Cambium Networks™

USER INTERFACE GUIDE

QoE Appliance

Release 4.17

Reservation of Rights

Cambium reserves the right to make changes to any products described herein to improve reliability, function, or design, and reserves the right to revise this document and to make changes from time to time in content hereof with no obligation to notify any person of revisions or changes. Cambium recommends reviewing the Cambium Networks website for the latest changes and updates to products. Cambium does not assume any liability arising out of the application or use of any product, software, or circuit described herein; neither does it convey license under its patent rights or the rights of others. It is possible that this publication may contain references to, or information about Cambium products (machines and programs), programming, or services that are not announced in your country. Such references or information must not be construed to mean that Cambium intends to announce such Cambium products, programming, or services in your country.

Copyrights

This document, Cambium products, and 3rd Party software products described in this document may include or describe copyrighted Cambium and other 3rd Party supplied computer programs stored in semiconductor memories or other media. Laws in the United States and other countries preserve for Cambium, its licensors, and other 3rd Party supplied software certain exclusive rights for copyrighted material, including the exclusive right to copy, reproduce in any form, distribute and make derivative works of the copyrighted material. Accordingly, any copyrighted material of Cambium, its licensors, or the 3rd Party software supplied material contained in the Cambium products described in this document may not be copied, reproduced, reverse engineered, distributed, merged or modified in any manner without the express written permission of Cambium. Furthermore, the purchase of Cambium products shall not be deemed to grant either directly or by implication, estoppel, or otherwise, any license under the copyrights, patents or patent applications of Cambium or other 3rd Party supplied software, except for the normal non-exclusive, royalty free license to use that arises by operation of law in the sale of a product.

Restrictions

Software and documentation are copyrighted materials. Making unauthorized copies is prohibited by law. No part of the software or documentation may be reproduced, transmitted, transcribed, stored in a retrieval system, or translated into any language or computer language, in any form or by any means, without prior written permission of Cambium.

License Agreements

The software described in this document is the property of Cambium and its licensors. It is furnished by express license agreement only and may be used only in accordance with the terms of such an agreement.

High Risk Materials

Cambium and its supplier(s) specifically disclaim any express or implied warranty of fitness for any high-risk activities or uses of its products including, but not limited to, the operation of nuclear facilities, aircraft navigation or aircraft communication systems, air traffic control, life support, or weapons systems (“High Risk Use”).

This product is not restricted in the EU. Any High Risk is unauthorized, is made at your own risk and you shall be responsible for any and all losses, damage or claims arising out of any High-Risk Use.

Contents

Chapter 1: Overview

This chapter contains the following sections:

- Audience
- Conventions

Audience

This guide is intended for network operators, IT managers and administrators of a QoE, and in general for anyone requiring a detailed description of QoE functionalities. A high level overview of the QoE platform and features can be found in https://www.cambiumnetworks.com.

When appropriate, a link is provided to documents with low-level information.

Conventions

Bold font is used in UI labels, adding a > symbol for the menu navigation. For example **Status** > **System** > **System Information**.

Command line interactions are displayed in different font and using bold for the command entered by the user:

```
QoE0:# bqnsh

root@QoE0# show interface management detail

Interface: en0o1

IP address: 192.168.0.121/24

Default gateway: 192.168.0.1

Nameserver: n/a
```

Chapter 2: Deploying QoE in the network

This chapter contains the following sections:

Choosing the right deployment location

Most of the QoE functionality requires seeing the subscribers' individual IP addresses (for example, to limit each subscriber maximum rate). For that reason, it is important to deploy the QoE in a network position where there is no NAT between the QoE and the subscribers. There can be still NAT between the QoE and the subscribers, but in those cases, the rate and shaping limitations, and the subscriber metrics, applies to the NAT IP address. TCP Optimization also works if there is NAT, but it does not benefit from per-subscriber adaptations.

Ensure that traffic through the QoE is symmetric. If the downlink traffic for any subscriber is going through the QoE, then all the downlink traffic and the corresponding uplink traffic for that subscriber must go through the QoE. Otherwise, the QoE does not able to limit the rate or do shaping for that subscriber (it does not see all the traffic), and the TCP optimization blocks downloads towards that subscriber (it does not see that some traffic is acknowledged). Care must therefore be taken if there are redundant links for those links going through the QoE, so that only in case of failure are redundant links bypassing the QoE used. If there is load balancing among the links, all the load-balanced links must go through the QoE.

Regarding whether it is better to place the QoE closer or further from the subscribers, you must first consider only places where there is no asymmetric traffic and, preferably, no NAT towards the subscribers. Then, from a TCP Optimization perspective, it is better that the most difficult hops (for example, a low-quality wireless backhaul link) are between the QoE and the subscribers, because the TCP Optimization then helps you over that difficult hop. In fact, a QoE can be installed on both ends of a very challenging transmission hop (like a satellite link), but you should also look at deployment options that result in a minimum number of QoE nodes. Figure 1 shows how to choose the right deployment location.

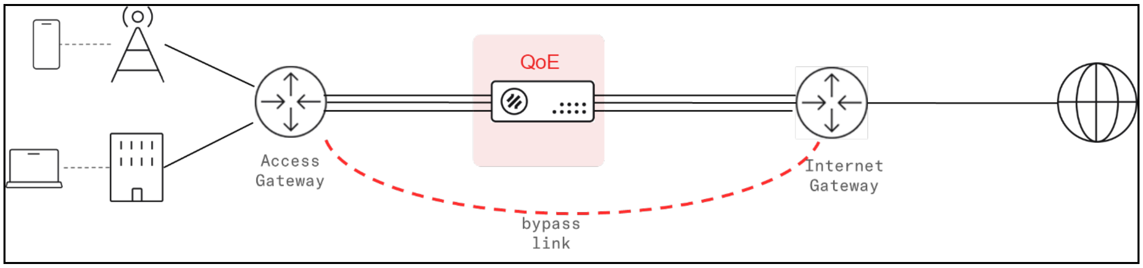

Figure 1: *Choosing the right deployment location*

It is recommended that a bypass path is established between the neighboring nodes of the QoE (Access and Internet gateways in the diagram above), so if there is a failure in the active link or the QoE, the traffic is automatically steered through the bypass. See the section Setting up a bypass.

Network requirements

Management interface requirements

- **One Management IP address**, with mask and default gateway (for example, 10.10.1.47/24 with default gateway 10.10.1.1).
- **Remote access to the management IP address** for TCP ports 22 (SSH) and 443 (HTTPS) from the Bequant support public IP address. This access can be set up with port forwarding rules in the router accessing the management network.
- Access from the QoE management IP address to the IP addresses and ports of the **Bequant license manager** (contact Bequant support for details about those IP addresses and ports).
- Access from the QoE management IP address to the **NTP servers** (UDP port 123) configured.

Data path requirements

The steering of traffic though the QoE must be **bidirectional** (all uplink and downlink traffic for the selected subscribers must go through the same QoE).

The supported traffic standards are:

- IEEE 802.1Q (VLAN)
- IEEE 802.1ad (QinQ)
- IEEE 802.3ad (LACP)
- TETF RFC2516 (PPPoE). PPP compressed traffic is not supported.
- IETF RFC 3032 (MPLS): VPLS and pseudowires are not supported

Traffic of supported type is automatically inspected by the QoE and does not require any special configuration. Non-supported traffic is transparently forwarded without any further processing.

The maximum MTU supported is 2026. Packets exceeded this MTU maximum value are discarded.

Aggregate links through the QoE

Since the two ports of each wire are bridged transparently through the QoE, aggregate links can go through the QoE over several wires transparently. The LACP, Cisco Etherchannel and Mikrotik/Linux bonding aggregate links are established between the end points before and after the QoE, as if the QoE was not there.

The aggregated links going through the QoE, are no longer physical end-to-end links, so the link monitoring on both ends of the aggregate should not rely on electrical means (like **mii monitoring** in Mikrotik) and should be based on message exchanges, like LACP messages (preferably with fast-lacp), or ARP (in case of Mikrotik active-backup bonds).

Setting up a bypass

Bypass link

A bypass link can be set up at layer-2 (for example, Mikrotik's **active-backup** link bonding, or an active-backup LACP setup) or layer-3 (for example, OSPF or BGP dynamic routing). Since the links are established directly between the two neighboring nodes, transparently with the QoE in the middle, the link monitoring mechanism should not be electrical (for example, MII), but based on the messages (for example, ARP or fast LACP). Figure 2 shows the bypass link in the QoE network.

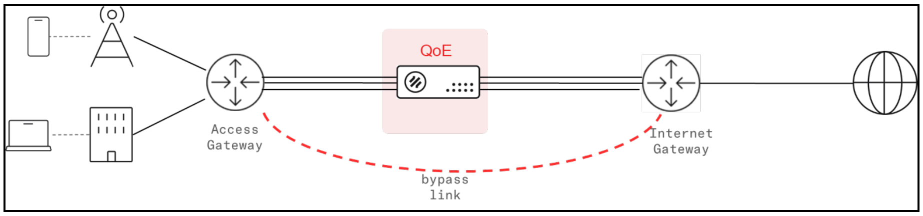

Figure 2: *Bypass link*

External bypass device

An external bypass device is connected to the external links and to the QoE. The device triggers an internal bypass if it detects the QoE is down (monitoring takes place through a USB connection between the QoE server and the bypass device). It is enabled selecting **Normal** in **Configuration > Interfaces > Bypass**. It is also possible to force the bypass from the QoE with the option **Forced Bypass** in the same screen. Figure 3 shows the external bypass device.

> **Note**
>
> Only Niagara devices are supported in this release.

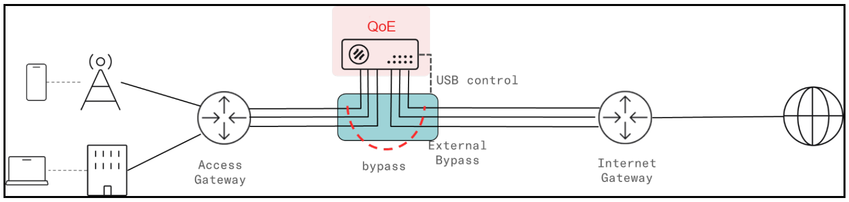

Figure 3: *External bypass device*

Network card with internal bypass

A network card with an internal bypass can be used in QoE server. Currently, only Silicom devices based on **Intel XL710-BM1** and **X540** are supported. They are controlled accessing through ssh to bqnsh.

The CLI command to see the status of the card is:

```
bqnadm@bqn0# show interface bypass

MASTER SLAVE POWER-CFG POWER-LIVE STEER-CFG STEER-LIVE

en0s3f0 en0s3f1 on on enabled enabled
```

This example shows the default configuration, that triggers the bypass if there is a power outage (POWER-CFG on and, when there is no bypass, it steers the traffic through the QoE Software (STEER-CFG on)). POWER-LIVE and STER-LIVE reports the current status on the network card and it should be equal to the configured state. The interface pair must match a wire.

To activate the bypass manually (for example to perform a planned maintenance task on the QoE server), make the following configuration changes in the master interface:

```
bqnadm@bqn0# configure

bqnadm@bqn0(config)# interface en0s3f0

bqnadm@bqn0(config-iface)#

no bypass steer

bqnadm@bqn0(config-iface)# root

bqnadm@bqn0(config)# commit
```

```
bqnadm@bqn0(config)# end
bqnadm@bqn0#
```

To steer the traffic back through the QoE software (for exapmle, after the maintenance task is complete), the changes are reverted:

```
bqnadm@bqn0# configure
bqnadm@bqn0(config)# interface en0s3f0
bqnadm@bqn0(config-iface)# bypass steer
bqnadm@bqn0(config-iface)# root
bqnadm@bqn0(config)# commit
bqnadm@bqn0(config)# end
bqnadm@bqn0#
```

You can disable the triggering of the bypass (for example, an external bypass path based on dynamic routing is used):

```
bqnadm@bqn0# configure
bqnadm@bqn0(config)# interface en0s3f0
bqnadm@bqn0(config-iface)# no bypass power
bqnadm@bqn0(config-iface)# root
bqnadm@bqn0(config)# commit
bqnadm@bqn0(config)# end
bqnadm@bqn0#
```

Connecting the QoE to the network

Management port

Connect the power supply and the management cable to the OAM port. If the server was installed from the iso, the management interface is the one selected during the installation process. Normally, the management interface is configured in the first integrated interface (the leftmost in the server motherboard, named as `en0o1`). Figure 4 shows the QoE management network.

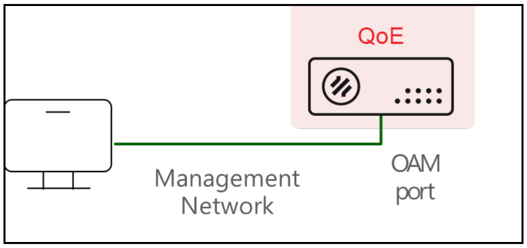

Figure 4: *Management network*

Switch ON the QoE server and access the management UI by following the steps described in Login section.

The QoE factory management IP address is **192.168.0.121/24**, with default gateway **192.168.0.1**. The management user is **QoEadm**. If the server has been installed from scratch, the IP address and password will be those provided during the installation process.

If you need to change the management IP address or default gateway, refer to Changing the management IP addresssection.

Verify date and time

Verify the date, time and time-zone in the **Administration > System Date > Set Date & Time** page are correct.

Verify connection to NTP servers and license manager

Access to the Cambium Networks license manager is OK if the **License Manager** icon in the dashboard is green. Contact Cambium Networks Support for details about license manager IP addresses and ports.

Access to NTP servers can be checked in **Administration > System Date > NTP Servers**, where some of the NTP entries should be in connected state. The NTP servers may be changed (you can eliminate some of the default servers and/or add new ones). NTP server takes some time to synchronize. So they may not be synchronized initially, but after few hours it is synchronized.

Network interface mapping

To do the mapping of the physical network ports to the interface names shown in the QoE UI, perform the following steps:

1. Connect one physical interface at a time (for example, connect it to a free switch port or to a port in the QoE already mapped).
2. Navigate to **Status > Interfaces > Link State** and select the interface that changes the link state to up.
3. Ensure that you reload the page every time after changing the physical connection, as the page does not reload automatically.

Connect subscriber traffic ports

Th subscriber traffic interfaces are paired with wires and internally connected, so the traffic received by one interface is sent to the other one and vice-versa. Each **wire** in use needs the interfaces to be connected on the access and Internet sides. It is important to connect the interfaces correctly (otherwise, the system performance will be affected).

Navigate to **Configuration > Interfaces > Data Wires** and use the network interface mapping to locate the ports of each wire.

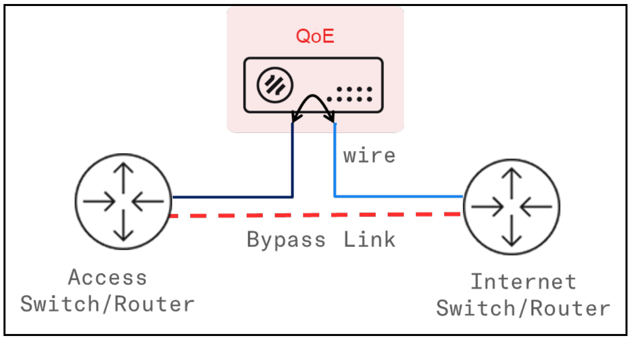

The network interfaces in use must be up and with the link detected. Navigate to **Status > Interfaces > Link State** to monitor the network interfaces from the UI .

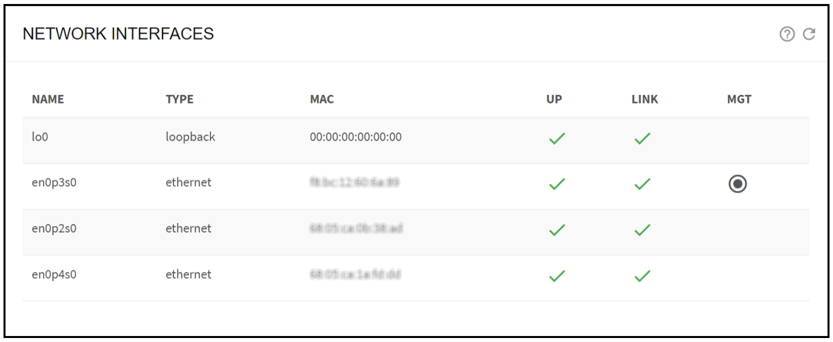

NETWORK INTERFACES

NAME	TYPE	MAC	UP	LINK	MGT
lo0	loopback	00:00:00:00:00:00	✓	✓	
en0p3s0	ethernet	[redacted]	✓	✓	◉
en0p2s0	ethernet	[redacted]	✓	✓	
en0p4s0	ethernet	[redacted]	✓	✓	

Once the traffic is steered, you can see its instant value in **Status** > **Interfaces > Throughput**.

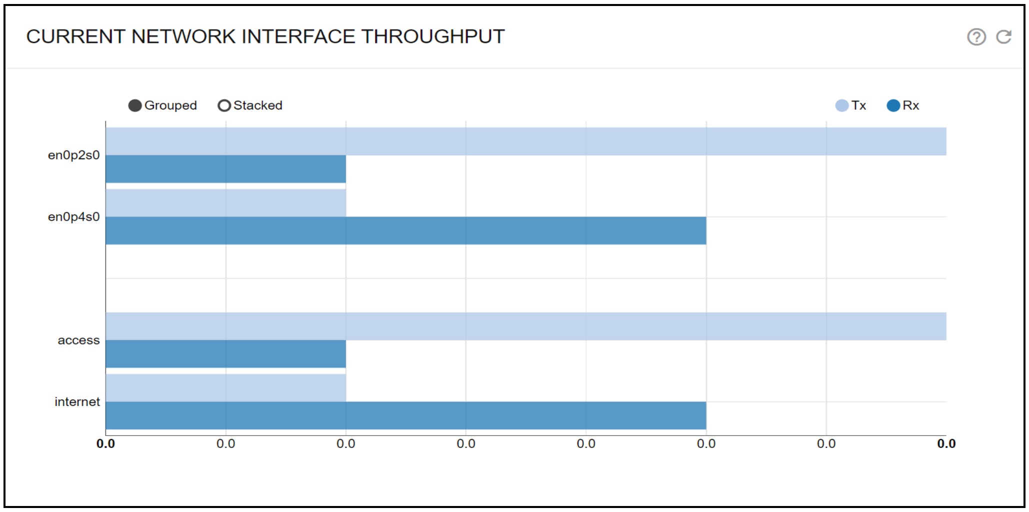

CURRENT NETWORK INTERFACE THROUGHPUT

The installation is complete.

Chapter 3: Initial Configuration

The QoE has a web-based Graphical User Interface (GUI) to perform the most common management tasks. Desktop browsers Chrome, Firefox, Safari and Microsoft Edge are supported (MS Explorer is not supported).

To open the UI contextual help, click ? icon on the page for which help is required.

This chapter contains the following sections:

- Logging in
- Changing the time
- Changing the management IP address
- Management interface firewall
- Wire configuration

Logging in

Login page

To access the UI, visit the URL: **https://oam-ip**, where **oam-ip** is the management IP address (192.168.0.121 by default).

The QoE uses a self-signed certificate, and the browser signals it as unsecure. Ignore the warning and go to the web page.

Enter username and password (default is admin and cambium).

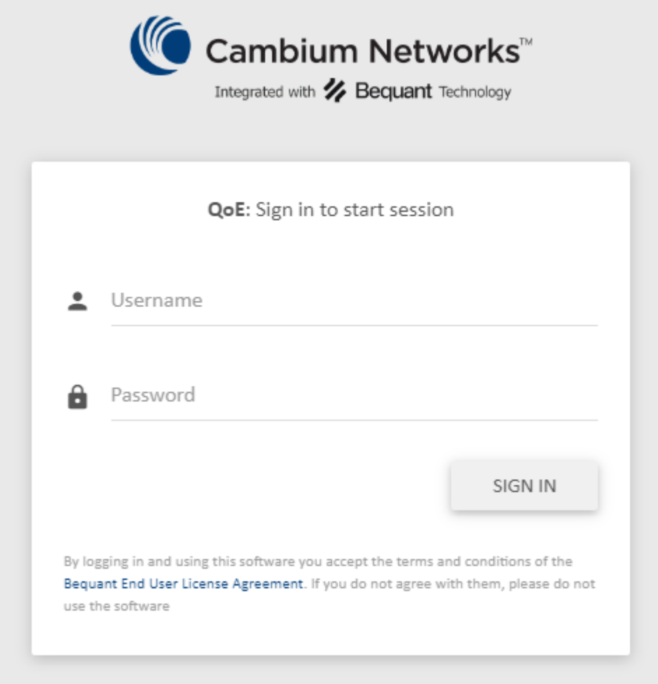

QoE: Sign in to start session

By logging in and using this software you accept the terms and conditions of the Bequant End User License Agreement. If you do not agree with them, please do not use the software

Attention

You cannot use the **root** username to login to the UI.

The dashboard page

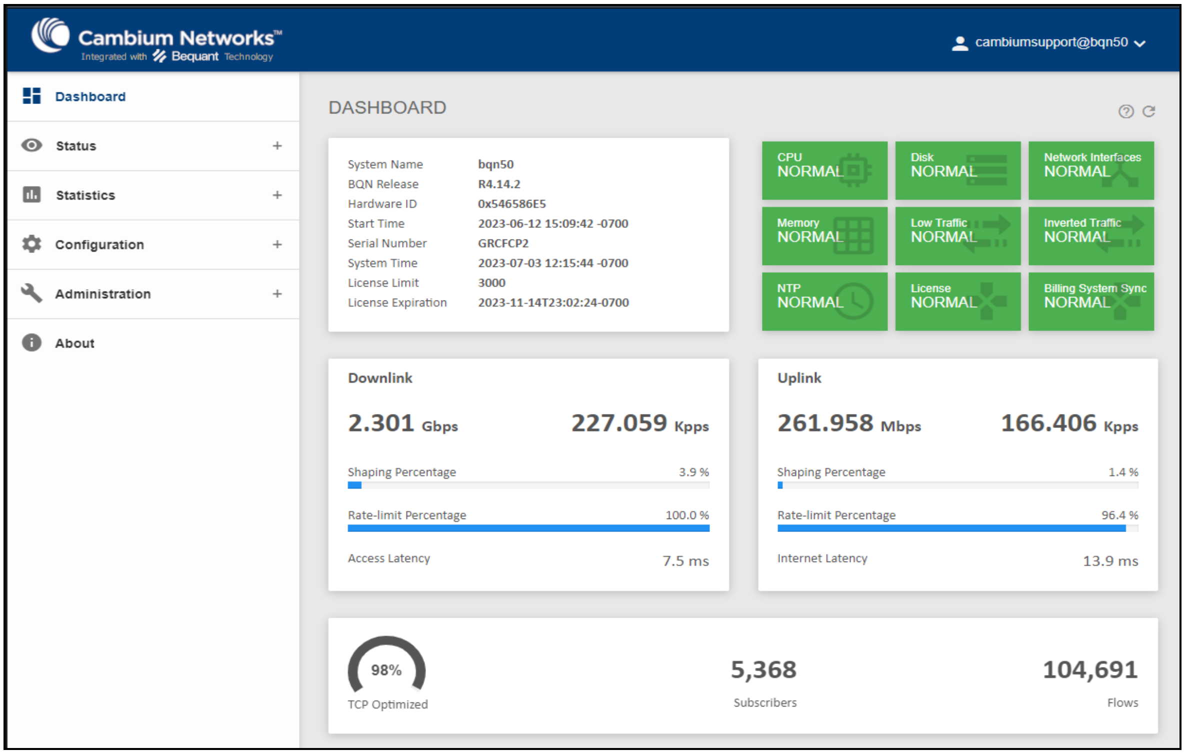

System Name	bqn50
BQN Release	R4.14.2
Hardware ID	0x546586E5
Start Time	2023-06-12 15:09:42 -0700
Serial Number	GRCFCP2
System Time	2023-07-03 12:15:44 -0700
License Limit	3000
License Expiration	2023-11-14T23:02:24-0700

CPU NORMAL
Disk NORMAL
Network Interfaces NORMAL
Memory NORMAL
Low Traffic NORMAL
Inverted Traffic NORMAL
NTP NORMAL
License NORMAL
Billing System Sync NORMAL

Downlink

2.301 Gbps 227.059 Kpps

Shaping Percentage 3.9 %

Rate-limit Percentage 100.0 %

Access Latency 7.5 ms

Uplink

261.958 Mbps 166.406 Kpps

Shaping Percentage 1.4 %

Rate-limit Percentage 96.4 %

Internet Latency 13.9 ms

98% TCP Optimized

5,368 Subscribers

104,691 Flows

The home page has a lateral menu, a dashboard, and a small summary of the system information.

The dashboard shows all the icons in **Green** color. The **Network Interfaces** icon does not be in Green color until all the configured **wires** are connected (if there are interfaces not being used in any of the configured **wires**, it remains in Orange color) and the icon **Low Traffic** does not be in Green until the traffic flows through the QoE. Some icons, takes to a window with more information about the QoE status. See Other External Subscriber Data for steps to take when an icon is not in normal state.

Tables

Many UI pages contain data tables to display the information. Figure 5 is an example for table.

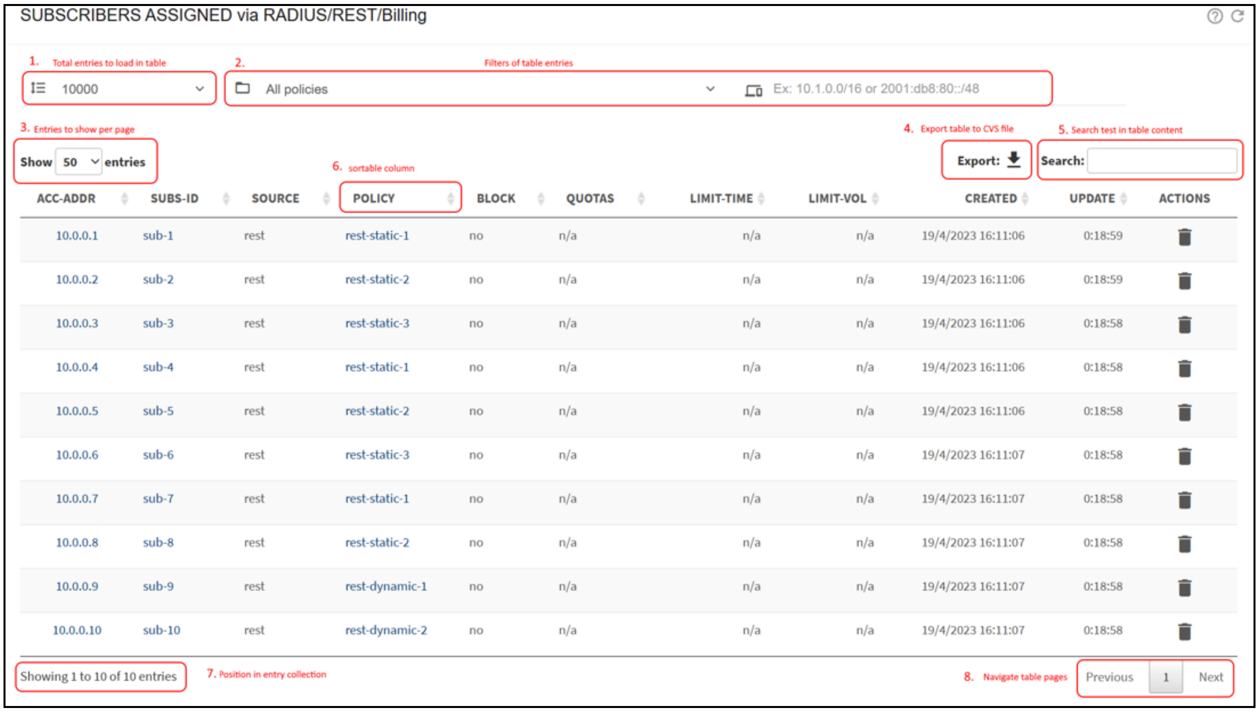

SUBSCRIBERS ASSIGNED via RADIUS/REST/Billing

1. Total entries to load in table
10000
2. Filters of table entries
All policies
Ex: 10.1.0.0/16 or 2001:db8:80::/48
3. Entries to show per page
Show 50 entries
4. Export table to CVS file
Export:
5. Search test in table content
Search:
6. sortable column

ACC-ADDR	SUBS-ID	SOURCE	POLICY	BLOCK	QUOTAS	LIMIT-TIME	LIMIT-VOL	CREATED	UPDATE	ACTIONS
10.0.0.1	sub-1	rest	rest-static-1	no	n/a	n/a	n/a	19/4/2023 16:11:06	0:18:59	🗑
10.0.0.2	sub-2	rest	rest-static-2	no	n/a	n/a	n/a	19/4/2023 16:11:06	0:18:59	🗑
10.0.0.3	sub-3	rest	rest-static-3	no	n/a	n/a	n/a	19/4/2023 16:11:06	0:18:58	🗑
10.0.0.4	sub-4	rest	rest-static-1	no	n/a	n/a	n/a	19/4/2023 16:11:06	0:18:58	🗑
10.0.0.5	sub-5	rest	rest-static-2	no	n/a	n/a	n/a	19/4/2023 16:11:06	0:18:58	🗑
10.0.0.6	sub-6	rest	rest-static-3	no	n/a	n/a	n/a	19/4/2023 16:11:07	0:18:58	🗑
10.0.0.7	sub-7	rest	rest-static-1	no	n/a	n/a	n/a	19/4/2023 16:11:07	0:18:58	🗑
10.0.0.8	sub-8	rest	rest-static-2	no	n/a	n/a	n/a	19/4/2023 16:11:07	0:18:58	🗑
10.0.0.9	sub-9	rest	rest-dynamic-1	no	n/a	n/a	n/a	19/4/2023 16:11:07	0:18:58	🗑
10.0.0.10	sub-10	rest	rest-dynamic-2	no	n/a	n/a	n/a	19/4/2023 16:11:07	0:18:58	🗑

Showing 1 to 10 of 10 entries

Figure 5: *Data table*

The following are the most important features that are highlighted in Figure 5:

1. The drop-down on the upper-left defines the number of entries to load in the table (10,000 in the example). By default, the value of each table is selected to load all entries. If the default is too low, then select a bigger number.

2. To the right of the total entries, you can filter the elements to the table content. In the example above, entries can be filtered by policy and by subscriber.

3. The **Show entries** define the size of the page.

4. The **Export** button generates a CVS file with the table content. The file contains column labels in first row and then one line per table row.

5. The **Search** field shows the table entries containing the sub-string specified in the text field.

6. Columns with arrows are sortable. To sort by a column, click on that column label for ascending/descending order.

7. A label at the bottom-left corner displays the list of entries.

8. The buttons at the bottom-right corner navigates to the table pages.

While exporting the tables to CVS, comma and commands are placed as follows:

- If a field contains a comma (“,”) the whole field is enclosed in double quotes. For example: a,b is exported as “a,b".

- If a field contains one double quote, it is replaced by two double quotes. For example: a”b is exported as a””b.

- If both a comma and a double quote are present, both rules are applied. For example, “a,b” is exported as “””a,b”””.

Changing the time

This page is used to change the system date and time if it is necessary. To change the time, navigate to **Administration > System Date > Set Date & Time**.

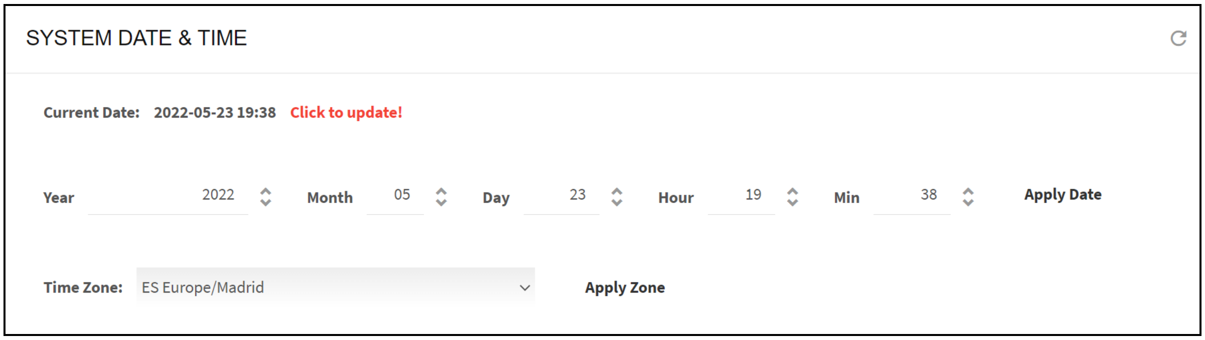

SYSTEM DATE & TIME

Current Date: 2022-05-23 19:38 Click to update!

Year	2022	**Month**	05	**Day**	23	**Hour**	19	**Min**	38	**Apply Date**

Time Zone: ES Europe/Madrid **Apply Zone**

Click **Apply Date** to apply the changes to the local date and time, and **Apply Zone** to change the time zone. It is possible to browse through the list of time zones pressing the initials of the country of Interest (for example, ES for Spain).

Changing the management IP address

To change the settings of the management interface, navigate to **Configuration** > **Interfaces > Management** from the menu. The IP settings includes IP address and mask, the default gateway and the VLAN identifier (if any). An optional DNS server IP address can be configured (it is required if server names are used in the integration with a billing system).

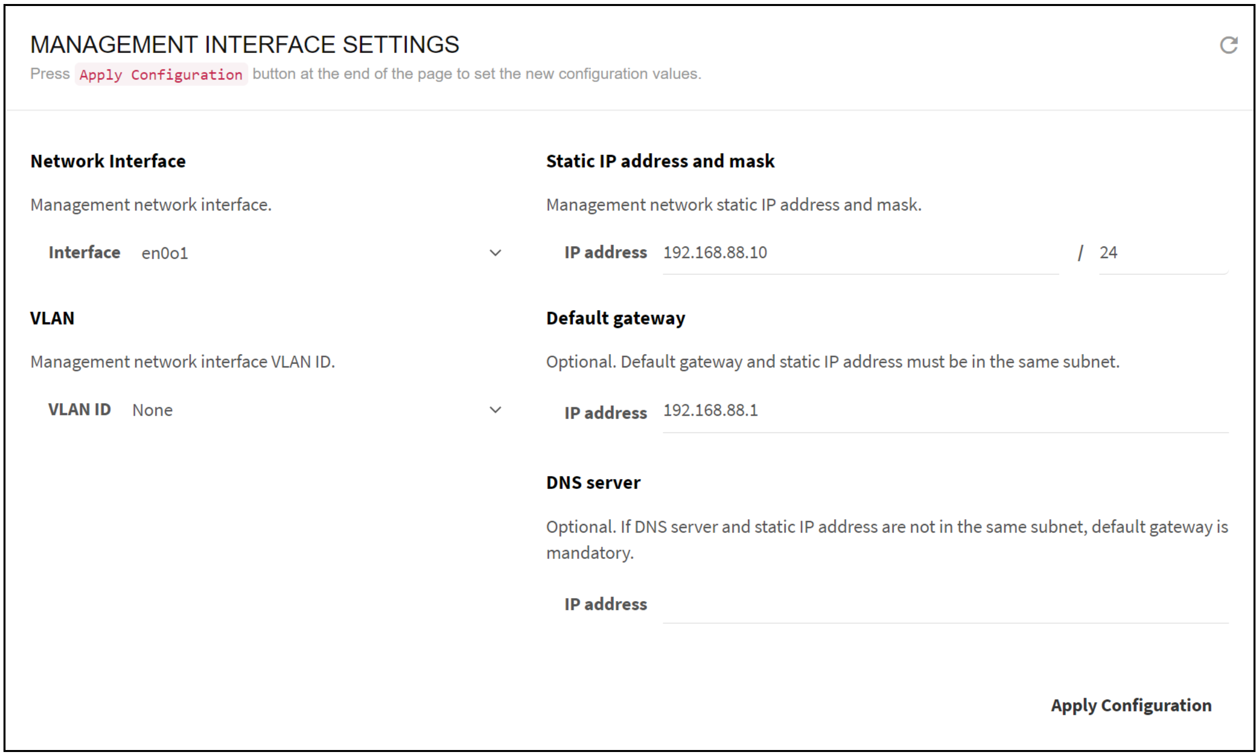

MANAGEMENT INTERFACE SETTINGS

Press `Apply Configuration` button at the end of the page to set the new configuration values.

Network Interface

Management network interface.

Interface en0o1

VLAN

Management network interface VLAN ID.

VLAN ID None

Static IP address and mask

Management network static IP address and mask.

IP address 192.168.88.10 / 24

Default gateway

Optional. Default gateway and static IP address must be in the same subnet.

IP address 192.168.88.1

DNS server

Optional. If DNS server and static IP address are not in the same subnet, default gateway is mandatory.

IP address

Apply Configuration

Note

Do not change the network interface used for management, unless it is recommended by the Cambium Networks support personnel.

When the new settings are completed, click **Apply Configuration** to commit the changes. Connecting back to the node may require access from the new subnet and logging back to the UI.

Management interface firewall

To set up the management interface firewall, which applies only to the management interface (not to the interfaces configured in **wires**) select on the lateral menu **Configuration**> **Interfaces > Management Firewall.** This shows the IP address ranges allowed to access the management interface. By default, no IP address ranges are configured, and all are allowed.

To add an allowed IP address range, click ⋮ icon and click **Add IP Address Range....** Once one IP address range is allowed, the firewall is enabled, and all incoming connections from IP addresses not part of the configured IP address ranges are blocked. It is therefore important to include an IP address range that covers the IP address from which we are accessing the UI and also the subnet of the management IP address (the UI includes them in the suggested list). Other IPs that interact with the management interface, such as RADIUS/REST clients, a billing system, and the NTP server, should also be included.

To disable the firewall, click **Delete** icon next to each entry and remove all the entries. After deleting all the entries, click **Apply**. It is important not to click **Apply** before all the entries are deleted, because a premature **Apply** keeps the firewall active and it may prevent you from accessing the server, if the entry that covers your IP is not present.

