Scrutinizer Documentation

Version 19.1.0

Plixer

October 04, 2021

Scrutinizer

Welcome to the online manual.

Please visit our online webcasts page which includes quick overviews (i.e. 2 - 5 minutes each) of specific features.

Important: Don't struggle, contact Plixer support!

CHAPTER 1

Deployment guides

1.1 Virtual Appliance deployment guide

1.1.1 What you need to know about deploying a Plixer Scrutinizer Virtual Appliance

The Plixer Scrutinizer Virtual Appliance can be obtained from Plixer or your local reseller. It is downloaded as an all-in-one virtual appliance, which can be deployed on an ESXi v5.5 and above or Hyper-V 2012 hypervisor.

- You will need to obtain an appliance license or evaluation license from Plixer or your local reseller for the Plixer Scrutinizer Virtual Appliance to function properly.
- It is recommended to give the Plixer Scrutinizer virtual machine NIC a static MAC address to prevent the machine ID from changing. This is especially important in clustered virtual environments where the VM can change hosts and MAC addresses. If the MAC address changes, the VM will need a new license key.

- The Plixer Scrutinizer Virtual Appliance is deployed on a hypervisor server. It will use 100GB of disk space, 16GB of RAM, and 1 CPU with 4 cores.
- The performance you get out of a Plixer Scrutinizer Virtual Appliance will be directly dependent on the hardware on which it's deployed. It's recommended to dedicate, not share, all the resources that are allocated to the Plixer Scrutinizer virtual machine. This is especially important for the Plixer Scrutinizer datastores. In environments with high volumes of NetFlow data, Plixer Scrutinizer will require dedicated datastores which are discussed in further detail later in this document. Plixer Scrutinizer hardware appliances are recommended for deployments with an exceedingly high volume of flow as they are designed to handle the highest flow rates.
- With the default of 100GB of disk space, you can store up to 1 month of NetFlow v5 data from 25 devices at 1,500 flows a second. If you're planning on exceeding this volume of flow data, or if you need to store data for longer than 30 days, there are detailed steps indicated below that will show you how to *expand the amount of disk space* allocated to the appliance.
- To enable the ability to shut down the Plixer Scrutinizer Virtual Appliance through vSphere, *install VMware Tools* using the instructions in this document. Using the "Power -> Off" method will result in database corruption.

1.1.2 System requirements

The Plixer Scrutinizer Virtual Appliance has the following requirements:

Component	Minimum Specifications (for trial installations)	Recommended Specifications (for production environments)
RAM	16GB	64GB
Disks	100GB	1+ TB 15K RAID 0 or 10 configuration
Processor	1 CPU 4 cores 2GHz+	2 CPUs 8 Cores 2GHz+
Operating System	ESXi 5.5+, Hyper-V 2012, KVM 14	ESXi 6+, Hyper-V 2012, KVM 16

1.1.3 Plixer Scrutinizer OVF deployment on ESX

1. Download the latest Plixer Scrutinizer Virtual Appliance
2. Using VMware vSphere, or vCenter, connect to the ESX host where you will deploy the appliance

3. Right-click a host you would like to deploy the appliance on. Choose the Deploy OVF Template menu option.

4. Select "Local file" and browse to the downloaded Plixer Scrutinizer OVF file and the Plixer Scrutinizer VMDK file, then click "Next".

5. Give your Plixer Scrutinizer VA a name and press "Next".

6. Select an ESX to deploy the machine on if your host is not already selected and press "Next".

7. Review the details of the virtual machine and press "Next".

8. Select your datastore, set your disk format to "Thin Provision" and press "Next".

> **Note:** Be sure to read the *Optimizing Plixer Scrutinizer Datastores* section to obtain the best performance and collection rates.

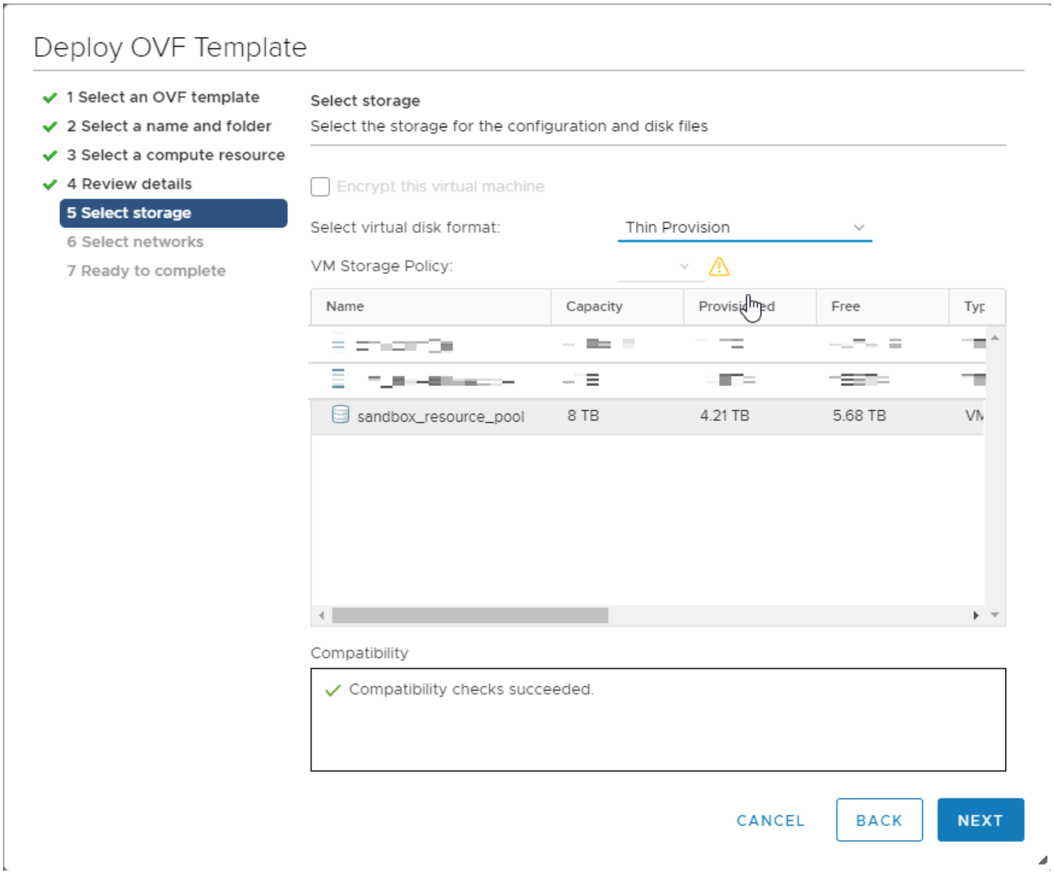

Deploy OVF Template

✓ 1 Select an OVF template
✓ 2 Select a name and folder
✓ 3 Select a compute resource
✓ 4 Review details
5 Select storage
6 Select networks
7 Ready to complete

Select storage

Select the storage for the configuration and disk files

□ Encrypt this virtual machine

Select virtual disk format: Thin Provision

VM Storage Policy: ⚠

Name	Capacity	Provisioned	Free	Typ
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
[redacted]	[redacted]	[redacted]	[redacted]	[redacted]
sandbox_resource_pool	8 TB	4.21 TB	5.68 TB	VM

Compatibility

✓ Compatibility checks succeeded.

9. Select the network to be used by the Plixer Scrutinizer Virtual Appliance.
10. A summary of the options you chose will appear. Click "Finish" and it will import the Plixer Scrutinizer Virtual Appliance. This can take a few moments.
11. Before powering on the Plixer Scrutinizer virtual machine, it's important to set a static MAC address for licensing purposes. Right-click on the Plixer Scrutinizer VM and select "Edit Settings. . . "

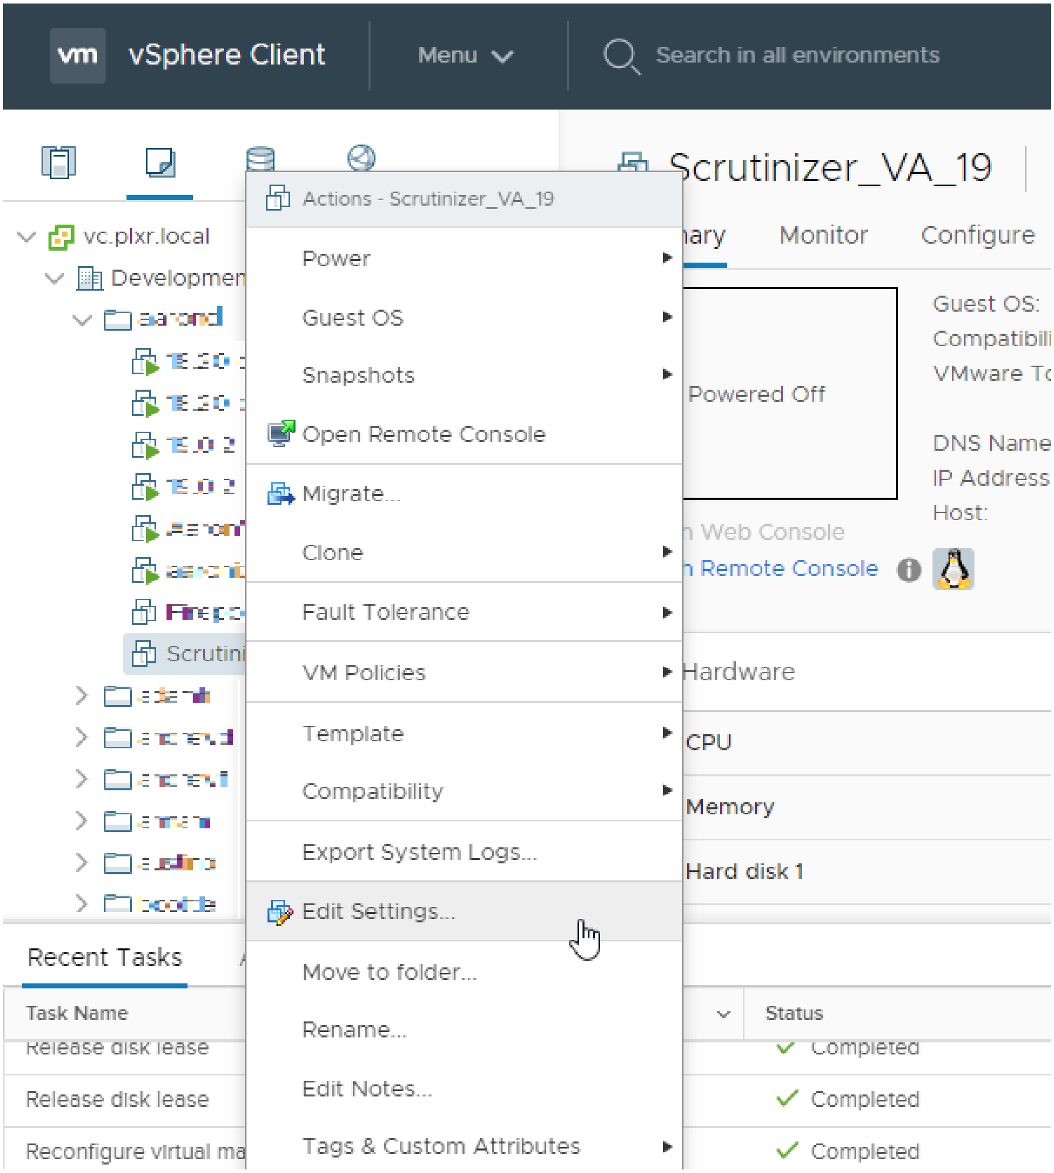

12. Select the Network adapter, set the MAC Address to Manual, enter in a unique MAC Address, and then proceed to the next step.

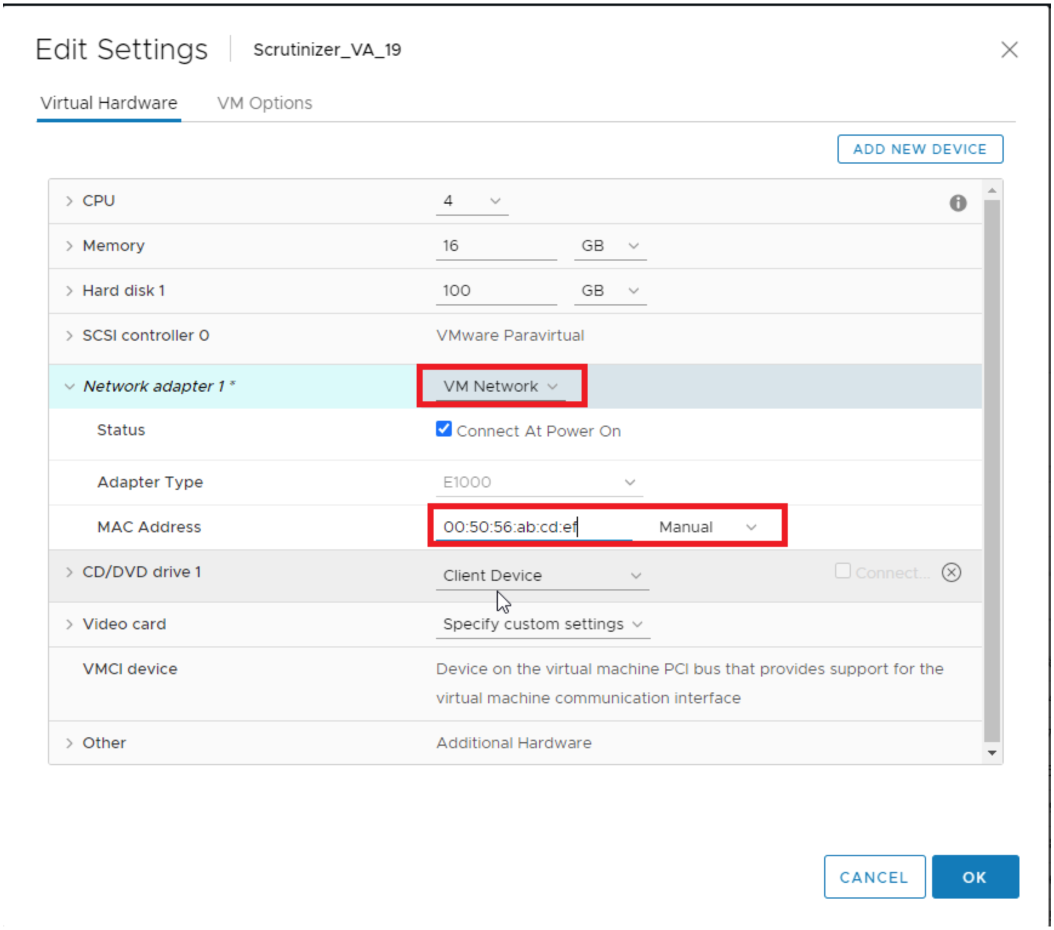

CPU	4	
Memory	16	GB
Hard disk 1	100	GB
SCSI controller 0	VMware Paravirtual	
Network adapter 1 *	VM Network	
Status	☑ Connect At Power On	
Adapter Type	E1000	
MAC Address	00:50:56:ab:cd:ef	Manual
CD/DVD drive 1	Client Device	□ Connect...
Video card	Specify custom settings	
VMCI device	Device on the virtual machine PCI bus that provides support for the virtual machine communication interface	
Other	Additional Hardware	

13. The next step is to allocate and dedicate resources to the Plixer Scrutinizer virtual machine. For evaluation purposes, the Plixer Scrutinizer OVF grabs 1 CPU with 4 cores, 16GB of RAM, and 100GB of disk space.

    When deploying the Plixer Scrutinizer Virtual Appliance it's recommended to increase the resources to meet the recommended *system requirements* listed earlier in this document. Since all installs will vary, more resources may be required.

    Increase the CPU, and Memory settings as necessary (see *system requirements section* for more detail).

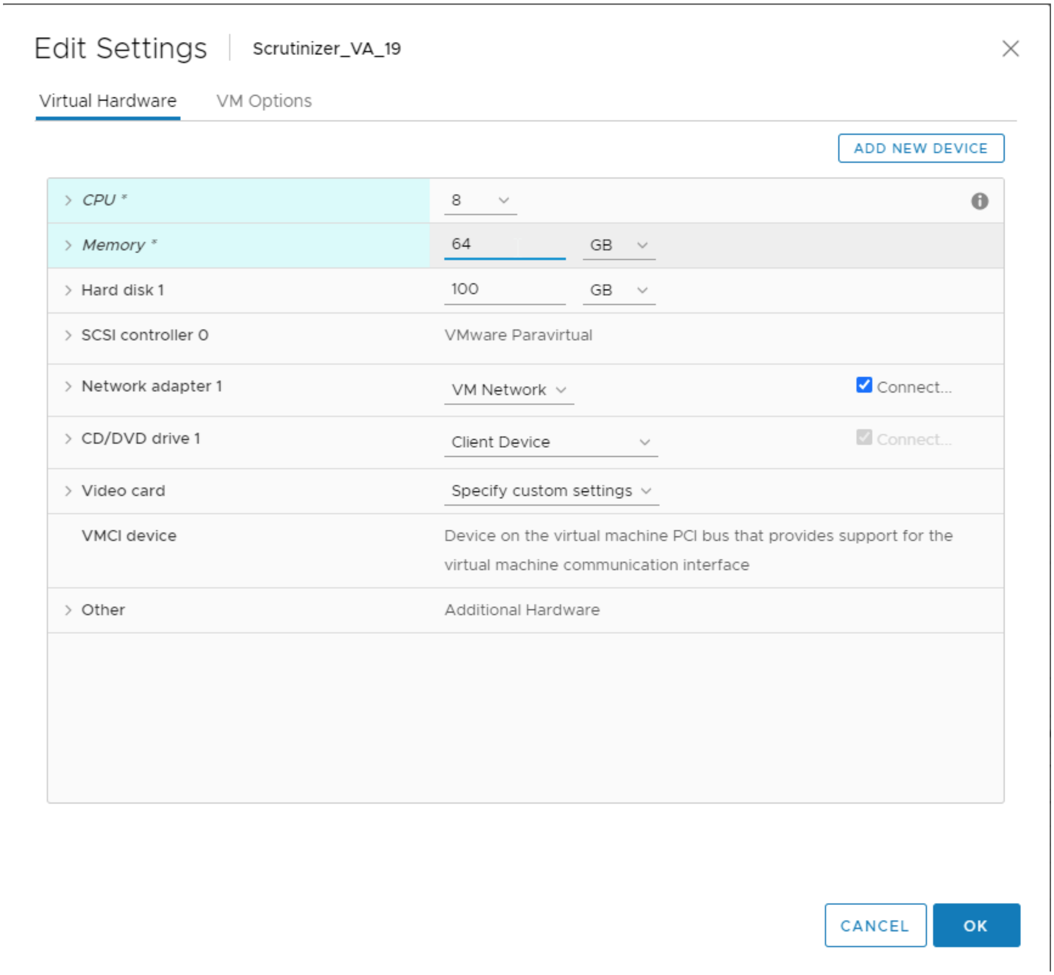

Edit Settings | Scrutinizer_VA_19

> CPU *	8		
> Memory *	64	GB	
> Hard disk 1	100	GB	
> SCSI controller 0	VMware Paravirtual		
> Network adapter 1	VM Network		☑ Connect...
> CD/DVD drive 1	Client Device		☑ Connect...
> Video card	Specify custom settings		
VMCI device	Device on the virtual machine PCI bus that provides support for the virtual machine communication interface		
> Other	Additional Hardware		

14. Next, expand the CPU and Memory sections. Under both,, set the “Shares” value to High and set the “Reservation” maximum value to the number of resources dedicated to the virtual machine. Now press “OK”.

Note: The amount of RAM in the screenshot below is on a small test ESX server, so it won’t match a production install.

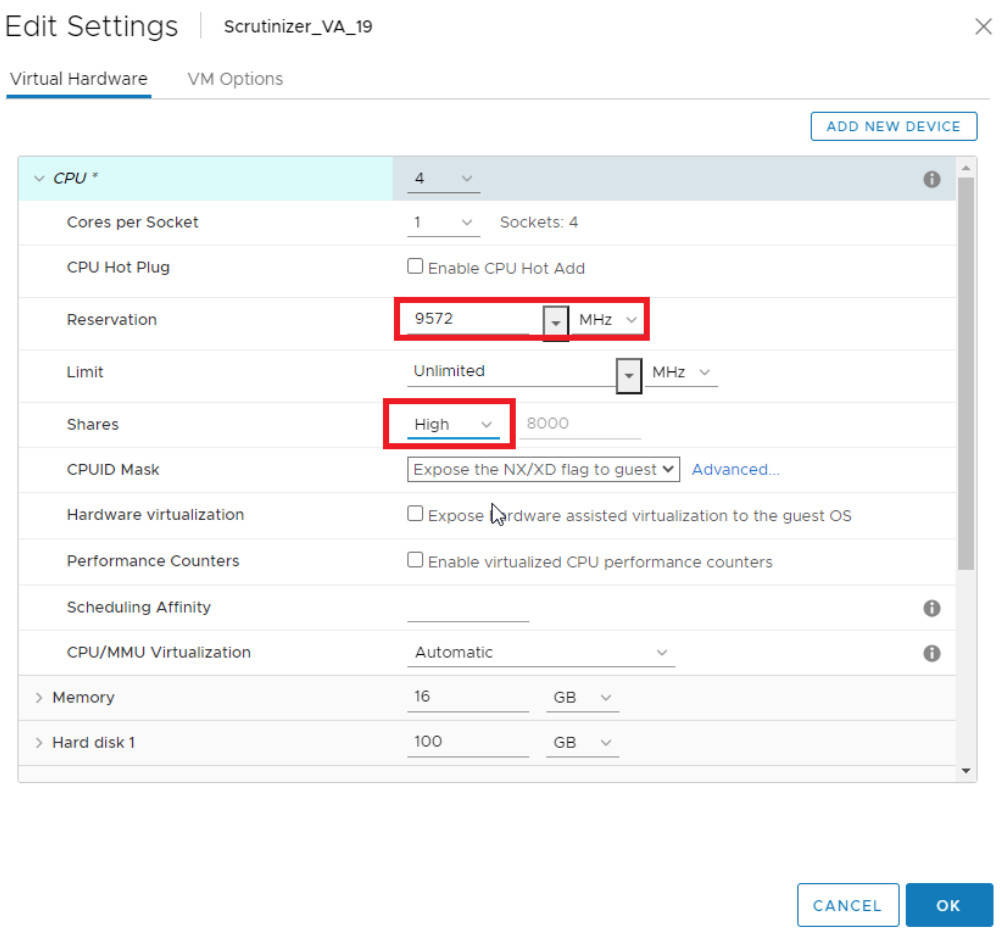

Edit Settings | Scrutinizer_VA_19

CPU *	4
Cores per Socket	1 Sockets: 4
CPU Hot Plug	□ Enable CPU Hot Add
Reservation	9572 MHz
Limit	Unlimited MHz
Shares	High 8000
CPUID Mask	Expose the NX/XD flag to guest Advanced...
Hardware virtualization	□ Expose hardware assisted virtualization to the guest OS
Performance Counters	□ Enable virtualized CPU performance counters
Scheduling Affinity	
CPU/MMU Virtualization	Automatic
Memory	16 GB
Hard disk 1	100 GB

15. Right-click on the Plixer Scrutinizer virtual machine and power it on.

16. Click the console preview window and select “Open Remote Console”. A new window will open and you can then login to the Plixer Scrutinizer Virtual Appliance using **plixer/scrutinizer**.

Note: The server will perform a quick setup and immediately reboot.

17. Log in to the server again and answer the provided questions. Press “Enter” and the server will reboot to apply the necessary settings.

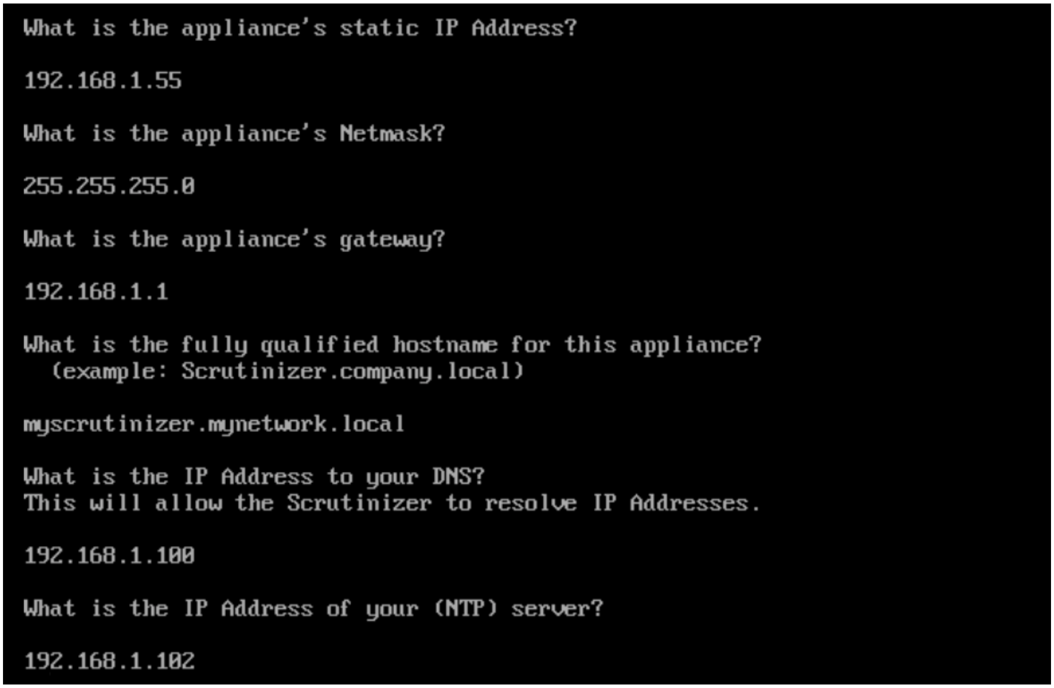

```
What is the appliance's static IP Address?

192.168.1.55

What is the appliance's Netmask?

255.255.255.0

What is the appliance's gateway?

192.168.1.1

What is the fully qualified hostname for this appliance?
  (example: Scrutinizer.company.local)

myscrutinizer.mynetwork.local

What is the IP Address to your DNS?
This will allow the Scrutinizer to resolve IP Addresses.

192.168.1.100

What is the IP Address of your (NTP) server?

192.168.1.102
```

18. Now log in to the Plixer Scrutinizer web interface in your web browser and apply the necessary license key.

Upgrading the Virtual Machine Hardware Version for ESXi

The Plixer Scrutinizer Virtual Appliance is built on Virtual Machine Hardware Version 11 to maintain backwards compatibility with older ESX hypervisors. If you're running vSphere 6.0 or 6.5 you can take advantage of the newer feature sets by upgrading the Virtual Machine Hardware Version as indicated below.

1. While the virtual machine is powered off, in vSphere (or vCenter), right-click on the virtual machine and under the "Compatibility" menu, select "Upgrade VM Compatibility".

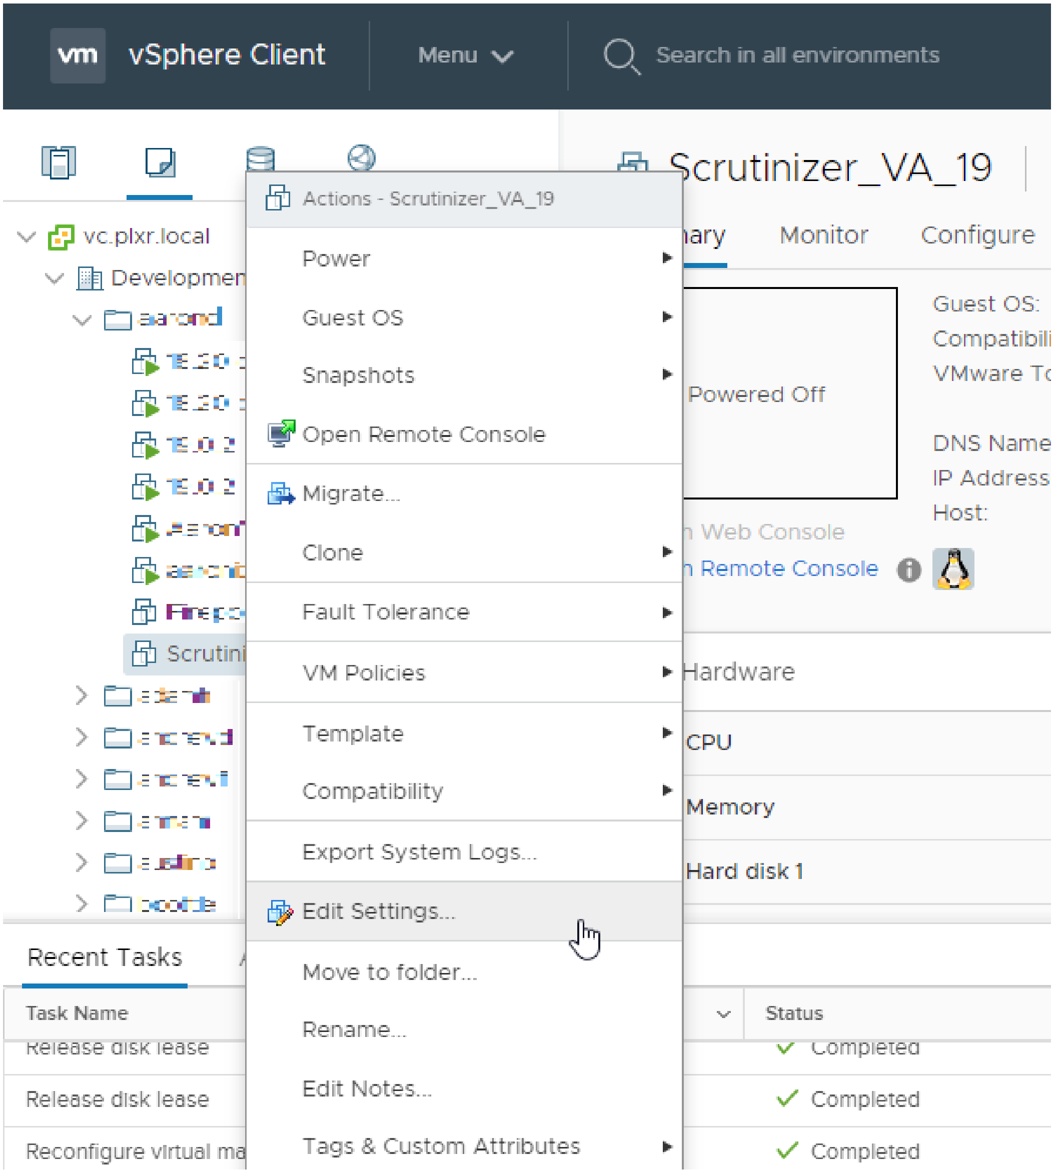

2. Next, power on the virtual machine

Installing VMware Tools for ESXi

After you have gone through the initial Plixer Scrutinizer configuration, you should enable VMware Tools on the appliance. VMware Tools is not installed by default because each version of ESX comes with a different VMware Tools package.

1. Log in to the appliance as the *plixer* user. Use the password you set in the initial deployment.

2. Launch the interactive scrut_util:

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

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

```
SCRUTINIZER> enable vmwaretools
```

4. Once the command completes successfully, type *exit* or *quit* to terminate the interactive prompt.

> **Important:** Installing VMware Tools allows you to properly shut down the Plixer Scrutinizer virtual machine from within vSphere by going to Power > Shut Down Guest.
>
> When shutting down the Plixer Scrutinizer virtual machine, DO NOT select Power > Power Off, as it will result in database corruption. Powering off a virtual machine is equivalent to unplugging a physical computer.

Expanding the database size for ESXi

Depending on the volume of NetFlow data that will be sent to the Plixer Scrutinizer appliance, you may need to expand the size of the database. Expanding the size of the database is a multi-stage process. If you have any questions, please contact Plixer support .

1. Power off the Plixer Scrutinizer virtual machine by logging in and issuing the “sudo shutdown -h now” command.

2. Add an additional hard drive to your Plixer Scrutinizer Virtual Appliance by right-clicking on the Plixer Scrutinizer virtual machine and going to “Edit Settings. . . ”

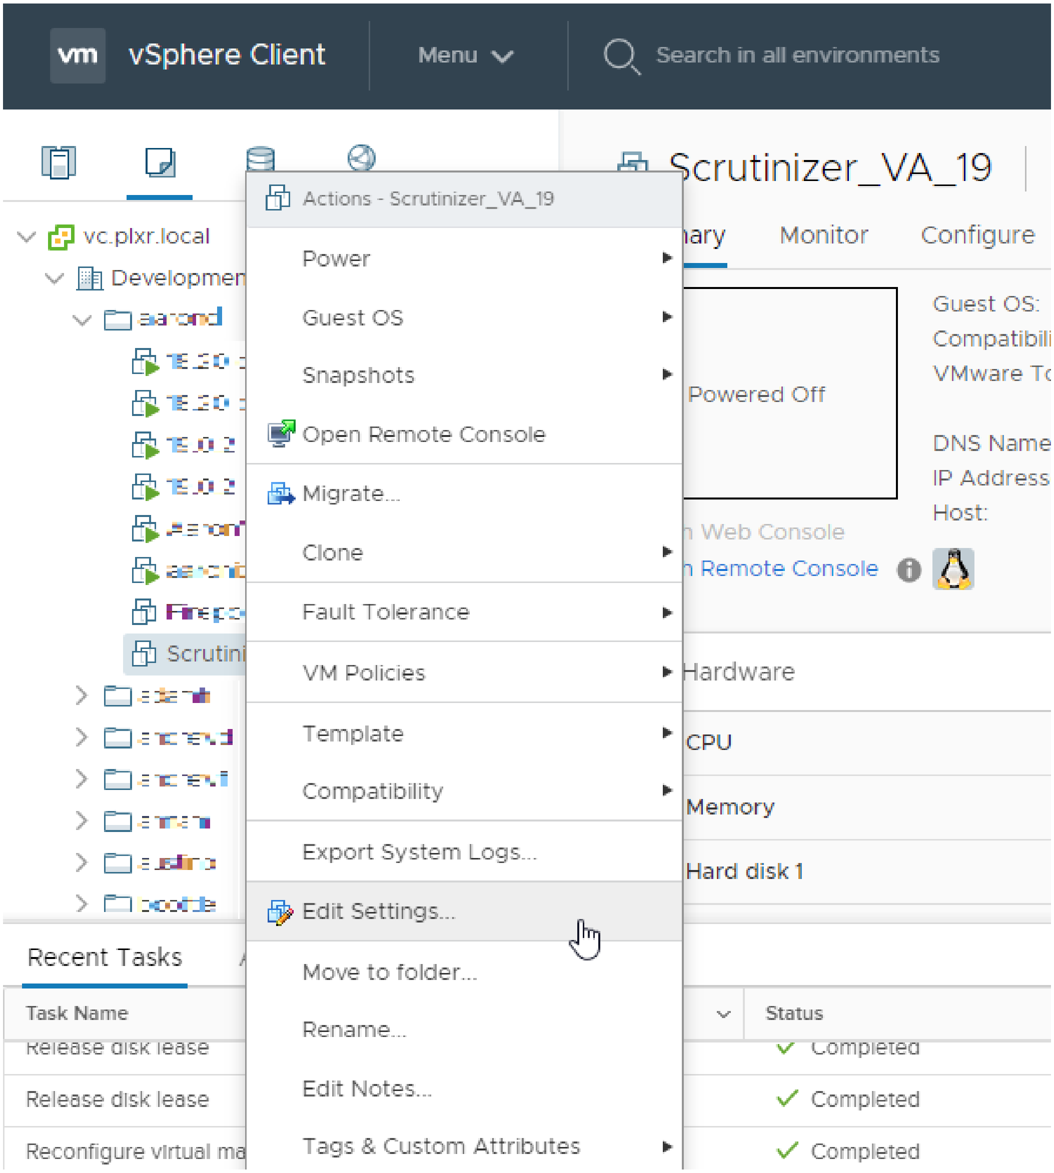

3. Click the “New Device” dropdown and select “New Hard Disk”.

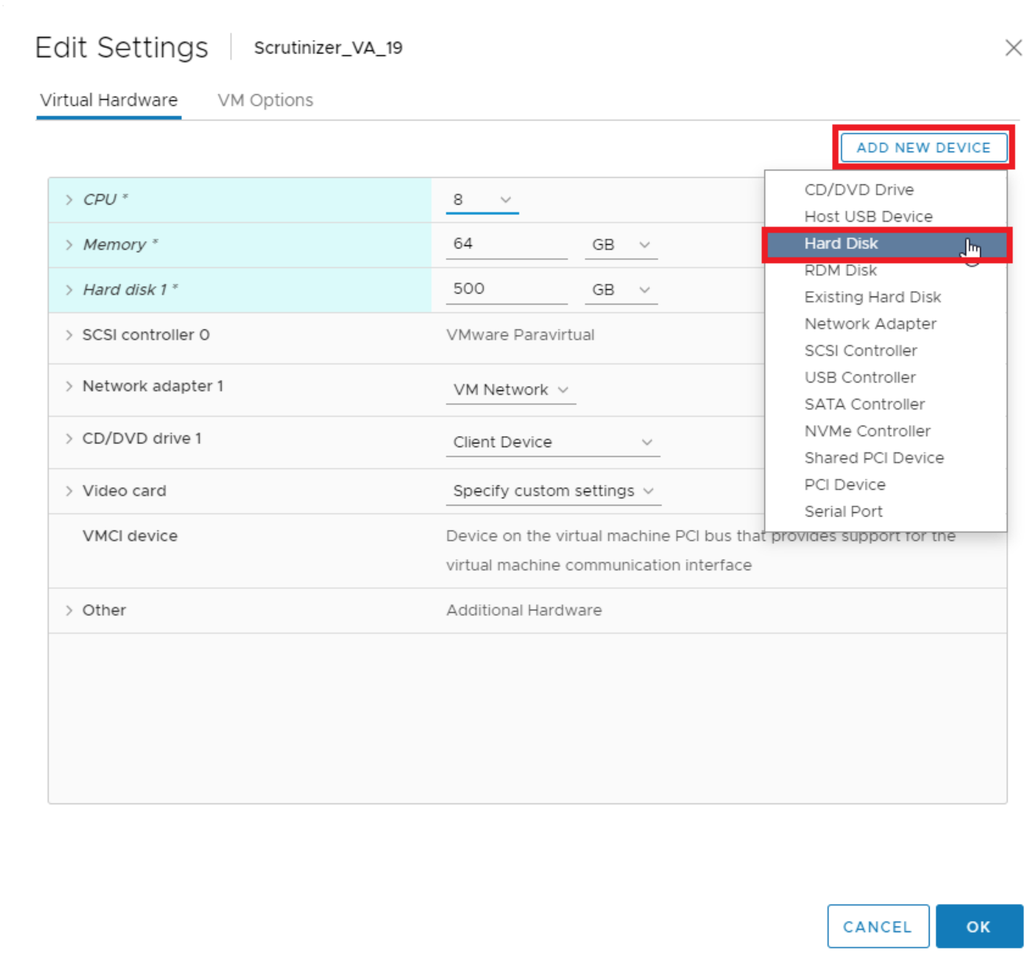

CPU * 8

Memory * 64 GB

Hard disk 1 * 500 GB

SCSI controller 0 VMware Paravirtual

Network adapter 1 VM Network

CD/DVD drive 1 Client Device

Video card Specify custom settings

VMCI device Device on the virtual machine PCI bus that provides support for the virtual machine communication interface

Other Additional Hardware

4. Expand the New Hard disk settings. Choose the type of Disk Provisioning and alter the Capacity of the disk size. Press “OK”.

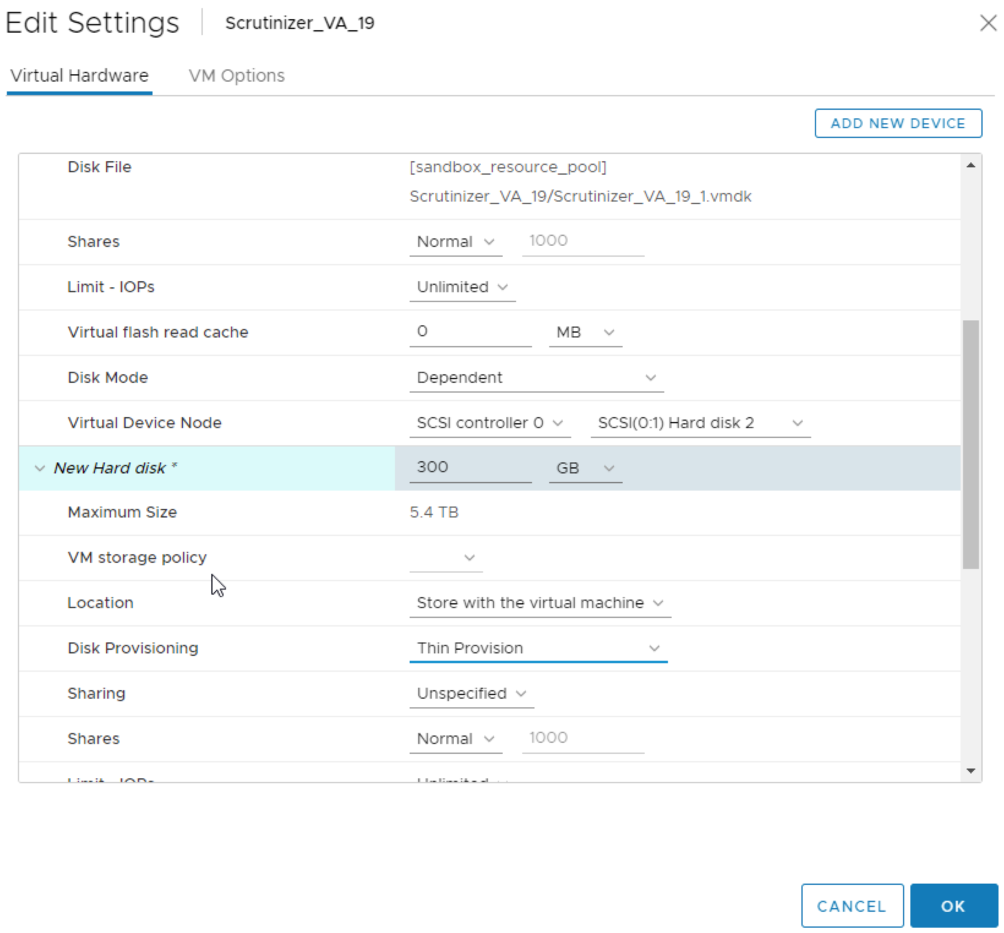

Edit Settings | Scrutinizer_VA_19

Disk File	[sandbox_resource_pool] Scrutinizer_VA_19/Scrutinizer_VA_19_1.vmdk
Shares	Normal 1000
Limit - IOPs	Unlimited
Virtual flash read cache	0 MB
Disk Mode	Dependent
Virtual Device Node	SCSI controller 0 SCSI(0:1) Hard disk 2
New Hard disk *	300 GB
Maximum Size	5.4 TB
VM storage policy	
Location	Store with the virtual machine
Disk Provisioning	Thin Provision
Sharing	Unspecified
Shares	Normal 1000
Limit - IOPs	Unlimited

5. Power on the virtual machine by right-clicking on the Plixer Scrutinizer virtual machine in vSphere. Mouse over to “Power” -> “Power On”.

6. Now that the new hard drive is added, we have to resize the volume group, the partition volume, and the file system so that Plixer Scrutinizer can use the newly allocated space.

   - Start by logging in to the Plixer Scrutinizer Virtual Appliance as the ‘plixer’ user
   - Start the Scrutinizer interative utility by running ‘scrut_util’
   - Type ‘show diskspace’ to view the current size of the database, which is mounted on /var/db. This is the current size of disk before we add the new space.

- Type ‘show partitions’ and make note of the disk in use for the newly added space.

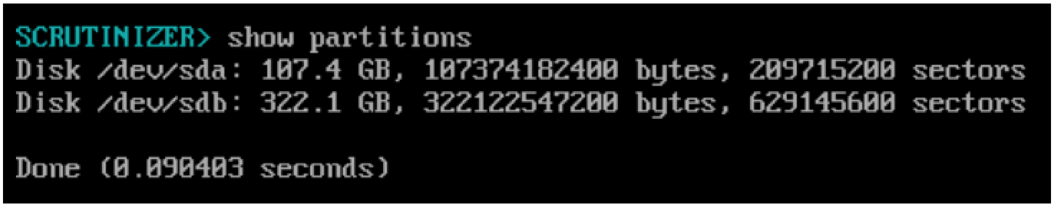

```
SCRUTINIZER> show partitions
Disk /dev/sda: 107.4 GB, 107374182400 bytes, 209715200 sectors
Disk /dev/sdb: 322.1 GB, 322122547200 bytes, 629145600 sectors

Done (0.090403 seconds)
```

7. Now that we know the disk to use, we can run a command to use the newly added space. There will be an interactive prompt to follow. One of the questions asked is if you have taken a backup of your data before proceeding.

   - Type ‘set partitions /dev/sd[from above]’

   - In the example in this guide, /dev/sdb is the correct partition.

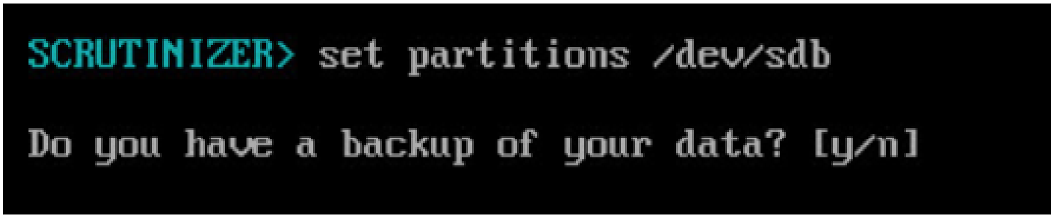

```
SCRUTINIZER> set partitions /dev/sdb

Do you have a backup of your data? [y/n]
```

   - Confirm that the new diskspace was added to the volume group.

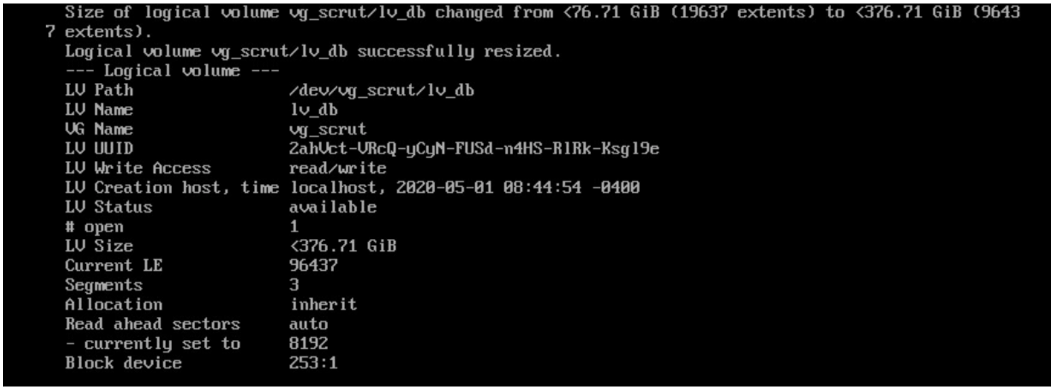

```
  Size of logical volume vg_scrut/lv_db changed from <76.71 GiB (19637 extents) to <376.71 GiB (9643
 7 extents).
  Logical volume vg_scrut/lv_db successfully resized.
  --- Logical volume ---
  LV Path                /dev/vg_scrut/lv_db
  LV Name                lv_db
  VG Name                vg_scrut
  LV UUID                2ahVct-VRcQ-yCyN-FUSd-n4HS-R1Rk-Ksg19e
  LV Write Access        read/write
  LV Creation host, time localhost, 2020-05-01 08:44:54 -0400
  LV Status              available
  # open                 1
  LV Size                <376.71 GiB
  Current LE             96437
  Segments               3
  Allocation             inherit
  Read ahead sectors     auto
  - currently set to     8192
  Block device           253:1
```

   - The next step will be automatic, please be patient. When it’s finished, you can run ‘show diskspace’ and see the new size of the files system mounted on /var/db

1.1.4 Plixer Scrutinizer deployment on Hyper-V

1. Download the latest Plixer Scrutinizer Virtual Appliance
2. Unzip the file on your Hyper-V server
3. Open Hyper-V Manager and select Import Virtual Machine

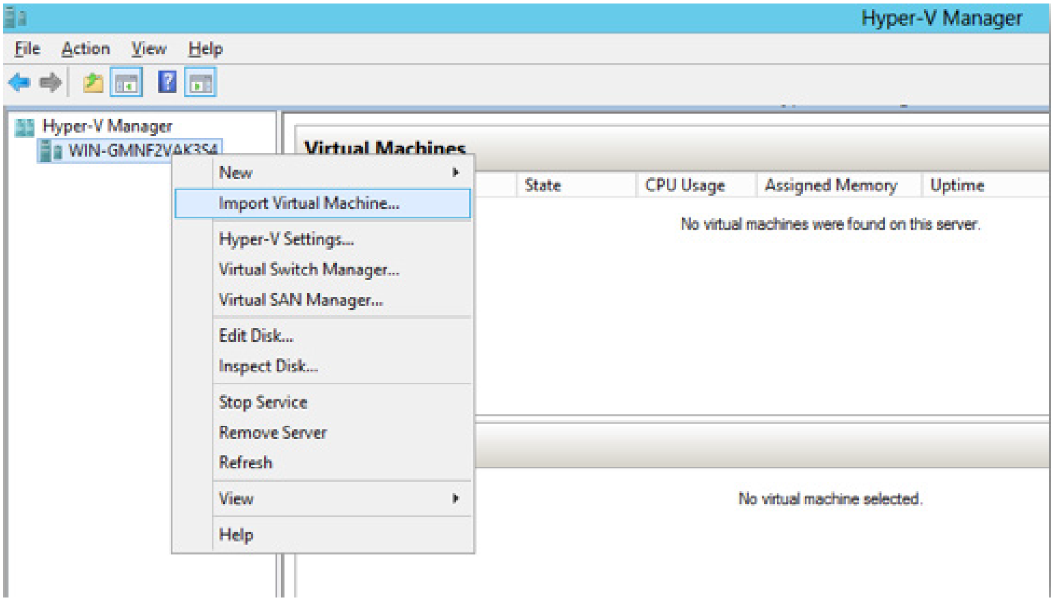

4. Specify the Scrutinizer_Hyper-V folder

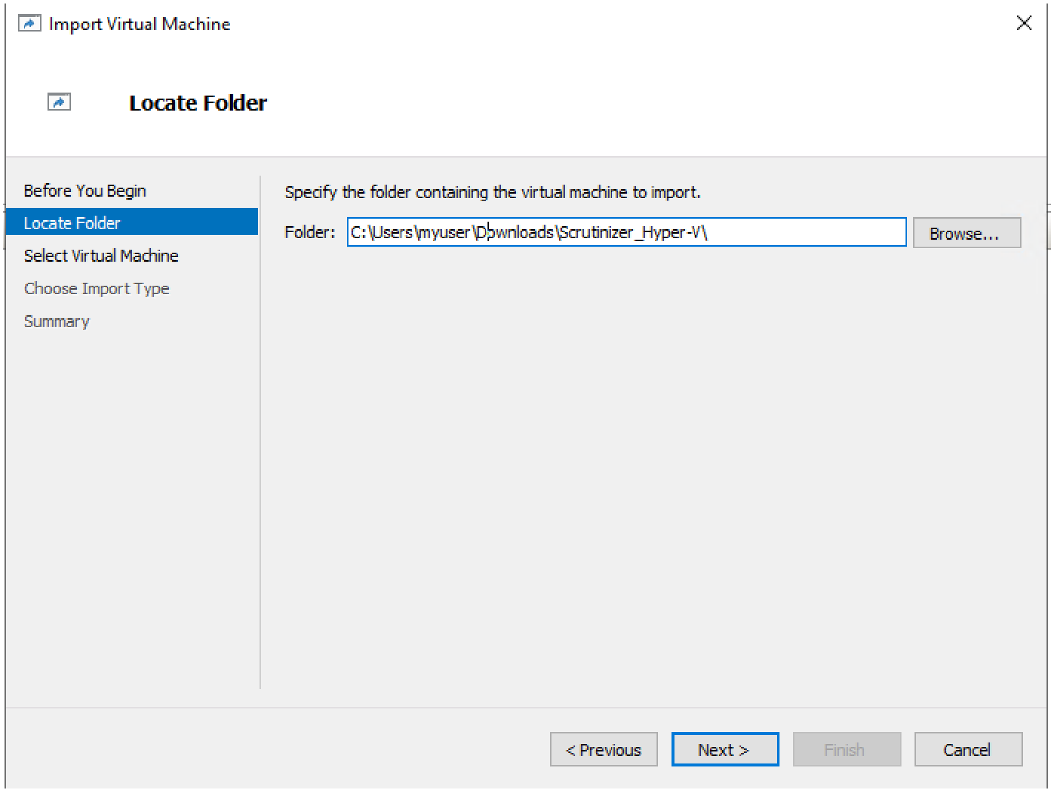

Specify the folder containing the virtual machine to import.

Folder: C:\Users\myuser\Downloads\Scrutinizer_Hyper-V\

5. Select the Virtual Machine

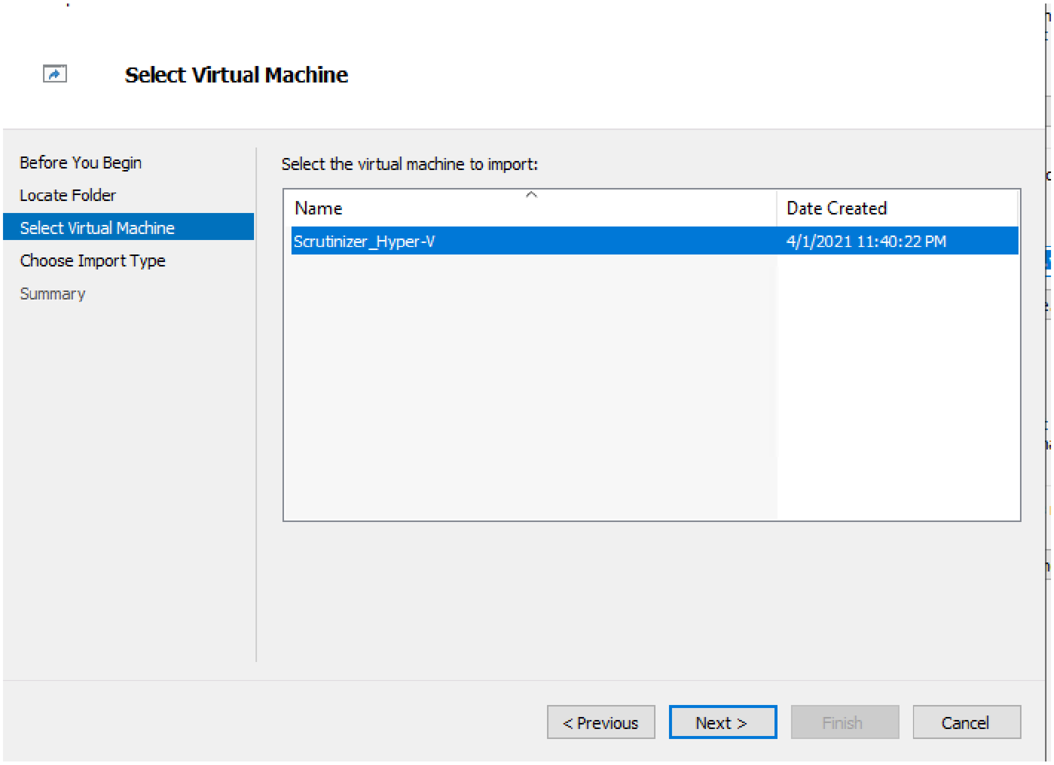

Select the virtual machine to import:

Name	Date Created
Scrutinizer_Hyper-V	4/1/2021 11:40:22 PM

6. Choose Import Type

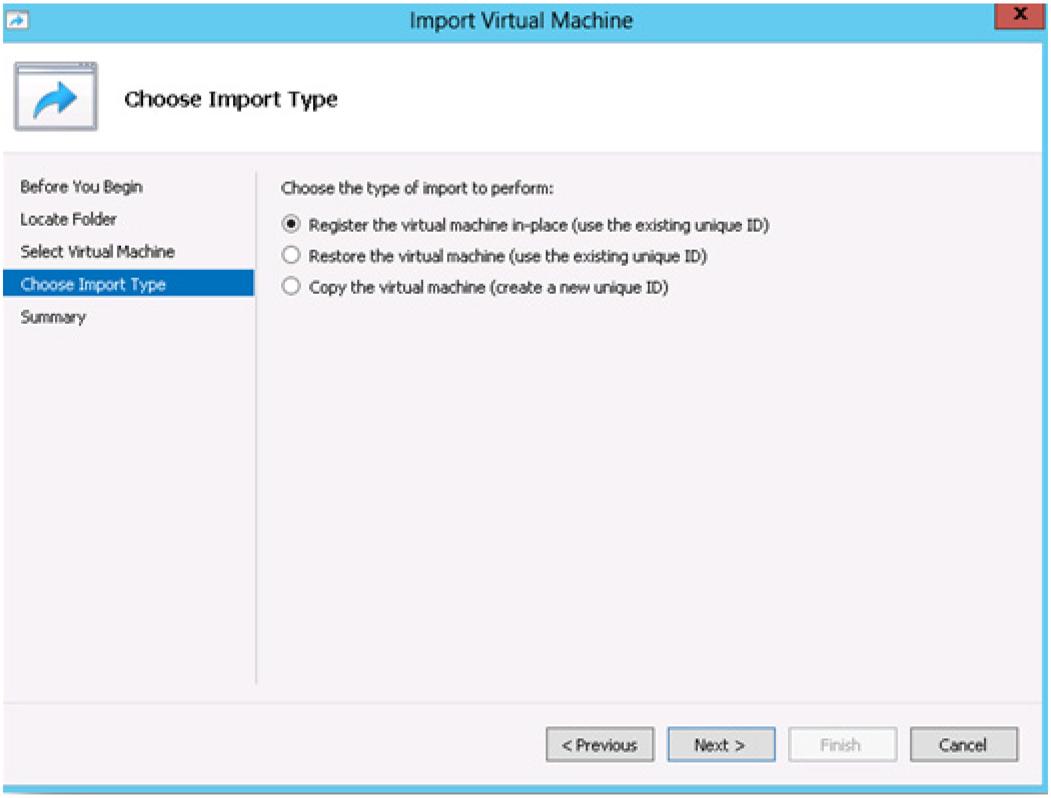

Import Virtual Machine

Choose Import Type

Choose the type of import to perform:

● Register the virtual machine in-place (use the existing unique ID)

○ Restore the virtual machine (use the existing unique ID)

○ Copy the virtual machine (create a new unique ID)

7. Go to Settings

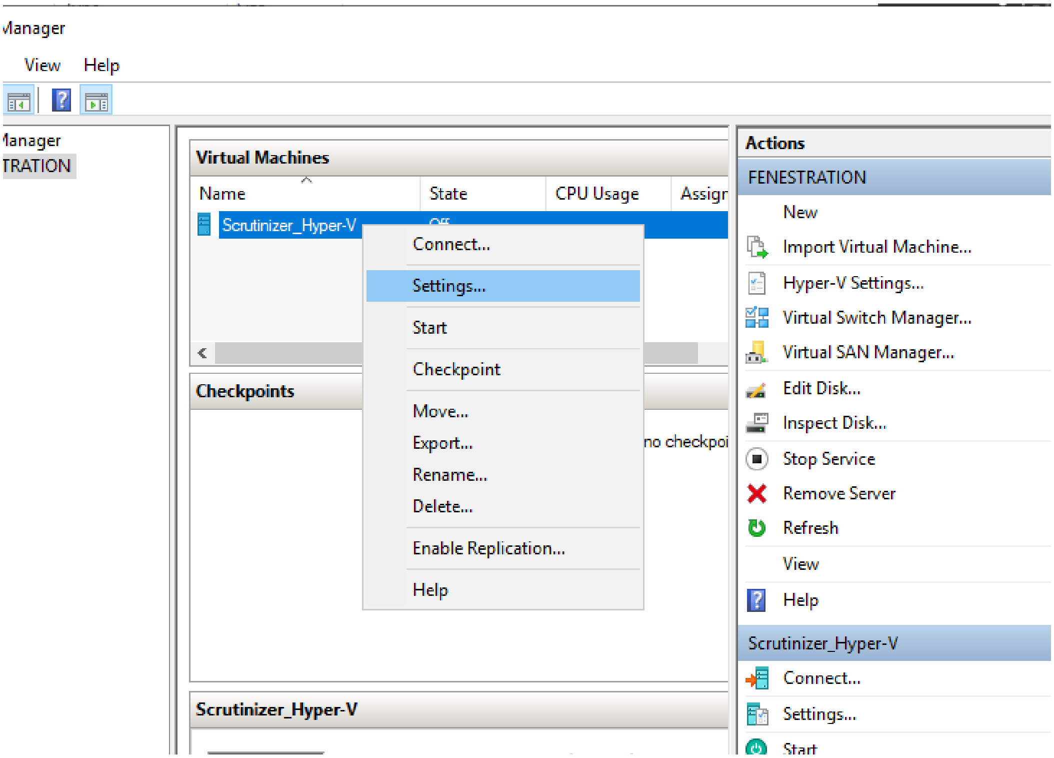

8. Make sure the memory is set to 16GB.

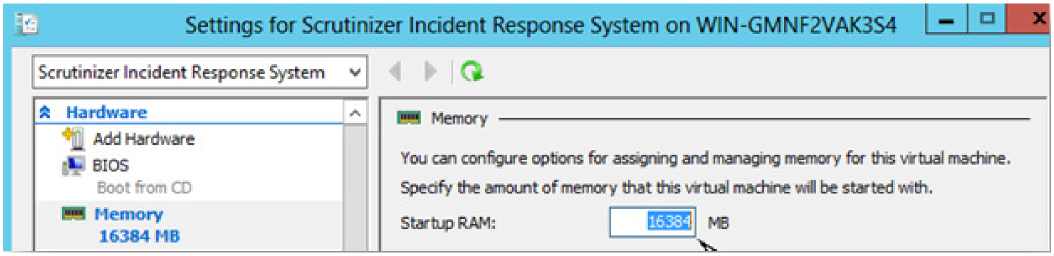
Settings for Scrutinizer Incident Response System on WIN-GMNF2VAK3S4
Scrutinizer Incident Response System
Hardware
Add Hardware
BIOS
Boot from CD
Memory
16384 MB

Memory

You can configure options for assigning and managing memory for this virtual machine.

