-IP

<http://search.cpan.org/~pevans/IO-Socket-IP-0.37/lib/IO/Socket/IP.pm>

Copyright (c) 2010-2013 Paul Evans

IO-Socket-SSL

<http://search.cpan.org/~sullr/IO-Socket-SSL/>

Copyright (c) 1999-2002 Marko Asplund, Copyright (c) 2002-2005 Peter Behroozi, Copyright (C) 2006-2014 Steffen Ullrich

JSON

<http://search.cpan.org/~makamaka/JSON/>

Copyright (c) 2005-2013 by Makamaka Hannyaharamitu

JSON::XS

<http://search.cpan.org/~mlehmman/JSON-XS/>

Copyright (c) 2008 Marc Lehmann

libwww-perl

<http://search.cpan.org/dist/libwww-perl/>

Copyright (c) 1995-2009 Gisle Aas, 1995 Martijn Koster, 2002 James Tillman, 1998-2004 Graham Barr, 2012 Peter Marschall.

libxml-perl

<http://perl-xml.sourceforge.net/libxml-perl/>

Copyright (c) 2001-2003 AxKit.com Ltd., 2002-2006 Christian Glahn, 2006-2009 Petr Pajas

Log::Log4perl

<http://search.cpan.org/~mschilli/Log-Log4perl/>

Copyright (c) 2002-2013 Mike Schilli and Kevin Goess

LWP::MediaTypes

<http://search.cpan.org/~gaas/LWP-MediaTypes-6.02/lib/LWP/MediaTypes.pm>

Copyright (c) 1995-1999 Gisle Aas.

Net::Flow

<http://search.cpan.org/~acferen/Net-Flow-1.003/lib/Net/Flow.pm>

Copyright (c) 2007-2008 NTT Information Sharing Platform Laboratories

Net-HTTP

<http://search.cpan.org/~oalders/Net-HTTP-6.17/lib/Net/HTTP.pm>

Copyright (c) 2001-2003 Gisle Aas.

Net-LibIDN

http://search.cpan.org/~thor/Net-LibIDN/_LibIDN.pm

Copyright (c) 2003-2009, Thomas Jacob

Net-SNMP Perl

<http://search.cpan.org/~dtown/Net-SNMP-v6.0.1/>

Copyright (c) 2001-2009 David M. Town

Net-SSLeay

<http://search.cpan.org/~mikem/Net-SSLeay/>

Copyright (c) 1996-2003 Sampo Kellomaki, Copyright (C) 2005-2006 Florian Ragwitz, Copyright (c) 2005 Mike McCauley

Perl

<http://www.perl.org>

Copyright (c) 1993-2005, by Larry Wall and others.

Perl Object Environment

<http://search.cpan.org/~rcaputo/POE-1.367/lib/POE.pm>

Copyright (c) 1998-2013 Rocco Caputo

perl-digest-sha1

<http://search.cpan.org/~gaas/Digest-SHA1-2.13/SHA1.pm>

Copyright (c) 2003-2008 Mark Shelor

perl-File-Listing

https://centos.pkgs.org/7/centos-x86_64/perl-File-Listing-6.04-7.el7.noarch.rpm.html

Copyright (c) 1996-2010, Gisle Aas

perl-ldap

<http://ldap.perl.org>

Copyright (c) 1997-2004 Graham Barr

perl-REST-Client

<https://centos.pkgs.org/6/epel-i386/perl-REST-Client-272-1.el6.noarch.rpm.html>

Copyright (c) 2008 - 2010 by Miles Crawford

perl-XML-NamespaceSupport

<http://search.cpan.org/~perigrin/XML-NamespaceSupport-1.11/lib/XML/NamespaceSupport.pm>

Copyright (c) 2001-2005 Robin Berjon.

Pod-Escapes

<http://search.cpan.org/~neilb/Pod-Escapes/>

Copyright (c) 2001-2004 Sean M. Burke

Pod-Simple

<http://search.cpan.org/~dwheeler/Pod-Simple-3.26/lib/Pod/Simple.pod>

Copyright (c) 2002 Sean M. Burke.

TimeDate

<http://search.cpan.org/dist/TimeDate/>

Copyright (c) 1995-2009 Graham Barr.