Wire configuration

A wire is a network interface pair that process the traffic of the subscriber. To configure the wires, navigate to **Configuration** > **Interfaces > Data Wires**.

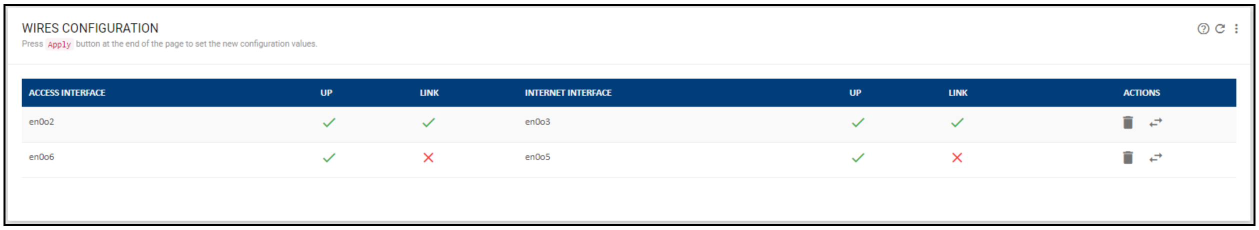

WIRES CONFIGURATION

Press Apply button at the end of the page to set the new configuration values.

ACCESS INTERFACE	UP	LINK	INTERNET INTERFACE	UP	LINK	ACTIONS
en0o2	✓	✓	en0o3	✓	✓	🗑 ⇆
en0o6	✓	✗	en0o5	✓	✗	🗑 ⇆

Wires are directional, with the first network interface connected to the access towards the subscribers and the second interface on the Internet side. If any mistake happened while connecting the ports, click on the ⇆ icon and swap the interfaces.

To add a wire, click on the ⋮ icon and select **Add Wire.....** A form allows you to select the access and Internet interfaces (the form lists the available interfaces).

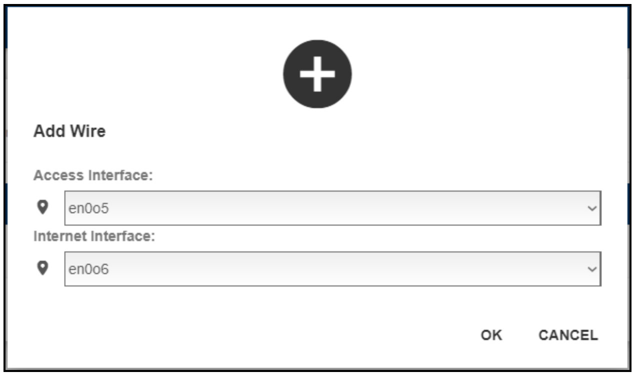

Add Wire

Access Interface: en0o5

Internet Interface: en0o6

To remove a wire, click on the 🗑 icon. To modify a wire, remove the existing wire and then, before applying the changes, add the wire with the change.

Click **Apply Configuration** to apply the changes.

> **Caution**
>
> Do not delete the wires, unless it is recommended by the Cambium Networks support personnel. Wrong configurations may lead to loss of the service.

IPMI Configuration

Some servers have a lights-out module for general management (power on/off, hardware monitoring, and so on). We recommend using the server own setting tools (For example, BIOS section), but if this is not possible (for example, the server is already powered and with traffic), the QoE CLI allows a basic setup.

To do it, access to the QoE shell is needed. The following example sets a static IP address 192.168.0.120/24 with default gateway 192.168.0.1 and creates a new user bequant.

```
bqnadm@bqn0# system ipmi lan static

bqnadm@bqn0# system ipmi lan 192.168.0.120/24 192.168.0.1

bqnadm@bqn0# system ipmi user add bequant

bqnadm@bqn0# system ipmi user passwd bequant

New password:

Retype new password:

bqnadm@bqn0#
```

Remote Access using Port Forwarding

To allow access from the internet to the QoE server in a secure way, a port forwarding rule should be configured, restricted to only a small set of source IP addresses. The rest of the section describes the steps to follow in a Mikrotik router to do port forwarding of HTTPS (port 443). Similar steps can be followed for SSH (port 22).

Navigate to **IP > Firewall > Address Lists** and create an address list with all the source IP addresses that will be allowed access to the QoE server (in our example, we call it **bqn-allowed-src**).

In **IP > Firewall > NAT**, create a destination NAT rule:

1. In the **General** tab, set **Chain** to **dstnat**.
2. In the **General** tab, set **Dst. Address** to the public IP address used to access the QoE from the internet.
3. In the **General** tab, set **Protocol** to TCP.
4. In the **General** tab, set **Dst. Port** to the public IP address used to access the QoE from the internet.
5. In the **Advanced** tab, set **Src. Address List** to the name of the previously created list (**bqn-allowed-src** in our example).
6. In the **Action** tab, set **Action** to **dst-nat**.
7. In the **Action** tab, set **To Address** to the IP address of the QoE server (for example, 192.168.0.121).
8. In the **Action** tab, set **To Ports** to **443**, the HTTPS port of the QoE server.

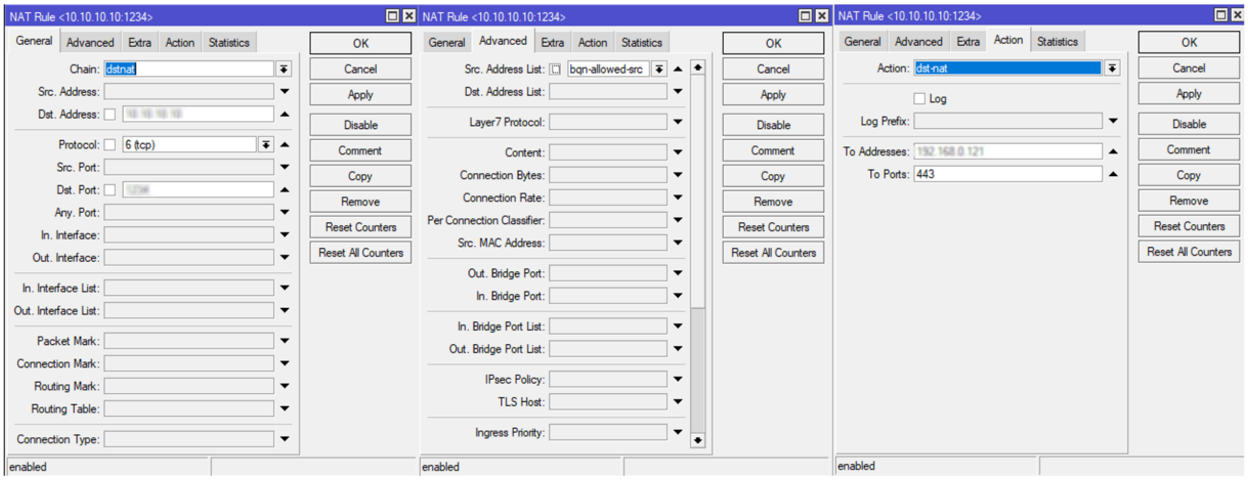

NAT Rule <10.10.10.10:1234>

General

Chain: dstnat

Dst. Address: □ [redacted]

Protocol: □ 6 (tcp)

Dst. Port: □ [redacted]

enabled

NAT Rule <10.10.10.10:1234>

Advanced

Src. Address List: □ bqn-allowed-src

enabled

NAT Rule <10.10.10.10:1234>

Action

Action: dst-nat

□ Log

To Addresses: [redacted]

To Ports: 443

enabled

Chapter 4: Management Tasks

This chapter contains the following sections:

- Updating the software
- Generating a diagnostic file
- Back up the configuration
- Monitoring with SNMP
- Traffic captures
- Logs
- Software bypass

Updating the software

Update within same major release

To update within the same major release (for example, from R4.7.1 to R4.8.3), only the QoE package to be updated (for example, bqn-R4.8.3.bpkg). This update contains the following steps:

1. Installing the software.
2. Activating the software.

> **Note**
>
> The software activation involves a traffic interruption of some seconds. So it is recommended to activate the software when the throughput is low.

To install the software, navigate to **Administration > Software**, click on the ⋮ menu icon and select **Install...**. A file selector pops up to choose the package, that it is transferred to the QoE server and installed.

To activate the software, navigate to **Administration > Software**, click on the ⟳ cycle icon of the package (highlighted in red in the picture below). After a few seconds, login to the QoE. During this time, the traffic flow is interrupted.

Figure 6 shows the software status page.

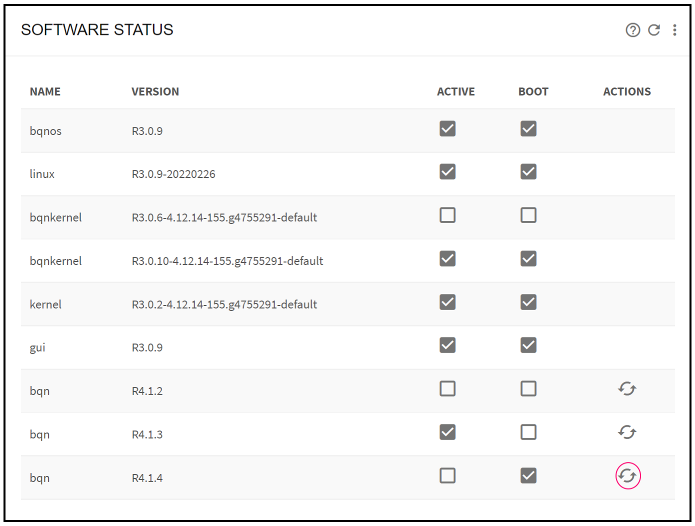

SOFTWARE STATUS

NAME	VERSION	ACTIVE	BOOT	ACTIONS
bqnos	R3.0.9	☑	☑	
linux	R3.0.9-20220226	☑	☑	
bqnkernel	R3.0.6-4.12.14-155.g4755291-default	□	□	
bqnkernel	R3.0.10-4.12.14-155.g4755291-default	☑	☑	
kernel	R3.0.2-4.12.14-155.g4755291-default	☑	☑	
gui	R3.0.9	☑	☑	
bqn	R4.1.2	□	□	⟳
bqn	R4.1.3	☑	□	⟳
bqn	R4.1.4	□	☑	⟳

Figure 6: *The software status page*

Update across major releases

To migrate to a new major release (for example, from R3 to R4), update the platform packages (bqnos, kernel, bqnkernel, linux and gui) in addition to the QoE package. Platform packages require a reboot to be activated.

To update the software across major releases, perform the following steps:

1. Install new qoeos, wait for one minute and reboot.
2. Install kernel and qoekernel, and reboot.
3. Install Linux and reboot.
4. Install gui and reboot.
5. Install QoE and reload.

The server must be placed out of the traffic path, as server reboots involve service losses.

Finally, the old configuration may require migration to the new release. Then remove deprecated commands accessing the QoE server through SSH and running the commands:

```
bqnadm@bqn# configure

bqnadm@bqn(config)# clear config undefined

bqnadm@bqn(config)# commit

%WARN: Verify configuration after deleting undefined commands

bqnadm@bqn(config)# end

bqnadm@bqn# exit
```

Generating a diagnostic file

A diagnostic file can be generated from **Administration > Diagnostic** if requested by [Cambium Networks Support Site](#).

The file is placed in the download folder of the corresponding browser.

Backup and restoring the configuration

To save the server configuration to a local file, navigate to **Administration > Backup > Save**.

To load a previously saved backup, navigate to **Administration > Backup > Load**.

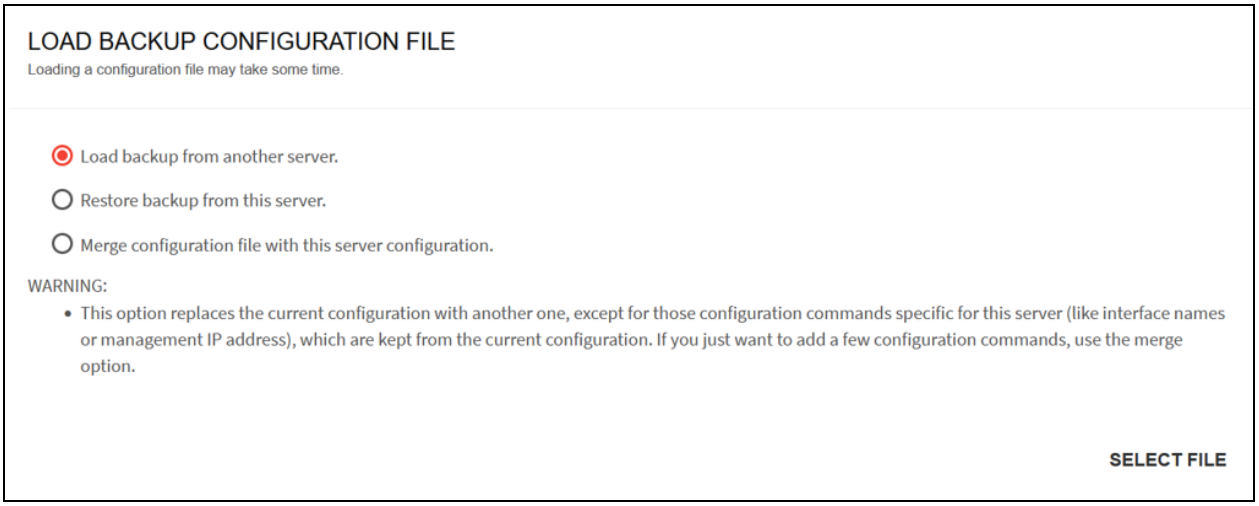

LOAD BACKUP CONFIGURATION FILE

Loading a configuration file may take some time.

● Load backup from another server.

○ Restore backup from this server.

○ Merge configuration file with this server configuration.

WARNING:

- This option replaces the current configuration with another one, except for those configuration commands specific for this server (like interface names or management IP address), which are kept from the current configuration. If you just want to add a few configuration commands, use the merge option.

There are three options:

- **Load backup from another server**: The purpose of this option is to bring a common configuration to a number of servers, with the same policies and API clients. When this option is used, only those generic sections of the configuration are loaded. Server-specific parts (management interface configuration, data wires, licenses, and API QoE own IP address) are left untouched.

- **Restore backup from this server**: This option is used to recover a previous state of this server configuration. In this option, the configuration is completely replaced by the backup one, so it is important that the configuration comes from this very server, otherwise the server may be left unreachable.

- **Merge configuration file with this server configuration**: This option is useful to carry only some configuration sections, for example a set of policy rules. In this option, the whole configuration in the file is added to the current configuration. Only generic configuration sections and of them, only the ones strictly needed should be included in the merge file, in order to avoid conflicts. For example, a configuration can have only one management default gateway, so merging a file containing a default gateway will replace the old one.

Select the load option and press **SELECT FILE**. Select the file to load and press **Open**. A dialogue will inform of the operation outcome (if OK or if an error is found). If there is an error, no modification is done to the server configuration (load operation is atomic).

Monitoring with SNMP

The QoE supports the following SNMP v2c alarms (traps):

- **Cpu**: Excessive server CPU load.

- **Memory-dpdk**: Excessive memory usage in DPDK packet processing.

- **Memory-pool**: Excessive memory usage in QoE general memory pool.

- **Disk**: File system full or almost full.

- **Process**: Some mandatory processes down.
- **Traffic-uplink**: No traffic in the uplink direction.
- **Traffic-downlink**: No traffic in the downlink direction.
- **Traffic-low**: Low traffic (uplink and downlink directions combined).
- **Traffic-inverted**: Uplink throughput higher than downlink throughput (possibly because some wires are inverted, with access port connected to the Internet and vice versa).
- **Wire**: No wires configured or some wires down.
- **License-available**: No license defined or invalid.
- **License-expiration**: License has expired.
- **License-usage**: Server throughput is above the license capacity.
- **Time**: No NTP server configured or not reachable.
- **QoEmgr**: QoE remote management system not reachable.

These alarms are related to the dashboard shown in QoE home page. For more information, refer to Troubleshooting section.

To configure the SNMP agent, navigate to **Administration** > **SNMP**. Figure 7 shows the SNMP page.

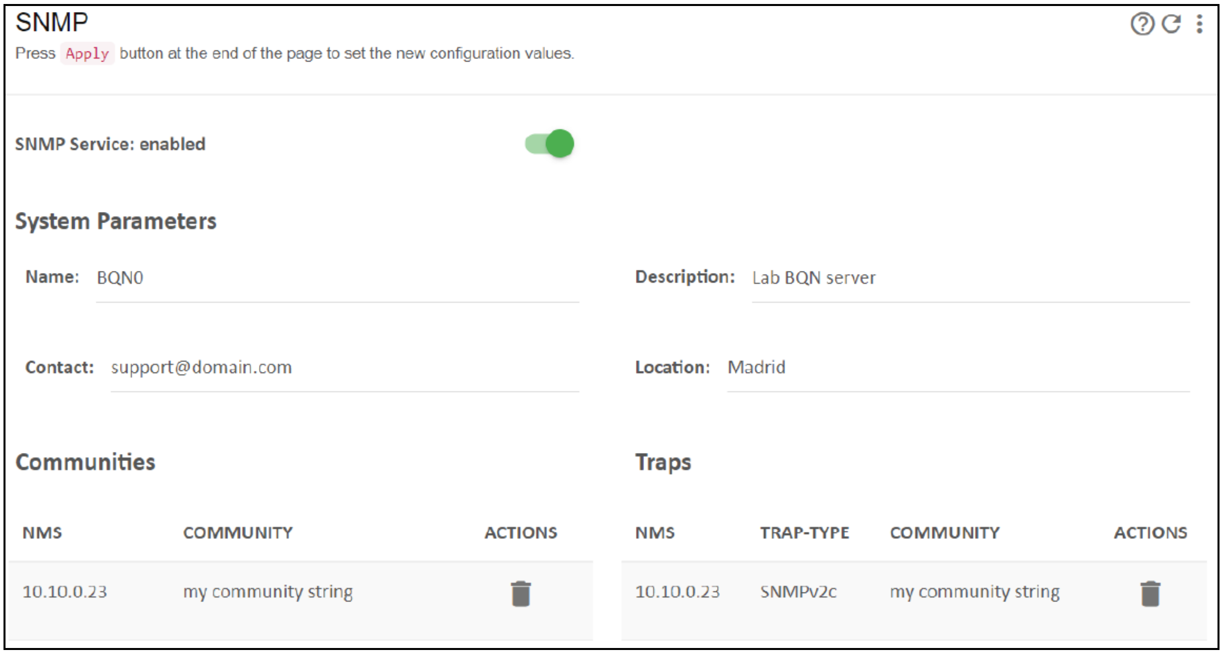

SNMP

Press Apply button at the end of the page to set the new configuration values.

SNMP Service: enabled

System Parameters

Name: BQN0

Description: Lab BQN server

Contact: support@domain.com

Location: Madrid

Communities

NMS	COMMUNITY	ACTIONS
10.10.0.23	my community string	🗑

Traps

NMS	TRAP-TYPE	COMMUNITY	ACTIONS
10.10.0.23	SNMPv2c	my community string	🗑

Figure 7: *The SNMP page*

The QoE SNMP also exports some system statistics. To get the QoE MIB files, visit Cambium Networks Support Site.

Traffic captures

QoE is used to get traffic captures in pcap format from any of its network interfaces. Those captures are used to troubleshoot issues in the QoE server, but also somewhere else in the network, as direct traffic captures are many times difficult to obtain in other network nodes.

Traffic captures are indicated with a glass icon next to the interface name. It is available in the following pages:

- Status > Interfaces > Link State
- Status > Interfaces > Data Wires

Click the icon to display a dialogue with the capture options.

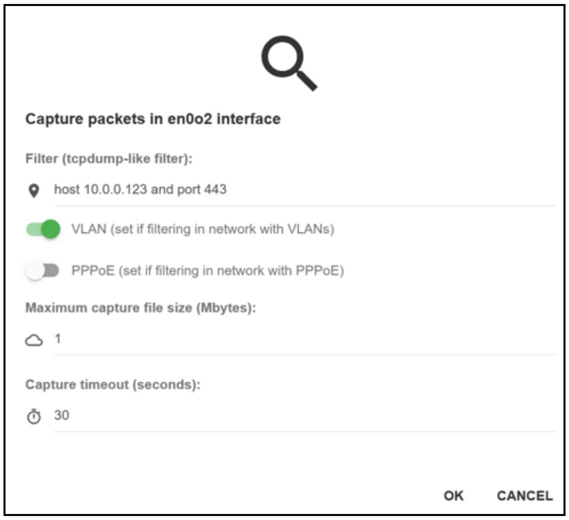

Capture packets in en0o2 interface

Filter (tcpdump-like filter):

host 10.0.0.123 and port 443

VLAN (set if filtering in network with VLANs)

PPPoE (set if filtering in network with PPPoE)

Maximum capture file size (Mbytes):

1

Capture timeout (seconds):

30

Figure 8: *Traffic captures*

The filter field accepts the format of `tcpdump` filters. If it is optional and empty, all the traffic are captured. The following are some filter examples:

- Traffic involving an IP address: `host 10.0.0.23`
- TCP traffic involving an IP address and port: `tcp and host 10.0.0.23 and port 443`
- UDP traffic involving an IP subset to a specific Internet address: `udp and net 10.0.0.0/24 and host 8.8.8.8`

If the network has VLANs and/or PPPoE, the corresponding toggle switch must be set for the filter to work. In the previous figure, the network has VLANs.

Maximum capture file size (up to a maximum of 500 MB). The capture stops if the maximum file size is reached.

Capture timeout (600 seconds maximum). The capture stops if the maximum time gets elapse. It can also be stopped before getting elapses by pressing the **CANCEL**.

The capture is limited by a maximum size and timeout (whatever happens first). The reason is to protect the system, because traffic captures has a performance impact.

> **Note**
>
> To reduce the performance impact on the system, use the smallest size and duration which meets the requirement.

After the capture is complete, a pcap file is generated at the download folder. The file can be inspected using a traffic tool supporting the pcap format, for example *wireshark*.

Logs

The UI displays two types of logs to debug the complex issues:

- **OS log messages**: Navigate to **Administration** > **Logs** > **System**.
- **Kernel log messages** (output of dmesg command): Navigate to **Administration** > **Log** > **Kernel**.

It is possible to request more log lines and to export the log entries to a local file. Figure 9 shows the System Logs page.

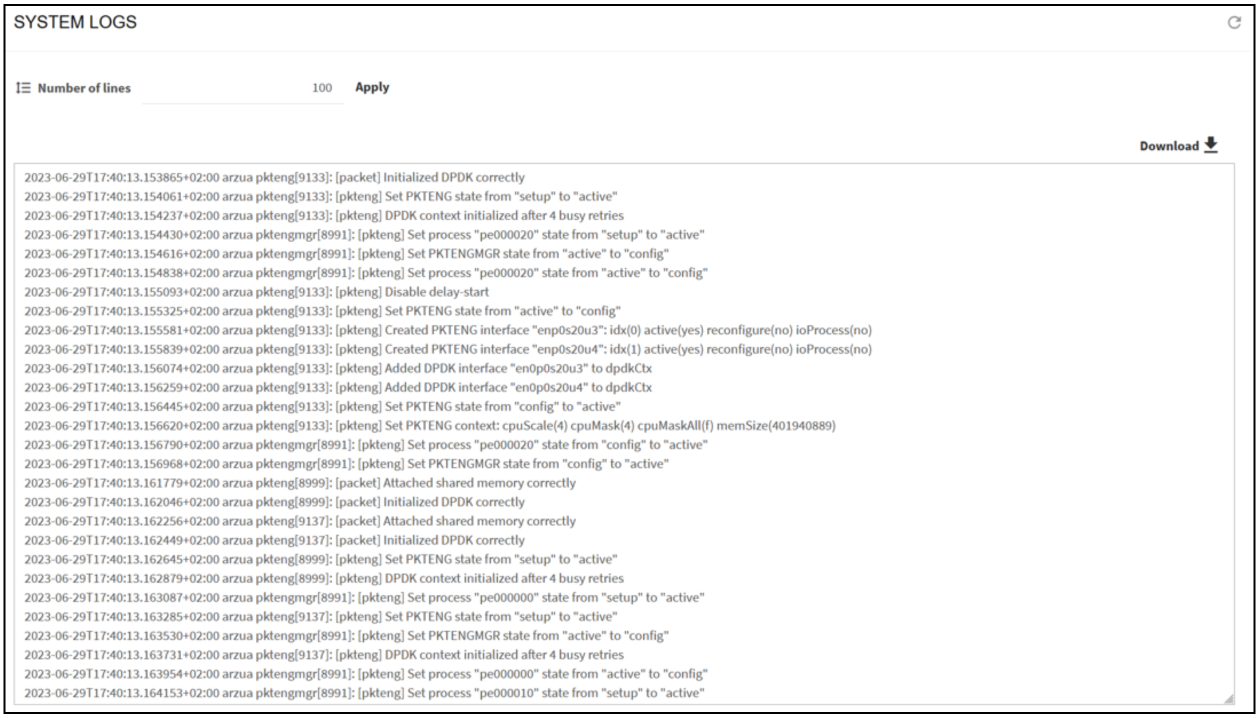

```
2023-06-29T17:40:13.153865+02:00 arzua pkteng[9133]: [packet] Initialized DPDK correctly
2023-06-29T17:40:13.154061+02:00 arzua pkteng[9133]: [pkteng] Set PKTENG state from "setup" to "active"
2023-06-29T17:40:13.154237+02:00 arzua pkteng[9133]: [pkteng] DPDK context initialized after 4 busy retries
2023-06-29T17:40:13.154430+02:00 arzua pktengmgr[8991]: [pkteng] Set process "pe000020" state from "setup" to "active"
2023-06-29T17:40:13.154616+02:00 arzua pktengmgr[8991]: [pkteng] Set PKTENGMGR state from "active" to "config"
2023-06-29T17:40:13.154838+02:00 arzua pktengmgr[8991]: [pkteng] Set process "pe000020" state from "active" to "config"
2023-06-29T17:40:13.155093+02:00 arzua pkteng[9133]: [pkteng] Disable delay-start
2023-06-29T17:40:13.155325+02:00 arzua pkteng[9133]: [pkteng] Set PKTENG state from "active" to "config"
2023-06-29T17:40:13.155581+02:00 arzua pkteng[9133]: [pkteng] Created PKTENG interface "enp0s20u3": idx(0) active(yes) reconfigure(no) ioProcess(no)
2023-06-29T17:40:13.155839+02:00 arzua pkteng[9133]: [pkteng] Created PKTENG interface "enp0s20u4": idx(1) active(yes) reconfigure(no) ioProcess(no)
2023-06-29T17:40:13.156074+02:00 arzua pkteng[9133]: [pkteng] Added DPDK interface "en0p0s20u3" to dpdkCtx
2023-06-29T17:40:13.156259+02:00 arzua pkteng[9133]: [pkteng] Added DPDK interface "en0p0s20u4" to dpdkCtx
2023-06-29T17:40:13.156445+02:00 arzua pkteng[9133]: [pkteng] Set PKTENG state from "config" to "active"
2023-06-29T17:40:13.156620+02:00 arzua pkteng[9133]: [pkteng] Set PKTENG context: cpuScale(4) cpuMask(4) cpuMaskAll(f) memSize(401940889)
2023-06-29T17:40:13.156790+02:00 arzua pktengmgr[8991]: [pkteng] Set process "pe000020" state from "config" to "active"
2023-06-29T17:40:13.156968+02:00 arzua pktengmgr[8991]: [pkteng] Set PKTENGMGR state from "config" to "active"
2023-06-29T17:40:13.161779+02:00 arzua pkteng[8999]: [packet] Attached shared memory correctly
2023-06-29T17:40:13.162046+02:00 arzua pkteng[8999]: [packet] Initialized DPDK correctly
2023-06-29T17:40:13.162256+02:00 arzua pkteng[9137]: [packet] Attached shared memory correctly
2023-06-29T17:40:13.162449+02:00 arzua pkteng[9137]: [packet] Initialized DPDK correctly
2023-06-29T17:40:13.162645+02:00 arzua pkteng[8999]: [pkteng] Set PKTENG state from "setup" to "active"
2023-06-29T17:40:13.162879+02:00 arzua pkteng[8999]: [pkteng] DPDK context initialized after 4 busy retries
2023-06-29T17:40:13.163087+02:00 arzua pktengmgr[8991]: [pkteng] Set process "pe000000" state from "setup" to "active"
2023-06-29T17:40:13.163285+02:00 arzua pkteng[9137]: [pkteng] Set PKTENG state from "setup" to "active"
2023-06-29T17:40:13.163530+02:00 arzua pktengmgr[8991]: [pkteng] Set PKTENGMGR state from "active" to "config"
2023-06-29T17:40:13.163731+02:00 arzua pkteng[9137]: [pkteng] DPDK context initialized after 4 busy retries
2023-06-29T17:40:13.163954+02:00 arzua pktengmgr[8991]: [pkteng] Set process "pe000000" state from "active" to "config"
2023-06-29T17:40:13.164153+02:00 arzua pktengmgr[8991]: [pkteng] Set process "pe000010" state from "setup" to "active"
```

Figure 9: *The System Logs page*

Software bypass

You can make some traffic to go through the QoE transparently, without processing by QoE. These traffics are captured in one of the network interfaces and relayed transparently to its peer interface in the same wire. The QoE software has no impact on such traffic.

The following traffic can be configured to be bypassed:

- IP traffic v4
- IP traffic v6
- Traffic with some specific VLAN tags
- Traffic without a VLAN tag (for example, untagged).

To bypass some traffic, navigate to **Configuration** > **Optimization Settings** and enable the corresponding option. Figure 10 shows the Software bypass options.

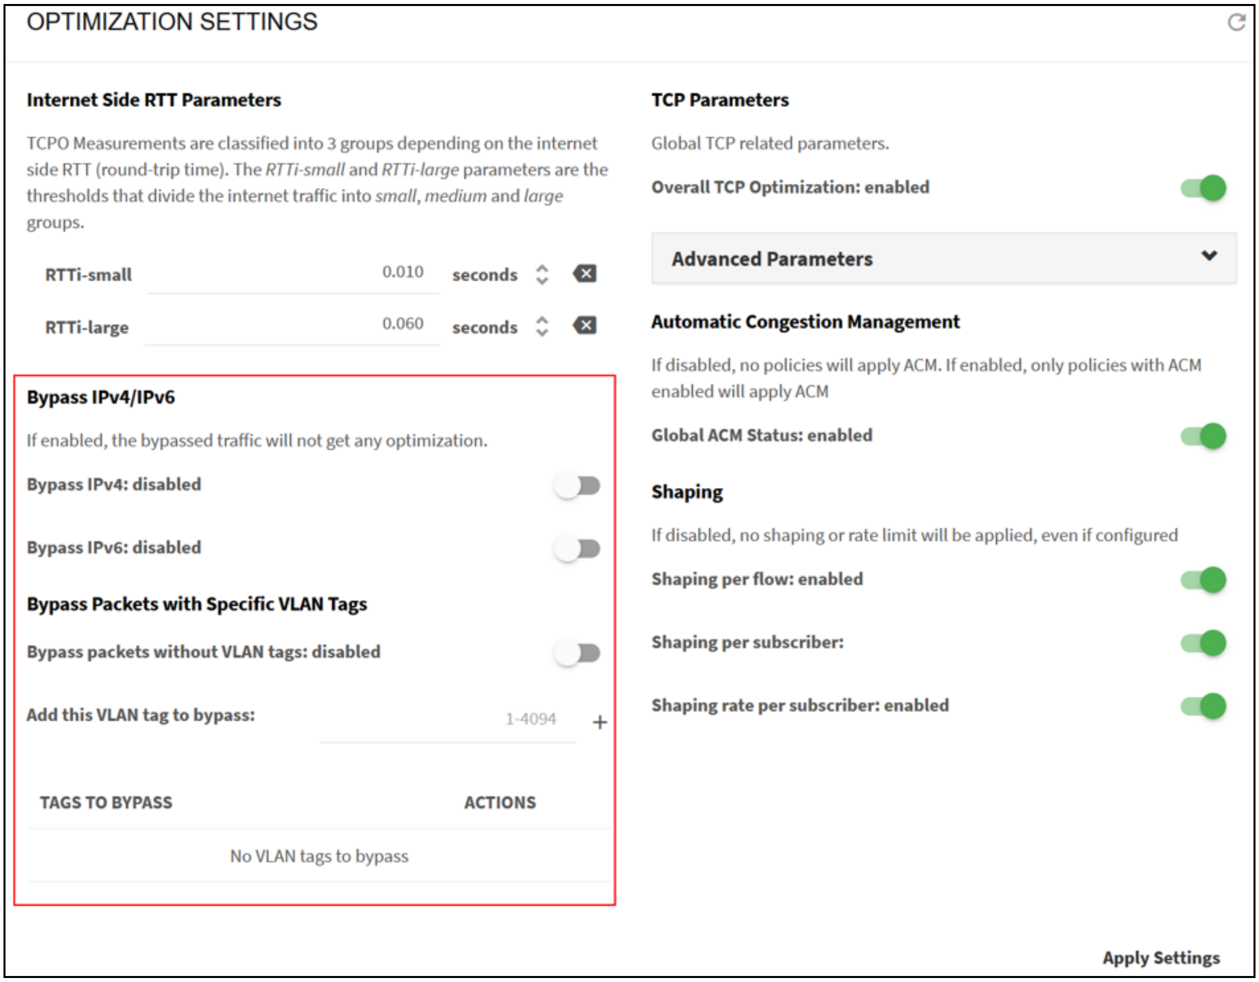

OPTIMIZATION SETTINGS

Internet Side RTT Parameters

TCPO Measurements are classified into 3 groups depending on the internet side RTT (round-trip time). The *RTTi-small* and *RTTi-large* parameters are the thresholds that divide the internet traffic into *small*, *medium* and *large* groups.

RTTi-small 0.010 seconds

RTTi-large 0.060 seconds

Bypass IPv4/IPv6

If enabled, the bypassed traffic will not get any optimization.

Bypass IPv4: disabled

Bypass IPv6: disabled

Bypass Packets with Specific VLAN Tags

Bypass packets without VLAN tags: disabled

Add this VLAN tag to bypass: 1-4094 +

TAGS TO BYPASS	ACTIONS
No VLAN tags to bypass	

TCP Parameters

Global TCP related parameters.

Overall TCP Optimization: enabled

Advanced Parameters

Automatic Congestion Management

If disabled, no policies will apply ACM. If enabled, only policies with ACM enabled will apply ACM

Global ACM Status: enabled

Shaping

If disabled, no shaping or rate limit will be applied, even if configured

Shaping per flow: enabled

Shaping per subscriber:

Shaping rate per subscriber: enabled

Figure 10: *Software bypass options*

Note

The bypassed traffic does not benefit from QoE features. It is not optimized, no metrics are recorded, and no policies are applied.

System users

The QoE system has two types of users:

- **Administrators:** With unrestricted access to the node functionality, including configuration changes and software installation. By default, a user named **bqnadm** is created with administration profile.
- **Operators:** With access only to data visualization. By default, a user named **bqnop** is created with operator profile.

An administrator can create, delete or modify system users from **Administration** > **Users**.

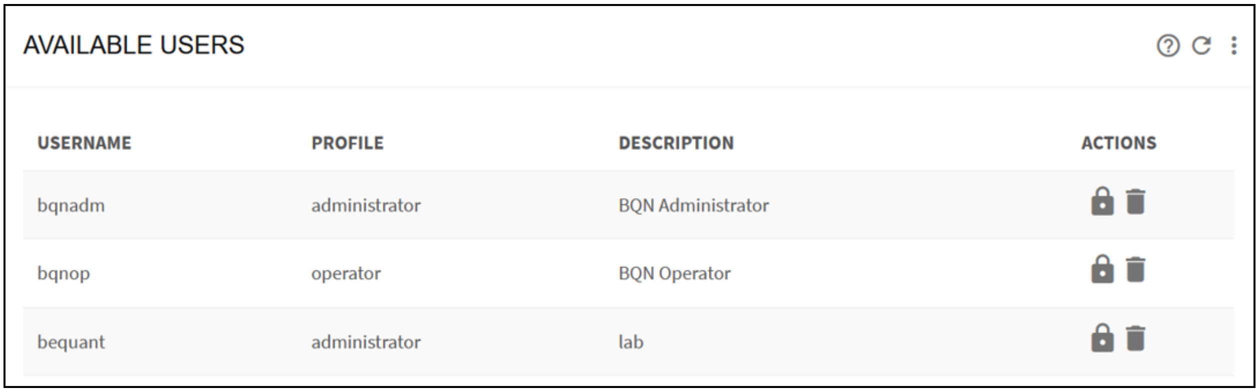

AVAILABLE USERS

USERNAME	PROFILE	DESCRIPTION	ACTIONS
bqnadm	administrator	BQN Administrator	
bqnop	operator	BQN Operator	
bequant	administrator	lab	

Figure 11: *The Available Users page*

Secure setup

Session timeout

A configurable timeout disconnects a GUI session after a time of inactivity.

To enable the inactivity timeout, navigate to **Administration** > **General Settings**, and set the value in seconds in GUI inactivity timeout and press **Apply**. Figure 12 shows the General Settings page.

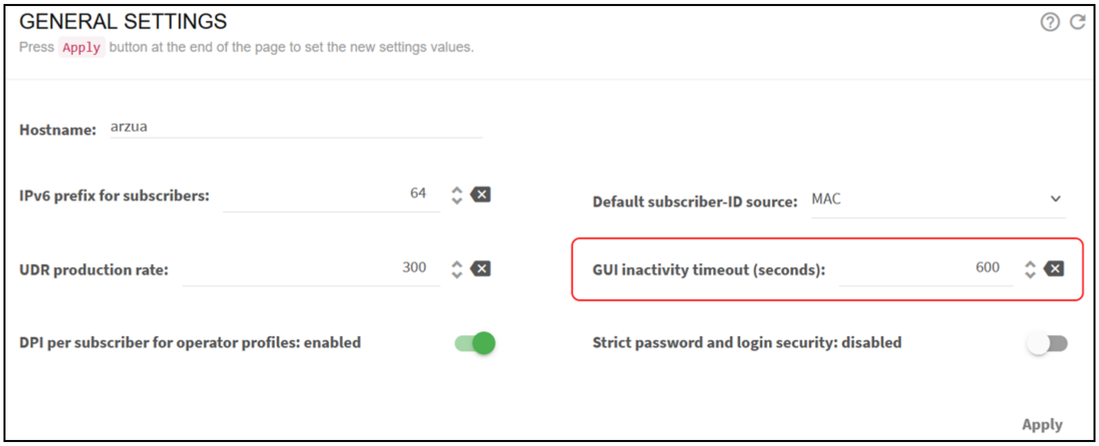

GENERAL SETTINGS

Press Apply button at the end of the page to set the new settings values.

