Replicator Documentation

Release 18.1

Plixer

Jul 12, 2018

Flow Replicator

Welcome to the on-line manual. *Click Here* for online troubleshooting or FAQs. This manual is also available in .pdf format.

Important: Don't struggle, contact Plixer support!

CHAPTER 1

What is a Flow Replicator

1.1 Overview

Many routers, servers, and other systems can only send messages to a single log management system. The Flow Replicator™ allows a single stream of log data to be transparently replicated to multiple destinations.

By configuring these network devices to send their log and flow data to the Flow Replicator, users can control which management system(s) receive the replicated data.

A Flow Replicator eliminates the limitation of sending log and flow data to a small number of management systems.

1.2 The Basic Concepts

The Flow Replicator relies on Profiles that contain a list of devices (exporters) sending or streaming data to management systems (collectors). A packet is received by the Flow Replicator on a particular UDP port. The Flow Replicator references a list of profiles to determine if the data received from an exporter should be forwarded on to one or more collectors.

- Collectors : A collector is a SIEM, Flow Collector, SNMPTrap Receiver, or other Network Management System that actively receives data from networked devices.
- Exporters : An exporter is a networked device such as a router, switch, or server that generates different types of data and is capable of sending that data to a collector.
- Profiles : A profile contains exporter(s), a listening port, collector(s), and sending UDP port.

1.2.1 How Profiles Work

```
+--------------------------------+-----------------------------------------+
| MyProfile                      | IN PORT 2002 -> OUT PORT 9996
+--------------------------------+-----------------------------------------+
```

(continues on next page)

(continued from previous page)

```
Policies                    Exporters     ->    Collectors
(include) 10.3.1.1/32       10.3.1.1            10.11.1.165
```

When a packet is received from the exporter 10.3.1.1 on port 2002 it is replicated to the collector 10.11.1.165 on port 9996. Collectors interpret that packet's origin as 10.3.1.1 and not the replicator.

Profiles can contain multiple exporters and collectors.

```
+-------------------------------+----------------------------------------+
| distdev-63                    | IN PORT 2002 -> OUT PORT 2055
+-------------------------------+----------------------------------------+

 Policies                     Exporters     ->    Collectors
 (include) 10.1.2.18/32       10.1.2.18           10.1.10.63
 (include) 10.4.1.1/32        10.4.1.1            10.1.4.203
 (include) 10.9.1.254/32      10.9.1.254          10.30.11.23
 (include) 192.168.0.17/32    192.168.0.17
```

If a packet is received from the exporters 10.1.2.18, 10.4.1.1, 10.9.1.254, or 192.168.0.7 on port 2002 it is replicated to 10.1.10.63, 10.1.4.203, and 10.30.11.23 on port 2055.

1.2.2 Policies

A policy is used to determine if a particular exporter should be included or excluded from the profile. Administrators add policies using CIDR notation to include or exclude exporters in a profile.

```
+-------------------------------+----------------------------------------+
| steady-replays                | IN PORT 2002 -> OUT PORT 9996
+-------------------------------+----------------------------------------+

 Policies                   Exporters     ->    Collectors
 (include) 10.25.5.0/24     10.25.5.122         10.1.10.1
                            10.25.5.123
                            10.25.5.10
                            10.25.5.29
                            10.25.5.30
```

The replicator will automatically replicate packets from any exporter matching 10.25.5.x on port 2002 to collector 10.1.10.1 on port 9996.

Exclude policies can be used with include policies to exclude one or more exporters.

```
+-------------------------------+----------------------------------------+
| steady-replays                | IN PORT 2002 -> OUT PORT 9996
+-------------------------------+----------------------------------------+

 Policies                     Exporters     ->    Collectors
 (include) 10.25.5.0/24       10.25.5.122         10.1.10.1
 (exclude) 10.25.5.10/32      10.25.5.123
                              10.25.5.29
                              10.25.5.30
```

The above example is identical to the previous example except it has an exclude policy. In this case, any exporter matching 10.25.5.x except 10.25.5.10 will be replicated to the collector 10.1.10.1.

CHAPTER 2

Installation

There are two types of appliances available. A valid or evaluation key is required with either install. A key can be obtained from Plixer or a local reseller.

2.1 Hardware Appliance

Once the hardware appliance is installed in a network rack, power it on and follow the steps below.

1. Using an SSH client, remotely login using the username root and password replicator. The hardware appliance will perform a quick setup and immediately reboot.

```
CentOS release 6.5 (Final)
Kernal 2.6.32-431.3.1.e16.x86_64 on an x86_64

localhost login: root
Passord: _
```

2. Login to the hardware appliance again using the username root and password replicator. Input the answers to the configuration questions. The hardware appliance will reboot to apply the necessary settings.

```
Last Login: Tue Feb 11 11:067:45 on tty1

*******************************
Replicator Virtual Appliance
Initial Configuration
*******************************

What is the appliances static IP Address?
10.1.15.128

What is the appiances Netmask?
255.255.0.0
```

(continues on next page)

(continued from previous page)

```
What is the appliances gateway?
10.1.1.1

What is the hostname for this appliance?
replicatorVA_
```

3. Login to the hardware appliance command line with the replicator username and password configured in the previous step. Apply the license key by logging into the web UI or by issuing the license set command on the CLI:
   (a) In the new window, under "license=", paste in your license key.
   (b) Press CTRL+x to save.

The replicator is now ready for configuration.

2.2 Virtual Appliance - ESX

The Replicator Virtual Appliance (RVA) is packaged as an all-in-one virtual machine template known as an OVF template.

For VMware deployments, ESX/ESXi 5 or higher is required. VMware Tools will be required to shut down the RVA through the VMware vSphere Client.

2.2.1 System Requirements

Component	Recommended Specifications
RAM	2GB
Disks	100GB
Processor	2 CPU 2 Core 2GHz+
Operating System	ESXi5+

2.2.2 Deploying the OVF Template

1. Connect to the ESX host using VMware vSphere, or vCenter.
2. Select File then Deploy OVF Template
3. Select Deploy from File, browse to the OVF Template, and click Next.
4. Review the OVF template details and click Next.
5. Define the name of the Replicator Virtual Appliance and click Next.
6. Select a datastore and click Next.
7. Select the disk format and click Next.
8. Select the Network Mapping and click Next.
9. Review the Virtual Settings and click Finish to import the OVF Template.
10. Right click on the Flow Replicator virtual machine and power it on.
11. Navigate to the Console tab and login using the username root and password replicator. The virtual appliance will perform a quick setup and immediately reboot.

12. Login to the hardware appliance again using the username root and password replicator. Input the answers to the configuration questions. The virtual appliance will reboot to apply the necessary settings.
13. Login to the hardware appliance command line with the replicator username and password configured in the previous step. Apply the license key by logging into the web UI or by issuing the license set command on the CLI:
    (a) In the new window, under "license=", paste in your license key.
    (b) Press CTRL+x to save.

The replicator is now ready for configuration.

2.2.3 Installing VMware Tools

VMware Tools are not required for proper function of the virtual appliance. However, there are certain advantages to deploying it on each virtual appliance. See VMware's documentation for more details.

VMware Tools are not installed by default because each version of ESX installs a different VMware Tools package. A script is included with the Virtual Replicator to simplify the install process.

1. In the VMware vSphere Client, right click on the Replicator virtual machine and select Guest, then Install/Upgrade VMware Tools.
2. Login to the console of the Replicator Virtual Appliance as the root user and run the command /home/replicator/conf/vmwareToolsInstall.sh

2.2.4 Upgrading the Virtual Machine Hardware Version

The Replicator Virtual Appliance is built on Virtual Machine Hardware Version 7 to maintain backwards compatibility with ESXi 5 hypervisors.

While the virtual machine is powered off, in vSphere (or vCenter) right click on the virtual machine and select Upgrade Virtual Hardware.

2.3 Virtual Machine – Hyper-V

2.3.1 System Requirements

```
Component   Recommended Specifications
RAM         2GB
Disks       100GB
Processor   2 CPU 2 Core 2GHz+
```

2.3.2 Importing Virtual Machine

1. Download the latest Plixer Flow Replicator
2. Unzip the file on your Hyper-V server
3. Open Hyper-V Manager and select Import Virtual Machine
4. Specify the Replicator System Folder
5. Select the Virtual Machine

6. Choose the import type
7. Go to Settings
8. Select your Network Adapter and assign it to the appropriate Virtual Switch
9. Expand the Network Adapter section, select Advanced Features, set the MAC Address to Static, enter in a unique MAC Address, and then press "OK".
10. Start the Virtual Machine.
11. Right Click on the Virtual Machine and click Connect to login to the Plixer Flow Replicator using root/replicator. The server will perform a quick setup and immediately reboot.

