
Scrutinizer Documentation

Version 19.6.1

Plixer

July 03, 2025

VERSION 19.6.1

| | | |
|----------|---|------------|
| 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 | 27 |
| 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 | 47 |
| 3.6 | Importing data | 48 |
| 3.7 | Plixer ML Engine | 57 |
| 3.8 | Environment sizing | 65 |
| 4 | Use Cases | 83 |
| 4.1 | NetOps Use Cases | 84 |
| 4.2 | SecOps Use Cases | 105 |
| 5 | Features and Functionality | 125 |
| 5.1 | Plixer Scrutinizer web interface | 125 |
| 5.2 | Data aggregation | 353 |
| 5.3 | Machine learning | 355 |

| | | |
|----------|------------------------------------|------------|
| 6 | Advanced Services | 359 |
| 6.1 | Backups | 359 |
| 6.2 | Certificate management | 370 |
| 6.3 | Integrations | 373 |
| 6.4 | Interactive CLI | 418 |
| 6.5 | Plixer Scrutinizer APIs | 442 |
| 6.6 | Reverse-path filtering | 469 |
| 6.7 | Streaming to data lakes | 472 |
| 6.8 | Upgrades and updates | 472 |
| 7 | Additional Resources | 491 |
| 7.1 | Appendices | 491 |
| 7.2 | Changelog | 645 |
| 7.3 | FAQ | 682 |
| 7.4 | Functional IDs | 687 |
| 7.5 | Localization | 688 |
| 7.6 | Glossary | 688 |
| 7.7 | Third-party attributions | 696 |
| 7.8 | Plixer Technical Support | 720 |

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 contains guides for the installation/deployment and initial setup of the different types of Plixer Scrutinizer appliances.

Note: Prior to deploying Plixer Scrutinizer and other Plixer One platform products, ensure that firewall rules are correctly configured based on [this table](#).

2.1 Hardware appliance

After removing the Plixer Scrutinizer hardware appliance from its packaging, verify that all accompanying accessories (rackmount kit, appliance-locking bezel and keys, and power cord) are included. The appliance can be mounted in a standard 19-inch rack or cabinet.

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. Refer to the port labels to identify the ports to be used on the rear panel of the appliance:
 - 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
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.
 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 setup.

Once the Plixer Scrutinizer hardware appliance has been set up and cabled, proceed to [configuring the appliance](#).

Note:

- The Ethernet port pairs are configured for adapting load balancing (bonding mode 6).
 - The iDRAC virtual console can also be used for the appliance's [initial setup](#).
-

2.2 Virtual appliances

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

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.

Note: Scrutinizer virtual appliance packages are also available for download from the Plixer Customer Portal.

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 AWS (AMI)

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

Deployment guide - AWS

After subscribing to the service via the AWS Marketplace product page, deploy the Plixer Scrutinizer AMI by creating/launching a new EC2 instance with the following configuration:

- **Names and tags:** Configure the name, resource types, and optional tags for the instance.
- **Application and OS images:** Select the Plixer Scrutinizer AMI from the *My AMIs* tab.
- **Instance type:** Select *C5.2xlarge* for flow rates up to 10,000 flows per second (contact [Plixer Technical Support](#) for assistance if the expected flow volume exceeds that).
- **Key pair:** Select or create a new key pair to assign to the instance.
- **Network settings:** Select the VPC, subnet, and security group to assign the instance to.

Important: Because an active instance's primary private IP address cannot be released, we recommend deploying the AMI with two NICs and using the secondary as the collection interface.

- **Storage:** Leave the the size of the root volume (`/dev/xvda/`) at the default 100 GB.
- **Advanced details:** Set **Shutdown behavior** to *Stop* and **Termination protection** to *Enabled*.

After the instance has been launched, access the Plixer Scrutinizer web interface via the instance's primary private or public IP address, and then proceed with [adding a license](#).

Note: Use the following command to SSH to the server as the `plixer` user after the instance has been launched:

```
ssh -i PATH_TO_KEY/key.pem plixer@SCRUTINIZER_IP
```

Expanding database size for AWS

To expand the database size for a Plixer Scrutinizer AMI, create one or more additional EBS volumes in the same *availability zone* and attach them to the instance.

These volumes can then be made available to Plixer Scrutinizer by following [these instructions](#).

Note: `set partitions` will need to be run from the `scrut_util/SCRUTINIZER>` prompt for each additional drive attached to the instance:

```
set partitions
```

Adding instance resources

Follow these steps to change the Plixer Scrutinizer instance type to increase *CPU and RAM allocations*:

1. SSH to the instance as the `plixer` user and stop all services via *scrut_util*:

```
SCRUTINIZER> services all stop
```

2. Power off the OS:

```
shutdown -h now
```

3. Stop the instance. If an Elastic IP was assigned, note the instance ID and Elastic IP address beforehand.
4. Change the instance type and restart the instance following this guide.
5. Verify that a new public DNS (IPv4), Private DNS, and Private IPs have been assigned. The Elastic IP address should also be re-assigned to the instance ID if necessary.

After the instance has been reconfigured, SSH to the Plixer Scrutinizer IP address as the `plixer` user and run the following *scrut_util command* to re-tune the system:

```
SCRUTINIZER> set tuning
```

2.2.2 ESXi

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

Deployment guide - ESXi

To deploy the Plixer Scrutinizer virtual appliance in ESXi, 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

Deploying the OVF template

1. Contact [Plixer Technical Support](https://files.plixer.com/PACKAGE_PATH_AND_FILENAME) and use the link they provide (https://files.plixer.com/PACKAGE_PATH_AND_FILENAME) to download the latest VMware virtual appliance package.

Note: The latest VMware virtual appliance package is also available for download from the Plixer Customer Portal.

2. Extract the contents of the package to a location on the ESXi server.
3. In vSphere or vCenter, right-click the host to deploy the appliance to and select *Deploy OVF Template* from the menu.
4. Select *Local file* and browse to the Plixer Scrutinizer OVF and VMDK files before clicking *Next*.
5. Provide a name for the Plixer Scrutinizer virtual appliance and continue to follow the deployment wizard.
6. When prompted, select the datastore, set the disk format to *Thick Provision* and click *Next*.
7. After selecting the network to be used by the virtual appliance, verify the configuration in the summary before clicking *Finish* to import the Plixer Scrutinizer virtual appliance. This may take a few moments.
8. 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 [environment sizing guides](#) if necessary.
 - d. Click *OK* to save the current configuration and return to the previous page.
9. Right-click on the Plixer Scrutinizer virtual machine to power it on.
10. After the appliance boots up, click the console preview window and select *Open Remote Console*.

From the console, you can now log in 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.

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 via the power menu in VMware Tools.
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:

`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 Plexier 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 the SCRUTINIZER> prompt. The new size of the database should reflect the additional space added with the new hard disk.

2.2.3 Hyper-V

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

Deployment guide - Hyper-V

To deploy the Plexier Scrutinizer virtual appliance in Hyper-V, 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

Deploying the Hyper-V virtual appliance

1. Contact [Plixer Technical Support](#) and use the link they provide (https://files.plixer.com/PACKAGE_PATH_AND_FILENAME) to download the latest Hyper-V virtual appliance package.

Note: The latest Hyper-V virtual appliance package is also available for download from the Plixer Customer Portal.

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

From the console, you can now log in 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 *these steps to configure the hard drive for Plixer Scrutinizer*.

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.

2.2.4 KVM

The Plixer Scrutinizer virtual appliance must be deployed in KVM 16 or higher.

Deploying the KVM virtual appliance

1. Contact [Plixer Technical Support](#) and use the link they provide to download the latest KVM virtual appliance package:

```
curl -k -o PACKAGE_FILENAME.tar.gz https://files.plixer.com/PACKAGE_PATH/  
↳PACKAGE_FILENAME.tar.gz
```

Note: The latest KVM virtual appliance package is also available for download from the Plixer Customer Portal.

2. Create a directory for the install:

```
mkdir /kvm/scrutinizer_vm/
```

3. Extract the contents of the package to the new directory:

```
sudo tar xvzf PACKAGE_FILENAME.tar.gz -C /kvm/scrutinizer_vm/
```

4. Run the installation script in the new directory:

```
cd /kvm/scrutinizer_vm/PACKAGE_FILENAME
sudo ./install-kvm-scrut.sh
```

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

After the 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*.

2.2.5 Proxmox

To deploy the Plixer Scrutinizer virtual appliance in Proxmox, follow these steps:

Deploying the KVM virtual appliance

1. Contact *Plixer Technical Support* and use the link they provide to download the latest VMware virtual appliance package:

```
curl -k -o PACKAGE_FILENAME.tar.gz https://files.plixer.com/PACKAGE_PATH/
PACKAGE_FILENAME.tar.gz
```

2. Extract the contents of the file and upload the *.vmdk file to a location that can be accessed by Proxmox on the Proxmox server (e.g., `/var/lib/vz/template/`).
3. Convert the vmdk disk image to a Proxmox-compatible format:

```
qemu-img convert -f vmdk -O qcow2 FILENAME.vmdk Plixer_Scrutinizer.qcow2
```

4. Create a new virtual machine in Proxmox with the following configuration:

- BIOS: OVMF (UEFI)
- SCSI controller: VMware PVSCSI
- Network adapter: E1000
- CPU/memory: [Recommended sizing](#)
- Add a new EFI disk with default sizing

5. Import the disk via the CLI:

```
qm importdisk 100 /var/lib/vz/template/Plixer_Scrutinizer.qcow2 local -  
↪ zfs
```

6. Attach the imported disk to the virtual machine:

```
qm set 100 -scsi0 local-zfs:vm-101-disk-1
```

7. Delete the unused disk and start the VM.

After the Plixer Scrutinizer virtual appliance has been deployed, log in using the credentials `plixer:scrutinizer` via the console and proceed with the [basic configuration process](#).

Note:

- When attaching the imported disk, verify that its name matches what's displayed in the GUI.
 - The syntax in the instructions above should be modified to match the actual VMID and disk numbers used.
-

2.2.6 Optimizing datastores

Plixer Scrutinizer environments with larger flow rates can be significantly 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 with the [recommended capacity](#).

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

The Plixer ML Engine is deployed as part of Plixer One Enterprise environments, enabling AI-driven monitoring and detection capabilities in Plixer Scrutinizer.

This section contains full deployment guides for the different types of appliances available.

Note: The guides below only apply to v19.5 of the Plixer ML Engine, which can only be paired with Plixer Scrutinizer 19.6.0 and above. Deployment instructions for v19.4 of the engine can be found in the Plixer Scrutinizer 19.5.3 manual. Contact [Plixer Technical Support](#) for assistance.

2.3.1 Pre-deployment preparations

The following preparatory steps should be completed before starting the deployment procedure for any type of Plixer ML Engine appliance.

Deploying the Plixer Machine Learning VM

Use the template obtained with the Plixer One Enterprise license to deploy the Plixer ML VM locally.

This VM will function as a separate deployment host and includes all prerequisite resources. The Plixer ML Engine environment will be deployed and managed from this VM.

Note:

- For vSphere multi-node deployments, the template should be added to vSphere. See the [vSphere multi-node cluster deployment guide](#) for more information.
 - [PuTTY] *VT100 line drawing even in UTF-8 mode* (under Settings > Window > Translation) must be enabled for the setup wizard to be displayed correctly. Requested details can be pasted into the prompts/dialogs using **Shift+Insert**.
-

License and engine registration

Creating an authentication token

The Plixer ML Engine will require an authentication token for Plixer Scrutinizer, which can be created as follows:

1. In the Plixer Scrutinizer web interface, navigate to **Admin > Resources > ML Engines**.
2. Click the **+** button to add a new ML engine.
3. From the dropdown, select the type of ML engine paired with your Plixer Scrutinizer environment:
 - *Single VM*
 - *Amazon AWS*
 - *Azure*
 - *vSphere multi VM Cluster*
4. Enter a name to assign to the engine, and then click **Save**.

After returning to the main view, click the name of the new ML engine and save/copy the primary reporter address and authentication token shown in the tray. These will be required during the Plixer ML Engine deployment process.

Confirming SSH credentials

To complete the appliance setup process, the Plixer ML Engine will need to establish an SSH session with the primary reporter/server of the Plixer Scrutinizer environment. As such, the IP address of the primary reporter and the `plixer` user password will need to be provided.

If a private SSH key is required, verify that the public key is configured on the primary Plixer Scrutinizer reporter/server under `/home/plixer/.ssh/authorized_keys`. The private key should also be accessible from the machine hosting the Plixer ML Engine virtual appliance, as the path to the key will need to be entered during the appliance setup process.

2.3.2 AWS (EKS)

After completing the *pre-deployment preparations*, follow the instructions below to set up the necessary infrastructure and deploy the Plixer ML Engine in AWS.

Additional prerequisites for AWS

- AWS IAM user secret access key ID and secret access key
- A VPC with two subnets for the deployment

Note:

- The Plixer ML VM (the deployment host) deployed as part of the [pre-deployment preparations](#) will have all software prerequisites (Docker, Terraform, etc.) preinstalled.
 - The setup scripts include an option to automatically set up a new VPC and will prompt the user to enter the necessary information.
 - For existing VPCs, the following requirements must be met:
 - The VPC must have a DHCP option set with the option to use **AmazonProvidedDNS** for its domain name servers.
 - The VPC must have two private subnets on separate Availability Zones (AZs).
 - If the subnets cannot access the Internet (no NAT gateway attached), set `airgap_install` in `/home/plixer/common/kubernetes/aws.tfvars` to `TRUE`.
 - For additional information on Amazon EKS VPC and subnet requirements and considerations, see [this article](#).
-

Hybrid cloud deployments

When pairing a Plixer ML Engine in AWS with an on-prem Plixer Scrutinizer environment, one of the following methods should be used to enable connectivity between the two before starting the deployment process.

AWS Site-to-Site VPN

Follow these instructions to create an AWS Site-to-Site VPN connection to allow communication between the two deployments.

Direct access via public IP

A public IP address can be used to allow external access to the on-prem Plixer Scrutinizer deployment. However, this will expose the Plixer Scrutinizer environment to the Internet via ports **5432**, **22**, and **443**.

The public IP address must be entered when prompted by the setup scripts. The Internet gateway IP must also be manually added to the Plixer Scrutinizer `pg_hba.conf` file to allow access to Postgres.

After the file has been modified, run the the following command on the Plixer Scrutinizer server to reload the configuration:

```
psql -c "SELECT pg_reload_conf()"
```

Deploying the Plixer ML Engine

Follow these instructions to set up the necessary infrastructure and deploy the Plixer ML Engine:

1. Log in to the Plixer ML VM image using `plixer:plixer`.
2. Accept the EULA, and then configure network settings for the host.
3. SSH to the Plixer ML VM image using the `plixer` credentials set in step 2, and then wait for the setup wizard/scripts to start automatically.
4. Enter the *infrastructure deployment parameters* as prompted.

Note: The requested details are automatically saved to `/home/plixer/common/kubernetes/aws.tfvars`, which also contains *other default parameters* for deploying the Plixer ML Engine Kubernetes cluster. If there are issues with the infrastructure deployment, contact *Plixer Technical Support* for assistance before making changes to the file.

5. Wait as the Kubernetes cluster is deployed (may take several minutes), and then enter the Plixer Scrutinizer SSH credentials when prompted.

After the scripts complete running, navigate to *Admin > Resources > ML Engines* and wait for the engine to show as *Deployed* under its *Deploy Status*. Refresh the page if the status has not updated after a few minutes.

Terraform configuration

The following table lists all required and optional variables in `/home/plixer/common/kubernetes/aws.tfvars`, which are used when deploying the Kubernetes infrastructure for the Plixer ML Engine.

Note: Contact *Plixer Technical Support* before making changes to this file.

| Field name | Description |
|----------------------|--|
| cluster_name | REQUIRED: Name to identify the ML engine cluster/deployment; can only contain the characters a to z (in lowercase), 0 to 9 , and - . |
| 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_name | 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. |
| 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 | OPTIONAL: If airgap_install = TRUE, this bool controls whether or not to create an S3 endpoint in the VPC. |
| create_ecr_endpoint | OPTIONAL: If airgap_install = TRUE, this bool controls whether or not to create an ECR endpoint in the VPC. |
| create_ssm_endpoint | 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. |

2.3.3 Azure (AKS)

After completing the *pre-deployment preparations*, follow the instructions below to set up the necessary infrastructure and deploy the Plixer ML Engine in Azure.

Additional prerequisites for Azure

- Credentials for the Azure user account that will be used for deployment
- A VNet with one subnet for the deployment

Note:

- The Plixer ML VM (the deployment host) deployed as part of the *pre-deployment preparations* will have all software prerequisites (Docker, Terraform, etc.) preinstalled.
 - The Azure user account must be assigned the **owner** role to allow a role to be assigned to the AKS cluster user.
 - VNet details for infrastructure deployment can be defined using the `vnet_addresses` and `new_subnet_cidr` fields in */home/plixer/common/kubernetes/azure.tfvars*.
-

Hybrid cloud deployments

When pairing a Plixer ML Engine in Azure with an on-prem Plixer Scrutinizer environment, one of the following methods should be used to enable connectivity between the two before starting the deployment process.

Azure site-to-site (S2S) VPN

Follow these instructions to create a site-to-site VPN connection to allow communication between the two deployments.

Direct access via public IP

A public IP address can be used to allow external access to the on-prem Plixer Scrutinizer deployment. However, this will expose the Plixer Scrutinizer environment to the Internet via ports **5432**, **22**, and **443**.

The public IP address must be entered when prompted by the `01_azure_infrastructure.sh` and `setup.sh` scripts. The Internet gateway IP must also be manually added to the Plixer Scrutinizer `pg_hba.conf` file to allow access to Postgres.

After the file has been modified, run the the following command on the Plixer Scrutinizer server to reload the configuration:

```
psql -c "SELECT pg_reload_conf()"
```

Deploying the Kubernetes infrastructure

1. Log in to the Plixer ML VM image using `plixer:plixer`.
2. Accept the EULA, and then configure network settings for the host.
3. SSH to the Plixer ML VM image using the `plixer` credentials set in step 2.
4. Exit the automated setup wizard by pressing **Ctrl + C**.
5. Start the Azure CLI and run the following to set up the client and log in:

```
az login
```

6. Define the *infrastructure deployment parameters* in `/home/plixer/common/kubernetes/azure.tfvars` (as described in the file).

Note: `azure.tfvars` may also include fields/variables with factory-defined values (e.g., `kube_version`) for deploying the Plixer ML Engine Kubernetes cluster. Contact [Plixer Technical Support](#) for assistance before making changes to any default value.

7. Navigate to `/home/plixer/common/kubernetes` and run the Kubernetes cluster deployment script:

```
./01_azure_infrastructure.sh
```

8. Verify that the infrastructure was successfully deployed (may take several minutes):

```
kubectl get nodes
```

After confirming the Kubernetes cluster has been correctly deployed, proceed to deploying the Plixer ML Engine.

Deploying the Plixer ML Engine

Once the Kubernetes cluster has been deployed, follow these steps to deploy the Plixer ML Engine:

1. Navigate to the `/home/plixer/ml` directory on the deployment host.
2. Run the Plixer ML Engine deployment script and follow the prompts to set up the appliance:

```
./setup.sh
```

3. When prompted, enter the following Plixer Scrutinizer environment details:

- *Authentication token*
- Primary reporter IP address
- *SSH credentials*

After the script completes running, navigate to [Admin > Resources > ML Engines](#) and wait for the engine to show as *Deployed* under its *Deploy Status*. Refresh the page if the status has not updated after a few minutes.

Terraform configuration

The following table lists all required and optional variables in `/home/plixer/common/kubernetes/azure.tfvars`, which are used when deploying the Kubernetes infrastructure for the Plixer ML Engine.

| Field name | Description |
|---------------------|--|
| cluster_name | REQUIRED: Name to identify the ML engine cluster/deployment; can only contain the characters a to z (in lowercase), 0 to 9 , and - . |
| 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_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. |

2.3.4 Local (single node)

After completing the [pre-deployment preparations](#), follow the instructions below to set up the necessary infrastructure and deploy a local, single-node Plixer ML Engine:

1. Log in to the Plixer ML VM image using `plixer:plixer`.
2. Accept the EULA, and then configure network settings for the host.
3. SSH to the Plixer ML VM image using the `plixer` credentials set in step 2, and then wait for the setup wizard/scripts to start automatically.
4. Enter the following when prompted:

- [Authentication token](#)

- Primary reporter IP address
- [Plixer Scrutinizer SSH credentials](#)

After the scripts complete running, navigate to [Admin > Resources > ML Engines](#) and wait for the engine to show as *Deployed* under its *Deploy Status*. Refresh the page if the status has not updated after a few minutes.

2.3.5 vSphere multi-node cluster

After completing the [pre-deployment preparations](#), follow the instructions below to set up the necessary infrastructure and deploy the Plixer ML Engine in vSphere.

Additional prerequisites for vSphere deployment

- The Plixer ML Engine template must be available in vSphere. Note the path to the template as it will need to be entered when deploying the engine.
- The deployment process will require credentials for a vSphere user with permissions to create VMs and resource groups.

Note: The Plixer ML VM template includes all software prerequisites (Docker, Terraform, etc.).

Deploying the Plixer ML Engine

Follow these instructions to set up the necessary infrastructure and deploy the Plixer ML Engine:

1. Log in to the Plixer ML VM image using `plixer:plixer`.
2. Accept the EULA, and then configure network settings for the host.
3. SSH to the Plixer ML VM image using the `plixer` credentials set in step 2, and then wait for the setup wizard/scripts to start automatically.

4. Enter the *infrastructure deployment parameters* as prompted.

Note: The requested details are automatically saved to `/home/plixer/common/kubernetes/vsphere.tfvars`, which also contains *other default parameters* for deploying the Plixer ML Engine Kubernetes cluster. If there are issues with the infrastructure deployment, contact *Plixer Technical Support* for assistance before making changes to the file.

5. Wait as the Kubernetes cluster is deployed (may take several minutes), and then enter the Plixer Scrutinizer SSH credentials when prompted.

After the scripts complete running, navigate to *Admin > Resources > ML Engines* and wait for the engine to show as *Deployed* under its *Deploy Status*. Refresh the page if the status has not updated after a few minutes.

Terraform configuration

The following table lists all required and optional variables in `/home/plixer/common/kubernetes/vsphere.tfvars`, which are used when deploying the Kubernetes infrastructure for the Plixer ML Engine.

Note: Contact *Plixer Technical Support* before making changes to this file.

| Field name | Description |
|-----------------------|---|
| create_hosts | Whether or not to create vSphere hosts. If FALSE, then the IPs in vm_master_ips 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_proxy | If set to TRUE and offline_install is also TRUE, then the script will attempt to proxy any downloads required for RKE2 Kubernetes setup through the host that ./01_vsphere_infrastructure.sh is running on. |

2.3.6 Expanding storage

Follow the steps below to extend a volume on a Plexier ML Engine appliance.

1. Add/attach a new hard disk to the hardware appliance or VM, and then restart the machine.
2. Navigate to [Admin > Resources > ML Engines](#) in the web interface and wait for the engine's deployment status to switch to *Deployed*.

Note: The ML engine can take up to 30 minutes to fully restart when under heavy load. Refresh the **ML Engines** page every few minutes until the engine is shown as *Deployed*.

3. Log in or SSH to the host using the credentials `plexer:plexer`.
4. Determine the device name of the new disk (usually `/dev/sdb`):

```
lsblk
```

5. Extend the volume that requires additional disk space:

```
/home/plexer/ml/tools/mladmin.sh --extend <DEVICE> <VOLUME>
```

Where `DEVICE` is the device name of the new disk and `VOLUME` is one of the following:

- `root` - root partition
- `sibyl` - models partition (`/SibylData`)
- `db` - database partition (`/var/db`)
- `zookeepers` - Kafka ZooKeeper partition (`/var/kafka/zookeepers`)
- `brokers` - Kafka brokers partition (`/var/kafka/brokers`)

When done, the selected partition will be extended by the full capacity of the newly added disk.

2.4 Initial configuration

After the Plexer Scrutinizer hardware or virtual appliance has been deployed and powered on, proceed with the steps below to configure 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 preliminary setup before rebooting itself.

After the reboot, log in again to start the initial setup 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 `https://IP_ADDRESS_ENTERED` 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

Once the Plixer Scrutinizer web interface is accessible, log in as the `admin` user with the password configured during the *initial appliance setup* to add/register an active license.

Note:

- For AWS AMI deployments, the default password for the web interface `admin` user is the instance ID of the Plixer Scrutinizer instance, which can be copied from the **Instance Summary** view of the AWS console.
- Passwords for the `admin` user and other user accounts can be changed from the [Admin > Users & Groups > User Accounts](#) page at any time.

A Plixer One or Plixer Scrutinizer license key can be obtained by contacting [Plixer Technical Support](#) and providing them with the *Machine ID* displayed under [Admin > Plixer > Scrutinizer Licensing](#). This key should then be pasted into the *License Key* field on the same page and saved.

After a license key has been added, the **Scrutinizer Licensing** page will display details for the active license (validity, appliance/server counts, etc.) and can be used to update the license key when needed.

2.4.3 Changing the default admin password

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

2.4.4 Configuring SSL

As part of the initial setup 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.

This self-signed certificate can later be replaced with a CA-signed certificate if desired.

Note: To learn more about additional certificate-related functions, see [this page](#).

Installing a CA-signed SSL certificate

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

To avoid 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.
2. After acquiring the CA-signed SSL certificate, stop the Apache service:

```
sudo systemctl stop httpd
```

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

```
sudo systemctl start httpd
```

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.

Note: The Plixer Scrutinizer [AMI](#) also uses a self-signed certificate by default. This can be replaced with a new certificate by starting the *scrut_util interactive utility* and running the following from the SCRUTINIZER> prompt:

```
set ssl on
```

Enabling/disabling SSL

If needed, SSL support can be disabled (and later re-enabled) by running the following from the SCRUTINIZER> prompt:

```
set ssl <off|on>
```

The `set ssl on` command can also be used to create a new certificate (with new details) and overwrite the current one.

CONFIGURATION GUIDES

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

3.1 Alarms and events

Plixer 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 Plixer 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

Plixer Scrutinizer automatically manages alarm and event data based on the following life cycle:

1. Plixer 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 the alarm and event data are 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/events are reported.

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

Status

Sets the policy to one of three states:

| Set-ting | Generates Events | Alarm Moni-tor | Stored Database in | Notifications by Pro-file(s) |
|------------------|------------------|----------------|--------------------|------------------------------|
| Active | Yes | Yes | Yes | Yes |
| Store | Yes | No | Yes | Yes |
| Inac-tive | No | No | No | No |

Hint: Setting nonessential 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.

Weight

Assigns each event/violation under a policy a numerical weight for calculating 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

Refer to [this section](#) of the documentation for further information on individual alarm policies, including default timeout settings.

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.

Table 1: Configuration Checklist

| Configuration step | Function/Benefit |
|---------------------------|---|
| Deploy Instance | Deploy the Plixer Scrutinizer hardware/physical or virtual appliance in your environment. |
| Appliance Setup Wizard | From the appliance terminal, 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). |
| SMTP Server | Configure an SMTP server to enable email notifications for alarms 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. |
| 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 for reporting, filtering, and inclusion/exclusion management. |
| External Authentication | Improve your security posture and simplify user management by leveraging AD, LDAP, SSO, Radius, and/or TACACS for user authentication. |
| Data History | Modify historical data retention settings to support your organization's forensics and archiving needs. |
| Security Groups | Populate the default security groups (Firewalls, Core Exporters, Edge Exporters, Defender Probes) to automatically enable flow analytics algorithms and other features for similar devices. |
| Exclusion IP Groups | Verify that the DNS servers, Public WiFi, Network Scanners, DNS Servers, DHCP Servers, and SNMP Pollers IP groups are correctly populated to automatically define recommended exclusions for flow analytics algorithms. |
| Flow Analytics Inclusions | Define additional inclusions and exclusions (including custom security groups and IP groups) necessary for specific flow analytics algorithms. |
| Device/Mapping Groups | Organize exporters into groups to quickly find flow data sources, enable group report filters, and generate customizable network maps. |
| Dashboards | Create/customize one or more dashboards to consolidate frequently accessed information 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. |
| Saved Report Thresholds | Add thresholds to saved reports to proactively watch for specific traffic/behaviors and trigger alarms (and notification profiles/actions) when the specified conditions are met. |
| Schedule Emailed Reports | Set up scheduled email reports to automatically re-run and send important reports as emails to any inbox. |
| Notifications | Create notification profiles that can be assigned to alarm policies to automatically send emails, forward details to your SIEM, or run |

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 required number of Plixer Scrutinizer hardware or virtual appliances following the appropriate *deployment guides* and complete the *initial appliance setup* process.
2. Start an SSH session as the `plixer` user with the appliance that will be used as the primary reporter for the cluster.
3. Launch the *scrut_util* interactive CLI by running:

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

4. At the `SCRUTINIZER>` prompt, register each additional appliance as a remote collector:

```
SCRUTINIZER> set registercollector APPLIANCE_IP
```

5. 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. The web interface for the cluster can be accessed using the IP address or hostname of the primary reporter.

Note:

- When registering remote collectors, it is highly recommended that one appliance/collector should also be assigned the *secondary reporter role*.

```
set registercollector APPLIANCE_IP secondary
```

This appliance can later be promoted to function as the primary reporter (using the `set selfreporter scrut_util` command) if the cluster's original primary reporter becomes unavailable.

- To avoid potential bottlenecks in distributed configurations that include hardware appliances, 10 Gb networking is strongly recommended. If the appliances are geographically dispersed, the WAN link should also support 10G.
-

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 |

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

Certificate management

Run [these scripts](#) to generate certificate signing requests (CSRs) and install the signed certificates to remote nodes in a distributed cluster.

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_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:

- A new license key is not required when promoting a secondary reporter to primary status. The promoted reporter will operate normally with the old license until it expires. However, it cannot register new appliances as collectors and secondary reporters.
- If the original primary Plixer Scrutinizer reporter in a high-availability configuration becomes permanently unavailable, follow these steps to point the Plixer FlowPro probe to the new primary reporter.

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**, do the following:

1. Configure all exporters sending flows to A to also send flows to Y.
2. In the web interface, navigate to **Admin > Resources > Exporters**, and then verify that the selected exporters are correctly sending flows to both collectors.
3. From the **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*.

And if remote collector A is removed from the cluster configuration, it cannot be added back.

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 Plexier Replicator

Plexier Replicator can simplify the process of setting up backup collectors by replicating flows data and forwarding it to multiple destination collectors.

View the Plexier Replicator online documentation or contact *Plexier 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.

Note: To learn more about individual algorithms, see *this appendix section*.

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 applications, it may include FA algorithms that may not apply to certain network configurations. 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 Algorithms](#) page lists *all FA algorithms* and shows whether they are currently enabled (default) or disabled.

Note: Most FA algorithms can also be tuned through *additional settings*, allowing them to be adapted to specific monitoring and detection requirements.

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 page.

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 Algorithms](#) page.

Because alarm-triggering 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 Algorithms](#) 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 settings

Plixer Scrutinizer's flow analytics functions can be further adapted to more unique network and/or security requirements through the configuration options below.

Global settings

The following global settings ([Admin > Settings > Flow Analytics Settings](#)) can be used to enable or configure additional FA-based features:

| Setting | Description |
|------------------------------|---|
| Auto-Enable Defender | When checked, Plixer FlowPro Defender is automatically enabled for algorithms that support it. |
| Jitter by Interface | 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 | Sets the latency value used for record highlighting in <i>Status</i> reports (Default: 75 ms) |
| Share Violations | When checked, allows the system to share details of cyber attacks coming from Internet IP addresses with the Plixer 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 | Controls whether <i>Top X</i> FA algorithms are applied to all exporters or need to be configured individually |

Algorithm settings

In addition to inclusions and exclusions, most FA algorithms have additional settings that control how they are applied to collected flow data. These settings include thresholds for adjusting detection sensitivity and traffic directionality inclusion/exclusion options.

For a full list of algorithm settings, see [this table](#).

Custom reputation lists

The *Host Reputation* FA algorithm is capable of using custom lists in conjunction with Plexier 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-triggering 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.
 - To enhance response/resolution workflows, [create one or more notification profiles](#) and associate them with the appropriate alarm policies.
-

Further tuning

After the initial setup and testing have been completed, flow analytics functions can be further adapted to an environments monitoring and detection requirements through [global](#) and [individual algorithm](#) settings.

3.5 Device groups

Plixer Scrutinizer supports multiple user-defined entity grouping schemes, which can further enhance the way teams monitor, visualize, and derive insights from network data.

IP groups

IP groups can be used to categorize similar (e.g., device type, ownership/department, geolocation, etc.) flow-exporting devices for use in [reports](#), filters, and [FA algorithm exclusion rules](#). Plixer Scrutinizer factory configuration includes default IP groups that should be populated as part of tailoring the system to the environment.

IP group definitions can be created/managed from the [Admin > Definitions > IP Groups](#) page.

Mapping groups

Mapping groups consist of devices that have been grouped together for the purpose of network mapping. Network maps will show network topology up to the interface level (i.e., not including endpoints) and can be tailored to a wide range of use cases using *customizable elements*.

The *Monitor > Network Maps* page is the primary interface for customizing and viewing network maps, while additional management options for mapping groups and map objects can be accessed via their respective pages under *Admin > Settings*.

Security groups

Security groups are device groups that can be used to enable one or more FA algorithms for *exporters of the same type*. The Plixer Scrutinizer factory configuration includes predefined security groups, which can be populated to automatically enable the recommended algorithms for the indicated device type.

Security groups can be created/managed from the *Admin > Alarm Monitor > Security Groups* page.

3.6 Importing data

Plixer Scrutinizer leverages a variety of user-customizable entity/resource labels, definitions, and groupings as part of its data aggregation and reporting functions. These details can be manually configured via the respective *admin views* or imported as a batch operation using the `import` utility.

This section covers the syntax, requirements, and other relevant information for each type of import operation.

Note: The `import` utility can be accessed via the *SCRUTINIZER> interactive prompt* or directly from the shell. The direct shell syntax can also be included in scripts to automatically update Plixer Scrutinizer's databases.

3.6.1 ACL information

To import custom ACL information from a file, execute the following from the *scrut_util interactive shell* (SCRUTINIZER> prompt):

```
import aclfile
```

Direct shell/script syntax

```
scrut_util --import aclfile
```

File requirements

- The file must contain the exact output when the command `show access-list` is run on the exporter.
- The file should be named `acl_file.txt` and saved to the `/home/plixer/scrutinizer/files/` directory.

3.6.2 Application definitions

To import a list of [application definition rules](#), execute the following from the [scrut_util interactive shell](#) (SCRUTINIZER> prompt):

```
import applications <PATH/FILE> [reset]
```

Direct shell/script syntax

```
scrut_util --import applications --file <PATH/FILE> [--reset]
```

File requirements

The file to be imported must be a CSV file.

Using the file `/home/plixer/scrutinizer/files/ipgroup_import.csv` for application rule definitions is recommended.

Definition format

Each application-rule pairing should be in a single line, following the format:

```
'APPLICATION NAME',RULE
```

Additional notes

- Rules can be defined as any of the following:
 - Subnets
 - Single IP address
 - IP address ranges
 - Wildcard masks
 - Child rules (must be defined first)
 - Port and protocol
- For an application definition to be valid, it must include **at least** one port rule **and** one rule of any other type. The import file may include applications that do not meet this requirement, but they will not be considered a *defined application* by Plixer Scrutinizer.
- Passing the `reset` option will delete all existing application definitions/rules before the import operation.
- If the `reset` option is not used, imported rules will be added to the specified application if it already exists.
- Each import operation supports up to 100,000 application rule definitions.

Definition examples

Rule types:

```
'Application subnet rule',10.0.0.0/8
'Application single IP rule',10.1.1.1
'Application IP range rule',10.0.0.1-10.0.0.42
'Application wildcard mask rule',10.0.0.1/0.255.255.0
'Parent application with a child rule', 'My Child Application Rule'
'Application port and protocol rule',0-65535/256
```

3.6.3 ASN definitions

To import a list of *custom ASN definitions*, execute the following from the *scrut_util interactive shell* (SCRUTINIZER> prompt):

```
import asns <PATH/FILE> [DELIMITER]
```

Direct shell/script syntax

```
scrut_util --import asns --file <PATH/FILE> [--delimiter <DELIMITER>]
```

File requirements

The file to be imported must be a CSV file, and the path provided must be relative to the `home/plixer/scrutinizer/` directory. The file's name should only include lower-case letters.

Definition format

Each ASN definition should be in a single line, following the format:

```
'AS_NUMBER',AS NAME,AS Description,IP_NETWORK(S)
```

Additional notes

- The optional DELIMITER parameter can be used to replace `` ` ` (space) for separating individual IP networks if the contents of the import file are formatted differently.
- , (comma) cannot be used as a custom delimiter, as it is reserved for separating elements in the definition.

Definition examples

```
213,My ASN,what a great autonomous system,10.0.0.0/8 192.168.0.0/16
214,Your List,this system is only meh,11.0.0.0/8
```

3.6.4 Custom hostnames

To import a list of *custom hostname assignments*, execute the following from the *scrut_util interactive shell* (SCRUTINIZER> prompt):

```
import hostfile
```

Direct shell/script syntax

```
scrut_util --import hostfile
```

File requirements

The file should be named `hosts.txt` and saved to the `/home/plixer/scrutinizer/files/` directory.

Definition format

Each definition should be in a single line, following the format:

```
IPv4orIPv6ADDRESS HOSTNAME DESCRIPTION
```

Additional notes

- This command will alter the Plexier Scrutinizer database tables and should be used with caution.
- The description element in the definition is optional.

Definition example

```
10.1.1.4    my.scrutinizer.rocks    The best software in my company
```

3.6.5 Device GPS details

To import a list of *device/object* latitude and longitude details for a specified *geographical network map*, execute the following from the *scrut_util interactive shell* (SCRUTINIZER> prompt):

```
import csv_to_gps <PATH/FILE> <GROUP NAME|GROUP_ID> [create_new] [FORMAT]
```

Direct shell/script syntax

```
scrut_util --import csv_to_gps --file <PATH/FILE> --group <GROUP_NAME|GROUP_ID> [--create_new] [--file_format <FORMAT>]
```

File requirements

The file to be imported must be a CSV file, and the path provided must be relative to the `home/plixer/scrutinizer/` directory.

Definition format

Each set of details should be in a single line, following the format:

```
IP_ADDRESS,LATITUDE,LONGITUDE
```

Additional notes

- The imported GPS details are only assigned to objects for the specified device/mapping group. If the devices are assigned to other groups, they will retain the GPS details configured for those groups.
- The optional `FORMAT` parameter can be used to override the default `ip,lat,lng` element formatting in case the contents of the import file are formatted differently (e.g., `ip,lng,lat`).
- If the `create_new` option is used, *objects will be created* for devices in the import file that are not currently assigned to the specified device group.

Definition examples

```
10.169.1.3,37.7749,122.4194
192.168.6.1,40.7128,74.0059
```

3.6.6 Device/mapping group assignments

To import a list of *device/mapping group assignments*, execute the following from the *scrut_util interactive shell* (SCRUTINIZER> prompt):

```
import csv_to_membership <PATH/FILE> <TYPE> [FORMAT]
```

Direct shell/script syntax

```
scrut_util --import csv_to_membership --file <PATH/FILE> --grouptype <TYPE> [-  
↪-file_format <FORMAT>]
```

File requirements

The file to be imported must be a CSV file, and the path provided must be relative to the `home/plixer/scrutinizer/` directory.

Definition format

Each assignment should be in a single line, following the format:

```
IP_ADDRESS, GROUP_NAME
```

Additional notes

- The `TYPE` parameter specifies the device/mapping group type for any groups that will be created as part of the import operation. Valid values are `plixer` (for spatial maps) and `google` for geographical maps.
- The optional `FORMAT` parameter can be used to override the default `ipaddr, group` element formatting in case the contents of the import file are formatted differently (e.g., `group, ipaddr`).

Definition examples

```
10.169.1.3,Routers
192.168.6.1,Firewalls
```

3.6.7 Interface details

To import a list of custom *interface details* to use for displaying utilization, threshold alerts, and other Plixer Scrutinizer functions, execute the following from the *scrut_util interactive shell* (`SCRUTINIZER>` prompt):

```
import ifinfo <PATH/FILE> [DELIMITER]
```

Direct shell/script syntax

```
scrut_util --import ifinfo --file <PATH/FILE> [--delimiter <DELIMITER>]
```

File requirements

The file to be imported must be saved to the `home/plixer/scrutinizer/files/` directory.

Definition format

Each set of details should be in a single line, following the format:

```
INBOUND_SPEED,OUTBOUND_SPEED,NAME,HOST_IP,INDEX_NUMBER
```

Additional notes

- This command will alter the Plixer Scrutinizer database tables and should be used with caution.
- The optional `DELIMITER` parameter can be used to replace `,` (comma) for separating elements in each set of details if the contents of the import file are formatted differently.

Definition examples

```
10000000,10000000,WAN_Interface_1,192.168.1.2,2
20000000,20000000,WAN_Interface_1,192.168.1.4,11
```

3.6.8 IP group inclusions

To import a list of *IP group inclusion definitions*, execute the following from the *scrut_util interactive shell* (SCRUTINIZER> prompt):

```
import ipgroups <PATH/FILE> [reset]
```

Direct shell/script syntax

```
scrut_util --import ipgroups --file <PATH/FILE> [--reset]
```

File requirements

The file to be imported must be a UTF8-encoded CSV file.

Definition format

Each inclusion definition should be in a single line, following the format:

'IP GROUP NAME', INCLUSION_RULE

Additional notes

- Rules can be defined as any of the following:
 - Subnets
 - Single IP address
 - IP address ranges
 - Wildcard masks
 - Child groups (must be defined first)
- Passing the `reset` option will delete all existing IP group definitions before the import operation.
- If the `reset` option is not used, IP addresses covered by an imported inclusion rule will be added to the specified IP group if it already exists.
- Because each line can contain only one rule, an IP group containing multiple single IP addresses will need to be defined using a separate definition/line for each address. Multiple rules in separate lines for the same IP group are also supported.
- Each import operation supports up to 100,000 IP group inclusion definitions.

Definition examples

Rule types:

```
'Subnet Group', 10.0.0.0/8
'Single IP Group', 10.1.1.1
'IP Range Group', 10.0.0.1-10.0.0.42
'Wildcard Mask Group', 10.0.0.1/0.255.255.0
'Parent/Child Group', 'My Subnet'
```

Multiple single addresses:

```
'Sales', 10.1.1.1
'Sales', 192.168.3.4
'Sales', 10.3.1.2
```

Multiple non-single-IP rules:

```
'New IP Group',10.0.0.1-10.0.0.42  
'New IP Group','My Subnet'
```

3.7 Plexier ML Engine

When deployed as part of a Plexier One Enterprise environment, the Plexier ML Engine applies anomaly and threat detection techniques to the network data collected by Plexier Scrutinizer.

Note: To learn more about Plexier One Enterprise licensing options, contact [Plexier Technical Support](#).

This configuration guide introduces the capabilities of the Plexier ML Engine and provides further information on managing the various settings controlling its functions and behavior.

3.7.1 About the Plexier ML Engine

Once *deployed* and configured, the engine is able to ingest flow data through Plexier Scrutinizer and *apply multiple machine learning techniques* to identify potentially problematic activity on the network.

The Plexier ML Engine has several key functions that enable intelligent, multi-layered anomaly and threat detection in a Plexier One Enterprise deployment:

Comprehensive network behavior modeling

Leveraging the large volumes of flow data collected by Plexier Scrutinizer, the engine is capable of building behavioral models encompassing network activity at any scale. It can then learn to recognize deviations and suspicious activity, such as data accumulation/exfiltration, tunneling, and lateral movement, that may indicate an attack on the network.

Accessible behavioral insights for network assets

After being alerted to anomalous behavior, network and security teams can drill down into the associated hosts, IP address groups, and/or exporter interfaces to better understand the details of their involvement in the reported detection.

Highly configurable ML modeling

The Plexier ML Engine monitors network activity based on user-customizable *dimensions* and *inclusion/exclusion rules*. Consistently repeated traffic patterns, asset/group importance, and data seasonality are all taken into consideration as well, resulting in models that are uniquely tailored to each environment.

ML-based malware detection

Using pre-trained classification models, the engine is able to recognize generic activity patterns that are associated with common classes of malware, including command and control, remote access trojans, and exploit kits. This adds another layer of protection to further reduce risk and mean time to resolution (MTTR) when threats are detected.

Continuous observation and learning

As it ingests additional flow data, the Plixer ML Engine updates its behavior models based on a schedule that defines weekdays, weeknights, and weekends to account for changes in legitimate activity patterns and improve recognition of advanced threats that attempt to disguise their behavior.

Configuration options

After being deployed, the Plixer ML Engine adapts its monitoring functions to the assets and traffic data observed to be available in the Plixer Scrutinizer environment. This allows it to function optimally in common enterprise scenarios. However, the engine can also be further tailored to more unique network and security requirements.

Dimensions and inclusions

To ensure that its behavior models represent only relevant network activity, the Plixer ML Engine supports custom dimension definitions and inclusion/exclusion rules.

- *ML dimensions*: Define traffic (based on protocol and port used) by host or exporter for ingestion by the engine
- *ML rules*: Define inclusion or exclusion rules for network assets to be monitored by the engine

The default configuration for the Plixer ML Engine includes recommended dimension definitions, which are used to automatically select suitable data sources as inclusions. After the engine is deployed and set up, the *ML Dimensions* and *ML Rules* management views under *Admin > Alarm Monitor* in the web interface should be compared against the monitoring requirements for the environment and updated if necessary.

Global settings

The global settings under *Admin > Settings* in the Plixer Scrutinizer web interface can be used to configure parameters for certain ML functions and behaviors across all engines in an environment.

- *ML AD Users*: Add/edit Microsoft Azure account credentials for AD Users UEBA integration
- *ML Alerts*: Manage thresholds for alerts related to engine vitals and sensitivity for Microsoft Office 365 detections

- *ML Data Limits*: Manage model and host count limits to use for training and prediction
- *ML Training Schedule*: Manage behavior model seasonality settings

With the exception of adding an Azure account, leaving the above settings at their default values is recommended for new Plexier ML Engine deployments.

Engine settings

After an engine has been *fully deployed*, its resource utilization settings can be modified via the tray in the *ML engine management page*.

These settings are applied on a per-engine basis and can be used to optimize pod counts and process resource utilization based on an engine's expected workload.

3.7.2 Managing inclusion and exclusion rules

To support more diverse network topologies, the Plexier ML Engine can be uniquely tailored to its environment using custom rules defining inclusions and exclusions for its functions. These rules can be managed from the *Admin > Alarm Monitor > ML Rules* view of the Plexier Scrutinizer web interface.

Inclusion rules

An inclusion rule defines either a network address (hosts/subnets) or exporter interface as a network data source for the Plexier ML Engine. Each rule also includes a sensitivity setting (see below) that is applied to the asset specified.

Malware detection, which uses pre-trained classification models to recognize generic malware behaviors, can also be enabled for individual inclusions.

Inclusion sensitivity

An inclusion's sensitivity setting can be used to tune the engine's tolerance for behavioral deviations for the host/subnet or exporter interface. Lowering the sensitivity setting for an asset will cause even minor deviations to be reported as detections, resulting in a higher volume of alarms. Conversely, increasing the sensitivity will allow for greater deviation, which translates to fewer detections reported.

When defining inclusions, the sensitivity setting should be left at its default value. After a period of 7 days (recommended), if too many unwarranted detection alarms are triggered, the sensitivity can be increased to the next level.

Exclusion rules

Exclusion rules can be used to ignore one or more ML-driven detections for traffic originating from a specified source and/or bound for a specified destination.

If expected traffic/activity triggers alarms, one or more exclusion rules should be created to exempt the sources and/or destination addresses from the detections being reported.

Recommendations

As part of the Plixer ML Engine's initialization, inclusion rules are automatically created for the 20 most suitable network assets (hosts and and exporters/interfaces) based on its *default dimension definitions*. If necessary, additional rules should be created to cover all assets associated with critical/sensitive network activity ("crown jewel" assets) and hard-to-monitor traffic (e.g., IoT devices, operational technology, etc.).

The following resources are examples of network assets that are highly recommended for inclusion:

- AD servers
- DB servers
- DBS servers
- DHCP servers
- Web servers
- Source code repositories
- Object repositories
- FTP servers

If there are assets whose typical behavior is being reported as anomalous/suspicious, exclusion rules should be defined to exempt the traffic from superfluous detections.

3.7.3 Managing dimensions

The Plixer ML Engine's feature dimension list defines the protocols and ports to be observed on the network assets defined by its *inclusion/exclusion rules*. These dimensions are used by the engine to build its behavior models, which are used to report *asset behavior insights*, as well as deliver anomaly and threat alerts via the *Plixer Scrutinizer alarm monitor*.

Dimensions are managed from the *Admin > Alarm Monitor > ML Dimensions* view of the Plixer Scrutinizer web interface.

Dimension configuration

An ML dimension is defined by the following parameters:

- Inclusion/asset type the dimension applies to (host/subnet or exporter interface)
- Template field to use for grouping (sourceipaddress or destinationipaddress, host/subnet dimensions only)
- Aggregation method to use (octetdeltacount or packetdeltacount)
- Traffic port used

Note: A feature dimension is only observed for traffic associated with the type of *inclusion* (host/subnet or exporter interface) it was defined for.

Dimensions can be configured to apply to all or only internal traffic matching the definition. They can also be disabled and re-enabled as necessary.

Recommendations

Once deployed, the Plixer ML Engine defaults to Plixer's recommended dimension definitions, which are based on the traffic in typical enterprise environments.

These default definitions should be reviewed and, if necessary, additional dimensions should be defined to monitor critical network services that are most often the target of attacks, such as:

- Authentication - Kerberos, NTLM
- Domain services - LDAP, DNS, DHCP
- File sharing services - SMB, NFS, CIFS
- Remote connectivity - SSH, Telnet, RDP, VNC, FTP
- Email protocols - SMTP, POP3
- Inter-process communication - ICMP
- Application protocols - HTTP, HTTPS
- Others - DB services, third-party APIs (especially those that connect to the Internet)

3.7.4 Global ML settings

The global ML settings under [Admin > Settings](#) can be used to customize various machine learning function/behavior settings, including training parameters and alert sensitivities/thresholds.

AD Users

The Plixer ML Engine is also able to ingest user activity data and access logs and alert users to anomalous behavior through user and entity behavior analytics (UEBA) detections.

UEBA alerts for Active Directory users can be enabled by adding the credentials for a Microsoft Azure account that is configured to store AD user sign-in logs under [Admin > Settings > ML AD Users](#).

Alerts

There are three categories of alert settings that can be adjusted under [Admin > Settings > ML Alerts](#):

Microsoft Office 365 alerts

These sensitivity values adjust the magnitude of deviation from typical behavior that will trigger the corresponding alerts. A higher value allows for greater deviation, resulting in fewer alerts for the corresponding activity.

- *Logon Sensitivity*: Unusual volumes of Office 365 login events
- *Unique Source Sensitivity*: Traffic coming from unusual numbers of unique hosts
- *Unique Location Sensitivity*: Traffic coming from unusual numbers of unique locations

Like [inclusion sensitivities](#), these values should only be adjusted after assessing the accuracy of alarms/detections.

System vitals alerts

These thresholds control alerts and other actions related to high utilization of the Plixer ML Engine's resources.

- *CPU/RAM/Disk Alert Threshold*: Percentages at which a high utilization alert for the corresponding resource is triggered

- *Disk Reclaim Threshold*: Disk utilization percentage at which the Plixer ML Engine will attempt to delete old indexes from Elasticsearch

Initially, these thresholds should be left at their default values. If *alarms* are triggered, run an *ML Engine CPU*, *ML Engine Memory*, and/or *ML Engine Storage* report to assess whether threshold(s) need to be increased (for temporary spikes) or additional resources should be allocated to the engine (for sustained high utilization).

Kafka lag thresholds

These thresholds manage the amount of latency tolerated by the Kafka engine before the corresponding lag alert is triggered.

- *Kafka Netflow Lag Threshold*: Alerts for flow ingestion latency
- *Kafka K-means Lag Threshold*: Alerts for prediction latency
- *Kafka Alerts Lag Threshold*: Alerts triggered by automated process reconnaissance
- *Kafka Training Data Lag Threshold*: Alerts for behavior modeling latency
- *Kafka UEBA Lag Threshold*: Alerts for user and entity behavior analytics (UEBA) data latency

If *alarms* are triggered, run an *ML Engine Kafka Lag* report to determine whether there is a need to scale up the engine's resources.

Data limits

The Plixer ML Engine's data limit settings manage the maximum numbers of behavior models and hosts used for network/user activity patterns and prediction. The initial values set are based on the engine's default resource configuration, but they can be adjusted under [Admin > Settings > ML Data Limits](#).

If there are *alarms* associated with these limits, the engine may need to be provisioned with additional resources to sustain the current volume of inclusions.

Note: To check the utilization for the current model limit, run an *ML Engine Model Count* report.

Training schedule

The settings under [Admin > Settings > ML Training Schedule](#) determine the seasonality applied when the Plixer ML Engine ingests traffic data, allowing it to distinguish between network activity during and outside of an organization's hours of operation.

The engine defaults to business hours of 8 am to 6 pm, from Monday to Friday. These settings can be changed after deployment if necessary.

3.7.5 ML engine settings

The [engine management page](#) provides access to configuration options for individual ML engines in the environment. These settings relate primarily to resource utilization for an engine's core processes and can be used to tailor resource allocations by service/process to each engine's expected workload.

The following settings can be accessed by selecting *Settings* in the engine configuration tray:

- *Ingestion Replica Count*: Number of pods to deploy for the ingestion service
- *Train Anomaly Detection Replica Count*: Number of pods to deploy for the anomaly detection training service
- *Ingestion Minimum CPU*: Minimum number of CPU cores that can be dedicated to the ingestion service
- *Ingestion Maximum CPU*: Maximum number of CPU cores that can be dedicated to the ingestion service
- *Ingestion Minimum Memory*: Minimum amount of memory (in GB) that can be dedicated to the ingestion service
- *Ingestion Maximum Memory*: Maximum amount of memory (in GB) that can be dedicated to the ingestion service
- *Elasticsearch memory*: Amount of memory (in GB) to dedicate to Elasticsearch
- *Elasticsearch Minimum CPU*: Minimum number of CPU cores that can be dedicated to Elasticsearch
- *Elasticsearch Maximum CPU*: Maximum number of CPU cores that can be dedicated to Elasticsearch

Note:

- The Kibana UI can be enabled from the same tray and will be deployed alongside Elasticsearch if toggled on.
 - Collectors assignments can also be configured on a per-engine basis from the main configuration tray.
-

3.8 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.

3.8.1 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 Plixer 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 Plixer 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 | | | | | | | | |
| 66 | | 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 | 3. Configuration Guides | | |
| | | | | | | | 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 |
| Basic Tuple Analysis | 5.85 | 3.3 | <ul style="list-style-type: none"> • <i>DNS Hits</i> • <i>FIN Scan</i> • <i>Host Reputation</i> • <i>ICMP Destination Unreachable</i> • <i>ICMP Port Unreachable</i> • <i>Large Ping</i> • <i>Odd TCP Flags Scan</i> • <i>P2P Detection</i> • <i>Packet Flood</i> • <i>Ping Flood</i> • <i>Ping Scan</i> • <i>Reverse SSH Shell</i> • <i>RST/ACK Detection</i> • <i>Slow Port Scan</i> • <i>SYN Scan</i> • <i>TCP Scan</i> • <i>Network Transports</i> • <i>UDP Scan</i> • <i>XMAS Scan</i> |
| Application Analysis | 0.25 | 0.1 | <ul style="list-style-type: none"> • <i>Protocol Misdirection</i> |
| Worm Analysis | 0.5 | 0.2 | <ul style="list-style-type: none"> • <i>Lateral Movement Attempt</i> • <i>Lateral Movement</i> |
| FlowPro DNS Exfiltration Analysis | 0.5 | 0.2 | <ul style="list-style-type: none"> • <i>DNS Command and Control Detection</i> • <i>DNS Data Leak Detection</i> |

Note:

- 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.
 - 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 > System Performance](#) 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.

3.8.2 Storage

The [Admin > Resources > System Performance](#) page of the web interface summarizes disk utilization for individual collectors in a Plexier 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 Plexier Scrutinizer collector's disk use and provides instructions for anticipating additional storage needs.

Data retention

Plexier Scrutinizer's [data history settings](#) can be used to adjust how long Plexier Scrutinizer stores [aggregated flow data](#), alarm/event details, and other data. 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.

For more accurate and detailed projections of disk space requirements based on specific data retention settings, the following [database size calculator](#) can be accessed from the data history settings tray:

Test data retention times

48

168

30

30

60

✕

✓

1 MIN (HRS)

5 MIN (HRS)

30 MIN (DAYS)

2 HR (DAYS)

12 HR (WEEKS)

Predicted HD utilization based on current settings

| COLLECTOR | 1 MIN | 5 MIN | 30 MIN | 2 HR | 12 HR | DATA SIZE | DISK SIZE |
|---------------|-------|-------|--------|------|-------|-----------|-----------|
| 10.42.100.155 | 860MB | 164MB | 190MB | 60MB | 210MB | 1.48GB | 2.09GB |
| 10.42.100.156 | 519MB | 69MB | 88MB | 37MB | 240MB | 953.17MB | 1.35GB |
| 10.42.100.157 | 564MB | 99MB | 119MB | 46MB | 254MB | 1.08GB | 1.52GB |

Current HD utilization per interval

| COLLECTOR | 1 MIN | 5 MIN | 30 MIN | 2 HR | 12 HR | DATA SIZE | DISK SIZE |
|---------------|-------|-------|--------|------|-------|-----------|-----------|
| 10.42.100.155 | 806MB | 153MB | 173MB | 56MB | 47MB | 1.24GB | 61GB |
| 10.42.100.156 | 484MB | 64MB | 81MB | 36MB | 88MB | 753.05MB | 61GB |
| 10.42.100.157 | 516MB | 92MB | 109MB | 44MB | 95MB | 855.56MB | 61GB |

The calculator shows both current and predicted disk usage for each historical flow data interval based on the retention times entered. Details are shown by collector, with total predicted usage and total storage currently available also included.

Note:

- More detailed storage utilization information can be accessed by drilling into a collector from the [Admin > Resources > System Performance](#) page.
- 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.

Auto-trimming

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.

Auto-trimming 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.

Host indexing

When *host indexing* is enabled, it may become necessary to allocate additional storage, *CPU cores, and RAM* to Plixer Scrutinizer collectors.

Host to host indexing can have a significant impact on disk utilization due to the two types of records stored:

- 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 *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.

3.8.3 Distributed cluster sizing

Distributed configurations consisting of one primary reporting server and multiple remote collectors allow Plixer 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 Plixer 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:

- The CPU core and RAM allocations above are exclusive of the base resource requirements for the virtual appliance.
 - 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.
-

3.8.4 Plixer ML Engine

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

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

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

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 |
| 76 | | 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 |
| 78 | | 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 |
| 80 | | 3. Configuration Guides | | | | | | | | | | |
| | 70k | 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 or Enterprise), 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 Enterprise, 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 or Enterprise) 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 *user groups*) 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 Enterprise** 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 then 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 or Enterprise) 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 > 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 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 or Enterprise), 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 then 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 or Enterprise), 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 *user groups*.

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 or Enterprise)—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 user group access</i> | Shows or hides the dashboard for user groups |

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 user group.

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 user groups:
 - a. Navigate to **Admin > Users and Groups > User Groups** 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 user group 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 user group.

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

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 Enterprise, 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 Enterprise 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 Enterprise 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 Enterprise 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 Enterprise. Contact [Plixer Technical Support](#) to learn more.

Workflows

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

4.1.8 Capacity planning

Through Plixer Scrutinizer, Plixer One Enterprise 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 Enterprise 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 Enterprise 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 Enterprise 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 or Enterprise) 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 AWS, Azure, or OCI virtual network, 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*, *Plixer ML Engine inclusion rules*, 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 flow logs

Setting up Azure flow log ingestion in Plixer Scrutinizer requires an Azure Blob Storage container that is *correctly configured* and receiving log data from the virtual network. 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
- All Details
- Resource IDs

Hint: To view only report types that apply to Azure 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 Enterprise 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 Enterprise'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 Plixer 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 Enterprise can provide general malware detections and extract additional value from the same flow data.

Overview

The Plixer 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 Plixer 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 Plixer 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 *Plixer 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 Enterprise'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 Plixer 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 within 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 Enterprise 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 Enterprise 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 Enterprise 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 Enterprise 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 Plexier 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 Plixer 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 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 Enterprise 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 Enterprise 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 Plixer ML Engine) as violators.

Note:

- 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/source](#) it belongs to.
 - Because the *Lateral Movement* FA algorithm references existing lateral movement attempts for its detections, its scope can be customized by [specifying traffic coverage](#) (*external to internal*, *internal to external*, or *internal to internal*) for the *Lateral Movement Attempt* algorithm. E.g., if internal-to-internal traffic is disabled for the *Lateral Movement Attempt* algorithm, there will be no detections for internal-to-internal traffic under the *Lateral Movement* algorithm.
-

Workflows

The following workflows show how lateral movement detections in Plixer 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 Enterprise 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 ping 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 accessed by pointing any supported browser to https://SCRUTINIZER_ADDRESS/ui/, after the *server has been deployed and set up*.

This section introduces the different pages and views of the web interface and provides detailed instructions for leveraging their associated functions.

5.1.1 UI overview

The Plixer Scrutinizer web interface enhances NetOps and SecOps workflows by providing access to a comprehensive feature set designed to transform raw flow data into fully contextualized intelligence for modern network teams.

Pages in the web interface are divided into four general categories that correspond to the most essential NetOps and SecOps workflows, with a set of admin views/trays for environment configuration and management.

Hint:

- Plixer Scrutinizer users can toggle between the persistent header bar and the collapsible sidebar using the *Slim Navigation* option in the **Admin > Users & Groups > User Accounts > Preferences** tray.
 - Click the **Help (?)** button in the header of any page to access the Plixer Scrutinizer online documentation at any time.
-

| <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 if the host(s) has been seen on the network and investigate activity linked to it | <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.

5.1.2 Monitor

The **Monitor** views of the Plixer Scrutinizer web interface consist of high-level overviews that can serve as ideal starting points for various NOC and SOC workflows.

This section discusses the different functions associated with each of the three main **Monitor** views.

Alarm Monitor

The **Alarm Monitor** page is Plixer Scrutinizer's main interface for monitoring alarm policy violations. The page is divided into three subtabs to support varied avenues for investigating performance issues and suspicious activity.

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

Policies

The **Monitor > Policies** tab/view is the default Alarm Monitor view and can be used to investigate alarms within the specified time period based on the alarm policy violated.

The overview table can be set to include any of the following columns via the **Available Columns** button:

- **Severity:** Distribution of individual events under the policy based on severity
- **Risk:** Aggregated risk level
- **Events:** Total number of violating events under the policy
- **Violators:** Total number of hosts observed as violators under the policy
- **Targets:** Total number of hosts observed as targets under the policy
- **First Observed:** Timestamp of the first violating event within the specified time period
- **Last Observed:** Timestamp of the most recent violating event within the specified time period
- **Category:** Policy category
- **Technology:** Plixer One component where the alarm originated

The host counts in the **Violators** and **Targets** columns also function as shortcuts to pivot to the [Hosts](#) view with a filter for the policy applied.

Note:


- Risk information requires [Plixer Endpoint Analytics](#) integration to be enabled. To learn more about Plixer Endpoint Analytics integration in Plixer Scrutinizer, see [this section](#) of this documentation.
 - For a full list of alarm policy categories and violation descriptions, see [this table](#).
-

Editing policy settings

To edit the [settings](#) of the policy for an active alarm, select **Edit Policy** from the three-dot menu in the list/table.

This will open the settings tray in the [alarm policy management view](#), where the policy's weight, timeout, and state can be modified. [Notification profiles](#) can also be created and assigned to the policy from this tray.

Inspecting hosts

Clicking the  icon in the *Violators* or *Targets* column of the table opens a tray listing violating and targeted hosts involved in the alarm. This tray can be used to select one or more hosts to apply as filters or view [alarm details](#) for any of the hosts involved.


Alternatively, clicking on the host count in the *Violators* or *Targets* column opens the Alarm Monitor [Hosts tab](#) with a filter for the policy applied.

The tray also includes toggles to hide/show system policy violations and acknowledged events in the active alarm list.

Managing exclusions

To add or remove exclusions for an active alarm policy, select **Manage Exclusions** from the three-dot menu in the list/table.

For *Scrutinizer* alarm policies (indicated in the *Technology* column), this will open the [FA algorithm management view](#), from where exclusions can be added to or removed from the algorithm driving the policy. For *Plixer Machine Learning* policies, the option will open the [ML rules management view](#) instead.

Individual hosts can also be added to FA algorithm or ML detection exclusion lists by opening the violators/targets tray and clicking the  icon for one or more hosts.

Alarm summary

Clicking on a policy in the main list opens the summary/details view for the alarm, which includes a chart/timeline summarizing observation details and a list of artifacts for separate events/violations under the same policy.

The following visualizations can be selected from the *View* dropdown:

- **Events Scatter Plot** - Shows distribution of the events and observations
- **Events Timeline** (default) - Shows the individual events and their durations in a timeline for the specified time period
- **Entities** - Shows observation distribution among top violators, IP groups, and targets

Note: Plixer Scrutinizer aggregates continuous or consecutive observations within the policy's *Timeout* setting as a single event. See [this page](#) on the alarm/event life cycle for further details.

Event list

The event list of the alarm summary view can be used to drill into the artifacts for discrete events/violations within the specified time period. The summary table lists total number of observations aggregated as well as the basic details (severity, hosts, etc.) for each event.

Hint: Mouse over the graph icon in the event list for additional shortcuts/options (varies by policy).

Click on an artifact to open a tray containing the [full details](#) for the event:

- Severity
- Start/end timestamps
- Most recent event message generated
- All hosts observed as targets
- All hosts observed as violators
- All events with matching violating criteria

In the tray, clicking on the link icon for target or violator opens the [host details view](#), where the details for all alarms associated with the host can be investigated. Details for other events with the same violating criteria (based on the alarm policy) can also be viewed in a secondary tray by clicking the view (eye) icon.

Auto-Investigate policy

The *Auto-Investigate* alarm policy reports sequential incident/event chains wherein each targeted host becomes the next violator in the sequence. Each chain includes all discrete events starting from the initial incident and ends when the target cannot be confirmed as the next violator.

When an Auto-Investigate alarm is active, its summary view will list all incident chains (aggregated by the initial violating host) instead of individual events.

Investigation details

Clicking the microscope icon in the list/table opens the investigation subview for the selected initial violator, which can be used to inspect the following information for all incident chains linking back to it:

- All incident chains with the same initial violator, including violators, targets, and exact timelines
- Visualized links between violators, policies, and targets
- Event distribution over time
- Event, target, and violator counts for all policies violated
- Number of policy violations, linked event violator counts (including itself), and roles for all hosts

The policy and host lists also link back to their respective Alarm Monitor views for further investigation and cross-referencing.

Hosts

The **Monitor > Policies** tab can be used to investigate alarms within the specified time period based on a target or violating host.

The overview table can be set to include any of the following columns via the **Available Columns** button:

- **Severity:** Distribution of individual events under the policy based on severity
- **Behavior:** Host behavior information (Click the icon to view behavior summary or drill into the [host behavior subview](#).)
- **Risk:** Endpoint risk level (Click the icon to view endpoint details.)

- **Country/Group:** IP group or country associated with the host
- **As Target:** Total number of events with the host as a target
- **As Violator:** Total number of events with the host as a violator
- **Policies:** Total number of policy violations involving the host as a target or violator
- **First Observed:** Timestamp of the first violating event involving the host within the specified time period
- **Last Observed:** Timestamp of the most recent violating event involving the host within the specified time period

The three-dot icon/menu can be used to access the host information summary tray or pivot to any report supported by the host.

Note:

- *Behavior information* requires a Plixer One Enterprise license.
 - Risk information requires *Plixer Endpoint Analytics* integration to be enabled. To learn more about Plixer Endpoint Analytics integration in Plixer Scrutinizer, see *this section* of this documentation.
 - The **Country/Group** column will display IP groups for internal hosts and countries for external addresses. Addresses can be designated as internal or external as part of *IP group definitions*.
-

Host details

Clicking on a hostname/address in the main list opens the host details page, which includes an overview pane and three (four if the host is an exporter) subviews with detailed insights related to the host's activity.

Note: If *Plixer Endpoint Analytics* integration is enabled, the overview pane will include a section with additional endpoint information and a link to the corresponding Plixer Endpoint Analytics view.

Traffic

The host traffic subview can be used to inspect a host's activity based on its communications with other hosts and/or IP groups.

This subview visualizes activity data for the host using the following charts:

- An activity timeline showing the inbound (green) and outbound (blue) rates over the specified time period in an activity timeline
- A traffic distribution chart of source IP groups where this host is the destination
- A traffic distribution chart representing the host's activity by *defined application* used
- A traffic distribution chart of destination IP groups where this host is the source

Each chart also includes a shortcut button to run a filtered report to break down the host's activity in greater detail.

Behavior

The host behavior subview can be used to investigate a host that has been observed by the Plixer ML Engine to be exhibiting anomalous behavior.

Host behavior insights for the selected *ML dimension* are summarized in the following:

- A timeline showing the deviation criteria (e.g., bytes, IP address count, etc.), magnitude (based on the host's typical activity patterns), and threshold for the selected dimension
- A table/list of timestamps and details for individual behavior deviations

To see behavior information for a different feature dimension, use the dropdown and select another dimension with an anomalous behavior count.

Further investigation is recommended for hosts with deviation magnitudes exceeding the indicated threshold.

Note:

- Behavior data will only be available for hosts that are covered by the Plixer ML Engine's *inclusion rules* and have exhibited anomalous behavior.
 - Behavior modeling and other Plixer ML Engine functions require a Plixer One Enterprise license. Contact *Plixer Technical Support* to learn more.
-

Alarms

The host alarms subview can be used to investigate alarms in which the host was involved as a target and/or violator.

This subview includes two overviews of all *unacknowledged alarms* associated with the host:

- A timeline showing individual *events* by alarm policy violated
- A summary table (similar to the main Alarm Monitor *policies view*) with details for all policies with violations involving the host

Drilling in from the summary table opens the *alarm details* view for the policy, where event artifacts can be inspected individually.

Interfaces

The host interfaces subview consists of a table listing all interfaces on a flow-exporting device along with their inbound and outbound activity details.

Note: Inbound and outbound activity details use rates by default. If *custom interface speed* has been assigned to an interface, utilization will be used instead.

To show highwater activity (inbound or outbound) details for an interface, hover over the corresponding information (i) icon in the table. Shortcuts to run reports or drill into interface traffic/behavior can be accessed from the three-dot menu.

Additional options

To support workflow efficiency, the host details page header includes buttons to access the following functions:

- Changing the time period/range covered
- Pivoting to any supported report type filtered on the current host
- Viewing additional details and information from integrated sources (**Learn more** button)
- Applying filters (alarms and interfaces subviews only)

ATT&CK

The **Monitor** > **ATT&CK** tab can be used to investigate events based on the tactic, technique, and sub-technique assigned by the MITRE ATT&CK framework.

Events are plotted in a timeline, where the user is able to drill into them individually to open a tray containing the following:

- MITRE ATT&CK tactic and technique information, with links to the relevant MITRE ATT&CK knowledge base articles
- Shortcuts to the **Policies** or **Hosts** Alarm Monitor tab with filters for the event's details applied
- Basic event information

The page also includes the MITRE ATT&CK Enterprise Matrix, with technique classifications highlighted to match the corresponding events in the timeline.