Hostname: arzua

IPv6 prefix for subscribers: 64

Default subscriber-ID source: MAC

UDR production rate: 300

GUI inactivity timeout (seconds): 600

DPI per subscriber for operator profiles: enabled

Strict password and login security: disabled

Figure 12: *The General Settings page*

The inactivity timeout will be applied to new sessions.

Strong user passwords

By default, any value is valid while setting a user password. It is possible to strengthen the system security by adding some complexity to the user passwords. To set the password, navigate to **Administration** > **General Settings** and enable **Strict password and login security**.

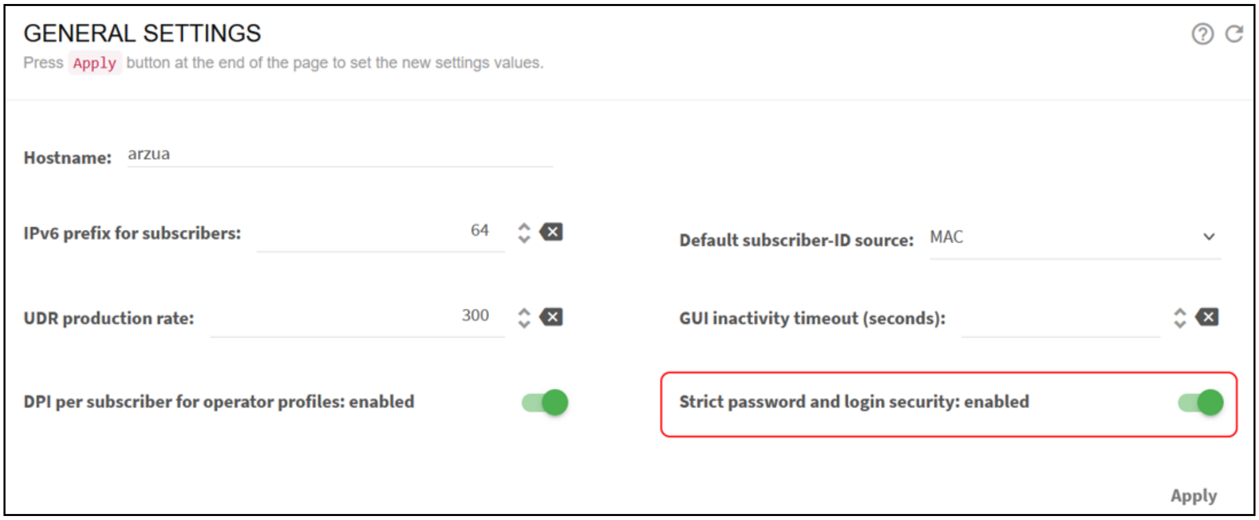

GENERAL SETTINGS

Press Apply button at the end of the page to set the new settings values.

Hostname: arzua

IPv6 prefix for subscribers: 64

Default subscriber-ID source: MAC

UDR production rate: 300

GUI inactivity timeout (seconds):

DPI per subscriber for operator profiles: enabled

Strict password and login security: enabled

Figure 13: *The General Settings page*

If strict password switch is enabled, then the following minimum password complexity is enforced (and password change rejected if not met):

- Length of at least eight characters
- At least one lowercase letter
- At least one uppercase letter
- At least one digit
- At least one special character
- The username cannot be part of the password, either straight or in reverse form. For example, if user is **bqnadm**, passwords **Bqnadm6?** and **Mdanqb6?** are rejected.

In addition, the password must pass the **pam_cracklib** simplicity test. This test rejects the following poor passwords:

- Dictionary words
- Palindromes (for example Af16-61fA)
- Same consecutive characters (for example ...aaa...)
- Too long monotonic sequence (for example ...123... or ...abc...)
- Less than five differences with the old password.

Also, the account is blocked for five minutes after five consecutive failed login attempts. Root is excluded from this policy to avoid denial of service attacks.

Management Interface Firewall

To set up the **Management Interface firewall**, which only applies to the management interface (not to the interfaces configured in wires) select on the lateral menu **Configuration** > **Interfaces** > **Management Firewall**. This shows the IP address ranges allowed to access the management interface. By default, no IP address ranges are configured and all are allowed.

To add an allowed IP address range, click and press the **Add IP Address Range...**. Once one IP address range is allowed and the Firewall is enabled, all incoming connections from IP addresses which are not part of the configured IP address ranges are blocked. It is therefore important to include an IP address range that covers the IP address from which we are accessing the GUI and also the subnet of the management IP address (the GUI includes them in the suggested list). Other IPs that interact with the management interface, such as RADIUS/REST clients, a billing system and the NTP server, should also be included.

To disable the Firewall, remove all entries pressing the delete icon next to each entry, and once all entries are deleted, click the **Apply** button. It is important not to press **Apply** before all entries have been deleted, because a premature **Apply** keeps the firewall active and it may prevent you from accessing the server, if the entry covering your IP is not present.

Hide Per-subscriber service information to operators

It is possible to configure the system, so users with operator profile do not see the following information:

- In **Subscriber dashboard**, the DPI service details
- In **Subscriber dashboard**, in the active flow table, the DOMAIN column.
- • In **Statistics** > **DPI Service Analysis** > **Hourly Volume Per Service**, the IP/Subscriber ID filter.
- • In **Statistics** > **DPI Service Analysis** > **Total Volume Per Service**, the IP/Subscriber ID filter.

To disable access to that information, go as administration to **Administration** > **General Settings** and disable the switch **DPI per subscriber for operator profiles** as shown in Figure 14.

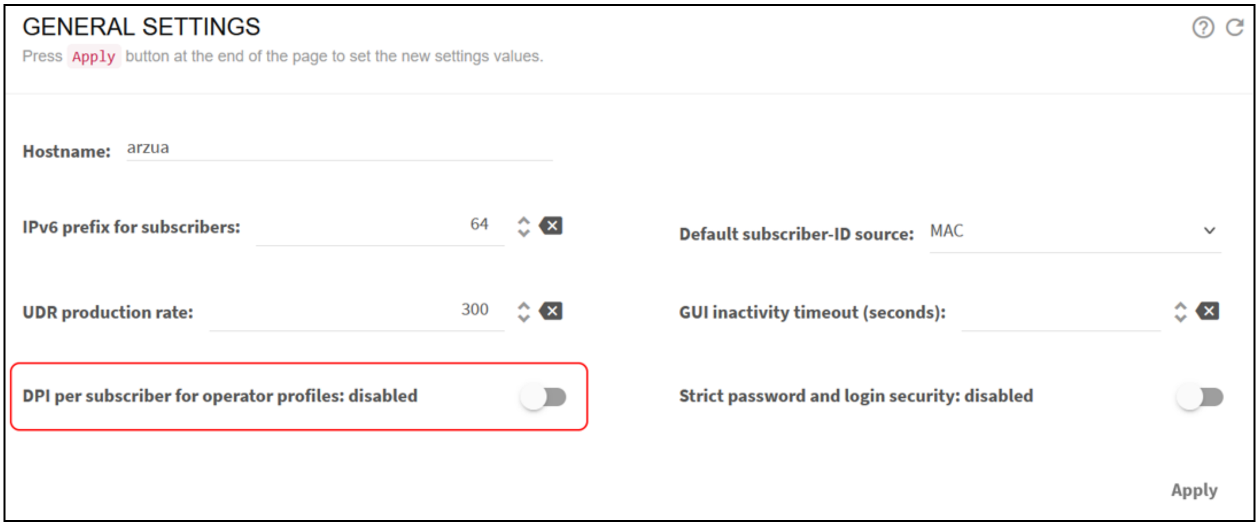

GENERAL SETTINGS

Press Apply button at the end of the page to set the new settings values.

Hostname: arzua

IPv6 prefix for subscribers: 64

Default subscriber-ID source: MAC

UDR production rate: 300

GUI inactivity timeout (seconds):

DPI per subscriber for operator profiles: disabled

Strict password and login security: disabled

Figure 14: *The General Settings page*

Audit log

The system keeps an audit log with a register of the most relevant actions performed in the system. The files are in */opt/bqn/var/audit* and readable only by the **root** user.

The current audit file is called audit and old audit files are compressed with gzip and named with the Unix epoch time of rotation (for example **audit-1688636727.gz**). Old files are kept for 182 days.

Each file row is an audit entry with the following fields:

- **Time:** Date and time of the event, in format YYYY-mm-ddTHH:MM:SS+UTC-Offset
- **Type:** Type of action: access, config, software, system, users

- **Author:** Name of the system user performing the action, in the cases where it is available
- **Description:** Action description.

The following are some of the registered actions:

- Access to the system
- Users created/deleted
- User password modifications
- Configuration changes
- Software updates
- System reboot or shutdown
- Time local/zone changes

Chapter 5: TCP Optimization

This chapter contains the following sections:

- TCPO metrics
- Configuring the TCP optimization

TCPO metrics

The QoE accelerates TCP traffic to raise the effective speed of data transfers and improve the user quality of experience.

Navigate to **Status > TCP Optimization > Speed & Acceleration** to view the average speed values of non-accelerated TCP connections (no TCPO), accelerated ones (TCPO), and the percentage of speed increase during the last 24 hours (n/s if no statistically significant difference found because of variability or not enough samples). The following metrics are shown in the TCP speed and acceleration page:

- TCP average speeds of all network traffic.
- TCP average speeds of main services.
- TCP average speeds per Internet latency (between the QoE and the Internet). Up to three latency ranges are considered. Latency thresholds are configurable in **Configuration** > **TCPO/ACM Settings** (latencies change from network to network and the default values may not be suitable).

Latency thresholds are configurable in **Configuration > TCPO/ACM Settings** (latencies change from network to network and the default values may not be suitable). Figure 15 shows the TCP speed and acceleration page.

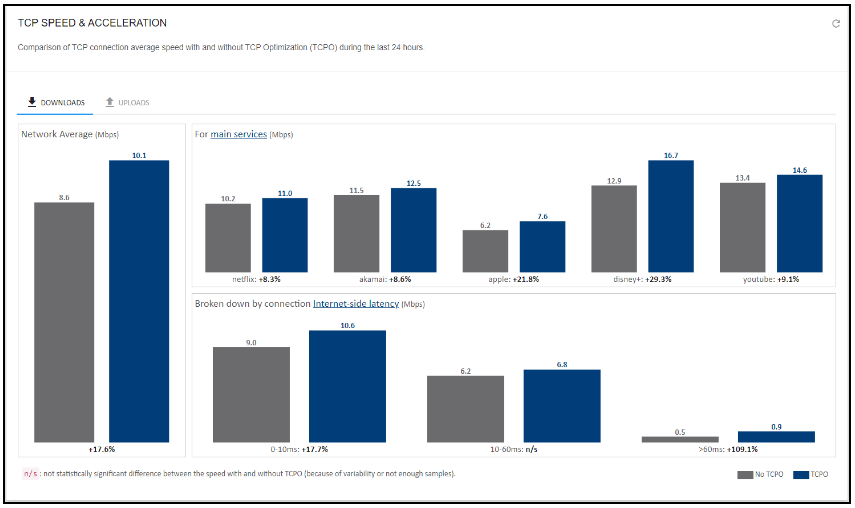

TCP SPEED & ACCELERATION

Comparison of TCP connection average speed with and without TCP Optimization (TCPO) during the last 24 hours.

n/s : not statistically significant difference between the speed with and without TCPO (because of variability or not enough samples).

Figure 15: *The TCP speed and acceleration page*

> **Note**
>
> There may be categories or graph periods without any information, due to no traffic or little traffic occasionally. To view the information, extend the period of calculation (1 day by default).

To view the speed evolution over time, navigate to **Statistics** > **TCP Optimization** > **TCP Speed**. Figure 16 shows the speed evolution over time graph.

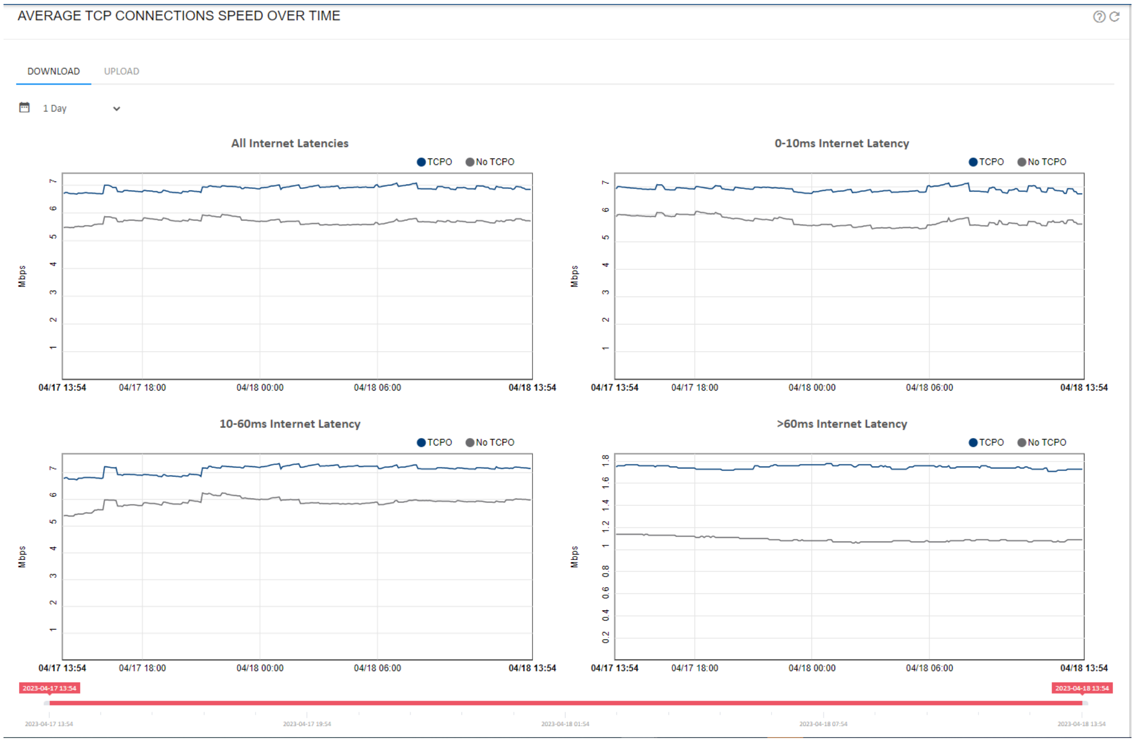

Figure 16: *Speed evolution over time*

Note

Sometimes, due to no or little traffic, there might be categories or graph periods without information.

To view the acceleration evolution overtime, navigate to **Statistics** > **TCP Optimization** > **TCP Acceleration**. Figure 17 shows the acceleration evolution overtime graph.

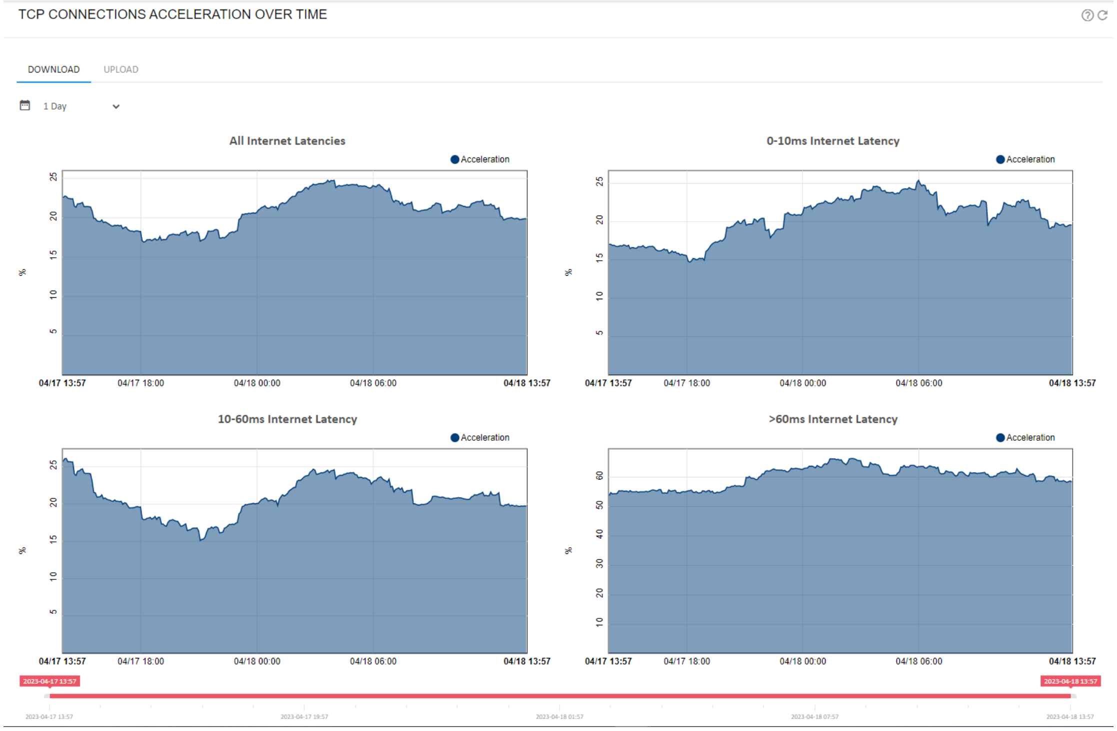

Figure 17: *Acceleration evolution overtime*

Configuring the TCP optimization

To configure TCP Optimization parameters, navigate to **Configuration > TCPO/ACM Settings**.

The latency thresholds can be changed in RTTi-small and RTTi-large. To view the Internet latencies of network, navigate to **Statistics > System > Latency** and **Statistics > DPI Service Analysis > Latency per Service**.

Latency thresholds only affect the way the metrics are classified and displayed, it does not impact TCP optimization, which adjusts continuously to each individual connection characteristics, latencies included. Figure 18 shows the TCPO/ACM settings page.

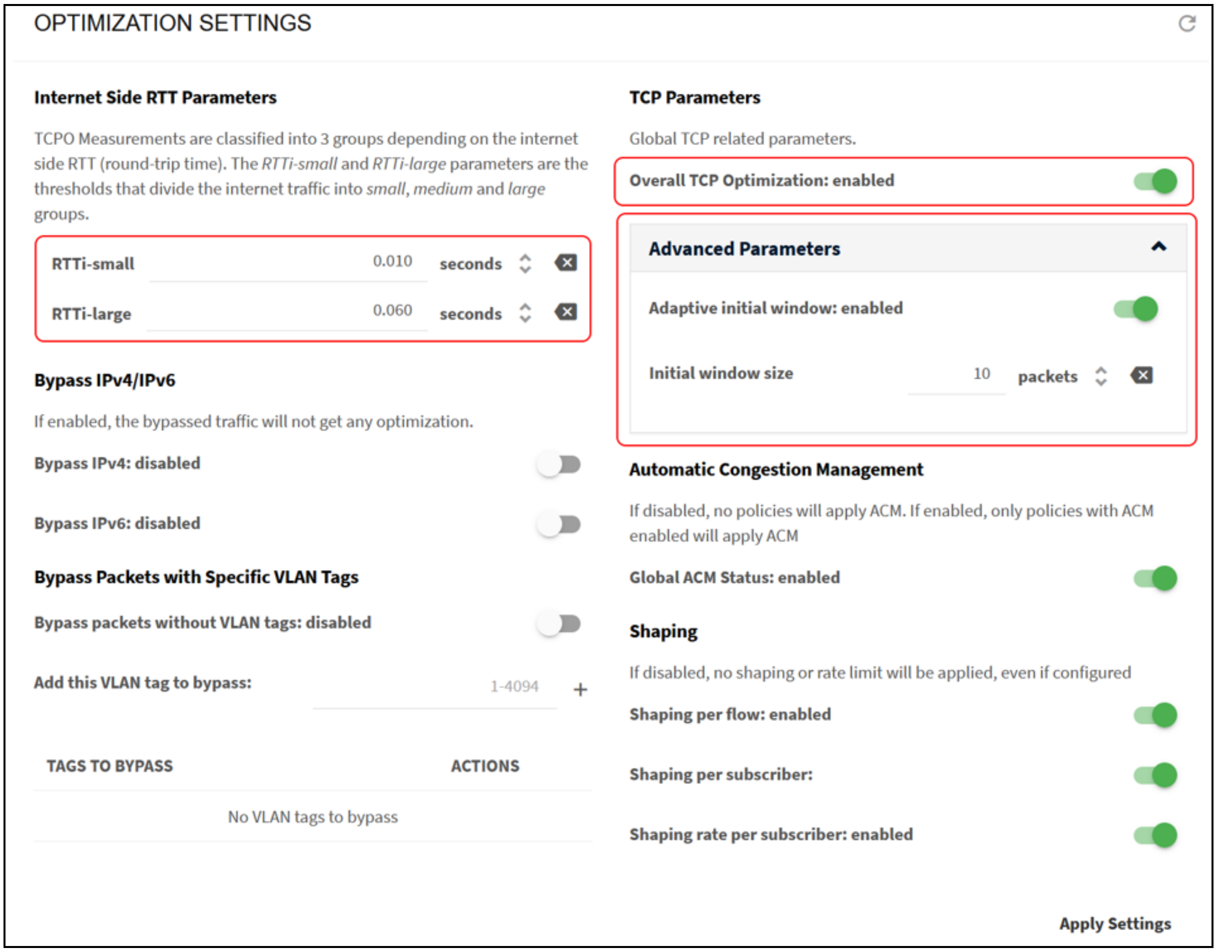

OPTIMIZATION SETTINGS

Internet Side RTT Parameters

TCPO Measurements are classified into 3 groups depending on the internet side RTT (round-trip time). The *RTTi-small* and *RTTi-large* parameters are the thresholds that divide the internet traffic into *small*, *medium* and *large* groups.

RTTi-small 0.010 seconds

RTTi-large 0.060 seconds

Bypass IPv4/IPv6

If enabled, the bypassed traffic will not get any optimization.

Bypass IPv4: disabled

Bypass IPv6: disabled

Bypass Packets with Specific VLAN Tags

Bypass packets without VLAN tags: disabled

Add this VLAN tag to bypass: 1-4094 +

TAGS TO BYPASS	ACTIONS
No VLAN tags to bypass	

TCP Parameters

Global TCP related parameters.

Overall TCP Optimization: enabled

Advanced Parameters

Adaptive initial window: enabled

Initial window size 10 packets

Automatic Congestion Management

If disabled, no policies will apply ACM. If enabled, only policies with ACM enabled will apply ACM

Global ACM Status: enabled

Shaping

If disabled, no shaping or rate limit will be applied, even if configured

Shaping per flow: enabled

Shaping per subscriber:

Shaping rate per subscriber: enabled

Figure 18: *TCPO/ACM settings page*

To disable all TCP optimization in the QoE, regardless of what the flow policy dictates, switch off the toggle **Overall TCP Optimization**. This is done to disable TCPO temporally, while debugging an issue or during maintenance tasks.

If there is a NAT between the QoE and the end users, the **Adaptive Initial Window** should be OFF (ON by default), because the adaptive algorithm that tries to find the best initial window per subscriber is no longer meaningful when a lot of real subscribers sit behind a NAT IP address that the QoE is treating as an individual subscriber.

It is also possible to modify the TCP initial window from its default of 10 packets. Changes to the initial window are only recommended for paths with very high latencies, such as a satellite link.

Chapter 6: Automatic Congestion Management (ACM) Optimization

When the subscriber rate limits are unknown, the QoE can automatically detect them using machine learning. Then the QoE becomes the bandwidth management element, and the network can benefit from QoE reduced latencies and losses. The ACM also detects congestions below the subscriber rate limit and helps when rate limits are known.

This chapter contains the following sections:

- Metrics
- Configuring the ACM optimization

Metrics

The overall reduction of latency and losses is achieved through the Metrics. To check the metrics, navigate to **Statistics** > **Congestion** > **ACM and Congestion**. Figure 19 shows the ACM Metrics.

The following charts are displayed:

- **Traffic at Max Speed & under Congestion** shows the percentage of traffic that is running at the maximum speed or near the maximum speed, along with the percentage of traffic suffering congestion and the percentage of traffic limited by the ACM.
- **Latency Reduction with ACM** contains the reduction in milliseconds of the access latency achieved.
- **Retransmission Reduction with ACM** depicts the reduction in the packet retransmission percentage obtained by the ACM.

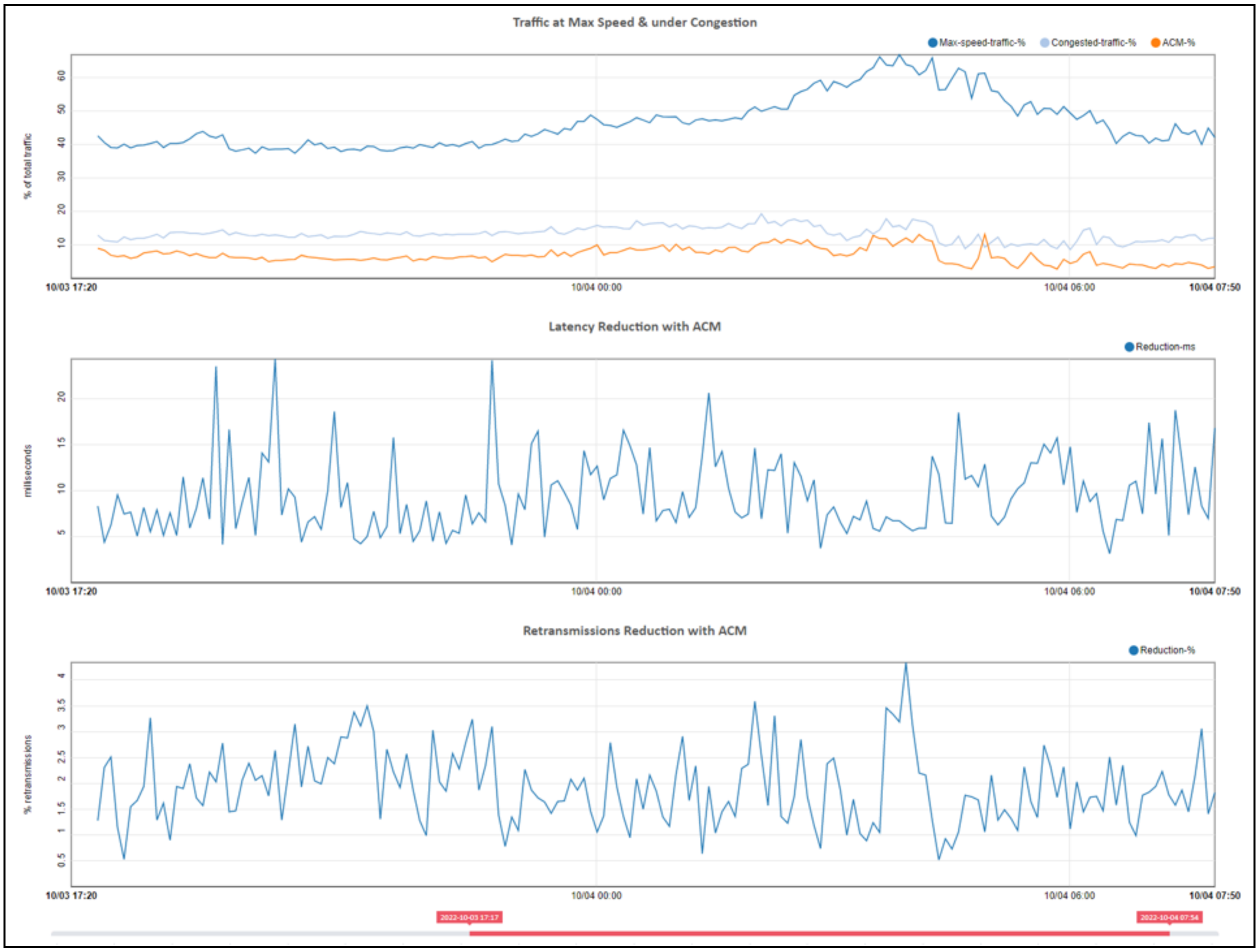

Figure 19: *ACM Metrics*

Also, these charts are available per Subscriber. Navigate to **Status** > **Subscribers**, click **Subscriber ID** or IP address to go to the subscriber dashboard.

Configuring the ACM optimization

ACM improves the network quality without fine tuning of the configuration. So it is recommended to keep it enabled at all times. To disable the ACM for all policies, navigate to **Configuration** > **Optimization Settings**and disable **Global ACM Status**. Figure 20 shows the Optimization Settings page.

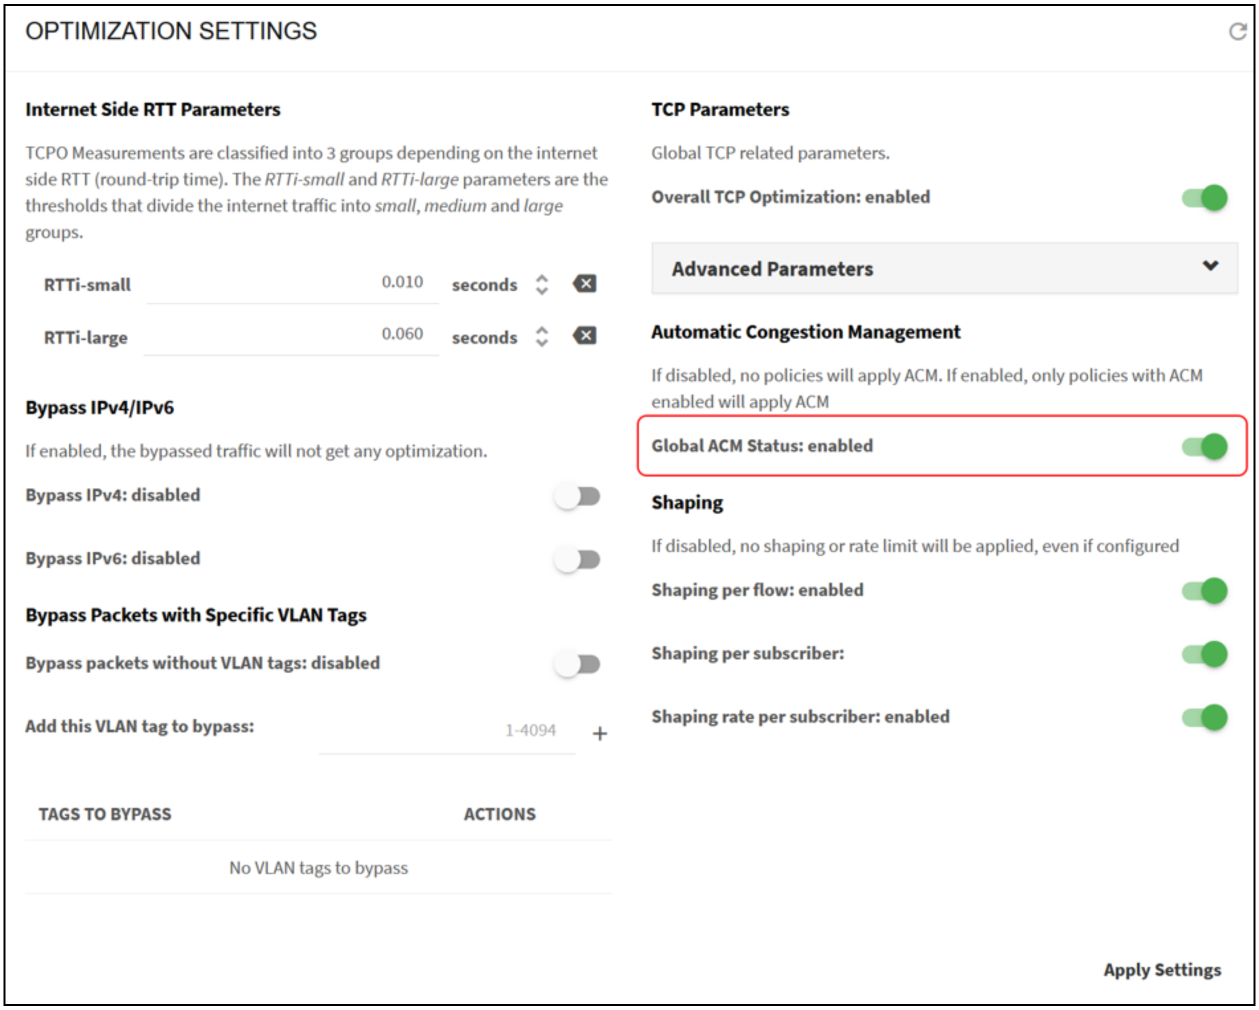

OPTIMIZATION SETTINGS

Internet Side RTT Parameters

TCPO Measurements are classified into 3 groups depending on the internet side RTT (round-trip time). The *RTTi-small* and *RTTi-large* parameters are the thresholds that divide the internet traffic into *small*, *medium* and *large* groups.

RTTi-small 0.010 seconds

RTTi-large 0.060 seconds

Bypass IPv4/IPv6

If enabled, the bypassed traffic will not get any optimization.

Bypass IPv4: disabled

Bypass IPv6: disabled

Bypass Packets with Specific VLAN Tags

Bypass packets without VLAN tags: disabled

Add this VLAN tag to bypass: 1-4094 +

TAGS TO BYPASS	ACTIONS
No VLAN tags to bypass	

TCP Parameters

Global TCP related parameters.

Overall TCP Optimization: enabled

Advanced Parameters

Automatic Congestion Management

If disabled, no policies will apply ACM. If enabled, only policies with ACM enabled will apply ACM

Global ACM Status: enabled

Shaping

If disabled, no shaping or rate limit will be applied, even if configured

Shaping per flow: enabled

Shaping per subscriber:

Shaping rate per subscriber: enabled

Apply Settings

Figure 20: *The Optimization Settings page*

Chapter 7: Network Visibility

This chapter contains the following sections:

Subscriber dashboard

The subscriber dashboard includes a comprehensive set of useful information about the current and past performance of the subscriber. It is used to analyze and diagnose the issues reported in subscriber data access.

To view the subscriber dashboard, click on the Subscriber IP address or Subscriber ID from the following windows:

- Status > Subscribers
- Status > Flows > Per Subscriber
- Status > RADIUS/REST/Billing > Subscribers

Use the right-side scroll bar to browse the subscriber dashboard information. Figure 21 shows the subscriber dashboard.

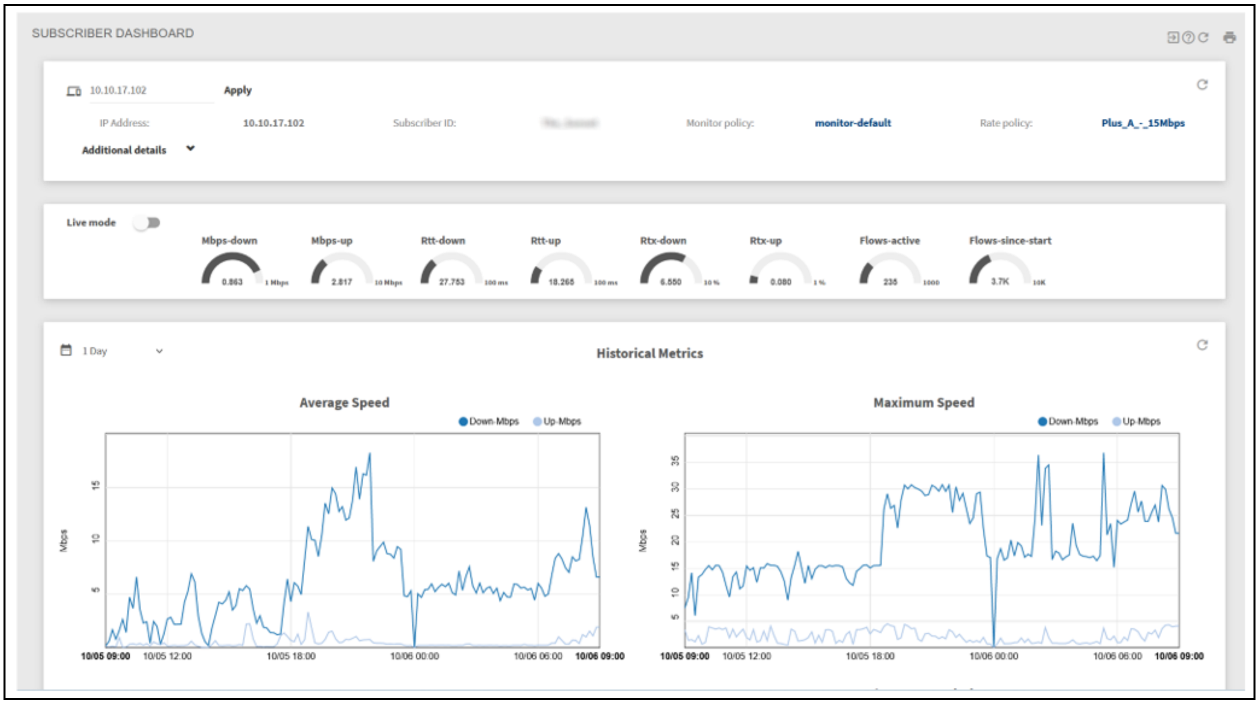

SUBSCRIBER DASHBOARD

10.10.17.102 Apply

IP Address: 10.10.17.102 | Subscriber ID: [redacted] | Monitor policy: monitor-default | Rate policy: Plus_A_-_15Mbps

Additional details

Live mode

Mbps-down 0.863 (1 Mbps) | Mbps-up 2.817 (10 Mbps) | Rtt-down 27.753 (100 ms) | Rtt-up 18.265 (100 ms) | Rtx-down 6.550 (10 %) | Rtx-up 0.080 (1 %) | Flows-active 235 (1000) | Flows-since-start 3.7K (10K)

1 Day

Historical Metrics

Figure 21: *Subscriber dashboard*

The dashboard contains the following information:

- Subscriber main session parameters (including the rate and monitor policies applied, IP and subscriber ID, and so on. Click on **Additional details** to see more information.
- Dial icons with a summary of the latest metrics.
- Charts with the evolution over time of the subscriber main metrics:
  - average and maximum speeds
  - latencies
  - packet retransmissions (losses)
  - traffic at high speeds, suffering congestion and limited by ACM
  - reductions in latency and losses due to ACM
  - number of flows (created per minute and active).
- Chart with the subscriber application usage over time.
- Internet latencies experienced by the subscriber for mostly used services. Clicking on the chart icons at the end of a service entry, will take you to the histogram with the latency distribution and to the histogram evolution over time.
- Table with the active flows of the subscriber.