Specify the amount of memory that this virtual machine will be started with.

Startup RAM: 16384 MB

9. Select your Network Adapter and assign it to the appropriate Virtual Switch.

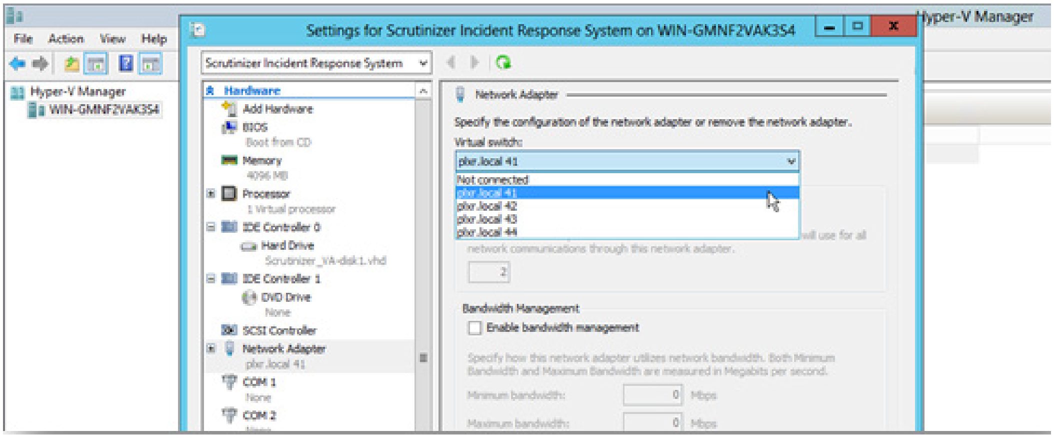

Network Adapter

Specify the configuration of the network adapter or remove the network adapter.

Virtual switch: plxr.local 41

- Not connected
- plxr.local 41
- plxr.local 42
- plxr.local 43
- plxr.local 44

...will use for all network communications through this network adapter.

2

Bandwidth Management

□ Enable bandwidth management

Specify how this network adapter utilizes network bandwidth. Both Minimum Bandwidth and Maximum Bandwidth are measured in Megabits per second.

Minimum bandwidth: 0 Mbps

Maximum bandwidth: 0 Mbps

10. Expand the Network Adapter section, select Advanced Features, set the MAC Address to Static, enter a unique MAC Address, and then press "OK".

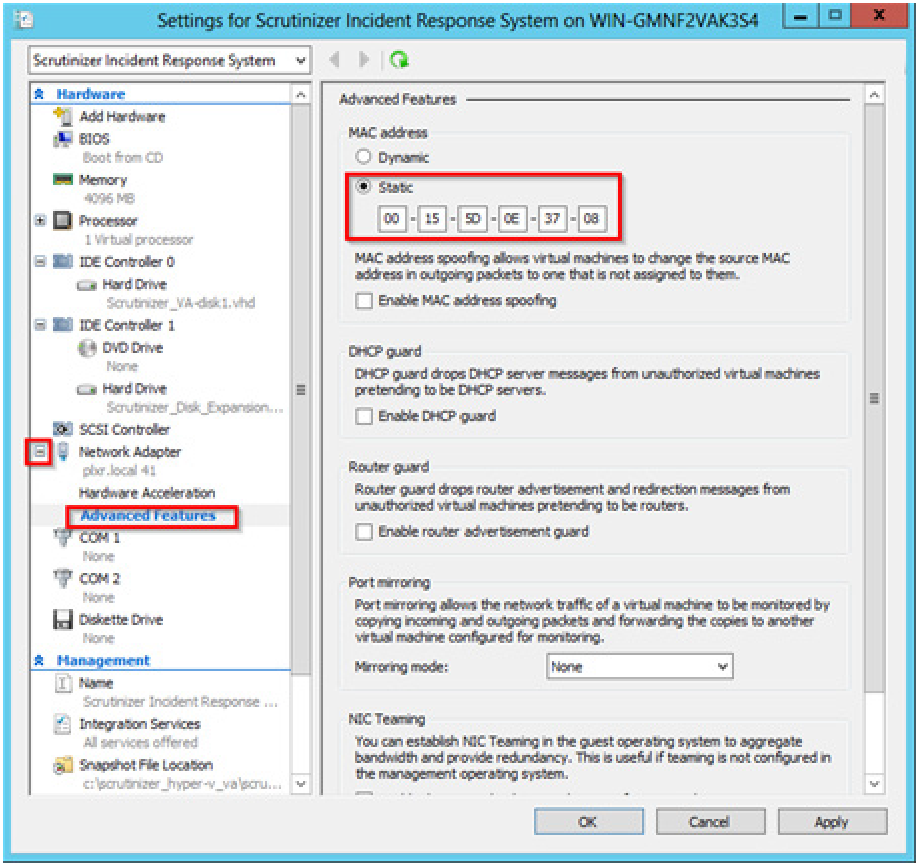

Advanced Features

MAC address

○ Dynamic

● Static

00 - 15 - 5D - 0E - 37 - 08

MAC address spoofing allows virtual machines to change the source MAC address in outgoing packets to one that is not assigned to them.

□ Enable MAC address spoofing

DHCP guard

DHCP guard drops DHCP server messages from unauthorized virtual machines pretending to be DHCP servers.

□ Enable DHCP guard

Router guard

Router guard drops router advertisement and redirection messages from unauthorized virtual machines pretending to be routers.

□ Enable router advertisement guard

Port mirroring

Port mirroring allows the network traffic of a virtual machine to be monitored by copying incoming and outgoing packets and forwarding the copies to another virtual machine configured for monitoring.

Mirroring mode: None

NIC Teaming

You can establish NIC Teaming in the guest operating system to aggregate bandwidth and provide redundancy. This is useful if teaming is not configured in the management operating system.

11. Start the Virtual Machine.

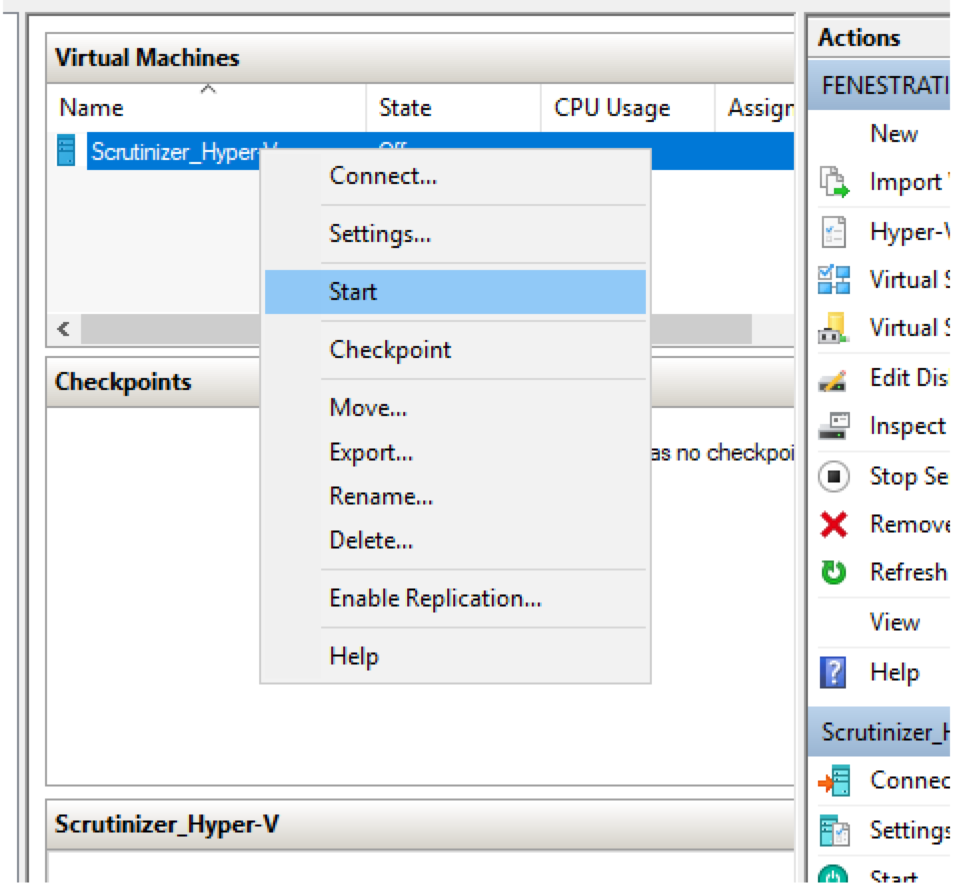

12. Right-click on the Virtual Machine and click Connect to log in to the Plixer Scrutinizer Virtual Appliance using plixer/scrutinizer.

13. Log in to the server again and answer the provided questions. Press "Enter" and the server will reboot to apply the necessary settings.

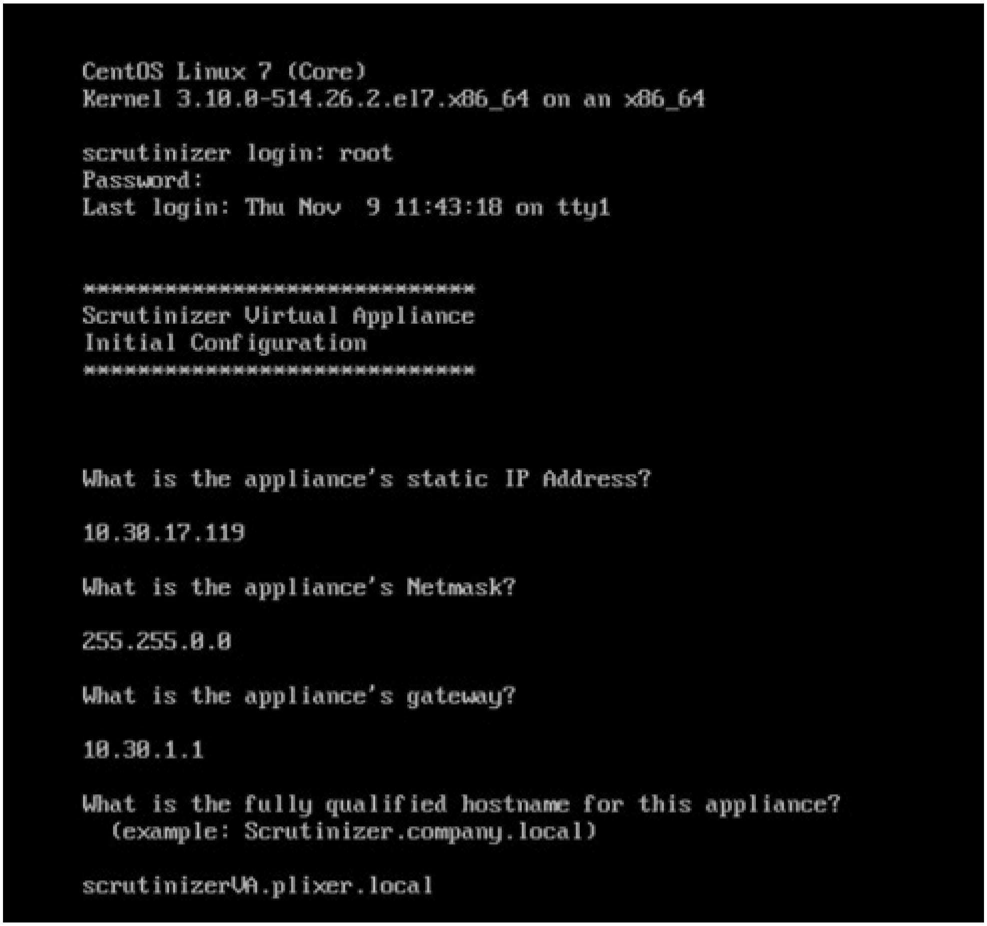

```
CentOS Linux 7 (Core)
Kernel 3.10.0-514.26.2.el7.x86_64 on an x86_64

scrutinizer login: root
Password:
Last login: Thu Nov  9 11:43:18 on tty1


*****************************
Scrutinizer Virtual Appliance
Initial Configuration
*****************************



What is the appliance's static IP Address?

10.30.17.119

What is the appliance's Netmask?

255.255.0.0

What is the appliance's gateway?

10.30.1.1

What is the fully qualified hostname for this appliance?
  (example: Scrutinizer.company.local)

scrutinizerVA.plixer.local
```

14. Now log in to the Plixer Scrutinizer web interface in your web browser and apply the necessary license key.

Expanding the database size for Hyper-V

Depending on the volume of NetFlow data that will be sent to the Plixer Scrutinizer appliance, you may need to expand the size of the database. Expanding the size of the database is a multi-stage process. If you have any questions, please contact your support representative.

1. Power off the Plixer Scrutinizer virtual machine by logging in and issuing the “shutdown -h now” command.

2. In the Hyper-V Manager, right-click on the Plixer Scrutinizer virtual machine and select “Settings”.

3. Next, select the IDE Controller and click “Add” to a hard drive.

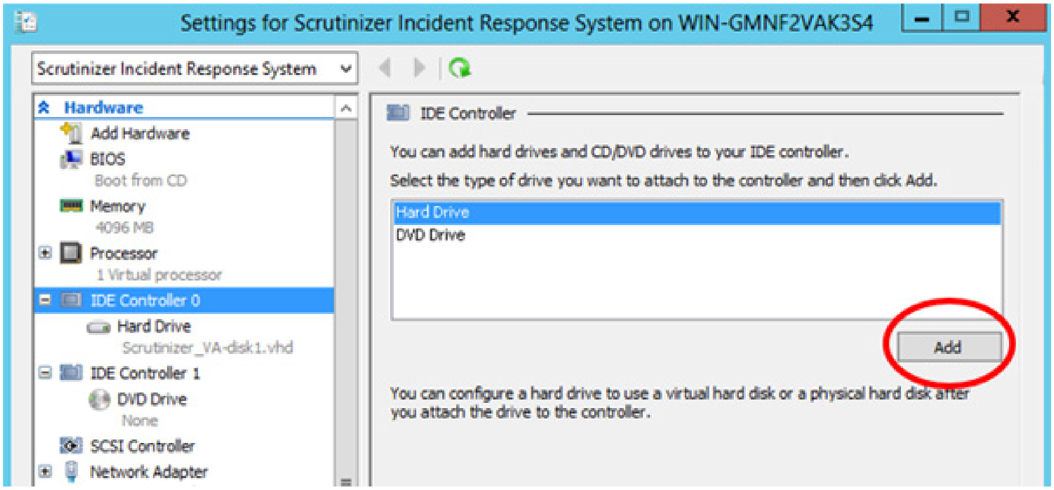

Settings for Scrutinizer Incident Response System on WIN-GMNF2VAK3S4

IDE Controller

You can add hard drives and CD/DVD drives to your IDE controller.

Select the type of drive you want to attach to the controller and then click Add.

You can configure a hard drive to use a virtual hard disk or a physical hard disk after you attach the drive to the controller.

4. Under Virtual hard disk, select “New”.

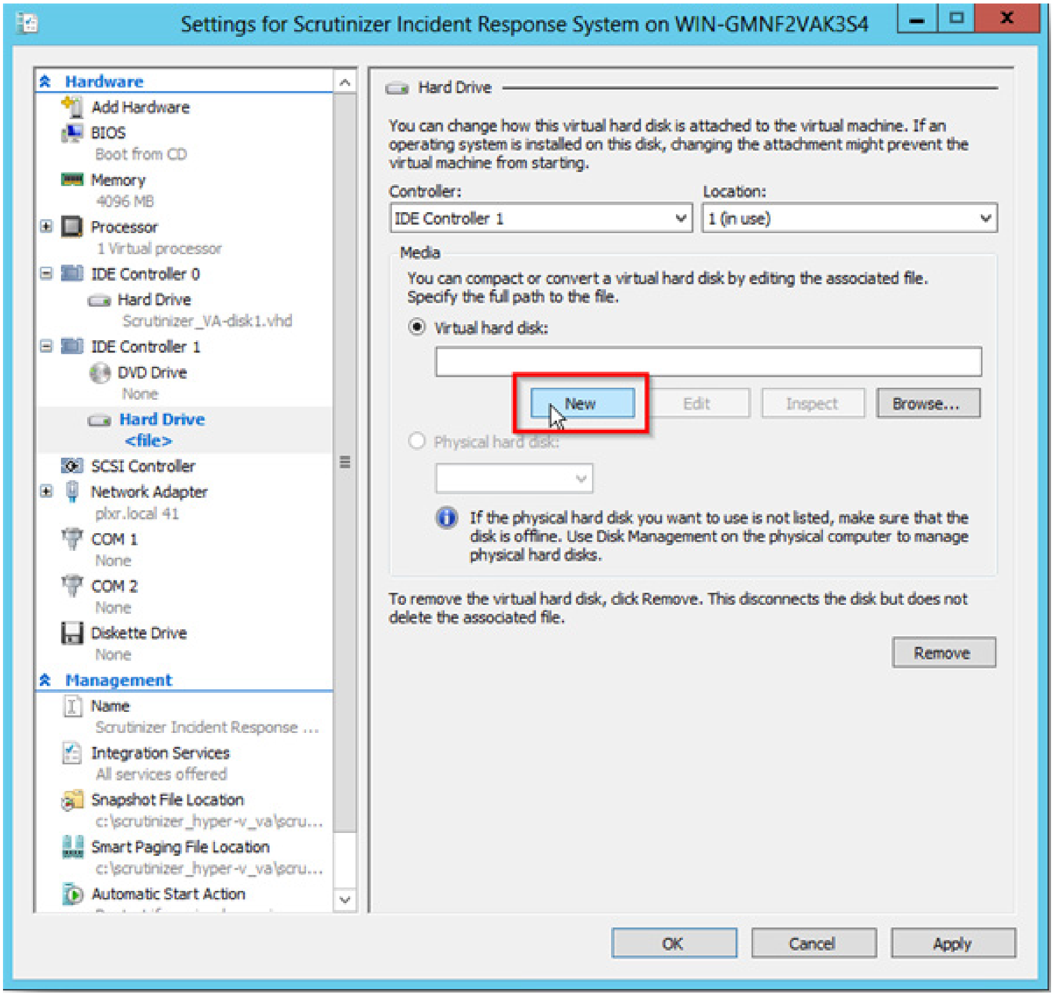

Settings for Scrutinizer Incident Response System on WIN-GMNF2VAK3S4

Hardware
- Add Hardware
- BIOS
  Boot from CD
- Memory
  4096 MB
- Processor
  1 Virtual processor
- IDE Controller 0
  - Hard Drive
    Scrutinizer_VA-disk1.vhd
- IDE Controller 1
  - DVD Drive
    None
  - **Hard Drive**
    <file>
- SCSI Controller
- Network Adapter
  plxr.local 41
- COM 1
  None
- COM 2
  None
- Diskette Drive
  None

Management
- Name
  Scrutinizer Incident Response ...
- Integration Services
  All services offered
- Snapshot File Location
  c:\scrutinizer_hyper-v_va\scru...
- Smart Paging File Location
  c:\scrutinizer_hyper-v_va\scru...
- Automatic Start Action

Hard Drive

You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.

Controller: IDE Controller 1

Location: 1 (in use)

Media

You can compact or convert a virtual hard disk by editing the associated file. Specify the full path to the file.

● Virtual hard disk:

New | Edit | Inspect | Browse...

○ Physical hard disk:

If the physical hard disk you want to use is not listed, make sure that the disk is offline. Use Disk Management on the physical computer to manage physical hard disks.

To remove the virtual hard disk, click Remove. This disconnects the disk but does not delete the associated file.

Remove

OK | Cancel | Apply

5. On the New Virtual Hard Disk Wizard, select “Next”.

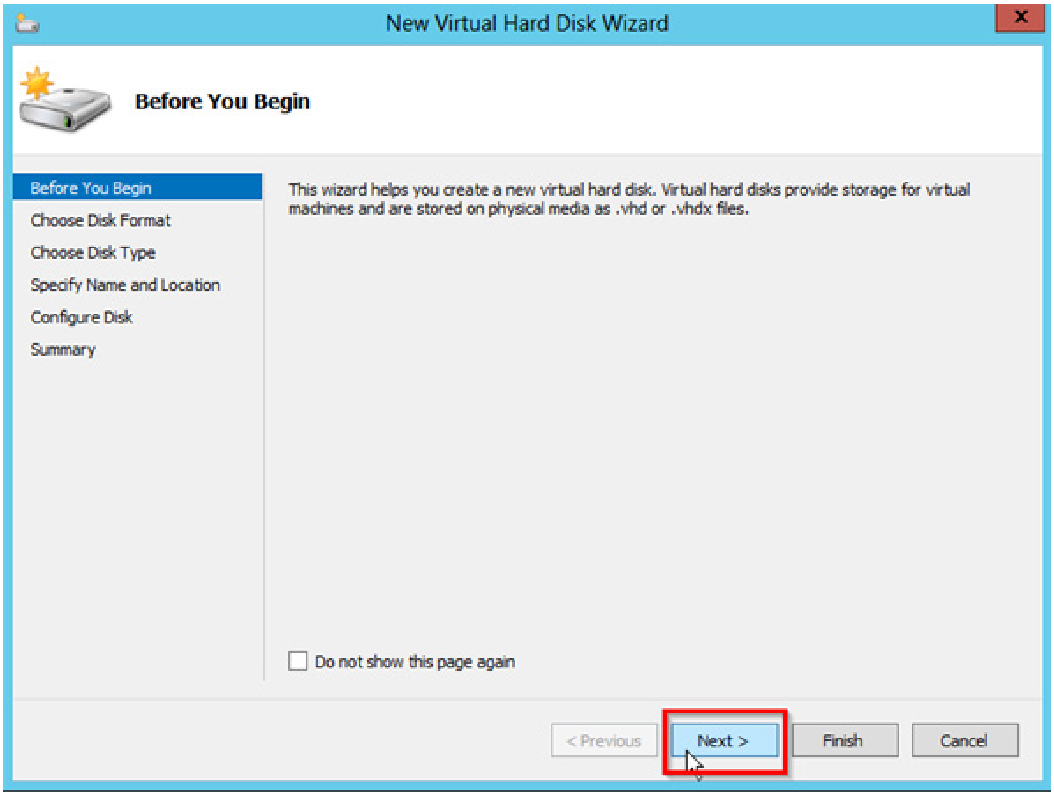

This wizard helps you create a new virtual hard disk. Virtual hard disks provide storage for virtual machines and are stored on physical media as .vhd or .vhdx files.

□ Do not show this page again

6. On the Choose Disk Format page, select VHDX. It's common for Plixer Scrutinizer VMs to expand past 2TB of disk space, so VHD is not recommended.

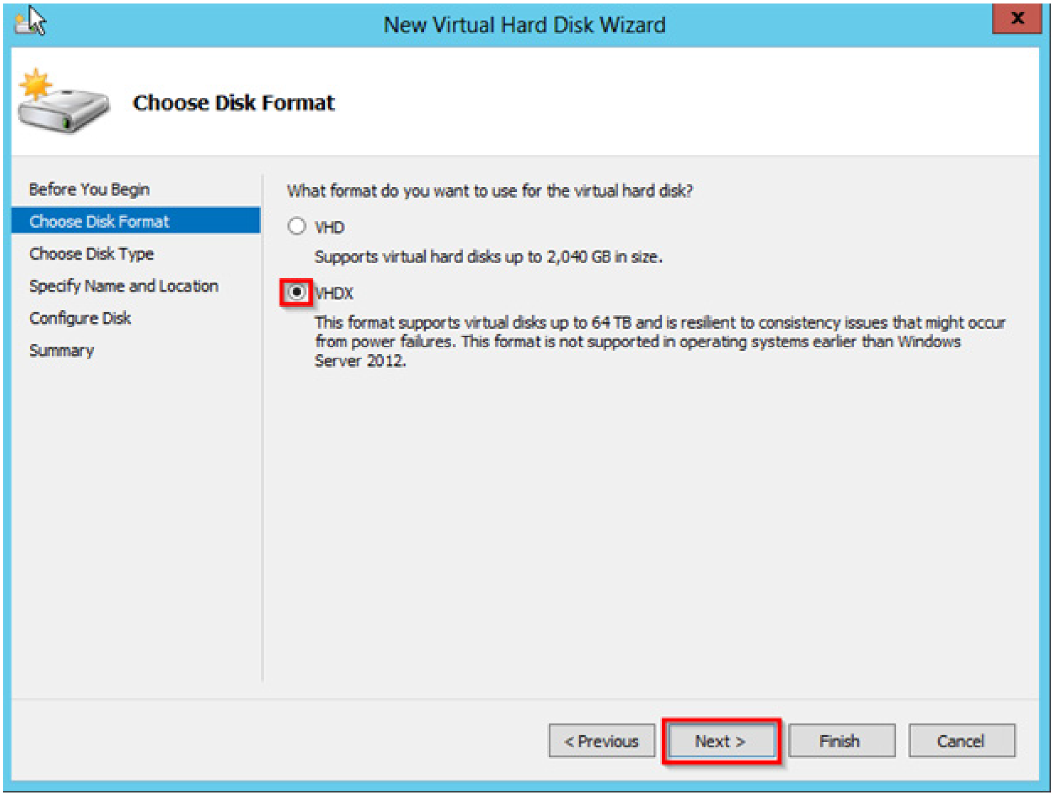

New Virtual Hard Disk Wizard

Choose Disk Format

What format do you want to use for the virtual hard disk?

○ VHD

Supports virtual hard disks up to 2,040 GB in size.

● VHDX

This format supports virtual disks up to 64 TB and is resilient to consistency issues that might occur from power failures. This format is not supported in operating systems earlier than Windows Server 2012.

7. On the Choose Disk Type page, select your preferred disk type and then press "Next".

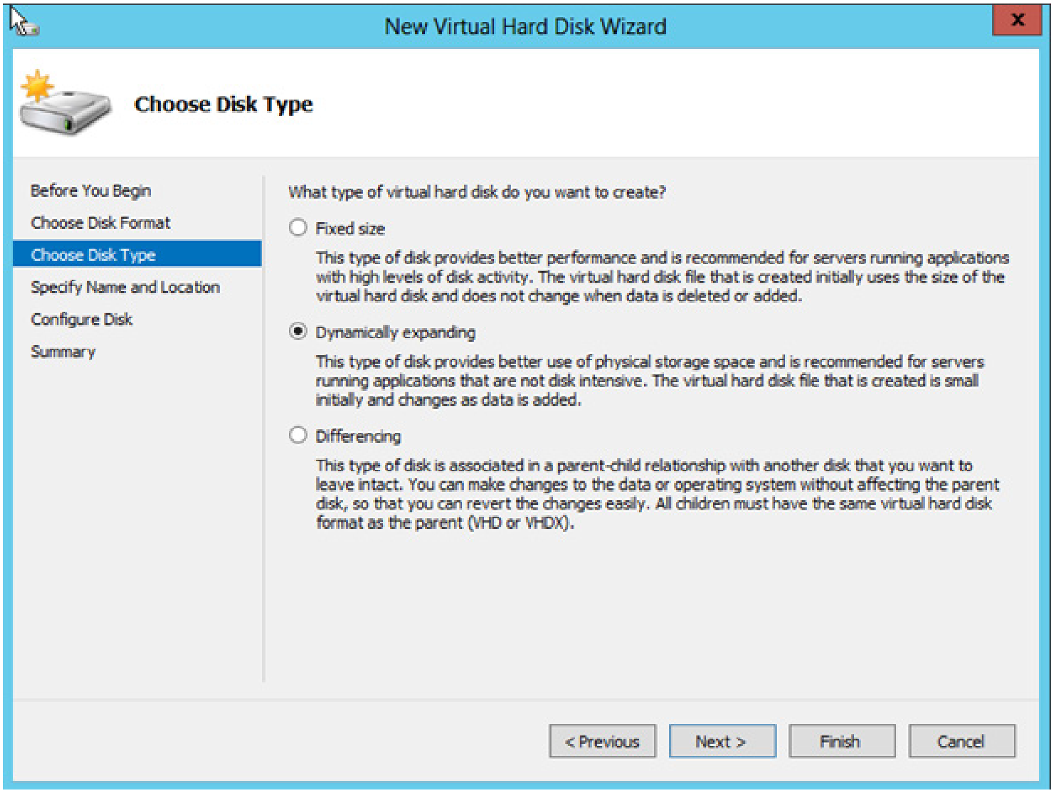

Choose Disk Type

What type of virtual hard disk do you want to create?

○ Fixed size

This type of disk provides better performance and is recommended for servers running applications with high levels of disk activity. The virtual hard disk file that is created initially uses the size of the virtual hard disk and does not change when data is deleted or added.

● Dynamically expanding

This type of disk provides better use of physical storage space and is recommended for servers running applications that are not disk intensive. The virtual hard disk file that is created is small initially and changes as data is added.

○ Differencing

This type of disk is associated in a parent-child relationship with another disk that you want to leave intact. You can make changes to the data or operating system without affecting the parent disk, so that you can revert the changes easily. All children must have the same virtual hard disk format as the parent (VHD or VHDX).

8. On the Specify Name and Location page, give your VHDX a name and then select the location for the virtual disk.

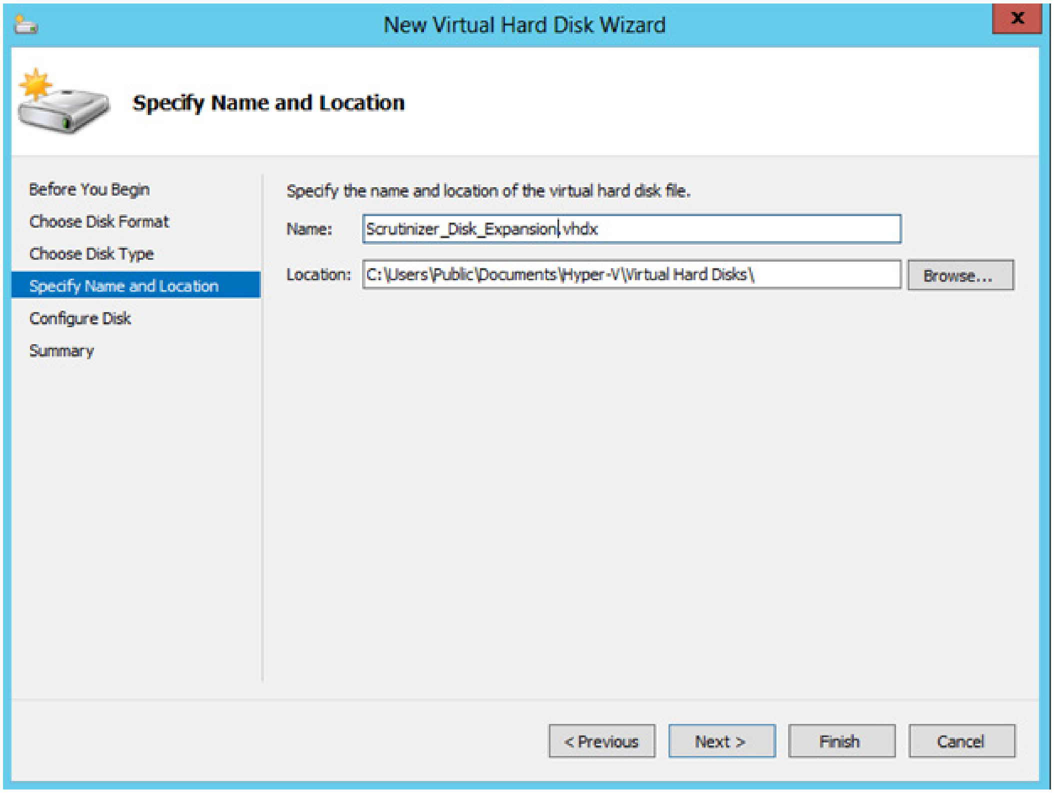

Specify the name and location of the virtual hard disk file.

Name: Scrutinizer_Disk_Expansion.vhdx

Location: C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks\

9. Set the size of the new virtual disk and then press "Next".

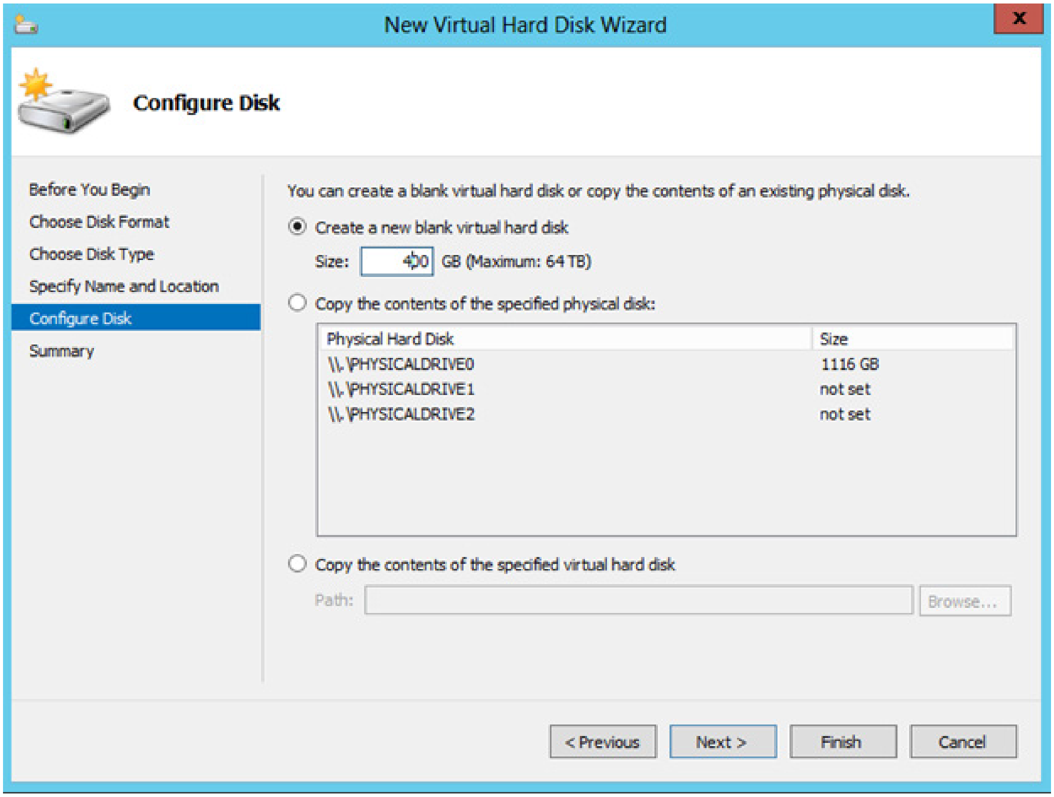

Configure Disk

You can create a blank virtual hard disk or copy the contents of an existing physical disk.

● Create a new blank virtual hard disk

Size: 400 GB (Maximum: 64 TB)

○ Copy the contents of the specified physical disk:

Physical Hard Disk	Size
\\.\PHYSICALDRIVE0	1116 GB
\\.\PHYSICALDRIVE1	not set
\\.\PHYSICALDRIVE2	not set

○ Copy the contents of the specified virtual hard disk

Path:

10. Review the new disk settings and then click “Finish”.

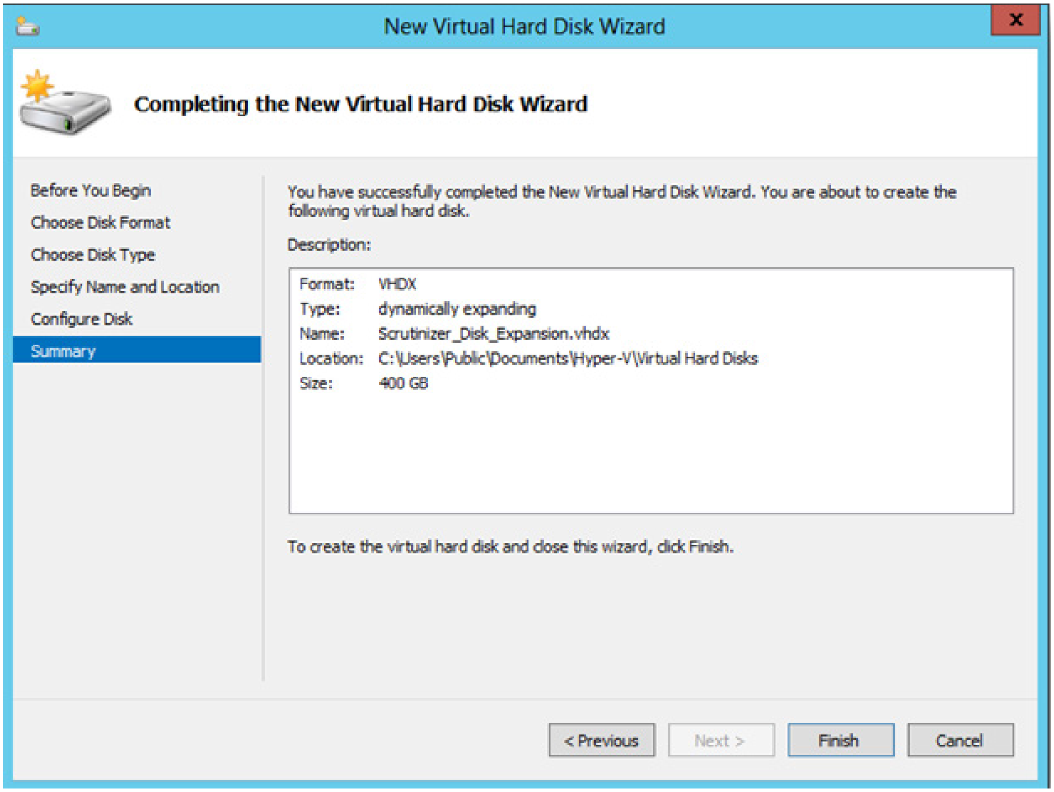

Completing the New Virtual Hard Disk Wizard

You have successfully completed the New Virtual Hard Disk Wizard. You are about to create the following virtual hard disk.

Description:

```
Format:   VHDX
Type:     dynamically expanding
Name:     Scrutinizer_Disk_Expansion.vhdx
Location: C:\Users\Public\Documents\Hyper-V\Virtual Hard Disks
Size:     400 GB
```

To create the virtual hard disk and close this wizard, click Finish.

11. Power on the Virtual Machine.
12. Follow from *step 6* onward under the “Expanding the database size for ESX” section of this manual.

1.1.5 Plixer Scrutinizer deployment on KVM

1. Create a directory for your install

   ```
   mkdir kvm/Scrut_VM_Guide/
   ```

2. Download the latest Plixer Scrutinizer Virtual Appliance to your KVM install

   Command line example:

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

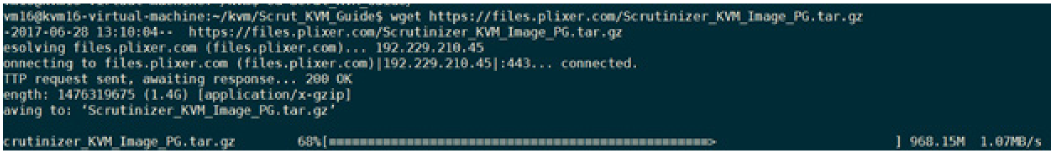

```
vm16@kvm16-virtual-machine:~/kvm/Scrut_KVM_Guide$ wget https://files.plixer.com/Scrutinizer_KVM_Image_PG.tar.gz
-2017-06-28 13:10:04--  https://files.plixer.com/Scrutinizer_KVM_Image_PG.tar.gz
esolving files.plixer.com (files.plixer.com)... 192.229.210.45
onnecting to files.plixer.com (files.plixer.com)|192.229.210.45|:443... connected.
TTP request sent, awaiting response... 200 OK
ength: 1476319675 (1.4G) [application/x-gzip]
aving to: 'Scrutinizer_KVM_Image_PG.tar.gz'

crutinizer_KVM_Image_PG.tar.gz        68%[==================================================>                    ] 968.15M  1.07MB/s
```

Note: Contact support for latest image if the URL above does not work.

3. Unzip the file on your KVM server to your new folder.

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

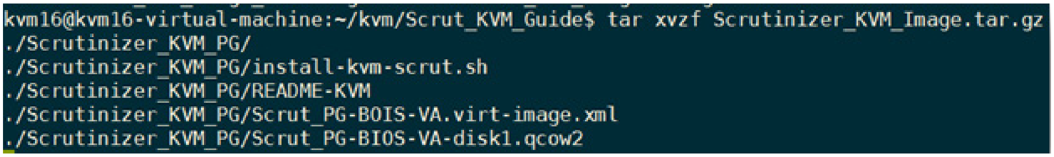

```
kvm16@kvm16-virtual-machine:~/kvm/Scrut_KVM_Guide$ tar xvzf Scrutinizer_KVM_Image.tar.gz
./Scrutinizer_KVM_PG/
./Scrutinizer_KVM_PG/install-kvm-scrut.sh
./Scrutinizer_KVM_PG/README-KVM
./Scrutinizer_KVM_PG/Scrut_PG-BOIS-VA.virt-image.xml
./Scrutinizer_KVM_PG/Scrut_PG-BIOS-VA-disk1.qcow2
```

4. Run your script to install Plixer Scrutinizer

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

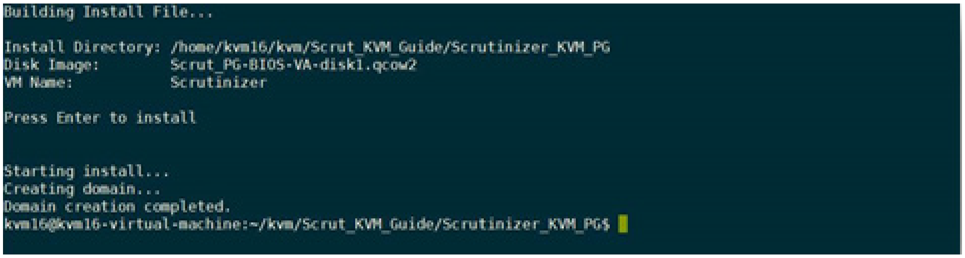

```
Building Install File...

Install Directory: /home/kvm16/kvm/Scrut_KVM_Guide/Scrutinizer_KVM_PG
Disk Image:        Scrut_PG-BIOS-VA-disk1.qcow2
VM Name:           Scrutinizer

Press Enter to install


Starting install...
Creating domain...
Domain creation completed.
kvm16@kvm16-virtual-machine:~/kvm/Scrut_KVM_Guide/Scrutinizer_KVM_PG$
```

At this point, you should see that your machine has been created from the image we deployed:

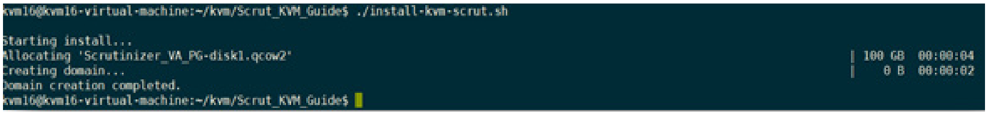

```
kvm16@kvm16-virtual-machine:~/kvm/Scrut_KVM_Guide$ ./install-kvm-scrut.sh

Starting install...
Allocating 'Scrutinizer_VA_PG-disk1.qcow2'                                                       | 100 GB  00:00:04
Creating domain...                                                                               |    0 B  00:00:02
Domain creation completed.
kvm16@kvm16-virtual-machine:~/kvm/Scrut_KVM_Guide$
```

5. Lastly, we just need to log in to the machine now that it is deployed. Run this command to get to the console:

   ```
   virsh console Scrutinizer
   ```

You will be prompted to log in; the default credentials are plixer/scrutinizer. The machine will reboot and you will be asked to log in again. This time you will be presented with a shell script asking for networking information. Follow the on-screen instructions and celebrate!

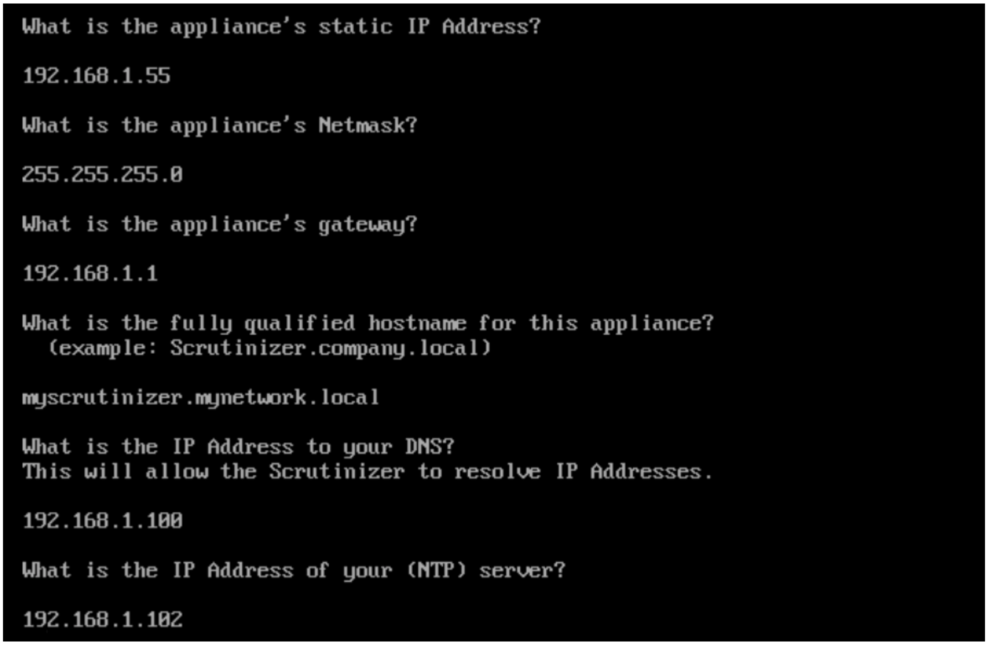

```
What is the appliance's static IP Address?

192.168.1.55

What is the appliance's Netmask?

255.255.255.0

What is the appliance's gateway?

192.168.1.1

What is the fully qualified hostname for this appliance?
  (example: Scrutinizer.company.local)

myscrutinizer.mynetwork.local

What is the IP Address to your DNS?
This will allow the Scrutinizer to resolve IP Addresses.

192.168.1.100

What is the IP Address of your (NTP) server?

192.168.1.102
```

1.1.6 Optimizing Plixer Scrutinizer datastores

Due to the nature of NetFlow, large deployments require a very high volume of disk I/O. For the best performance, the Plixer Scrutinizer Virtual Appliance should be deployed on a dedicated 15,000RPM RAID 10 datastore, with the amount of disk space that is required to meet your history setting requirements; 1.8 TB of disk space in RAID 10 is the recommended datastore deployment size.

If Plixer Scrutinizer is deployed on shared drives, such as a storage area network (SAN) or network-attached storage (NAS), then collection rates cannot be guaranteed as the collection rates will directly depend on what other applications are also using the same disk I/O.

In high flow volume environments, if you cannot get dedicated datastores, it's recommended to use a Plixer Scrutinizer Hardware Appliance for the dedicated resources and higher collection rates.

1.1.7 FAQ

Q: I got an UNEXPECTED INCONSISTENCY error when trying to power on the Plixer Scrutinizer Virtual Appliance. What do I do now? A: This error indicates that the clock on the ESX server is not set correctly and is in the past. As a result, the disk checks fail which does not allow the virtual machine to start. To resolve this, set your ESX host to sync with an NTP server and then redeploy the Plixer Scrutinizer OVF.

Q: How do I stop/start the services? A: Run the following commands (stop|start means type one OR the other): | service plixer_flow_collector stop|start | service plixer_syslogd stop|start | service httpd stop|start | service plixer_db stop|start

Q: I have a German 'QWERTZ' keyboard layout, how come I keep getting password failures when logging into the appliance for the first time? A: On the German 'QWERTZ' keyboard layout, the Z and Y keys are switched. You'll need to login with the password 'scrutiniyer'.

1.2 AMI deployment guide

1.2.1 What you need to know about Plixer Scrutinizer AMI

The latest Plixer Scrutinizer AMI can be obtained from Plixer or your local reseller. Please contact support if you do not already have the Plixer Scrutinizer AMI. You will need to know your AWS Region and your AWS account ID so the AMI link can be shared with you.

- Plixer Scrutinizer will deploy without a license key. A license key will be required for Scrutinizer operation. A key can be requested from Plixer upon the first login to the WebUI.
- Contact Plixer technical support to discuss the recommended instance type for production environments. Choose **c4.2xlarge** if the expected flow rate is under 10,000 flows per second.
- Decide on the VPC and security group rules that fit the needs of your organization. You will need to specify these in the deployment process.
- Deploy AMIs with two NICs and use the secondary NIC as the collection point. AWS does not let you release the primary private IP address of an instance unless you terminate the instance itself.
- The default size of the root partition size is 100GB. This should be left unchanged. If you need more disk space, you will need to add another disk. With the default of 100GB of disk space, you can store up to 1 month of NetFlow v5 data from 25 devices at 1,500 flows per second.
- Do not lose the SSH key that you will be asked to create in the deployment process. This key is the only way to access the server via SSH.

1.2.2 Pre-deployment checklist

Please provide a technical support engineer with the following information:

- Amazon account number
- The region you are planning to deploy an instance in
- Expected flow rate

1.2.3 Deploying AMI

1. Open the Amazon AWS console. From the console navigation bar, select the region that you instructed Plixer's Support to place the AMI in. Navigate to the EC2 Dashboard. Click the **Launch Instance** button.
2. In the navigation pane, click **My AMIs** to display the list of AMIs available to you in the region. Make sure to check the **Shared with me** checkbox. Select the Plixer Scrutinizer AMI from the list.
3. Choose **c4.2xlarge** for an Instance Type if the expected flow rate is under 10,000 flows per second. Contact Plixer technical support to review the recommended instance types for the production environments with higher flow rates.
4. Navigate to the **Configure Instance Details** section. Set the **Shutdown Behavior** to **Stop** and enable **Termination Protection**.
5. Select the **Network** and **Subnet** you would like to assign to the instance from the drop-down menus. Then assign the IP addresses to the AMI.

Important: It is highly recommended to deploy AMIs with two NICs and use the secondary NIC as the collection point. AWS does not let you release the primary private IP address of an instance unless you terminate the instance itself.

6. In the **Add Storage** section, the **Root Volume** should not be increased. The partition */dev/xvda/* should be set to 100GB. Make sure the **Delete on Termination** box is checked.

Important: If you need more storage, attach additional disks **after** the instance is running. Please see the *Adding storage to AMI* section or contact Plixer support for assistance.

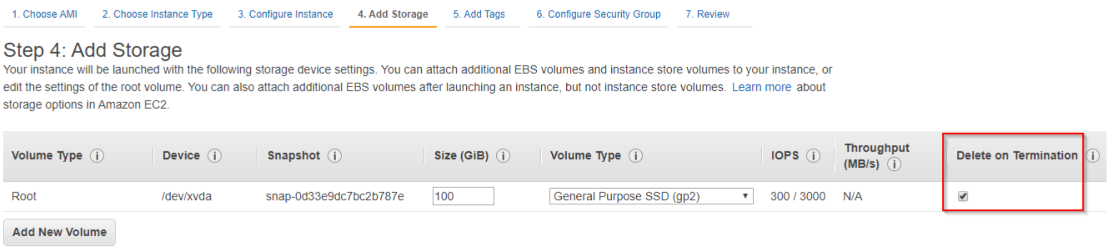

Step 4: Add Storage

Your instance will be launched with the following storage device settings. You can attach additional EBS volumes and instance store volumes to your instance, or edit the settings of the root volume. You can also attach additional EBS volumes after launching an instance, but not instance store volumes. Learn more about storage options in Amazon EC2.

Volume Type	Device	Snapshot	Size (GiB)	Volume Type	IOPS	Throughput (MB/s)	Delete on Termination
Root	/dev/xvda	snap-0d33e9dc7bc2b787e	100	General Purpose SSD (gp2)	300 / 3000	N/A	☑

7. [Optional] We recommend using **Tags** to categorize your AWS resources, for example, by purpose, owner, or environment. Each tag consists of a key and an optional value, both of which you define. Using a consistent set of tag keys makes it easier for you to manage your resources. You can search and filter the resources based on the tags you add.

8. Create a new **Security Group** or assign the instance to an existing group. A security group is a set of firewall rules that control the traffic for your instance. You can add rules to allow specific traffic to reach your instance. For example, if you want to set up a web server and allow internet traffic to reach your instance, add rules that allow unrestricted access to the HTTP and HTTPS ports.

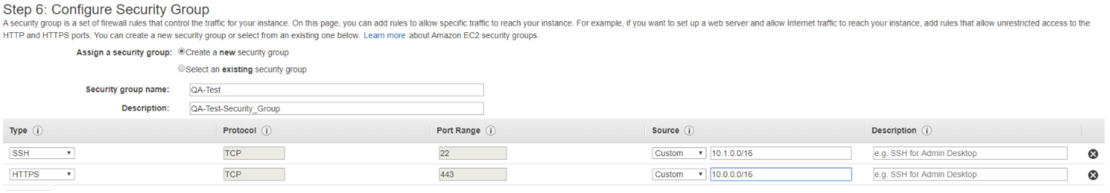

Step 6: Configure Security Group

A security group is a set of firewall rules that control the traffic for your instance. On this page, you can add rules to allow specific traffic to reach your instance. For example, if you want to set up a web server and allow Internet traffic to reach your instance, add rules that allow unrestricted access to the HTTP and HTTPS ports. You can create a new security group or select from an existing one below. Learn more about Amazon EC2 security groups.

Assign a security group: ● Create a **new** security group
○ Select an **existing** security group

Security group name: QA-Test

Description: QA-Test-Security_Group

Type	Protocol	Port Range	Source		Description
SSH	TCP	22	Custom	10.1.0.0/16	e.g. SSH for Admin Desktop
HTTPS	TCP	443	Custom	10.0.0.0/16	e.g. SSH for Admin Desktop

9. Navigate to the **Review Launch** section. Verify your instance launch details. You can go back to edit changes for each section. Click **Launch** to assign a key pair to your instance and complete the launch process.

10. Once it is launched, you can SSH into the server as the *ec2-user*. Then you can run the *sudo su -* command to become the root user.

```
ssh -i /PATH/TO/KEY/KEY.pem ec2-user@SCRUTINIZERIPADDRESS
```

1.2.4 Accessing Plixer Scrutinizer user interface

To access the application via a web browser, use its primary private or public IP address. After accepting the user agreement, log in as an *admin* user. The default password is the AMI instance ID.

For example, an instance with the instance id of i-07d159db2cd771515 can be accessed by using the following URL: https://172.18.1.11 with the username *admin* and password *i-07d159db2cd771515*:

Important: The default user interface login credentials are **user:** admin and **password:** the instance ID of your AMI instance.

1.2.5 Changing passwords

After deploying your new Plixer Scrutinizer instance, we recommend changing the user interface login credentials to meet your security requirements.

1. Navigate to the **Admin** tab within Plixer Scrutinizer's web console.
2. Click on the *admin* user.
3. Enter your desired password and click the **Save** button.

In some environments, additional security measures are placed onto passwords for all applications. If your environment requires a more complex password for Plixer Scrutinizer accounts, Plixer has a password complexity setting you can enable for Plixer Scrutinizer.

Navigate to **Admin > Settings > System Preferences** and check **Enforce Password Complexity**.

When checked, creating new users' passwords will require:

- length of at least 8 characters;
- one capital letter;
- one number;
- one special character.

1.2.6 Adding storage to AMI

1. Get the correct availability zone for the instance you will be adding the disk space to, from the **Availability Zone** column on the Instance page.

2. Navigate to the **Volumes** page and click on the **Create Volume** button in the top left. On the **Create Volume** page, create the new volume with the desired size and in the correct Availability Zone.

Important: ALWAYS choose GP2 General Purpose SSD.

3. Right-click on the new volume and select **Attach Volume**. Start entering the name of the instance and select it from the drop-down menu. Once you have the instance selected, you will need to change the name of the volume to "dev/xvdb". Click **Attach** when finished.

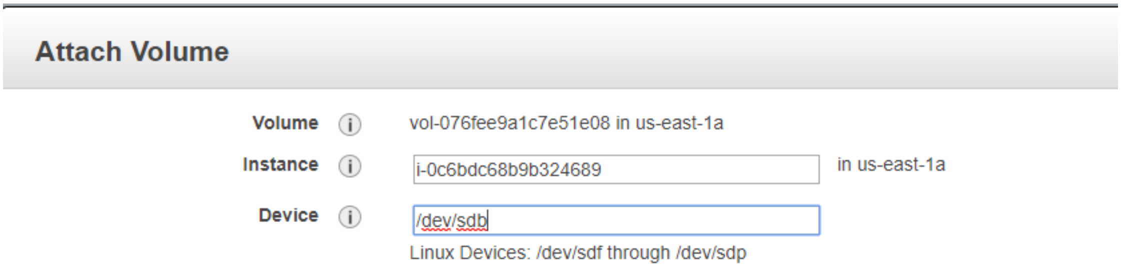

Attach Volume

Volume: vol-076fee9a1c7e51e08 in us-east-1a

Instance: i-0c6bdc68b9b324689 in us-east-1a

Device: /dev/sdb

Linux Devices: /dev/sdf through /dev/sdp

4. To configure the OS, SSH into the instance. Make sure the new volume was attached. Run the *sudo fdisk -l* command to find the new volume.