2.4 Virtual Machine – KVM

2.4.1 System Requirements

```
Component   Recommended Specifications
RAM         2GB
Disks       100GB
Processor   2 CPU 2 Core 2GHz+
```

2.4.2 Importing Virtual Machine

1. Create a directory for your install

   mkdir kvm/Scrut_VM_Guide/

2. Download the latest Replicator Virtual Appliance to your KVM install

   Command Line Example:

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

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 Replicator_KVM.tar.gz
```

4. Run your script to install Replicator

```
sudo ./install.sh
```

At this point the machine has been created from the image that was deployed.

5. Lastly, we just need to log into the machine now that it has been deployed. Run this command to get to the console.

```
virsh console Replicator
```

You will be prompted to login, default credentials are root/replicator. The machine will reboot and you will be asked to login again. This time, you will be presented with a shell script asking for networking information. Follow the on screen instructions and celebrate!

CHAPTER 3

Getting Started

3.1 Wizard

The web interface contains a wizard to assist users in the initial configuration process. The process consists of three steps.

1. Apply a license key
2. Create a profile
3. Verify packets are inbound

3.2 Enabling HTTPS

It is recommended that HTTPS be enabled on your replicator to secure any information sent through the browser. This is done on install or by using the script “enable_ssl.sh” in the /home/replicator/conf/ directory.

3.2.1 Logging in for the first time

The default login user name and password are admin. The password can be updated using the interactive mode command password webui.

Once LDAP server information is entered in the settings tab, users will also be able to log in using valid LDAP credentials. The LDAP password can be updated using the interactive mode command ldapadminpass.

Future versions of the replicator will provide local multiple user accounts and roles. The current version provides all users with an administrator role account to update and maintain any configuration via the web interface.

3.3 Applying a License Key

Once logged in, the replicator will detect if a license key is present. If not, the license wizard will appear. Otherwise, the replicator will redirect to the next step in the wizard if a profile doesn't exist. If everything is minimally configured, the replicator displays the Dashboard Tab.

A license key can be requested directly from plixer or a reseller. Evaluation keys are available for testing and evaluating purposes.

3.4 Create a Profile

Once a valid license key is configured, the replicator detects if it has at least one profile. A profile defines how packets are received and sent from the replicator.

If there are no profiles configured, the replicator will ask a few questions and get the user started quickly.

A profile requires the port packets will come in, the port packets should go out, a policy to match incoming IP Addresses on the defined port, and a collector to send the packets.

In many cases, users typically want to match all incoming IP addresses. In this case, the policy of 0.0.0.0/0 can be used.

3.5 Verify Packets are inbound

The Streams Tab and Dashboard Tab provide users with instant feedback of packet activity.

If it is the first time configuring the replicator, the last step of the wizard will redirect the user to the Streams Tab.

For more information on the web interfaces see the section on the Web Interface.

3.6 Interactive Mode

Using an SSH Client, ssh to the Flow Replicator and log in as the replicator user using the password configured during the installation process.

```
[root@demo ~]# ssh replicator@10.1.4.66
replicator@10.1.4.66's password:
Last login: Tue Jul 15 16:55:17 2014 from scrutinizer.plxr.local

Plixer Replicator (TM) v17.1.17.1771
[2016-06-03 09:22:51 -0400 (Fri, 03 Jun 2016)]
Copyright (C) 2012 - 2017 Plixer International, Inc. All rights reserved.
Replicate Anything!
Need an IPFIX Collector? Download Scrutinizer at https://www.plixer.com

Machine ID : 6YZ6XEPTJA6VG749B
Licensed Version : 16.6
Licensed Type  : eval (standalone/primary)
Expiration : Thu May 18 2017

License expires in 317 day(s) (eval)
REPLICATOR>
```

The REPLICATOR> prompt indicates the Flow Replicator is ready for commands.

Before beginning any configuration of the Flow Replicator, configure one or more networked device (exporter) to send flow or log data to the Flow Replicator. Note the UDP port(s) used to send data during the configuration process. For example: A Cisco router is configured to send NetFlow v9 using port 9996 to the Flow Replicator. 9996 will be used during the profile creation process.

It is important to send data to the Flow Replicator first. Once the exporters are sending data to the Flow Replicator, profiles can be configured to replicate the data to the appropriate collectors. The profile command is used to manage profiles on the Flow Replicator. The required parameters for creating a profile are name, listening port, and sending port.

```
profile <add|update> name listen_port send_port
```

In the above example, the Cisco router is sending data to the Flow Replicator on port 9996. The Scrutinizer Flow Collector is listening for flows on port 2055. A profile is created to listen on port 9996 and send on port 2055 as follows:

```
REPLICATOR> profile add maine 9996 2055

Success: Profile 'maine' has been added and enabled.

Done in 0.316286 secs

REPLICATOR> show profile maine

+--------------------------------+-----------------------------------------+
| maine                          | IN PORT 9996 -> OUT PORT 2055
+--------------------------------+-----------------------------------------+

 Policies                      Exporters     ->    Collectors
 -                             -                   -

+--------------------------------+-----------------------------------------+
Done in 0.01249 secs
```

Next, add the exporter (i.e. Cisco router) to the profile using the exporter command.

```
REPLICATOR> exporter add 10.1.1.1 maine

Success: Exporter [10.1.1.1] -> Profile [maine]

Done in 0.183711 secs

REPLICATOR> show profile maine
+--------------------------------+-----------------------------------------+
| maine                          | IN PORT 9996 -> OUT PORT 2055
+--------------------------------+-----------------------------------------+

Policies                      Exporters     ->    Collectors
(include) 10.1.1.1/32         -                   -

+--------------------------------+-----------------------------------------+
Done in 0.014918 secs
```

As soon as the Flow Replicator detects traffic from 10.1.1.1 on port 9996, the exporter will show up in the Exporters column. This indicates that the Flow Replicator is actively replicating to the collectors specified.

Lastly, add a collector to the profile.

```
REPLICATOR> collector add 10.1.4.20 maine

Success: Collector [10.1.4.20] -> Profile [maine]

Done in 0.439209 secs

+--------------------------------+----------------------------------------+
| maine                          | IN PORT 9996 -> OUT PORT 2055          |
+--------------------------------+----------------------------------------+

 Policies                      Exporters     ->    Collectors
 (include) 10.1.1.1/32         10.1.1.1            10.1.4.20

+--------------------------------+----------------------------------------+
Done in 0.011346 secs
```

The profile is complete. Within moments, the replicated traffic can be verified within the collector's interface.

CHAPTER 4

Normal Operations

4.1 Replication

During normal operation the Flow Replicator will replicate incoming packets to all configured collectors in enabled profiles. At any time, executing show realtime will display the exporters in and out packet rates and totals.

4.2 Alarming

The Flow Replicator is actively tracking the number of packets received, packets sent, and the state of any exporter and collector. An alarm is generated and a syslog is sent if an exporter stops sending packets or a collector becomes unreachable.

By default, the Flow Replicator is configured to stop replicating traffic to collectors that are considered offline. Replication will resume once the collector is reachable.

4.2.1 Dropped packets

The Flow Replicator examines the netstat details of each interface and each direction (Rx and Tx) once every minute. When the OS reports there are interface drops a syslog alarm is sent to the server configured in the Notifications section of the settings tab.

The counting system used for the web interface that tracks the number of packets traversing the Replicator uses tcpdump instead of netstat and does a hard cut off on a timed basis. This means in the web interface there may be packets counted inbound that haven't been counted outbound yet. The counters increment close to realtime, but not instantly. The web interface metrics are a good reference point for packet activity, but not for drops. In addition to device state and dropped packets, the Flow Replicator will send a notification if CPU is high or processes were terminated abnormally.

The following settings control alarming capabilities in the Flow Replicator.

- downDisplayHour : The number of hours before an incoming stream is automatically acknowledged as being down. Default is 24 hours.

- flowStopAlert : The number of minutes an incoming stream must stop or a collector is unreachable before it is considered down.
- highCPUThreshold : Send alerts about the CPU when it exceeds this percentage. Default is 90%
- noRepWhenDown : If ping is enabled and a collector is unreachable, stop replicating data to that device. Replication will continue when the collector begins to respond to pings.
- notificationSent : Send Replicator Alert and Notification Syslogs to the SERVER and Port specified.
- pingCollectors : If enabled, the Plixer replicator will routinely check the configured collectors for availability.

Use the setting CLI command to affect the global behavior of alarming.

Plixer's Scrutinizer Incident Response System includes policies for all possible alarms from the Flow Replicator.

4.2.2 Reporting

The show or list command has several different options to generate reports based on live data. Additionally, the Flow Replicator can export replication statistics as IPFIX to a Flow Collector.

- metricsSent : Export Replicator Statistics and Metrics to an IPFIX Collector on the specified Collector IP and Port Number.

Use the setting command to manage IPFIX metrics.

A profile can be set up to send IPFIX metrics to multiple collectors by configuring the metricsSent option to send metrics back to the Flow Replicator on a certain port (e.g. 10.1.4.66:2003)