Hint: Click on a technique cell in the matrix to view the policies violated in the **Policies** tab.

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

Applying filters

To further facilitate monitoring and investigation, the Plixer Scrutinizer Alarm Monitor views support multiple approaches to applying filters to the Alarm Monitor views.

Time range filter

The Alarm Monitor views can be set to show alarm/event information for either a custom date and time range or a specified *Last X* period (last 15 minutes, last 24 hours, last week, etc.).

To view data for a different period, click the **Time Range** (calendar) button and configure the range to apply.

Hint: When a custom range is specified, click the up/down arrows to automatically adjust the dates to cover the same period of time.

Card/chart filters

By default, the **Policies** and **Hosts** tabs use sparkline cards to summarize severity distribution across policies or hosts. These cards can be clicked to apply a filter for policy violations or hosts matching the selected severity.

Other visualization types (timelines and connection diagrams) showing different event details (events, alarm policy category, etc.), can be selected from the **View** dropdown and used to quickly apply the corresponding filter.

Advanced filters

Clicking the **Filters** button opens a tray where one or more filters can be manually configured.

The following filtering options are available:

- Policy
- Severity
- Risk
- Hosts
- Violators
- Targets
- Category (of alarm policy)

To apply a filter, expand the filter option/section, and select the criteria to use. Multiple options and criteria can be applied at the same time.

Note:

- The *Risk* filter is only available when the [Plixer Endpoint Analytics integration](#) is enabled. To learn more about Plixer Endpoint Analytics integration in Plixer Scrutinizer, see [this section](#) of this documentation.
 - The filter options tray also includes an option to show policies and hosts associated with events that have already been [acknowledged](#).
 - When exporting alarm/event data (via the **Options** button/tray), use the *Export CSV (All)* option to ignore any filters currently applied.
-

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. 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 event(s) will be hidden from all Alarm Monitor views.

Dashboards

Plixer Scrutinizer further enhances diverse network and security workflows through user-configurable dashboards, which can be configured and accessed via the **Monitor > Dashboards** page of the web interface.

Teams can set up and save any number of fully customized dashboards, allowing for the use of purpose-built views to address even the most unique monitoring or investigative requirements.

This section discusses the features and functions accessed via the **Monitor > Dashboards** tab/section of the web interface, and includes detailed guides for the creation, customization, and management of dashboards.

Creating a new dashboard

Creating a dashboard in the **Monitor > Dashboards** page allows users to enhance various network and security workflows by enabling tailored views to meet even the most specific monitoring or investigative needs. It also allows users to switch between different unique views to segregate monitoring requirements and workflows. Creating a dashboard is a highly recommended step when [setting up a new Plixer Scrutinizer environment](#).

To create a new dashboard, follow these steps:

1. Click the **Dashboard Options** icon (gear icon), and then click **Add New Dashboard**.
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. (Optional) Select the gadgets to add to the new dashboard from the **Gadget Selection** list.

Note: Gadgets can also be added or removed later as needed.

6. When done, click **Save** to create the new dashboard with the selected gadgets, if any.

Once the dashboard has been created, it will replace the current view and can be accessed via **Dashboard Options > All Dashboards** at any time.

Hint: To use an existing dashboard as a base/template instead, create a copy from [Edit Dashboard](#) or [Dashboard Options](#).

Editing/customizing dashboards

Existing dashboards can be modified or further customized by clicking the **Edit Dashboard** icon (pencil icon). A dashboard cannot be modified if it has been set to read-only. If you wish to edit an existing dashboard but retain a copy of its current state, you can choose to copy the dashboard before making any changes. See the Copying dashboards section for more information.

Adding/removing gadgets

New and existing gadgets can be added to a dashboard depending on the specific tasks and workflows that the user needs. To learn more about the different types of gadgets that can be added, refer to the [dashboard gadgets](#) section.

To add an existing gadget, follow these steps:

1. Click the **Edit Dashboard** icon (pencil icon).
2. In the *Gadgets* section, click the **Add Existing Gadget** icon (+).

3. Tick the checkbox beside the gadget(s) that you want to add.

Hint: Use the search field to enter the name of the gadget that you want to add, or use the dropdown menu to filter and view the existing gadgets according to type.

To create a new gadget, refer to the [external gadgets](#) section.

Gadgets can also be removed at any time. To do this, click the trash icon by the right hand side of the gadget that you want to delete/remove.

Editing gadget layout

The **Edit Gadget Layout** feature allows users to customize the arrangement and appearance of visual elements within the current dashboard. Clicking the four-arrow icon enables the **Edit Gadget Layout** where you can drag and drop the gadgets to rearrange their positions and resize them. Clicking the icon a second time exits the **Edit Gadget Layout** mode.

Additional options

When editing a dashboard, you can change the dashboard name via the **Modify/Rename** section. The dashboard can also be set as default or read-only in this section. The read-only option locks the dashboard's current configuration to prevent unintended changes, especially when it is shared among multiple users.

The **Options** section enables users to copy the current dashboard, [create a new dashboard](#), or remove the current dashboard.

Viewing dashboards

The **Monitor > Dashboards** view displays the current default dashboard. Switching to a different dashboard can be done by clicking the **Dashboard List** (meter) icon. This opens the **All Dashboards** tray which displays all the existing dashboards. Clicking either the dashboard name or the link icon automatically changes the dashboard in the current view.

The **All Dashboards** tray can also be accessed via the **Dashboard Options** (gear icon).

Default dashboard

The default dashboard is what the user sees upon first opening the **Monitor > Dashboards** page. This can be set via the main **Dashboards** page or via the **Admin** page.

To set a dashboard as default in the main **Dashboards** page:

1. Click on a dashboard name in the **All Dashboards** tray, and then enter *edit* mode.
2. Click **Modify/Rename**, and then tick the **Default Dashboard** checkbox.
3. Click **Save**.

Setting a default dashboard via the **Admin** page enables an admin to set a user's default view. To do this:

1. Go to *Admin > Users & Groups > User Accounts*.
2. Click the three-dot menu, and then click **Edit User**.
3. Click **Preferences**, and then select a dashboard from the *Default Dashboard* dropdown.
4. Click **Save**.

Note: The default dashboard is set only for the current user.

Refresh dashboard

Dashboard gadgets automatically update at regular intervals. A countdown showing the time until the next automatic refresh is also displayed when mousing over the gadget tile. By default, refresh interval is set to 5 minutes. This can be changed in the *Display Options* when *editing a gadget*.

To manually refresh a gadget, click the **Refresh Dashboard** icon.

Full screen

The **Full Screen** option expands the dashboard to occupy the entire screen, maximizing the space available for viewing data visualizations and insights. This feature can be ideal for large displays in NOCs.

At the bottom of the screen, there are three available options that allow you to exit full screen mode, refresh the data of all the gadgets displayed in the dashboard, and *edit the gadget layout*. The dashboard name is displayed beside these three options.

Dashboard management

The **Dashboard Options** tray can be accessed from the main **Monitor > Dashboards** view and provides access to dashboard and gadget management functions as well as user/user group access controls.

Deleting dashboards

To delete an existing dashboard, select *All Dashboards* in the tray, and then click the corresponding delete icon in the secondary tray.

A dashboard can also be deleted from the [dashboard configuration tray](#) by expanding the **Options** section and then selecting *Remove This Dashboard*.

Users that do not have the [Dashboard Admin permission](#) cannot delete dashboards created by other users. Dashboards that are set as default or read-only cannot be deleted.

Copying dashboards

The *Copy This Dashboard* option saves a copy of the current dashboard under a specified name. This function can be used to make modifications to an existing dashboard configuration while retaining the dashboard's current state.

This option can also be accessed from [dashboard edit/settings tray](#).

Deleting gadgets

To delete an existing gadget, select *All Gadgets*, and then click the corresponding delete icon in the secondary tray.

To find specific gadgets more quickly, use the dropdown to display gadgets by category.

Managing user and user group access

Users with the [Dashboard Admin permission](#) can manage access to dashboards by user or by user group from the **Dashboard Options** tray.

To grant a user or user group access to one or more existing dashboards:

1. Click *User Dashboards/User Group Dashboards*.
2. In the secondary tray, select the user/user group from the dropdown.

3. Use the checkboxes to select all dashboards the user/user group should be granted access to.
4. Click the **Save** button to save any changes made.

Once granted access to a dashboard, a user/user group member will be able to view and copy a dashboard. However, they cannot edit or delete dashboards created by other users.

Note:

- To fully access a dashboard, a user must also be granted *access to the dashboard's gadgets* through their user group.
 - A user can only manage dashboard access for other members of their user group(s). Additionally, they are only able to grant access to dashboards that have been created by members of their user group(s).
-

Dashboard gadgets

Each Plexier Scrutinizer dashboard can be tailored to a specific task, workflow, or user through its gadget configuration. When selecting gadgets, they are divided into the following categories:

- Custom
- Flow Analytics
- Flow Reports
- Maps
- Plexier
- Top n
- Vendor Reports
- Vitals

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.

Gadget types

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.

For new installs, the *Welcome* dashboard is displayed by default with the following gadgets:

- **Configuration Checklist** - Displays the configuration status (in percentage) of the Plixer Scrutinizer environment.
- **Quick Start** - Provides helpful links for understanding the basic functions of the Dashboards page.
- **Enabled Exporters** - Displays the total percentage of exporters enabled in the Plixer Scrutinizer environment.
- **Alarm Monitor** - Displays a graphic summary of alarms generated within the last 24 hours, categorized by severity.
- **Contact Us** - Provides Plixer's contact information for additional support.

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. The report gadget can be edited from the **Monitor > Dashboards** view. Enter the dashboard edit mode, and then click the pencil icon by the right hand side of the report gadget.

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

Network maps

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.

Custom gadgets

There are two types of custom/external gadgets that can be added to any dashboard:

- **iframe:** Gadgets that display another webpage on the dashboard, including external sites.
- **interfaces:** Allow users to add a custom-configured display showing interface speeds and traffic to the dashboard.

To add a custom gadget, do the following:

1. Enter dashboard edit mode, and then click **Create New Gadget**.
2. Select the type of gadget.
3. Depending on the type of gadget selected, provide the following details:
 - Name/label for the gadget
 - Gadget URL
 - Refresh interval for the gadget (in minutes)
 - Display options
 - Exporters
 - Interfaces
4. Click **Save To Dashboard**.

Once a gadget has been added, 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.

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.

Editing a gadget

1. Click the **Edit Dashboard** icon (pencil icon).
2. In the *Gadgets* section, click the pencil icon beside the gadget you wish to edit to open the gadget tray.
3. Configure the necessary gadget settings.
4. Click **Save**.

Managing gadget access via user groups

Access to gadgets can be managed via [user group permissions](#).

1. Navigate to **Admin > Users & Groups > User Groups**.
2. Click the three-dot menu beside the user group name.
3. Click **Edit User Group**.
4. In the user group tray, click the pencil icon for *Dashboard Gadgets*.
5. Then, either select **All Gadgets** to grant the user group access to all dashboard gadgets or select only the specific gadgets that the user group can access.

Changes are automatically saved upon selecting the gadget.

Network maps

The **Monitor > Network Maps** page is the interface for viewing/monitoring, creating, and customizing network maps. Management views for mapping groups and objects can also be accessed from this page.

This section contains guides and information on the **Network Maps** tab/section of the web interface, as well as further details related to Plixer Scrutinizer's network mapping functions.

Network mapping overview

Plixer Scrutinizer's integrated network mapping functions allow users to create dynamic, highly-customizable topology visualizations that can greatly enhance network monitoring and management workflows.

Network maps can be created as one of two types:

- **Spatial maps** allow *map objects* to be manually positioned in any layout. Custom connections, objects, and backgrounds (e.g., wiring cabinet, office floor plan, etc.) can also be used to increase the level of detail.
- **Geographical maps** use the longitudinal and latitudinal coordinates associated with map objects to automatically position them on a global map. Geomaps can help identify devices with issues, even when there are multiple topologies dispersed across different physical locations. Coordinates can be entered via the *mapping object management view*.

Existing maps can be viewed from the main **Network Maps** page and/or added to *dashboards*.

Important:

- Geographical maps require a Google Maps browser API key, which can be entered in the options tray of the *mapping group management view*.
 - Access to the Internet is required for Google Maps geolocation requests. If Plixer Scrutinizer is unable to reach the Internet normally, a Google Maps proxy server can be configured under *Admin > Settings > Google Maps Proxy Server*.
-

Mapping/map groups

Network maps are populated by assigning *objects* to *mapping groups*. Maps/groups can be created from either the main **Network Maps** page or the mapping group management view.

Groups are populated as part of *creating a new map*, but they can also be manually defined from the mapping group management view.

Note: Mapping groups are a separate *grouping scheme* from *IP groups*.

Object membership for existing groups can be modified at any time, and the map will automatically be updated the next time it refreshes.

For further details on mapping groups, see the page on *mapping group management*.

Mapping/map objects

Each mapping group can contain any number of objects of the following types:

- Devices/exporters
- Other mapping groups
- *Custom map objects* (spatial maps only)

Map objects can be added or reconfigured while in *map edit mode* (objects assigned to current map/group only) or from the mapping object management view.

For further details on mapping objects, see the page on *mapping object management*.

Connections

Connections are used to add links between objects in network maps and can serve the following functions:

- Show basic association between objects
- Display status/activity between interfaces
- Run a saved report for the connected objects

Each map can be configured with any number of connections, allowing users to tailor maps to their monitoring needs.

To learn more about adding and configuring connections, see *this section*.

Creating a new map

New network maps/mapping groups can be created from the main **Monitor > Network Maps** page or the [mapping group management view](#).

Note: After a spatial map/group is first populated, map objects will be stacked on top of each other until they are manually repositioned in [map edit mode](#). Objects in new geographical maps will be similarly clustered, unless they have had an address or GPS coordinates associated beforehand.

To add/create a new network map, follow these steps:

1. Navigate to **Monitor > Network Maps**, and then click the **Add** button.
2. In the **New Group** tray, select whether to create a spatial map or geomap.

Important: Geographical maps require a valid Google Maps browser API key to be displayed correctly. If no key has been added, it can be entered in the field provided when creating a new geomap. An API key can also be added via the [mapping group management view](#), in the options (gear) tray.

3. Enter a name (required) and description (optional) for the map.
4. Click the **Apply** button to automatically open the map configuration tray.
5. [Optional] In the **Add Object** secondary tray, use the checkboxes to select the devices and/or mapping groups to add as objects to the new map.

Note: Objects are added to the map in real time as they are selected.

6. [Optional] Add one or more [custom objects](#) to the new map by clicking the + button under *Objects* in the primary tray.
7. [Optional] Add one or more [connections](#) between objects in the new map by clicking the + button under *Connections* in the primary tray.
8. [Optional] Add a [background image](#) for the map by clicking *Background* in the primary tray and either selecting one of the provided images or uploading a custom background.
9. [Optional] Apply [optional map settings](#) to the new map under *Settings* in the primary tray.
10. Close the tray to return to the previous view.

Once a map has been created, it can be reconfigured or [further customized](#) at any time, through the optional steps described above. The map configuration tray can be accessed while in [map edit mode](#) or by selecting *Settings* in the three-button menu in the [mapping group management view](#).

Viewing network maps

When navigating to the main **Monitor > Network Maps** page, the default map for the *current user* is displayed. To bring up a different map, click the *All Maps* button and select it from the list. Network maps can also be accessed by clicking on the map/group name in the *mapping group management view*.

The main map view includes the shortcuts/buttons to the following views and functions:

| | |
|------------------------|---|
| Mapping Groups | Opens the <i>mapping group management view</i> |
| Mapping Objects | Opens the <i>mapping object management view</i> |
| Add | Opens the <i>new map/group tray</i> |
| Refresh | Updates the map to reflect most recent collected data |
| Report Menu | Opens a tray from which any <i>report</i> applicable to the group can be run |
| Edit Map | Switch to <i>map edit mode</i> |
| Options | <div>Configure the following global network map settings:</div> <ul style="list-style-type: none">- Refresh interval (in minutes)- Set connections to show rate or utilization- Use resolved hostnames instead of IP addresses as object labels |

Exporter/device map objects can be clicked to inspect rate or utilization by interface (via the *Explore > Exporters* view). Child groups can also be clicked to drill into their maps.

To run a report, modify *object properties* (including *location information*), or create a connection from the object, right-click an exporter or group object at any time.

Note:

- Full screen mode can be used to monitor network maps on larger displays. The zoom level can also be adjusted as needed using the corresponding buttons.
- To access data for all objects assigned to a network map, a user must also be granted *access to all included devices and/or interfaces*.

Map selection

The **All Maps** tray can be used to quickly switch between existing maps.

Clicking the reports (graph) icon opens the **Available Reports** tray, which allows the user to quickly pivot to any report applicable to the group. Hierarchical display for the map/group list can also be toggled on or off as needed.

Hint: If *Map Hierarchy* is enabled, expand a parent group to see all maps/groups added to it.

The *View All Mapping Groups* and *View All Mapping Objects* links can be used to navigate directly to the *mapping group* and *mapping object* management views.

Adding maps to dashboards

Network maps can be added to dashboards, where they can be viewed alongside other *dashboard gadgets*.

To learn more about dashboards and gadgets, see *this section* of the manual.

Map edit mode

Once a map/group has been created, it can be further reconfigured/customized by switching to map edit mode in the main network map view. This allows map objects in spatial maps to be freely repositioned.

The map editing toolbar provides access to the following functions:

- Configure *map settings*
- Inspect/edit object membership
- Auto-align objects
- Edit object layering (*bring to front, send to back*, etc.)

After making changes, click the save button to update the map/group. To exit edit mode and return to the main view, click the **Edit Map** button again.

Map customization

Network maps in Plixer Scrutinizer can be uniquely tailored to any type of environment/topology using several customizable elements.

Map edit mode

Once a map/group has been created, it can be further re-configured/customized by switching to map edit mode in the main network map view. This allows map objects in spatial maps to be freely repositioned.

The map editing toolbar provides access to the following functions:

- Configure *map settings*
- *Mapping objects*
- Select and drag objects
- Inspect/edit object membership
- Auto-align objects
- Edit object layering (*Bring to front, Send to back, etc.*)

After making changes, click the save button to update the map/group.

Additional shortcuts can also be accessed by right-clicking any object. Clicking the **Edit Map** button a second time exits map edit mode.

Note:

- To assign location information to a geographical map object, left-click on the object to enable editing and select *GPS Location* in the secondary tray.
 - Location information for an object is unique for each geographical map/group it is a member of.
-

Custom objects

Custom objects are non-device/non-group objects that can be displayed as icons (similar to regular map objects) or text boxes in a network map.

To add a new custom object, follow these steps:

1. Navigate to the [mapping object management view](#) and click the **+** button.
2. Use the dropdown to select the object type to create (*Icon* or *Text Box*), and then configure the following properties for the object:

| Icon object properties | |
|------------------------|---|
| Icon | Icon graphic to use for the object in the network map |
| Color | Color to apply to the icon |
| Size | Size of the icon |
| Weight | Icon variant to use |
| Label | Label to display for the object in the network map |
| Link | Complete URL of page to open when object is clicked (for example, http://www.plixer.com) |
| Description | Custom description to display in the tray object list and the mapping object management view |

| Text box object properties | |
|----------------------------|---|
| Label | Label to display for the object in the network map |
| Type | Determines whether the selected color is applied to the text or the background |
| Link | Complete URL of page to open when object is clicked (for example, http://www.plixer.com) |
| Description | Custom description to display in the tray object list and the mapping object management view |
| Shape | Shape to use as the background for the text |
| Dimensions | Dimensions to apply to the selected shape (length, width, radius, etc.) |
| Color | Color to apply to the selected shape |

3. Verify that the correct details have been entered, and then click the **Apply** button to save the new custom object.

Once a custom object has been created, it can be added to any map/group at any time. It can also be repositioned while in [map edit mode](#) and used as an endpoint in [map connections](#) (line and saved report connections only).

Note:

- After adding a custom object to a geographical map, switch to map edit mode and left-click on the object to assign location information to it.
- Custom objects can also be defined while in *map edit mode* or from the *mapping group management view*, under *Objects* in the map settings/configuration tray.

Connections

To add links showing relationships or network activity between objects in a network map, define one or more connections for the group.

To define one or more connections for a map/group, follow these steps:

1. Navigate to the *mapping group management view* and select *Settings* from the three-dot menu for the map to add connections to.
2. In the tray, click the + button under *Connections*.
3. In the secondary tray, use the two *Endpoint* dropdowns to select the objects to link with the connection.
4. Select a connection type from the *Type* dropdown, and then configure the required properties:

| Connection Type | Function | Properties |
|---------------------|--|--|
| Interface | Displays activity/utilization between the two endpoints on the specified interface | Interface/instance |
| Line | Static line linking the two objects | <ul style="list-style-type: none"> • Color • Label (optional) |
| Saved Report | Runs a specified <i>saved report</i> when clicked and changes color if thresholds are configured | <ul style="list-style-type: none"> • Saved report to run • Yellow, orange, and red thresholds (optional) |

5. Verify that the correct details have been entered and click the **Apply** button.

After a connection has been added to a map, mouse over the connection to view additional details.

Note:

- If a map/group has existing connections, they can be edited or deleted by clicking the corresponding icons in the list.
 - Connections can also be added to a map by opening the tray while in *edit mode*.
-

Interface connections

Interface connections are dynamic links whose visual properties indicate real-time status and traffic/activity between objects:

- The arrow **orientation** indicates the directionality of the highest traffic volume (as indicated in the connection label).
- The object **closer to the arrow** is the device/interface on which activity is metered (.e., the object that the activity displayed is *inbound to* or *outbound from*).
- The connection's **color** indicates one of the following:
 - Green: Active
 - Yellow/orange/red: Utilization reaching global thresholds configured under *Admin > Settings > Thresholds*
 - Blue: No bandwidth statement available
 - Grey: No traffic
 - Dashed grey: No flow data received in the last 5 minutes

Note: Utilization percentages can only be displayed for interfaces whose speeds are known (via SNMP or a *custom setting*).

Saved report connections

Saved report connections can be configured to run any saved report that applies to the two endpoints defined in the connection.

The yellow, orange, and red thresholds configured for the connection apply to the total of the rightmost/calculated column of the report. These thresholds are independent of *report thresholds*, which can be added to any saved report to trigger alarms.

Backgrounds

Spatial maps support the use of custom backgrounds (e.g., wiring cabinets, office floor plans, etc.), which can be uploaded under *Background*, in the map settings/configuration tray.

This can be done while in *map edit mode* or from the *mapping group management view*.

Managing mapping groups

The **Network Maps > Mapping Groups** view can be used to create, configure, and manage mapping groups.

It lists all existing maps/groups, alongside the following details for each one:

- Status
- Type (spatial or geographic)
- Description (if set)
- Timestamp when the group was last modified
- User that created the map

Clicking on a group name displays the map in the main **Network Maps** page. The map settings/configuration tray and other shortcuts can be accessed from the three-dot menu.

Global mapping group/page settings

The **Options** tray (gear button) contains global network map settings, as well as options for the mapping group management view, including:

- Default map for the current user
- Google Maps browser API key and TLD
- Map refresh interval in minutes
- Enable/disable hierarchy view in the mapping group management view (independent of the *map selection tray* toggle in the main **Network Maps** view)

Bulk actions

When one or more maps/groups are selected, the following batch operations can be performed via the **Bulk Actions** tray:

- Add or remove objects for all selected groups (shows objects common to all selected groups)
- Run a report filtered on all devices/objects included in the selected groups
- Delete the selected groups
- Clear selection

Creating a new map/group

New maps can be created from the mapping group management view by clicking the **+** button and entering the map type and name for the new group. After the group has been created, it can be further customized (membership, connections, etc.) via the configuration tray.

Clicking a group name will display the map in the main **Network Maps** view, where it can be further configured in *edit mode*.

Configuration tray

Selecting *Settings* from the three-dot menu opens the configuration tray for that map/group, from where the map/group can be reconfigured/customized at any time.

The tray is divided into the following main sections (also accessible in *map edit mode* in the main **Network Maps** view):

| | |
|--------------------|---|
| Settings | Inspect/edit <i>general settings</i> |
| Objects | Manage object membership or add custom objects to the group |
| Connections | Define or manage connections for the group |
| Background | Upload/select a background for the network map |

In addition, the configuration tray includes shortcuts for the following actions/functions:

- Run a specified report filtered on the devices/objects included in the group
- Create a duplicate of the selected map
- View the map in the main **Network Maps** view
- Set the selected map as the default map for the current user

Settings

The *Settings* section of the configuration tray contains the following general map settings/options:

| | |
|---------------------------|--|
| Name | Name identifying the map/group |
| Auto-add devices | Automatically add devices with resolved hostnames matching the specified regular expressions (RegEx) |
| Truncate labels on | Shortens map object labels by omitting the entered string |
| Description | Optional description to add to the map/group (can also be viewed by clicking the i icon in member object lists) |
| Pass status | When enabled, the current group's status will be reflected in its map icon in parent maps. |

Objects

When expanded, the *Objects* section of the tray shows all map objects currently assigned to the group. Objects can also be reconfigured or deleted by clicking the corresponding icon in the list.

To manage object membership for the current group, click the edit (pencil) button and select/deselect devices, groups, or *custom objects* in the secondary tray.

Connections

To learn more about map connections, see [this section](#) of the map customization guide.

Backgrounds

To learn more about adding custom backgrounds to spatial maps, see [this section](#) of the map customization guide.

Managing mapping objects

The **Network Maps > Mapping objects** view can be used to manage mapping object properties and group membership. New custom objects can also be defined from this view.

The main view lists all map objects currently assigned to at least one network map/group, alongside the following details:

- Icon assigned (reflects device or group availability)
- Type (exporter, user-created/custom, or map/group)
- Status
- Link that will be opened when the object is clicked (custom objects only)
- Number of maps/groups the object has been assigned to
- Timestamp when the group was last modified

Clicking on a group name opens the object properties/configuration tray, from where the object can be edited or assigned to maps/groups.

Bulk actions

When one or more maps/groups are selected, the following batch operations can be performed via the **Bulk Actions** tray:

- Add selected objects to one or more maps/groups
- Remove selected objects from one or more groups they share
- Add GPS location details to all selected objects

Object properties

Exporter and group object icons can be customized through the following properties:

- Icon
- Color
- Size
- Weight (icon variant)
- Label
- Link
- Description

These properties are applied to the object icon across all maps/groups an object has been assigned to.

Group membership

To manage group membership for the current object, expand the *Current Groups* section of the object properties tray.

Clicking the + button opens a secondary tray, where the object can be assigned to or removed from one or more maps/groups.

GPS location

Objects assigned to geographical maps are automatically positioned based on their location.

To enter GPS coordinates for an object, click on *GPS Location* in the object properties tray. An address can also be entered instead, for which coordinates will automatically be obtained via GPS lookup.

Custom objects

To learn more about adding/defining custom objects, see [this section](#) of the map customization guide.

5.1.3 Explore

The **Explore** views of the web interface can be used to quickly look up information on exporters, hosts, and other entities (users, applications, etc.) in the Plexier Scrutinizer environment.

This section covers the different functions and types of information that can be accessed via **Explore** views of the web interface.

Exporters

The **Explore > Exporters** tab can be used to look up information for all devices sending flows to Plixer Scrutinizer collectors.

The main view lists device status, traffic information, and other details either *by interface* (default) or *by exporter* and provides access to a summary tray for drilling into the corresponding *alarm and host views*. The left-hand *mapping/device group pane* can be used to apply filters and manage mapping group *settings*, *membership* and *connections*.

Interfaces view

In the *By Interface* view, the table lists the associated exporter as well as inbound and outbound activity details for each interface. A status icon indicates whether the exporter is available (green) or offline (red).

The following options can be accessed by clicking the exporter address/hostname, interface name, or three-dot menu in the table:

- **Reports:** Run any *report* supported by the exporter
- **Information:** Shows general interface information and links to the *Admin > Interfaces management view* filtered on the interface
- **Exporter:** Opens the **Alarms** subtab of the *host details view* for the exporter
- **View Interface:** Opens to the *host details view* for the interface
- **View Exporter Alarms:** Opens the *Alarm Monitor > Hosts view* filtered on the exporter
- **Reset Highwater Inbound:** Resets highwater mark data for inbound traffic
- **Reset Highwater Outbound:** Resets highwater mark data for outbound traffic
- **Reset Highwater Both:** Resets highwater mark data for both inbound and outbound traffic

Note:

- The *Inbound* and *Outbound* columns will display utilization percentage for any interfaces whose speeds are known (via SNMP or a *custom setting*). Otherwise, actual rates (in b/s) will be shown instead. Visualization options can also be manually set in the **Options** tray.
- The bulk actions tray, which contains options to run applicable reports and reset highwater values, can be accessed after one or more exporters or interfaces are selected using the checkboxes.

Exporters view

The **By Exporter** view lists exporter hostnames/addresses alongside the following details:

- Current status of the exporter (green: available, red: offline)
- Number of *mapping groups* the exporter is assigned to
- Number of interfaces associated with the exporter
- Average packets per second over the last 12 hours
- Average flows per second over the last 12 hours
- Timestamp of the most recent flow received from the exporter

In this view, the following options can be accessed by clicking the exporter address/hostname or three-dot menu in the table:

- **Reports:** Run any *report* supported by the exporter
- **Information:** Shows general exporter information and links to the *Admin > Exporters management view* filtered on the exporter
- **Exporter:** Opens the **Alarms** subtab of the *host details view* for the exporter
- **Interfaces:** Switches to the **By Interface** view filtered on the exporter
- **Tags:** View/manage custom tags for the device
- **Mapping:** Edit *object icon properties, mapping group membership, or location details* for the exporter
- **Admin:** Opens the *Admin > Exporters management view* (no filters applied)
- **View Exporter Alarms:** Opens the *Alarm Monitor > Hosts view* filtered on the exporter

Note:

- Click the details in the *Groups* and *Interfaces* columns of the table to quickly access the corresponding options in the tray.

- In the **By Exporter** view, the bulk actions tray contains options to run reports, add custom tags, and edit mapping details for all selected exporters.
-

Mapping group pane

The mapping group pane lists all current *mapping/device groups* and provides quick access to the following functions:

- Run any report supported by the group's *devices/exporters*
- View the network map for the group
- Apply a filter for the group's exporters or interfaces to the main list/table (click the filters button for additional options)
- Create a duplicate of the selected network map

In addition, the *Modify* option opens a tray where the *settings*, *membership*, *connections* or *background* for the network map can be modified.

Entities

The **Explore > Entities** tab can be used to look up and inspect the individual data entities—both user-defined and discovered—monitored by Plixer Scrutinizer as part of network activity.

The page is divided into separate subtabs displaying the following details for each entity type:

| | |
|--|--|
| Username | <ul style="list-style-type: none"> • Host associated with the observation • Data source • Machine name (if available) • Timestamp when the username was first seen on the host • Timestamp when the username was last seen on the host |
| Applications Defined | <ul style="list-style-type: none"> • Number of exporters the application was observed on • Total number of flows with data associated with the application • Average packet rate for activity involving the application • Average data transfer rate for activity involving the application |
| Hosts - Sources/Destinations/Pairs | <ul style="list-style-type: none"> • Source and/or destination IP address(es)/hostname(s) • Number of exporters the source, destination, or pair was observed on • Total number of flows with data associated with the host(s) • Average packet rate for activity involving the host(s) • Average data transfer rate for activity involving the host(s) |
| Autonomous Systems - Sources/Destinations/Pairs | <ul style="list-style-type: none"> • Source and/or destination autonomous system(s) • Number of exporters the source, destination, or pair was observed on • Total number of flows with data associated with the autonomous system(s) • Average packet rate for activity involving the autonomous system(s) |
| 5.1. Plexer Scrutinizer web interface | <ul style="list-style-type: none"> • Average data transfer rate for activity involving the autonomous system(s) |
| IP Groups - Sources/Destinations/Pairs | |

Clicking on an entity in any subtab opens a summary page (similar to the [host traffic subview](#)) that contains visualizations of the entity's activity as well as [report](#) shortcuts for deeper investigations.

Note: Shortcut links to manage [application definitions](#), [protocol exclusions](#), and [FA algorithm exclusion rules](#) are included in the corresponding subtabs.

Search

The **Explore > Search** tab 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](#).

The following are the available details displayed in the search results:

- Host
- 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, Plexier 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 Algorithms](#).
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
6. (Optional) Enter a database path in the *Host to Host Database* field to enable host pair indexing. To disable the feature, leave it blank.
7. 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 host indexing is enabled, additional resources may need to be allocated to the Plexier Scrutinizer collectors as described [here](#).

Host index population from historical data

If host indexing is not immediately enabled after Plexier Scrutinizer is deployed, the database can be back-filled at a later date using historical data.

To populate the host index database from historical tables, follow these steps:

1. SSH to the Plexier Scrutinizer server as the `plexier` user.
2. Stop the host index service:

```
sudo systemctl stop scrutinizer-host-index
```

3. Run the following to populate the database using the specified historical data tables and time range/window:

```
host_index --db_config --verbose --populate_from_history --  
→ table_interval=INTERVAL_TABLE --date_start="<START_DATE_TIME>  
→ " --date_end="<END_DATE_TIME>"
```

where:

- `START_DATE_TIME` and `END_DATE_TIME` must be formatted as `YYYY-MM-DD HH:MM`, with the time in 24-hour format (leading zeroes should be omitted).
- `INTERVAL_TABLE` is an integer that specifies the [aggregation interval tables](#) and should be set to 1, 5, or 30.

Note:

- If the time element is omitted from `END_DATE_TIME`, data from the end date specified will be excluded from the operation.
 - The utility can also be used to repopulate the host index database in case of data corruption. However, it is highly recommended to contact [Plexier Technical Support](#) for assistance with restoring data.
-

5.1.4 Investigate

The **Investigate** views of the web interface provide access to Plixer Scrutinizer's collaborative investigation and ML-powered forecasting (requires Plixer One Enterprise) functions.

This section introduces the **Collections** and **Forecasts** views and includes detailed guides for their associated functions.

Collections

Collections are bundles of one or more alarms, events, and/or reports that have been compiled and assigned to a specific user for further review and analysis.

Once created, a collection can be annotated and reassigned, allowing multiple users (e.g., NetOps and SecOps) to share workloads and collaborate in investigations.

In the web interface, all collection-related functions can be 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**.

Along with each collection's name, the table also shows the following details:

- Indicator that shows the current active collection (green checkmark)
- User who created the collection
- Date and time the collection was created
- 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, and then 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](#).
- **Filtering options** - Click the filter button to view available filtering options for the list.
- **Options** - Click the gear icon to view the available options for the list.

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 then click the **Delete** button.

Collection management

The collection management menu can be accessed from either of the following:

Alarm monitor view

1. Navigate to either **Alarm Monitor > Policies** or **Alarm Monitor > Hosts** tab.
2. In the **Alarm Policy** or **Host** list, hover over the star icon, and then select **Manage Collections**.

Current report view

1. Navigate to **Reports > Run Report**, and then *create/run a new report*.

2. After the report is run, hover over the star icon, and then select **Manage Collections**.

Creating a new collection

To create a new collection, click the **Add New Collection (+)** button, and then enter a unique name for the collection. Select a user to assign the collection to, and then click the **+** button to save the collection.

Note: The name and user fields must both be filled to create a new collection.

Once the collection has been successfully created, it will be added to the list in the management menu.

Setting the active collection

To set/change the current active collection, open the management menu, and then select the collection from the list. The green checkmark beside the collection name indicates that it is the current active collection. 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

1. Click the star button to open **Manage Collections** menu.
2. Click the button a second time (after it turns into an add (+) button).

This automatically adds the alarm, event, or report to the active collection.

To remove the item from the active collection, click the star button, and then click the button a second time (after it turns into a minus (-) button).

Forecasts

When paired with the Plixer ML Engine, Plixer Scrutinizer is able to use the aggregated flow data of a specified report to generate forecasts of future network activity and/or resource utilization.

Important: Forecasts require an active **Plixer One Enterprise** license. To learn more about licensing options, contact [Plixer Technical Support](#).

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

To generate a forecast, a [report](#) must first be run to define the scope of data for extrapolation.

The following data elements in a report will be used to generate the forecast:

- Hosts
- Data points
- Time period covered
- Filters applied

In the [results/output](#), click the **Forecast** button, and then enter a name to save the new forecast under. The main [Investigate > Forecasts](#) page will automatically be displayed after the forecast is created.

Note: The amount of time it takes to generate a forecast varies, depending on the amount of data that needs to be processed.

Forecast horizon and seasonality customization

By default, Plixer Scrutinizer applies a recommended forecast horizon and seasonality based on the volume of data sampled in the report used.

To manually define the horizon and seasonality instead, the filename for the forecast should be formatted as follows:

```
[forecast_name] ? <horizon_integer> <time_unit> with [no|auto|null] season [↩<season_integer> <time_unit>]
```

Natural language is also supported, so a forecast titled:

```
VPN Usage ? for 3 months with a season of 14 days
```

will generate a forecast with projected values for 3 months (after the end of the [report time range/window](#)) and a seasonality of 14 days.

Viewing Forecasts

All previously created/saved forecasts can be accessed from the **Investigate** > **Forecasts** page. Forecasts that are marked *Complete* under the **Status** column are ready to view.

Clicking on a forecast opens a detailed view with two sections:

Forecast timeline

The forecast timeline plots the data aggregated by the base report (solid lines) and shows the extrapolations (broken lines) up to the *horizon* of the forecast. Hovering over a line will show the upper and lower bounds of potential deviation (highlighted region), as well as additional details for the data element used to aggregate the data (hosts, applications, etc.).

The timeline can be viewed as either a line or step graph.

Inbound events

In addition to the timeline, the forecast details view includes a table listing the following information for each host, application, etc.:

- Rank (based on the forecast's calculated data)
- Date and time when the calculated data is expected to reach the expected maximum value
- Expected maximum value of the calculated data
- Upper bound for deviation in the calculated data's expected maximum

When applicable, the table links directly to the relevant **Explore** summary page for each element. The base report for the forecast can also be re-run at any time by clicking the **View Report** button.

Forecast management

The main **Investigate > Forecasts** page can be used to access forecasts after they are created and includes the following details for each forecast:

- ID number assigned to the forecast
- Forecast name/filename
- Name of the report used for the forecast (click to re-run)
- Forecast creator
- Current status of the forecast (*Initializing -> Starting -> Data Retrieval -> Processing -> Strategy Selection -> Learning -> Prediction -> Complete*)
- Timestamp when the forecast became ready to view

Note: In some cases, it may take up to several minutes for the Forecasting task to progress from *Initializing* to *Complete*.

Updating forecasts

Clicking the refresh icon reinitializes the forecast using the most up-to-date dataset for the base report's *time window/range* settings.

Note: Forecasts based on reports with a custom date and time range (i.e., not *Last X*) can also be refreshed but will result in the same projections. To obtain an updated forecast, re-run the report with adjusted date and time settings, and then generate a new forecast.

Deleting Forecasts

To delete one or more forecasts, select the forecasts using the checkboxes and then click the *Delete* button to permanently delete them.

5.1.5 Reports

The **Reports** views of the Plixer Scrutinizer web interface are used to create, run/view, and manage reports. Advanced features, such as defining *custom report thresholds*, setting up *scheduled email reports*, and creating *forecasts* (requires Plixer One Enterprise), can also be accessed from these views.

This section comprises detailed guides for leveraging the various functions related to reports in Plixer Scrutinizer.

How reports work

Reports are fully customizable network data aggregations that enable complete transparency for any asset or activity on the network.

When a report is *run*, traffic data is collated based on the configured time window, sources/devices, and filters before being grouped by the criteria defined in the report type (e.g., source-destination pairs, applications, etc.). The results are then displayed in the *output view*, where the modified settings can be applied as required by the current resolution or investigation.

Primary report settings

Plixer Scrutinizer uses the following primary settings to run/generate a report, all of which must be defined when *creating a new report configuration*:

Data sources/devices

When a report is run, Plixer 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 type

The base type of a report determines how network metadata from the selected observation points is aggregated (i.e., *by X*).

A report type can be selected when *creating a new report configuration* or as part of *refining a report's configuration*.

Report types are grouped according to their functional parameters to facilitate report type selection. The *Recommended*, *Recent* (last 16 report types run), and *Designed Reports* groups can also be used to quickly find frequently run report types.

Hint: For further detailed on available report types, refer to [this table](#).

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 backwards or forwards.

Graph type

The report output view includes multiple options for visualizing the aggregated data, and users are able to freely switch between any graph supported by the current report type.

Report graphs are further explained in the section on the [report output view](#).

Additional filters

To supplement the primary settings, additional filters can be applied to further limit or expand the scope of data covered by a report. These filters can be added when creating a new report configuration and/or redefined as needed from the [output view](#).

Custom reports

To learn more about creating custom reports, see the [Report Designer](#) topic in the [Classic UI](#) section of this documentation.

Creating/running reports

The **Reports > Run Report** page is the starting point for creating/running new report configurations.

New reports

To create/run a new report from the **Reports > Run Report** page:

1. Select between the two options to start creating a report:

| | |
|---------------------------|--|
| Select Devices | Select one or more devices before choosing a supported report type |
| Select Report Type | Select a <i>report type</i> before choosing data sources/devices. |

2. In the next step, select the type or devices for the report:

- Report type: Use the dropdown to select a category, and then select the report type to run.
- Devices: Check the devices under *Available Devices* and use the arrow buttons to add them to the *Selected Devices* list.

Note: Only supported report types or eligible devices (based on the selection(s) made in the previous step) are displayed.

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.

Saving reports

After a report is created and run, the configuration can be saved by clicking the save (disk) button in the *output view*.

Once a report has been saved, it can be re-run at a later time (either as-is or with *modified settings*). Saved reports can also be used to set up *custom thresholds to trigger alarms* and *scheduled email reports*.

To learn more about accessing and managing saved reports, see *this section*.

Hint: Access to specific reports and/or report folders can be defined as part of *user group permissions* from the *Admin > Users & Groups > User Groups* page.

Running reports via URL

To quickly run an all-devices *Host to Host* pair report, with a filter for a specified IP address (FILTER_IP), a URL in the following format can be used:

`https://SCRUTINIZER_ADDRESS/ui/reports/run-report/search/el/FILTER_IP`

Note: Plexier Scrutinizer will also accept a FILTER_IP in hex format but only if the IP address belongs to an exporter.

Report output

The output of a report will mainly consist of two classes of data: the grouping criteria/entities(sources/destinations, IP groups, users, etc.) and their aggregated activity data.

After a report completes running, the results are displayed in both graph and table formats in the output view, where the reports original settings can continuously be *refined* to create the visibility required for the current task.

Graph details and functions

Each report type supports multiple interactive graph options to visualize the data for the top ten grouping entities based on their activity. An *Others* entity, which combines the aggregated activity data for all entities outside the top ten, is also included.

The **Graph** dropdown allows the user to quickly switch between the available visualizations directly from the output view. Additional details for any entity or activity can be viewed by hovering over the corresponding graph element.

Table/list details

The output view table functions as both a summary of the report results and a legend for the graph. The columns to the left (without the sorting arrows) list report type's grouping entities, while the right-hand columns are used for the aggregated activity details. Traffic values can be displayed as average rates or totals (for the entire time range) by selecting the corresponding global setting in the *Options tray*.

Clicking on an entity in any grouping criteria column (e.g., source, application, or destination in a *Conversations App* report) opens a tray from where any supported report type can be run.

Hint:

- Timeline graphs (line, step, stacked bar, etc.) can be used to apply a new time range to the current report. To do this, click on the graph once, and then click and drag to highlight the new range to use.
- To hide the graph for the current report, click the **Hide** button in the header.
- Individual cells in the grouping criteria columns of the table can be dragged to the left into *inclusion* and *exclusion* dropzones to configure additional filters for the current report (click the **Apply** button in the tray when done).

Filters tray

Clicking the **Report Filters** button in the output view opens a tray where the filters for the current report can be redefined.

To add a new filter, do the following:

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 required details for the filter (varies by filter type).
5. Click the **Add** button.
6. In the primary tray, click the **Apply** button to re-run the report with the new filter(s) applied.

Existing filters can be modified by clicking the edit (pencil) button or removed by clicking the delete (trash bin) button.

Note: The data sources/devices that were initially selected for the report can also be modified via the filters tray.

Additional options

Clicking the **Options** button in the header opens a tray containing the following option submenus:

| | |
|------------------|--|
| Global | <p><i>Data:</i> Toggle between rates or totals in report results.</p> <p><i>Data Source:</i> Specify an <i>aggregation/roll-up table</i> to use for reports.</p> <p><i>Data Units:</i> Toggle between bits or bytes in report results.</p> <p><i>Interfaces:</i> Enable/disable grouping report results by interface.</p> <p><i>Data Mode:</i> Toggle between <i>summary and forensic</i> flow data to run reports.</p> <p><i>Show Others:</i> Enable/disable including the <i>Others</i> grouping entity in report results.</p> <p><i>Show Host Names:</i> Toggle between host IP addresses and hostnames in report results.</p> <p><i>Rows:</i> Select the number of grouping categories to include in report results.</p> |
| Table | <p><i>Peak:</i> Show/hide additional column for peak activity details.</p> <p><i>95th:</i> Show/hide additional column for 95th percentile activity details.</p> <p><i>Values:</i> Toggle between formatted/rounded and raw calculated activity data in the report table.</p> |
| Threshold | Configure a <i>custom threshold</i> for the current report. |
| Details | <p><i>Collectors:</i> View expanded details for the collectors associated with the data sources of the current report.</p> <p><i>Exporters:</i> View expanded details for the exporters/data sources used for the current report.</p> <p><i>Report JSON:</i> View the report JSON (for reporting API calls)</p> |

Note:

- Toggle on **Display Advanced Options** in the tray to access the *Data Mode* and *Values* settings.
 - If the *Rows* setting is increased beyond 10, additional grouping criteria/entities will be displayed in gray in the graph.
 - Use the **Copy to clipboard** button to quickly copy the report JSON to your clipboard.
-

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
- Data sources (devices/interfaces)
- 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 current task.

The following table lists the additional filters that can be applied to reports:

| Type | Description | Parameter(s) | Option(s) |
|--|---|--|------------------------|
| <i>Applications</i> | Filters 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 | Filters results for the selected IP group (defined under Admin > Definitions > IP | | 181 |

- *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.

TCP Flags filters

In the *Report Type* dropdown, you can run a TCP Flags report to retrieve information about the TCP flags set in TCP packets observed during a network analysis or packet capture.

To run the report, do the following:

1. Navigate to the **Reports > Run Report** page.
2. Select one of the two starting points to create a report.

Note: For more information, refer to the [Creating/running reports](#) section.

3. In the **Report Type** dropdown menu, select **Designed Reports**, and then select **TCP Flags**.
4. Configure the following settings:
 - Time Window
 - Display Type
5. In the **Additional Filters** field, select **Advanced Filters**.
6. In the **Select Element** field, select **tcpcontrolbits**.
7. Select **Equal** in the **Select Comparison** field, type in *SYN*, and then click **Add**.
8. Click **Run Report**.

Note: Setting this filter generates a TCP Flag report specifically about the SYN (Synchronize) flag in TCP packets observed during a network analysis or packet capture.

Refining report results

After a report is run, the [output view](#) can be used to further investigate any entity or activity included in the report results.

Sample use cases and workflows for reports can be found in [this section](#) of this documentation.

Switching between graphs

After a report has been run, the **Graph** dropdown allows the user to freely switch between the different graph and chart types supported by the report type.

This allows teams to highlight different aspects of a report's results as needed for their resolution or investigation.

Modifying the time range

The current report can be re-run to cover a different time range of flow data, allowing teams to inspect activity for the same grouping criteria at different points in time.

The period of time covered by the current report configuration can be adjusted via the time range selector in the main output view or by highlighting (click and drag) an area in any timeline graph.

Editing filters

Once a report completes running, its initial [filter configuration](#) can be modified to highlight activity for specific grouping entities.

In the main output view, click the **Filters** button to [add, modify, and/or remove filters](#). Additional filters can also be defined by dragging entities from the table's grouping criteria columns into the corresponding dropzones on the left side of the page. After the new filter configuration has been set up, click the **Apply** button in the tray to re-run the report.

Pivoting to different report types

The **Report Type** dropdown in the main output view can be used to run a different report type using the current data sources, filters, and other settings. This function can be used when additional context is required to further investigate a host or activity on the network.

Additionally, a different report type can be filtered for a specific entity in any of the table's grouping criteria columns. This is done by clicking on the entity and selecting the report to run in the **Available Reports** tray.

Managing saved reports

Reports that have been previously created and saved can be re-run from the **Reports > Saved Reports** subtab. This page also functions as the management view for saved reports.

Saved report list

To re-run a saved report, click on the report name in the main view of the **Saved Reports** subtab. Filters, including *report folders*, can be applied to the list, and it can be displayed in a tabular list or as individual tiles.

Both viewing modes indicate whether the following functions have been enabled or configured for each saved report:

- *Custom threshold*
- *Dashboard gadget*
- *Scheduled email*
- Added to dashboard(s) as a gadget (count)

In addition, the list mode table also indicates the report type, the last-run timestamp, and the creator of each report.

Deleting saved reports

To delete one or more saved reports, select the report(s) using the checkboxes and select *Delete* in the bulk actions tray.

Report folders

After a report has been saved, it can be assigned to one or more user-created folders.

Report folders can be used to organize/filter reports in the **Saved Reports** view. They can also be used to simplify report access management through *user group permissions*.

Creating report folders

New folders can be created from the **Saved Reports** view as follows:

1. Click the report folders button.
2. In the **Report Folders** tray, click the add (+) button.
3. Enter a name for the new report folder in the secondary tray.
4. Click the **Save** button.

Once created, the report folder will be added to the list in the **Report Folders** tray.

Note: Existing report folders cannot be renamed. However, a new folder with the desired name can be created and populated with the same saved reports.

Adding saved reports to folders

There are three ways to assign saved reports to folders:

- When entering a name to save a report, use the dropdown to select a folder to assign it to (*Unfolded* saves the report without adding it to any folders).
- In the **Report Folders** tray, click the edit (pencil) icon to make changes to the membership list of the selected folder.
- From the main **Saved Reports** view, select one or more saved reports using the checkboxes, and then use the *Move to folder* option in the **Bulk Actions** menu/tray.

Folder management

By default, the main **Saved Reports** view lists all saved reports accessible by the current user. To view only reports assigned to a specific folder instead, open the **Report Folders** tray and select the folder using the link icon.

The following functions can also be accessed via the folder list:

- Edit folder membership (edit/pencil icon)
- Delete folder (delete/bin icon)

Exporting reports

After a report is run, the results can be exported in PDF or CSV format from the **Export** (share button) tray in the output view.

Hint: PDF and or CSV copies of a report can also be attached to *email reports*.

Email reports

Once an *email server has been configured*, reports can be forwarded to any email address to provide external access to network data.

Email reports 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 reports

After any report is run, the results can be sent to one or more specified email addresses.

To send an email report, select *Email Report* in the export options tray (share button), and then enter the following details:

- Sender email address
- Recipient email address(es)
- Subject (optional)
- Message (optional)

Tick the appropriate checkbox(es) to attach PDF and/or CSV copies of the report results, if desired, and then click **Send**. A message confirming that the email report has been sent will be displayed.

Scheduled reports

Saved reports can be scheduled to run at specified intervals and sent to one or more recipients, enabling continuous network monitoring from any email inbox.

Hint: Configure a *last X time window* for a report to send/receive regular updates for any type of network metadata.

To set up a scheduled email report for a report:

1. Create, run, and save the report.

Note: Scheduled reports filtered on a specific date/time range will send either the same or no output when they are re-run.

2. In the output/results view, click the share button to open the export options tray.
3. Select *Schedule Report*.
4. In the secondary tray, enter/configure the following details:
 - A name for the scheduled report (used in the email subject line and for *scheduled report management*)
 - Recipient email address(es)
 - Frequency and exact minute on the hour that the email report should be re-run and sent
5. [Optional] Tick the appropriate checkbox(es) to attach PDF and/or CSV copies of the report results.
6. [Optional] Select additional reports to include in the scheduled email.
7. Click the **Save** button to save the scheduled email report configuration.

Once set up, a scheduled report will continue to be re-run and emailed at the scheduled intervals until it is disabled or deleted.

Scheduled report management

The **Reports > Scheduled Reports** subtab is the management view for scheduled report configurations. Scheduled reports can be created, reconfigured, and deleted from this page.

The table/list shows all current scheduled email reports and includes the following information for each configuration:

- Name/email subject
- Schedule details (frequency, time, day or date)
- Expected execution/run time
- Timestamp of the last run/email
- Configured recipient email addresses

One or more filters can also be applied to show only scheduled reports that match the defined criteria.

Creating/editing scheduled reports

New scheduled report configurations can be created from the management view, without having to run the saved report(s) beforehand. This can facilitate setting up multiple email configurations for reports that have been previously run/saved.

To create a new scheduled report from the management view, click the add (+) button and follow [these instructions](#), starting from step 4. Configurations can also be modified at any time by clicking the saved report name/subject to open the settings tray.

Deleting scheduled reports

To delete one or more scheduled reports that are no longer needed, use the checkboxes in the main view to select them, and then select *Delete* from the bulk actions tray.

Scheduled reports can also be temporarily disabled by ticking the *Disable* checkbox.

Report thresholds

Saved reports can be used to set up custom thresholds to alert network and security teams to specified network behavior. Report thresholds are applied to the calculated/aggregated columns of reports (either per row or total) and will trigger *alarms* based on the options configured.

Note: To deliver *Report Threshold Violation* alarms, Plixer Scrutinizer automatically re-runs reports with custom thresholds in the background every 5 minutes. As such, having a large number of active report thresholds—particularly *total reports* (as opposed to rate)—may result in performance issues. The total number of concurrent report processes that can be run at a time for threshold checks can also be adjusted under *Admin > Settings > Reporting*.

To add a report threshold to a report, follow these steps:

1. After the report is run (and saved), click the gear button to open the options tray and select **Threshold**.
2. Select whether the threshold should be applied per row or to the total of the calculated column.
3. Select the appropriate comparison operator (\geq or \leq) for the desired criteria.
4. Enter the desired threshold (value and prefix).

To disable a report threshold, re-run the report and click the delete (X) button in the **Filters** tray of the output view.

Report gadgets

Reports can be added to *dashboards* as gadgets, enabling continuous active monitoring of any specified network traffic/activity.

To create/configure a dashboard gadget for a report, follow these steps:

1. Run the report (new or saved).
2. In the output/results view, open the export options tray and select **Add to Dashboard** (or **Edit Gadget**, if the gadget was previously configured).
3. Enter a name for the gadget. If the report has not been saved, it will be saved under the name entered.

4. Select a dashboard to add the gadget to from the **Dashboard Tab** dropdown. Select *Don't send to dashboard* to manually add the report gadget to dashboards at a later time.
5. In the **Type** dropdown, select whether the gadget should show the report graph only, the table only, or both.
6. [*Graph* or *Graph & Table*] Select the gadget graph type and the report column to sort by.
7. [*Table* or *Graph & Table*] Use the checkboxes to select the columns to display in the gadget table.
8. [Optional] Expand the **Display Options** section of the tray to modify the default layout and behavior of the gadget.
9. Click the **Save** button to save the gadget configuration.

After a report gadget has been configured/saved, it will be included in the list of available gadgets when *creating* or *editing* a dashboard.

Note:

- To view a report in a dashboard, the current user must be granted access to both the report and the dashboard(s) through their *user group*.
- If the default gadget name for a saved report is changed, a new saved report will automatically be created under that name. If the gadget is renamed multiple times, the saved reports are still created, but only the most recent name change is applied to the gadget.

Creating forecasts

With the Plixer One Enterprise solution, Plixer Scrutinizer is able to further leverage the data aggregated by a report to generate a forecast of future traffic/activity.

A forecast can be generated after running any report by clicking the **Save Forecast** button. It can then be viewed via the main *Investigate > Forecasts* page.

To learn more about creating, viewing, and managing forecasts, see *this section* of this documentation.

To create a new forecast, click the **Save Forecast** button in the report output/results view.

Adding reports to collections

To add a report to the [active collection](#), open the **Manage Collections** menu (star icon), and then click the icon a second time (after it turns into an add (+) button).

To add a report to a the current [active collection](#):

1. Run the report.
2. In the results/output view, click the star button to open the collections menu.
3. Click the button a second time (after it turns into a + button).

If the report was previously added to the active collection, clicking a second time (- button) will remove it. To add the report to a different collection, select *Manage Collections* and then set that collection as active, before following the same steps.

Hint: A report can be included in multiple collections.

Once added, a report can be re-run directly from the [collection summary page](#).

5.1.6 Admin

The **Admin** views of the Plexer 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 Plixer Scrutinizer environment. It is the default view opened when clicking on the **Admin** text in the web interface header.

This page comprises the following interactive dashboard gadgets:

| | |
|----------------------------------|--|
| System: CPU | <i>Displays system performance metrics in timelines or charts</i> Click on a metric to switch views. Click on the <i>Vitals</i> icon to view server health. |
| Storage: Free Disk System | <i>Displays available storage per collector</i> Click on a storage element to switch views. Click on the <i>Vitals</i> icon to view OS health. |
| Services: Collector | <i>Displays the status of system services per collector</i> Hover over a chart element to view additional details. Click on the <i>Vitals</i> icon to view exporter health. |
| Configuration Status | <i>Shows the overall configuration progress for Plixer Scrutinizer and can be expanded to show the detailed configuration checklist</i> Click on a configuration item to view its current status and accept/decline the item. Click the <i>Launch</i> icon to open the relevant documentation page for an item, or hover over the <i>Dependencies</i> icon to see other related or required configuration items. |
| User Activity | Shows activity for individual users in a timeline |

Note:

- Click the **X** button to close the expanded tables for the vitals gadgets. To collapse the configuration checklist, click the progress bar a second time.
 - A configuration status dashboard gadget is also included in the default **Welcome *dashboard*** for Plixer Scrutinizer installs.
-

Vitals LEDs

Three notification LEDs for system vitals are persistent across all admin pages/views and can be used to monitor the general health of the Plixer Scrutinizer environment.

These LEDs correspond to the following system components/functions, from left to right:

- Server
- Software
- Exporter

Hovering over an LED will display additional details related to the component's current status. Each LED also functions as a shortcut to return to the admin dashboard with the corresponding vitals gadget expanded.

Admin Menu tray

The **Admin Menu** tray is the main access point for administrative functions in Plixer Scrutinizer. The tray can be opened from any admin page/view by clicking on the three-dot button.

The admin tray search field supports lookahead searching and can be used to quickly find settings, configuration views, or help descriptions that match the entered string.

Note: Admin views marked with a [-> are still only accessible via the Classic UI of the web interface.

Settings

The **Admin > Settings** page provides access to global settings for Plixer Scrutinizer's core functions and behavior, organized under the subcategories listed in the table below.

Click on a setting/subcategory below to learn more:

| | |
|--|---|
| Alarm Notifications | Configure global alarm message options and <i>Flow Inactivity</i> and <i>Interface Threshold Violation</i> alarm settings |
| Collector | Configure global collector settings and low resource fallback options |
| DNS | Set DNS cache retention duration and resolution attempt timeout |
| Data History | Set alarm and flow data history retention durations |
| Flow Analytics Settings | Configure global settings and auto-enable FlowPro Defender for appropriate algorithms |
| Global Authentication Settings | Configure user session and login security options (See also: user and user group settings) |
| Google Maps Proxy Server | Configure proxy server settings for Google Maps requests |
| Login Banner | Add a custom message to the Plixer Scrutinizer login page |
| ML AD Users | Configure Azure account info for integrating AD Users with Machine Learning (for UEBA alerts) |
| ML Alerts | Manage alarm thresholds for Plixer ML Engine vitals and Office 365 detection sensitivities |
| ML Data Limits | Set model and host/subnet limits for user and network behavior learning |
| ML Training Schedule | Set business hours for network behavior observation and modeling |
| Mapping Groups | Define and manage device groups for network mapping |
| Mapping Objects | Define custom map objects and manage object/group object properties |
| Reporting | Customize Plixer Scrutinizer reporting engine functions |
| System Preferences | Configure general Plixer Scrutinizer environment preferences/settings |
| System/New User Default | Set up default preferences/settings for new users |
| Thresholds | Customize color thresholds for displaying utilization |

Alarm Notifications

The **Admin > Settings > Alarm Notifications** tray contains the following settings:

| | |
|---------------------------------------|--|
| Hostnames | Enable to display device, target, and violator hostnames (when available) instead of IP addresses in alarm messages |
| Flow Inactivity | Enables <i>Flow Inactivity</i> alarms for devices that have not received flows in the last 30 minutes |
| Alarm Many Crop | Maximum number of devices, targets, and violators to display in alarm messages |
| Interface Threshold Violations | Enables <i>Interface Threshold Violation</i> alarms when utilization (in or out) for any interface exceeds the <i>Threshold - Utilization</i> value specified under Admin > Settings > System Preferences tray |

Hint: [Notification profiles](#) can be assigned to the *Flow Inactivity* and *Interface Threshold Violation* [alarm policies](#) to trigger [custom notification actions](#) for violations.

Important: If flow inactivity and interface threshold violation notifications are disabled from this tray, *Flow Inactivity* and *Interface Threshold Violation* alarm policy violations will not be reported or saved, even if the policies are set to the [Active or Store](#) state.

Collector Settings

The **Admin > Settings > Collector Settings** tray contains the following settings:

| | |
|--|--|
| Resolve Hosts at Collection Time | Forces 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:</i> Enabling this feature may result in significant latency at high flow volumes. For assistance, contact Plixer Technical Support . |
| Auto SNMP Update | Enables re-discovery of SNMP devices at 1:00 am every day. |
| Low Resource Fallback Cooldown Period | Amount of time (in seconds) to wait between low resource fallback “stages” (to prevent unwarranted feature or exporter pausing) |
| Low Resource Fallback Exporter Chunk Size | Number of exporters to pause or resume as a group when required for low resource fallback or recovery |
| Allowed Flow Rate Multiplier Percent | Multiplier/percentage of maximum supported flow rate that will not immediately trigger low resource fallback to accommodate brief spikes in flow rates <i>*Note:</i> Sustained flow rates exceeding 100% of the rated limit may result in stability issues. |
| Low Resource Fallback Mode | Select one of three modes to define Plixer Scrutinizer’s low resource fallback behavior |
| Listener Port | Ports that will be used to listen for NetFlow or sFlow traffic (separate by comma) |

Important: In distributed environments, these settings will be applied to all collectors in the cluster.

DNS

The **Admin > Settings > DNS** tray contains the following settings:

| | |
|----------------------------|--|
| DNS Cache Retention | Number of days to retain DNS names (0-365, 0 = never retain) |
| DNS Timeout | Maximum time (in seconds) to wait for DNS name resolution |

Data History

The **Admin > Settings > Data History** tray contains the following settings:

| | |
|--|--|
| Auto-Acknowledge Alarms | Number of days before alarms/events are automatically <i>acknowledged</i> |
| Alarm Retention Days | Maximum number of days that alarm/event data will be retained |
| Alarm Retention Size | Maximum amount of disk space (in MB) that can be used for alarm/event data before older records are deleted |
| Audit Log Keep Duration | Number of months audit logs will be retained |
| Auto History Trimming | Enables <i>automatic trimming</i> of older historical records based on the specified <i>Minimum Percent Free Disk before Trimming</i> setting (Overrides history retention settings) |
| Days of DNS Request Data | Number of days (0 - 365) to retain DNS request data |
| Minimum Percent Free Disk before Trimming | Minimum amount of free storage to maintain when <i>Auto History Trimming</i> is enabled |
| Flow Historical 1 Min Avg | Number of hours to retain <i>1-minute summary tables (totals)</i> of conversation data, as well as alarm/event data |
| Flow Historical 5 Min Avg | Number of hours to retain 5-minute summary tables (averages) of conversation data |
| Flow Historical 30 Min Avg | Number of days to retain 30-minute summary tables (averages) of conversation data |
| Flow Historical 2 Hr Avg | Number of days to retain 2-hour summary tables (averages) of conversation data |
| Flow Historical 12 Hr Avg | Number of weeks to retain 12-hour summary tables (averages) of conversation data |
| Flow Maximum Conversations | Number of top conversations to save for busy devices |

Note:

- When *Auto History Trimming* is enabled, 1m and 5m historical tables are trimmed to maintain the value specified in *Minimum Percent Free Disk Space before Trimming*. Automatic trimming is also used to retain a similar level of historical data for all configured exporters.
 - Assigning a value of 0 to historical flow data retention settings under *Data History* will not disable retention of the corresponding data table.
-

Disk calculator

Clicking the calculator icon in the data history settings tray opens the database size calculator, which can be used to view current and predicted storage use based on a specified set of conversation history retention settings.

In the calculator, enter the desired retention time for each flow data history interval (1m, 5m, etc.), and then click the check button. Current and predicted disk usage for each interval will then be displayed by collector, along with the predicted total disk space required for the current retention settings.

Note:

- Disk usage for other elements/functions, such as system metadata, alarm/event data retention, and host indexing are factored into these calculations. A 10% buffer for the operating system is also included.
 - All calculations/predictions are based on the system's current settings and collection parameters (flow volume/rate, templates, etc.).
-

Flow Analytics Settings

The following global settings for Flow Analytics can be modified from the **Admin > Settings > Flow Analytics Settings** tray:

| | |
|------------------------------|---|
| Auto Enable Defender | Enables automatic inclusion of FlowPro Defenders for the appropriate FA algorithms |
| Jitter by Interface | Packet delay variance (in ms) threshold used for record highlighting in <i>Status</i> reports |
| Latency | Latency threshold (in ms) used for record highlighting in <i>Status</i> report |
| Share Violations | Share violation details for cyber attacks originating from Internet IP addresses with Plexier to continuously improve host reputation records |
| Top Algorithm Devices | Sets whether <i>Top X</i> algorithms are automatically run against all devices or only manually defined inclusions |

Hint: Configuration options for individual Flow Analytics algorithms can be accessed from the [Flow Analytics Configuration](#) page.

Global Authentication Settings

The **Admin > Settings > Global Authentication Settings** tray contains the following global settings related to user credentials and logins:

| | |
|---------------------------------|---|
| Failed Login Max | Maximum number of failed logins before a user account is locked (0 = disabled) |
| Failed Login Window | Length of time (in minutes) within which failed logins will count towards the maximum allowed |
| Minimum Unique Passwords | Number of previous passwords that a local Plexier Scrutinizer user cannot reuse when changing their password |
| Session Timeout | Maximum time (in minutes) a Plexier Scrutinizer web session can be idle before the user is forcibly logged out (0 = disabled) |

Google Maps Proxy Server

The **Admin > Settings > Google Maps Proxy Server** tray is used to configure a proxy server to allow Plexier Scrutinizer to access the Internet and make Google Maps geolocation requests.

The following details must be provided:

- Username and password for authentication with the proxy server
- Proxy domain name
- Port used by the proxy server
- IP address or hostname (absolute URL) of the proxy server to use for geolocation requests

Login banner

Text entered in the following fields of the **Admin > Settings > Login Banner** tray will be displayed at the specified location on the web interface login page:

- Above the username input field
- Below the **Login** button

ML AD Users

The **Admin > Settings > ML AD Users** tray is used to add a Microsoft Azure account to enable [AD Users UEBA integration](#). The account must be configured to store Active Directory user sign-in logs.

After entering the account name and key, click the **Apply** button to save the details and enable UEBA detections/alerts.

ML Alerts

The **Admin > Settings > ML Alerts** tray can be used to adjust the CPU/RAM/DISK utilization and Kafka streaming latency alarm thresholds for the Plixer ML Engine. [Sensitivity](#) settings for detections related to Office 365 activity can also be modified from this tray.

After making changes, click the **Apply** button to save and apply the new settings.

For further details, see [this section](#) of the Plixer ML Engine configuration guide.

ML Data Limits

The **Admin > Settings > ML Data Limits** tray can be used to modify the limits for the number of models and the number of included hosts/subnets used by the Plixer ML Engine for learning network and user behavior patterns and making predictions.

After making changes, click the **Apply** button to save and apply the new settings.

For further details, see [this section](#) of the Plixer ML Engine configuration guide.

Note: Increasing any of the model or IP maximums in this tray may require allocating additional resources to the Plixer ML Engine appliance.

ML Training Schedule

The **Admin > Settings > ML Training Schedule** tray is used to set the business hours used for seasonality in the network behavior being observed by the Plixer ML Engine.

After entering the necessary details, click the **Apply** button to save and apply the new business hours.

For further details, see [this section](#) of the Plixer ML Engine configuration guide.

Note: The business hours used for network behavior seasonality are separate from the business hours applied when running reports, which are defined under [Admin > Settings > Reporting](#).

Mapping Groups

The **Admin > Mapping Groups** page can be used to create, configure, and manage device groups for network maps.

To learn more about the this page's functions, see this section on [mapping group management](#).

Mapping Objects

The **Admin > Mapping Objects** page can be used to manage mapping object properties and group membership. New custom objects can also be defined from this view.

To learn more about the this page's functions, see this section on [mapping object management](#).

Reporting

The **Reporting** tray contains the following options/settings, which are used to control how reports are run, displayed, and managed:

| | |
|---|--|
| Push Data Aggregation | Enable to automatically aggregate data when temp tables are pushed to the primary report (distributed environments only). |
| Business Hours End | Ending hour of the business day (as an integer, e.g., 5 pm -> 17) to use in reports. |
| Business Hours Start | Starting hour of the business day (as an integer) to use in reports. |
| CSV Include All Rows | Enable to include all rows in report CSVs (instead of only the selected <i>top X</i>). |
| CSV Repository | Path to use for saving exported CSVs. |
| Max Aggregations from Data Source | Maximum number of aggregations from a single data source. |
| Target Graph Intervals | Maximum number of intervals that Plexier Scrutinizer will aim to plot in graphs. |
| Limit All Device Report Results | Limits the number of results returned when running <i>All Devices</i> reports to this value (0 = no limit). |
| Maximum Raw Flow Exporters | Maximum number of exporters/devices allowed as filters for a <i>raw flows</i> report. |
| Max Reports per Interval | Maximum number of <i>scheduled email reports</i> that can be set to run within the same minute. |
| Max Reports per Email | Maximum number of reports that can be sent in a single scheduled email report. Note: Including too many reports in the same scheduled email may result in timeouts. |
| Max Report Processes | Maximum number of subprocesses (by time or by device) that a report will be divided into to reduce running time. |
| Display Others on Top | Allow or prevent report graphs from displaying other traffic above or below the top 10 results. |
| Display Raw MAC Addresses in Reports | When enabled, raw MAC addresses are displayed alongside other details in report results. |
| Use Alternative Times | If the <code>observationtimesseconds</code> field is included in a flow template, Plexier Scrutinizer will use it in place of <code>intervaltime</code> for reporting. |
| Use Host Index | Enables the use of the host index to limit the number of exporters/devices checked for <i>Group</i> and <i>All Devices</i> reports. |
| Saved Report Threshold Processes | Number of processes to fork when running <i>report threshold</i> checks. |
| Re-use Temp Tables | When enabled, reports will use temp tables whenever possible. |
| Report Caching Timeout | Number of minutes available reports will be kept cached. |
| Always Display Totals | Enable to always show totals in <i>Status</i> report results tables, even if the graph is set to show rates. |

Note: The times entered for *Business Hours End* and *Business Hours Start* do not affect the seasonality of the Plixer ML Engine's behavior monitoring/modeling functions for Plixer One Enterprise.