5. Launch the scrut_util interactive prompt and run the *set partitions /dev/xvdb* command to add the new disk.

```
[ec2-user\@supportami ~]$ sudo su
[root\@supportami ec2-user]# cd /home/plixer/scrutinizer/bin/
[root\@supportami bin]# ./scrut_util


SCRUTINIZER> set partitions /dev/xvdb
```

6. Run the 'show diskspace' command to check that the space was added and */var/db* is now its own mount point.

```
SCRUTINIZER> show diskspace
```

Hint: If you are adding more than one new drive to a single AWS instance, you will need to run the *set partitions* command for each drive.

1.2.7 How to add resources to a Plixer Scrutinizer AMI

As needs change, an instance can become overutilized. If this is the case, the size can be modified. For example, if an instance named t2.micro is resource-strapped, it can be changed to an m3.medium instance. When an instance is resized, an instance type that is compatible with the configuration of the instance must be selected. If the instance type is not compatible with the instance configuration, the application must be migrated to a new instance.

1. SSH into the instance.

2. Stop all the services using the interactive Plixer Scrutinizer utility:

```
[ec2-user\@supportami ~]$ sudo su
[root\@supportami ec2-user]# cd /home/plixer/scrutinizer/bin/
[root\@supportami bin]# ./scrut_util


SCRUTINIZER> services all stop
```

3. Power off the operating system:

```
[root\@supportami bin]# shutdown -h now
```

4. Open the Amazon EC2 console. In the navigation pane, choose **Instances**, then select the instance.

[EC2-Classic] If the instance has an associated Elastic IP address, write down the Elastic IP address and the instance ID shown in the details pane.

5. Choose **Actions> Instance State> Stop**

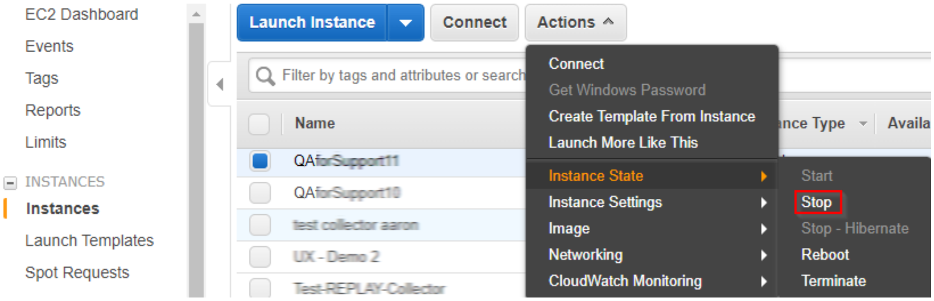

In the confirmation dialog box, choose **Yes, Stop**. It can take a few minutes for the instance to stop.

[EC2-Classic] When the instance state becomes stopped, the Elastic IP, Public DNS (IPv4), Private DNS, and Private IPs fields in the details pane are blank to indicate that the old values are no longer associated with the instance.

6. With the instance still selected, choose **Actions> Instance Settings**, then click **Change Instance Type.** Note that this action is disabled if the instance state is not stopped.

7. From **Instance Type**, select the instance type desired. If the desired instance type does not appear in the list, then it is not compatible with the configuration of the instance (for example, because of virtualization type).

8. To restart the stopped instance, select the instance, choose **Actions> Instance State**, then **Start**.

In the confirmation dialog box, choose **Yes, Start**. It can take a few minutes for the instance to enter the running state.

[EC2-Classic] When the instance state is running, the Public DNS (IPv4), Private DNS, and Private IPs fields in the details pane contain the new values that we assigned to the instance. If an instance had an associated Elastic IP address, it must be reassociated as follows:

In the navigation pane, choose Elastic IPs. Select the Elastic IP address that was written down in the previous step before the instance was stopped. Choose **Actions> Associate address**. From **Instance,** select the instance ID that was written down before the instance was stopped, then click on the **Associate** button.

9. SSH into the instance and use the interactive utility to retune:

```
SCRUTINIZER> set tuning
```

CHAPTER 2

System administration

2.1 Backups

This section explains the local and remote backup methods for Scrutinizer Postgres Databases.

Important: It is critical that you understand the fundamentals of Postgres databases. If you have any questions or concerns, please contact support for assistance.

Please select a backup method from the methods below:

2.1.1 Local Backup

Local backup is created on the same server that is hosting the Scrutinizer database.

1. Add the following line to the /var/db/big/pgsql/data/pg_hba.conf file to allow local replication connections:

```
# Allow replication connections from localhost, by a user with the
# replication privilege.
local   replication     postgres                                peer
```

2. Run the command below to reload the database configuration:

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

3. Use the pg_basebackup utility to back up the PostgreSQL database:

```
sudo su postgres -c "cd ~; pg_basebackup -D backup -Ft -z -P"
```

Hint: When taking a database backup, “waiting for checkpoint” is to be expected. However, you can force a checkpoint by running the following command:

```
psql -c "CHECKPOINT"
```

4. Your backup should consist of three tarballs in the /var/lib/pgsql/backup directory. All three are critical in restoration.

```
ls -l /var/lib/pgsql/backup
-rw-r--r-- 1 postgres postgres     2025 Sep 24 18:20 16385.tar.gz
-rw-r--r-- 1 postgres postgres 40170653 Sep 24 18:20 base.tar.gz
-rw------- 1 postgres postgres  1453827 Sep 24 18:20 pg_wal.tar.gz
```

Important: It is recommended to transfer backup files to an FTP server so they will not be lost in the event of a failure.

2.1.2 Remote Backup

This method assumes a second “destination” host for the backup that has at least enough disk space to contain the backup and that the same version of Scrutinizer and Postgres have already been installed.

1. On the source host, run this command to create a ROLE for remote replication:

```
psql -c  "CREATE ROLE remote_rep WITH REPLICATION LOGIN ENCRYPTED PASSWORD
↪'think_about_it'"
```

2. To allow remote connections for the ROLE on the source host, add this line to the pg_hba.conf. Replace the $DESTINATION_IP with the IP of the destination server and reload the configuration.

```
sudo tee -a /var/db/big/pgsql/data/pg_hba.conf << EOF
host replication remote_rep $DESTINATION_IP/32 md5
EOF
psql -c "SELECT pg_reload_conf()"
```

3. On the destination host, populate the $SOURCE_IP with the source server. Use the pg_basebackup utility to back up the PostgreSQL database. You will be prompted for the password used in Step 1.

```
sudo su postgres -c "cd ~; pg_basebackup -h $SOURCE_IP -U remote_rep -D␣
↪backup -Ft -z -P"
```

4. Your backup should consist of three tarballs in the /var/lib/pgsql/backup directory. All three are critical in restoration.

```
sudo ls -l /var/lib/pgsql/backup
-rw-r--r--. 1 postgres postgres        501 Sep 24 13:23 16385.tar.gz
-rw-r--r--. 1 postgres postgres 4146859412 Sep 24 13:32 base.tar.gz
-rw-------. 1 postgres postgres   12316287 Sep 24 13:32 pg_wal.tar.gz
```

2.1.3 Restoring from a Backup

The instructions below assume a fresh Scrutinizer instance was deployed on the same Scrutinizer and postgres version as the backup was taken on. Upload the backup files to the /var/lib/pgsql/backup/ directory.

Important: You MUST be restoring to the same version of PostgreSQL that was used during the backup.

1. Stop all Scrutinizer services.

```
sudo service scrutinizer stop
```

2. Delete the current database files.

```
sudo rm -Rf /var/db/big/pgsql/data/* /var/db/fast/pgsql_tmp/*
```

3. Restore the base.tar.gz backup files.

```
sudo mv /var/lib/pgsql/backup/base.tar.gz /var/db/big/pgsql/data/
sudo tar -zxvf /var/db/big/pgsql/data/base.tar.gz -C /var/db/big/pgsql/data
```

4. Restore The WAL.

```
sudo mv /var/lib/pgsql/backup/pg_wal.tar.gz /var/db/big/pgsql/data/pg_wal/
sudo tar -zxvf /var/db/big/pgsql/data/pg_wal/pg_wal.tar.gz -C /var/db/big/
↪pgsql/data/pg_wal
```

5. Restore The 'fast' TABLESPACE. Keep in mind that the number for the fast TABLESPACE will vary from system to system. In my example, it's 16385.

```
sudo mv /var/lib/pgsql/backup/16385.tar.gz /var/db/fast/pgsql_tmp/
sudo tar -zxvf /var/db/fast/pgsql_tmp/16385.tar.gz -C /var/db/fast/pgsql_
↪tmp
```

6. Start The PostgreSQL service. If it doesn't start, please contact Plixer support before continuing.

```
sudo service plixer_db start
```

7. It's common for database passwords to be different across a backup and restore. Run the following commands and reset the database passwords to ensure they are in sync:

```
scrut_util
set password scrutdb
services all restart
```

8. Register the collector with the new IP address. By moving the database, it is possible for the IP to change on the destination host. We need to ensure the collector service is aware of the new IP.

```
scrut_util
set selfregister reset
```

9. Check SSL Configuration. If you are not using SSL for the web interface on the source or destination servers, you can skip this step.

```
scrut_util
set ssl off
set ssl on
```

10. Apply a new license key. When changing hosts, Scrutinizer needs a new license key. In the web interface of the restoration host, navigate to Admin -> Settings -> Licensing, copy the Machine ID and contact Plixer support for a new license. Once you receive the new license key, in the web interface of the restoration host, navigate to Admin -> Settings -> Licensing and apply the new license.

11. Before deleting the backup, it is recommended to verify the restored system is collecting new flow data and that you can report on the restored historical data.If you changed the IP, you may need to point your exporters to the IP address of the new collector.

12. Once you're confident the database has been restored, remove the backup tarballs from the server.

```
sudo rm /var/db/fast/pgsql_tmp/16385.tar.gz var/db/big/pgsql/data/pg_wal/
↪pg_wal.tar.gz /var/db/big/pgsql/data/base.tar.gz
```

2.2 Changelog

For more details on the new features below, reference the Plixer website and Scrutinizer documentation.

KEY: ACTION: (Ticket Number) description

Ex. ADDED: (1640) Thresholds based on outbound traffic

Change Log History

Version 19.1.0 - May 2021

Scrutinizer

ADDED: (187): Scrutinizer services not required to run as root
ADDED: (261): Client - Server reports
ADDED: (516): Encrypt stored keys
ADDED: (733): Copy to clipboard button to api json tab
ADDED: (786): Option to toggle Show System Policies
ADDED: (883): Expanded and reworked Host Index and H2H Search
ADDED: (898): Target / Violator views and filtering in Alarm Monitor
ADDED: (948): Show Host Names and Show Acknowledged Events for Alarms
ADDED: (1971): Include collector IP address in all vitals reports for grouping and filtering
ADDED: (2053): Refactor Alarms backend for better performance
ADDED: (2060): Flexible notification policies based on event criteria
ADDED: (2111): Autoreplicate support for multiple replicators (encrypt multiple passwords)
ADDED: (2231): Ability to set Alarm policies to inactive or store
ADDED: (2361): root login disabled on new deployments
ADDED: (2374): Cisco SDWan (Viptela) integration updated to support version 20

FIXED: Addressed various security issues
FIXED: (313): Mapping: add checks and errors for duplicate map connections
FIXED: (724): Sorting by bytes does not account for units in Entity Views
FIXED: (793): New UI reports do not display Host Names
FIXED: (805): PDF Export of Summary Reports Top N and Overview failure
FIXED: (893): Classic View option from user menu doesn't work
FIXED: (939): Fix scrolling issues for Exporter Details list in Report Settings
FIXED: (1586): Alarms takes too long to load and acknowledge
FIXED: (1798): Reverse DNS exclusions for alarms
FIXED: (1970): Reparser crash when Linux ARP cache filled
FIXED: (1977): Adding a notification profile to a saved report threshold doesn't work
FIXED: (2030): Child Groups not enforced for FA exclusion
FIXED: (2090): Vitals process crashing with extremely high MFSNs in flow streams
FIXED: (2214): Custom URL Dashboard Gadgets not working
FIXED: (2217): Valid licenses with Expired PNI/PSI eval's prevent the upgrade from running

FIXED: (2235): Stream bloat on heavily loaded systems could cause disk space problems
FIXED: (2250): Running out of file descriptors on heavily loaded systems
FIXED: (2273): Invalid certificates in distributed upgrades
FIXED: (2279): TopN views are not always populated
FIXED: (2300): LDAP login takes too long with a very large list of security groups
FIXED: (2307): P2P Alarm report link not working
FIXED: (2336): Improve handling of truncated sFlow sampled headers
FIXED: (2346): Flow collection doesn't resume at the end of a network outage
FIXED: (2358): Set webui_timeout not working
FIXED: (2379): Scheduled report tasks called wrong binary name after upgrade
FIXED: (2382): IP exclusion only checking source IP for RST/ACK and Host Reputation
FIXED: (2393): Fix incorrect or missing sFlow interface numbers for instances above 63
FIXED: (2401): AES key not syncing on upgrade affecting SNMP, AWS, and other credentials needed on a collector
FIXED: (2414): License Exceeded alarm detail shows no data in Alarm Monitor
FIXED: (2457): Addressed CVE-2021-28993

Version 19.0.2 - January 2021

Scrutinizer

FIXED: (2075) Disabling User Does Not Invalidate Session
FIXED: (2076) Input validation needed in some forms
FIXED: (2080) Session cookie value stored in local storage
FIXED: (2118) Postgres log noise from unnecessary scheduled analytics command
FIXED: (2198) Distributed upgrade issue coming from 19.0.0
FIXED: (2202) pg_cron memory leak
FIXED: (2205) Fresh v19.0.1 OVA does not use the 19.0.1 repository

Version 19.0.1 - December 2020

Scrutinizer

ADDED: (12) DDOS: Support IPv6

ADDED: (377) Add AWS Role Based Authentication for use in AWS

ADDED: (940) Allow AWS flowlog polling at 1m frequency

ADDED: (1235) Enforce password policy on password change and restrict from using last four values

ADDED: (1459) Summary Reports added to new UI

ADDED: (1539) Add "scrut_util –show datasize" to enumerate DB schemas and their disk usage.

ADDED: (1633) Define Allegro IEs

ADDED: (1890) Support for new format of VPC flow logs

ADDED: (1891) Provide descriptions for AWS entity IDs

ADDED: (1899) Add Velocloud 4.0 IEs (tcpRttMs and tcpRetransmits)

ADDED: (1992) Document new AWS integration requirements

FIXED: (54) Mapping: Show Utilization only works for percent

FIXED: (304) Not excluding protocols by default

FIXED: (696) Secondary reporters show incorrect clock drift

FIXED: (739) Apache HTTP Server 2.4.0 - 2.4.39 Remote Open Redirect Vulnerability in mod_rewrite

FIXED: (765) Cannot Filter on S3 Bucket Element aws_account_id in a designed report

FIXED: (1065) Internal Server Error when emailing PDF report name includes /

FIXED: (1316) Unable to Exclude IP address from DDoS algorithim

FIXED: (1480) Collector log error sflow buffer overrun at ./protocol/sflow/buffer.hpp line 146

FIXED: (1482) VPC Flow Logs should be cleaned up more aggressively

FIXED: (1579) The plixer.idp.login_url field appears to be vestigial

FIXED: (1592) Other Options > GeoIP links not working

FIXED: (1660) Login banners are not working

FIXED: (1728) Interface names with special characters cause errors when triggering thresholds

FIXED: (1734) Alarm when disabling algorithms or ML stream

FIXED: (1743) Group Labels retain original input on Maps Dashboard Widget

FIXED: (1744) Host2host and host index lookups to work in distributed setup

FIXED: (1796) pgbouncer wont start after yum update

FIXED: (1797) Some reports were unable to display in percent interface view

FIXED: (1812) Reparser freezes on error during minutely exporter status updates

FIXED: (1813) No drillp-down into Connection on Maps

FIXED: (1817) Reparser memory leak in sFlow parser

FIXED: (1840) Devices blue after upgrade to version 19

FIXED: (1842) ServiceNow Integration doesn't work when server response is too large

FIXED: (1879) Reporting: No Data for Timeframe automatically sends to start report wizard

FIXED: (1911) Sliding windows falling behind after upgrade to v19

FIXED: (1912) Fix rollup issue for droppedPacketDeltaCount<unsigned64>

FIXED: (1917) Closing the report modal doesn't keep the report open

FIXED: (1918) Entity Views: sorting by bytes does not account for units

FIXED: (1920) Using LDAP user is authenticated but never added to a group when group list was too long

FIXED: (1930) Unable to disable unlicensed FA features

FIXED: (1941) Unrecognized key type: AWSLogs/xxxxxxxxxxx/ inc/lib/Plixer/Scrutinizer/awss3.pm line 547

FIXED: (1942) Awss3.pm:373 – get_flowlogs() encountered an error while processing s3_connection_list: Invalid data Invalid data(unknown) for aws_account_id

FIXED: (1945) get_flowlogs() encountered an error while processing s3_connection_list: Invalid data (-) @ 1084 for transform

FIXED: (1946) Alarm Report data interval default empty for large time frame events

FIXED: (1969) NetFlow v5 sampling crashes postgres

FIXED: (1981) Too many open files

FIXED: (1984) multicast send failure 22: Invalid argument

FIXED: (1988) CEF notifications missing 'Device Version'

FIXED: (1994) Set 'ssl_prefer_server_ciphers' by default

FIXED: (2002) Missing sflow records after an upgrade

FIXED: (2021) Report values as rates in tables are incorrect after drilling in on a graph

FIXED: (2029) Distributed: AWS S3 secret failing when assigned to remote collector

FIXED: (2068) The application is running a vulnerable version of Apache

FIXED: (2069) The application is running a vulnerable version of Perl

FIXED: (2070) XSS Vulnerability in old UI mechanism to create groups

FIXED: (2072) Local file inclusion

FIXED: (2111) Autoreplicate support for multiple replicators (encrypt multiple passwords)

FIXED: (2071) Formula injection vulnerability in the ability to create third-party CrossCheck methods

Scrutinizer UI

ADDED: (652) Entities: Hosts: Anomaly Chart

ADDED: (692) Summary Reports: Filtering

FIXED: (657) Report filter descriptions don't always fill in
FIXED: (685) Dashboards not deleted
FIXED: (688) Drilling into Policy from Collection loses consistency vs Monitor View
FIXED: (693) Apache httpd: CWE-345: Insufficient verification of data authenticity
FIXED: (744) Reporting: Summary reports not stretching on page
FIXED: (765) Stop 'topping' the graphs

Machine Learning Engine

ADDED: (338) Add ML Engine metrics to Vitals reports
ADDED: (419) Support high availability
ADDED: (446) Support Zerologon detection
ADDED: (447) Support SIGRed detection

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.

ADDED: (9) New workflow-based user interface
ADDED: (12) DDOS: Support IPv6
ADDED: (370) Address data encryption in Scrutinizer
ADDED: (371) Initial Collections implementation
ADDED: (476) magicbus_fdw: Avro serialization

ADDED: (481) Advanced threat intelligence feeds

ADDED: (717) SNMP Enterprise MIB support for Viptela

ADDED: (727) Support for new VeloCloud information elements

ADDED: (740) Use tenant_id for db ROLE

ADDED: (780) Require a license key for free mode

ADDED: (781) Support for content updates

ADDED: (782) Streaming support for customer data lakes

ADDED: (783) Host to host flow connection search

ADDED: (784) Plixer Replicator integration

ADDED: (874) Update the Silverpeak IPFIX information elements

ADDED: (903) Advanced security algorithms

ADDED: (1006) STIXV1 IP watchlist import

ADDED: (1007) STIXV2 IP watchlist import

ADDED: (1008) TAXII 2 feed support for IP indicators

ADDED: (1142) Domain reputation checking

ADDED: (1144) JA3 fingerprinting support

ADDED: (1152) Machine learning for security-specific events

ADDED: (1153) Machine learning for network-specific events

ADDED: (1215) New licensed features

ADDED: (1256) ML forecasting in Scrutinizer

ADDED: (1258) ServiceNow integration

ADDED: (1411) CEF notification action

FIXED: (541) Failed "system updates" report "no updates available"

FIXED: (614) scrut_util.exe –collect asa_acl gives error Use of uninitialized value $debug in concatenation

FIXED: (636) Saved Reports Folder changes are not audited

FIXED: (749) Insecure Direct Object Reference

FIXED: (767) Vitalser Memory Leak

FIXED: (820) Define missing Cisco IEs (unknown_9_20000)

FIXED: (865) Define the unknown_elements for Viptela IPFIX exports

FIXED: (1196) scrut_util –collect db_size is timing out

Version 18.20 - April 2020

ADDED: (496) Optimized sFlow collection
ADDED: (2073) New VeloCloud information elements
ADDED: (2154) Security updates
ADDED: (2164) SNMP Enterprise MIB support for Viptela
ADDED: (2165) Updated Silverpeak IPFIX information elements
ADDED: (2176) CentOS 7 : kernel update
ADDED: (2177) PostgreSQL security release 10.12
ADDED: (2190) Change default eval key to 14 days

FIXED: (2156) sFlow traffic discrepancies
FIXED: (2167) Saved report dashboard gadgets always display in totals
FIXED: (2179) Reporting issues when 0 byte flows are excluded
FIXED: (2196) Fixed issue with totals when both ingress and egress flows are exported

---

Version 18.18 - December 2019

ADDED: (1939) New VeloCloud reports
ADDED: (2036) Set admin password to instance_id for AMIs
ADDED: (2039) Add SSO authentication method to the manual
ADDED: (2051) Many updates, improvements, and clarifications in documentationi
ADDED: (2124) New Viptela reports
ADDED: (2133) Option template based descriptions for VeloCloud LinkUUID

FIXED: (421) Create scheduled reports was also requiring admin tab permission
FIXED: (1441) Auto refreshing pages would prevent session timeout
FIXED: (1405) Resolve timeout for FA reverse DNS exlusions wasn't using setting from admin tab
FIXED: (1536) We now exclude 0 byte flows biFlow records for reporting and FA
FIXED: (1756) Protocol exclusions were not audited

FIXED: (1816) 255 character limitation for 'Security Groups Allowed' when configuring LDAP integration

FIXED: (1936) Improved column naming in some VeloCloud reports

FIXED: (1985) Resolve a harmless UDP receive buffer error

FIXED: (1992) Viptela reports would sometimes not show all vEdge hosts

FIXED: (2030) Session timeout based on backend activity, not frontend activity

FIXED: (2040) PDF report displays no data when data is present

FIXED: (2041) Expand Disk scrut_util commands now support NVME drives

FIXED: (2106) If an IdP certificate is not provided, SAMLRequests should be unsigned

FIXED: (2107) SSO - Submitting metadata XML via the admin view form incorrectly parses out tags

FIXED: (2041) Fixed memory leak in vitalser

Version 18.16 - September 2019

ADDED: (16) Viptela SD-WAN reports

ADDED: (270) Permission configuration on a role basis

ADDED: (378) Changed AWS Flow Log collection to use S3 buckets and added support for multiple regions and customer IDs

ADDED: (550) VeloCloud SD-WAN reports

ADDED: (569) Service Now Notification support

ADDED: (826) Appliance self migration from CentOS 6 to CentOS 7

ADDED: (891) Ability to Add/remove/update Defined Applications via the API

ADDED: (897) Single-Sign-On support through SAML 2.0

ADDED: (937) Alarm when authentication tokens will expire in 30 days or have expired

ADDED: (992) Deleting an exporter doesn't block collection

ADDED: (1099) Removed device specific status notifications

ADDED: (1171) Audit logs can now be expired after a configurable duration

ADDED: (1205) FDW option to Database migrator for faster PostgreSQL migrations

ADDED: (1254) Flow inactivity alarms are now checked across a distributed cluster and are per exporter rather than per interface

ADDED: (1425) Support for Fortinet application names

ADDED: (1735) Support Nokia (formerly 'Alcatel-Lucent') IPFIX

ADDED: (1832) Support for Gigamon Application Intelligence

FIXED: (185) Schedule emails will now use the theme from Admin > Settings > System Preferences
FIXED: (308) The ability to use an auth token with any URL
FIXED: (636) UTF8 issue with Japanese characters in email alert notifications
FIXED: (700) 'Truncate map labels' was grabbing an extra character sometimes
FIXED: (753) Addressed an issue with flow class sequence numbers with distributed upgrades
FIXED: (841) Removed admin restriction on running group level reports
FIXED: (846) Clarify several log error messages, and reduce their volume
FIXED: (900) Some Scrutinizer custom gadgets break the ability to add any gadget for all users
FIXED: (1066) AMI: set partitions doesn't remount pg_stat_tmp as a RAM drive
FIXED: (1079) Issue where deleted exporters may not be cleared out of LED stats table
FIXED: (1082) Issue where system updates could revert a setting causing "Panic: Can't find temp dir" errors and the interface failing to load
FIXED: (1085) Higher default timeouts for collect asa_acl task
FIXED: (1117) Issue with special characters in PRTG integration
FIXED: (1120) Warnings when an exporter sends the same multiplier data two different ways as long as what it sends is consistent
FIXED: (1132) UNION SELECT errors in migrator
FIXED: (1140) Autofilling IP on host search from report tables
FIXED: (1142) Scheduled reports last sent time used incorrect
FIXED: (1145) SQL GROUP BY ERROR in the collector log
FIXED: (1158) Issue with Auto SNMP Update not disabling all SNMP calls
FIXED: (1209) PostgreSQL logs using too much disk space
FIXED: (1229) Special characters in notification profile breaks threshold's 'save & edit policy' option
FIXED: (1231) Added stray columnar file check and alarm policy
FIXED: (1239) Monitor association of /var/db/fast and RAM spools
FIXED: (1249) Issue with running yum update on AWS EC2 instances
FIXED: (1272) Issue with load time of Admin > Host names view
FIXED: (1297) Defined application changes now realized on distributed collectors w/o a collector restarts
FIXED: (1314) Issue with alarm details and FQDN data for clusters using DB encryption
FIXED: (1322) DB disk usage stats did not always expire on distributed installs
FIXED: (1385) Collect support files includes the PostgreSQL log
FIXED: (1392) Allow snmpSystem details longer than 255 characters
FIXED: (1422) Errors from set tuning when two changes require a collector restart
FIXED: (1431) Getting Internal Server Error (500) when trying to access Maps > CrossCheck and Service Level Reports

FIXED: (1440) Some administrative changes for authentication did not generate audit events

FIXED: (1447) Addressed issue with ASA ACL collection when the reporter can not communicate with all firewalls

FIXED: (1458)* Issue with LDAP/TACACS usernames being case sensitive

FIXED: (1489) LDAP authentication was not failing over to try other servers

FIXED: (1506) Backup method documentation on docs.plixer.com

FIXED: (1527) Advanced TCP flag filters using strings would generate log noise

FIXED: (1536) Improved performance of Persistent Flow Risk algorithm

FIXED: (1542) Developer tasks_view hours filter causes Internal Server Error (500)

FIXED: (1544) Dashboards with multiple saved report gadgets cause oops errors

FIXED: (1553) Reporting across migrated data and new data doesn't use the migrated totals tables

FIXED: (1556) Migrated totals tables have the wrong scrut_templateid

FIXED: (1588) Peak values being less then the total values in the volume -> traffic volume reports

FIXED: (1599) Some English values in foreign language themes were out of date

FIXED: (1632) New reparser performance

FIXED: (1663) Migration from 16.3 mysql to 18.14 removed dashboard gadget permissions

FIXED: (1668) LDAP group checking was using sAMAccountName instead of the value specified in the configuration page

FIXED: (1691) Map object icons change colors based on polling availability

FIXED: (1731) The default group was not being set correctly for new users

FIXED: (1733) Payload size preventing CSV rendering of reports

FIXED: (1789) Saved reports belonging to users that no longer exist would not show up in report folders

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

ADDED: (873) Now including cstore table conversion script in utils

ADDED: (951) Improved default work_mem settings

FIXED: (640) DB process needs priority over other processes when system runs out of memory

FIXED: (676) Acknowledging Multiple Pages of an Alarm, acknowledges all alarms

FIXED: (714) ‘unhandled multicast message’ in the collector log

FIXED: (778) Report Designer not saving added row

FIXED: (780) Drilling into Palo Alto User Report generates a blank pop up

FIXED: (784) Top Interfaces summarization timing out with high interface count

FIXED: (790) Issue when upgrading from version 16.7

FIXED: (793) Issue where exporters sending bad timestamps would freeze spool file processing

FIXED: (832) “Save password” error when navigating from group membership

FIXED: (849) Large number of DrDOS violations could crash process

FIXED: (850) Error when changing exporter status

FIXED: (851) Backup exporters count against licensing even if same IP is already active

FIXED: (872) Interface thresholds would only violate if there was both inbound and outbound traffic

FIXED: (894) IP group detection not working for v6 addresses

FIXED: (895) Cleanup logging for sFlow exports from Cumulus Router

FIXED: (896) Not all interface names are collected from FireSIGHT

FIXED: (903) Issue with business hours ending at midnight

FIXED: (904) First time LDAP authentication would fail if local authentication is disabled

FIXED: (956) Scheduled reports attaching wrong pdf to email

FIXED: (963) Drilling in on an interval from volume reports could display the wrong timeframe

FIXED: (971) A slow connection could impact API latency LED for other collectors

FIXED: (990) Issue with NTP daemon not starting automatically on some installs

FIXED: (1004) Updated DRDOS thresholds to be ratios instead of fixed packet counts

FIXED: (1009) TACACS authentication would work if disabled but configured

FIXED: (1019) Issue with the scale APM outbound jitter was displayed in

FIXED: (1063) Reparser could not connect to the DB with a space in the password

FIXED: (1130) One exporter not collecting when at maximum license count for exporters

---

Version 18.12.14 - January 2019

ADDED: (10) Realtime DDOS and DRDOS detection before data is written to disk

ADDED: (87) FQDN reports are back and better performing

ADDED: (105) Interface threshold checks are now done once a minute and check one minute of data
ADDED: (111) FireSIGHT integration includes username support
ADDED: (112) FireSIGHT integration includes interface names
ADDED: (274) Group reports now include members of child groups
ADDED: (299) "User Accounts" permission to allow restriction of Scrutinizer user account creation
ADDED: (447) Added option to disable CrossCheck threshold notifications

FIXED: (132) Faster report CSV generation
FIXED: (167) FireSIGHT integration detects connection loss and attempts to reconnect to FirePOWER
FIXED: (177) Top interfaces values were understated for sFlow exporters sending multiple totals flows per minute
FIXED: (263) PostgreSQL log rotation
FIXED: (267) Rate values for Trend reports are now based on graph interval
FIXED: (301) Link Back Host set to the wrong port on a deployed AMI
FIXED: (319) Installer no longer displays post install script errors
FIXED: (26415) Add Audit messages when connections to LDAP servers fail
FIXED: (26768) Fixed username filtering when name is based on IPv6 address
FIXED: (26874) Faster Defined Application tagging

Version 18.9 - September 2018

FIXED: (26874) Fixed issue with multiple defined applications on the same IP
FIXED: (26511) Improved contrast for some icons in dark themes
FIXED: (26536) System user was counting against licensing limits
FIXED: (26550) Fixed issue with top N gadgets and exporters only sending egress flows
FIXED: (26557) Fixed the Analytics Violation Overview link on the Alarms tab
FIXED: (26579) Fixed issue using Gmail to send emails
FIXED: (26587) Fixed issue with emailing table views
FIXED: (26600) Fixed issue with TopN subnets gadget and SAF aggregation
FIXED: (26602) Fixed issue with editing designed reports
FIXED: (26613) Backslash in LDAP passwords caused issue on upgrade
FIXED: (26619) Fixed issue with map labels in dashboards

FIXED: (26629) Multiple subnet filters issue in MySQL
FIXED: (26632) Fixed issue with threshold details not being cleared out when switching reports
FIXED: (26650) Fixed issue editing designed reports with some manufactured columns in them
FIXED: (26652) Fixed issue with interface permissions in mapping
FIXED: (26655) Fixed issue with row limiting in CSV files
FIXED: (26699) Fixed issue with flow vitals when packets contain multiple flow sets for the same template
FIXED: (26731) Reporting: Top 10 rows on any page are now color coded as the graph
FIXED: (26735) Postgres installs - improved reporting temp table performance

---

Version 18.7 - July 2018

ADDED: (23542) Added QRadar Integration
ADDED: (26194) Changed dashboard gadget behavior to improve usability and clearly display gadget titles
ADDED: (26310) Numerous improvements to the manual

FIXED: (24546) Flickering issue with report graphs when loading a report
FIXED: (25156) Formatting issues in Maps Tab alerts
FIXED: (25504) Double tooltip when mousing over report graph
FIXED: (26042) Audits from IPv6 hosts are now correctly received and recorded
FIXED: (26298) Issues with input parameters for the Users API
FIXED: (26317) Optimized rollups
FIXED: (26318) Decreased time necessary to run upgrades
FIXED: (26342) Links from alarms heatmap were not working
FIXED: (26345) Tuning would too aggressively set roller memory
FIXED: (26350) Addressed upgrade issue related to DB locking
FIXED: (26358) Improved dashboard gadget behavior based on customer feedback
FIXED: (26360) Reparser: Fix understatement of NetFlow v9 flow volume in vitals report
FIXED: (26370) AWS instances would not upgrade if on Postgres 9.5
FIXED: (26371) Maps couldn’t be saved in dashboard gadgets
FIXED: (26372) Could not generate PDFs of reports in Japanese

FIXED: (26373) Fixed issue with Japanese characters in emailed reports
FIXED: (26395) Other Options > Search link not working
FIXED: (26399) Peaks in totals tables were 5 minute byte counts rather than 1 minute byte counts
FIXED: (26406) Forensic filters were not forcing change to forensic data
FIXED: (26431) Fixed filtering on AS number under Admin > Definitions > Autonomous Systems
FIXED: (26451) Fixed issue with making dashboards visible to a user group

* This is the last supported release for the CentOS 6 and MariaDB platforms

---

Version 18.6 - June 2018

ADDED: (9911) Test button for LDAP/RADIUS/TACACS setup
ADDED: (15154) Ability to acknowledge alarms with any combination of filters
ADDED: (16826) scrut_util command to disable ping for devices that have not responded
ADDED: (17589) Manufactured columns can be included in the report designer
ADDED: (18291) Full back button support
ADDED: (19981) Automatically detect which SNMP credentials to use for exporters
ADDED: (20068) Ability to manage interface details via API
ADDED: (21522) Ability to filter on a port range
ADDED: (21744) All interface reports now account for metering on each interface in the report
ADDED: (21770) Host -> AS -> Host reports for additional BGP reporting
ADDED: (22220) Major release upgrade to PostgreSQL 9.6 and 10
ADDED: (22773) scrut_util command to enable/disable ipv6
ADDED: (23267) User can be locked out after n failed login attempts
ADDED: (23478) Full foreign datastore support in collection and rollups
ADDED: (23924) Ability to exclude domain names from flow analytics
ADDED: (24134) Ability to edit URLs for custom gadgets
ADDED: (24164) Milliseconds now included with formatted timestamps where applicable
ADDED: (24297) Columnar store support for AWS Scrutinizers
ADDED: (24452) Ability to customize the login page
ADDED: (24600) Improved support for configuration of multiple LDAP servers and domains
ADDED: (24661) Ability to grant dashboards to other users / groups

ADDED: (24781) Default PostgreSQL datastore is columnar. Better disk space utilization and IO performance.
ADDED: (24948) Performance improvements for flow class lookups
ADDED: (25077) Support IPv4-mapped IPv6 addresses in subnet and ipgroup filters (PostgreSQL)
ADDED: (25216) Report IP Group with protocol and defined applications
ADDED: (25289) Support for Flowmon probe elements
ADDED: (25396) DrDoS detection for memcached and CLDAP attacks
ADDED: (26187) Ability to schedule operating system updates

FIXED: (12972) Flow metrics vitals times now align with ingestion time
FIXED: (22530) Ungrouped now visible by non-admin users
FIXED: (22588) Tidy up loose ends when deleting exporters. Deleted exporters will stay deleted.
FIXED: (22654) Stop showing disabled exporters in the exporters LED
FIXED: (24107) Some timezones were duplicated in the selector
FIXED: (24115) Latency reports per exporter
FIXED: (24659) Addressed issue reporting on multiple interfaces with different metering configured
FIXED: (24703) Issue with generating PDF with device group filters
FIXED: (24790) Restrict PaloAlto username collection to only internal IPs
FIXED: (24875) Donut/Pie Graph not available in Top -> Interfaces report
FIXED: (24893) Map interface utilization arrows always pointed in the same direction
FIXED: (24899) 'cancel report' button truly cancels backend reporting requests.
FIXED: (24993) Device menu in Google maps
FIXED: (25027) Cleaned up log noise from Cisco ISE data collection
FIXED: (25111) Scheduled reports font issue on AWS
FIXED: (25317) Remove memcached external exposure CVE-2017-9951
FIXED: (25323) FlowPro APM jitter report
FIXED: (25399) Audit report times now display as clients timezone
FIXED: (25419) Addressed CVE-2014-8109
FIXED: (25660) Issue with Queue Drops >> Queue Drops By Hierarchy

---

Version 17.11 - November 2017

ADDED: (24685) Support for Oracle cloud

FIXED: (24500) Vitals errors when a user with a long UID is created

FIXED: (24560) Save button for filters would go away if field was selected, but not changed

FIXED: (24586) Localhost Unlicensed after upgrade to 17.10

FIXED: (24616) Collector appears down after Daylight Savings Time change

FIXED: (24647) Potential short gap in rollups after collector restart

---

Please reference our End of Life Policy for details regarding the end of life schedule. For more information on Scrutinizer, please reference the online documentation or visit our website.

2.3 Check vulnerabilities

It is important to keep the Scrutinizer software and Operating System up to date to patch any known vulnerabilities.

After applying all the system updates it is common for vulnerability scanners that only look at package version numbers to report that vulnerabilities still exist when they have already been patched. This is a result of backporting security patches.

Backporting

Backporting has a number of advantages for customers, but it can create confusion when it is not understood. Customers need to be aware that just looking at the version number of a package will not tell them if they are vulnerable or not. For example, stories in the press may include phrases such as "upgrade to Apache httpd 2.0.43 to fix the issue," which only takes into account the upstream version number. This can cause confusion as even after installing updated packages from a vendor, customers will not have the latest upstream version. They will instead have an older upstream version with backported patches applied.

Don't trust the version number

Some security scanning and auditing tools make decisions about vulnerabilities based solely on the version number of components they find. This results in false positives as the tools do not take into account backported security fixes.

Look for the CVE number

Backported security fixes can be manually verified by looking for the CVE number in the package changelog.

Here is an example:

# rpm –q –changelog httpd | grep CVE-2017-3169

- Resolves: #1463197 - CVE-2017-3169 httpd: mod_ssl NULL pointer dereference

To fix the false positives generated by security scanners that only look at version numbers, something called OVAL definitions (machine-readable versions of the advisories) are supplied for third-party vulnerability tools. These can be used to determine the status of vulnerabilities, even when security fixes have been backported. In doing this, we hope to remove some of the confusion surrounding backporting and make it easier for customers to always keep up to date with the latest security fixes.

OVAL definitions can be downloaded from oval.cisecurity.org/repository

2.4 Data aggregation

The system can save and roll up data in one of two ways:

- *Traditional mode* (creates forensic tables). This is the default mode for version 17.x and prior.
- *SAF Mode* (creates both summary tables and forensic tables). This is the default mode for version 18.x and higher. Summary tables are condensed flows that represent 100% of the traffic, but by summarizing the data as explained below, reporting is much faster.

Upgrades

When existing customers upgrade from version 17.x and prior to 18.x versions, the aggregation method will stay in Traditional (Forensic) Mode. New installs default to SAF mode. Contact technical support for instructions on how to change this to SAF mode. When running a report and switching between the two types of data, this is referred to as the “Data Mode”.

Traditional mode

Traditional mode creates forensic tables. In this mode, the collector saves 100% of all data in raw format to the 1 minute forensic tables for each router. Every hour it creates a new 1 minute interval table per router. Every 5 minutes, it creates *higher intervals* using the smaller intervals. This process is called “rollups”.

When the rollups occur for 5 min, 30 min, 2 hr, 12 hr, 1 day and 1 week, two tables are created:

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

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

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

The totals tables are not used when:

- Reporting across all interfaces of a device
- A report is run on multiple interfaces from different devices regardless of filters
- A non interface filter has been applied to the report (e.g. IP address)
- When a report uses more than one template
- Looking at 1 minute intervals in a report. One minute intervals contain 100% of all data exported for the template as no rollups have occurred. As a result, no totals tables are created for 1 minute intervals.

The totals tables are only being used when:

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

Note: Only the top 1000 (default) conversations are saved in the rollups by default. A conversation is defined as 1 or more flows in a time frame that match based on values in a tuple (source and destination IP address, ingress and egress interfaces, commonPort, etc.). If the collection server has the available disk space, try increasing the *Maximum Conversations* under **Admin Tab -> Settings -> Data History** to 10,000 and see if it improves the accuracy. Don't configure it right away for the maximum. Instead try carefully increasing the number of conversations saved over a few days. Some reports may render more slowly when the maximum conversations is increased, this is the result of the tables being larger.

SAF mode (Summary and Forensic)

SAF mode creates both summary and forensic tables, but only the summary tables are rolled up. summary tables contain a subset of forensic columns and are useful for fast and long term reporting. It can act as a search index before drilling into the forensic data.

In SAF mode, the collector saves 100% of all data in aggregated format to the 1 minute summary tables for each router. By default, with summary data, flows are aggregated based on a tuple that includes the common port. Fields that prevent aggreggation such as the source and destination transport ports are dropped. Kept fields include, but are not limited to:

- intervalTime
- commonport
- ingressInterface
- egressInterface
- sourceIpAddress
- destinationIpAddress
- octetDeltaCount
- octetDeltaCount_rev
- packetDeltaCount
- packetDeltaCount_rev
- flowDirection
- applicationId

- protocolIdentifier

The summary aggregation logic used to create the above can be modified. Contact technical support for details on how to modify these settings.

This dramatically reduces table size and maintains accuracy. Every hour the collector creates a new 1 minute interval summary table per router. Every 5 min, 30 min, 2 hr, 12 hr (no 1 day or 1 week) it creates higher intervals using the smaller intervals. This process is called "rollups".

When the rollups occur on summary data, two tables are created in exactly the same way as outlined for traditional (forensic) Mode. The 2nd table is the totals table.

In SAF mode:

- Nearly all reports leverage the summary tables.
- Only vendor specific reports and a few other reports that require elements such as source and destination ports go back to the Forensic data.
- Since forensic tables in SAF mode are not rolled up, totals tables are not created for forensic tables.

A word about sFlow

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

> **Warning:** If the sflow exporting device (e.g. switch) is exporting multiple templates for different flows, 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 found to be the case, use the filters to select a single template.

2.5 Support for distributed collectors

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

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

Network ports used

The following network ports are used in communications between servers in a distributed environment.

These ports are used for two-way communication between the collector(s) and the reporter(s):

- TCP port 22
- TCP port 80 (or 443)
- TCP port 6432 and 5432

The following port is used for communication from the collector(s) to the reporter(s):

- UDP port 514

2.6 Configuring Scrutinizer for dual/multi-homing

Scrutinizer is configured with Reverse-Path Filtering which is disabled by default. It allows the collector to receive flows from IP addresses that it does not know how to route to. This functionality exists in order to allow collectors to receive non-local traffic that may have been forwarded by a proxy or replication appliance and is intended ONLY to be used when Scrutinizer:

1. is in a secure environment
2. has a single interface

Plixer recommends that in multi-interface/multi-homed situations and where strict networking practices are required, best practices defined in RFC 3704 should be observed. This is required in order to ensure that spoofed/forged packets are not used to generate responses that are sent out another interface.

It is recommended the following steps be taken:

1. Enable Reverse Path Forwarding –

   > Edit the /etc/sysctl.conf and modify the line “net.ipv4.conf.default.rp_filter = 1” to “net.ipv4.conf.default.rp_filter = 0”

2. If the goal is to not restart networking after editing this file, simply run the command “sysctl net.ipv4.conf.default.rp_filter = 0” and this should turn it on (editing the file is still required).
3. Make sure the routing tables contain routes to all networks that will be receiving flows. Failure to properly configure routing will result in the inability to collect flows from non-local address spaces.

Virtual Routing and Forwarding (VRF) mode

In certain circumstances it might be best to isolate the routing tables from the management network. Common cases for this are either security requirements or the management network may overlap with IP addresses on the collection side interfaces.

In these cases, separate routing tables can be created in order to isolate management traffic to the management interface. Collection and polling traffic will then only impact their respective interfaces.

Configure separate routing tables

In this example, Scrutinizer has 2 interfaces: eth0 and eth1. Each will have its own routing table. One called “plixer”, another called “public”. Notice that 2 default gateways are configured. This would not be possible under a standard configuration with a single routing table. Each maintains separate isolated IP networks which are not routable to one another on the Scrutinizer box.

1. Add the 2 routing tables to the file named “/etc/iproute2/rt_tables”. The 2 lines will result in the file looking like this:

```
#
# reserved values
#
255 local
254 main
253 default
0 unspec
#
# local
#
#1 inr.ruhep
1 public
2 plixer
```

2. Create file /etc/sysconfig/network-scripts/route-eth0 containing the following:

   default via 172.16.2.20 table plixer

3. Create file /etc/sysconfig/network-scripts/route-eth1 containing the following:

   default via 10.1.1.251 table public

4. Interfaces do not need any special configuration. Here is an example based on the above configuration.

```
/etc/sysconfig/network-scripts/ifcfg-eth0
DEVICE="eth0"
BOOTPROTO="none"
HWADDR=""
NM_CONTROLLED="yes"
ONBOOT="yes"
BOOTPROTO="none"
PEERDNS=no
TYPE="Ethernet"
NETMASK=255.255.255.0
IPADDR=172.16.2.7


/etc/sysconfig/network-scripts/ifcfg-eth1
DEVICE="eth1"
BOOTPROTO="none"
HWADDR=""
NM_CONTROLLED="yes"
ONBOOT="yes"
```

(continues on next page)

(continued from previous page)

```
BOOTPROTO="none"
PEERDNS=no
TYPE="Ethernet"
NETMASK=255.255.0.0
IPADDR=10.1.4.190
GATEWAY=10.1.1.251
```

5. Restart the server in order to make sure networking is reset and operating properly.

6. Verify that iptables are configured to accept or deny the traffic desired on each interface.

2.7 Functional IDs

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

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

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

Interactive usage allows a user to inherit all privileges that were granted to that ID. **Non-interactive usage** is used internally by the system only and cannot be assigned to a user.

Privileged access level is defined as an ID with elevated entitlements such as a system administrator or super user. **Non-privileged** access allows just the access level required for the intended functions of the ID.

2.8 Getting started

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

How to enable NetFlow or sFlow on various hardware.

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

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

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

2.9 Language translations

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

1. **Admin Tab >Definitions >Language**
2. Select a language and make updates. Notice the pagination at the bottom.
3. Languages are saved as /home/plixer/scrutinizer/files/localize_LANGUAGENAMEHERE.xls

This feature can be used to make changes to nearly all text in the interface. For example, by modifying the following key types with custom text:

- 'loginCustomText'
- 'loginCustomTitle'

A message can be displayed to people logging in.

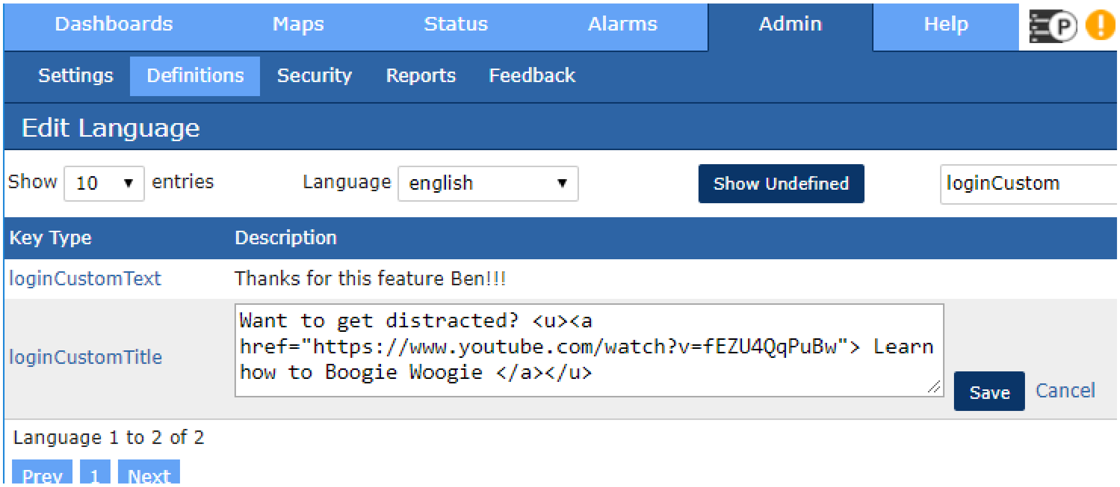

Edit Language

Show 10 entries Language english loginCustom

Key Type	Description
loginCustomText	Thanks for this feature Ben!!!
loginCustomTitle	Want to get distracted? <u><a href="https://www.youtube.com/watch?v=fEZU4QqPuBw"> Learn how to Boogie Woogie </a></u>

Language 1 to 2 of 2

2.10 Licensing

The licensing page is where you can enter the license key, view the current licensing status for local install, and get the Machine ID to provide to Plixer customer service to obtain a new license.

The licensing page contains two sections: sections.

Details

Product Type – This field indicates the licensing level
Status – Licensing status: valid, expired, due to expire
Days Left – Days remaining for valid license
Customer ID – Plixer generated customer identification
Machine ID – Unique server ID used by Plixer Customer Service to generate the license key
Server Count – Total number of servers that are licensed (includes both reporters and collectors)
Reporter Count – Number of reporters that are licensed (primary and secondary)
Exporter Count – Number of exporters that are licensed
Deployed Servers – Number of actual servers
Enabled Exporters – Number of actual exporters

Plixer Network Intelligence - This field indicates if the PNI license is applied
Plixer Security Intelligence - This field indicates if the PSI license is applied

License key

The text box in this section is where the license key received from Plixer customer service is entered.

Example License Key:

Nb4JuYhx5R1Uv60ipcnvuuEBsK2HBcZym3t3XwuVZKkFB08HfoVHQgBGDfOPmN+he
bMENz+1vUuz1+vpAKg3S9C5M1rLb52HYyadaVIYWJgY8mqwsEdMPk968Bs7qFPfta
QekdGn4IgCEDqnzWlUccQcTD54vnyqR9SSySH9zIxqUebDPxOUjcdt6eK9gWSHPm3

TVX2Nz1mPQagpRPXeYZFVLL593yIZwBtEmSBQDiGaAiC0MG3h4pINox0u9OWBcOdr
ZKQjTIUNRWF2NzMw/JBfAzFxdxWUGicfyaRjNsutNmj6RB9gupmLhPzJdchgHLZg7
+vj4PRSRM0n0g==

2.11 Replicator load balancing

The load balancing feature provides integration between a Scrutinizer distributed cluster and a Replicator appliance. Load balancing for multiple replicators is covered in the advanced configuration section.

2.11.1 Setting up load balancing

Load balancing requires you add credentials in the UI of Scrutinizer for your replicator, and running the autoreplicate binary from the command line of the reporter twice. The initial run will generate a config file, the second will start the processing of exporters.

Default Configuration

1. Fill out the replicator credentials to connect to, navigate Admin > Settings > Plixer Replicator
2. SSH as the plixer user into your Scrutinizer primary reporter.
3. Run the following command (Note this is NOT run inside interactive CLI):

```
scrut_util --autoreplicate
```

When you launch the command for the first time, it will create a configuration file called */home/plixer/scrutinizer/files/autoreplicate.conf*. The file will be pre-populated with the necessary entries as JSON.

4. Edit the */home/plixer/scrutinizer/files/autoreplicate.conf* file for your setup. More information about the configuration file is detailed in the Configuration section

5. Rerun the *scrut_util –autoreplicate* command. It will create a replicator policy for each collector and add exporters that are in the seed_profile to collector profiles up to collector_capacities. After you run the auto-replicate command, Scrutinizer will create an alarm that will contain the configuration details.

Note: If you would like Replicator to forward the traffic from all the exporters to Scrutinizer, add a policy for "0.0.0.0/0" to your seed profile.

If you are running in the default mode as scrut_util from the bin directory, the configuration file will be /home/plixer/scrutinizer/files/autoreplicate.conf

Advanced Configuration

Advanced options are available such as running multiple autoreplicate binaries and pooling.

With the help of Plixer Support, additional binaries can be created and a directory setup to allow for replicator pooling.

If the binary you are running is named 'replicator_util' or the directory path it is in contains 'replicator_pools' , the configuration file will be created / referenced in the same directory as the binary you are running.

Additional parameters added to the configuration file will take precendence over what was set in the UI

Those optional parameters are:

replicator_host
replicator_pass

replicator_seed_profile
replicator_receive_port
replicator_send_port

Passwords can be AES256 encrypted in the conf file or as plain text. To avoid using plain text password, generate an encrypted value use this command, then put the output in the configuration file as replicator_pass:

```
--autoreplicate --encrypt YOUR_PASSWORD
```

2.11.2 Editing the configuration file

Here are a few helpful tips for modifying the /home/plixer/scrutinizer/files/autoreplicate.conf file.

1. When editing the collector_capacities section, make sure you have an entry for each collector in the cluster. If you add a new collector, update the configuration accordingly. A policy is automatically created and managed for each collector in the Scrutinizer distributed cluster.

2. The exporters value is the maximum number of exporters that should be sending flows to a collector.

3. The flow_rate setting controls the maximum flows per second the collector should receive.

4. The seed_profile stands for the replicator profile that you add all exporters to on the replicator appliance. Any exporters added to a designated policy on a replicator will be sent to a collector in the distributed cluster.

Important: The collector_capacities setting limits the number of exporters. If you exceed the capacity, exporters will be removed.

2.11.3 FAQs

Q: If a new collector is added to the Scrutinizer cluster does Scrutinizer automatically start balancing exporters to it? A: No. Configuration file needs to be updated. You can manually add the collector to the configuration file and then subsequently run the scrut_util –autoreplicate (or wait for next scheduled* run, if configured.)

Q: How does auto-replicate handle a collector going offline? Can it detect a down collector and re-balance the cluster to send flows elsewhere?

A: No. It does not automatically detect down collectors.

Q: The collectors within my cluster are unbalanced. Why is auto-replicate not distributing the load evenly across all of the collectors? A: Every exporter sends flows at a different rate; that rate changes throughout the day. When exporters are added to a collector, the assumed rate of 200 flows/second, is used. As the real rate is observed, the exporter will only be moved when a collector is pushed outside of its defined limits.

Q: How does auto-replicate determine where new exporters will be assigned? A: Exporter placement is determined by a round-robin of all collectors defined in the autoreplicate.conf file, until a collector is over its limits.

Q: When a collector is over provisioned, how does Scrutinizer determine which collector the exporter will be moved to? A: Exporter placement is determined by a round-robin of all collectors defined in the autoreplicate.conf file, until a collector is over its limits.

Q: My auto-replicate collector threshold is set to 40,000 flows per second but I'm seeing times every day where the collector is getting more than 40,000 flows/s. Why are exporters not being moved to other collectors when the peak threshold is being crossed? A: Flowrate is based on a 24-hour average.

Q: How many Replicators does this functionality support? A: One replicator can be defined in the Scrutinizer UI, however unlimited Replicators are supported with manual configuration. See the Advanced Configuration section of the manual regarding autoreplicate.

Q: How many Replicator seed profiles does this functionality support? A: One per autoreplicate configuration. See Advanced Configuration section for details.

Q: How many Replicator unique listening ports does this functionality support? A: One per autoreplicate configuration. See Advanced Configuration section for details.

Q: The manual says to run the scrut_util -autoreplicate command. Why do I need to do this manually and do I always need to do this manually? A: Autoreplicate is not scheduled to run by default. Users have the option to cron this to run whenever they would like.

Q: How often does auto-replicate run? A: Autoreplicate is not scheduled to run by default. Users have the option to cron this to run whenever they would like.

Q: Can I statically assign some exporters to some collectors with auto-replicate? A: No. For static exporters create a new replicator profile, as you normally would, to send additional exporters to your collector.

Q: How do exporters that are sending directly to Scrutinizer and not going through a Replicator impact auto-balancing? A: They are referred to as 'Rogue Exporters' and will count against the collectors exporter count and flow limits.

Q: What happens if all of the collectors are exceeding their defined collections rates and new exporters are added to the replicator? Where do they go? A: Those exporters will not be placed anywhere.

Q: Where does Scrutinizer log the auto-replicate changes that are being made? A: Output can be found at /home/plixer/scrutnizer/files/logs/ inside a epoch-stamped file which will contain a final output file after the autoreplicate command completes.

Q: Does auto-replicate take into account MFSN's when balancing? A: No it does not.

Q: Does old profiles are automatically removed after the configuration changes? A: No it does not. The UI profile descriptions says not to edit by default, but it is safe to delete profiles no longer relevant to your config file.

2.12 Meta data collection

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

How often it is collected

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

What meta data is collected

- **Vendor:** identifies if the Scrutinizer installer is branded as Plixer or an OEM vendor.
- **smtpfromEmail:** The email address configured as the administrator in the server preferences.

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

How is the data transferred

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

What Plixer does with the data

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

How do I turn it off?

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

Uncheck the option **Share Health Statistics**, scroll to the bottom, and click **Save**.

2.13 Migration utility

The Plixer Scrutinizer database migration utility is designed to ease the process of moving a Scrutinizer install between hosts. As with any kind of critical database work, a complete back-up should be made before attempting any kind of migration.

Important: Please contact Plixer technical support for assistance with migrations.

2.14 Predicting disk needs

How much disk space do I need to keep n days of data?

The disk space required to store a day of data is a function of:

- how many exporters are sending data;
- what data templates each exporter is sending;
- the cardinality of the data in each exporter/template.

To help understand how much disk space is needed, Plixer Scrutinizer includes details abou disk space that is being used as well as predictions based on your current settings.These can be configured via the **Admin > Settings > Data History** page. At the top of this interface are the desired data retention settings. Below the settings is the method by which data is being aggregated:

- Every flow sent to Plixer Scrutinizer is stored in its original form in one-minute buckets (1m). That is the minute the flow was exported if the exporter's clock is within one minute of the Plixer Scrutinizer collector clock. If the clocks are off, it is the minute the flow was collected.
- 1m records are "rolled up" or aggregated into higher intervals to allow fast long term trending 1m -> 5m -> 30m -> 2h -> 12h. Rollups are limited to the top N conversations and ordered by bytes to determine what will be kept.
- **Traditional rollups:** Every element in the original flow template will be in the higher interval templates. This takes more disk space, and for some elements, the higher interval data has little value.

- **Summary and Forensic (SAF) rollups:** Any template with the required information elements will be aggregated into a new template definition containing only common elements (srcIP, dstIP, bytes, packets, etc.). This allows for all common reports to be run (for example, country, IP Group, and AS by IP are based on the src/dst IPs) while storing data more efficiently.

Next is a table displaying how much disk space is being used for each data interval. Finally, there is a table showing how much disk space would be needed based on the current settings and the previously collected data.