Figure 22 shows the evolution over time of average and maximum speed, number of flows and access latency. For the access latency, the network average is also included as a reference, to determine of the subscriber is below or above the network average quality.

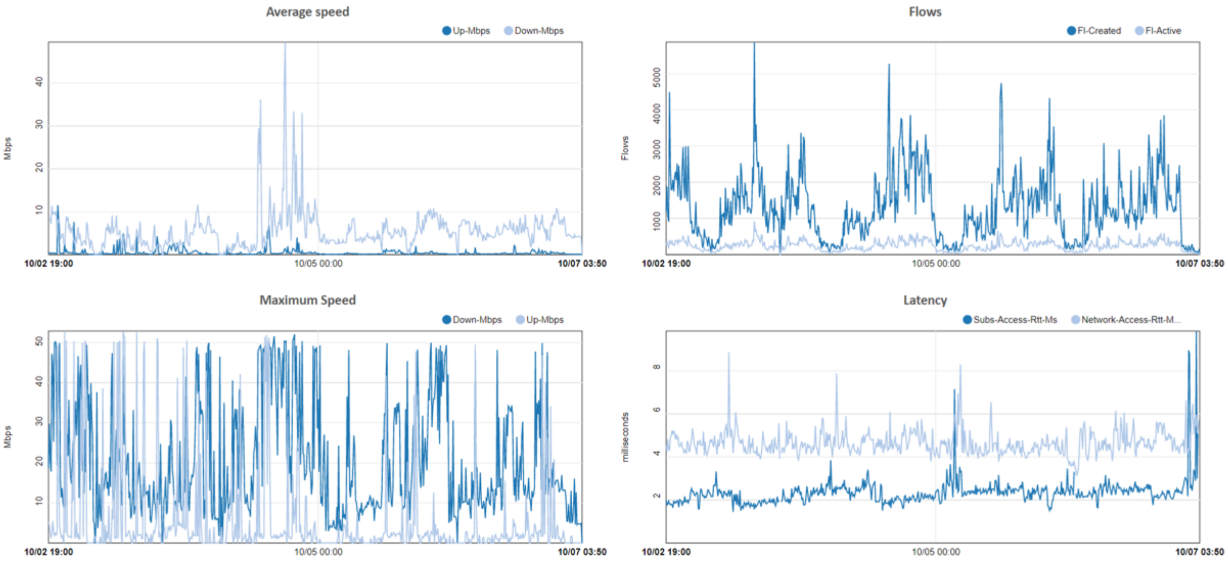

Figure 22: *Average and maximum speed*

Figure 23 shows the packet retransmission on the access side, also including the network average as a reference. The last three charts contain information about the Automatic Congestion Management (ACM) feature for the subscriber.

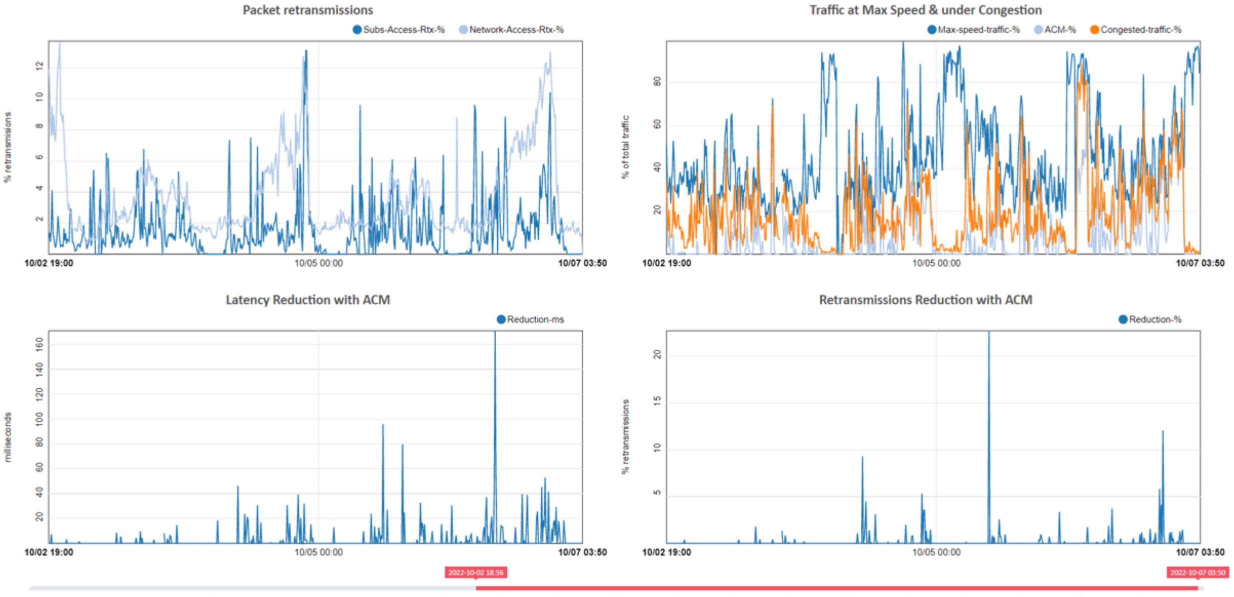

Figure 23: *Packet retransmission*

To monitor the subscriber more closely, set the **Live mode** switch to ON. Charts will appear with the current evolution of the main metrics. A flow table will appear and will be continuously refreshed. Flows are ordered in the table with active flows first, and among them, the faster flows. Not null speeds are highlighted in blue to facilitate the identification of active flows. You can display the life charts with more vertical space clicking on the Layout ↕ icon.

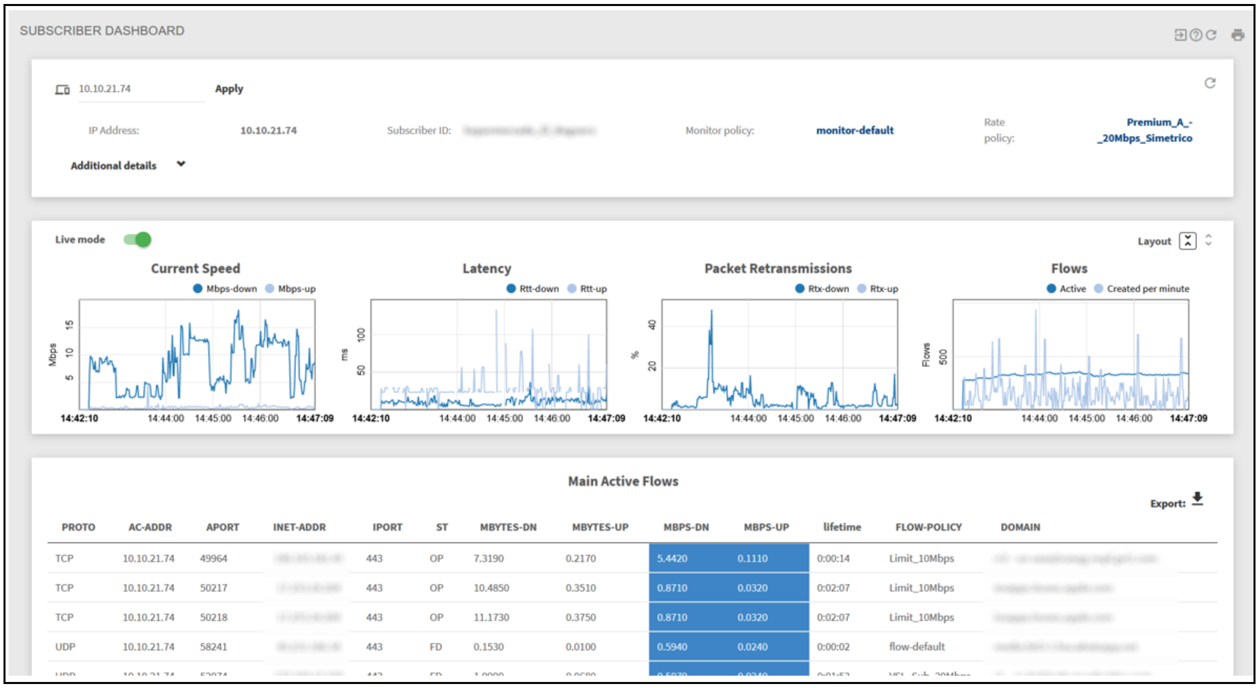

10.10.21.74 Apply

IP Address: 10.10.21.74 Subscriber ID: [redacted] Monitor policy: monitor-default Rate policy: Premium_A_-_20Mbps_Simetrico

Additional details

Main Active Flows

PROTO	AC-ADDR	APORT	INET-ADDR	IPORT	ST	MBYTES-DN	MBYTES-UP	MBPS-DN	MBPS-UP	lifetime	FLOW-POLICY	DOMAIN
TCP	10.10.21.74	49964	[redacted]	443	OP	7.3190	0.2170	5.4420	0.1110	0:00:14	Limit_10Mbps	[redacted]
TCP	10.10.21.74	50217	[redacted]	443	OP	10.4850	0.3510	0.8710	0.0320	0:02:07	Limit_10Mbps	[redacted]
TCP	10.10.21.74	50218	[redacted]	443	OP	11.1730	0.3750	0.8710	0.0320	0:02:07	Limit_10Mbps	[redacted]
UDP	10.10.21.74	58241	[redacted]	443	FD	0.1530	0.0100	0.5940	0.0240	0:00:02	flow-default	[redacted]
UDP	10.10.21.74	52074	[redacted]	443	FD	1.0000	0.0680	0.5070	0.0240	0:01:53	VSL_Sub_20Mbps	[redacted]

The live mode will end when the **Live mode** switch is set to OFF or after 15 minutes.

Subscriber Identifier (ID)

The Subscriber ID field is supported to facilitate the identification of a subscriber session. The subscriber ID can be used when requesting metrics to obtain historical information even if the subscriber is changed IP address over time.

The following are the possible sources for subscriber ID:

- **MAC access address** (this is the default). In some networks, the MAC address might be the same for all subscribers (for example if all traffic is coming from the same router port) but in other networks, it may identify the subscriber access points.
- **Subscriber access IP address**. To configure the IP address to fill the subscriber ID. Navigate to **Administration** > **General Settings** > **Default subscriber-ID source**.
- **An external system**: An external system can use the QoE REST API, RADIUS or one of our Billing. For more information, refer to *QoR REST API Guide*, *QoE RADIUS Guide*, and Billing Systems section.

Subscriber QoE Metrics page

Status > **Subscribers** > **QoE Metrics** gives access to the list of active subscribers with their metrics (if the subscriber of interest is not listed, type the IP address on the filter field). Figure 24 shows the active subscribers status page.

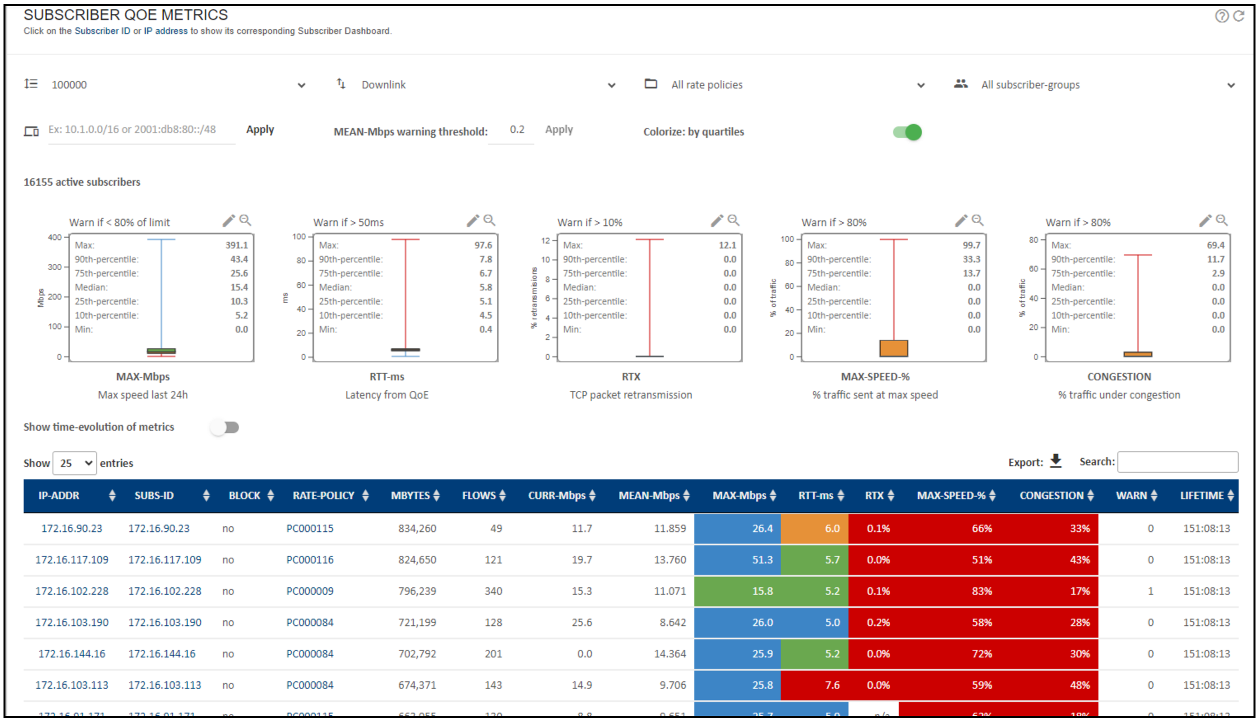
SUBSCRIBER QOE METRICS

Click on the Subscriber ID or IP address to show its corresponding Subscriber Dashboard.

100000 · Downlink · All rate policies · All subscriber-groups

Ex: 10.1.0.0/16 or 2001:db8:80::/48 Apply

MEAN-Mbps warning threshold: 0.2 Apply

Colorize: by quartiles

16155 active subscribers

	MAX-Mbps	RTT-ms	RTX	MAX-SPEED-%	CONGESTION
	Warn if < 80% of limit	Warn if > 50ms	Warn if > 10%	Warn if > 80%	Warn if > 80%
Max:	391.1	97.6	12.1	99.7	69.4
90th-percentile:	43.4	7.8	0.0	33.3	11.7
75th-percentile:	25.6	6.7	0.0	13.7	2.9
Median:	15.4	5.8	0.0	0.0	0.0
25th-percentile:	10.3	5.1	0.0	0.0	0.0
10th-percentile:	5.2	4.5	0.0	0.0	0.0
Min:	0.0	0.4	0.0	0.0	0.0
	Max speed last 24h	Latency from QoE	TCP packet retransmission	% traffic sent at max speed	% traffic under congestion

Show time-evolution of metrics

Show 25 entries · Export: · Search:

IP-ADDR	SUBS-ID	BLOCK	RATE-POLICY	MBYTES	FLOWS	CURR-Mbps	MEAN-Mbps	MAX-Mbps	RTT-ms	RTX	MAX-SPEED-%	CONGESTION	WARN	LIFETIME
172.16.90.23	172.16.90.23	no	PC000115	834,260	49	11.7	11.859	26.4	6.0	0.1%	66%	33%	0	151:08:13
172.16.117.109	172.16.117.109	no	PC000116	824,650	121	19.7	13.760	51.3	5.7	0.0%	51%	43%	0	151:08:13
172.16.102.228	172.16.102.228	no	PC000009	796,239	340	15.3	11.071	15.8	5.2	0.1%	83%	17%	1	151:08:13
172.16.103.190	172.16.103.190	no	PC000084	721,199	128	25.6	8.642	26.0	5.0	0.2%	58%	28%	0	151:08:13
172.16.144.16	172.16.144.16	no	PC000084	702,792	201	0.0	14.364	25.9	5.2	0.0%	72%	30%	0	151:08:13
172.16.103.113	172.16.103.113	no	PC000084	674,371	143	14.9	9.706	25.8	7.6	0.0%	59%	48%	0	151:08:13

Figure 24: *The active subscribers status page*

To facilitate the identification of subscribers with quality of experience issues, the **WARN** column displays a score with the number of metrics above their warning threshold. Sort the table so the highest WARN values are shown first. In the above example, the first three subscribers have access latency issues (values above a threshold set to 10 ms in this particular network). Also, the third customer has an issue with packet losses. The rest of the listed subscribers expend most of the time at their maximum speed, indicating that they should upgrade to a higher plan.

To avoid highlighting subscribers that are simply inactive, a threshold is applied to the mean speed (**MEAN-Mbps warning threshold**). If a customer does not reach at least this mean speed, then customer does not get a warning score and it will not be highlighted. The threshold can be changed, but it is not required.

Box plots

To facilitate the monitoring of subscriber's access quality, five box plot charts summarize the distribution of these key metrics:

- MAX-Mbps: Maximum speed in Mbps
- RTT-ms: Access latency in milliseconds
- RTX: Packet loss rate in percentage
- MAX-SPEED: Percentage of the traffic going at or near maximum speed
- CONGESTION: Percentage of the traffic suffering congestion.

A box plot is a summary of the distribution of a set of values. Figure 25 shows the maximum and minimum values, the median (value separating the lower half and higher half), and the first and third quartiles (values separating the lower 25% and the higher 25%, respectively):

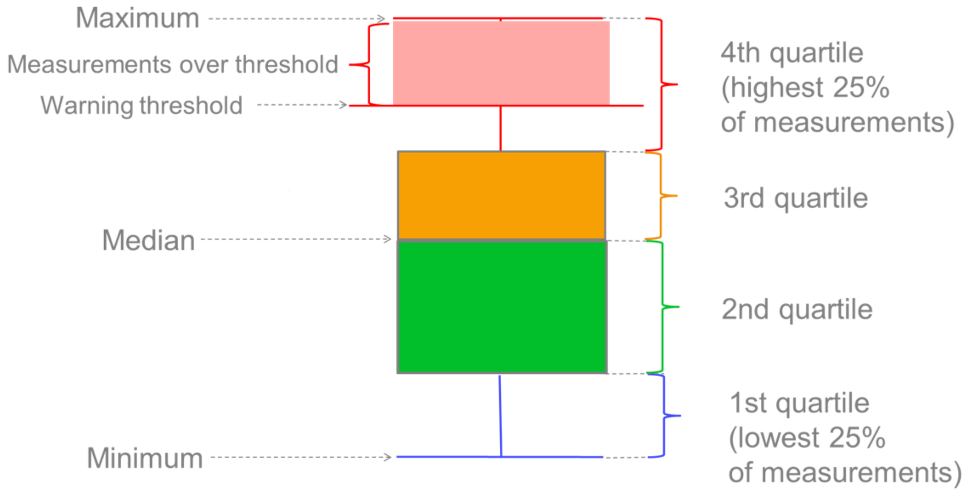

Figure 25: *Box plots*

The exact values of the percentiles are shown on the right-hand side of a box plot.

Warning thresholds

A configurable threshold helps to identify which subscribers are experiencing problems. For example, if the RTT threshold is set at 21ms, the box plot shades in red all values between 21ms and the maximum. In the table, the RTT values above the threshold is also shaded in red, making very easy to spot the affected customers. The table can sort the metric column to show the biggest values first (click on the column label to do that).

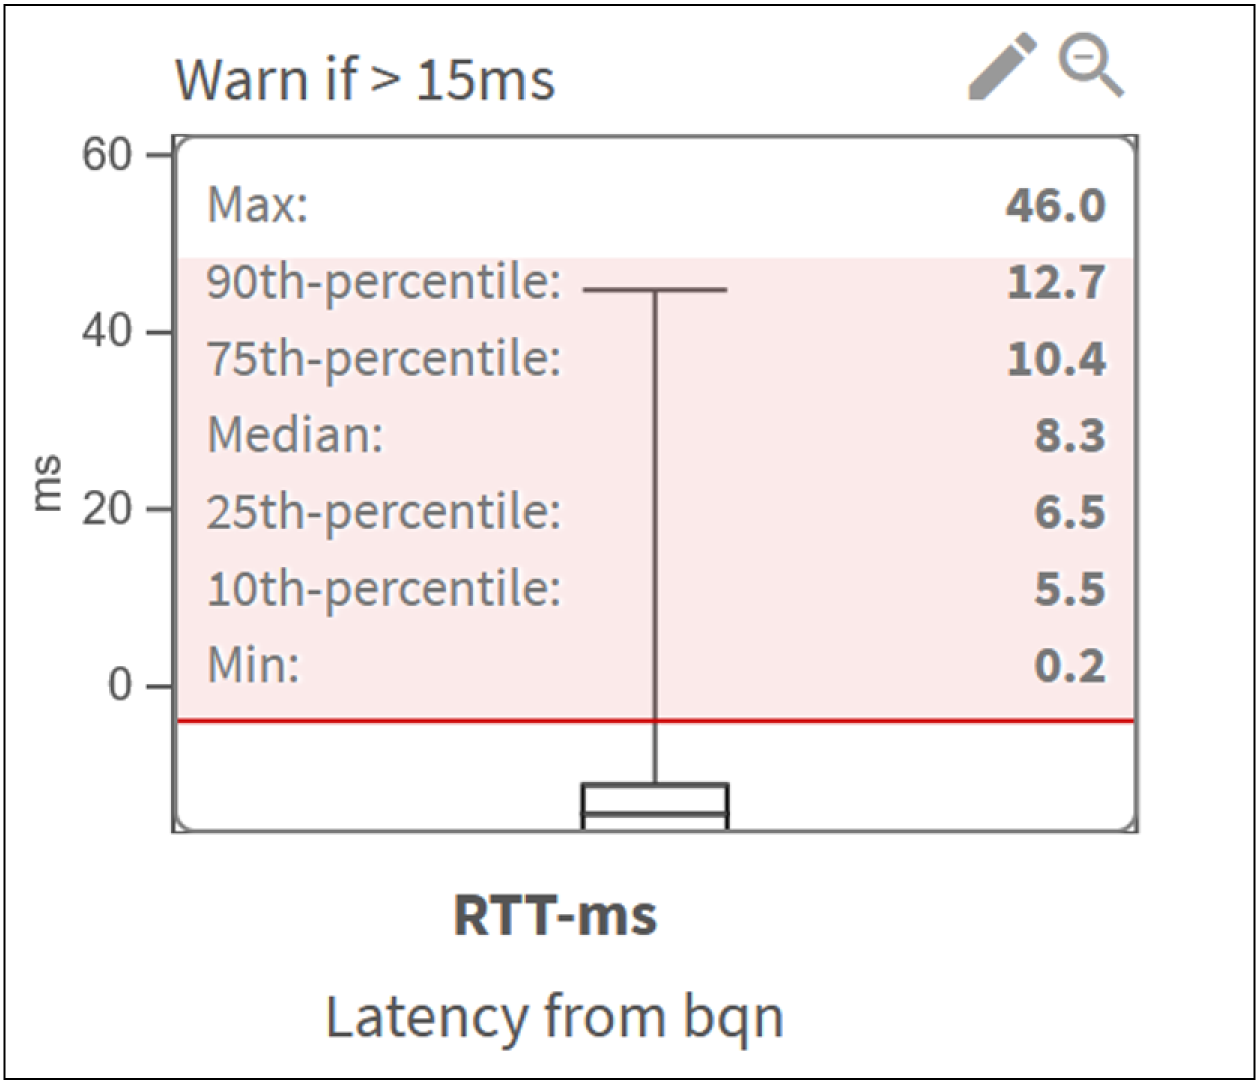

Figure 26: *TCP packet retransmission*

The value of the warning threshold is shown on top of the box plot. To adjust the value, click the **Edit** icon at the top-right of the box plot.

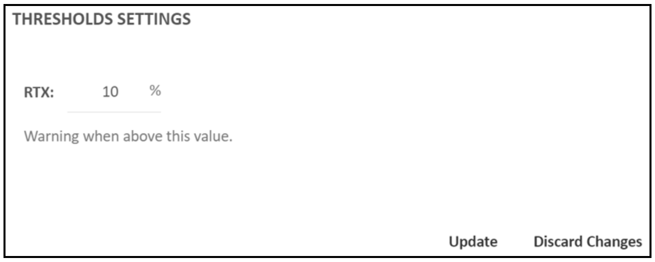
THRESHOLDS SETTINGS

RTX: 10 %

Warning when above this value.

Highlight percentiles

To view the quartile of a value, set the colorize switch at the top-right of the page to ON. Figure 27 shows the Subsriber QoE metrics.

The following are the quartile codes:

- First quartile: Blue
- Second quartile (up to median): Green
- Third quartile: Orange
- Forth quartile (from 3rd up to maximum): Red

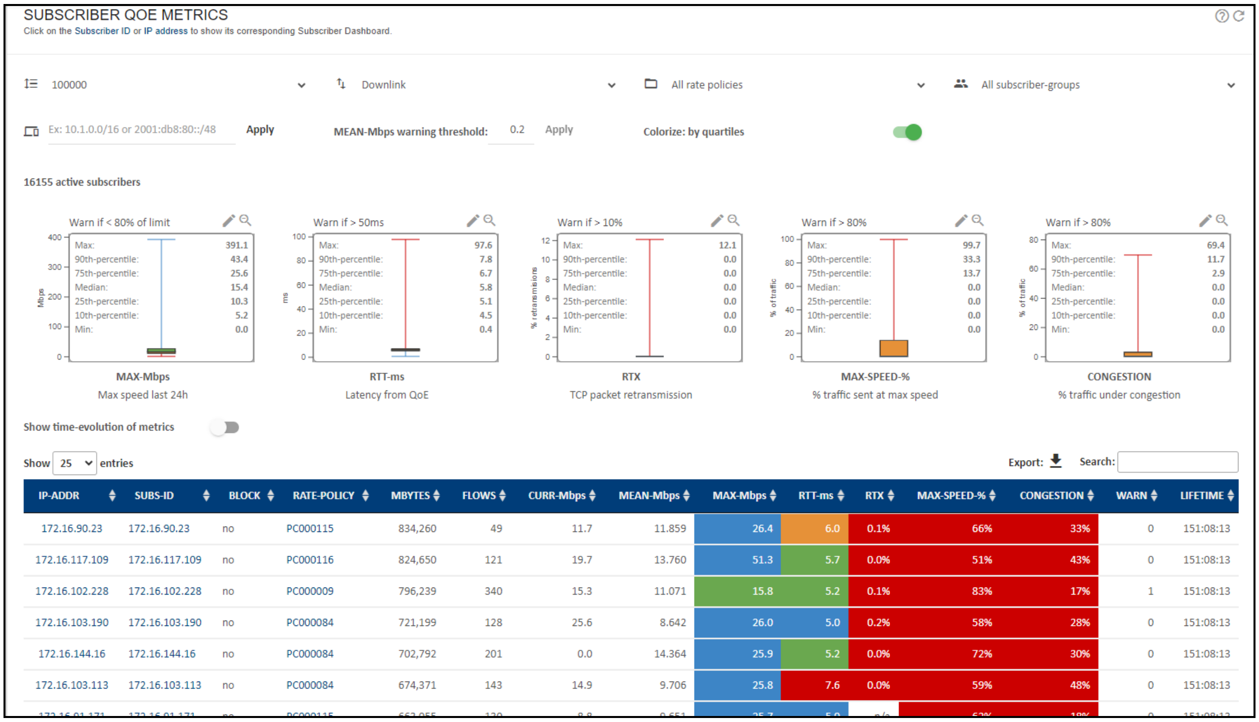

SUBSCRIBER QOE METRICS

Click on the Subscriber ID or IP address to show its corresponding Subscriber Dashboard.

100000 | Downlink | All rate policies | All subscriber-groups

Ex: 10.1.0.0/16 or 2001:db8:80::/48 Apply

MEAN-Mbps warning threshold: 0.2 Apply

Colorize: by quartiles

16155 active subscribers

	MAX-Mbps	RTT-ms	RTX	MAX-SPEED-%	CONGESTION
	Warn if < 80% of limit	Warn if > 50ms	Warn if > 10%	Warn if > 80%	Warn if > 80%
Max:	391.1	97.6	12.1	99.7	69.4
90th-percentile:	43.4	7.8	0.0	33.3	11.7
75th-percentile:	25.6	6.7	0.0	13.7	2.9
Median:	15.4	5.8	0.0	0.0	0.0
25th-percentile:	10.3	5.1	0.0	0.0	0.0
10th-percentile:	5.2	4.5	0.0	0.0	0.0
Min:	0.0	0.4	0.0	0.0	0.0
	Max speed last 24h	Latency from QoE	TCP packet retransmission	% traffic sent at max speed	% traffic under congestion

Show time-evolution of metrics

Show 25 entries

Export: Search:

IP-ADDR	SUBS-ID	BLOCK	RATE-POLICY	MBYTES	FLOWS	CURR-Mbps	MEAN-Mbps	MAX-Mbps	RTT-ms	RTX	MAX-SPEED-%	CONGESTION	WARN	LIFETIME
172.16.90.23	172.16.90.23	no	PC000115	834,260	49	11.7	11.859	26.4	6.0	0.1%	66%	33%	0	151:08:13
172.16.117.109	172.16.117.109	no	PC000116	824,650	121	19.7	13.760	51.3	5.7	0.0%	51%	43%	0	151:08:13
172.16.102.228	172.16.102.228	no	PC000009	796,239	340	15.3	11.071	15.8	5.2	0.1%	83%	17%	1	151:08:13
172.16.103.190	172.16.103.190	no	PC000084	721,199	128	25.6	8.642	26.0	5.0	0.2%	58%	28%	0	151:08:13
172.16.144.16	172.16.144.16	no	PC000084	702,792	201	0.0	14.364	25.9	5.2	0.0%	72%	30%	0	151:08:13
172.16.103.113	172.16.103.113	no	PC000084	674,371	143	14.9	9.706	25.8	7.6	0.0%	59%	48%	0	151:08:13

Figure 27: *Subsriber QoE metrics*

Active subscriber table

The information provided is as follows:

- **IP-ADDR**: IP address of the subscriber.
- **SUBS-ID**: Subscriber identifier. MAC address by default, though it can be filled using external system information (see RADIUS and Billing sections).
- **RATE-POLICY**: Name of the rate policy being applied to this subscriber.
- **TOTAL-MBYTES**: Total traffic volume of this subscriber session, in megabytes.
- **ACTIVE-FLOWS**: Total number of traffic flows (mainly TCP connections and UDP flows) of this subscriber that are active (exchanging traffic).
- **CURR-Mbps**: Current speed in Mbps.
- **MAX-Mbps**: Maximum speed in Mbps over one day period.
- **RTT-ms**: Minimum access latency in milliseconds over one day period.
- **RTX**: Average percentage of packet losses over one day period.
- **MAX-SPEED**: Percentage of traffic with a speed close to maximum speed, over one day period.

- **CONGESTION**: Percentage of traffic suffering congestion, over one day period.
- **LIFETIME**: Duration so far of this subscriber session.

Click subscriber IP address or ID to view the subscriber dashboard, with historical data up to three months.

Metrics over time

To see the temporal evolution of a metric, set the switch **Show time-evolution of metrics**.

Select the metric to be showed (Average speed by default).

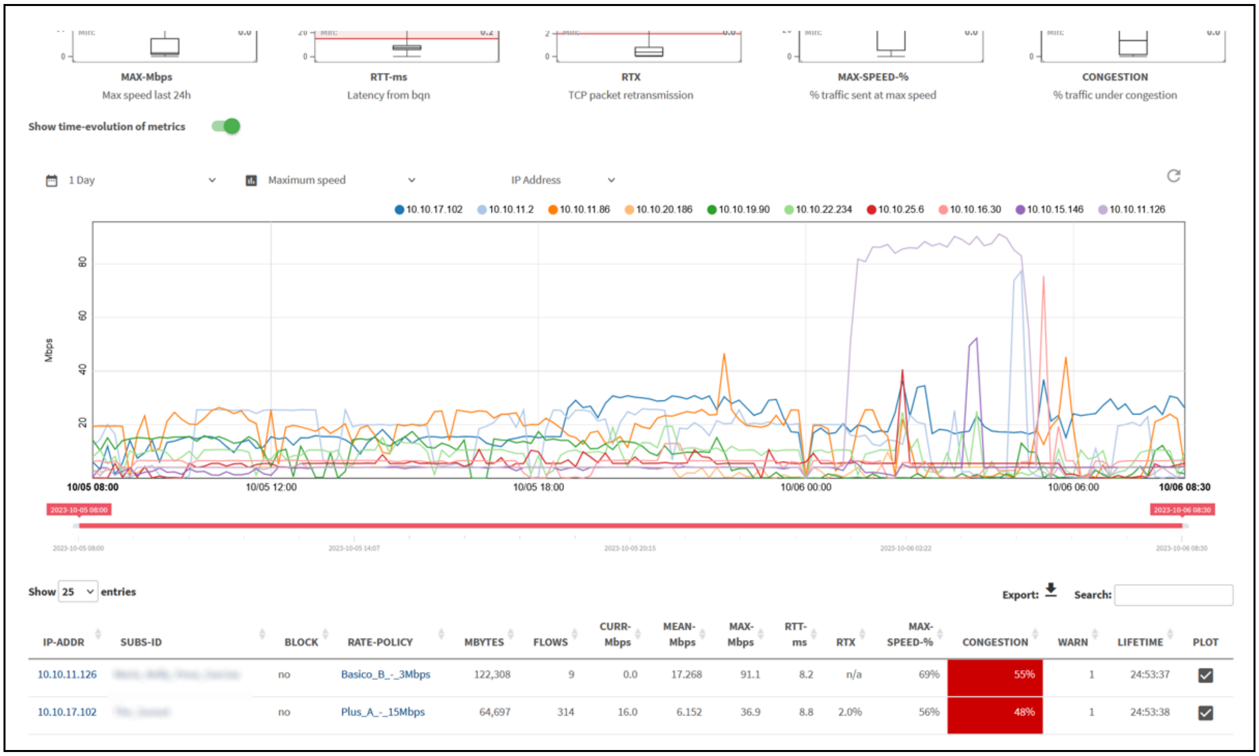

IP-ADDR	SUBS-ID	BLOCK	RATE-POLICY	MBYTES	FLOWS	CURR-Mbps	MEAN-Mbps	MAX-Mbps	RTT-ms	RTX	MAX-SPEED-%	CONGESTION	WARN	LIFETIME	PLOT
10.10.11.126	[redacted]	no	Basico_B_-_3Mbps	122,308	9	0.0	17.268	91.1	8.2	n/a	69%	55%	1	24:53:37	☑
10.10.17.102	[redacted]	no	Plus_A_-_15Mbps	64,697	314	16.0	6.152	36.9	8.8	2.0%	56%	48%	1	24:53:38	☑

Figure 28: *Metrics over time*

The chart will show by default the first 10 subscribers of the metrics table, and a **PLOT** column in the table indicates the subscribers included in the chart. Up to 30 subscribers can be shown at the same time selecting their plot tick box and refreshing the chart (reload icon in the upper right of the chart). Unselecting the plot tick will remove that subscriber from the chart.

The chart will show metrics of the IPs or subscriber IDs filtered in the metrics table, so you can use rate policy or subscriber group selectors to see their particular metrics evolution.

Right-clicking on an IP or ID of the chart legend will take you to the subscriber's dashboard.

Top subscribers by usage

Statistics > Subscribers Analysis > Hourly Volume displays the subscriber IP addresses with the largest traffic consumption over time. The usage is displayed per IP address or by Subscriber ID using the respective selection. **Statistics > Subscribers Analysis >Total Volume** displays the total in the given period. Figure 29 shows the usage per IP address.

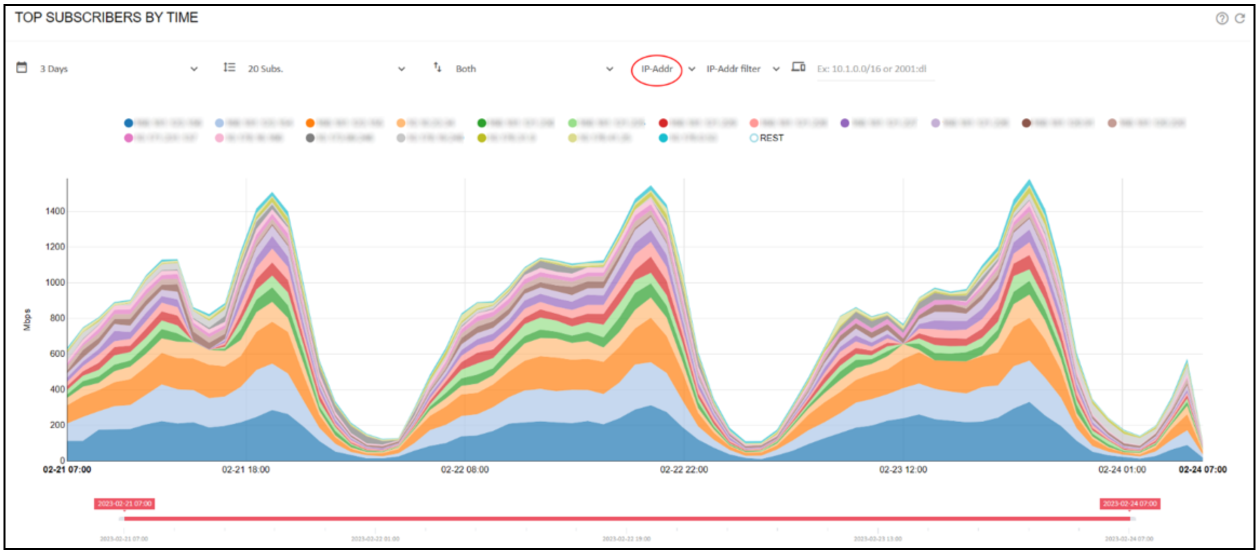

TOP SUBSCRIBERS BY TIME

Figure 29: *Top subscribers by time page*

You can access the subscriber dashboard page by right-clicking on the IP address or subscriber ID in the chart legend.