Types::Serialiser

<http://search.cpan.org/~mlehmann/Types-Serialiser-1.0/Serialiser.pm>

Terms of Perl - No Copyright Author - Marc Lehmann

URI

<http://search.cpan.org/~ether/URI/>

Copyright (c) 1998 Graham Barr, 1998-2009 Gisle Aas

WWW-RobotRules

<http://search.cpan.org/~gaas/WWW-RobotRules-6.02/lib/WWW/RobotRules.pm>

Copyright (c) 1995, Martijn Koster, 1995-2009, Gisle Aas

XML-LibXML

<http://search.cpan.org/~shlomif/XML-LibXML/>

Copyright (c) 2001-2003 AxKit.com Ltd., 2002-2006 Christian Glahn, 2006-2009 Petr Pajas

XML-SAX

<http://search.cpan.org/~grantm/XML-SAX/>

No Copyright listed - Terms of Perl

Xml-sax-base

<http://search.cpan.org/~grantm/XML-SAX-Base-1.08/BuildSAXBase.pl>

No Copyright listed - Terms of Perl

yaml-perl-pm

<http://search.cpan.org/dist/YAML-Perl/>

Copyright (c) 2001, 2002, 2005. Brian Ingerson., Copyright (c) 2005, 2006, 2008. Ingy döt Net., Some parts Copyright (c) 2009 Adam Kennedy

Artistic 2.0 License

NetPacket::

<http://search.cpan.org/~cganesan/NetPacket-LLC-0.01/>

Copyright (c) 2001 Tim Potter and Stephanie Wehner., Copyright (c) 1995 - 1999 ANU and CSIRO on behalf of the participants in the CRC for Advanced Computational Systems ('ACSys').

BSD 2-Clause Simplified License

JabberWerxC

<https://github.com/cisco/JabberWerxC>

Copyright (c) 2010-2013 Cisco Systems, Inc.

BSD 3-Clause License

Babel

<http://babel.pocoo.org/>

Copyright (c) 2007 - 2008 Edgewall Software

Crypt-DES

<http://search.cpan.org/~dparis/Crypt-DES/>

Copyright (c) 1995, 1996 Systemics Ltd, Modifications are Copyright (c) 2000, W3Works, LLC

D3.js

<http://d3js.org/>

Copyright (c) 2010-2014 2010-2017 Mike Bostoc

Jinja2

<http://jinja.pocoo.org/>

Copyright (c) 2008 - 2011 Armin Ronacher, Copyright 2007-2011 by the Sphinx team, 2006 - 2010 the Jinja Team, Copyright 2010, John Resig, Copyright 2010, The Dojo Foundation

libevent

<http://libevent.org/>

Copyright (c) 2000-2007 Niels Provos, Copyright (c) 2007-2012 Niels Provos and Nick Mathewson

MarkupSafe

<http://github.com/mitsuhiko/markupsafe>

Copyright (c) 2010 by Armin Ronacher

memcached

<http://code.google.com/p/memcached/>

Copyright (c) 2000 - 2003 Niels Provos, Copyright (c) 2003, Danga Interactive, Inc.

Netcast

<http://freshmeat.sourceforge.net/projects/netcast>

Copyright (c) Stanislaw Pasko

Net-SNMP

<http://www.net-snmp.org/>

Copyright: See licenses/net-snmp.txt

PhantomJS

<http://phantomjs.org/>

Copyright (c) 2011 Ariya Hidayat

pyasn1

<http://sourceforge.net/projects/pyasn1/>

Copyright (c) 2005-2017, Ilya Etingof

RequireJS

<http://requirejs.org/>

Copyright (c) 2010-2012, The Dojo Foundation

Scala

<http://www.scala-lang.org/>

Copyright (c) 2002-2010 EPFL, Lausanne, unless otherwise specified

SNMP::Info

<http://freshmeat.net/projects/snmp-info>

Copyright (c) 2002-2003, Regents of the University of California, Copyright (c) 2003-2010 Max Baker and SNMP::Info Developers

strace

<http://sourceforge.net/projects/strace/>

Copyright (c) 1991, 1992 Paul Kranenburg, Copyright (c) 1993 Branko Lankester, Copyright (c) 1993 Ulrich Pegelow, Copyright (c) 1995, 1996 Michael Elizabeth Chastain, Copyright (c) 1993, 1994, 1995, 1996 Rick Sladkey, Copyright (c) 1998-2001 Wichert Akkerman, Copyright (c) 2001-2017 The strace developers

sudo

<http://www.sudo.ws/sudo/>

Copyright (c) 1994-1996, 1998-2018 Todd C. Miller

uthash

<http://sourceforge.net/projects/uthash/>

Copyright (c) 2008-2017 Troy D. Hanson

Yahoo! User Interface Library

<http://developer.yahoo.com/yui>
Copyright (c) 2007, Yahoo! Inc.

yuicompressor

<http://developer.yahoo.com/yui/compressor/>
Copyright (c) 2013 Yahoo! Inc.