```
+--------------------------------+-----------------------------------------+
| replicator_metrics             | IN PORT 2003 -> OUT PORT 2056
+--------------------------------+-----------------------------------------+

 Policies                      Exporters     ->    Collectors
 (include) 0.0.0.0/0           10.1.4.66           10.1.10.1
                                                   10.1.4.20

+--------------------------------+-----------------------------------------+
Done in 0.00897 secs
```

Plixer's Scrutinizer Incident Response System supports different reports for the Flow Replicator. For additional information on reporting, reference the section in this manual on the show command.

CHAPTER 5

Server Maintenance

5.1 Hardware Failure

If any hardware malfunctions occur, contact technical support for assistance.

5.2 Applying Security Patches

Although efforts are made to minimize the risk for security breaches on the appliance, updates to core OS components may be applied. It is recommended that updates are not installed unless technical support advises or assists. For more information, contact technical support.

5.3 Upgrades

Customers are entitled to upgrades provided that maintenance is active. For further instructions, contact technical support.

5.4 Backing up the Flow Replicator

The Flow Replicator database can be backed up and restored using the backup and restore Interactive Mode commands.

CHAPTER 6

Web Interface

6.1 Logging In

When navigating to the replicator using a modern web browser, the user will need to authenticate. The default login user name and password are admin. The password can be updated using the interactive mode command password webui.

Once LDAP server information is entered in the settings tab, users will also be able to log in using valid LDAP credentials. The LDAP password can be updated using the interactive mode command ldapadminpass.

After successfully logging in, the Dashboard Tab is displayed.

6.2 Dashboard Tab

The Dashboard Tab provides real time visualization feedback of the overall operations of the replicator. Information such as total packets inbound, packets outbound, bits inbound, and bits outbound are provided in both capacity charts (donut charts) and live trends (line graphs) displaying the last 20 data points of traffic. The time span of each trend will vary depending on the overall refresh rate of the live data.

Below the trends is a summary of total exporters, collectors, exporter/collector pairs, profiles, and last sample of traffic totals.

6.3 Streams Tab

The streams tab provides real time status of exporters sending packets to the replicator. New streams will highlight in green, and streams that stopped will highlight in red. The user can summarize all stopped streams by checking the alarm only box. A filter box is provided to quickly narrow the list down based on a keyword match.

By clicking the [check] link, the replicator will attempt to resolve the name of the exporter using DNS.

Additional details such as Last Status check, current status, and traffic details are displayed.

6.4 Exporters Tab

The exporters tab provides details related to the devices sending packets to the replicator. If there are any alarm conditions from the exporters, the total number of alerts will be highlighted by the Exporter Tab.

6.4.1 Exporters in Profiles

Any exporter that is in an alarm state is highlighted in red. The total number of exporters in profiles with alarms are highlighted beside the sub tab. The same checkbox and filter is available on each tab to quickly find the desired exporter(s).

6.4.2 Exporters Not in a Profile

The exporter tab contains two sub tabs. Exporters in Profiles (explained above) and Exporters Not in a Profile. The replicator will show users which incoming exporter streams are currently not associated with a profile.

The sub tab helps users quickly find and configure streams that are not replicated to any collectors. Exporters highlighted in red indicate a stream that was previously known has stopped. The user can click the [check] link to attempt to resolve the DNS name of that exporter. Clicking the [Add] link will pop up a dialog box to select a profile to associate this stream.

Additional details such as packets and bits received are displayed.

6.5 Collector Tab

The collector tab displays real time data related to collectors. The number beside the tab represents the number of collectors in an alarm state.

Any row highlighted in red indicates an issue with that collector. The user has many options to manage collectors.

- Add details, or a note, to a collector by clicking the [edit] link
- Set a threshold to alarm if a certain packet rate is exceeded. This can be done by clicking the [set] link in the status column.
- [Delete] a collector from all profiles
- [Add] the collector to a profile
- [Remove] the collector from a profile
- [View] a profile (this will leave the current tab)
- Click [Check] to try and resolve the DNS name

6.6 Profiles Tab

Profiles define how packets come in and go out of the replicator. The profiles tab lists all configured profiles and their current state. The number beside the tab represents the number of profiles in an alarm state.

Real time data such as inbound and outbound traffic is shown as well as total policies, exporters, and collectors. Profiles highlighted in red indicate an issue with the policy. [edit] a policy will provide details on why that policy is in an alarm state.

The following actions are available from the profile tab:

- [Delete] the profile which will stop packets to the collectors configured for the defined exporter streams
- [Edit] the current profile
- Create a new profile by clicking [New]

6.6.1 Creating a new Profile

Clicking the [new] link will display the new profile wizard. A profile consists of a name, an In Port, and an Out Port. The ports defined are UDP ports that direct the replicator on how to listen for and send packet traffic.

In most cases, users want to replicate traffic from any exporter sending packets on the in port. However if this isn't desired, users can uncheck the match all and specify a network and CIDR (e.g. 192.168.2.0/24) to match a certain set of exporters.

Multiple policies can be later added to a profile. See the next section for more information.

After the information has been completed correctly, click save.

6.6.2 Editing a Profile

Clicking the [edit] link next to a profile will display its current configuration. The name, description, real time data, policies, exporters, and collectors are displayed.

Profiles are not considered complete until there are at least one exporter and one collector configured. From this interface the following actions are available:

- [all] returns the user to the list of profiles
- [new] launches the new profile wizard
- [delete] permanently removes the profile
- [edit] launches the edit profile modal to modify the name, the in port, and out port, or the description modal to modify the description of the profile
- [add policy] launches the add policy modal
- [add collector] launches the add collector modal
- [add exporter] launches the add exporter modal
- [remove] removes the policy, exporter, or collector depending on which one is clicked

6.7 Settings Tab

The settings tab is where users can modify the configuration of the replicator. Additional settings are available in the interactive prompt. Certain features can be enabled and disabled by checking or unchecking the enabled checkbox in the setting column. The settings column also contains the current value of the option. The description column explains what the option does.

If the user enters an invalid value, the replicator will provide feedback informing the user what entry values would be valid.

6.8 Status LED

The Status LED lists the state of all services and profiles and provides a status of each. It is a high level view of the health of the replicator's processes.

CHAPTER 7

Interactive Mode Commands

At any time, running the command help, help <command>, <command> ?, or ? will display help in the interface.

7.1 acknowledge

The Flow Replicator actively monitors the state of exporters and collectors. If either one is in an alarm state, the acknowledge command can be used to stop the Flow Replicator from sending notifications about unavailable resources.

- *acknowledge <exporter|collector> ip_address:port*

```
REPLICATOR> acknowledge exporter 10.1.1.2
REPLICATOR> acknowledge collector 10.1.1.1:2055
```

7.1.1 backup

Creates a backup of the database in /home/replicator/backups/<filename>. Lists the files in /home/replicator/backups with the date they were last accessed. Backups can be restored using the restore command. Backup names can not contain any spaces.

- *backup [filename]*
- *show backups [filename]*
- *restore [filename]*

```
REPLICATOR> backup replicator_backup
REPLICATOR> show backups
REPLICATOR> restore replicator_backup
```

7.2 collector

The collector command is used to add or remove collectors from profiles.

- *collector <add\remove> collector_ip profile*

```
REPLICATOR> collector add 10.1.1.1 maineStreet
REPLICATOR> collector remove 10.1.1.1 maineStreet
```

The allremove directive will remove the specified collector IP address from all profiles.

- *collector allremove collector_ip*

```
REPLICATOR> collector allremove 10.1.1.1
```

Collectors receive replicated packets. Some collectors may not be able to handle high volume. Use this option to set or remove a packet per second threshold.

- *collector threshold collector_ip threshold*

```
REPLICATOR> collector threshold 10.1.1.1 100000
```

7.3 exporter

The exporter command is used to add or remove exporters from profiles.

- *exporter <add\remove> exporter_ip profile*

```
REPLICATOR> exporter add 10.1.1.2 maineStreet
REPLICATOR> exporter remove 10.1.1.2 maineStreet
```

The allremove directive will remove the specified exporter IP address from all profiles.

- *exporter allremove exporter_ip*