Statistics > **Subscribers Analysis** > **Total Volume** displays the total in the selected period. Figure 30 shows the usage per subscriber ID.

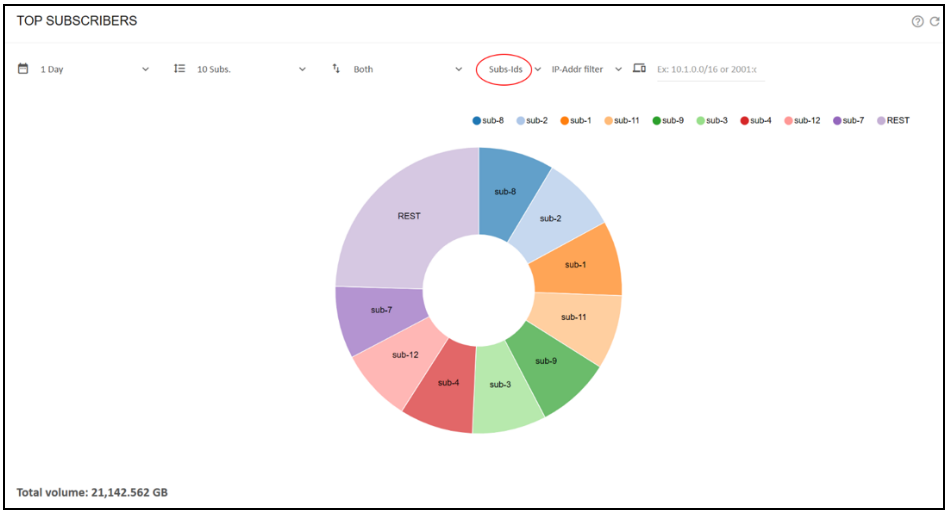

TOP SUBSCRIBERS

Total volume: 21,142.562 GB

Figure 30: *The top subscribers page*

You can access the subscriber dashboard page by clicking on the IP address or subscriber ID in the pie chart.

Traffic and subscribers per policy

Statistics > Subscribers > Per Policy shows the split of the subscribers into the different rate policies:

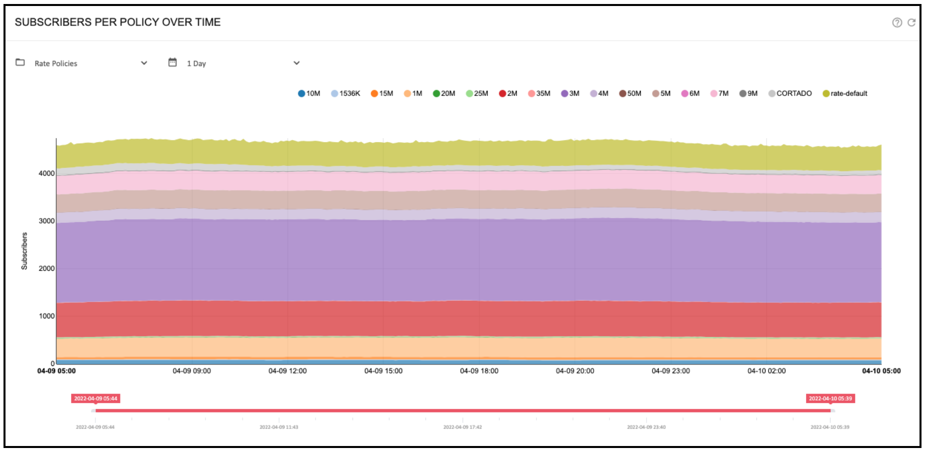

Figure 31: *Subscribers per policy over time page*

Statistics > Subscribers > Per Policy displays the split of the traffic volume into the different rate policies:

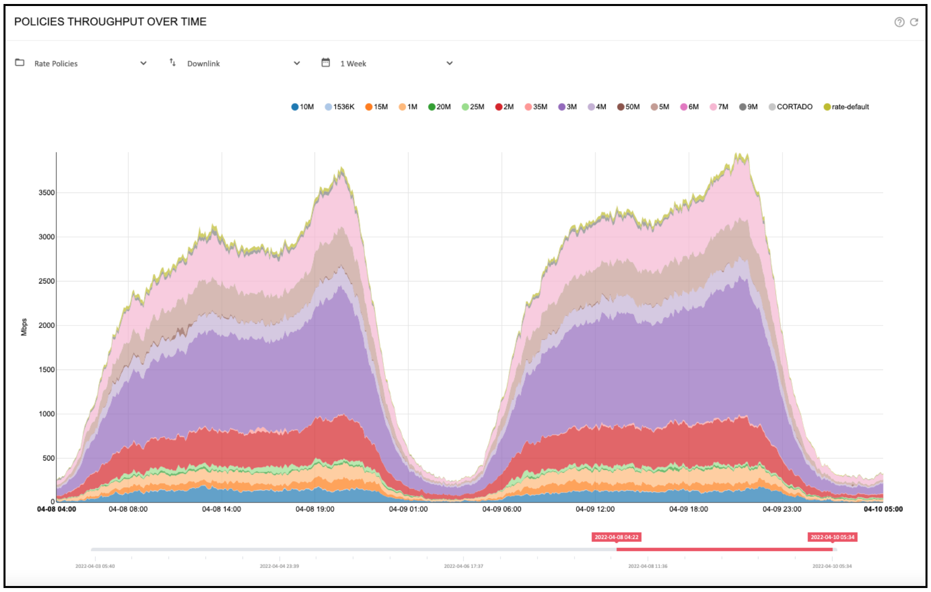

Figure 32: *Policies throughput over time page*

Statistics > **Subscribers** > **Per Policy** displays the split of the traffic volume into the different flow policies:

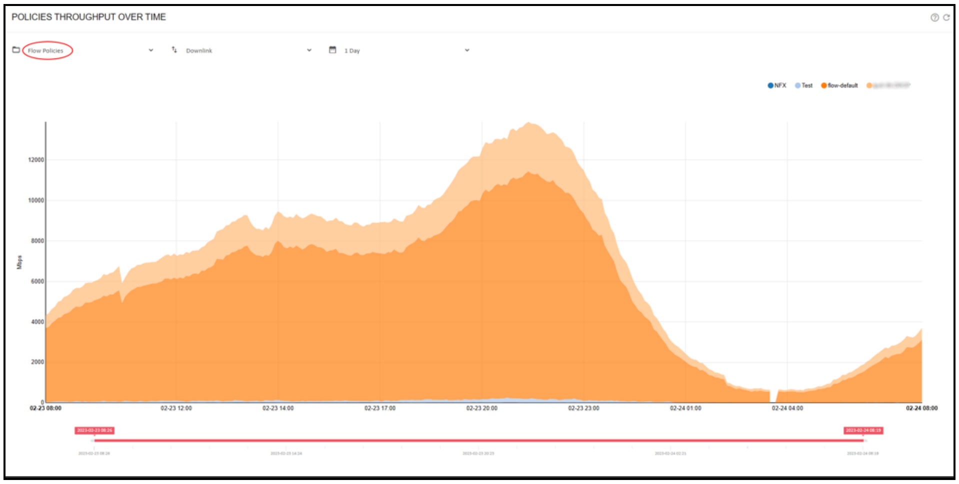

POLICIES THROUGHPUT OVER TIME

Flow Policies · Downlink · 1 Day

Traffic per service

QoE displays overall traffic composition per service under **Statistics > Service Analysis > Total Volume per Service**. There is information about traffic composition under **Statistics > Service Analysis** . Figure 33 shows DPI total statistics.

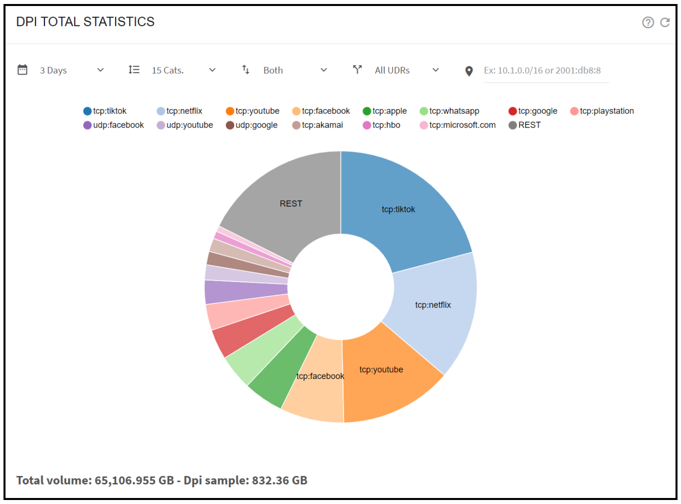

DPI TOTAL STATISTICS

3 Days · 15 Cats. · Both · All UDRs · Ex: 10.1.0.0/16 or 2001:db8:8

Total volume: 65,106.955 GB - Dpi sample: 832.36 GB

Figure 33: *DPI total statistics*

The evolution over time can be obtained from **Statistics > Service Analysis > Hourly Volume per Service**. Figure 34 shows DPI hourly statistics.

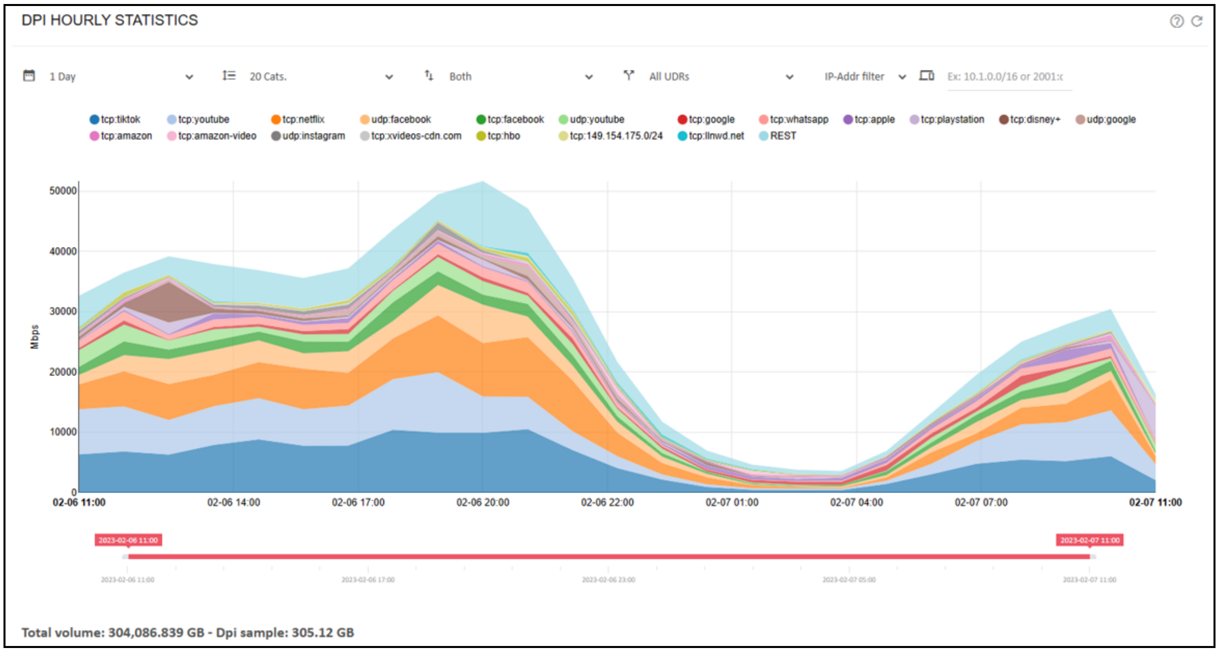

DPI HOURLY STATISTICS

1 Day · 20 Cats. · Both · All UDRs · IP-Addr filter · Ex: 10.1.0.0/16 or 2001:c

Total volume: 304,086.839 GB - Dpi sample: 305.12 GB

Figure 34: *DPI hourly statistics*

By default, all DPI samples are considered by the reporting (*All UDRs*), both the samples generated automatically by the QoE and those generated by monitoring policies. If the monitoring policies generate so many samples that they can cause a bias in the reporting, then select **Only auto UDRs** and exclude those samples. This is case if several subscribers have a monitoring policy that generates UDRs for all their traffic and make them over-represented in the traffic sample mix.

Latency per service

To view the latency per service, navigate to **Statistics > Service Analysis > Latency per Service**. Figure 35 shows Average Internet latency per service page.

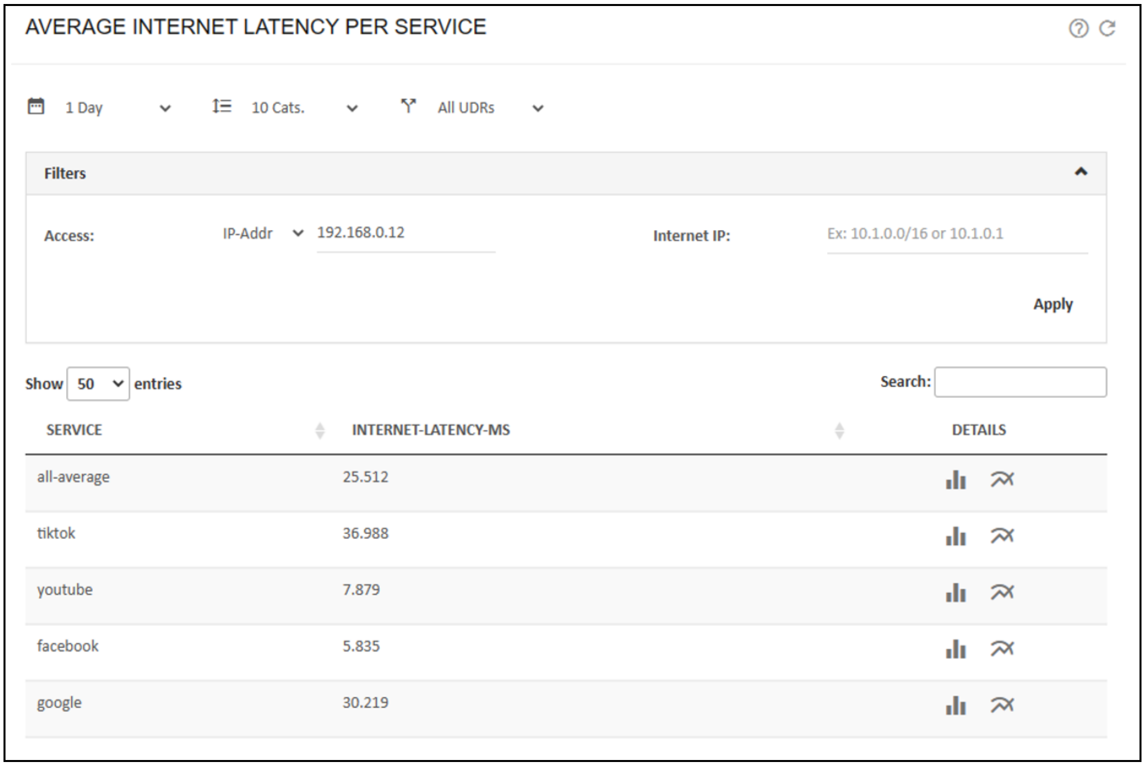
AVERAGE INTERNET LATENCY PER SERVICE

1 Day | 10 Cats. | All UDRs

Filters

Access: IP-Addr 192.168.0.12

Internet IP: Ex: 10.1.0.0/16 or 10.1.0.1

Apply

Show 50 entries

Search:

SERVICE	INTERNET-LATENCY-MS	DETAILS
all-average	25.512	
tiktok	36.988	
youtube	7.879	
facebook	5.835	
google	30.219	

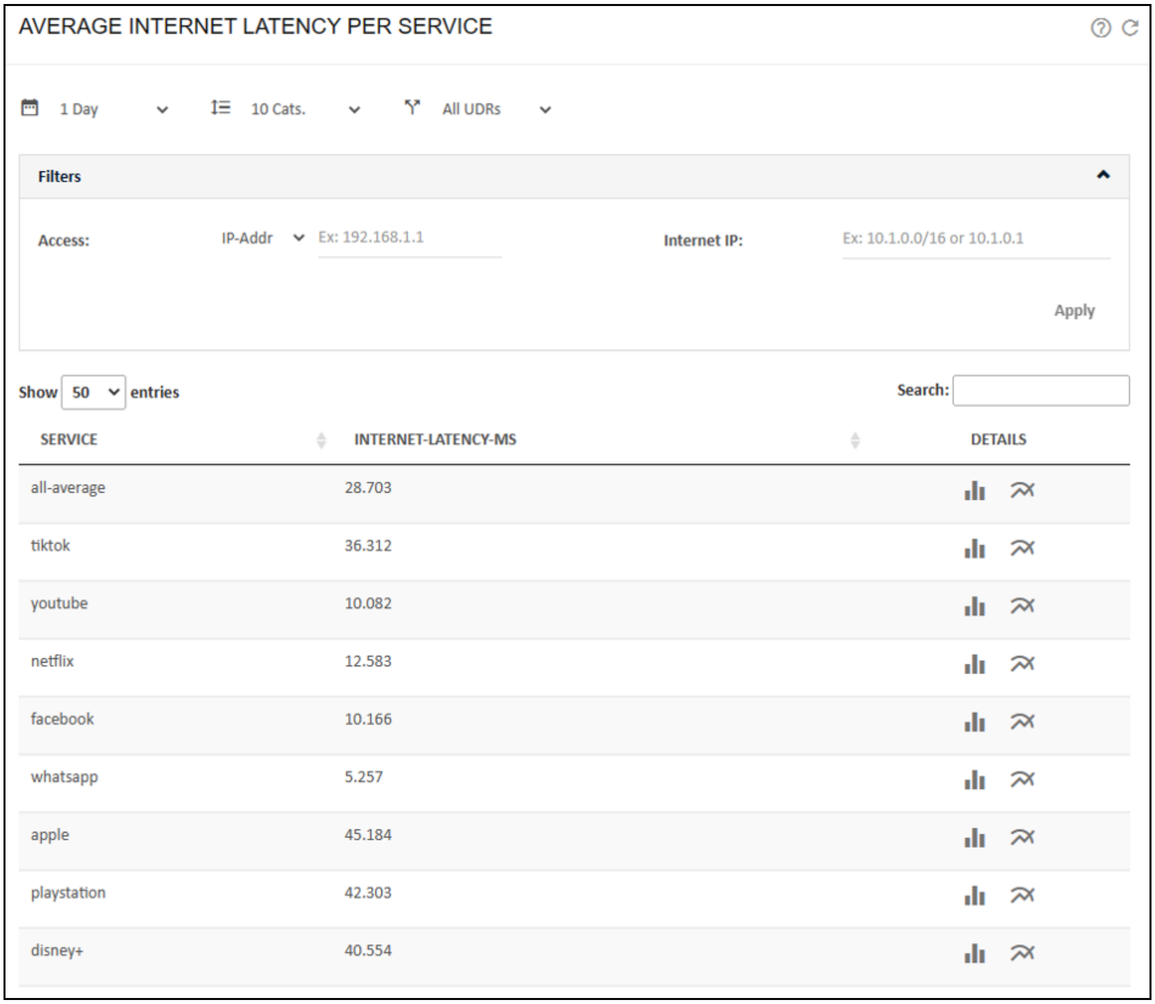

AVERAGE INTERNET LATENCY PER SERVICE

1 Day | 10 Cats. | All UDRs

Filters

Access: IP-Addr Ex: 192.168.1.1 Internet IP: Ex: 10.1.0.0/16 or 10.1.0.1

Apply

Show 50 entries Search:

SERVICE	INTERNET-LATENCY-MS	DETAILS
all-average	28.703	📊 〰
tiktok	36.312	📊 〰
youtube	10.082	📊 〰
netflix	12.583	📊 〰
facebook	10.166	📊 〰
whatsapp	5.257	📊 〰
apple	45.184	📊 〰
playstation	42.303	📊 〰
disney+	40.554	📊 〰

Figure 35: *Average Internet latency per service page*

To view the latency distribution of a particular service, click on the bar chart icon (📊) in the DETAILS column. Figure 36 shows **Distribution of Internet Latencies**.

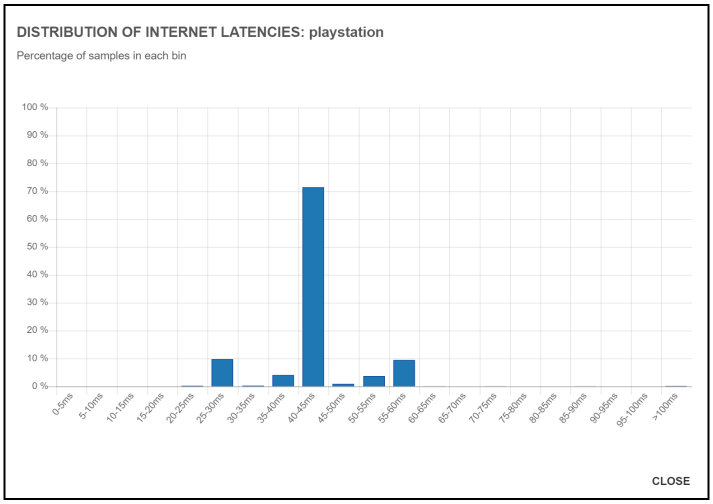

Figure 36: *Distribution of Internet latencies*

The distribution shows the percentage of measurements falling into the different intervals of the range of latency values. In this example, over 70 percent of latencies are between 40 ms and 45 ms.

To see the evolution over time of those Internet latencies, click on the dial icon in the **DETAILS** column. Figure 37 shows **Distribution of Internet Latencies Over Time**.

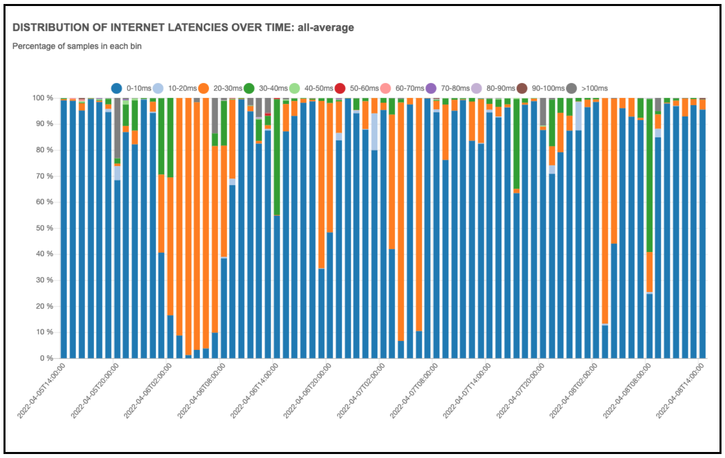

Figure 37: *Distribution of Internet latencies over time*

For each time interval, the chart shows the relative percentages of each latency range. In the example, most latencies are in the 0-10 ms range (dark blue) but during the night hours most latencies fall in the 20-30 ms range (orange).

This latency per service analysis, whether globally or over time, can be further refined by filtering with a subscriber address IP address (or range), or by an Internet-side IP address (or range), which provides the information on latency of different services coming from different providers, and specifically for certain subscribers. Figure 38 shows **Average Internet Latency per Service**.

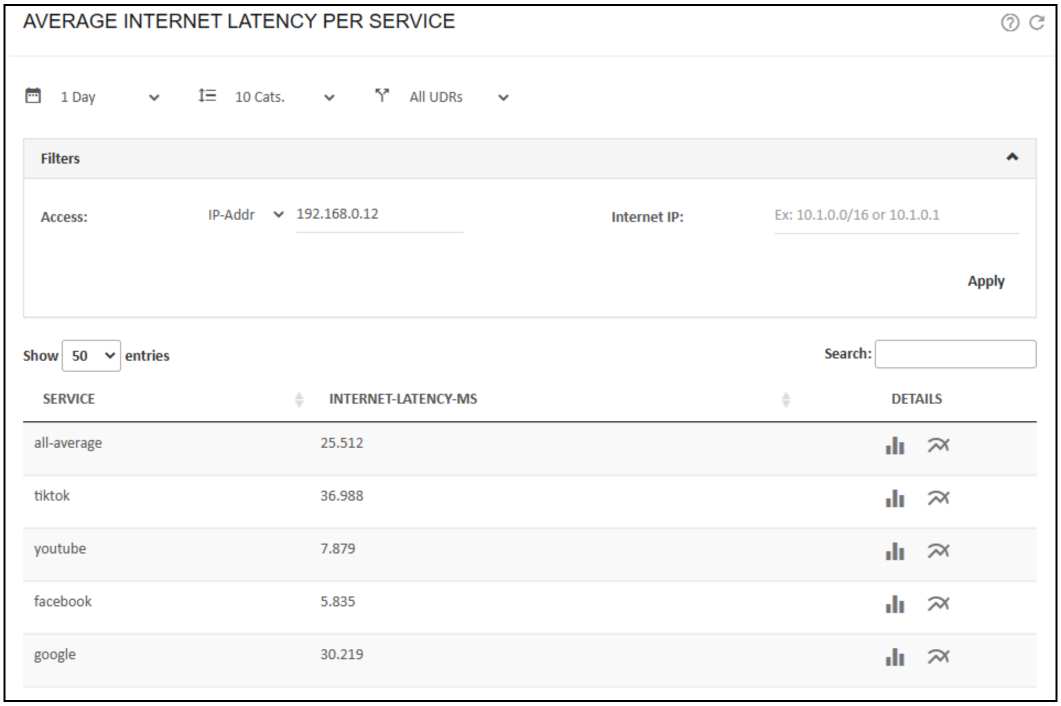

AVERAGE INTERNET LATENCY PER SERVICE

1 Day | 10 Cats. | All UDRs

Filters

Access: IP-Addr 192.168.0.12

Internet IP: Ex: 10.1.0.0/16 or 10.1.0.1

Apply

Show 50 entries

Search:

SERVICE	INTERNET-LATENCY-MS	DETAILS
all-average	25.512	
tiktok	36.988	
youtube	7.879	
facebook	5.835	
google	30.219	

Figure 38: *Average Internet latency per service*

Main subscribers per service

The user can identify the main IP addresses of a given service, both subscriber ID or addresses (access-side IPs) and addresses of content servers over the Internet.

In **Statistics > DPI Service Analysis > Total Volume Per Service**, click on the pie sector of the service whose main IPs you want to display. A histogram of the main IP addresses is shown, with server IP addresses at the top and subscriber addresses at the bottom. Figure 39 shows the percentage of the total service volume of that IP address. The user can also view the top subscriber IDs by selecting the required option at the bottom of the page.

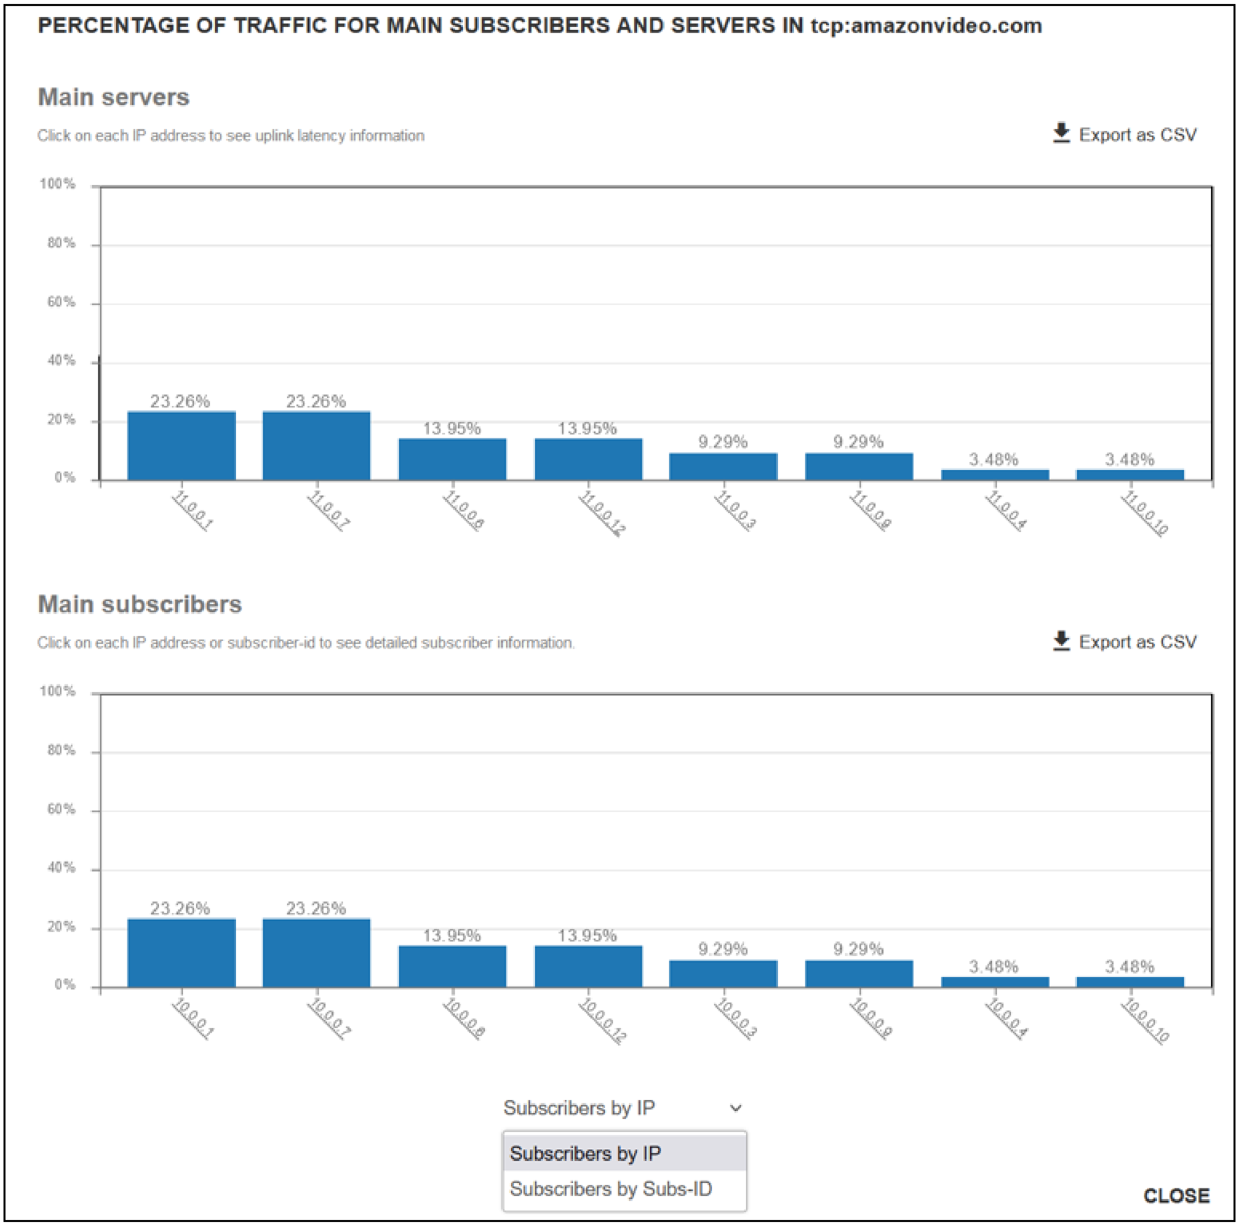

PERCENTAGE OF TRAFFIC FOR MAIN SUBSCRIBERS AND SERVERS IN tcp:amazonvideo.com

Main servers

Click on each IP address to see uplink latency information

Main subscribers

Click on each IP address or subscriber-id to see detailed subscriber information.

Figure 39: *Percentage of traffic*

Addresses are clickable: server IPs will lead to latency distribution for that server IP and subscriber IPs or IDs will go to the subscriber dashboard.

To obtain a CVS file with the server or subscriber volume percentages, click on the **Export** icon.

Overall traffic metrics

Statistics > Throughput > Overview page displays the temporal evolution of total traffic throughput adding both directions and all *wires*.

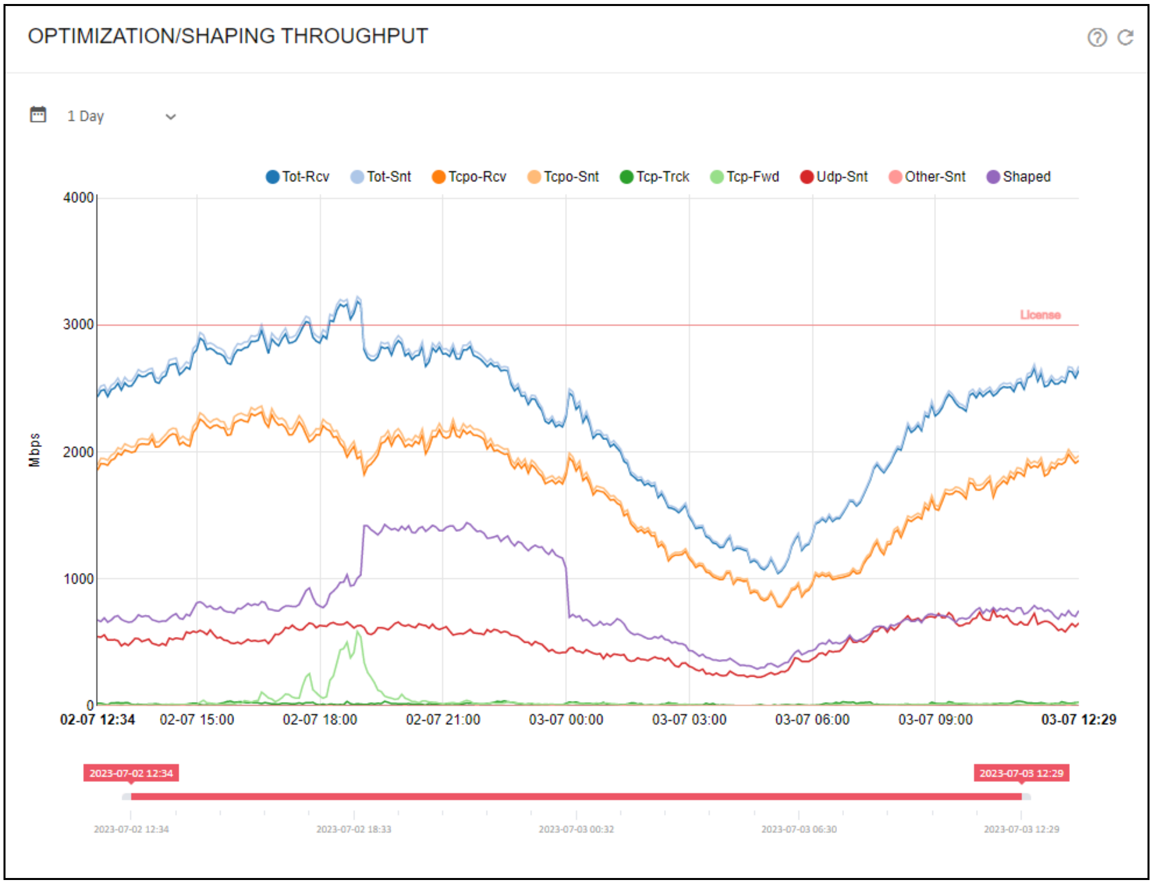

Figure 40: *The throughput over time page*

Statistics > **Throughput** > **Protocols** shows the temporal evolution of traffic throughput. The following are the parts of this traffic throughput:

- Direction: downlink (DN) and uplink (UP).
- IP version: IPV4, IPV6 or No IP.
- L4 Protocol type: TCP, UDP or Other IP protocol.
- Bypass traffic.

Figure 41 shows The Throughput overview page.

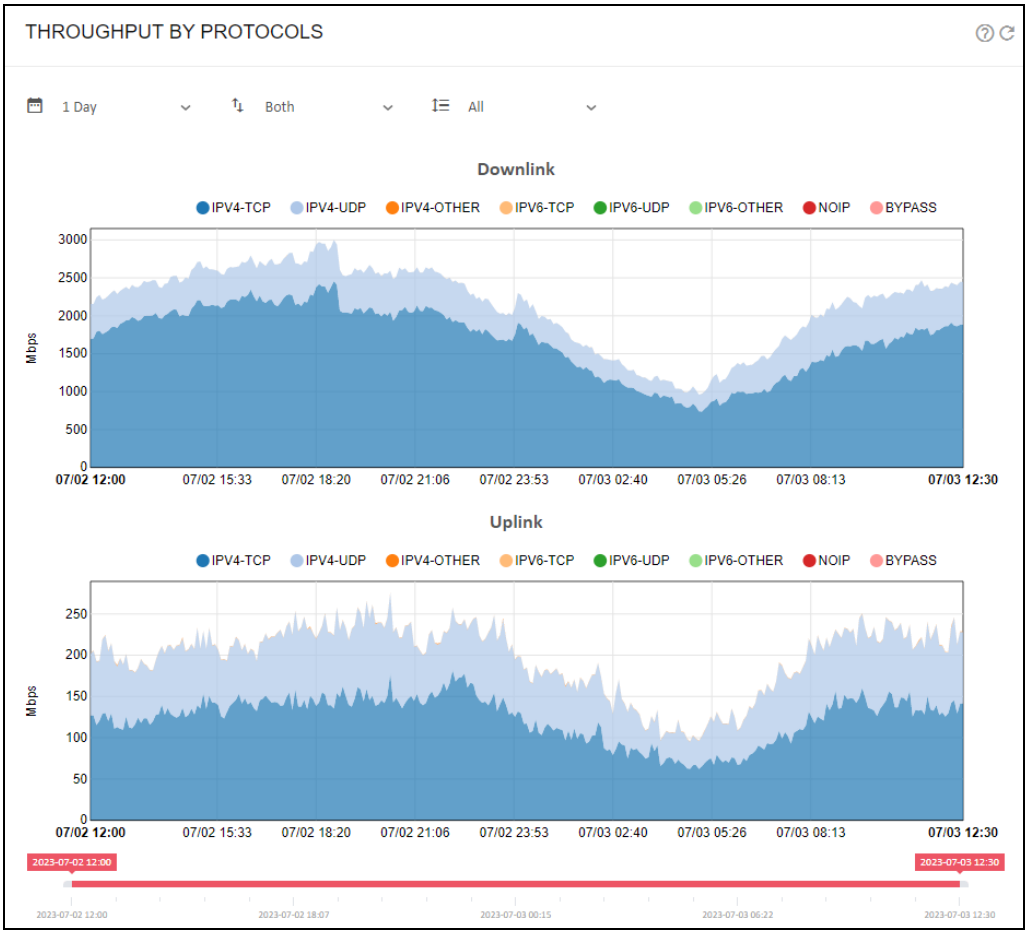

THROUGHPUT BY PROTOCOLS

1 Day

Both

All

Figure 41: *The Throughput overview page*

The evolution over time per network interface is available in **Statistics > Throughput > Interfaces**. Figure 42 shows network interface throughput over time.