CDDL 1.0 License

Java Servlet API

<http://java.sun.com/products/servlet/index.jsp>
Copyright (c) 1997-2003 Oracle and/or its affiliates

JAX-RS Specification

<https://java.net/projects/jax-rs-spec>
Copyright (c) 1996-2014 Oracle and/or its affiliates

Jersey

<http://jersey.java.net/>
Copyright (c) 2010-2016 Oracle and/or its affiliates, 2000-2011 INRIA, France Telecom, 2004-2011 Eugene Kuleshov,

jsr250-api

<https://jcp.org/aboutJava/communityprocess/final/jsr250/index.html>
Copyright (c) 1999-2013 Oracle and/or its affiliates.

CDDL 1.1 License

HK2

<https://javaee.github.io/hk2/>
Copyright (c) 2010-2017 Oracle and/or its affiliates.

CURL License

cURL

<http://curl.haxx.se>

Copyright (c) 1998 - 2013, Daniel Stenberg

GPL & MIT Licenses

coResizable 1.6

<http://www.bacubacu.com/colresizable/>

Copyright (c) 2012 Alvaro Prieto Lauroba

jQuery Accordion

<http://docs.jquery.com/UI/Accordion>

Copyright (c) 2007 Jörn Zaefferer

jQuery Ajaxmanager

<http://github.com/aFarkas/Ajaxmanager>

Copyright (c) 2010 Alexander Farkas

jQuery Autocomplete

<http://bassistance.de/jquery-plugins/jquery-plugin-autocomplete/>

Copyright (c) 2009 Jörn Zaefferer

jQuery blockUI

<http://malsup.com/jquery/block/>

Copyright (c) 2007-2013 M. Alsup

jQuery Checkboxes

<https://github.com/SamWM/jquery-Plugins>

Copyright (c) 2006-2008 Sam Collett

jQuery Form

<http://malsup.com/jquery/form/>

Copyright (c) 2017 jquery-form

jQuery Select Boxes

<https://github.com/SamWM/jQuery-Plugins>

Copyright (c) 2006-2008 Sam Collett

GPL 2.0 License

CSSTidy

<http://csstidy.sourceforge.net>

Copyright (c) 2005, 2006, 2007 Florian Schmitz

Filesystem in Userspace

<http://fuse.sourceforge.net/>

Copyright (c) 1989, 1991 Free Software Foundation, Inc.

filterlist.js

<http://www.barelyfitz.com/projects/filterlist/index.php>

Copyright (c) 2003, Patrick Fitzgerald

Iotop

<http://freshmeat.net/projects/iotop>

Copyright (c) 2007, 2008 Guillaume Chazarain, 2007 Johannes Berg

jQuery Pagination

https://github.com/gbirke/jquery_pagination

Copyright (c) Gabriel Birke

libdbi-drivers

<http://freshmeat.net/projects/libdbi-drivers>

Copyright (c) 2001-2007, David Parker, Mark Tobenkin, Markus Hoenick

Nmap Security Scanner

<http://nmap.org/>

Copyright (c) 1996–2016 Insecure.Com LLC

sshpas

<http://freshmeat.net/projects/sshpas>

sysstat

<http://sebastien.godard.pagesperso-orange.fr/>

Copyright (c) 1999-2009 Sebastien Godard

GPL 3.0 License

Ansible

<http://www.ansible.com/>

Copyright (c) 2017, Ansible Project

MariaDB

<http://mariadb.org/>

Copyright (c) The MariaDB Foundation

LGPL 2.1 License

DHTMLGoodies

<http://www.dhtmlgoodies.com/index.html?page=termsOfUse>

Copyright (c) 2005 - 2007 Alf Magne Kalleland, www.dhtmlgoodies.com

Dynarch DHTML Calendar

<http://www.dynarch.com/jscal/>

Copyright (c) 2002 - 2005 Mihai Bazo

jFeed

<https://github.com/jfhovinne/jFeed>

Copyright (c) 2007-2011 Jean-François Hovinne

dual mit/gpl

libmspack

<http://freshmeat.net/projects/libmspack>
Copyright (c) 1991, 1999, 2003-2004 Stuart Caie

Open Virtual Machine Tools

<http://open-vm-tools.sourceforge.net>
Copyright (c) 2010-2015 VMware, Inc. All rights reserved.

paramiko

<https://github.com/paramiko/paramiko/>
Copyright (c) 2003-2009 Robey Pointer

whatever_hover

https://github.com/jasoncheow/whatever_hover/
Copyright (c) 2005 - Peter Nederlof