```
REPLICATOR> exporter allremove 10.1.1.2
```

The noprofile directive will list all exporters actively sending packets to the Flow Replicator that are not configured in any profiles.

It's recommended to either add these exporters to a profile or configure them to stop sending packets to the Flow Replicator.

7.4 exporter noprofile

```
REPLICATOR> exporters noprofile

+-------------------------------------------------------------------------+
| 10.1.73.1         Wed Jul 16 11:06:37 2014         1 packet(s)
| 10.1.29.60        Wed Jul 16 11:06:38 2014         2 packet(s)
| 10.202.0.103      Wed Jul 16 11:06:39 2014         5 packet(s)
| 10.200.10.1       Wed Jul 16 11:06:39 2014        32 packet(s)
| 172.20.124.41     Wed Jul 16 11:06:33 2014         1 packet(s)
+-------------------------------------------------------------------------+
Done in 0.035998 secs
```

This list contains the IP Address of the exporter, the last time stamp a packet was received, and the number of packets counted since the last packet summary. It is possible to have 0 packet(s) for exporters that export data infrequently.

7.5 license

The license command is used to manage the Flow Replicator license key.

To generate a license key, Plixer or the reseller will need the Flow Replicator's unique machine ID. The machine ID is displayed when issuing the license check command. The following command can be used to show licensing details.

- *license <check|status>*

```
REPLICATOR> license check

      Machine ID : 5YZ6XEPV66C766369M8DBN2A
Licensed Version : 3.1
   Licensed Type : valid
      Expiration : Thu Jul 28 2016

License expires in 730 day(s)
```

The license key can be configured on the Flow Replicator using the license set command.

- *license <set|update>*

```
REPLICATOR> license set
```

When applying the license key, it must be one continuous string without any line feeds or carriage returns on the same line as the license=

```
[replicator]
engine=sqlite
dbname=/home/replicator/html/db/replicator.db
user=
pass=
license=Nb7RuYhxJWxUv9u+nTdHCnRj5R9EiXQv5qDS9WO41jC4XBBYkErNZ6Q+Oi+Q+6uGwfaQJZO6QzE3wjgWsf2CfqlCp3Sd0
 ↪txz6yhFurK7Cz4Js1kuraTt96Q1pRru9zCk5gUxbNjISzI3BlY75eMMDddTFv2XKJRxzDe8CK8N1Ov4Okkod1gx9tWW2xFToAJr
 ↪HgVYMahQgDjPHhbuq2ft2HA1iuhRZU2q0Bt8TbSy+6CmvKLe7tSqht5V9bSLYQSdaJ1/
 ↪gntqAJaa4dGG4fBGmDgK30zLkC+OEFm402axzCmQ==
```

In the new window, under "license=" paste in your license key. Press CTRL+x to save.

Issuing either the license check or show status command will verify the key is properly installed. Contact technical support to acquire a new license key.

7.6 notate

The notate command can be used to add a description to a profile or IP Address. The description does not require enclosure quotation marks.

- *notate <profile|ip> <profile_name|ip_address> description*

```
REPLICATOR> notate profile maineStreet A fun and happy place
REPLICATOR> notate ip 10.1.1.2 my awesome router
```

This description will show up in various reports generated by the show command.

7.7 password

The password command will change the password used for the replicator user in interactive mode and the admin user in the web interface.

- *password <interactive|webui>*

```
REPLICATOR> password interactive
(current) UNIX password:
New password:
Retype new password:

Successful password changes will be applied to the next log in.
```

This password is used when logging in interactively or to the web interface.

7.8 policy

The policy command manages what exporters are automatically included or excluded in profiles. Policy inclusion policies are checked first, then exclusion policies. Policies are defined in subnet/cidr notation.

- *policies <add|remove> subnet/cidr profile <include|exclude>*

```
REPLICATOR> policy add 192.168.0.0/16 maineStreet include
REPLICATOR> policy add 192.168.2.0/24 maineStreet exclude
REPLICATOR> policy remove 192.168.2.0/24 maineStreet
```

The include/exclude option is only required if using the add directive.

Collectors are not affected by policies.

7.9 profile

The profile command is used to add, update, remove, enable, disable, and rename profiles.

The name, listening port, and sending port are required when adding or updating a profile.

- *profile <add|update> name listen_port send_port*

```
REPLICATOR> profile add maineStreet 2002 2055
REPLICATOR> profile update maineStreet 2003 2056
```

Removing a profile will also remove any policies assigned to it. However, other profiles will remain unmodified.

Disabling a profile will keep its settings, policies, exporters, and collectors intact. However, replication will not occur.

- *profile <remove|disable|enable> name*

```
REPLICATOR> profile remove maineStreet
REPLICATOR> profile disable maineStreet
REPLICATOR> profile enable maineStreet
```

Profiles can also be renamed with the rename directive. Only the name of the profile will be updated. Use profile update to change other details such as the sending port or listening port.

- *profile rename old_name new_name*

```
REPLICATOR> profile rename maineStreet streetOfMaine
```

7.9.1 Singularity (spoofing)

Singularity mode will replicate packets as the replicator IP instead of the original exporter's IP. This allows users to combine packets from multiple exporters into a single exporter IP.

```
REPLICATOR> profile singularity maineStreet <enable|disable>
```

By default, the replicator will replicate packets using the original sources of those packets as the exporter.

7.10 rebuild

The rebuild command is only necessary when replication services are down and the administrator wishes to rewrite the internal configuration. Otherwise, the Flow Replicator manages all configurations as real time changes are detected.

```
REPLICATOR> rebuild
```

The sampcfg command is an alias to the rebuild command.

7.11 role

The role command is used when setting up a fault tolerant environment.

- *role set ha master <ip_address>*
- *role set ha off*
- *role set ha on <priority> <virtual_ip> <ifname> <master|backup>*
- *role set primary*
- *role set secondary <primary_replicator_ip:listener_port> [timeout]*
- *role test <ha|secondary>*

```
REPLICATOR> role set master <ip_address>
REPLICATOR> role set ha off
REPLICATOR> role set ha on 101 10.1.4.223 eth0 master
REPLICATOR> role set primary
REPLICATOR> role set secondary 10.1.4.66:2002 10
REPLICATOR> role test secondary
```

Reference the section on fault tolerance for more information.

7.12 setting

The setting command manages the global configuration for the Flow Replicator. Features can be enabled, disabled, and set.

- *setting set name value*

- *setting <enable|disable> name*

```
REPLICATOR> setting set metricsSent 10.1.4.66:2003
REPLICATOR> setting disable convertSyslog
```

Use the show setting command to get a list of settings in the global configuration.

7.13 show

The show command generates reports based on configuration settings and real time data.

- *show alarm [filter]*
- *show asset [filter]*
- *show collector [filter]*
- *show config*
- *show exporter [filter]*
- *show profile [filter]*
- *show realtime [filter]*
- *show setting [filter]*
- *show status*

Most show commands also have a [filter] option which will only display details that match the filter. The entire report is displayed if no filter is included.

The list and sh commands are aliases to the show command.

7.13.1 alarm

Lists exporters that have stopped sending data to the replicator and collectors that are no longer reachable by the replicator.

- *show alarm [filter]*

```
REPLICATOR> show alarm
REPLICATOR> show alarm 10.1.4
```

7.13.2 asset

Generates a report detailing IP addresses, whether the IP address is an exporter and/or collector, dns names, and descriptions.

- *show asset [filter]*

```
REPLICATOR> show asset
REPLICATOR> show asset plxr.local
```

7.13.3 collector

Generates a report showing the collector(s) IP address, dns name, description, and which profiles currently include the collector(s).

- *show collector [filter]*

```
REPLICATOR> show collector
REPLICATOR> show collector 10.1.1.1
```

7.13.4 config

Lists all commands necessary to rebuild all profile settings.

- *show config*

```
REPLICATOR> show config
```

7.13.5 exporter

Generates a report showing the exporter(s) IP address, dns name, description, and which profiles are currently including the exporter(s).

- *show exporter [filter]*

```
REPLICATOR> show exporter
REPLICATOR> show exporter 10.1.2.5
```

7.13.6 profile

Lists profiles and all policies, exporters, and collectors associated.

- *show profile [filter]*

```
REPLICATOR> show profile
REPLICATOR> show profile maineStreet
```

7.13.7 realtime

Peers into the live stream and shows statistics of exporters, in and out statistics, and CPU usage.

- *show realtime [filter]*

```
REPLICATOR> show realtime
REPLICATOR> show realtime 192.168
```

Press CTRL+C to exit the realtime report.

7.13.8 setting

Displays a list of all global configuration settings, the current values, and whether they are enabled or disabled.

- *show setting [filter]*

```
REPLICATOR> show setting
REPLICATOR> show setting metric
```

7.13.9 status

Lists all replicator services and licenses, and shows the status of each.

- *show status*