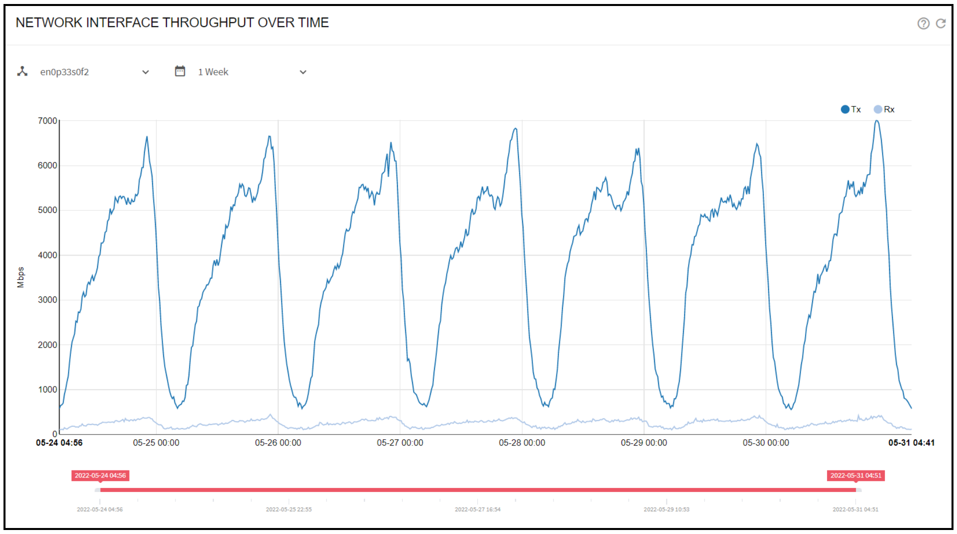

NETWORK INTERFACE THROUGHPUT OVER TIME

en0p33s0f2

1 Week

Figure 42: *Network interface throughput over time*

It is possible to check how much traffic is being processed according to each of the configured policies. For **Subscriber Flows** policies, it can be checked in **Statistics > Throughput > Subscriber Flows Policies** and similarly for **Subscriber Rate Policies** and **Subscriber Monitoring Policies**.

The chart in **Statistics** > **System** > **Latencies** displays the access RTT (RTT-Access). It is the average across all flows of the minimum value per flow. Figure 43 shows latency over time.

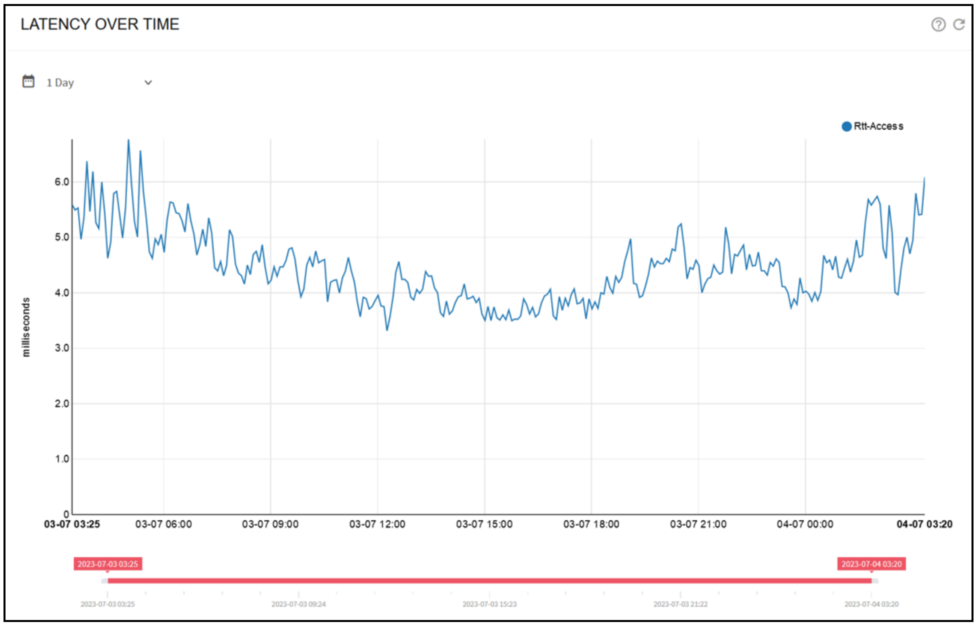

Figure 43: *Latency over time*

Also, **Statistics > System > Retransmissions** displays the average retransmission percentages in downlink and uplink directions. Figure 44 shows average TCP retransmissions over time.

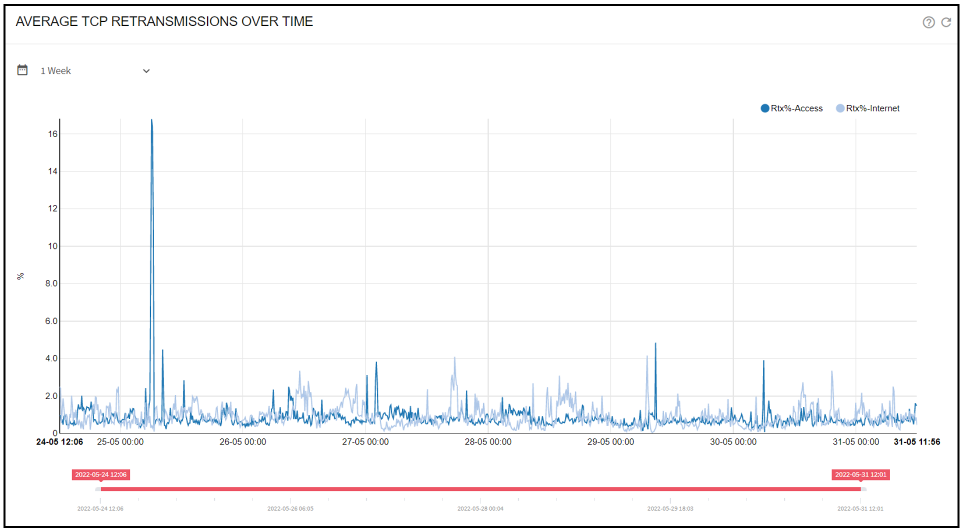

Figure 44: *Average TCP retransmissions over time*

To see the number of flows per policy and per protocol, click on **Statistics > Flow > Per Policy** and **Statistics > Flow > Per Protocol** respectively. Figure 45 shows active flows over time.

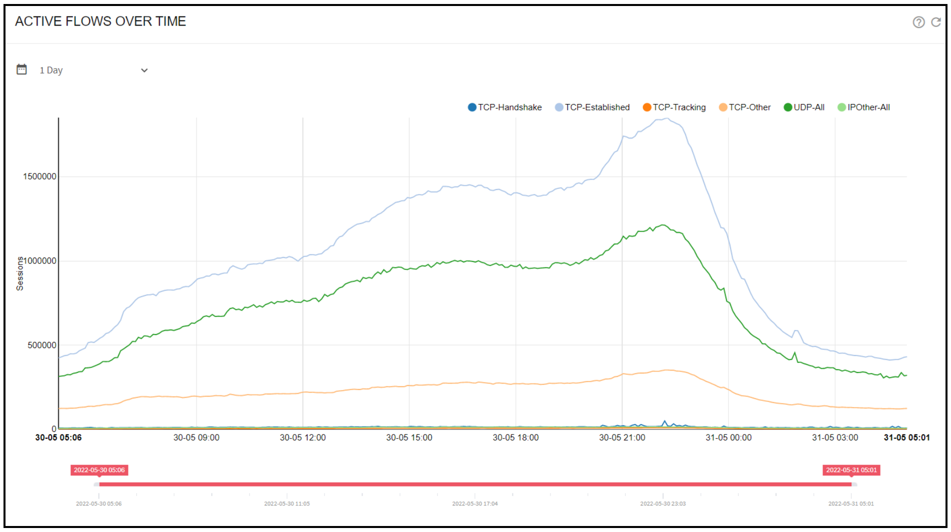

Figure 45: *Active flows over time*

You can also view the instantaneous number of flows per protocol in **Status > Flows > Per Protocol** and per subscriber in **Status > Flows > Per Subscriber.** Figure 46 shows active flows.

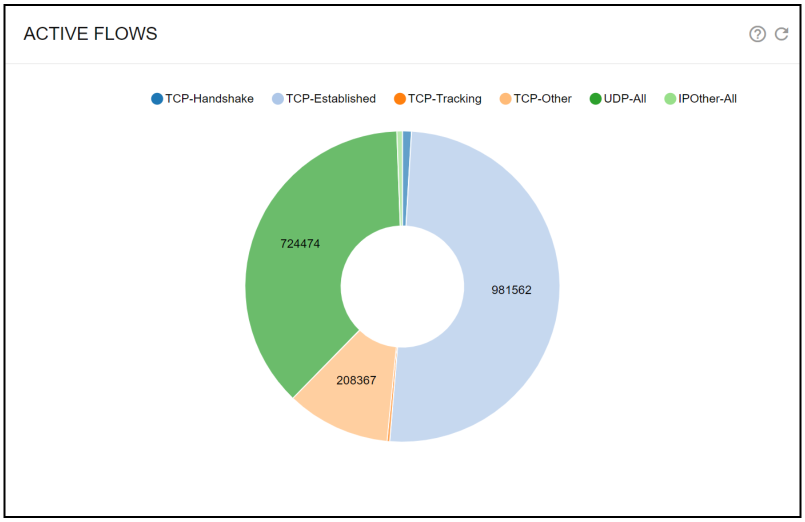

Figure 46: *Active flows*

DoS

The QoE detects Denial of Service (DoS) attacks. To perform this, DoS thresholds must be configured in **Configuration** > **DoS**:

- **Downlink failed handshake rate**. SYNs per second without an answer in the direction towards the subscribers (initialized from the Internet). A typical value is 50.

- **Uplink failed handshake rate**. SYNs per second without an answer initialized by a subscriber. A typical value is 50 Mbps.

- **Minimum rate**. Minimum speed rate that can be considered a volumetric attack. The exact value depends on the network speed, but a typical value is 50 Mbps.

- **Multiplier of subscriber rate policy**. If the subscriber has a known rate policy, a threshold is defined as multiplier * downlink limit. A typical multiplier is 3. For example, a subscriber with a 20 Mbps plan has a DoS threshold of 3 * 20 = 60 Mbps. shows DoS settings page.

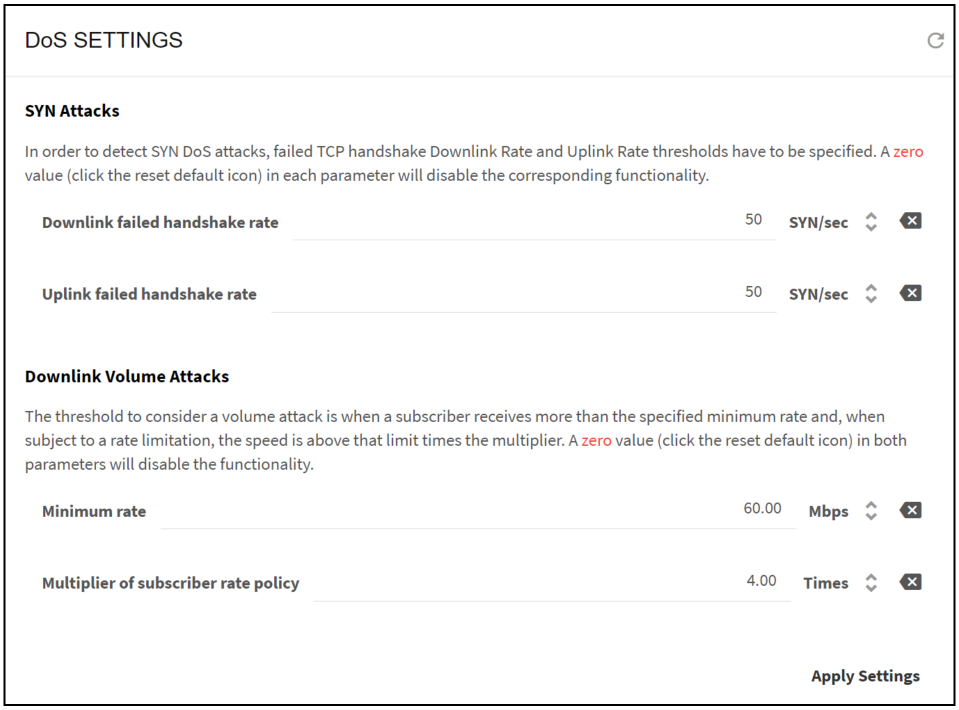

DoS SETTINGS

SYN Attacks

In order to detect SYN DoS attacks, failed TCP handshake Downlink Rate and Uplink Rate thresholds have to be specified. A zero value (click the reset default icon) in each parameter will disable the corresponding functionality.

Downlink failed handshake rate	50	**SYN/sec**
Uplink failed handshake rate	50	**SYN/sec**

Downlink Volume Attacks

The threshold to consider a volume attack is when a subscriber receives more than the specified minimum rate and, when subject to a rate limitation, the speed is above that limit times the multiplier. A zero value (click the reset default icon) in both parameters will disable the functionality.

Minimum rate	60.00	**Mbps**
Multiplier of subscriber rate policy	4.00	**Times**

Figure 47: *DoS settings page*

The DoS events are shown in **Statistics > DoS Attacks**. In **DoS Attacks Over Time**, the DoS attack events are displayed showing its type, its duration and parameters such as the affected subscriber IP and the main IP contributing to the attack. shows DoS attacks over time page.

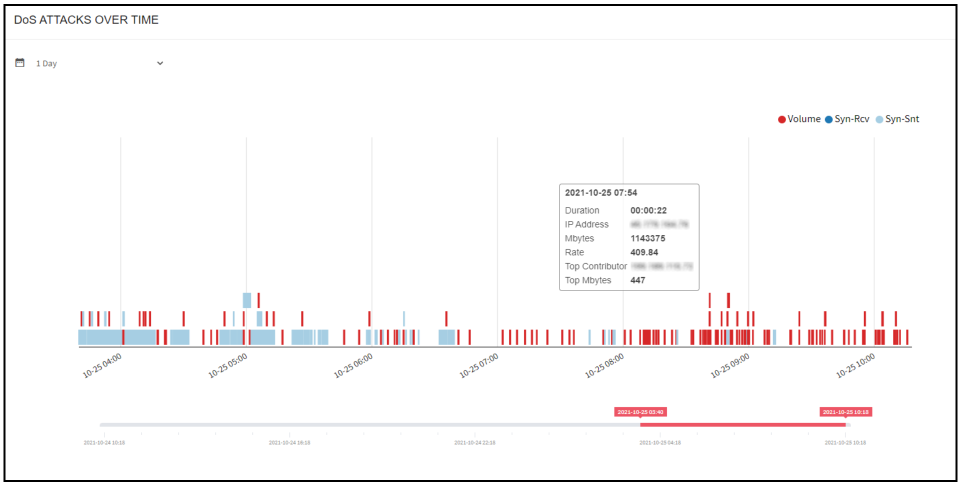

DoS ATTACKS OVER TIME

1 Day

2021-10-25 07:54

Duration	00:00:22
IP Address	[redacted]
Mbytes	1143375
Rate	409.84
Top Contributor	[redacted]
Top Mbytes	447

Figure 48: *DoS attacks over time page*

In **Details of DoS Attacks** all DoS events are listed, with information about the time, event type, IP address affected, the direction of the attack (ingress or egress) and its duration. In **SYN Attacks** can be found attacks of SYN type, with the number of failed SYN and its rate per second. In **Volume Attacks** there is a list of volumetric attacks, with information of the traffic volume and its average rate.

Chapter 8: Introduction to Policies

To manage in a flexible way the product many features, QoE uses policies. Policies define the actions to perform on the traffic (such as traffic optimization, rate limitation, and generation of metrics), along with the action parameters (for example a speed limit).

There are three kinds of policies:

- **Flow policies**, to act on IP flows (for example, a TCP connection or a UDP flow).
- **Rate policies**, associated with subscriber sessions.
- **Monitoring policies**, also associated with subscriber sessions.

A subscriber session is defined as all the traffic from the same IPv4 address on the access side, or, in the case of IPv6, from the same /64 subnet. For more information, see Policy enforcement section.

Every flow is assigned a flow policy. Every subscriber is assigned a rate policy and a monitoring policy. Because a subscriber has many flows, the flows may be assigned to different flow policies.

Through flow policies, you can control the following functionalities:

- TCP Optimization (TCPO).
- Shaping per subscriber: Limit to the combined speeds of all the flows assigned to the flow policy, for that subscriber. For example, if the limit is 12 Mbps, four flows of the same subscriber can have 3 Mbps each.
- Shaping per flow: Speed limit of a flow assigned to the flow policy. For example, a limit of 5 Mbps prevent any flow under that policy to exceed those 5 Mbps.
- Total blocking of the traffic under this policy.
- Block only the incoming connections from Internet of specific traffic types.
- Quota counting: Decides this traffic volume counts when checking the volume quota.
- Quota limitation: Steps to perform if the quota is reached, whether traffic is blocked or slow down to some specified speed.

Through rate policies, you can control the following functionalities:

- Limit the total network speed of a subscriber.
- ACM optimization.

Through monitoring policies, you can control the following functionalities:

- The amount of sampling if collecting DPI information for a subscriber (whether automatic or base on some explicit sampling percentage).

Policies are defined as part of the QoE configuration, along with rules and profiles that decide the policy type to apply depending on the traffic characteristics.

Additionally, rate policies can be managed from an external system, by creating dynamically and assigning to the subscribers. The QoE supports the following APIs to integrate with external systems:

- RADIUS
- QoE REST API
- Integrations with many billing vendors.

The rate policies from an external system always take precedence over those rate policies configured in the QoE, that are used as a fallback. That is for the subscribers without a policy assignment from the external system.

This chapter contains the following sections:

- Policy enforcement
- Check the policy of a subscriber
- Check the subscribers of a policy

Policy enforcement

QoE enforces the policies on subscriber sessions. A subscriber session is all traffic of a distinctive IP address on the access side for one single IPv4 address or one IPv6 subnet. For example, a policy with rate limits apply to the total throughput of that distinctive IP address.

> **Note**
>
> If there is a NAT between the QoE server and the real subscribers, subscribers whose IP addresses are translated to the same IP address is considered as the same subscriber.

The default IPv6 subnet is /64. To change the default IPv6 subnet, navigate to **Administration** > **General Settings** and edit the field **IPv6 prefix for subscribers**.

To evaluate the subscriber rate and monitoring rules to choose the policies to applied, a new subscriber session is identified when the first packet from an access IP address is received.

Check the policy of a subscriber

To check the rate policies applied to a subscriber, navigate to **Status** > **Subscribers** > **Subscriber Attributes**. It lists the subscribers, with the applied policy in **RATE-POLICY** column. Figure 49 shows the Subscriber attributes page.

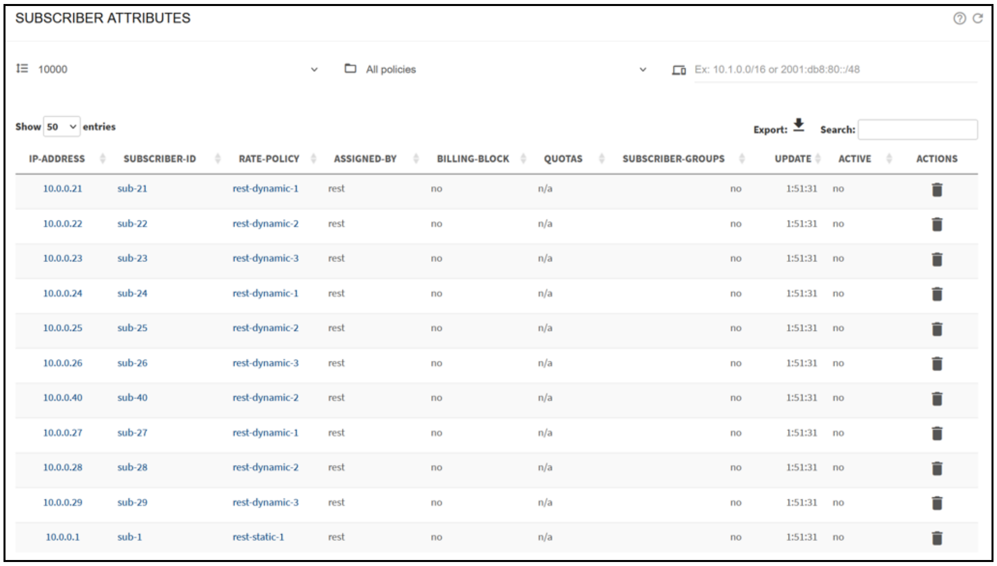

IP-ADDRESS	SUBSCRIBER-ID	RATE-POLICY	ASSIGNED-BY	BILLING-BLOCK	QUOTAS	SUBSCRIBER-GROUPS	UPDATE	ACTIVE	ACTIONS
10.0.0.21	sub-21	rest-dynamic-1	rest	no	n/a	no	1:51:31	no	🗑
10.0.0.22	sub-22	rest-dynamic-2	rest	no	n/a	no	1:51:31	no	🗑
10.0.0.23	sub-23	rest-dynamic-3	rest	no	n/a	no	1:51:31	no	🗑
10.0.0.24	sub-24	rest-dynamic-1	rest	no	n/a	no	1:51:31	no	🗑
10.0.0.25	sub-25	rest-dynamic-2	rest	no	n/a	no	1:51:31	no	🗑
10.0.0.26	sub-26	rest-dynamic-3	rest	no	n/a	no	1:51:31	no	🗑
10.0.0.40	sub-40	rest-dynamic-2	rest	no	n/a	no	1:51:31	no	🗑
10.0.0.27	sub-27	rest-dynamic-1	rest	no	n/a	no	1:51:31	no	🗑
10.0.0.28	sub-28	rest-dynamic-2	rest	no	n/a	no	1:51:31	no	🗑
10.0.0.29	sub-29	rest-dynamic-3	rest	no	n/a	no	1:51:31	no	🗑
10.0.0.1	sub-1	rest-static-1	rest	no	n/a	no	1:51:31	no	🗑

Figure 49: *The Subscriber attributes page*

The **ASSIGNED-BY** column indicates the origin of the policy. That is QoE configured rules, RADIUS, QoE REST API or a billing system.

Click on the subscriber IP address or Subscriber ID to view the Subscriber dashboard (for more information, see Network visibility, Subscriber Dashboad sections).

At the top of the page, there are fields to filter the subscribers by policy, source of the policy assignment or IP address.

To view the active flows of a subscriber, navigate to **Status** > **Flows** > **Details**. The policy applied to each flow in **FLOW-POLICY** column, along with other information is displayed. Figure 50 shows the Flows per subscriber page.

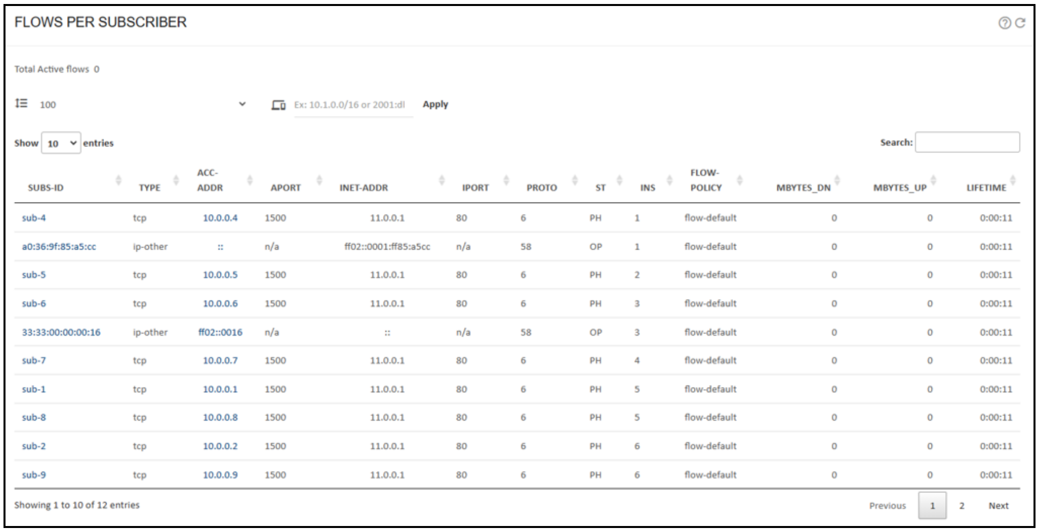

FLOWS PER SUBSCRIBER

Total Active flows 0

SUBS-ID	TYPE	ACC-ADDR	APORT	INET-ADDR	IPORT	PROTO	ST	INS	FLOW-POLICY	MBYTES_DN	MBYTES_UP	LIFETIME
sub-4	tcp	10.0.0.4	1500	11.0.0.1	80	6	PH	1	flow-default	0	0	0:00:11
a0:36:9f:85:a5:cc	ip-other	::	n/a	ff02::0001:ff85:a5cc	n/a	58	OP	1	flow-default	0	0	0:00:11
sub-5	tcp	10.0.0.5	1500	11.0.0.1	80	6	PH	2	flow-default	0	0	0:00:11
sub-6	tcp	10.0.0.6	1500	11.0.0.1	80	6	PH	3	flow-default	0	0	0:00:11
33:33:00:00:00:16	ip-other	ff02::0016	n/a	::	n/a	58	OP	3	flow-default	0	0	0:00:11
sub-7	tcp	10.0.0.7	1500	11.0.0.1	80	6	PH	4	flow-default	0	0	0:00:11
sub-1	tcp	10.0.0.1	1500	11.0.0.1	80	6	PH	5	flow-default	0	0	0:00:11
sub-8	tcp	10.0.0.8	1500	11.0.0.1	80	6	PH	5	flow-default	0	0	0:00:11
sub-2	tcp	10.0.0.2	1500	11.0.0.1	80	6	PH	6	flow-default	0	0	0:00:11
sub-9	tcp	10.0.0.9	1500	11.0.0.1	80	6	PH	6	flow-default	0	0	0:00:11

Showing 1 to 10 of 12 entries

Figure 50: *The Flows per subscriber page*

Check the subscribers of a policy

You can view the number of subscriber IP addresses are under each policy. To view number of subscriber IP addresses, navigate to **Status** > **Policies**.

To check flow policies, navigate to **Status** > **Policies** > **Flow Policies**. Figure 51 shows the Flow policies page.

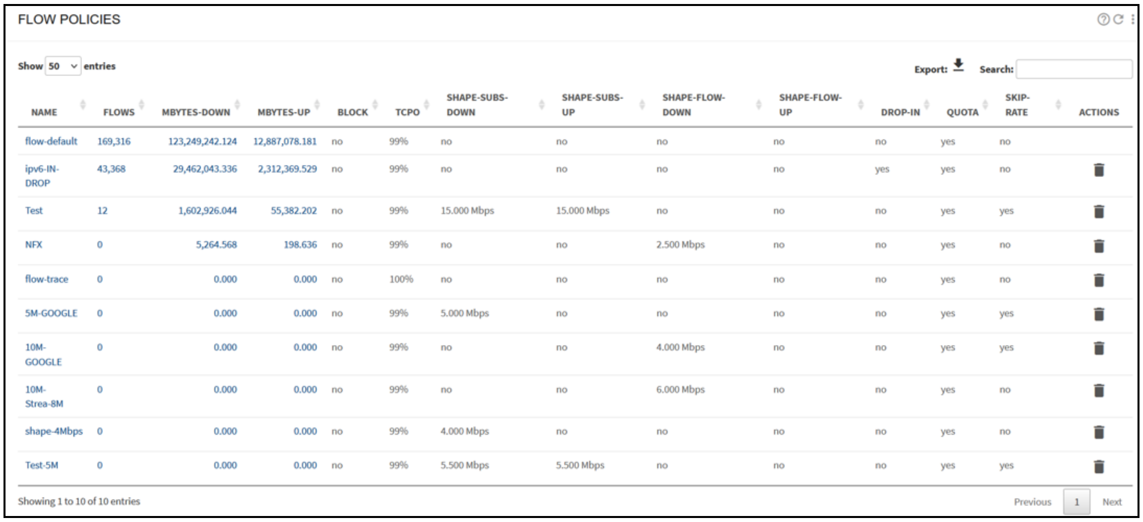

FLOW POLICIES

NAME	FLOWS	MBYTES-DOWN	MBYTES-UP	BLOCK	TCPO	SHAPE-SUBS-DOWN	SHAPE-SUBS-UP	SHAPE-FLOW-DOWN	SHAPE-FLOW-UP	DROP-IN	QUOTA	SKIP-RATE	ACTIONS
flow-default	169,316	123,249,242.124	12,887,078.181	no	99%	no	no	no	no	no	yes	no	
ipv6-IN-DROP	43,368	29,462,043.336	2,312,369.529	no	99%	no	no	no	no	yes	yes	no	🗑
Test	12	1,602,926.044	55,382.202	no	99%	15.000 Mbps	15.000 Mbps	no	no	no	yes	yes	🗑
NFX	0	5,264.568	198.636	no	99%	no	no	2.500 Mbps	no	no	yes	no	🗑
flow-trace	0	0.000	0.000	no	100%	no	no	no	no	no	yes	no	🗑
5M-GOOGLE	0	0.000	0.000	no	99%	5.000 Mbps	no	no	no	no	yes	yes	🗑
10M-GOOGLE	0	0.000	0.000	no	99%	no	no	4.000 Mbps	no	no	yes	yes	🗑
10M-Strea-8M	0	0.000	0.000	no	99%	no	no	6.000 Mbps	no	no	yes	no	🗑
shape-4Mbps	0	0.000	0.000	no	99%	4.000 Mbps	no	no	no	no	yes	no	🗑
Test-5M	0	0.000	0.000	no	99%	5.500 Mbps	5.500 Mbps	no	no	no	yes	yes	🗑

Showing 1 to 10 of 10 entries

Figure 51: *The Flow policies page*

Click on a policy name to view the policy definition and click on the **FLOWS** counter displays a list of flows associated to that policy.

To check rate policies, navigate to **Status** > **Policies** > **Rate Policies**. Figure 52 shows the Rate policies page.

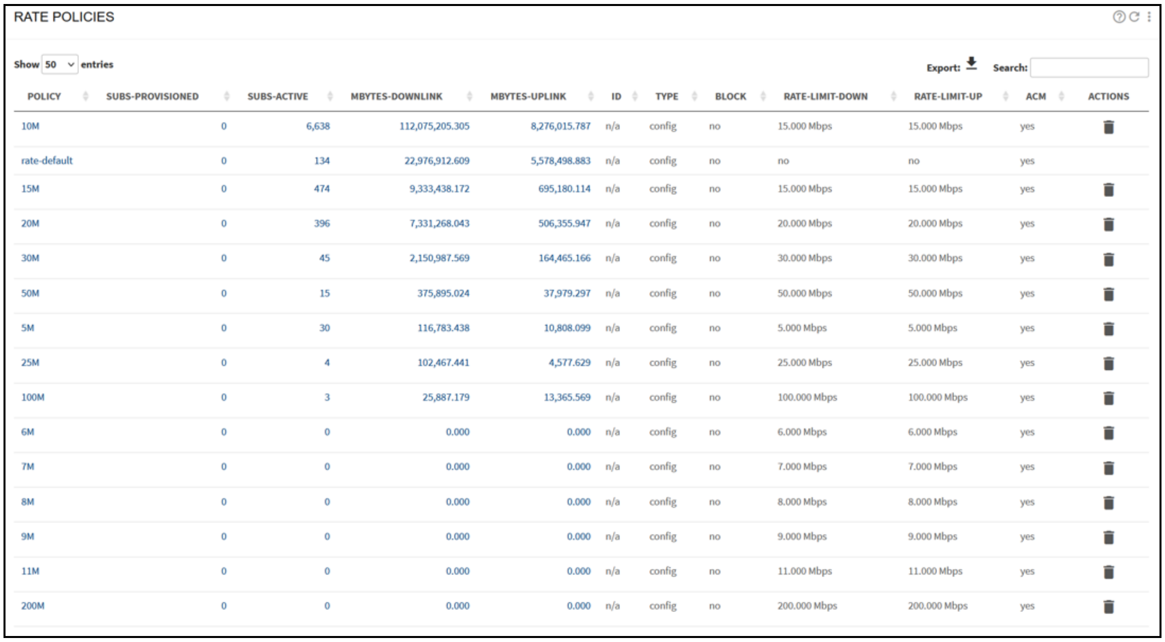

RATE POLICIES

POLICY	SUBS-PROVISIONED	SUBS-ACTIVE	MBYTES-DOWNLINK	MBYTES-UPLINK	ID	TYPE	BLOCK	RATE-LIMIT-DOWN	RATE-LIMIT-UP	ACM	ACTIONS
10M	0	6,638	112,075,205.305	8,276,015.787	n/a	config	no	15.000 Mbps	15.000 Mbps	yes	🗑
rate-default	0	134	22,976,912.609	5,578,498.883	n/a	config	no	no	no	yes	
15M	0	474	9,333,438.172	695,180.114	n/a	config	no	15.000 Mbps	15.000 Mbps	yes	🗑
20M	0	396	7,331,268.043	506,355.947	n/a	config	no	20.000 Mbps	20.000 Mbps	yes	🗑
30M	0	45	2,150,987.569	164,465.166	n/a	config	no	30.000 Mbps	30.000 Mbps	yes	🗑
50M	0	15	375,895.024	37,979.297	n/a	config	no	50.000 Mbps	50.000 Mbps	yes	🗑
5M	0	30	116,783.438	10,808.099	n/a	config	no	5.000 Mbps	5.000 Mbps	yes	🗑
25M	0	4	102,467.441	4,577.629	n/a	config	no	25.000 Mbps	25.000 Mbps	yes	🗑
100M	0	3	25,887.179	13,365.569	n/a	config	no	100.000 Mbps	100.000 Mbps	yes	🗑
6M	0	0	0.000	0.000	n/a	config	no	6.000 Mbps	6.000 Mbps	yes	🗑
7M	0	0	0.000	0.000	n/a	config	no	7.000 Mbps	7.000 Mbps	yes	🗑
8M	0	0	0.000	0.000	n/a	config	no	8.000 Mbps	8.000 Mbps	yes	🗑
9M	0	0	0.000	0.000	n/a	config	no	9.000 Mbps	9.000 Mbps	yes	🗑
11M	0	0	0.000	0.000	n/a	config	no	11.000 Mbps	11.000 Mbps	yes	🗑
200M	0	0	0.000	0.000	n/a	config	no	200.000 Mbps	200.000 Mbps	yes	🗑

Figure 52: *The Rate policies page*

Click on a policy name to view the policy definition and click on the **SUBS-ACTIVE** counter to view a list of subscribers associated to that policy.

To check rate policies, navigate to **Status** > **Policies** > **Monitoring Policies**. Figure 53 shows the Monitoring policies page.

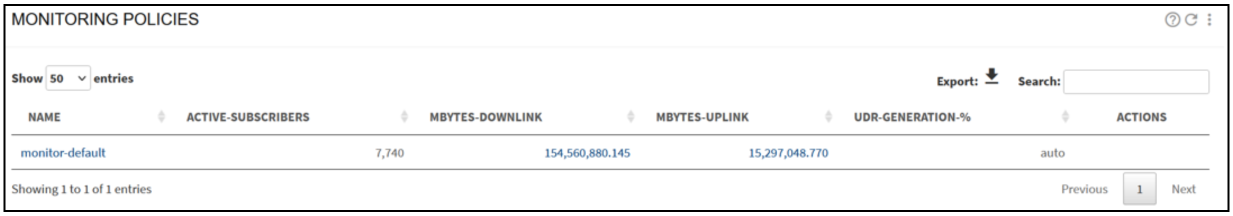

MONITORING POLICIES

NAME	ACTIVE-SUBSCRIBERS	MBYTES-DOWNLINK	MBYTES-UPLINK	UDR-GENERATION-%	ACTIONS
monitor-default	7,740	154,560,880.145	15,297,048.770	auto	

Showing 1 to 1 of 1 entries

Figure 53: *The Monitoring policies page*

Click on a policy name to view the policy definition and click on the **ACTIVE-SUBSCRIBERS** counter to view a list of subscribers associated to that policy.

> Note
>
> - **Status** > **Policies** > **Rate Policies** combines the information in **Status** > **Radius/REST/Billing** > **Policies and Status** > **Policies** from the previous versions.
> - **Status** > **Policies** > **Flow Policies and Status** > **Policies** > **Monitoring Policies** contain the information in **Status** > **Policies** from the previous versions.
>
> .