System Preferences

The **Admin > Settings > System Preferences** tray contains the following general settings:

| | |
|--|---|
| Disable File Upload | When enabled, files cannot be uploaded to the Plixer Scrutinizer server |
| Maximum Uploaded File Size in Bytes | Sets the maximum size allowed for uploaded files |
| Inactivity Threshold | Sets the number of minutes that the Explore > Exporters > By Interface view will display inbound and outbound activity details for inactive interfaces |
| Threshold - Utilization | Sets the interface utilization percentage (in or out) that will trigger an <i>Interface Threshold Violation</i> alarm |
| Inactive Expiration | Sets the number of hours (1 to 168) before an inactive interface is removed from the Explore > Exporters > By Interface view |
| LDAP Group Membership | Sets the schedule for syncing users between local user groups and LDAP Security Groups with the same name (Options: <i>On Login, Nightly, Both, Disabled</i>) |
| Version Checking | When enabled, Plixer Scrutinizer will automatically connect to the Internet and check for updates |

Note: Interfaces that have been inactive past the *Inactivity Threshold* setting but not longer than the *Inactive Expiration* setting will be displayed with **0.00 b/s** in their inbound and outbound columns.

Important: For users to be synced between a *local user group* and an LDAP Security Group, the two groups must have the exact same name, including any capitalization and punctuation.

System/New User Defaults

The settings/options **Admin > Settings > System/New User Defaults** tray can be used to define the default system preferences applied for new user accounts:

| | |
|------------------------------|--|
| Disable Welcome Modal | When ticked/enabled, hides the “Welcome to Scrutinizer” model for new user logins |
| Language | Sets the default system language for new users |
| Theme | Sets the default system theme for new users |
| Slim Navigation | When ticked/enabled, uses a theme with slimmed-down containers and icons for navigation elements |

Hint: Users can set their own language and theme by navigating to [Admin > Users & Groups > User Accounts](#) page and editing the **Preferences** for their username/account.

Note: Technical support (including this documentation) is only available in English.

Thresholds

The settings in the **Admin > Settings > Thresholds** tray can be used to adjust the percentage thresholds used to highlight interface utilization in different colors.

The default threshold values are as follows:

- Yellow: 51%
- Orange: 76%
- Red: 90%

Note: These values are also used to highlight [map connections](#) representing interfaces. Connections representing saved reports can have their color thresholds defined separately.

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.

Click on a setting/subcategory below to learn more:

| | |
|--|---|
| <i>Applications</i> | Define custom applications using IP address and port rules |
| <i>Autonomous Systems (AS)</i> | View autonomous system (AS) numbers/properties |
| <i>Host Names</i> | Define custom hostname-to-IP mappings and static subnet labels for reporting |
| <i>IP Groups</i> | Define rule-based IP range/subnet groups for reporting |
| <i>MAC Addresses</i> | Add and manage custom MAC address labels |
| <i>Protocol Exclusions</i> | Define protocol exclusion rules for reporting |
| <i>Type of Service</i> | Add custom labels for Type of Service (ToS) and Differentiated Services Code Point (DSCP) values in reports (<i>ToS Family</i> must first be set under Admin > Settings > Reporting) |
| <i>Well Known Ports</i> | Add and manage well-known port definitions |

Note: This category includes views/pages under the **Admin > Definitions** tab of the Plixer Scrutinizer Classic UI.

IP groups

IP groups are user-defined device groupings that can be leveraged when running reports, applying filters, or defining exclusions for FA algorithms.

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.
2. Enter a name for the group.
3. Select whether the group is internal or external from the *IP Group Type* dropdown.
4. Click **Save**.
5. In the main view, click the newly created IP group to open the configuration tray.
6. Expand the **Rules** section of the tray, and then click the (+) button to add a new rule.
7. In the secondary tray, select the rule type (IP address, subnet, etc.) to add.
8. Enter the details required for the rule in the additional fields.
9. Click **Add** to save the rule.

Steps 6 - 9 can be repeated as needed to define any number of membership rules for the IP group. Settings for existing IP groups can be further modified at any time.

Note:

- If there are overlapping host sets between IP groups, a host will automatically be assigned to the group whose rules define the narrowest range of addresses.
- The locality (internal or external) designations have multiple uses, including specifying traffic directionality (e.g., internal->internal, external->internal, etc.) for FA detections and defining inclusion and exclusion filters for report data sources. They also allow teams to quickly identify addresses as being internal or external to the organization when viewing [host details](#).

Bulk actions

When one or more IP groups are selected using the checkboxes, the following batch operations become available via the *Bulk Actions* button:

- Adding new rules to all selected IP groups
- Deleting all selected IP groups

Users & Groups

The **Admin > Users & Groups** category provides access to settings, options, and 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.

Click on a setting/subcategory below to learn more:

| | |
|--|---|
| Auditing Logs | View logs of Plexier Scrutinizer web interface user actions |
| Authentication Providers | Add and configure third-party authentication methods/servers |
| Authentication Settings | Configure global options for local and third-party authentication methods |
| Authentication Tokens | Add and manage user authentication tokens |
| User Accounts | Manage user accounts and preferences |
| User Groups | Set up local user groups and manage access to features and resources |

Auditing Logs

The **Admin > Users & Groups > Auditing Logs** page displays logs of Plexier Scrutinizer web interface user actions.

The main table of the *Auditing Logs* page includes the following details for each activity log:

- **Timestamp** - Date and time of the user activity
- **Message** - Description of the user activity
- **IP Address** - Local IP address of the device/machine used
- **Operating System** - Operating system of the device/machine used
- **Category** - Category or related page where the user action occurs (e.g. settings, admin, dashboards)
- **User Agent** - Web browser used
- **Username** - Username of the person who performed the activity

Time range filter

The *Auditing Logs* view can be set to show information for either a custom date and time range or a specified *Last X* period (last 15 minutes, last 24 hours, last week, etc.).

To view data for a different period, click the **Time Range** (calendar) button and configure the range to apply.

Hint: When a custom range is specified, click the up/down arrows to automatically adjust the dates to cover the same period of time.

Advanced filters

Clicking the **Filters** button opens a tray where one or more filters can be manually configured.

The following filtering options are available:

- Message
- IP Address
- Operating System
- Category
- User Agent
- Username

To apply a filter, expand the filter option/section, and select the criteria to use. Multiple options and criteria can be applied at the same time.

Note: When exporting activity logs (via the **Options** button/tray), use the *Export CSV (All)* option to ignore any filters currently applied.

Authentication Providers

The **Admin > Users & Groups > Authentication Providers** page can be used to set up and manage additional authentication methods/servers for the Plixer Scrutinizer web interface.

Adding a new authentication server

To set up a new authentication server, follow these steps:

1. In the main view, click the **+** button.
2. Select the authentication method to set up.
3. Enter the required details for the server in the secondary tray.
 - [Single sign-On](#)
 - [LDAP](#)
 - [RADIUS](#)
 - [TACACS+](#)
4. Click the **Save** button.

Once saved, the authentication server will be added to the list/table in the main view. To edit the details for a server, click on its name and make the necessary changes in the configuration tray.

Single sign-on

After [adding Plixer Scrutinizer to the IdP application list](#), the following details must be entered in the Plixer Scrutinizer web interface:

| | |
|----------------------------|---|
| Name | Name identifying the SSO service configuration in Plexier Scrutinizer |
| IdP Identifier URL | IdP-provided redirect URL for user authentication |
| Entity ID | https://<SCRUTINIZER_SERVER_IP> |
| Assertion URL | https://SCRUTINIZER_SERVER/fcgi/scrut_fcgi.fcgi?rm=usergroups&action=sso_response |
| Audience Value | https://<SCRUTINIZER_SERVER_IP> |
| Name Attribute | User identifier attribute passed by the IdP |
| Groups Attribute | User group identifier attribute passed by the IdP |
| IdP Metadata URL | URL to access the IdP metadata XML (only required if metadata XML cannot be downloaded) |
| IdP Metadata XML | Metadata XML downloaded from the IdP |
| Signing Certificate | Path to the SAML signing certificate obtained from the IdP (e.g., /home/plexier/scrutinizer/azure.cert) |

Note: If no group identifier attribute is provided, only the name attribute will be referenced for authentication.

Adding Plexier Scrutinizer as an SAML application

To set up SSO authentication for Plexier Scrutinizer, it should first be added to the application list of the SAML 2.0 SSO platform.

Note: Plexier Scrutinizer can be integrated into any SAML 2.0 SSO platform. For further information, contact [Plexier Technical Support](#) or refer to the provider's documentation.

Azure AD FS

To add Plexier Scrutinizer as an enterprise application to Azure AD FS, follow these steps:

1. Log in to the Azure portal as a global administrator.

2. Go to **Enterprise Applications > New Application**, and then select *Create your own application*.
3. Enter a name for the application (e.g., Plixer Scrutinizer).
4. Select the option to create a non-gallery application, and then click the **Create** button.
5. Under **Getting Started** in the application overview, click on *Set up single sign on*, and then select *SAML* as the single sign-on method on the next page.
6. In the next step, configure the following details under **Basic SAML Configuration** (other fields should be left blank):
 - **Identifier (Entity ID)**: `https://<SCRUTINIZER_SERVER_IP>/`
 - **Reply URL (Assertion Consumer Service URL)**: `https://<SCRUTINIZER_SERVER_IP>/fcgi/scrut_fcgi.fcgi?rm=usergroups&action=sso_response`
 - **Sign on URL**: `https://<SCRUTINIZER_SERVER_IP>/`
7. From the previous SAML SSO setup page, obtain the following for the Plixer Scrutinizer [SSO configuration form/tray](#):
 - **IdP Identifier URL**: Azure AD Identifier (under **Set up application**)
 - **Name Attribute**: Source attribute for the *Unique User Identifier (Name ID)* claim name (under **Attributes & Claims**)
 - **Groups Attribute**: Source attribute for the `http://schemas.microsoft.com/ws/2008/06/identity/claims/groups` claim name (under **Attributes & Claims**)
 - **IdP Metadata URL/XML**: Copy the *App Federation Metadata URL* or download the *Federation Metadata XML* (under **SAML Certificates**)
 - **Signing Certificate**: Download the x64/Base64 certificate (under **SAML Certificates**)

After completing all configuration steps in Azure and Plixer Scrutinizer, users and/or groups should be added to the Plixer Scrutinizer application to enable SSO authentication for the web interface.

Okta

To add Plixer Scrutinizer as a direct-access application to an Okta org, follow these steps:

Note: The instructions below are specific to the Okta Classic Engine. For Identity Engine users, click [here](#).

1. From the **Admin Console**, navigate to the **Applications** page.
2. Click *Create App Integration*, and then select *SAML 2.0* as the sign-in method.
3. After clicking **Next**, enter the general information for the application integration.
4. In the next step, enter the following SAML configuration details (other fields should be left blank):
 - **Single sign on URL:** `https://<scrutinizer_server>/fcgi/scrut_fcgi.fcgi?rm=usergroups&action=sso_response`
 - **Audience URL:** `https://<SCRUTINIZER_SERVER_IP>/`
5. When done, click **Finish**, and then select the Plixer Scrutinizer application from the **Applications** page.
6. Click on the **Sign On** tab, and then click *Identity Provider metadata* under the **SAML 2.0** section of the **Settings** page.
7. Obtain the details required for the Plixer Scrutinizer [SSO configuration form/tray](#).
8. Download the active signing certificate from the **SAML Signing Certifications** section.

After completing all configuration steps in Okta and Plixer Scrutinizer, users and/or groups should be added to the Plixer Scrutinizer application to enable SSO authentication for the web interface.

LDAP

When adding an LDAP authentication server, the following details must be entered:

| | |
|--|--|
| LDAP Server | IP address or hostname of the LDAP server |
| LDAP Port | TCP port used on the LDAP server |
| Domain | Domain used for authentication at the login page (e.g., <code>example.plixer.com</code>) |
| Administrator Password | Password to use in conjunction with the <i>Administrator DN</i> |
| Administrator DN | Distinguished name (DN) string to use (e.g., <code>CN=Example,OU=SampleUser,DC=PLIXER,DC=com</code>) |
| LDAP Server CA Certificate File | [Optional] Full path to the LDAP server's CA-signed certificate (must be in PEM format) |
| Certificate Verification | Select <i>Require</i> to use the specified certificate for verification with the server |
| ID Attribute | Attribute to use for verifying provided usernames (<code>sAMAccountName</code> (default), <code>UserPrincipalName</code> , and <code>UID</code> are supported) |
| Searchbase | Groups (semicolon-delimited if more than one) to search for authorized users (e.g., <code>OU=Example,DC=PLIXER,DC=com</code>) |
| Security Groups Allowed | [Optional] Security groups users must be assigned to for authentication (e.g., <code>CN=ExampleGroupName,OU=Securitygroups,OU=Applications,DC=PLIXER,DC=com</code>) |
| SSL Protocol | SSL/TLS protocol to use (if LDAPS is configured) |
| Timeout | Timeout (in seconds) for LDAP authentication requests |

Group syncing

When LDAP is enabled and a local user group 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 user group 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* user group (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* user group (if they were a member when they logged in).

Important: This feature requires the names of the local user group and the LDAP security group to be an *exact* match, including any capitalization and/or punctuation.

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

When adding a RADIUS server for authentication, the following details must be entered:

| | |
|-----------------------|---|
| RADIUS Server | IP address or hostname of the RADIUS server |
| RADIUS Timeout | Timeout (in seconds) for RADIUS authentication requests |
| Shared Secret | Shared secret for the RADIUS server |

TACACS+

When adding a TACACS+ server for authentication, the following details must be entered:

| | |
|------------------------|--|
| Pre-shared Key | Pre-shared secret/key for the TACACS+ server |
| TACACS+ Port | TCP port used on the TACACS+ server (Default: 49) |
| TACACS+ Server | IP address or hostname of the TACACS+ server |
| TACACS+ Timeout | Timeout (in seconds) for TACACS+ authentication requests |

Authentication Settings

The **Admin > Users & Groups > Authentication Settings** page is used to manage the following global options for each of *Plixer Scrutinizer's supported authentication methods*:

- Enable/disable the authentication method
- Default local group for new users created/added via the authentication method
- User access exception rules

To edit the settings for an authentication method, select the method in the main view and make the desired changes in the configuration tray. Settings are applied to all servers for the same authentication method.

Note:

- If an authentication method is disabled, users without credentials associated with a different method will not be able to access the web interface.
 - The authentication method associated with a user account can be changed from the *Admin > Users & Groups > User Accounts* view.
 - Additional settings related to user logins/authentication can be found under *Admin > Settings > Security*.
-

Authentication Tokens

The **Admin > Users & Groups > Authentication Tokens** can be used to add and manage authentication tokens, which can be used to grant external applications permissions based on a specified user account. Authentication tokens also allow applications to access the web interface without having to include the username and password in the URL.

Creating a new token

To create a new authentication token, click the **+** button in the main view, and then configure the following details in the tray:

- Expiration date
- User account whose permissions should be enabled by the token

When done, click the **Generate Token** button. The token string can be copied from the configuration tray or from the main list/view.

Token management

The main view of the **Authentication Tokens** page lists the following details for all existing tokens:

- Status (active/inactive)
- Token string
- Expiration date
- Timestamp when the token details were last modified

To modify the settings for a token, click on the string and make the necessary changes in the configuration tray.

To delete one or more existing tokens, select the tokens using the checkboxes, and then click the **Delete** button.

User Accounts

The **Admin > Users & Groups > User Accounts** page is the configuration and management view for Plixer Scrutinizer user accounts.

The main view/table of the page lists the following details for all existing users/accounts:

- **User Groups:** Current number of user groups the user is assigned to
- **Authentication Method:** Authentication type/method associated with the user account
- **Last Activity:** Timestamp of the most recent web interface activity logged for the user

Clicking on a username opens the [account configuration tray](#), where the current settings for the account can be modified. The details in the *User Groups* and *Authentication Method* columns also function as shortcuts to edit those settings.

User account settings

The account configuration tray is divided into five sections:

- **Preferences:** Set web interface preferences, including default views, display options, and timezone
- **User Group Membership:** Add/remove the user to/from user groups
- **Password:** Change the user's password
- **Authentication Method:** Edit the authentication type/method associated with the user account
- **Authentication Token:** Create/manage user account authentication tokens

Users can edit their preferences or change their password at any time. However, only the `admin` user and users assigned to [user groups](#) with the appropriate [permissions](#) will have access to all account settings/options.

Note: The *Locked* authentication method is automatically applied to an account that has exceeded the [maximum number of failed logins allowed](#). To unlock the account, select the previous authentication method used from the dropdown.

Creating a new local account

To create a new user account, click the **+** button in the main **User Accounts** view, and then enter the desired username and password in the fields provided. The new user must also be assigned to an existing user group via the dropdown.

When done, click the **Save** button to create the user account. Preferences and other settings for the account can be edited at a later time.

Bulk actions

When one or more user accounts are selected using the checkboxes, clicking the **Bulk Actions** button allows group membership changes to be applied to multiple groups at once.

Existing user accounts can also be deleted via the same tray/menu.

User Groups

The **Admin > Users & Groups > User Groups** page is the configuration and management page for local Plixer Scrutinizer user groups.

The main view/table of the page lists the following basic details for all existing user groups:

- **Members:** Number of users assigned to the group
- **Features:** Number of features/permission sets enabled for the group
- **Devices:** Number of devices/exporters that can be accessed
- **Interfaces:** Number of device interfaces that can be accessed
- **Groups:** Number of device/mapping groups that can be accessed
- **Saved Reports:** Number of saved reports that can be accessed
- **Dashboard Gadgets:** Number of dashboard gadgets that can be accessed
- **Third-Party Links:** Number of third-party integration links that can be accessed

Clicking on a user group name opens a configuration tray where the group's *access privileges* can be configured.

Creating a new user group

To create a new user group, click the + button in the main **User Groups** view, and then enter a name for the user group in the tray. An existing user group to use as a template for the new group must also be selected from the dropdown.

When done, click the **Save** button to create the user group. The group's name and access privileges can be modified at a later time via the configuration tray.

Managing group membership

To add/remove one or more users to a user group, click on the user group name to open the configuration tray, and then click the edit (pencil) icon for **Members**.

In the secondary tray, use the checkboxes to select members to assign to the group. Changes are automatically saved as they are made.

Managing user group access

To manage access to resources, functions, and network assets for members of a user group, click on the edit (pencil) icon for the corresponding category below.

Note:

- The search field can be used to quickly find resources, functions, or assets, use the search field in the secondary tray for the category.
 - When one or more user groups are selected using the checkboxes, clicking the **Bulk Actions** button allows access settings to be applied to multiple groups at once. User groups can also be deleted from this tray.
-

Dashboard Gadgets

The dashboard gadget access list is used to manage the gadgets that can be added to *dashboards* by group members. The selected gadgets can also be viewed by group members through *any other dashboards they have access to*.

The group should also be granted access to the *Dashboard User* feature set (see below) to allow members to create and view dashboards.

Devices

The device access list grants the group access to the status and other basic activity details for the selected network devices. The devices are also made available for use in functions that leverage the information, such as *network maps*.

Features

The feature access list is used to manage permissions for groups of related web interface functions or feature sets. Access can also be enabled using granular permissions for individual functions by toggling on the *Use Advanced* option in the secondary tray.

For a full list of features sets and individual permissions, see *this page*.

Groups

The groups access list is used to manage viewing access to existing device/mapping groups.

The group should also be granted access to the *Maps User* feature set to enable access to the main **Network Maps** page.

Interfaces

The interface access list grants the group access to all data for the selected interfaces and any hosts associated with them. The interfaces are also made available for use in functions that leverage interface data, such as *creating/running reports* and network maps.

Saved Reports

The saved report access list is used to manage access to saved reports for the group. Access can be enabled by individual saved report or by *report folder*.

To allow members to run reports, the group should also be granted access to the *Reporting User* feature set.

Integrations

The **Admin > Integrations** category provides access to the configuration views for the various third-party integrations that can be enabled in Plexier Scrutinizer.

Click on an integration type below to learn more:

| | |
|------------------------------|--|
| <i>3rd Party Integration</i> | Enable/disable and configure third-party integrations for <i>Explore > Exporters</i> view |
| <i>ASA ACL Descriptions</i> | Add/edit ASA firewall credentials for ACL description retrieval |
| <i>Email Server</i> | Configure SMTP server settings for email <i>notifications</i> and <i>reports</i> |
| <i>Flow Log Ingestion</i> | Configure and manage <i>Azure</i> , <i>AWS</i> , <i>OCI</i> , or <i>GCP</i> flow log ingestion sources |
| <i>STIX-TAXII</i> | Add and manage STIX-TAXII threat intelligence feeds |
| <i>ServiceNow</i> | Configure and manage ServiceNow instances for incident/ticket generation via <i>notifications</i> and <i>collections</i> |
| <i>Viptela Settings</i> | Enable/disable and configure <i>Viptela integration</i> for Cisco vManage devices |

Flow log ingestion

Plixer Scrutinizer can be configured to ingest flow logs from cloud data sources, enabling seamless visibility between on-prem and cloud-based assets.

Data sources are added from the **Admin > Integrations > Flow Log Ingestion** page as follows:

1. Click the **+** button to open the configuration tray for a new data source:
2. Select the service/type of data source to be added.
3. Enter the *required details* in the secondary tray.
4. [Optional] Click **Test** to verify that Plixer Scrutinizer can access the data source.
5. Click **Save** to save the data source configuration.

Once flows originating from a cloud data source are being ingested, any exporters reported—either as part of flow contents or in attached metadata—will be added to Plixer Scrutinizer. These devices can then be used similarly to regular exporters in Plixer Scrutinizer’s functions (e.g., *reports*, *network maps*, *Security Groups*, etc.).

Hint: To delete one or more data source configurations, select them using the checkboxes and use the *Delete Integrations* option in the **Bulk Actions** tray.

For further information and additional set-up steps for specific cloud providers, see the corresponding sections below:

- *Amazon Web Services VPC flow log ingestion*
- *Azure flow log ingestion*
- *Oracle Cloud Infrastructure Streaming flow log ingestion*
- *Google Cloud Platform VPC flow log ingestion*

Alarm Monitor

The **Admin > Alarm Monitor** category covers the configuration and management views for functions related to events/detections and alert delivery.

Click on a settings subcategory below to learn more:

| | |
|---|--|
| Alarm Policies | Reconfigure, enable/disable, and assign notification profiles to alarm policies |
| Flow Analytics Algorithms | Reconfigure, enable/disable, and add inclusions/exclusions to FA algorithms |
| ML Dimensions | Define traffic for the Plixer ML Engine to monitor for behavior modeling |
| ML Rules | Define subnet, host, or interface inclusion/exclusion rules for Plixer ML Engine observation |
| Notification Profiles | Create and manage profiles to assign notification actions by alarm policy |
| Security Groups | Create and manage IP address security groups to define FA algorithm inclusions |

Alarm policies

The **Admin > Alarm Monitor > Alarm Policies** page can be used to enable/disable, inspect, or reconfigure individual alarm policies.

Hint: For detailed information about individual alarm policies, refer to [this section](#) of the documentation.

The main view lists the following details and settings for each policy:

| | |
|---------------------------------|---|
| Status | Current state the policy is set to (green: <i>Active</i> , blue: <i>Store</i> , grey: <i>Inactive</i>) |
| Flow Analytics Algorithm | FA algorithm driving detections for the policy |
| Category | Type/nature of detections reported under the policy |
| Violations | Current number of active violations of the policy |
| Exporters | Number of exporters defined as inclusions for the associated FA algorithm |
| Timeout | Amount of time (in seconds) that must pass before the next observed violation is counted as a new event |
| Weight | Value used to calculate <i>severity</i> when violations are reported in the Alarm Monitor views |

Filters can be applied to quickly find specific policies, and the table can be exported for external use.

Modifying policy settings

To view additional details (including message format, variables, and event/artifact criteria) about an alarm policy or make changes to its configuration, open the configuration tray by clicking on the policy.

In the *Information & Settings* section of the tray, click the **Edit** (pencil) icon to modify any of the following settings:

- Weight
- Timeout
- Status

The secondary tray also shows the message format for reporting violations and lists all message variables used. It also contains the exact criteria used for aggregating individual observations as the same event/artifact.

Hint: When one or more alarm policies are selected via the checkboxes, the **Bulk Actions** button can be used to apply the same configuration changes to all selected policies.

Adding custom notifications

The *Current Notifications* section of the tray can be used to manage *notification profiles* for the selected policy.

To assign a new/additional notification profile to the alarm policy:

1. Click the + button.
2. In the secondary tray, use the dropdown to select the notification profile to assign (or click the + button to *create a new profile*).

3. Customize notification behavior using the following settings:

| | |
|----------------------------|--|
| Frequency | <p>Specifies how often the actions defined in the notification profile are triggered (with any configured filters applied):</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 criteria 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: Use the Alarm Monitor page to drill down into the Policy > Event > Observations view to see which details should be applied as filters for notifications.

4. Click **Apply** to assign the notification profile with the current settings.

An alarm policy can be assigned multiple notification profiles, which will be triggered based on the frequency setting and filters configured for each profile. The same notification profile can also be added multiple times using different frequency settings and filters.

Hint: In the main view, the three-dot menu for alarm policies also includes shortcuts to create, inspect, or assign notification profiles for the policy.

Flow Analytics algorithms

The **Admin > Alarm Monitor > Flow Analytics Algorithms** page can be used to enable/disable, inspect, or *reconfigure individual FA algorithms*.

The main view consists of a graph showing total duration of observations/detections and a table listing the following details and settings for each algorithm:

| | |
|-------------------|--|
| Status | Current <i>state</i> of the algorithm (green: <i>Active</i> , grey: <i>Inactive</i>) |
| Exporters | Number of exporters defined as <i>inclusions</i> for the algorithm |
| Groups | Number of <i>security groups</i> defined as inclusions for the algorithm |
| Exclusions | Number of <i>exclusions</i> (IP addresses, subnets, IP groups, etc.) defined for the algorithm |
| Policies | Number of <i>alarm policies</i> associated with the algorithm |

Filters can be applied to quickly find specific algorithms, and the table can be exported for external use.

Algorithm configuration

To view or make changes to the current settings of an FA algorithm, open the configuration tray by clicking on the algorithm.

Inclusions

When defining inclusions for an algorithm, exporters can be added individually or through *security groups*:

1. Expand the *Exporters* or *Security Groups* section of the tray and click the edit (pencil) button.
2. In the secondary tray, use the checkboxes to select the exporters or security groups to add to the inclusion list.

Hint: Use the search box/field to quickly find specific exporters or security groups.

3. Close the trays to return to the main view.

Algorithm inclusion lists can be edited at any time. Exporters or security groups can also be removed by clicking the delete (trash bin) icon after expanding the corresponding section in the configuration tray.

Exclusions

To define traffic to be exempted from monitoring using a specific FA algorithm, add exclusions as follows:

1. Expand the *Exclusions* section of the configuration tray, and then click the + button.
2. In the secondary tray, use the dropdown to select the type of exclusion to add.
3. Enter the details/criteria (based on the type) for the exclusion.
4. Click the **Apply** button to save the exclusion.
5. Repeat the steps as necessary to add all necessary exclusions.

Exclusions can be added or removed at any time. To delete an exclusion, click the delete (trash bin) icon after expanding the *Exclusions* section of the configuration tray.

Hint: Assign devices with similar Flow Analytics requirements to an IP group to quickly add them to any algorithm's exclusion list using the *Child Group* exclusion type. The default *DNS Servers*, *Public WiFi*, *Network Scanners*, and *SNMP Pollers* IP groups are already defined as exclusions where necessary and only need to be populated after Plixer Scrutinizer is deployed.

Algorithm settings

To modify how an algorithm is applied to collected flow data, click *Settings* to access additional settings for the algorithm. After making desired changes, click **Apply** to save the new settings or **Defaults** to revert to default values.

For a full list of additional settings by algorithm, see [this table](#).

Enabling/disabling algorithms

To optimize performance and resource utilization, FA algorithms that are not applicable to the current Plixer Scrutinizer environment can be disabled.

This is done using the enable/disable toggle in the configuration tray. The [Admin > Settings > System Preferences](#) view can also be used to disable algorithms with similar applications as part of predefined *feature sets*.

Bulk actions

When one or more algorithms are selected using the checkboxes, the following batch configuration actions can be accessed via the **Bulk Actions** button:

- Adding sources/inclusions (exporters and/or security groups) to all selected algorithms
- Disable or enable all selected algorithms

For further details on FA algorithms and configuration recommendations, see the [configuration guide for Flow Analytics](#).

ML dimensions

The **Admin > Alarm Monitor > ML Dimensions** page is the management view for the *feature dimensions* covered by the Plixer ML Engine's network behavior models.

Note: The Plixer ML Engine is part of the Plixer One Enterprise solution. Contact [Plixer Technical Support](#) to learn more.

The page's main view/table lists the following details for all dimensions currently defined:

| | |
|----------------------|--|
| Status | Current state the dimension is set to (green: <i>Enabled</i> , grey: <i>Disabled</i>) |
| Protocol | Communication protocol |
| Port | Communication port |
| Internal Only | Option to interrogate only internal communications |
| Used For | Type of inclusion/source the dimension is applied to |
| Aggregation | Flow template field used for data aggregation |
| Grouped By | Flow template field used to group observed flow data |
| Created By | User ID of dimension creator |
| Last Modified | Date and time the definition was last modified |

Clicking on a dimension opens the details/settings tray, where the dimension can be enabled/disabled and configured to only apply to internal traffic.

Adding a new dimension

Additional feature dimensions can be defined from the **ML Dimensions** management view as follows:

1. In the main view, click the + button to open the *Add Dimension* tray.
2. Select which inclusion type the dimension should apply to (hosts/subnets or exporter interfaces).
3. In the secondary tray, fill in the form with the following information:
 - A name for the dimension
 - Flow template field to use for grouping (can only be changed for host dimensions)

- Aggregation method/field
 - Communication protocol to monitor
 - Port to monitor
4. [Optional] To monitor only internal traffic for the dimension, toggle on *Internal Only*.
 5. [Optional] To add the dimension in a disabled state, use the *Enabled* toggle.
 6. Verify that the details and settings entered are correct and then click the *Add* button.

Once added, host dimensions (prefixed with CLIENT-) and exporter dimensions (prefixed with NET-) will be included in the main table/view. Settings for existing dimensions can be edited at any time by clicking on them to open the configuration tray.

Deleting dimensions

To delete feature dimensions, select one or more dimensions using the checkboxes in the list/table, and then select the *Delete* option in the bulk actions tray.

Alternatively, feature dimensions can instead be disabled (either individually or as a bulk action) to retain the definitions.

Dimensions can also be disabled and re-enabled from the bulk actions tray if the definitions need to be retained for future use.

Note: The **Bulk Actions** button is only available when one or more items are selected in the main table/view.

ML rules

The **Admin > Alarm Monitor > ML Rules** page is the management view for *inclusion and exclusion rules* for the Plixer ML Engine.

Inclusions and exclusions are managed in separate subtabs.

Note: The Plixer ML Engine is part of the Plixer One Enterprise solution. Contact *Plixer Technical Support* to learn more.

Managing inclusion rules

The **Inclusions** tab defaults to the *By Host* subview, which lists the following details for all current host/-subnet inclusions:

| | |
|----------------------|--|
| Status | Current state the inclusion is set to (green: <i>Enabled</i> , grey: <i>Disabled</i>) |
| CIDR | CIDR number |
| # HOST(s) | Number of hosts included in the subnet |
| Sensitivity | <i>Sensitivity setting</i> for the inclusion |
| Detections | Optional malware detections (green: <i>Enabled</i> , grey: <i>Disabled</i>) |
| Last Modified | Date and time the rule was last modified |

The *By Exporter* subview (accessible via the dropdown) lists the following details for all current exporter interface inclusions:

| | |
|----------------------|--|
| Status | Current state the inclusion is set to (green: <i>Enabled</i> , grey: <i>Disabled</i>) |
| Sensitivity | <i>Sensitivity setting</i> for the inclusion |
| Last Modified | Date and time the rule was last modified |

Adding an inclusion rule for a host or subnet

Additional host inclusion rules can be defined from the *By Host* subview as follows:

1. Click the add (+) button to open the **Add ML Host** tray.
2. Enter the network address and select the appropriate netmask for the host/subnet to be added.
3. Select the *sensitivity setting* for the inclusion.
4. [Optional] Enable threat detection using pre-trained algorithms for the host/subnet with the *Malware Detections* toggle.
5. [Optional] To add the inclusion rule in a disabled state, use the *Enabled* toggle.
6. Click the **Save** button to save the rule configuration.

Once created, new host inclusion rules will be added to the list in the *By Host* subview under the network address specified. Settings for existing host inclusion rules can be modified at any time by clicking on the edit (pencil) icon in the details/configuration tray.

Adding an inclusion rule for an exporter Interface

Additional exporter inclusion rules can be defined from the *By Exporter* subview as follows:

1. Click the add (+) button to open the **Add ML Exporter** tray.
2. Select the exporter to add from the *Network* dropdown.
3. Select the *sensitivity setting* for the inclusion.
4. [Optional] To add the inclusion rule in a disabled state, use the *Enabled* toggle.
5. Click the **Save** button to save the rule configuration.

Once created, new exporter inclusion rules will be added to the list in the *By Exporter* subview under the exporter interface specified. Settings for existing exporter inclusion rules can be modified at any time by clicking on the edit (pencil) icon in the details/configuration tray.

Deleting inclusion rules

Inclusion rules can be deleted from either subview by selecting one or more rules in the list/table, and then selecting the *Delete* option in the bulk actions tray.

Alternatively, inclusion rules can instead be disabled (either individually or as a bulk action) to retain the definitions.

Note: The **Bulk Actions** button is only available when one or more items are selected in the main table/view.

Managing exclusion rules

The **Exclusions** tab lists the following details for all current exclusion rules:

| | |
|----------------------|---|
| Source | Source address |
| Host(s) | Hosts included in the source address |
| Destination | Destination address |
| Host(s) | Hosts included in the destination address |
| Detections | Number of detections ignored for the rule |
| Last modified | Date and time the rule was last modified |

Adding an exclusion rule

An exclusion rule can be defined from the **Exclusions** tab as follows:

1. Click the add (+) button to open the **Add Exclusion** tray.
2. Configure the source network address.
3. Configure the destination network address.
4. Under *Detections*, select the detections that should be ignored for the specified traffic.
5. Click the **Save** button to save the rule configuration.

Once created, new exclusion rules will be added to the list in the main view of the **Exclusions** tab. Settings for existing exclusion rules can be modified at any time by clicking on the edit (pencil) icon in the details/configuration tray.

Note: 0.0.0.0/0 can be used as the source or the destination to exempt all incoming/outgoing traffic to/from the paired address from the selected ML detections.

Deleting exclusion rules

Exclusion rules can be deleted from the main **Exclusions** list/table by selecting one or more rules, and then selecting the *Delete* option in the bulk actions tray.

Alternatively, exclusion rules can instead be disabled (either individually or as a bulk action) to retain the definitions.

Note: The **Bulk Actions** button is only available when one or more items are selected in the main table/view.

Notification profiles

The **Admin > Alarm Monitor > Notification Profiles** page can be used to add, edit, and manage notification profiles, which can be used to add custom notifications to alarm policies.

Once created, a notification profile can be assigned to one or more alarm policies from the [Admin > Alarm Monitor > Alarm Policies](#) page. All notification actions defined in the profile will automatically be triggered whenever the policy is violated.

Note: Notification actions are only triggered if the alarm policy it's assigned to is set to [Active or Store](#). The FA algorithm associated with the policy must also be enabled.

Creating a notification profile

To create a new notification profile, click the + button, enter a name (can be changed later) for it in the provided field, and click the **Save** button.

Hint: The notification profile management page can also be accessed directly from the tray when [configuring notifications for an alarm policy](#).

Once saved, the profile will be added to the main view list and can be further configured.

Adding notification actions to a profile

To add [notification actions](#) to an existing 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. Enter the additional details (based on the action type) in the provided fields.

Hint: Use the [listed variables](#) to include additional details in notification messages or as arguments in custom scripts.

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.

To define additional actions in the same profile, repeat the steps as needed. Each notification profile can be configured with any number of actions in any combination.

Hint: To add notifications for [custom report thresholds](#), set up a notification profile and assign it to the *Report Threshold Violation* alarm policy via the [Admin > Alarm Monitor > Alarm Policies](#) page.

Bulk actions

When one or more profiles are selected using the checkboxes in the main view, an action can be added to all selected profiles via the **Bulk Actions** button.

Notification profiles can also be deleted this way.

Notification actions

Each notification profile can be configured with any number of notification actions, all of which will be triggered when the associated alarm policy is violated.

Click on a notification action type below for additional details and configuration steps:

| | |
|-----------------------------------|--|
| Email | Email event details to one or more specified users |
| Logfile | Output event details to a logfile |
| Syslog | Forward event details to a specified host via syslog |
| SNMP Trap | Create an SNMP trap to report event details to a specified host |
| Script | Run any custom script and optionally use variables to pass event details as arguments |
| Auto Acknowledge | Automatically acknowledge alarms/events under any specified policy (overrides data history setting) |
| ServiceNow Ticket | Create a ServiceNow ticket (with an optional API JSON script) for a configured ServiceNow instance |
| CEF | Use a CEF notification to send event details to a specified host |

Hint: Notification profiles can include multiple configurations of the same notification action type.

Email

The email notification action can be used to automatically send email alerts when events are reported under the associated alarm policy.

Important: To configure email notifications and [email reports](#), an email server must first be set up under [Admin > Integrations > Email Server](#).

To add an email notification to a notification profile, follow these steps:

1. Click the notification profile to open the configuration tray.
2. Under *Actions*, click the + button.

3. In the secondary tray, select *Email* from the action type dropdown.
4. Enter one or more recipient email addresses (comma-separated) in the *To* field.
5. Enter a subject to use in the emails in the *Subject* field.
6. [Optional] Enter a custom email notification message in the *Message* field.

Hint: The default `%m` variable in message field passes the raw event message generated by the alarm policy triggering the notification. This can be replaced with a custom message using any of the [variables supported by the policy](#).

7. Click the **Add** button to save the action configuration to the profile.

Once added, the email notification will be triggered following the [alarm policy's settings for the notification profile](#).

Note: All emails sent by Plexier Scrutinizer, such as alarm notifications and [scheduled email reports](#), will be shown as coming from the address configured in **Admin > Settings > Email Server**.

Logfile

The logfile notification action saves event details to a specified logfile, which can be used for external tracking, investigation, and archival.

Note: Logfiles are saved to `home/plexier/scrutinizer/files/logs`.

To add a logfile action to a notification profile, follow these steps:

1. Click the notification profile to open the configuration tray.
2. Under *Actions*, click the **+** button.
3. In the secondary tray, select *Logfile* from the action type dropdown.
4. In the *File Name* field, enter the name of the file to save the logs to.

Important: Do not include the path when entering the logfile name.

5. [Optional] Enter a custom log message in the *Message* field.

Hint: The default %m variable in message field passes the raw event message generated by the alarm policy triggering the notification. This can be replaced with a custom message using any of the *variables supported by the policy*.

6. Click the **Add** button to save the action configuration to the profile.

Once added, the logfile notification will be triggered following the *alarm policy's settings for the notification profile*.

Syslog

The syslog notification action can be used to send syslog messages containing event details to a specified logging server.

To add a syslog notification to a notification profile, follow these steps:

1. Click the notification profile to open the configuration tray.
2. Under *Actions*, click the + button.
3. In the secondary tray, select *Syslog* from the action type dropdown.
4. Enter the IP address or hostname of the destination logging server in the *Host* field.
5. Enter the UDP port to use on the destination logging server in the *UDP Port* field.
6. Use the dropdowns to select the *severity level* and *type/facility code* to assign to the syslog message.
7. [Optional] Enter a custom log message in the *Message* field.

Hint: The default %m variable in the message field passes the raw event message generated by the alarm policy triggering the notification. This can be replaced with a custom message using any of the *variables supported by the policy*.

- Click the **Add** button to save the action configuration to the profile.

Once added, the syslog notification be triggered following the *alarm policy's settings for the notification profile*.

Syslog priority levels

Plixer Scrutinizer uses the following keyword mappings to assign a priority level to a syslog notification message:

| Keyword | Priority Level |
|---------|----------------|
| emerg | 0 |
| alert | 1 |
| crit | 2 |
| err | 3 |
| warning | 4 |
| notice | 5 |
| info | 6 |
| debug | 7 |

Types/facility codes

Plixer Scrutinizer supports the following keyword mappings for assigning facility codes to a syslog notification messages:

| Keyword | Facility Code |
|-----------------|---------------|
| auth | 4 |
| authpriv | 10 |
| cron | 9 |
| daemon | 3 |
| ftp | 11 |
| kern | 0 |
| lpr | 6 |
| mail | 2 |
| news | 7 |
| syslog | 5 |
| user | 1 |
| uucp | 8 |
| local0 - local7 | 16 - 23 |

SNMP trap

The SNMP trap notification action can be used to automatically create SNMP traps to send event details to a specified SNMP manager.

To add an SNMP trap action to a notification profile, follow these steps:

1. Click the notification profile to open the configuration tray.
2. Under *Actions*, click the **+** button.
3. In the secondary tray, select *SnmptTrap* from the action type dropdown.
4. Using the provided fields, enter the following details for the trap:
 - The IP address or hostname of the destination *Host*
 - The *UDP Port* to use on the destination host (default: 162)
 - The *Community String* to use for authentication on the destination host
5. [Optional] Enter a custom message in the *Message* field.

Hint: The default *variables* in the message field will pass the basic event details, as well as the raw event message generated by the alarm policy. This can be replaced with a custom message as long as the variables used are supported by the policy.

6. Click the **Add** button to save the action configuration to the profile.

Once added, the SNMP trap action will be triggered following the *alarm policy's settings for the notification profile*.

Script

Script notifications allow for more advanced alerts through the use of custom scripts. They can be used to run virtually any scriptable action(s) when the notification profile is triggered by the associated policy.

Hint: Additional configuration steps are required to set up script notifications, but they allow for the most flexibility and sophistication among the different notification action types.

To add a script action to a notification profile, follow these steps:

1. Click the notification profile to open the configuration tray.
2. Under *Actions*, click the + button.
3. In the secondary tray, select *Script* from the action type dropdown.
4. Enter the name of the script file to run in the *Script* field.
5. [Optional] Enter any variables or strings to use as arguments in the *Command Line Arguments* field.

Hint: The default `%m` variable in the arguments field passes the raw event message generated by the alarm policy triggering the notification. This can be replaced any other strings or *variables supported by the policy*.

6. Click the **Add** button to save the action configuration to the profile.

Once added, the script action will be triggered following the *alarm policy's settings for the notification profile*.

Additional notes

- Only script files saved to `/home/plixer/scrutinizer/files` can be run as notification actions.
- When adding variables and strings to the arguments field, use quotation marks (") to enclose terms that should be passed as a single argument.
- Script files must be assigned the appropriate permissions to be run.

Note: When setting a script notification action, using the **Test** button runs the script as the `apache` user. When the notification profile is triggered by an alarm policy, Plexier Scrutinizer will run the script as the `plixer` user.

Auto-acknowledge

The auto-acknowledge notification action can be used to automatically *acknowledge events* for any specified alarm policy (including the policy triggering the action).

To add an auto-acknowledge action to a notification profile, follow these steps:

1. Click the notification profile to open the configuration tray.
2. Under *Actions*, click the **+** button.
3. In the secondary tray, select *Auto Acknowledge* from the action type dropdown.
4. Select the policy to automatically acknowledge from the second dropdown.
5. Click the **Add** button to save the action configuration to the profile.

Once added, the auto-acknowledge action will be triggered following the *alarm policy's settings for the notification profile*.

ServiceNow - Ticket

The ServiceNow notification action can be used to automatically create tickets for a specified *ServiceNow instance* when the notification profile is triggered.

Important: Before configuring ServiceNow notification actions, set up at least one ServiceNow instance via *Admin > Integrations > ServiceNow*. For further details, see the section on *ServiceNow integration*.

To add a ServiceNow notification to a notification profile, follow these steps:

1. Click the notification profile to open the configuration tray.
2. Under *Actions*, click the **+** button.
3. In the secondary tray, select *ServiceNow* from the action type dropdown.
4. Use the dropdown to select the ServiceNow instance to create the ticket for.
5. Under *Short Description*, enter a message to use in the the ticket's Short Description field.

6. Under *Description*, enter a message to use in the ticket's *Description* field.

Note: Any *variables* included in either description field will pass their values when the corresponding fields in the ticket are populated. The raw event message generated by policy triggering the notification can also still be sent using `%m` variable.

7. Use the dropdowns to select urgency and impact levels for the ticket.
8. [Optional] Use the *API JSON* field to send alternate/additional details in the API call by either re-defining any of the default keys used or defining additional keys to include.

Hint: When keys matching the defaults sent with API calls are defined in the API JSON field, the new values will overwrite the defaults. Any new keys defined will be appended to the API call.

9. Click the **Add** button to save the action configuration to the profile.

Once added, the ServiceNow notification will be triggered following the *alarm policy's settings for the notification profile*.

Important: To be able to fully populate all corresponding fields in the ServiceNow ticket, the *user configured for the selected instance* must be provisioned with the `sn_incident_write` permission.

Sample JSON key definition

Key definitions in the API JSON field should follow the following format:

```
{
  "extra_data": "my data: %m"
}
```

CEF

The CEF notification action uses CEF (Common Event Format) syslog messages to forward alarm/event details to external applications.

To add a CEF notification to a notification profile, follow these steps:

1. Click the notification profile to open the configuration tray.
2. Under *Actions*, click the **+** button.
3. In the secondary tray, select *CEF* from the action type dropdown.
4. Enter the IP address or hostname of the host to send the CEF syslog message to.
5. Enter the port UDP port to use on the destination host.
6. Click the **Add** button to save the action configuration to the profile.

Once added, the CEF notification will be triggered following the [alarm policy's settings for the notification profile](#).

CEF message mapping

Based on the standard CEF message format (CEF:Version|Device Vendor|Device Product|Device Version|Signature ID|Name|Severity|Extension), Plixer Scrutinizer uses the following mapping for the first seven (prefix) keys:

Prefix keys

The first seven keys of the CEF message will use the following standard mappings across all alarm policies:

| Key | Value |
|----------------|-------------------------------|
| Version | 1 |
| Device Vendor | Plixer |
| Device Product | Scrutinizer |
| Device Version | \${SCRUTINIZER_VERSION} |
| Signature ID | \${EVENT_POLICY_LANGKEY} |
| Name | \${EVENT_POLICY_NAME} |
| Severity | \${EVENT_SEVERITY_AS_INTEGER} |

Extension keys

Because the CEF message is automatically generated using the event message of the alarm policy violated, the extension keys included will vary based on what details/fields are reported under the policy.

The following table lists all mappings that may be used for event details in the CEF message:

| CEF Key | Event Key |
|---------|--------------------|
| app | app_proto |
| cnt | hits |
| dpt | dst_port |
| dst | target |
| duser | target_username |
| dvc | devices |
| end | last_ts |
| proto | protocol |
| spt | src_port |
| src | violation |
| start | first_ts |
| suser | violation_username |

Note: By default, Plixer Scrutinizer maps the `dst` and `src` CEF keys to the `target` and `violation` event keys exclusive to Plixer Scrutinizer's *Report Threshold Violation* alarm policy. These are **not** same general `targets` and `violators` keys that are common to all events. This is to support a specific use case for report thresholds.

Sample CEF message sent by Plixer Scrutinizer:

```
CEF:1|Plixer|Scrutinizer|${SCRUTINIZER_VERSION}|${EVENT_POLICY_
  ↳LANGKEY}|${EVENT_POLICY_NAME}|${EVENT_SEVERITY_AS_INTEGER}|dvc=$
  ↳{EVENT_DEVICES} start=${EVENT_FIRST_TS} end=${EVENT_LAST_TS} cnt=$
  ↳{EVENT_HITS}
```

To learn more about the customization of Plixer Scrutinizer CEF key mappings, contact [Plixer Technical Support](#).

Variables in notifications

When defining a [notification action](#), the message sent can be customized to include additional event details passed through variables.

Note: The default `%m` variable used in notification messages will pass the event message generated by the alarm policy triggering the notification. Message formats by policy can be viewed via the [policy management page](#).

The following table lists the variables available for use in notification messages or custom scripts:

| | |
|-------------------------------|--|
| <code>%m</code> | Event message generated by the alarm policy triggering the notification |
| <code>%pol</code> | Alarm policy violated to trigger the notification |
| <code>%v</code> | IP address(es) of violating host(s) reported in the event |
| <code>%url</code> | URL to the relevant saved report (only available for the <i>Report Threshold Violation</i> alarm policy) |
| <code>%h</code> | IP address of the host (i.e., Plexier Scrutinizer server/reporter) sending the notification |
| <code>%v_resolved</code> | Resolved hostnames of violator addresses |
| <code>%id</code> | The log identifier for the event that triggered the notification |
| <code>%h_resolved</code> | Resolved hostname of the address sending the notification |
| <code>%violation_users</code> | Username(s) associated with violating host(s) |
| <code>%time</code> | Timestamp of the event/violation that triggered the notification |
| <code>%p</code> | Protocol used in the violation, if applicable |
| <code>%t</code> | IP addresses of the host(s) targeted in the violation, if applicable |
| <code>%tactic_id</code> | MITRE ATT&CK framework ID of the tactic under which the violation is classified * |
| <code>%tactic_name</code> | MITRE ATT&CK tactic under which the violation is classified * |
| <code>%target_users</code> | Username(s) associated with targeted host(s) |
| <code>%technique_id</code> | MITRE ATT&CK framework ID of the technique associated with the violation * |
| <code>%technique_name</code> | MITRE ATT&CK technique associated with the violation * |
| <code>%category</code> | Alarm policy category of the violated policy |

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

Security groups

The **Admin > Alarm Monitor > Security Groups** page can be used to create, edit, and manage security groups, which allow similar devices to be quickly added to FA algorithm inclusion lists.

Once an FA algorithm is added to a security group, it is enabled for all exporters assigned to that group. Changes to group membership are also automatically applied to the inclusion lists of associated algorithms.

Hint: The default *Firewalls*, *Core Exporters*, *Edge Exporters*, and *Defender Probes* security groups are predefined as inclusions for the recommended FA algorithms and need only be populated with the specified device type.

New security groups can be added by clicking the **+** button in the main view. Membership and enabled algorithms for existing groups can be edited at any time.

Adding exporters to a security group

To add exporters to an existing security group, follow these steps:

1. Click the security group to open the configuration tray.
2. Expand the *Active Exporters* of the tray and click the **Add** button.

Hint: To remove an exporter from the group, click the **Delete** icon in the list.

3. In the secondary tray, use the checkboxes to select the exporters to add.
4. Click the **Add** button to assign all selected exporters to the group.

Any algorithms enabled for the security group will automatically be enabled for the new exporters added.

Hint: To add exporters to multiple security groups, select the groups in the main view and click the **Bulk Actions** button.

Enabling algorithms for a security group

To enable FA algorithms for an existing security group, follow these steps:

1. Click the security group to open the configuration tray.
2. Expand the *Algorithms* of the tray and click the **Add** button.

Hint: To disable an algorithm for the group, click the **Delete** icon in the list.

3. In the secondary tray, use the checkboxes to select the algorithms to enable.
4. Click the **Add** button to enable all selected algorithms for the group.

New algorithms added will automatically be enabled for all exporters in the security group.

Hint: To enable algorithms for multiple security groups, select the groups in the main view and click the **Bulk Actions** button.

Reports

The **Admin > Reports** category includes management views for report-related functions.

| | |
|--------------------------------|---|
| Flow Report Thresholds | Define <i>custom report thresholds</i> to trigger alarms and/or notifications |
| Report Designer | Create custom report type configurations |
| Report Folders | Create and manage <i>folders</i> to organize <i>saved reports</i> |
| Scheduled Email Reports | Set up and manage <i>scheduled email report</i> configurations |

Note: *Report threshold*, *folder*, and *scheduled email report* management options can also be accessed from main Reports views of the web interface.

Plixer

The **Admin > Plixer** category provides access to configuration and management views for Plixer product integrations:

Click a product/settings subcategory below to learn more:

| | |
|--|--|
| Plixer Endpoint Analytics | Configure and enable/disable Plixer Endpoint Analytics integration |
| Plixer FlowPro Licensing | Add a Plixer FlowPro license key and view license details (only for Plixer FlowPro 20.0.0 and above) |
| Plixer Replicator | Configure and enable/disable Plixer Replicator integration |
| Plixer Scrutinizer Licensing | Add a Plixer Scrutinizer license key and view license details |

Plixer Endpoint Analytics

The **Admin > Plixer > Endpoint Analytics** page is used to configure and enable/disable Plixer Endpoint Analytics integration in Plixer Scrutinizer.

Hint: To learn more about Plixer Endpoint Analytics integration in Plixer Scrutinizer, see [this section](#) of this documentation.

The following details must be entered:

| | |
|-----------------|--|
| Host | IP address or configured hostname of the Plixer Endpoint Analytics appliance |
| Password | Password to use for authentication with Plixer Endpoint Analytics |
| Port | Port to use to communicate with Plixer Endpoint Analytics |
| Protocol | Protocol to use to communicate with Plixer Endpoint Analytics |
| Username | Username to use for authentication with Plixer Endpoint Analytics |

After the fields have been filled in with the required information, click **Save**.

Hint: Click the **Defaults** button to revert to the default settings used by Plixer Endpoint Analytics deployments.

Plixer FlowPro Licensing

The **Admin > Plixer > FlowPro Licensing** page is used to add or update a Plixer FlowPro license after one or more appliances have been deployed.

Hint: To learn more about Plixer FlowPro integration in Plixer Scrutinizer, see [this section](#) of this documentation.

To learn more about licensing options or obtain an active key, contact [Plixer Technical Support](#).

The following license details are also displayed on this page:

- Product/license type
- License status
- Validity duration left
- Customer ID
- Machine ID
- Number of FlowPro probes supported by the license
- Number of FlowPro probes currently deployed
- Number of FlowPro probes currently registered

To add or update a license key, paste the new key into the *License Key* field, and then click **Save**.

Plixer Replicator

The **Admin > Plixer > Plixer Replicator** page is used to configure and enable/disable Plixer Replicator integration in Plixer Scrutinizer.

The following details must be entered:

| | |
|------------------------|---|
| Password | Password to use for authentication with the Plixer Replicator appliance |
| Receive Port | Receiving port on the Plixer Replicator appliance |
| Replicator Host | IP address or hostname of the Plixer Replicator appliance |
| Seed Profile | Plixer Replicator profile containing exporters for auto-replication |
| Send Port | Port to use for sending flows to Plixer Scrutinizer |

After the fields have been filled in with the required information, click **Save** to automatically configure the Plixer Replicator appliance.

Hint: Click the **Defaults** button to revert to the default settings used by Plixer Replicator deployments.

Plixer Scrutinizer licensing

The **Admin > Plixer > Scrutinizer Licensing** page is used to add or update the Plixer One or Plixer Scrutinizer license key, after *system has been deployed and set up*.

To learn more about licensing options or obtain an active key, contact [Plixer Technical Support](#).

The following license details are also displayed on this page:

- Product/license type
- License status
- Validity duration left
- Customer ID
- Machine ID
- Number of Plixer Scrutinizer servers/appliances supported by the license
- Number of reporting servers supported by the license
- Number of exporters supported by the license
- Number of servers (reporter + remote collectors) currently deployed
- Number of exporters currently enabled
- Plixer One Enterprise license status

To add or update a license key, paste the new key into the *License Key* field, and then click **Save**.

Note:

- This admin category includes pages/views from the **Admin > Settings** section of the Plixer Scrutinizer Classic UI.
 - Additional licensing may be required to enable integration with certain Plixer components. Contact [Plixer Technical Support](#) to learn more.
-

Resources

The **Admin > Resources** category provides access to pages/views for monitoring and managing Plixer Scrutinizer features and elements in the environment.

Click on a settings subcategory below to learn more:

| | |
|--|--|
| <i>Collectors</i> | Manage Plixer Scrutinizer collectors and Plixer ML Engines in the environment |
| <i>Exporters</i> | Manage and add protocol exclusions to flow-exporting devices in the environment |
| <i>FlowPro Capture Rules</i> | Define and manage packet capture rules for FlowPro Probes |
| <i>FlowPro Probes</i> | Manage Plixer FlowPro probes sending data to Plixer Scrutinizer collectors |
| <i>Interfaces</i> | Manage Plixer Scrutinizer settings and SNMP credentials for individual interfaces |
| <i>ML Engines</i> | Manage hose settings for Plixer Machine Learning Engine |
| <i>SNMP Credentials</i> | Manage SNMP credential sets for polling exporters in the environment |
| <i>System Performance</i> | View current and predicted resource utilization for individual Plixer Scrutinizer collectors |

Collectors

The **Admin > Resources > Collectors** page is used to access management functions for Plixer Scrutinizer flow collectors and Plixer ML Engine deployments.

Managing collectors

The **Collectors** tab lists the following details for each Plixer Scrutinizer collector currently deployed:

| | |
|------------------------|---|
| Rank | Number assigned to the collector as part of a distributed cluster |
| Collector | Collector's IP address or hostname |
| Status | Collector's current operational status |
| Exporter count | Number of exporters sending flows to the collector |
| First flow time | Timestamp of the first flow received by the collector |
| Last flow time | Timestamp of the most recent flow received by the collector |
| 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 |

To change the information displayed, click the *Available Columns* button and select the details to show. Full details can also be viewed in a tray by clicking the collector IP address or hostname.

Deleting collectors

To remove one or more collectors from the environment, select them using the checkboxes, and then use the *Delete* option in the **Bulk Actions** tray.

Exporters

The **Admin > Resources > Exporters** page is used to inspect and manage all flow-exporting devices in the Plixer Scrutinizer environment.

The page's main view/table lists the following details for all exporters/interfaces:

- Availability (icon):
 - Green: Up on all collectors
 - Red: Down on all collectors
 - Yellow: Flows not being received *as expected*
 - Grey: Exporter is disabled, unlicensed, or configured as a backup
- Exporter's configured name, IP address, or hostname
- Flow format and version
- Current status:
 - *Enabled*: Flows are being collected as normal
 - *Backup*: Enabled as a backup but does not count towards license limit
 - *Unresourced Enabled*: Disabled due to low resources but counts towards license; Automatically set to *Enabled* once resources become available
 - *Unresourced Backup*: Disabled due to low resources and does not count towards license; Automatically set to *Backup* once resources become available
 - *Disabled*: Manually disabled and does not count towards license

- *Unlicensed*: Disabled due to license exporter limit
- Average missed flow sequence numbers (MFSNs) per second
- Average flow rate
- Number of interfaces
- Average packet rate
- Timestamp of first flow received from the exporter (hidden by default)
- Timestamp of most recent flow received from the exporter
- IP address or hostname

Note:

- The default sorting order for the table/list is by flow rate (flows/s).
 - To change what details are displayed in the table, click the **Available Columns** button and select the columns to display.
 - Rates and other details displayed are relative to the specified collector
 - The *Unresourced Enabled* and *Unresourced Backup* states are automatically applied by Plixer Scrutinizer as part of its low resource fallback functions. However, they can also be manually set to prioritize disabling specific exporters before others, when the system becomes underprovisioned.
-

Configuring exporter settings

Clicking on an exporter in the table opens a settings tray containing the following actions/functions:

- Add or edit a custom exporter name
- View collectors receiving flows from the exporter
- View available interfaces and any *custom interface speeds* configured
- Add or edit *protocol exclusions*

- View *FA algorithms currently being applied*
- View or edit the *SNMP credentials* used to poll the device
- Add or view tags associated with the device

Adding protocol exclusions

To ensure that only relevant traffic data is collected from an exporter, one or more protocol exclusions can be defined as follows:

1. After opening the settings tray for an exporter, click the edit (pencil) icon for **Protocol Exclusions**.
2. In the secondary tray, verify that the correct exporter is selected under *Device*.
3. Use the *Interface* dropdown to select the interface/instance to apply the protocol exclusion to.
4. Use the *Protocol* dropdown to select the protocol to exclude from collection.
5. Click the **Add** button to save the protocol exclusion for the exporter.

Protocol exclusions are saved by exporter and will be applied, even if a device is set to send flows to a new/different collector.

Exporter management

The exporter management/configuration functions become available via the **Bulk Actions** tray, when one or more exporters are selected using the checkboxes in the list/table:

- Change exporter status (e.g., enabled/disabled, backup, unresourced enabled/backup)
- Set SNMP credentials to use for polling
- Change polling method (IP address, hostname, or disabled)
- Enable/disable *Ignore Flow Duration* option for selected exporters
- Enable/disable *Ignore MFSNs* option for selected exporters
- Enable/disable *Ignore Outage* option for selected exporters
- Poll selected exporters to update saved SNMP information

- Delete selected exporters

Additional exporter settings

Certain exporters may not send flow data according to Plixer Scrutinizer's expected patterns, which will be indicated by a yellow availability icon in the main list/table.

The following settings can be toggled on to allow for such irregular behavior:

- *Ignore Flow Duration* - Should be enabled for devices that export flows less frequently than the recommended once every minute
- *Ignore MFSNs* - Should be enabled for devices that do not send flow sequence numbers correctly, resulting in MFSNs
- *Ignore Outage* - Should be enabled for devices when intermittency is expected

FlowPro capture rules

The **Admin > Resources > FlowPro Capture Rules** page is used to define and manage selective packet capture rules for Plixer FlowPro probes.

To learn more about selective packet capturing, see this page in the Plixer FlowPro documentation.

Adding a new rule

To define a new rule, click the add (+) button in the main capture rules view, and then configure the following details in the tray:

- **Name:** A name for the capture rule
- **Client IP:** Client/destination IP address of packets to capture
- **Server IP:** Server/source IP address of packets to capture
- **Well-Known Port:** Well-known port to monitor for packets
- **Max Packets:** Maximum number of packets to capture
- **Stops On:** End date for capturing packets
- **Retain Until:** End date for retaining captured packet data

Captures can be downloaded by clicking **Download PCAP** for events under the *FlowPro Event Capture* policy in the Alarm Monitor views.

Note:

- Packets will start being captured as soon as a rule is saved (if enabled). Rules with captured data will be indicated by a check in the *Data* column.
 - If the capture download link does not work, navigate to `https://FLOWPRO_MGMT_IP:8080/` and clear the certificate error before trying again.
 - The timezones configured on the Plixer Scrutinizer server and the Plixer FlowPro probe must be the same for the *Stops On* rule to be correctly observed.
 - Capture rules can also be created and managed via [API](#).
-

Rule management

Once the maximum number of packets has been captured, or the defined end date has been reached, a rule will automatically be disabled. Inactive rules will be marked with a yellow indicator in the main view/table instead of green (enabled/active).

To continue capturing packets, click on the rule name, make the necessary changes (*Max Packets* or *Stops On*) in the configuration tray, and then re-enable the rule.

Rules that are no longer needed can instead be deleted. To do this, use the checkboxes to select one or more rules to be deleted, and then use the *Delete* option in the **Bulk Actions** menu/tray.

FlowPro Probes

The **Admin > Resources > FlowPro Probes** page is used to add/register, configure, and manage Plixer FlowPro probes/appliances (v20.0.0+ only).

The main view of this page lists all registered probes along with the following details for each one:

- Name assigned to the probe
- IP address of the probe's MGMT interface
- IP address of server/collector used for the probe

- License support status
- APM key registered for the probe
- FlowPro APM status
- FlowPro Defender status
- Authentication token used by the probe

Clicking on a probe name opens a [configuration tray](#) where settings for that probe can be configured.

Adding a probe

To add/register a new probe, click the add (+) button in the main view, and then configure the following details in the tray:

- **Name:** A name for the probe
- **IP Address:** IP address to be assigned to the probe's MGMT interface
- **Collector:** Plixer Scrutinizer server or remote collector to be assigned to the probe

The *Default NIDS Rules* option can also be toggled on to apply NIDS rules from open-source threat feeds for network event reporting.

Important: This step must be completed before the corresponding probe appliance is deployed. After the appliance's first boot sequence, the IP address assigned to the MGMT interface must match the IP address entered in Plixer Scrutinizer.

Note:

- Additional probes can be registered and deployed if supported by the current license. Check the [Plixer FlowPro licensing page](#) for details.
 - If default NIDS rules are disabled, the probe will only send basic IPFIX observations unless custom rules are manually added.
-

Probe configuration/management

After a probe has been registered, the following settings/options can be modified via the configuration tray:

- Probe/appliance name
- MGMT interface IP address
- Collector assigned
- Registered APM key
- Enabled features by interface

Important: If a probe is deleted or the IP address registered for its MGMT interface is changed, the corresponding appliance will need to be re-deployed to assign the new IP address.

To delete/deregister one or more probes, select them using the checkboxes in the main view, and then select *Delete* in the **Bulk Actions** menu.

Interfaces

The **Admin > Resources > Interfaces** page is used to manage interface settings for flow-exporting devices in the Plixer Scrutinizer environment.

The page's main view/table lists the following details for each device instance (if available or configured):

- *Configured name*, IP address, or hostname of the device
- Instance name
- Custom description
- ifAlias
- ifName
- ifDescr
- ifSpeed
- Custom inbound interface speed
- Custom outbound interface speed

- Metering directionality

The page's **Options** tray also includes additional toggles to show/hide inactive interfaces and make hidden interfaces visible.

Note:

- Select *Information* in the three-dot menu to view basic details for a device. Selecting *Summary of Device* opens the **Admin > Resources > Exporters** view filtered on the device.
 - To change what details are displayed in the table, click the **Available Columns** button and select the columns to display.
 - Custom descriptions and interface speeds are only used by Plixer Scrutinizer (displaying utilization, threshold alerts, etc.). They are not applied to the device.
-

Interface Settings

Clicking on an instance name opens a settings tray where the following details can be configured for the interface:

- Custom description
- Custom inbound speed
- Custom outbound speed
- Hide/show setting
- SNMP credentials

To hide an instance in the UI, select *Yes* in the **Hidden** dropdown or tick the *Hide* checkbox in the main view. These instances can be made visible by toggling on *Show interfaces hidden in the UI* in the **Options** tray.

Hint: The above settings can also be applied to multiple interfaces by selecting the instances in the main view and making the configuration changes in the **Bulk Actions** tray.

ML Engines

The **Admin > Resources > ML Engines** page is used to manage host settings for each Plixer ML engine deployed in the Plixer Scrutinizer environment.

The page's main view/table lists the following details for each device instance (if available or configured):

- ML engine name
- Hostname
- Type
- Engine status
- Deploy status
- *Authentication token*
- Last modified

Adding an ML engine

To add a new ML engine, click the add (+) button in the main view, and then configure the following details in the tray:

- **Name:** A name for the ML engine
- **Type:** Type of ML engine paired with your Plixer Scrutinizer environment
 - *Single VM*
 - *Amazon AWS*
 - *Azure*
 - *vSphere multi VM Cluster*

After clicking **Save**, click the newly-added ML engine to open the tray and make sure to take note of the generated auth token and the primary reporter IP. Once done, proceed to deploying the ML engine. For instructions on how to deploy an ML engine, see the *[Plixer ML Engine deployment guide](#)*.

ML engine management

Clicking on an ML engine name opens a settings tray where the following details can be configured for the ML engine:

- Name
- Ingestion replica count
- Train anomaly detection replica count
- Ingestion minimum CPU
- Ingestion maximum CPU
- Ingestion minimum memory
- Ingestion maximum memory
- Elasticsearch memory
- Elasticsearch minimum CPU
- Elasticsearch maximum CPU
- Enable/disable Kibana

Note: The settings tray can also be accessed by clicking the three-dot menu beside the ML engine name, and then clicking **Settings**. For more information, see the [ML engine settings](#) section.

One or more ML engines can be deleted by selecting them using the checkboxes, and then using the *Delete* option in the **Bulk Actions** tray.

The settings tray is designed for configuring the settings of a specific ML engine. To configure settings across multiple ML engines, go to one of the following:

- [Admin > Settings > ML AD Users](#)
- [Admin > Settings > ML Alerts](#)
- [Admin > Settings > ML Data Limits](#)
- [Admin > Settings > ML Training Schedule](#)

SNMP credentials

The **Admin > Resources > SNMP Credentials** page can be used to add/manage sets of SNMP credentials for use with devices/exporters in the Plexier Scrutinizer environment.

Once defined/saved, credentials can be assigned to one or more specified exporters from the [exporters management view](#). SNMP v1, v2, and v3 are all 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 then click **Save**.

Saved credentials can also be edited at any time by clicking on their name in the main view table. To delete one or more credential sets, tick their checkboxes 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 the Plexier Scrutinizer environment.

The page is divided into a graph/timeline and a summary table listing current allotment and utilization details for each collector.

Utilization timeline

The timeline can display the following utilization details (select from the dropdown) for all collectors for the past 24 hours:

- CPU utilization (%)
- Available memory (GB)
- Host index size (%)
- Alarm database size (%)

To highlight utilization and view general information for a specific collector, hover over its line in the graph.

Collector utilization details

Drilling down into a collector from the summary table opens a more detailed view with the following information:

- Current total vs. predicted utilization based on current disk capacity
- Current vs. predicted maximum disk utilization, based on current flow volume and [data retention settings](#)
- Current disk utilization per [roll-up interval](#) vs. predicted maximum, based on the number of days the data is configured to be stored

The default **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.

The **Feature Resources** summary/management view for the collector and [recommended resource allocations](#) tables can also be accessed via the **Chart** dropdown.

Feature Resources

The **Feature Resources** view can be used to inspect and manage resource usage by feature set.

The page's main view lists all available feature sets, alongside the following details:

- Current state (green: active, grey: inactive)
- *Importance*
- Number of active alarms indicating resource issues for the feature set
- Expected CPU core usage per collector
- Expected RAM usage per collector
- Number of FA algorithms associated with the feature set
- Number of alarm policies associated with the feature set

Users are also able to toggle between graphs showing algorithms, policies, CPUs, or RAM per feature set.

Enabling/disabling feature sets

To allow teams to better adapt Plixer Scrutinizer's functions to monitoring and resource requirements, related *FA algorithms* and their associated *alarm policies* can be disabled/deactivated by feature set instead of individually disabling them via the respective management pages.

Clicking on a feature set name opens the information tray, where it can be activated or deactivated via a toggle. All FA algorithms and alarm policies included in the feature set are also listed in this tray.

Important: Deactivating services may result in loss of functionality and/or other issues. Contact *Plixer Technical Support* for assistance.

Low resource fallback modes

When the total expected resource utilization results in the current allocations falling below the *recommended values* for the observed exporter count and flow rate, Plexier Scrutinizer can automatically pause certain functions as low resource fallback.

There are two *low resource fallback modes* that can be enabled:

- LRF_mode_pauseFeatureSets - Pause feature sets before pausing exporters
- LRF_mode_pauseExporters - Pause only exporters

When low resource fallback becomes necessary, feature sets are paused based on their *importance* value (lowest first, 100 = never paused).

Features and/or exporters will automatically be resumed when the configured CPU core and RAM allocations can support additional computational load.

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, Plexier Scrutinizer will also continuously attempt to resume paused feature sets.

Additional low resource fallback settings

The following settings under *Admin > Settings > Collector* can be modified to further 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/chunk
- Flow rate multiplier/percentage for accommodating brief, recoverable spikes

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.

The Classic UI remains accessible either via the URL `https://scrutinizer_ip/oldui/` or by toggling the appropriate setting in the **Admin > Users & Groups > User Accounts > Preferences** tray.

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 user group are available to other users in the same user group. 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 user group 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 user group 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 user group 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 **User Group Dashboards** option. This requires the **Dashboard Admin** permission.
- User Group Dashboards:

Use this option to select the dashboards a user group 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 user group, 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 > User Groups:**

1. Choose Dashboard Gadgets

- Click the “Dashboard Gadgets” value for the user group you want to change
- Uncheck “All Dashboard Gadgets”
- Move the individual gadgets the selected user group 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 user group 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 Plexier 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 Plexier Scrutinizer UI. To learn more about upgrading to the latest version of Plexier 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:** 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:
- Device Status: Lists all of the flow sending devices with corresponding details.
- 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.
- **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. 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 the poller and 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 > Integrations > 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.