```
REPLICATOR> show status

+-------------------------------------------------------------------------+
| Converting Syslog              ACTIVE
| Replicating Port 2002          ACTIVE
| Replicating Port 515           ACTIVE
| Replicating Port 2003          ACTIVE
| Replicator Vitalizer           ACTIVE
| Replicator License             License expires in 730 day(s)
| Replicator Heartbeat           ACTIVE
+-------------------------------------------------------------------------+
+ process checking took 0.079859 secs

Done in 0.270207 secs
```

7.14 snoop

The snoop command can be used to verify that packets are being received by or sent from the replicator for a certain IP address.

- *snoop <ip_address>*

```
REPLICATOR> snoop 10.1.1.1
```

Press CTRL+C to exit the snoop command.

7.15 system

The system command is used to change state of the replicator. The directive change is used to change the host name or IP address.

- *system <change|restart|shutdown>*

```
REPLICATOR> system change
REPLICATOR> system restart
REPLICATOR> system shutdown
```

- *system virtualip enable <ifname> <virtual_ip>*

- *system virtualip disable <ifname> <virtual_ip>*

```
REPLICATOR> system virtualip enable eth0 10.1.4.223
REPLICATOR> system virtualip disable eth0 10.1.4.223
```

These commands allow users to add or remove a virtual IP address for High Availability (HA) configurations. For more information see the role command.

CHAPTER 8

Advanced Configuration

8.1 Converting syslogs to IPFIX

The Flow Replicator is capable of converting syslogs into IPFIX. A global configuration setting convertSyslog specifies what UDP port to convert.

```
REPLICATOR> show setting convertSyslog

+------------------------------+------------------------------+----------+
| convertSyslog                | 514                          | Enabled
|
|  When enabled, syslogs sent to the specified port will be converted to
|  IPFIX and sent to the sender port(s) in the profile.
+------------------------------+------------------------------+----------+
Done in 0.168892 secs
```

By default, this functionality is disabled.

```
+------------------------------+----------------------------------------+
| syslog                        | IN PORT 514 -> OUT PORT 9995
+------------------------------+----------------------------------------+

 Policies                     Exporters     ->    Collectors
 (include) 0.0.0.0/0          10.1.1.242          10.1.10.1
                              10.1.1.249          10.1.4.101
                              10.1.1.252          10.1.4.19
                              10.12.1.98          10.1.4.222
                              10.3.1.1            10.1.4.93
                              192.168.21.254
                              24.39.1.172

+------------------------------+----------------------------------------+
Done in 0.009376 secs
```

If convertSyslog is disabled, a syslog received from one of the exporters on port 514 is replicated as a syslog to all the collectors on port 9995.

If convertSyslog is enabled and set to 514, the same syslog received from one of the exporters on port 514 is converted and replicated as IPFIX to all the collectors on port 9995. The syslog will not be replicated as a syslog.

8.2 Fault Tolerance

A Second Flow Replicator (SFR) can be set up to provide a fault tolerant environment in case the Primary Flow Replicator (PFR) goes offline.

A fault tolerant environment will function with either a virtual or hardware appliance. There are two methods available for Fault Tolerance Environments.

The first method (i.e. traditional method) is used when two replicators exist on different subnets. The second method (i.e. High Availability) can be used when two replicators are on the same subnet. The first method, if desired, can be used when replicators are on the same subnet.

8.3 Traditional Configuration

8.3.1 How it works

The secondary flow replicator (SFR) actively monitors the state of the primary flow replicator (PFR) and frequently synchronizes its database with the settings from the primary.

When a Flow Replicator is in secondary mode, it will no longer maintain its current configuration and any current configuration is lost. At any time, the SFR will not allow any configuration changes. With the exception of the role and show commands, all profile and global configuration must be changed on the PFR.

If the SFR determines that the PFR is offline, the SFR attempts to contact collectors. If the collectors are reachable, the SFR takes over replication based on the configured profiles in the last known good synchronized database. If the collectors are not reachable, the SFR maintains SFR status until a collector or the PFR is reachable.

Additionally, the SFR continues to monitor the state of the PFR. When the PFR is reachable, replication on the SFR stops. The SFR synchronizes any updates to the PFR database and begins the process of monitoring the state of the PFR.

8.3.2 Requirements

The following requirements must be met to set up a Fault Tolerant Environment:

- A Primary and Secondary Flow Replicator
- Each exporter configured to send data to both the Primary and Secondary Flow Replicator IP Addresses

A Primary and Secondary Server can be a mix of Virtual and Hardware Appliances. Configure the primary with the desired profiles, global configuration settings, and verify the flow collectors are receiving the replicated data.

Note: The secondary flow replicator requires a fault tolerance license key. To acquire a fault tolerance license key contact plixer directly.

8.3.3 Configuring a secondary flow replicator

Deploy a second Flow Replicator based on the installation instructions earlier in this guide. The role command is used when setting up a fault tolerant environment. Run the following command on the secondary flow replicator:

- *role set secondary <primary_replicator_ip:listener_port> [timeout]*

Replace the <*primary_replicator_ip*> with the IP Address of the primary server and <*listener_port*> with a port on which the primary is actively listening for packets.

The timeout represents the number of consecutive missed polls before the secondary takes the role as the primary. The default is 2 polls if timeout is not passed in.

```
REPLICATOR> role set secondary 10.1.4.66:2002 3
```

8.3.4 Testing a secondary replicator

Once the flow replicator has been successfully set to secondary, the last step is to verify connectivity to the primary.

- *role test secondary*

```
REPLICATOR> role test secondary
```

The fault tolerant environment is active if the test runs successfully.

8.3.5 Turning off Fault Tolerance

At anytime, a secondary Flow Replicator can become an independent Flow Replicator by issuing the role set primary command.

```
REPLICATOR> role set primary
```

8.4 High Availability

8.4.1 How it works

Packets are sent to a virtual IP Address that both the Primary and Secondary Replicator are configured to listen for traffic. The primary and secondary configurations each contain a priority so the replicators are aware of their role.

When the primary replicator fails, the packets are instantly picked up by the secondary and forwarded to the collectors configured as part of the replicator profiles.

Once the primary is active again, the primary starts forwarding the packets to the configured collectors. This method can be used to set up redundancy beyond a primary and secondary role.

8.4.2 Requirements

The following requirements must be met to set up a Fault Tolerant Environment:

- A Primary and Secondary Flow Replicator
- Each exporter configured to send data to a single virtual IP Address

A Primary and Secondary Server can be a mix of Virtual and Hardware Appliances. Configure the primary with the desired profiles, global configuration settings, and verify the flow collectors are receiving the replicated data.

8.4.3 Configuring High Availability

1. On the master replicator, run the following commands to add the virtual IP address:

```
REPLICATOR> system virtualip enable eth0 <virtual_ip_here>
```

2. Next, declare the IP address of the master:

```
REPLICATOR> role set ha master <primary_ip>
```

3. Next, execute the following to enable High Availability on the primary:

```
REPLICATOR> role set ha on 101 <virtual_ip_here> eth0 master
```

4. The final step for configuring the master replicator is to edit /etc/plixer.ini and delete the line that starts with "primaryip="
5. On the secondary replicator, run the following commands to add the virtual IP address:

```
REPLICATOR> system virtualip enable eth0 <virtual_ip_here>
```

6. Next, declare the IP address of the master:

```
REPLICATOR> role set ha master <primary_ip>
```

7. Lastly, execute the following to enable High Availability on the secondary:

```
REPLICATOR> role set ha on 100 <virtual_ip_here> eth0 backup
```

8.4.4 Testing a secondary replicator

1. Run the following command on both the primary and secondary replicators:

```
REPLICATOR> role test ha
```

2. Purposely shutdown the primary replicator and ping the virtual IP. If a response is received, high availability has been properly configured.
3. Restart the primary replicator.

8.4.5 Turning off Fault Tolerance

To turn off High Availability Fault Tolerance, simply run the following command on the primary and secondary replicators.