Chapter 9: Configured Policies

Locally configured policies are selected using rules which combine profiles with the policies:

- **Profiles** classify the traffic according to some conditions (for example, an access profile identifies all the traffic from subscribers within a set of IP address ranges).
- **Rules** relate policies and profiles (for example, a rule may specify that some specific access profile is limited by a rate policy. That is the subscribers whose IP addresses are in some subnet has a specific rate limit).

This chapter contains the following sections:

Profiles

Profiles are configured from the menu option **Configuration > Profiles**. Profiles classify the traffic. Determines the policy applied to a subscriber or a flow when used by policy rules. There are different profile types, according to the properties being used for traffic classification. The latest version supports the following profile types:

- **Interface Profile**: Identifies the flows or subscribers whose first data packet comes in through a network interface in the profile.
- **VLAN Profile:** Identifies the flows or subscribers whose first data packet uses a VLAN tag within the set of VLAN tags (or the absence of any tag) of the profile.
- **Policy Rate Profile**: Identifies the name of the subscriber rate policy. The profile may contain patterns with wildcards. For example, a policy rate profile containing pattern **premium-*** matches subscriber traffic with rate policies named **premium-gold** and **premium-platinum**.
- **Internet Profile**: Identifies the flows involving an IP address on the Internet side, contained in the set of IP address ranges of the profile. Optionally, Internet-side ports can also be specified (for example, port 80).

- **Access Profile**: Identifies the flows or subscribers involving an IP address on the Access side, contained in the set of IP address ranges of the profile. Optionally, access-side ports can also be specified.

- **Subscriber group Profile**: Identifies the name of the subscriber group. The profile may contain patterns with wildcards. For example, a subscriber group name containing pattern wireless* matches subscriber traffic with rate policies named wireless-north and wireless-south.

- **Subscriber ID Profile**: Identifies the subscriber ID. The profile may contain patterns with wildcards. For example, a subscriber ID pattern 100 will match subscriber traffic with IDs 100123 and 100435.

- **Time Profile**: defines time ranges. Optionally, ranges can be restricted to only some days of the week.

- **Throughput Profile**: identifies all the flows which were created while the total downlink traffic through the QoE was above the threshold specified by the profile.

- **DPI (Deep Packet Inspection) Profile**: identifies the flows that have a DPI domain that matches one of the domain patterns (signatures) part of the profile. There are a set of pre-defined DPI signatures, which include the signatures of popular applications (like the most important video-streaming apps or the most common software updates). To add custom signatures, see end of this section.

Interface profile

An interface profile contains a list of network interfaces part of a data wire. It is true, if the first packet is received by one of the interfaces of the profile. A network interface can only be part of one profile at the same time.

In Figure 54, a profile is defined for one of the two wires of the QoE server and another profile for a second wire.

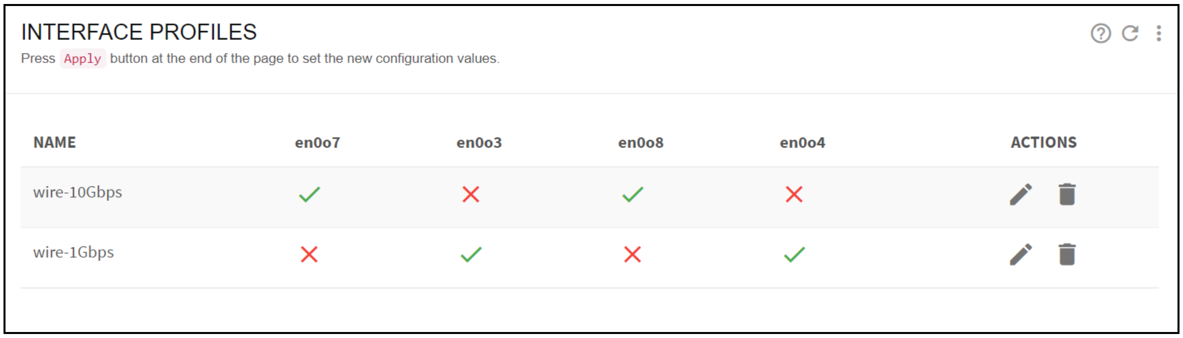

INTERFACE PROFILES

Press `Apply` button at the end of the page to set the new configuration values.

NAME	en0o7	en0o3	en0o8	en0o4	ACTIONS
wire-10Gbps	✓	✗	✓	✗	✎ 🗑
wire-1Gbps	✗	✓	✗	✓	✎ 🗑

Figure 54: *The interface profiles page*

VLAN profile

A VLAN profile is a list of VLAN tags. The profile is true, if the traffic has any of the VLAN tags defined by the profile. Figure 55 shows defines two profiles for two network areas and a third profile for traffic without VLAN tag.

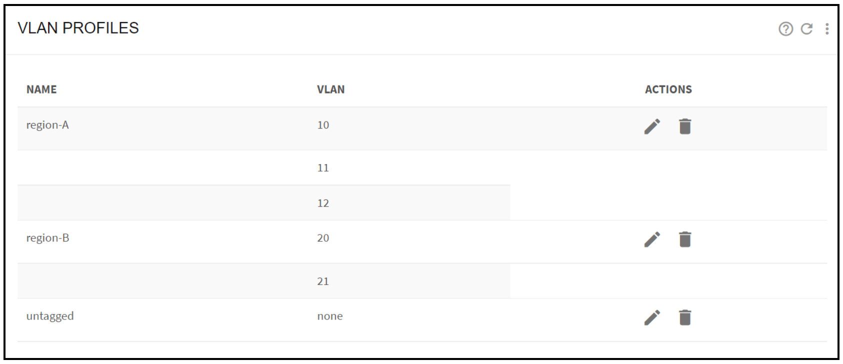

VLAN PROFILES

NAME	VLAN	ACTIONS
region-A	10	✎ 🗑
	11	
	12	
region-B	20	✎ 🗑
	21	
untagged	none	✎ 🗑

Figure 55: *The VLAN profiles page*

Policy rate profiles

This profile is used to select the flow policies based on the rate policy of the subscriber. It is a list of subscriber rate policy names, or patterns with wildcards. Figure 56 shows the policy rate profiles page.

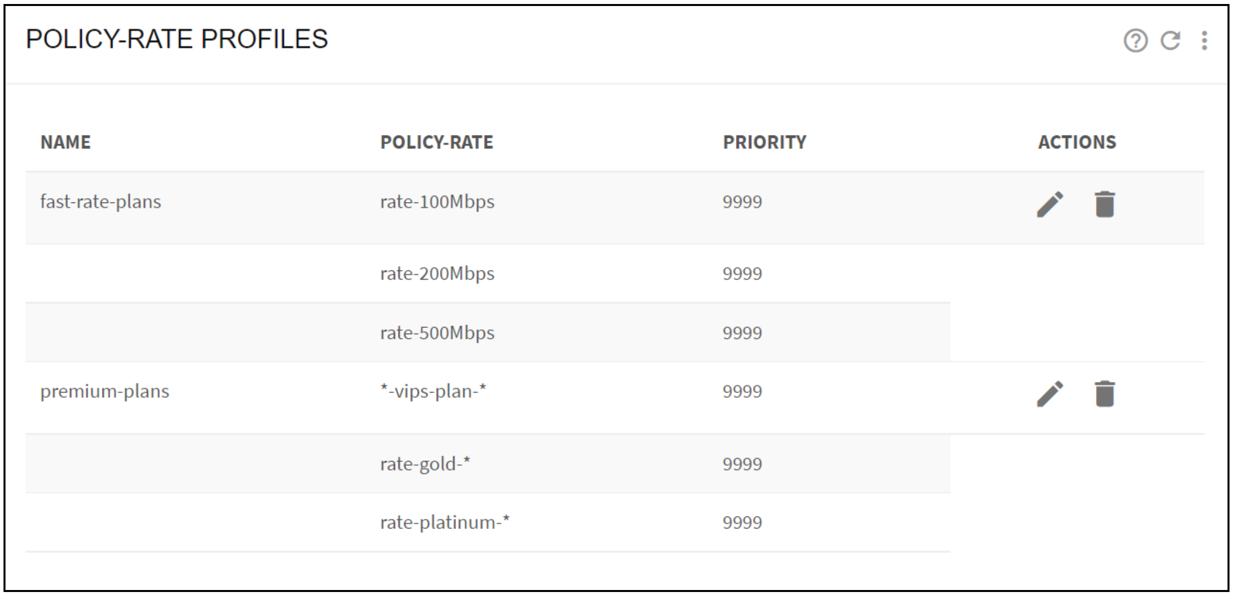

POLICY-RATE PROFILES

NAME	POLICY-RATE	PRIORITY	ACTIONS
fast-rate-plans	rate-100Mbps	9999	✎ 🗑
	rate-200Mbps	9999	
	rate-500Mbps	9999	
premium-plans	*-vips-plan-*	9999	✎ 🗑
	rate-gold-*	9999	
	rate-platinum-*	9999	

Figure 56: *The policy rate profiles page*

Internet and access profiles

Internet and access profiles are a list of IP addresses. IP address ranges are also possible. An Internet profile defines a list of addresses on the Internet side (content sever addresses) and an access profile defines addresses on the access side (subscribers addresses). Figure 57 shows the access profiles page.

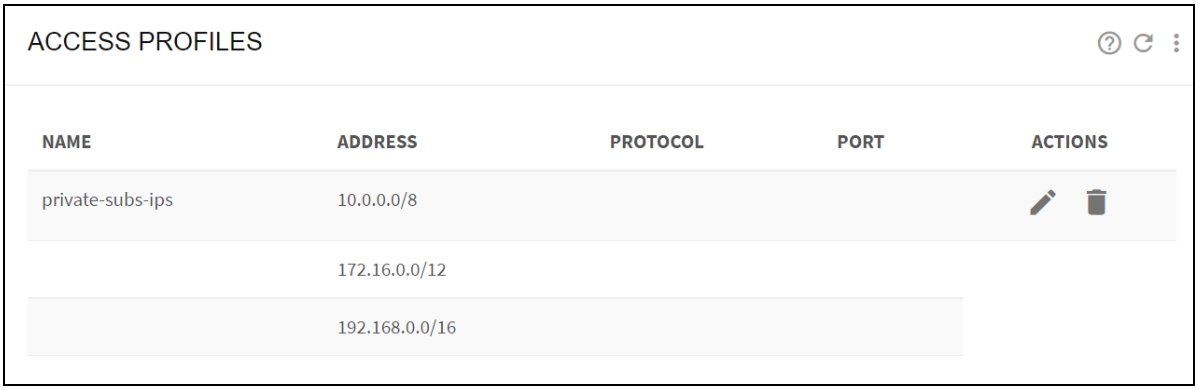

ACCESS PROFILES

NAME	ADDRESS	PROTOCOL	PORT	ACTIONS
private-subs-ips	10.0.0.0/8			✎ 🗑
	172.16.0.0/12			
	192.168.0.0/16			

Figure 57: *The access profiles page*

Optionally, it is possible to define a TCP or UDP port. To match any IP address and a given port, use 0.0.0.0/0 or ::/0 as the address range. Figure 58 shows the Internet profiles page.

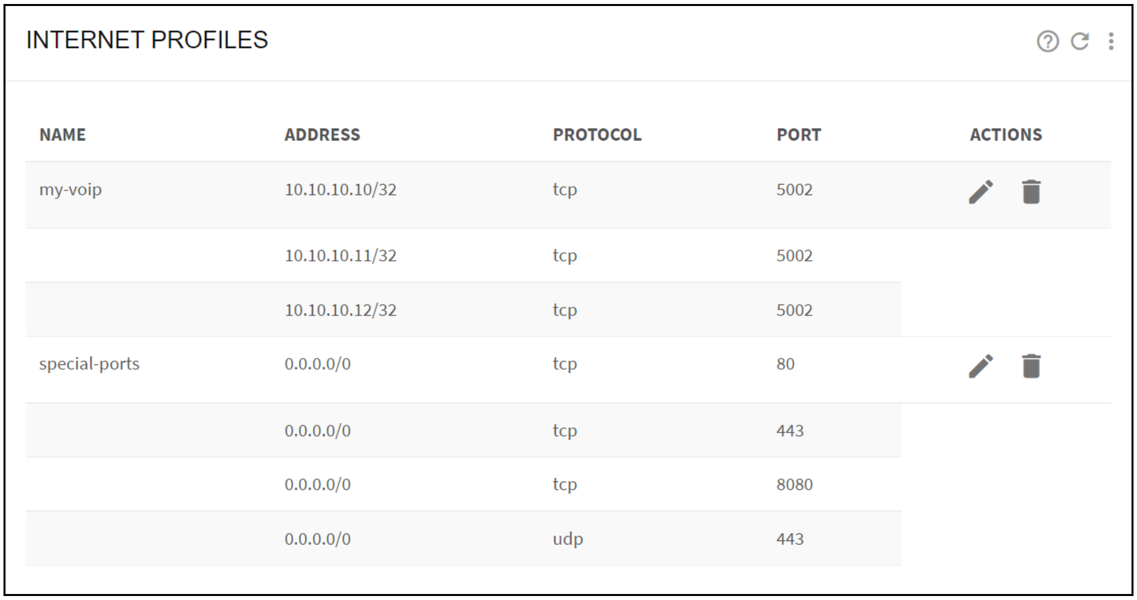

INTERNET PROFILES

NAME	ADDRESS	PROTOCOL	PORT	ACTIONS
my-voip	10.10.10.10/32	tcp	5002	✎ 🗑
	10.10.10.11/32	tcp	5002	
	10.10.10.12/32	tcp	5002	
special-ports	0.0.0.0/0	tcp	80	✎ 🗑
	0.0.0.0/0	tcp	443	
	0.0.0.0/0	tcp	8080	
	0.0.0.0/0	udp	443	

Figure 58: *The Internet profiles page*

It is possible to load an IP address list from a text file. The text file format is one IP address or address range per line. To load a file, edit the profile and select, in the upper right menu, the option **Replace Using File...** (the profile mirrors the file address list) or the option **Merge Using File...** (the content of the file is added to the profile existing addresses).

Subscriber group profile

This profile is used to select Flow, Rate or Monitor Policies based on the group of the subscriber. It is a list of Subscriber Group names, or patterns with wildcards. Figure 59 shows the Subscriber group profiles page.

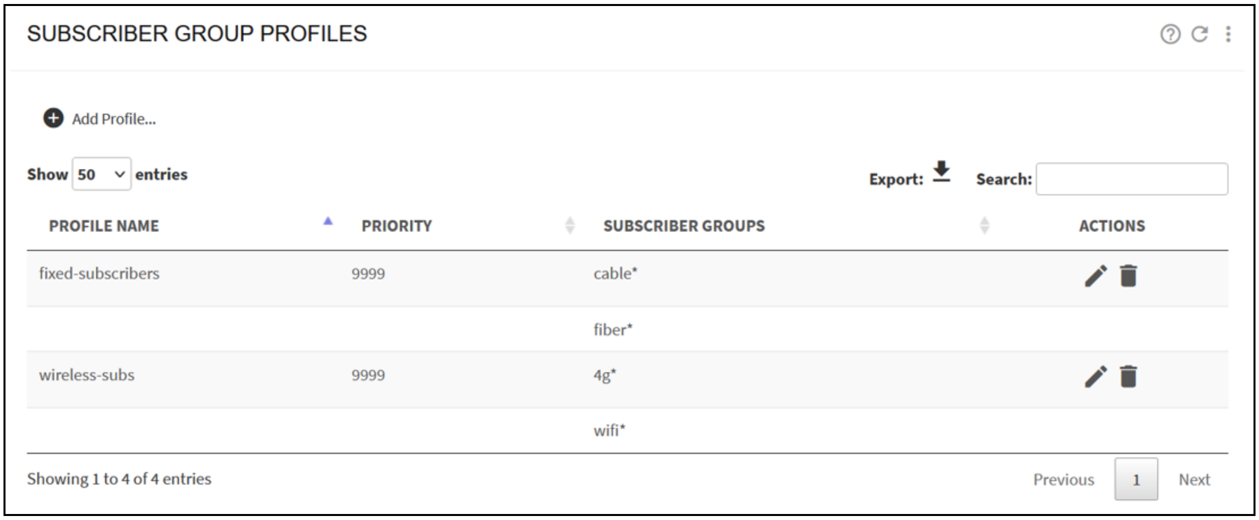

SUBSCRIBER GROUP PROFILES

PROFILE NAME	PRIORITY	SUBSCRIBER GROUPS	ACTIONS
fixed-subscribers	9999	cable*	
		fiber*	
wireless-subs	9999	4g*	
		wifi*	

Showing 1 to 4 of 4 entries

Figure 59: *The Subscriber group profiles page*

Subscriber ID profile

This profile is used to select Flow, Rate or Monitor Policies based on the ID of the subscriber. It is a list of Subscriber IDs, or patterns with wildcards. Figure 60 shows the Subscriber ID profiles page.

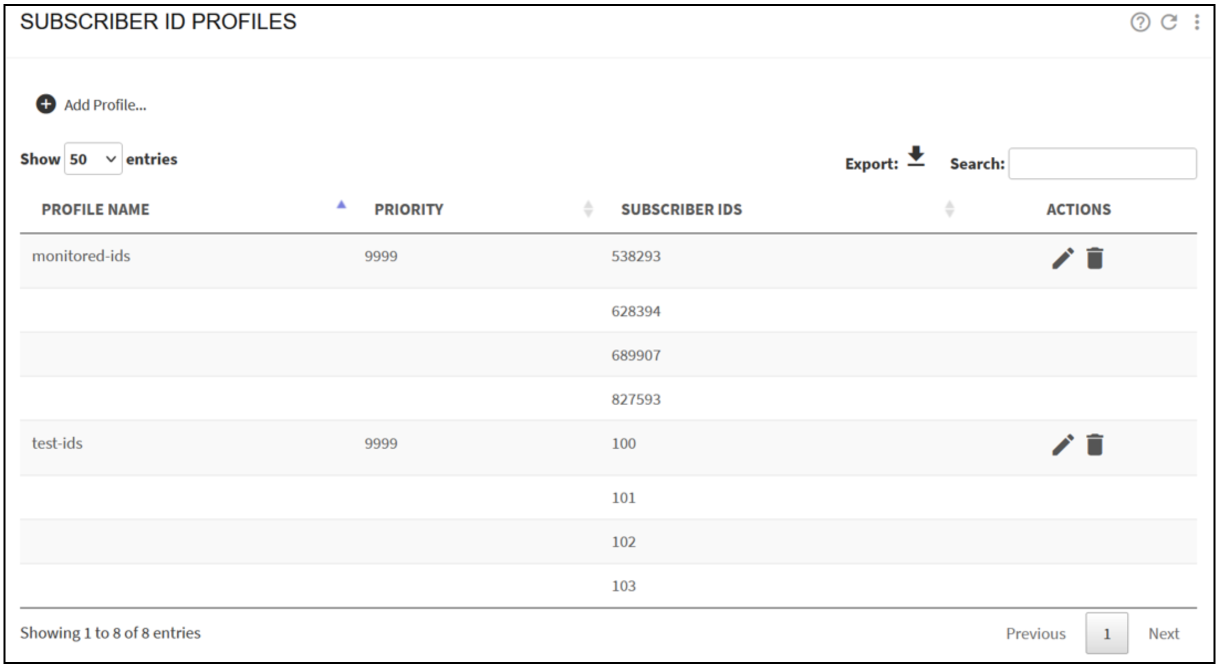

SUBSCRIBER ID PROFILES

PROFILE NAME	PRIORITY	SUBSCRIBER IDS	ACTIONS
monitored-ids	9999	538293	
		628394	
		689907	
		827593	
test-ids	9999	100	
		101	
		102	
		103	

Showing 1 to 8 of 8 entries

Figure 60: *The Subscriber ID profiles page*

Time profiles

Activates the rule during a period. A time profile is a list of time ranges, and it is true if any of the ranges is true. Ranges within the same profile cannot overlap. Figure 61 shows the time profiles page.

A range can apply to all days of the week or just to selected days.

The following example shows:

- A rush hour profile at the end of any day (note how we define the 20:00 – 1:00 interval as two separate periods).
- A weekend profile during Saturday and Sundays.
- A time profile for working hours, with two ranges, valid only from Monday to Friday.

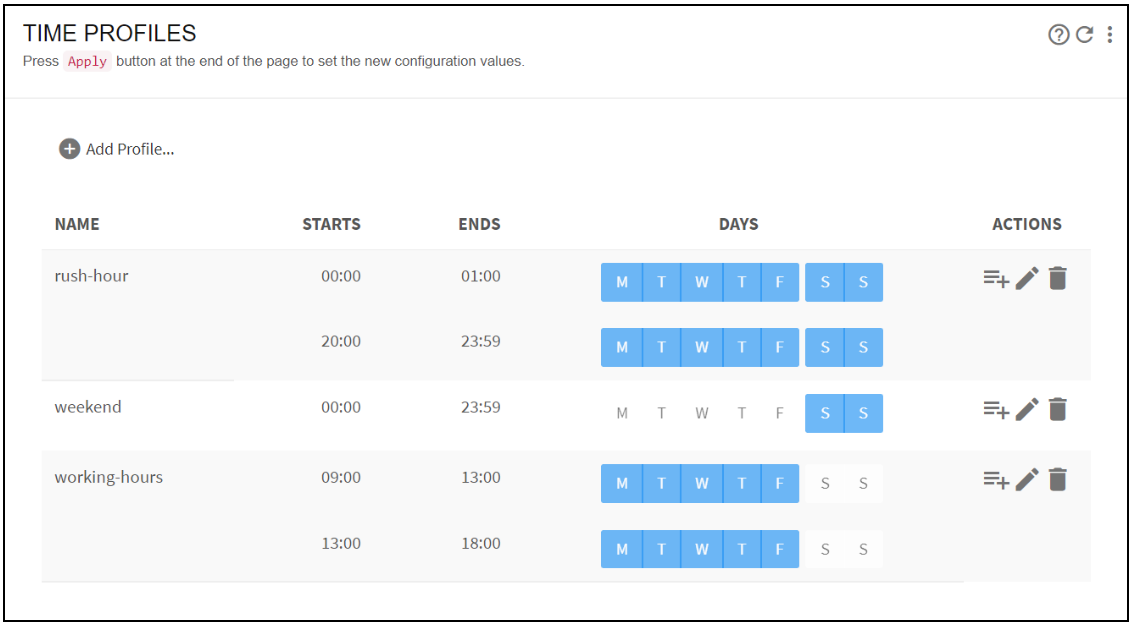

TIME PROFILES

Press `Apply` button at the end of the page to set the new configuration values.

NAME	STARTS	ENDS	DAYS	ACTIONS
rush-hour	00:00	01:00	**M** **T** **W** **T** **F** **S** **S**	≡+ ✎ 🗑
	20:00	23:59	**M** **T** **W** **T** **F** **S** **S**	
weekend	00:00	23:59	M T W T F **S** **S**	≡+ ✎ 🗑
working-hours	09:00	13:00	**M** **T** **W** **T** **F** S S	≡+ ✎ 🗑
	13:00	18:00	**M** **T** **W** **T** **F** S S	

Figure 61: *The time profiles page*

Throughput profiles

A throughput profile defines a threshold of the total downlink traffic through the QoE. It is true when the throughput is exceeded. The following example defines a threshold of 9 Gbps. Figure 62 shows the throughput profiles page.

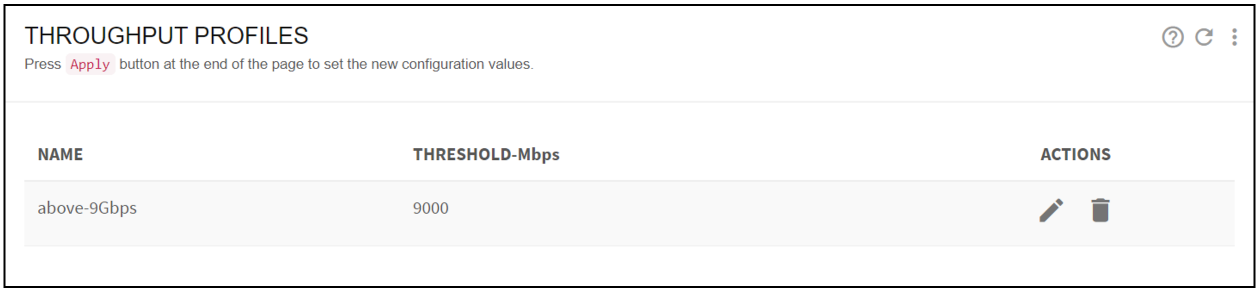

THROUGHPUT PROFILES

Press `Apply` button at the end of the page to set the new configuration values.

NAME	THRESHOLD-Mbps	ACTIONS
above-9Gbps	9000	✎ 🗑

Figure 62: *The throughput profiles page*

DPI profiles

A DPI profile is a collection of signatures to identify in the information obtained through deep-packet-inspection. The signatures can have different types depending on the DPI information:

- HTTP-Host: the hostname in HTTP traffic.
- HTTPS-SNI: the Service-Name-Indication in HTTPS traffic.
- QUIC-SNI: the Service-Name-Indication in QUIC traffic.

- QUIC-MVFST: presence of MVFST traffic (a QUIC variant).

- P2P-FILESHARING: presence of P2P traffic (BitTorrent).

- SPEEDTEST-OOKLA: presence of Speedtest traffic.

Signatures of the type HTTP-Host, HTTPS-SNI and QUIC-SNI must have a pattern indicating the expected content of the DPI information. Patterns can contain up to two **wildchars. Examples: “prefix*”, “*suffix”, “*substring*”.**

A set of pre-defined signatures can be loaded to the DPI profile selecting in the menu **Add Predefined Signatures...**

The following are some of the predefine sets:

- **Video streaming:** YouTube, Netflix, Facebook, Instagram, Amazon Video, HBO, Hulu, Apple TV, Disney+, Twitch, Tiktok, Peacock TV, Pluto TV, Roku, Filmin, DAZN, Magis, Perseo

- **Software downloads:** MS windows, Mac OS and Android updates, PlayStation, Xbox and Steam downloads

- **Speed tests:** Ookla, fast.com, cloudflare, waveform, m-lab, nperf.

A DPI profile can be filled with custom signatures using the option **Add Custom Signature....**

Custom signatures can be loaded from a file using the option **Add Signature File...** A signature file is a text file with one line per signature, with the following format:

```
<pattern> <pattern-type>
```

where:

- **<pattern>** is a domain with wildcards (example: *domain-one.com)

- **<pattern-type>** is one of the following values: HTTPS-SNI, HTTP-host, QUIC-SNI

For example:

*domain-one.com HTTPS-SNI

*domain-one.com HTTP-host

prefix*domain-two.com HTTPS-SNI

Subscriber identification

Policy control

QoE enforces policies on subscriber sessions. A subscriber session is all traffic of a distinctive IP address on the access side. That is one single IPv4 address or one IPv6 subnet. For example, a policy with rate limits apply those limits to the total throughput of that distinctive IP address.

> Note
>
> If there is a NAT between the QoE server and the real subscribers, subscribers whose IP addresses are translated to the same IP address would be considered as the same subscriber.

The default IPv6 subnet is /64 (to change it, navigate to **Administration > General Settings** and edit **IPv6 prefix for subscribers**field).

> Note
>
> If there is a NAT between the QoE server and the real subscribers, subscribers whose IP addresses are translated to the same IP address would be considered as the same subscriber.

A new subscriber is identified if the first packet from an IP address is received. Then the subscriber rate and monitoring rules are evaluated, to choose the policies to be applied.

Subscriber ID

The subscriber ID field is supported to facilitate the identification of a subscriber session. The subscriber ID can be used when requesting metrics to obtain historical information even when the subscriber has changed IP address over time.

The following are possible sources for the subscriber ID:

- **MAC access address** (this is the default): In some networks, the MAC address may be the same for all subscribers (for example if all traffic is coming from the same router port) but in other networks, it may identify the subscriber access points.
- **Subscriber access IP address**: To configure the IP address and to fill the subscriber ID, navigate to **Administration** > **General Settings** > **Default subscriber-ID source**.

As part of the subscriber POST/PUT operation, a subscriber ID can be provided (For more information, see *QoE Appliance RESTAPI Guide*)

If External billing system is configured, then the Subscriber ID can be selected based on the type of billing system as follows:

RADIUS:

- Username
- Calling-Station-ID

Powercode billing system:

- Customer-ID
- Equipment-ID
- MAC address
- Name

Powercode-API:

- Customer-ID
- Name

Azotel billing system:

- Customer-ID
- Name
- Nickname

Sonar billing system:

- Customer-ID
- Name

Splynx billing system:

- Customer-ID
- Username

- Login

To configure a parameter from an external billing system, navigate to **Configuration** > **RADIUS/REST/Billing** > **Billing System**.

Subscriber flow policies

When a new flow is created, a subscriber flow policy is assigned to it, which specifies treating all the flows within that subscriber. The following are the actions that can be defined in a subscriber flow policy:

- **TCP Optimization**: It improves TCP traffic performance. It specifies whether to apply optimization to TCP traffic. It is recommended to enable the TCP Optimization (the default value).

- **Shaping per subscriber**: It limits the combined speed of a subscriber flows to a given value. It is possible to limit in the downlink and/or uplink direction. The limit applies to all flows matching the policy belonging to the same subscriber. For example, if video streaming flows are assigned to a flow policy with a 6 Mbps limit, and the subscriber has three video streaming flows at the same time, the three flows will share the 6 Mbps limit (getting around 2 Mbps each). It is possible to define bursts that allow flows to exceed temporally the limit (see the end of this section).

- **Shaping per flow**: It limits the speed of one flow to a given value. It is possible to limit in the downlink and/or uplink direction. The limit applies to any flow matching the policy. For example, if video streaming flows are assigned to a per flow 2 Mbps limit, a video flow cannot exceed those 2 Mbps. Shaping per flow can be combined with shaping per subscriber: for example, if there is a per subscriber 6 Mbps limit and a 2 Mbps per flow, a subscriber with four flows will have them limited to the 6 Mbps maximum (around 1.5 Mbps each). Per flow shaping has no burst option. Because per-flow shaping is not applied per subscriber, it can be used even when there is a NAT between the QoE and the end subscribers.

- **Block**: It blocks all flows falling in the blocking policy, in both directions, and does not let it proceed. It should be used with care, to avoid affecting traffic different to the one intended.

- **Drop incoming connections**: It blocks only flows that are started from the Internet side, but not flows started from the subscribers.

- **Skip subscriber rate limitation**: The traffic from flows getting this policy are no longer be affected by the rate limitation specified in the rate policy for this subscriber. They get only the rate limitation specified by this flow policy (if any).

These policies are configured from the menu option **Configuration > Subscriber Flows**, and select the **POLICIES** tab.

Shaping per subscriber

The following example defines a downlink speed limit of 10 Mbps, an uplink speed limit of 8 Mbps, and bursts of 3 seconds of double the normal speed. Figure 63 shows the Edit subscriber flow policy page.

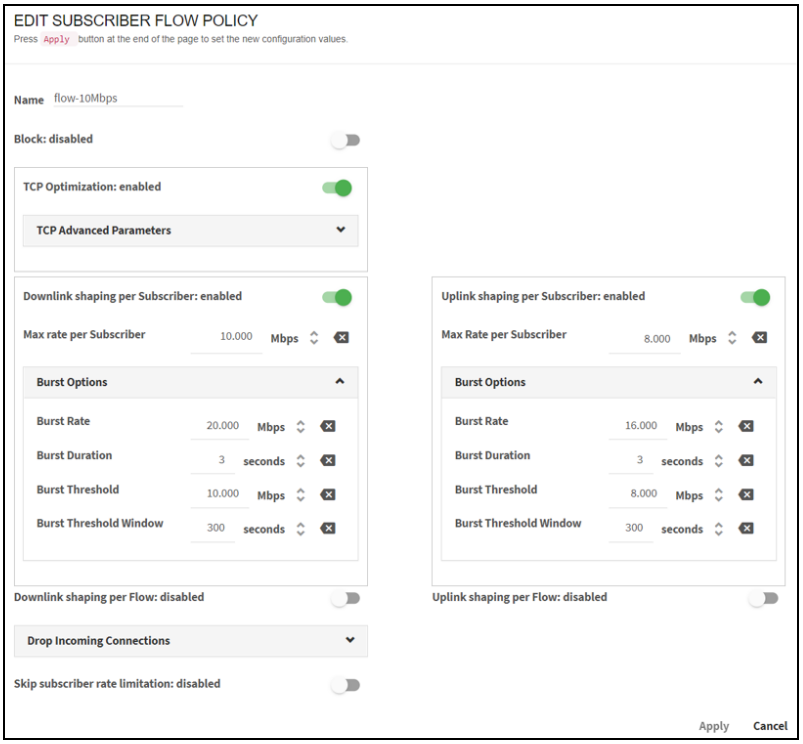

EDIT SUBSCRIBER FLOW POLICY

Press Apply button at the end of the page to set the new configuration values.

Name: flow-10Mbps

Block: disabled

TCP Optimization: enabled

Downlink shaping per Subscriber: enabled

Max rate per Subscriber: 10.000 Mbps

Burst Rate: 20.000 Mbps

Burst Duration: 3 seconds

Burst Threshold: 10.000 Mbps

Burst Threshold Window: 300 seconds

Uplink shaping per Subscriber: enabled

Max Rate per Subscriber: 8.000 Mbps

Burst Rate: 16.000 Mbps

Burst Duration: 3 seconds

Burst Threshold: 8.000 Mbps

Burst Threshold Window: 300 seconds

Downlink shaping per Flow: disabled

Uplink shaping per Flow: disabled

Skip subscriber rate limitation: disabled

Figure 63: *The Edit subscriber flow policy page*

Burst options

Bursts are configured under **Burst Options** of the appropriate direction. Burst policy is defined by the following parameters:

- **Burst Rate**: The maximum rate during the burst, typically bigger than the normal shaping maximum rate (for example, allow burst of 20 Mbps for flows normally limited to 10 Mbps).
- **Burst Duration**: The duration of the burst, that is time taken for the burst rate can be sustained.
- **Burst Threshold**: An average speed that, if exceeded, prevent a new burst from happening. It is the way to control when a new burst can be granted. For example, for a 10 Mbps limit with 20 Mbps bursts, a 5 Mbps burst threshold requires the subscriber flows to drop the speed to half its normal limit before allowing a new burst.
- **Burst Threshold Window**: Period (in seconds) used to compute the average speed that is checked versus the threshold. The longer window, the bigger weight of past subscriber activity on the decision of grating a new burst.

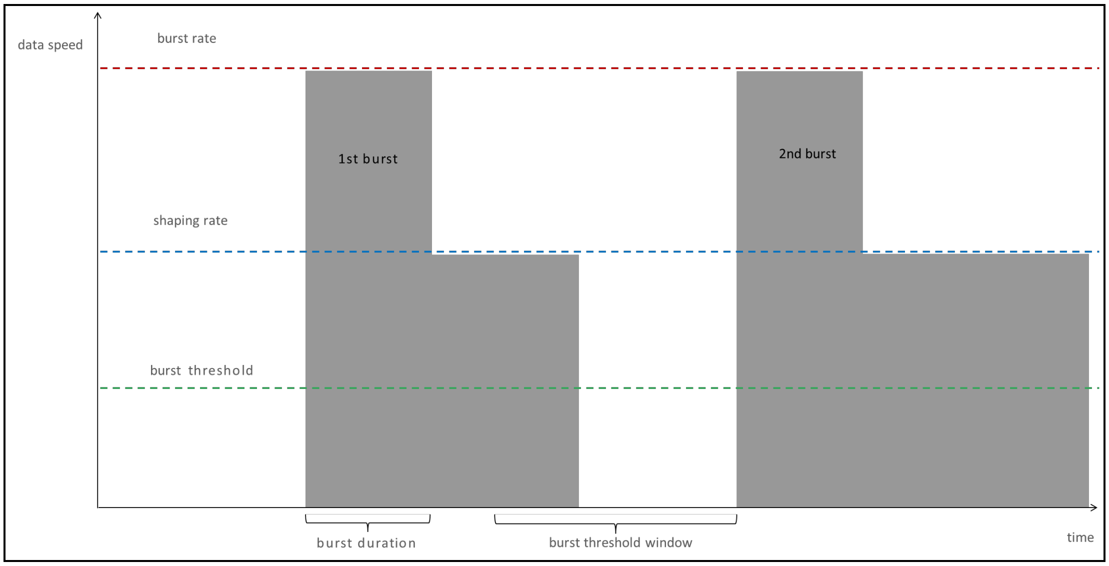

Figure 64: *Burst options*

Shaping per flow

It is possible to add a shaping per subscriber. Per flow and per subscriber shaping limits act at the same time, per flow shaping limiting the speed of individual flows and subscriber shaping limiting the combined flow speed per subscriber.

The following example is a policy with a limit per flow of 4 Mbps in either direction. Figure 65 shows shaping per flow.