LGPL 3.0 License

GNU Libidn

<http://www.gnu.org/software/libidn/>
Copyright (c) 2004-2012 Simon Josefsson

MIT License

Argparse4j

<http://argparse4j.sourceforge.net/>
Copyright (c) 2011, 2015, Tatsuhiro Tsujikawa

Backbone.js

<https://github.com/jashkenas/backbone>
Copyright (c) 2010-2017 Jeremy Ashkenas, DocumentCloud Copyright (c) 2013 Charles Davison, Pow Media Ltd

base2

<http://code.google.com/p/base2/>
copyright (c) 2007-2009, Dean Edwards

c3.js
<http://c3js.org/>
Copyright (c) 2013 Masayuki Tanaka

Cocktail.js
<https://github.com/onsi/cocktail>
Copyright (c) 2012 Onsi Fakhouri

d3pie.js
<http://d3pie.org/>
Copyright (c) 2014-2015 Benjamin Keen

dshistory.js
<http://code.google.com/p/dshistory/>
Copyright (c) Andrew Mattie

Expat
<http://expat.sourceforge.net>
Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd

Flotr2
<https://github.com/HumbleSoftware/Flotr2>
Copyright (c) 2012 Carl Sutherland

gridstack.js
<http://troolee.github.io/gridstack.js/>
Copyright (c) 2014-2016 Pavel Reznikov, Dylan Weiss

hoverIntent
<http://cherne.net/brian/resources/jquery.hoverIntent.html>
Copyright (c) 2011 Brian Cherne

httplib2
<https://github.com/jcgregorio/httplib2>
Copyright (c) 2006 by Joe Gregorios, Thomas Broyer, James Antills, Xavier Verges Farreros, Jonathan

Feinbergs, Blair Zajacs, Sam Rubys, Louis Nyffeneggert, Dan-Haim, 2007 Google Inc.

JOpt Simple

<http://jopt-simple.sourceforge.net/>

Copyright (c) 2004-2015 Paul R. Holser, Jr

jQuery

<http://jquery.com/>

Copyright (c) 2007 - 2011, John Resig

jQuery Fixed Header Table

<http://fixedheadertable.com>

Copyright (c) 2013 Mark Malek

jQuery Form Plugin

<https://github.com/malsup/form>

Copyright (c) Mike Alsup

jQuery Live Query

<https://github.com/brandonaaron/livequery>

Copyright (c) 2010 Brandon Aaron

jQuery Migrate

<https://plugins.jquery.com/migrate/>

Copyright (c) jQuery Foundation and other contributors

jQuery Plugin: Superfish

<https://superfish.joelbirsch.co/>

Copyright (c) 2008 Joel Birch

jQuery Plugin: tablesorter

<http://tablesorter.com/docs/>

Copyright (c) 2014 Christian Bach

jQuery Plugin: Treeview

<http://bassistance.de/jquery-plugins/jquery-plugin-treeview/>

Copyright (c) 2007 Jörn Zaefferer

jQuery qtip.js

<http://craigsworks.com/projects/qtip/>

Copyright (c) 2009 Craig Thompson

jQuery UI

<http://jqueryui.com/>

Copyright (c) 2014, 2015 jQuery Foundation and other contributors

jQuery Validation Plugin

<http://bassistance.de/jquery-plugins/jquery-plugin-validation/>

Copyright (c) Jörn Zaefferer

jQuery-metadata

<https://github.com/jquery-orphans/jquery-metadata>

Copyright (c) 2001-2010. Matteo Bicocchi (Pupunzi)