```
REPLICATOR> role set ha off
REPLICATOR> system virtualip disable eth0 <virtual_ip_here>
```

8.5 Closed Networks with No Gateway

This configuration currently requires root access and a reboot of the Replicator.

In the */etc/sysctl.conf file*, there is a setting called *net.ipv4.all.rp_filter*. When this value is set to 1 (i.e. net.ipv4.all.rp_filter = 1), all packets that come into an interface that the host doesn't have a route are filtered (dropped).

For example: A user has a closed network of 192.168.250.x with no gateway. The packets sent to the Replicator that were not 192.168.0.0/24 are dropped and never replicated to any configured collectors.

The solution is to edit the */etc/sysctl.conf* file and change *net.ipv4.all.rp_filter = 1* to *net.ipv4.all.rp_filter = 0*. A reboot is required.

CHAPTER 9

Application Programming Interface

9.1 Authentication

Users must authenticate before using any of the API calls. This is done through the resource URL /api/1/login. To log out manually, use the resource URL /api/1/logout. Log out is not required to end sessions. Sessions are automatically expired after 30 minutes.

9.1.1 Log In

Used to start a new API session.

GET request

Note: Legacy API – Deprecated- recommend using POST

Resource URL

https://{[}replicator{]}/api/1/login/{[}sha2]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[sha2]	The interactive (CLI) mode's password converted to sha2 512-bit hex

Example Request

https://10.30.17.131/api/1/login/098f6bcd4621d373ca098f6bcd4621d563cade4e832629b4f6d098f4bcd4621d373cade4e831627b4f6e4e

POST request

Used to start a new API session.

Resource URL

POST https://{[}replicator{]}/api/1/login/

```
{
   "user" : "admin"
   "passwd": "[passwd]"
}
```

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[passwd]	The sha512 hash of the password for admin

Example Response

```
{
  "path": "login successful.",
  "result": "success"
}
```

9.1.2 Log Out

Used to manually stop an API session.

Resource URL

https://{[}replicator{]}/api/1/logout

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance

Example Request

https://10.30.17.131/api/1/logout

Example Response

```
{
  "path": "logout successful.",
  "result": "success"
}
```

9.2 Acknowledge

Used to acknowledge collectors and exporters in an alarm state.

Resource URL

https://{[}replicator{]}/api/1/acknowledge/{[}entity{]}/{[}ip:port]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[entity]	valid options are collector or exporter
[ip:port]	The IP Address and Port of the [entity] to acknowledge

Example Request

https://10.30.17.131/api/1/acknowledge/collector/10.1.10.4:2055

Example Response

```
{
   "description": "Success: collector '10.1.10.4:2055' acknowledged.",
   "result": "success"
}
```

9.3 Collector

Perform actions to manage collectors in profiles.

RESOURCE URL

Use this resource to manage a collector for an individual profile.

https://{[}replicator{]}/api/1/collector/{[}action{]}/{[}ip{]}/{[}profile]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[action]	valid actions are add or remove
[ip]	The IP address of the collector
[profile]	The profile to perform the [action] for the [ip]

Example Request

https://10.30.17.131/api/1/collector/add/10.1.10.60/myprofile

Example Response

```
{
  "description": "Success: Collector [10.1.10.60] -> Profile [myprofile]",
  "result": "success"
}
```

RESOURCE URL

Use this resource to remove a collector from all profiles.

https://{[}replicator{]}/api/1/collector/allremove/{[}ip]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[ip]	The IP of the collector to remove from all profiles

Example Request

https://10.30.17.131/api/1/collector/allremove/10.1.10.4

Example Response

```
{
  "description": "Success: Collector [10.1.10.4] <- All Profiles",
  "result": "success"
}
```

9.4 DNSCheck

Performs a DNS check on an IP Address

RESOURCE URL

Use this resource to manually perform a DNS Name resolve on an IP address.

https://{[}replicator{]}/api/1/dnscheck/{[}ip]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[ip]	The IP address to resolve

Example Request

https://10.30.17.131/api/1/dnscheck/10.1.1.3

Example Response

```
{
  "dnscheck": {
  "addr": "10.1.1.3",
  "expiry": {
  "epoch": 1462801036,
  "time": "Mon May 9 09:37:16 2016"
  },
  "name": "newexp.plxr.local",
  "resolveTime": 0.04485,
  "status": "success"
  }
}
```

9.5 Exporter

Perform actions to manage exporters in profiles.

RESOURCE URL

Use this resource to manage an exporter for an individual profile.

https://{[}replicator{]}/api/1/exporter/{[}action{]}/{[}ip{]}/{[}profile]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[action]	valid actions are add or remove
[ip]	The IP address of the exporter
[profile]	The profile to perform the [action] for the [ip]

Example Request

https://10.30.17.131/api/1/exporter/add/10.1.1.1/myprofile

Example Response

```
{
  "description": "Success: Exporter [10.1.1.1] -> Profile [myprofile]",
  "result": "success"
}
```

RESOURCE URL

Use this resource to remove an exporter from all profiles.

https://{[}replicator{]}/api/1/exporter/allremove/{[}ip]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[ip]	The IP of the exporter to remove from all profiles

Example Request

https://10.30.17.131/api/1/exporter/allremove/10.1.1.4

Example Response

```
{
   "description": "Success: Exporter [10.1.1.4] <- All Profiles",
   "result": "success"
}
```

RESOURCE URL

This resource identifies exporters sending data to the replicator that are not members of any profiles.

https://{[}replicator{]}/api/1/exporter/noprofile/{[}filter]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[filter]	The filter to apply against the list of exporters in no profiles. To get the entire list, pass a 0.

Example Request

https://10.30.17.131/api/1/exporter/noprofile/0

Example Response

```
{
  "noprofile": {
    "10.30.17.131": {
      "in_o_delta": 0,
      "in_o_rate": "0.0",
      "in_p_delta": 0,
      "in_p_rate": "0.0",
      "lastseen": "Tue Apr 19 09:55:18 2016",
      "out_o_delta": 1493,
      "out_o_rate": "49.8",
      "out_p_delta": 6,
```

(continues on next page)

(continued from previous page)

```
      "out_p_rate": "0.2",
      "unixtime": 1461074118
    },
    "total": 1
  }
}
```

9.6 License

Check the status of the current license

RESOURCE URL

This resource returns current license details.

https://{[}replicator{]}/api/1/license/check

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance

Example Request

https://10.30.17.131/api/1/license/check

Example Response

```
{
    "daysLeft" : "365 day(s)",
    "expiration" : "Thu May 18 2017",
    "licensedType" : "eval",
    "licensedVersion" : "16.6",
    "machineID" : "6YZ6XEPTA66JA6VHFPXG749B",
    "role" : "failover"
}
```

9.7 Notate

Users can add descriptions for profiles, collectors, and exporters.

RESOURCE URL

https://{[}replicator{]}/api/1/notate/{[}entity{]}/{[}identity{]}/{[}description]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[entity]	valid options are profile or ip.
[identity]	Use a profile name when [entity] profile is specified, or the IP address of a collector or exporter when [entity] ip is specified.
[description]	the description of the entity. Use standard ASCII characters the are URI compatible.

Example Request

https://10.30.17.131/api/1/notate/profile/myprofile/My_Fantastic_Description

Example Response

```
{
    "description": "Success: Profile 'myprofile' has a new description",
    "result": "success"
}
```

9.8 Policies

Manage policies that are associated to profiles.

RESOURCE URL

https://{[}replicator{]}/api/1/policies/{[}action{]}/{[}network{]}/{[}cidr{]}/{[}profile{]}/{[}incexc]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[action]	valid options are add or remove
[network]	the address to the network (e.g. 172.17.0.0)
[cidr]	the CIDR to the network specified (e.g. 16)
[profile]	the name of the profile to perform the [action]
[incexc]	valid options are include or exclude

Example Request

https://10.30.17.131/api/1/policies/add/10.1.20.0/16/myprofile/include

Example Response

```
{
    "description": "Success: Policy [10.1.20.0/16] -> Profile [myprofile]",
    "result": "success"
}
```

9.9 Profile

Manage characteristics and behaviors of profiles.

RESOURCE URL

Use this resource to rename an existing profile.

https://{[}replicator{]}/api/1/profile/rename/{[}current{]}/{[}new]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[current]	The name of the profile to rename
[new]	The new name of the profile

Example Request

https://10.30.17.131/api/1/profile/rename/oldprofile/myprofile

Example Response

```
{
    "description": "Success: Profile 'oldprofile' is now 'myprofile'",
    "result": "success"
}
```

RESOURCE URL

Use this resource to toggle the singularity flag.

https://{[}replicator{]}/api/1/profile/singularity/{[}name{]}/{[}action]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[name]	the name of the profile to toggle singularity
[action]	valid options are enable and disable

Example Request

https://10.30.17.131/api/1/profile/singularity/myprofile/enable

Example Response

```
{
    "description": "Success: Profile 'myprofile' singularity enabled",
    "result": "success"
}
```

RESOURCE URL

Use this resource to enable, disable, or remove a profile

https://{[}replicator{]}/api/1/profile/{[}action{]}/{[}name]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[action]	valid options are enable, disable, and remove
[name]	the name of the profile to perform the [action]

Example Request

https://10.30.17.131/api/1/profile/disable/myprofile

Example Response