What if the disks can't support the settings?

Disks have a 10% available threshold. When that is passed, 1m and 5m historical tables will be trimmed until disk utilization falls back under that threshold. If a single exporter was coming in for years and then 50 more added, that single export history will be trimmed until all exporters have the same history.

Note: If all flows are between the same two IPs, the data can be stored much more efficiently than if each flow is a unique pair of IPs.

2.15 Interactive CLI

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

To launch the interactive utility, run:

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

This will open the Scrutinizer prompt:

```
SCRUTINIZER>
```

To close the interactive prompt, type 'exit':

```
SCRUTINIZER> exit
Exiting...
[plixer@Scrutinizer ~] #
```

2.15.1 Modes of operation

The scrut_util utility mode of operation.

1) Interactive:

SSH as the plixer user and type *scrut_util* to launch the interactive utility and enter the commands in the SCRUTINIZER> prompt.

2.15.2 Help function

To display the list of the available commands, run:

```
SCRUTINIZER> help
```

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

```
SCRUTINIZER> help show
```

For help with specific extended commands (for example, the "show groups" command) type:

```
SCRUTINIZER> help show groups
```

2.15.3 Commands

Following are the available top level commands:

- *aws*
- *check*
- *ciscoise*
- *clean*
- *collect*

- *convert*
- *counteract*
- *delete*
- *disable*
- *download*
- *enable*
- *endace*
- *expire*
- *export*
- *import*
- *moloch*
- *optimize*
- *remove*
- *repair*
- *rotate*
- *services*
- *set*
- *show*
- *snoop*
- *system*
- *unlock*
- *update*
- *upload*
- *version*

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

aws

Manages AWS flow logs integration with Plixer Scrutinizer.

Command	Description
awssync	Sync IDs and descriptions from AWS.

check

Runs a test or check against the command provided.

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

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

ciscoise

Manage CiscoISE Node Integration with Scrutinizer.

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

clean

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

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

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

collect

Manually collect data that is useful for Scrutinizer.

Command	Description
collect asa_acl	Manually collects ASA ACL information from Cisco ASA Devices. This task is scheduled and routinely executed as part of normal operations.
collect baseline	Manually collects baseline data and checks for alarms. This task is scheduled and routinely executed as part of normal operations.
collect dbsize	Collects database size information.
collect elk <elk_ip>	Manually collects data from Scrutinizer and sends it to the configured ELK server. Reference the Elasticsearch / Kibana (ELK) Integration guide for more detailed information on the ELK integration.
collect optionsummary	Manually process flow option data collected by Scrutinizer. This information is routinely processed automatically.
collect pcap <in_sec> [<host>]	Collects a packet capture on the interfaces of the Scrutinizer server. Requires a timeout (in seconds) and an optional host name in IP format to further filter the capture.
collect snmp	Manually collects SNMP data that is used during Scrutinizer's operations. This process is automatically scheduled by Scrutinizer to run regularly.
collect splunk <splunk_ip> <port>	Manually collect data from Scrutinizer and send it over to the configured Splunk server. Reference the Scrutinizer for Splunk Application integration guide for more information
collect supportfiles	Collects various log files and server configuration data used by Plixer support to troubleshoot server issues.
collect topology	Collects various types of data from devices and Scrutinizer to help Scrutinizer understand the topology layout of the network.
collect useridentity	Manually process user identity data collected by Scrutinizer. This information is routinely processed automatically.

convert

This operation converts all encrypted information stored in Scrutinizer to use AES 256 encryption.

Command	Description
converttoaes	Converts all encrypted information stored in Scrutinizer to use AES 256 encryption. Warning: The command will alter database tables in Scrutinizer. Please use with caution.

counteract

Third-party integration support for ForeScout CounterACT servers.

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

delete

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

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

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

disable

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

Command	Description
disable baseline <exporter_ip>	Disables all baselines for the specified <exporter_ip>. The historical data will not be deleted. However, it will expire based on Scrutinizer's historical data settings. Warning: This command will alter the behavior of Scrutinizer baseline functionality. Please use with caution.
disable elk http://<ip:port>	Disables ELK (Elasticsearch, Logstash, and Kibana) flows from Scrutinizer to the URL specified. Reference the Elasticsearch / Kibana (ELK) Integration guide for more detailed information on the ELK integration. Warning: This command will alter the behavior of Scrutinizer functionality. Please use with caution.
disable ipv6	Disables ipv6 in sysctl.conf for all interfaces. Warning: This command will alter the behavior of Scrutinizer functionality. Please use with caution.
disable splunk http://<ip:port>	Disables Splunk flows from Scrutinizer to the URL specified. Reference the Scrutinizer for Splunk Application integration guide for more information on the Scrutinizer for Splunk integration. Warning: This command will alter the behavior of Scrutinizer functionality. Please use with caution.
disable user <username>	Removes a login account with access to the interactive utility for Scrutinizer server. Warning: This command will alter the behavior of Scrutinizer functionality. Please use with caution.
disable unresponsive	Disables ping for exporters that have not responded. Warning: This command will alter the behavior of Scrutinizer functionality. Please use with caution.
disable hypervtools	Disables Hyper-V Integration Tools for a Virtual Appliance running on Hyper-V. Warning: This command will alter the behavior of Scrutinizer functionality. Please use with caution.
disable vmwaretools	Disables vmwaretools for a Virtual Appliance running on VMware. Warning: This command will alter the behavior of Scrutinizer functionality. Please use with caution.

download

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

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

enable

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

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

Command	Description
enable baseline <exporter_ip> default	Enables default or custom baselines (manual) based on elements from NetFlow and IPFIX templates. Baselining has several parameters available to customize the specific baseline data to collect with the 'manual' option.
enable baseline <exporter_ip> manual <pri_element[,sec_element]> <element> <AVG|COUNT|MIN|MAX|STD|SUM> <dailyhr|busday|sameday>	Enables default or custom baselines (manual) based on elements from NetFlow and IPFIX templates. Baselining has several parameters available to customize the specific baseline data to collect with the 'manual' option.
enable custom_algorithm <identifier> "<algoname>"	Reference the Flow Analytics Custom Algorithms section for specific information on how to configure custom algorithms and create alarm policies.
enable dbpool <pool_port>	Enables database connection pooling for Postgresql.
enable ipv6	Enables ipv6 in sysctl.conf for all interfaces.
enable perl_support	Installs additional perl packages to assist with with custom scripting.
enable splunk http://<ip:port> <syslog_port> [<indexer>]	Enables splunk flows from Scrutinizer to the URL specified. The Scrutinizer for Splunk App is required on the Splunk Server. Visit plixer.com for more information.
enable user <username> <security_level>	Creates a new login account with access to the interactive utility for Scrutinizer server maintanance. The <security_level> switch is for disabling commands that can alter Scrutinizer's functionality. The security levels are: 1 - Disable commands that can stop data collection. 2 - Disable the ability to remove integrations or stop data collection. 3 - These users can only collect information about Scrutinizer and the operating system.
enable hypervtools	Enables Hyper-V Integration Tools for a Virtual Appliance running on Hyper-V. Running enable hypervtools wil also upgrade an already existing install of the hyperv-daemons.
enable vmwaretools	Enables vmwaretools for a Virtual Appliance running on VMware. Running enable vmwaretools will also upgrade an already existing install of the vmware agent.

endace

Third-party integration support for Endace probes.

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

expire

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

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

Command	Description
expire alarms	Expires alarm history from the threatsoverview and fa_transports_violations tables as specified in the Data History Flow Historical 1 Min Avg preference.
expire bulletinboards	Purges alarm bulletin board events older then the number of days defined by Scrutinizer's history settings.
expire dnscache	Purges DNS cache older then the number of days defined by Scrutinizer's history settings.
expire history [trim]	Expires flow data as defined by Scrutinizer's history settings. If the optional 'trim' mode is passed, Scrutinizer will trim older data to make more space on the hard disk.
expire ifinfo	Purges old and outdated interface information.
expire inactiveflows	Expires interfaces from the interface view that have stopped sending flows. Entries are expired based on the number of hours specified in the Scrutinizer System Preferences. (Admin -> Settings -> System Preferences -> Inactive Expiration)
expire orphans	Purges alarm orphan events older then the number of days defined by Scrutinizer's history settings.
expire templates	Expires flow template meta data for templates that haven't been seen in 30 days.

export

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

Command	Description
export langtemplate <lang_name>	The <lang_name> parameter is required. If the language exists, then it will create a CSV file that shows the english and <lang_name> keys. If the language does not exist, a blank template will be created. The language file resides at /home/plixer/scrutinizer/files/pop_languages_<lang_name>_template.csv
export peaks_csv <file> <interval> <dir> <date_range> [<group_id>]	Exports a CSV file listing interfaces and peak values based on criteria specified. Valid options for are specified as raw minutes (1, 5, 30, 120, 720, 1440, 10080). Directory must exist as a subdirectory of Scrutinizer's home directory. If specifying /home/plixer/scrutinizer/temp, then use temp as the directory. The valid <ranges> are Last24hours, LastFifteenMinutes, LastFiveMinutes, LastFortyfiveMinutes, LastFullHour, LastHour, LastMonth, LastSevenDays, LastTenMinutes, LastThirtyDays, LastThirtyMinutes, LastThreeDays, LastTwentyMinutes, LastWeek, LastYear, ThisMonth, ThisWeek, ThisYear, Today, or Yesterday. <group_id> is optional. To see a list of group_ids use show groups.

import

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

Command	Description
import aclfile	Imports ACL information from a file. The file must reside at. /home/plixer/scrutinizer/files/acl_file.txt. The format is a direct output of SHOW ACCESS-LIST directly on the exporter.
import applications <path/file> [reset]	Import application rules from a CSV file. It is recommended to use this file and path for the applications import csv file. /home/plixer/scrutinizer/files/application_import.csv A reset option can be passed which will remove all application rules before the bulk import. Expected format is one named application and one application rule per line. Supported rule types are subnet, single IP, IP range, wildcard, port, and child rules. Child Applications must be declared before being used in a parent Application's rule set. Valid application rule syntax is: "subnet rule",10.0.0.0/8 "single ip rule",10.1.1.1 "range rule",10.0.0.1-10.0.0.42 "wildcard rule",10.0.0.1/0.255.255.0 "parent/child rule","my subnet" "ports and protocols",0-65535/256 Applications must have at least one port rule and one of the IP rule types defined above. Applications not defined this way will be imported, but may not be tagged properly in flow data. For example, the first application in this import file is valid while the second is not. The second application does not have at least one port rule: 'My first Application',10.0.0.0/8 'My first Application',0-65535/6 'My second Application',11.0.0.0/8 Up to 100,000 individual application rules are supported.
import asns <path/file> [<delimiter>]	Imports custom asn definitions from a csv file. The is a required field. The path should be specified from after the /home/plixer/ scrutinizer/ directory. The is an optional parameter and defaults to " " (i.e. space). The csv file name must be all lowercase and requires these elements, in this order: AS Number,AS Name,AS Description,IP Network(s) The fields are comma delimited, whereas the optional parameter applies specifically to the IP Network(s) element. A comma cannot be used for the IP Network(s) delimiter. Example File: 213,my_list,what a great autonomous system,10.0.0.0/8 192.168.0.0/16 214,your_list,meh its an okay system,11.0.0.0/8 Example Command: SCRUTINIZER > import asns files/custom_asn.import

import csv_to_gps <csv_file> <group_name|group_id> [<create_new>] [<file_format>]	Uploads latitude and longitude locations of devices from a csv file and imports them into an existing Google map. The csv file must be located in the '/home/plixer/scrutinizer' directory. If the csv file is in '/home/plixer/scrutinizer/files/', enter 'files/[name_of_file]' as the file name. The csv file format is 'ip,latitude,longitude'. If the csv file format is different, specify that layout as the <file_format> command parameter. For example, "ip,lng,lat" 10.169.1.3,37.7749,122.4194 192.168.6.1,40.7128,74.0059 Provide either the group ID or group name in the arguments. The group_id can be determined by running show groups. Using the optional <create_new> parameter will add new objects if the IP address does not already exist. Example command: SCRUTINIZER> import csv_to_gps import_gps.import 3 Example command with <create_new> and different file format SCRUTINIZER> import csv_to_gps import_gps.import 3 create_new ip,lng,lat
import csv_to_membership <csv_file> <grouptype> [<file_format>]	Imports group definitions from a csv file. The csv file must be located in the '/home/plixer/scrutinizer' directory. If the csv file is in '/home/plixer/scrutinizer/files/', enter 'files/[name_of_file]' as the file name. The <grouptype> field refers to the map type that will be created if the group in the csv file does not already exist and can be either 'flash' or 'google'. The default csv file format is ipaddr,group. If the csv file format is different, specify that layout as <file_format> command parameter. EXAMPLE group,ipaddr 10.169.1.3,Routers 192.168.6.1,Firewalls
import hostfile	Imports a custom hosts.txt file that contains a list of IP Addresses and hostnames. The file format is: IPv4orIPv6Address HostName Optional Description Example: 10.1.1.4 my.scrutinizer.rocks The Best Software in my company The file must be located at /home/plixer/scrutinizer/files/hosts.txt.

import ipgroups [<path/file>] [reset]	Import ipgroup rules from a csv file. It is recommended to use this file for the ipgroups import csv file: /home/plixer/scrutinizer/files/ip_group.import A reset option can be passed which will remove all ipgroup rules before the bulk import. Each line of the file is an individual ipgroup with the name of the group as the first field and the rules of the group separated by a space in the second field. Supported rule types are subnet, single ip, ip range, wildcard and child rules. Any child groups must already exist in Scrutinizer or be declared in the import file BEFORE it can be used as a rule in another group. Valid ipgroup rule syntax is: 'subnet rule',10.0.0.0/8 'single ip rule',10.1.1.1 'range rule',10.0.0.1-10.0.0.42 'wildcard rule',10.0.0.1/0.255.255.0 'parent/child rule','my subnet' Up to 100,000 individual IpGroup rules are supported.

moloch

Third-party integration support for Moloch probes.

Command	Description
moloch <onloff> <moloch_ip> <moloch_port>	Manages integration with Moloch probes. The <moloch_port> parameter is optional.

optimize

Run various optimization tasks.

> **Warning:** These commands will alter database tables in Scrutinizer. Please use with caution.

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

remove

Removes a configured setting from the system.

Command	Description
remove address ipv6	Removes any IPv6 address configured, but there has to be an IPv4 address set up. Use the set myaddress command to change the addresses configured. Warning: This command will alter Scrutinizer's operations. Please use with caution.

repair

Runs various database check and repair commands.

Command	Description
repair business_hour_saved_reports	Saved reports prior to 15.5 that were saved with business hours will require a manual check and repair. This command converts older saved reports with business hours specified to the newer format.
repair database <db_name> <db_pass>	Repairs errors for the database specified.
repair history_tables	Fixes history tables that have the wrong col type for octetdeltacount. It may be updated in the future to address other issues.
repair policy_priority_order	With some professional services and automated policy creation, some policy IDs have been known to get out of whack (or duplicated). This function fixes that.
repair range_starts	Fixes history tables that may not have a start time that helps identify the range of data within the individual history tables. NOTE: This command may take a long time to complete. Only execute under the direction of technical support.

rotate

Rotates Scrutinizer's keys and certificates.

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

Command	Description
rotatekeys	Creates a new encryption key and re-encrypts all encrypted fields in the database.
rotatedbcerts	Creates new database certificates used for authentication.

services

Manages the Scrutinizer services.

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

Command	Description
services <service|all> <action>	Starts, stops, or restarts the specified service (or all services).

set

Modifies certain behaviors on how Scrutinizer authenticates and performs operations.

set columnmoniker <old_element> <new_element> [<element_list>]	Occasionally it is necessary to rename an information element. This is no problem for datareceived after the name has changed. However, if that element used in any reports it will no longer be possible to report on the historical data.columnmoniker takes 3 parameters. Two parameters are required: the <old_element>name and <new_element> name. The third optional parameter is list of info_elementsthat must also exist in the flow template to restrict renaming. This list can be one or moreelements separated by commas (e.g. elementname1,elementname2)Warning: This command should only be used under the instructions of technical support.
set dns	Modifies system file to manage list of dns servers. This command will remove any precon-figured dns servers. Use show dns to see what is currently configured.
set hostinfo <ip_address> <fqhn>	Sets the local machine name to the fully qualified host name provided Ensures that/etc/hosts is configured to resolve between the given <fqhn> and <ip_address>.
set httpd <port>	Changes the web port of non-ssl installs for the Scrutinizer WebUI. Use set ssl to changethe SSL port.
set myaddress <ip_address> <netmask> <gateway>	Changes the IPv4 address of the current Scrutinizer server. After entering the new IPinformation, you will be asked if the address provided is correct. Once you answer ‘yes’to the question, you will lose connection to the ssh session. Running this command from aconsole connection is advised.All fields are required.If you have multiple IP addresses on the Scrutinizer server or you have enabled encrypteddatabase communication, please contact Plixer support for assistance.Warning: This command will alter Scrutinizer’s operations. Please use with caution.

set myaddress <ipv6_address/cidr> <gateway>	Changes the IPv6 address of the current Scrutinizer server. After entering the new IPinformation, you will be asked if the address provided is correct. Once you answer 'yes'to the question, you will lose connection to the ssh session. Running this command from aconsole connection is advised.All fields are required. If you are setting an IPv6 address, netmask is not needed but cidrmust be added to the IP. You must provide the new IP address, netmask and gateway forIPv4 addresses. If you have multiple IP addresses on the Scrutinizer server or you haveenabled encrypted database communication, please contact Plixer support for assistance.Warning: This command will alter Scrutinizer's operations. Please use with caution.
set ntp	Modifies system file to manage list of ntp servers.
set partitions <partition_name> [extend]	Use this command to expand the operating system diskspace for hardware and virtual ap-pliances.If this is a virtual appliance and you expanded the existing disk, add the [extend] option.NOTE:make a backup before using this command.Warning: This command will alter Scrutinizer's operations. Please use with caution.
set password plixer	Resets the CentOS 'plixer' user's password.
set password webui <user>	Modifies the webui password for the specified user.
set permissions	Resets file and directory permissions to what is expected by Scrutinizer. Warning: This command will alter Scrutinizer's operations. Please use with caution.
set registercollector <collector_ip> [secondary]	Manually register a collector for distributed use. This command must be run from the primary reporting server. Adding 'secondary' to the command will register the collector as a secondary reporting server. Before registering a collector on AWS, make sure you have the key to collector used during deployment. This key must be available on the primary reporter. Warning: This command will alter Scrutinizer's operations. Please use it with caution.
set reportmenu	Manually recreates the report menu. NOTE: The report menu is automatically maintained based on the flows received.
set salt <salt>	Setting a salt value will allow users to mask certain machine characteristics from any license key generated.
set selfregister [reset]	Manually registers this Scrutinizer server to identify itself for both stand-alone or distributed functionality.
set selfreporter	Promotes this Scrutinizer Server to a reporter.

set sshcollectorkeys	Generates a new SSH key pair, and distribute it to all active, registered machines. Any previous SSH key pairs will be overwritten unconditionally, making this suitable for resynchronizing SSH access should problems arise. This enables future functionality to perform upgrades and other maintenance operations en masse.
set serverpref <serverpref> <value>	Changes the value of the serverpref setting. Use with caution.
set ssl <on|off> [ecc]	Enables or disables SSL support in Scrutinizer. It only works with the local Apache server bundled with Scrutinizer. Please reference the System/SSL section for detailed configuration instructions.
set timezone <timezone>	Sets the server's time zone. To see a list of time zones, run show tzlist
set tuning	This command will alter some operating system and Scrutinizer settings in these database tables: plixer.exporters and plixer.serverprefs; and these files: sysctl.conf, postgresql.conf , and plixer.ini.
set voip <on|off>	Toggles the predefinition of VoIP port ranges on or off.
set webui_timeout <seconds>	Resets the timeout for the WebUI. This command must be run on all all collectors/reporters. Warning: This command will alter how Scrutinizer and/or users access data. Please use with caution.
set yum_proxy <host> <port> <user>	Used to set up yum proxy setting in the yum configuration file. This command will remove any previously configured proxy servers. All fields are required. Once all fields are entered on the command line, a prompt for the users password will appear. To see what proxy servers are currently configured, use show yum_proxy

show

Shows various details about the Scrutinizer Server.

Command	Description
show alarms [filter]	Displays a list of alarms ordered by timestamp, descending.
show custom_algorithms	Displays a list of custom algorithms available and whether they are enabled. For information on managing custom algorithms, reference the Flow Analytics Custom Algorithm section.
show diskspace	Displays details about available storage.
show dns	Displays a list of DNS servers currently used to resolve hostnames. Use the set dns command to change the list of DNS servers.
show exporters [filter]	Displays a list of exporters that are currently sending data to Scrutinizer based on the supplied filter (if any).
show extalarms [filter]	Displays a list of alarms with extended json data ordered by timestamp, descending.
show groups	Displays a list of groups currently configured on this Scrutinizer server.
show interfaces [filter]	Displays a list of interfaces that are currently sending data to Scrutinizer based on the supplied filter (if any).
show ipaddresses	Displays the current ip address(es) on this Scrutinizer server.
show metering [filter]	Displays a list based on the supplied filter (if any) of matching exporter IPs and how each interface is metered (i.e. ingress and/or egress).
show ntp	Displays a list of NTP servers currently used to sync time.
show partitions	Displays a list of partitions on the current Scrutinizer Appliance. This command is only available for Hardware and Virtual Appliances. Use show diskspace if looking for diskspace per volume (or partition).
show pcaplist	List what current pcap files have been created and their sizes. Pcaps can be removed using the clean pcap command.
show serverpref [filter]	Displays serverprefs and their current values. The filter parameter is optional to narrow the serverprefs to match the string provided.
show task [name]	Displays a list of tasks currently configured in Scrutinizer. The name parameter is optional to narrow the task names to match the string provided.
show timezone	Displays the current timezone of this Scrutinizer Server. Use set timezone command to modify the timezone.
show tzlist [filter]	Displays the list of timezones.
show unknowncolumns	List info elements from exporters that are unknown to Scrutinizer. Don't fret! Give the list to Plixer and support will be added for it!
show yum_proxy	Displays the currently configured yum proxy settings. To change these settings, use set yum_proxy

Note: If after running the show command the results are long, 'q' can be typed in to quit and return to the SCRUTINIZER> prompt.

snoop

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

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

system

Scrutinizer system level functions.

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

Command	Description
system <restart|shutdown> system update [schedule|unschedule]	Performs system level functions such as rebooting, shutting down, or applying operating system level patches. To enable daily scheduled operating system updates, run the 'system update schedule' command. This will run the system update command every day at a random time. This time is selected outside of the 'business hours' set in Admin > Settings > Reporting. An alert is sent to Scrutinizer describing what time this command will run. To change the time, simply run the 'system update schedule' command again. A new time will be selected. To disable daily scheduled operating system updates, run the 'system update unschedule' command. If operating system patches are applied, all Scrutinizer services will be restarted and could cause a minute of missed data.

unlock

Unlocks accounts that have exceeded the maximum failed login attempts.

Command	Description
unlock <username> [<auth_method>]	Unlocks accounts that have exceeded the maximum failed login attempts set by the Scrutinizer administrator, and are locked out from authentication. By default, the user account will be set to local authentication. To specify another auth_method, use 'ldap', 'radius', or 'tacacs'.

update

Updates Scrutinizer product.

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

Command	Description
update Scrutinizer	Performs Scrutinizer product updates that are pulled from Plixer repositories. This command must be run from the primary reporting server. It will update Scrutinizer accross all collectors in the cluster. Warning: If operating system patches are applied, all Scrutinizer services be restart and could cause a minute of missed data.

upload

Uploads files for troubleshooting purposes.

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

version

Displays Scrutinizer version.

Command	Description
version	Shows version information about Scrutinizer.

2.16 Security updates

Security updates can be run on demand or scheduled to run on a daily basis. The commands to perform the security updates are listed below. Plixer recommends scheduled updates be enabled to ensure maximum protection. The update process will reach out to a plixer repository at files.plixer.com to pull down updates. Those updates have been applied to our QA servers internally and determined to be stable before being posted to our repository.

All patches and updates, including vulnerabilities and major upgrades, are included in the system updates. An audit event will be logged to the Scrutinizer Alarms table whenever the 'system update' command is run. Each update will also be listed within the audit event. In the event of a problem with the security updates, yum history can be used to roll back updates. If a proxy server is required, it can be configured within the yum.conf file.

The **system update** command runs yum update using https. Firewall policies will need to allow traffic to files.plixer.com on TCP port 443 from your Scrutinizer servers. Cryptographic verification of the downloaded update files is provided by yum.

The **system update** command will need to be run on each server in a distributed environment. It can be run directly from the command line and also from within the *Interactive scrut_util* utility. If operating system patches are applied, all Scrutinizer services will be restarted. This can cause a minute of missed data.

To pull down updates from plixer.com:

```
scrut_util --system update
```

To schedule the updates to be pulled from plixer.com on a daily basis:

```
scrut_util --system update --schedule
```

A random hour/minute is chosen to run the update. This time is selected outside of the 'business hours' set in the **Admin > Settings > Reporting** page. An alert is sent to Scrutinizer and can be viewed in the **Audit Event** policy in the Alarms tab.

To cancel the daily update schedule:

```
scrut_util --system update --unschedule
```

Interactive scrut_util

The following command syntax is used from within the *interactive mode* of the scrut_util utility. Running scrut_util from command line will open the interactive prompt.

To pull down updates from plixer.com:

```
SCRUTINIZER> system update
```

To schedule the updates to be pulled from plixer.com on a daily basis:

```
SCRUTINIZER> system update schedule
```

To cancel the daily update schedule:

```
SCRUTINIZER> system update unschedule
```

2.17 Sizing your environment

2.17.1 Overview

A single Plixer Scrutinizer collector instance can scale up to 100,000 fps sustained with spikes up to 200,000 fps, collecting from up to 500 exporters per collector. A distributed cluster can scale up to 50 collectors. That allows for sustained 5Mfps (spikes to 10Mfps) from up to 25000 exporters.

Flows per second	Flows per minute	Flows per hour	Flows per day
10K	600K	36M	864M
100K	6M	360M	8.6B
5M	300M	18B	432B

Processing 8.64 billion records a day will naturally require more than our minimum system specifications would allow for. This document will help you determine what resources are required.

Keep in mind there are many more factors than are outlined here, therefore requirements for some instances will vary.

2.17.2 Plixer Scrutinizer

This section contains information on Plixer Scrutinizer sizing.

Minimum specs

Our minimum system specifications are based on a max of 5kfps and 25 exporters. As system loads increase required resources will increase.

The big three

The big three variables are: CPUs, Memory, and Disk. Processing huge data volumes requires large amounts of all.

CPU: The requirements for CPU most closely correlates with the number of exporters coming in.

Memory: The requirements for memory most closely correlate with flow volume.

Disk: Disk IO closely correlates with flow rate. Disk size requirements will be a function of an organizations data retention needs.

Important: All recommendations are based off of DEDICATED resources and that SHARED CPUs, RAM, and disk may not perform up to the recommended levels.

If you are streaming to ML or an external data lake, collected flow rates will be 25% less or CPUs and RAM need to be 25% higher.

Single instance guidelines

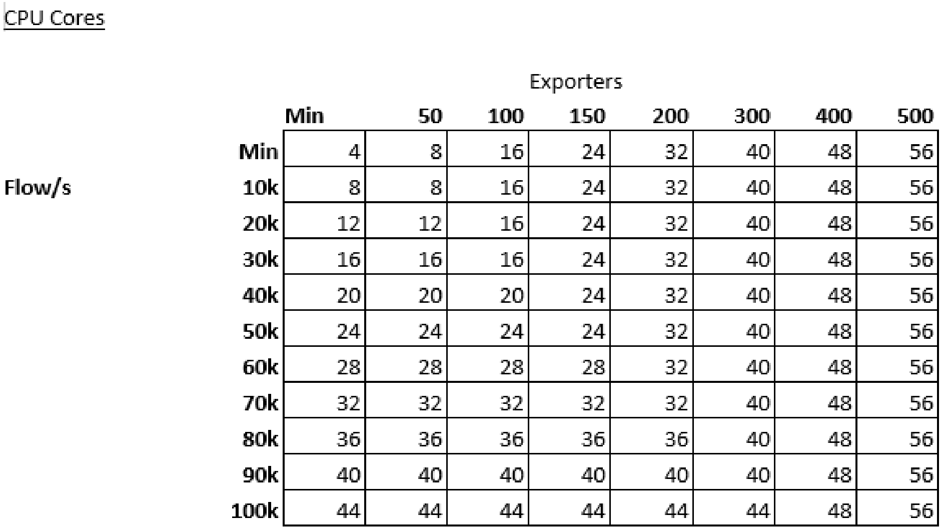

CPU Cores

Exporters

		Min	50	100	150	200	300	400	500
	Min	4	8	16	24	32	40	48	56
Flow/s	10k	8	8	16	24	32	40	48	56
	20k	12	12	16	24	32	40	48	56
	30k	16	16	16	24	32	40	48	56
	40k	20	20	20	24	32	40	48	56
	50k	24	24	24	24	32	40	48	56
	60k	28	28	28	28	32	40	48	56
	70k	32	32	32	32	32	40	48	56
	80k	36	36	36	36	36	40	48	56
	90k	40	40	40	40	40	40	48	56
	100k	44	44	44	44	44	44	48	56

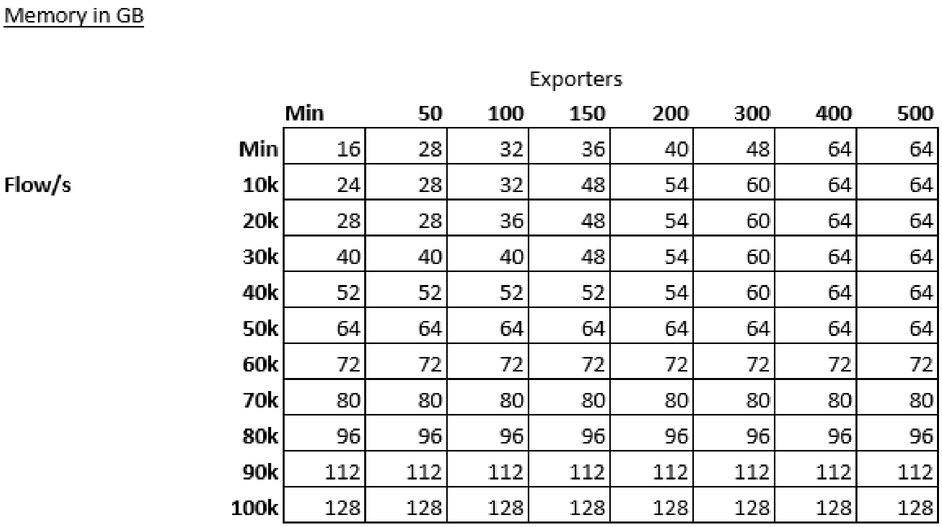

Memory in GB

Exporters

		Min	50	100	150	200	300	400	500
	Min	16	28	32	36	40	48	64	64
Flow/s	10k	24	28	32	48	54	60	64	64
	20k	28	28	36	48	54	60	64	64
	30k	40	40	40	48	54	60	64	64
	40k	52	52	52	52	54	60	64	64
	50k	64	64	64	64	64	64	64	64
	60k	72	72	72	72	72	72	72	72
	70k	80	80	80	80	80	80	80	80
	80k	96	96	96	96	96	96	96	96
	90k	112	112	112	112	112	112	112	112
	100k	128	128	128	128	128	128	128	128

Important: Plixer Scrutinizer is an IO intensive product. We recommend 15K drives or SSDs in RAID 10 for the best performance.

Reporter resources

With a distributed Scrutinizer deployment a number of servers work in concert. The reporter(s) act as the coordinator for all servers and therefore require resources in proportion with the number of servers.

Minimum CPUs: 2x servers in a cluster

Recommended CPUs: 4x servers in a cluster

Minimum Memory: 2GB per server in a cluster

Recommended Memory: 4 GB per server in a cluster

Take for example a distributed cluster with 10 collectors plus a dedicated reporter where the reporter is not collecting any external flow data. That reporter still has minim specs of 20 cores (we recommend 40) and 20 GB of RAM (we recommend 40 GB).

Additional considerations

Disk IO: In virtualized environments disk configurations and performance characteristics can vary greatly. Plixer Scrutinizer is a disk intensive application and avoiding waiting on disk is critical. There are too many factors that go into load on disk:

- Size in bytes of each flow record
- Cardinality of flow data
- Aggregation method selected

Features enabled: Overall load on a system will vary greatly depending on which features are being utilized and at what levels. Some of the features that can impact resource needs are:

- Number of Flow Analytics algorithms enabled and how many data sources are enabled
- Number of configured report thresholds
- Number of scheduled reports

Important: All flows are not the same: Performance will vary greatly depending on the size and complexity of the flows being collected.

- The simplest flow configuration is NetFlow v5, where each flow record is 48 bytes on the wire (excludes headers and Plixer enhancements. Bytes on disk will be different).
- More complex IPFIX templates can be well over 200 bytes per flow and include come complex structures like variable length strings that require more CPU to decode.

Note: Multiple templates matter: Multiple flow templates can add load like an additional exporter would.

- If an exporter is sending the same flows in two templates, for example sending both ingress and egress metered flows, the load on the system for one exporter feels just like two exporters.
- Option templates are small amounts of data sent infrequently so system impact is minimal. Recommended specs assume each exporter will be sending an option template.
- This document uses the measure of "exporter", because it simplifies things in almost all cases. If an exporter is sending additional template(s) with flow records it is safes to count that exporter as 2+ exporters.

2.17.3 Plixer Machine Learning Engine

This section contains information on Plixer Machine Learning Engine sizing.

Important: For PSI an "asset" is a host, for PNI an "asset" is an exporter interface.

CPU

Rows are flow per second (FPS), columns are number of assets supported. Measurement in number of cores.

	50	100	150	200	250	300	350	400	450	500
10k	8	12	16	20	24	28	32	36	40	44
20k	12	14	18	22	26	30	34	38	42	46
30k	16	18	20	24	28	32	36	40	44	48
40k	20	22	24	26	30	34	38	42	46	50
50k	24	26	28	30	32	36	40	44	48	52
60k	28	30	32	34	36	38	42	46	50	54
70k	32	34	36	38	40	42	46	50	54	56
80k	36	38	40	42	44	46	50	54	56	56
90k	40	42	44	46	48	52	54	56	56	56
100k	44	46	48	50	52	54	56	56	56	56

Memory

Rows are FPS, columns are number of assets supported. Measurements in GB.

	50	100	150	200	250	300	350	400	450	500
10k	40	80	112	136	160	184	208	232	256	256
20k	80	112	136	160	184	208	232	244	256	288
30k	112	136	160	184	208	232	244	256	288	320
40k	136	160	184	208	232	244	256	288	320	352
50k	160	184	208	232	244	256	288	320	352	384
60k	184	208	232	244	256	288	320	352	384	416
70k	208	232	244	256	288	352	352	384	448	448
80k	232	256	288	320	352	384	416	448	480	480
90k	256	288	320	352	384	416	448	480	512	512
100k	256	288	320	352	384	416	448	480	512	512

Disk

Rows are FPS, columns are number of assets supported. Measurements in TB.

	50	100	150	200	250	300	350	400	450	500
10k	.2	.4	.6	.8	1	1.2	1.4	1.6	1.8	2
20k	.4	.6	.8	1	1.2	1.4	1.6	1.8	2	2.2
30k	.6	.8	1	1.2	1.4	1.6	1.8	2	2.2	2.4
40k	.8	1	1.2	1.4	1.6	1.8	2	2.2	2.4	2.6
50k	1	1.2	1.4	1.6	1.8	2	2.2	2.4	2.6	2.8
60k	1.2	1.4	1.6	1.8	2	2.2	2.4	2.6	2.8	3
70k	1.4	1.6	1.8	2	2.2	2.4	2.6	2.8	3	3.2
80k	1.6	1.8	2	2.2	2.4	2.6	2.8	3	3.2	3.4
90k	1.8	2	2.2	2.4	2.6	2.8	3	3.2	3.4	3.6
100k	2	2.2	2.4	2.6	2.8	3	3.2	3.4	3.6	3.6

2.18 Configuring SSL in Scrutinizer

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

Important: Scrutinizer AMIs come with a self-signed certificate. Disable SSL prior to creating a new certificate request.

To open the interactive scrut_util, use the following command:

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

The Scrutinizer prompt will then display:

```
SCRUTINIZER>
```

2.18.1 Enabling SSL

To enable SSL, at the Scrutinizer prompt, enter:

```
SCRUTINIZER> set ssl on
```

Warning: This command will alter Scrutinizer's operations. Please use with caution. Scrutinizer will then issue the following prompt for these mandatory fields:

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

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

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

Enter the city of the organization. ex: Kennebunk

Enter the name of the organization. ex: Plixer

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

Enter the server name or IP of the Scrutinizer server.

> ex: 1.2.3.4 or scrutinizer.company.com

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

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

Note: The optional argument 'ecc' can be used if you would like to generate a 256b Elliptical Curve public/private key pair.

2.18.2 Creating a signed certificate

Important: Scrutinizer AMIs come with a self-signed certificate. To create a new certificate request, disable SSL using the interactive scrut_util *set ssl off* command.

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

```
SCRUTINIZER> services httpd stop
```

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

```
SCRUTINIZER> services httpd start
```

2.18.3 Disabling SSL

To disable SSL, at the Scrutinizer prompt, enter:

```
SCRUTINIZER> set ssl off
```

2.19 Streaming support for customer data lakes

Important: The configuration is done via the database. Contact Plixer for assistance with setting this up.

2.20 System LEDs

The system LEDs are shown in the upper right hand corner providing the status on server health and other various critical operations of the flow collection and reporting architecture. Preceding the LEDs is the letter **P** indicating that this is a Primary reporting server. In a single server installation of Scrutinizer, it will always be **P**. In a distributed server environment, this letter can alternatively be **S** for the Secondary reporting server. The LED's color and the icon indicate the status of the LED:

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

The LED status is based on the most critical level of severity of any item within the detail of the LED modal. All columns within each LED modal are both sortable and searchable.

2.20.1 Scrutinizer Server Health

P [x] [ ] [ ]
Scrutinizer Server Health: [status]

This LED provides vital server statistical information such as CPU utilization and memory and disk storage availability. The information below is available per Server, with the color change occurring at the threshold levels defined.

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

	Description	Thresholds		
		Green	Orange	Red
CPU	CPU Utilization percentage	<=60%	<=85%	>85%
Memory	Available Memory	>8GB	>=4GB	<4GB
Free Memory	Available Free memory	>=2GB	>=1GB	<1GB
Free Disk DB*	Available Free Disk space for database	>10%	>4%	>2.5%
Free DB%	Percent of disk storage that is still available for database	N/A	N/A	N/A
DB Latency	Server to server database latency (in milliseconds)	<250ms	>=250ms	N/A
API Latency	Server to server reporting latency (in milliseconds)	<1000ms	>=1000ms	N/A
Clock Drift	Server to server clock difference (in seconds)	0s	<>0s	N/A

- **Free Disk DB**

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

  If disk space hits 2.5% of available space, the collector will stop saving flows. In the event that the collector stops, a utility can be run that expires historical data to free up space. Go to Admin>Settings>Data History, and adjust the current retention settings.

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

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

At the scrut_util prompt, run:

> **SCRUTINIZER>** expire history

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

> **SCRUTINIZER>** services plixer_flow_collector start

2.20.2 Scrutinizer Software Health

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

The following information is available per Server from this LED:

- **Role** – primary/secondary/collector (secondary and collector used in distributed server environments)

- **API** – Reporting (web) interface, Up/Green means the web server is running. Down/Red means the web server is down.

- **Database** – Up/Green means the database is running. Down/Red means the database is down.

- **Collector** - Up/Green means the collector service is running. Down/Red means the collector service is down.

- **Alarms** – Up/Green means the Alarms/syslog service is running. Down/Red means that the Alarms service is down.

- **Version** – current running Scrutinizer version

- **Flow Rate** – flows per second

- **MFSNs – Missed Flow Sequence Numbers** (see Scrutinizer exporter health LED for more information on MFSNs)

2.20.3 Scrutinizer Exporter Health

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

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

- To run the collector from the command prompt (i.e. CLI) type:

  /home/plixer/scrutinizer/bin/scrut_collector.exe

Note: The output from this command is for internal use only.

Information available per exporter:

- **Collector** – IP Address of the server collecting the flows
- **Exporter** – IP Address/name of exporter
- **Flows** – flow collection rate (green=flows/orange=none)
- **Packets** – packet rate (green=packets/orange=none)
- **MFSN** – This led turns orange if Missed Flow Sequence Numbers exceed 300 in an interval.

  If only one or a few of the flow sending devices report high MFSNs, it is likely the network or the flow exporting device that is dropping or skipping flows. If all devices report high MFSNs, it is likely to be the collector that is dropping flows. To improve performance, make sure the server hardware meets the minimum requirements. Visit the *Vitals dashboard* for trending details of the server.

- **Max Flow Duration** –

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

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

Learn more about the ip flow-cache timeout commands.

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

2.21 Third-party licenses

Certain open source or other third-party software components are integrated and/or redistributed Scrutinizer software. The licenses are reproduced here in accordance with their licensing terms, these terms only apply to the libraries themselves, not Scrutinizer software. Copies of the following licenses can be found in the licenses directory at /home/plixer/scrutinizer/files/licenses/

2.21.1 Licensed under 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

2.21.2 Licensed under Artistic 1.0 License

business–isbn

https://github.com/briandfoy/business-isbn/

Copyright (c) 2001-2013, Brian D Foy

Common-Sense

http://search.cpan.org/~mlehmann/common-sense/

Terms of Perl - No Copyright Author - Marc Lehmann

Compress-Raw-Zlib

http://search.cpan.org/~pmqs/Compress-Raw-Zlib/

Copyright (c) 2005-2009 Paul Marquess.

Compress-Zlib

http://search.cpan.org/~pmqs/IO-Compress-2.066/lib/Compress/Zlib.pm

Copyright (c) 1995-2009 Paul Marquess.

crypt-ssleay

https://github.com/gisle/crypt-ssleay/

Copyright (c) 2006-2007 David Landgren, Copyright (c) 1999-2003 Joshua Chamas, Copyright (c) 1998 Gisle Aas, Copyright (c) 2010-2012 A. Sinan Unur

DBD-mysql

http://search.cpan.org/dist/DBD-mysql/

Large Portions Copyright (c) 2004-2013 Patrick Galbraith, 2004-2006 Alexey Stroganov, 2003-2005 Rudolf Lippan, 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

2.21.3 Licensed under 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 theparticipants in the CRC for Avanced Computational Systems ('ACSys').

2.21.4 Licensed under BSD 2-Clause Simplified License

JabberWerxC

https://github.com/cisco/JabberWerxC

Copyright (c) 2010-2013 Cisco Systems, Inc.

2.21.5 Licensed under 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.

2.21.6 Licensed under 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.

2.21.7 Licensed under CDDL 1.1 License

HK2
https://javaee.github.io/hk2/
Copyright (c) 2010-2017 Oracle and/or its affiliates.

2.21.8 Licensed under CURL License

cURL
http://curl.haxx.se
Copyright (c) 1998 - 2013, Daniel Stenberg

2.21.9 Licensed Dually under 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

2.21.10 Licensed under 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

sshpass
http://freshmeat.net/projects/sshpass

sysstat
http://sebastien.godard.pagesperso-orange.fr/
Copyright (c) 1999-2009 Sebastien Godard

2.21.11 Licensed under GPL 3.0 License

Ansible
http://www.ansible.com/
Copyright (c) 2017, Ansible Project

MariaDB
http://mariadb.org/
Copyright (c) The MariaDB Foundation

2.21.12 Licensed under 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

2.21.13 Licensed under LGPL 3.0 License

GNU Libidn

http://www.gnu.org/software/libidn/

Copyright (c) 2004-2012 Simon Josefsson

2.21.14 Licensed under 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

2.21.15 Licensed under 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

2.21.16 Licensed under Mozilla Public License 1.1

Rhino
https://github.com/mozilla/rhino

2.21.17 Licensed under 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.

2.21.18 Licensed under Oracle BCL License

Oracle Java
http://www.oracle.com/technetwork/java/index.html
Copyright (c) 1993 - 2015, Oracle and/or its affiliates.

2.21.19 Licensed under 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

2.21.20 Licensed under 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.

2.22 Upgrade guide

Important: Please contact Plixer technical support for assistance with upgrades.

CHAPTER 3

Machine learning

3.1 ML Engine AMI deployment guide

3.1.1 What you need to know about ML Engine AMIs

The latest ML Engine AMI can be obtained from Plixer or your local reseller. Please contact support if you do not already have the Plixer ML Engine AMI. You will need to know your AWS Region and your AWS account ID so the AMI link can be shared with you.

- Contact Plixer technical support to discuss the recommended instance type for production environments.
- Decide on the VPC and security group rules that fit the needs of your organization. You will need to specify these in the deployment process.
- Do not lose the SSH key that you will be asked to create in the deployment process. This key is the only way to access the server via SSH.

3.1.2 Pre-deployment checklist

Please provide a technical support engineer with the following information:

- Amazon account number;
- Region you are planning to deploy an instance in;
- Expected flow rate.

3.1.3 Deploying AMI

Important: Please contact Plixer for assistance with installing and configuring an ML Engine AMI.

3.2 ML Engine Virtual Appliance deployment guide

3.2.1 What you need to know about deploying a ML Engine Virtual Appliance

The ML Engine Virtual Appliance can be obtained from Plixer or your local reseller. It is downloaded as an all-in-one virtual appliance which can be deployed on an ESXi v5.5 and above or Hyper-V 2012 hypervisor.

- You will need to obtain an appliance license or evaluation license from Plixer or your local reseller in order for the ML Engine Virtual Appliance to function properly.
- The ML Engine Virtual Appliance is deployed on a hypervisor server.
- The performance you get out of a ML Engine Virtual Appliance will be directly dependent on the hardware on which it's deployed. It's recommended to dedicate, not share, all the resources that are allocated to the ML Engine virtual machine. This is especially important for the ML Engine datastores. In environments with high volumes of NetFlow data, ML Engine will require dedicated datastores which are discussed in further detail later in this document. ML Engine hardware appliances are recommended for deployments of exceedingly high volume of flow as they are designed to handle the highest flow rates.

  Important: Please contact Plixer for assistance with deploying an ML Engine Virtual Appliance.

3.3 Forecasting with Plixer Network Intelligence

Learning from the network metadata collected by Plixer Scrutinizer, forecasting provides insight into future network behavior. Data is selected via the Plixer Scrutinizer reporting interface and is supported with all report types. This allows for the use of reporting filters to select specific classes of traffic.

3.3.1 Forecast horizon

A forecast horizon determines how far into the future we want to forecast. In the initial release of forecasting, the horizon is determined by the volume of input data, ensuring statistically sound results. Future releases will allow for advanced configuration options where the user can balance forecast length and forecast variance.

Forecast horizons are determined as follows:

If the input data is between 0 and 3 days:

- No seasonality applied, likely result is mean reversion unless there is a legitimately strong AC relation
- Forecast horizon = 1/3rd the original report width, 1/4 of the total view with forecast appended

If the input data is between 3 and 14 days:

- Daily seasonality selected
- Forecast horizon = floor round # of seasons (days) in report / forecast_horizon_ratio (2 by default)

If the input data is between 14 and 45 days:

- Weekly seasonality selected
- Forecast horizon = floor round # of seasons (weeks) in report / forecast_horizon_ratio (2 by default)

If input data is 45+ days:

- Monthly seasonality selected (likely to be deprecated for weekly season)
- Forecast horizon = floor round # of seasons (months) in report / forecast_horizon_ratio (2 by default)

3.3.2 Creating a forecast

Forecast creation begins by creating a report that represents the traffic classes to forecast. The quantity being forecast will be the primary data column from the initial report (bytes, packets, counts, milliseconds of latency, etc). When selecting the time bounds of the report, try to provide as much history as possible, while excluding regions behavior should not be considered (i.e., a period of 0 utilization before a new application was implemented/observed in NetFlow).

The data source for the report should be as granular as possible, 30m is optimal for most mid-to-long term forecasts (days/weeks ahead), 2hr is optimal for long term forecasts (months/quarters). High-resolution data 1m/5m should be used for non-seasonal "real-time" monitoring. Application performance metrics such as stream jitter or various qualities of TCP latency should be monitored on a LastX time interval with the highest possible resolution reporting performance allows.

When the appropriate data has been selected, click the **Plan Position Indicator Plot** icon, located in the upper right hand corner of the reporting interface.

When prompted, give the forecast a title & submit. At this point you will be forwarded to the forecast menu (**Investigate>Forecasts**) forecast menu (**Investigate>Forecasts**)

The forecast menu lists all forecasting tasks. The primary view features the forecast id#s, the forecast title, a link to the source report used as input data, created by, the status of the forecast service in relation to that task, and data-ready date.

The refresh icon allows the user to update the forecast from the forecast menu. Input data reports with dynamic (LastX) time selections will remain up to date on refresh.

3.3.3 Forecast menu (Investigate>Forecasts)

The forecast menu lists all forecasting tasks. The primary view features the forecast id#s, the forecast title, a link to the source report used as input data, created by, the status of the forecast service in relation to that task, and data-ready date.

The refresh icon allows the user to update the forecast from the forecast menu. Input data reports with dynamic (LastX) time selections will remain up to date on refresh.

3.3.4 Forecasting service status

Forecasting jobs are processed within the Plixer ML Engine, remote to Plixer Scrutinizer. The status field of the forecast menu displays the current status of a task on the remote system.

1. **Initializing:** Plixer Scrutinizer has made the task available for consumption by the forecasting service
2. **Starting:** The forecasting service as accepted the task and queued it to run
3. **Data Retrieval:** The Plixer Scrutinizer reporting API is queried by forecasting, to collect input data.
4. **Processing:** Input data received and pre-processed for model training
5. **Strategy Selection:** Optimal method for producing forecast chosen based on input data characteristics
6. **Learning:** The forecasting model is learned from provided input data
7. **Prediction:** A prediction is made for the forecast horizon and results are returned to Plixer Scrutinizer
8. **Complete:** The remote forecasting service has provided results and closed this session

Clicking on any forecast title will open the forecast display.

3.3.5 Forecast display

The forecast display shows a plot of the forecast and the input data. Each line segment will active on mouse hover, showing the critical boundaries of the forecast in the shaded region. This screen is intended to help users visually verify the forecast provided.

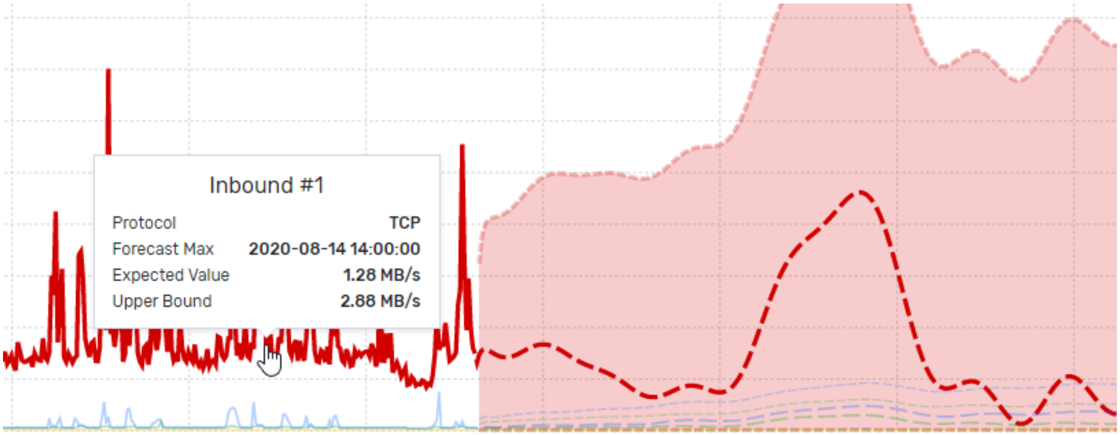

Inbound #1

Protocol	TCP
Forecast Max	2020-08-14 14:00:00
Expected Value	1.28 MB/s
Upper Bound	2.88 MB/s

The data provided in the tools tips and the table below provide insight into the time interval of highest expected usage, easily allowing for the identification of hot spots in network time.

CHAPTER 4

Integration guides

4.1 Advanced Threat Intelligence Feed

The Advanced Threat Intelligence Feed is the premium functionality offered as part of Plixer's Security Intelligence (PSI). It expands coverage of malicious IPs and domains to protect your network.

Important: Additional licensing is required for this feature. Contact Plixer support for assistance.

The capabilities and expanded coverage are automatically enabled and updated. No further action by the administrator is required.

4.2 Amazon Web Services flow logs

4.2.1 Overview

The integration between Amazon Web Services (AWS) and Plixer Scrutinizer provides insight into network traffic destined for AWS, such as top AWS users, top AWS applications, as well as overall traffic load of AWS hosted applications. After configuring Amazon Web Services Flow Log integration, the following reports become available in Plixer Scrutinizer:

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

4.2.2 Prerequisites

The following information is required to configure AWS flow logs integration:

1. AWS IDs and secrets with full access permission to the FlowLog S3 buckets and necessary permissions to collect descriptions OR Plixer Scrutinizer needs to be running in AWS where the EC2 instance is assigned a role with those permissions.

Hint: The VPC(s) you want to monitor need to be configured to send flow logs to the S3 buckets Plixer Scrutinizer is configured to read.

Note: These S3 buckets are solely intended for Plixer Scrutinizer's use. Plixer Scrutinizer will delete the logs from the buckets as it collects them.

2. The region that hosts the S3 bucket.
3. It is strongly recommended to include the two log fields listed below:

log-status

This is a version 2 default log field. Customized VPC logs without the field will not be collected until it is re-added.

Important: VPC flow logs that do not include this field will be discarded.

vpc-id

This is a version 3 custom log field. Users running the default version 2 logs will need to add it to start log collection. Plixer Scrutinizer will attempt to generate exporter IDs using alternate data if vpc-id is not present. These fallbacks will usually generate more exporters than using vpc-id. In addition to potentially exceeding license limits on exporter counts, egregious exporters may also impair Plixer Scrutinizer's flow collection rates and reporting times.

4. Exporter IDs will change for the upgraded deployments. v19.0.1 Plixer Scrutinizer sources exporter IDs from the vpc-id log field and the AWS account number in the directory path. Prior to v 19.0.1, exporter IDs were determined by the AWS S3 bucket name. They will stop exporting after an upgrade.
5. Installations with several VPCs will have several new exporters. The new VPC-based exporters can potentially exceed license limits.

Note: Review the status of your exporters via the Admin>Definitions>Manage Exporters page if you are not collecting data.

4.2.3 Configuring AWS flow logs

1. Navigate to the **Admin > Settings > AWS Flow Logs S3** page.
2. Click “Add” to create a new flow log source in Plixer Scrutinizer. A single S3 bucket can comprise data from several of the sources we consider to be exporters.
3. Provide a unique name for the Flow Log source.
4. Select the collector that will communicate with AWS to receive data for this bucket.
5. Enter the bucket name, region, ID, and Secret.
6. Save the entry.

4.2.4 Enabling IAM role-based authentication for S3 buckets

1. Navigate to the **Admin > Settings > AWS Flow Logs S3** page of your Plixer Scrutinizer AMI deployment.
2. Click “Add” to create a new flow log source.
3. Provide a unique name for the flow log source.
4. Select the collector that will communicate with AWS to receive data for this bucket.
5. Enter the bucket name and its region.
6. Check the box to enable IAM role-based authentication
7. Save the entry.

The necessary policy permissions that a role must have are as follows:

```
{ "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:::<S3BUCKET>"
             }
  \]
}
```

Hint: <S3BUCKET> should be replaced with the name of the bucket you are using. The “Version” element is an AWS-provided version string that refers to the version of the policy specification it adheres to, i.e. only change this value to something that you know AWS will support and only if necessary.

4.2.5 Importing descriptions for AWS entity IDs

The AWS entity IDs import functionality provides reporting descriptions and filtering for any AWS entity identifiers, such as interface-id, vpc-id, subnet-id, and instance-id. With the feature configured, a user will not need to use the AWS console or documentation to determine the meaning of the identifier.

1. The following user permissions are required:

```
ec2:DescribeInstances
ec2:DescribeSubnets
ec2:DescribeVpcs
ec2:DescribeNetworkInterfaces
```

For example, the policy below can be applied to a user/IAM role:

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Sid": "VisualEditor0",
            "Effect": "Allow",
            "Action": [
                "ec2:DescribeInstances",
                "ec2:DescribeNetworkInterfaces",
                "ec2:DescribeVpcs",
                "ec2:DescribeSubnets"
            ],
            "Resource": "*"
        }
    ]
 }
```

2. SSH into the Plixer Scrutinizer standalone or a primary reporter appliance in the distributed cluster. Run the interactive CLI command to sync IDs and descriptions from AWS.

```
SCRUTINIZER> awssync
AWS entities synced!
```

3. Launch an AWS specific report. The identifiers will be replaced with their appropriate descriptions.

Note: The awssync task is scheduled to run hourly.

4.2.6 Helpful tips

Amazon flow logs are updated every 1 minute or every 10 minutes. Choosing every 10 minutes will result in data spikes and longer data update times. There will, however, be less processing load on the Plixer Scrutinizer server. Choosing every 1 minute will result in better reporting granularity.

If you are not seeing an exporter:

- check the collector log for errors;
- go to the AWS interface and make sure you see flow logs in the configured bucket;
- edit the S3 profile in Plixer Scrutinizer and use the "Test" button to make sure the configuration is correct;
- make sure the exporter is not disabled under **Admin > Definitions > Manage Exporters**.

4.3 Elasticsearch / Kibana (ELK) integration

4.3.1 Overview

What does ELK integration with Scrutinizer do?

- From within Scrutinizer, searches for IP Addresses in Kibana can be launched.
- From Kibana, users can view *Scrutinizer Vitals* and *FA TopN gadget details*.

4.3.2 What is ELK?

Elasticsearch - A searching service to look through the stored data collected by Logstash.

Logstash - A means to collect logs and events (like syslogs) and filter them in a specific way to be stored for later analysis.

Kibana - Front-end to present data by creating dashboards and visualizations, similar to Scrutinizer's dashboards and gadgets.

4.3.3 Integration prerequisites

Kibana 4.2

Scrutinizer v15.12+

Note: The following configuration instructions apply to Scrutinizer v16.7 and later. If an earlier version of Scrutinizer is installed, contact plixer for assistance.

4.3.4 Kibana searches from Scrutinizer reports and alarms

The following steps will walk through how to search Kibana's database from Scrutinizer reports.

1. Select a Scrutinizer report that includes IP addresses

   a. Select a host of interest and click on it

   b. Select 'Other Options' from the Reports menu

   c. Select 'Kibana (ELK)'

2. The IP address and report timeframe are passed to Kibana's search engine for detailed Kibana reporting.

For more detail (from Kibana) on Scrutinizer Alarms, follow these steps.

1. Go to the *Alarms tab* in Scrutinizer and select either:

   a. Bulletin Board by Violator : Select a violator

   b. Bulletin Board by Policy : Select the policy

2. In the Bulletin Board Events view that opens, click on the dropdown arrow to the left of the Message column for the alarm that was selected.

3. Select 'Kibana (ELK)' from the Available Options menu and the Violator's IP address and timeframe of the violation are passed to Kibana.

   - If the Violator's IP address and an alarm time (not timeframe) are being passed, then 30 minutes before and after the alarm time is searched.

4.3.5 Scrutinizer reporting from within Kibana

Within the Kibana (ELK) integration with Scrutinizer, dashboards can be setup which include:

- Scrutinizer Vitals information
- Flow Analytics TopN Algorithms

The Scrutinizer Vitals dashboard in Kibana can include:

- CPU
- Memory
- Disk Usage
- Flows per collector
- Status per collector

Dashboards created with the TopN Algorithm gadgets from Flow Analytics include:

- Top Applications
- Top Countries
- Top Rev 2nd lvl Domains (Top reverse 2nd level domains)
- Top Flows
- Top Hosts
- Top Jitter
- Top Networks

4.3.6 How to configure ELK integration with Scrutinizer

There are two components to the ELK Integration.

1. Preparing Scrutinizer
2. Importing the necessary files into Kibana's UI

Preparing Scrutinizer

Note: Flow Analytics must be enabled and collecting statistics for the Top X Algorithms.

1. SSH into the Scrutinizer server as the plixer user.
2. Use the interactive scrut_util command:

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

  **SCRUTINIZER>** enable elk http://<ip:port>

where <ip:port> is the ELK server's IP address and port.
```

1. After a few moments, Scrutinizer will begin to send events to ELK.

- To test the data export, from within the scrut_util shell, run:

```
collect elk <elk_ip>
```

- To disable the data export, run:

```
disable elk http://<ip:port>
```

Preparing Kibana

Integrating ELK with Scrutinizer displays details in Kibana that have been collected and processed by Scrutinizer. For more information, visit Plixer's Elasticsearch / Kibana Integration page.