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).
- **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 Plexier 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. $TI = violations * threat\ multiplier$.
- **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 country of the IP. 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 are 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** - 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 Plixer 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 Plixer Technical Support](#) for assistance.
- **Create New Board** enables the user to create new or delete existing 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 [user group permissions](#).

- **Flow Analytics Configuration**

The overall status of all algorithms and the total runtime and count of violations across all algorithms. For more information, see the [Flow Analytics Configuration](#) section.

- **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 Plixer 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 Plixer Scrutinizer server, do not require any internet access.
- **Google:** Useful for displaying network devices geographically.

Highlights:

- Flow devices can be added to more than 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 [user group permissions](#)

Objects

Objects come in four formats:

- **Devices:** are manually selected part of creating a new map.
 - **Label:** Label to assign to the device.
 - **Poll Using:** Select IP Address, Hostname or Disable Polling.
 - **Notification:** Select a notification profile which will be triggered when the threshold is exceeded.
 - **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".
- **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 is based on 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
 - 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 Plexer 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 *collect optionsummary* scrut_util command 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.
- **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).

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 Plexier 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 Plexier 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 Plexier Scrutinizer.
- **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 Plexier 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 Plexier 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 user group name. If it does, they're added to the Plixer Scrutinizer user group 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 user group 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 user group 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* user group (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* user group (if they were a member when they logged in).

Important: This feature requires the names of the local user group and the LDAP security group to be an *exact* match, including any capitalization and/or punctuation.

LDAP servers

To provide LDAP server redundancy, 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 Plixer Scrutinizer

- Copy the Azure AD Identifier value.

Note: The values and the certificate you copied will be required to complete the Plixer Scrutinizer configuration.

This completes the Azure configuration. You should now assign users or groups to the **Plixer Scrutinizer** application in Azure ADFS so that they can successfully authenticate.

Plixer Scrutinizer configuration

Now that you have the required information from Azure's configuration, you can set up Plixer Scrutinizer's authentication. Log into Plixer 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 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. “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 user group 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/plixer/scrutinizer/azure.cert” |

4. Click **Save** to save the configuration.

A new row will appear in Plixer 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 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 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 drop-down 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>/cgi/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 https://<scrutinizer_server>/ |
| Assertion URL | Enter in the format of https://<scrutinizer_server>/fcgi/scrut_fcgi.fcgi?rm=usergroups&action=sso_res |
| Audience Value | Enter in the format of https://<scrutinizer_server>/ |
| 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 “/home/plixer/scrutinizer/okta.cert” |

- 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 **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 [user group 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.

User group permissions

Users are assigned to user groups. User groups are granted permissions. Users inherit permissions from all the user groups 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 user group that individual users can be assigned to. Give the group a name and apply a template from another user group that has similar permissions to the new user group. After creating an account, find the new user group on the left and click it to modify.

[Click here](#) for a special note regarding Plexier Scrutinizer user groups 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 user groups it is a member of. To view all the user groups 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 user group. A user can be a member of multiple user groups and inherit all applicable permissions.

Features

Permissions control features the user group should have access to within Plexer Scrutinizer. Permissions can restrict product features entirely for a user group or specific features can be accessed based on your user group membership.

Features include:

- Which tab the members of the user group should be able to see,
- Administrative permissions the user group 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 user group 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 user groups to use predefined roles or manually specifying features. A user group 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-Admin-istrator | ackBBEvent alarmSettings almDelete LogalotPrefs NotificationManager Policy-Manager |
| Alarm-sUser | alarmsTab |
| Dash-board-Admin-istrator | dashboardAdmin |
| Dash-board-User | createDashTabs myViewTab |
| Map-sAdmin-istrator | mappingGroupConfiguration mappingObjectConfiguration |
| Map-sUser | adminTab allLogalotReports mapsTab reportFilters statusTab |
| Re-portin-gAdmin-istrator | ApplicationGroups asnames deleteReport HostNames protocolExclusions report-Settings tos viptelaSettings wkp |
| Re-porting-PowerUser | reportFolders ReportDesigner saveReport scheduledReports srCreate |
| Re- | runReport |

- **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 user group will need to have access to run.
- **Dashboard gadgets** selects the gadgets that the user group will need to be able to add to dashboards.
- **Third-party links** controls the vendor third-party integrations that the user group will be able to integrate with.
- **Bulletin boards** manages the bulletin boards that the user group will need to be able to access in the Alarms tab.

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

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 Plexier Scrutinizer's advanced functions and includes configuration guides as well as additional background information related to their use.

6.1 Backups

The Plexier 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 Plexier 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 Plexier 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.

6.1.1 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`.

Note: The default runmode of `backup.sh` *saves the backup file locally*. Due to the size of full backup files, however, the remote method outlined below is recommended.

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:

```
df -h /var/db | awk '!/^Filesystem/ {print "Space Required: "$3}'
```

4. Run `backup.sh` as the `plixer` user, with the mounted remote directory set as the backup file location:

```
BACKUPDIR=/mnt/backup ~plixer/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 `plixer` user.

Note:

- A *second Plexier Scrutinizer instance* can be used as the remote backup host, provided it has sufficient disk space available and is running the same Plexier Scrutinizer version as the instance to be backed up. However, doing so is only recommended for redundancy.
 - If no remote hosts are available, backups can be *saved locally on the same Plexier Scrutinizer instance*. However, this will limit the amount of storage available for system functions and is not recommended.
-

For further details or assistance with issues, contact *Plexier Technical Support*.

Backing up additional files

When creating a full backup of a Plexier 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 --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 RESTORED_INSTANCE_IP
```

- 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:

```
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:

```
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 Plexier 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 Plexier 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.
-

6.1.2 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 | <code>./conf.dump</code> |
| PGHOST | IP address or hostname of the PostgreSQL database | <code>localhost</code> |
| PGUSER | Role/user used to connect to PGHOST | <code>plixer</code> |
| PGDATABASE | The database to access at PGHOST | <code>plixer</code> |

Backing up configuration data

To create a backup of a Plixer Scrutinizer server's current configuration data, follow these steps:

1. Stop the `httpd` and `plixer_flow_collector` services:

```
sudo /bin/systemctl stop httpd
sudo /bin/systemctl stop plixer_flow_collector
```

2. Run the backup script.

To save the backup file to the default location:

```
~/scrutinizer/database/utils/scrut_conf_dump.sh
```

To use a custom location/filename:

```
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
```

3. Restart the stopped services:

```
sudo /bin/systemctl start httpd
sudo /bin/systemctl start plixer_flow_collector
```

Restoring configuration data

To restore configuration data to a Plixer Scrutinizer server from a backup file, follow these steps:

1. Stop the `httpd` and `plixer_flow_collector` services:

```
sudo /bin/systemctl stop httpd
sudo /bin/systemctl stop plixer_flow_collector
```

2. Run the restore script.

To restore from the default backup location/file:

```
~/scrutinizer/database/utils/scrut_conf_dump.sh
```

To restore from a specified location/file:

```
PGHOST=SCRUTINIZER_IP
DUMP=/tmp/CONF_BACKUP_DIR/CONF_BACKUP.dump ~/scrutinizer/database/utils/
↪scrut_conf_restore.sh
```

3. Restart the stopped services:

```
sudo /bin/systemctl start httpd
sudo /bin/systemctl start plixer_flow_collector
```

4. Resync the access table:

```
psql -c "SELECT setval(pg_get_serial_sequence('plixer.access', 'access_id'
↪'), COALESCE(max(access_id) + 1, 1), false) FROM plixer.access;"
```

Note: `scrut_conf_restore.sh` should only 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.

6.1.3 Data migration

The [scrut_util](#) migration command/utility can be used to perform the following migration operations between a source and a destination Plexier Scrutinizer server:

- Migrate configuration data from a Plexier Scrutinizer 19.4.0 (CentOS) server to v19.5.0+ (Oracle Linux) server
- Migrate historical data from a Plexier Scrutinizer 18.20+ server to a v19.5.0+ server
- Migrating configuration data from a source Plexier Scrutinizer 19.5.0+ server to a destination server on the same version

Migrating both configuration and historical data to a Plexier Scrutinizer 19.5.0+ server requires the source to be on v19.4.0. Historical data can be migrated directly to a v19.5.0+ server from any server on v18.20 or higher.

Important: Because the steps outlined here can potentially result in an irrecoverable state for the source and/or destination Plexier Scrutinizer server, it is highly recommended to contact [Plexier Technical Support](#) for assistance. In case of issues, provide the Plexier support engineer with the migration utility logs (`/var/log/migrate.log`) for troubleshooting.

6.2 Certificate management

This section contains additional instructions/guides related to certificate management in Plixer Scrutinizer.

For further information or assistance with these functions, contact [Plixer Technical Support](#).

6.2.1 Certificate rotation/regeneration

The following certificate rotation utilities can be run to re-issue certificates and keys to address database communication issues:

Note:

- The optional DAYS flag can be used to set an expiration date for the certificate(s) regenerated by each utility. Once they expire, the same command can be run again to re-issue certificates with new expiry dates.
 - With the exception of `set ssl on`, these commands *cannot* be run from the SCRUTINIZER> prompt.
-

| | |
|---|---|
| <div data-bbox="221 254 714 335"><pre>scrut_util --rotatecerts [--days →<DAYS>] [--reset] [--verbose]</pre></div> | <p>Regenerates all certificates on all nodes, including any ML engines, with an optional expiration of the number of DAYS specified</p> <p><i>If the reset flag is included, the CA and Apache web server certificates on the primary reporter will also be regenerated.</i></p> |
| <div data-bbox="221 571 714 652"><pre>scrut_util --httpdcerts [--days →<DAYS>] [--csr] [--verbose]</pre></div> | <p>Regenerates the Apache web server certificate and key with an optional expiration of the number of DAYS specified</p> <p><i>The csr flag can be used to create a certificate signing request (CSR) using the current private key and user-configured subject in /home/plixer/scrutinizer/files/scrutinizer.csr instead of generating a new certificate.</i></p> |
| <div data-bbox="221 977 714 1024"><pre>SCRUTINIZER> set ssl on</pre></div> | <p>Generates a new self-signed certificate (/etc/pki/tls/certs/ca.crt), private key (/etc/pki/tls/private/ca.key), and CSR (/etc/pki/tls/private/ca.csr) with the details entered</p> <p><i>The provided syntax must be run from the SCRUTINIZER> prompt. To run the command directly from the shell, use</i></p> <pre>scrut_util --set ssl --toggle [on off] --port <TCP_PORT> --country <COUNTRY> --state <STATE/PROVINCE> --city <CITY/LOCALITY> --org <ORG_NAME> --email <CONTACT_EMAIL> --name <COMMON_NAME> --keysize [1024 2048 4096]</pre> |
| <div data-bbox="221 1582 714 1693"><pre>scrut_util --mlcerts --ip <IP_ ADDRESS> [--days <DAYS>] [-- →install] [--verbose]</pre></div> | <p>Regenerates TLS certificates and private keys on the Plixer ML Engine node with the specified IP_ADDRESS with an optional expiration of the number of DAYS specified</p> |

6.2.2 Wildcard certificates

If a signed wildcard certificate and key were generated with a passphrase, the passphrase must be removed from the private key to allow Plixer Scrutinizer to use the pair.

1. Copy the private key file (*.key) to /etc/pki/tls/private/.
2. Re-generate the key without a passphrase (replace ORIGINAL with the filename of the key):

```
openssl rsa -in /etc/pki/tls/private/ORIGINAL.key -out /etc/pki/tls/  
private/new.key
```

3. When prompted, enter the passphrase used for the original key.

This will create a new, unencrypted key named `new.key` in `/etc/pki/tls/private/`, which must be renamed to `ca.key`. If the key pair was originally created without a passphrase, it need only be renamed after being copied into the correct directory.

6.2.3 Full chain certificates

A full chain certificate or chain of trust can be created as follows:

1. Create the file `ca_chain.crt` under `/etc/pki/tls/CA/`.
2. Copy the contents of the intermediate CA .crt file into `ca_chain.crt`.
3. Copy the contents of the root CA .crt file into `ca_chain.crt` (after the intermediate CA).
4. Add the following line to `/home/plixer/scrutinizer/files/conf/httpd-plixer.conf`:

```
SSLCertificateChainFile /etc/pki/tls/CA/ca_chain.crt
```

5. Restart the httpd service:

```
sudo /bin/systemctl restart httpd
```

After the restart, the full chain certificate will be in use.

6.2.4 CA-signed distributed cluster certificates

To generate CSRs and install the signed keys for a distributed cluster, run the following scripts:

Note:

- These scripts should be run from the distributed cluster’s primary reporter as the `plexer` user and rely on Plexier Scrutinizer’s default SSH connectivity.
- `scrut_util --rotatecerts --reset` (see above) can be used if either of these scripts causes unexpected issues or DB connection errors. However, *any existing signed certificates will be lost.*

| | |
|--|---|
| <pre>/home/scrutinizer/files/generate_ ↪ requests.sh</pre> | <p>Generates certificate signing requests (CSRs) for all TLS keys in a distributed cluster</p> <p><i>CSRs are saved to subdirectories in <code>/tmp/request</code> with <code>apache_server.csr</code> being the signing request for the primary reporter’s web server.</i></p> |
| <pre>/home/scrutinizer/files/install_ ↪ certs.sh</pre> | <p>Installs signed TLS certificates to all nodes in a distributed cluster</p> <p><i>.cer files should be saved to <code>/tmp/signed</code> following the path and filename conventions used by <code>generate_requests.sh</code> for the signing requests. The Certificate Authority’s root certificate should be saved as <code>ca.cer</code>.</i></p> |

6.3 Integrations

Plexier 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.

6.3.1 Plixer Replicator

When Plixer Replicator is deployed as part of a *distributed environment*, Plixer Scrutinizer can leverage the **Auto-Replicate** feature to automatically assign exporters to collectors based on their capacities and current workloads. With Auto-Replicate enabled, all exporters can be configured to send flows to Plixer Replicator, and replication profiles are automatically created for all collectors in the distributed cluster.

Note: See [this section](#) of the FAQ or the Plixer Replicator online documentation to learn more about Plixer Replicator's functions. For information on licensing options, contact [Plixer Technical Support](#).

Setting up auto-replication on Plixer Scrutinizer is divided into three main processes:

- [Creating the seed profile](#)
- [Registering Plixer Replicator and enabling Auto-Replicate](#)
- [Editing the auto-replication configuration file](#)

Creating the seed profile

The *seed profile* on the Plixer Replicator instance functions as a list of exporters that Plixer Scrutinizer should add to the auto-replication profiles it maintains.

The seed profile can be created as follows:

1. Deploy the Plixer Replicator instance and complete the basic configuration steps as described here.
2. Create a new profile to use as the seed profile (or use the profile that was created during the basic configuration process).
3. Add a policy to include all exporters to the profile.
4. [Optional] If there are exporters whose flows should not be auto-replicated, add policies to exclude those exporters to the profile.

Once the seed profile has been set up, it will need to be registered with Plixer Scrutinizer to complete the auto-replication configuration.

Registering the seed profile

After the seed profile is created on Plixer Replicator, log in to Plixer Scrutinizer and register the profile as follows:

1. In the Plixer Scrutinizer web interface, navigate to **Admin > Plixer > Plixer Replicator**.
2. Fill in the form with the following details:
 - `replicator` user password
 - Listening/in port (for receiving flows) on the Plixer Replicator instance
 - IP address or hostname of Plixer Replicator instance (must start with `https://`)
 - Name of the seed profile created for auto-replication
 - Sending/out port (for replicated flows) on the Plixer Replicator instance
3. Tick the *Enable* checkbox, and then click the **Save** button to save the seed profile configuration.
4. After the seed profile has been registered, SSH to the distributed cluster's primary reporter as the `plixer` user.
5. Run the following command **once** to create the auto-replication configuration file:

```
scrut_util --autoreplicate
```

Running `scrut_util --autoreplicate` for the first time will create `/home/plixer/scrutinizer/files/autoreplicate.conf`, where individual collector properties and additional auto-replication settings can be configured.

Note:

- Because Plixer Scrutinizer manages exporter-to-collector assignments for auto-replication, the seed profile should **not** have a collector defined.
 - If additional devices are deployed to the environment, `scrut_util --autoreplicate` should be run again to update Plixer Scrutinizer's list of auto-replication exporters.
 - Instead of using a policy to include all exporters in an environment, devices can also be added individually to the seed profile. In such cases, the seed profile should be updated before running `scrut_util --autoreplicate` again, when additional exporters are deployed.
-

Editing the configuration file

After the `autoreplicate.conf` file has been created, it should be reviewed and edited to configure the parameters Plexier Scrutinizer will use to manage auto-replication.

Afterwards, `scrut_util --autoreplicate` should be run a second time to start the auto-replication service.

Note:

- Any time the configuration file is modified in any way, `scrut_util --autoreplicate` should be run again to update Plexier Scrutinizer's auto-replication configuration.
- The `--verbose` flag can also be added to verify that the correct auto-replication settings have been applied in Plexier Scrutinizer.
- Once auto-replication is active, Plexier Scrutinizer will only reassign an exporter if its current collector exceeds the configured exporter or flow rate capacity. Setting up a *cronjob* to run `scrut_util --autoreplicate` at regular intervals can help maintain a balanced load across available collectors.

For example:

```
##*/3 * * * * sudo -i -u plixer /home/plixer/scrutinizer/bin/scrut_util --  
↪autoreplicate --verbose 2>&1 /home/plixer/scrutinizer/files/logs/  
↪autoreplicate.log
```

This will run `scrut_util --autoreplicate` daily at 3 am, with logs saved to `/home/plixer/scrutinizer/files/logs/autoreplicate.log`

Collector limits

The configuration file will contain an entry defining the exporter and flow rate capacities for each remote collector in the cluster in the following format:

```
"COLLECTOR_ADDRESS" : {  
  "exporters" : MAX_EXPORTER_COUNT,  
  "flow_rate" : MAX_FLOW_RATE  
}
```

These values are used by Plexier Scrutinizer to determine how exporters are assigned and can be edited as needed. Once auto-replication has been started, an exporter will only be reassigned if its current collector exceeds its configured exporter or flow rate capacity.

If additional collectors are deployed to the cluster at a later time, they should also be added to the configuration file.

Additional auto-replication settings

The following global auto-replication settings can also be edited in the configuration file:

- `replicator_host`: Full address of the Plexier Replicator instance (must be formatted as `https://REPLICATOR_IP`)
- `replicator_pass`: Password set for the `replicator` user account
- `replicator_seed_profile`: Auto-replication seed profile name
- `replicator_receive_port`: Listening/in port
- `replicator_send_port`: Sending/out port

Because the values for these settings in the configuration file take precedence over those entered via the Plexier Scrutinizer web interface, the configuration file can be used to update auto-replication parameters at any time.

Password encryption

For added security, the plain text password in the configuration file can be replaced with its AES256-encrypted value.

To generate a password's encrypted value, run:

```
scrut_util --autoreplicate --encrypt YOUR_PASSWORD
```

The output can then be assigned to `replicator_pass` in `autoreplicate.conf`.

Using multiple configurations

In certain scenarios, it may be necessary to use multiple auto-replicate configurations in the same environment.

These scenarios include:

- Deploying multiple Plixer Replicator instances in the same Plixer Scrutinizer cluster
- Setting up different listening/in and sending/out combinations for specific exporters
- Creating collector groups for auto-replication

Sample multi-config setup

The following instructions outline the steps for setting up the directories and files for two `autoreplicate.conf` files:

1. Create and configure a [seed profile](#) for each auto-replicate configuration on Plixer Replicator (e.g., `Seed_Profile_1` and `Seed_Profile_2`).
2. SSH to the primary Plixer Scrutinizer server/reporter as the `plixer` user, and then create a `replicator_pools` directory:

```
cd /home/plixer/scrutinizer/  
mkdir replicator_pools
```

3. Under `replicator_pools`, create a subdirectory for each auto-replication configuration:

```
mkdir replicator_pools/config_1  
mkdir replicator_pools/config_2
```

4. Copy `/home/plixer/scrutinizer/bin/scrut_util` to each configuration subdirectory, and then rename it to `replicator_util`:

```
cp /home/plixer/scrutinizer/bin/scrut_util /home/plixer/scrutinizer/  
↪replicator_pools/config_1/replicator_util  
cp /home/plixer/scrutinizer/bin/scrut_util /home/plixer/scrutinizer/  
↪replicator_pools/config_2/replicator_util
```

5. Create the `autoreplicate.conf` file in each directory:

```
touch replicator_pools/config_1/autoreplicate.conf
touch replicator_pools/config_2/autoreplicate.conf
```

6. Use the following template to configure auto-replication in each `autoreplicate.conf`:

```
"collector_capacities" : {
  "COLLECTOR_1_ADDRESS" : {
    "exporters" : MAX_EXPORTER_COUNT,
    "flow_rate" : MAX_FLOW_RATE
  },
  "COLLECTOR_2_ADDRESS" : {
    "exporters" : MAX_EXPORTER_COUNT,
    "flow_rate" : MAX_FLOW_RATE
  }
},
"replicator_enable" : 1,
"exporter_timeout" : 1800,
"per_port_profile_name" : 1,
"replicator_host" : REPLICATOR_ADDRESS,
"replicator_pass" : REPLICATOR_PASSWORD,
"replicator_receive_port" : LISTENING_PORT,
"replicator_send_port" : SENDING_PORT,
"replicator_seed_profile" : SEED_PROFILE_NAME
```

Each file must be edited to reference the corresponding seed profile (e.g., `Seed_Profile_1` or `Seed_Profile_2`) and the collectors to which exporters in that profile can be assigned. [Additional parameters](#) should also be defined for the configuration, if necessary.

7. Run `replicator_util --autoreplicate` for each configuration file.

```
replicator_pools/config_1/replicator_util --autoreplicate --verbose
replicator_pools/config_2/replicator_util --autoreplicate --verbose
```

8. [Optional] Create a cronjob for each auto-replication configuration to reset exporter-to-collector assignments at regular intervals:

```
*/3 * * * * sudo -i -u plixer /home/plixer/scrutinizer/files/replicator_
↪pools/config_1/replicator_util --autoreplicate --verbose 2>&1 /home/
↪plixer/scrutinizer/files/logs/autoreplicate.log
```

(continues on next page)

(continued from previous page)

```
#!/3 * * * * sudo -i -u plixer /home/plixer/scrutinizer/files/replicator-  
→pools/config_2/replicator_util --autoreplicate --verbose 2>&1 /home/  
→plixer/scrutinizer/files/logs/autoreplicate.log
```

Auto-replication will start after autoreplicate is run for each configuration file.

Note:

- To further organize auto-replication configurations, the directory structure can be expanded as needed, as long as the directories containing the autoreplicate.conf files are created under a directory called replicator_pools (e.g., /home/plixer/scrutinizer/autoreplicate/data_center_A/collector_group_X/replicator_pools/configuration_1).
 - Any time a seed profile or configuration file is modified, replicator_util --autoreplicate must be run in the corresponding directory to update Plixer Scrutinizer's auto-replication configurations.
 - When setting up collector groups for auto-replication, using a unique listening/in port for exporters for each group is highly recommended.
-

6.3.2 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**.

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](#).

6.3.3 Plexier FlowPro

Plixer Scrutinizer is able to use Plexier FlowPro as a data source, enabling seamless transfer and analysis of network traffic data.

Configuration guide

After setting up a Plexier FlowPro, configure integration in Plexier Scrutinizer as follows:

1. Navigate to **Admin > Plexier > FlowPro Licensing**.
2. Paste the Plexier FlowPro license key.
3. Click **Save**.

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

When configured and integrated, Plexier Scrutinizer is able to monitor and provide insights to network performance in real time, diagnose issues, and quickly identify the root causes.

The following Plexier 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. |

Packet capture rules

Plixer FlowPro can enable targeted traffic sampling in Plixer Scrutinizer through customizable selective packet capture rules. Capture rules can be defined from [Admin > Resources > FlowPro Capture Rules](#) in the web interface or [via API request](#).

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](#).

6.3.4 Flow log ingestion

Plixer Scrutinizer can be configured to ingest AWS, Azure, and/or OCI virtual private network flow logs for traffic monitoring and analysis, as well as to enable additional functionality specific to each platform.

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 |

This section covers the prerequisites and setup/configuration steps for AWS VPC flow log ingestion.

Setting up the AWS S3 storage bucket

Before setting up AWS VPC flow log ingestion in Plixer 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 Plixer 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 Plixer Scrutinizer.
- Versioning should be disabled.

Note:

- When upgrading from older versions of Plixer 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 Plixer 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, Plixer 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 Plixer Scrutinizer to become current more quickly.
-

Configuring AWS VPC flow log ingestion in Plixer Scrutinizer

To add an S3 bucket 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 + icon and select *AWS VPC FlowLogs* in the tray.

3. In the secondary tray, configure the bucket details as follows:

- Enter a name to identify the bucket/source by.
- Select the Plixer Scrutinizer servers to use as log downloader(s) and collector(s) for the bucket (the primary reporter of a distributed cluster is not recommended for either role).
- 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; must have full access to the bucket)

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:

- After a bucket configuration has been saved, click on the name assigned to it in the main view to open the settings tray, and use the **Test** button to confirm that Plixer Scrutinizer is able to establish a connection to the bucket with the credentials entered.
 - 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 > 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, and a single bucket can have multiple downloaders and collectors.
-

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 flow logs

Once flow log ingestion for Azure-based resources has been enabled, Plixer Scrutinizer can monitor and run reports on traffic traversing assets in the cloud.

Once flow data for network resources on Azure 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 configured rules |
| Flow Decisions Count | Flow count aggregation for each 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 |
| All Details | Aggregation based on full range of flow details, including the rule and application associated with the traffic |
| Resource IDs | Aggregation based on resource IDs |

This section covers the prerequisites and setup/configuration steps for Azure flow log ingestion.

Note: Plixer Scrutinizer supports both VNet and NSG flow logs on Azure.

Setting up the Azure blob storage container

Before setting up Azure flow log ingestion in Plixer Scrutinizer, the Azure Storage blob container(s) that will be used should be configured as follows:

- Set the virtual networks 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 as a flow log source, Plixer Scrutinizer will regularly 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 Azure flow log ingestion in Plixer Scrutinizer

To add an Azure Storage blob container as a flow log source in Plixer Scrutinizer, follow these steps:

1. Navigate to **Admin > Integrations > Flow Log Ingestion** in the web interface.
2. Click the + icon, and then select *Azure FlowLogs* in the tray.
3. In the secondary tray, configure the container details as follows:
 - Enter a name to identify the bucket/source by.
 - Enter the container name:
 - For NSG flow logs, this will typically follow the format of `insights-logs-networksecuritygroupflowevent`
 - For VNet flow logs, this will typically follow the format of `insights-logs-flowlogflowevent`
 - Select the collector(s) to assign to the container from the dropdown (the primary reporter of a distributed cluster is not recommended).
 - Enter the storage account name and key to use to access the container (in most cases, the service URL host name without `.blob.core.windows.net/` or another domain)
 - 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 Azure virtual network will also be added to the device lists for Plixer Scrutinizer's various functions ([Flow Analytics](#), [network maps](#), [reports](#), etc.).

Note:

- After a container configuration has been saved, click on the name assigned to it in the main view to open the settings tray, and use the **Test** button to confirm that Plixer Scrutinizer is able to establish a connection to the container with the credentials entered.
 - To verify that the Azure flow log source has been successfully added, look for an exporter whose hostname matches the virtual network in the **Explore > Exporters > By Exporters** view or the **Admin > Resources > Exporters** page (after ~1 hour).
-

For assistance with any issues, consult the [troubleshooting guide](#) or contact [Plixer Technical Support](#).

Google Cloud Platform VPC flow logs

With GCP Virtual Private Cloud (VPC) flow log ingestion enabled, Plixer Scrutinizer is able to monitor and report on traffic data associated with GCP VPC assets.

This section covers the prerequisites and setup/configuration steps for GCP VPC flow log ingestion.

Setting up the Google Cloud Pub/Sub topic and subscription

Plixer Scrutinizer uses the GCP Pub/Sub messaging service as an ingestion source for VPC flow logs.

To set up the Pub/Sub topic that will receive the log entries to be ingested, follow these steps:

Note: To ensure seamless access between components/services, it is highly recommended to set everything up under the project where flow logs will originate.

1. Enable and configure VPC Flow Logs for the target resources.
2. Next, navigate to the **Pub/Sub Topics** page and create a new topic with message retention enabled and set to at least one hour (other topic settings can be configured as desired).
3. After the topic has been created, go to the **Subscriptions** page and create a pull subscription for the new topic (note the Subscription ID for later use).
4. Next, go to the **Log Router** page and create a sink to route the log entries to the newly created topic and configure any inclusion/exclusion filters necessary.
5. After adding the Pub/Sub topic as a sink, navigate to the **Service Accounts** page and select a service account associated with the sink/topic.

6. Under the *Keys* tab, click the **Add Key** button and select *JSON* to download a file containing the credentials required to subscribe to the Pub/Sub topic.

Once the above steps have been completed, verify that log entries are being correctly routed to the Pub/Sub topic, and then proceed to configuring ingestion in Plixer Scrutinizer.

Configuring GCP VPC flow log ingestion in Plixer Scrutinizer

Once the Pub/Sub topic is receiving log entries and the subscription has been set up, it can be added to Plixer Scrutinizer as follows:

1. Navigate to **Admin > Integrations > Flow Log Ingestion** in the web interface.
2. Click the **+** icon, and then select *Google Cloud Platform* in the tray.
3. In the secondary tray, configure the subscription details as follows:
 - Enter a name to identify the source by.
 - Select the Plixer Scrutinizer servers to use as the log downloader(s) and collector(s) (the primary reporter of a distributed cluster is not recommended for either role).
 - Enter the GCP project ID associated with the topic subscription.
 - Enter the subscription name/ID used.
 - Enter/paste the contents of the service account key JSON file.
4. Click the **Save** button to add the subscription with the current settings.

Note:

- After a subscription configuration has been saved, click on the name assigned to it in the main view to open the settings tray, and use the **Test** button to confirm that Plixer Scrutinizer is able to establish a connection with the credentials entered.
 - To verify that an GCP VPC flow log source has been successfully added, look for an exporter whose hostname matches the GCP 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 through the topic subscription) and the *flow collector* (collects and processes the downloaded logs). A different Plixer Scrutinizer server can be used for each role, and a single subscription can have multiple downloaders and collectors.
-

Oracle Cloud Infrastructure VCN flow logs

With OCI Virtual Cloud Network (VCN) flow log ingestion enabled, Plixer Scrutinizer is able to monitor and report on traffic associated with specified Oracle Virtual Network Interface Cards (VNICs).

This section covers the prerequisites and setup/configuration steps for OCI VCN flow log ingestion.

Setting up the OCI flow log stream

VCN flow log ingestion in Plixer Scrutinizer uses the OCI streaming service as the log data source. After being downloaded from a stream, the log data is forwarded to one or more specified collectors as regular flows.

To set up the flow log stream, follow these steps:

1. Create a new stream in any stream pool to publish the flow logs to.
2. Enable flow logs for the VCN, subnet, or VNICs.
3. Configure a new service connector as follows:
 - Source: Compartment, log group, and name associated with the logs enabled in step 2.
 - Target: Compartment and name associated with the stream created in step 1.
4. Create/provision an IAM group with the use `stream-pull` permission and add a user to the group (or select an existing user).
5. Generate an API signing key pair for the user and download the private key as described here.
6. Get the private key fingerprint using this command.

Verify that the flow logs are correctly being published to the stream, and then proceed to configuring Plixer Scrutinizer to download/ingest the log data.

Note: If the key pair was not generated via the OCI console, the public key will need to be uploaded for the user.

Configuring OCI VCN flow log ingestion in Plixer Scrutinizer

Once the OCI stream has been successfully configured, it can be added to Plixer Scrutinizer as a flow log source as follows:

1. Navigate to **Admin > Integrations > Flow Log Ingestion** in the web interface.
2. Click the **+** icon, and then select *Oracle Cloud Streams* in the tray.
3. Enter the following details in the secondary tray:
 - Enter a name to identify the stream/source by.
 - Select the Plixer Scrutinizer servers to use as log downloader(s) and collector(s) for the stream (the primary reporter of a distributed cluster is not recommended for either role).
 - Enter the URL for the stream pool containing the flow log stream.
 - Enter the OCID of the stream receiving the VCN flow logs.
 - Enter the OCID of the OCI tenancy.
 - Enter the OCID of the user to be used to access the streams (must have the required permissions).
 - Enter the fingerprint of the private API signing key generated for the user.
 - Enter the passphrase associated with the private key (leave blank if no passphrase was used when the key was generated)
 - Enter the private key in PEM format.
 - Enter the name of the home region of the tenancy.
4. Click the **Save** button to add the stream with the current settings.

Once added, the stream will be listed in the main [Admin > Integrations > Flow Log Ingestion](#) view under the configured name. An exporter associated with VCN will also be added to the device lists for Plixer Scrutinizer's various functions ([Flow Analytics](#), [network maps](#), [reports](#), etc.).

Note:

- After a stream configuration has been saved, click on the name assigned to it in the main view to open the settings tray, and use the **Test** button to confirm that Plixer Scrutinizer is able to establish a connection to the stream with the credentials entered.
- To verify that an OCI VCN flow log source has been successfully added, look for an exporter whose hostname matches the VCN 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 stream) and the *flow collector* (collects and processes the downloaded logs). A different Plixer Scrutinizer server can be used for each role, and a single stream can have multiple downloaders and collectors.

Troubleshooting

MFSNs and a buildup of log files in flow log source containers 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 (starting with CPU cores) 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.

Note:

- In distributed deployments, it is recommended to start with a 1:1 pairing of sources and collectors.
 - The *Unresourced - Enabled* status in the [Admin > Resources > Exporters](#) view is another indication that flow log sources are being temporarily disabled/paused due to insufficient resources.
-

If the **Admin > Resources > Exporters** view does not list exporters that are associated with the virtual network(s) set up for flow ingestion, do the following:

1. Navigate to **Admin > Integrations > Flow Ingestion**, open the configuration tray for the collector it was assigned to, and then use the **Test** button to verify that the correct details were entered.

Note: The **Test** button only checks if the communication with the data source works.

2. Verify that flow logs are correctly being sent to the bucket or container.
3. Check the collector log file in `/home/plixer/scrutinizer/files/logs/` for errors.
4. Check `awss3_log.json` (AWS), `azure_log.json` (Azure), or `ocist_log.json` for possible source-side issues.

Note: The **Admin > Resources > Exporters** view also displays exporters that have been disabled. Because each AWS, Azure, or OCI flow log source counts as an exporter, one or more sources may be disabled automatically (in last-in/first-out order) if the exporter count limit of the current license is reached.

6.3.5 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 EndaceProbe-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 become 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 then click **Save**.
3. Click the **Create Client (+)** button, and then enter the IP address of the Plixer Scrutinizer collector.
4. [OPTIONAL] Enter a password.
5. Locate the Plixer Scrutinizer client in the list, and then click **Download** 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 `/etc/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:

- The Plixer Scrutinizer eStreamer client configuration will automatically be updated whenever `firesight.ini` is modified.
 - Editing the provided `firesight.ini` file is recommended, but a new file can also be created in the same directory. A sample file (`firesight.ini.sample`) can be found in `/home/plixer/scrutinizer/files`.
 - 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 of the updated FireSIGHT eStreamer client certificate.
7. `pkcs12_password` is the certificate password, or blank if a password wasn't specified.
8. `fs_bind_addr` is the eStreamer client address registered with FireSIGHT (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.
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](#).

Grafana

Integrating the Grafana plugin with Plexier Scrutinizer allows users to monitor systems, applications, and infrastructure through dashboards, charts, and graphs.

To setup the Grafana integration, do the following:

1. Deploy the default Grafana server.
2. Start the server with the default (Production) settings.
3. Adjust the `default.ini` file to run the server in development mode.

Note: If you are deploying a non-default server, edit the `custom.ibj` file instead.

4. Start the server.

For assistance in getting or setting up the Grafana plugin, contact [Plexier Technical Support](#).

PRTG

When PRTG integration is enabled, users can view PRTG-based device information when inspecting exporters in the Plexier Scrutinizer web interface.

To set up PRTG integration in Plexier Scrutinizer, navigate to **Admin > Integrations > 3rd Party Integration** 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 **Save**.

Once configured, the option to view PRTG details becomes 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 available 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 **Save**.

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 Plixer Scrutinizer collector log for errors.
- Verify the credentials you entered in Plixer Scrutinizer are correct.
- Use the **Test** button on the **Viptela Settings** page to confirm that the Plixer 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 Plexier Scrutinizer must be assigned the `sn_incident_write` role.

3. Verify that the details entered are correct, and then click **Save**.

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 > Integrations > 3rd Party Integration**, and then do the following:

1. Select **SolarWinds** from the dropdown, and then 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 **Save**.

Once configured, the option to view SolarWinds details will be available from the **Integrations** menu when inspecting exporters.

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 set up 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 then 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. Click **Select Nodes** to assign the property to at least one existing node, and then click 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. Click **Submit**.

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 Plixer Scrutinizer in the widget.

Splunk

Splunk integration enables the inspection of Plixer Scrutinizer flow and event data in the Splunk dashboard via the **Scrutinizer for Splunk** app.

This allows teams already using Splunk to seamlessly leverage Plixer Scrutinizer's flow collection and analysis capabilities and quickly jump between the two platforms as needed.

Note: Splunk integration requires Plixer Scrutinizer 19.6.0 or higher. The **Scrutinizer for Splunk** app expects both the *Splunk Enterprise* server and *Splunk Forwarder* client software to be pre-installed in the customer environment.

Configuring Splunk integration in Plixer Scrutinizer

To set up Plixer Scrutinizer for Splunk integration, follow these steps:

1. SSH to the Plixer Scrutinizer server as the `plixer` user.
2. Launch the interactive CLI:

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

3. At the `SCRUTINIZER>` prompt, run the following:

```
SCRUTINIZER> enable splunk http://<SPLUNK_SERVER_IP:PORT> <SYSLOG_PORT>  
→<SPLUNK_FORWARDER_IP>
```

The default Splunk server port is 8000 (if port 80 is used, no port number is required after the server IP address). The default listening port (`SYSLOG_PORT`) on the Splunk Forwarder is 1514.

After the command is run, Plixer Scrutinizer will begin sending flow and event data once the next flow analytics collection and detection cycle is complete.

Installing the Scrutinizer for Splunk app

After configuring Plixer Scrutinizer to send data to Splunk, the **Scrutinizer for Splunk** app can be installed as follows:

1. Download the **Scrutinizer for Splunk** app:

```
REPO_HOST=files.plixer.com  
curl -k -o scrutinizer.spl https://$REPO_HOST/plixer-repo/scrutinizer/19.  
→6.1/util/scrutinizer.spl
```

If an *offline repo host* was used to install or upgrade Plixer Scrutinizer, `REPO_HOST` can be set to the IP address of that host.

2. Log into Splunk.
3. Go to **Apps > Manage Apps** in the Splunk dashboard.
4. Click the **Install app from file** button, and then select the `scrutinizer.spl` file downloaded in step 1.
5. After the app is installed, locate the **Scrutinizer for Splunk** app in the **Manage Apps** menu and click *View Objects*.

6. Select *Default*, and then replace the default IP address (10.42.100.142) with the address of the Plixer Scrutinizer server to connect to Splunk.
7. Click the **Save** button to save the new address.

When done, return to the dashboard and access the **Scrutinizer for Splunk** app from the **Apps** menu. Data and graphs should begin to be filled in after a few minutes.

Note:

- If no data appears in the Splunk UI after 5-10 minutes, restart the Splunk service on the Splunk Server by running:

```
sudo /opt/splunk/bin/splunk restart
```

Data should start to appear on the *Scrutinizer Vitals* page in the Splunk UI.

- To upgrade the **Scrutinizer for Splunk** app instead, tick the *Upgrade app* checkbox when selecting the `scrutinizer.spl` in the install dialog.

Visit <https://www.plixer.com> to learn more or contact [Plixer Technical Support](#) for further assistance.

Disabling Splunk integration

To disable Splunk integration in Plixer Scrutinizer:

1. SSH to the Plixer Scrutinizer server as the `plixer` user, and then launch the interactive CLI:
2. At the `SCRUTINIZER>` prompt, run the following:

```
SCRUTINIZER> disable splunk http://<SPLUNK_SERVER_IP:PORT>
```

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 in the Plexier Scrutinizer web interface, follow these steps:

1. Navigate to **Admin > Integrations > STIX-TAXII**, and then click **Add** 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 **Save**.
 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, Plexier 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 Plexier AD Users utility to retrieve username data and forward it to Plexier Scrutinizer as IPFIX flows.

The Plexier AD Users utility reads a Windows event log file, continually parses authentication events, and sends event data to an IPFIX collector (Plexier 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 Plexier Scrutinizer. This feature helps avoid duplicate events being sent to the collector or a lapse in the authentication events processed should the program restart.

Plexier AD Users 2.0.0

Plexier AD Users 2.0.0 allows integration with the Plexier ML Engine 19.5.0. 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 Plexier ML Engine will remove duplicates of both AD and Azure events for model generation.

Configuring the servers

User permissions

By default, the Plexier 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 Plexier 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

To configure your Azure account info with the Plixer ML Engine, do the following:

1. Navigate to **Admin > Settings > ML AD Users** in the Plixer Scrutinizer web UI to configure the Azure storage account name and Azure storage account key.
2. Edit the `ad-users.yml` config file and use the newly generated ML auth token for the `ml.scrut_auth_token` value found under **Admin > Resources > ML Engines**.
3. Restart `ad-users` through Windows Server hosting `ad-users`.

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 | Required? | Description | Example Value | Default Value |
|-----------------|-----------|--|---|-----------------------------------|
| chunking | 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_time | 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 |
| path | Yes | Path to the Windows event log. Use <i>ForwardedEvents.evtx</i> if forwarding events, or <i>Security.evtx</i> if running directly on AD server. | C:\Windows\System32\WindowsSystem32\eventlog\ForwardedEvents.evtx | |
| collector | Yes | This is the Plixer Scrutinizer collector to send the flows to. The format must be IP:port. | 127.0.0.1:2055 | 127.0.0.1:2055 |
| exporter | No | The IP address of the Windows server running AD Users. To specify your own value, use the format IP:port. | 8.8.8.8:9996 | (Local IP address with port 9996) |

log section

| Name | Re-quired? | Description | Example Value | Default Value |
|--------------|------------|---|---------------------------------|--|
| name | no | Name of the log file if not running as a service. | C:\ad-users\Plixer\ad-users.log | (ad-users.log in executable directory) |
| level | no | Log level to use for file logging if not running as a service. | info | debug |
| max_size | no | Maximum size (in MB) of the log file before it gets rotated, if not running as a service. | 100 | 5 |
| max_backups | no | Maximum number of old log files to retain, if not running as a service. | 10 | 0 |
| max_age_days | no | Maximum number of days to retain old log files, if not running as a service. | 14 | 0 |

ml section

| Name | Required? | Description | Example Value | Default Value |
|------------------|--|---|---------------|---------------|
| send_ad_to_ml | No | Flag for sending AD data to the Plixer ML Engine, in addition to Plixer Scrutinizer. | True | False |
| scrut_auth_token | No (only required if send_ad_to_ml is set to True) | 24-character authentication token generated from Plixer Scrutinizer. | • | • |
| ml_engine_ip | No (only required if send_ad_to_ml is set to True) | IP address of the Plixer ML Engine to send to the Active Directory and/or Azure data via Kafka. | 0.0.0.0 | 0.0.0.0 |

azure section

| Name | Required? | Description | Example Value | Default Value |
|---------------------------|---|---|---|---|
| send_azure_to_ml | No | Flag for sending Azure data to the Plixer ML Engine. | True | False |
| interactive_container | No (only required if send_azure_to_ml is set to True) | Name of the configured interactive sign-in log storage container in Azure. | insights-logs-signinlogs | insights-logs-signinlogs |
| non_interactive_container | No (only required if send_azure_to_ml is set to True) | Name of the configured non-interactive sign-in log storage container in Azure. | insights-logs-noninteractive-signinlogs | insights-logs-noninteractive-signinlogs |
| max_log_age_minutes | No | Logs terminated within the last [number] minutes set will be processed. | 15 | 0 |
| account_filter_list | No | Used to monitor sign-in events of admin Azure accounts (emails), and then send these events to the ML Engine. Separate multiple email addresses with a comma. | [user1@plixer.com, user2@plixer.com, user3@plixer.com] | |
| app_filter_list | No | Used to monitor sign-in events of applications, and then send these events to the ML Engine. Separate multiple applications with a comma. | [Microsoft Teams, Microsoft Office, Office 365 Exchange Online] | [Microsoft Teams, Microsoft Office, Office 365 Exchange Online] |

Starting the service

1. Open **Services**, and then right-click on **Plixer AD Users**.
2. Select **Properties**, and then in the **General** tab, set the startup type to **Automatic (Delayed Start)**.
3. Go to the **Recovery** tab, and then set all three failure options to **Restart the Service**.
4. Click **OK** to save.

Verifying the setup

Checking log files

If running Plixer AD Users as a service, the application log in Event Viewer will show the program's log messages. At startup, there will be a few info messages indicating everything was configured properly and the program has started event monitoring. After that, there will only be error log messages if any errors occur or if the service is stopped. If the service restarts, the startup info messages will be logged again.

If running Plixer AD Users in command prompt, use command-line argument `run`. Log messages will be written to the log file as well as the console (stdout).

```
C:\ad-users> ad-users.exe run
Detected 'run' program argument, not running as a service
Config: Trying to get Azure credentials via ML Engine API at x.x.x.x
Config: GET https://x.x.x.x:30888/azure_auth response code=200
{"level":"info","time":"2024-04-08T12:18:16-07:00","message":
  ↳ "Successfully set config values: chunking=10; path=C:\\Windows\\
  ↳ System32\\winevt\\Logs\\ForwardedEvents.evtx; exporter=;
  ↳ collector=x.x.x.x:2055; flush wait seconds=60; send AD to ML=true;
  ↳ send Azure to ML=true"}
{"level":"info","buffer_size":67108864,"time":"2024-04-08T12:18:16-
  ↳ 07:00","message":"SetWriteBuffer"}
{"level":"info","time":"2024-04-08T12:18:16-07:00","message":
  ↳ "Successfully set collector: x.x.x.x:2055 and exporter: x.x.x.
  ↳ x:9996 endpoints"}
{"level":"info","time":"2024-04-08T12:18:16-07:00","message":
  ↳ "Successfully set ML Engine endpoint x.x.x.x and Kafka producer
  ↳ for topic x-topic"}
{"level":"info","time":"2024-04-08T12:18:16-07:00","message":
  ↳ "Successfully set up access to Azure containers x-interactive and
  ↳ y-noninteractive"}
{"level":"info","time":"2024-04-08T12:18:17-07:00","message":
  ↳ "Successfully opened Windows events file: C:\\Windows\\System32\\
  ↳ winevt\\Logs\\ForwardedEvents.evtx"}
{"level":"info","time":"2024-04-08T12:18:17-07:00","message":"Last
  ↳ sent record ID was saved as 12345"}
{"level":"info","time":"2024-04-08T12:18:17-07:00","message":
  ↳ "Starting event monitoring"}
```

Checking Plixer Scrutinizer for IPFIX flows

In the Plixer Scrutinizer UI **Explore > Entities > Usernames** table, AD Users authentication events will start populating. Plixer AD Users sends the IP address (no IPv6 support currently), logon type (logon or logoff), domain, username, and machine name of the authentication event.

If usernames aren't showing up as expected, double-check that you have enough exporters licensed for the number of exporters enabled. You can see how many exporters you have licensed in the Plixer Scrutinizer UI under **Admin > Plixer > Scrutinizer Licensing > Exporter Count** and **Enabled Exporters**. You can also view specific exporters under **Admin > Resources > Manage Exporters**.

The Plixer AD Users machine will count as an exporter since it is sending flows with username data to Plixer Scrutinizer.

Export spreading

The config values for `chunking` and `flush_wait_seconds` should mitigate any issues from too many events being exported at a time: `chunking` allows for a given number of events to be queued in the buffer then sent all at once, and `flush_wait_seconds` will flush the buffer periodically to avoid events sitting in the queue for too long when fewer authentication events are logged in a minute than the set chunking value.

However, if working with a Plixer Scrutinizer setup where too many Active Directory authentication events at a time is a concern, you can prevent *Netflow export storms* by enabling *export spreading*. Follow the instructions here for your performance monitor configuration.

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.x, 2.x, and 3.x.

Enabling ERS

Before setting up Cisco ISE username reporting in Plexier 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 Plexier Scrutinizer

1. SSH into the Plexier Scrutinizer server as the `plexier` user and run `/home/plexier/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](#).

6.4 Interactive CLI

The Plixer Scrutinizer interactive CLI utility provides access to system-level functions, such as admin operations, configuration/maintenance routines, and integration management.

The interactive prompt (SCRUTINIZER>) is accessed by establishing an SSH session with the Plixer Scrutinizer server and running:

```
scrut_util
```

Note: Most **scrut_util** commands can also be executed using direct shell syntax, which allows them to be used in scripts to automate maintenance tasks. Run the following from the shell to view the equivalent syntax for the top-level interactive commands listed below:

```
scrut_util --help [COMMAND]
```

Select a command below to view details about its function and usage.

6.4.1 awssync

The **awssync** command can be used to sync IDs and descriptions from AWS when AWS flow log ingestion is enabled.

Syntax

```
awssync
```

6.4.2 check

The `check` commands can be used to run a check/test against the resource, setting, or function specified by the option used.

Options and syntax

Note: The collector should be stopped before running any of the `history_index` commands.

| | |
|--|--|
| <code>check activeif</code> | Checks for active flows based on interface details and returns the last timestamp and number of interfaces that received flows |
| <code>check collectorclass <CLASS_↵ ↵[SUBSYSTEM]></code> | Returns running state details for the specified collector CLASS or, if provided, the specified SUBSYSTEM of that class <i>This command is used by Plixer Technical Support for troubleshooting.</i> |
| <code>check data_last_written</code> | Returns activity details for collected flow data written to the database |
| <code>check dist_info</code> | Returns distributed cluster configuration details for the Plixer Scrutinizer server |
| <code>check hdtest <TRIES></code> | Tests hard drive performance by running a write-delete operation either 10 times (default) or, if provided, the number of times specified by the TRIES parameter and returns details for the amount of time taken |
| <code>check heartbeat <database api></code> | Test and returns information on internal communications with the specified resource type |
| <code>check history_index</code> | Checks the history index and returns historical activity information for the 1m interval aggregation table |
| <code>check history_index_empty_tables</code> | Checks the history index and returns a list of tables with zero rows (collector should be stopped first) <i>To delete empty tables, use the delete command instead.</i> |
| <code>check history_index_orphans</code> | Checks the history index and returns a list of entries for which a table does not actually exist |
| 420 | 6. Advanced Services <i>To delete orphan entries, use the delete command instead.</i> |

6.4.3 ciscoise

The `ciscoise` commands are used to manage Cisco Identity Services Engine (ISE) node integration in Plixer Scrutinizer.

Options and syntax

| | |
|---|---|
| <pre>ciscoise add <IP_ADDRESS> <TCP_ ↪PORT> <ISE_USER></pre> | <p>Adds a Cisco ISE node with the specified IP_ADDRESS, TCP_PORT, and ISE_USER (must have API access) to queue to acquire user identities for all active sessions</p> <p><i>The ISE_USER password will also need to be entered after this command is run.</i></p> |
| <pre>ciscoise check</pre> | <p>Tests node polling and returns the results</p> <p><i>This command can be used to verify that Plixer Scrutinizer is able to collect user identity information.</i></p> |
| <pre>ciscoise kick <ISE_ID> <IP_ADDRESS> ↪ [MAC_ADDRESS]</pre> | <p>Kicks the ISE_ID off the ISE node at the specified IP_ADDRESS and optional MAC_ADDRESS, forcing re-authentication</p> |
| <pre>ciscoise nodelist</pre> | <p>Returns a list of all Cisco ISE nodes currently configured</p> |
| <pre>ciscoise poll</pre> | <p>Forces a poll of all Cisco ISE nodes and returns the results</p> |
| <pre>ciscoise remove <IP_ADDRESS></pre> | <p>Removes the Cisco ISE node with the specified IP_ADDRESS from Plixer Scrutinizer</p> |
| <pre>ciscoise update <IP_ADDRESS> <TCP_ ↪PORT> <ISE_USER></pre> | <p>Updates the current configuration of the Cisco ISE node with the specified IP_ADDRESS to use the provided TCP_PORT and ISE_USER</p> <p><i>The ISE_USER password will also need to be entered after this command is run.</i></p> |

6.4.4 clean

The `clean` commands are used to manually execute housekeeping processes that are automatically run at regular intervals.

Options and syntax

| | |
|-----------------------------|---|
| <code>clean all</code> | Immediately executes all scheduled housekeeping tasks |
| <code>clean baseline</code> | Resets all configured baselines to the default values <i>Historical data will not be deleted but will still expire following the configured data retention settings.</i> |
| <code>clean database</code> | Purges all temporary database entries |
| <code>clean ifinfo</code> | Purges <i>ifinfo</i> entries that do not have matching entries in <i>activeif</i> |
| <code>clean old_logs</code> | Purges old log files that are set to the <i>backup</i> status |
| <code>clean tmp</code> | Purges all temporary files created by the graphing engine |

6.4.5 collect

The `collect` commands are used to manually execute collection processes for data that can be used in various Plexer Scrutinizer functions. Many of these processes are run automatically at regular intervals.

Options and syntax

| | |
|---|--|
| <code>collect asa_acl</code> | Immediately polls Cisco ASA devices to collect ASA ACL information |
| <code>collect baseline</code> | Collects baseline data and checks for alarms/events |
| <code>collect dbsize</code> | Collects database size information |
| <code>collect elk <IP_ADDRESS></code> | Collects data from Plixer Scrutinizer and forwards it to the ELK server using the IP_ADDRESS specified |
| <code>collect optionsummary</code> | Initiates processing of flow option data collected by Plixer Scrutinizer |
| <code>collect snmp</code> | Immediately polls SNMP devices to collect data used by Plixer Scrutinizer |
| <code>collect splunk <IP_ADDRESS> <PORT></code> | Collects data from Plixer Scrutinizer and forwards it to the Splunk server using the IP_ADDRESS and PORT specified |
| <code>collect supportfiles</code> | Collects various logs and configuration data that can be used by <i>Plixer Technical Support</i> for troubleshooting |
| <code>collect topology</code> | Collects device data to help Plixer Scrutinizer understand the network's topological layout |
| <code>collect useridentity</code> | Initiates processing of user identity data collected by Plixer Scrutinizer |

6.4.6 convert

The convert command is used to convert different types of data and information.

Syntax and options

| | |
|-------------------------|--|
| <div>converttoaes</div> | Converts all encrypted information stored by Plexier Scrutinizer to use AES 256 encryption |
|-------------------------|--|

6.4.7 delete

The delete commands are used to delete database entries or tables from the Plexier Scrutinizer system.

Options and syntax

Note:

- These commands will permanently delete data and should be used with caution.
- The collector should be stopped before running any of the `history_index` commands.

| | |
|---|---|
| <div>delete custom_algorithm <FILENAME></div> | <p>Permanently deletes the custom algorithm with the specified FILENAME</p> <p><i>FILENAME should not include the algorithm file's .pm extension.</i></p> |
| <div>delete history_index_empty_tables</div> | Deletes all tables with zero rows from the history index |
| <div>delete history_index_orphans</div> | Deletes all history index entries for which a table does not actually exist |
| <div>delete history_table_orphans</div> | Deletes all tables that do not have a history index entry |

6.4.8 disable

The `disable` commands are used to disable specific functions/features in Plixer Scrutinizer.

Options and syntax

Note: These commands can alter Plixer Scrutinizer functionality and should be used with caution.

| | |
|--|---|
| <code>disable baseline <IP_ADDRESS></code> | <p>Disables all baselines for the exporter with the specified IP_ADDRESS</p> <p><i>Historical data associated with the exporter will not be deleted but will still expire following the configured data retention settings.</i></p> |
| <code>disable elk http://<IP:PORT></code> | Disables ELK flows from Plexier Scrutinizer to the URL specified by IP:PORT |
| <code>disable ipv6</code> | Disables IPv6 for all interfaces in <code>sysctl.conf</code> |
| <code>disable splunk http://<IP:PORT></code> | Disables Splunk flows from Plexier Scrutinizer to the URL specified by IP:PORT |
| <code>disable ssh_root_login</code> | <p>Prohibits the superuser root account from logging into a Linux shell directly from outside hosts</p> <p><i>Instead of allowing remote root SSH login, it is recommended to instead log in as the plexier user and use sudo for maintenance tasks. This command will not affect root logins from a physical or virtual console.</i></p> |
| <code>disable unresponsive</code> | Disables ping of exporters that have been flagged as unresponsive |
| <code>disable user <USERNAME></code> | Disables the specified USERNAME account with <code>scrut_util</code> access (e.g., for server maintenance) |

6.4.9 enable

The `enable` commands are used to enable/configure specific functions in Plixer Scrutinizer.

Options and syntax

Note: These commands can alter Plixer Scrutinizer functionality and should be used with caution.

| | |
|--|---|
| <pre>enable baseline <IP_ADDRESS> default</pre> | <p>Enables default baselines for the exporter with the specified IP_ADDRESS</p> |
| <pre>enable baseline <IP_ADDRESS> manual <PRIMARY[, SECONDARY]> ELEMENT avg count min max std sum dailyhr busday sameday></pre> | <p>Enables a custom baseline with the following parameters for the exporter with the specified IP_ADDRESS:</p> <ul style="list-style-type: none">• PRIMARY - IPFIX element to be included in the baseline (e.g., sourceIPv4Address, applicationName, etc.)• SECONDARY - Optional secondary IPFIX element to be included in the baseline• ELEMENT - Corresponding numeric IPFIX element for the primary and secondary elements to be used to determine the baseline (e.g., packetDeltaCount, octetDeltaCount, etc.)• AVE COUNT MIN MAX STD SUM - Selects between average (AVE), flow count (COUNT), minimum value (MIN), maximum value (MAX), standard deviation (STD), or sum (SUM) for measuring the specified ELEMENT• dailyhr busday sameday - Selects between daily (dailyhr), daily on business days (busday), or same day weekly (sameday) for baseline comparison <p><i>When baselining IP addresses, IP groups should be defined for the address ranges and subnets to be included in the baseline. This will prevent addresses that may only talk once from triggering false positives.</i></p> |
| <pre>enable custom_algorithm <FILENAME> <NAME></pre> | <p>Enables the custom algorithm FILENAME in the flow analytics engine under the specified NAME</p> <p><i>FILENAME should not include the .pm extension of the algorithm file (must be saved to scrutinizer/files/algorithms/).</i></p> |
| 6.4. Interactive CLI | 429 |
| <pre>enable elk http://<IP:PORT></pre> | <p>Enables ELK flows from Plexier Scrutinizer to the URL specified by IP:PORT</p> |

6.4.10 endace

The endace commands are used to manage EndaceProbe integration.

Options and syntax

| | |
|--|---|
| <div><pre>endace add <IP_ADDRESS> <PORT> ↪<USER> <PASSWORD></pre></div> | <p>Enable integration with an EndaceProbe with the specified IP_ADDRESS, PORT, and Endace USER:PASSWORD</p> <p><i>The default port used by an EndaceProbe is ``443``.</i></p> |
| <div><pre>endace remove <IP_ADDRESS></pre></div> | <p>Remove the EndaceProbe with the specified IP_ADDRESS</p> |
| <div><pre>endace update <IP_ADDRESS> <PORT> ↪<USER> <PASSWORD></pre></div> | <p>Update EndaceProbe integration settings with the specified IP_ADDRESS, PORT, and Endace USER:PASSWORD</p> |

6.4.11 expire

The expire commands are used to delete expired historical data following the configured [data retention settings](#).

Options and syntax

Note: These commands will permanently delete data and should be used with caution.

| | |
|------------------------------------|---|
| <code>expire dnscache [all]</code> | Purges expired DNS cache data (based on the <i>Days of DNS Request Data</i> setting) or, if the <code>all</code> option is included, all DNS cache data |
| <code>expire history [trim]</code> | Purges expired flow data (based on <i>Flow Historical X Avg</i> settings) and also deletes older data until the <i>Minimum Percent Free Disk Before Trimming</i> is reached if the <code>trim</code> option is included |
| <code>expire inactiveflows</code> | Removes expired inactive interfaces (based on the <i>Inactive Expiration system preference setting</i>) from interface views |
| <code>expire templates</code> | Purges flow template metadata for templates that haven't been observed for 30 days |

6.4.12 export

The export commands are used to dump data from Plixer Scrutinizer for external use.

Options and syntax

| | |
|--|--|
| <pre>export langtemplate <LANG_NAME></pre> | <p>If the language exists, creates a CSV file with the English and LANG_NAME keys and saves it as home/plixer/scrutinizer/files/pop_languages_LANGNAME_template.csv</p> |
| <pre>export peaks_csv <FILENAME_ ↵ INTERVAL PATH RANGE [GROUP_ID]></pre> | <p>Exports a CSV list of interfaces and peak values under the specified PATH (must be relative to the home/plixer/scrutinizer/ directory) and FILENAME following the criteria passed:</p> <ul style="list-style-type: none"> • INTERVAL - Roll-up interval table to use (in raw minutes; must be 1, 5, 30, 120, or 720) • RANGE - Time and date range to cover: <ul style="list-style-type: none"> • Last24Hours • Last15Minutes • Last5Minutes • LastFortyfiveMinutes • LastFullHour • LastHour • LastMonth • LastSevenDays • LastTenMinutes • LastThirtyDays • LastThirtyMinutes • LastThreeDays • LastTwentyMinutes • LastWeek • LastYear • ThisMonth • ThisWeek • ThisYear • Today • Yesterday • GROUP_ID - If used, limits the list to interfaces included in the corresponding IP group <p><i>For a list of IP group IDs, use the show groups command.</i></p> |

6.4.13 import

The `import` commands are used to import various types of data (labels, definitions, groupings, etc.) for use in Plixer Scrutinizer's functions.

For further information, see [this guide on importing data](#).

6.4.14 moloch

The `moloch` command is used to enable or disable integration for the Moloch probe using the specified `IP_ADDRESS` and `PORT`.

Syntax

```
moloch <on|off> <IP_ADDRESS [PORT]>
```

6.4.15 optimize

The `optimize` commands are used to manually execute optimization processes that are automatically run at regular intervals.

Options and syntax

Note: These commands will modify database tables in Plixer Scrutinizer and should be used with caution.

| | |
|---|---|
| <pre>optimize common</pre> | Optimizes tables that are commonly inserted and deleted to improve database performance |
| <pre>optimize database <DATABASE></pre> | Optimizes only tables in the specified DATABASE |

6.4.16 remove

The `remove address ipv6` command is used to delete the current IPv6 address assigned to the server.

Syntax

Note:

- The IPv6 address can only be removed if there is an IPv4 address assigned. To edit IP address settings, use the *set myaddress* command.
- This command will alter Plixer Scrutinizer functionality and should be used with caution.

```
remove address ipv6
```

6.4.17 repair

The `repair` commands are used to run various repair processes related to Plixer Scrutinizer functions and databases.

Options and syntax

Note: These commands will modify database tables in Plixer Scrutinizer and should be used with caution.

| | |
|---|--|
| <code>repair business_hour_saved_reports</code> | Converts saved reports with business hours that were created in older Plexier Scrutinizer versions (15.5 and below) to the latest format with the same business hours |
| <code>repair history_tables</code> | Repairs history tables that have the wrong <i>col</i> type for <code>octetDeltaCount</code> <i>This command is not used for PostgreSQL installations.</i> |
| <code>repair policy_priority_order</code> | Repairs irregularities in alarm policy IDs (e.g., duplication) |
| <code>repair range_starts</code> | Repairs history tables without the start time used to identify the range of data they contain <i>This repair process may take some time to complete and should only be executed under the direction of Plexier Technical Support.</i> |

6.4.18 rotate

The `rotate` commands are used to replace the keys and certificates used by Plexier Scrutinizer in its functions.

Options and syntax

Note:

- These commands will alter Plexier Scrutinizer functionality and should be used with caution.
- `rotatecerts` can only be run using direct shell/script syntax and not from the `SCRUTINIZER>` prompt (as shown below).

| | |
|--|---|
| <pre>rotatekeys</pre> | Creates a new encryption key and re-encrypts all encrypted fields in the database |
| <pre>scrut_util --rotatecerts [--days ↪<DAYS>] [--reset] [--verbose]</pre> | <p>Regenerates all certificates on all nodes (including any Plexier ML Engine deployments) with an optional expiration date in the specified number of DAYS</p> <p><i>If the --reset flag is included, the CA certificate on the primary reporter and the Apache web server certificate will also be regenerated.</i></p> |

6.4.19 services

The services command is used to stop, start, or restart all services.

Syntax

```
services all <stop|start|restart>
```

Note: To stop, start, or restart specific services, run the following from the shell instead:

```
sudo systemctl <stop|start|restart> <SERVICE>
```

6.4.20 set

The set commands are used to manage settings/behaviors related to authentication, networking, and general operation for the Plexier Scrutinizer server.

Options and syntax

Note: These commands can alter Plixer Scrutinizer functionality and should be used with caution.

| | |
|---|---|
| <pre>set columnmoniker <OLD_NAME> <NEW_NAME> [ELEMENT_LIST]</pre> | <p>Replaces an information element's OLD_NAME with the specified NEW_NAME</p> <p><i>If the optional ELEMENT_LIST of one or more elements (comma-delimited) is included, renaming will be limited to flow templates that also include those elements.</i></p> <p><i>This command should only be run under the direction of Plixer Technical Support.</i></p> |
| <pre>set dns</pre> | <p>Allows use the user to enter one or more new DNS servers for hostname resolution</p> <p><i>The operation will overwrite the system's previous DNS server list.</i></p> |
| <pre>set hostinfo <IP_ADDRESS> <FQHN></pre> | <p>Assigns the specified FQHN (fully qualified host-name) to the current Plixer Scrutinizer appliance and configures resolution for the provided IP_ADDRESS</p> |
| <pre>set leds_threshold</pre> | <p>Resets the LED warning threshold to 10% of the total storage available on the appliance's data partition</p> <p><i>When combined with the Auto History Trimming settings, this function can help prevent Plixer Scrutinizer from using up all available storage.</i></p> |
| <pre>set myaddress <IPv4_ADDRESS> <NETMASK> <GATEWAY> set myaddress <IPv6_ADDRESS/CIDR> <GATEWAY></pre> | <p>Assigns the specified IPv4/IPv6_ADDRESS, CIDR/NETMASK, and GATEWAY to the current appliance</p> <p><i>After the provided IP information has been confirmed to be correct, the previous address of the same type will be overwritten.</i></p> <p><i>Because an SSH session will automatically be terminated after the new IP address is assigned,</i></p> |
| 438 | <p><i>it is recommended to run this command from a console connection.</i></p> |

6.4.21 show

The `show` commands are used to view various details, settings, and other functional elements for the Plexier Scrutinizer server/environment.

Options and syntax

| | |
|---------------------------------------|--|
| <code>show custom_algorithms</code> | Displays a list of all custom algorithms saved to scrutinizer/files/algorithms/ and their current state |
| <code>show datasize</code> | Displays a breakdown of database storage sizes by schema |
| <code>show diskspace</code> | Displays storage allocation and utilization details |
| <code>show dns</code> | Displays a list of all DNS servers used for host-name resolution |
| <code>show exporters [FILTER]</code> | Displays a list of exporters sending data to collectors (using the specified FILTER if included) |
| <code>show groups</code> | Displays a list of all current device/mapping groups |
| <code>show interfaces [FILTER]</code> | Displays a list of interfaces sending data to collectors (using the specified FILTER if included) |
| <code>show ipaddresses</code> | Displays all IP addresses assigned to the current Plexier Scrutinizer appliance |
| <code>show metering [FILTER]</code> | Displays a list of interfaces by exporter and their metering direction (using the specified device IP address FILTER if included) |
| <code>show partitions</code> | Displays partition information for the current Plexier Scrutinizer appliance |
| <code>show task [FILTER]</code> | Displays a list of all tasks currently configured in Plexier Scrutinizer (using the specified task name FILTER if included) |
| <code>show timezone</code> | Displays the timezone configured for the current Plexier Scrutinizer appliance |
| <code>show tzlist [FILTER]</code> | Displays a list of timezones that can be configured for the Plexier Scrutinizer appliance (via the set timezone command) |
| <code>show unknowncolumns</code> | Displays a list of exporter information elements that are unrecognized by Plexier Scrutinizer |

6.4.22 snoop

The snoop commands are used to listen for traffic at the interface level.

Options and syntax

| | |
|--|--|
| <code>snoop interface <INTERFACE> <PORT></code> | Listens for traffic on the specified INTERFACE and PORT |
| <code>snoop ipaddress <IP_ADDRESS> <PORT></code> | Listens for traffic on the specified IP_ADDRESS and PORT |

6.4.23 system

The system command is used to reboot or shut down the system.

Syntax

```
system <restart|shutdown>
```

6.4.24 unlock

The unlock command is used to unlock a locked USER account (due to *failed login attempts*).

If no authentication method is specified (ldap, radius, or tacacs) the account defaults to local authentication.

Syntax

```
unlock <USER> [ldap|radius|tacacs]
```

6.4.25 upload

The upload supportfiles command is used to upload the log and configuration data package (after running the *collect supportfiles* command) for use by *Plixer Technical Support*.

Syntax

```
upload supportfiles
```

6.4.26 version

The version command is used to show version information for Plexier Scrutinizer.