```
{
    "description": "Success: Profile myprofile has been set to 'disable'",
    "result": "success"
}
```

RESOURCE URL

Use this resource to create or update an existing profile.

https://{[}replicator{]}/api/1/profile/{[}action{]}/{[}name{]}/{[}listeningport{]}/{[}sendingport]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[action]	valid options are add and update
[name]	The name of the profile
[listeningport]	The UDP port for this profile to listen for incoming packets
[sendingport]	The UDP port for this profile to send packets out

Example Request

https://10.30.17.131/api/1/profile/add/myprofile/2055/4739

Example Response

```
{
    "description": "Success: Profile 'myprofile' has been added and enabled.",
    "result": "success"
}
```

9.10 Rebuild

Force the replicator to rebuild its configuration immediately instead of waiting for the replicator to automatically do it.

RESOURCE URL

https://{[}replicator{]}/api/1/rebuild

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance

Example Request

https://10.30.17.131/api/1/rebuild

Example Response

```
{
    "description": "rebuild request submitted",
    "result": "success"
}
```

9.11 Role

When a replicator is configured as a secondary/backup, send a test to verify its configured properly.

RESOURCE URL

https://{[}replicator{]}/api/1/role/test/secondary

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance

Example Request

https://10.30.17.131/api/1/role/test/secondary

Example Response

```
{
    "description": "!!! This replicator is the primary !!!",
    "result": "error"
}
```

9.12 Setting

Manage global settings of the replicator.

RESOURCE URL

Use this resource to set values for specified settings

https://{[}replicator{]}/api/1/setting/set/{[}name{]}/{[}value]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[name]	the exact name of the setting to modify
[value]	the value to set the setting specified in [name]

Example Request

https://10.30.17.131/api/1/setting/set/metricssent/10.1.2.3:2055

Example Response

```
{
    "description": "Success: setting 'metricssent' has been set to '10.1.2.3:2055'",
    "result": "success"
}
```

RESOURCE URL

Use this resource to enable or disable the specified setting.

https://{[}replicator{]}/api/1/setting/{[}action{]}/{[}name]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[action]	valid options are enable or disable
[name]	the exact name of the setting to modify

Example Request

https://10.30.17.131/api/1/setting/disable/metricssent

Example Response

```
{
    "description": "Success: setting 'metricssent' has been set to 'disable'",
    "result": "success"
}
```

9.13 Show

Shows configuration and realtime information from the replicator.

RESOURCE URL

Use this resource to view the current configuration of the replicator.

https://{[}replicator{]}/api/1/show/config

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance

Example Request

https://10.30.17.131/api/1/show/config

Example Response

```
{
 "api": {
  "collector": [
   "/api/1/collector/add/10.30.1.20/benchmark-20",
   ...
  ],
  "notate": [
   "/api/1/notate/ip/10.1.4.101/ej-win2012 install test machine (fresh)",
   ...
  ],
  "policy": [
   "/api/1/policies/add/10.1.1.252/32/buildqa/include",
   ...
  ],
  "profile": [
   "/api/1/profile/add/frandev/2002/2055",
   ...
  ]
 },
 "cli": {
  "collector": [
   "collector add 10.30.1.20 benchmark-20",
   ...
  ],
  "notate": [
   "notate ip 10.1.4.101 ej-win2012 install test machine (fresh)",
   ...
  ],
  "policy": [
   "policy add 10.1.1.252/32 buildqa include",
   ...
  ],
  "profile": [
   "profile add frandev 2002 2055",
   ...
  ]
 }
}
```

RESOURCE URL

Use this resource to view the current status of the replicator.

https://{[}replicator{]}/api/1/show/status

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance

Example Request

https://10.30.17.131/api/1/show/status

Example Response

```
{
   "converting syslog" : "active",
   "ipfixify system metrics" : "active",
   "replicating port 2002" : "active",
   "replicating port 2003" : "active",
   "replicating port 2055" : "active",
   "replicator api" : "active",
   "replicator license" : "259 day(s)",
   "replicator monitor" : "active",
   "replicator vitalizer" : "active",
   "result" : "success",
   "version" : "v17.12.19.2255"
}
```

RESOURCE URL

Use this resource to see realtime data from the replicator.

https://{[}replicator{]}/api/1/show/realtime/{[}filter]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[filter]	Currently filter is not supported. Pass in 0 to get the entire state of statistics.

Example Request

https://10.30.17.131/api/1/show/realtime/0

Example Response

```
{
 "collector": {
  "10.1.4.19": {
   "in": {
    "octets": {
     "delta": 0,
     "rate": 0
    },
     "packets": {
     "delta": 0,
     "rate": 0
    }
   }
  },
  "10.1.4.20": {
```

(continues on next page)

(continued from previous page)

```
   ...
  },
},
"exporter": {
 "10.30.17.131": {
  "awareness": {
   "last_epoch": 1461175921,
   "last_timestamp": "2016-04-20 14:12:01"
  },
  "in": {
   "octets": {
    "delta": 0,
    "rate": "0.0"
   },
   "packets": {
    "delta": 0,
    "rate": "0.0"
   }
  },
  "out": {
   "octets": {
    "delta": 0,
    "rate": "0.0"
   },
   "packets": {
    "delta": 0,
    "rate": "0.0"
   }
  },
  "profiles": [
   "benp"
  ]
 }
},
"pair": {
 "10.30.17.131 -> 10.1.4.19": {
  "octets": {
   "delta": 0,
   "rate": 0
  },
  "packets": {
   "delta": 0,
   "rate": 0
  }
 },
 "10.30.17.131 -> 10.1.4.20": {
  ...
 },
 "profile": {
  "benp": {
   "in": {
    "octets": {
     "delta": 0,
     "rate": 0
    },
    "packets": {
     "delta": 0,
```

(continues on next page)

(continued from previous page)


```
      "rate": 0
     }
    },
    "out": {
     "octets": {
      "delta": 0,
      "rate": 0
     },
    "packets": {
     "delta": 0,
     "rate": 0
    }
   }
  }
},
"stats": {
 "totals": {
  "collectors": 2,
  "exporters": 1,
  "pairs": 2,
  "profiles": {
   "disabled": 3,
   "enabled": 67,
   "total": 70
  }
 }
},
"system": {
 "cpu": "0"
}
```


RESOURCE URL

Use this resource to show various configured information from the replicator.

https://{[}replicator{]}/api/1/show/{[}entity{]}/{[}filter]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[entity]	Valid options are alarms, assets, collectors, config, exporters, profile, realtime, settings, and status
[filter]	A custom keyword match filter can be specified. If no filter is necessary, pass 0.

Example Request

https://10.30.17.131/api/1/show/collector/10.1.10.1

Example Response


```
{
 "collector": {
  "10.1.10.1": {
   "acknowledged": [
    9996
   ],
   "description": "erpdev",
   "in_profiles": [
    "steady-replays"
```


(continues on next page)

(continued from previous page)

```
   ],
   "ip": "10.1.10.1",
   "name": null,
   "status": {
    "unreachable_port": [
     "9996"
    ],
   }
   "threshold": 10000
  }
 }
}
```

9.14 Threshold

Set thresholds on collectors to warn when replicated packets are exceeding the current packet per second replicated.

RESOURCE URL

https://{[}replicator{]}/api/1/collector/threshold/{[}collector{]}/{[}threshold]

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance
[collector]	The IP address of the collector
[threshold]	The threshold to set. It represents packets per second. Set this to a number higher then 0. If 0 is set, the threshold is removed.

Example Request

https://10.30.17.131/api/1/collector/threshold/10.1.5.2/10000

Example Response

```
{
  "description": "Success: Collector [10.1.5.2] threshold set to 10000",
  "result": "success"
}
```

9.15 Version

Returns version information about the replicator.

RESOURCE URL

https://{[}replicator{]}/api/1/version

Parameter	Description
[replicator]	The hostname or IP address of the Replicator Appliance

Example Request

https://10.30.17.131/api/1/version

Example Response