1. After enabling the ELK integration on Scrutinizer, refresh the index on Logstash in order to get Scrutinizer's fields to show up. In Kibana, go to Indices > Logstash. Click on the Reload field list icon at the center top of the screen.

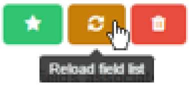

2. Download the Kibana Integration Plugin from the Elasticsearch / Kibana Integration page. Extract the files from the scrutinizer-elk.zip file.

3. In Kibana, go to **Settings > Objects > Import > Visualizations** and navigate to the elk-scrutinizer-visualizations.json file extracted in Step 2a above and click Open.

4. Go to **Settings > Objects > Import > Dashboards** and navigate to the elk-scrutinizer-dashboards.json file extracted in Step 2a above and click Open.

5. The Kibana dashboards and visualizations are now all imported and events have been configured to be coming from Scrutinizer.

6. In the **Visualize** tab, scroll to the bottom and filter for a specific visualization. Typing Scrutinizer in the filter for example, will show all of the Scrutinizer visualizations.

7. In the **Dashboard** tab, navigate around the various Scrutinizer dashboards imported.

4.4 Endace probe integration

4.4.1 Overview

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

With this integration, Plixer Scrutinizer allows the user to quickly filter down to to certain flow data related to the issue, then the Endace probes can be selected from the reporting menus to download just the packets related to the specific flow data observed in Plixer Scrutinizer.

4.4.2 Setting up Endace packet capture integration

In order to configure Plixer Scrutinizer to download the packet captures from Endace probes, the probe must first be added (enabled) via the interactive scrut_util utility on the Plixer Scrutinizer appliance.

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

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

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

- Add a probe:

```
SCRUTINIZER> endace add
```

- Remove a probe:

```
SCRUTINIZER> endace remove
```

- Change/update a probe:

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

4.4.3 Accessing Endace probes in Plixer Scrutinizer

There are three ways to access the probes from within Plixer Scrutinizer:

1. *Vendor Specific Menu*
2. *Violated Alarms*
3. *Reports with IP Addresses*

Vendor specific menu

From the **Status tab > Vendor Specific** menu, select an **Endace Probes** option. This option allows the user to access the Endace device without being in a report. This is especially handy when the users already knows the IP addresses, protocol, and ports. Take the following steps to access Endace packet captures via the Endace Probes option.

1. Select the Status tab
2. Select the Vendor specific menu.
3. Select the Endace probes menu entry.
4. Complete the following fields:

- Initiator IP
- Target IP
- Protocol
- Initiator and Target Port fields are optional

1. Click Search.
2. Download the pcap and open it in the desired packet analyzer.

Violated alarms

You can get more details regarded the violations from the Alarms. This can be done be using the following instructions to access Endace reports from violated alarms.

1. Select the Alarms tab
2. If looking for a specific Alarm type

a. Select **Views > Bulletin Board by Policy**.

b. Select the **Policy Violated** desired to retrieve the packet details

c. Or expand the **Violators list** for that policy and select the violator.

3. If looking for a specific violator

a. Select **Views > Bulletin Board by Violator**;

b. Select the violator address to get the packet details,

c. Or expand the **Policies Violated** list and select the Policy that packet details are desired for.

4. In the **Bulletin Board Events** page, click the dropdown arrow between the Board Name and Message columns
5. Select **Endace Probes**. Any relevant details from the alarm are pre-populated. This is useful because the actual packets from the conversations that triggered alarms become available.
6. Click **Search**.
7. Download the pcap and open in your favorite packet analyzer.

Reports with IP addresses

The reports with IP addresses option allows the user to select the source or destination IP Address (or DNS Name) from a report and get information from the Endace probes.

To investigate the conversation further, launch the **Flow** report.

1. Start within a report that includes source and destination IP addresses.
2. Select an IP address from the report.
3. Select **Other Options** from the drop-down menu.
4. Select **Endace Probes**. Any relevant details from the conversation are pre-populated. This is useful because the user can get to the actual packets from the conversation.
5. Click **Search**.
6. Download the pcap and open it in the packet analyzer.

4.5 Cisco's FireSIGHT eStreamer client

4.5.1 Overview

Cisco FireSIGHT Management Center manages network security and operational functions for Cisco ASA with Firepower Services and Cisco Firepower network security appliances. Configuring the FireSIGHT eStreamer client to send flows to Plixer Scrutinizer will make the following flow reports available:

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

Important: The minimum supported version of eStreamer is 5.4.

4.5.2 Registering Plixer Scrutinizer with FireSIGHT

In the configuration example below, Plixer Scrutinizer collector's IP address is 10.30.11.5. 10.1.2.70 is the FireSIGHT eStreamer IP address.

1. Log into the FireSIGHT Defense Center.

For Firepower v5.4: navigate to System > Local > Registration:

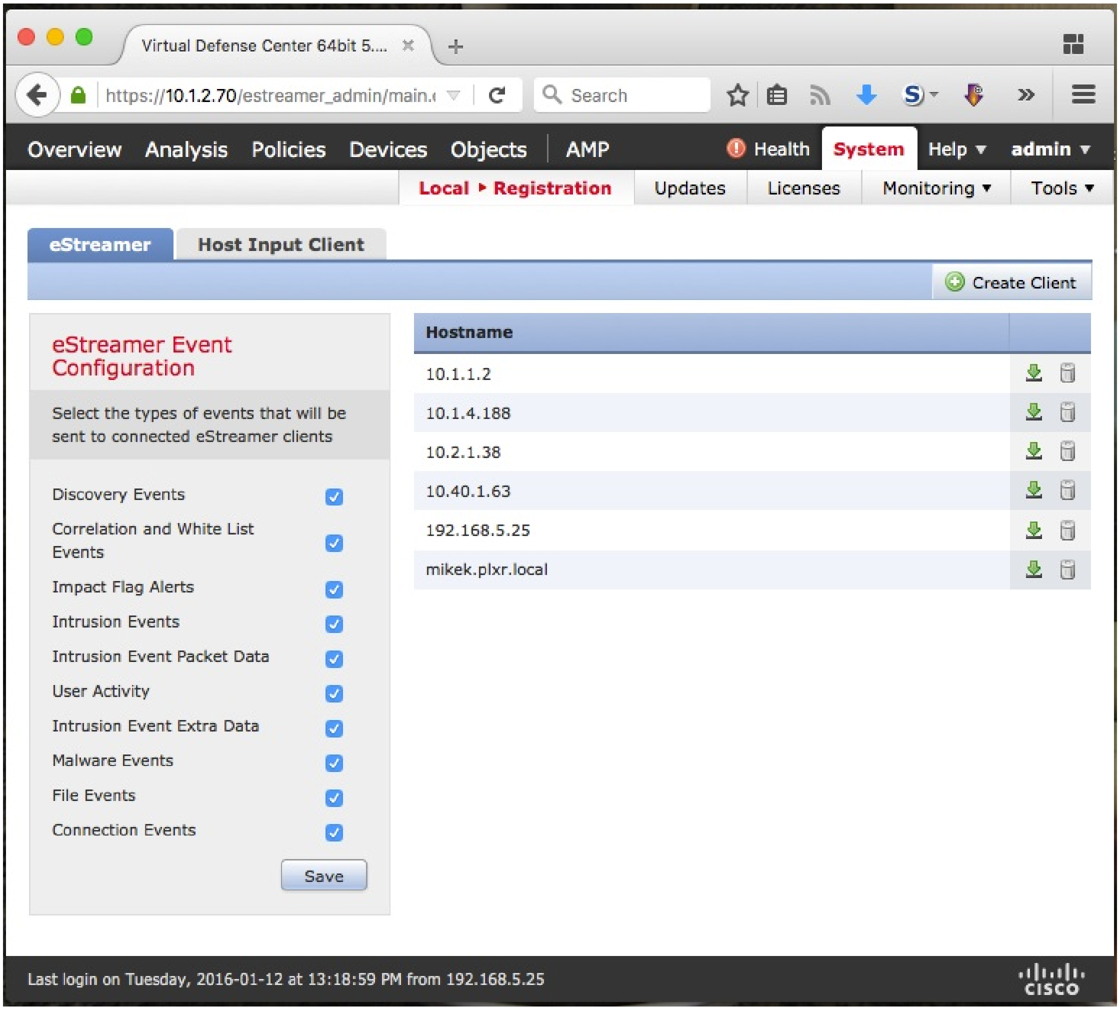

eStreamer Event Configuration

Select the types of events that will be sent to connected eStreamer clients

Discovery Events	☑
Correlation and White List Events	☑
Impact Flag Alerts	☑
Intrusion Events	☑
Intrusion Event Packet Data	☑
User Activity	☑
Intrusion Event Extra Data	☑
Malware Events	☑
File Events	☑
Connection Events	☑

Hostname
10.1.1.2
10.1.4.188
10.2.1.38
10.40.1.63
192.168.5.25
mikek.plxr.local

Last login on Tuesday, 2016-01-12 at 13:18:59 PM from 192.168.5.25

For Firepower v6.x: navigate to System > Integration > eStreamer:

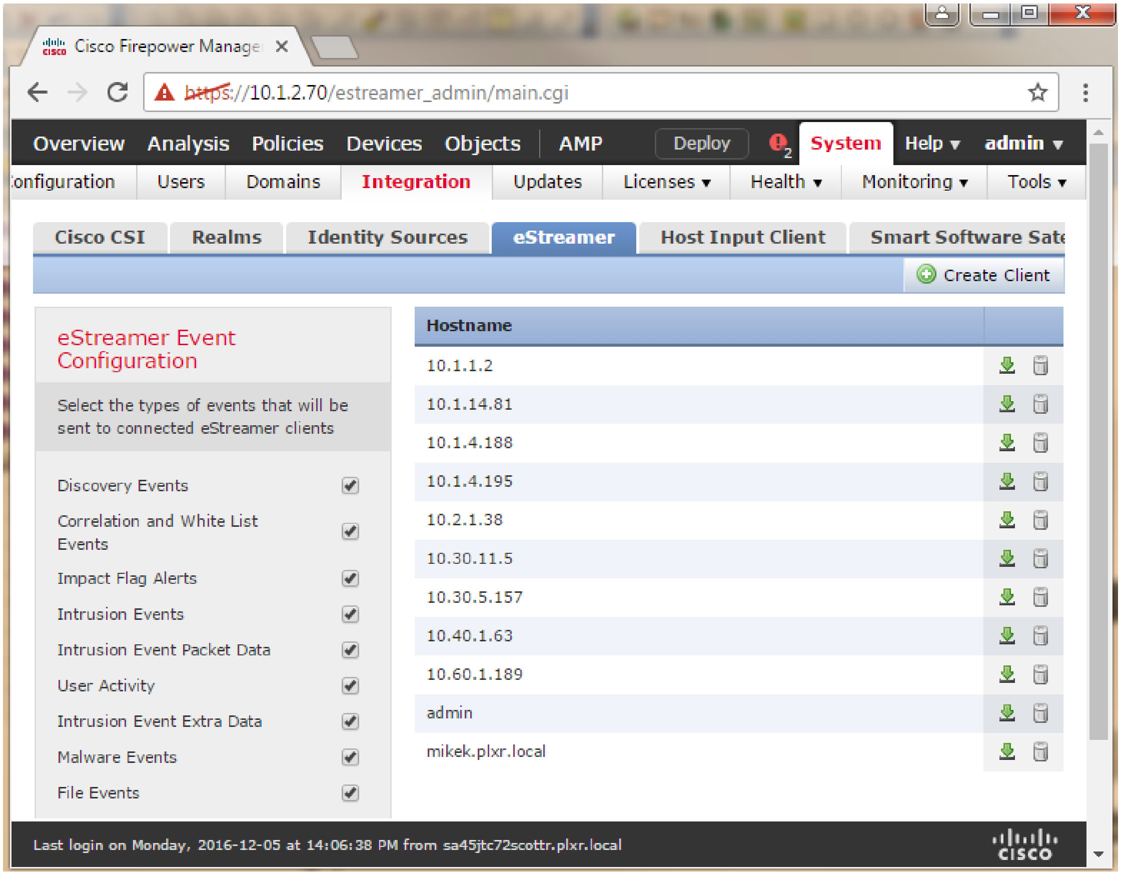
Cisco Firepower Manage
https://10.1.2.70/estreamer_admin/main.cgi
Overview Analysis Policies Devices Objects AMP
Deploy
System Help ▾ admin ▾
Configuration Users Domains Integration Updates Licenses ▾ Health ▾ Monitoring ▾ Tools ▾
Cisco CSI Realms Identity Sources eStreamer Host Input Client Smart Software Sate
Create Client

eStreamer Event Configuration

Select the types of events that will be sent to connected eStreamer clients

Discovery Events ☑

Correlation and White List Events ☑

Impact Flag Alerts ☑

Intrusion Events ☑

Intrusion Event Packet Data ☑

User Activity ☑

Intrusion Event Extra Data ☑

Malware Events ☑

File Events ☑

Hostname
10.1.1.2
10.1.14.81
10.1.4.188
10.1.4.195
10.2.1.38
10.30.11.5
10.30.5.157
10.40.1.63
10.60.1.189
admin
mikek.plxr.local

Last login on Monday, 2016-12-05 at 14:06:38 PM from sa45jtc72scottr.plxr.local

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

2. Enable all eStreamer Events, and click the Save button at the bottom of the list. Wait for the page to refresh. It may not give any other indication that a change has been made.

3. Click on the (+) Create Client button on the right.

4. Enter the Scrutinizer collector's IP address.

5. Enter a password (optional). If a password is entered, make sure to remember it. It will be needed in a later step.

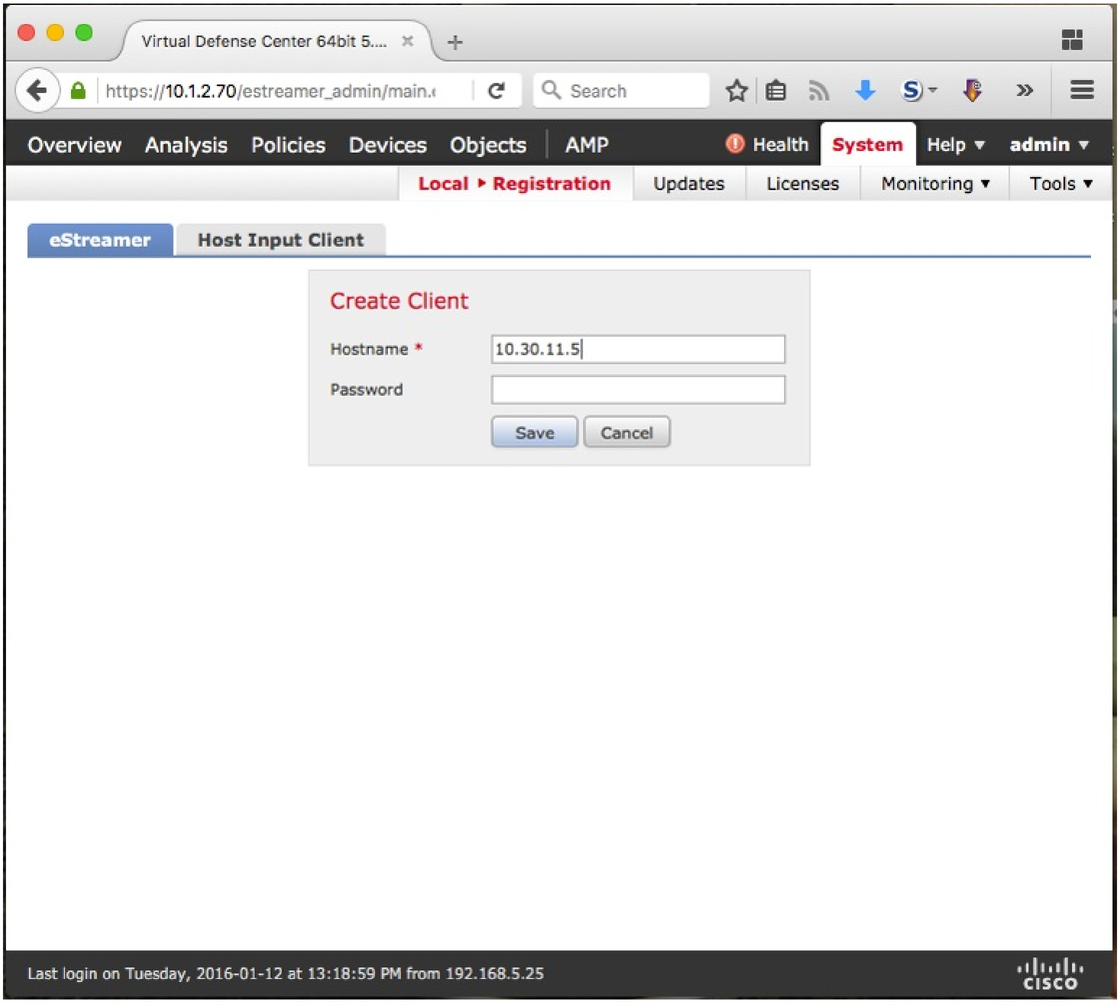
Virtual Defense Center 64bit 5....
https://10.1.2.70/estreamer_admin/main.
Search
Overview Analysis Policies Devices Objects AMP
Health System Help admin
Local ▸ Registration Updates Licenses Monitoring Tools
eStreamer Host Input Client

Create Client

Hostname * 10.30.11.5

Password

Save Cancel

Last login on Tuesday, 2016-01-12 at 13:18:59 PM from 192.168.5.25

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

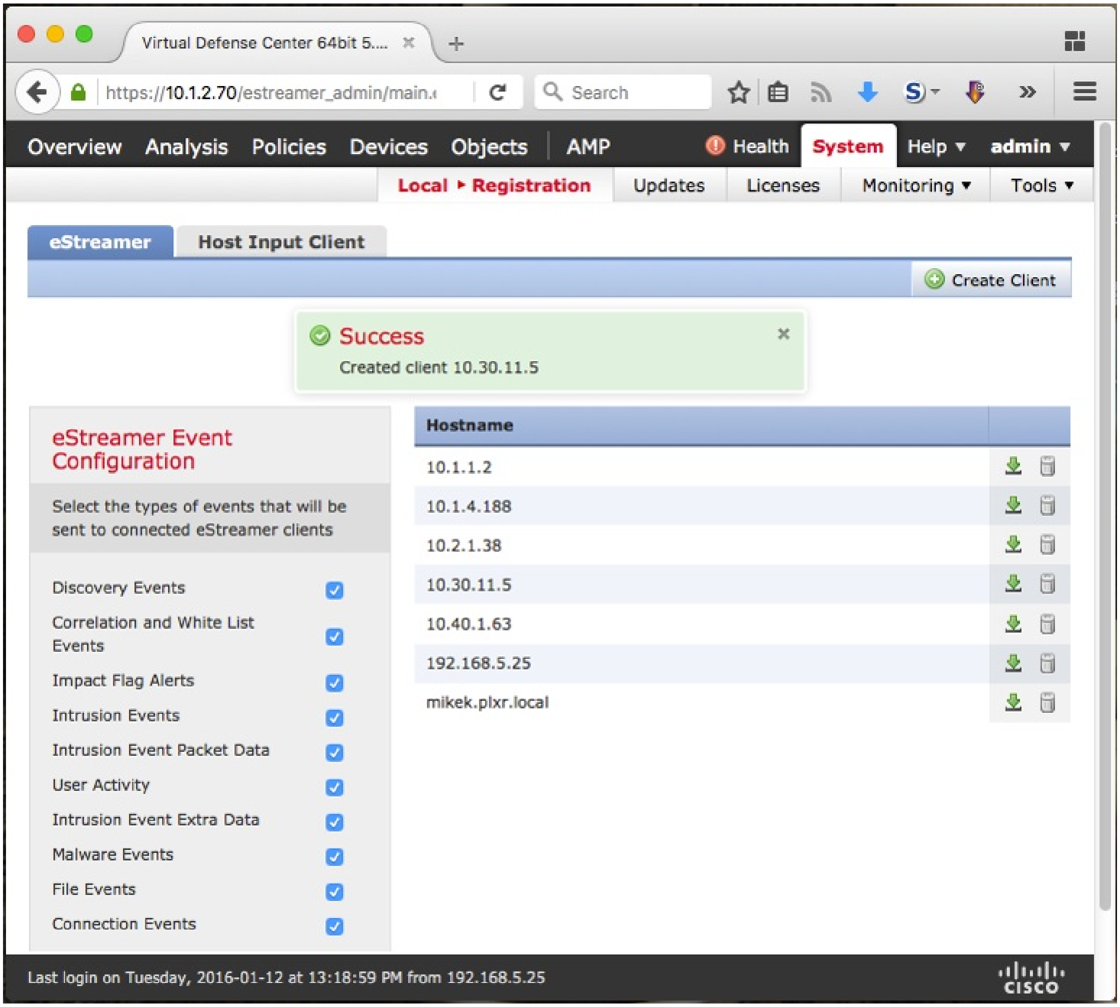

Success

Created client 10.30.11.5

eStreamer Event Configuration

Select the types of events that will be sent to connected eStreamer clients

- Discovery Events ☑
- Correlation and White List Events ☑
- Impact Flag Alerts ☑
- Intrusion Events ☑
- Intrusion Event Packet Data ☑
- User Activity ☑
- Intrusion Event Extra Data ☑
- Malware Events ☑
- File Events ☑
- Connection Events ☑

Hostname
10.1.1.2
10.1.4.188
10.2.1.38
10.30.11.5
10.40.1.63
192.168.5.25
mikek.plxr.local

Last login on Tuesday, 2016-01-12 at 13:18:59 PM from 192.168.5.25

7. License Plixer Scrutinizer's eStreamer client. Upload the client certificate to the */home/plixer/scrutinizer/files/* directory on the Plixer Scrutinizer appliance.

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

4.5.3 Configuring Plixer Scrutinizer's eStreamer client

1. SSH into the Plixer Scrutinizer collector server and configure the client.

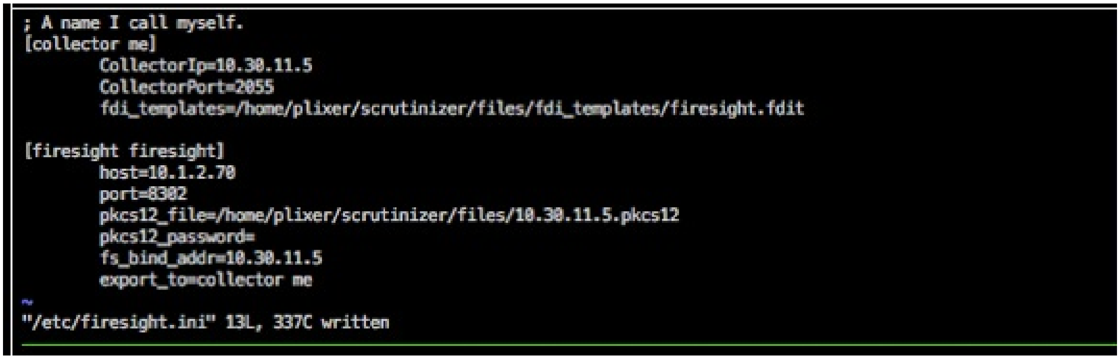

```
; A name I call myself.
[collector me]
        CollectorIp=10.30.11.5
        CollectorPort=2055
        fdi_templates=/home/plixer/scrutinizer/files/fdi_templates/firesight.fdit

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

2. Create or edit /etc/firesight.ini similar to the example above. Change the settings to reflect your network. There is the */home/plixer/scrutinizer/files/firesight.ini* file that you can edit and move to the /etc/ directory.

Note: Plixer Scrutinizer's eStreamer client will reconfigure itself every time a change is saved to *firesight.ini*.

3. The eStreamer client will export flows to the collector at CollectorIP and CollectorPort.

4. fdi_templates is the path where the export templates are defined. Use the location provided in the example.

5. The eStreamer client will connect to the FireSIGHT at the firesight host and port.

6. pkcs12_file is the location FireSIGHT certificate was updated.

7. pkcs12_password is the certificate password, or blank if a password wasn't specified.

8. fs_bind_addr is the eStreamer client address registered with FireSIGHT (Plixer Scrutinizer collector IP address). It must be a bindable address that can route to the eStreamer service.

9. export_to tells the eStreamer client which collector or collectors will receive exported flows.

Important: There can be more than one collector and/or firesight, but they must have different names. A single collector can receive flows from multiple firesights. A firesight exporter can send flows to multiple collectors.

10. Edit the /home/plixer/scrutinizer/env/local_env file. Change the following line:

```
exporter PLIXER_NO_FIRESEER=1
```

to

```
exporter PLIXER_NO_FIRESEER=0
```

Save the changes.

11. Restart the flow collector:

```
service plixer_flow_collector restart
```

12. Wait for flows which should be observed in Plixer Scrutinizer within a minute. Contact technical support for assistance with troubleshooting.

4.6 Grafana integration

This datasource provied the ability to visualize Scrutinizer's collected data within the Grafana interface.

4.6.1 Datasource setup

The datasource provides the ability to visualize the data collected by Scrutinizer within the Grafana interface.

- *Getting Started*
- *Adding Datasource to Grafana*
- *Setting up Scrutinizer Datasource*

Getting started

Currently we are working with Grafana to make this plugin available on their plugin repo. At the moment, you can install the plug by cloning it into the plugins directory. Depending on the OS you are running, the path will vary.

Path:

- Linux : '*/var/lib/grafana/plugins*'
- Windows : '*/data/plugins*' *

When inside the directory, run the following command:

'git clone https://github.com/plixer/scrutinizer-datasource.git'

Once the datasource is cloned, restart the Grafana server and proceed with adding the datasource to Grafana. You may need to create the **data/plugins** directoryon a Windows system.

Adding datasource to Grafana

1. In the Grafana user interface navigate to **Settings** and select **Data Sources** in the **Configuration** section.

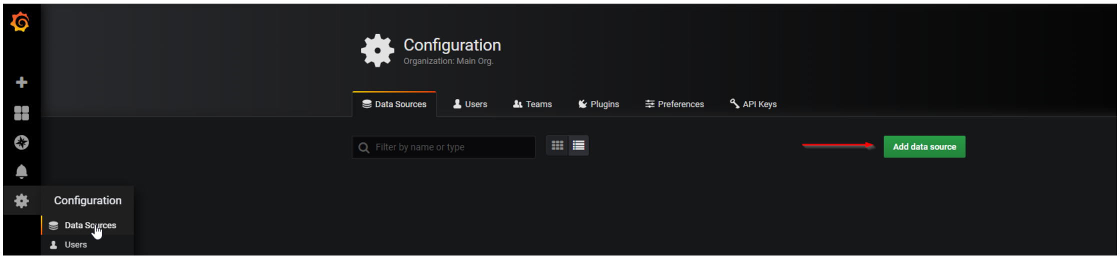

2. Select the **Plixer Scrutinizer Datasource**.

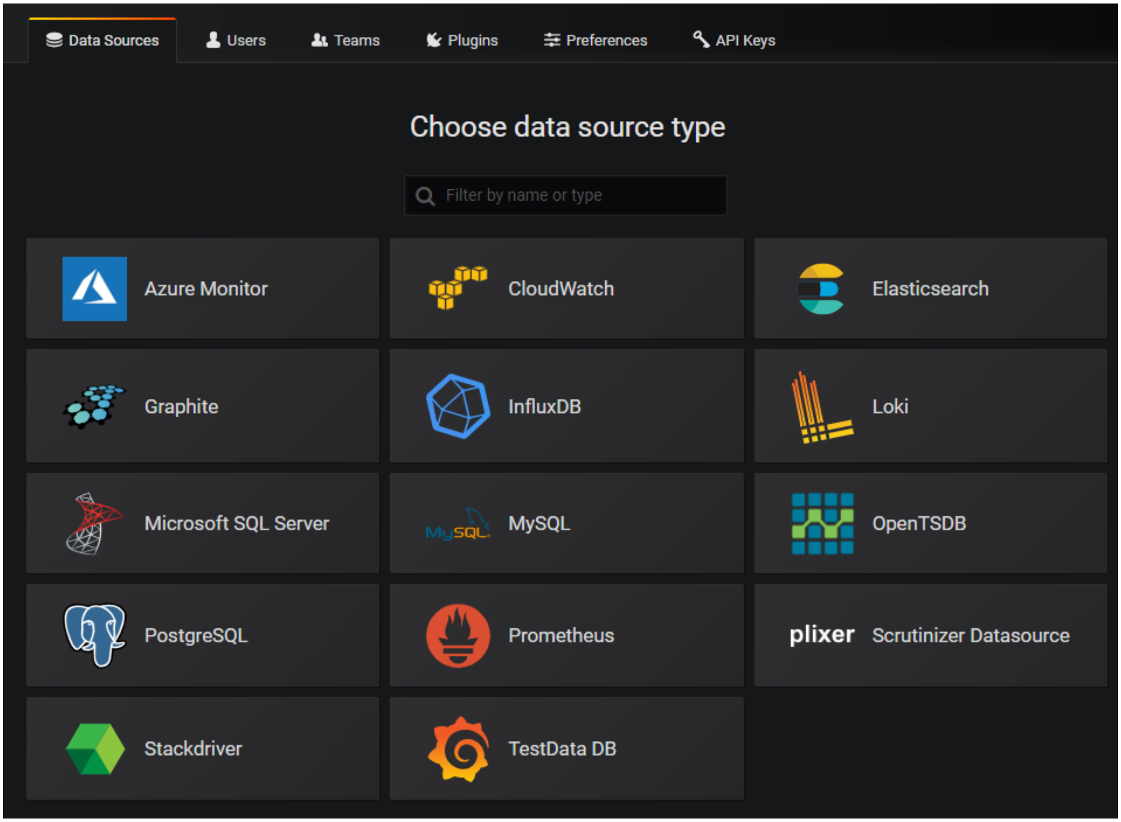

Setting up Scrutinizer datasource

1. Fill out the required fields.

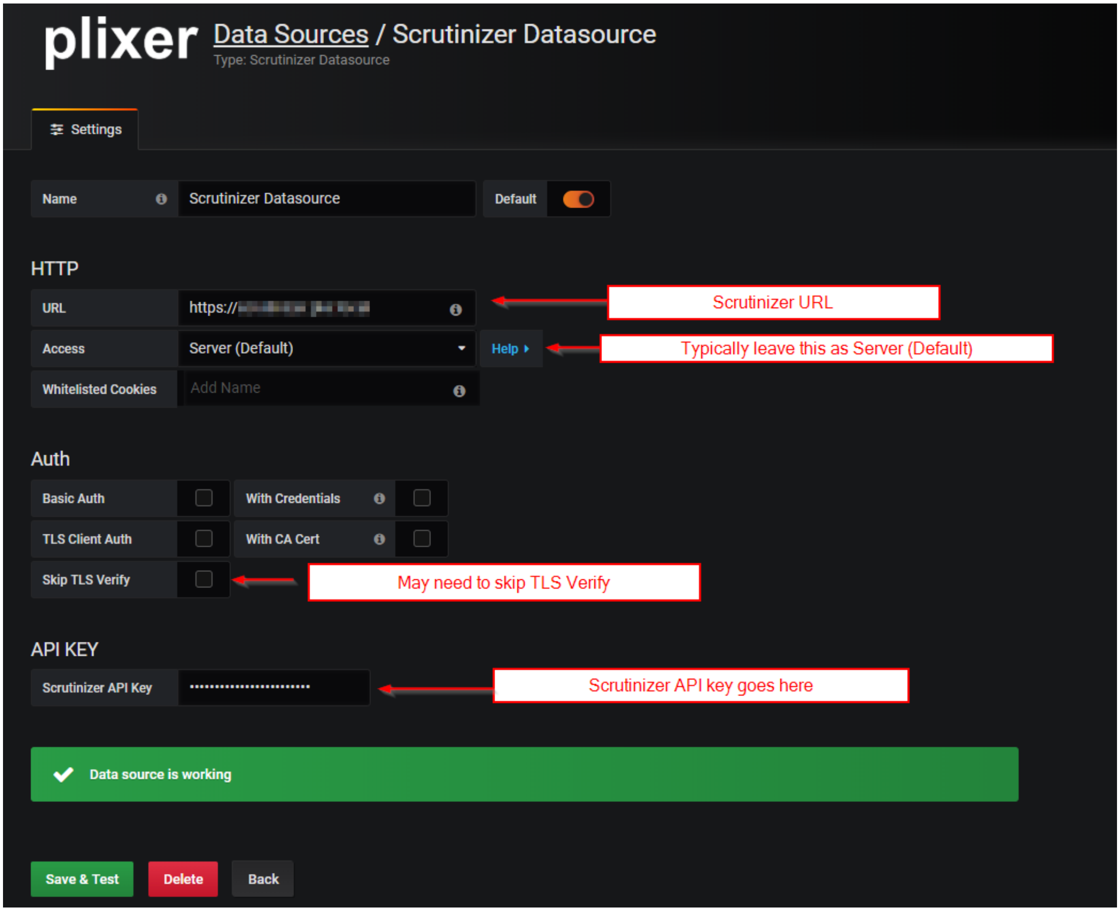

Data Sources / Scrutinizer Datasource

Type: Scrutinizer Datasource

Name: Scrutinizer Datasource

Default: ●

HTTP

URL: https://[redacted]

Scrutinizer URL

Access: Server (Default)

Typically leave this as Server (Default)

Whitelisted Cookies: Add Name

Auth

Basic Auth □ With Credentials □

TLS Client Auth □ With CA Cert □

Skip TLS Verify □

May need to skip TLS Verify

API KEY

Scrutinizer API Key: ••••••••••••••••••••••••

Scrutinizer API key goes here

✓ Data source is working

2. You can generate an authentication token in Scrutinizer via the **Admin Tab - > Security - > Authentication Token** page.

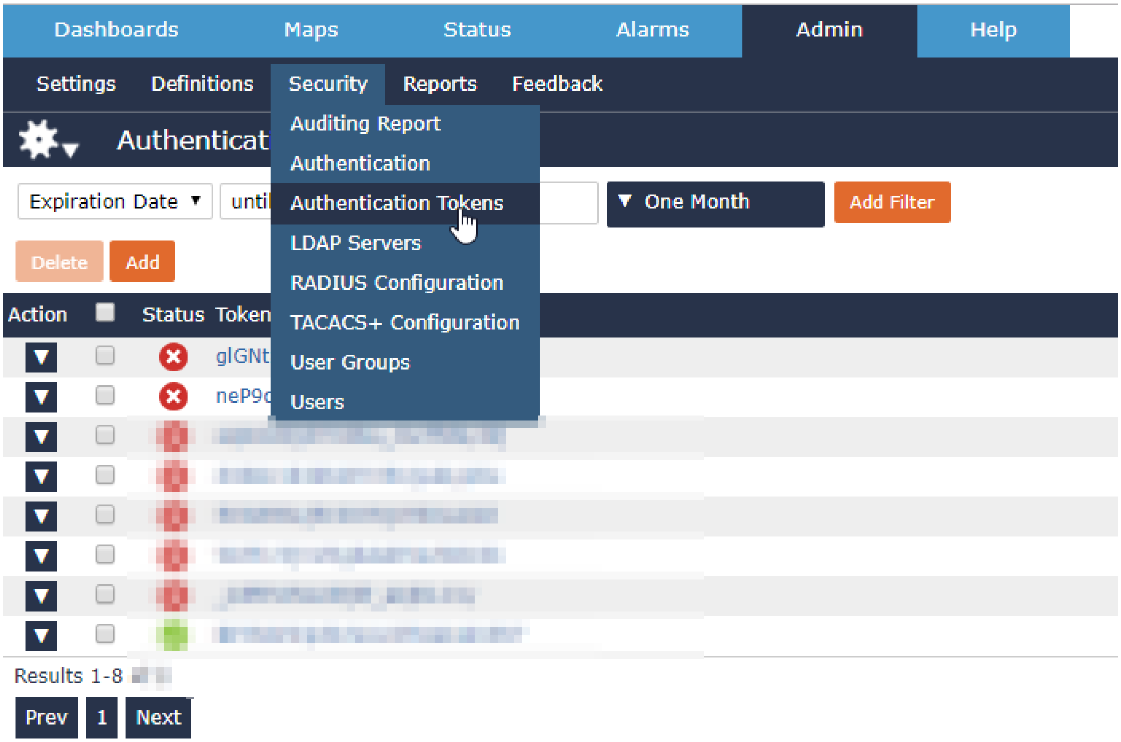

If you are getting an error, it can be caused by a self-signed SSL certificate.

If so, make sure you check off the ability to skip TLS verify.

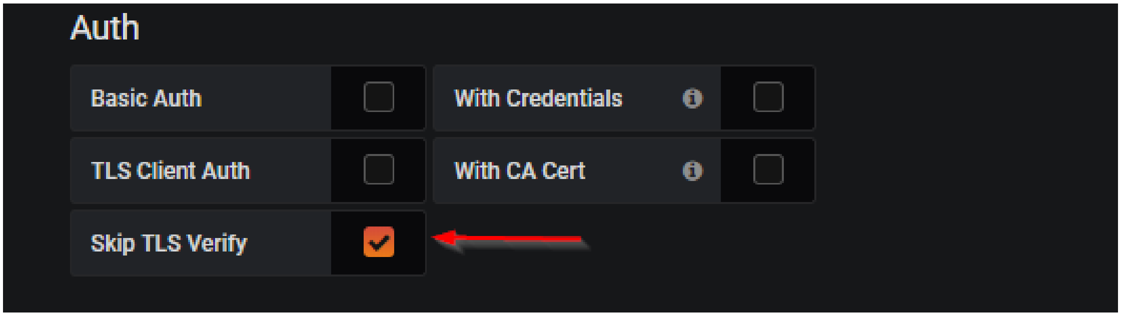

Auth

Basic Auth	□	With Credentials	□
TLS Client Auth	□	With CA Cert	□
Skip TLS Verify	☑		

4.6.2 Runnings reports

Reporting is done primarily by populating the drop-down menus.

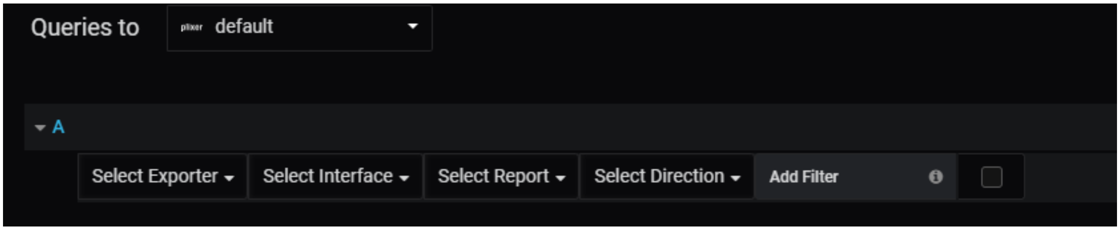

Queries to: default

The report will only rendor if you have selected something from each drop-down or you have selected to add a filter.

The data is going to be represented in bits / second. Make sure Grafana is displaying it this way by updating the visualization field:

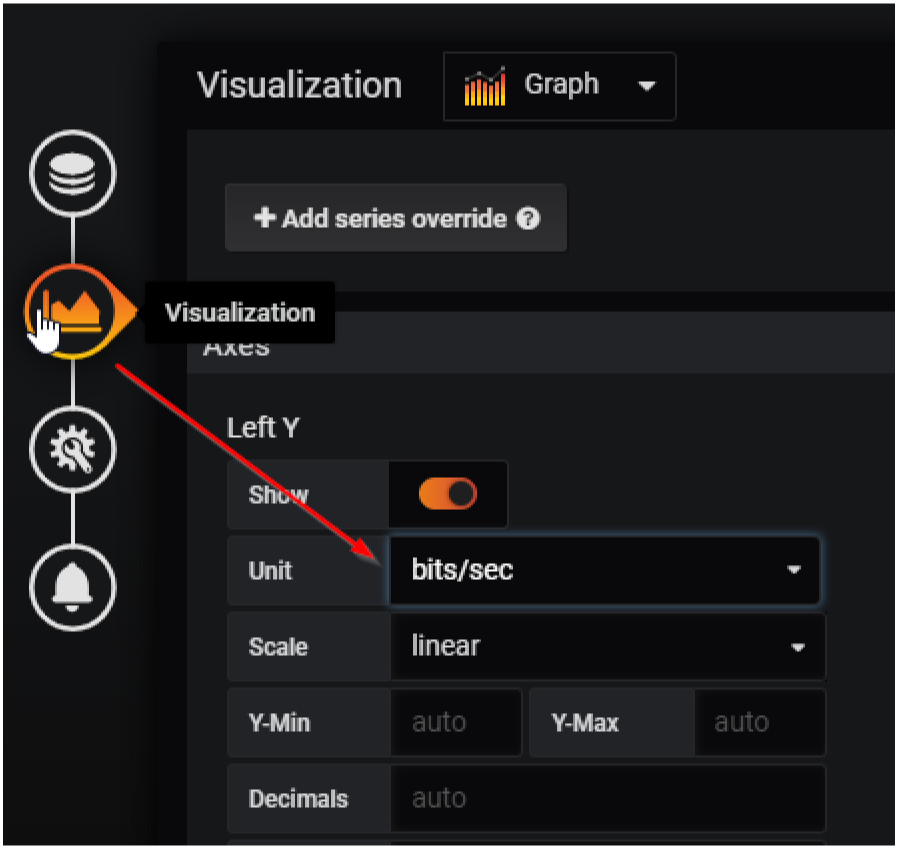

Visualization: Graph

Left Y

Show: ●

Unit: bits/sec

Scale: linear

Y-Min: auto

Y-Max: auto

Decimals: auto

Adding filters

Scrutinizer offers an extremly powerful filtering engine which can also be leveraged by the datasource. The easiest way to pass filters to the datasource is to first build them within Scrutinizer to get a feel for the format.

For example, to see what JSON would be needed to pass a filter for host 10.60.1.240 and application TCP 443, build the filters in Scrutinizer and then look at the *Report JSON (API)*.

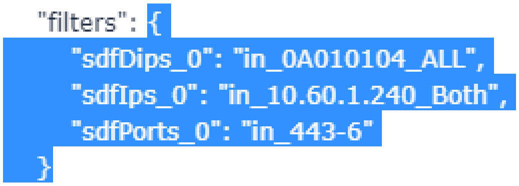

```
"filters": {
    "sdfDips_0": "in_0A010104_ALL",
    "sdfIps_0": "in_10.60.1.240_Both",
    "sdfPorts_0": "in_443-6"
}
```

Then copy from the opening bracket to the closing bracket and paste the filter within Grafana. Next select **Apply Filter**.

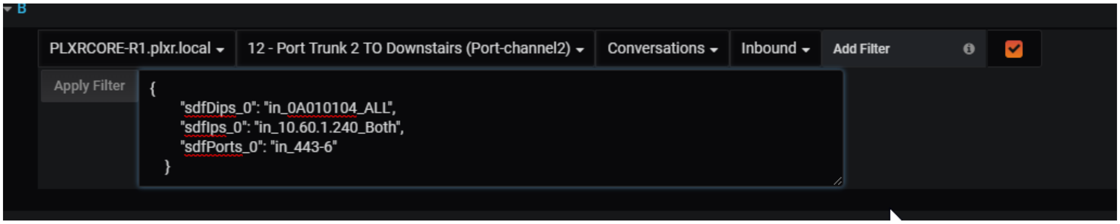

```
{
    "sdfDips_0": "in_0A010104_ALL",
    "sdfIps_0": "in_10.60.1.240_Both",
    "sdfPorts_0": "in_443-6"
  }
```

Keep in mind the “sdfDips_0” will be igored. This references the device in Scrutinizer you were looking at when you built the filters. You can leave it in if it’s easier to paste that way, or you can remove it. Both will yield same results.

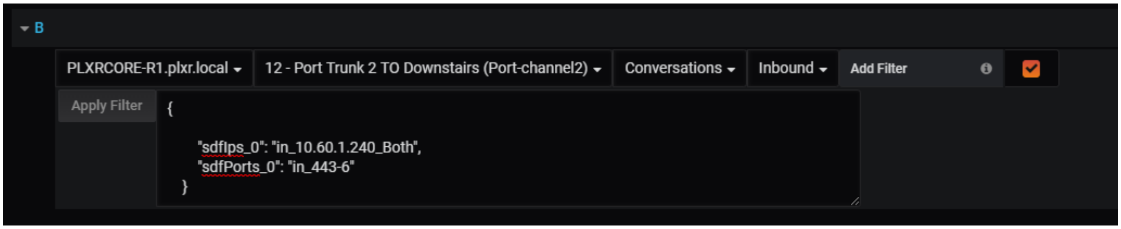

```
{

    "sdfIps_0": "in_10.60.1.240_Both",
    "sdfPorts_0": "in_443-6"
  }
```

Specifying a report

The **Select Report** drop-down menu comes with a few reccomended reports. If you are familiar with Scrutinizer you know that there are hundreds of different reports that are available to choose from.

Instead of populating all of those reports in the dropdown, you have the ability to put any report name you like in the box. To find out the report names, run the report in Scrutinizer and look at the report JSON. Once you have the report name, copy and paste it into the box and press Enter.

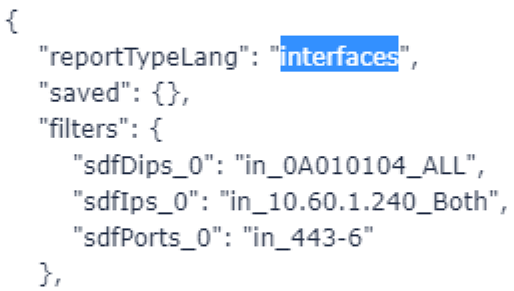

```
{
   "reportTypeLang": "interfaces",
   "saved": {},
   "filters": {
      "sdfDips_0": "in_0A010104_ALL",
      "sdfIps_0": "in_10.60.1.240_Both",
      "sdfPorts_0": "in_443-6"
   },
```

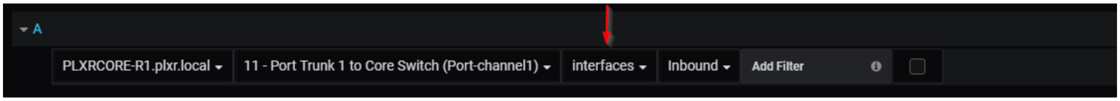

4.6.3 Report JSON (API)

Any report within Scrutinizer provides an option to view the JSON data that is passed from the front end to the back end in order to render the reports.

When in a report, click **Filters/Details** to open a module where you can select the Report JSON (API) tab.

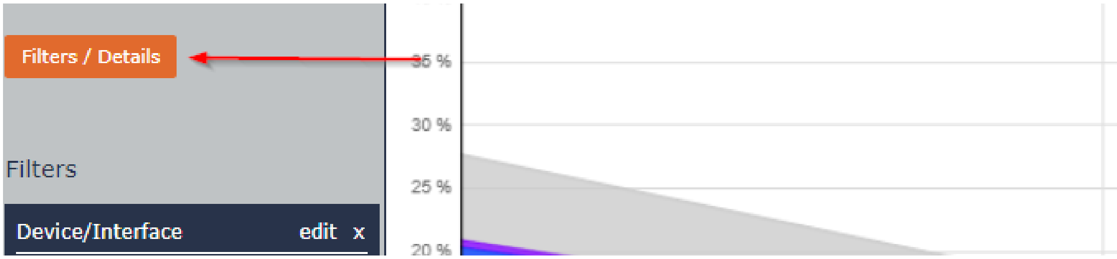

Experimenting with this tab will be very helpful when *Adding filters* or *Specifying a report* that are not available by default.

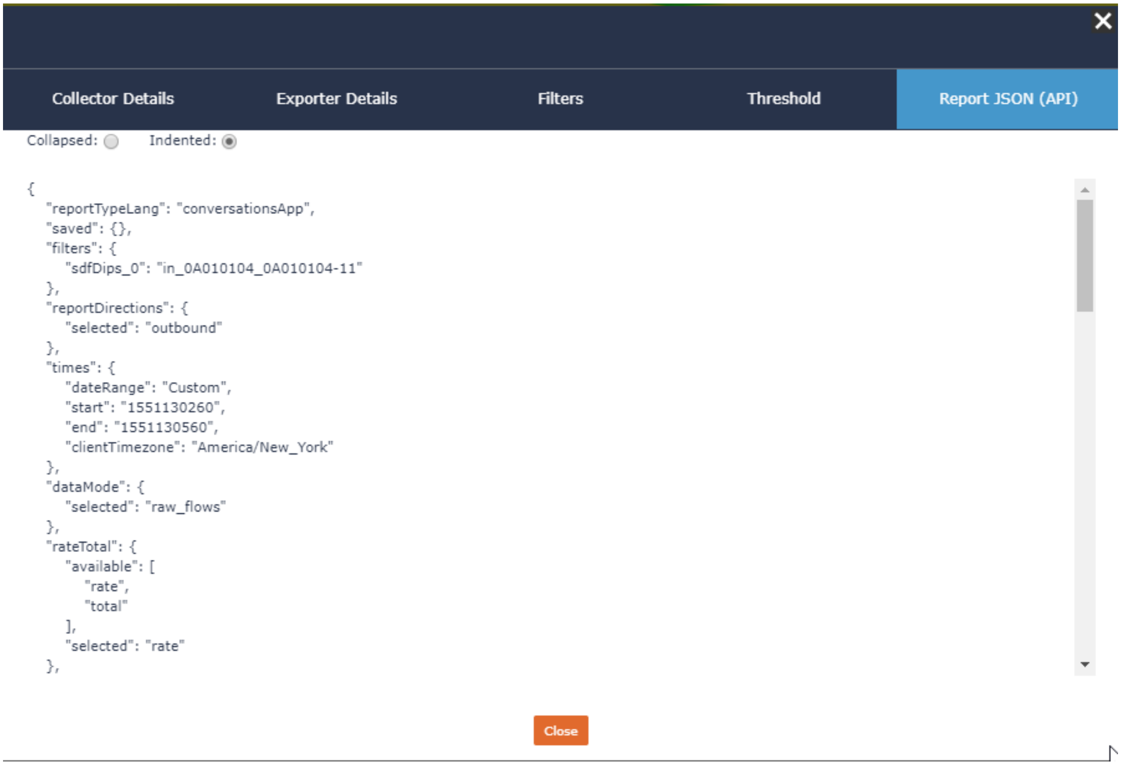

Collapsed: ○ Indented: ●

```
{
  "reportTypeLang": "conversationsApp",
  "saved": {},
  "filters": {
    "sdfDips_0": "in_0A010104_0A010104-11"
  },
  "reportDirections": {
    "selected": "outbound"
  },
  "times": {
    "dateRange": "Custom",
    "start": "1551130260",
    "end": "1551130560",
    "clientTimezone": "America/New_York"
  },
  "dataMode": {
    "selected": "raw_flows"
  },
  "rateTotal": {
    "available": [
      "rate",
      "total"
    ],
    "selected": "rate"
  },
```

4.7 Plixer Replicator load balancing

The load balancing feature provides integration between a Plixer Scrutinizer distributed cluster and a Plixer Replicator appliance. This integration will create a “seed” profile, as well as a profile for each collector in a Plixer Scrutinizer cluster Plixer Replicator. Any exporters added to that profile (via policy, UI, or API) will be automatically assigned to collector policies as long as resources are available.

Note: For additional details about Plixer Replicator, please contact Plixer support.

4.7.1 Setting up Plixer Replicator integration

1. Navigate to **Admin > Settings > Plixer Replicator**.
2. Mark the checkbox to enable the Plixer Replicator integration.
3. Enter the Plixer replicator admin user password in the password field, then re-enter it in the confirmation field that will appear.
4. Select the Plixer Replicator receive port that flows will come into the Plixer Replicator on.
5. Provide the IP or hostname of the Plixer Replicator deployment.

Important: The Replicator host URL must include http:// or https://.

6. Enter a name for the Plixer Replicator seed profile which will include the exporters that will be auto-replicated.
7. Select the Plixer Replicator port to send flows to Plixer Scrutinizer on.
8. Save the entry.
9. SSH into the Plixer Scrutinizer primary reporter appliance. Run the following command:

```
scrut_util --autoreplicate
```

10. After the command completes, log into the Plixer Replicator user interface and assign the exporters you'd like to load balance between collectors to the seed profile.
11. On the Plixer Scrutinizer appliance, re-run the *scrut_util –autoreplicate* command. The output will show the exporters being assigned to the Plixer Scrutinizer collectors.

Important: Please contact Plixer support for assistance with the integration.

4.8 ServiceNow bi-directional integration

4.8.1 Overview

Plixer Scrutinizer introduces bi-directional integration with ServiceNow, enabling NetOps to streamline the process of creating troubleshooting tickets. In addition, NetOps can share the “collected” network- and end-device-related data that is associated with any incident. This provides context into why the ticket was opened and eliminates the need to duplicate investigative effort. NetOps can now track security-related tickets, allowing the team to demonstrate their value in keeping the business safe.

Important: Additional licensing is required for this feature. Contact Plixer support for assistance.

4.8.2 Configuring ServiceNow integration

1. Navigate to the Admin > Settings > ServiceNow page.

2. Click “Add” to create a new ServiceNow integration profile.

3. Provide a unique name for the ServiceNow instance.

4. Enter the instance URL and login credentials.

5. Save the entry.

Note: Once configured, the ServiceNow instance name will appear as available under **Collections** and **Notifications.**

4.8.3 Collections

When a collection is associated with ServiceNow, an incident will be created in ServiceNow with details linking back to the Plixer Scrutinizer collection. Once an incident has been created for a collection, details about the incident status can be viewed from within Plixer Scrutinizer.

4.8.4 Notifications

ServiceNow integration instances can be associated with alarm policies so that when a policy is violated a ServiceNow incident will be created.

4.9 Scrutinizer for Splunk application

4.9.1 Overview

What does the Scrutinizer for Splunk application do?

- From within Scrutinizer, searches for IP addresses can be launched in Splunk.
- From Splunk, *Scrutinizer Vitals* and *FA TopN gadget details* can be viewed.

4.9.2 Splunk searches from Scrutinizer reports and alarms

The following are the steps necessary to search Splunk's database from Scrutinizer reports.

1. Select a Scrutinizer report that includes IP addresses
   a. Select a host of interest and click on it
   b. Select **Other Options** from the Reports menu
   c. Select **Search Splunk**
2. The IP address and report timeframe are passed to Splunk's search engine for detailed Splunk reporting.

For further details (from Splunk) on Scrutinizer Alarms, follow these steps.

1. Go to the *Alarms tab* in Scrutinizer and select either:
   a. Bulletin Board by Violator,
   b. Bulletin Board by Policy.
2. In the Bulletin Board Events view that opens, click on the dropdown arrow to the left of the Message column for the alarm selected.
3. Select **Search Splunk** from the Available Options menu and the Violator's IP address and timeframe of the violation are passed to Splunk. If the Violator's IP address and an alarm time (not timeframe) are

   being passed, then 30 minutes before and after the alarm time is searched.

4.9.3 Scrutinizer reporting from within Splunk

With the Scrutinizer for Splunk application, dashboards can be setup which include:

- Scrutinizer Vitals information
- Flow Analytics TopN Algorithms

The Default Dashboard view for the **Scrutinizer - Splunk Application** is the Vitals information which includes:

- CPU
- Memory
- Disk Usage
- Flows per collector
- Status per collector

Other Splunk menus can include links to the Scrutinizer tabs:

- Dashboards
- Maps
- Status and Reports
- Alarms
- Admin and Settings

Additionally, panels can be created which are based on the TopN Algorithm gadgets from Flow Analytics:

- Top Applications
- Top Countries
- Top Rev 2nd lvl Domains (Top reverse 2nd level domains)
- Top Flows
- Top Hosts

- Top Jitter
- Top Networks
- and more…

Note: Clicking on any entities in the graphs or detail in a table report will run a Splunk search for that detail and time range. As mentioned earlier, those searches can also be initiated directly from *Scrutinizer reports or alarms*.

A menu can be added which consists of useful links on our website:

- plixer.com
- Configure NetFlow
- Internet Threat Center
- Latest Blogs
- Plixer Support

4.9.4 How to configure Splunk integration with Scrutinizer

There are two components to the Splunk integration.

- Preparing Scrutinizer
- Installing the Scrutinizer for Splunk application on Splunk

Preparing Scrutinizer

Note: Flow Analytics must be enabled and collecting statistics for the Top X Algorithms.

The following configuration instructions apply to Scrutinizer v16.7 and later. If an earlier version of Scrutinizer is installed, contact Plixer directly for assistance.

1. Log on to the Scrutinizer server with administrative permissions
2. To *enable the Scrutinizer/Splunk integration*, from the command line:

```
/home/plixer/scrutinizer/bin/scrut_util
**SCRUTINIZER>** enable splunk http\://<ip:port> <syslog port>
```

where<ip:port> is Splunk IP address and port, <syslog port> is the port used to send syslogs to Splunk.

After a few moments, Scrutinizer will begin to export data to Splunk.

- To manually *collect data from Scrutinizer* and send to Splunk, run:

  SCRUTINIZER> collect splunk <splunk_ip> <port>

- To *disable the data export*, run:

  SCRUTINIZER> disable splunk http://<ip:port>

Installing the Scrutinizer for Splunk application on Splunk

The Scrutinizer for Splunk App displays details that have been collected and processed by Scrutinizer. For more information, visit Plixer's Splunk integration page.

1. Download the Splunk plugin using the link at the bottom of the Splunk Integration page.
2. Log into Splunk
3. Select **Apps > Manage Apps**
4. Click **Install app from file**
5. Click **Choose File** and locate the PlixerScrutinizerForSplunk.spl file within the Splunk plugin file downloaded in Step 1. If upgrading the plugin, click the **Upgrade App** checkbox below the **Choose File** button.
6. Click the **Upload** button.
7. Follow the onscreen instructions and restart Splunk.
8. Navigate to the "**Apps > Manage Apps** menu.
9. Locate the **Scrutinizer for Splunk** app in the list below and click the **View Objects** link associated to the application.
10. Locate the **Default** link and click it.
11. Replace the text "http://ADD_SCRUTINIZER_ADDR_HERE" with the link to the desired Scrutinizer server. (e.g. https://10.1.1.12:88)
12. Click the **Save** button and then access the Scrutinizer application in the Apps menu.

After a few minutes you will start seeing data in Splunk.

4.10 Third-party integrations

Third-party integration from Scrutinizer can be enabled by navigating to **Admin > Definitions > 3rd Party Integration**. Select an integration in the dropdown menu and enter the integration URL of the desired integration server in the URL box. Uncheck the disable checkbox, otherwise the integration will not be visible in the **Device Explorer**.

Some integrations, such as PRTG and Solarwinds, require object ID translations. Please review the setup instructions if these types of integrations are going to be enabled.

4.10.1 PRTG integration

PRTG integration can be enabled from the **Admin > Definitions > 3rd Party Integration** menu. Additional fields are required to connect to the PRTG API.

1) Select **PRTG** in the dropdown menu.

2) Uncheck the 'disabled' checkbox.

3) Fill out the additional fields that appear:

Field	Description
Protocol	The default is https, but http can be entered if PRTG is set up to use it
IP	The IP address of the PRTG server
Port	The default is 443, but another port can be entered if PRTG is configured to use it
User	The login user to connect to PRTG API
Password	The login users password to connect to PRTG API

4) Click the **Save** button.

Note: Defaulted fields assume that the PRTG instance is running on HTTPS. If the local install is running on a different port or using a different protocol the settings can be adjusted to match what is configured on the PRTG server (PRTG Administration Tool > Web Server Tab)

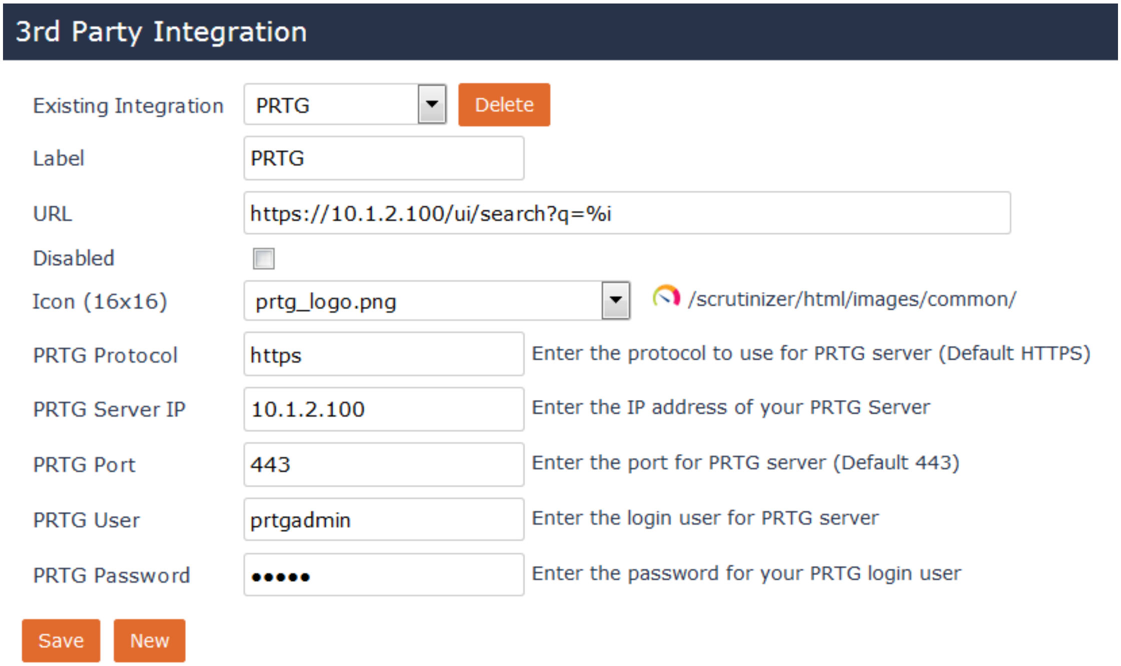

3rd Party Integration

Existing Integration	PRTG ▾ Delete	
Label	PRTG	
URL	https://10.1.2.100/ui/search?q=%i	
Disabled	☐	
Icon (16x16)	prtg_logo.png ▾	/scrutinizer/html/images/common/
PRTG Protocol	https	Enter the protocol to use for PRTG server (Default HTTPS)
PRTG Server IP	10.1.2.100	Enter the IP address of your PRTG Server
PRTG Port	443	Enter the port for PRTG server (Default 443)
PRTG User	prtgadmin	Enter the login user for PRTG server
PRTG Password	●●●●●	Enter the password for your PRTG login user

Now that the PRTG third-party integration is enabled, its icon will be displayed in the Device Tree for each exporter on the Status Tab and the Maps tab. By clicking on the PRTG icon, a browser is launched. PRTG will then display the device statistics for the exporter selected in Scrutinizer.

4.10.2 Solarwinds integration

The Solarwinds integration can be enabled from the ** Admin > Definitions > 3rd Party Integration** menu. Additional fields are required to connect to the Solarwinds API.

1) Select **Solarwinds** in the dropdown menu.

2) Uncheck the 'disabled' checkbox.

3) Fill out the additional fields that appear:

Field	Description
IP	The IP address of the Solarwinds server
User	The login user to connect to Solarwinds API
Password	The login users password to connect to Solarwinds API
Port	The default is 17778, but another port can be entered if Solarwinds is configured to use it

4) Click the **Save** button.

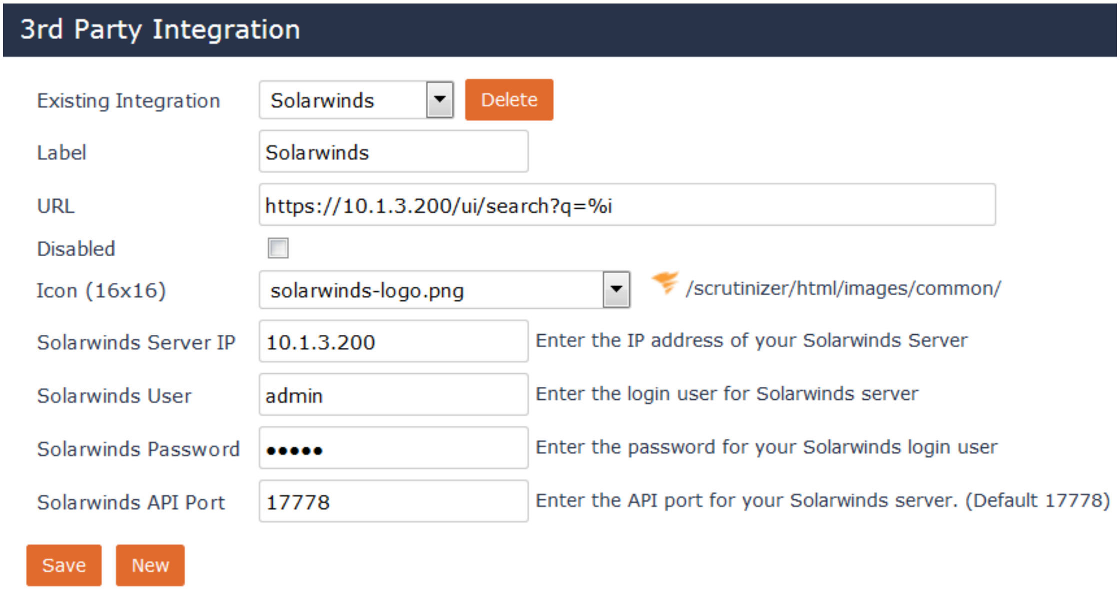

3rd Party Integration

Existing Integration	Solarwinds	
Label	Solarwinds	
URL	https://10.1.3.200/ui/search?q=%i	
Disabled	□	
Icon (16x16)	solarwinds-logo.png	/scrutinizer/html/images/common/
Solarwinds Server IP	10.1.3.200	Enter the IP address of your Solarwinds Server
Solarwinds User	admin	Enter the login user for Solarwinds server
Solarwinds Password	•••••	Enter the password for your Solarwinds login user
Solarwinds API Port	17778	Enter the API port for your Solarwinds server. (Default 17778)

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.

Now that the Solarwinds third-party integration is enabled, its icon will be displayed in the Device Tree for each exporter on the Status and Maps tab. Clicking on the Solarwinds icon will launch a browser. Solarwinds will then display the device statistics for the exporter selected in Scrutinizer.

Solarwinds Integration into Scrutinizer

You can also pivot from the **Node Details** page in Solarwinds NPM into a Scrutinizer report by following these steps:

Note: This integration was written with Solarwinds NPM 12.2. It is not guaranteed to work on older installations.

1) Navigate to **Settings > All Settings**. Under the section **Node & Group Management**, choose **Manage Custom Properties**.

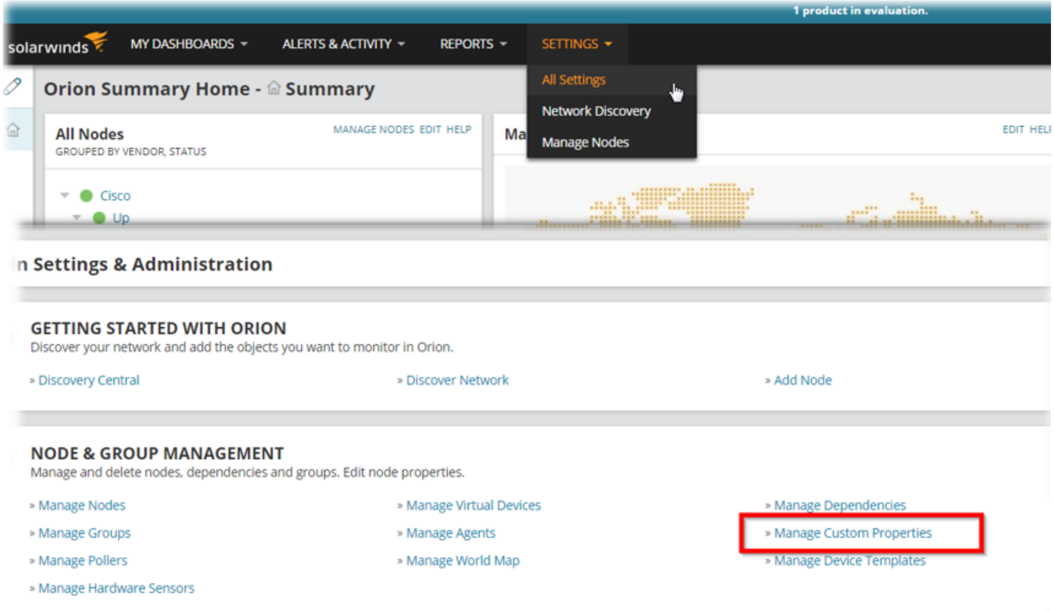

n Settings & Administration

GETTING STARTED WITH ORION

Discover your network and add the objects you want to monitor in Orion.

» Discovery Central » Discover Network » Add Node

NODE & GROUP MANAGEMENT

Manage and delete nodes, dependencies and groups. Edit node properties.

» Manage Nodes » Manage Virtual Devices » Manage Dependencies

» Manage Groups » Manage Agents » Manage Custom Properties

» Manage Pollers » Manage World Map » Manage Device Templates

» Manage Hardware Sensors

2) Click **Add Custom Property**. Select **Nodes** from the dropdown list:

- Fill out the name and description fields for the property (eg. Scrutinizer).
- Click Next.

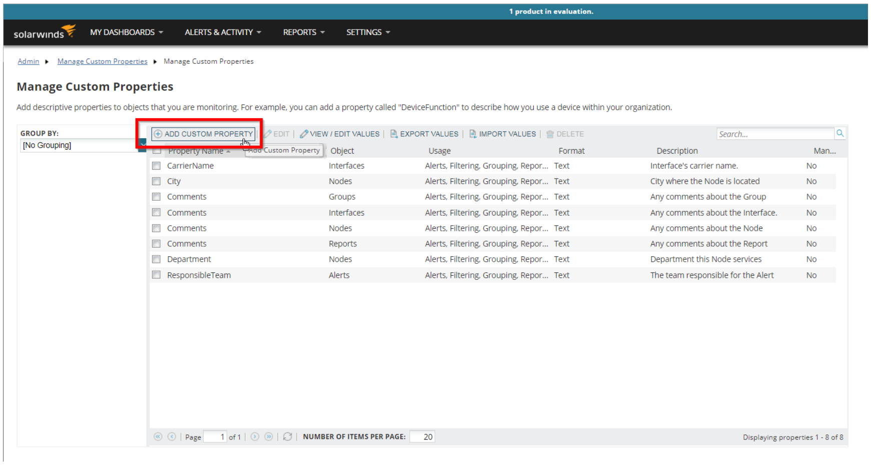

Admin ▸ Manage Custom Properties ▸ Manage Custom Properties

Manage Custom Properties

Add descriptive properties to objects that you are monitoring. For example, you can add a property called "DeviceFunction" to describe how you use a device within your organization.

GROUP BY: [No Grouping]

	Property Name	Object	Usage	Format	Description	Man...
□	CarrierName	Interfaces	Alerts, Filtering, Grouping, Repor...	Text	Interface's carrier name.	No
□	City	Nodes	Alerts, Filtering, Grouping, Repor...	Text	City where the Node is located	No
□	Comments	Groups	Alerts, Filtering, Grouping, Repor...	Text	Any comments about the Group	No
□	Comments	Interfaces	Alerts, Filtering, Grouping, Repor...	Text	Any comments about the Interface.	No
□	Comments	Nodes	Alerts, Filtering, Grouping, Repor...	Text	Any comments about the Node	No
□	Comments	Reports	Alerts, Filtering, Grouping, Repor...	Text	Any comments about the Report	No
□	Department	Nodes	Alerts, Filtering, Grouping, Repor...	Text	Department this Node services	No
□	ResponsibleTeam	Alerts	Alerts, Filtering, Grouping, Repor...	Text	The team responsible for the Alert	No

Page 1 of 1 | NUMBER OF ITEMS PER PAGE: 20 Displaying properties 1 - 8 of 8

3) Now to assign this property to your existing nodes

- Click the button labeled **Select Nodes** to specify your exporters. You may select all or just a few.
- Use the **Add** arrow to move your selected exporters.
- Choose **Select Nodes** when you are finished.

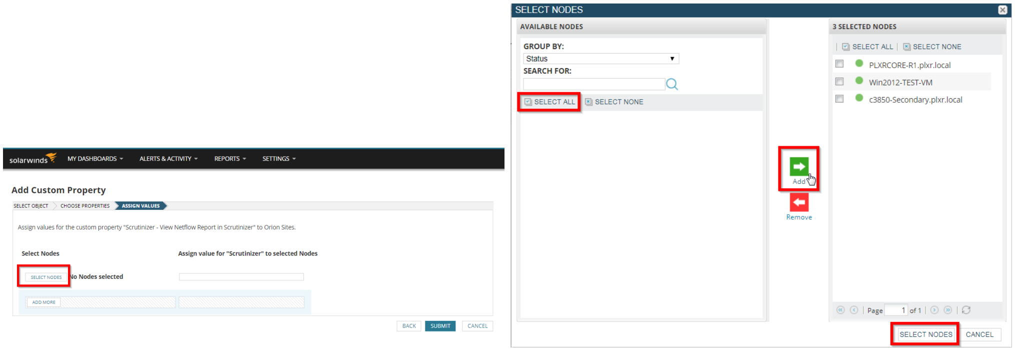

Add Custom Property

Assign values for the custom property "Scrutinizer - View Netflow Report in Scrutinizer" to Orion Sites.

Select Nodes

No Nodes selected

Assign value for "Scrutinizer" to selected Nodes

GROUP BY: Status

PLXRCORE-R1.plxr.local

Win2012-TEST-VM

c3850-Secondary.plxr.local

4) Now to add the value for all these nodes. This is where we specify the Scrutinizer info.

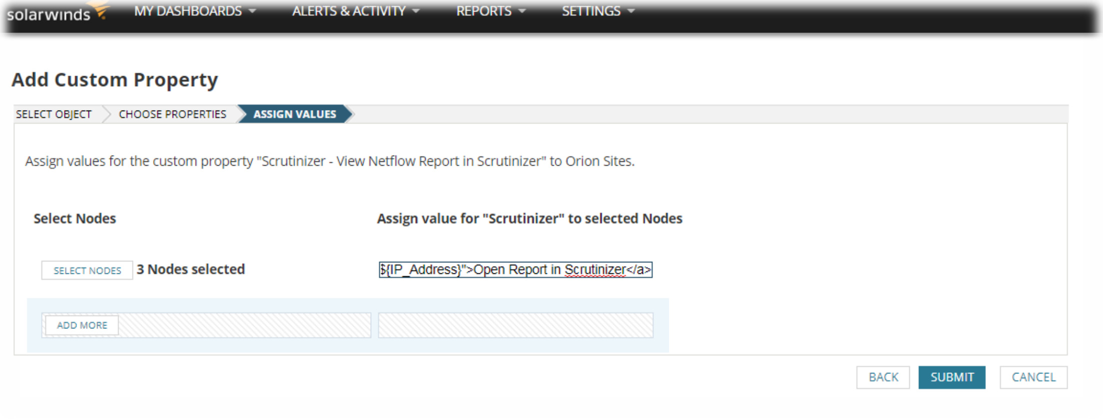

Add Custom Property

Assign values for the custom property "Scrutinizer - View Netflow Report in Scrutinizer" to Orion Sites.

Select Nodes

Assign value for "Scrutinizer" to selected Nodes

3 Nodes selected

${IP_Address}">Open Report in Scrutinizer</a>

- Fill in the value box with the following code:

```
<a href="http://SCRUTINIZER_IP_ADDRESS/search.html?el=${IP_Address}&
↪reportType=conversations">
        <img src="https://cdn.plixer.com/wp-content/uploads/2016/09/
↪scrutinizer_logo-300x49.png" height="49" width="300"></img>
</a>
```

Tip: You can specify the default report type that opens in Scrutinizer by editing the code block in step 4. To specify another report type by the API report name for example, "Conversations WKP" is known as `conversations` you just need to replace the `reportType` field in the url.

- Specify the IP address of your server in the above URL where it says **SCRUTINIZER_IP_ADDRESS**.
- Click **Submit**.

5) You will now have a custom properties widget on your node details page for all your selected hosts. Clicking the Scrutinizer logo will open a report in Scrutinizer.

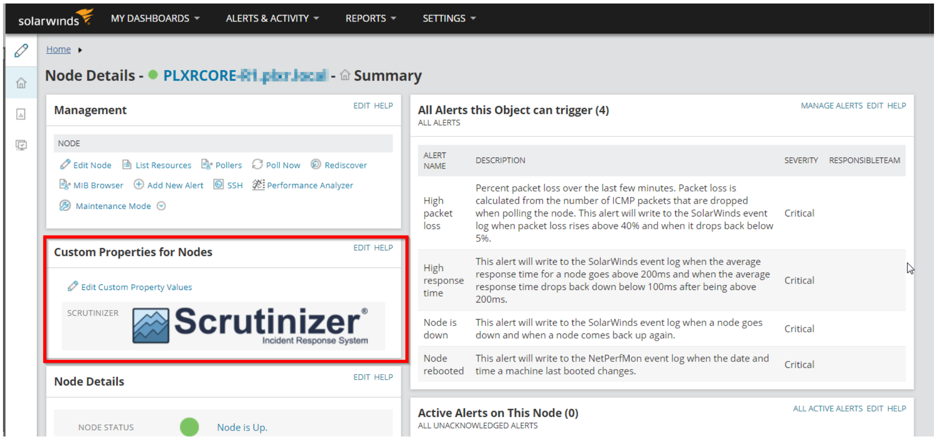

Node Details - PLXRCORE-[redacted] - Summary

Management

Custom Properties for Nodes

SCRUTINIZER

Node Details

NODE STATUS Node is Up.

All Alerts this Object can trigger (4)

ALL ALERTS

ALERT NAME	DESCRIPTION	SEVERITY	RESPONSIBLETEAM
High packet loss	Percent packet loss over the last few minutes. Packet loss is calculated from the number of ICMP packets that are dropped when polling the node. This alert will write to the SolarWinds event log when packet loss rises above 40% and when it drops back below 5%.	Critical	
High response time	This alert will write to the SolarWinds event log when the average response time for a node goes above 200ms and when the average response time drops back down below 100ms after being above 200ms.	Critical	
Node is down	This alert will write to the SolarWinds event log when a node goes down and when a node comes back up again.	Critical	
Node rebooted	This alert will write to the NetPerfMon event log when the date and time a machine last booted changes.	Critical	

Active Alerts on This Node (0)

ALL UNACKNOWLEDGED ALERTS

4.11 Viptela SD-WAN

4.11.1 Overview

The integration between Viptela and Scrutinizer provides insight into SD-WAN for Viptela such as:

- Performance Reports
- vEdge Health Metrics
- Policy Events

After configuring Viptela integration with Scrutinizer, the following Viptela specific reports are available in Scrutinizer:

- Carrier Performance
- Transport Performance
- Tunnel Performance
- Application Performance
- Status All Components
- vEdge Health
- SLA Events
- Policies Added
- Policies Removed

Some reports are found under the vManage exporter, while others are placed under the vEdge devices.

4.11.2 Prerequisites

A minimum of Scrutinizer v18.16 is required for the configuration to successfully complete and for the Viptela reports to become available.

Note: The user configured in Scrutinizer to connect to Viptela API must have full read access.

4.11.3 Viptela configuration in Scrutinizer

1. Navigate to Admin > Settings > Viptela Settings
2. Mark the checkbox to enable the Viptela integration
3. Provide the IP or hostname of the Viptela vManage NMS
4. Enter the user password in the password field, then re-enter it in the confirmation field that will appear
5. Select the Viptela port that will communicate with Scrutinizer, default is 8443
6. Select the protocol to communicate over. Scrutinizer will use HTTPS by default
7. Save the entry

4.11.4 Frequently asked questions

Q: It isn't working, how can I see what is going on?
A: The Viptela collection process runs under the umbrella of the plixer_flow_collector daemon.

- Check the collector log for errors.
- Verify the credentials you entered in Scrutinizer are correct.
- Use the Test button to confirm that Scrutinizer user can access Viptela SD-WAN API.

Q: What permissions does my Viptela user need?
A: The user configured in Scrutinizer to connect to Viptela API must have full read access.

4.12 STIX-TAXII feeds

4.12.1 Overview

The Structured Threat Information eXchange (STIX) is an industry-standard file format for the exchange of threat information between organizations and platforms. The Trusted Automated eXchange of Indicator Information (TAXII)is a protocol that allows the transmission of threat information, primarily in STIX format, between systems and organizations. Importing threat intelligence information, such as IP indicators, in STIX format, via the TAXII protocol from a remote source, enhances Plixer Scrutinizer's existing IP detection capabilities.

Important: Additional licensing is required for this feature. Contact Plixer support for assistance.

4.12.2 Setting up STIX imports via CLI

To configure the STIX import, collect IP or domain watchlists in the STIX format (v1 or v2).The name of the file will become the category. Place the files into the */home/plixer/scrutinizer/files/threats* directory so that the application will automatically import them.

Important: Plixer Scrutinizer supports .stix, .stix1 or .stixv1 as the extension for v1 (XML) or .stix2 or .stxv2 for v2 (JSON).

4.12.3 Configuring STIX-TAXII feeds

1. Navigate to the **Admin>Settings>STIX-TAXII** page and click the Add button to create a new feed.
2. Fill out the following fields:

- the Feed Name;
- the API Root (NOT the Discovery URL);
- the Collection ID;
- the username and password.

3. Save the entry.
4. Use the Test button to confirm the Scrutinizer user can access the feed.
5. After you complete the setup, Plixer Scrutinizer will attempt to pull the lists from the TAXII server every time the process of downloading hostreputation lists runs. Alerting should happen automatically.

4.12.4 Helpful 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.

4.13 User name reporting - Active Directory integration

4.13.1 Overview

The Active Directory integration provides lists of user names along with domain, datasource, first seen and last seen details. It also allows to search across all flows for user names.

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

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

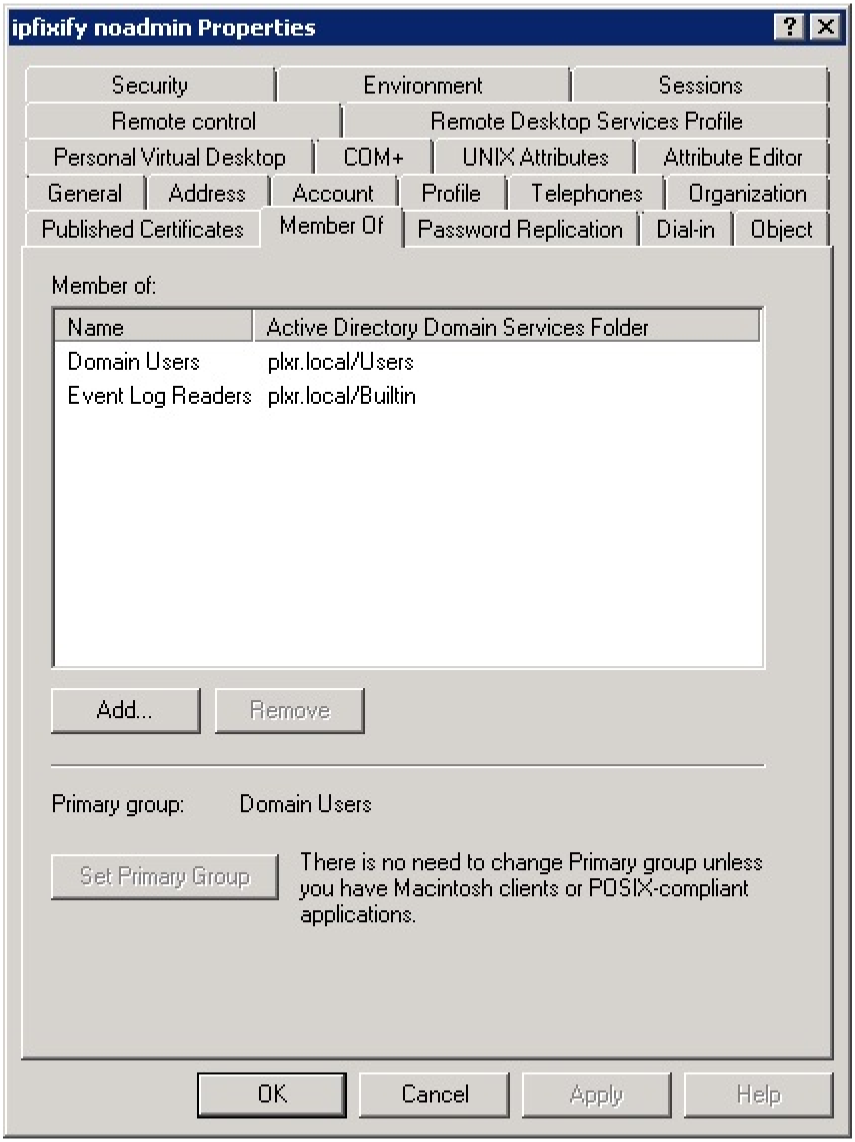

Name	Active Directory Domain Services Folder
Domain Users	plxr.local/Users
Event Log Readers	plxr.local/Builtin

Primary group: Domain Users

There is no need to change Primary group unless you have Macintosh clients or POSIX-compliant applications.

2. Provide WMI Permissions:

a. Login to the Domain Controller as an administrator.

b. Go to Start -> Run.

c. Type wmimgmt.msc.

d. Right click on WMI Control (Local) and select *Properties*.

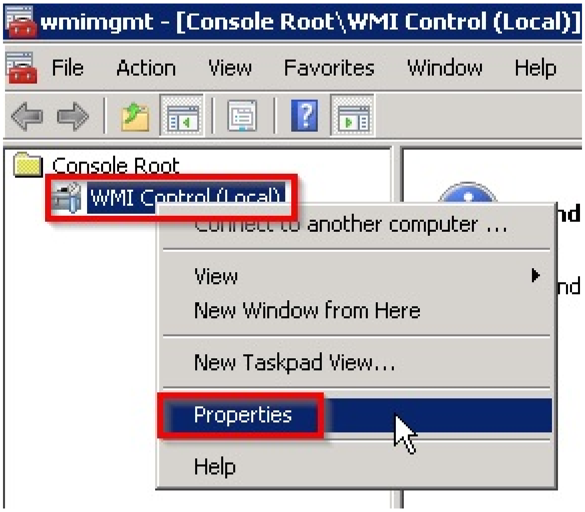

e. Go to the Security tab, click on *Root*, then select *Security*.

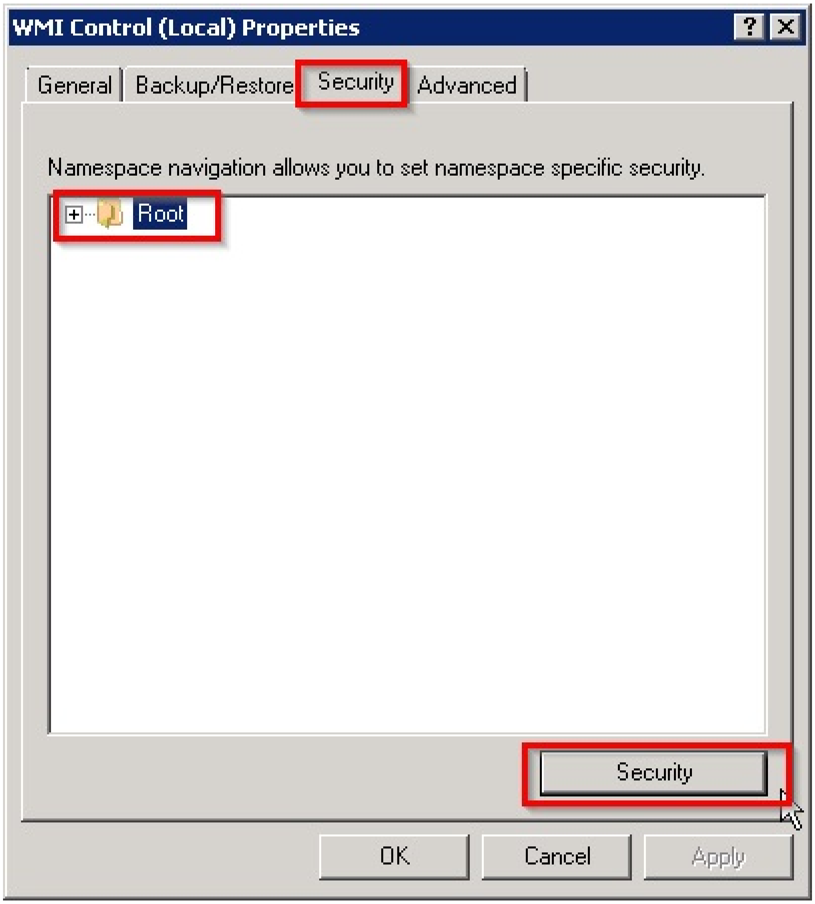

WMI Control (Local) Properties
? ✕
General
Backup/Restore
Security
Advanced
Namespace navigation allows you to set namespace specific security.
Root
Security
OK
Cancel
Apply

f. In the next popup, select *Advanced.*

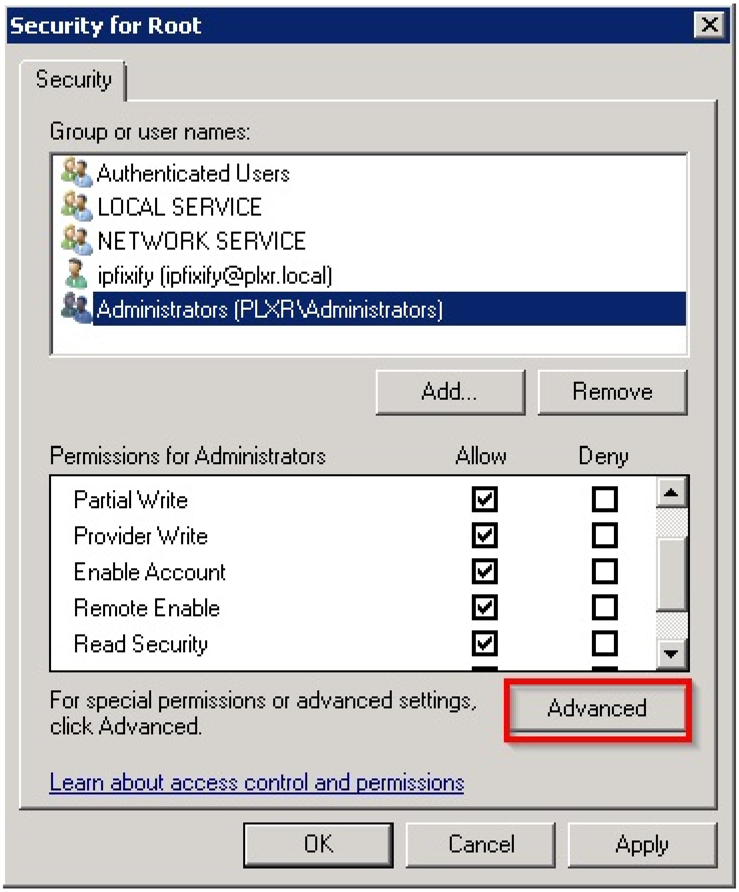

Group or user names:

- Authenticated Users
- LOCAL SERVICE
- NETWORK SERVICE
- ipfixify (ipfixify@plxr.local)
- Administrators (PLXR\Administrators)

Permissions for Administrators

	Allow	Deny
Partial Write	☑	□
Provider Write	☑	□
Enable Account	☑	□
Remote Enable	☑	□
Read Security	☑	□

For special permissions or advanced settings, click Advanced.

Learn about access control and permissions

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

h. Under the *Apply to:* section, make sure it is configured for *This namespace and subnamespaces*. Give the user *Enable Account* and *Remote Enable Allow* privileges. Apply these changes by pressing OK in each of the popup windows.

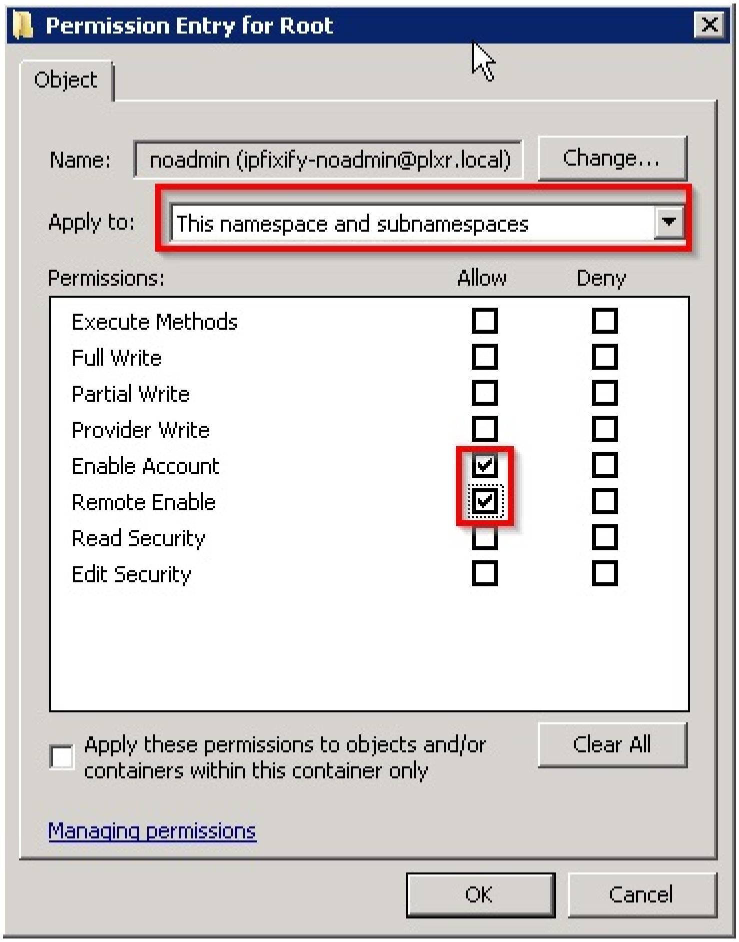

Name: noadmin (ipfixify-noadmin@plxr.local)

Apply to: This namespace and subnamespaces

Permissions:	Allow	Deny
Execute Methods	□	□
Full Write	□	□
Partial Write	□	□
Provider Write	□	□
Enable Account	☑	□
Remote Enable	☑	□
Read Security	□	□
Edit Security	□	□

□ Apply these permissions to objects and/or containers within this container only

4.13.3 Enabling Logon/Logoff Audit policies on the domain controller

1. Modify the default domain policy for domain controllers and enable the following group policies:

    a. Expand Computer Configuration -> Policies -> Windows Settings -> Security Settings -> Advanced Audit Policy Configuration -> Audit Policies -> Logon/Logoff and then enable success and failure for *Audit Logoff* and *Audit Logon*.

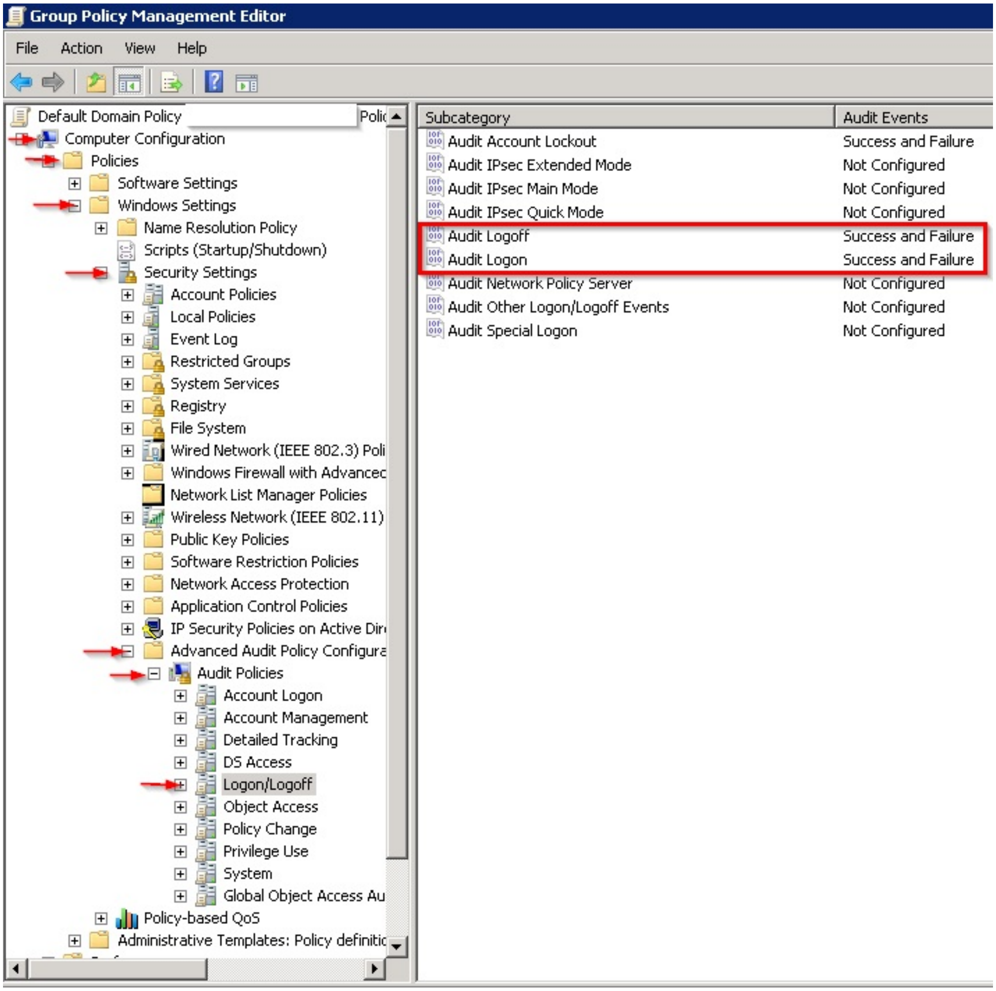

Subcategory	Audit Events
Audit Account Lockout	Success and Failure
Audit IPsec Extended Mode	Not Configured
Audit IPsec Main Mode	Not Configured
Audit IPsec Quick Mode	Not Configured
Audit Logoff	Success and Failure
Audit Logon	Success and Failure
Audit Network Policy Server	Not Configured
Audit Other Logon/Logoff Events	Not Configured
Audit Special Logon	Not Configured

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

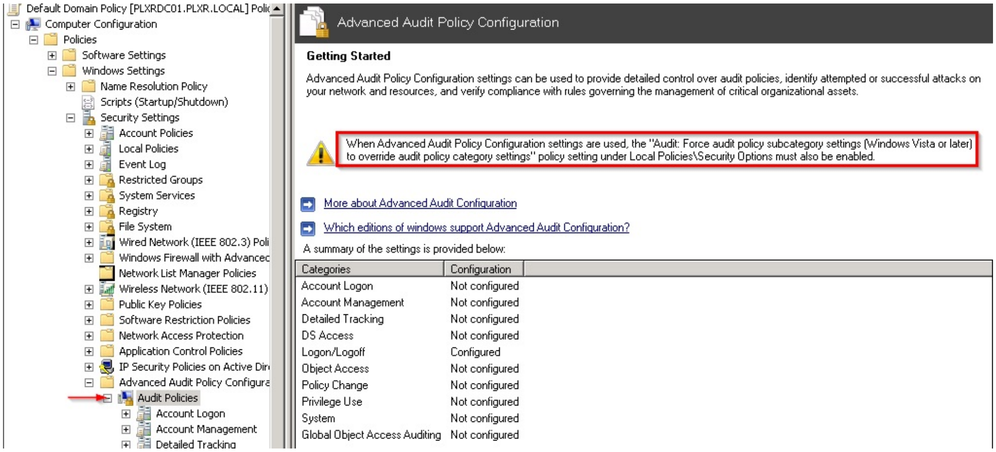
Default Domain Policy [PLXRDC01.PLXR.LOCAL] Polic
Computer Configuration
Policies
Software Settings
Windows Settings
Name Resolution Policy
Scripts (Startup/Shutdown)
Security Settings
Account Policies
Local Policies
Event Log
Restricted Groups
System Services
Registry
File System
Wired Network (IEEE 802.3) Poli
Windows Firewall with Advancec
Network List Manager Policies
Wireless Network (IEEE 802.11)
Public Key Policies
Software Restriction Policies
Network Access Protection
Application Control Policies
IP Security Policies on Active Dir
Advanced Audit Policy Configura
Audit Policies
Account Logon
Account Management
Detailed Tracking

Advanced Audit Policy Configuration

Getting Started

Advanced Audit Policy Configuration settings can be used to provide detailed control over audit policies, identify attempted or successful attacks on your network and resources, and verify compliance with rules governing the management of critical organizational assets.

When Advanced Audit Policy Configuration settings are used, the "Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings" policy setting under Local Policies\Security Options must also be enabled.

More about Advanced Audit Configuration

Which editions of windows support Advanced Audit Configuration?

A summary of the settings is provided below:

Categories	Configuration
Account Logon	Not configured
Account Management	Not configured
Detailed Tracking	Not configured
DS Access	Not configured
Logon/Logoff	Configured
Object Access	Not configured
Policy Change	Not configured
Privilege Use	Not configured
System	Not configured
Global Object Access Auditing	Not configured

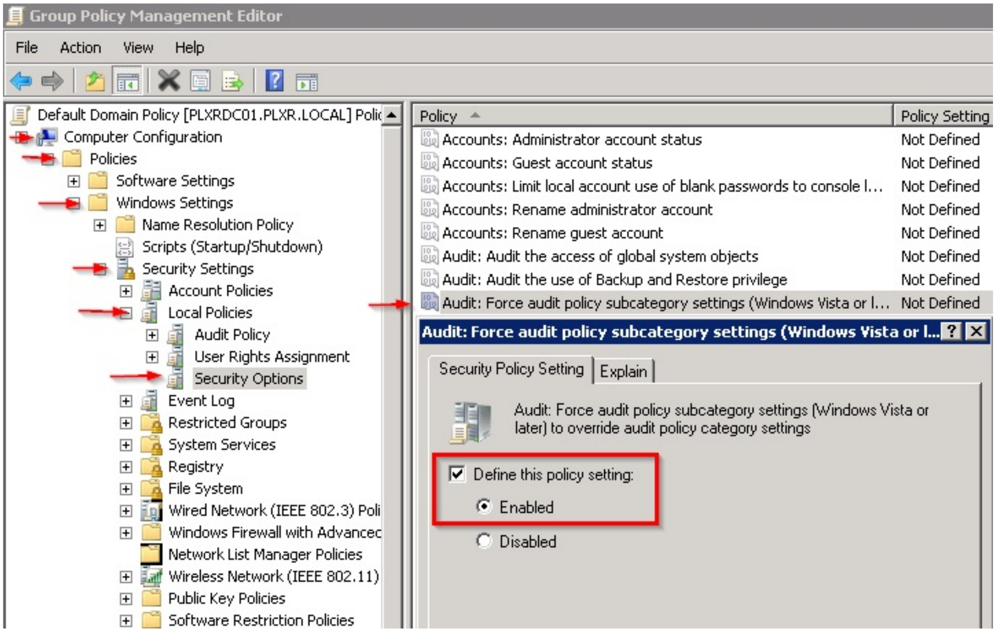

Policy	Policy Setting
Accounts: Administrator account status	Not Defined
Accounts: Guest account status	Not Defined
Accounts: Limit local account use of blank passwords to console l...	Not Defined
Accounts: Rename administrator account	Not Defined
Accounts: Rename guest account	Not Defined
Audit: Audit the access of global system objects	Not Defined
Audit: Audit the use of Backup and Restore privilege	Not Defined
Audit: Force audit policy subcategory settings (Windows Vista or l...	Not Defined

Audit: Force audit policy subcategory settings (Windows Vista or l...

Audit: Force audit policy subcategory settings (Windows Vista or later) to override audit policy category settings

☑ Define this policy setting:

● Enabled

○ Disabled

4.13.4 Setting up IPFIXify on a Windows computer

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

2. Rename ipfixify-template.cfg to ipfixify.cfg, open the file in a text editor.

Enter the NetFlow collector's IP and port:

```
collector=NetFlowIP:port
```

Enter the IP address of the domain controller. For each additional domain controller, add another member line:

```
member=DCip
```

Set this value to yes if the goal is to collect username data:

```
usernamesOnly=yes
```

3. Configure the IPFIXify user credentials

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

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

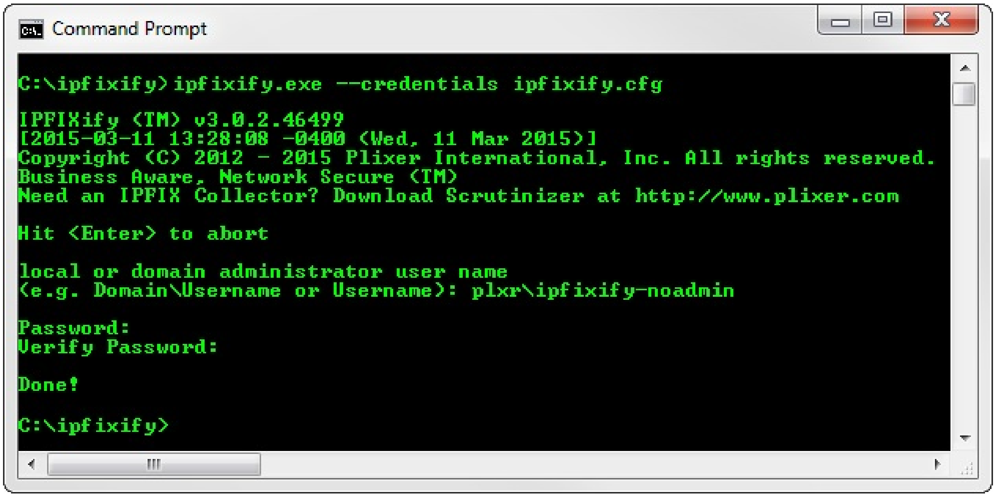

```
C:\ipfixify>ipfixify.exe --credentials ipfixify.cfg

IPFIXify (TM) v3.0.2.46499
[2015-03-11 13:28:08 -0400 (Wed, 11 Mar 2015)]
Copyright (C) 2012 - 2015 Plixer International, Inc. All rights reserved.
Business Aware, Network Secure (TM)
Need an IPFIX Collector? Download Scrutinizer at http://www.plixer.com

Hit <Enter> to abort

local or domain administrator user name
(e.g. Domain\Username or Username): plxr\ipfixify-noadmin

Password:
Verify Password:

Done!

C:\ipfixify>
```

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

Hint: PSTools.zip download: https://technet.microsoft.com/en-us/sysinternals/bb897553. Before PsExec.exe will function, the user must accept the agreement.

Hold down Shift and right-click on PsExec. In the menu, select "Run as different user"

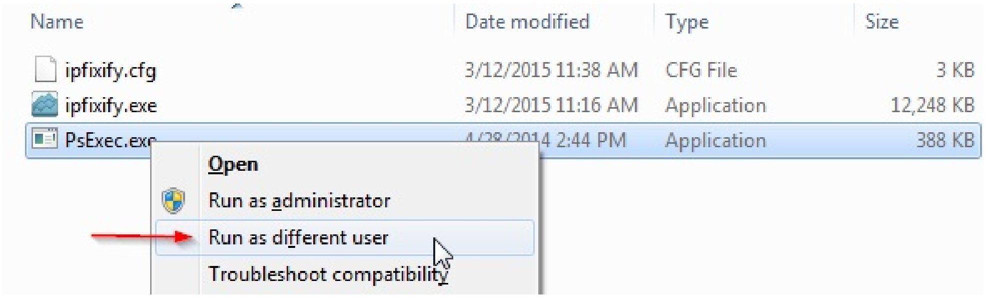

Name	Date modified	Type	Size
ipfixify.cfg	3/12/2015 11:38 AM	CFG File	3 KB
ipfixify.exe	3/12/2015 11:16 AM	Application	12,248 KB
PsExec.exe	4/28/2014 2:44 PM	Application	388 KB

Type in the IPFIXify user and password and press enter. If the user does not have access to the directory that PsExec.exe is in, this will fail. The IPFIXify user must be granted access to the directory that PsExec.exe and ipfixify.exe are in.

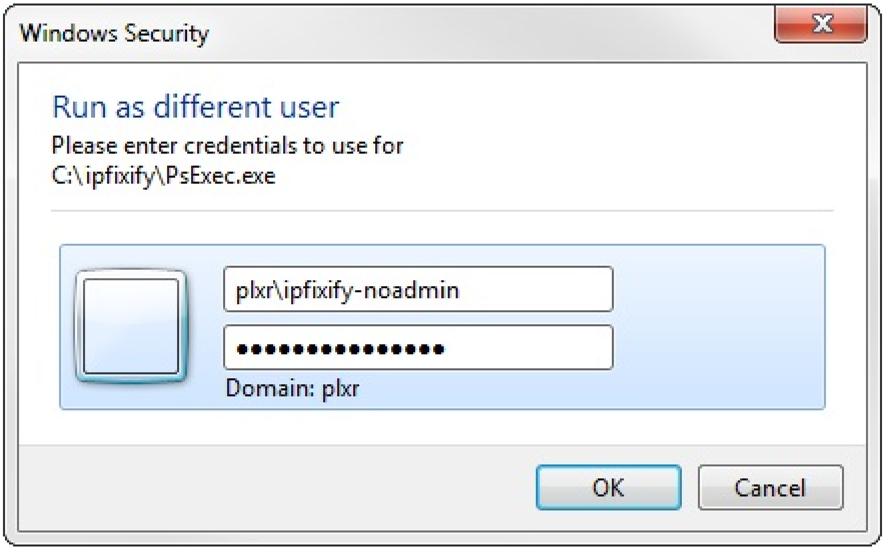

Please enter credentials to use for
C:\ipfixify\PsExec.exe

plxr\ipfixify-noadmin

●●●●●●●●●●●●●●●●

Domain: plxr

Agree to the PsExec EULA:

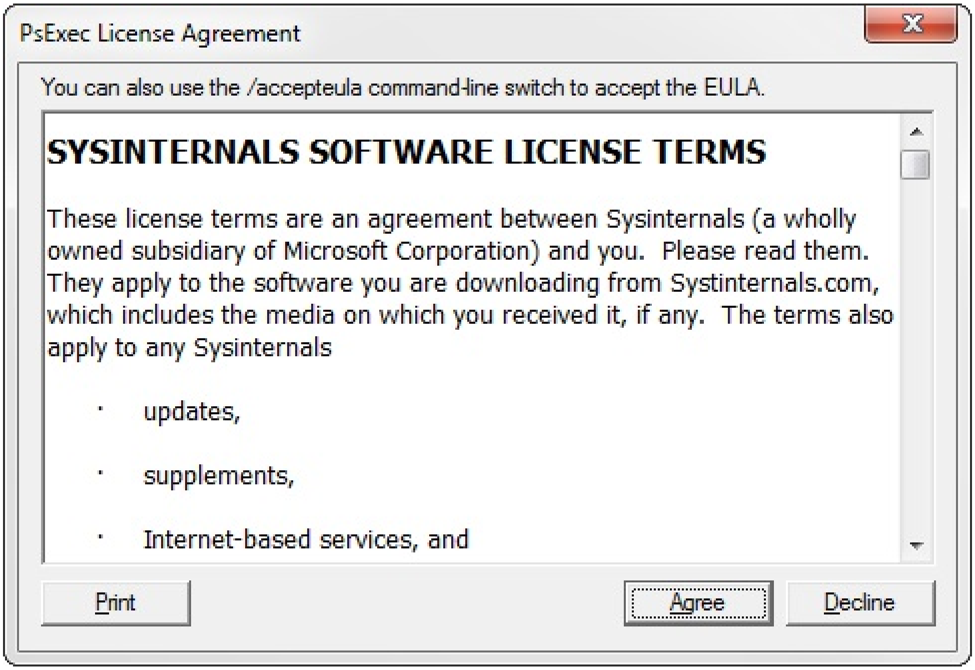

PsExec License Agreement

You can also use the /accepteula command-line switch to accept the EULA.

SYSINTERNALS SOFTWARE LICENSE TERMS

These license terms are an agreement between Sysinternals (a wholly owned subsidiary of Microsoft Corporation) and you. Please read them. They apply to the software you are downloading from Systinternals.com, which includes the media on which you received it, if any. The terms also apply to any Sysinternals

- updates,
- supplements,
- Internet-based services, and

5. From an Administrative command prompt, run the following command to verify that IPFIXify has all the permissions to poll the domain controller:

```
ipfixify.exe --sysmetrics --config “C:\ipfixify\ipfixify.cfg”␣
↪-psexec=“C:\ipfixify\PsExec.exe” -permtest IPofDC
```

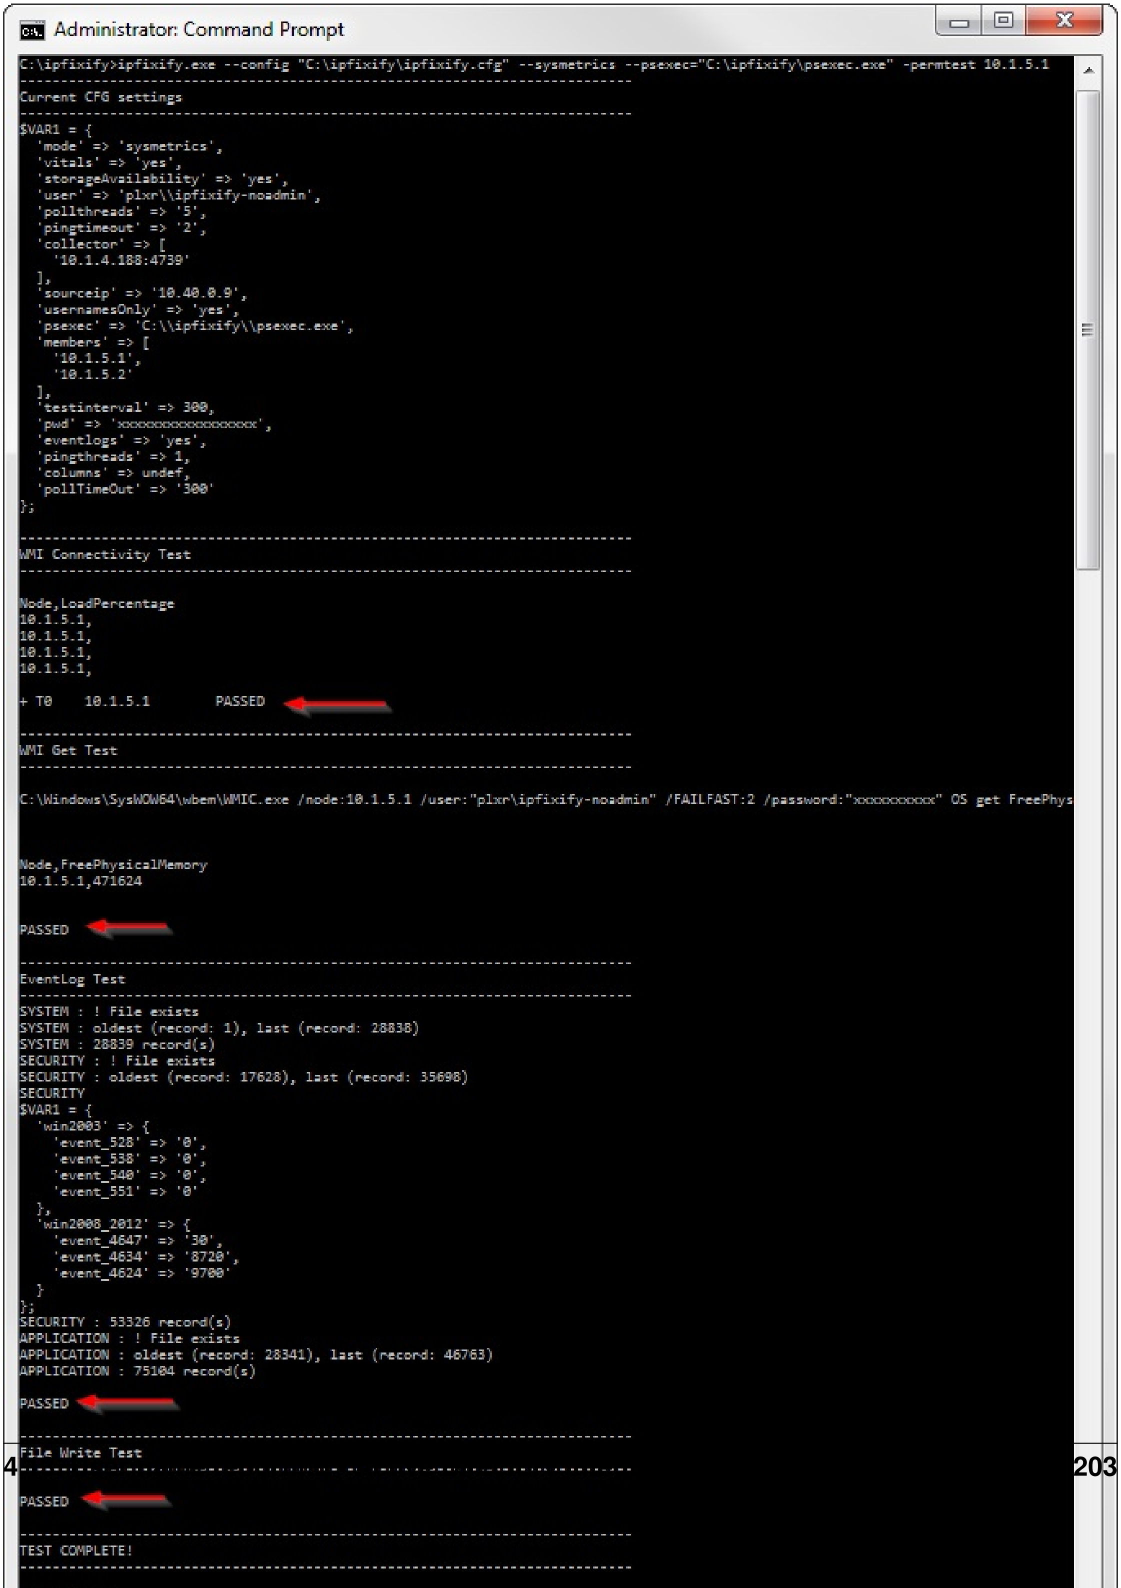

```
C:\ipfixify>ipfixify.exe --config "C:\ipfixify\ipfixify.cfg" --sysmetrics --psexec="C:\ipfixify\psexec.exe" -permtest 10.1.5.1
---------------------------------------------------------------------------
Current CFG settings
---------------------------------------------------------------------------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

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

Node,LoadPercentage
10.1.5.1,
10.1.5.1,
10.1.5.1,
10.1.5.1,

+ T0    10.1.5.1        PASSED

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

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



Node,FreePhysicalMemory
10.1.5.1,471624


PASSED

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

PASSED

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

PASSED

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

6. If all the tests passed, set up IPFIXify to run as a service. In an administrative command prompt, execute the following command:

```
ipfixify.exe --install auto --name "Scrutinizer Username Collection" --
↪config "C:\ipfixify\ipfixify.cfg" --sysmetrics --psexec=
↪"C:\ipfixify\PsExec.exe"
```

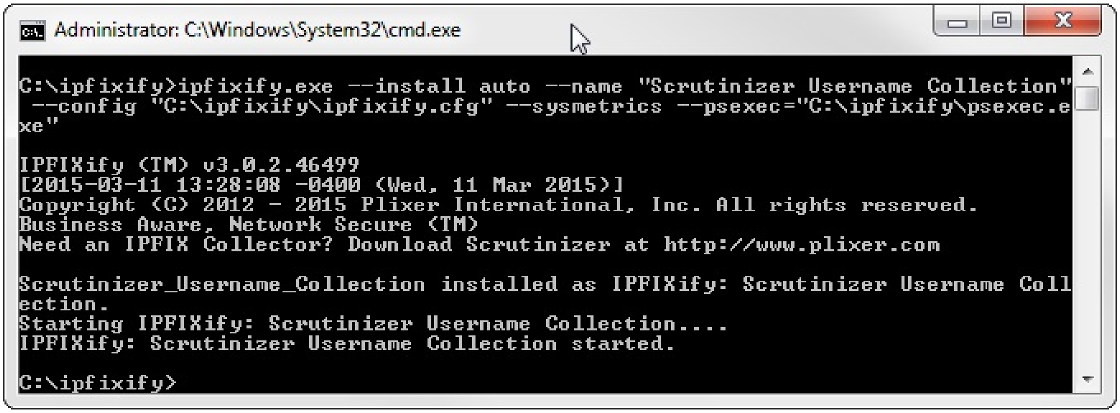

```
C:\ipfixify>ipfixify.exe --install auto --name "Scrutinizer Username Collection"
 --config "C:\ipfixify\ipfixify.cfg" --sysmetrics --psexec="C:\ipfixify\psexec.e
xe"

IPFIXify (TM) v3.0.2.46499
[2015-03-11 13:28:08 -0400 (Wed, 11 Mar 2015)]
Copyright (C) 2012 - 2015 Plixer International, Inc. All rights reserved.
Business Aware, Network Secure (TM)
Need an IPFIX Collector? Download Scrutinizer at http://www.plixer.com

Scrutinizer_Username_Collection installed as IPFIXify: Scrutinizer Username Coll
ection.
Starting IPFIXify: Scrutinizer Username Collection....
IPFIXify: Scrutinizer Username Collection started.

C:\ipfixify>
```

7. Configure the IPFIXify service to log on as the IPFIXify user.

   a. Go to Start -> Run -> and type “services.msc”

   b. Find the service named “IPFIXIfy: Scrutinizer Username Collection”, right click on it and select *Properties*.