Syntax

```
version
```

6.5 Plexier Scrutinizer APIs

Plexier Scrutinizer supports API access for the following function sets:

6.5.1 Capture rule configuration

Selective packet capture (requires *Plexier FlowPro*) rules can be added via API, which requires the following fields:

- authToken - Admin authentication token *generated by Plexier Scrutinizer* (required for API access)
- rm - flowpro_capture_rules (runmode corresponding to the function set being accessed)
- name - Name to assign to the new capture rule
- server_ip - Packet source/server IP address or CIDR
- client_ip - Packet destination/client IP address or CIDR
- max_packets - Maximum number of packets to capture
- stops_on - End date/time for capturing packets as UNIX epoch timestamp
- well_known_port - Well-known port to monitor for packets
- retention_hours - Duration to store captured packet data
- enabled - State to add the rule in (1: enabled; 0: disabled)
- action - add (adds/creates a new capture rule as defined in the request)

Request example

Below is an example of an API call to create a new packet capture rule.

```
curl --location 'https://<SCRUTINIZER_ADDRESS>/fcgi/scrut_fcgi.fcgi' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=flowpro_capture_rules' \
--form 'name="LDAP Traffic 2"' \
--form 'server_ip="1.1.1.1/32"' \
--form 'client_ip="2.2.2.2/32"' \
--form 'max_packets="1000"' \
--form 'stops_on="1743048000"' \
--form 'well_known_port="393605"' \
--form 'retention_hours="168"' \
--form 'enabled="1"' \
--form 'action=add'
```

6.5.2 IP group management

The following fields are required for all IP group management API requests:

- **authToken** - Admin authentication token *generated by Plixer Scrutinizer* (required for API access)
- **rm** - ipgroups (runmode corresponding to the function set being accessed)
- **action** - One or more of the following *actions* to be initiated by the request:
 - *saveRule* - Creates an IP group with the specified rule
 - *update* - Modifies an existing IP group
 - *loadTreeRootFast* - Loads a condensed list of all IP group names and IDs
 - *search* - Searches for an IP group by name
 - *loadRules* - Loads a list of all rule definitions for an IP group
 - *deleteRule* - Removes a rule from an IP group
 - *delete* - Deletes an IP group

- deleteAll - Deletes all IP group definitions from Plixer Scrutinizer

Rule definitions

Use the following JSON object formats to pass IP group inclusion rule definitions in requests:

| Rule type | JSON |
|-----------------------|---|
| Single IP address(es) | <pre>[{ "type": "ip" "address": "<IP_ADDRESS_1>" }, { "type": "ip" "address": "<IP_ADDRESS_2>" }]</pre> |
| IP address range | <pre>[{ "type": "range" "sip": "<START_IP>" "eip": "<END_IP>" }]</pre> |
| Subnet | <pre>[{ "type": "network" "address": "<ADDRESS>" "mask": "<SUBNET_MASK>" }]</pre> |
| Wildcard mask | <pre>[{ "type": "wildcard" "address": "<ADDRESS>" "mask": "<WILDCARD_MASK>" }]</pre> |
| Child group | <pre>[{ "type": "child" "child_id": "<CHILD_IPGROUP_ID>" }]</pre> |

Request examples

Below are additional details and request examples for actions that can be included in an IP group management API call.

saveRule

The following additional fields can be passed in the request when creating new IP groups using the saveRule action:

- new_fc - Specifies a name for the new IP group
- added - Specifies a JSON array of one or more *inclusion rule definitions* to add to the new IP group

API request example

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/  
↳fcgi/scrut_fcgi.fcgi' \  
--form 'authToken=<AUTH_TOKEN>' \  
--form 'rm=ipgroups' \  
--form 'action=saveRule' \  
--form 'new_fc=UK Data Center' \  
--form 'added=[  
    {  
        "type": "ip",  
        "address": "10.30.10.1"  
    }  
]'
```

Returned JSON object

```
{
  "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 Data Center",
  "rule_id": 506588,
  "total": 1
}
```

update

The following additional fields can be passed in requests to add, replace, or remove a rule definition from an IP group using the `update` action:

- `name` - Replaces the current name of the IP group if included
- `added` - Specifies a JSON array of one or more *inclusion rule definitions* to add to the new IP group
- `updated` - Specifies a JSON array of rules (based on the included `rule_id` field) that will be overwritten with the new definitions provided
- `removed` - Specifies a JSON array of rule IDs to be deleted

API request example

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/
↳fcgi/scrut_fcgi.fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=ipgroups' \
--form 'action=update' \
--form 'fc_id=169000006' \
--form 'name=Renamed Group' \
--form 'added=[
    {
        "type": "ip",
        "address": "10.1.4.66"
    }
]' \
--form 'updated=[
    {
        "rule_id": "84",
        "type": "ip",
        "address": "192.1.0.0"
    }
]' \
--form 'removed=[114]'
```

search

The following additional fields can be passed in requests to search IP groups for a partial string or full name using the `search` action:

- `name` - IP group name or string to search for
- `fc_name_comp` - Specifies the comparison operator to search with (`like` or `notlike`)
- `page` - Specifies the number of pages of results to load (default: one page)
- `maxRows` - Specifies the maximum number of results per page in the response

API request example

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/  
→fcgi/scrut_fcgi.fcgi' \  
--header 'Content-Type: application/json' \  
--form 'authToken=<AUTH_TOKEN>' \  
--form 'rm=ipgroups' \  
--form 'action=search' \  
--form 'name=UK' \  
--form 'fc_name_comp=like' \  
--form 'page=1' \  
--form 'maxRows=10'
```

deleteRule

The deleteRule action deletes a rule definition with the specified rule_id from its IP group.

API request example

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/  
→fcgi/scrut_fcgi.fcgi' \  
--header 'Content-Type: application/json' \  
--form 'authToken=<AUTH_TOKEN>' \  
--form 'rm=ipgroups' \  
--form 'action=deleteRule' \  
--form 'rule_id=506588'
```

Returned JSON object

```
{  
  "fc_id": 16900006,  
  "success": 1,  
  "myrules": "",  
  "rule_id": 506588,  
  "total": 0  
}
```

delete

The delete action deletes the IP group(s) specified by the group id array passed in the request's json field.

API request example

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/
↳fcgi/scrut_fcgi.fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=ipgroups' \
--form 'action=delete' \
--form 'json=[
  {
    "id": "16900032",
  }
]'
```

Returned JSON object

```
{
  "processedCount": 1,
  "removed": [
    "16900032"
  ]
}
```

6.5.3 Reporting

The following fields are required for all IP group management API requests:

- authToken - Admin authentication token *generated by Plixer Scrutinizer* (required for API access)
- rm - report_api (runmode corresponding to the function set being accessed)

- `action` - `get` (runs the report defined in the request)
- `rpt_json` - JSON object defining the *parameters* of the report to be run
- `data_requested` - Specifies the *elements of the report* to be included in the response

Report parameters

Each report API request must specify the parameters for the report using the following elements of the `rpt_json` object:

| Object Element/Field | Report Parameter | Available Options | Example |
|----------------------|--|---|---|
| reportTypeLang | <i>Report type</i> | conversations: Conversations WKP host2host: Host to host ipGroupGroup: IP group to IP group applications: Applications defined country2country: Country to country | <pre>"reportTypeLang" ↳ ":" ↳ "conversations" ↳ "</pre> |
| filters | <i>Exporter and interface filter</i> | in_<EXPORTER_IP_HEX>_ALL: Includes all interfaces on the specified exporter in_<EXPORTER_IP_HEX>_<HEX_INTERFACE_N>: Includes interface index N | <pre>"filters": { ↳ "sdfDips_0" ↳ ": "in_" ↳ "0A190101_ALL" }</pre> |
| reportDirections | Traffic directionality (relative to interfaces included) | inbound or outbound | <pre>↳ "reportDirections" ↳ ": { ↳ "selected" ↳ ": "inbound" }</pre> |
| times | <i>Report time range/window</i> and time zone to display dates in (use <code>scrut_util --show tzlist</code> for a list of valid timezones) | LastFiveMinutes LastTenMinutes LastFifteenMinutes LastTwentyMinutes LastThirtyMinutes LastFortyFiveMinutes LastHour LastFullHour LastThreeDays LastSevenDays | <pre>"times": { ↳ "dateRange" ↳ ":" ↳ "LastFiveMinutes" ↳ "clientTimezone" ↳ ": "America/" ↳ "New York" }</pre> |
| 452 | | | 6. Advanced Services |

Response data

The `data_requested` field specifies how to format the graph and table of the report output.

JSON object example:

```
{
  "inbound": {
    "graph": "none",
    "table": {
      "query_limit": {
        "offset": 0,
        "max_num_rows": 10
      }
    }
  }
}
```

Note: The directionality specified in the `data_requested` object must match the `reportDirections` field.

Request example

The following API call runs a default report against all interfaces of the specified device for the last 5 minutes:

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/
↳fcgi/scrut_fcgi.fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=report_api' \
--form 'action=get' \
--form 'rpt_json={
  "reportTypeLang": "conversations",
  "filters": {
    "sdfDips_0": "in_0A190101_ALL"
  },
}
```

(continues on next page)

(continued from previous page)

```
"reportDirections": {
  "selected": "inbound"
},
"times": {
  "dateRange": "LastFiveMinutes"
  "clientTimezone": "America/New_York"
},
"dataMode": {
  "selected": "saf"
},
"rateTotal": {
  "selected": "total"
},
"dataGranularity": {
  "selected": "auto"
},
"bbp": {
  "selected": "bits"
},
"reportDirections": {
  {
    "type": "ip",
    "address": "10.30.10.1"
  }
}
}' \
--form 'data_requested'=
{
  "inbound": {
    "graph": "none",
    "table": {
      "query_limit": {
        "offset": 0,
        "max_num_rows": 10
      }
    }
  }
}
}'
```

Returned JSON object

The following condensed response shows the typical structure of the object returned for a report API request:

```
{
  "report": {
    "request_id": "0xed184820e4b611eab58f1fc02130f7f9",
    "table": {
      "inbound": {
        "totalRowCount": 1,
        "footer": [],
        "columns": [],
        "rows": []
      }
    },
    "time_details": {},
    "exporter_details": {},
    "graph": {}
  }
}
```

Field details:

| | | |
|--|---------------|--|
| table (will include separate data for inbound and outbound if applicable) | columns | elementName: Name of the data element in the column format: Formatting details for data in the column label: Table header label |
| | rows | rawValue: Unformatted value (as returned from the database) label: Formatted value including bits, bytes, or percent |
| | footer | [0]: Represents the Others data for a calculated column, which is the sum of the data in all rows not included in the table [1]: Represents Total for a calculated column, which is the sum of the data in all included rows plus the Others value for the same column |
| | totalRowCount | Integer specifying the total number of rows available |
| graph | all | Includes data for all graph types |
| | pie | Values for graphing table data as a pie chart |
| | timeseries | Values for graphing table data as a line graph |
| | none | Includes only default graph (pie) data |

6.5.4 User account management

The following fields are required for all user account management API requests:

- **authToken** - Admin authentication token *generated by Plexier Scrutinizer* (required for API access)
- **rm - user_api** (runmode corresponding to the function set being accessed)
- **action** - One or more of the following *actions* to be initiated by the request:
 - *createUser* - Creates one or more new Plexier Scrutinizer *user accounts* with the option to assign each to *user groups*
 - *delUsers* - Deletes one or more user accounts
 - *createUsergroup* - Creates one or more user groups with the option to add users to each group
 - *delUsergroups* - Deletes one or more user groups
 - *membership* - Adds and/or removes users to or from specified user groups
 - *prefs* - Edits preferences for a single user
 - *permissions* - Edits permissions for one or more user groups
 - *changeUsername* - Renames an existing user account

Request examples

Below are additional details and request examples for actions that can be included in an user account management API call.

createUser

Creating user accounts using the **createUser** action requires an additional **json** field containing an array (**users**) of the following:

- **name** - Username for the account
- **pass** - Password for the account
- **membership** - Array of one or more user group IDs to assign the user account to

API request example

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/
↳fcgi/scrut_fcgi.fcgi' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=user_api' \
--form 'action=createUser' \
--form 'json=
{
  "users": [
    {
      "name": "NewAdmin",
      "pass": "NewAdminPassword",
      "membership": [1]
    },
    {
      "name": "NewGuest",
      "pass": "NewGuestPassword",
      "membership": [2]
    }
  ]
}'
```

Note: User group IDs are stored in the `plexer.usergroups` table. By default, 1 is the administrators group and 2 is the guest users group.

Returned JSON object

```
{
  "data": [
    {
      "id": 3,
      "name": "NewAdmin"
    },
    {
      "id": 4,
```

(continues on next page)

(continued from previous page)

```

        "name": "NewGuest"
    }
]
}

```

delUser

Deleting user accounts using the `delUser` action requires an additional `json` field containing an array (`delUsers`) of the usernames and/or user IDs of the accounts to be deleted:

API request example

```

curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/
→fcgi/scrut_fcgi.fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=user_api' \
--form 'action=delUser' \
--form 'json=
{
    "delUsers": [
        11,
        "NewGuest",
        207
    ]
}'

```

Returned JSON object

```

{
    "data": [
        "Deleting user id 11 (1 matched)",
        "Deleting user named 'NewGuest' (1 matched)",
        "Deleting user id 207 (0 matched)"
    ]
}

```

createUsergroup

Creating user groups using the `createUsergroup` action requires an additional `json` field containing an array (`usergroups`) of the following:

- `name` - User group name
- `template_usergroup` - Existing user group ID of the existing group to use as the template for the new user group
- `users` - Array of usernames or user IDs to be added to the group (if an empty array is passed, an empty user group will be created)

API request example

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/
↳fcgi/scrut_fcgi.fcgi' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=user_api' \
--form 'action=createUsergroup' \
--form 'json=
{
  "usergroups": [
    {
      "name": "My Group",
      "template_usergroup": 1,
      "users": [1,"AnotherUser"]
    },
    {
      "name": "Other Group",
      "template_usergroup": 2,
      "users": ["MyUser",2]
    }
  ]
}'
```

Returned JSON object

```
{
  "data": [
    {
      "id": 5,
      "name": "My Group",
      "members": [1, "AnotherUser"]
    },
    {
      "name": "Other Group",
      "error": "A usergroup already exists with that name"
    }
  ]
}
```

delUsergroups

Deleting user groups using the `delUsergroups` action requires an additional `json` field containing an array (`delUsergroups`) of the names and/or IDs of the user groups to be deleted.

API request example

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/
→fcgi/scrut_fcgi.fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=user_api' \
--form 'action=delUsergroups' \
--form 'json=
{
  "delUsergroups": [
    3,
    "My User Group"
  ]
}
```

Returned JSON object

```
{
  "data": [
    "Deleting usergroup id 3 (1 matched)",
    "Deleting usergroup named 'My User Group' (0 matched)",
  ]
}
```

membership

Editing user group membership using the `membership` action requires an additional `json` field containing `add` and/or `remove` arrays to specify the usernames/user IDs and user groups to add/remove them to/from.

API request example

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/
↳fcgi/scrut_fcgi.fcgi' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=user_api' \
--form 'action=membership' \
--form 'json=
{
  "membership":
  {
    "add": [
      {
        "user_id": 13,
        "usergroup_id": 5
      },
      {
        "user_name": "NewUser",
        "usergroup_name": "Other Group"
      }
    ],
    "remove": [
      {
```

(continues on next page)

(continued from previous page)

```

        "user_name": "USER3",
        "usergroup_id": 4
      }
    ]
  }
}'

```

Returned JSON object

```

{
  "data":
    "added": [
      "User 13 added to usergroup 5",
      "User 14 added to usergroup 3"
    ],
    "removed": [
      "User 15 removed from usergroup 4"
    ]
}

```

prefs

The `prefs` action modifies one or more user preferences for a single user account and requires an additional `json` field containing an array (`prefs`) of all preference changes.

API request example

```

curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/
→fcgi/scrut_fcgi.fcgi' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=user_api' \
--form 'action=prefs' \
--form 'json=

```

(continues on next page)

(continued from previous page)

```
{
  "user_id": 11,
  "prefs": [
    {
      "pref": "statsTopn",
      "setting": 10
    },
    {
      "pref": "language",
      "setting": "english"
    }
  ]
}
```

Returned JSON object

```
{
  "data": {
    "updated": [
      "statusTopn updated to 10 for user_id 11",
      "language updated to english for user_id 11"
    ],
    "errors": []
  }
}
```

permissions

The `permissions` action updates permissions for one or more user groups and requires an additional `json` field containing all user groups names/IDs and permission changes as `add` and `remove` arrays.

The following table lists all available `permission_type` and `seccode` (for use with the “plexer” `permission_type`) options in the request:

| | | |
|-----------------|---------------------|--|
| permission_type | device | IP address of a device in hex (e.g. '0A010107') |
| | interface | IP address of a device in hex and the interface index separated by |
| | group | Group ID of a mapping/device group from <code>plexer.groups</code> |
| | report | <code>saved_id</code> of a saved report from <code>reporting.saved_reports</code> |
| | gadget | <code>gadget_id</code> of a dashboard gadget from <code>plexer.dash_gadgets</code> |
| | thirdparty | ID of a third-party link from <code>plexer.third_party</code> |
| | plexer | Permission code corresponding to different functions/sections with |
| seccode | 3rdPartyIntegration | Create, edit, and delete third-party integration links |
| | ackBBEvent | Acknowledge alarms |
| | adminTab | Access the Admin tab/section |
| | alarmSettings | Configure alarm notifications |
| | alarmsTab | Access the Alarm Monitor tab/section |
| | allDevices | Access the status of all devices and their interfaces |
| | allGadgets | Access all gadgets created by any user |
| | allGroups | Access all mapping/device groups |
| | allInterfaces | Report on interfaces for any device |
| | allLogalotReports | All Logalot reports |
| | allReportFolders | Access all saved report folders |
| | allReports | Access all saved reports created by any user |
| | allThirdparty | Access all configured third-party links |
| | almDelete | Permanently delete alarms |
| | ApplicationGroups | Configure application groups |
| | asnames | Configure AS names |
| | auditing | Access auditing reports containing logs of Plexier Scrutinizer user |
| | auth | Manage external authentication tokens |
| | Authentication | Manage external authentication types |
| | authLdapServers | Manage LDAP server configuration for Plexier Scrutinizer authen |
| | awsSettings | AWS configuration |
| | changeUserPasswords | Change passwords for other users without needing their credential |
| | createDashTabs | Create new dashboards |
| | createUsers | Create new local Plexier Scrutinizer user accounts |
| | dashboardAdmin | Manage all dashboards created by any user |
| | DataHistory | Configure data history/retention settings |
| | deleteReport | Delete saved reports regardless of owner |
| | deleteUsers | Delete local Plexier Scrutinizer user accounts |
| | DeviceDetails | Edit device interface details |
| | EmailNotifications | Configure the mailserver for Plexier Scrutinizer reports and email |
| | faExclusions | Configure flow analytics exclusions |
| | fa_mgmt_link | Configure flow analytics thresholds and settings |
| | feedbackForm | Access the link to send feedback to Plexier |

Table 1 – continued from previous page

| | |
|----------------------------|---|
| FlowAnalyticsSettings | Access global flow analytics settings |
| helpTab | Access the Help tab/section |
| HostNames | Edit hostname information |
| IPGroups | Configure Plixer Scrutinizer IP groups |
| language | Create and edit language localization settings |
| licensing | Configure Plixer Scrutinizer product licensing and features |
| LogalotPrefs | Configure global alarm settings |
| MACAddresses | Configure device MAC address information |
| ManageCollectors | Manage devices collecting flow data for Plixer Scrutinizer |
| ManageExporters | Manage devices exporting flow data to Plixer Scrutinizer |
| mappingGroupConfiguration | Create and edit mapping/device groups |
| mappingObjectConfiguration | Create and edit mapping objects |
| mapsTab | Access the Network Maps page |
| myViewTab | Access the Dashboards page |
| NotificationManager | Manage alarm notifications |
| PolicyManager | Manage alarm policies |
| protocolExclusions | Edit protocol exclusions for flow reports |
| proxySettings | Configure proxy server settings in Plixer Scrutinizer |
| radiusConf | Manage RADIUS server configuration for Plixer Scrutinizer authentication |
| ReportDesigner | Design new custom report types |
| reportFilters | Update the filters used in reports |
| reportFolders | Manage saved report folders |
| reportSettings | Reporting engine configuration options |
| runReport | Run flow reports |
| saveReport | Name and save flow reports |
| scheduledReports | Create, edit, and delete scheduled email reports |
| sf_asa_acls | Configure ASA ACL descriptions |
| SNMPCredentials | Manage SNMP credentials for polling device information |
| srCreate | Schedule saved reports to be emailed on a regular basis |
| sso | Add, delete, and edit Identity Provider configurations for Single Sign-On |
| statusTab | Access the Status tab |
| syslogNotifications | Syslog server configuration |
| SystemPreferences | Administrative access to global Plixer Scrutinizer preferences |
| tacacsConf | Manage TACACS+ server configuration for Plixer Scrutinizer authentication |
| tos | Edit TOS configuration |
| userAccounts | Admin access to the user management page |
| usergroups | Manage Scrutinizer user groups |
| viewUserIdentity | View identity and access information relevant to GDPR restrictions |
| viptelaSettings | Configure Viptela settings |
| Vitals | View Plixer Scrutinizer server vitals reports |

Table 1 – continued from previous page

| | | |
|--|-----|------------------------|
| | wkp | Edit WKP configuration |
|--|-----|------------------------|

API request example

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/
→fcgi/scrut_fcgi.fcgi' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=user_api' \
--form 'action=permissions' \
--form 'json=
{
  "permissions":
  {
    "add": [
      {
        "usergroup_name": "Dashboarders",
        "permission_type": "device",
        "seccode": "0A010107"
      }
    ],
    "remove": [
      {
        "usergroup_name": "ReadOnlyReporters",
        "permission_type": "plexer",
        "seccode": "allGadgets"
      }
    ]
  }
}
```

Returned JSON object

```
{
  "data":
  {
    "errors": []
  }
}
```

(continues on next page)

(continued from previous page)

```
    "updated": [
      "Added device permission 0A010107 to usergroup 26",
      "Removed plixer permission allGadgets from usergroup 27"
    ]
  }
}
```

changeUsername

The changeUsername action is used to edit the name of an existing user account and requires an additional json field specifying the account (by oldname or user_id) and the new name.

API request example

```
curl --location --insecure --request POST 'https://<SCRUTINIZER_ADDRESS>/
↳fcgi/scrut_fcgi.fcgi' \
--form 'authToken=<AUTH_TOKEN>' \
--form 'rm=user_api' \
--form 'action=changeUsername' \
--form 'json=
{
  "changeUsername":
  {
    "oldname": "MyUser",
    "newname": "OpSCT"
  }
}'
```

Returned JSON object

```
{
  "data":
  {
    "message": "User MyUser successfully renamed to OpSCT"
  }
}
```

6.6 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.

6.6.1 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.

6.6.2 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.

6.7 Streaming to data lakes

Plixer Scrutinizer supports data streaming to customer data lakes.

For assistance with the configuration process, contact [Plixer Technical Support](#).

6.8 Upgrades and updates

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.

6.8.1 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-install.run.sha256 -L https://files.plixer.com/  
↳plixer-repo/scrutinizer/19.6.1/scrutinizer-install.run.sha256
```

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 the 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.

6.8.2 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 upgrades may include additional functionality, performance enhancements, and/or QoL improvements, in addition to implementing fixes for certain types of issues.

Latest release

After completing the recommended [pre-upgrade preparations](#), follow the instructions below to upgrade Plixer Scrutinizer to the latest version.

Note:

- Only deployments on v19.5.3 (v19.5.4 for AWS AMIs) can be upgraded directly to v19.6.0 and beyond. For older versions, follow the steps in [these guides](#) to upgrade to the required Plixer Scrutinizer 19.5.x release before upgrading to the latest version.
- If the Plixer Scrutinizer server being upgraded does not have Internet access, an internal NTP server can be configured by running the following:

```
sed -i -e '/^pool/aserver NTP_ADDRESS' -e 's/^pool/#&/' /etc/chrony.conf
```

- After Plixer Scrutinizer has been upgraded to v19.6.0 or higher, any Plixer ML Engine deployments must also be [upgraded to v19.5.0 or the latest available version](#). This will enable the ML management and configuration options integrated into the Plixer Scrutinizer UI.
- To obtain a copy of the AMI installer, contact [Plixer Technical Support](#).

Contact [Plixer Technical Support](#) for any concerns or assistance.

Online upgrades

To download and install the latest version upgrade for Plixer Scrutinizer, follow these steps:

1. SSH to the primary reporter as the `plixer` user and start a new tmux session (to maintain the upgrade session if the SSH connection is lost):

```
tmux new -s upgrade
```

2. Download the installer for the latest version:

```
cd /tmp
curl -o scrutinizer-install.run https://files.plixer.com/plixer-repo/
↳scrutinizer/19.6.1/scrutinizer-install.run
```

3. Download the checksum file and validate the integrity of `scrutinizer-install.run`:

```
curl -o scrutinizer-install.run.sha256 https://files.plixer.com/plixer-
↳repo/scrutinizer/19.6.1/scrutinizer-install.run.sha256
cat scrutinizer-install.run.sha256
sha256sum scrutinizer-install.run
```

4. Set the correct permissions for the installer:

```
sudo chmod 755 scrutinizer-install.run
```

5. Run the installer as the `plixer` user:

```
./scrutinizer-install.run
```

6. [Distributed cluster upgrades only] When prompted for the authentication method to use for remote collectors in the cluster, enter either `existing` (recommended) or `passwords`.
7. After the installer finishes running, execute the following heartbeat checks to verify communication between nodes:

```
scrut_util --check heartbeat --type database
scrut_util --check heartbeat --type api
```

If the heartbeat checks are successful, the upgrade is complete.

Offline upgrades

To upgrade Plixer Scrutinizer collectors/servers that are unable to access <https://files.plixer.com>, an offline repository can be created on the primary reporter.

To set up the offline repository and start the upgrade process, follow these steps:

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/19.6.1_offline.tgz to a computer with Internet access.
3. Download https://files.plixer.com/plixer-repo/scrutinizer/19.6.1_offline.tgz.sha256 and validate the checksum of the *.tar* file.
4. Start an SSH session with the primary reporter as the `plixer` user.
5. Confirm that the primary reporter has at least 84 GB of free disk space under `/var/db/big`:

```
df -h
```

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. Copy the *.tar* file to the new directory (where `REPO_HOST_IP` is the IP address of the primary reporter/offline repository host):

```
scp 19.6.1_offline.tgz plixer@REPO_HOST_IP:/var/db/big/offline/19.6.1_  
↪offline.tgz
```

8. Extract the contents of the file to the same directory (this will exhaust all available storage on the appliance):

```
sudo tar -zxvf /var/db/big/offline/19.6.1_offline.tgz -C /var/db/big/  
↪offline
```

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 using the offline repository host IP:

```
curl -o scrutinizer-install.run -L -k https://REPO_HOST_IP/plixer-  
↳ repo/scrutinizer/19.6.1/scrutinizer-install.run
```

11. Download the checksum file and validate the integrity of the `scrutinizer-install.run` file.

```
curl -o scrutinizer-install.run.sha256 -k https://REPO_HOST_IP/  
↳ plixer-repo/scrutinizer/19.6.1/scrutinizer-install.run.sha256  
  
cat scrutinizer-install.run.sha256  
  
sha256sum scrutinizer-install.run
```

12. Set the correct permissions for the installer.

```
sudo chmod 755 scrutinizer-install.run
```

13. Run the installer as the `plixer` user, replacing `x.x.x.x` with the IP address or hostname of the offline repository host:

```
REPO_HOST=x.x.x.x  
LOCAL_REPO_BASEDIR=/var/db/big/offline  
./scrutinizer-install.run -- -k
```

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`.

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.

Upgrading pre-19.5.x deployments

The Plixer Scrutinizer 19.5.0 upgrade includes the [migration to Oracle Linux 9](#), which will be required for all new versions/releases going forward. Deployments on older versions must first be upgraded to the latest v19.5.x release before being upgraded further.

Select the appropriate guide below for instructions to complete the required upgrade(s).

- [Upgrading to v19.5.3 from v19.4.0 or above](#)
- [Upgrading to v19.5.4 from v19.4.0 or above](#) (AWS AMI only)
- [Upgrading to 19.4.0 from previous versions](#) (required to upgrade to v19.5.3/v19.5.4)

Upgrading to v19.5.3

Follow the steps outlined below to upgrade a Plixer Scrutinizer deployment on v19.4.0 or above to v19.5.3.

To upgrade an AWS AMI from 19.4.0 to v19.5.4, follow [this guide](#) instead. For older versions, refer to [this guide](#) to upgrade to v19.4.0 before proceeding.

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 SSH login is enabled on the Plixer Scrutinizer server, it will be disabled as part of the upgrade.
 - If upgrading from v19.5.0 or above, proceed directly to [upgrading to Plixer Scrutinizer 19.5.3](#).
 - 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.
-

For assistance or clarifications, contact [Plixer Technical Support](#).

Upgrade process

The process of upgrading a v19.4.0 Plixer Scrutinizer server to v19.5.3 involves the following steps:

- [Backing up the current install's 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). **This only applies if upgrading from v19.4.0.**
- Downloading and running the Plixer Scrutinizer v19.5.3 installation script (`scrutinizer-install.run`)
- Verifying that the current install's data has been successfully migrated after v19.5.3 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](#).

Distributed cluster upgrades

Nodes in distributed environments must be reverted to standalone appliances before being individually upgraded to v19.5.3:

1. Navigate to [Admin > Resources > Collectors](#) and delete all remote collectors.
2. SSH to each remote collector as the `plixer` user and register it as a standalone appliance:

```
scrut_util --set selfregister --reset
```

3. Verify that each appliance is now running in standalone mode (no other addresses under `collector_ips`):

```
scrut_util --check dist_info
```

When done, proceed with the OS migration and v19.5.3 upgrade for each node, and then *rebuild the distributed cluster*.

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://$REPO_HOST/plixer-repo/scrutinizer/19.5.
↪3/olmigrate.run
curl -k -o olmigrate.run.sha256 https://$REPO_HOST/plixer-repo/
↪scrutinizer/19.5.3/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 Plixer 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 Plixer Scrutinizer 19.5.3

Once the appliance is on the *new OS*, Plixer Scrutinizer can be upgraded to v19.5.3 as follows:

1. Change directories to /tmp:

```
cd /tmp/
```

2. Download the Plixer Scrutinizer v19.5.3 installation script and its checksum file:

```
REPO_HOST=files.plixer.com
curl -k -o scrutinizer-install.run https://$REPO_HOST/plixer-repo/
↪scrutinizer/19.5.3/scrutinizer-install.run
curl -k -o scrutinizer-install.run.sha256 https://$REPO_HOST/plixer-repo/
↪scrutinizer/19.5.3/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 Plixer Scrutinizer v19.5.3:

```
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 cluster has been reestablished.*

If the heartbeat checks are successful, then the Plixer Scrutinizer appliance has been successfully upgraded to v19.5.3.

Offline upgrades to v19.5.3

The following instructions for setting up an offline repo are intended for upgrading to Plixer Scrutinizer v19.5.3 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 19.5.3 and its checksum file:

```
curl -o /var/db/big/offline/19.5.3_offline.tgz https://files.plixer.com/
↳plixer-repo/scrutinizer/19.5.3_offline.tgz
curl -o /var/db/big/offline/19.5.3_offline.tgz.sha256 https://files.
↳plixer.com/plixer-repo/scrutinizer/19.5.3_offline.tgz.sha256
```

5. Validate the integrity of `19.5.3_offline.tgz`:

```
sha256sum -c /var/db/big/offline/19.5.3_offline.tgz.sha256
```

6. Extract the offline tar file:

```
cd /var/db/big/offline
tar xvf 19.5.3_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 v19.5.4 (AMI only)

Follow the steps outlined below to upgrade a Plixer Scrutinizer AMI on v19.4.0 or above to v19.5.4.

For older versions, refer to [this guide](#) to upgrade to v19.4.0 before proceeding.

Note:

- The upgrade will take **at least one hour to complete**.
 - The `plixer` user SSH password will be needed during the upgrade.
 - If root SSH login is enabled on the Plixer Scrutinizer server, it will be disabled as part of the upgrade.
 - The new v19.5.4 instance must be in the same availability zone as the original v19.4.0 machine. Volumes outside the current availability zone will not be accessible from the AWS console.
-

Distributed cluster upgrades

Nodes in distributed environments must be reverted to standalone appliances before being individually upgraded to v19.5.4:

1. Navigate to [Admin > Resources > Collectors](#) and delete all remote collectors.
2. SSH to each remote collector as the `plixer` user and register it as a standalone appliance:

```
scrut_util --set selfregister --reset
```

3. Verify that each appliance is now running in standalone mode (no other addresses under `collector_ips`):

```
scrut_util --check dist_info
```

When done, proceed with upgrading each node as described below, and then [rebuild the distributed cluster](#).

For assistance or clarifications, contact [Plixer Technical Support](#).

Upgrade process

The process of upgrading a Plixer Scrutinizer 19.4.0 Plixer Scrutinizer AMI to v19.5.4 involves the following steps:

- Backing up the current Plixer Scrutinizer install by taking a VM snapshot
- Deploying a new v19.5.4 AMI appliance
- Copying the `dbexport.sh` file from the new v19.5.4 appliance to the current v19.4.0 appliance
- Detaching the storage volume from the v19.4.0 instance (using `dbexport.sh` and running as the `root` user)
- Attaching the storage volume to the new v19.5.4 instance (using `dbimport.sh` and running as the `root` user)
- Verifying that the v19.4.0 data has been successfully migrated after v19.5.4 is installed

Expanding storage

AMI deployments will require additional storage to be upgraded to v19.5.4.

To verify whether the Plixer Scrutinizer 19.4.0 AMI instance is running on the default sizing, run the following:

```
df -h
```

If the output does not list a line that includes `vg_scrut-lv_db`, contact [Plixer Technical Support](#) for assistance with expanding storage before proceeding with the [upgrade to v19.5.4](#).

Upgrading to Plixer Scrutinizer 19.5.4

1. Copy the following file from the new v19.5.4 appliance to your current v19.4.0 appliance:

```
/home/plixer/scrutinizer/files/dbimport/dbexport.sh
```

2. Run the following command to make `dbexport.sh` executable:

```
sudo chmod +x dbexport.sh
```

- SSH to the 19.4.0 appliance as the `plixer` user, and then navigate to the location where `dbexport.sh` was saved.
- Run the script to prepare the 19.4.0 storage volume to be detached:

```
sudo ./dbexport.sh exportdb
```

- Shut down the Plixer Scrutinizer v19.4.0 instance.
- In the AWS EC2 management page, navigate to the *Volumes* page.
- In the *Volume Management* page, select the storage volume, click the **Actions** menu, and then select **Detach volume**. It may take a minute for the storage volume to go from *In use* to *Available*.
- Once the detached storage volume(s) are marked as *Available* (it may take several minutes), attach it to the Plixer Scrutinizer v19.5.4 instance. Refer to STEP 6 of the [storage expansion instructions](#).
- After the storage volume(s) have been attached, SSH to the 19.5.4 instance as the `plixer` user.
- Run the following to import and set up the database on the 19.5.4 instance.

```
/home/plixer/scrutinizer/files/dbimport/dbimport.sh importdb  
↪ <device name for the storage volume, e.g. /dev/xvdg>
```

Use the `lsblk` and `show partitions` commands to get the correct partition/device name to use. Once the script completes running, Plixer Scrutinizer will run a self-register reset that requires user input for verification.

- Add a new license key to fully activate your Plixer Scrutinizer v19.5.4 instance. You can refer to [this guide on how to add a license](#).

Note:

- If there are multiple volumes listed after `dbexport.sh` completes running, all volumes will need to be detached from the v19.4.0 instance and attached to the v19.5.4 instance.
- At the end of the output from `dbexport`, the volumes that are part of the volume group for the database are listed. If the volume group contains more than one volume, the output will list all of those volumes, which you will need to detach and then attach to the Plixer Scrutinizer v19.5.4 instance.
- When you first log in to the v19.5.4 UI to add a new Plixer Scrutinizer license, you must use the UI admin password for the v19.4.0 AWS instance. Alternatively, you can reset the UI admin password in `scrut_util` first.

Upgrading to 19.4.0

Pre-v19.4.0 Plixer Scrutinizer deployments must first be upgraded to v19.4.0 before being upgraded to v19.5.3 (or v19.5.4 for AWS AMI appliances), which includes the migration to Oracle Linux 9.

Follow [these instructions](#) to download the v19.4.0 installer (replace 19.6.1 with 19.4.0 in the download URLs) and apply the update. Once done, proceed with [upgrading to the latest v19.5.x release](#).

Note:

- When upgrading an appliance that was previously upgraded from v18.20, the installer script will ask whether to delete the `data.old` backup created during that upgrade. Since a more recent backup [should be created](#) before the current upgrade process, this file can safely be deleted.
 - If a distributed cluster is being upgraded from v18.20, the prompt to create a new *Plixer control key* should be left blank unless encrypted keys are required. Additionally, passwords should be selected in the next step, when prompted for the login method to use for remote collectors.
-

Plixer ML Engine

Review/complete the [recommended upgrade preparations](#), and then follow these steps to upgrade a Plixer ML Engine deployment to the latest version:

1. SSH to the Plixer ML Engine VM (i.e., the host used for management/deployment) as the `plixer` user.
2. Download the installer for the latest version:

```
curl -o plixer-machine-learning-update.run https://files.plixer.com/  
↳scripts/plixer-machine-learning/release/19.5.0/plixer-machine-learning-  
↳update.run
```

3. Download the checksum file and validate the integrity of `plixer-machine-learning-update.run`:

```
curl -o plixer-machine-learning-checksums.txt https://files.plixer.com/  
↳scripts/plixer-machine-learning/release/19.5.0/plixer-machine-learning-  
↳checksums.txt  
cat plixer-machine-learning-checksums.txt  
sha256sum plixer-machine-learning-update.run
```

4. Set the correct permissions for the installer:

```
chmod +x plixer-machine-learning-update.run
```

5. Run the installer as the `plixer` user:

```
STAGE="release"  
VERSION="19.5.0"  
STAGE=$STAGE ./plixer-machine-learning-update.run
```

After the installer script completes running, `setup.sh` will automatically be run to pull in any configuration changes and redeploy pods with new images.

Note: If any changes were previously made to `pxi-settings.yaml`, `azure.tfvars`, `aws.tfvars`, or `vsphere.tfvars`, the file(s) will be retained even if the upgrade package includes a newer version of the file. The updated file will instead be saved with a `.dpk-dist` extension, and any necessary edits should be migrated before it is used to overwrite the old configuration/tfvars file.

Once the upgrade process is complete, wait for the **rke2-server** service to restart. This sequence can be monitored by running:

```
journalctl -xeu rke2-server -f
```

Additional notes for Plixer ML Engine upgrades from v19.4.0 to v19.5.0

- Plixer Scrutinizer 19.6.0 includes new management/configuration functions for the Plixer ML Engine, requiring all attached engine deployments to also be upgraded from v19.4.0 to v19.5.0 or higher.
- After Plixer Scrutinizer is upgraded to v19.6.0 or higher, all previous settings related to attached Plixer ML Engine deployments will be reset. Engines will need to be *re-registered* (but not re-deployed) via the Plixer Scrutinizer web interface before being upgraded to v19.5.0.
- When upgrading from v19.4.0 to v19.5.0, `setup.sh --reconfigure` will automatically be run (instead of `setup.sh` as described above) to initiate the new configuration process and collect all required information (including the authentication token generated by Plixer Scrutinizer).
- If the Plixer ML Engine is deployed as a standalone VM, new Docker images will be downloaded (may take several minutes) after the package updates. This step is skipped for cloud deployments.

6.8.3 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.

6.8.4 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 Appendices

This section contains additional references/guides for Plixer Scrutinizer's functional elements.

7.1.1 Alarm policies

Refer to the following tables for information on alarm policies, violations, and event messages.

Alarm policy list

The table below contains general information for all alarm policies available in Plixer Scrutinizer.

| Category | Policy | Technology | License | Description |
|--|-----------------------------------|-------------------------|-----------------------|---|
| Collection > Data Staged > Local Data Staging | Data Accumulation | Plixer Machine Learning | Plixer One Enterprise | A host is accumulating data from various internal sources in preparation to exfiltrate |
| Command and Control > Application Layer Protocol > DNS | DNS Command and Control Detection | Scrutinizer | Plixer One Core | This algorithm monitors the use of DNS TXT messages traversing the network perimeter as detected by FlowPro Defender. DNS TXT messages provide a means of sending information into and out of your protected network over DNS, even when you have blocked use of an external DNS server. This technique is used by malware as a method of controlling compromised assets within your network and to extract information back out. Additionally, some legitimate companies also use this method to communicate as a means to 'phone home' from their applications to the developer site. The algorithm will detect inbound, outbound, and bidirectional communications using DNS TXT messages. Thresholds may be set based either on the number of DNS TXT messages or the number of bytes observed in the DNS TXT messages within a three-minute period. The default setting is for any detected traffic to alarm, and alarm aggregation defaults to 120 minutes. To suppress alarms from authorized applications in your network, you may add the domain generating the alarm message to the 'trusted.domains' list on FlowPro Defender. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|---|-------------------------|-----------------------|---|
| Command and Control > Application Layer Protocol > DNS | DNS Hits | Scrutinizer | Plixer One Core | Triggers an alarm when a host initiates an excessive number of DNS queries. This identifies hosts that perform an inordinate number DNS lookups. To do this, set the flow threshold to a large value that reflects normal behavior on your network. The default threshold is 2500 DNS flows in three minutes. Either the source or destination IP address can be excluded from triggering this alarm. |
| Command and Control > Application Layer Protocol > DNS | DNS Server Detection | Scrutinizer | Plixer One Core | When used with FlowPro Defender, detects new DNS Servers being used on or by your network through analysis of the DNS packets being exchanged between the client and the server. Exclude DNS servers that are authorized for use on the network. |
| Command and Control > Custom Command and Control Protocol | Detection of a non-standard protocol or event | Plixer FlowPro Defender | Plixer One Enterprise | Detects non-standard protocols or events (e.g. use of deprecated or rarely used protocols) |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|---------------------------------|--------------------------|-----------------------|--|
| Command and Control > Custom Command and Control Protocol | Generic Protocol Command Decode | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects generic protocol command decodes (e.g. malformed DHCP options) |
| Command and Control > Data Obfuscation > Protocol Impersonation | Protocol Misdirection | Scrutinizer | Plixer One Enterprise | Identifies when the type of traffic doesn't match the port being used. |
| Command and Control > Dynamic Resolution | BotNet Detection | Scrutinizer | Plixer One Core | This alarm is generated when a large number of unique DNS name lookups have failed. When a DNS lookup fails, a reply commonly known as NXDOMAIN is returned. By monitoring the number of NXDOMAINs detected as well as the DNS name looked up, behavior normally associated with a class of malware that uses Domain Generation Algorithms (DGAs) can be detected. The default threshold is 100 unique DNS lookup failures (NXDOMAIN) messages in three minutes. Either the source or destination IP address can be excluded from triggering this alarm. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|---|---|-----------------------|---|
| Command and Control > Dynamic Resolution | Domain Observed Used for C2 Detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects domains known to be used for malware command and control |
| Command and Control > Encrypted Channel | Encrypted traffic alert | Plixer Machine Learning, Plixer Flow-Pro Defender | Plixer One Enterprise | Detects anomalous encrypted network traffic |
| Command and Control > Non-Standard Port | Malware Command and Control Activity Detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects malware communicating with an external command and control server |
| Command and Control > Non-Standard Port | ML Engine command and control alert | Plixer Machine Learning | Plixer One Enterprise | Detect traffic signatures that are similar to those of well known banking trojans (Dridex, Emotet, Quakbot, Trickbot) |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|--------------------------------------|-------------------------|-----------------------|--|
| Command and Control > Proxy > External Proxy | Tunneling through external DNS host | Plixer Machine Learning | Plixer One Enterprise | Detect when an external host is being used as a DNS proxy tunnel to another host |
| Command and Control > Proxy > External Proxy | Tunneling through external ICMP host | Plixer Machine Learning | Plixer One Enterprise | Detect when an external host is being used as an ICMP proxy tunnel to another host |
| Command and Control > Proxy > External Proxy | Tunneling through external SSH host | Plixer Machine Learning | Plixer One Enterprise | Detect when an external host is being used as an SSH proxy tunnel to another host |
| Command and Control > Proxy > Internal Proxy | Tunneling through internal DNS host | Plixer Machine Learning | Plixer One Enterprise | Detect when an internal host is being used as a DNS proxy tunnel to another host |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|--------------------------------------|-------------------------|-----------------------|---|
| Command and Control > Proxy > Internal Proxy | Tunneling through internal ICMP host | Plixer Machine Learning | Plixer One Enterprise | Detect when an internal host is being used as an ICMP proxy tunnel to another host |
| Command and Control > Proxy > Internal Proxy | Tunneling through internal SSH host | Plixer Machine Learning | Plixer One Enterprise | Detect when an internal host is being used as an SSH proxy tunnel to another host |
| Command and Control > Remote Access Software | ML Engine remote access trojan alert | Plixer Machine Learning | Plixer One Enterprise | Detect traffic signatures that are similar to those associated with remote access trojans |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|-------------------|-------------|-----------------|--|
| Command and Control > Web Service > Bidirectional Communication | Domain Reputation | Scrutinizer | Plixer One Core | Domain reputation provides much more accurate alarming with a dramatic decrease in the number of false positive alarms as compared to IP based Host Reputation. The domain list is provided by Plixer and is updated each hour and currently contains over 400,000 known bad domains. FlowPro Defender performs the actual monitoring, and when it detects a domain with poor reputation, it passes the information to Scrutinizer for additional processing. The default setting is for any detected traffic to alarm, and alarm aggregation defaults to disabled so that all DNS lookups observed will result in a unique alarm. To suppress alarms from authorized applications in your network, you may add the domain generating the alarm message to the 'Trusted Domain' list on FlowPro Defender. See the discussion on FlowPro Defender for additional details. |
| Command and Control > Web Service > Bidirectional Communication | Host Reputation | Scrutinizer | Plixer One Core | This algorithm maintains a current list of active Tor nodes that you should monitor. Some malware families use Tor for Command and Control communications. White-list your users who are authorized to use Tor and regard other uses as suspicious. This algorithm will also monitor any IP address lists that you provide as a custom list as described in the 'Custom List' section that follows. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|---------------------------|-------------|-----------------------|---|
| Command and Control > Web Service > Bidirectional Communication | Host Watchlist | Scrutinizer | Plixer One Enterprise | Identifies hosts that have violated internal host watchlist |
| Command and Control > Web Service > Bidirectional Communication | NetFlow Domain Reputation | Scrutinizer | Plixer One Core | A blacklisted domain has been detected in NetFlow traffic |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|--------------------------|-------------------------|-----------------------|--|
| Credential Access > Adversary-in-the-Middle > DHCP Spoofing | Rogue DHCP Service | Plixer Machine Learning | Plixer One Enterprise | Find rogue DHCP services that may not be known or desired on a network |
| Credential Access > Adversary-in-the-Middle > DHCP Spoofing | Rogue LDAP Service | Plixer Machine Learning | Plixer One Enterprise | Find rogue LDAP services that may not be known or desired on a network |
| Credential Access > Brute Force | Breach Attempt Detection | Scrutinizer | Plixer One Core | This algorithm is examining flow behaviors that may indicate a brute force password attack on an internal IP address. This is accomplished by examining the flow, byte, and packet counts being exchanged in short-duration completed flows between one source and one destination, with specific behaviors observed for common attack vectors such as SSH, LDAP and RDP. If the number of flows that match these characteristics exceeds the alarm threshold, an alarm will be raised. The default flow count threshold is 100. Either IP address can be excluded from triggering this alarm. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|-----------------------------------|-------------------------|-----------------------|---|
| Credential Access > Brute Force > Password Cracking | Zerologon | Plixer Machine Learning | Plixer One Enterprise | Detect traffic signatures that are similar to those associated with Zerologon malware |
| Credential Access > Brute Force > Password Guessing | Brute-force RDP (Client-side) | Plixer Machine Learning | Plixer One Enterprise | Detects a client trying to gain access to RDP via brute force attack |
| Credential Access > Brute Force > Password Guessing | Brute-force RDP (Server-side TCP) | Plixer Machine Learning | Plixer One Enterprise | Detects a server experiencing an RDP (tcp) brute force attack |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|-----------------------------------|-------------------------|-----------------------|--|
| Credential Access > Brute Force > Password Guessing | Brute-force RDP (Server-side UDP) | Plixer Machine Learning | Plixer One Enterprise | Detects a server experiencing an RDP (udp) brute force attack |
| Credential Access > Brute Force > Password Guessing | Brute-force SSH (Client-side) | Plixer Machine Learning | Plixer One Enterprise | Detects a client trying to gain access to SSH via brute force attack |
| Credential Access > Brute Force > Password Guessing | Brute-force SSH (Server-side) | Plixer Machine Learning | Plixer One Enterprise | Detects a server experiencing a SSH brute force attack |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|--------------------------------------|---|-----------------------|---|
| Credential Access > Brute Force > Password Guessing | SMB Brute-force Attempt | Plixer Machine Learning, Plixer Flow-Pro Defender | Plixer One Enterprise | Detects a client trying to gain access to an SMB server via brute force password guessing |
| Credential Access > Credential Dumping | Successful Credential Theft Detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects successful attempts at stealing user credentials |
| Defense Evasion > Non-Application Layer Protocol | A client was using an unusual port | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects when a client is using an unusual port for a given well-known protocol (e.g. a client sending HTTP requests over a non-standard port) |
| Defense Evasion > Obfuscated Files or Information | A suspicious filename was detected | Plixer Flow-Pro Defender | Plixer One Enterprise | A suspicious filename is detected that is often related to known malware families |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--------------------------------------|----------------------------------|--------------------------|-----------------------|--|
| Discovery > Network Service Scanning | Detection of a Network Scan | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects network scanning activities (e.g. a large number of requests to different ports on a single machine or multiple machines) |
| Discovery > Network Service Scanning | FIN Scan (Internal) | Scrutinizer | Plixer One Core | Alerts when a FIN scan is detected. FIN scans are often used as reconnaissance prior to an attack. They are considered to be a 'stealthy scan' as they may be able to pass through firewalls, allowing an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Discovery > Network Service Scanning | ICMP Port Unreachable (Internal) | Scrutinizer | Plixer One Core | This alarm is generated when a large number of ICMP destination unreachable messages have been sent to the suspect IP address. This may happen as a result of scanning activity, misconfiguration, or network errors. ICMP Destination Unreachable is a message that comes back from a destination host or the destination host gateway to indicate that the destination is unreachable for one reason or another. The default threshold is 100 destination unreachable messages. Either the source or destination IP address can be excluded from triggering this alarm. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|------------------------------|-------------|-----------------|---|
| Discovery > Network Service Scanning | NULL Scan (Internal) | Scrutinizer | Plixer One Core | Alerts when a NULL scan is detected. NULL scans are a TCP scan with all TCP Flags cleared to zero. This scan is often used as reconnaissance prior to an attack. They are considered to be a ‘stealthy scan’ as they may be able to pass through firewalls, allowing an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Discovery > Network Service Scanning | Odd TCP Flags (Internal) | Scrutinizer | Plixer One Core | Alerts when a scan is detected using unusual TCP Flag combinations. These types of scans may allow an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Discovery > Network Service Scanning | RST/ACK Detection (Internal) | Scrutinizer | Plixer One Core | Alerts when a large number of TCP flows containing only RST and ACK flags have been detected being sent to a single destination. These flows indicate that a connection attempt was made on the host sending the RST/ACK flow, and was rejected. This algorithm may detect other scan types used by an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--------------------------------------|---------------------------|-------------|-----------------------|---|
| Discovery > Network Service Scanning | Slow Port Scan (Internal) | Scrutinizer | Plixer One Enterprise | Detects when a large number of ports have been probed on the target machine over a long period of time. This alert could indicate malicious activity or reconnaissance for lateral movement. |
| Discovery > Network Service Scanning | SYN Port Scan (Internal) | Scrutinizer | Plixer One Core | Alerts when a SYN scan is detected. SYN scans are a TCP scan with the TCP SYN Flag set. This scan is often used as reconnaissance prior to an attack as it is fast and somewhat stealthy. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Discovery > Network Service Scanning | TCP Half-Open (Internal) | Scrutinizer | Plixer One Core | Alerts when a SYN scan is detected. SYN scans are a TCP scan with the TCP SYN Flag set. This scan is often used as reconnaissance prior to an attack as it is fast and somewhat stealthy. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Discovery > Network Service Scanning | TCP Scan (Internal) | Scrutinizer | Plixer One Core | Alerts when a possible TCP scan is detected from an exporter that does not provide TCP Flag information. These types of scans may allow an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|--|--------------------------|-----------------------|--|
| Discovery > Network Service Scanning | UDP Scan (Internal) | Scrutinizer | Plixer One Core | Alerts when a possible UDP scan is detected. These types of scans may allow an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. NOTE: if your policy allows P2P traffic on your network, then you will likely want to exclude the allowed host(s) or disable this alarm as it will often detect P2P control traffic as a UDP Scan violation. |
| Discovery > Network Service Scanning | Xmas Scan (Internal) | Scrutinizer | Plixer One Core | Alerts when a XMAS scan is detected. XMAS scans are a TCP scan with the FIN, PSH, and URG TCP Flags set. This scan is often used as reconnaissance prior to an attack. They are considered to be a 'stealthy scan' as they may be able to pass through firewalls, allowing an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Discovery > Remote System Discovery | Device Retrieving External IP Address Detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects devices retrieving their external IP addresses (e.g. a device making a request to whatismyip services, commonly used in malware recon and exfiltration) |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|-------------------------------------|---|-------------------------|-----------------------|---|
| Discovery > Remote System Discovery | ICMP Destination Unreachable (Internal) | Scrutinizer | Plixer One Core | This alarm is generated when a large number of ICMP destination unreachable messages have been sent to the suspect IP address. This may happen as a result of scanning activity, misconfiguration, or network errors. ICMP Destination Unreachable is a message that comes back from a destination host or the destination host gateway to indicate that the destination is unreachable for one reason or another. The default threshold is 100 destination unreachable messages. Either the source or destination IP address can be excluded from triggering this alarm. |
| Discovery > Remote System Discovery | Lateral Movement Behavior | Plixer Machine Learning | Plixer One Enterprise | Detect a host moving laterally inside a network during a Reconnaissance phase |
| Discovery > Remote System Discovery | Ping Scan (Internal) | Scrutinizer | Plixer One Enterprise | Alerts when a host is suspected of performing a ping scan. A ping scan uses ICMP Echo Requests (ping) to discover what IPs are in use on a network. The behavior is commonly demonstrated by attackers attempting to find targets for compromise or lateral movement. |
| Discovery > Remote System Discovery | SYN IP Scan (Internal) | Scrutinizer | Plixer One Core | Alerts when a SYN scan is detected. SYN scans are a TCP scan with the TCP SYN Flag set. This scan is often used as reconnaissance prior to an attack as it is fast and somewhat stealthy. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|--------------------------|-------------------------|-----------------------|---|
| Discovery > Remote System Discovery | Worm Activity | Plixer Machine Learning | Plixer One Enterprise | Network traffic patterns appear to indicate a worm malware propagating throughout the network |
| Discovery > System Network Connections Discovery | Lateral Movement Attempt | Scrutinizer | Plixer One Enterprise | Identifies behavior from a host which could be attempted lateral movement. |
| Endpoint Data | Endpoint Analytics Info | Endpoint Analytics | Plixer One Enterprise | Informational messages from Endpoint Analytics |
| Execution > Command and Scripting Interpreter | Reverse SSH Shell | Scrutinizer | Plixer One Enterprise | Identifies possible reverse SSH tunnels to external destinations. A reverse SSH tunnel allows an external entity access to internal, protected resources via use of an established outbound SSH connection. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|-------------------------------|--------------------------|-----------------------|--|
| Execution > Exploitation for Client Execution | Exploit Kit Activity Detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects known exploit kit activities |
| Execution > Exploitation for Client Execution | SIGRed Exploit Attempt | Plixer Machine Learning | Plixer One Enterprise | Detect malformed DNS query responses which could be used as an exploit via SigRED |
| Execution > System Services | A system call was detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects when a potential system call was made (e.g. x86 shellcode found in a network payload) |
| Execution > System Services | Executable code was detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects when executable binary shellcode is detected in a network payload |
| Execution > User Execution > Malicious File | ML Engine exploit kit alert | Plixer Machine Learning | Plixer One Enterprise | Detect traffic signatures that are similar to those associated with RigEK + Ramnit exploit kit |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|---------------------------|-------------------------|-----------------------|---|
| Execution > User Execution > Malicious Link | Blocked Malicious Domains | Plixer Machine Learning | Plixer One Enterprise | A known malicious domain has been blocked by Plixer DNS proxy |
| Exfiltration > Exfiltration Over Alternative Protocol | Data Exfiltration | Plixer Machine Learning | Plixer One Enterprise | A host is exfiltrating large amounts of data to an external host |
| Exfiltration > Exfiltration Over Alternative Protocol | DNS Data Leak Detection | Scrutinizer | Plixer One Core | This algorithm monitors the practice of encoding information into a DNS lookup message that has no intention of returning a valid IP address or making an actual connection to a remote device. When this happens, your local DNS server will fail to find the DNS name in its cache, and will pass the name out of your network to where it will eventually reach the authoritative server for the domain. At that point, the owner of the authoritative server can decode the information embedded in the name, and may respond with a 'no existing domain' response, or return a non-routable address. Flow-Pro Defender uses proprietary detection algorithms to identify suspicious DNS names that may contain encoded data, and passes this information to Scrutinizer where it is processed by the DNS Data Leak algorithm. Thresholds may be set based either on the number of suspicious DNS names or the number of bytes observed in the suspicious DNS name within a three-minute period. The default setting is for any detected traffic to alarm, and alarm aggregation defaults to 120 minutes. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|---|---|-----------------------|--|
| Flow-Pro Event Captured | Flow-Pro Event Capture | Plixer Flow-Pro Defender | Plixer One Enterprise | A user defined FlowPro capture rule. |
| Forecast Events | Forecast Anomaly | Plixer Machine Learning | Plixer One Enterprise | An anomaly outside the range of a network forecast has been detected |
| Impact > Data Encrypted for Impact | Ransomware Behavior | Plixer Machine Learning, Plixer Flow-Pro Defender | Plixer One Enterprise | Detects a client accessing an SMB share and potentially encrypting files |
| Impact > Endpoint Denial of Service > Application or System Exploitation | Detection of a Denial of Service Attack | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects Denial of Service (DoS) attacks |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|-------------------|--------------------------|-----------------------|---|
| Impact > End-point Denial of Service > Application or System Exploitation | Large Ping | Scrutinizer | Plixer One Enterprise | Alerts on the observance of unusually large ICMP Echo Request (ping) packets. This alert could indicate malicious activity within the network including possible Denial of Service (DoS) attempts. |
| Impact > Network Denial of Service | DDoS | Scrutinizer | Plixer One Core | Identifies generic Distributed Denial of Service (DDoS) attacks targeted at your protected network space. Refer to the DRDoS algorithm for detection of the more common Distributed Reflection DoS attacks. Note that DDoS algorithm may take a lot of time depending on the exporters selected. There are four settings which are used to adjust the sensitivity of the DDoS detection algorithm: DDoS Packet Deviation (10) and DDoS Bytes Deviation (10) - These settings control how similar the flows associated with the attack must be. The standard deviation of the byte count and packet counts associated with the flows must be less than this setting for DDoS attacks that are not reflection attacks. Reflection attacks ignore these settings. DDoS Packets(4) controls the minimum number of packets each source must have sent to be registered as a DDoS attack. The sensitivity can be reduced by increasing this setting to six or higher. DDoS Unique Hosts controls the threshold for the minimum number of hosts that have sent flows that match the other characteristics required to trigger the alarm. |
| Impact > Network Denial of Service | Denial of Service | Plixer Flow-Pro Defender | Plixer One Enterprise | A known threat vector has been observed that indicated a DoS attempt has been successful |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|------------------------------------|--|--------------------------|-----------------------|---|
| Impact > Network Denial of Service | DR-DoS | Scrutinizer | Plixer One Core | Identifies Distributed Reflection Denial of Service (DRDoS) attacks targeted at your protected network space. DRDoS attacks are often launched by a BotNet, and ‘reflection attacks’ have become the most common form of DoS attack. Scrutinizer may identify attacks against your network as ‘reflection attacks’ if they meet the criteria. DRDoS attacks are detected by an imbalance in the number of queries sent to external UDP services often used for DRDoS attacks and the number of replies observed. If the number of replies exceeds the number of requests by the threshold, then a DRDoS alarm is triggered. |
| Impact > Network Denial of Service | Packet Flood | Scrutinizer | Plixer One Enterprise | Alerts when a packet flood is detected. A packet flood is characterized as a large volume of small sized packets intended to overwhelm the target’s ability to process legitimate traffic. |
| Impact > Network Denial of Service | Ping Flood | Scrutinizer | Plixer One Enterprise | Alerts when a ping flood is detected. A ping flood is characterized as a large volume of ICMP Echo requests intended to overwhelm the target’s ability to process legitimate traffic. |
| Impact > Resource Hijacking | Crypto Currency Mining Activity Detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects cryptocurrency mining activities (e.g. traffic to known mining pools) |
| Impact > Resource Hijacking | ML Engine coin miner alert | Plixer Machine Learning | Plixer One Enterprise | Detect traffic signatures that are similar to those associated with XMRig coin miner |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--------------------------------------|------------------------------------|--------------------------|-----------------------|--|
| Indicators of Compromise | Bogon Attempt | Scrutinizer | Plixer One Enterprise | Alerts if traffic to or from unallocated public IP space is detected |
| Indicators of Compromise | Bogon Connection | Scrutinizer | Plixer One Enterprise | Alerts if traffic to or from unallocated public IP space is detected |
| Indicators of Compromise | Denied Flows Firewall | Scrutinizer | Plixer One Core | Triggers an alarm for internal IP addresses sending to external IP addresses that cause greater than the threshold of denied flows. The default threshold is set to 5 denied flows. Either the source or destination IP address can be excluded from triggering this alarm. |
| Indicators of Compromise | P2P Detection | Scrutinizer | Plixer One Core | Peer to Peer (P2P) traffic such as BitTorrent are identified by this algorithm. The default threshold is a P2P session involving over 100 external hosts, which will detect most P2P applications. However, there are several P2P applications that are stealthier, so you may want to experiment with lower thresholds or periodically lower the threshold to about 20 to determine if other 'low and slow' P2P traffic is on your network. |
| Initial Access > Drive-by Compromise | Possibly Unwanted Program Detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects potentially unwanted programs (e.g. various spyware applications) |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|--|--------------------------|-----------------------|---|
| Initial Access > Exploit Public-Facing Application | Access to a potentially vulnerable web application | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects when there is access to a potentially vulnerable web application (e.g. an apache ?M=D directory list attempt) |
| Initial Access > Exploit Public-Facing Application | Web Application Attack | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects when a possible web application attack occurs (e.g. a SQL injection attack on a web application or shellcode found in URI) |
| Initial Access > Phishing | Targeted Malicious Activity was Detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Fires when targeted malicious activity is detected (e.g. Advanced Persistent Threats (APTs) that try to remain undetected on a network) |
| Initial Access > User Execution | A Network Trojan was detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects known network Trojans. Plixer default rules contain over 10,000 different trojan detections out of the box |
| Initial Access > User Execution | Possible Social Engineering Attempted | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects possible social engineering attempts (e.g. a phishing email, fake tech support landing pages, etc.) |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---|---|--------------------------|-----------------------|--|
| Initial Access > Valid Accounts | An attempted login using a suspicious username was detected | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects a suspicious network login, such as a TELNET root login |
| Initial Access > Valid Accounts | Attempted User Privilege Gain | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects attempts to gain user-level privileges (e.g. a non-admin user trying to gain admin privileges) |
| Initial Access > Valid Accounts | Attempt to login by a default username and password | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects attempts to login to services using known default credentials (e.g. login attempts with username admin and password admin) |
| Lateral Movement > Exploitation of Remote Services | Lateral Movement | Scrutinizer | Plixer One Enterprise | Identifies successful lateral movement. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|---------------------------------------|---|--------------------------|-----------------------|---|
| Lateral Movement > Remote Services | Decode of an RPC Query | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects decoded Remote Procedure Call (RPC) portmap activity |
| ML Engine Malware Detection | ML Engine malware alert | Plixer Machine Learning | Plixer One Enterprise | Detect traffic signatures that are similar to those of well known malware |
| Privilege Escalation > Valid Accounts | Attempted Denial of Service | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects attempts to make a machine or network resource unavailable (e.g. a sudden surge in traffic from various sources) |
| Privilege Escalation > Valid Accounts | Successful Administrator Privilege Gain | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects when administrator-level access has been successfully gained. For example, a new user created with admin privileges |
| Privilege Escalation > Valid Accounts | Successful User Privilege Gain | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects when user-level privileges have been successfully gained (e.g. Metasploit Meterpreter activity detected) |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|---|--------------------------|-----------------------|---|
| Privilege Escalation > Valid Accounts | Unsuccessful User Privilege Gain | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects when an attempt to gain user-level privileges is unsuccessful (e.g. RPC rlogin login failure) |
| Reconnaissance > Active Scanning > IPs | Attempted Information Leak | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects attempts to gain unauthorized access to information (e.g. a request for a list of all users or data) |
| Reconnaissance > Active Scanning > IPs | ICMP Destination Unreachable (External) | Scrutinizer | Plixer One Core | This alarm is generated when a large number of ICMP destination unreachable messages have been sent to the suspect IP address. This may happen as a result of scanning activity, misconfiguration, or network errors. ICMP Destination Unreachable is a message that comes back from a destination host or the destination host gateway to indicate that the destination is unreachable for one reason or another. The default threshold is 100 destination unreachable messages. Either the source or destination IP address can be excluded from triggering this alarm. |
| Reconnaissance > Active Scanning > IPs | Ping Scan (External) | Scrutinizer | Plixer One Enterprise | Alerts when a host is suspected of performing a ping scan. A ping scan uses ICMP Echo Requests (ping) to discover what IPs are in use on a network. The behavior is commonly demonstrated by attackers attempting to find targets for compromise or lateral movement. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|----------------------------------|--------------------------|-----------------------|--|
| Reconnaissance > Active Scanning > IPs | SYN IP Scan (External) | Scrutinizer | Plixer One Core | Alerts when a SYN scan is detected. SYN scans are a TCP scan with the TCP SYN Flag set. This scan is often used as reconnaissance prior to an attack as it is fast and somewhat stealthy. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Reconnaissance > Active Scanning > Ports | FIN Scan (External) | Scrutinizer | Plixer One Core | Alerts when a FIN scan is detected. FIN scans are often used as reconnaissance prior to an attack. They are considered to be a 'stealthy scan' as they may be able to pass through firewalls, allowing an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Reconnaissance > Active Scanning > Ports | ICMP Port Unreachable (External) | Scrutinizer | Plixer One Core | This alarm is generated when a large number of ICMP destination unreachable messages have been sent to the suspect IP address. This may happen as a result of scanning activity, misconfiguration, or network errors. ICMP Destination Unreachable is a message that comes back from a destination host or the destination host gateway to indicate that the destination is unreachable for one reason or another. The default threshold is 100 destination unreachable messages. Either the source or destination IP address can be excluded from triggering this alarm. |
| Reconnaissance > Active Scanning > Ports | Information Leak | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects when a limited information leak has occurred (e.g. possible Ipconfig information was detected in an HTTP response) |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|------------------------------|--------------------------|-----------------------|--|
| Reconnaissance > Active Scanning > Ports | Large Scale Information Leak | Plixer Flow-Pro Defender | Plixer One Enterprise | Detects when a large scale information leak has occurred (e.g. a full Wordpress DB has been exported as XML) |
| Reconnaissance > Active Scanning > Ports | NULL Scan (External) | Scrutinizer | Plixer One Core | Alerts when a NULL scan is detected. NULL scans are a TCP scan with all TCP Flags cleared to zero. This scan is often used as reconnaissance prior to an attack. They are considered to be a 'stealthy scan' as they may be able to pass through firewalls, allowing an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Reconnaissance > Active Scanning > Ports | Odd TCP Flags (External) | Scrutinizer | Plixer One Core | Alerts when a scan is detected using unusual TCP Flag combinations. These types of scans may allow an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|------------------------------|-------------|-----------------------|---|
| Reconnaissance > Active Scanning > Ports | RST/ACK Detection (External) | Scrutinizer | Plixer One Core | Alerts when a large number of TCP flows containing only RST and ACK flags have been detected being sent to a single destination. These flows indicate that a connection attempt was made on the host sending the RST/ACK flow, and was rejected. This algorithm may detect other scan types used by an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Reconnaissance > Active Scanning > Ports | Slow Port Scan (External) | Scrutinizer | Plixer One Enterprise | Detects when a large number of ports have been probed on the target machine over a long period of time. This alert could indicate malicious activity or reconnaissance for lateral movement. |
| Reconnaissance > Active Scanning > Ports | SYN Port Scan (External) | Scrutinizer | Plixer One Core | Alerts when a SYN scan is detected. SYN scans are a TCP scan with the TCP SYN Flag set. This scan is often used as reconnaissance prior to an attack as it is fast and somewhat stealthy. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Reconnaissance > Active Scanning > Ports | TCP Half-Open (External) | Scrutinizer | Plixer One Core | Alerts when a SYN scan is detected. SYN scans are a TCP scan with the TCP SYN Flag set. This scan is often used as reconnaissance prior to an attack as it is fast and somewhat stealthy. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|----------------------|-------------|-----------------|--|
| Reconnaissance > Active Scanning > Ports | TCP Scan (External) | Scrutinizer | Plixer One Core | Alerts when a possible TCP scan is detected from an exporter that does not provide TCP Flag information. These types of scans may allow an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |
| Reconnaissance > Active Scanning > Ports | UDP Scan (External) | Scrutinizer | Plixer One Core | Alerts when a possible UDP scan is detected. These types of scans may allow an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. NOTE: if your policy allows P2P traffic on your network, then you will likely want to exclude the allowed host(s) or disable this alarm as it will often detect P2P control traffic as a UDP Scan violation. |
| Reconnaissance > Active Scanning > Ports | Xmas Scan (External) | Scrutinizer | Plixer One Core | Alerts when a XMAS scan is detected. XMAS scans are a TCP scan with the FIN, PSH, and URG TCP Flags set. This scan is often used as reconnaissance prior to an attack. They are considered to be a ‘stealthy scan’ as they may be able to pass through firewalls, allowing an attacker to identify additional information about hosts on your network. The default threshold is 100 unique scan flows in three minutes. Internal IP addresses that are allowed to scan your internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|--|-------------------------|-----------------------|---|
| Resource Development > Acquire Infrastructure > DNS Server | Rogue DNS Service | Plixer Machine Learning | Plixer One Enterprise | Find rogue DNS services that may not be known or desired on a network |
| Resource Development > Compromise Accounts | Azure user logged on from many hosts | Plixer Machine Learning | Plixer One Enterprise | Authentications from more hosts than normal in the past 30 minutes |
| Resource Development > Compromise Accounts | Azure user logged on from many locations | Plixer Machine Learning | Plixer One Enterprise | More locations authenticated from in the past 30 minutes than normal |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|--------------------------------------|-------------------------|-----------------------|---|
| Resource Development > Compromise Accounts | Azure user logged on many times | Plixer Machine Learning | Plixer One Enterprise | More authentications than normal in the past 30 minutes |
| Resource Development > Compromise Accounts | New user using elevated logon | Plixer Machine Learning | Plixer One Enterprise | New LDAP user logging in with elevated privileges |
| Resource Development > Compromise Accounts | Office 365 user logged in many times | Plixer Machine Learning | Plixer One Enterprise | More authentications than normal in the past 30 minutes |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|--|-------------------------|-----------------------|---|
| Resource Development > Compromise Accounts | Office 365 user logged on from many hosts | Plixer Machine Learning | Plixer One Enterprise | Authentications from more hosts than normal in the past 30 minutes |
| Resource Development > Compromise Accounts | Office 365 users logged on from many locations | Plixer Machine Learning | Plixer One Enterprise | More locations authenticated from in the past 30 minutes than normal |
| Resource Development > Compromise Accounts | Privileged user logged on from many hosts | Plixer Machine Learning | Plixer One Enterprise | LDAP Authentications from more hosts than normal in the past 30 minutes |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|--|--------------------------------------|-------------------------|-----------------------|---|
| Resource Development > Compromise Accounts | Privileged user logged on many times | Plixer Machine Learning | Plixer One Enterprise | More LDAP authentications than normal in the past 30 minutes |
| Security Events | Auto Investigate | Scrutinizer | Plixer One Core | This algorithm correlates potential sequences of events into overall security incidents using the event policy classes, targets, and violators. |
| System | Access and Audit Events | Plixer One System | Plixer One Core | All user access and activity can be logged and reviewed |
| System | Bad Exporter Flow | Plixer One System | Plixer One Core | An exporter sent a flow record with invalid values |
| System | Bad Exporter Packet | Plixer One System | Plixer One Core | An exporter sent a packet with invalid values |
| System | Bad Exporter Template | Plixer One System | Plixer One Core | An exporter sent a template with invalid values |
| System | Collector Alert | Plixer One System | Plixer One Core | Warnings about collector status |
| System | Collector Message | Plixer One System | Plixer One Core | Informational messages from collectors |

continues on next page

Table 1 – continued from previous page

| Cate- gory | Policy | Tech- nol- ogy | Li- cense | Description |
|-----------------------|-----------------------------------|-------------------------------|-----------------------|---|
| System | Con- figu- ration Alert | Plixer One System | Plixer One Core | Warnings about Scrutinizer configuration |
| System | Cstore Strays | Plixer One System | Plixer One Core | Scrutinizer has detected orphaned history files |
| System | Diskspace Alert | Plixer One System | Plixer One Core | Scrutinizer is running low on disk space |
| System | Event Queue Alert | Scruti- nizer | Plixer One Core | Event Queue Alert |
| System | Ex- porter Ig- nored | Plixer One System | Plixer One Core | Flows were received from an exporter that is not enabled for collection |
| System | Ex- porter Paused | Plixer One System | Plixer One Core | Exporter has been paused due to Low Resources |
| System | Ex- porter Re- sumed | Plixer One System | Plixer One Core | Exporter has been resumed after Low Resources |
| System | Fea- ture Set Paused | Plixer One System | Plixer One Core | Feature Set has been paused due to Low Resources |
| System | Fea- ture Set Re- sumed | Plixer One System | Plixer One Core | Feature Set has been resumed after Low Resources |
| System | Flow Col- lection Paused | Plixer One System | Plixer One Core | Flow Collection Paused due to Low Resources |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|-----------------|------------------------------------|-------------------|-----------------------|--|
| System | Flow Collection Re-summed | Plixer One System | Plixer One Core | Flow Collection Resumed |
| System | Flow Inactivity | Plixer One System | Plixer One Core | Flow Inactivity alarms when flows have not been seen in 30 minutes. |
| System | Flow Rate Limit Changed | Plixer One System | Plixer One Core | Flow Rate Limit Changed |
| System | Flows Limited - Licensing | Plixer One System | Plixer One Core | Flows were limited due to licensing restrictions |
| System | Hardware Resources Exceeded | Plixer One System | Plixer One Core | Hardware Resources Exceeded |
| System | Heartbeat Alert | Plixer One System | Plixer One Enterprise | Warnings about API or DB heartbeats in a distributed environment |
| System | Host Index Disk Availability Error | Plixer One System | Plixer One Core | Disk space allocated to host indexing is full and indexing has been paused. Manage Host Index disk allocation under Admin > Alarm Monitor > Flow Analytics Configuration > Host Indexing |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|----------|-------------------------------|-------------------------|-----------------------|--|
| System | Host Index Disk Space Error | Plixer One System | Plixer One Core | Disk space allocated to host indexing is full and indexing has been paused. Manage Host Index disk allocation under Admin > Alarm Monitor > Flow Analytics Configuration > Host Indexing |
| System | Host Index Disk Space Warning | Plixer One System | Plixer One Core | Disk space allocated to host indexing is close to full. Manage Host Index disk allocation under Admin > Alarm Monitor > Flow Analytics Configuration > Host Indexing |
| System | Kafka Lag | Plixer Machine Learning | Plixer One Enterprise | ML data stream processing has fallen behind |
| System | ML Engine alert | Plixer Machine Learning | Plixer One Enterprise | The ML Engine has reached its maximum number of models it can process. Increase pod's maximum in Admin -> Settings -> ML Data Limits. |
| System | ML Engine Down | Plixer Machine Learning | Plixer One Enterprise | The ML Engine is not responding to heartbeat status checks |
| System | ML models still building | Plixer Machine Learning | Plixer One Enterprise | The ML Engine needs to start building models for the current schedule, but the last schedule isn't finished yet. The replica count config values should be increased. |
| System | ML Service Alert | Plixer Machine Learning | Plixer One Enterprise | The ML Engine has found some required services to not be available |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|----------|------------------------|-------------------------|-----------------------|--|
| System | Run-time Over-run | Plixer One System | Plixer One Core | A scheduled task is taking longer than the time allotted |
| System | Scheduled Task Error | Plixer One System | Plixer One Core | An error occurred while processing a scheduled task |
| System | Setup Problem | Plixer Machine Learning | Plixer One Enterprise | An issue was detected during the setup process |
| System | Stream Deactivated | Plixer Machine Learning | Plixer One Enterprise | Stream has been deactivated |
| System | Stream Reactivated | Plixer Machine Learning | Plixer One Enterprise | Stream has been reactivated |
| System | System Capacity | Plixer Machine Learning | Plixer One Enterprise | The ML Engine is low on resources |
| System | TLS Certificate Expiry | Plixer One System | Plixer One Core | Alert when a TLS certificate is about to expire |
| System | Token Expiration | Plixer One System | Plixer One Core | An API token has expired |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|------------|-------------------------------|-------------------------|-----------------------|---|
| Thresholds | Forecast Task Complete | Plixer Machine Learning | Plixer One Enterprise | Flows were limited due to licensing restrictions |
| Thresholds | Forecast Task Error | Plixer Machine Learning | Plixer One Enterprise | An error occurred while processing a network forecast |
| Thresholds | Forecast Task Starting | Plixer Machine Learning | Plixer One Enterprise | Forecasting backend status update |
| Thresholds | Interface Threshold Violation | Scrutinizer | Plixer One Core | Alerts when an interface exceeds a utilization threshold |
| Thresholds | IP Address Violations | Scrutinizer | Plixer One Core | This algorithm compares the traffic from included exporters against a list of allowed subnets. If both source and destination addresses are outside of an allowed subnet, an alarm will be triggered. A common use of this algorithm is to identify unknown or unauthorized internal network addresses that are communicating with the public Internet. |
| Thresholds | Medianet Jitter Violations | Scrutinizer | Plixer One Core | This algorithm compares the jitter values as reported by the Medianet flows to the threshold defined by the user in the Settings section of this algorithm. The default threshold is 80ms. |
| Thresholds | Report Threshold Violation | Scrutinizer | Plixer One Core | Alerts when a saved report exceeds its configured threshold |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|----------------------------|--------------------------------------|-------------------------|-----------------------|---|
| Unexpected Network Traffic | Plixer Network Intelligence Anomaly | Plixer Machine Learning | Plixer One Enterprise | Anomalous behavior detected by Plixer Network Intelligence |
| Unexpected Network Traffic | Plixer Security Intelligence Anomaly | Plixer Machine Learning | Plixer One Enterprise | Anomalous behavior detected by Plixer Security Intelligence |
| Unexpected Network Traffic | Source Equals Destination | Scrutinizer | Plixer One Enterprise | Alerts when traffic is observed that has the same source and destination addresses. This alarm commonly occurs due to misconfigurations within a network, but may also indicate possible malicious activity. |
| Unexpected Network Traffic | Suspicious Host Communication | Plixer Machine Learning | Plixer One Enterprise | This alert signifies that the mlEngine's deep learning model, which applies GraphSAGE on network data, detected an anomalous communication pattern between host X and host Y over the specified protocol. It flags deviations from established network behaviors, such as unusual traffic volumes or protocol activities, which may indicate potential security threats like malware operations or unauthorized access. This detection is based on analyzing the relationships and data flow patterns in the network to identify outliers that could compromise security. |
| Unexpected Network Traffic | Suspicious Host Communication | Plixer Machine Learning | Plixer One Enterprise | This alert signifies that the mlEngine's deep learning model, which applies GraphSAGE on network data, detected an anomalous communication pattern between host X and host Y over the specified protocol. It flags deviations from established network behaviors, such as unusual traffic volumes or protocol activities, which may indicate potential security threats like malware operations or unauthorized access. This detection is based on analyzing the relationships and data flow patterns in the network to identify outliers that could compromise security. |

continues on next page

Table 1 – continued from previous page

| Category | Policy | Technology | License | Description |
|----------------------------|---------------------|-------------|-----------------|--|
| Unexpected Network Traffic | Unapproved Protocol | Scrutinizer | Plixer One Core | An unapproved protocol was detected in Netflow traffic |

Event details

The table below lists the default timeout settings and details reported for alarm policy violations in Plixer Scrutinizer.

| Name | Criteria | Alarm Keys | Timeout (s) | Message |
|--|--------------------|----------------------------------|-------------|---|
| Access and Audit Events | violators, message | violators, message | 300 | %{VIOLATORS} %{MESSAGE} |
| Access to a potentially vulnerable web application | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| A client was using an unusual port | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--|-----------|----------------------------------|--------------|--|
| An attempted login using a suspicious user-name was detected | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| A Network Trojan was detected | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| A suspicious file-name was detected | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| A system call was detected | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| Attempted Denial of Service | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|---|----------------|--|--------------|--|
| At-tempted Information Leak | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| At-tempted User Privilege Gain | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| At-tempt to login by a default user-name and password | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| Auto Investigate | first_violator | violators, targets, host_count, policy_count, chain_count, event_count, start_epoch, end_epoch | 86400 | The host %{FIRST_VIOLATOR} was seen in %{CHAIN_COUNT} event chains involving %{POLICY_COUNT} policies, %{HOST_COUNT} directly involved hosts, and %{EVENT_COUNT} events. |
| Azure user logged on from many hosts | user_id | user_id, total_hosts | 300 | In the last 30 minutes, %{USER_ID} has attempted to authenticate from %{TOTAL_HOSTS} hosts, which is more hosts than normal. Hosts performing authentication(s) are %{VIOLATORS} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--|------------------------|---|--------------|--|
| Azure user logged on from many locations | user_id | user_id, total_locations | 300 | In the last 30 minutes, %{USER_ID} has attempted to authenticate from %{TOTAL_LOCATIONS} different locations, which is more than normal. Locations performing authentication(s) are %{VIOLATORS} |
| Azure user logged on many times | user_id | user_id, total_auths | 300 | In the last 30 minutes, %{USER_ID} has attempted %{TOTAL_AUTHS} authentications, which is more authentications than normal. Hosts performing authentication(s) are %{VIOLATORS} |
| Bad Exporter Flow | violators, reason_text | reason_text, reason_num, repetition, sequence, set_id, source_id, violators, devices | 3600 | Exporter %{VIOLATORS} sent a bad flow (source %{SOURCE_ID}, sequence %{SEQUENCE}, set %{SET_ID}): %{REASON_TEXT} |
| Bad Exporter Packet | violators, reason_text | reason_text, reason_num, repetition, violators, devices | 3600 | Exporter %{VIOLATORS} sent a bad packet: %{REASON_TEXT} |
| Bad Exporter Template | violators, reason_text | reason_text, reason_num, repetition, sequence, source_id, template_id, violators, devices | 3600 | Exporter %{VIOLATORS} sent a bad template #%{TEMPLATE_ID} (source %{SOURCE_ID}, sequence %{SEQUENCE}): %{REASON_TEXT} |
| Blocked Malicious Domains | violators | violators, targets, domain | 300 | %{VIOLATORS} is accessing blocked domain %{DOMAIN} |
| Bogon Attempt | violators | violators, targets, devices | 3600 | Connections to a bogon network, %{TARGETS}, were seen on %{DEVICES} by %{VIOLATORS} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|-----------------------------------|-----------------------|---|--------------|---|
| Bogon Connection | violators | violators, targets, devices | 3600 | Inbound traffic from a bogon network was seen going to %{TARGETS} on %{DEVICES} by %{VIOLATORS} |
| BotNet Detection | violators | violators, targets, devices, nxcount | 3600 | Internal IP %{VIOLATORS} performed %{NXCOUNT} unique DNS lookups using DNS server(s) %{TARGETS} that returned a No Existing Domain (NXDOMAIN) message as seen on %{DEVICES} exporter(s). This may indicate the presence of malware on %{VIOLATORS} that uses a domain generation algorithm (DGA) to communicate with malware C&C servers. |
| Breach Attempt Detection | violators, breachtype | devices, violators, breachtype, targets | 900 | Detected %{BREACHTYPE} breach by: %{VIOLATORS} with targets: %{TARGETS} |
| Brute-force RDP (Client-side) | violators | violators, targets | 300 | %{VIOLATORS} is attempting a RDP brute force attack on %{TARGETS} |
| Brute-force RDP (Server-side TCP) | targets | violators, targets | 300 | %{TARGETS} is receiving a RDP (tcp) brute force attack from %{VIOLATORS} |
| Brute-force RDP (Server-side UDP) | targets | violators, targets | 300 | %{TARGETS} is receiving a RDP (udp) brute force attack from %{VIOLATORS} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--|----------------------|---|--------------|--|
| Brute-force SSH (Client-side) | violators | violators, targets | 300 | %{VIOLATORS} is attempting a SSH client brute force attack on %{TARGETS} |
| Brute-force SSH (Server-side) | targets | violators, targets | 300 | %{TARGETS} is receiving a SSH server brute force attack from %{VIOLATORS} |
| Collector Alert | error | process, process_id, devices, violators, error | 300 | %{PROCESS}(%{PROCESS_ID}) %{DEVICES} encountered %{ERROR} on %{VIOLATORS} |
| Collector Message | event_type, priority | process, process_id, message, event_type, violators | 300 | %{PROCESS}(%{PROCESS_ID}) on %{VIOLATORS} reported %{EVENT_TYPE}: %{MESSAGE} |
| Configuration Alert | event_type, priority | process, process_id, message, event_type, violators | 300 | %{PROCESS}(%{PROCESS_ID}) reported %{EVENT_TYPE} by %{VIOLATORS}: %{MESSAGE} |
| Crypto Currency Mining Activity Detected | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Cstore Strays | devices | count | 86400 | Found and removed: %{COUNT} stray cstore files on: %{DEVICES} |
| Data Accumulation | violators | violators, targets, total_data | 300 | In the last 30 minutes, %{VIOLATORS} accumulated %{TOTAL_DATA} bytes from %{TARGETS} |
| Data Exfiltration | violators | violators, targets, total_data | 300 | In the last 30 minutes, %{VIOLATORS} exfiltrated %{TOTAL_DATA} bytes to %{TARGETS} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|---|-----------|--|--------------|---|
| DDoS | targets | attacker_count, bytes_std_dev, duration, flow_count, packets_std_dev | 300 | Possible Inbound DDoS Attack: Within %{DURATION} seconds %{ATTACKER_COUNT} external hosts generated a combined total of %{FLOW_COUNT} flows having bytes within %{BYTES_STD_DEV} standard deviations and packets within %{PACKETS_STD_DEV} standard deviations. |
| De-code of an RPC Query | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Denial of Service | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Denied Flows Firewall | violators | devices, violators, target_count, flowcount | 900 | IP %{VIOLATORS} had %{FLOW-COUNT} connection attempts to %{TARGET_COUNT} external IP addresses denied by the firewall as seen on %{DEVICES} exporter(s) |
| Detection of a Denial of Service Attack | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Detection of a Network Scan | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--|---------------------------------------|--|--------------|---|
| Detection of a non-standard protocol or event | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| Device Retrieving External IP Address Detected | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| Disk Space Alert | disk_error, disk_partition, violators | process, process_id, disk_error, disk_partition, message | 300 | %{PROCESS}(%{PROCESS_ID}) The disk partition “%{DISK_PARTITION}” is “%{DISK_ERROR}”. %{MESSAGE} |
| DNS Command and Control Detection | violators | violators, targets, devices | 900 | Possible Command and Control (C&C) Activity. DNS TXT messages are being exchanged between asset %{VIOLATORS} and %{TARGETS} as seen on the %{DEVICES} exporter(s) |
| DNS Data Leak Detection | violators | violators, totaltextlength, dnsname | 900 | DNS lookups initiated from asset: %{VIOLATORS} using complex domain name: %{DNSNAME} containing a high number of domain levels and a total of: %{TOTALTEXTLENGTH} characters. |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--------------------------------------|--------------------|--|--------------|--|
| DNS Hits | violators | violators, flowcount, threshold | 900 | Internal IP %{VIOLATORS} performed %{FLOWCOUNT} DNS lookups in the last 5 minutes exceeding the treshold of %{THRESHOLD} |
| DNS Server Detection | violators | violators, client_count, flowcount, devices | 900 | %{CLIENT_COUNT} IP address(es) initiated %{FLOWCOUNT} DNS lookups to IP address %{VIOLATORS} as seen on %{DEVICES} exporter(s) |
| Domain Observed Used for C2 Detected | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Domain Reputation | violators, dnsname | violators, dnsname, category | 900 | IP %{VIOLATORS} performed a DNS lookup on a black-listed domain: %{DNSNAME} in the %{CATEGORY} category |
| DR-DoS | targets, port_name | devices, attacker_count, duration, packet_in_count, packet_io_ratio, packet_out_count, port, port_name | 900 | Possible Inbound DRDoS Attack from common port %{PORT} (%{PORT_NAME}): Within %{DURATION} seconds %{ATTACKER_COUNT} violators generated a combined total of %{PACKET_IN_COUNT} inbound packets in response to %{PACKET_OUT_COUNT} outbound request packets, for a ratio of %{PACKET_IO_RATIO} inbound packets per outbound packet. |
| Encrypted traffic alert | violators | violators, ja3, ja3s, reason, severity | 300 | ML generated an encrypted traffic alert for %{VIOLATORS}: %{REASON} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|-------------------------------|--------------------------------|---|--------------|---|
| End-point Analytics Info | violators | violators, macaddress, risk_score, location | 300 | Host %\{VIOLATORS\} has MAC address %\{MACADDRESS\}, has a risk score of %\{RISK_SCORE\}, and has location %\{LOCATION\}. |
| Event Queue Alert | violators, type | threshold, value | 300 | Event queue on host: %\{VIOLATORS\} has breached %\{TYPE\} threshold: %\{THRESHOLD\} with value: %\{VALUE\} |
| Executable code was detected | violators | violators, targets, devices, msg | 900 | %\{DEVICES\} observed %\{MSG\} from %\{VIOLATORS\} targeting %\{TARGETS\} |
| Exploit Kit Activity Detected | violators | violators, targets, devices, msg | 900 | %\{DEVICES\} observed %\{MSG\} from %\{VIOLATORS\} targeting %\{TARGETS\} |
| Exporter Ignored | devices, violators, reason_num | reason_text, repetition, violators | 3600 | Discarding flows from exporter %\{VIOLATORS\}: %\{REASON_TEXT\} |
| Exporter Paused | violators, exporter_id | | 1 | Exporter %\{EXPORTER_ID\} paused on collector %\{VIOLATORS\} due to insufficient resources. See the feature sizing interface for more details. |
| Exporter Resumed | violators, exporter_id | | 1 | Exporter %\{EXPORTER_ID\} resumed on collector %\{VIOLATORS\} due to additional available resources. See the feature sizing interface for more details. |
| Feature Set Paused | violators, feature_set | | 1 | Feature set %\{FEATURE_SET\} paused on collector %\{VIOLATORS\} due to insufficient resources. See the feature sizing interface for more details. |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|---------------------------|------------------------|---|--------------|--|
| Feature Set Re-summed | violators, feature_set | | 1 | Feature set <code>%{FEATURE_SET}</code> resumed on collector <code>%{VIOLATORS}</code> due to additional available resources. See the feature sizing interface for more details. |
| FIN Scan (External) | violators | devices, violators | 900 | A FIN Scan was seen on <code>%{DEVICES}</code> by <code>%{VIOLATORS}</code> |
| FIN Scan (Internal) | violators | devices, violators | 900 | A FIN Scan was seen on <code>%{DEVICES}</code> by <code>%{VIOLATORS}</code> |
| Flow Collection Paused | violators | | 60 | Flow collection paused on collector <code>%{VIOLATORS}</code> due to hardware and/or configuration change. See the feature sizing interface for more details. |
| Flow Collection Re-summed | violators | new_flow_rate | 60 | Flow collection resumed at <code>%{NEW_FLOW_RATE}</code> flows/sec on collector <code>%{VIOLATORS}</code> . |
| Flow Inactivity | violators, collector | last_flow | 1200 | Exporter <code>%{VIOLATORS}</code> stopped sending flows to the <code>%{COLLECTOR}</code> collector. The last flow was received <code>%{LAST_FLOW}</code> . If this is expected, set the exporter to disabled or delete it in manage exporters to stop these alarms. |
| Flow-Pro Event Capture | devices, capture_name | violators, targets, devices, capture_name, lookup | 900 | Traffic captured for <code>%{CAPTURE_NAME}</code> from <code>%{VIOLATORS}</code> to <code>%{TARGETS}</code> seen on <code>%{DEVICES}</code> |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|---------------------------|--|---|--------------|---|
| Flow Rate Limit Changed | violators | new_flow_rate | 60 | Flow collection rate limit changed to %{NEW_FLOW_RATE} flows/sec on collector %{VIOLATORS} due to hardware and/or configuration change. See the feature sizing interface for more details. |
| Flows Limited - Licensing | devices, violators, reason_num | reason_text | 60 | Collector %{VIOLATORS} license exceeded: %{REASON_TEXT} |
| Forecast Anomaly | devices, interfaces, applications, types | forecast_id, devices, interfaces, target_quantity, observed_value, mean, forecast_start_time, forecast_end_time | 300 | Forecast: %{FORECAST_ID} found %{INTERFACES} on %{DEVICES} observed value: %{OBSERVED_VALUE} %{TARGET_QUANTITY} is outside forecast for interval %{FORECAST_START_TIME}-%{FORECAST_END_TIME}, Expected Value: %{LOWER_CONF} <= %{MEAN} <= %{UPPER_CONF} |
| Forecast Task Complete | devices, interfaces, applications, type | forecast_id | 60 | Forecast: %{FORECAST_ID} complete, results available |
| Forecast Task Error | devices, interfaces, applications, type | forecast_id, error_stage, error | 60 | Forecast: %{FORECAST_ID} resulted in an error during %{ERROR_STAGE}. Message: %{ERROR} |
| Forecast Task Starting | devices, interfaces, applications, type | forecast_id | 60 | Forecast: %{FORECAST_ID} received by forecasting module |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|------------------------------------|---------------------------|---|--------------|--|
| Generic Protocol Command Decode | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Hardware Resources Exceeded | violators | drop_rate, flow_limit_period | 60 | Collector %{VIOLATORS} incoming flow rate exceeds hardware recommendations. %{DROP_RATE} flows per second dropped over the last %{FLOW_LIMIT_PERIOD} seconds. See the feature sizing interface for more details. |
| Heartbeat Alert | heartbeat_type, violators | process, process_id, heartbeat_type, devices, violators | 300 | %{PROCESS}(%{PROCESS_ID}) %{HEARTBEAT_TYPE} heartbeat failed from %{DEVICES} to %{VIOLATORS} |
| Host Index Disk Availability Error | violators | threshold, current | 300 | Host Indexing service has reached disk storage volume limit of %{THRESHOLD} percent in use, Currently %{CURRENT} percent in use. Stopping processing and starting garbage collection until under threshold. |
| Host Index Disk Space Error | violators | threshold, current | 300 | Host Indexing service has reached disk space usage: %{CURRENT}MB, threshold: %{THRESHOLD}MB. Stopping processing and starting garbage collection until under threshold. |
| Host Index Disk Space Warning | violators | threshold, current | 300 | Host Indexing service has reached disk space usage: %{CURRENT}MB, over 75% of threshold: %{THRESHOLD}MB |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|---|--------------------|--|--------------|---|
| Host Reputation | violators, targets | violators, targets, devices, category_note | 3600 | IP %{VIOLATORS} sent traffic to a suspect %{CATEGORY_NOTE} at IP address %{TARGETS} as seen on the %{DEVICES} exporter(s) |
| Host Watchlist | violators | devices, violators, port, protocol | 900 | Host Watchlist - %{DEVICES} saw watchlisted host %{VIOLATORS} communicating from %{PROTOCOL} %{PORT} |
| ICMP Destination Unreachable (External) | violators | flowcount, violators | 900 | External IP %{VIOLATORS} triggered %{FLOWCOUNT} ICMP Destination Unreachable flows within 5 minutes |
| ICMP Destination Unreachable (Internal) | violators | flowcount, violators | 900 | Internal IP %{VIOLATORS} triggered %{FLOWCOUNT} ICMP Destination Unreachable flows within 5 minutes |
| ICMP Port Unreachable (External) | violators | flowcount, violators | 900 | External IP %{VIOLATORS} triggered %{FLOWCOUNT} ICMP Protocol Unreachable flows within 5 minutes |
| ICMP Port Unreachable (Internal) | violators | flowcount, violators | 900 | Internal IP %{VIOLATORS} triggered %{FLOWCOUNT} ICMP Protocol Unreachable flows within 5 minutes |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|-------------------------------|-------------------------------------|--|--------------|---|
| Information Leak | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Interface Threshold Violation | violators, interface_name, instance | exporter, interface_name, instance, threshold, violation, graphStart, graphEnd | 900 | Interface %{EXPORTER}: %{INTERFACE_NAME} exceeded the threshold of %{THRESHOLD} %{VIOLATION} |
| IP Address Violations | violators | devices, violators, targets | 900 | Traffic on %{DEVICES} between %{VIOLATORS} and %{TARGETS} is outside of allowed subnets |
| Kafka Lag | topic_lagged | topic_lagged, messages_behind | 660 | ML Kafka topic %{TOPIC_LAGGED} is lagging %{MESSAGES_BEHIND} messages behind |
| Large Ping | violators | violators, targets, devices, threshold, avg_ping_size | 900 | Unexpected ICMP Echo traffic seen from violator %{VIOLATORS} to target %{TARGETS} on exporter %{DEVICES} with an average packet size of %{AVG_PING_SIZE} Bytes which violates the threshold of %{THRESHOLD} Bytes |
| Large Scale Information Leak | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Lateral Movement | violators, targets, worm_type | devices, targets, violators | 1200 | %{WORM_TYPE} lateral movement detected on %{DEVICES}, from %{VIOLATORS} to %{TARGETS} |
| Lateral Movement Attempt | violators, worm_type | devices, violators, targets, worm_type, dst_port | 1200 | %{WORM_TYPE} lateral movement attempt detected on %{DEVICES} from %{VIOLATORS} to %{TARGETS} over port %{DST_PORT} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|---|-------------------|---|--------------|---|
| Lateral Movement Behavior | violators | violators | 300 | %{VIOLATORS} is exhibiting lateral movement behavior |
| Malware Command and Control Activity Detected | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Medianet Jitter Violations | violators | targets, violators, jitter | 420 | Jitter values of %{JITTER}ms between %{VIOLATORS} and %{TARGETS} exceeds threshold |
| ML Engine alert | violators, source | source, threshold | 300 | ML service %{SOURCE} has reached threshold %{THRESHOLD}, throttling until next run |
| ML Engine coin miner alert | violators | violators, family, probability, threshold | 300 | ML detected %{VIOLATORS} generating malicious traffic related to %{FAMILY} malware family (%{PROBABILITY}% match, threshold set to %{THRESHOLD}%) |
| ML Engine command and control alert | violators | violators, family, probability, threshold | 300 | ML detected %{VIOLATORS} generating malicious traffic related to %{FAMILY} malware family (%{PROBABILITY}% match, threshold set to %{THRESHOLD}%) |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--------------------------------------|-------------------|---|--------------|--|
| ML Engine Down | host | host, violators | 300 | ML Engine %{HOST} is not responding to pings |
| ML Engine exploit kit alert | violators | violators, family, probability, threshold | 300 | ML detected %{VIOLATORS} generating malicious traffic related to %{FAMILY} malware family (%{PROBABILITY}% match, threshold set to %{THRESHOLD}%) |
| ML Engine malware alert | violators | violators, family, probability, threshold | 300 | ML detected %{VIOLATORS} generating malicious traffic related to %{FAMILY} malware family (%{PROBABILITY}% match, threshold set to %{THRESHOLD}%) |
| ML Engine remote access trojan alert | violators | violators, family, probability, threshold | 300 | ML detected %{VIOLATORS} generating malicious traffic related to %{FAMILY} malware family (%{PROBABILITY}% match, threshold set to %{THRESHOLD}%) |
| ML models still building | violators | violators, schedule | 300 | ML is still building models for schedule %{SCHEDULE}, but the next schedule is currently expected to start. Increase replica count values in the config. |
| ML Service Alert | service_name | service_name, unavailable, expected | 300 | ML service %{SERVICE_NAME} has %{UNAVAILABLE}/%{EXPECTED} instances unavailable |
| Net-Flow Domain Reputation | violators, domain | violators, domain, category | 900 | Internal IP %{VIOLATORS} performed a lookup of %{DOMAIN}, categorized as %{CATEGORY} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--------------------------------------|-----------|--|--------------|---|
| New user using elevated logon | user_id | user_id | 300 | A new user, %{USER_ID}, is logging in with elevated privileges. Hosts performing login(s) are %{VIOLATORS} |
| NULL Scan (External) | violators | devices, violators, flowcount, threshold | 900 | A NULL scan was seen on %{DEVICES} by %{VIOLATORS} in %{FLOWCOUNT} flows violating the threshold of %{THRESHOLD} |
| NULL Scan (Internal) | violators | devices, violators, flowcount, threshold | 900 | A NULL scan was seen on %{DEVICES} by %{VIOLATORS} in %{FLOWCOUNT} flows violating the threshold of %{THRESHOLD} |
| Odd TCP Flags (External) | violators | devices, violators, flags, flowcount | 900 | Odd TCP flags (%{FLAGS}) were seen in %{FLOWCOUNT} flows on %{DEVICES} by %{VIOLATORS} |
| Odd TCP Flags (Internal) | violators | devices, violators, flags, flowcount | 900 | Odd TCP flags (%{FLAGS}) were seen in %{FLOWCOUNT} flows on %{DEVICES} by %{VIOLATORS} |
| Office 365 user logged in many times | user_id | user_id, total_auths | 300 | In the last 30 minutes, %{USER_ID} has attempted %{TOTAL_AUTHS} authentications, which is more authentications than normal. Hosts performing authentication(s) are %{VIOLATORS} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--|-----------|--|--------------|--|
| Office 365 user logged on from many hosts | user_id | user_id, total_hosts | 300 | In the last 30 minutes, %{USER_ID} has attempted to authenticate from %{TOTAL_HOSTS} hosts, which is more hosts than normal. Hosts performing authentication(s) are %{VIOLATORS} |
| Office 365 users logged on from many locations | user_id | user_id, total_locations | 300 | In the last 30 minutes, %{USER_ID} has attempted to authenticate from %{TOTAL_LOCATIONS} different locations, which is more than normal. Locations performing authentication(s) are %{VIOLATORS} |
| P2P Detection | violators | devices, violators, dst_host_count, dst_port_count | 900 | P2P traffic to %{DST_HOST_COUNT} destinations using %{DST_PORT_COUNT} distinct port(s) was seen on %{DEVICES} from %{VIOLATORS} |
| Packet Flood | violators | devices, violators, targets, count | 3600 | Packet flood seen from %{VIOLATORS} to %{TARGETS} comprising of %{COUNT} small packets in a minute by devices: %{DEVICES} |
| Ping Flood | violators | devices, violators, targets, count | 3600 | Ping flood seen from %{VIOLATORS} to %{TARGETS} comprising of %{COUNT} pings in a minute by devices: %{DEVICES} |
| Ping Scan (External) | violators | devices, violators, count | 3600 | Ping scan seen from %{VIOLATORS} to %{COUNT} hosts by devices: %{DEVICES} |
| Ping Scan (Internal) | violators | devices, violators, count | 3600 | Ping scan seen from %{VIOLATORS} to %{COUNT} hosts by devices: %{DEVICES} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|---|---------------------------------------|---------------------------------------|--------------|--|
| Plixer Network Intelligence Anomaly | violators, interface_id, anomaly_type | violators, interface_id, anomaly_type | 300 | Exporter %{VIOLATORS} is generating anomalous %{ANOMALY_TYPE} traffic on interface %{INTERFACE_ID} |
| Plixer Security Intelligence Anomaly | violators, anomaly_type | violators, anomaly_type | 300 | %{VIOLATORS} is generating anomalous %{ANOMALY_TYPE} traffic |
| Possible Social Engineering Attempted | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Possibly Unwanted Program Detected | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Privileged user logged on from many hosts | user_id | user_id, total_hosts | 300 | In the last 30 minutes, %{USER_ID} has attempted to authenticate from %{TOTAL_HOSTS} hosts, which is more hosts than normal. Hosts performing authentication(s) are %{VIOLATORS} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--------------------------------------|------------------------------|---|--------------|---|
| Privileged user logged on many times | user_id | user_id, total_auths | 300 | In the last 30 minutes, %{USER_ID} has attempted %{TOTAL_AUTHS} authentications, which is more authentications than normal. Hosts performing authentication(s) are %{VIOLATORS} |
| Protocol Misdirection | violators | violators, traffic_type, port, targets | 3600 | Mismatched traffic type of %{TRAFFIC_TYPE} to port %{PORT} from %{VIOLATORS} to %{TARGETS} |
| Ransomware Behavior | violators | violators, targets, file_count, files | 900 | Observed a possible ransomware encryption attack from %{VIOLATORS} targeting SMB share %{TARGETS}. %{FILE_COUNT} files were both read and written to, including files: %{FILES} |
| Report Threshold Violation | saved_report, row_identifier | saved_report, row_identifier, violation, graphStart, graphEnd, src_port, dst_port, violator, violator_username, target, target_username, protocol, app_proto, url | 420 | The report %{SAVED_REPORT} %{ROW_IDENTIFIER} has exceeded its threshold %{VIOLATION} |
| Reverse SSH Shell | violators | origin_bytes, bytes_per_packet | 3600 | Possible reverse SSH tunnel from %{VIOLATORS} to %{TARGETS} seen by devices: %{DEVICES} based on %{ORIGIN_BYTES} origin bytes and %{BYTES_PER_PACKET} average origin bytes per packet |
| Rogue DHCP Service | violators | violators, targets | 300 | %{VIOLATORS} is hosting a rogue DHCP service contacted by %{TARGETS}. If this is expected behavior, please add the DHCP server IP address to the DHCP Servers IP group |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|------------------------------|----------------------|--|--------------|---|
| Rogue DNS Service | violators | violators, targets | 300 | %{VIOLATORS} is hosting a rogue DNS service contacted by %{TARGETS}. If this is expected behavior, please add the DNS server IP address to the DNS Servers IP group |
| Rogue LDAP Service | violators | violators, targets | 300 | %{VIOLATORS} is hosting a rogue LDAP service contacted by %{TARGETS}. If this is expected behavior, please add the LDAP server IP address to the LDAP Servers IP group |
| RST/ACK Detection (External) | violators | violators, flowcount, targets | 900 | Anomalous Behavior - Possible - RST/ACK Replies Observed Host %{TARGETS} received %{FLOW-COUNT} packets from %{VIOLATORS} without observing any other flags |
| RST/ACK Detection (Internal) | violators | violators, flowcount, targets | 900 | Anomalous Behavior - Possible - RST/ACK Replies Observed Host %{TARGETS} received %{FLOW-COUNT} packets from %{VIOLATORS} without observing any other flags |
| Run-time Over-run | process | process, process_id, threshold, duration, action | 300 | %{PROCESS}(%{PROCESS_ID}) ran for %{DURATION} seconds and exceeded the configured run-time of %{THRESHOLD} seconds (%{ACTION}) |
| Scheduled Task Error | violators, task_name | task_id, command, error_code, start_time, run_time | 300 | A scheduled task on collector %{VIOLATORS}, %{TASK_NAME} (ID %{TASK_ID}) returned error code: %{ERROR_CODE} running: “%{COMMAND}”. It started at %{START_TIME} AND ran for %{RUN_TIME} seconds. View the collector log and/or run the task manually for more details. |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|---------------------------|-----------|--|--------------|---|
| Setup Problem | issue | message | 900 | %{MESSAGE} |
| SI-GRRed Exploit Attempt | violators | violators, targets | 300 | %{VIOLATORS} is targeting a SI-GRRed attack on %{TARGETS} |
| Slow Port Scan (External) | violators | devices, violators, targets | 3600 | %{VIOLATORS} is port scanning %{TARGETS} on %{DEVICES} |
| Slow Port Scan (Internal) | violators | devices, violators, targets | 3600 | %{VIOLATORS} is port scanning %{TARGETS} on %{DEVICES} |
| SMB Brute-force Attempt | violators | violators, targets, failed_logins, usernames | 900 | Observed a possible SMB brute force attack from %{VIOLATORS} targeting SMB share %{TARGETS}. %{FAILED_LOGINS} failed logins observed including usernames: %{USERNAMES} |
| Source Equals Destination | violators | devices, violators | 900 | Traffic with source and destination of %{VIOLATORS} was seen on %{DEVICES} |
| Stream Deactivated | stream | size, threshold | 900 | The stream: %{STREAM} has breached its configured threshold: %{THRESHOLD} with total size: %{SIZE} and has been deactivated. |
| Stream Reactivated | stream | minutes, size, threshold | 900 | The stream: %{STREAM} with total size: %{SIZE} below its configured threshold: %{THRESHOLD} has been reactivated after having been deactivated for: %{MINUTES} minutes. |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|---|-----------|-----------------------------------|--------------|---|
| Successful Administrator Privilege Gain | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| Successful Credential Theft Detected | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| Successful User Privilege Gain | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed from %{VIOLATORS} targeting %{TARGETS} |
| Suspicious Host Communication | violators | violators, targets, protocol_name | 300 | Based on how these hosts and those around them normally communicate, the communication between %{VIOLATORS} and the host(s) %{TARGETS} on protocol %{PROTOCOL_NAME} is unexpected. Use the explore event traffic link to view these communications in detail. |
| Suspicious Host Communication | violators | violators, targets, protocol | 300 | Based on how these hosts and those around them normally communicate, the communication between %{VIOLATORS} and the host(s) %{TARGETS} on protocol %{PROTOCOL} is unexpected. Use the explore event traffic link to view these communications in detail. |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--------------------------|------------|---|--------------|--|
| SYN IP Scan (External) | violators | devices, violators, targets, scanned_host_count, scanned_port_count, host_thresh, port_thresh | 900 | A SYN IP Scan by %{VIOLATORS} seen scanning %{SCANNED_HOST_COUNT} hosts which exceeds the threshold of %{HOST_THRESH} and %{SCANNED_PORT_COUNT} ports per host exceeding the threshold of %{PORT_THRESH} |
| SYN IP Scan (Internal) | violators | devices, violators, targets, scanned_host_count, scanned_port_count, host_thresh, port_thresh | 900 | A SYN IP Scan by %{VIOLATORS} seen scanning %{SCANNED_HOST_COUNT} hosts which exceeds the threshold of %{HOST_THRESH} and %{SCANNED_PORT_COUNT} ports per host exceeding the threshold of %{PORT_THRESH} |
| SYN Port Scan (External) | violators | devices, violators, targets, scanned_host_count, scanned_port_count, host_thresh, port_thresh | 900 | A SYN Port Scan by %{VIOLATORS} seen scanning %{SCANNED_HOST_COUNT} hosts which exceeds the threshold of %{HOST_THRESH} and %{SCANNED_PORT_COUNT} ports per host exceeding the threshold of %{PORT_THRESH} |
| SYN Port Scan (Internal) | violators | devices, violators, targets, scanned_host_count, scanned_port_count, host_thresh, port_thresh | 900 | A SYN Port Scan by %{VIOLATORS} seen scanning %{SCANNED_HOST_COUNT} hosts which exceeds the threshold of %{HOST_THRESH} and %{SCANNED_PORT_COUNT} ports per host exceeding the threshold of %{PORT_THRESH} |
| System Capacity | vital_type | vital_type, value | 300 | ML is using %{VALUE} percent of its %{VITAL_TYPE} capacity |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--|-----------|--|--------------|--|
| Targeted Malicious Activity was Detected | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| TCP Half-Open (External) | violators | devices, violators, targets, packets_per_port, scanned_port_count, pkt_thresh, port_thresh | 900 | A possible SYN Half Open Attack by %{VIOLATORS} seen targeting %{TARGETS}. Port count of %{SCANNED_PORT_COUNT} exceeded the threshold of %{PORT_THRESH} and flows per port of %{PACKETS_PER_PORT} exceed the threshold of %{PKT_THRESH}. |
| TCP Half-Open (Internal) | violators | devices, violators, targets, packets_per_port, scanned_port_count, pkt_thresh, port_thresh | 900 | A possible SYN Half Open Attack by %{VIOLATORS} seen targeting %{TARGETS}. Port count of %{SCANNED_PORT_COUNT} exceeded the threshold of %{PORT_THRESH} and flows per port of %{PACKETS_PER_PORT} exceed the threshold of %{PKT_THRESH}. |
| TCP Scan (External) | violators | devices, violators, port_count, dst_count | 900 | A TCP Scan was seen on %{DEVICES} by %{VIOLATORS} scanning %{DST_COUNT} IPs and %{PORT_COUNT} ports |
| TCP Scan (Internal) | violators | devices, violators, port_count, dst_count | 900 | A TCP Scan was seen on %{DEVICES} by %{VIOLATORS} scanning %{DST_COUNT} IPs and %{PORT_COUNT} ports |
| TLS Certificate Expiry | violators | days | 86400 | TLS certificates on nodes: %{VIOLATORS} will expire in %{DAYS} days. Contact Plexier Support or see <i>scrut_util -help certs</i> . |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|--------------------------------------|----------------------|---------------------------------|--------------|--|
| Token Expiration | username, expires_on | username, expires_on, status | 86400 | An authentication token for %{USERNAME} %{STATUS} on %{EXPIRES_ON} |
| Tunneling through external DNS host | violators | violators, targets, tunnel_type | 300 | %{VIOLATORS} is tunneling external DNS traffic through %{TARGETS} |
| Tunneling through external ICMP host | violators | violators, targets, tunnel_type | 300 | %{VIOLATORS} is tunneling external ICMP traffic through %{TARGETS} |
| Tunneling through external SSH host | violators | violators, targets, tunnel_type | 300 | %{VIOLATORS} is tunneling external SSH traffic through %{TARGETS} |
| Tunneling through internal DNS host | violators | violators, targets, tunnel_type | 300 | %{VIOLATORS} is tunneling internal DNS traffic through %{TARGETS} |
| Tunneling through internal ICMP host | violators | violators, targets, tunnel_type | 300 | %{VIOLATORS} is tunneling internal ICMP traffic through %{TARGETS} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|-------------------------------------|-----------|---|--------------|---|
| Tunneling through internal SSH host | violators | violators, targets, tunnel_type | 300 | %{VIOLATORS} is tunneling internal SSH traffic through %{TARGETS} |
| UDP Scan (External) | violators | devices, violators, dst_count, port_count | 900 | A UDP Scan was seen on %{DEVICES} by %{VIOLATORS} scanning %{DST_COUNT} IPs and %{PORT_COUNT} ports |
| UDP Scan (Internal) | violators | devices, violators, dst_count, port_count | 900 | A UDP Scan was seen on %{DEVICES} by %{VIOLATORS} scanning %{DST_COUNT} IPs and %{PORT_COUNT} ports |
| Unapproved Protocol | protocol | protocol_name, devices | 900 | Unapproved network transport: %{PROTOCOL_NAME} was seen on: %{DEVICES} |
| Unsuccessful User Privilege Gain | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Web Application Attack | violators | violators, targets, devices, msg | 900 | %{DEVICES} observed %{MSG} from %{VIOLATORS} targeting %{TARGETS} |
| Worm Activity | violators | violators | 300 | %{VIOLATORS} is exhibiting worm behavior |
| Xmas Scan (External) | violators | devices, violators | 900 | An Xmas Scan was seen on %{DEVICES} by %{VIOLATORS} |

continues on next page

Table 2 – continued from previous page

| Name | Criteria | Alarm Keys | Time out (s) | Message |
|----------------------|-----------|--------------------|--------------|--|
| Xmas Scan (Internal) | violators | devices, violators | 900 | An Xmas Scan was seen on %{DEVICES} by %{VIOLATORS} |
| Zerologon | violators | violators, targets | 300 | %{VIOLATORS} is targeting a Zerologon attack on %{TARGETS} |

7.1.2 Flow analytics algorithms

Refer to the following tables for general information, recommended applications, and configuration options for FA algorithms.

FA algorithm list

The table below contains general information and recommended applications for all flow analytics algorithms available in Plixer Scrutinizer.

| Algorithm | Function |
|-----------------------------------|--|
| Bogon Traffic | Alerts if traffic to or from an unallocated public IP space is detected |
| BotNet Detection | Alerts when a large number of unique DNS name lookups have failed |
| Breach Attempt Detection | Alerts when flow behaviors that may indicate a brute force password attack on |
| DDoS Detection | Alerts when a Distributed Denial of Service (DDoS) attack targeting the protec |
| Denied Flows Firewall | Alerts when the number of denied flows from an internal to an external IP addr |
| DNS Command and Control Detection | Alerts when the volume or size of DNS TXT messages at the network perimet |
| DNS Data Leak Detection | Alerts when the volume or size of messages with suspicious DNS names exce |
| DNS Hits | Alerts when a host initiates an excessive number of DNS queries |
| DNS Server Detection | Alerts when a new DNS is detected based on packet exchanges between clients |
| Domain Reputation | Alerts when traffic associated with a suspicious domain (based on a list mainta |
| DRDoS Detection | Alerts when a Distributed Reflection Denial of Service attack targeting the pro |
| FIN Scan | Alerts when a FIN scan is detected |
| Flow Reports Thresholds | Alerts when a <i>custom threshold configured for a saved report</i> is exceeded |
| Host Indexing | Monitors traffic to maintain an index of hosts seen on the network that includes |
| Host Reputation | Monitors traffic to maintain a list of active, non-whitelisted Tor nodes |

| Algorithm | Function |
|------------------------------|--|
| Host Watchlist | Alerts when a host violating a user-defined IP address blacklist is detected |
| ICMP Destination Unreachable | Alerts when a large number of <i>ICMP Destination Unreachable</i> messages are sent |
| ICMP Port Unreachable | Alerts when a large number of <i>ICMP Port Unreachable</i> messages are sent |
| Incident Correlation | Alerts when multiple Indicator of Compromise (IOC) events for a single host are detected |
| IP Address Violations | Alerts when a flow containing a non-authorized IP address as the source or destination is observed |
| JA3 Fingerprinting | Alerts when software sending suspicious encrypted traffic based on TLS handshake |
| Large Ping | Alerts when an unusually large ICMP Echo Request (ping) is observed |
| Lateral Movement | Alerts when successful lateral movement is observed |
| Lateral Movement Attempt | Alerts when behavior that may indicate attempted lateral movement is observed |
| Medianet Jitter Violations | Alerts when jitter values reported by a Medianet flow exceed the configured threshold |
| Multicast Violations | Alerts when multicast traffic volume exceeds the configured threshold |
| NetFlow Domain Reputation | Alerts when a DNS lookup from a blacklisted IP is reported via NetFlow |
| Network Transports | Alerts when traffic over unapproved transport protocols is observed |
| NULL Scan | Alerts when a NULL scan is detected |
| Odd TCP Flags Scan | Alerts when a scan using unusual TCP flag combinations is detected |
| P2P Detection | Alerts when a P2P session with a host count exceeding the configured threshold is observed |
| Packet Flood | Alerts when a packet flood is detected |
| Persistent Flow Risk | Alerts when a persistent flow is detected |
| Persistent Flow Risk - ASA | Alerts when a persistent flow matching a specified 5-tuple is detected |
| Ping Flood | Alerts when a ping flood is detected |
| Ping Scan | Alerts when a host suspected of performing a ping scan is observed |
| Protocol Misdirection | Alerts when traffic not matching the port being used is detected |
| Reverse SSH Shell | Alerts when potential reverse SSH tunnels to external destinations are detected |
| RST/ACK Detection | Alerts when the system observes a large number of TCP flows containing a RST/ACK flag |
| Slow Port Scan | Alerts when the system observes a large number of ports on the same host being scanned |
| Source Equals Destination | Alerts when traffic with the same host and destination is observed |
| SYN Scan | Alerts when a SYN scan is detected |
| TCP Scan | Alerts when a potential TCP scan is detected from an Exporter that does not support it |
| Top Applications | Monitors application traffic |
| Top Autonomous Systems | Monitors traffic to and from autonomous systems |
| Top Countries | Monitors traffic by country |
| Top Hosts | Monitors traffic by host |
| Top IP groups | Monitors traffic by IP group |
| UDP Scan | Alerts when a potential UDP scan is detected |
| XMAS Scan | Alerts when a XMAS scan is detected |

Algorithm settings

The table below lists the additional settings that can be used to tune individual FA algorithm behavior.

Table 4: Algorithm Settings

| Algorithm Name | Setting | Description |
|-----------------------------------|----------------------------------|--|
| Auto Investigate | Candidate Limit | The maximum number of Violator->Policy->Target links to review for correlation. |
| Auto Investigate | Chain Max | The maximum number of Violator->Policy->Target chains that will be considered for deduplication. |
| Auto Investigate | Length Limit | The maximum length of any chain of Violator->Policy->Target links. |
| BotNet Detection | Threshold | Number of unique No Existing Domain (NXDOMAIN) replies within a three-minute period to trigger alarm |
| DDoS Detection | DDoS Bytes Deviation | Maximum number of bytes allowed in a single standard deviation to trigger (default 10) |
| DDoS Detection | DDoS Packet Deviation | Maximum number of packets allowed in a single standard deviation to trigger (default 10) |
| DDoS Detection | DDoS Packets | Number of packets each source must have sent to be counted |
| DDoS Detection | DDoS Unique hosts | Minimum number of unique hosts participating in a DDoS attack |
| Denied Flows Firewall | Denied Threshold | The number of denied flows from a single host within a three-minute period to trigger an event |
| DNS Command and Control Detection | DNS Command and Control attempts | DNS Command and Control attempts within a three-minute period to trigger alarm |
| DNS Command and Control Detection | DNS Command and Control bytes | DNS Command and Control bytes within a three-minute period to trigger alarm |
| DNS Data Leak Detection | DNS Data Leak attempts | DNS Data Leak attempts within a three-minute period to trigger alarm |
| DNS Data Leak Detection | DNS Data Leak bytes | DNS Data Leak bytes within a three-minute period to trigger alarm |
| DNS Hits | Flow Threshold | The number of DNS requests within a three-minute period to trigger an event |
| DNS Server Detection | Flow threshold to trigger alarm | Number of properly formatted DNS request packets sent to the specified IP address to trigger alarm |
| DRDoS Detection | CharGen (UDP 19) | Enable/Disable Distributed Reflection DoS (DRDoS) Attack Detection |

continues on next page

Table 4 – continued from previous page

| Algorithm Name | Setting | Description |
|-----------------|---|---|
| DRDoS Detection | DNS (UDP 53) | Enable/Disable Distributed Reflection DoS (DR-DoS) Attack Detection |
| DRDoS Detection | Flow Imbalance Threshold | How many inbound packets per outbound packet to trigger a DRDoS alarm |
| DRDoS Detection | LDAP (UDP 389) | Enable/Disable Distributed Reflection DoS (DR-DoS) Attack Detection |
| DRDoS Detection | Memcached (UDP 11211) | Enable/Disable Distributed Reflection DoS (DR-DoS) Attack Detection |
| DRDoS Detection | NetBIOS Name Server (UDP 137) | Enable/Disable Distributed Reflection DoS (DR-DoS) Attack Detection |
| DRDoS Detection | NTP (UDP 123) | Enable/Disable Distributed Reflection DoS (DR-DoS) Attack Detection |
| DRDoS Detection | Quote of the Day (UDP 17) | Enable/Disable Distributed Reflection DoS (DR-DoS) Attack Detection |
| DRDoS Detection | RPC Portmap (UDP 111) | Enable/Disable Distributed Reflection DoS (DR-DoS) Attack Detection |
| DRDoS Detection | Sentinel (UDP 5093) | Enable/Disable Distributed Reflection DoS (DR-DoS) Attack Detection |
| DRDoS Detection | SNMP (UDP 161,162) | Enable/Disable Distributed Reflection DoS (DR-DoS) Attack Detection |
| DRDoS Detection | SSDP (UDP 1900) | Enable/Disable Distributed Reflection DoS (DR-DoS) Attack Detection |
| DRDoS Detection | Trivial File Transfer Protocol (UDP 69) | Enable/Disable Distributed Reflection DoS (DR-DoS) Attack Detection |
| FIN Scan | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| FIN Scan | Flow Threshold | The number of FIN flows from a single host within a three-minute period to trigger an event |
| FIN Scan | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| FIN Scan | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |
| Host Indexing | Days of host index data retention | The host index entries last seen more than this many days ago will be trimmed. |
| Host Indexing | Host Index Database | File path of Host Index. *Background service must be restart from CLI after update. Service will start clean in new location. |
| Host Indexing | Host Indexing Domain Socket | File path of Host Indexing Domain Socket |

continues on next page

Table 4 – continued from previous page

| Algorithm Name | Setting | Description |
|------------------------------|----------------------------------|---|
| Host Indexing | Host Index Max Disk Space | Maximum combined disk space threshold for host indexing (in MB). Warning events sent at 75%, indexing temporarily suspended at 100% until record expiration frees space. |
| Host Indexing | Host Index Sync Interval Minutes | The sync interval in minutes for each index update |
| Host Indexing | Host to Host Database | File path of Host-to-Host Index. Leave blank to disable Host-to-Host indexing. *Background service must be restart from CLI after update. Service will start clean in new location. |
| Host Indexing | Window Limit | The maximum number of records considered on each index update |
| Host Reputation | Aggregate Timeout | Aggregate similar alarms until there are no new alarms for over N minutes (default 2 hours = 120 minutes, zero to disable aggregation) |
| Host Reputation | Threshold | Number of bytes (octets) within a three-minute period to trigger alarm |
| ICMP Destination Unreachable | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| ICMP Destination Unreachable | Flow Threshold | The number flows from a single host triggering an ICMP Destination Unreachable reponse within a three-minute period |
| ICMP Destination Unreachable | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| ICMP Destination Unreachable | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |
| ICMP Port Unreachable | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| ICMP Port Unreachable | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| ICMP Port Unreachable | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |
| ICMP Port Unreachable | Threshold | The number flows from a single host triggering an ICMP Port Unreachable reponse within a three-minute period |
| IP Address Violations | Threshold | Number of bytes (octets) within a three-minute period to trigger alarm |
| Large Ping | Size Threshold | Average packet threshold for determining a large ping packet. |

continues on next page

Table 4 – continued from previous page

| Algorithm Name | Setting | Description |
|----------------------------|---------------------------------|---|
| Lateral Movement Attempt | Backdoor Threshold | Number of destination hosts on backdoor ports to trigger alert |
| Lateral Movement Attempt | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| Lateral Movement Attempt | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| Lateral Movement Attempt | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |
| Lateral Movement Attempt | IOT Threshold | Number of destination hosts on IOT ports to trigger alert |
| Lateral Movement Attempt | Remote Access Threshold | Number of destination hosts on remote access ports to trigger alert |
| Lateral Movement Attempt | Windows Remote Access Threshold | Number of destination hosts on Windows remote access ports to trigger alert |
| Medianet Jitter Violations | Jitter by Interface | The millisecond variation in packet delay caused by queuing, contention and/or serialization effects on the path through the network. Default = 80 ms. This is also used for record highlighting in Status reports. |
| Multicast Violations | Threshold | Number of bytes (octets) within a three-minute period to trigger alarm |
| NULL Scan | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| NULL Scan | Flow Threshold | The number of flows from a single host within a three-minute period to trigger an event |
| NULL Scan | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| NULL Scan | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |
| Odd TCP Flags Scan | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| Odd TCP Flags Scan | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| Odd TCP Flags Scan | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |
| Odd TCP Flags Scan | Threshold | The number of flows from a single host with odd TCP flags within a three-minute period to trigger an event |
| P2P Detection | Threshold | Number of distinct destination IPs in a three-minute period to trigger alarm |

continues on next page

Table 4 – continued from previous page

| Algorithm Name | Setting | Description |
|----------------------------|---------------------------------|--|
| Packet Flood | Packet Size Threshold | The Maximum average packet size to be considered a flood packet |
| Packet Flood | Packet threshold | The number of packets that should be observed within a three-minute period to trigger an event |
| Persistent Flow Risk | Active Flow Threshold (hours) | How long should a flow be active before an alarm is triggered |
| Persistent Flow Risk | Aggregate Timeout | Aggregate similar alarms until there are no new alarms for over N minutes (default 2 hours = 120 minutes, zero to disable aggregation) |
| Persistent Flow Risk | Inactive Flow Threshold (hours) | How long should a flow be inactive before it no longer is considered the same flow |
| Persistent Flow Risk | PCR Threshold | The ratio of traffic where 1 is a pure upload and -1 is a pure download. Set to 0 to disable |
| Persistent Flow Risk - ASA | Active Flow Threshold (hours) | How long should a flow be active before an alarm is triggered |
| Persistent Flow Risk - ASA | Aggregate Timeout | Aggregate similar alarms until there are no new alarms for over N minutes (default 2 hours = 120 minutes, zero to disable aggregation) |
| Persistent Flow Risk - ASA | Inactive Flow Threshold (hours) | How long should a flow be inactive before it no longer is considered the same flow |
| Persistent Flow Risk - ASA | PCR Threshold | The ratio of traffic where 1 is a pure upload and -1 is a pure download. Set to 0 to disable |
| Ping Flood | Ping Flood Threshold | Minimum number of pings from a host to a distinct destination in a minute that should trigger |
| Ping Scan | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| Ping Scan | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| Ping Scan | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |
| Ping Scan | Ping Scan Host Threshold | Minimum number of distinct hosts that a violator must ping to trigger |
| Reverse SSH Shell | Packet Size Threshold | Maximum average packet size in the SSH session that should be considered for triggering the alert |
| Reverse SSH Shell | Reverse Shell Threshold | The maximum number of outbound bytes on an SSH connection that should be considered for triggering the alert |
| RST/ACK Detection | External to Internal | Enable/Disable Scan Detection in the direction indicated |

continues on next page

Table 4 – continued from previous page

| Algorithm Name | Setting | Description |
|-------------------|----------------------------|---|
| RST/ACK Detection | Flow Threshold | The number of flows from a single host within a three-minute period to trigger an event |
| RST/ACK Detection | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| RST/ACK Detection | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |
| Slow Port Scan | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| Slow Port Scan | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| Slow Port Scan | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |
| SYN Scan | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| SYN Scan | Half-Open packet per port | The number of packets per dst port to be considered a half-open flood |
| SYN Scan | Half-Open port count | The number of distinct destination ports to be considered a half-open flood |
| SYN Scan | Host Scan Hosts | The number of distinct destination hosts to be considered a host scan |
| SYN Scan | Host Scan Ports | The number of distinct destination ports to be considered a host scan |
| SYN Scan | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| SYN Scan | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |
| SYN Scan | Port Scan Hosts | The number of distinct destination hosts to be considered a port scan |
| SYN Scan | Port Scan Ports | The number of distinct destination ports to be considered a port scan |
| TCP Scan | Destination Host Threshold | Number of distinct destination hosts to trigger alarm |
| TCP Scan | Destination Port Threshold | Number of distinct destination ports to trigger alarm |
| TCP Scan | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| TCP Scan | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| TCP Scan | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |

continues on next page

Table 4 – continued from previous page

| Algorithm Name | Setting | Description |
|----------------|----------------------|--|
| UDP Scan | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| UDP Scan | Host threshold | The number of hosts scanned within a three-minute period that will trigger an event |
| UDP Scan | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| UDP Scan | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |
| UDP Scan | Port threshold | The number of ports per host scanned within a three-minute period that will trigger an event |
| XMAS Scan | External to Internal | Enable/Disable Scan Detection in the direction indicated |
| XMAS Scan | Flow Threshold | The number of flows from a single host within a three-minute period to trigger an event |
| XMAS Scan | Internal to External | Enable/Disable Scan Detection in the direction indicated |
| XMAS Scan | Internal to Internal | Enable/Disable Scan Detection in the direction indicated |

7.1.3 User permissions

See below for a full list of features/permission sets and individual permissions that can be granted to users through [user groups](#).

Table 5: Alarms Administrator

| Permission | Description |
|----------------------------------|---|
| Acknowledge Bulletin Board Event | Ability to acknowledge events on Alarms tab bulletin boards |
| Delete Alarms | Permission to permanently delete alarms |

Table 6: Alarms User

| Permission | Description |
|------------|-----------------------|
| Alarms Tab | Access the Alarms tab |

Table 7: Dashboard Administrator

| Permission | Description |
|-------------------------|---|
| Dashboard Administrator | Manage all dashboards created by any user |

Table 8: Dashboard User

| Permission | Description |
|-------------------|---------------------------|
| Create Dashboards | Create new Dashboards |
| Dashboards Tab | Access the Dashboards tab |

Table 9: Maps Administrator

| Permission | Description |
|-------------------------------|---|
| Mapping Groups Configuration | Define and manage device groups for network mapping |
| Mapping Objects Configuration | Define custom map objects and manage object/group object properties |

Table 10: Maps User

| Permission | Description |
|------------|---------------------|
| Maps Tab | Access the Maps tab |

Table 11: Reporting Administrator

| Permission | Description |
|--------------------------------|---|
| Application Groups | Define custom applications using IP address and port rules |
| AS Names Configuration | View autonomous system (AS) numbers/properties |
| Delete Reports | Ability to delete saved reports regardless of owner |
| Host Names Configuration | Define custom hostname-to-IP mappings and static subnet labels for reporting |
| TOS Configuration | Add custom labels for Type of Service (ToS) and Differentiated Services Code Point (DSCP) values in reports |
| Well-known Ports Configuration | Edit WKP Configuration |

Table 12: Reporting Power User

| Permission | Description |
|--------------------------------|---|
| Add/Edit Report Filters | Permission to update the filters used in Status Tab reports |
| Report Designer | Design custom report type configurations |
| Report Folders | Create and manage folders to organize saved reports |
| Save Reports | Ability to name and save flow reports |
| Scheduled Report Administrator | Set up and manage scheduled email report configurations |
| Schedule Emailed Reports | Schedule a saved report to be emailed on a regular basis |

Table 13: Reporting User

| Permission | Description |
|-------------|-----------------------------|
| Run Reports | Ability to run flow reports |
| Status Tab | Access the Status Tab |

Table 14: System Administrator

| Permission | Description |
|--|--|
| Admin | Access Scrutinizer's administrative functions |
| Alarm Notifications | Configure alarm notifications |
| Alarm Settings | Configure global alarm message options and Flow Inactivity and Interface Threshold Violation alarm settings |
| ASA ACL Descriptions | Add/edit ASA firewall credentials for ACL description retrieval |
| Authentication Tokens | Add and manage user authentication tokens |
| Authentication Types | Manage external authentication types |
| AWS Configuration | AWS configuration |
| Change User Passwords | The ability to change the passwords of other users without needing their credentials |
| Collectors | Manage Plexier Scrutinizer collectors and Plexier ML Engines in the environment |
| Configure SMTP server settings for email notifications and reports | Configure the mailserver Scrutinizer will use to send reports and emails |
| Create Users | The ability to create new local Scrutinizer user accounts |
| Data History Settings | Set alarm and flow data history retention durations |
| Delete Users | The ability to delete local Scrutinizer user accounts |
| Enable/disable and configure third-party integrations for Explore > Exporters view | Create, edit, and delete third-party integration links |
| Endpoint Analytics | Configure and enable/disable Plexier Endpoint Analytics integration |
| Enforce Session Timeout | If the system preference for user activity timeout is set, members of user groups with this permission will be timed-out of the UI according to that setting |
| Exporters | Manage and add protocol exclusions to flow-exporting devices in the environment |
| Flow Analytics Configuration | Configure Flow Analytics thresholds and settings |
| Flow Analytics Exclusions | Configure Flow Analytics exclusions |
| Flow Analytics Settings | Configure global settings and enable/disable FlowPro Defender for FA algorithms |
| Flow Log Ingestion | Third-party Flow Log source configuration |

continues on next page

Table 14 – continued from previous page

| Permission | Description |
|------------------------------------|--|
| Google Maps Proxy Server Settings | Configure proxy server settings for Google Maps requests |
| Host Indexing | Host Indexing settings |
| Interface Details Configuration | Edit device interface details |
| IP Groups Configuration | Define rule-based IP range/subnet groups for reporting |
| LDAP Server Configuration | Manage LDAP server configuration used for Scrutinizer authentication |
| MAC Addresses Configuration | Add and manage custom MAC address labels |
| Notification Manager | Create and manage profiles to assign notification actions by alarm policy |
| Plixer Network Intelligence | Add exporters for machine learning monitoring |
| Plixer Replicator | Configure and enable/disable Plixer Replicator integration |
| Plixer Security Intelligence | Configure network for machine learning monitoring |
| Policy Manager | Reconfigure, enable/disable, and assign notification profiles to alarm policies |
| Protocol Exclusions | Define protocol exclusion rules for reporting |
| RADIUS Server Configuration | Manage RADIUS server configuration used for Scrutinizer authentication |
| Reporting Configuration | Customize Plixer Scrutinizer reporting engine functions |
| Scrutinizer Audit Report | View logs of Plixer Scrutinizer user actions |
| Scrutinizer Language Configuration | Create and edit language localization settings |
| Scrutinizer Product Licensing | Add a Plixer Scrutinizer license key and view license details |
| Scrutinizer System Preferences | Configure general Plixer Scrutinizer environment preferences/settings |
| ServiceNow | Configure and manage ServiceNow instances for incident/ticket generation via notifications and collections |
| Single Sign-On Configuration | Add, Delete, and Edit Identity Provider configuration for Scrutinizer's Single Sign-On Integration |
| SNMP Credentials | Manage SNMP credential sets for polling exporters in the environment |
| STIX-TAXII | Add and manage STIX-TAXII threat intelligence feeds |
| Syslog Server Settings | Syslog server configuration |
| TACACS+ Server Configuration | Manage TACACS+ server configuration used for Scrutinizer authentication |
| User Accounts | Manage user accounts and preferences |
| User Groups | Set up local user groups and manage access to features and resources |
| View User Identity Information | View identity and access information relevant to GDPR restrictions |
| Viptela Settings | Viptela Settings |

continues on next page

Table 14 – continued from previous page

| Permission | Description |
|---------------|--|
| Vitals Report | View the Scrutinizer server vitals reports |

7.1.4 Report types

The following table lists all Plexer Scrutinizer *report types* alongside their data aggregation parameters.

Table 15: Amazon AWS

| Report | Description |
|-------------------------------|--|
| Action | A grouping of Action trending Flows, Packets, Bytes. Information Elements: aws_action, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Action with Interface | A grouping of Action, Interface trending Flows, Packets, Bytes. Information Elements: aws_action, aws_interface, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Action with Interface and Dst | A grouping of Destination, Action, Interface trending Flows, Packets, Bytes. Information Elements: destinationipaddress, aws_action, aws_interface, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Action with Interface and Src | A grouping of Source, Action, Interface trending Flows, Packets, Bytes. Information Elements: sourceipaddress, aws_action, aws_interface, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Availability Zones | A grouping of Availability Zone trending Flows, Packets, Bytes. Information Elements: aws_az_id, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Dst Service | A grouping of Destination Service trending Flows, Packets, Bytes. Information Elements: aws_pkt_destination_service, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Interface | A grouping of Interface trending Flows, Packets, Bytes. Information Elements: aws_interface, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Pair Interface | A grouping of Source, Interface, Destination trending Flows, Packets, Bytes. Information Elements: sourceipaddress, aws_interface, destinationipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Pair Interface Action | A grouping of Source, Interface, Action, Destination trending Flows, Packets, Bytes. Information Elements: sourceipaddress, aws_interface, aws_action, destinationipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Src Service | A grouping of Source Service trending |
| 7.1. Appendices | Flows, Packets, Bytes. Information Elements: aws_pkt_source_service, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Src Service-Dst Service | A grouping of Source Service, Destination |

Table 16: AppFlow

| Report | Description |
|----------------------------------|--|
| Application | A grouping of Application trending Count, Packets, Bytes. Information Elements: appflow_applicationid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Application RTT | A grouping of Application, Destination trending Packets, Bytes, RTT. Information Elements: appflow_applicationid, destinationipaddress, octetdeltacount, packetdeltacount, tcprtt. |
| Connections | A grouping of Src Port, Source, Connection, Destination, Dst Port trending RTT, Count, Packets, Bytes. Information Elements: sourcetransportport, sourceipaddress, connectionid, destinationipaddress, destinationtransportport, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount, tcprtt. |
| HTTP Request Cookie | A grouping of Transaction ID, HTTP Request Cookie trending Flow Count, Bytes. Information Elements: transactionid, httprequestcookie, octetdeltacount, plixeraggregatedrecordcount. |
| HTTP Response Length | A grouping of Source, Src Port, Destination, Dst Port trending Count, Avg. Length. Information Elements: sourceipaddress, sourcetransportport, destinationipaddress, destinationtransportport, httpresponselen, plixeraggregatedrecordcount. |
| HTTP Response Time to First Byte | A grouping of Source, Src Port, Destination, Dst Port trending Count, Avg. Time. Information Elements: sourceipaddress, sourcetransportport, destinationipaddress, destinationtransportport, httpresponsetimetofirstbyte, plixeraggregatedrecordcount. |
| HTTP Response Time to Last Byte | A grouping of Source, Src Port, Destination, Dst Port trending Count, Avg. Time. Information Elements: sourceipaddress, sourcetransportport, destinationipaddress, destinationtransportport, httpresponsetimetolastbyte, plixeraggregatedrecordcount. |
| HTTP Status | A grouping of HTTP Status Code, Source, Src Port, Destination, Dst Port trending Count, Bytes. Information Elements: httpresponsestatus, sourceipaddress, sourcetransportport, destinationipaddress, destinationtransportport, octetdeltacount, plixeraggregatedrecordcount. |
| Request Host | A grouping of HTTP Request Host trending Count, Packets, Bytes. Information Elements: httprequesthost, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |

Table 17: Astaro

| Report | Description |
|---------------------------|---|
| afcprotocol Conversations | A grouping of Source, afcprotocol, Destination trending Packets, Bytes. Information Elements: sourceipaddress, afcprotocol, destination-ipaddress, octetdeltacount, packetdeltacount. |
| Top afcprotocol | A grouping of afcprotocol trending Packets, Bytes. Information Elements: afcprotocol, octetdeltacount, packetdeltacount. |

Table 18: Azure

| Report | Description |
|--------------------------------|---|
| Azure NSG All Details | A grouping of Rule Name, Application, Flow Decision, Flow State trending Packets, Bytes, Count. Information Elements: nsg_rulename, applicationid, nsg_flowdecision, nsg_flowstate, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Azure NSG Flow Decisions | A grouping of Flow Decision, Application trending Packets, Bytes, Count. Information Elements: nsg_flowdecision, applicationid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Azure NSG Flow Decisions Count | A grouping of Flow Decision trending Packets, Bytes, Count. Information Elements: nsg_flowdecision, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Azure NSG Flow States | A grouping of Flow State, Application trending Packets, Bytes, Count. Information Elements: nsg_flowstate, applicationid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Azure NSG Flow States Count | A grouping of Flow State trending Packets, Bytes, Count. Information Elements: nsg_flowstate, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Azure NSG Resource IDs | A grouping of Resource ID, Rule Name trending Packets, Bytes, Count. Information Elements: nsg_resourceid, nsg_rulename, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Azure VNET All Details | A grouping of Rule Name, Application, Flow State trending Packets, Bytes, Count. Information Elements: vnet_rulename, applicationid, vnet_flowstate, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Azure VNET Flow States | A grouping of Flow State, Application trending Packets, Bytes, Count. Information Elements: vnet_flowstate, applicationid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Azure VNET Flow States Count | A grouping of Flow State trending Packets, Bytes, Count. Information Elements: vnet_flowstate, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Azure VNET Resource IDs | A grouping of Target Resource ID, Rule Name trending Packets, Bytes, Count. Information Elements: vnet_targetresourceid, vnet_rulename, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |

Table 19: Barracuda

| Report | Description |
|---|---|
| Bind and Conn | A grouping of Bind IP , Bind Port, Conn IP, Conn Port trending Flows, Bytes. Information Elements: bindipv4address, bindtransportport, connipv4address, conntransportport, octetdeltacount, plixeraggregatedrecordcount. |
| FW Rule | A grouping of FW Rule trending Flows, Packets, Bytes. Information Elements: fwrule, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Logop | A grouping of Logop trending Flows, Packets, Bytes. Information Elements: logop, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Pair with Rule and Reason | A grouping of Source IP, Destination IP, FW Rule, Reason trending Flows, Bytes. Information Elements: sourceipaddress, destinationipaddress, fwrule, reasontext, octetdeltacount, plixeraggregatedrecordcount. |
| Pair with Rule, Reason, Service & Traffic | A grouping of Source IP, Destination IP, FW Rule, Reason, Service, Traffic Type trending Flows, Bytes. Information Elements: sourceipaddress, destinationipaddress, fwrule, reasontext, servicename, traffictype, octetdeltacount, plixeraggregatedrecordcount. |
| Reason | A grouping of Reason trending Flows, Packets, Bytes. Information Elements: reasontext, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Rule, Reason, Service, Traffic & Logop | A grouping of FW Rule, Reason, Service, Traffic Type, Logop trending Flows, Bytes. Information Elements: fwrule, reasontext, servicename, traffictype, logop, octetdeltacount, plixeraggregatedrecordcount. |
| Service | A grouping of Service trending Flows, Packets, Bytes. Information Elements: servicename, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Source, Bind, Conn, & Destination | A grouping of Source IP, Bind IP , Conn IP, Destination IP trending Flows, Bytes. Information Elements: sourceipaddress, bindipv4address, connipv4address, destinationipaddress, octetdeltacount, plixeraggregatedrecordcount. |
| 7.1.1. Appendices | A grouping of Traffic Type trending Flows, Packets, Bytes. Information Elements: traffictype, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |

Table 20: Chassis

| Report | Description |
|----------------|--|
| Line Card | A grouping of Line Card trending Pkts, Bytes. Information Elements: linecardid, octetdeltacount, packetdeltacount. |
| Line Card Port | A grouping of Interface, Line Card, Port trending Pkts, Bytes. Information Elements: exportinterface, linecardid, portid, octetdeltacount, packetdeltacount. |

Table 21: Cisco AnyConnect

| Report | Description |
|------------------------|--|
| Dest Host IP & Name | A grouping of Destination, Dst Host Name trending Flows, Bytes In. Information Elements: destinationipaddress, nvzflowdestinationhostname, octetdeltacount, plixeraggregatedrecordcount. |
| DNS suffix | A grouping of DNS Suffix trending Flows, Bytes In. Information Elements: nvzflowdnssuffix, octetdeltacount, plixeraggregatedrecordcount. |
| Loggedin Source | A grouping of Logged User, Source trending Flows, Bytes. Information Elements: nvzflowloggedinuser, sourceipaddress, octetdeltacount, plixeraggregatedrecordcount. |
| Loggedin Source & DNS | A grouping of Logged User, Source, DNS Suffix trending Flows, Bytes. Information Elements: nvzflowloggedinuser, sourceipaddress, nvzflowdnssuffix, octetdeltacount, plixeraggregatedrecordcount. |
| Pair with Host Details | A grouping of Logged User, Source, Destination, Dst Host Name trending Flows, Bytes. Information Elements: nvzflowloggedinuser, sourceipaddress, destinationipaddress, nvzflowdestinationhostname, octetdeltacount, plixeraggregatedrecordcount. |
| Parent Process Details | A grouping of Parent Proc. Acct., Parent Proc. Name, Parent Proc. Hash trending Flows, Bytes. Information Elements: nvzflowparentprocessaccount, nvzflowparentprocessname, nvzflowparentprocesshash, octetdeltacount, plixeraggregatedrecordcount. |
| Process Details | A grouping of Process Name, Process Hash trending Flows, Bytes. Information Elements: nvzflowprocessname, nvzflowprocesshash, octetdeltacount, plixeraggregatedrecordcount. |
| Process to Host | A grouping of Parent Proc. Acct., Destination trending Flows, Bytes In. Information Elements: nvzflowparentprocessaccount, nvzflowdestinationhostname, octetdeltacount, plixeraggregatedrecordcount. |
| Source with Process | A grouping of Source, Logged User, Process Name, System Type trending Flows, Bytes. Information Elements: sourceipaddress, nvzflowloggedinuser, nvzflowprocessname, nvzflowsystemtype, octetdeltacount, plixeraggregatedrecordcount. |
| 7.1. Appendices | |
| Station Name & Dst IP | A grouping of STA Name, Destination trending Flows, Bytes In. Information Elements: sourceipaddress, destinationipaddress, octetdeltacount, plixeraggregatedrecordcount. |

Table 22: Cisco AVC

| Report | Description |
|-----------------------------------|--|
| EzPM: Host Jitter by SSRC (Dst) | A grouping of Destination, DSCP, SSRC trending % Pkt Loss, TEPL, Jitter. Information Elements: destinationipaddress, ipdiffservcodepoint, trans_rtp_ssrc, ciscopktlostpercent, rtp_jitter_mean_sum, trans_pkt_lost_count. |