```
{
    "apiversion": 1,
    "bestipfixcollector": "Need an IPFIX Collector? Download Scrutinizer at https://
↪www.plixer.com",
    "build": "2016-04-01 08:34:05 -0400 (Fri, 01 Apr 2016)",
    "copyright": "Copyright (C) 2012 - 2018 Plixer - All rights reserved.",
    "name": "Plixer Replicator (TM) v16.4.1.1429 ",
    "yeswecan": "Replicate Anything!"
}
```

CHAPTER 10

Third Party Software Attributions

10.1 Licenses Directory

Required License Documentation can be found in /home/replicator/files/licenses

10.2 Third Party Attributions

Certain open source or other third-party software components are integrated and/or redistributed Replicator software. The licenses are reproduced here in accordance with their licensing terms, these terms only apply to the libraries themselves, not Replicator software.

10.2.1 Backbone.js

https://github.com/jashkenas/backbone/blob/master/LICENSE Copyright (c) 2010-2017 Jeremy Ashkenas, DocumentCloud Licensed under the MIT License – see Licenses Directory

10.2.2 C3.js

https://github.com/c3js/c3/blob/master/LICENSE Copyright (c) 2013 Masayuki Tanaka Licensed under the MIT License – see Licenses Directory

10.2.3 D3.js

https://github.com/d3/d3/blob/master/LICENSE Copyright (c) 2010-2014 2010-2017 Mike Bostoc Licensed under the BSD 3-clause License – see Licenses Directory

10.2.4 Hogan.js

https://github.com/twitter/hogan.js/blob/master/LICENSE Copyright (c) 2011 Twitter, Inc. Licensed under the Apache License 2.0 – see Licenses Directory

10.2.5 JQuery

https://jquery.org/license/ Copyright jQuery Foundation and other contributors, https://jquery.org This software consists of voluntary contributions made by many individuals. For exact contribution history, see the revision history available at https://github.com/jquery/jquery Licensed under the MIT License – see Licenses Directory

10.2.6 JQuery.floatThread.js

https://github.com/mkoryak/floatThead/blob/master/LICENSE Copyright (c) 2012-2017 Misha Koryak Licensed under the MIT License – see Licenses Directory

10.2.7 jsSHA

https://github.com/Caligatio/jsSHA/blob/master/LICENSE Copyright (c) 2008-2017 Brian Turek Licensed under the BSD 3-clause License – see Licenses Directory

10.2.8 JustGage

https://github.com/toorshia/justgage/blob/master/LICENSE Copyright (c) 2012-2015 Bojan Djuricic Licensed under the MIT License – see Licenses Directory

10.2.9 Raphaël

https://github.com/DmitryBaranovskiy/raphael/blob/master/license.txt Copyright © 2008-2013 Dmitry Baranovskiy, Sencha Labs Licensed under the MIT License – see Licenses Directory

10.2.10 UDP Samplicator

https://github.com/sleinen/samplicator/blob/master/COPYING Copyright (c) 2000-2015 Simon Leinen Licensed under the GNU GPL 2.0 – see Licenses Directory

10.2.11 Underscore.js

https://github.com/jashkenas/underscore/blob/master/LICENSE Copyright (c) 2009-2017 Jeremy Ashkenas, DocumentCloud and Investigative Reporters & Editors Licensed under the MIT License – see Licenses Directory

CHAPTER 11

Troubleshooting

11.1 Support

Technical support is available, provided maintenance is active.

11.2 Frequently Asked Questions

Q) I've configured my router to send flows to the replicator, but the replicator says it hasn't heard traffic from this device.

A) There may be a firewall or other network intrusion system preventing traffic from reaching the Flow Replicator. You can verify if the Flow Replicator is seeing traffic from that device by using the snoop <ip_of_network_device> command.

Q) Is there a way to get a list of exporters that are not in any profiles but still sending packets to the replicator?

A) Use the exporter noprofiles command:

```
REPLICATOR> exporters noprofile

+------------------------------------------------------------------------+
| 10.1.73.1         Wed Jul 16 11:06:37 2014         1 packet(s)
| 10.1.29.60        Wed Jul 16 11:06:38 2014         2 packet(s)
| 10.202.0.103      Wed Jul 16 11:06:39 2014         5 packet(s)
| 10.200.10.1       Wed Jul 16 11:06:39 2014        32 packet(s)
| 172.20.124.41     Wed Jul 16 11:06:33 2014         1 packet(s)
+------------------------------------------------------------------------+
Done in 0.035998 secs
```

Q) I want to send syslog notifications and IPFIX metrics generated by the replicator to multiple collectors but the setting seems to only support 1 collector. How do I configure it for multiple collectors?

A) This is possible by configuring the setting to send the notifications and/or IPFIX metrics back to the Flow Replicator. Then, create a profile to send the data on to one or more collectors. However, since one setting is for syslogs and the

other setting is for IPFIX, two different profiles are required. Assuming the Flow Replicator's IP is 10.1.4.66, below is an example of what both profiles would look like.

```
+-------------------------------+----------------------------------------+
| replicator_notifications      | IN PORT 514 -> OUT PORT 514
+-------------------------------+----------------------------------------+

 Policies                   Exporters     ->    Collectors
 (include) 0.0.0.0/0        10.1.4.66           10.1.10.1
                                                10.1.4.19
                                                10.1.4.20
                                                10.1.4.222
                                                10.1.4.93
                                                10.1.4.94

+-------------------------------+----------------------------------------+

+-------------------------------+----------------------------------------+
| replicator_metrics            | IN PORT 2003 -> OUT PORT 2056
+-------------------------------+----------------------------------------+

 Policies                   Exporters     ->    Collectors
 (include) 0.0.0.0/0        10.1.4.66           10.1.10.1
                                                10.1.4.19
                                                10.1.4.20
                                                10.1.4.222
                                                10.1.4.93
                                                10.1.4.94

+-------------------------------+----------------------------------------+
```

Q) Can a profile have the same listening and sending port?

A) Yes. However, the Flow Replicator will perform additional checking when adding collectors and exporters to verify a loop isn't introduced.

Q) What does it mean when I'm trying to add an exporter or collector and the Flow Replicator won't add it because it would create a loop?

A) A loop is when flows from a device are exported from and sent back to itself repeatedly. This requires that a device is both an exporter and collector. If the profile is configured with the same listening and sending ports, the packets will be sent back and forth continuously.

Q) How do I change the root password on the Flow Replicator?

A) Log into the Flow Replicator as the root user and issue the passwd command.

Q) How do I change the hostname and IP Address of the Flow Replicator?

A) There are two options to change the hostname and IP Address. First, log in as root and run the /home/replicator/conf/sethostname.sh command. Or, login as replicator and issue the system change command.

CHAPTER 12

Flow Replicator Change Log

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

KEY: ACTION: (Bug Ticket Number) description

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

12.1 Change Log History

12.1.1 Version 18.5 - 5/31/2018

(Bug 25633) Updated the EULA
(Bug 25257) The show config output now has quotes around profile names
(Bug 25292) Fixed an issue where Apache fails to start if SSL enabled/setup either from install script or enable_SSL.sh
(Bug 25653) Updated Licensing checks
(Bug 25770) Fixed an issue where an install script pointed to a previous version
(Bug 25828) Added the new online manual from docs.plixer.com
(Bug 25832) Fixed an issue where policy remove profile include/exclude gives internal error (500) but still removes profile
(Bug 26000) Fixed an issue where the refresh countdown timer would default to 1 day

12.1.2 Version 18.1 - 1/30/2018

New Features

Replicator now supports LDAP Authentication
Profiles now display in Alphabetical Order
Replicator now creates install and upgrade logs
Backup Process can now be run from the CLI
API calls can be made using https
Profile names can now use upper case letters and spaces
Added the ability restrict snoop command by port
The current Replicator Version is now displayed under the Status LED
Replicator database now runs on Postgres

Bug Fixes

(bug 22918) Improved performance and responsiveness of the Web Interface
(bug 23820) Replicator now replicates SNMP traps on low port numbers
(bug 23824) Exporters no longer false alarm at start up
(bug 23828) Exporters, no longer sending, and not in profiles no longer show up as 'exporters not in profiles' indefinitely
(bug 23904) Search filters are no longer lost on page refresh
(bug 23977) Fixed issue that could cause phantom collector alarms
(bug 24299) Upgrades no longer reset the Web Interface password

---

12.1.3 Version 17.6 - 7/14/2017

New Features

A Shiny, New, Live Data Web Interface
Fully supported and Documented API
Receive packets from multiple interfaces
replicate packets from multiple exporters as a single exporter IP

Bug Fixes

(Bug 19835) Licensing says expired one day before expiration date
(Bug 19396) ICMP drops have been added to iptables
(Bug 19285) semicolon at end of command yields unexpected results
(Bug 20848) Removing a policy doesn't remove the associated exporters
(Bug 21086) Replicator falsely reporting more packets inbound then out

12.1.4 Version 16.9 - 10/3/2016

New Features

A Shiny, New, Live Data Web Interface
Fully supported and Documented API
Receive packets from multiple interfaces
replicate packets from multiple exporters as a single exporter IP

Bug Fixes

(Bug 19835) Licensing says expired one day before expiration date
(Bug 19396) ICMP drops have been added to iptables
(Bug 19285) semicolon at end of command yields unexpected results
(Bug 20848) Removing a policy doesn't remove the associated exporters
(Bug 21086) Replicator falsely reporting more packets inbound then out