jQuery-mousewheel

<https://github.com/brandonaaron/jquery-mousewheel>

Copyright (c) 2011 Brandon Aaron

Logalot

<https://www.npmjs.com/package/logalot>

Copyright (c) Kevin Mårtensson

Moment Timezone

<http://momentjs.com/timezone/>

Copyright (c) JS Foundation and other contributors

Moment.js

<http://momentjs.com/>

Copyright (c) JS Foundation and other contributors

pbox.js

<http://www.ibegin.com/labs/>

Python Six

<https://pypi.python.org/pypi/six/>

Copyright (c) 2010-2015 Benjamin Peterson

are therefore Copyright (c) 2001, 2002, 2003 Python Software Foundation

PyYAML

<http://pyyaml.org/wiki/PyYAML>

Copyright (c) 2006 Kirill Simonov

Raphael

<https://github.com/DmitryBaranovskiy/raphael>

Copyright (c) 2008-2013 Dmitry Baranovskiy, Copyright (c) 2008-2013 Sencha Labs

setuptools

<https://github.com/pypa/setuptools>

Copyright (C) 2016 Jason R Coomb

Simple AJAX Code-Kit

<https://github.com/abritinthebay/simpleajaxcodekit>

Copyright (c) 2005 Gregory Wild-Smith

simplejson

<https://github.com/simplejson/simplejson>

Copyright (c) 2008, Bob Ippolito

SLF4j

<http://www.slf4j.org>

Copyright (c) 2004-2017 QOS.ch

sqlify

<https://www.npmjs.com/package/sqlify>

Copyright (c) 2017 Vajahath Ahmed

Underscore JS

<http://underscorejs.org/>

Copyright (c) 2009-2015 Jeremy Ashkenas, DocumentCloud and Investigative Reporters & Editors

wickedpicker.js

<http://github.com/wickedRidge/wickedpicker>

Copyright (c) 2015-2016 Eric Gagnon

MIT Old Style License

c-ares

<http://c-ares.haxx.se/>

Copyright (c) 1998, 2009 by the Massachusetts Institute of Technology., Copyright (c) 2004 - 2011, Daniel Stenberg with many contributors

Mozilla Public License 1.1

Rhino

<https://github.com/mozilla/rhino>

OpenSSL License & SSLeay License (conjunctive)

OpenSSL

<http://www.openssl.org>

Copyright (c) 1998-2011 The OpenSSL Project, Copyright (C) 1995-1998 Eric Young. This product includes software written by Tim Hudson. Copyright (C) 1998-2011 The OpenSSL Project.

Oracle BCL License

Oracle Java

<http://www.oracle.com/technetwork/java/index.html>

Copyright (c) 1993 - 2015, Oracle and/or its affiliates.

PostgreSQL License

PostgreSQL

<http://www.postgresql.org/>

Portions Copyright (c) 1996-2018, The PostgreSQL Global Development Group

Portions Copyright (c) 1994, The Regents of the University of California

Unicode, Inc. License Agreement

International Components for Unicode (ICU)

<http://www.icu-project.org/>

Copyright (c) 2010 Yahoo Inc., Copyright (c) 1996-2012, International Business Machines Corporation and Others.

7.6.2 Pliker Machine Learning

Apache Software License

Cython

<https://cython.org/>

Copyright (c) Robert Bradshaw, Stefan Behnel, Dag Seljebotn, Greg Ewing, et al.

asyncpg

<https://github.com/MagicStack/asyncpg>

Copyright (c) MagicStack Inc

python-dateutil

<https://github.com/dateutil>

Copyright (c) Gustavo Niemeyer

requests

<https://docs.python-requests.org/en/latest/>

Copyright (c) MMXVIX. A Kenneth Reitz Project. Kenneth Reitz

BSD License

idna

<https://github.com/kjd/idna>

Copyright (c) Kim Davies

joblib

<https://joblib.readthedocs.io/en/latest/>

Copyright (c) Gael Varoquaux

numpy

<https://numpy.org/>

Copyright (c) 2021 NumPy. All rights reserved. Travis E. Oliphant et al.

pandas

<https://pandas.pydata.org/>

patsy

<https://github.com/pydata/patsy>

Copyright (c) Nathaniel J. Smith

scikit-learn

<https://scikit-learn.org/stable/>

scipy

<https://scipy.org/>

Copyright (c) 2021 SciPy.

statsmodels

<https://www.statsmodels.org>

Copyright (c) 2009-2019, Josef Perktold Skipper Seabold, Jonathan Taylor, statsmodels-developers

LGPL GNU License

chardet

<https://github.com/chardet/chardet>

Copyright (c) Daniel Blanchard

MIT License

pmdarima

<http://alkaline-ml.com/pmdarima/>

© Copyright 2017-2021, Taylor G Smith

pytz

<https://github.com/stub42/pytz>

Copyright (c) Stuart Bishop

six

<https://github.com/benjaminp/six/tree/65486e4383f9f411da95937451205d3c7b61b9e1>

Copyright (c) Benjamin Peterson

urllib3

<https://github.com/urllib3/urllib3>

Copyright (c) Andrey Petrov

Mozilla Public License 2.0

certifi

<https://certifi.io/en/latest/>

Copyright (c) 2020 Kenneth Reitz

7.7 Plixer Technical Support

Plixer Technical Support is available with an active maintenance contract. Contact our support team at:

- +1 (207) 324-8805 ext 4
- <https://www.plixer.com/support/>

For further questions, check out the [FAQ page](#) or contact [Plixer Technical Support](#).

Looking for documentation specific to the Classic UI? [Click here](#).