| EzPM: Host Jitter by SSRC (Src) | A grouping of Source, DSCP, SSRC trending % Pkt Loss, TEPL, Jitter. Information Elements: sourceipaddress, ipdiffservcodepoint, trans_rtp_ssrc, ciscopktlostpercent, rtp_jitter_mean_sum, trans_pkt_lost_count. |
| EzPM: Host Jitter (Dst) | A grouping of Destination, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: destinationipaddress, ipdiffservcodepoint, ciscopktlostpercent, rtp_jitter_mean_sum, trans_pkt_lost_count. |
| EzPM: Host Jitter (Src) | A grouping of Source, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: sourceipaddress, ipdiffservcodepoint, ciscopktlostpercent, rtp_jitter_mean_sum, trans_pkt_lost_count. |
| EzPM: Host to Host Jitter | A grouping of Source, DSCP, Destination trending % Pkt Loss, TEPL, Max Jitter, Jitter. Information Elements: sourceipaddress, ipdiffservcodepoint, destinationipaddress, ciscopktlostpercent, rtp_jitter_mean_sum, trans_pkt_lost_count. |
| EzPM: Host to Host Jitter by SSRC | A grouping of Source, DSCP, Destination, SSRC trending % Pkt Loss, TEPL, Jitter. Information Elements: sourceipaddress, ipdiffservcodepoint, destinationipaddress, trans_rtp_ssrc, ciscopktlostpercent, rtp_jitter_mean_sum, trans_pkt_lost_count. |
| EzPM: Jitter by Interface | A grouping of Exporter, in Int trending % Pkt Loss, Jitter. Information Elements: plixerexporter, ingressinterface, ciscopktlostpercent, rtp_jitter_mean_sum. |
| EzPM: Metadata Jitter | A grouping of Application trending % Pkt Loss, TEPL, Jitter. Information Elements: application-tag, ciscopktlostpercent, rtp_jitter_mean_sum, trans_pkt_lost_count. |

continues on next page

Table 22 – continued from previous page

| Report | Description |
|-------------------------------|---|
| EzPM: Metadata Jitter by DSCP | A grouping of Application, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: applicationtag, ipdiffservcodepoint, ciscopktlostpercent, rtp_jitter_mean_sum, trans_pkt_lost_count. |
| EzPM: Metadata-Session Jitter | A grouping of Application, globalsectionid, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: applicationtag, globalsectionid, ipdiffservcodepoint, ciscopktlostpercent, rtp_jitter_mean_sum, trans_pkt_lost_count. |
| EzPM: Metadata-Version Jitter | A grouping of Application, Application Version Id, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: applicationtag, app_version_id, ipdiffservcodepoint, ciscopktlostpercent, rtp_jitter_mean_sum, trans_pkt_lost_count. |
| Host and URI | A grouping of Server, HTTP Host, URI Stat trending Packets, Bytes. Information Elements: serveripv4address, subapplicationvalue, uristatistic, octetdeltacount, packetdeltacount. |
| Jitter by Interface | A grouping of Exporter, in Int trending % Pkt Loss, Jitter. Information Elements: plixerexporter, ingressinterface, ciscopktlostpercent, trans_rtp_jitter. |
| NBAR: Metadata Jitter | A grouping of Application trending % Pkt Loss, TEPL, Jitter. Information Elements: applicationtag, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| NBAR: Metadata Jitter by DSCP | A grouping of Application, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: applicationtag, ipdiffservcodepoint, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| NBAR: Metadata RTT | A grouping of Application trending RTT. Information Elements: applicationtag, trans_rtt. |
| NBAR: Metadata RTT by DSCP | A grouping of Application, DSCP trending RTT. Information Elements: applicationtag, ipdiffservcodepoint, trans_rtt. |

continues on next page

Table 22 – continued from previous page

| Report | Description |
|---|--|
| NBAR: Metadata-Session Jitter | A grouping of Application, globalseSSID, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: applicationtag, globalseSSID, ipdiffservcodepoint, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| NBAR: Metadata-Session RTT | A grouping of Application, globalseSSID, DSCP trending RTT. Information Elements: applicationtag, globalseSSID, ipdiffservcodepoint, trans_rtt. |
| NBAR: Metadata-Version Jitter | A grouping of Application, Application Version Id, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: applicationtag, app_version_id, ipdiffservcodepoint, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| NBAR: Metadata-Version RTT | A grouping of Application, Application Version Id, DSCP trending RTT. Information Elements: applicationtag, app_version_id, ipdiffservcodepoint, trans_rtt. |
| Percent Re-Transmitted By Interface | A grouping of Exporter, in Int trending Packets, RTX Pkts, AVG % RTX. Information Elements: plixerexporter, ingressinterface, clientretransmissionspackets, packetdeltacount, retransmittedpercent. |
| PfA: Application Avg. Transaction Duration (cs) | A grouping of Application trending Avg Trans. Duration. Information Elements: applicationtag, avc_avgtranstime. |
| PfA: Application Delay | A grouping of Application trending App Delay. Information Elements: applicationtag, avc_avgsvrresptime. |
| PfA: Application Performance | A grouping of Application, DSCP trending Dest Hosts, RTX Pkts, App Delay. Information Elements: applicationtag, ipdiffservcodepoint, avc_avgsvrresptime, clientretransmissionspackets, plixeraggregatedrecordcount. |
| PfA: Connections App Delay | A grouping of Client, Application, Server trending App Delay. Information Elements: sourceipaddress, applicationtag, destinationipaddress, avc_avgsvrresptime. |

continues on next page

Table 22 – continued from previous page

| Report | Description |
|---|--|
| PfA: Connections App Delay (cs) | A grouping of Client, Application, Server trending App Delay. Information Elements: clientip4address, applicationtag, serverip4address, avc_avgsvrresptime. |
| PfA: Connections App Trans. Duration (cs) | A grouping of Client, Application, Server trending Avg Trans. Duration. Information Elements: clientip4address, applicationtag, serverip4address, avc_avgtranstime. |
| PfA: Connections Nwk Delay | A grouping of Client, Application, Server trending Nwk Delay. Information Elements: sourceipaddress, applicationtag, destinationipaddress, avc_avgnwktime. |
| PfA: Connections Nwk Delay (cs) | A grouping of Client, Application, Server trending Nwk Delay. Information Elements: clientip4address, applicationtag, serverip4address, avc_avgnwktime. |
| PfA: Connections Trans. Duration | A grouping of Client, Application, Server trending Avg Trans. Duration. Information Elements: sourceipaddress, applicationtag, destinationipaddress, avc_avgtranstime. |
| PfA: Connections Trans. Duration (cs) | A grouping of Client, Server trending Avg Trans. Duration. Information Elements: clientip4address, serverip4address, avc_avgtranstime. |
| PfA: Conversations | A grouping of Client, Server, App trending Packets, RTX Pkts, % RTX, Clnt Delay, Svr Delay, App Delay, TXN Delay. Information Elements: sourceipaddress, destinationipaddress, applicationtag, avc_avgcltnwktime, avc_avgsvrnwktime, avc_avgsvrresptime, avc_avgtranstime, clientretransmissionspackets, packetdeltacount, retransmittedpercent. |
| PfA: Conversations (cs) | A grouping of Client, Server, App trending Packets, RTX Pkts, % RTX, Clnt Delay, Svr Delay, App Delay, TXN Delay. Information Elements: clientip4address, serverip4address, applicationtag, avc_avgcltnwktime, avc_avgsvrnwktime, avc_avgsvrresptime, avc_avgtranstime, clientretransmissionspackets, packetdeltacount, retransmittedpercent. |

continues on next page

Table 22 – continued from previous page

| Report | Description |
|--|--|
| PfA: Conversations NBAR with QoS | A grouping of Client, Application, Server, QoS Policy trending Packets, Bytes. Information Elements: sourceipaddress, applicationtag, destinationipaddress, policymaphierarchy, octetdeltacount, packetdeltacount. |
| PfA: Conversations NBAR with QoS (cs) | A grouping of Client, Application, Server, QoS Policy trending Packets, Bytes. Information Elements: clientipv4address, applicationtag, serveripv4address, policymaphierarchy, octetdeltacount, packetdeltacount. |
| PfA: Conversation Trans. Duration vs Retransmission | A grouping of Client, Application, Server trending Packets, RTX Pkts, % RTX, Avg Trans. Duration. Information Elements: sourceipaddress, applicationtag, destinationipaddress, avc_avgtranstime, clientretransmissionspackets, packetdeltacount, retransmittedpercent. |
| PfA: Conversation Trans. Duration vs Retransmission (cs) | A grouping of Client, Application, Server trending Packets, RTX Pkts, % RTX, Avg Trans. Duration. Information Elements: clientipv4address, applicationtag, serveripv4address, avc_avgtranstime, clientretransmissionspackets, packetdeltacount, retransmittedpercent. |
| PfA: HTTP Host | A grouping of HTTP Host trending Packets, Bytes. Information Elements: httphostname, octetdeltacount, packetdeltacount. |
| PfA: HTTP Host By Client | A grouping of Client, HTTP Host trending Packets, Bytes. Information Elements: sourceipaddress, httphostname, octetdeltacount, packetdeltacount. |
| PfA: HTTP Host By Client (cs) | A grouping of Client, HTTP Host trending Packets, Bytes. Information Elements: clientipv4address, httphostname, octetdeltacount, packetdeltacount. |
| PfA: HTTP Host Trans. Duration by Host | A grouping of Client, HTTP Host trending Avg Trans. Duration. Information Elements: sourceipaddress, httphostname, avc_avgtranstime. |
| PfA: HTTP Host Trans. Duration by Host (cs) | A grouping of Client, HTTP Host trending Avg Trans. Duration. Information Elements: clientipv4address, httphostname, avc_avgtranstime. |

continues on next page

Table 22 – continued from previous page

| Report | Description |
|--|--|
| PfA: Network Delay by Application (LL) | A grouping of App trending Client Delay, Server Delay, Network Delay. Information Elements: applicationtag, sumlonglivedclientnwkttime, sumlonglivednwkttime, sumlonglivedservernwkttime. |
| PfA: Network Delay by Client (LL) | A grouping of Client trending Client Delay, Server Delay, Network Delay. Information Elements: clientip4address, sumlonglivedclientnwkttime, sumlonglivednwkttime, sumlonglivedservernwkttime. |
| PfA: Network Delay by Server (LL) | A grouping of Server trending Client Delay, Server Delay, Network Delay. Information Elements: serverip4address, sumlonglivedclientnwkttime, sumlonglivednwkttime, sumlonglivedservernwkttime. |
| PfA: Root Cause Delay (cs) | A grouping of Client, App, CBQoS, DSCP, Server trending Packets, RTX Pkts, % RTX, CND, SND, App Delay, Nwk Delay. Information Elements: clientip4address, applicationtag, policymaphierarchy, ipdiffservcodepoint, serverip4address, avc_avgcltnwkttime, avc_avgnwkttime, avc_avgsvrnwkttime, avc_avgsvrresptime, clientretransmissionspackets, packetdeltacount, retransmittedpercent. |
| PfA: Root Cause Delay w/ Users | A grouping of Client, App, CBQoS, DSCP, Server trending Packets, RTX Pkts, % RTX, CND, SND, App Delay, Nwk Delay. Information Elements: sourceipaddress, applicationtag, policymaphierarchy, ipdiffservcodepoint, destinationipaddress, avc_avgcltnwkttime, avc_avgnwkttime, avc_avgsvrnwkttime, avc_avgsvrresptime, clientretransmissionspackets, packetdeltacount, retransmittedpercent. |

continues on next page

Table 22 – continued from previous page

| Report | Description |
|-------------------------------------|---|
| PfA: Root Cause Delay w/ Users (cs) | A grouping of Client, User Name(s), App, CBQoS, DSCP, Server trending Packets, RTX Pkts, % RTX, CND, SND, App Delay, Nwk Delay. Information Elements: clientipaddress, clientipname, applicationtag, policymaphierarchy, ipdiffservcodepoint, serveripaddress, avc_avgcltnwkttime, avc_avgnwkttime, avc_avgsvrnwkttime, avc_avgsvrresptime, clientretransmissionspackets, packetdeltacount, retransmittedpercent. |
| PfA: Server All Apps Delay | A grouping of Server trending App Delay. Information Elements: destinationipaddress, avc_avgsvrresptime. |
| PfA: Server All Apps Delay (cs) | A grouping of Server trending App Delay. Information Elements: serveripaddress, avc_avgsvrresptime. |
| PfA: Server Application Delay | A grouping of Server, Application trending App Delay. Information Elements: destinationipaddress, applicationtag, avc_avgsvrresptime. |
| PfA: Server Application Delay (cs) | A grouping of Server, Application trending App Delay. Information Elements: serveripaddress, applicationtag, avc_avgsvrresptime. |
| PfA: Top QoS Policies | A grouping of QoS Policy trending Packets, Bytes. Information Elements: policymaphierarchy, octetdeltacount, packetdeltacount. |
| PfM: Application Latency | A grouping of Application trending TEPL, RTT. Information Elements: applicationid, trans_event_pkt_lost_count, trans_rtt. |
| PfM: Connections RTT | A grouping of src Port, Source, Destination, dstPort trending TEPL, RTT. Information Elements: sourcettransportport, sourceipaddress, destinationipaddress, destinationtransportport, trans_event_pkt_lost_count, trans_rtt. |
| PfM: Domain RTT (Dst) | A grouping of Destination Domain, DSCP trending TEPL, RTT. Information Elements: dstdomain, ipdiffservcodepoint, trans_event_pkt_lost_count, trans_rtt. |

continues on next page

Table 22 – continued from previous page

| Report | Description |
|--|--|
| PfM: Domain RTT (Src) | A grouping of Source Domain, DSCP trending TEPL, RTT. Information Elements: srcdomain, ipdiffservcodepoint, trans_event_pkt_lost_count, trans_rtt. |
| PfM: Host Jitter by SSRC (Dst) | A grouping of Destination, DSCP, SSRC trending % Pkt Loss, TEPL, Jitter. Information Elements: destinationipaddress, ipdiffservcodepoint, trans_rtp_ssrc, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| PfM: Host Jitter by SSRC (Src) | A grouping of Source, DSCP, SSRC trending % Pkt Loss, TEPL, Jitter. Information Elements: sourceipaddress, ipdiffservcodepoint, trans_rtp_ssrc, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| PfM: Host Jitter (Dst) | A grouping of Destination, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: destinationipaddress, ipdiffservcodepoint, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| PfM: Host Jitter (Src) | A grouping of Source, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: sourceipaddress, ipdiffservcodepoint, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| PfM: Host RTT (Dst) | A grouping of Destination, DSCP trending TEPL, RTT. Information Elements: destinationipaddress, ipdiffservcodepoint, trans_event_pkt_lost_count, trans_rtt. |
| PfM: Host RTT (Src) | A grouping of Source, DSCP trending TEPL, RTT. Information Elements: sourceipaddress, ipdiffservcodepoint, trans_event_pkt_lost_count, trans_rtt. |
| PfM: Host to Host Application Jitter by SSRC | A grouping of Source, Application, DSCP, Destination, SSRC trending % Pkt Loss, TEPL, Byte Rate, Jitter. Information Elements: sourceipaddress, applicationtag, ipdiffservcodepoint, destinationipaddress, trans_rtp_ssrc, app_media_byte_rate, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |

continues on next page

Table 22 – continued from previous page

| Report | Description |
|-------------------------------------|---|
| PfM: Host to Host Jitter | A grouping of Source, DSCP, Destination trending % Pkt Loss, TEPL, Max Jitter, Jitter. Information Elements: sourceipaddress, ipdiffservcodepoint, destinationipaddress, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| PfM: Host to Host Jitter by SSRC | A grouping of Source, DSCP, Destination, SSRC trending % Pkt Loss, TEPL, Jitter. Information Elements: sourceipaddress, ipdiffservcodepoint, destinationipaddress, trans_rtp_ssrc, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| PfM: Host to Host RTP Type and Rate | A grouping of Source, Destination, SSRC, Type trending Byte Rate, Jitter. Information Elements: sourceipaddress, destinationipaddress, trans_rtp_ssrc, trans_rtp_payload_type, app_media_byte_rate, trans_rtp_jitter. |
| PfM: Host to Host RTT | A grouping of Source, DSCP, Destination trending TEPL, RTT. Information Elements: sourceipaddress, ipdiffservcodepoint, destinationipaddress, trans_event_pkt_lost_count, trans_rtt. |
| PfM: Latency by Interface | A grouping of Exporter, in Int trending RTT. Information Elements: plixerexporter, ingressinterface, trans_rtt. |
| PfM: Subnet Jitter (Dst) | A grouping of Dst Subnet, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: dstnetwork, ipdiffservcodepoint, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| PfM: Subnet Jitter (Src) | A grouping of Src Subnet, DSCP trending % Pkt Loss, TEPL, Jitter. Information Elements: srcnetwork, ipdiffservcodepoint, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| PfM: Subnet RTT (Dst) | A grouping of Dst Subnet, DSCP trending TEPL, RTT. Information Elements: dstnetwork, ipdiffservcodepoint, trans_event_pkt_lost_count, trans_rtt. |
| PfM: Subnet RTT (Src) | A grouping of Src Subnet, DSCP trending TEPL, RTT. Information Elements: srcnetwork, ipdiffservcodepoint, trans_event_pkt_lost_count, trans_rtt. |

continues on next page

Table 22 – continued from previous page

| Report | Description |
|------------------------------|---|
| PfM: Subnet to Subnet Jitter | A grouping of Src Subnet, DSCP, Dst Subnet trending % Pkt Loss, TEPL, Jitter. Information Elements: srcnetwork, ipdiffserv-codepoint, dstnetwork, ciscopktlostpercent, trans_event_pkt_lost_count, trans_rtp_jitter. |
| PfM: Subnet to Subnet RTT | A grouping of Src Subnet, DSCP, Dst Subnet trending TEPL, RTT. Information Elements: srcnetwork, ipdiffservcodepoint, dstnetwork, trans_event_pkt_lost_count, trans_rtt. |
| PfM: VoIP | A grouping of Source, Application, DSCP, CBQoS, RTP Payload, Destination, SSRC trending TEPL, Jitter. Information Elements: sourceipaddress, applicationtag, ipdiffservcodepoint, policymaphierarchy, trans_rtp_payload_type, destinationipaddress, trans_rtp_ssrc, trans_event_pkt_lost_count, trans_rtp_jitter. |
| PfR: Active Jitter | A grouping of Border Router, out Int, Traffic Class trending Jitter. Information Elements: ipv4_br_addr, egressinterface, pfrtraffic-class, trans_rtp_jitter. |
| PfR: Active Jitter OOP | A grouping of Border Router, out Int, Traffic Class, PfR Status trending Jitter. Information Elements: ipv4_br_addr, egressinterface, pfrtraffic-class, pfr_status, trans_rtp_jitter. |
| PfR: Active mos | A grouping of Border Router, out Int, Traffic Class trending % Thresh. Information Elements: ipv4_br_addr, egressinterface, pfrtrafficclass, rtpworstmos100. |
| PfR: Active mos OOP | A grouping of Border Router, out Int, Traffic Class, PfR Status trending % Thresh. Information Elements: ipv4_br_addr, egressinterface, pfrtrafficclass, pfr_status, rtpworstmos100. |
| PfR: Active OOP | A grouping of Border Router, out Int, Traffic Class, PfR Status trending Count. Information Elements: ipv4_br_addr, egressinterface, pfrtraffic-class, pfr_status, plixeraggregatedrecordcount. |

continues on next page

Table 22 – continued from previous page

| Report | Description |
|-----------------------------|--|
| PfR: Active RTT | A grouping of Border Router, out Int, Traffic Class trending RTT. Information Elements: <code>ipv4_br_addr</code> , <code>egressinterface</code> , <code>pfrtrafficclass</code> , <code>trans_rtt</code> . |
| PfR: Active RTT OOP | A grouping of Border Router, out Int, Traffic Class, PfR Status trending RTT. Information Elements: <code>ipv4_br_addr</code> , <code>egressinterface</code> , <code>pfrtrafficclass</code> , <code>pfr_status</code> , <code>trans_rtt</code> . |
| PfR: OOP Detail | A grouping of MasterController, Border Router, out Int, Traffic Class, PfR Status trending Count. Information Elements: <code>plixerexporter</code> , <code>ipv4_br_addr</code> , <code>egressinterface</code> , <code>pfrtrafficclass</code> , <code>pfr_status</code> , <code>plixeraggregatedrecordcount</code> . |
| PfR: OOP Master Controllers | A grouping of MasterController, PfR Class trending Count. Information Elements: <code>plixerexporter</code> , <code>pfrtrafficclass</code> , <code>plixeraggregatedrecordcount</code> . |
| PfR: Passive OOP | A grouping of Border Router, out Int, Traffic Class, PfR Status trending Count. Information Elements: <code>ipv4_br_addr</code> , <code>egressinterface</code> , <code>pfrtrafficclass</code> , <code>pfr_status</code> , <code>plixeraggregatedrecordcount</code> . |
| PfR: Passive RTT | A grouping of Border Router, out Int, Traffic Class trending RTT. Information Elements: <code>ipv4_br_addr</code> , <code>egressinterface</code> , <code>pfrtrafficclass</code> , <code>trans_rtt</code> . |
| PfR: Passive RTT OOP | A grouping of Border Router, out Int, Traffic Class, PfR Status trending RTT. Information Elements: <code>ipv4_br_addr</code> , <code>egressinterface</code> , <code>pfrtrafficclass</code> , <code>pfr_status</code> , <code>trans_rtt</code> . |

Table 23: Cisco CTS

| Report | Description |
|-------------------------|---|
| ctsDestination Group | A grouping of ctsdestinationgrouptag trending Packets, Bytes. Information Elements: ctsdestinationgrouptag, octetdeltacount, packetdeltacount. |
| ctsGroups Connections | A grouping of src Port, Group Tag, ctsdestinationgrouptag, dst Port trending Packets, Bytes. Information Elements: sourcettransportport, ctssourcegrouptag, ctsdestinationgrouptag, destinationtransportport, octetdeltacount, packetdeltacount. |
| ctsGroups Conversations | A grouping of Group Tag, Well Known, ctsdestinationgrouptag, Rate trending Packets, Bytes. Information Elements: ctssourcegrouptag, commonport, ctsdestinationgrouptag, rate, octetdeltacount, packetdeltacount. |
| ctsGroups Grouped Flows | A grouping of src Port, Group Tag, Type Of Service, ctsdestinationgrouptag, dst Port trending Packets, Bytes. Information Elements: sourcettransportport, ctssourcegrouptag, ipclassofservice, ctsdestinationgrouptag, destinationtransportport, octetdeltacount, packetdeltacount. |
| ctsSource Group | A grouping of Group Tag trending Packets, Bytes. Information Elements: ctssourcegrouptag, octetdeltacount, packetdeltacount. |
| ctsSrcGrp to ctsDstGrp | A grouping of Group Tag, ctsdestinationgrouptag trending Packets, Bytes. Information Elements: ctssourcegrouptag, ctsdestinationgrouptag, octetdeltacount, packetdeltacount. |

Table 24: Cisco FW

| Report | Description |
|-------------|--|
| ACL to ACL | A grouping of Ingress ACL, Egress ACL trending Flows. Information Elements: nf_f_ingress_acl_id, nf_f_egress_acl_id, plixer-aggregatedrecordcount. |
| Egress ACL | A grouping of Egress ACL trending Flows. Information Elements: nf_f_egress_acl_id, plixer-aggregatedrecordcount. |
| Ingress ACL | A grouping of Ingress ACL trending Flows. Information Elements: nf_f_ingress_acl_id, plixer-aggregatedrecordcount. |

Table 25: Cisco HSL

| Report | Description |
|----------------------------------|---|
| Classes | A grouping of Class, Packets trending Bytes. Information Elements: classid, packetdeltacount, octetdeltacount. |
| Destination-Event | A grouping of Destination, Firewall Event, Extended Event Code, Zone Pair trending Flows. Information Elements: destinationipaddress, firewalllevent, fw_ext_event, zonepair_id, plixeraggregatedrecordcount. |
| Host to Host Events | A grouping of Source, Destination, Firewall Event, Extended Event Code, Zone Pair trending Flows. Information Elements: sourceipaddress, destinationipaddress, firewalllevent, fw_ext_event, zonepair_id, plixeraggregate-drecordcount. |
| Host to Host Events by VRF | A grouping of In VRF, Source, Destination, Out VRF, Firewall Event, Extended Event Code trending Flows. Information Elements: ingressvrfid, sourceipaddress, destinationipaddress, egressvrfid, firewalllevent, fw_ext_event, plixeraggregate-drecordcount. |
| Host to Host with Zone and Class | A grouping of Source, Class, Zone Pair, Destination trending Bytes. Information Elements: sourceipaddress, classid, zonepair_id, destinationipaddress, octetdeltacount. |
| Source-Event | A grouping of Source, Firewall Event, Extended Event Code, Zone Pair trending Flows. Information Elements: sourceipaddress, firewalllevent, fw_ext_event, zonepair_id, plixeraggregatedrecordcount. |
| Zone Pair | A grouping of Zone Pair trending Bytes. Information Elements: zonepair_id, octetdeltacount. |
| Zone Pair and Class | A grouping of Zone Pair, Class trending Bytes. Information Elements: zonepair_id, classid, octetdeltacount. |
| Zone Pair Volume | A grouping of Zone Pair trending Flows. Information Elements: zonepair_id, plixeraggregatedrecordcount. |

Table 26: Cisco IWAN

| Report | Description |
|-----------------------------|---|
| IWAN Bandwidth Usage | A grouping of Source Site, Path Tag ID, Interface Description trending BW In, Speed In, BW Out, Speed Out. Information Elements: source_site_id, path_tag_id, interfacedescription, egress_bw, ingress_bw, maxof_egress_bw, maxof_ingress_bw. |
| IWAN Route Changes | A grouping of Site, BR, Path Tag ID, IWAN Circuit trending Routes Changed. Information Elements: source_site_id, ipv4_br_addr, path_tag_id, interfacedescription, plixeraggregatedrecordcount. |
| IWAN Site to Site Bandwidth | A grouping of BR Router, Src Site, Dst Site, Dst Prefix, Interface ID trending Packets, Avg Bits. Information Elements: ipv4_br_addr, source_site_id, destination_site_id, destination_site_prefix, egressinterface, octetdeltacount, packetdeltacount. |
| IWAN Traffic Control Alerts | A grouping of Source Site, Destination Site, Interface Description, Interface ID, BR Addr, Path Tag ID, Status trending One way delay, AVG Jitter, PKT Loss, Bytes Lost. Information Elements: source_site_id, destination_site_id, interfacedescription, egressinterface, ipv4_br_addr, path_tag_id, oer_unreach, one_way_delay, rtp_jitter_inter_arrival_mean, trans_pkt_lost_rate, trns_cnt_bytes_lost_rate. |

Table 27: Cisco SLT

| Report | Description |
|----------------|--|
| Event | A grouping of l2l3switchevent trending Count, Packets, Bytes. Information Elements: l2l3switchevent, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Event-Extevent | A grouping of l2l3switchevent, l2l3switchextevent trending Count, Packets, Bytes. Information Elements: l2l3switchevent, l2l3switchextevent, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Int | A grouping of Exporter, ingressphysicalinterface trending Count, Packets, Bytes. Information Elements: plixerexporter, ingressphysicalinterface, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Int-Vlan-Event | A grouping of Exporter, ingressphysicalinterface, vlanid, l2l3switchevent trending Count, Packets, Bytes. Information Elements: plixerexporter, ingressphysicalinterface, vlanid, l2l3switchevent, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Vlan | A grouping of Exporter, vlanid trending Count, Packets, Bytes. Information Elements: plixerexporter, vlanid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |

Table 28: Cisco VQM

| Report | Description |
|--------------------------|--|
| Connections eMOS | A grouping of Source, Src Port, Destination, Dest Port trending Frame Rate, eMOS Score. Information Elements: sourceipaddress, sourcettransportport, destinationipaddress, destinationtransportport, videoemosscore, vqmframerate. |
| Connections eMOS Detail | A grouping of Source, Src Port, Destination, Dest Port trending Frame Rate, eMOS Pkt Lost, eMOS Compression, eMOS Score. Information Elements: sourceipaddress, sourcettransportport, destinationipaddress, destinationtransportport, videoemosscore, vqmemoscompressionbitstream, vqmemospacketlostbitstream, vqmframerate. |
| Destination eMOS | A grouping of Destination trending Frame Rate, eMOS Score. Information Elements: destinationipaddress, videoemosscore, vqmframerate. |
| Destination eMOS Detail | A grouping of Destination trending Frame Rate, eMOS Pkt Lost, eMOS Compression, eMOS Score. Information Elements: destinationipaddress, videoemosscore, vqmemoscompressionbitstream, vqmemospacketlostbitstream, vqmframerate. |
| Host to Host eMOS | A grouping of Source, Destination trending Frame Rate, eMOS Score. Information Elements: sourceipaddress, destinationipaddress, videoemosscore, vqmframerate. |
| Host to Host eMOS Detail | A grouping of Source, Destination trending Frame Rate, eMOS Pkt Lost, eMOS Compression, eMOS Score. Information Elements: sourceipaddress, destinationipaddress, videoemosscore, vqmemoscompressionbitstream, vqmemospacketlostbitstream, vqmframerate. |
| Source eMOS | A grouping of Source trending Frame Rate, eMOS Score. Information Elements: sourceipaddress, videoemosscore, vqmframerate. |
| Source eMOS Detail | A grouping of Source trending Frame Rate, eMOS Pkt Lost, eMOS Compression, eMOS Score. Information Elements: sourceipaddress, videoemosscore, vqmemoscompressionbitstream, vqmemospacketlostbitstream, vqmframerate. |

Table 29: Client Server

| Report | Description |
|--------------------------|---|
| Client | A grouping of Client IP trending sum_plxr_client_bytes, sum_plxr_server_bytes. Information Elements: plxr_client_ip, plxr_client_bytes, plxr_server_bytes. |
| Client Apps | A grouping of Client IP, Application ID trending sum_plxr_client_bytes, sum_plxr_server_bytes. Information Elements: plxr_client_ip, applicationid, plxr_client_bytes, plxr_server_bytes. |
| Client Server | A grouping of Client IP, Server IP trending sum_plxr_client_bytes, sum_plxr_server_bytes. Information Elements: plxr_client_ip, plxr_server_ip, plxr_client_bytes, plxr_server_bytes. |
| Client Server Apps | A grouping of Client IP, Application ID, Server IP trending Client, Server. Information Elements: plxr_client_ip, applicationid, plxr_server_ip, plxr_client_bytes, plxr_server_bytes. |
| Client Server Apps Flags | A grouping of Client IP, Application ID, Server IP trending TCP Flags, Client, Server. Information Elements: plxr_client_ip, applicationid, plxr_server_ip, plxr_client_bytes, plxr_server_bytes, tcpcontrolbits. |
| Client Server Flags | A grouping of Client IP, Server IP trending TCP Flags, sum_plxr_client_bytes, sum_plxr_server_bytes. Information Elements: plxr_client_ip, plxr_server_ip, plxr_client_bytes, plxr_server_bytes, tcpcontrolbits. |
| Server | A grouping of Server IP trending sum_plxr_client_bytes, sum_plxr_server_bytes. Information Elements: plxr_server_ip, plxr_client_bytes, plxr_server_bytes. |
| Server Apps | A grouping of Server IP, Application ID trending sum_plxr_client_bytes, sum_plxr_server_bytes. Information Elements: plxr_server_ip, applicationid, plxr_client_bytes, plxr_server_bytes. |

Table 30: Counts

| Report | Description |
|--------------------------------|---|
| Clients | A grouping of Client trending Flows. Information Elements: clientipaddress, plixeraggregatedrecordcount. |
| Destination | A grouping of Destination trending Flows. Information Elements: destinationipaddress, plixeraggregatedrecordcount. |
| Initiator Group with Dst Port | A grouping of Source IP Group, Well Known Port, Destination IP Group, Destination Port trending Packets, Bytes, Flows. Information Elements: srcipgroup, commonport, dstipgroup, destinationtransportport, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Internal External Destinations | A grouping of Destination trending Unique Hosts. Information Elements: dstinternal, destinationipaddress. |
| Internal External Pairs | A grouping of Source, Destination trending Unique Srcs, Unique Dsts. Information Elements: srcinternal, dstinternal, destinationipaddress, sourceipaddress. |
| Internal External Sources | A grouping of Source trending Unique Hosts. Information Elements: srcinternal, sourceipaddress. |
| Pairs | A grouping of Source, Destination trending Flows. Information Elements: sourceipaddress, destinationipaddress, plixeraggregatedrecordcount. |
| Pair Source post NAT | A grouping of Source, Src Post NAT, Destination trending Flows. Information Elements: sourceipaddress, postnatsourceipaddress, destinationipaddress, plixeraggregatedrecordcount. |
| Pair Source post NAT and NAP | A grouping of Source, Src Post NAT, Src Port, Src NAP Port, Dst Port, Destination trending Flows. Information Elements: sourceipaddress, postnatsourceipaddress, sourcetransportport, postnapsourceport, destinationtransportport, destinationipaddress, plixeraggregatedrecordcount. |
| Protocol | A grouping of Protocol trending Flows. Information Elements: protocolidentifier, plixeraggregatedrecordcount. |
| Servers | A grouping of Server trending Flows. Information Elements: serveripaddress, plixeraggregatedrecordcount. |
| Source | A grouping of Source trending Flows. Information Elements: sourceipaddress, plixeraggregatedrecordcount. |
| VRFID with NAT and Src | A grouping of In VRFID, NAT Event, NAT Pool, NAT Source, NAT Destination, NAT Information, NAT |

660

7. Additional Resources

Table 31: Designed Reports

| Report | Description |
|---------------------|--|
| Issue 4657 | A grouping of Source IP Address, sourcetransportport, Destination IP Address, destinationtransportport trending avg_connectionapplicationdelay, avg_connectionclienttoserverresponsedelay, avg_connectionnetworktoclientdelay, sum_connectionclientpacketretransmissioncount. Information Elements: sourceipaddress, source-transportport, destinationipaddress, destinationtransportport, connectionapplicationdelay, connectionclientpacketretransmissioncount, connectionclienttoserverresponsedelay, connectionnetworktoclientdelay. |
| Issue 4657a | A grouping of Source IP Address, sourcetransportport, Destination IP Address, destinationtransportport trending avg_connectionapplicationdelay, avg_connectionclienttoserverresponsedelay, avg_connectionnetworktoclientdelay, sum_connectionclientpacketretransmissioncount. Information Elements: sourceipaddress, source-transportport, destinationipaddress, destinationtransportport, connectionapplicationdelay, connectionclientpacketretransmissioncount, connectionclienttoserverresponsedelay, connectionnetworktoclientdelay. |
| Source Packets Size | A grouping of Source IP Address trending Avg. Information Elements: sourceipaddress, avgpacketsize. |

Table 32: Destination Reports

| Report | Description |
|---------------------------------|---|
| Autonomous System by IP | A grouping of Destination AS trending Packets, Bytes. Information Elements: dstipas, octetdeltacount, packetdeltacount. |
| Autonomous System by Tag | A grouping of Dst AS trending Packets, Bytes. Information Elements: bgpdestinationasnumber, octetdeltacount, packetdeltacount. |
| Autonomous System by Tag (Peer) | A grouping of bgpNextadjacentasnumber trending Packets, Bytes. Information Elements: bgpNextadjacentasnumber, octetdeltacount, packetdeltacount. |
| Countries | A grouping of Destination Country trending Packets, Bytes. Information Elements: dstcountry, octetdeltacount, packetdeltacount. |
| Countries with AS | A grouping of Dest Country, Dest AS, Hosts (Dst) trending Flows, Packets, Bytes. Information Elements: dstcountry, dstipas, sourceipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Customer VLAN | A grouping of postdot1qcustomervlanid trending Flows, Packets, Bytes. Information Elements: postdot1qcustomervlanid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Destination w/Flags | A grouping of Destination IP Address, tcpcontrolbits trending Packets, Bytes, Flows. Information Elements: destinationipaddress, tcpcontrolbits, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Dest. IP Groups | A grouping of Destination IP Group trending Packets, Bytes. Information Elements: dstip-group, octetdeltacount, packetdeltacount. |
| dot1q VLAN | A grouping of postdot1qvlanid trending Flows, Packets, Bytes. Information Elements: postdot1qvlanid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Dst IP - Src AS | A grouping of Exporter, Destination IP Address, Src AS trending Packets, Bytes. Information Elements: plixerexporter, destinationipaddress, bgpsourceasnumber, octetdeltacount, packetdeltacount. |
| Host Flows | A grouping of Destination trending Hosts (Source), Packets, Flows. Information Elements: destinationipaddress, packetdeltacount, plixeraggregatedrecordcount, sourceipaddress, packetdeltacount, plixeraggregatedrecordcount. |
| Hosts | A grouping of Destination trending Packets, Bytes. Information Elements: destinationipaddress, octetdeltacount, packetdeltacount. |

Table 33: Exinda

| Report | Description |
|--------------------------|--|
| Application - App Group | A grouping of Exporter, Application, GROUP, TRAFFIC_CLASS trending Avg Srv Del, RTT, NULL. Information Elements: plixerexporter, applicationtag, ex_app_group_name, ex_traffic_class, ex_rtt, ex_server_delay, octetdeltacount. |
| Application Detail | A grouping of Application trending Avg. AQS, Packets, NULL. Information Elements: applicationtag, ex_aqs, octetdeltacount, packetdeltacount. |
| Application Group | A grouping of Exporter, GROUP, TRAFFIC_CLASS trending Avg Srv Del, RTT, NULL. Information Elements: plixerexporter, ex_app_group_name, ex_traffic_class, ex_rtt, ex_server_delay, octetdeltacount. |
| Application Performances | A grouping of Application trending Avg. AQS, Bytes Lost, Nwk. Delay, Srv. Delay, RTT. Information Elements: applicationtag, ex_aqs, ex_bytes_lost, ex_net_delay, ex_rtt, ex_server_delay. |
| Destination User | A grouping of Exporter, dst_user trending Packets, NULL. Information Elements: plixerexporter, ex_user_id_dst, octetdeltacount, packetdeltacount. |
| Extra Info | A grouping of Exporter, EXTRA_INFO_ID, TRAFFIC_CLASS trending Bytes Lost, Avg Srv Del, RTT, NULL. Information Elements: plixerexporter, ex_extra_info_id, ex_traffic_class, ex_bytes_lost, ex_rtt, ex_server_delay, octetdeltacount. |
| Pair by Policy | A grouping of Exporter, Source, Destination, Policy trending Packets, NULL. Information Elements: plixerexporter, sourceipaddress, destinationipaddress, ex_policy_id, octetdeltacount, packetdeltacount. |
| Pair Latency | A grouping of Source, Destination, TRAFFIC_CLASS trending Avg Srv Del, RTT, NULL. Information Elements: sourceipaddress, destinationipaddress, ex_traffic_class, ex_rtt, ex_server_delay, octetdeltacount. |
| Pair, Ports and Latency | A grouping of Source, Src Port, Dst Port, Destination trending Avg Srv Del, RTT, NULL. Information Elements: sourceipaddress, sourceport, destinationipaddress, destinationport, destinationtransportport, destinationipaddress, ex_rtt, ex_server_delay, octetdeltacount. |
| 7.1. Appendices | 603 |
| Pair, V-IP Detail | A grouping of Source, Destination, TRAFFIC_CLASS trending Avg Srv Del, RTT, NULL. Information Elements: sourceipaddress, destinationipaddress, ex_traffic_class, ex_rtt, ex_server_delay, octetdeltacount. |

Table 34: FirePOWER

| Report | Description |
|-------------------------------|--|
| App Internet HTTP Host | A grouping of Application, FS App, HTTP Host trending Flows, Bytes. Information Elements: applicationname, firesight_application, firesight_http_host, octetdeltacount, plixeraggregatedrecordcount. |
| Application E-Zone & Sub Type | A grouping of Application, FS App, Egress Zone, Event Subtype, Event Type trending Flows. Information Elements: applicationname, firesight_application, firesight_egress_zone, firesight_event_subtype, firesight_event_type, plixeraggregatedrecordcount. |
| Application I-Zone & Sub Type | A grouping of Application, FS App, Ingress Zone, Event Subtype, Event Type trending Flows. Information Elements: applicationname, firesight_application, firesight_ingress_zone, firesight_event_subtype, firesight_event_type, plixeraggregatedrecordcount. |
| Firewall List | A grouping of Firewall trending Flows, Packets, Bytes. Information Elements: firesight_sensor_ipv6, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Ingress and Egress Zones | A grouping of Ingress Zone, Egress Zone, Event Type trending Flows. Information Elements: firesight_ingress_zone, firesight_egress_zone, firesight_event_type, plixeraggregatedrecordcount. |
| User App HTTP Host | A grouping of Source IP, Username, Application, FS App, HTTP Host trending Flows, Bytes. Information Elements: sourceipaddress, username, applicationname, firesight_application, firesight_http_host, octetdeltacount, plixeraggregatedrecordcount. |
| User App HTTP URL | A grouping of Source IP, Username, Application, FS App, FS URL trending Flows. Information Elements: sourceipaddress, username, applicationname, firesight_application, firesight_http_url, plixeraggregatedrecordcount. |
| User Application | A grouping of Source IP, Username, Application, FS App trending Flows, Bytes. Information Elements: sourceipaddress, username, applicationname, firesight_application, octetdeltacount, plixeraggregatedrecordcount. |
| Web App and Source IP | A grouping of Web Application, Application, Source IP trending Flows. Information Elements: firesight_web_application, applicationname, sourceipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |

Table 35: Firewall Events

| Report | Description |
|-------------------------|--|
| Destination-Event | A grouping of Destination, Firewall Event trending Flows. Information Elements: destinationipaddress, firewallevnt, plixeraggregatedrecordcount. |
| Destination-Event-Ext | A grouping of Destination, Firewall Event, Extended Event trending Flows. Information Elements: destinationipaddress, firewallevnt, nf_f_fw_ext_event, plixeraggregatedrecordcount. |
| Event-Ext-ACL | A grouping of Firewall Event, Extended Event, Ingress ACL, Egress ACL trending Flows. Information Elements: firewallevnt, nf_f_fw_ext_event, nf_f_ingress_acl_id, nf_f_egress_acl_id, plixeraggregatedrecordcount. |
| Firewall Events | A grouping of Firewall Event trending Count. Information Elements: firewallevnt, plixeraggregatedrecordcount. |
| Firewall Events by Host | A grouping of Source, Firewall Event trending Count. Information Elements: sourceipaddress, firewallevnt, plixeraggregatedrecordcount. |
| Firewall Events Ext | A grouping of FW Event Ext trending Flows. Information Elements: fw_ext_event, plixeraggregatedrecordcount. |
| Pairs-Event | A grouping of Source, Destination, Firewall Event trending Flows. Information Elements: sourceipaddress, destinationipaddress, firewallevnt, plixeraggregatedrecordcount. |
| Pairs-Event Ext | A grouping of Source, Destination, FW Event Ext trending Flows. Information Elements: sourceipaddress, destinationipaddress, fw_ext_event, plixeraggregatedrecordcount. |
| Pairs-Event-Ext | A grouping of Source, Destination, Firewall Event, Extended Event trending Flows. Information Elements: sourceipaddress, destinationipaddress, firewallevnt, nf_f_fw_ext_event, plixeraggregatedrecordcount. |
| Protocol-Event | A grouping of Protocol, Firewall Event trending Flows. Information Elements: protocolidentifier, firewallevnt, plixeraggregatedrecordcount. |
| Protocol-Event-Ext | A grouping of Protocol, Firewall Event, Extended Event trending Flows. Information Elements: protocolidentifier, firewallevnt, nf_f_fw_ext_event, plixeraggregatedrecordcount. |
| 7.1. Appendices | 605 |
| Source-Event | A grouping of Source, Firewall Event trending Flows. Information Elements: sourceipaddress, firewallevnt, plixeraggregatedrecordcount. |

Table 36: FlowPro APM Reports

| Report | Description |
|-----------------------------------|---|
| Application Latency | A grouping of L7 App trending Client, Server, Appl. Information Elements: l7_proto_name, appl_latency_ms, client_nw_delay_ms, server_nw_delay_ms. |
| App Priority & Latency | A grouping of L7 App, Priority trending Client, Server, Appl. Information Elements: l7_proto_name, ipclassofservice, appl_latency_ms, client_nw_delay_ms, server_nw_delay_ms. |
| Defined Application Latency | A grouping of Application trending Appl, Client, Server, Packets, Bytes. Information Elements: applicationid, appl_latency_ms, client_nw_delay_ms, octetdeltacount, packetdeltacount, server_nw_delay_ms. |
| Host Jitter | A grouping of Source trending Pkt Loss, Jitter, Packets, Bytes. Information Elements: sourceipaddress, octetdeltacount, packetdeltacount, rtp_in_jitter, rtp_in_pkt_lost. |
| Host Jitter By SSRC (Dst) | A grouping of Destination, SSRC, Codec trending Pkt Loss, Jitter, Packets, Bytes. Information Elements: destinationipaddress, rtp_ssrc, rtp_out_payload_type, octetdeltacount, packetdeltacount, rtp_out_jitter, rtp_out_pkt_lost. |
| Host Jitter By SSRC (Src) | A grouping of Source, SSRC, Codec trending Pkt Loss, Jitter, Packets, Bytes. Information Elements: sourceipaddress, rtp_ssrc, rtp_in_payload_type, octetdeltacount, packetdeltacount, rtp_in_jitter, rtp_in_pkt_lost. |
| Hosts Latency (Dst) | A grouping of Destination trending Appl, Client, Server, Packets, Bytes. Information Elements: destinationipaddress, appl_latency_ms, client_nw_delay_ms, octetdeltacount, packetdeltacount, server_nw_delay_ms. |
| Hosts Latency (Src) | A grouping of Source trending Appl, Client, Server, Packets, Bytes. Information Elements: sourceipaddress, appl_latency_ms, client_nw_delay_ms, octetdeltacount, packetdeltacount, server_nw_delay_ms. |
| Host to Host Jitter All by SSRC | A grouping of Source, Src Payload, SSRC, Destination, Dst Payload trending Src Pkt Loss, Src Jitter, Dst Pkt Loss, Dst Jitter, Packets, Bytes. Information Elements: sourceipaddress, rtp_in_payload_type, rtp_out_payload_type, octetdeltacount, packetdeltacount, rtp_in_jitter, rtp_in_pkt_lost, rtp_out_jitter, rtp_out_pkt_lost. |
| 606 | 717 |
| Host to Host Jitter By SSRC/Codec | A grouping of Source, Destination, SSRC, Codec trending Pkt Loss, Jitter, Packets, Bytes. Information Elements: sourceipaddress, destinationipaddress, rtp_ssrc, rtp_payload_type, octetdeltacount, packetdeltacount, rtp_in_jitter, rtp_in_pkt_lost, rtp_out_jitter, rtp_out_pkt_lost. |

Table 37: FlowPro Defender Reports

| Report | Description |
|------------------------------|---|
| Alert > All Details | A grouping of Category, Signature, Source, Destination trending Observation Count. Information Elements: nids_category, nids_signature, sourceipaddress, destinationipaddress, plixeraggregatedrecordcount. |
| Alert > Category | A grouping of Category trending Observation Count. Information Elements: nids_category, plixeraggregatedrecordcount. |
| Alert > Category & Signature | A grouping of Category, Signature trending Observation Count. Information Elements: nids_category, nids_signature, plixeraggregatedrecordcount. |
| DNS > Auth | A grouping of Auth Rname trending Observation Count. Information Elements: dns_soa_rname, plixeraggregatedrecordcount. |
| DNS Client Latency | A grouping of Client trending DNS Requests, Latency. Information Elements: dnsnxclientip4address, dnsresolvetime, plixeraggregatedrecordcount. |
| DNS Client / Server Latency | A grouping of Client, Responding DNS Svr trending DNS Requests, Latency. Information Elements: dnsnxclientip4address, dnsnxserverip4address, dnsresolvetime, plixeraggregatedrecordcount. |
| DNS Domain Reputation | A grouping of Source, QName, Resolved Address, DNS Server, Threat Category trending Count. Information Elements: sourceipaddress, dnsname, dnsresolvedip4address, dnsnxserverip4address, reputationcategoryid, plixeraggregatedrecordcount. |
| DNS Exfiltration | A grouping of Source, Destination, QName, DNS Text trending Length, Count. Information Elements: sourceipaddress, destinationipaddress, dnsname, dnstext, dnstextlength, plixeraggregatedrecordcount. |
| DNS Query Refused | A grouping of Client, DNS Server, FQDN trending Lookup Time. Information Elements: dnsnxclientip4address, dnsnxserverip4address, dnsname, flowstartseconds. |

continues on next page

Table 37 – continued from previous page

| Report | Description |
|-------------------------------|---|
| DNS > RCodes | A grouping of Rcode trending Observation Count. Information Elements: dnsrcode, plixeraggregatedrecordcount. |
| DNS Request Latency | A grouping of Client, QName, Resolved to, Responding DNS Svr trending Latency. Information Elements: dnsnxclientip4address, dnsname, dnsresolvedip4address, dnsnxserverip4address, dnsresolvetime. |
| DNS > Requests | A grouping of Request trending Observation Count. Information Elements: dns_rrname, plixeraggregatedrecordcount. |
| DNS Request Timeout | A grouping of Client, DNS Query Name trending Count. Information Elements: dnsnxclientip4address, dnsname, plixeraggregatedrecordcount. |
| DNS Server Failure | A grouping of Client, DNS Server, FQDN trending Lookup Time. Information Elements: dnsnxclientip4address, dnsnxserverip4address, dnsname, flowstartseconds. |
| DNS Server Latency | A grouping of Responding DNS Svr trending DNS Requests, Latency. Information Elements: dnsnxserverip4address, dnsresolvetime, plixeraggregatedrecordcount. |
| DNS Server Responding Details | A grouping of DNS Server, Client, FQDN, Resolved Address trending Resolve Count. Information Elements: dnsnxserverip4address, dnsnxclientip4address, dnsname, dnsresolvedip4address, plixeraggregatedrecordcount. |
| DNS Server Responding Summary | A grouping of DNS Server trending Number of Clients, Unique Lookup Count, Minimum Resolution Time. Information Elements: dnsnxserverip4address, dnsnxclientip4address, dnsresolvetime, plixeraggregatedrecordcount. |
| File Info > All File Details | A grouping of Source, Destination, File Name, MD5 Checksum, SHA256 Checksum trending Bytes. Information Elements: sourceipaddress, destinationipaddress, filename, md5_file_checksum, sha256_file_checksum, file_size_octets. |

continues on next page

Table 37 – continued from previous page

| Report | Description |
|----------------------------------|--|
| File Info > CheckSums | A grouping of MD5 Checksum, SHA256 Checksum trending File Size. Information Elements: md5_file_checksum, sha256_file_checksum, file_size_octets. |
| File Info > Filename & CheckSums | A grouping of File Name, MD5 Checksum, SHA256 Checksum trending File Size. Information Elements: filename, md5_file_checksum, sha256_file_checksum, file_size_octets. |
| HTTP > All Details | A grouping of Source, Destination, Request Host, Request Target, User Agent, Content Type, Request Method, Status Code trending Total Payload. Information Elements: sourceipaddress, destinationipaddress, httprequesthost, httprequesttarget, httpuseragent, httpcontenttype, httprequestmethod, httpstatuscode, ippayloadlength. |
| HTTP > Content Type | A grouping of Content Type, Request Method, Status Code trending Total Payload. Information Elements: httpcontenttype, httprequestmethod, httpstatuscode, ippayloadlength. |
| HTTP > Request Target | A grouping of Request Target trending Total Payload. Information Elements: httprequesttarget, ippayloadlength. |
| HTTP > User Agent | A grouping of User Agent trending Observation Count. Information Elements: httpuseragent, plixeraggregatedrecordcount. |
| NX-FQDN | A grouping of FQDN trending DNS Clients, Resolve Count. Information Elements: dnsnxqname, dnsnxclientip4address, plixeraggregatedrecordcount. |
| SMB > File Details | A grouping of Source, Destination, Command, Status, File Name, Operation, Permissions, Accessed, Modified, File Size trending Observed Count. Information Elements: sourceipaddress, destinationipaddress, smb_command, smb_status, smb_filename, smb_disposition, smb_access, smb_accessed_time, smb_modified_time, smb_file_size, plixeraggregatedrecordcount. |

continues on next page

Table 37 – continued from previous page

| Report | Description |
|--------------------------------------|--|
| SMB > NTLMSSP Authentication Details | A grouping of Source, Destination, User, Host, Domain, Status, Version trending Observed Count. Information Elements: sourceipaddress, destinationipaddress, smb_ntlmssp_user, smb_ntlmssp_host, smb_ntlmssp_domain, smb_status, smb_ntlmssp_version, plixeraggregatedrecordcount. |
| SNMP > All Details | A grouping of Community, User, Vars, PDU Type trending Observation Count. Information Elements: mrtgsnmpcommunity, snmp_usm, snmp_var, snmp_pdu_type, plixeraggregatedrecordcount. |
| SNMP > Community | A grouping of Community trending Observation Count. Information Elements: mrtgsnmpcommunity, plixeraggregatedrecordcount. |
| SNMP > PDU Type | A grouping of PDU Type trending Observation Count. Information Elements: snmp_pdu_type, plixeraggregatedrecordcount. |
| SNMP > User | A grouping of User trending Observation Count. Information Elements: snmp_usm, plixeraggregatedrecordcount. |
| SNMP > Version | A grouping of Version trending Observation Count. Information Elements: mrtgsnmpversion, plixeraggregatedrecordcount. |
| Src and # of DNS servers | A grouping of Client, User Name(s) trending # of DNS servers. Information Elements: dnsnxclientip4address, dnsclientname, dnsnxserverip4address. |
| Src and # of NX 2LD | A grouping of Client, User Name(s), DNS Server trending NX Replies. Information Elements: dnsnxclientip4address, dnsclientname, dnsnxserverip4address, dnsqname2ld. |
| Src and # of NX 3LD | A grouping of Client, User Name(s), DNS Server trending NX Replies. Information Elements: dnsnxclientip4address, dnsclientname, dnsnxserverip4address, dnsqname3ld. |
| Src and # of NX Replies | A grouping of Client, User Name(s) trending NX Responses. Information Elements: dnsnxclientip4address, dnsclientname, dnsnxqname. |

continues on next page

Table 37 – continued from previous page

| Report | Description |
|------------------|---|
| Src with NX 2LD | A grouping of Client, User Name(s), 2nd Level Domain, DNS Server trending Count. Information Elements: dnsnxclientip4address, dnsclientname, dnsqname2ld, dnsnxserverip4address, plixeraggregatedrecordcount. |
| Src with NX 3LD | A grouping of Client, User Name(s), 3rd Level Domain, DNS Server trending Count. Information Elements: dnsnxclientip4address, dnsclientname, dnsqname3ld, dnsnxserverip4address, plixeraggregatedrecordcount. |
| Src with NX FQDN | A grouping of Client, User Name(s), DNS Query Name, DNS Server trending Count. Information Elements: dnsnxclientip4address, dnsclientname, dnsnxqname, dnsnxserverip4address, plixeraggregatedrecordcount. |
| Top 2LD Requests | A grouping of 2nd Level Domains trending Clients Requesting, Resolve Count. Information Elements: request2ld, dnsnxclientip4address, dnsresolvedip4address. |
| Top 3LD Requests | A grouping of 3rd Level Domains trending Clients Requesting, Resolve Count. Information Elements: request3ld, dnsnxclientip4address, dnsresolvedip4address. |

Table 38: FQDN Reports

| Report | Description |
|----------------------------|--|
| Destination FQDN | A grouping of Destination, FQDN trending Lookups. Information Elements: destinationipaddress, dst_fqdn, fqdn_lookup_count. |
| Host to Host with Dst FQDN | A grouping of Source, Destination, Dst FQDN trending Lookup. Information Elements: sourceipaddress, destinationipaddress, dst_fqdn, fqdn_lookup_count. |

Table 39: Gigamon

| Report | Description |
|--------------------------|--|
| App Intel - DNS | A grouping of App, Src IP, Dst IP, Query, Response, Query Type trending Count. Information Elements: applicationid, sourceipaddress, destinationipaddress, dnsqueryname, gigamondnsresponseipv4address, gm_dns_networkservice_host_type, plixer-aggregatedrecordcount. |
| App Intel - FTP | A grouping of App, Src IP, Dst IP, Filename, User, Pass, File Size trending Bytes. Information Elements: applicationid, sourceipaddress, destinationipaddress, gm_ftp_fileserver_filename, gm_ftp_fileserver_login, gm_ftp_fileserver_password, gm_ftp_fileserver_filesize, octetdeltacount. |
| App Intel - HTTP | A grouping of App, Src IP, Dst IP, User Agent, HTTP Method, Host, URI, Referer, User Agent trending Bytes. Information Elements: applicationid, sourceipaddress, destinationipaddress, httpuseragent, gm_http_web_method, gm_http_web_host, gm_http_web_uri, gm_http_web_referer, httpstatuscode, octetdeltacount. |
| App Intel - SMB | A grouping of App, Src IP, Dst IP, File, SMB Version, NTLM User, NTLM Workstation trending Bytes. Information Elements: applicationid, sourceipaddress, destinationipaddress, gm_smb_fileserver_filename, gm_smb_fileserver_version, gm_smb_fileserver_ntlm_user, gm_smb_fileserver_ntlm_workstation, octetdeltacount. |
| App Intel - SMTP | A grouping of App, Src IP, Dst IP, Recipient, Sender, Subject, Attachment trending Bytes. Information Elements: applicationid, sourceipaddress, destinationipaddress, gm_smtp_mail_receiver, gm_smtp_mail_sender, gm_smtp_mail_subject, gm_smtp_mail_attach_filename, octetdeltacount. |
| Destination Name and URL | A grouping of Destination, User Name(s), URL trending Count. Information Elements: destinationipaddress, dstipname, gigamonhttprequrl, plixeraggregatedrecordcount. |

612 DNS All Details

717 Additional Resources

A grouping of Src IP, Dst IP, DNS Request, IP Returned, Authority Name trending Count. Information Elements: sourceipaddress, destinationipaddress, dnsqueryname, gigamondnsresponseipv4address, gigamondnsauthorityname, plixer-

Table 40: Honeynet

| Report | Description |
|-----------------------------|--|
| Adversary and State | A grouping of Adversary, State trending Count. Information Elements: sourceipaddress, connectionstate, plixeraggregatedrecordcount. |
| Adversary and String | A grouping of Adversary, String trending Count. Information Elements: sourceipaddress, comments, plixeraggregatedrecordcount. |
| Adversary, String and State | A grouping of Adversary, String, State trending Count. Information Elements: sourceipaddress, comments, connectionstate, plixeraggregatedrecordcount. |
| Forensic with Start | A grouping of Start Time, Source, String, State trending Count. Information Elements: flowstart-milliseconds, sourceipaddress, comments, connectionstate, plixeraggregatedrecordcount. |
| State | A grouping of State trending Count. Information Elements: connectionstate, plixeraggregatedrecordcount. |
| Strings | A grouping of String trending Count. Information Elements: comments, plixeraggregatedrecordcount. |
| Strings and State | A grouping of String, State trending Count. Information Elements: comments, connectionstate, plixeraggregatedrecordcount. |

Table 41: HTTP

| Report | Description |
|-----------------------------|---|
| Host to Host Request Volume | A grouping of Source, Destination trending Requests, Packets, Bytes. Information Elements: httprequesthost, destinationipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| HTTP User Agent | A grouping of Source, User Agent trending Flow Count, Bytes. Information Elements: httprequesthost, httpuseragent, octetdeltacount, plixeraggregatedrecordcount. |
| User Agent | A grouping of pm_cisco_httpuseragent trending Count, Packets, Bytes. Information Elements: pm_cisco_httpuseragent, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |

Table 42: Juniper

| Report | Description |
|-------------------------|--|
| Application Performance | A grouping of Application trending Uplink Pkts, Downlink Pkts, Retrans Uplink, Retrans Downlink, Smooth RTT Up, Smooth RTT Down. Information Elements: application-name, downlinkpackets, retranscppedacketsdownlink, retranscppedacketsuplink, smoothrttdownlink, smoothrttuplink, uplinkpackets. |
| Host and Num Inst | A grouping of Host, Num Inst 1, Num Inst 2, Num Inst 3, Num Inst 4, Num Inst 5 trending Flows. Information Elements: host, numinstances_1, numinstances_2, numinstances_3, numinstances_4, numinstances_5, plixeraggregatedrecordcount. |
| Host and Status Code | A grouping of Host, Status Code 1, Status Code 2, Status Code 3, Status Code 4, Status Code 5 trending Flows. Information Elements: host, statuscode_1, statuscode_2, statuscode_3, statuscode_4, statuscode_5, plixeraggregatedrecordcount. |
| Host DNS Response Time | A grouping of Source trending Flows, Max DNS Resp., Avg DNS Resp.. Information Elements: host, dnsresponsetime, plixeraggregatedrecordcount. |
| HTTP Details | A grouping of Host, Method, Referrer, Response Code, URI trending Flows. Information Elements: host, http_method, http_referrer, http_responsecode, http_uri, plixeraggregatedrecordcount. |
| HTTP Method | A grouping of HTTP Method trending Uplink Pkts, Downlink Pkts, Uplink Octets, Downlink Octets, Flows. Information Elements: http_method, downlinkoctets, downlinkpackets, plixeraggregatedrecordcount, uplinkoctets, uplinkpackets. |
| HTTP Referrer | A grouping of HTTP Referrer trending Uplink Pkts, Downlink Pkts, Uplink Octets, Downlink Octets, Flows. Information Elements: http_referrer, downlinkoctets, downlinkpackets, plixeraggregatedrecordcount, uplinkoctets, uplinkpackets. |
| HTTP Response Code | A grouping of Response Code trending Uplink Pkts, Downlink Pkts, Uplink Octets, Downlink Octets, Flows. Information Elements: http_responsecode, downlinkoctets, downlinkpackets, plixeraggregatedrecordcount, uplinkoctets, uplinkpackets. |
| 614 | 7. Additional Resources |
| HTTP URI | A grouping of URI trending Uplink Pkts, Downlink Pkts, Uplink Octets, Downlink Octets, Flows. Information Elements: http_uri, downlinkoctets, downlinkpackets, plixeraggregatedrecordcount, uplinkoctets, uplinkpackets. |

Table 43: Keysight Reports

| Report | Description |
|--------------------------|---|
| App with Latency | A grouping of Application trending RTT, Bytes. Information Elements: applicationid, latency, octetdeltacount. |
| Browsers | A grouping of Browser trending Packets, Bytes. Information Elements: browsername, octetdeltacount, packetdeltacount. |
| Connections with Latency | A grouping of Source IP, Source Port, Destination IP, Destination Port trending RTT. Information Elements: sourceipaddress, sourceport, destinationipaddress, destinationport, latency. |
| Conversation App Latency | A grouping of Source IP, Application, Destination IP trending RTT, Bytes. Information Elements: sourceipaddress, applicationid, destinationipaddress, latency, octetdeltacount. |
| Device and Location | A grouping of OS Name, Source, City, Country trending Packets, Bytes. Information Elements: osdevicename, sourceipaddress, sourcecityname, sourcecountryname, octetdeltacount, packetdeltacount. |
| Encryption | A grouping of Source, Destination, connencrypttype, encryptioncipher, encryptionkeylength trending Packets, octets. Information Elements: sourceipaddress, destinationipaddress, connencrypttype, encryptioncipher, encryptionkeylength, octetdeltacount, packetdeltacount. |
| L7 Application | A grouping of L7 Application trending Packets, Bytes, Flows. Information Elements: l7applicationname, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| OS | A grouping of OS Name trending Packets, Bytes. Information Elements: osdevicename, octetdeltacount, packetdeltacount. |
| OS Device Name | A grouping of OS Name, Source trending Packets, Bytes. Information Elements: osdevicename, sourceipaddress, octetdeltacount, packetdeltacount. |
| Source City | A grouping of City trending Packets, Bytes. Information Elements: sourcecityname, octetdeltacount, packetdeltacount. |
| Source Country | A grouping of Country trending Packets, Bytes. Information Elements: sourcecountryname, octetdeltacount, packetdeltacount. |
| 7.1. Appendices | 615 |

Table 44: NAT

| Report | Description |
|---------------------|--|
| All Details | A grouping of Source, Src Port, NAT Src IP, NAT Src Port, NAT Dst Port, NAT Dst IP, Dst Port, Destination trending Flows, Bytes. Information Elements: sourceipaddress, sourcetransportport, postnatsourceipv4address, postnaptsourcetransportport, postnaptdestinationtransportport, postnatdestinationipv4address, destinationtransportport, destinationipaddress, octetdeltacount, plixeraggregatedrecordcount. |
| Destination Details | A grouping of Destination, Dst Port, NAT Dst IP, NAT Dst Port trending Flows, Bytes. Information Elements: destinationipaddress, destinationtransportport, postnatdestinationipv4address, postnaptdestinationtransportport, octetdeltacount, plixeraggregatedrecordcount. |
| Dst Translations | A grouping of Destination, Post Dst IP trending Packets, Bytes. Information Elements: destinationipaddress, postnatdestinationipv4address, octetdeltacount, packetdeltacount. |
| Post Connections | A grouping of in Int, Post Src Port, Post Src IP, Post Dst IP, post , out Int trending Packets, Bytes. Information Elements: ingressinterface, postnaptsourcetransportport, postnatsourceipv4address, postnatdestinationipv4address, postnaptdestinationtransportport, egressinterface, octetdeltacount, packetdeltacount. |
| Post Host to Host | A grouping of in Int, Post Src IP, Post Dst IP, out Int trending Packets, Bytes. Information Elements: ingressinterface, postnatsourceipv4address, postnatdestinationipv4address, egressinterface, octetdeltacount, packetdeltacount. |
| Source Details | A grouping of Source, Src Port, NAT Src Port, NAT Src IP trending Flows, Bytes. Information Elements: sourceipaddress, sourcetransportport, postnaptsourcetransportport, postnatsourceipv4address, octetdeltacount, plixeraggregatedrecordcount. |
| Src Translations | A grouping of Source, Post Src IP trending Packets, Bytes. Information Elements: sourceipaddress, postnatsourceipv4address, octetdeltacount, packetdeltacount. |
| Translations | A grouping of Source, Post Src IP, Post Dst IP trending Packets, Bytes. Information Elements: sourceipaddress, destinationipaddress, postnatsourceipv4address, postnatdestinationipv4address, octetdeltacount, packetdeltacount. |

Table 45: NBAR Reports

| Report | Description |
|----------------------------|---|
| Application Categories | A grouping of Application Category trending Packets, Bytes. Information Elements: ciscoapp-categoryname, octetdeltacount, packetdeltacount. |
| Application Compression | A grouping of Application trending % Pkt Comp, % Octet Comp. Information Elements: applicationtag, percentoctetcompression, percentpacketcompression. |
| Application Groups | A grouping of Application Group trending Packets, Bytes. Information Elements: ciscoapp-groupname, octetdeltacount, packetdeltacount. |
| Applications | A grouping of Application trending Packets, Bytes. Information Elements: applicationtag, octetdeltacount, packetdeltacount. |
| Application Sub Categories | A grouping of Application Sub Category trending Packets, Bytes. Information Elements: ciscosubappcategoryname, octetdeltacount, packetdeltacount. |
| Conversations | A grouping of Source, Application, Destination trending Packets, Bytes. Information Elements: sourceipaddress, applicationtag, destinationipaddress, octetdeltacount, packetdeltacount. |

Table 46: Overlay Network

| Report | Description |
|------------------------------|---|
| Destination Hosts by Network | A grouping of Network ID, Network Type, Destination trending Count, Packets, Bytes. Information Elements: overlay_net_id, overlay_net_type, destinationipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Network ID and Type | A grouping of Network ID, Network Type trending Count, Packets, Bytes. Information Elements: overlay_net_id, overlay_net_type, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Source Hosts by Network | A grouping of Network ID, Network Type, Source trending Count, Packets, Bytes. Information Elements: overlay_net_id, overlay_net_type, sourceipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |

Table 47: Pair Reports

| Report | Description |
|------------------------|--|
| AS to AS by IP | A grouping of Source AS, Destination AS trending Packets, Bytes. Information Elements: srcip, dstip, octetdeltacount, packetdeltacount. |
| AS to AS by Tag | A grouping of Src AS, Dst AS trending Packets, Bytes. Information Elements: bgpsourceasnumber, bgpdestinationasnumber, octetdeltacount, packetdeltacount. |
| AS to AS by Tag (Peer) | A grouping of bgpprevadjacentasnumber, bgpNextadjacentasnumber trending Packets, Bytes. Information Elements: bgpprevadjacentasnumber, bgpNextadjacentasnumber, octetdeltacount, packetdeltacount. |
| Avg Pkt Size | A grouping of Source, Destination trending Avg. Pkt. Size, Packets, NULL. Information Elements: sourceipaddress, destinationipaddress, avgpacket-size, octetdeltacount, packetdeltacount. |

continues on next page

Table 47 – continued from previous page

| Report | Description |
|--------------------------|---|
| Client to Server | A grouping of Client, Server trending Packets, Bytes. Information Elements: clientipaddress, serveripaddress, octetdeltacount, packetdeltacount. |
| Connections By Bytes | A grouping of src Port, Source, Protocol, Destination, dst Port trending Packets, Bytes. Information Elements: sourcettransportport, sourceipaddress, protocolidentifier, destinationipaddress, destinationtransportport, octetdeltacount, packetdeltacount. |
| Connections By Flows | A grouping of src Port, Source, Protocol, Destination, dst Port trending Flows. Information Elements: sourcettransportport, sourceipaddress, protocolidentifier, destinationipaddress, destinationtransportport, plixeraggregatedrecordcount. |
| Connections w/ Obsrv Pt. | A grouping of Source, src Port, Destination, dst Port, Obsrv Pt trending Packets, Sum of Sq. Octets. Information Elements: sourceipaddress, sourcettransportport, destinationipaddress, destinationtransportport, observationpointid, octetdeltasumofsquares, packetdeltacount. |
| Conversations App | A grouping of Source, Application, Destination trending Packets, Bytes. Information Elements: sourceipaddress, applicationid, destinationipaddress, octetdeltacount, packetdeltacount. |
| Conversations w/Flags | A grouping of Source IP Address, Well Known Port, tcpcontrolbits, Destination IP Address trending Packets, Bytes, Flows. Information Elements: sourceipaddress, commonport, tcpcontrolbits, destinationipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Conversations WKP | A grouping of Source, Well Known, Destination trending Packets, Bytes. Information Elements: sourceipaddress, commonport, destinationipaddress, octetdeltacount, packetdeltacount. |
| Conv IP Groups | A grouping of Source IP Group, Well Known, Destination IP Group trending Packets, Bytes. Information Elements: srcipgroup, commonport, dstipgroup, octetdeltacount, packetdeltacount. |

continues on next page

Table 47 – continued from previous page

| Report | Description |
|-----------------------|--|
| Country to Country | A grouping of Source Country, Destination Country trending Packets, Bytes. Information Elements: srccountry, dstcountry, octetdeltacount, packetdeltacount. |
| Customer VLAN to VLAN | A grouping of postdot1qcustomervlanid, dot1qcustomervlanid trending Flows, Packets, Bytes. Information Elements: postdot1qcustomervlanid, dot1qcustomervlanid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| dot1q VLAN to VLAN | A grouping of postdot1qvlanid, dot1qvlanid trending Flows, Packets, Bytes. Information Elements: postdot1qvlanid, dot1qvlanid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Flow End Reason | A grouping of Source, Src Port, Dst Port, Destination, Flow End Reason trending Packets, Bytes. Information Elements: sourceipaddress, sourcetransportport, destinationtransportport, destinationipaddress, flowendreason, octetdeltacount, packetdeltacount. |
| Forensic Audit | A grouping of Flow Start, Source, Destination, Common Port, Protocol trending Pkts, Bytes. Information Elements: flowstartmilliseconds, sourceipaddress, destinationipaddress, commonport, protocolidentifier, octetdeltacount, packetdeltacount. |
| Grouped Flows (DSCP) | A grouping of src Port, Source, DSCP, Destination, dst Port trending Packets, Bytes. Information Elements: sourcetransportport, sourceipaddress, ipdiffservcodepoint, destinationipaddress, destinationtransportport, octetdeltacount, packetdeltacount. |
| Grouped Flows (TOS) | A grouping of src Port, Source, Type Of Service, Destination, dst Port trending Packets, Bytes. Information Elements: sourcetransportport, sourceipaddress, ipclassofservice, destinationipaddress, destinationtransportport, octetdeltacount, packetdeltacount. |

continues on next page

Table 47 – continued from previous page

| Report | Description |
|-------------------------|---|
| Host - AS by IP - Host | A grouping of Source, Src AS, Dst AS, Destination trending Flows, Packets, Bytes. Information Elements: sourceipaddress, srcipas, dstipas, destinationipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Host - AS - Host | A grouping of Source, Src AS, Dst AS, Destination trending Flows, Packets, Bytes. Information Elements: sourceipaddress, bgpsourceasnumber, bgpdestinationasnumber, destinationipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Hosts with Country | A grouping of Source, Source Country, Destination, Destination Country trending Packets, Bytes. Information Elements: sourceipaddress, srccountry, destinationipaddress, dstcountry, octetdeltacount, packetdeltacount. |
| Host to Host | A grouping of Source, Destination trending Packets, Bytes. Information Elements: sourceipaddress, destinationipaddress, octetdeltacount, packetdeltacount. |
| Host to Host ICMP | A grouping of Source, Code, Type, Destination trending Count. Information Elements: sourceipaddress, icmpcodeipv4, icmptypeipv4, destinationipaddress, plixeraggregatedrecordcount. |
| Host to Host L2 | A grouping of Source, Destination trending Packets, L2 Octets. Information Elements: sourceipaddress, destinationipaddress, layer2octetdeltacount, packetdeltacount. |
| Host to Host Sum of Sq. | A grouping of Source, Destination trending Packets, Sum of Sq. Octets. Information Elements: sourceipaddress, destinationipaddress, octetdelta-sumofsquares, packetdeltacount. |
| Host to Host w/Flags | A grouping of Source IP Address, tcpcontrolbits, Destination IP Address trending Packets, Bytes, Flows. Information Elements: sourceipaddress, tcpcontrolbits, destinationipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |

continues on next page

Table 47 – continued from previous page

| Report | Description |
|-----------------------------|--|
| Host To Host With Next Hop | A grouping of Source, Destination, Next Hop trending packet, octet. Information Elements: sourceipaddress, destinationipaddress, ipnexthopipv4address, octetdeltacount, packetdeltacount. |
| IP Groups with Apps Defined | A grouping of Src Group, Protocol, Application, Dst Group trending Packets, Bytes. Information Elements: srcipgroup, protocolidentifier, applicationid, dstipgroup, octetdeltacount, packetdeltacount. |
| IP Group to IP Group | A grouping of Source IP Group, Destination IP Group trending Packets, Bytes. Information Elements: srcipgroup, dstipgroup, octetdeltacount, packetdeltacount. |
| MAC to MAC Routed | A grouping of Source MAC, Post Source MAC, Destination MAC, Post Destination MAC trending Flows, Packets, Bytes. Information Elements: sourcemacaddress, postsourcemacaddress, destinationmacaddress, postdestinationmacaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| MAC to MAC Switched | A grouping of Source MAC, Destination MAC trending Packets, Bytes. Information Elements: sourcemacaddress, destinationmacaddress, octetdeltacount, packetdeltacount. |
| Rev 2nd lvl Domain pairs | A grouping of Src Rev 2nd lvl Domain, Dst Rev 2nd lvl Domain trending Packets, Bytes. Information Elements: srcdomain, dstdomain, octetdeltacount, packetdeltacount. |
| Subnet to Subnet | A grouping of Src Subnet, Dst Subnet trending Packets, Bytes. Information Elements: srcnetwork, dstnetwork, octetdeltacount, packetdeltacount. |
| TOS to TOS | A grouping of Type of Service, Post Type of Services trending Packets, Bytes. Information Elements: ipclassofservice, postipclassofservice, octetdeltacount, packetdeltacount. |

continues on next page

Table 47 – continued from previous page

| Report | Description |
|--------------|--|
| VLAN to VLAN | A grouping of postvlanid, vlanid trending Flows, Packets, Bytes. Information Elements: postvlanid, vlanid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |

Table 48: Palo Alto Networks

| Report | Description |
|--------------------------------|--|
| Applications | A grouping of appid_pa trending Packets, Bytes. Information Elements: appid_pa, octetdeltacount, packetdeltacount. |
| CloudGenix Exporter Path Stats | A grouping of Exporter, Path ID trending Down Jitter, Up Jitter, Down Loss, Up Loss, Down MOS, Up MOS, RTT Latency. Information Elements: plixerexporter, cgnxlqmpathidentifier, cgnxlqmdownlinkjittermilliseconds, cgnxlqmdownlinkmos, cgnxlqmdownlinkpacketloss, cgnxlqmrttlatencymilliseconds, cgnxlqmu-plinkjittermilliseconds, cgnxlqmu-plinkmos, cgnxlqmu-plinkpacketloss. |
| CloudGenix Path Stats | A grouping of Path ID trending Down Jitter, Up Jitter, Down Loss, Up Loss, Down MOS, Up MOS, RTT Latency. Information Elements: cgnxlqmpathidentifier, cgnxlqmdownlinkjittermilliseconds, cgnxlqmdownlinkmos, cgnxlqmdownlinkpacketloss, cgnxlqmrttlatencymil-liseconds, cgnxlqmu-plinkjittermilliseconds, cgnxlqmu-plinkmos, cgnxlqmu-plinkpacketloss. |
| Users | A grouping of userid_pa trending Packets, Bytes. Information Elements: userid_pa, octetdelta-count, packetdeltacount. |

Table 49: Procera Reports

| Report | Description |
|---|--|
| APN and Base Service | A grouping of Access Point Name, Base Service trending Bytes. Information Elements: procer-aapn, procerabaseservice, octetdeltacount. |
| Base Service RTT | A grouping of Base Service trending Internal RTT, External RTT. Information Elements: procerabaseservice, proceraexternalrtt, procerainternalrtt. |
| Content Categories | A grouping of Content Categories trending External RTT, Bytes. Information Elements: procer-acontentcategories, octetdeltacount, proceraexternalrtt. |
| HTTP Content Type, Language, and Location | A grouping of Content Type, Language, Location trending Bytes. Information Elements: pro-cerahttpcontenttype, procerahttplanguage, pro-cerahttplocation, octetdeltacount. |
| HTTP Location, Referrer, and Request Method | A grouping of Location, referer, Request Method trending Bytes. Information Elements: procer-ahttplocation, procerahttppreferer, procerahttppre-questmethod, octetdeltacount. |
| HTTP URL, Response Status, User Agent | A grouping of procerahttpurl, Response Status, User Agent trending Bytes. Information Ele-ments: procerahttpurl, procerahttpresponsestatus, procerahttpuseragent, octetdeltacount. |
| Incoming Destination Details | A grouping of Destination IP Address trending Drops, Latency, Packets, Bytes. Informa-tion Elements: destinationipaddress, pro-ceraincomingoctets, proceraincomingpackets, proceraincomingshapingdrops, proceraincoming-shapinglatency. |
| Incoming Source Details | A grouping of Source IP Address trending Drops, Latency, Packets, Bytes. Information El-ements: sourceipaddress, proceraincomingoctets, proceraincomingpackets, proceraincomingshap-ingdrops, proceraincomingshapinglatency. |
| Outgoing Destination Details | A grouping of Destination IP Address trending Drops, Latency, Packets, Bytes. Informa-tion Elements: destinationipaddress, pro-ceraoutgoingoctets, proceraaoutgoingpackets, proceraaoutgoingshapingdrops, proceraaoutgoing-shapinglatency. |
| Outgoing Source Details | A grouping of Source IP Address trending Drops, Latency, Packets, Bytes. Information El-ements: sourceipaddress, proceraaoutgoingpackets, proceraaoutgoingshap-ingdrops, proceraaoutgoingshapinglatency. |
| Property and Service | A grouping of property, service trending In Ext. Count, Latency, Out Ext. Count, Out Latency. |

Table 50: Queue Drops

| Report | Description |
|--------------------------|---|
| Queue Drops By Hierarchy | A grouping of Policy Map Hierarchy, Policy QoS Queue Index trending Flows, Q Drops. Information Elements: policymaphierarchy, policyqosqueueindex, plixeraggregatedrecordcount, plixer_qos_queue_drops. |
| Queue Drops By Index | A grouping of Policy QoS Queue Index trending Flows, Q Drops. Information Elements: policyqosqueueindex, plixeraggregatedrecordcount, plixer_qos_queue_drops. |

Table 51: Replicator

| Report | Description |
|--------------------|---|
| CPU | A grouping of Replicator trending Min, Avg, Max. Information Elements: sourceipaddress, plixercpuutilizationpercent. |
| Profile Statistics | A grouping of Profile trending Pkts In, Pkts Out, Bytes In, Bytes Out. Information Elements: observationdomainname, octetdeltacount, packetdeltacount, postoctetdeltacount, postpacketdeltacount. |

Table 52: Riverbed

| Report | Description |
|------------------------------|--|
| Conversations RTT | A grouping of in Int, Source, Application, Destination, out Int trending RTT. Information Elements: ingressinterface, sourceipaddress, applicationid, destinationipaddress, egressinterface, tcpconnectionrtt_rvbd. |
| FE Type RTT | A grouping of FE Type trending Retrans Bytes, Retrans Pkts, RTT, Packets, Bytes. Information Elements: fetype_rvbd, octetdeltacount, packetdeltacount, tcpconnectionrtt_rvbd, tcppacketretransmissioncount_rvbd, tcpretransmissionbytecount_rvbd. |
| FE Type RTT and Source | A grouping of Source, FE Type trending Retrans Pkts, RTT, Packets, Bytes. Information Elements: sourceipaddress, fetype_rvbd, octetdeltacount, packetdeltacount, tcpconnectionrtt_rvbd, tcppacketretransmissioncount_rvbd. |
| FE Type RTT and Visibility | A grouping of FE Type, Visibility trending Retrans Pkts, RTT, Packets, Bytes. Information Elements: fetype_rvbd, visibility_rvbd, octetdeltacount, packetdeltacount, tcpconnectionrtt_rvbd, tcppacketretransmissioncount_rvbd. |
| Inner Connection IPs and RTT | A grouping of Source, Destination, IC CFE IP, IC SFE IP trending RTT. Information Elements: sourceipaddress, destinationipaddress, innerconnectioncfeipv4address_rvbd, innerconnectionsfeipv4address_rvbd, tcpconnectionrtt_rvbd. |
| Non Optimized Traffic | A grouping of Source, Destination, Common Port, Destination trending Packets, Bytes. Information Elements: sourceipaddress, destinationipaddress, commonport, passthroughreason_rvbd, octetdeltacount, packetdeltacount. |
| Pair RTT and Retrans | A grouping of Source, Destination trending Retrans Pkts, Retrans Bytes, RTT, Packets, Bytes. Information Elements: sourceipaddress, destinationipaddress, octetdeltacount, packetdeltacount, tcpconnectionrtt_rvbd, tcppacketretransmissioncount_rvbd, tcpretransmissionbytecount_rvbd. |
| Pair RTT with Ports | A grouping of Source, Src Port, Dst Port, Destination trending Retrans Pkts, RTT, Packets, Bytes. Information Elements: sourceipaddress, sourcetransportport, destinationtransportport, destinationipaddress, octetdeltacount, packetdeltacount, tcpconnectionrtt_rvbd, tcppacketretransmissioncount_rvbd. |
| 626 | 7. Additional Resources |
| Retransmissions | A grouping of in Int, Source, Destination, out Int trending Pckt Retrans, Bytes Retrans. Information Elements: ingressinterface, sourceipaddress, applicationid, destinationipaddress, egressinterface, tcpconnectionrtt_rvbd, tcppacketretransmissioncount_rvbd, tcpretransmissionbytecount_rvbd. |

Table 53: Saisei

| Report | Description |
|---------------------------|---|
| Dropped Pkts per Int | A grouping of Ingress Int, Distress, Egress Class trending Dropped Octets, Octets, Dropped Packets, Packets. Information Elements: ingressinterface, distress, egressflowclass, droppedoctetttotalcount, droppedpackettotalcount, octetdeltacount, packetdeltacount. |
| Dropped Pkts per User | A grouping of User, Distress, Egress Class trending Dropped Octets, Octets, Dropped Packets, Packets. Information Elements: username, distress, egressflowclass, droppedoctetttotalcount, droppedpackettotalcount, octetdeltacount, packetdeltacount. |
| Forensic Audit | A grouping of User, Application, Egress Class, Flow Start, Flow End trending RTT. Information Elements: username, applicationname, egressflowclass, flowstartmilliseconds, flowendmilliseconds, rttestimate. |
| Pair with Dropped Pkts | A grouping of Source IP, Destination IP, Distress, Egress Class trending Dropped Octets, Octets, Dropped Packets, Packets. Information Elements: sourceipaddress, destinationipaddress, distress, egressflowclass, droppedoctetttotalcount, droppedpackettotalcount, octetdeltacount, packetdeltacount. |
| Pair with Retrans & RTT | A grouping of Source IP, Destination IP, Distress, Egress Class trending Retransmits, Retransmit Events, RTT. Information Elements: sourceipaddress, destinationipaddress, distress, egressflowclass, retransmissiondeltacount, retransmissioneventdeltacount, rttestimate. |
| Retransmits & RTT per Int | A grouping of Ingress Int, Distress, Egress Class trending Retransmits, Retransmit Events, RTT. Information Elements: ingressinterface, distress, egressflowclass, retransmissiondeltacount, retransmissioneventdeltacount, rttestimate. |

Table 54: SonicWALL Reports

| Report | Description |
|----------------------|--|
| App Conv | A grouping of Source, SonicWALL Application, Destination trending Packets, Bytes. Information Elements: sourceipaddress, swapp, destinationipaddress, octetdeltacount, packetdeltacount. |
| Applications | A grouping of SonicWALL Application trending Packets, Bytes. Information Elements: swapp, octetdeltacount, packetdeltacount. |
| Available Memory | A grouping of Exporter trending Available Memory. Information Elements: plixerexporter, mem_avail_ram. |
| CPU Avg. Utilization | A grouping of Core ID trending AVG Util. Information Elements: core_stat_core_id, core_stat_core_util. |
| CPU Max. Utilization | A grouping of Core ID trending MAX Util. Information Elements: core_stat_core_id, core_stat_core_util. |
| Intrusions | A grouping of SonicWALL Intrusion trending Packets, Bytes. Information Elements: flow_to_ips_id, octetdeltacount, packetdeltacount. |
| Spyware | A grouping of SonicWALL Spyware trending Packets, Bytes. Information Elements: flow_to_spyware_id, octetdeltacount, packetdeltacount. |
| Urls | A grouping of SonicWALL URL trending Packets, Bytes. Information Elements: swurl, octetdeltacount, packetdeltacount. |
| User Details | A grouping of SonicWALL User, swuserip, swuserauthtype, swuserdomain trending Packets, Bytes. Information Elements: swuser, swuserip, swuserauthtype, swuserdomain, octetdeltacount, packetdeltacount. |
| Users | A grouping of SonicWALL User trending Packets, Bytes. Information Elements: swuser, octetdeltacount, packetdeltacount. |
| Virus | A grouping of SonicWALL Virus trending Packets, Bytes. Information Elements: flow_to_virus_id, octetdeltacount, packetdeltacount. |
| VoIP Conversations | A grouping of swinitcallid, swrespcallid trending Jitter, Pkt Loss, Packets, Bytes. Information Elements: swinitcallid, swrespcallid, octetdeltacount, packetdeltacount, swvoiplostpkts. |
| VoIP Initiators | A grouping of swinitcallid trending Jitter, Pkt Loss, Packets, Bytes. Information Elements: swinitcallid, octetdeltacount, packetdeltacount, swvoiplostpkts. |

Table 55: Source Reports

| Report | Description |
|---------------------------------|---|
| Autonomous System by IP | A grouping of Source AS trending Packets, Bytes. Information Elements: srcipas, octetdeltacount, packetdeltacount. |
| Autonomous System by Tag | A grouping of Src AS trending Packets, Bytes. Information Elements: bgpsourceasnumber, octetdeltacount, packetdeltacount. |
| Autonomous System by Tag (Peer) | A grouping of bgpprevadjacentasnumber trending Packets, Bytes. Information Elements: bgpprevadjacentasnumber, octetdeltacount, packetdeltacount. |
| Countries | A grouping of Source Country trending Packets, Bytes. Information Elements: srccountry, octetdeltacount, packetdeltacount. |
| Countries with AS | A grouping of Source Country, Source AS, Hosts (Dst) trending Flows, Packets, Bytes. Information Elements: srccountry, srcipas, destinationipaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Customer VLAN | A grouping of dot1qcustomervlanid trending Flows, Packets, Bytes. Information Elements: dot1qcustomervlanid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| dot1q VLAN | A grouping of dot1qvlanid trending Flows, Packets, Bytes. Information Elements: dot1qvlanid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Host Flows | A grouping of Source trending Hosts (Destination), Packets, Flows. Information Elements: sourceipaddress, destinationipaddress, packetdeltacount, plixeraggregatedrecordcount. |
| Host Pkt Length | A grouping of Source trending Length MIN, Length MAX, Length AVG . Information Elements: sourceipaddress, iptotallength. |
| Hosts | A grouping of Source trending Packets, Bytes. Information Elements: sourceipaddress, octetdeltacount, packetdeltacount. |
| ICMP | A grouping of Source, Code, Type trending Count. Information Elements: sourceipaddress, icmpcodeipv4, icmptypeipv4, plixeraggregatedrecordcount. |
| L2 Octets | A grouping of Source trending Packets, L2 Octets. Information Elements: sourceipaddress, |
| 7.1. Appendices | layer2octetdeltacount, packetdeltacount. 629 |
| MAC | A grouping of Source MAC trending Flows, Packets, Bytes. Information Elements: sourcemacaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |

Table 56: Stormshield

| Report | Description |
|--------------------|---|
| Top Url Categories | A grouping of Url Category trending Packets, Bytes. Information Elements: netasqurlcategory, octetdeltacount, packetdeltacount. |
| Top Urls | A grouping of Url trending Packets, Bytes. Information Elements: netasqurl, octetdeltacount, packetdeltacount. |
| Top Users | A grouping of User trending Packets, Bytes. Information Elements: username, octetdeltacount, packetdeltacount. |

Table 57: Top Reports

| Report | Description |
|------------------------|---|
| Applications Defined | A grouping of Application trending Packets, Bytes. Information Elements: applicationid, octetdeltacount, packetdeltacount. |
| Availability By IP | A grouping of Destination IP Address trending Availability. Information Elements: destination-ipaddress, state. |
| Clients | A grouping of Client trending Packets, Bytes. Information Elements: clientipv4address, octetdeltacount, packetdeltacount. |
| DSCP | A grouping of DSCP trending Packets, Bytes. Information Elements: ipdiffservcodepoint, octetdeltacount, packetdeltacount. |
| Exporters | A grouping of Exporter trending Bytes. Information Elements: plixerexporter, octetdeltacount. |
| ICMP Type IPv4 | A grouping of ICMP Type Code trending Count, Packets, NULL. Information Elements: icmp-typecodeipv4, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| ICMP Type IPv6 | A grouping of ICMP Type Code trending Count, Packets, NULL. Information Elements: icmp-typecodeipv6, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| IGMP Type | A grouping of IGMP Type trending Count, Packets, NULL. Information Elements: igmptype, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Interface Compression | A grouping of Exporter, Outbound Interface trending % Pkt Comp, % Octet Comp. Information Elements: plixerexporter, egressinterface, percentoctetcompression, percentpacketcompression. |
| Interface-IP-MAC | A grouping of in Int, Source, Source MAC trending Flows, Packets, Bytes. Information Elements: ingressinterface, sourceipaddress, sourcemacaddress, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Interfaces | A grouping of Exporter, Inbound Interface, Interface Speed trending Bytes, % Util. Information Elements: plixerexporter, ingressinterface, inif-speed, interfacepercent, octetdeltacount. |
| Multicast Destinations | A grouping of Destination trending Pkts, Bytes. Information Elements: destinationipaddress, octetdeltacount, packetdeltacount. |
| 7.1. Appendices | 631 |
| Multicast Pairs | A grouping of Source, Destination trending Pkts, Bytes. Information Elements: sourceipaddress, destinationipaddress, octetdeltacount, packet- |

Table 58: UEBA Reports

| Report | Description |
|------------------------|---|
| Azure User Logins | A grouping of User, Source, Success, Location trending Count. Information Elements: username, sourceipaddress, ipfixifyloginstatename, ipfixifylogsource, plixeraggregatedrecordcount. |
| LDAP User Logins | A grouping of User, Source, Admin trending Count. Information Elements: username, sourceipaddress, ipfixifylogintypename, plixeraggregatedrecordcount. |
| Office 365 User Logins | A grouping of User, Source, Application, Success, Location trending Count. Information Elements: username, sourceipaddress, application-name, ipfixifyloginstatename, ipfixifylogsource, plixeraggregatedrecordcount. |

Table 59: VeloCloud Reports

| Report | Description |
|--------------------------|--|
| Application Flow Path | A grouping of Application, destinationuuid, vcflopath trending Flows. Information Elements: applicationtag, destinationuuid, vcflopath, plixeraggregatedrecordcount. |
| Application Link Policy | A grouping of Application, destinationuuid, vclinkpolicy trending Flows. Information Elements: applicationtag, destinationuuid, vclinkpolicy, plixeraggregatedrecordcount. |
| Application Policies | A grouping of Application, vclinkpolicy, vcroute-type, Traffic Type trending Packets, Bytes. Information Elements: applicationtag, vclinkpolicy, vcroutetype, vctractictype, octetdeltacount, packetdeltacount. |
| Application Priority | A grouping of Application, destinationuuid, vcpriority trending Flows. Information Elements: applicationtag, destinationuuid, vcpriority, plixeraggregatedrecordcount. |
| Application Route Type | A grouping of Application, destinationuuid, vcroutetype trending Flows. Information Elements: applicationtag, destinationuuid, vcroutetype, plixeraggregatedrecordcount. |
| Application Traffic Type | A grouping of Application, destinationuuid, Traffic Type trending Flows. Information Elements: applicationtag, destinationuuid, vctractictype, plixeraggregatedrecordcount. |
| Conv Dst Edge | A grouping of Source, Application, Destination, destinationuuid trending Flows, Bytes. Information Elements: sourceipaddress, applicationtag, destinationipaddress, destinationuuid, octetdeltacount, plixeraggregatedrecordcount. |
| Dst Edge | A grouping of destinationuuid trending Flows, Bytes. Information Elements: destinationuuid, octetdeltacount, plixeraggregatedrecordcount. |
| Flow Path | A grouping of vcflopath trending Flows, Bytes. Information Elements: vcflopath, octetdeltacount, plixeraggregatedrecordcount. |
| Interface Jitter | A grouping of ingressinterface trending count-distinct_destinationipaddress, avg_avgjittertxms. Information Elements: ingressinterface, avgjittertxms, destinationipaddress. |
| Interface Latency | A grouping of ingressinterface trending Unique Dsts, Avg Latency. Information Elements: ingressinterface, avglatencytxms, destinationipaddress. |
| 7.1. Appendices | |
| Interface Metrics | A grouping of ingressinterface trending Avg Latency, avg_avgjittertxms, avg_avglosstxpct. Information Elements: ingressinterface, avglatencytxms, avgavglosstxpct, destinationipaddress. |

Table 60: Viptela Reports

| Report | Description |
|--------------------------|---|
| Health | A grouping of Device Name, Device Model, System IP trending Memory Used, CPU System(%), Disk Used. Information Elements: vtla_host_name, vtla_device_model, vtla_system_ip, vtla_cpu_system, vtla_disk_used, vtla_mem_used. |
| Local Color Performance | A grouping of vEdge Host, Local Color trending Avg. Latency, Avg. Loss, Avg. Jitter. Information Elements: vtla_vdevice_host_name, vtla_local_color, vtla_mean_jitter, vtla_mean_latency, vtla_mean_loss. |
| Policies Added | A grouping of vEdge, Policies Added trending Record Count. Information Elements: vtla_host_name, vtla_policies_added, plixeraggregatedrecordcount. |
| Policies Removed | A grouping of vEdge, Policies Removed trending Record Count. Information Elements: vtla_host_name, vtla_policies_removed, plixeraggregatedrecordcount. |
| Remote Color Performance | A grouping of vEdge Host, Remote Color trending Avg. Latency, Avg. Loss, Avg. Jitter. Information Elements: vtla_vdevice_host_name, vtla_remote_color, vtla_mean_jitter, vtla_mean_latency, vtla_mean_loss. |
| SLA Events | A grouping of Event ID, vEdge, Policies Added, Policies Removed trending Event Count. Information Elements: vtla_id, vtla_host_name, vtla_policies_added, vtla_policies_removed, plixeraggregatedrecordcount. |
| Tunnel Applications | A grouping of vEdge Host, Local Color, Remote System, Remote Color, Application trending Packets, Bytes. Information Elements: vtla_host_name, vtla_local_color, vtla_remote_system_ip, vtla_remote_color, vtla_application, octetdeltacount, packetdeltacount. |
| Tunnel Performance | A grouping of vEdge Host, Local Color, Remote System, Remote Color trending Avg. Latency, Avg. Loss, Avg. Jitter. Information Elements: vtla_vdevice_host_name, vtla_local_color, vtla_remote_system_ip, vtla_remote_color, |
| 634 | vtla_mean_jitter, vtla_mean_loss. 7. Additional Resources |
| vEdge Host Performance | A grouping of vEdge Host trending Avg. Latency, Avg. Loss, Avg. Jitter. Information Elements: vtla_vdevice_host_name, vtla_mean_jitter, |

Table 61: Vitals

| Report | Description |
|-----------------|--|
| CPU | A grouping of Server trending CPU. Information Elements: plixercomponentipaddress, plixercpuutilizationpercent. |
| CPU per Process | A grouping of Process trending min, avg, max. Information Elements: processcommandline, processpercentcpu. |
| Data Ages | A grouping of Source IP Address, timingtest trending Sent. Information Elements: sourceipaddress, timingtest, dataagesseconds. |
| Database | A grouping of Server trending txid, Connections By Bytes, Queries, Timed Checkpoints, Requested Checkpoints, Shared Buffers, Buffers Written. Information Elements: plixercomponentipaddress, buffers_allocd, buffers_written, checkpoints_requested, checkpoints_timed, plixerdbconnections, plixerdbquestions, postgresql_txid. |
| Database | A grouping of Server trending Connections By Bytes, Read Req, Write Req, Cache Free, Queries, Threads, Buffers Used. Information Elements: plixercomponentipaddress, plixerdbconnections, plixerdbkeybufferused, plixerdbkeyreadreq, plixerdbkeywritereq, plixerdbqcache.freemem, plixerdbquestions, plixerdbthreadsconnected. |
| Dir Sizes | A grouping of Collector, Directory trending Bytes. Information Elements: plixercomponentipaddress, plixerstoragedrive, plixerstorageused-bytes. |
| Disk Requests | A grouping of Collector, Drive trending Backlog, Request Wait, Read Merges/Sec, Read Requests/Sec, Request Sectors/Sec, Write Octets/Sec, Write Requests/Sec. Information Elements: plixercomponentipaddress, hddlabel, plixerdiskaveragerequestbacklog, plixerdiskaveragerequestwait, plixerdiskreadrequestmergesps, plixerdiskreadrequestsp, plixerdiskrequestsectorsps, plixerdiskwriterequestmergesps, plixerdiskwriterequestsp. |

continues on next page

Table 61 – continued from previous page

| Report | Description |
|-----------------------------|---|
| Disk Utilization | A grouping of Collector, Drive trending % Utilization, Read Wait, Write Wait, Read Octets/Sec, Write Octets/Sec. Information Elements: plixercomponentipaddress, hddlabel, plixerdiskaveragepercentutilization, plixerdiskaverageawait, plixerdiskaveragewait, plixerdiskreadoctetsps, plixerdiskwriteoctetsps. |
| Distributed Heartbeat | A grouping of Exporter, Plexier Server, Type, Status trending Time. Information Elements: plixercomponentipaddress, ipv4polled, plixerheartbeattype, plixerheartbeatstatus, plixereventdurationmilliseconds. |
| Distributed Synchronization | A grouping of Source, Destination, Caller, DB Table trending Avg Time, Records. Information Elements: syncsourceipv4addr, syncdestinationipv4addr, plixersubroutine, plixertablename, plixereventdurationmilliseconds, plixerrowcount. |
| Event Queue Statistics | A grouping of Collector, DB Table trending Data Age, Total Rows, Disk Used. Information Elements: plixercomponentipaddress, plixertablename, plixerdataageseconds, plixerrowcount, plixerstorageusedbytes. |
| FA Counts | A grouping of Collector, algorithm trending Min, Avg, Max. Information Elements: plixercomponentipaddress, faalgorithmid, faviolationcount. |
| FA Times | A grouping of Collector, algorithm trending Min Dur., Avg Dur., Max Dur.. Information Elements: plixercomponentipaddress, faalgorithmid, plixereventdurationmilliseconds. |
| Flow Metrics/Collector | A grouping of Collector trending MFSN, Packets, Flows. Information Elements: plixercomponentipaddress, plixerflowcount, plixerflowpacketcount, plixermsncount. |
| Flow Metrics/Exporter | A grouping of Collector, Exporter trending MFSN, Packets, Flows. Information Elements: plixercomponentipaddress, plixerexporterid, plixerflowcount, plixerflowpacketcount, plixermsncount. |

continues on next page

Table 61 – continued from previous page

| Report | Description |
|--------------------------------|--|
| Flow Metrics/Port | A grouping of Collector, Port trending MFSN, Packets, Flows. Information Elements: plixercomponentipaddress, plixerlisteningport, plixerflowcount, plixerflowpacketcount, plixermfsncount. |
| Frozen XIDs Age | A grouping of plixercomponentipaddress trending max_postgresql_datfrozenxid_age. Information Elements: plixercomponentipaddress, postgresql_datfrozenxid_age. |
| Frozen XIDs Age by DB | A grouping of plixercomponentipaddress, postgresql_datname trending max_postgresql_datfrozenxid_age. Information Elements: plixercomponentipaddress, postgresql_datname, postgresql_datfrozenxid_age. |
| Memory | A grouping of Server trending Available. Information Elements: plixercomponentipaddress, plixermemavailablebytes. |
| Memory per process | A grouping of Process trending Shared, Resident, Virtual. Information Elements: processcommandline, processresidentmemorysize, processsharedmemorysize, processvirtualmemorysize. |
| ML Engine Heartbeat | A grouping of Exporter, Plexer Server, Type, Status trending Response Time, Data Age. Information Elements: plixercomponentipaddress, plixerexporteripv6address, plixerheartbeattype, plixerheartbeatstatus, dataagesecounds, plixereventdurationmilliseconds. |
| ML Engine Index Document Count | A grouping of ML Engine, Elasticsearch Index trending Avg Count. Information Elements: plixercomponentipaddress, plixermlelasticsearchindexname, plixermlelasticsearchindexcount. |
| ML Engine Kafka Lag | A grouping of ML Engine, Kafka Topic trending Avg Lag. Information Elements: plixercomponentipaddress, plixermlkafkatopicname, plixermlkafkalag. |
| ML Engine Model Count | A grouping of ML Engine trending Avg Model Count. Information Elements: plixercomponentipaddress, plixermlmodelfilecount. |

continues on next page

Table 61 – continued from previous page

| Report | Description |
|------------------------|--|
| PG Lock Count | A grouping of Collector trending Locks. Information Elements: exporteripv4address, postgresql_locks. |
| Report Request Time | A grouping of Collector, reportrequestid, Report Type trending duration. Information Elements: plixercomponentipaddress, reportrequestid, reporttype, plixereventdurationmilliseconds. |
| Report Type Data Time | A grouping of Report Type trending Count, Min Dur., Avg Dur., Max Dur.. Information Elements: reporttype, plixeraggregatedrecordcount, plixereventdurationmilliseconds. |
| Report Type Query Time | A grouping of Report Type trending Count, Min Dur., Avg Dur., Max Dur.. Information Elements: reporttype, plixeraggregatedrecordcount, plixereventdurationmilliseconds. |
| Rollup Counts | A grouping of Exporter, Message Info trending Max Rows. Information Elements: plixerexporterid, message_info, plixerrowcount. |
| Rollup Data Ages | A grouping of Exporter, Template trending Min, Avg, Max. Information Elements: plixerexporterid, plixertemplateid, plixerdataagesseconds. |
| Spool Counts | A grouping of Collector, Directory trending Spool Mins. Information Elements: exporteripv4address, plixerstoragedrive, plixerspoolcount. |
| Storage | A grouping of Server, Drive/Mount trending Avail Bytes. Information Elements: plixercomponentipaddress, plixerstoragedrive, plixerstorageavailablebytes. |
| Stream Age | A grouping of Collector, Stream trending Min Age, Avg Age, Max Age. Information Elements: plixercomponentipaddress, plixertable-name, plixerdataagesseconds. |
| Stream Statistics | A grouping of Collector, Stream trending Data Age, Total Rows, Disk Used. Information Elements: plixercomponentipaddress, plixertable-name, plixerdataagesseconds, plixerrowcount, plixerstorageusedbytes. |

continues on next page

Table 61 – continued from previous page

| Report | Description |
|------------------------|---|
| Syslogs | A grouping of Agent trending Processed, Received. Information Elements: exporteripv4address, plixersyslogsprocessed, plixersyslogsreceived. |
| Task Runtime | A grouping of Collector, Task trending Count, Min Dur., Avg Dur., Max Dur.. Information Elements: plixercomponentipaddress, plixertaskname, plixeraggregatedrecordcount, plixereventdurationmilliseconds. |
| Totals / Rollups Times | A grouping of Exporter, Template, Event, Interval trending Min Rows, Avg Rows, Max Rows, Min Dur., Avg Dur., Max Dur.. Information Elements: plixerexporterid, plixertemplateid, plixereventid, plixerdstintervallength, plixereventdurationmilliseconds, plixerrowcount. |

Table 62: VMware DFW

| Report | Description |
|--------------------------------|--|
| Destination IP, vNIC, FW Event | A grouping of Destination, UUID, vNIC, FW Event, Rule ID trending Flow Count, Packets, Bytes. Information Elements: destinationipaddress, vmuuid, vnicindex, firewallevnt, ruleid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| Source IP, vNIC, FW Event | A grouping of Source, UUID, vNIC, FW Event, Rule ID trending Flow Count, Packets, Bytes. Information Elements: sourceipaddress, vmuuid, vnicindex, firewallevnt, ruleid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |
| UUID, vNIC, FW Event | A grouping of UUID, vNIC, FW Event, Rule ID trending Flow Count, Packets, Bytes. Information Elements: vmuuid, vnicindex, firewallevnt, ruleid, octetdeltacount, packetdeltacount, plixeraggregatedrecordcount. |

Table 63: VMware VDS

| Report | Description |
|----------------------|---|
| Pairs with Tenants | A grouping of Source, Src Tenant, Dst Tenant, Destination, vxLan ID trending Packets, Bytes. Information Elements: sourceipaddress, tenantsourceipv4, tenantdestipv4, destinationipaddress, overlay_net_id, octetdeltacount, packetdeltacount. |
| Tenant Conversations | A grouping of Src Tenant, Src Tenant Port, Tenant Protocol, Dst Tenant Port, Dst Tenant, vxLan ID trending Packets, Bytes. Information Elements: tenantsourceipv4, tenantsourceport, tenantprotocol, tenantdestport, tenantdestipv4, overlay_net_id, octetdeltacount, packetdeltacount. |
| Top Destination | A grouping of Destination, Dst Tenant, Egress Attribute, vxLan ID trending Packets, Bytes. Information Elements: destinationipaddress, tenantdestipv4, egressinterfaceattr, overlay_net_id, octetdeltacount, packetdeltacount. |
| Top Interfaces | A grouping of Ingress Interface, vxLan ID trending Packets, Bytes. Information Elements: ingressinterfaceattr, overlay_net_id, octetdeltacount, packetdeltacount. |
| Top Source | A grouping of Source, Src Tenant, Ingress Attribute, vxLan ID trending Packets, Bytes. Information Elements: sourceipaddress, tenantsourceipv4, ingressinterfaceattr, overlay_net_id, octetdeltacount, packetdeltacount. |

Table 64: Volume Reports

| Report | Description |
|------------------|--|
| Availability | A grouping of Time Stamp, Total trending . Information Elements: goodtime, total, . |
| Flow Volume | Flow rate. As a volume report, the table represents values per time bucket |
| Host Count (dst) | The number of distinct destination hosts. As a volume report, the table represents values per time bucket |
| Host Count (src) | The number of distinct source hosts. As a volume report, the table represents values per time bucket |
| Pair Volume | The number of distinct source/destination pairs. As a volume report, the table represents values per time bucket |
| Round Trip Time | A grouping of Time Stamp, Total trending . Information Elements: goodtime, total, . |
| Traffic Volume | Utilization in bits or bytes along with peak values and 95th percentile. As a volume report, the table represents values per time bucket |

Table 65: Wireless Reports

| Report | Description |
|---|--|
| Applications by Wireless Host | A grouping of Host(s), Application trending Packets, octets. Information Elements: staipv4address, applicationtag, octetdeltacount, packetdeltacount. |
| Applications by Wireless Host with DSCP | A grouping of Host(s), Application, DSCP, Post DSCP trending Packets, octets. Information Elements: staipv4address, applicationtag, ipdiffservcodepoint, postipdiffservcodepoint, octetdeltacount, packetdeltacount. |
| Applications Downstream | A grouping of Application trending Avg. Pkt. Size, Packets, octets. Information Elements: applicationtag, avgpacketsize, octetdeltacount, packetdeltacount. |
| Applications Upstream | A grouping of Application trending Avg. Pkt. Size, Packets, octets. Information Elements: applicationtag, avgpacketsize, octetdeltacount, packetdeltacount. |
| Clients per AP | A grouping of AP trending Clients. Information Elements: wtpmacaddress, stamacaddress. |
| Clients per SSID | A grouping of WLAN SSID trending Clients. Information Elements: wlanssid, stamacaddress. |
| Hosts by SSID | A grouping of Host(s), WLAN SSID trending Packets, octets. Information Elements: staipv4address, wlanssid, octetdeltacount, packetdeltacount. |
| Hosts with MAC | A grouping of Host(s), STA Mac Addr, AP Mac Addr trending Packets, octets. Information Elements: staipv4address, stamacaddress, wtpmacaddress, octetdeltacount, packetdeltacount. |
| Hosts with User Name | A grouping of Source, User Name(s) trending Packets, Bytes. Information Elements: staipv4address, staipname, octetdeltacount, packetdeltacount. |
| Host to Host with AP Mac | A grouping of Source, Destination, AP Mac Addr trending Packets, octets. Information Elements: sourceipaddress, destinationipaddress, wtpmacaddress, octetdeltacount, packetdeltacount. |
| Host to Host with SSID | A grouping of Source, Destination, WLAN SSID trending Packets, octets. Information Elements: sourceipaddress, destinationipaddress, wlanssid, octetdeltacount, packetdeltacount. |
| SSID List | A grouping of WLAN SSID trending Packets, octets. Information Elements: wlanssid, octetdeltacount, packetdeltacount. |
| Usage by SSID and AP | A grouping of AP MAC, WLAN SSID trending Packets, octets, Clients. Information Elements: wtpmacaddress, wlanssid, octetdeltacount, packetdeltacount, clientcount. |

Table 66: Ziften

| Report | Description |
|---------------------|--|
| App Details | A grouping of Application, Version, App Description, Internal Name, File Name, CMD, MD5 trending Flows, Bytes. Information Elements: zflowverproductname, zflowverproductversion, zflowverfiledescription, zflowverinternalname, zflowveroriginalfilename, zflowcommandline, zflowmd5, octetdeltacount, plixeraggregatedrecordcount. |
| Base File and User | A grouping of User Name, Base File, OS trending Flows, Bytes. Information Elements: username, zflowparentimagebasefilename, zflowosname, octetdeltacount, plixeraggregatedrecordcount. |
| Command Line by Src | A grouping of Source, Command Line, PID trending Flows, Bytes. Information Elements: sourceipaddress, zflowcommandline, zflowpid, octetdeltacount, plixeraggregatedrecordcount. |
| Machine Details | A grouping of Machine, User Name, MD5, OS Name, OS Version, Agent UUID trending Flows, Bytes. Information Elements: zflowmachinename, username, zflowmd5, zflowosname, zflowosversion, zflowagentguid, octetdeltacount, plixeraggregatedrecordcount. |
| Machines | A grouping of Machine trending Flows, Bytes. Information Elements: zflowmachinename, octetdeltacount, plixeraggregatedrecordcount. |
| MD5 | A grouping of Parent MD5, Parent Product Name, MD5, zflowverproductname trending Flows. Information Elements: zflowparentmd5, zflowparentverproductname, zflowmd5, zflowverproductname, plixeraggregatedrecordcount. |

7.1.5 Required ports

When deploying the Plixer One platform, refer to the table below and configure firewall rules as necessary.

| Source Component | Destination Component | Protocol | Port | Real-time |
|------------------------------|------------------------------|----------|------------------------------------|-----------|
| All | NTP | UDP | 123 | Time |
| All | DNS Server(s) | UDP | 53 | DNS |
| All Endpoints | Plixer Endpoint Analytics | UDP | 67 | DHCP |
| DNS Server(s) | All | UDP | 53 | DNS |
| Plixer Endpoint Analytics | Exporters | UDP | 161 | SNMP |
| Plixer Endpoint Analytics | SIEM | UDP | 514 | System |
| Plixer Endpoint Analytics | Active Directory Server(s) | TCP | 389,636 | LDAP |
| Plixer Endpoint Analytics | nba.plixer.com | TCP | 443 | Sign |
| Plixer Endpoint Analytics | Tenable IP | TCP | 443 | API |
| Plixer Endpoint Analytics | MS Defender | TCP | 443 | API |
| Exporters | Plixer Scrutinizer Collector | UDP | 2055,2056,4432,4739,9995,9996,6343 | Flow |
| Exporters | Plixer Scrutinizer Collector | UDP | 161 | SNMP |
| Exporters | Plixer Replicator | UDP | 2055,2056,4432,4739,9995,9996,6343 | Flow |
| Exporters | Plixer Endpoint Analytics | UDP | 162 | SNMP |
| Exporters | Plixer Endpoint Analytics | UDP | 161 | SNMP |
| Plixer FlowPro | Flow Collector | UDP | 2055 | Flow |
| Plixer FlowPro | Plixer Replicator | UDP | 2055 | Flow |
| Plixer FlowPro | nba.plixer.com | TCP | 443 | Sign |
| Plixer AD Users Server | Active Directory Server(s) | TCP | 135 | RPC |
| Plixer AD Users Server | Plixer Replicator | UDP | 2055 | Flow |
| Plixer AD Users Server | Plixer Scrutinizer Collector | UDP | 2055 | Flow |
| NTP Server | All | UDP | 123 | Time |
| RADIUS Server(s) | Plixer Endpoint Analytics | UDP | 1813 | RADIUS |
| Plixer Replicator | LDAP Server | TCP | 636 | User |
| Plixer Replicator | Plixer Scrutinizer Collector | UDP | 2055 | Flow |
| Plixer Scrutinizer Collector | Plixer Scrutinizer Reporter | TCP | 22,80,443,5432,6432 | Intra |
| Plixer Scrutinizer Collector | ML | TCP | 22,30404,32000-32002,30323 | Intra |
| Plixer Scrutinizer Collector | Exporters | ICMP | N/A | Up/ |
| Plixer Scrutinizer Collector | AWS S3 Bucket | TCP | 443 | AWS |
| Plixer Scrutinizer Collector | Azure Storage Account | TCP | 443 | Azu |
| Plixer Scrutinizer Collector | Viptela IP | TCP | 8443 | Vipt |
| Plixer Scrutinizer Collector | Exporters | UDP | 161 | SNMP |
| Plixer Scrutinizer Reporter | Plixer Scrutinizer Collector | TCP | 22,80,443,5432,6432 | Intra |
| Plixer Scrutinizer Reporter | ML | TCP | 22,30404,32000-32002,30323,31111 | Intra |
| Plixer Scrutinizer Reporter | Plixer Replicator | TCP | 22,443 | Intra |

Table 67 – continued from previous page

| | | | | |
|-----------------------------|------------------------------|---------|-----------|--|
| Plixer Scrutinizer Reporter | Plixer Endpoint Analytics | TCP | 443 | |
| Plixer Scrutinizer Reporter | Mail Server | TCP | 25,587 | |
| Plixer Scrutinizer Reporter | SIEM | UDP | 514 | |
| Plixer Scrutinizer Reporter | nba.plixer.com | TCP | 443 | |
| Plixer Scrutinizer Reporter | LDAP Server | TCP | 636 | |
| Plixer Scrutinizer Reporter | RADIUS Server | TCP | 1645,1812 | |
| Plixer Scrutinizer Reporter | TACACS+ Server | TCP | 49 | |
| User | Plixer Scrutinizer Reporter | TCP | 443 | |
| User | Endpoint Analytics | TCP | 443 | |
| User | Replicator | TCP | 443 | |
| User | Plixer Scrutinizer Reporter | TCP | 22 | |
| User | Plixer Scrutinizer Collector | TCP | 22 | |
| User | ML Engine | TCP | 22 | |
| User | Plixer Endpoint Analytics | TCP | 22 | |
| User | Plixer FlowPro Sensor | TCP | 22 | |
| User | Plixer Replicator | TCP | 22 | |
| User | ML Engine | TCP | 31112 | |
| User | ML Engine | TCP | 30880 | |
| User | ML Engine | TCP/UDP | 53 | |
| User | ML Engine | TCP | 80 | |

7.2 Changelog

Changelog entries are displayed in the following format:

Description (Ticket Number)

Ex. Thresholds based on outbound traffic (1640)

Note:

- For more information on Plixer Scrutinizer, visit www.plixer.com or contact [Plixer Technical Support](#).
 - Please refer to our End of Life Policy for EOL schedule details.
-

7.2.1 Plexer Scrutinizer changelogs

Select a version below to view the changelog for that release:

Plexer Scrutinizer Version 19.6.1 - June 2025

Fixes

Reparser crashes when sFlow is missing L2 header data in sample (4842)

PDF Report generation not working in 19.6.0 (4863)

Some combinations of protocol exclusions can result in collector crash (4866)

Device Group filter is not displayed correctly (4878)

New deployments don't default to slim navigation (4889)

Saving a user with some missing preferences results in losing all preferences (4915)

Unreadable map labels in dark theme (2754)

Plexer Scrutinizer Version 19.6.0 - March 2025

New Features

“DHCP Servers” and “LDAP Servers” IP Groups for ML Exclusion management

Ability to run reports menu from the host entity view

Ability to specify number of rows in a report gadget

Add support for Azure VNet Flow Logs

Add/Update support for Cisco VXlan IEs

Added additional CloudGenix / PaloAlto SDWAN reports

Additional support for additional Keysight

Admin UI for FlowPro Capture Rules

An interfaces tab to the host entity view when host is an exporter

Audit Report to Admin UI

Collect VPC Flow Logs from Google Cloud Platform

Direct links to Exporters and Interfaces in Explore

Disk space calculator as part of data settings under admin
Edit gadget features from Dashboards
External NAT filter
Full screen mode for Dashboards
Interactive configuration checklist
Interface entity view
LIKE/NOT LIKE filters to Alarm Monitor, Explore, and Admin views
Lollipop chart reporting graph type
MITRE ATT&CK dashboard gadget
ML behavior data in Alarm Monitor workflows
ML Exclusions Admin View
Move feature resources under system performance so it is per server
New Dashboard workflows
New report folder management workflows
New top interfaces dashboard gadget
Option to include full Interface names in reports
Oracle flow log ingestion
Reporting on VXLAN from sFlow samples
Reporting: Line Item gadget type
Ridgeline graph type
Ring Gauge reporting graph type
Scheduled Reports view
Slim navigation mode with vertical navigation bar
Support for TLS v1.3 with LDAP integration
Top Exporter report type
Top N dashboard gadgets
User behavior reports
User configurable horizon in report forecasting
Zabbix client package in our repositories

Enhancements

Add “Last Year” & “This Year” options to report custom time ranges
Add Exclusion workflow from alarm monitor - Front End
Add Google API key from mapping UI

- Add link to host entity view in “Other Options” under report menu
- Add link to user settings in the User menu
- Add option to report on all interfaces bi-directionally
- Admin menu search
- Always search for report types in all report groups
- Auto-expand active filter sections in tray
- Entities: Filter on click from host to alarm
- External Custom Gadget URLs preference
- FA Setting to exclude Internal or External communication for lateral movement
- General Components: Improved severity display
- General Components: Add plixTips when slimNav is collapsed and title tags when expanded
- Include technology in the Alarm monitor view somehow
- Informative detailed error messages when UI can’t communicate with a reporter
- Lateral Movement FA algorithms and preferences
- Mapping automatic grid layout
- Mapping Icons Updated
- Mapping workflows
- New mapping workflows
- Provide interface names in Sankey graph tooltips
- Recategorized system preferences
- Remember selected columns in alarm monitor
- Rename “Favorites” to “Recent” in report menu
- Report description tooltip issues
- Report menu search by description and information element
- Reports: Hidden graph button shouldn’t be there if we don’t have a graph available
- Reports: Make the Saved Reports title clickable
- Saved reports view
- sFlow 801.2ah header support
- Show active filter status for Alarm Monitor and Reporting
- Support newer versions of Cisco ISE for user name reporting
- Workflows in Manage Exporters
- Workflows in Manage Interfaces

Fixes

- Addressed various security issues

‘Client Server’ report failure while filtering for domain (4068)
Ability to edit Flow Analytics Configuration rules in the new UI (4179)
Ability to save a key with an unsupported feature_set (4091)
Adding IP Group with Subnet Rules didn’t save mask selection on initial save (4439)
Admin: Guest users need these routes inaccessible (2243)
Admin: Set LEDs to refresh every 30 seconds and on click (2220)
Admin: Users & Usergroups not respecting routing (2080)
Apache server version is shown in header responses (4326)
Apply button for Single Host-Index search not functioning (1981)
automatic template naming (3154)
Azure NSG exporter naming for distributed collection (3989)
Azure NSG flow log bi-flow support (3993)
Changes made in the oldUI manage interfaces tab are not saved (4301)
Cleaned up Host Index Max Disk Space error when increasing too much (4219)
Collections: No Results Found still shown after creating a collection (2582)
Collector won’t run if the reporter is down (4147)
Copying or refreshing reporting URLs displays error (2108)
Dashboards: map resize on gadget resize (2075)
Date selector doesn’t always allow for shifting the date forward (2236)
Destination AS filter fails on top interfaces report (4309)
Editing an applied filter does not provide enough space (1996)
Expire history failing in some cases after upgrade (4376)
Explore > Exporters > Interfaces should prefer ifAlias over ifDesc (4397)
Exported PDF report has Subscription ending message (1359)
External links using the search route (2105)
Flow Analytics Admin workflow issues (2153)
Flow Version is not visible in new UI (2196)
FlowHopper can’t find flow starting from sFlow (4711)
Full Interface name in reports, sporadically displays (4230)
General Components: Adjust ML graph to take in entire time period, not just data extents (2211)
General Components: app-table pagination skip buttons (2116)
General Components: Inconsistent table header behavior (2047)
Gigamon tcpcontrolbits exceeding smallint value (4374)
Grafana Plugin (443)
Having a “/” in a report name breaks the ability to run that report from an alarm (4327)
Hidden interfaces showing in reports (4100)
Internal Server Error when editing Network Map connections (4063)
Investigate > Host: Learn More button “Host Details” option has no function, related observations (2491)

IP Groups, adding a child group displays wrong selection (4440)
IP V6 import hostfile is broken (4142)
IP/DNS in Flow Analytics Configuration (2014)
IPv6 Exporters don't retain snmp configuration (2865)
Issues when graphing Silverpeak performance reports (microsecond values) (4657)
LDAP login slow with 100K+ group definitions (4038)
Less Than & Greater Than options missing from Advanced Filters (2299)
Manage Exporters & Explore Exporters – Slow / Not loading (4080)
Missing some country codes (4532)
Newline characters in report threshold alarm messages (4037)
Nightly clean all task is removing valid snmp credentials (4419)
Non-Admin users not able to run reports from alarm pages (4263)
Old Host Indexing “first seen” data is unavailable after upgrading. - Import From History Option (4011)
Other options menu opens new window (1965)
Out of file descriptor errors (3941)
Provide description of timeframes for top interfaces and exporters view in Explore (2224)
Recent and Recommended report groups need to show the group details (2536)
Report JSON link returns unnecessary data (2199)
Reports: Saving a report as ‘testSave’ results in ‘Test and Save’ in header name (2445)
Restore Manage Exporters view in the Classic UI (4317)
Saved Reports - Host filters get removed when pivoting (2322)
Scheduled email reports have the license subscription ending soon warning (4082)
Scheduled Traffic Volume reports revert to Line graph when set to Step (3962)
SSL langkey is blank in serverprefs after running set ssl on (4019)
Sync Primary taking too long in some cases (3384)
sysbench package should be installed (4682)
Targets CSV file has ‘violators’ in the name (2194)
Threats Domains temp directory is not always being cleaned out (4395)
Unable to zoom in on TopN report graph (2043)
Update certificate scripts for Oracle Linux (4468)
Usability issues with usergroup permissions in the new UI (1264)
User able to set ‘unlicensed’ in Manage Exporters (4362)
Usergroup Permissions do not carry over to new UI when editing saved reports (1316)
web certificate paths changed back to pre-19.5 location (4359)
When changing an alarm notification frequency to “Rate”, it reverts back to “Each Observation” (3906)

Deprecated

Remove “Additional notes” input from new object form

Plixer Scrutinizer Version 19.5.4 - November 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.
-

New Features

- Support for AWS OL9 AMI
- Support for additional Palo Alto Prisma information elements
- Support for additional Keysight information elements
- FlowPro 20.1 compatibility

Fixes

- Addressed various security issues
- CyberArk Dependencies missing (4497)
- Collector fails to start when the primary reporter is down (4552)
- ML Heartbeat can prevent registering of Plixer ML Engine (4556)

Plixer Scrutinizer Version 19.5.3 - October 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.
-

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)

Plixer Scrutinizer 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.

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 Scrutinizer 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.

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

Plixer Scrutinizer 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.

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)

Plixer Scrutinizer Version 19.4.0 - October 2023

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 Baselining 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)

Version 19.3.2 - September 2023

Plixer Scrutinizer

Fixes

- Addressed various security issues
- AWS Upgrade package dependency problem (3826)

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)

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)

Version 19.2.2 - September 2023

Plixer Scrutinizer

Fixes

- Addressed various security issues
- AWS Upgrade package dependency problem (3826)

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 user groups 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 disable 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)

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)

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):

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

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)

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)

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)

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)

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.

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)

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)

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)

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

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)

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)

7.2.2 Plexer ML Engine changelogs

Select a version below to view the changelog for that release:

Plixer ML Engine Version 19.5.0 - March 2025

New Features

Detect brute force activity using failed SMB logon attempts
Detect remote ransomware attacks using SMB read/write data
Support ability to modify AI Engine deployment settings from the user interface
Support custom thresholds for data accumulation detections
Support KVM
Support offline vSphere cluster deployments

Enhancements

Optimize processing to support larger workloads
Prompt user for deployment information rather than use configuration files
Support ability to modify AI engine application settings from the user interface
Support ability to modify AI Engine seasonality settings from the user interface
Support child groups for IP Group exclusions
Support extending expiration date of ssl certificate used by Kubernetes cluster
Verify application pods are updated during an online system update

Fixes

Addressed various security issues
Define violator and target correctly for brute force events (889)
Update rogue DHCP detection logic (954)

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
- Added LDAP rogue service detection
- Added support for Hyper-V deployments
- Added new alert type in Plixer Scrutinizer for Suricata TLS alerts
- Added 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)

7.3 FAQ

To quickly look up answers to frequently asked questions, select a category:

- [Plixer Scrutinizer](#)
- [Plixer Replicator integration](#)

Important: For additional questions or concerns, contact [Plixer Technical Support](#).

7.3.1 Plixer Scrutinizer

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 > Scrutinizer 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 > System/New User Defaults > 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) Use the [services](#) command at the [scrut_util](#) (SCRUTINIZER>) prompt to stop/start/restart all services or the [systemctl](#) command to manage individual services.

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 `plixer` 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 unlock it:

- In the web interface, go to **Admin > Users & Groups > User Accounts**. Select the locked user-name/account, click the **Authentication Method** tab in the *Edit User* tray, and then change the authentication method from 'locked' to the appropriate method.
- Use the [unlock](#) command at the [scrut_util](#) (SCRUTINIZER>) prompt.

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.
- Use the [expire history](#) command at the [scrut_util](#) (SCRUTINIZER>) prompt.

Q) How do I install/enable VMware Tools or the Hyper-V daemons package?

A) These packages are automatically installed and enabled as part of the initial configuration for new Plexier Scrutinizer deployments. If they were previously disabled, they can be re-enabled by running the following commands:

VMware Tools:

```
sudo systemctl --quiet enable vmtoolsd.service
sudo systemctl --quiet start vmtoolsd.service
```

Hyper-V daemons:

```
sudo systemctl --quiet enable hypervfcopyd.service hypervkvpd.  
↳service hypervvssd.service  
sudo systemctl --quiet start hypervfcopyd.service hypervkvpd.  
↳service hypervvssd.service
```

7.3.2 Plixer Replicator integration

Q: If a new collector is added to the distributed Plixer Scrutinizer cluster, will it automatically be included under auto-replication/load balancing?

A: After adding a new Plixer 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 Plixer Replicator detect the offline collector and automatically re-balance flows?

A: If a Plixer 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 Plexier Scrutinizer log any auto-replicate changes that are made?

A: These changes can be found under `/home/plexier/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 Plexier 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.

7.4 Functional IDs

The Plexier 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 Plexier Scrutinizer installation:

| System Component | Account/ID | Type | Access Level | Function |
|-------------------------|------------|-----------------|----------------|--|
| Operating system | root | Interactive | Privileged | Provides root access to the Plexier Scrutinizer OS, with unrestricted shell, SSH, and console access |
| | plexier | Interactive | Non-privileged | Primary user for the interactive <code>scrut_util</code> CLI utility and provides access to run all Plexier 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 | plexier | 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)

7.5 Localization

Plixer Scrutinizer supports translation of the web interface for localization purposes.

To add or modify translations of UI elements:

1. Navigate to **Admin > Settings > System/New User Defaults > 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`.

7.6 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.

7.6.1 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

Collectors

SIEMs, flow collectors, SNMP trap receivers, and other network management systems that capture, analyze, and report on flow data sent by exporters

EULA (End-User License Agreement)

A legal agreement between Plixer Scrutinizer and the user, outlining the terms and conditions, including usage rights, restrictions, and liability limitations

Exporters

Network devices, such as routers, switches, or servers that can send traffic/activity logs as flows to external systems, such as Plixer Replicator and Plixer Scrutinizer

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

7.6.2 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 un-tagged data sets

7.6.3 General networking terms

2LD (Second-level Domain)

Part of the naming convention for domain names. For example, in example.com, *example* is the second-level domain of the .com TLD (Top level domain)

3LD (Third-level Domain)

For example, in www.mydomain.com, *www* is the third-level domain

ACK (Acknowledgment Code)

A unique signal sent by a computer to show that it has successfully transmitted data

ACL (Access Control List)

A set of rules governing access to a particular object or system resource

Active Directory / AD

Proprietary directory service offered by Microsoft, which allows for centralized management of users, devices, and other IT assets

API (Application Programming Interface)

A software component that allows applications to share data and functionality

ARP (Address Resolution Protocol)

Protocol that maps a dynamic IP address to a physical machine's permanent MAC address in a local area network (LAN)

CA (Certification Authority)

A trusted entity that issues, signs, and stores digital certificates

CDP (Cisco Discovery Protocol)

Protocol used by Cisco devices to allow neighboring networking devices to learn about each other

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

Collector

SIEMs, Flow Collectors, SNMPTrap Receivers, or other network management systems that analyze data forwarded from networked devices

DHCP (Dynamic Host Configuration Protocol)

Network management protocol used to automatically assign IP addresses and other communication parameters to devices on an Internet protocol network

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

Egress

Traffic that exits a device or network

Endpoint

An entity (device, service, node, etc.) at the end of a network communication channel

Encapsulated Remote SPAN (ERSPAN)

Encapsulates mirrored traffic in GRE (Generic Routing Encapsulation) and sends it over Layer 3 networks

ESX (Elastic Sky X)

A pre-configured, ready-to-deploy virtual machine (VM) designed to run on VMware ESX or ESXi

Exporter

A networked device such as a router, switch, or server that generates data and sends it to the flow collector device

Fault tolerance

A system's ability to continue operating without interruptions in the event of hardware or software failure

FQDN (Fully Qualified Domain Name)

The complete address of a computer, host, or any other entity on the Internet

GRE (Generic Routing Encapsulation)

A tunneling protocol developed by Cisco Systems

Hyper-V

A pre-configured, ready-to-deploy virtual machine designed to run on Microsoft Hyper-V, typically packaged in VHD/VHDX format

ICMP (Internet Control Message Protocol)

A protocol used for devices within the network to determine possible network issues

Identity Provider (IdP)

A third-party entity and/or service that stores and manages identities and credentials for use by other websites, applications, or other digital resources

IP address

A unique numerical label assigned to a networked device

IPFIX (Internet Protocol Flow Information Export)

A protocol intended to collect and analyze the flow data from supported network devices

KVM (Kernel-based Virtual Machine)

A pre-configured virtual machine designed to run on KVM hypervisors, packaged in formats like QCOW2 or OVA for easy deployment in Linux-based virtualization environments

Latency

The latency of a network is the time it takes for a data packet to be transferred from its source to the destination

LDAP (Lightweight Directory Access Protocol)

An open, cross-platform protocol used to access and maintain directory services for assets in an Internet protocol network

LLDP (Link Layer Discovery Protocol)

A vendor-neutral protocol used by devices on IEEE 802 networks to advertise their identity, capabilities, and other information

MAC (Media Access Control) address

A unique hardware identifier typically assigned by manufacturers to network adapters and devices

MIB (Management Information Base)

A database that stores information used for managing a 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

Network interface

A (physical or software-based) point of connection between a network entity and the rest of the network

NIC (Network Interface Card)

Adapter that provides devices network connections, either wired or wireless

NID (Network Infrastructure Device)

Any device, such as an access point, router, or switch, that provide the means for entities to communicate with each other over a network

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

Open port

A TCP or UDP port that has been configured to accept packets

OUI (Organizationally Unique Identifier)

A unique 24-bit number in a MAC address that identifies the vendor or the manufacturer of the device

OVF (Open Virtualization Format)

An open source standard for packaging and distributing virtual machines and software applications

Packet

A block of data transmitted across a network

PDU (Protocol Data Unit)

An individual unit of information exchanged by entities on a network using the same protocol

PostgreSQL

An open-source relational database management system (RDBMS) that supports both SQL and JSON querying

PXE (Preboot Execution Environment)

A network booting protocol that allows computers to boot from a network rather than a local storage device like a hard drive or USB

RADIUS (Remote Authentication Dial-In User Service)

A client-server AAA (authentication, authorization, accounting) protocol used to manage remote user access to a network

Redundancy

The state of having duplicate or alternative services as backups to allow for continuous availability

REST API (Representational State Transfer Application Programming Interface)

A set of rules that allows systems to communicate over the web using standard HTTP methods

Router

A device that forwards or routes data packets to devices on a network

Server

A system or device that provides resources, data, services, or applications to other devices over a network

Single Sign-On (SSO)

Allows the integration of third-party authentication services for user access to the Plixer Endpoint Analytics web interface

SIP/RTP (Session Initiation Protocol/Real Time Protocol)

SIP is the control protocol, and RTP is the payload protocol used to send and receive Voice over IP (VoIP)

SNMP (Simple Network Management Protocol)

An IP network protocol used to collect data related to state and/or behavior from devices on a network

SNMP trap

An alert message that is initiated by an SNMP-enabled device to notify the management system of significant events or changes in status

Software agent

A persistent piece of software that performs certain actions and/or interacts with its environment on behalf of a user or another program

SPAN (Switched Port Analyzer)

A dedicated port on a switch that takes a mirrored copy of network traffic from within the switch to be sent to a destination

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

SSL (Secure Sockets Layer)

A protocol for establishing secure connections between networked devices

STIX (Structured Threat Information eXchange)

An industry-standard file format for the exchange of threat information between organizations and platforms

Suricata

A network threat detection engine used to analyze network traffic and identify potential security threats

Switch

A device that connects devices in a network and allows them to communicate with each other

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

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

TAXII (Trusted Automated eXchange of Indicator Information)

A protocol that allows the transmission of threat information, primarily in STIX format, between systems and organizations

TCP (Transmission Control Protocol)

A connection-oriented protocol that enables the bidirectional exchange of messages between devices on the same network

TLS handshake

The process that starts secure communication between a client and a server

TSIG (Transaction Signature)

A protocol that secures DNS packets and allows a Domain Name System to authenticate updates to the DNS database

TTL (Time To Live)

A field in the IP packet header that specifies the maximum number of hops (or router passes) a packet can take before being discarded

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

VoIP (Voice over Internet Protocol)

A technology that allows voice calls using an internet connection

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

Web server banner

A text-based greeting message, which includes information like open ports, services, and version numbers, returned by a web host

7.7 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/`.

7.7.1 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/~mlehmman/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/~mlehmann/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.joelbirch.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.

7.7.2 Plexer Machine Learning

Apache Software License

Cython

<https://cython.org/>

Copyright (c) Robert Bradshaw, Stefan Behnel, Dag Seljebotn, Greg Ewing, et al.

asynpg

<https://github.com/MagicStack/asynpg>

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

7.8 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](#) or contact *Plixer Technical Support*.

Looking for documentation specific to the Classic UI? [Click here](#).