   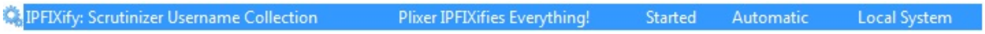

   IPFIXify: Scrutinizer Username Collection   Plixer IPFIXifies Everything!   Started   Automatic   Local System

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

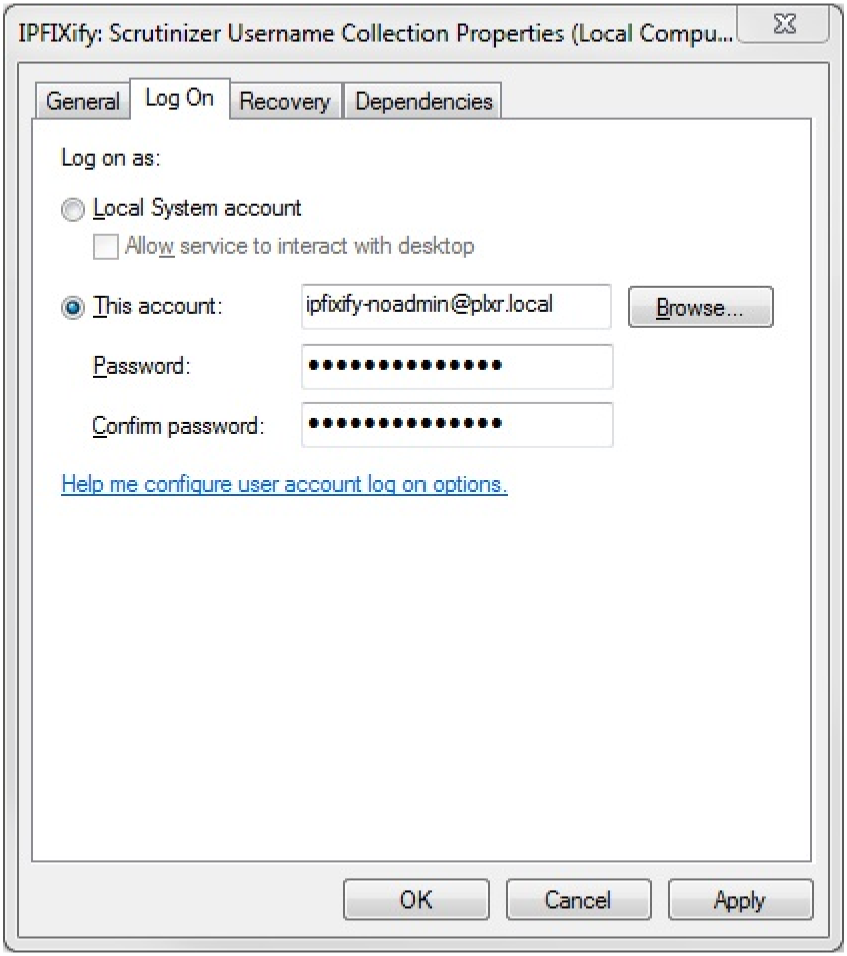

IPFIXify: Scrutinizer Username Collection Properties (Local Compu...

Log on as:

○ Local System account

□ Allow service to interact with desktop

● This account: ipfixify-noadmin@plxr.local

Password: •••••••••••••••

Confirm password: •••••••••••••••

Help me configure user account log on options.

d. Click OK. A popup will say, “the user has been granted the log on as a service right.” It means that the user will not maintain the log on as a service permission across reboots. Permission can be granted as outlined in this Microsoft document https://technet.microsoft.com/en-us/library/cc794944(v=ws.10).aspx

8. Wait a few minutes. You should start seeing user names in Plixer Scrutinizer.

Example IPFIXify configuration

```
[options]
; The IP Address/Hostname and port of the IPFIX Collector(s) multiple
; collectors can be specified on additional lines
; collector=IP:PORT (e.g. 10.1.4.19:4739)
collector=10.1.4.188:4739
; When accessing remote machines, use the supplied credentials this is
; encoded. So execute the following command to manage it
; ipfixify.exe --credentials=<PATH/TO/CFG>
credentials=6e6ff0a30ff3d13d0f9a38a753f52f44283f9a7dfd928511dbaf2f7af1446e57981dc4628c03
; Number of minutes between ping and WMI test of all members. The default
; is 60 minutes.
testinterval=5
;
↪ The number of seconds to try and ping a host during the process of verifying
; a member is reachable. If 0 is used, then the ping test is ignored.
pingtimeout=2
; The number of threads to gather data from the members who responded.
↪ If there
; is only a small list of members, then this can be a small number (e.g.
↪ 1 - 3).
; The more threads used, the more memory will be consumed by IPFIXify.
pollthreads=5
; If vitals is a true value, then CPU, Memory,
↪ and Number of processes running
; data is collected. To disable these statistics, comment out the following
; line.
vitals=yes
; If storageAvailability is a true value,
↪ then disk availability is collected.
; To disable these statistics, comment out the following line.
storageAvailability=yes
; If eventlogs is a true value, then System, Security, and Application
; Eventlogs are collected. To disable these statistics, comment out the
; following line.
eventlogs=yes
; usernamesOnly is used in conjunction with the eventlogs option.
↪  If username
; integration with Scrutinizer is the only goal, then this line should be␣
↪un-commented
usernamesOnly=yes
; If processLists is a true value,
↪ then running processes data is collected.
; To disable these statistics, comment out the following line.
;processLists = yes
; If processListCPU is a true value,
↪ then CPU per process data is collected.
```

(continues on next page)

(continued from previous page)

```
; To disable these statistics, comment out the following line.
;processListsCPU = yes
; If netstatDetails is a true value, then netstat details are collected.
; To disable these statistics, comment out the following line.
;netstatDetails = yes
; The list below contains the current hosts being polled by the IPFIXify
; Agent. One host or IP Address per line. It is recommended to use IP
; Addresses in case there are DNS issues.
member=10.1.5.1
member=10.1.5.2
```

4.14 User name reporting - Cisco ISE integration

4.14.1 Overview

User Name Reporting options for Cisco ISE include lists of user names, ability to search across all flows for user names, as well as a Cisco ISE option in the Other menu in reports to see which user has generated specific traffic.

4.14.2 Enabling ERS

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

Important: Supported versions of Cisco ISE are ISE 1.2, 1.3, 1.4, 2.0, 2.1 and 2.3.

1. On the Cisco ISE server, create a new user with the following permissions:
   - ERS Admin
   - ERS Operator
   - Super Admin
   - System Admin
2. To learn more about enabling ERS on Cisco ISE version 1.2, 1.3, 1.4, 2.0, 2.1, or 2.3, visit this page on Cisco's web site.
3. To test the configuration external to Scrutinizer, use POSTMAN to make a GET request with this URL:

```
https://[ISE_SERVER]/ise/mnt/Session/AuthList/null/null
```

Hint: When making a GET request using POSTMAN, navigate to the server with your browser, tell Chrome it is OK to use a bad certificate, and leave that window open.

4.14.3 Configuring Cisco ISE integration in Plixer Scrutinizer

1. Log into the Plixer Scrutinizer server with administrative permissions and run the following command to open the interactive CLI prompt:

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

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

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

This command adds a CiscoISE node to the queue to acquire user identity on all active sessions. The required parameters are the host address <*ise_ip*>, tcp port <*ise_tcp_port*>, and user <*ise_user*> that can access the API.

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

4.14.4 Other scrut_util options for CiscoISE

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

CHAPTER 5

Scrutinizer API

5.1 IP Groups API

The IP Groups API functionality is a simple way to add, remove, and edit IP Groups.

5.1.1 Prerequisite

When using the API the following will be used on all requests:

authToken The authentication token from Plixer Scrutinizer that allows access to API

rm The runmode for accessing the API. It is specific to each section of the product. *ipgroups* will be used for each of the following examples

action The list of available actions will change with each request. Below are the actions available within the *user_api* run mode

Action	Description
saveRule	Create a defined IP Group
update	Re-define/modify an existing IP Group
loadTreeRootFast	Load condensed list all IP Group names and IDs
search	Search for an IP Group by name
loadRules	View all rule definitions for an IP Group
deleteRule	Remove a rule from an IP Group
delete	Delete an IP Group
deleteAll	Delete all defined IP Groups from Scrutinizer

Rules

The IP Groups API provides the ability to define what makes the network traffic unique.

- **IP Host** - one or multiple IPs can be used to define a rule
- **IP Range** - use a range of IPs instead of multiple IP rules
- **IP Subnet** - use a network subnet and mask instead of multiple IP rules or ranges
- **All IPs** - specify every IP to be used in an IP Group definition
- **Wildcard Mask** - specify a wildcard mask that will define hosts available for examination
- **Child Group** - create a hierarchy of groups, with a child group being more specific than its parent.

Rule type examples

IP host

IP host - one or multiple IPs can be used to define a rule.

```
[
 {
  "type": "ip",
  "sip": "10.1.1.1"
  }
]
```

```
[
 {
  "type": "ip",
  "sip": "192.168.1.1"
 },
 {
  "type": "ip",
  "sip": "192.168.2.2"
 }
]
```

IP range

IP range - use a range of IPs instead of multiple IP rules

```
[
 {
  "type": "range",
  "sip": "10.1.1.1",
  "eip": "10.1.1.254"
 }
]
```

Note: the ‘range’ IP rule type requires a start IP (sip) and end IP (eip) to define the start and end of the range.

IP subnet

IP subnet - use a subnet and mask instead of multiple IP rules or ranges

```
[
 {
  "type": "network",
  "address": "192.168.0.0",
  "mask": "16"
 }
]
```

Note: A subnet rule uses the ‘network’ type. A subnet mask is required.

All IPs

All IPs - specify all IPs to be used in an IP Group definition

```
[
 {
  "type": "ipall",
  "all": 1
 }
]
```

Wildcard

Wildcard - specify a rule based on mask of bits that indicates which parts of an IP address are to be used for defining the IP Group hosts

```
[
 {
  "type": "wildcard",
  "address": "10.0.4.0",
  "mask": "0.255.0.255"
 }
]
```

Note: The example above will tag all hosts with the first octet of ‘10’ and the third octet of ‘4’. Therefore IPs such as 10.1.4.1, 10.2.4.250, 10.99.4.98, etc. would be included in the defined IP Group as the first and third octets match in the wildcard rule.

Child group

Child group - Nest IP Groups to create a hierarchy with child groups’ rules being more specific than their parent.

```
[
 {
  "type": "child",
  "child_id": "16900062"
 }
]
```

Important: A child group definition is based on the parent group. Define smaller and more distinct child groups, then create the parent group so that you can add the child group that already exists.

1. Create ***UK Datacenter*** and ***UK Office*** groups for their respective subnets/IPs.
2. Create a parent group ***UK*** that will include child groups *UK Datacenter* and *UK Office*.
3. Create ***Germany Datacenter*** and ***Germany Office*** groups for their respective subnets/IPs.
4. Create a parent group ***Germany*** that will include child groups *Germany Datacenter* and *Germany Office*.
5. Finally, create a parent group ***European Offices*** that will include child groups *UK* and *Germany*.

The workflow above will create the following hierarchy:

```
* European Offices         10.0.0.0/8
  ** UK                       10.30.0.0/16
        *** UK Datacenter           10.30.10.1/32
        *** UK Office               10.30.20.0/24
  ** Germany                   10.40.0.0/16
         *** Germany Datacenter      10.40.10.1/32
         *** Germany Office          10.40.20.0/24
```

5.1.2 Create an IP Group

When creating IP Groups with the API, use the **saveRule** action. There is 1 additional field that can be used:

new_fc Provide a name for the IP Group.

added Specify a Json array of rules to add to/define the IP Group.

Here is an example of how to use the **added** field for tagging a single IP address:

JSON object expected:

```
[
 {
  "type": "ip",
  "address": "10.1.4.66"
 }
]
```

Example API call

```
curl --location  --insecure --request POST '{{scrutinizer}}/fcgi/scrut_
↪fcgi.fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken={{authToken}}' \
--form 'rm=ipgroups' \
--form 'action=saveRule' \
--form 'new_fc=UK Data Center' \
--form 'added=[
        {
                "type": "ip",
                "address": "10.30.10.1"
        }
]'
```

JSON object returned:

```
{
  "removed": [],
  "updated": [],
  "added": [
      {
          "rule_id": 506588,
          "cid": null,
          "type": "ip",
          "address": "10.30.10.1"
      }
  ],
  "warnings": [],
  "fc_id": 16900006,
  "myrules": "IP Address:10.30.10.1",
  "fc_name": "UK Datacenter",
  "rule_id": 506588,
```

(continues on next page)

(continued from previous page)

```
  "total": 1
}
```

5.1.3 Update an IP Group

You can update or remove existing rules or add new ones with one request. There are four additional fields that can be optionally used with the **update** action:

name Optional. If specified, the name will be updated. Otherwise, the name of the IP Group remains unchanged.

added Optional. Specify a Json array of rules to add to/define the IP Group.

updated Optional. Specify a Json array of rules to modify the IP Group. The *rule_id* field must be defined to change a rule.

Important: When updating IP Groups, the rule type must remain the same. For example, you can not change an IP rule to a subnet rule. The workflow in that case is to remove the old rule type and create a new one.

removed Optional. Specify a Json array of rule IDs to be removed.

Note: You can leave any of the [name|added|updated|removed] fields empty.

Here is an example of the **added** field for tagging traffic for a single IP address on a specific port:

```
[
 {
  "type": "ip",
  "address": "10.1.4.66"
 }
]
```

Use the **updated** field to change an existing rules:

```
[
 {
  "rule_id": "84",
  "type": "network",
  "address": "10.30.0.0",
  "mask": "16"
 }
]
```

Here is an example of the **removed** field syntax:

```
[ 81, 82, 83 ]
```

Note: The result of the three examples above would do the following:

1. Create two new rules for the IP Group, an IP rule and a port rule.
2. Update the rule with ID 84 to change IPs to match on.
3. Remove rules 81, 82, and 83 that already existed in the IP Group definition.

Example API call

```
curl --location  --insecure --request POST '{{scrutinizer}}/fcgi/scrut_
↪fcgi.fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken={{authToken}}' \
--form 'rm=ipgroups' \
--form 'action=update' \
--form 'fc_id={{new_ipgroup_fcid}}' \
--form 'name=Renamed Group' \
--form 'added=[
        {
                "type": "ip",
                "address": "10.1.4.66"
        }
]' \
--form 'updated=[
```

(continues on next page)

(continued from previous page)

```
        {
                "rule_id": "84",
                "type": "ip",
                "address": "192.1.0.0"
        }
]' \
--form 'removed=[114]'
```

5.1.4 Search IP Groups

This feature is useful for searching Plixer Scrutinizer for IP Groups with a partial string or a full name. Search for IP Groups with comparisons such as " like 'UK Office' " or " notLike 'Email Server' ". There are four additional fields that you can use with the **search** action.

name The name of the IP Group (or string) to find in Plixer Scrutinizer

fc_name_comp Specify the comparison criteria for search results. Valid options are [like|notLike]

page Default is 1, which would load the first page of pagination.

maxRows The number of results per page returned in the API response.

Example API call

```
curl --location  --insecure --request POST '{{scrutinizer}}/fcgi/scrut_
↪fcgi.fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken={{authToken}}' \
--form 'rm=ipgroups' \
--form 'action=search' \
--form 'name={{search_term}}' \
--form 'fc_name_comp={{search_type}}' \
--form 'page=1' \
--form 'maxRows=10'
```

5.1.5 Delete entry from an IP Group

Use this command to remove a single IP from a specific IP Group. The rule_id is required and can be obtained by viewing the rules of an IP Group. The update action also has an optional 'removed' field that can be used for deleting rules from an IP Group.

Example API call

```
curl --location  --insecure --request POST '{{scrutinizer}}/fcgi/scrut_
↪fcgi.fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken={{authToken}}' \
--form 'rm=ipgroups' \
--form 'action=deleteRule' \
--form 'rule_id=506588'
```

JSON object returned:

```
{
  "fc_id": 16900006,
  "success": 1,
  "myrules": "",
  "rule_id": "506588",
  "total": 0
}
```

5.1.6 Delete IP Groups

You can remove more than one IP Group at a time by specifying more IDs in the array. There is an additional field that is required with the **delete** action:

json Array of IP Group IDs to be deleted

For example, here's how you define the **json** field for removing a single IP Group:

```
[
 {
  "id": "16900032"
 }
]
```

This is the **json** field for removing multiple IP Groups:

```
[
 {
  "id": "16900032"
 },
 {
  "id": "16900033"
 }
]
```

Example API call

```
curl --location  --insecure --request POST '{{scrutinizer}}/fcgi/scrut_
↪fcgi.fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken={{authToken}}' \
--form 'rm=ipgroups' \
--form 'action=delete' \
--form 'json=[
        {
        "id": "16900032"
        }
]'
```

JSON object returned:

```
{
  "processedCount": 1,
  "removed": [
      "16900006"
  ]
}
```

5.2 Reporting API

5.2.1 Overview

The reporting API is a simple way to access the report data via HTTP or the command line.

5.2.2 Prerequisites

When using the API, pass the following mandatory fields:

authToken The authentication token from Plixer Scrutinizer that allows access to API.

rm The runmode for accessing the API. It is specific to each section of the product. *report_api* will be used for each of the following examples.

action The list of available actions will change with each request. Below is a list of available actions within the *report_api* run mode:

Field	Description
get	the action used to execute flow reports

rpt_json An array of details to specify most of the options found in the gear menu of the UI, where you can specify report type, date range, etc. *See section below for more details* on how to fill out the rpt_json field. Every report has an API tab for obtaining JSON for it..

data_requested It is used to indicate to the API what part of the report we need. Reports return two data sets, one for rendering the graph and the other for rendering the table. Each one is divided into outbound and inbound.

While rpt_json tells Plixer Scrutinizer how to prepare the data (timeframe, filters, aggregation, etc.), *data_requested* tells Plixer Scrutinizer what data should be returned with the request (inbound, outbound, table, graph, or just a table name).

5.2.3 Available Parameters

rpt_json

The rpt_json field tells Plixer Scrutinizer how to prepare the data: report type, timeframe, filters, aggregation, etc.

JSON object expected (More options available, this is the minimal needed):

```
{
  "reportTypeLang": "conversations",
  "filters": {
      "sdfDips_0": "in_0A190101_ALL"
  },
  "reportDirections": {
      "selected": "inbound"
  },
  "times": {
      "dateRange": "LastFiveMinutes",
      "clientTimezone": "America/New_York"
  },
  "dataMode": {
      "selected": "saf"
  },
  "rateTotal": {
      "selected": "total"
  },
  "dataGranularity": {
      "selected": "auto"
  },
  "bbp": {
      "selected": "bits"
  }
}
```

Here is a breakdown of the options for each of the fields from above:

		Description
reportTypeLang		a language keycode that represents a r
		Report Lang
		conversations
		host2host
		ipGroupGroup
		applications
		country2country
		..etc..
filters		
	sdfDips_0	The exporter and interface to use to ge Hex]_[interface(s)]
		Value
		in_0A190101_ALL
		in_0A190101_0A190101_1

Table 1 – continued from previous p

reportDirections		
	selected	Specify which direction to query. Directic just use inbound.
		inbound
		outbound
times		You can specify just dateRange at a minin
	dateRange	Indicate timeframe to run report. Valid Op
		LastFiveMinutes
		LastTenMinutes
		LastFifteenMinutes
		LastTwentyMinutes
		LastThirtyMinutes
		LastFortyfiveMinutes
		LastHour
		LastFullHour
		LastThreeDays
		LastSevenDays
		LastThirtyDays
		Today
		Yesterday
		Last24Hours
		ThisWeek
		LastWeek
		ThisMonth
		LastMonth
		ThisYear
		LastYear
		Custom
	start	Epoch timestamp of the report start (Optic
		1597974180
	end	Epoch timestamp of the report start (Optic
		1597974480
	clientTimezone	Used to display dates local to your timezo
		America/New_York
		America/Los_Angeles
		..etc..
dataMode		The system can save and roll up data in or <*data_aggregation*> to learn more
	selected	
		Value
		saf

Table 1 – continued from previo

		traditional
rateTotal		specify whether to display data as a ra
	selected	
		Value
		rate
		total
dataGranularity		Data from each exporter is stored and retrieve data for report.
	selected	
		Value
		auto
		1m
		5m
		30m
		2h
		12h
bbp		This field determines how to display d
	selected	
		Value
		bits
		bytes
		percent

data_requested

The data_requested field tells Plixer Scrutinizer how to prepare the data for graphs, table pagination, etc.

Important: The direction specified in data_requested needs to match the reportDirections selected value in rpt_json. e.g. (inbound/inbound or outbound/outbound)

JSON object expected (More options available, this is the minimal needed):

```
{
    "inbound": {
        "graph": "none",
        "table": {
            "query_limit": {
```

(continues on next page)

(continued from previous page)

```
                "offset": 0,
                "max_num_rows": 10
            }
        }
    }
}
```

5.2.4 Running A Report

The following example is an API call to run a default report, over the last 5 minutes, on all interfaces of a device.

Note: In the example below, you will need to replace `{{scrutinizer_ip_address}}` with your Plixer Scrutinizer's IP address, as well as `{{authToken}}` with an Authentication Token that can be obtained from the UI.

Example API call

```
curl --location --request POST 'https://{{scrutinizer_ip_address}}/fcgi/
↪scrut_fcgi.fcgi' \
--header 'Content-Type: application/json' \
--form 'authToken={{authToken}}' \
--form 'rm=report_api' \
--form 'action=get' \
--form 'rpt_json=
{
    "reportTypeLang": "conversations",
    "filters": {
        "sdfDips_0": "in_0A190101_ALL"
    },
    "reportDirections": {
        "selected": "inbound"
    },
    "times": {
        "dateRange": "LastFiveMinutes",
        "clientTimezone": "America/New_York"
    },
    "dataMode": {
        "selected": "saf"
```

(continues on next page)

(continued from previous page)

```
    },
    "rateTotal": {
        "selected": "total"
    },
    "dataGranularity": {
        "selected": "auto"
    },
    "bbp": {
        "selected": "bits"
    }

}' \
--form 'data_requested=
{
    "inbound": {
        "graph": "none",
        "table": {
            "query_limit": {
                "offset": 0,
                "max_num_rows": 10
            }
        }
    }
}'
```

The above API call will get processed by the reporting engine. The server will return a JSON response.

JSON object returned (Note: Response condensed to show structure):

```
{
    "report": {
        "request_id": "0xed184820e4b611eab58f1fc02130f7f9",
        "table": {
            "inbound": {
                "totalRowCount": 1,
                "footer": [],
                "columns": [],
                "rows": []
            }
        },
        "time_details": {},
        "exporter_details": {},
        "graph": {}
    }
}
```

Here is a breakdown of the most important fields from the ‘report’ field/key from the above response:

Field		Description	
table		The table data is divided into inbound or outbound. For each direction the following are provided:	
	columns	An array of objects that represent the definition of each column in the table. Most of the data is used to create the html table in the browser	
		elementName	Name of the element that the data in the column is for
		format	Details about how the data in the column needs to be formatted
		label	Label used in the table header
	rows	Each collection in the array represents a row and each object in the collection represents a table cell for that row.	
		rawValue	The unformatted value as it is return from the database
		label	formatted value including bits bytes or percent
	footer		
		The footer[0] represents the ‘Others’ data for the operations columns (columns which value is the product of an arithmetic operation). That is, the amount for the data not included in the rows.	
		The footer[1] represents the ‘total’. Total for a column is equal to the sum of the rows for that column plus the ‘Others’ value for that column.	
	totalRowCount	integer representing the total number of rows available	
graph			
		all	return all available graph types
		pie	values used in drawing pie chart of table data
		timeseries	values used in drawing line graph chart of table data
		none	Will only return data for pie, as that is the default

5.3 User API

5.3.1 Overview

The user API functionality allows creating user accounts within Plixer Scrutinizer and adding them to user groups at the same time.

5.3.2 Prerequisites

When using the API, pass the following mandatory fields:

authToken The authentication token from Scrutinizer that allows access to API.

rm The runmode for accessing the API. It is specific to each section of the product. *user_api* will be used for each of the following examples.

action The list of available actions will change with each request. Below is a list of available actions within the *user_api* run mode:

5.3.3 Creating users

The **createUser** action allows creating users within the API. It calls for an additional field:

json An array of users each contains the name, password and group template that user will be a member of.

JSON object expected:

```
{
 "users":[
     {
         "name": "MyAdmin",
         "pass":"secretAdminPass",
         "membership":[2]
     },
     {
         "name": "MyGuest",
         "pass":"myGuestPass",
         "membership":[2]
     }
 ]
}
```

Field	Description
users	An array of all users to be created. Multiple users can be created at once. If only one user is needed, this will be an array of one.
name	The name for the new user account.
pass	The password for the new user account.
membership	An array of usergroup_ids from the plixer.usergroups table. The user will be added to these groups when the account is created. You can add your own usergroups. By default, there are two groups installed. Group 1 is the Administrators group. Group 2 is reservered for Guests.

JSON object returned:

```
{
 "data": [
     {
         "id": 3,
         "name": "MyAdmin"
     },
     {
         "id": 4,
         "name": "MyGuest"
     }
 ]
}
```

Field	Description
data	An array of responses for each user account that Scrutinizer attempted to create.
id	The new user_id of the user account that was created.
name	The name of the account created (by design this is identical to the name passed in).

Example API call

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
↪fcgi' \
 --form 'authToken={{authToken}}' \
 --form 'rm=user_api' \
 --form 'action=createUser' \
 --form 'json=
 {
     "users":[
         {
```

(continues on next page)

(continued from previous page)

```
            "name": "MyAdmin",
            "pass": "MyPass",
            "membership": [ 1 ]
        },
        {
            "name": "MyGuest",
            "pass": "OtherPass",
            "membership": [ 2 ]
        }
    ]
}'
```

Note: If Plixer Scrutinizer is using a self-signed certificate, add `--insecure` to the header options to tell curl to ignore it.

5.3.4 Deleting users

The **delUser** action allows the deletion of user accounts within Plixer Scrutinizer by ID or name. There is an additional field used with the action:

json An array of users where each contains the name, password and group template that user will be a member of.

JSON object expected:

```
{
 "delUsers":[
         11,
         "MyGuest",
         207
 ]
}
```

Field	Description
delUsers	An array of all users to be deleted. You can delete multiple users at once. If only one user is needed, this will be an array of one.
id/-name	The user_id of the user to be deleted. Alternatively, the name of the user can be used also.

JSON object returned:

```
{
 "data":[
     "Deleting user id 11 (1 matched)",
     "Deleting user named 'MyGuest' (1 matched)"
     "Deleting user id 207 (0 matched)",
 ]
}
```

Field	Description
data	An array of responses for each user account that Plixer Scrutinizer attempted to delete. The example includes a failure message when a user did not exist.

Example API call

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
↪fcgi' \
 --form 'authToken={{authToken}}' \
 --form 'rm=user_api' \
 --form 'action=delUsers' \
 --form 'json=
 {
     "delUsers":[
         11,
         "MyGuest",
         207
     ]
 }'
```

5.3.5 Creating Usergroups

You can use the **createUsergroup** action to create user groups and add members at the same time. It requires an additional field:

json An array of usergroups. Each entry contains the name of the usergroup, the id of the usergroup to use as a template, and the id or name of the users to be added to the group.

JSON object expected:

```
{
 "usergroups":[
      {
          "name":"GroupA",
          "template_usergroup":1,
          "users":[1,2]
      },
      {
          "name":"GroupB",
          "template_usergroup":2,
          "users":["MyUser","MyUser2"]
      }
 ]
}
```

Field	Description
usergroups	An array of responses for each usergroup that Scrutinizer attempted to create.
name	The name to be applied to the usergroup
template_usergroup	the id of the user group to use as a template for creating a new user group.
users	An array of all users to be added, by user ID or username. If only one user is needed, this will be an array of one. An empty array will create an empty user group

JSON object returned:

```
{
 "data":[
      {
          "id":5,
          "name":"GroupA",
          "members":["1","2"]
      },
      {
          "name":"GroupB",
          "error":"A usergroup already exists with that name"
      }
 ]
}
```

Field	Description
data	An array of responses for each user group that Plixer Scrutinizer attempted to create.
id	The new usergroups_id of the user group that was created.
name	The name of the user group created (by design this is identical to the name passed in)
members	An array of user IDs or user names for the members successfully added to the group
error	Any errors encountered during the creation of a particular user group

Example API call

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
↪fcgi' \
 --form 'authToken={{authToken}}' \
 --form 'rm=user_api' \
 --form 'action=createUsergroup' \
 --form 'json=
 {
     "usergroups":[
         {
             "name":"GroupA",
             "template_usergroup":1,
             "users":[1,2]
         },
         {
             "name":"GroupB",
             "template_usergroup":2,
             "users":["MyUser","MyUser2"]
         }
     ]
 }'
```

5.3.6 Deleting user groups

The **delUsergroups** action deletes user groups. It has an additional field:

json An array of usergroups. Each entry contains the name or ID of the usergroup to be deleted.

JSON object expected:

```
{
 "delUsergroups": [
     3,
     "My Usergroup"
 ]
}
```

Field	Description
delUsergroups	An array of responses for each user group that Scrutinizer attempted to delete
id/name	The usergroups_id or exact name of the user group to be deleted

JSON object returned:

```
{
 "data":[
     "Deleting usergroup named '3' (1 matched)",
     "Deleting usergroup id My Usergroup (0 matched)"
 ]
}
```

Field	Description
data	An array of responses for each usergroup that Plixer Scrutinizer attempted to delete.

Example API call

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
↪fcgi' \
 --form 'authToken={{authToken}}' \
 --form 'rm=user_api' \
 --form 'action=delUsergroups' \
 --form 'json=
 {
     "delUsergroups": [
         3,
         "My Usergroup"
     ]
 }'
```

5.3.7 Modifing group membership

The **membership** action changes user group membership. When adding or removing a user, use either “user_id” or “user_name”, but not both as shown below. There is an additional field that can be used with the **membership** action in the user API:

json Contains two arrays, “add” and “remove,” which have information on each membership change.

JSON object expected:

```
{
 "membership":
     {
     "add":[
         {
             "user_id": 13,
             "usergroup_id": 2
         },
         {
             "user_name":"USER2",
             "usergroup_name":"USERGROUP2"
         }
     ],
     "remove":[
         {
             "user_name":"USER3",
             "usergroup_id": 4
         }
     ]
 }
}
```

Field	Description
membership	Contains two arrays, “add” and “remove,” which have information on each membership change
user_id	Required for the user with preferences to change
user_name	An alternative to user_id. It can be the plain text name of the user.
usergroup_id	The ID from plixer.usergroups that the user will be added to or removed from.
usergroup_name	An alternative to usergroup_id. It can be the plain text name of the user group.

JSON object returned:

```
{
 "data":
     "added":[
         "User 13 added to usergroup 1",
         "User 14 added to usergroup 3",
     ],
     "removed":[
         "User 15 removed from usergroup 4"
     ]
}
```

Field	Description
data	An array of responses for membership updated
added	Contains an array of either statements of success or statements of errors explaining why the membership change failed
removed	Contains an array of either statements of success or statements of errors explaining why the membership change failed

Example API call

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
↪fcgi' \
 --form 'authToken={{authToken}}' \
 --form 'rm=user_api' \
 --form 'action=membership' \
 --form 'json=
 {
     "membership": {
         "add": [
             {
                 "user_id": 3,
                 "usergroup_id": 17
             },
             {
                 "user_id": "USER2",
                 "usergroup_id": "GROUPB"
             }
         ],
         "remove": []
     }
 }'
```

5.3.8 Editing user preferences

The **prefs** action makes changes to the user preferences for individual accounts. It contains an array of preferences and new settings.There is an additional field used with the **prefs** action in the user API:

json the user_id and array of prefs each contains the pref code and setting value to be modified.

JSON object expected:

```
{
 "user_id": 11,
 "prefs":[
     {
         "pref":"statusTopn",
         "setting":10
     },
     {
         "pref":"language",
         "setting":"english"
     }
 ]
}
```

Field	Description
user_id	Required for the user with preferences to change
prefs	An array of user preferences and setting values
pref	The Scrutinizer user preference to edit
setting	The value that will be set for the user_id specified

JSON object returned:

```
{
 "data": {
     "updated": [
         "statusTopn updated to 10 for user_id 11",
         "language updated to english for user_id 11"
     ],
     "errors": []
 }
}
```

Field	Description
data	An array of responses for each preference change updated or attempted
updated	Messages for any preference successfully changed
errors	Any errors encountered while changing preferences

Example API call

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
↪fcgi' \
 --form 'authToken={{authToken}}' \
 --form 'rm=user_api' \
 --form 'action=prefs' \
 --form 'json=
 {
     "user_id":11,
     "prefs":[
         {
             "pref":"statusTopn",
             "setting":10
         },
         {
             "pref":"language",
             "setting":"english"
         }
     ]
 }'
```

5.3.9 Changing permissions

The **permission** action makes changes to a user group's permissions. Users inherit permissions from their user group. There is an additional field used with the **permissions** action in the user API:

json An array of permissions each contains a usergroup identifier (name or ID), the security code and permission type to be modified.

JSON object expected:

```
{
 "permissions": {
     "add": [
         {
             "usergroup_name": "Dashboarders",
             "permission_type": "gadget",
             "seccode": "lLabelCPU"
         }
     ],
     "remove": [
         {
             "usergroup_name": "ReadOnlyReporters",
             "permission_type": "plixer",
             "seccode": "allGadgets"
         }
     ]
 }
}
```

Field	Description	
permissions	Contains two arrays, “add” and “remove,” which have information on each permission change	
add/remove	Each contains an array of objects of permissions and user groups	
usergroup_id	The ID from plixer.usergroups that the user will be added to or removed from	
usergroup_name	An alternative to usergroup_id, and can be the plain text name of the user grou	
permission_type	differentiates the types of permissions. Options include:	
	Permission	Description
	device	A hex representation of the IP address of a devic
	interface	Hyphenated, a hex representation of the IP addre
	group	The group ID of a map/device group from plixer
	report	The saved_id of a saved report from reporting.sa
	gadget	The gadget_id of a dashboard gadget, from plixe
	thirdparty	The id of a 3rd party link from plixer.third_party
	bboard	The bulletin board id from plixer.alm_bulletin_b
	plixer	A static string of permission codes we use for di
seccode	another field in plixer.usergroups_permissions and contains the permission. It will be different depending on the “type” (e.g. “1” for interface 1, “0A010107” for a device, etc.)	
	seccode	Description
	3rdPartyIntegration	Create, edit, and delete third-party integration lin
	ackBBEvent	Ability to acknowledge events on Alarms tab bu
	adminTab	Access the Admin Tab
	alarmSettings	Configure alarm notifications
	alarmsTab	Access the Alarms tab

Ta

	allBBoards	View all Alarms bulletin boards
	allDevices	The status of all devices and each of their int
	allGadgets	Every gadget created on the Dashboards tab,
	allGroups	Access to all maps/device groups created in S
	allInterfaces	Report on interfaces for any device
	allLogalotReports	All Logalot Reports
	allReportFolders	Permission to all saved report folders
	allReports	Saved reports created by any user
	allThirdparty	All configured third-party links will be availa
	almDelete	Permission to permanently delete alarms
	ApplicationGroups	Configure Application Groups
	asnames	Configure AS Names
	auditing	Access the Auditing report containing logs o
	auth	Manage external authentication tokens
	Authentication	Manage external authentication types
	authLdapServers	Manage LDAP server configuration used for
	awsSettings	AWS configuration
	changeUserPasswords	The ability to change the passwords of other
	createDashTabs	Create new Dashboards
	createUsers	The ability to create new local Scrutinizer us
	CrossCheck	View and edit CrossCheck configuration, wh
	crossCheckView	Access to the CrossCheck methods table vie
	dashboardAdmin	Manage all dashboards created by any user
	DataHistory	Configure settings that control how long Scr
	deleteReport	Ability to delete saved reports regardless of c
	deleteUsers	The ability to delete local Scrutinizer user ac
	DeviceDetails	Edit device interface details
	EmailNotifications	Configure the mailserver Scrutinizer will use
	faExclusions	Configure Flow Analytics exclusions
	fa_mgmt_link	Configure Flow Analytics thresholds and set
	feedbackForm	Access the link to send feedback to Plixer
	FlowAnalyticsSettings	Global Flow Analytics settings
	helpTab	Access the Help tab
	HostNames	Edit Host Name information
	IPGroups	Configure Scrutinizer IP Groups
	language	Create and edit language localization setting
	licensing	Configure Scrutinizer product licensing and
	LogalotPrefs	Configure global alarms settings
	MACAddresses	Configure device MAC Address information
	ManageCollectors	Manage the devices collecting flow data for S
	ManageExporters	Manage the devices exporting flow data to S
	mappingGroupConfiguration	Create and edit Maps/Groups

Table

	mappingObjectConfiguration	Create and edit Mapping Objects
	mapsTab	Access the Maps tab
	myViewTab	Access the Dashboards tab
	NotificationManager	Manage alarm notifications
	PolicyManager	Manage alarm policies
	protocolExclusions	Edit which protocols are discarded from flow rep
	proxySettings	Configure proxy server settings in Scrutinizer
	radiusConf	Manage RADIUS server configuration used for S
	ReportDesigner	Design new custom report types
	reportFilters	Permission to update the filters used in Status Ta
	reportFolders	Manage saved report folders
	reportSettings	Reporting engine configuration options
	runReport	Ability to run flow reports
	saveReport	Ability to name and save flow reports
	scheduledReports	Create, edit, and delete scheduled email reports
	sf_asa_acls	Configure ASA ACL descriptions
	SNMPCredentials	Manage SNMP credentials used to poll device in
	srCreate	Schedule a saved report to be emailed on a regul
	sso	Add, Delete, and Edit Identity Provider configur
	statusTab	Access the Status Tab
	syslogNotifications	Syslog server configuration
	SystemPreferences	Administrative access to global Scrutinizer prefe
	tacacsConf	Manage TACACS+ server configuration used fo
	tos	Edit TOS Configuration
	userAccounts	Access to the Users view on the Admin Tab, listi
	usergroups	Manage Scrutinizer usergroups
	viewUserIdentity	View identity and access information relevant to
	viptelaSettings	Viptela Settings
	Vitals	View the Scrutinizer server vitals reports
	wkp	Edit WKP Configuration

JSON object returned:

```
{
 "data": {
     "errors": [],
     "updated": [
         "Added gadget permission lLabelCPU to usergroup 26 ",
         "Removed plixer permission allGadgets from usergroup 27 "
     ]
 }
}
```

Field	Description
data	An array of responses for each permission change updated or attempted
updated	Messages for any sucessful changes to permissions
errors	An array of errors explaining why the permission change failed

Example API call

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
↪fcgi' \
 --form 'authToken={{authToken}}' \
 --form 'rm=user_api' \
 --form 'action=permissions' \
 --form 'json=
 {
     "permissions": {
         "add": [
             {
                 "usergroup_id": 23,
                 "permission_type": "plixer",
                 "seccode":   "statusTab"
             }
         ],
         "remove": []
     }
 }'
```

5.3.10 Changing user names

The **changeUsername** action allows editing the name of a user account. It requires an additional field:

json An array of user groups each containing the existing name of user, the new user name to be set. Alternatively, the user_id can be used instead of the oldname field.

JSON object expected:

```
{
 "changeUsername":
     {
         "user_id":"14",
         "newname":"OpSCT"
     }
}
```

Field	Description
user_id	The user ID of the account to be changed
oldname	An alternative to user ID, and contains the current name of the user
newname	Contains the name to which you wish to change this user

JSON object returned:

```
{
 "data":
      {
          "message": "User myUser successfully renamed to OpSCT"
      }
}
```

Field	Description
data	An array of responses for each preference change updated or attempted
message	Contains either a statement of success or an error explaining why the name change failed

Example API call

```
curl --location --insecure --request POST '{{scrutinizer}}/fcgi/scrut_fcgi.
↪fcgi' \
 --form 'authToken={{authToken}}' \
 --form 'rm=user_api' \
 --form 'action=changeUsername' \
 --form 'json=
 {
     "changeUsername":
         {
             "oldname":"myUser",
             "newname":"OpSCT"
         }
 }'
```

CHAPTER 6

Dashboards

6.1 Overview

Dashboards are used to create custom views of precisely what the user or group of users wants to see when they log in. Multiple unique dashboards can be created.

- With the right permissions, these dashboards are customizable per login account.
- All dashboards created by any user in a usergroup are available to other users in the same usergroup. The default is read-only access.
- Each dashboard can be manipulated and shared with others.
- The Read-only permission (check box) is used to grant others the ability to manipulate a shared dashboard.

6.2 Dashboard administration

In the upper left-hand corner of the dashboard there are three drop down menus.

1. **Gear with down arrow:**

    - If the user has permission, this option can be used to change the dashboard name.
    - Set the default dashboard when the Dashboard tab is clicked.
    - If the user has permission, the user can make a dashboard **Read-Only** to others whom will be viewing the same dashboard. Leaving unchecked allows them to change the dashboard which includes rearranging as well as adding and removing gadgets.
    - The user with ownership of the dashboard is also displayed with the dashboard ID. The dashboard ID can be accessed directly through a URL: https://<server>/dashboard/id/<dashboard_id>
    - A user wanting to modify a dashboard that doesn't have permission, can copy the dashboard and make changes to the copy. Copying a dashboard requires permission as well.

2. **Dashboard name:**

    - Use this menu to select the desired dashboard to view.
    - The default dashboard is displayed at the top of this menu.
    - A '*' after the dashboard name indicates that it is read-only.

3. **Configuration:**

    - Add a New Gadget: When viewing a dashboard, this option can be used to add additional gadgets. Select the category of gadgets in the drop down box at the top. To add gadgets, click on them.
    - Copy this Dashboard: Use this option to make a copy of the dashboard which can then be modified by the user. This requires either the "Create New Dashboards" or **Dashboard Admin** permission.
    - Create New Dashboard: If the user belongs to a usergroup that has permissions, this option can be used to create a new dashboard. This requires either the **Create New Dashboards** or **Dashboard Admin** permission.
    - Remove Dashboard: Use this option to remove a dashboard from the menu. Both read-only and user created dashboards can be removed and added back to the menu. This is done under **Configuration > User Dashboards**.

6.2.1 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 gadgetsat 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 viewof 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.

6.2.2 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 gagdet 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) previx to avoid a 404 error. Gadgets may not load if you specify HTTP content when Scrutinizer is using HTTPS.

6.2.3 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.

6.3 User and usergroup permissions

- User Dashboards:

Use this option to select the dashboards a user will have visible in their menu of available dashboards.

- Select the user from the drop down box at the top. To see other users, the user must be a member of a usergroup with the **Dashboard Admin** permission.
- Select the dashboards in the "Available" box and move them to the **Visible** box to grant permission.
- Notice the filter on the left. If granting permission to multiple users, administrators generally use the **Usergroup Dashboards** option. This requires the **Dashboard Admin** permission.
- Usergroup Dashboards:

Use this option to select the dashboards a usergroup will have visible in their menu of available dashboards. This feature requires that the User be a member of a User Group with **Dashboard Admin** permission.

- Select the user group from the dropdown box at the top.
- Select the dashboards in the **Available** box and move them to the **Visible** box to grant permission.

Note: Even if a dashboard is added to the menu for a specific user or for all users in a usergroup, individuals can still remove a dashboard from their menu.

Additional permission options can be found under:

- **Admin tab > Security > Users** to set the default dashboard the user will see when first opening the Dashboard tab.
- **Admin tab > Security > Usergroups**:

1. Choose Dashboard Gadgets
   - Click the "Dashboard Gadgets" value for the usergroup you want to change
   - Uncheck "All Dashboard Gadgets"
   - Move the individual gadgets the selected usergroup should be able to view from the "Deny" to "Allow" box
2. Choose Feature Access
   - Click the **Configure** link in the "*Features*" column for the usergroup you want to change.
   - With **Predefined** selected, add or remove the **Dashboard User** and **Dashboard Administrator** roles.
   - With **Advanced** selected, add or remove individual features like **Create Dashboards**.

6.4 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 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*.

CHAPTER 7

Status

7.1 Overview

7.1.1 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:

- Check Box: Check off the interfaces desired to include in a single report and then click the trend icon at the top of this column.

- Icon color status: The color is determined by *CrossCheck*. Mousing over the icon will provide polling details.

- Flow Version: Clicking on the version of flows received (e.g. N9, N5, I10) opens a report menu for the device which includes a Flow Stats report for the device.

- Interface: Clicking on the Interface will open the Report menu. Selecting a report from here will run an inbound/outbound (bidirectional) report for the last 24 hours in 30 minute intervals. The user can drill down from there.
- Arrow down menu: Clicking this presents a menu:
  - Reset the high watermark(s) in the Inbound/Outbound columns
  - *Interface Details*
  - Device Overview
- Inbound/Outbound: these columns represent utilization over the last 5 minutes. Clicking on them will prompt the user to run a report for the last 5 minutes in 1 minute intervals.

7.1.2 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 opitons.

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.
- Share your desktop: Launches GoToMeeting for remote desktop control.
- 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 reoccuring 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 a more flexible and allows for searching by host address, username, wireless host or SSID across some, or all, flow exporters for a specified timeframe.
- System: These are advanced reports used by engineering when trying to understand why something isn't working. In a future release, they will be moved to the Admin tab.
  - Available Reports: Lists the report and the number of templates received that contain the necessary elements.
  - Flow Report Thresholds: Lists all the reports that have been saved with a threshold.
  - Templates: These display the device and each flow (NetFlow v9/IPFIX) template exported from that device.
  - Vitals: These are reports on the system resources from the flow collection and reporting servers.
- Views:
- *CrossCheck*
- Device Status: Lists all of the flow sending devices with corresponding details. The color of the icons is determined by CrossCheck.
- Interfaces: This is the default view of the status tab before a report is run.
- SLA: Lists all of the flow sending devices which by default are being pinged by the collector. The response from each ping is used to determine the Response times and availability for each device polled.
- Usernames: This view displays any username information collected from exporters such as Cisco ASA, SonicWALL firewalls, or authentication servers such as Active Directory, RADIUS, etc.
- Vendor Specific: lists reports that will work ONLY if the collector is receiving the necessary templates from the flow exporters.

Left hand side menus provide three views:

- Device Expolorer: 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.

7.1.3 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 the group and give it a new name.
- Modify: Modify the membership of the objects in the group.
- CrossCheck: *View CrossCheck* for the devices in this group.
- SLA: View the Service Level Report for the flow exporting devices in this group.
- Show Interfaces: Show all active interfaces for the flow exporting devices in this group. The interface list will display in the main window of this screen.
- Exporters: Devices that are exporting flows show up in the left column. The color of the icon represents the selected primary status for the *object*. The sub icon represents the Fault Index value for the device in *CrossCheck*. Expland the flow exporter for the menu.
- Reports: Run a report on the flows coming from the device. :ref:‘Select a report <network_traffic_reporting>‘to display flow data. Selecting a report from here will run the report for ALL interfaces of the device resulting in the inbound traffic matching the outbound traffic. For this reason, this report is displayed inbound by default.The default timeframe for this report is Last 24 hours in 30 minute intervals.

- Interfaces: Displays a list of interfaces for the device. Click on an interface to run a report. Selecting an interface (or All Interfaces) from this list will open a report menu. Select and run a report for the last 24 hours. ALL Interfaces reports will default to Inbound as described above, selecting a single interface will report on both Inbound and Outbound.
- Properties: Modify the properties of the device.
- Device Overview: Provides the overall status of the device by leveraging data from CrossCheck, the poller and the alarms.
- Show Interfaces: Displays a list of all active interfaces for the device in the main window of the page.
- **Other Options**:
  - Alarms: Displays the outstanding alarms for the device.
  - Interface Details: launches the *Interface Details* view which lists SNMP details about the device including the interface speeds.
  - Flow Templates (Advanced): displays the templates (e.g. NetFlow v9, IPFIX, etc.) currently being received from the device.

7.1.4 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 a 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.

7.2 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.

7.2.1 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.

7.2.2 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)

7.2.3 Gear icon

Cliking 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 with be the SUM of all interface speeds. Inbound is calculated separately from outbound. The summed port speed is used for percent calculations. All interfaces are required to have a defined speed for percentage reports. If 'Percent' is selected in the drop down box, it represents the overall percent of the entire interface. The preceding percent column that can't be changed represents the percent of the overall bandwidth consumed.
- **Show Peak:** If 'Yes' is selected, a Peak column is added to the report. Peak values are the highest data point in the graph in the same interval the graph is reporting in.
- **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.

7.2.4 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.

7.2.5 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.

7.2.6 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. Fo**

        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 predefined list of items, a dropdown list will appear to select from, otherwise a textbox will be displayed for entering the filter data. If Source or Destination are applicable, another dropdown selector will appear for selecting Source, Destination, or Both. If it is a calculated column, a dropdown selector of numerical comparisons will appear.

3. The next option is to select whether this will be an Include or an Exclude filter. Include filters will only display flow data where the filter criteria is equal. Exclude filters will display everything except the filtering criteria.

4. When all options are completed, the Add Filter button will appear, allowing the new filter to be added to the existing filters. After adding the new filter, the Update Report button displays and clicking that button is the last step to 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 they 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<F5> 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.

7.2.7 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)

7.3 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 varients 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 WellKnownPort).

For example, take a flow with a source port of 5678 and a destination port of 1234. The collector will look at both ports (5678, 1234) and perform the following logic:

- Which port is lower: 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 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.

7.4 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.

7.5 Scheduling a report

Prerequisites

- The email server needs to be configured in *Admin > Settings > Email Server*
- One or more report(s) need to have been 'saved'

Schedule reports from the Status tab

- Either create a new saved report or select existing Saved Reports from left pane in Status tab, then select the saved report(s) from the list. Make sure the report is saved with a 'last' time frame (E.g. Last Seven Days).
- Click the 'clock' icon to Schedule an emailed report. It can be found at the top under Current report. It will launch the Schedule Report modal.
- Schedule Report modal
    - Email Subject: This field is mandatory and is auto filled with the Report name when coming from the Status tab. The subject of the email can be changed here.
    - PDF / CSV: Check these boxes to attach the report in PDF or CSV format.
    - Frequency and Time: This report will kick off on the current day:
        - Hourly: Specify the minute each hour that report(s) will run
        - Daily: Specify the hour, minute, and AM/PM that report(s) will run each day
        - Weekly: Specify the hour, minute, AM/PM, and day of week that report will run each week
        - Monthly: Specify the hour, minute, AM/PM, and day of month that report will run each month
    - Recipients: Enter the email address(es) of recipients here.This field is mandatory and must include at least one recipient's email address. Multiple email addresses may be separated by commas, semi-colons, or spaces, and may be entered all on one line, or on separate lines.

- Include/Exclude: This section shows which reports are in the scheduled report (Included) and which ones are not, but are available to add to this schedule (Excluded). At least one report must be in the Include section. By default, when scheduling from the Status tab, the saved report being viewed will be automatically included. Add more reports to a scheduled report by selecting from the Exclude list and clicking the double left arrows (<<) to move it to the Include list.
- Click 'Save' to add any selections to the Scheduled Report list.
- To monitor and manage the Scheduled Report go to Admin > Reports > Scheduled Reports.

Important: Make sure the report is saved with a 'last' time frame (E.g. Last Seven Days). If the frequency is set to 'Hourly' for example, a report will be emailed every hour which shows the last seven days. Also, in order to avoid excessive processing overhead, try to avoid scheduling multiple reports to run at the same time.

Managing scheduled reports

Scheduled reports can be managed at: *Admin > Reports > Scheduled Email Reports*. This page will list all existing scheduled reports. Columns in this page include:

- Action
- Edit Schedule – opens the Schedule Report modal allowing changes to any aspect of the scheduled report.
- Send Now – email this report on-demand
- Disable - checkbox
- Email Subject
- Schedule
  - Hourly
  - Daily
  - Weekly
  - Monthly

- Time – scheduled time for report
- Day of Week – scheduled day for report
- Day # - scheduled day of month for report
- Execute Time – the amount of time taken the last time the scheduled report has run
- Last Sent – time stamp for last time the scheduled report has run
- Recipients – email addresses configured to receive this scheduled report

Note: Email Subject and Included report 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 process. 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 schedulign 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.

7.6 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 Plixer 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.

7.7 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.

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.

7.8 User name 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.

7.9 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.

7.10 CrossCheck & Service Level reports

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

CrossCheck logic

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

After each 3rd Party Method completes, CrossCheck (i.e. mapping.cgi) updates the FI for each device. It queries the above list for any device that has violated theconfigurable FI threshold. For each device that exceeded the threshold,CrossCheck sends a syslog to the *alarm server* which will violate the "Exceeded CrossCheck Fault Index" policy. Any hosts that are up are removed from the xcheck_notifications table (see below).

Fault Index

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

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

CrossCheck table

The main CrossCheck table displays an inventory of all hosts being monitored as well as their status in each of the applications that are monitoring the device. The overall Fault Index (FI) on the far right provides the status of devices in the Scrutinizer maps and all of the 3rd party applications monitoring them. Click on the headings to sort. The query time frame is the last 1 minute by default. The buttons are explained below.

- CrossCheck Summary: Explained below.
- Thresholds: This sets the threshold at which the sub icon on a device changes color. Color thresholds are based on a percentage of the Fault Index threshold.
- 3rd Party Methods: Data collected from other applications that will impact the Fault Index (FI) for a device.

CrossCheck action:

- Device Overview
- Edit: Polling and Appearance: This launches the modify object view in the mapping utility.
- Run Report: Flow Report: This runs the flow volume report for the device.
- Host: List the IP address or DNS host name for the device. Mouse over for tool tip.
- Flow: Method used to determine if a device is sending flows.
- Poller: Method used to determine if the poller can ping the device.

- Optional Third Party Methods: Contact your vendor to create new methods (e.g PRTG, Solarwinds, etc.). When active, a new column appears.
- Fault Index (FI): The FI provides the overall status of a device across all 3rd Party Methods.

CrossCheck summary

The CrossCheck summary is an bar chart of the CrossCheck list. It displays the number of unique hosts found across all 3rd Party Methods as well as the Fault Index.

- Refresh: Set the interval. The default is every 5 minutes.
- Thresholds: This sets the threshold at which the sub icon on a device changes color. Color thresholds are based on percentage of the Fault Index threshold.
- Define Networks: Specify a subnet to group IP addresses found across all 3rd Party Methods.
- Overview: Provides a small window with the number of devices discovered in each 3rd Party Method as well as the overall total FI for the devices in the application.

Service Level Report

The Service Level Report is a dual purpose report listing device availability and response time. In the Service Level Report:

- Click on the headings to sort.
- The number to display (e.g. 25) followed by pagination.
- Click on the Poller to run a trend across all devices.
- Click on the Device to run an availability or response time report on the selected device.
- Click on the round trip time column value or the availability percent value to run a report.

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

7.11 Vitals reporting

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

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

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

- **% CPU per Process:** This report displays CPU percentage consumed per process on the server.
- **CPU:** Average CPU utilization for the Scrutinizer server(s).
- **CrossCheck Runtime:** Monitors runtimes for CrossCheck methods (processes).
- **Database:** Provides the following database metrics:
  - **Connections by Bytes:** Excessive connections can result in reduced performance. NOTE: other applications using the same database will cause this number to increase.
  - **Read Req:** The number of requests to read a key block from the cache. A high number requested means the server is busy.
  - **Write Req:** The number of requests to write a key block to the cache. A high number of requests means the server is busy.
  - **Cache Free:** The total amount of memory available to query caching. Contact 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.

CHAPTER 8

Maps

8.1 Overview

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*

8.1.1 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 locationed 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.

8.1.2 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

8.2 Groups

Groups are the foundation of all maps. Creating a new group creates a map. Flow sending devices that are not assigned to a map are placed in Ungrouped. There are two types of maps:

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

Highlights:

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

8.3 Objects

Objects come in four formats:

- **Devices:** are imported from *CrossCheck* or can be manually added. These objects change color based on the Fault Index and the threshold settings in CrossCheck. To remove a device that was imported from CrossCheck, the 3rd Party Method must be disabled or the device must be removed from the 3rd Party Method script or removed from the 3rd party application, else it will continue to be re-imported after deletion.
  - Label: If the device is imported from CrossCheck, this value is imported.
  - Poll Using: Select IP Address, Hostname or Disable Polling.
  - Notification: Select a Notification Profile which will be triggered when the *CrossCheck* Fault Index threshold is breached.
  - Icon: The default icon type and size can be modified.

- **Groups:** represent other maps and the status of devices in those maps. They are clickable and bring up the appropriate map.
- **Symbols:** represent devices in the maps that don't display a status. They can be assigned labels and made clickable to launch other applications and/or web pages.
- **Text Boxes:** can be placed on maps and generally contain text. Shapes, colors and size can all be defined. As well as the Label and a clickable link. Text boxes are for Plixer maps only. They cannot be placed on Google maps.

Note: To modify the Google address of an object, select a map the object is 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.i

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

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

- **Icon image:** shape of the icon
- Dependencies: are used to determine how and when the device is polled.
- **Membership:** Specify the groups / maps the object is a member of.

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

8.4 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.

8.5 Creating Plixer maps

When creating a Plixer map, the user is presented with the following options:

- **Settings**
  - **Name:** The name given to the map. It can be changed later.
  - **Pass Status:** If some maps are intended to be submaps, the status can be passed to another 'Parent' map that contains an icon representing the lower map. This in effect will cause the icon color status of the Parent to change. The status of an icon is determined by *CrossCheck* which considers multiple factors.
  - **Auto-add Devices (RegEx):** This option is used to add similar devices quickly using regular expressions. For example, if a number of IP addresses resolve to host names that all contain the text 'company.local', this can be entered here. When Save is clicked, all devices that resolve to a host name containing this text will automatically be added to the map.
  - **Truncate Map Labels on:** Sometimes the icon labels can contain excessive amounts of text. Often times, a portion of the trailing text on each icon can be omitted. Enter the text here that shouldn't be displayed.

- **Objects**
  - **Add/Remove objects:** Use this window to move Available objects from the right side to the Members section on the left. Multiple objects can be selected by holding down the shift or CTRL key. Use the filter on the left to quickly locate objects. The search can be performed by IP address or host name by clicking on the button below the filter.
  - **New:** Non-flow sending objects can be added to the maps. IP addresses are optional (E.g. text box).

Form fields for **Object Type > Icon** are:

- **Icon:** Use the arrow keys on the key board to scroll through the different icon options.
- **Label:** This names the object and displays in the maps and groups listings.
- **IP Address:** By default, the optional IP address is polled every 60 seconds.
- **Primary Status:** The Primary Status indicator is the largest colored portion of the icon.
- **Link:** The web site that is launched when the icon is clicked in the map.
- **Additional notes:** Help the user understand what the object represents.

Form fields for **Object Type > Text Box** are:

- **Label:** Name of the object, displays in the maps and groups listings.
- **Shape:** Select Rectangle, Circle, Polygon
- # of Sides: (Polygon only) Select from 3 - 10 sides for the polygon shape.
- **Height (px) / Width (px):** (Rectangle only) Define the height and width of the rectangle in pixels.
- **Radius (px): (Circle and Polygon):** Define the size of the shape in pixels.
- **Color:** Click on the box to open the color palette to choose the Text Box color.
- **Type:** Choose from Text or Background text box type.
- **Link:** The web site that is launched when the text box is clicked in the map.
- **Connections**

Connections change color based on the utilization settings found in **Admin tab > Settings > System Preferences** and require that an interface speed was collected from, or defined for, the device. Without an interface speed, the connection will stay blue. The arrow on a connection represents the highest flow direction in the last five minutes.

When a map is in view mode, clicking on a link will launch a report showing the last 5 minutes in the highest utilized direction (I.e. inbound or outbound). Connections using a Saved Report are based on the inbound value for that report.

- **Connections:** Click this button to list all of the configured connections with options to either delete or edit a connection. i
- **Create:** Links between devices and objects can be connected using interfaces, saved filters, or a simple black line.

    1. Select a 'From' device
    2. The Type of connection
    3. Fill out the additional options
    4. Then select the 'To' device or object
    5. Click Save
    6. Continue this process to represent the major connections on the network.

- **Background**

There are multiple options to represent the background of a map:

- Existing Map: select a map image from the dropdown selection list that is provided.
- Set background color: click the color in the square to select from the color pallet.
- Upload: create a custom background by transferring an image file to the Scrutinizer server. You can copy the new images directly to the *home/plixer/scrutinizer/files/map_backgrounds* directory.

Maps with background images autoscale to the size of the image. Very light, grayscale backgrounds are ideal as they allow the status of the icons to be visible. The images can be in .gif, .jpg, or .png format. The image size should be at least 800x600 pixels to allow room for icon positioning. Maps with background images autoscale to the size of the background image.

Important: By default, the maximum file size is 5 MB (5000000). You can adjust the setting as well as disable file uploads via the **Admin>Settings>System Preferences** page. The application will discard values below the minimum file size of 200KB (200000).

Laying out the Plixer map

After a new Plixer map has been created and objects added, the objects will be all clustered in the upper left hand corner. To start arranging the icons, the user must enter Edit Mode. When finished editing the map, the user should return to View Mode. Select a blank area on the map and click the RIGHT mouse button, then in the menu, select “Edit Mode”.

- **Gear Menu:** Use these options to set the refresh rate, to display either the IP Address or hostname on the icons, and to reset the zoom level.
- **Edit Mode:** Right click anywhere in the map and select Edit Mode to enter this mode.

The Edit Mode status is then clearly indicated at the top left of the map.In this mode, the icons can be selected with the mouse and dragged to different areas of the map for custom arrangement. Click the right mouse button and notice that several new options present themselves in the Mapping Menu.

- **Align:** (Applies to a group of selecte 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.

8.6 Creating Google maps

To set up the first Google map, an API key has to be generated. To apply a key, navigate to the **Admin > Maps and Device Groups > Global Settings” and paste it into the **Google Maps - Browser API Key** box.

Note: Google TLD” defaults to .com and should be changed if the install is located in a country that defaults to a top level domain other than .com. For example, in the U.K. change it to .co.uk

Modifying the Google maps involves launching most of the same options found in a *Plixer map*. There are a few exceptions such as no RIGHT mouse button menu which is reserved by Google for zooming out of the map.

Click on an icon with the LEFT mouse button to launch the menu with the following options:

- **Device Overview:** Launches the device overview including this information.

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

- **Create a connection** works as outlined in the *Plixer Maps* section.

- **GPS Location:**

  1. Placing a device in a specific location requires entering either a physical address or the GPS coordinates. Simply specifying a city in a country will also work.
  2. After entering an address, click (Resolve GPS) to ensure the address is resolved to the new GPS coordinates.
  3. Click Save.

- **Properties** work as outlined in the *Plixer Maps* section.

CHAPTER 9

Flow Analytics

9.1 Overview

Flow Analytics (i.e. FA) brings the following additional features to Scrutinizer:

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

Navigation

The navigation for FA is via gadgets in the Dashboard tab. The primary gadget "Flow Analytics Configuration" should be added to Dashboard. It can also be reached by navigating to Admin tab > Settings > Flow Analytics Configuration. Below are the primary utilities for configuring and observing the performance of Flow Analytics.

- *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 Dashboard Gadgets*: Used to visualize the results of the FA algorithms.
- Flow Analytics Settings: are explained at **Admin tab > Settings > Flow Analytics Settings**.

Aggregated alarms

Aggregated alarms combine alarms from events that are continuous (may last several hours) into a single alarm, displaying the original alarm time, the most recent alarm time, and the number of times that the alarm has triggered while it was active. Aggregated alarms will continue to collect matching alarm events until the **Aggregated Alarm Timeout** has expired with no new alarms. The **Aggregated Alarm Timeout** defaults to two hours (120 minutes). This value is controlled on the **Admin > Settings > Flow Analytics Configuration** page and is also configurable per algorithm. A value of zero will disable the alarm aggregation.

Security event algorithms

The security event FA algorithms provide alarms that are focused on providing actionable information while significantly minimizing false positives. IP Addresses are classified in two ways:

1. Internal IPs: are the IP addresses, or assets, that comprise your network.
2. External IPs: are the public Internet and other IP address spaces that are not under your administration.

Alarm messages identify both the source and destination of suspect activity as "external to internal", "internal to external", or "internal to internal", as well as providing additional details that are specific to the alarm type. Internal Addresses are defined as any addresses entered as your IP Groups plus the following list of IP Address blocks. These addresses are private, non-routable addresses per RFC 1918 and link-local addresses per RFC 3927:

> 10.0.0.0/8
> 172.16.0.0/12
> 192.168.0.0/16
> 169.254.0.0/16

All other addresses are treated as external IPs.

9.2 Algorithms and gadgets

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

9.2.1 FA gadgets that can be added to dashboards

- **Analytics Violation Overview:** Top Flow Analytics policy violation summary with violations counts for the Last 5 minutes, Last Hour, and Overall time.
- **Flow Analytics Summary:** The overall status of all algorithms and the total runtime and count of violations across all algorithms. Algorithms can be ordered alphabetically or by order of execution. The *FA Configuration* page can be opened by clicking the button at the top left of this gadget.
- **Flow Reports Thresholds:** Saved reports that are given a threshold to compare against every five minutes show up in this gadget.
- **Medianet Jitter Violation:** Jitter values as reported by the Medianet flows that exceed the threshold defined in this algorithm. The default threshold is 80ms.
- **Network Volume:** The scale of the traffic traversing through the core network. It lists the volume of unique traffic on the network for the last 5 minute vs. last 30 hours. Only include a few core routers/switches in this algorithm.
- **Policies Violated:** Last 24 hour report of top alarm policies violated with violations counts.
- **Recent Alarms/Recent Alarms by Violator:** Violations listing by policy and violator, with threat heat maps included.
- **Threats:** Summary report of top Flow Analytics algorithm violations.
- **Threat Index:** Last 24 hour report of top violators by threat index values.
- **Top Applications:** Top Applications across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Conversations:** Top Conversations across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.

- **Top Countries:** Top Countries across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Flows:** Top Flow sending end systems across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Hosts:** Top Hosts sending data across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Network Transports:** Top Transport Layer Protocols across selected flow exporting devices. Alarms trigger for protocols that appear that haven't been approved. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Rev 2nd lvl Domains:** Top reverse 2nd level domains across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Subnet Traffic:** Top IP Subnets across selected flow exporting devices. This algorithm is on by default across all flow exporting devices that are exporting the necessary fields.
- **Top Violators:** Last 24 hour report of top alarm violators.

Note: Some gadgets include algorithms that should only be run against core routers/switches. Watch the Flow Analytics Summary gadget for algorithms that are taking an excessive amount of time to run. Everything needs to finish in under 5 minutes (300 seconds).

9.2.2 FA algorithms that don't include gadgets

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

- **Bogon Traffic:** This algorithm alerts if traffic to or from unallocated public IP space is detected.
- **Breach Attempt Detection:** This algorithm is examining flow behaviors that may indicate a brute force password attack on an internal IP address. This is accomplished by examining the flow, byte, and packet counts being exchanged in short-duration completed flows between one source and one destination. Specific behaviors are observed for common attack vectors such as SSH, LDAP and RDP. If the number of flows that match these characteristics exceeds the alarm threshold, an alarm will be raised. The default flow count threshold is 100. Either IP address can be excluded from triggering this alarm. This algorithm is enabled by default across all flow exporting devices that are exporting the necessary fields.

- **DDoS Detection:** Identifies a Distributed Denial of Service attack targeting the protected network space. DDoS attacks are often launched by a BotNet, and *"reflection attacks"* are becoming more common.

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

- **DDoS Packet Deviation** (default: 10) and **DDoS Bytes Deviation** (default: 10): These settings control how similar the flows associated with the attack must be. The standard deviation of the byte count and packet counts associated with the flows must be less than this setting.
- **DDoS Flows** (default: 4) controls the minimum number flows used to identify attacking hosts. The sensitivity of the DDoS attack can be reduced by increasing this setting to six or higher.
- **DDoS Unique Hosts** (default: 200) controls the threshold for the minimum number of hosts that have sent flows that match the other characteristics required to trigger the alarm.
- **Denied Flows Firewall:** Triggers an alarm for internal IP addresses sending to external IP addresses that cause greater than the threshold of denied flows. The default threshold is set to 5 denied flows. Either the source or destination IP address can be excluded from triggering this alarm.
- **DNS Hits:** Triggers an alarm when a host initiates an excessive number of DNS queries. This identifies hosts that perform an inordinate number of DNS lookups. To do this, set the flow threshold to a large value that reflects normal behavior on the network. The default threshold is 2500 DNS flows in five minutes. Either the source or destination IP address can be excluded from triggering this alarm.
- **DRDoS Detection:** Identifies a Distributed Reflection Denial of Service attack targeted at the protected network space. DRDoS attacks are often launched by a BotNet, and "reflection attacks" are becoming more common. Scrutinizer may identify attacks against the network as "reflection attacks" if they meet the following criteria.

  Scrutinizer detects the following ten **Distributed Reflection Denial of Service (DRDoS) attacks**:

  - DNS
  - NTP
  - SNMP
  - SSDP
  - Chargen
  - NetBIOS Name Server

  - RPC Portmap
  - Sentinel
  - Quote of the Day
  - Trivial File Transfer Protocol

There is an option to enable or disable a specific reflection attack via **Admin > Settings > Flow Analytics Configuration > DRDoS Detection > Settings.**

- **FIN Scan:** Alerts when a FIN scan is detected. FIN scans are often used as reconnaissance prior to an attack. They are considered to be a "stealthy scan" as they may be able to pass through firewalls, allowing an attacker to identify additional information about hosts on the network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan the internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.
- **Host Reputation:** This algorithm maintains a current list of active Tor nodes that should be monitored. Some malware families use Tor for Command and Control communications. White-list users who are authorized to use Tor and regard other uses as suspicious. This algorithm will also monitor any IP address lists that can be provided as custom lists.

The Host Reputation algorithm also supports the use of custom lists where the user can add additional reputation lists to the system. A custom list of IP addresses can be imported into the Host Reputation algorithm. To do this, Host Reputation needs to be enabled and this list will be compared with traffic on the devices selected under **Configured Flow Analytics**.

To enable Host Reputation:

- Go to **Admin > Settings > FA Configuration**.
- Expand the **Host Reputation Monitor**.
- Make sure Disabled is not checked and exporters are being included.

To create custom lists:

- The IP Addresses need to be in a file with a single address on each line.
- The name of the file will be the Threat Category Name and the Alarm Policy Name.
- The file must have a .import file extension; for example: "CustomThreatList.import".

- The file must be placed in the */home/plixer/scrutinizer/files/threats/* directory. Once an hour, this file will be imported into Scrutinizer and used for the next hour of processing.
- To force a new file import to become active immediately, run:
  - */home/plixer/scrutinizer/bin/scrut_util*
  - SCRUTINIZER> *download hostreputationlists*
- After the import, the Alarms Policy can be modified to change the threat_multiplier from the default of 0.
- **Host Watchlist:** Identifies hosts that have violated an internal host watchlist.

To enable Host Watchlist:

- Go to **Admin > Settings > FA Configuration**.
- Expand the **Host Watchlist** algorithm.
- Make sure Disabled is not checked and exporters you would like to monitor are included.

To add custom lists:

- Create a csv file that contains a set of blacklisted IPs or CIDRs, one IP or CIDR per line in dotted format.
- To trigger alarms for a blacklisted host generating traffic on a specific protocol/port, add a port/protocol to the csv row.

Valid examples:

10.0.1.100,6

10.0.1.100,6,22

10.0.1.0/24,6

10.0.1.0/24,6,22

Invalid example:

10.0.1.100,,22

- Place the file into the */home/plixer/scrutinizer/files/watchlist/* directory. Once an hour, it will be imported into Plixer Scrutinizer and used for the next hour of processing.
- To force a new file import to become active immediately, run:
  - */home/plixer/scrutinizer/bin/scrut_util*
  - SCRUTINIZER> *download hostreputationlists*

Note: You can specify a protocol without a port, however specifying a port requires a protocol.

- **ICMP Destination Unreachable:** This alarm is generated when a large number of ICMP destination unreachable messages have been sent to the suspect IP address. This may happen as a result of scanning activity, misconfiguration, or network errors. ICMP Destination Unreachable is a message that comes back from a destination host or the destination host gateway to indicate that the destination is unreachable for one reason or another. The default threshold is 100 destination unreachable messages. Either the source or destination IP address can be excluded from triggering this alarm.
- **ICMP Port Unreachable:** This alarm is generated when a large number of ICMP port unreachable messages have been sent to the suspect IP address. This may happen as a result of scanning activity, misconfiguration, or network errors. ICMP Port Unreachable is a message that comes back from the destination host gateway to indicate that the destination port is unavailable for the transport protocol. The default threshold is 100 port unreachable messages. Either the source or destination IP address can be excluded from triggering this alarm.
- **Indicator Correlation:** This algorithm escalates multiple Indicator of Compromise (IOC) and security events for a single host to a new alarm on the security event BB. While a single IOC may be indicative of malware, it is much more likely to be a real security concern if there are multiple indicators. By default, This algorithm correlates multiple IOCs along with any events posted to the Security Event BB and issues an alarm for any host that has three or more entries in the IOC and Security Event bulletin boards. Each of the contributing algorithms will be listed in the alarm message. By default, three different algorithms are required, the threshold setting for Indicator Correlation can be adjusted.
- **IP Address Violation:** By default, this algorithm allows all subnets. Once subnets are defined, any flow that contains an IP address where either the source or destination IP address isn't in an allowed subnet, an event will trigger. In other words, if in a single flow both source and destination IP addresses are outside of the allowed subnets, an event will be triggered. A common use of this algorithm is to identify unknown or unauthorized internal network addresses that are communicating with the Internet.

- **Large Ping:** Alerts when unusually large ICMP Echo Request (ping) packets are observed. This alert could indicate malicious activity within the network, including possible Denial of Service (DoS) attempts.

- **Medianet Jitter Violations:** This algorithm compares the jitter values as reported by the Medianet flows to the threshold defined by the user in the Settings section of this algorithm. The default threshold is 80ms.

- **Multicast Violations:** Any multicast traffic that exceeds the threshold that isn't excluded will violate this algorithm. The default threshold is 1,000,000 and the minimum that can be set is 100,000.

- **NetFlow Domain Reputation:** Checks DNS lookups exported in NetFlow (Gigamon, Allegro, FlowMon) against a blacklist maintained on nba.plixer.com and cached locally. Upon observance of a domain lookup for a blacklisted IP, an alert is generated.

- **NULL Scan:** Alerts when a NULL scan is detected. NULL scans are a TCP scan with all TCP Flags cleared to zero. This scan is sometimes used as a reconnaissance tactic prior to an attack and is considered to be stealthy because often times it is able to pass through firewalls. Eluding a firewall makes it easier for an attacker to identify additional information about the hosts on the network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan the internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.

- **Odd TCP Flags Scan:** Alerts when a scan is detected using unusual TCP Flag combinations. These types of scans may allow an attacker to identify additional information about hosts on the network. The default threshold is 100 unique scan (aka flows) in five minutes. Internal IP addresses that are allowed to scan the internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.

- **P2P Detection:** Peer to Peer (P2P) traffic such as BitTorrent are identified by this algorithm. The default threshold is a P2P session involving over 100 external hosts, which will detect most P2P applications. However, there are several P2P applications that are stealthier. Experimenting with lower thresholds or periodically lowering the threshold to about 20 will allow the security admins to determine if other "low and slow" P2P traffic is on the network.

- **Packet Flood:** 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.

- **Persistent Flow Risk:** Alerts when a persistent flow is detected. Persistent flows are a strong indicator of VPNs, proxy traffic, remote desktop technologies, or other means of covert communication. The default threshold for a flow to be considered persistent is 12 hours. In addition to the temporal threshold an optional ratio threshold is available to identify the relationship of traffic as it pertains to ingressing or egressing the network. The default PCR threshold is set to .9, identifying persistent flows where the ratio indicates more traffic is destined outside the network.

- **Persistent Flow Risk - ASA:** IP communication matching a 5-tuple (external IP and port, internal IP and port, and common port) up for 12 hours or longer. The duration can be adjusted. This algorithm can identify VPNs or proxy traffic, remote desktop technologies, and other means of covert communication across various applications.

- **Ping Flood:** 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.

- **Ping Scan:** 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.

- **Protocol Misdirection:** Identifies when the type of traffic doesn't match the port being used.

- **Reverse SSH Shell:** Identifies possible reverse SSH tunnels to external destinations. A reverse SSH tunnel allows external entity access to internal, protected resources via the use of an established outbound SSH connection.

- **RST/ACK Detection:** Alerts when a large number of TCP flows containing only RST and ACK flags have been detected that are sending 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 the hosts on the network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan the 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.

- **Slow Port Scan:** 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.

- **Source Equals Destination:** Alerts when traffic that has the same source and destination addresses is observed. This alarm commonly occurs due to misconfigurations within a network, but may also indicate possible malicious activity.

- **SYN Scan:** Alerts when a SYN scan is detected. SYN scans are a TCP scan with the TCP SYN flag set. This scan is often used as reconnaissance prior to an attack as it is fast and somewhat stealthy. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan the internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.
- **TCP Scan:** Alerts when a possible TCP scan is detected from an exporter that does not provide TCP flag information. These types of scans may allow an attacker to identify additional information about hosts on the network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan the internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.
- **UDP Scan:** Alerts when a possible UDP scan is detected. These types of scans may allow an attacker to identify additional information about hosts on the network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan the 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 company policy allows P2P traffic on the network, then it is unwise to enable this alarm as it will often detect P2P control traffic as a UDP Scan violation.

- **Worm Attack:** Identifies possible worm behavior from a host. Worms are malicious software that replicates across hosts and can lead to further security risks, including data loss and botnet activity.
- **Worm Propagation:** Identifies when a worm has successfully replicated across hosts.
- **XMAS Scan:** Alerts when a XMAS scan is detected. XMAS scans are a TCP scan with the FIN, PSH, and URG TCP flags set. This scan is often used as reconnaissance prior to an attack. They are considered to be a "stealthy scan" as they may be able to pass through firewalls, allowing an attacker to identify additional information about hosts on the network. The default threshold is 100 unique scan flows in five minutes. Internal IP addresses that are allowed to scan the internal network, such as security team members and vulnerability scanners, should be entered into the IP exclusions list. Either the source or destination IP address can be excluded from triggering this alarm.

Note: By default, all of the scan algorithms are looking for "internal to internal" and "internal to external" scanning activity. Security admins have the option to control which scanning directions the different algorithms look for, including "external to internal", which would normally be used to monitor public facing IP addresses listed in an IP Group. Within each of the scanning algorithms, the settings screen provides a directional control option.

9.2.3 FA algorithms that require FlowPro Defender

BotNet Detection

(Formerly named NXDomain detection)

This alarm is generated when a large number of unique DNS name lookups have failed. When a 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 failure (NXDOMAIN) messages in five minutes. Either the source or destination IP address can be excluded from triggering this alarm.

DNS Command and Control

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

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

To suppress alarms from authorized applications in the network, the domain generating the alarm message can be added to to the "Trusted Domain" list on FlowPro Defender. See the *Trusted Domain List* discussion below.

DNS Data Leak

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

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

DNS Server Detection

The algorithm 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.

Domain Reputation

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 several hundred thousand known bad domains.

FlowPro Defender performs the actual monitoring, and when it detects a domain with a 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 which means that all DNS lookups observed will result in a unique alarm.

To suppress alarms from authorized applications in the network, the domain generating the alarm message can be added to the “Trusted Domain” list on FlowPro Defender. See the *Trusted Domain List* discussion below.

JA3 Fingerprinting

JA3 fingerprinting functionality leverages a TLS handshake’s unique characteristics to identify the software generating encrypted traffic by comparing it against a list of known signatures. If a positive match is made, Plixer FlowPro Defender will send the details of that connection to Plixer Scrutinizer.

Important: To effectively detect security threats, configure Plixer FlowPro Appliance to monitor external interfaces. Plixer FlowPro Defender licensing is required.

Contact technical support for assistance with configuration.

Malware Behavior Detection

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

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

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

Malware Domain Communications

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

1. Defender contains a list, updated every 10 minutes, of several hundred thousand known malware domains created by forensic analysis of the actual malware. These are very high confidence domains.
2. Defender monitors all of the DNS resolution requests, and generates an IOC (Indicator of Compromise) alert on detection of a match with a malware domain and saves the resolved "Malware IP Address". This only rates an "IOC" as a browser may "prefetch", or resolve an address, for all of the links on a web page. Browsers like Chrome do this to make the browsing experience feel faster. However, as yet, no connection to the malware site has been made.
3. Scrutinizer then examines all flows for any communications with the "Malware IP Address" resolved by Defender. On detection of any flows to or from that address, a connection to the malware site has been established, and a Malware Domain Communication alert is triggered.

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

9.2.4 Correlation algorithms

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

Below are the Correlation Algorithms available in Flow Analytics:

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

9.2.5 Trusted domain list

A "trusted domain list", often called a whitelist, is preconfigured on FlowPro Defender to suppress alarms involving specific domains. The default whitelist contains five entries. Add or remove entries as necessary to best fit the local environment.

- mcafee.com
- sophos.com
- sophosxl.net
- webcfs03.com
- apple.com

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

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

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

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

It is possible to have authorized software within the local networks that abuse the DNS to bypass firewalls for data communications. If this is the case, add the domain(s) involved with the software to the Trusted Domain list as described below. Once they have been configured for the local environment, any other traffic using DNS to communicate will be worth additional investigation.

To modify the trusted domain list:

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

9.2.6 Untrusted domain lists

FlowPro Defender supports both the use of a domain reputation list that is downloaded from Plixer, as well as allowing for the addition of a unique list.

Plixer domain reputation list

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

1. Log on to the FlowPro Defender
2. Enter: “edit plixer.ini”
3. To enable the list (default is enabled), set the value enableDomainReputationList=1

4. or, to disable the list, set the value enableDomainReputationList=0
5. enter control-x, and select “Y” to save changes
6. quit

User defined domain lists

The Plixer Domain Reputation list can be augmented by creating one or more lists that contain domains that the system should alarm on. The rules for the domain lists are:

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

To create a list of domains to detect domainReputation violations:

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

To enable a domain list:

1. Log into the FlowPro Defender appliance
2. show domainlists
3. Enter: “enable domain_list_name”
4. quit

To disable a domain list:

1. Log into the FlowPro Defender appliance.
2. show domainlists
3. Enter: “disable domain_list_name”
4. quit

9.2.7 Machine learning algorithms

Anomalous Behavior

Super set of Plixer Network Intelligence and Plixer Security Intelligence anomaly events.

> **Plixer Network Intelligence Anomaly**
>
> Anomalous behavior detected for an interface being monitored by Plixer Network Intelligence.
>
> **Plixer Security Intelligence Anomaly**
>
> Anomalous behavior detected for a host being monitored by Plixer Security Intelligence.

Important: All of the detections below are initiated by a host exhibiting anomalous behavior.

Brute-force Attempts

SSH and RDP login behavior is analysed for usage patterns that indicate brute force attempts to gain access.

Data Accumulation

This algorithm is characterized by a significant majority of data ingressing from a single host on the private network space defined by IP Groups.

Data Loss

Data Loss is characterized by a significant majority of data egressing the private network space defined by IP Groups.

DNS Tunneling

DNS Tunneling is characterized by a majority of data egressing the violator IP to a single destination IP over DNS.

ICMP Tunneling

ICMP Tunneling is characterized by a majority of data egressing the violator IP to a single destination IP over ICMP.

Worm Activity

The algorithm identifies worms attempting lateral movement via connections on specific ports to various hosts.

9.3 Algorithm activation strategy

Algorithm name	Internal / core routers	Edge routers	Public IP addresses defined in I
Bogon Traffic	No	Yes	Yes
BotNet Detection	FlowPro Defender	FlowPro Defender	N/A
Breach Attempt Detection	Yes	Yes	Yes
DDoS Detection	No	Yes	Yes
Denied Flows Firewall	Yes	No	No
DNS Command and Control	FlowPro Defender	FlowPro Defender	N/A
DNS Data Link	FlowPro Defender	FlowPro Defender	N/A
DNS Hits	Yes	Yes	Yes
Domain Reputation	FlowPro Defender	FlowPro Defender	N/A
DRDoS Detection	No	Yes	Yes
FIN Scan	Yes	Yes	No
Host Reputation	No	Yes	Yes
Host Watchlist	No	Yes	Yes
ICMP Destination Unreachable	Yes	No	No
ICMP Port Unreachable	Yes	No	No
IP Address Violations	Yes	Yes	Yes
JA3 Fingerprinting	FlowPro Defender	FlowPro Defender	N/A
Malware Behavior Detection	FlowPro Defender	FlowPro Defender	N/A
Multicast Violations	Yes	Yes	Yes
Large Ping	Yes	Yes	Yes
NetFlow Domain Reputation	Yes	Yes	Yes
Null Scan	Yes	Yes	No
Odd TCP Flags Scan	Yes	Yes	No
Packet Flood	Yes	Yes	Yes
Persistent Flow Risk	Yes	Yes	No
P2P Detection	Yes	Yes	No
Ping Flood	Yes	Yes	Yes
Ping Scan	Yes	Yes	Yes
Protocol Misdirection	Yes	Yes	Yes
Reverse SSH Shell	Yes	Yes	Yes
RST/ACK Detection	Yes	Yes	No
Slow Port Scan	Yes	Yes	Yes
Source Equals Destination	Yes	Yes	Yes
SYN Scan	Yes	Yes	No
TCP Scan	Yes	Yes	No
UDP Scan	Yes	Yes	No
Worm Attack	Yes	Yes	Yes
Worm Propagation	Yes	Yes	Yes
XMAS Scan	Yes	Yes	No

Algorithms for public IP addresses

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

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

If the primary concern is monitoring public assets, ensure that all public IP addresses are contained within an IP Group. Add the edge routers to most algorithms.

9.4 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 of events trigger a notification when a configurable threshold is breached.

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

In the practice of threat detection, reacting to any single odd behavior often leads to tail chasing because often times normal communications can lead to an occasional odd connection that triggers an event. The Threat Index reduces this problem. Also, different algorithms that increase the Threat Index have different multiplier weights as they are considered more suspicious behaviors. Modify the TI weight by *editing the Policy*.

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

9.5 Configuration

9.5.1 Setting up Flow Analytics

FA algorithms are executed sequentially. Most of them do not run until one or more NetFlow exporters are added to the individual algorithms.

To add exporters to an algorithm, visit **Admin > Settings > Flow Analytics Configuration** and click on an algorithm name listed in the table.

At the top of the **Flow Analytics Configuration table**, it displays the overall time to run all algorithms and the total count of violations across all algorithms.

FA configuration columns

- **Down Arrow Menu:** This action menu provides several options:
  - Modify the Exporters this Algorithm runs against: Many algorithms do not need to run against all exporters. Visit the *Algorithm strategy page* to learn more about types of flows to send to each algorithm.
  - Modify the **Hosts Excluded from violating this Algorithm**: Use this utility to exclude IP addresses and portions of hostnames that are triggering false positives.
  - View the **Run Time Trend for this Algorithm**: View a report that displays how long the algorithm takes the run each time it is executed.
  - View the **Violation Count Trend for this Algorithm**: View a report that indicates how frequently the algorithm is triggering for a matching event.
- **Mouse over columns to learn what they do.**
- **Round Icon:** This icon indicates the status of the algorithm using different colors. Mouse over the icon and the tool tip that appears will explain the status.
- **Name:** This is the name of the algorithm that is checking for abnormal behaviors. Click on the algorithm name to modify the settings, apply exporters or change the exclusions for the algorithm.
- **Time:** This is the amount of time the algorithm takes to run across all selected routers/switches.
- **Count:** This is the number of violations found the last time the algorithm ran. Click on the number to view graphs for longer time periods.

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

9.5.2 Optimizing Flow Analytics

Flow Analytics can be optimized in several different ways:

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

9.5.3 Excluding from algorithms

In an effort to reduce the false positives triggered by algorithms, IP addresses and portions of host names can be excluded from them. This feature can be found by visiting Admin tab > Settings > Flow Analytics Exclusions. Most exclusions are added when viewing an individual event for an alarm.

Exclusion Types

There are two ways to exclude hosts from triggering algorithms.

IP Address: Exclude one or more IP addresses from an individual

algorithm by:

- IP address

- IP range
- IP subnet
- Child: A child group is defined in *IP Groups*.

Reverse DNS Name: A portion of or all of a DNS resolved name can be entered.Entries are created when false positives occur. Use this interface to manage all of the entries.

Visit the **Alarms tab** and drill in on an alarm to see the individual events. To add an exclusion, click the Down Arrow Menu icon on the far left and select:

- Exclude IP x.x.x.x from this algorithm
- Exclude a portion of the reverse DNS name from this algorithm

The user can then use this interface to manage all of the the entries that will be made over time.

9.6 Custom algorithms

Developers and administrators can now create their own algorithms that Flow Analytics will execute routinely whenever the pre-packaged algorithms are executed. The level of complexity in a custom algorithm will vary based on the task(s) that are being performed. An indepth knowledge of Perl and Scrutinizer's database structure are required.

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

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

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

Creating a custom algorithm

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

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

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

Adding a custom algorithm

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

To open the Interactive scrut_util prompt, enter:

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

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

```
SCRUTINIZER> enable custom_algorithm customAlgo1 My_Custom_Algorithm
```

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

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

Testing a custom algorithm

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

To test a custom algorithm, execute the following command:

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

Configuring a custom algorithm

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

Settings available to custom algorithms are:

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

Running a custom algorithm in production

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

Creating a policy for a custom algorithm

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

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

1. Navigate to Alarms > Configuration > Policy Manager
2. Click New Policy
3. Give the policy a Name

4. Under Filter and to the right of Message, paste in **FA:userCustom_customAlgo1** in the box below where it says: Logical AND=&& Logical OR=||

5. Set the other options for this policy (for more details, reference the *Policy Manager* section of the documentation)

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

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

Disabling a custom algorithm

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

Deleting a custom algorithm

Custom Flow Analytics algorithms can be deleted from the Flow Analytics engine by executing the following interactive scrut_util command.

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

SCRUTINIZER> delete custom_algorithm custom_algo1
```

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

Display a list of custom algorithms

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

```
SCRUTINIZER> show custom_algorithms
```

9.7 FA Bulletin Boards

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

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

Policy Events

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

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

Indicators of Compromise (IOC)

The algorithms that post to the IOC bulletin board (BB) are those that indicate possible malware activity with insufficient confidence to initiate a security event alarm. This is a bulletin board that should be periodically reviewed to pick up on recent changes. If multiple IOCs are associated with a single host, these may generate an "Indicator Correlation Event" alarm that is posted to the Security Event BB. This is discussed in more detail below. The following algorithms are posted to the IOC Events BB by default:

- Breach Attempt Detection
- DNS Hits
- ICMP Port Unreachable
- ICMP Destination Unreachable
- Denied Flows

- Domain Reputation
- FIN Scan
- NULL Scan
- Odd TCP Flags Scan
- Persistent Flow Risk
- RST/ACK Scan
- SYN Scan
- TCP Scan
- UDP Scan
- XMAS Scan

Security Events

The algorithms that post to the Security Events bulletin board are high confidence detections that a host is compromised. Any events posted to this bulletin board should be investigated. The following algorithms are posted to the Security Events BB by default:

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

Note: Edit the policy for any algorithm to change the Bulletin Board from the default settings.

CHAPTER 10

Baselining

10.1 Overview

Baselining is analyzing the performance of the network by comparing current performance to historical data (also known as the “baseline”). For example, when measuring the current traffic from an interface, alerts will be sent if any traffic exceeds the baseline. Plixer Scrutinizer can be configured to baseline any NetFlow or IPFIX element and alert. By default, there are several baselines enabled and configured to collect data whenever a new exporter starts exporting flows.

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

- applicationTag (NBAR) and packetDeltaCount

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

Note: Baselining is disabled by default.

Using the interactive scrut_util command, baselining behavior can be modified whenever that baseline exceeds a higher than normal traffic pattern.

10.2 Baseline reports

You can monitor the progress of the baselining function by ruiing baseline reports:

1. In the **Status** tab, under **Device Explorer**, filter on **Scrutinizer**.
2. Click on **Scrutinizer**.
3. Navigate to **Reports > Baselines**.
4. This will provide a list of available auto-created baseline reports on your server so that you can create a report to view.

10.3 How baselining works

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

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

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

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

10.4 Configuring baselines via the interactive scrut_util

To enable the interactive scrut_util utility, run:

> /home/plixer/scrutinizer/bin/scrut_util

This will open the Plixer Scrutinizer prompt:

> **SCRUTINIZER>**

For more information on using this utility, reference the *interactive scrut_util* section.

10.5 Adding/removing default exporter baselines

Run the following interactive scrut_util command, *default baselines* to add/remove baselines from a specified exporter.

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

Adding enable baseline <exporter_ip> default

Removing disable baseline <exporter_ip>

10.6 Adding custom baselines

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

The following interactive scrut_util command can be used to add custom (manual) baselines to a specified exporter.

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

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

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

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

- Replace <exporter_ip> with the exporter to collect the specified baseline data.
- <pri_element> and (optionally <sec_element>) specify which IPFIX elements will be part of this baseline. This is how the data is grouped and calculated.
- <element> must be an IPFIX element with a numeric value that the <pri_element> and/or <sec_element> will be baselined on.
  - For example: to collect data on source addresses and bytes, the <pri_element> would be sourceipv4address and the <element> would be octetdeltacount.
  - Another example: To collect data off multiple elements such as source address and applicationtag based on packets; the <pri_element> would be sourceipv4address, the <sec_element> would be applicationtag, and the <element> would be packetdeltacount.
  - To find the numeric values for IPFIX elements in Plixer Scrutinizer, go to **Status > System > Templates**. Drill in on the Plixer ID to see all details regarding the elements in that template.
- <AVG|COUNT|MIN|MAX|STD|SUM> are options that are used to calculate how to measure the <element>
  - **AVG** = Average
  - **COUNT** = Flow Count
  - **MIN** = Minimum Value
  - **MAX** = Maximum value
  - **STD** = Standard Deviation

  - **SUM** = Sum

- Lastly, there are several ways baselines can be compared:
  - **dailyhr** = same time frame (e.g. 1a - 2a) for each day
  - **busday** = same time frame (e.g. 1a - 2a) for each business day, skipping weekends
  - **sameday** = same day and time each week (e.g. 1a - 2a Mondays)

When baselining IP addresses, IP Groups must be configured with the ranges and subnets of addresses to be included in the baseline. This reduces a number of false positives as by excluding addresses that may talk once.

10.7 Monitoring baseline processing

To get the task ID for the baseline task (to be used in the 'check task' command below), run the scrut_util <configuring_baselines>'command:

show task baseline

Displays information regarding the baseline task.

Example:

```
SCRUTINIZER> show task baseline


+-----------+---------+----------------+--------------------------+
| TASK_NAME | TASK_ID | EXECUTABLE     | ARGUMENTS                |
+-----------+---------+----------------+--------------------------+
| baseline  | 224     | scrut_util     | ["--collect","baseline"] |
+-----------+---------+----------------+--------------------------+
1 Result(s) Found

Done (0.010545 seconds)
```

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

check task <id>

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

Example:

```
SCRUTINIZER> check task 224


+---------------------+----------+------------+------------------------
↪------------------+
| START_TIME          | RUN_TIME | RETURN_VAL | INFO_STR                   ␣
↪                  |
+---------------------+----------+------------+------------------------
↪------------------+
| 2019-01-23 13:10:00 | 764      | 0          | baseline : scrut_util     -
↪-collect baseline |
| 2019-01-23 12:10:00 | 875      | 0          | baseline : scrut_util     -
↪-collect baseline |
| 2019-01-23 11:10:00 | 2018     | 0          | baseline : scrut_util     -
↪-collect baseline |
| 2019-01-23 10:10:00 | 691      | 0          | baseline : scrut_util     -
↪-collect baseline |
| 2019-01-23 09:10:00 | 752      | 0          | baseline : scrut_util     -
↪-collect baseline |
| 2019-01-23 08:10:00 | 637      | 0          | baseline : scrut_util     -
↪-collect baseline |
```

- RUN_TIME indicates how long the baselining task has run. It should complete in a few seconds.
- RETURN_VAL indicates if any error code was returned. It should be 0. If it is any other value,contact Plixer technical support for assistance.

10.8 Resetting baselines to defaults

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

clean baseline

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

CHAPTER 11

Alarms

11.1 Overview

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 responsibile 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.

11.2 Views menu

11.2.1 Bulletin Board by Policy

In the bulletin board by policy view, the alarms are grouped by policy violated. The heat map in the bulletin board by policy view displays the policies (e.g. threat algorithms) that are violated. Y axis = count, X axis = unique hosts. The bulletin board by policy table displays:

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

11.2.2 Bulletin board by violator

In the bulletin board by violator view, the alarms are grouped by violating IP address. The heat map in the bulletin board by policy view displays the hosts that are violating policies. Y axis = count, X axis = unique policies. The bulletin board by violator table introduces a few new columns that were not outlined above:

- **Country / Group** – If an IP is a public address we determine the IPs country. If it is not a public address we check to see if it is in a defined IP Group.
- **Users** – User is determined based on violator address. The lookup requires eventlog collection be configured. See *Username Reporting* for details.
- **Violator Address** - The IP and/or DNS associated with the violator. *Click on a Violator address* to see all of the alarm events generated by that address.
- **Other columns** - described above.

11.2.3 Notification queue

The notification queue lists the last 24 hours of notifications that were sent or that currently in queue and waiting for execution. The notification queue table displays:

- **Violator Address** - the IP and/or DNS associated with the violator
- **Policy** - The associated Policy.
- **Notification** - The name of the notification sent.
- **Alert Type** - The type of notification sent (see Notication 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.

11.2.4 Orphans

The orphans view lists messages that did not violate policies. From this view, new policies can be created to organize alarms. The Orphan Table displays:

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

11.2.5 Policy violation overview

This view lists the threats detected by Flow Analytics. It includes the policies and the corresponding violations that occured 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*.

11.3 Configuration menu

- **Alarm Notifications** allow checingk off the entries that the Scrutinizer administrator would like to trigger events for. Events are posted as policy violations in the Alarms tab.
- **Alarm Settings** optimize how notifications are triggered depending on the unique environment. Call support for assistance.
- **Create New Board** enables the user to create or delete new bulletin boards.

To modify the bulletin board that a policy posts to, visit *Admin tab > Definitions > Policy Manager* and edit the corresponding policy.

Note: Bulletin boards can have permissions assigned to them. More details regarding the permissions can be read about under Usergroup Permissions

- **Flow Analytics Configuration**

The overall status of all algorithms and the total runtime and count of violations across all algorithms. For more information on Flow Analytics Configuration, please go to *FA Configuration* and *Algorithm Activation Strategy*.

- **Flow Analytics Settings** brings the user to **Admin > Settings > Flow Analytics Settings**.
- **IP Groups** allows the user to exclude IP addresses, entire subnets or ranges of IPs, as well as child groups from violating specific alorithms.
- **Notification Manager** sets up notifications which can be triggered by policy violiations.
- **Policy manager** brings the user to **Admin tab > Definitions > Policy Manager** which lists all of the policies that can be triggered by events. Events are passed through the policies and matches occur based on content in the Message, Source Address or Syslog Alert Level. A policy can be configured to do one of three things with an alarm:
- Post it to a Bulletin Board (Alarms posted to a Bulletin Board will also be stored in history).
- Only store in history for reporting.
- Delete the alarm (It is not available in any way).

Policies also determine if a notification should be processed for an alarm by associating alarm messages with a notification profile. The Policy Manager table displays:

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

11.4 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 tends 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.

11.5 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.

11.6 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.

11.7 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.

11.7.1 Setting the global threshold

Scrutinizer relies 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

11.7.2 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>Policy Manager*
>
> 1. Enter ‘Interface Exceeded’ in the search field
> 2. Click the Search button
> 3. Then click on the ‘Scrutinizer: Interface Exceeded Threshold’ Policy when it’s displayed
> 4. Next, click on the Assign button, select the Notification profile previously set up (see below to set up a Notification profile), then click Save

11.7.3 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.

11.7.4 Creating the notification profile to use

Notification profiles can be created once and applied to multiple *policies*. Notification methods include Email, Logfile, Syslog, SNMPTrap, Script and Auto Acknowledge. Enter the necessary data and select additional details from the **Available Variables for Message** list to ensure the desired information is included in the alert.

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.

Alerts

- **Email:** send an email alert
  - Enter the email address the alert is destined for in the "To" field

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

- **Logfile:** add alert message to a file
  - Enter log file name with the absolute file path of:

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

Note: Log files must be placed at this location.

- **Syslog:** send syslog alert to Host address.

  Required fields are:

  - Host: Target server address
  - UDP Port: Target server port (default 514)

- **Priority**
  - Facility

- **SnmpTrap:** send snmptrap alert to Host address.

  Required fields are:

  - Host: Target server
  - Community String
  - UDP Port
  - Enterprise OID
  - Generic ID
  - Specific ID
  - Binding OID
  - From Host

- **Script:** trigger action defined in Script.

  Required fields are:

  - Script: /home/plixer/scrutinizer/files/{alert_script.sh}
  - **Note**: Script must be placed in this folder and absolute path must be included in the Script field.
  - Command-line Arguments: Variables to include in the script from the Available Variables list below.

- **Auto Acknowledge:** automatically acknowledge policy alarms
  - Policy To Acknowledge: select target policy from dropdown list

11.7.5 Available variables for message

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

11.7.6 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.

If any assistance is needed, please contact us.

CHAPTER 12

Admin

12.1 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:** Setup and modify known hosts. Use this option to statically assign host names to IP addresses that will not age out. It can also be used to label subnets in the 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 mas: Defines a wildcard mask for IP Addresses. Example: IP Address: 10.0.0.1, Wildcard Mask: 0.255.255.0
  - Child group: Include other (child) IP Groups in this parent group. Select a child group from the dropdown selection list of existing IP Groups.
- **Language:** Use this interface to update languages or create new translations.
- **MAC Addresses:** Lists MAC Addresses with labels as collected by the utility. It is scheduled to run nightly.

  - MAC address descriptions are collected from Cisco wireless LAN controllers via SNMP.
  - MAC address descriptions are collected from option templates that contain these two elements: 'stamacaddress' and 'username'.
  - Run the scrut_util *'collect optionsummary'* utility to force immediate collection.
  - Manually enter or edit MAC address information here.

- **Manage Collectors:** Provides details on the servers which are collecting flows for this Scrutinizer install. Multiple collectors will be listed if a distributed solution has been deployed.
  - Delete: This check box can be used to remove collector(s) from the list.
  - Collector: IP Address of the flow collector.
  - State: Current state of the flow collector - ONLINE or OFFLINE.
  - Exporter Count: Number of exporters 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 over write it.

    - Edit Protocol Exclusions: Used to tell the collector to drop flows on certain ports. This was build 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.

  - Check Box: 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. The Understanding NetFlow Traffic Volume blog explains this in more detail.

- **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.

12.2 Settings

- **Alarm Notifications:** Enable additional system alarms.
- **Alarm Settings:** Modify settings to optimize syslog and SMTP processing.
- **ASA ACL Descriptions:** Enter the username and password used to SSH into ASA firewalls to retrieve ACL descriptions (Appliance only).
- *AWS Configuration*: Set parameters for Amazon Web Services flow streaming configuration here.
- **CrossCheck:** Specify the thresholds for changing color and the syslog threshold that the Fault Index must reach to trigger a syslog.
- **Data History:** Specify how long each flow interval is saved.

- **Historical 1 Min Avg:** Saves 100% of all flows received. Make sure the server has enough disk space to save significant quantities of the raw flows. The 1 minute intervals consume the most disk space as it is not aggregated and flows are in raw format.
- **Historical 5 minute - 1 week Avg:** These intervals only save the specified Maximum Conversations after aggregation per interval.
- **Maximum Conversations:** Used when creating large intervals (e.g. 5 minute) from prior intervals (e.g. 1 minute). All flows are aggregated together per router. The top 1,000 (default) based on bytes are saved.

> **Note:** The default value for the Flow Maximum conversations field is 1,000 and the maximum value is 25,000.

-**Auto History Trimming:** This option allows for automatic database trimming when available disk space falls below 10% (with a minimum threshold of 10GB). Check the checkbox to activate this option. An alarm will also be generated to send an alert that the database is being trimmed (1 minute and 5 minute conversation database tables) and includes how much 1 minute and 5 minute data currently exists in the database (in hours).

Read more about topics related to this subject:

- *Data Aggregation*
- *System LEDs*

> **Note:** In a distributed collector environment, each collector will perform the database trimming independent of the other collectors. Auto History Trimming on/off applies to all of the collectors in the cluster, but the database trimming will only occur on the server(s) that fall below 10% of available disk space.

- **Email Server:** Necessary for on demand and scheduled emailed reports. Make sure the test is successful.
- *Flow Analytics Configuration:* Used to configure the algorithms and monitor their performance.
- Flow Analytics Exclusions: Used to manage the Flow Analytics IP Group and hostname exclusions.
- **Flow Analytics Settings:** Used to modify default settings of Flow Analytics relating to FlowPro Defender, jitter, latency, violations and top algorithms.

- **Licensing:** Displays the current licensing level, expiration date(s), and unique Machine ID for this installation. **The Machine ID is required by Plixer Customer Service for generating new license keys.** Once a new key is received, to activate the key, copy and paste the entire key in the License Key textbox. See the *System > Licensing page* for more information.
- *Mapping Groups:* Add and manage Map Groups.
- *Mapping Objects:* Add and manage Map Objects.
- **Proxy Server:** Setup the server to work with a proxy server.
- Reporting: Report settings configuration options.
- **Syslog Server:** Configure the syslog server, port and priority.
- **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.

12.3 Security

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

12.3.1 LDAP user authentication process

1. In the LDAP configuration, administrators provide credentials for an LDAP account with permission to see any users they'd like to permit access to.

a. This is the account that will be used to search for and authenticate users when they attempt to log in.

b. The searchbase defines the group that will be used to search for authorized users. This is a required field.

c. The scope of users in that searchbase who are allowed to authenticate can be limited in two ways:

- By specifying one or more Security Groups in the LDAP Configuration

- By specifying individual user account names in Security > Authentication > LDAP

1. 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.

2. A user attempts to log in. The system authenticates as the administrative account provided, then checks a searchbase specified by the Scrutinizer administrator for any account matching the username provided. Authentication with the sAMAccountName, UserPrincipalName, or uid attribute is supported.

3. If the LDAP server responds with an LDAP_REFERRAL code, Scrutinizer will check the referred server.

4. If the Scrutinizer administrator has specified multiple LDAP servers, it will check them all until successfully authentication succeeds or fails.

5. Once the user has successfully authenticated for the first time, Scrutinizer checks for any security group they're a member of which also exists in Scrutinizer with the same usergroup name. If it does, they're added to the Scrutinizer usergroup automatically.

Note: With LDAP enabled, if you create a usergroup in Scrutinizer that has the same name as a security Group on the LDAP server, when a Scrutinizer user logs in they will automatically be added or removed from the Scrutinizer usergroup to keep Scrutinizer synchronized with the LDAP security groups.

For example, if you have a “Scrutinizer Users” User Group in Scrutinizer and a “Scrutinizer Users” security group in LDAP, when that user logs into Scrutinizer with their LDAP credentials and is a member of the LDAP Security Group “Scrutinizer Users”, that user will automatically be added to the Scrutinizer “Scrutinizer Users” User Group.

If the LDAP user is not a member of the “Scrutinizer Users” security group in LDAP, when that user logs into Scrutinizer, they will be removed from the Scrutinizer “Scrutinizer Users” User Group.

12.3.2 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.

12.3.3 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.

12.3.4 Single sign-on

Scrutinizer-Azure ADFS SAML integration

To set up the **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**

Identifer (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 Scrutinizer**
  - Copy the Azure AD Identifier value.

Note: The valies and the certificate you copied will be required to complete the Scrutinizer configuration.

This completes the Azure configuration. You should now assign users or groups to the **Scrutinizer** application in Azure ADFS so that they can successfully authenticate.

Scrutinizer configuration

Now that you have the required information from Azure's configuration, you can set up Scrutinizer's authentication. Log into 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 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 usergroup names if the company’s IdP is set up to provide them.
IdP Metadata URL	Enter the “App Federation Metadata URL” link you previously copied
IdP Metadata XML	Optonal. 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 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 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 “Scrutinizer” application in Azure ADFS, they will be granted access. If the local 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.

Scrutinizer-Okta SAML integration

To enable single sign-on through Okta in Scrutinizer, you must first create the application in Okta.Launch the Okta Classic UI to perform the steps below. If you see "Developer Console" in a dropdown at the top of your page, click it to switch to Classic UI.

1. Select **Applications** in the navigation bar.
2. Click the **Add Application** button.
3. In the sidebar, pick the green **Create New App** button.
4. In the modal that appears, set **Platform** to **Web**, tick the **SAML 2.0** radio button,and then click **Create**.

Once a new application is created, you will see page 1 of its **General Settings:**

1. Enter **Scrutinizer** for the App name. Click Next.
2. Use the following format for **Single sign on URL:** https://<scrutinizer_server>/fcgi/scrut_fcgi.fcgi?rm=usergroups&action=sso_response
3. Set **Audience URI** to: https://<scrutinizer_server>/
4. Skip the other options and click Next, and then Finish.

You will be redirected to the **Sign On** settings page for the Scrutinizer application.

1. 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
2. Click **View Setup Instructions.**
3. Set the **Identity Provider Single Sign-On URL** to: https://<okta_server>/app/application_id_and_name/identifier/sso/saml
4. Use this link for the **Identity Provider Issuer:** http://www.okta.com/identifier
5. Click the **Download Certificate** button and save your okta.cert file. We will need to copy it to the Scrutinizer server later.

With the Okta configuration complete, you should now assign users or groups to the **Scrutinizer** application so that they will be able to successfully authenticate.

Scrutinizer configuration

Now that you have the required information from Okta's configuration, you can set up SSO authentication in Scrutinizer. Log into 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 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 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 Scrutinizer admin account is needed, or if other authentication methods are configured (e.g. LDAP or RADIUS), the login form can be accessed directly at "https://<scrutinizer_server>/login"

- **User Groups:** Specifies what a Group login account can access. More details regarding the permissions can be read about under Usergroup Permissions.
- **Users:** Configure login preferences for individual accounts. User Accounts must be a member of one or more User Groups. If no group is selected when a User Account is created, they are placed in the default (e.g. Guest) User Group. Permissions for a User Account are inherited from all the User Groups it is a member of.
- **User Account Lockout:** If a user has a specified amount of failed logins within a defined period of time, that user's account will be set to 'locked' status and will require a user with administrative permissions to unlock it.

  These settings are defined in **Admin > Settings > System Preferences** and include:

- Failed Login Max: the maximum failed logins allowed within the Failed Login Window time
- Failed Login Window: the number of minutes that the Failed Login Max value is matching against

  For example, with these settings:

  > Failed Login Max = 2 Failed Login Window = 5
  >
  > Two failed logins within a 5 minute timespan would cause that user account to be locked out.
  >
  > 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.

12.4 Managing devices and interfaces

You can make changes to the device and intreface 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 over write it.
- Edit Protocol Exclusions: Used to tell the collector to drop flows on certain ports. This was build 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.
- Check Box: 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 check box for SNMP Enabled. If SNMP Enabled is checked, the Watcher Service will attempt to poll and update SNMP information for the device. By default, the automatic SNMP discovery occurs once a night. The user can disable the automatic SNMP capability by unchecking **Auto SNMP Update** from the **Admin Tab > Settings -> System Preferences**.

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

- **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.

12.5 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.

12.5.1 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

12.5.2 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 countedonce.

- **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.

1. Stack 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.

12.6 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 exportering 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.

12.6.1 Usergroup permissions

Users are assigned to usergroups. Usergroups are granted permissions. Users inherit permissions from all the usergroups they are a member of. This functionality also serves as the basis for the enterprise focused multi-tenancy functionanlity.

- **New User Groups:** Is used to create a new usergroup that individual users can be assigned to. Give the group a name and apply a template from another Usergroup that has similar permissions to the new user group. After creating an account, find the new usergroup on the left and click it to modify.

  Click here for a special note regarding Scrutinizer usergroups and LDAP security groups.

- **Administrators:** This is the admin account and cannot be deleted. Users can be assigned to this group and inherit all of its permissions.
- **Guest:** This is the default guest account which cannot be deleted. Users can be assigned to this group and will have limited permissions.

Important: Permissions for an individual user account will be inherited from all usergroups it is a member of. To view all the usergroups a user account is a member of, visit **Admin tab > Security > Users** and click on a user account. Then open the **Group Membership** tab.

Members

Select the user accounts that will need to have access to this usergroup. A user can be a member of multiple usergroups and inherit all applicable permissions.

Features

Permissions control features the usergroup should have access to within Scrutinizer. Permissions can restrict product features entirely for a usergroup or specific features can be accessed based on your usergroup membership.

Features include:

- Which tab the members of the usergroup should be able to see,
- Administrative permissions the usergroup should have access to,
- Advanced features like acknowledging alarms, scheduling reports, adding/deleting users etc.

Clicking the **Configure** link in the **Features** column will provide a click and drag modal to adjust usergroup permissions. Inside that modal, on the left will two radio buttons with **Predefined** and **Advanced** labels. The following section describes the difference between the two modes, as you must chose one or the other per group.

12.6.2 Predefined roles vs advanced features

The features modal allows Usergroups to use predefined roles or manually specifiying features. A Usergroup must use either the Predefined Feature sets **or** the Advanced features that can be manually configured.

Important: You cannot configure manual permissions for a predefined set.

- **Advanced** - Manually configure all permissions available. Use Advanced to create custom feature sets.
- **Predefined roles** - Feature sets for common persona's like "ReportUser" or "DashboardAdministrator"

Predefined role	Underlying permissions
AlarmsAdministrator	ackBBEvent alarmSettings almDelete LogalotPrefs NotificationManager PolicyManager
AlarmsUser	alarmsTab
DashboardAdministrator	dashboardAdmin
DashboardUser	createDashTabs myViewTab
MapsAdministrator	mappingGroupConfiguration mappingObjectConfiguration
MapsUser	adminTab allLogalotReports mapsTab reportFilters statusTab
ReportingAdministrator	ApplicationGroups asnames deleteReport HostNames protocolExclusions reportSettings tos viptelaSettings wkp
ReportingPowerUser	reportFolders ReportDesigner saveReport scheduledReports srCreate
ReportingUser	runReport
SystemAdmin-	3rdPartyIntegration auditing auth Authentication authLdapServers awsSettings changeUserPasswords createUsers CrossCheck DataHistory deleteUsers DeviceDetails EmailNotifications fa_mgmt_link faExclusions feedbackForm Flow-

- **Device status** is used to grant permission to see the status of the device (i.e. Flow exporter). Device icons appear blue in maps if the **Device Group** permission is granted without this permission.
- **Interface statistics** grants permission to see the statistics of an interface.
- **Groups** are used to grant permission to see a group (i.e. map). Devices (i.e. flow exporters) appear blue and interfaces black unless permission is granted in **Device Status** and **Interface Statistics**.
- **Saved reports** allows to select the saved reports/ filters that the usergroup will need to have access to run.
- **Dashboard gadgets** selects the gadgets that the usergroup will need to be able to add to dashboards.
- **Third-party links** controls the vendor third-party integrations that the usergroup will be able to integrate with.
- **Bulletin boards** manages the Bulletin boards that the usergroup will need to be able to access in the Alarms tab.