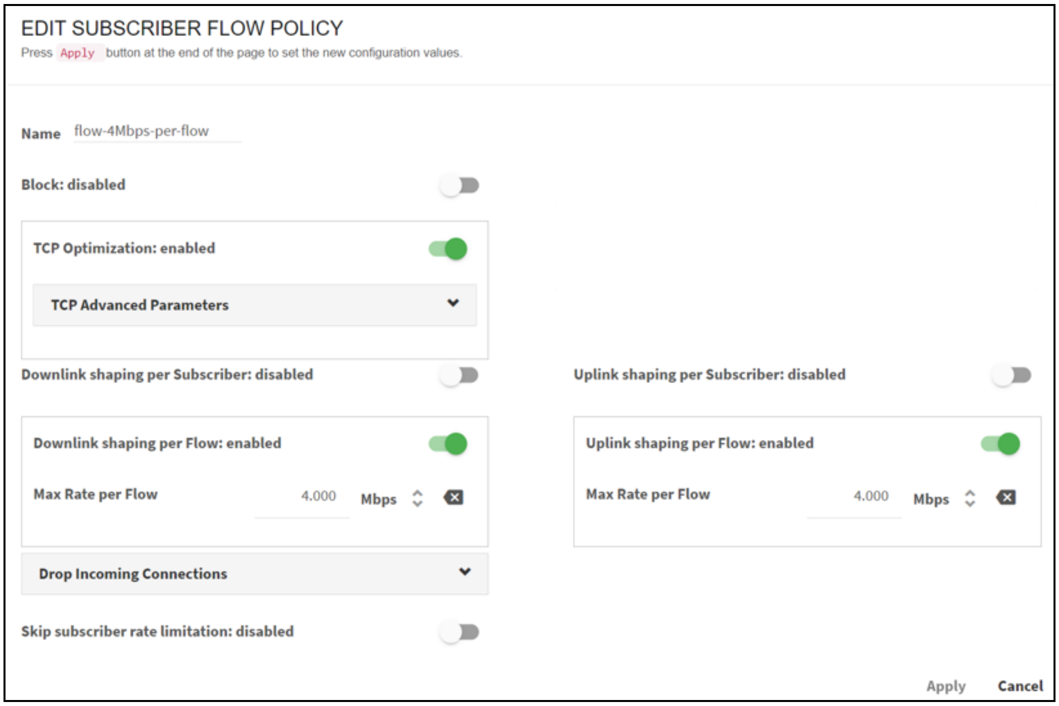

EDIT SUBSCRIBER FLOW POLICY

Press Apply button at the end of the page to set the new configuration values.

Name flow-4Mbps-per-flow

Block: disabled

TCP Optimization: enabled

TCP Advanced Parameters

Downlink shaping per Subscriber: disabled

Uplink shaping per Subscriber: disabled

Downlink shaping per Flow: enabled

Max Rate per Flow 4.000 Mbps

Uplink shaping per Flow: enabled

Max Rate per Flow 4.000 Mbps

Drop Incoming Connections

Skip subscriber rate limitation: disabled

Figure 65: *The Edit subscriber flow policy page*

It is possible to add a shaping per subscriber. Per flow and per subscriber shaping limits act at the same time, per flow shaping limiting the speed of individual flows and subscriber shaping limiting the combined flow speed per subscriber.

Blocking incoming traffic

It is possible to block incoming traffic, initiated from the Internet (TCP connections, UDP flows or other IP traffic like ICMP pings). To perform this, follow the **Drop Incoming Connections** section as part of a **Subscriber Flow** policy. The edit subscriber flow policy page shows the edit subscriber flow policy page.

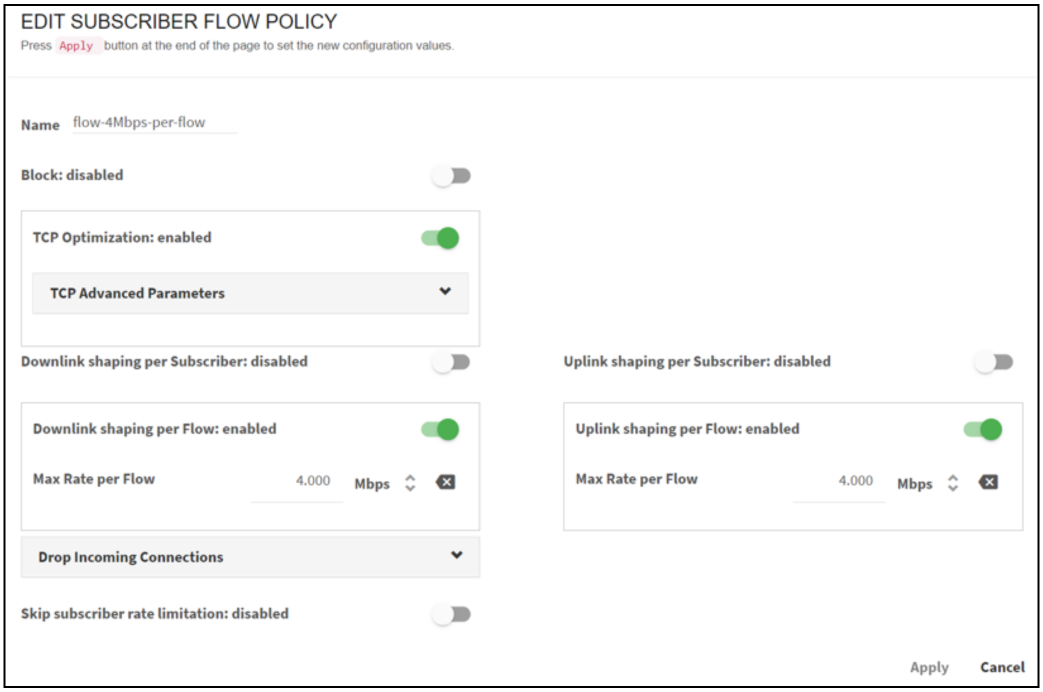

EDIT SUBSCRIBER FLOW POLICY

Press Apply button at the end of the page to set the new configuration values.

Name: flow-4Mbps-per-flow

Block: disabled

TCP Optimization: enabled

TCP Advanced Parameters

Downlink shaping per Subscriber: disabled

Uplink shaping per Subscriber: disabled

Downlink shaping per Flow: enabled

Max Rate per Flow 4.000 Mbps

Uplink shaping per Flow: enabled

Max Rate per Flow 4.000 Mbps

Drop Incoming Connections

Skip subscriber rate limitation: disabled

Figure 66: *The edit subscriber flow policy page*

Subscriber rate policies

Subscriber rate policies are applied per subscriber.

The following are possible actions:

- **Maximum downlink speed**: It is the maximum speed in the downlink direction for all traffic going towards the subscriber's IP address.
- **Maximum uplink speed**: It is the maximum speed in the uplink direction for all traffic coming from the subscriber's IP address.
- Under **Advanced Parameters**: You can find the same burst options as for subscriber flow policies.
- There is an **Automatic Congestion Management (ACM)** option, that detects congestion and select a rate limit automatically (enabled by default).

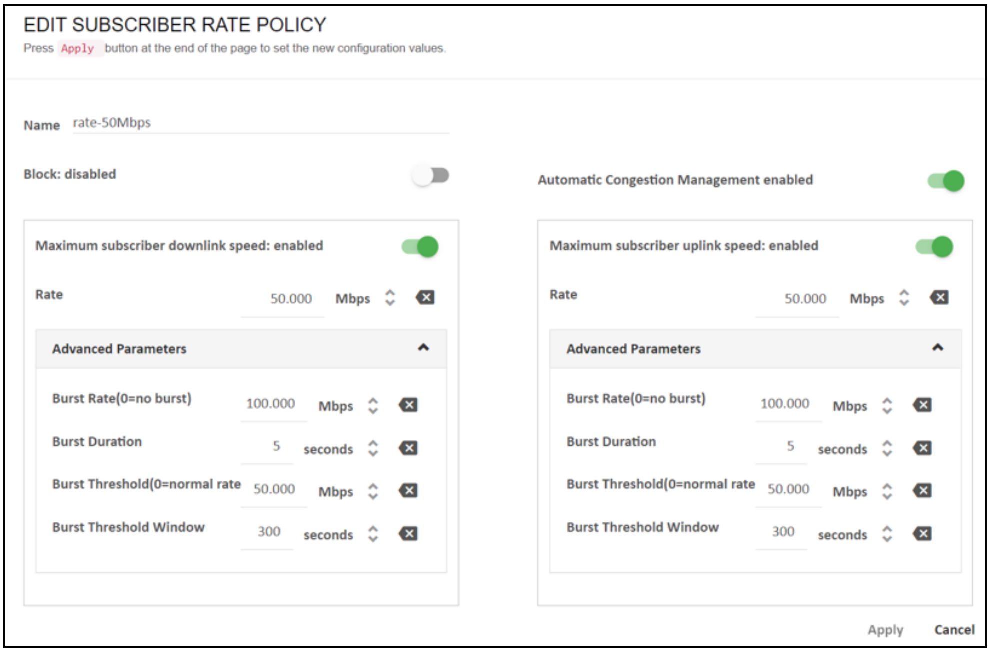

EDIT SUBSCRIBER RATE POLICY

Press Apply button at the end of the page to set the new configuration values.

Name rate-50Mbps

Block: disabled

Automatic Congestion Management enabled

Maximum subscriber downlink speed: enabled

Rate 50.000 Mbps

Advanced Parameters

Burst Rate(0=no burst) 100.000 Mbps

Burst Duration 5 seconds

Burst Threshold(0=normal rate 50.000 Mbps

Burst Threshold Window 300 seconds

Maximum subscriber uplink speed: enabled

Rate 50.000 Mbps

Advanced Parameters

Burst Rate(0=no burst) 100.000 Mbps

Burst Duration 5 seconds

Burst Threshold(0=normal rate 50.000 Mbps

Burst Threshold Window 300 seconds

Figure 67: *Edit subscriber rate policy page*

To configure the policies, navigate to **Configuration > Subscriber Rates** and select **POLICIES** tab. To identify traffic from the same subscribers, refer to Subscriber identification section.

Subscriber rate policies can also be created dynamically through the QoE external interfaces (RADIUS, REST API or an interface to an external billing system). In that case, the policy parameters, and the associations of the policy to the subscribers are controlled from the external interfaces and independent of QoE-configurated rules (configured rules are a fallback for subscribers without a policy assigned by the external interface). The external interfaces can assign subscribers both to rate policies defined though the external interface and to rate policies defined though the QoE UI.

Enable/Disable ACM optimization

To enable or disable ACM in a a configured policy, change the **Automatic Congestion Management** field in the **Subscriber Rate** window. shows the edit subscriber rate policy page.

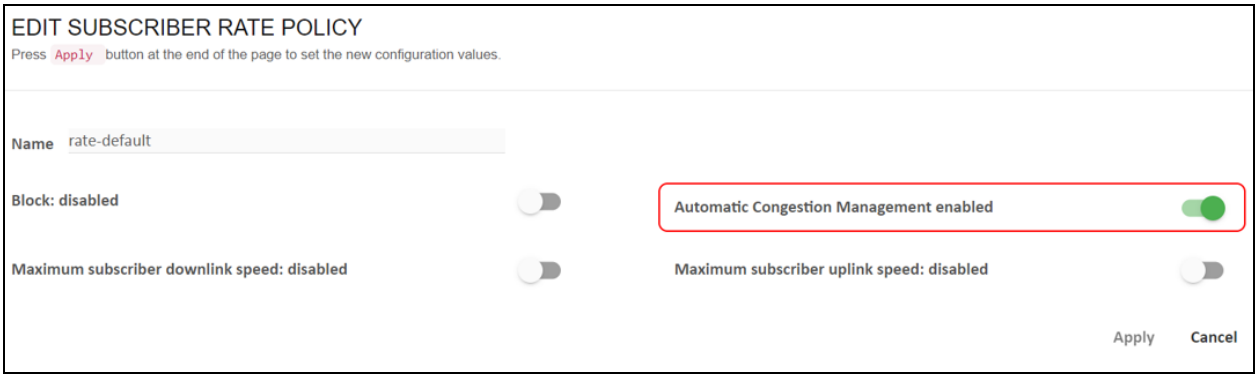

EDIT SUBSCRIBER RATE POLICY

Press Apply button at the end of the page to set the new configuration values.

Name rate-default

Block: disabled

Automatic Congestion Management enabled

Maximum subscriber downlink speed: disabled

Maximum subscriber uplink speed: disabled

Subscriber monitoring policies

Subscriber monitoring policies are applied per subscriber. They just determine the amount of sampling of DPI records for each subscriber. By default, the QoE comes with a monitor-default policy that applies to all subscribers, which automatically determines the number of DPI records (called UDRs, or Usage Data

Records) to produce, so that the DPI statistics are meaningful and do not take many CPU and disk to produce. The monitor-default is the recommended mode of operation.

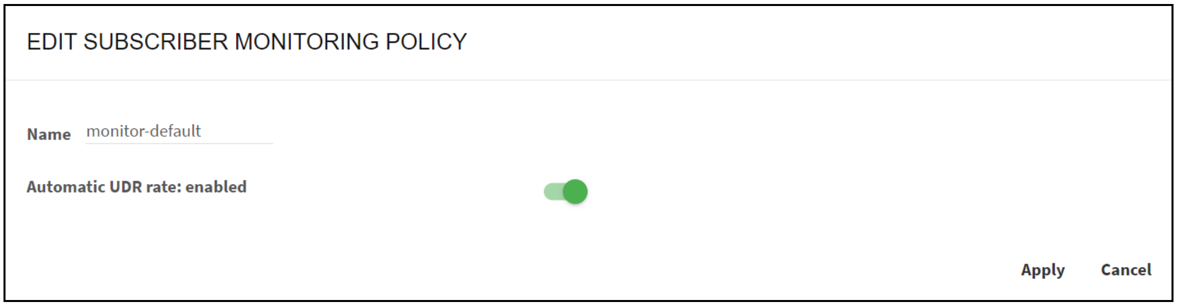
EDIT SUBSCRIBER MONITORING POLICY

Name monitor-default

Automatic UDR rate: enabled

It is possible to adjust the amount of UDRs for certain subscribers by creating a monitoring policy that specifies percentage of flows with UDRs (for example, 0% or 100%). UDRs for 100% of flows produce and accurate description of the subscriber usage statistics, but for performance reasons, it is recommended to set a 100% UDR monitoring policies only on a few subscribers at a time, because these policies can produce a huge number of records that can use all the available disk space and also can make the production of DPI statistics very slow and CPU-consuming.

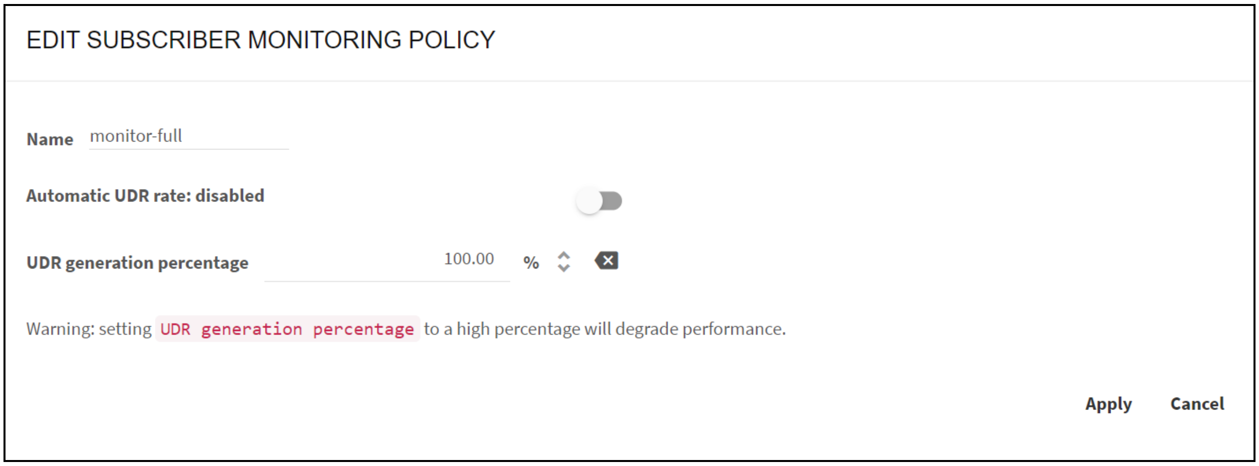
EDIT SUBSCRIBER MONITORING POLICY

Name monitor-full

Automatic UDR rate: disabled

UDR generation percentage 100.00 %

Warning: setting `UDR generation percentage` to a high percentage will degrade performance.

Rules

Rules specify the configured policies to be assigned to each subscriber and flow, based on the profiles in the rule.

There are independent sets of rules for each policy type. Subscriber flow rules select the appropriate configured subscriber flow policy for each flow, subscriber rate rules select the appropriate subscriber rate policy for each subscriber, and similarly, subscriber monitoring rules select the appropriate subscriber monitoring policy for each subscriber.

A rule can use one profile of each type (or, alternatively, use the **any** option, if the profile type is indifferent), and it defines one and only one policy to apply.

Every set of rules may have many rules, but only the one with the best match is selected for each flow or subscriber. To evaluate the rules in a way that maximizes performance, profiles are checked in order. This predefined order determines which rule is finally selected. A tree view of the rules helps in identifying which rule is selected in each case. For more information on the trees and the profile evaluation order, refer to Subscriber flows decision tree section .

Manually configured rule priorities are not used, because of the performance penalties they entail and the burden on the operator to keep priorities consistent.

The subscriber flow rules are configured in the menu option **Configuration > Subscriber Flows**, selecting the tab **RULES TREE-VIEW** or **RULES TABLE-VIEW.** Similarly, the subscriber rate rules are configurable from **Configuration > Subscriber Rates** and the subscriber monitoring rules from **Configuration > Subscriber Monitoring**.

Subscriber flows decision tree

The evaluation of subscriber flow rules is as follows: when a new traffic flow is created (for example, a TCP connection), the profiles in the subscriber flow rule set are checked against that new flow, a matching rule selected, and with it the flow policy to apply.

For efficiency, profiles are evaluated in this predefined order:

1. Interface
2. VLAN
3. Policy-rate
4. Internet
5. Access
6. Subscriber group
7. Subscriber ID
8. Time
9. Throughput
10. DPI

The profile evaluation order defines a decision tree, whose nodes are the different profiles and with policies as leaves. The tree determines which rule is finally selected, because a rule can be excluded if it belongs to a branch that the decision tree does not follow. It may be the case that a flow matches more than one rule. In that case, the profile type order is important: for example, a rule matching the Interface profile would have priority over the rule matching the VLAN profile, and so on in the order specified above.

If two rules have a match with the same profile type, the more restrictive profile would have priority. For example, a flow from a subscriber with IP address 192.168.0.1 would match both an access profile with the 192.168.0.0/24 range and another access profile with the 192.168.0.0/16 range, but the first rule, with a narrower range (/24 vs /16), and therefore more restrictive, is selected.

To facilitate the understanding of this order, the UI includes a graphic representation of the decision tree, where the better matching path leads to the selected policy (except when there is more than one match at the same profile level, when the most restrictive would win). To view the decision tree, navigate to **Configuration > Subscriber Flows > Rules** and select **RULES TREE-VIEW** tab.

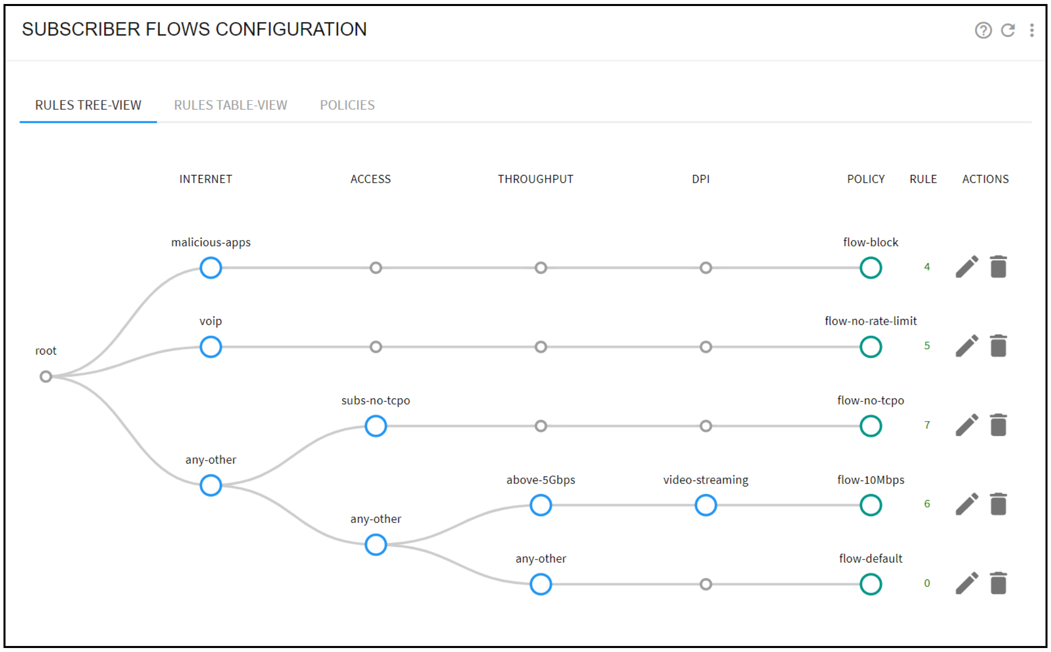

Figure 68: *The subscriber flows configuration page*

If there are common elements in two profiles of the same type and therefore a rule conflict, the decision tree gives the signal so the rules can be reviewed by the operator and the conflict corrected. In the following example, two access profiles have an overlap (they both contain the same IP range). This is signaled with a warning window and also, one of the conflicting rules has its number in yellow. Removing the overlap (for example making one of the profiles to have a more specific range), the conflict disappears.

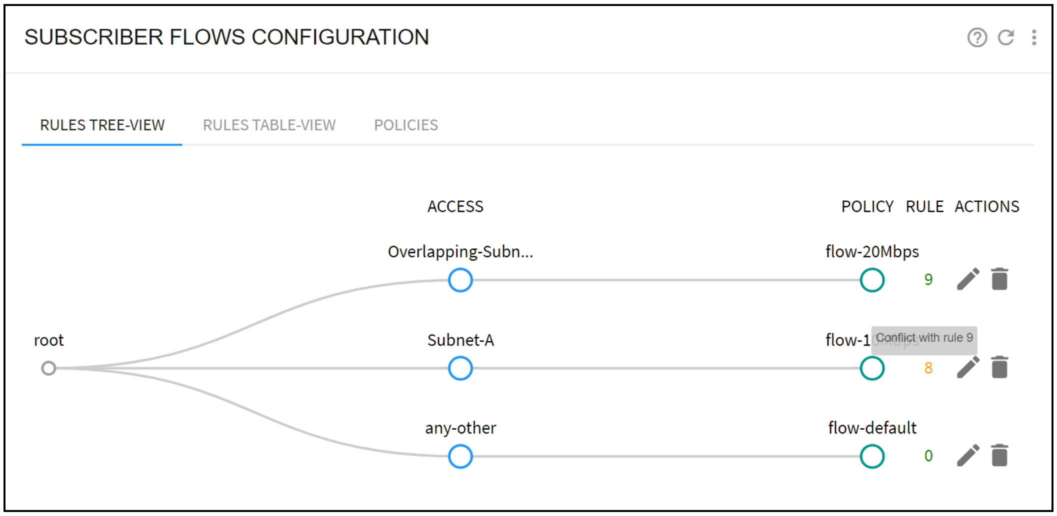

Figure 69: *The subscriber flows configuration page*

Subscriber rate decision tree

The evaluation of the subscriber monitoring rules happens whenever a new subscriber session is detected (when traffic from a new access IP address is received). The profiles are checked against the subscriber session and a matching rule found, that specifies the rate policy to apply. For efficiency, the profiles are evaluated in the following pre-determined order:

1. Interface
2. VLAN
3. Access
4. Subscriber group
5. Subscriber ID
6. Time

Other profile types cannot be used in subscriber rate rules. For example, Internet profiles and DPI profiles apply to some of the subscriber applications and not to others. Another example is throughput profile, because the rate rules are evaluated at the start of the subscriber session and the throughput changes continuously. Figure 70 shows the subscriber rate configuration page.

The decision tree is similar to the tree of subscriber flow rules. To view the tree of subscriber flow rules, navigate to **Configuration > Subscriber Rate > Rules** and select **RULES TREE-VIEW** tab.

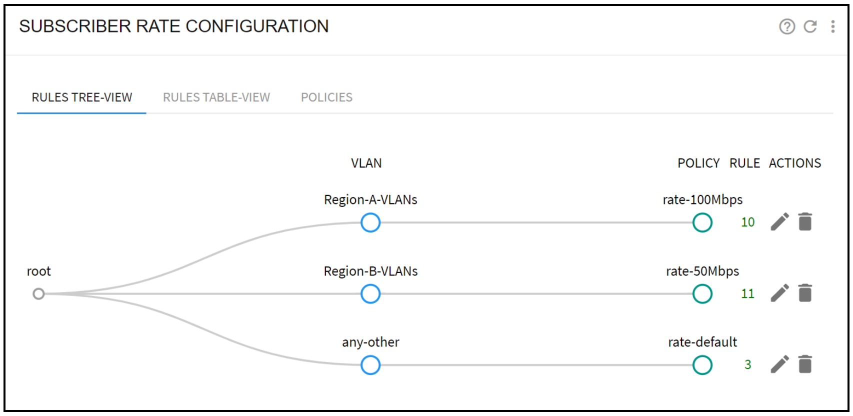

Figure 70: *The subscriber rate configuration page*

Subscriber monitoring decision tree

The evaluation of the subscriber monitoring rules happens whenever a new subscriber session is detected (that is when traffic from a new access IP address is received). The profiles are checked against the subscriber session and a matching rule found, that specifies the rate policy to apply. For efficiency, the profiles are evaluated in the following pre-determined order:

1. Interface
2. VLAN
3. Access
4. Subscriber group
5. Subscriber ID

Other profile types cannot be used in subscriber rate rules. For example, Internet profiles and DPI profiles apply to some of the subscriber applications and not to others. Another example is throughput profile, because the rate rules are evaluated at the start of the subscriber session and the throughput changes continuously.

The decision tree is similar to the tree of subscriber flow rules. To view the decision tree, navigate to **Configuration > Subscriber Monitoring > Rules** and select **RULES TREE-VIEW** tab.

Policy examples

The following are several common examples of policies:

- Implementing subscriber rate plans
- Limiting the speed of some applications
- Limiting the speed of some applications with NAT
- Services not limited by the subscriber rate
- Blocking applications
- Exclude traffic from TCP optimization

Implementing subscriber rate plans

The objective is to apply the speed limits in each subscriber's data plan.

The QoE applies these limits better than a conventional shaping element because, for TCP traffic (the most common), it does not need to discard packets. It uses independent queues per flow and that makes application latencies independent of each other, which greatly improves the experience of interactive applications. Figure 71 shows the queue structure, with a queue per flow and policy control at flow and subscriber levels.

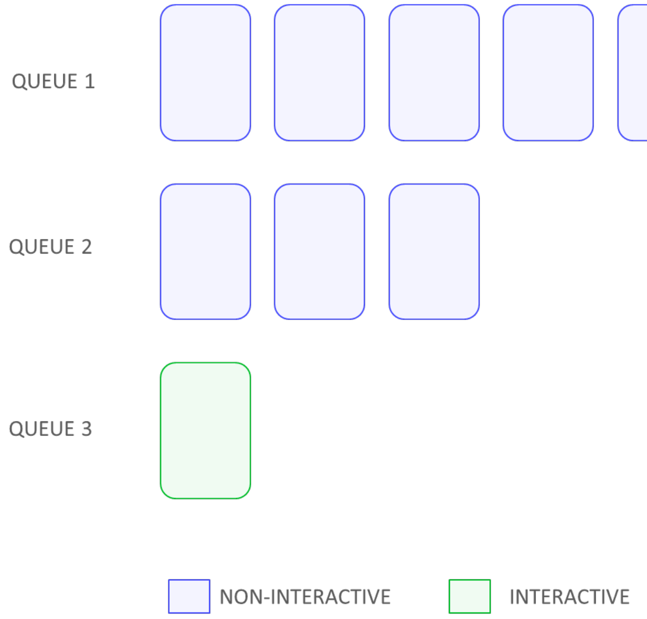

Figure 71: *Queue structure*

The easiest way to implement rate plans is to use the RADIUS, REST or Billing interfaces (see Monitoring with SNMP and REST API sections) interfaces. Rate policies then be assigned by an external system for every subscriber. The external system can directly create those policies, or it can assign rate policies configured from the QoE UI.

If an external system cannot be used, you can create one rate policy for each plan, an access policy with all the subscriber IP addresses (or ranges) assigned to that plan, and then a rule linking the corresponding access profiles and rate policies. Figure 72 is an example of the rule tree.

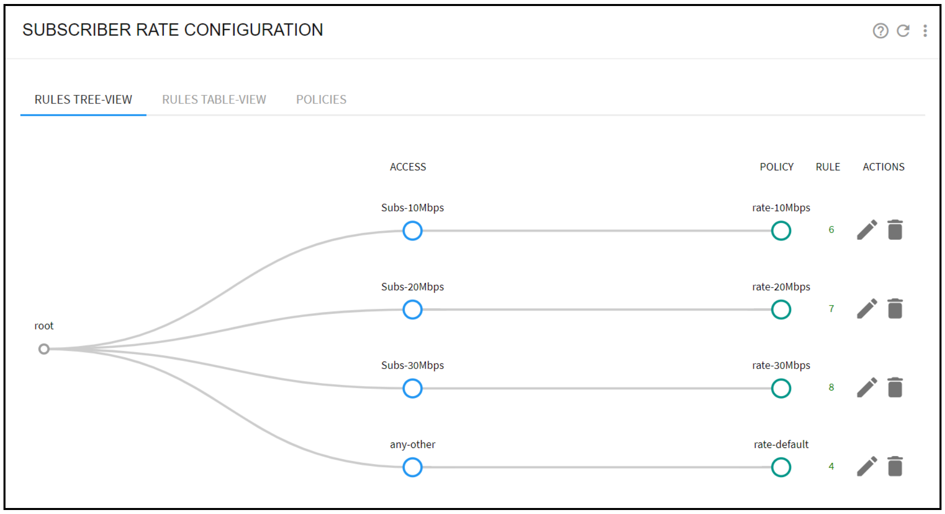

Figure 72: *Rule tree*

It is also possible to define the rules just for one test IP address that is used during a proof-of-concept to see the QoE performance as a bandwidth manager. Figure 73 shows the subscriber rate configuration page.

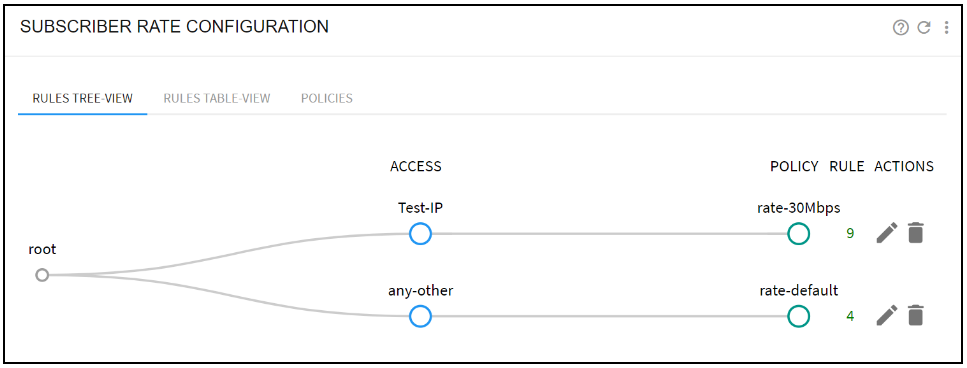

Figure 73: *Subscriber rate configuration*

Limiting the speed of some applications

This is used to reduce the network peak throughput to mitigate the congestion at the peak hour. To that end, a DPI profile is defined (video in the example) to identify the applications to limit. This example makes use of video streaming predefined signatures. To include them, in **Add DPI profile**, select **Add Predefined Signatures** and choose the **Video Streaming** predefined signature.

Also, a throughput profile is created with the traffic load from which to start limiting (**above-5Gbps** in this example). Then, a subscriber flow policy (*flow-10Mbps* in the example) is created with a downlink limit (*Downlink shaping per Subscriber*) set at 10 Mbps. Finally, the DPI profile, the throughput profile and the subscriber flow policy are tied together in a subscriber flow rule. Figure 74 shows the subscriber flows configuration page.

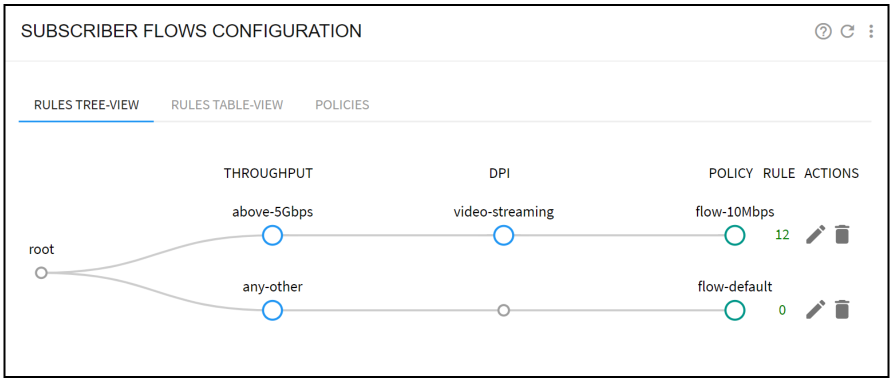
SUBSCRIBER FLOWS CONFIGURATION

RULES TREE-VIEW RULES TABLE-VIEW POLICIES

THROUGHPUT DPI POLICY RULE ACTIONS

root

above-5Gbps video-streaming flow-10Mbps 12

any-other flow-default 0

Figure 74: *Subscriber flows configuration*

Limiting the speed of some applications with NAT

This is used to reduce the network peak throughput to mitigate the congestion at the peak hour. In this scenario there is a NAT between the QoE server and the end subscribers and, therefore, a shaping per subscriber is not possible.

Use the same throughput and DPI profiles of the previous example, but use a per-flow shaping policy. Figure 75 shows the edit subscriber flow policy page.

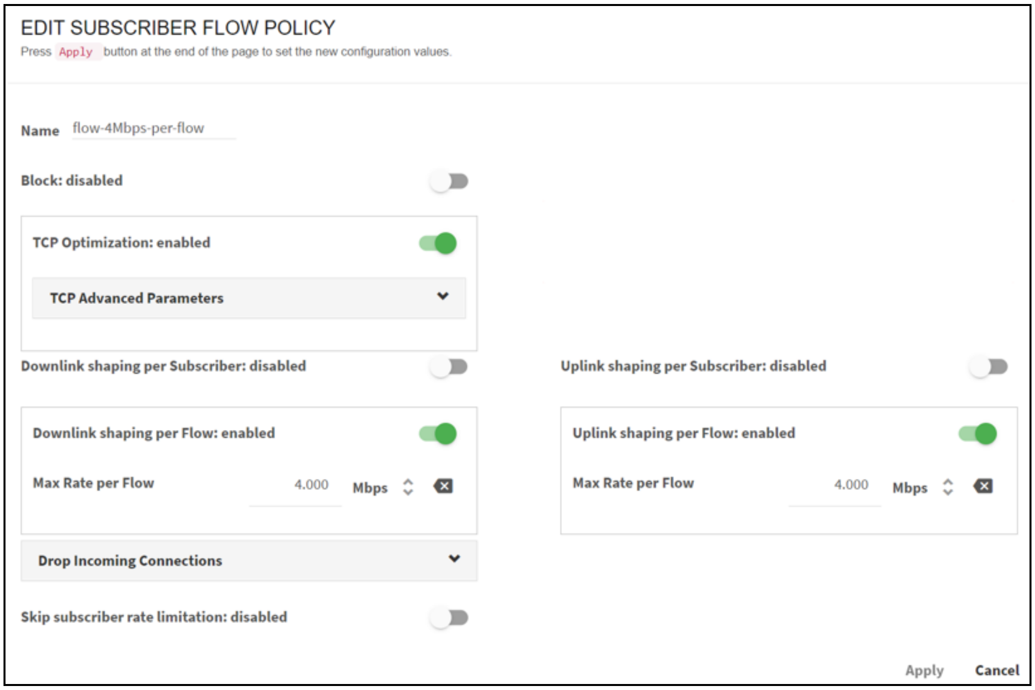
EDIT SUBSCRIBER FLOW POLICY

Press Apply button at the end of the page to set the new configuration values.

Name flow-4Mbps-per-flow

Block: disabled

TCP Optimization: enabled

TCP Advanced Parameters

Downlink shaping per Subscriber: disabled

Uplink shaping per Subscriber: disabled

Downlink shaping per Flow: enabled

Max Rate per Flow 4.000 Mbps

Uplink shaping per Flow: enabled

Max Rate per Flow 4.000 Mbps

Drop Incoming Connections

Skip subscriber rate limitation: disabled

Apply Cancel

Figure 75: *The edit subscriber flow policy page*

Services not limited by the subscriber rate

To preserve the quality of experience of some services by granting throughput to them even when the subscriber rate plan is fully used. For example, Voice over IP (VoIP) service hosted by the ISP. The policy is limited to subscribers with gold plans. It is created with an Internet profile (**voip**) with the IP address and port of the ISP-hosted VoIP service, a policy-rate profile and a flow policy (with **Skip subscriber rate limitation** selected to **On**). Next, the Internet profile and the policy-rate profiles are linked to the policy by a subscriber flow rule.

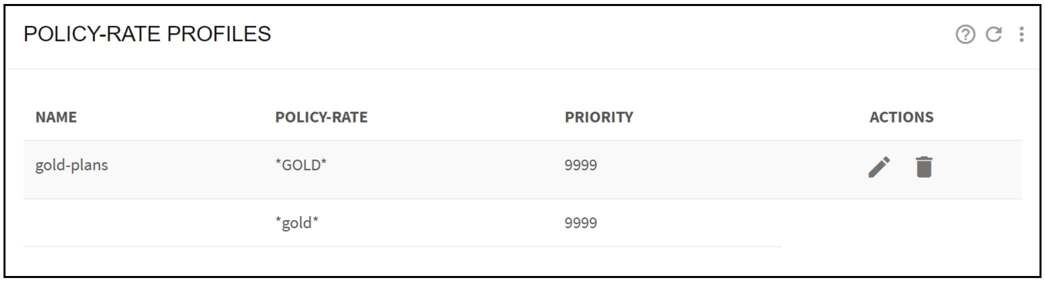

POLICY-RATE PROFILES

NAME	POLICY-RATE	PRIORITY	ACTIONS
gold-plans	*GOLD*	9999	
	gold	9999	

Figure 76: *The policy-rate profiles page*

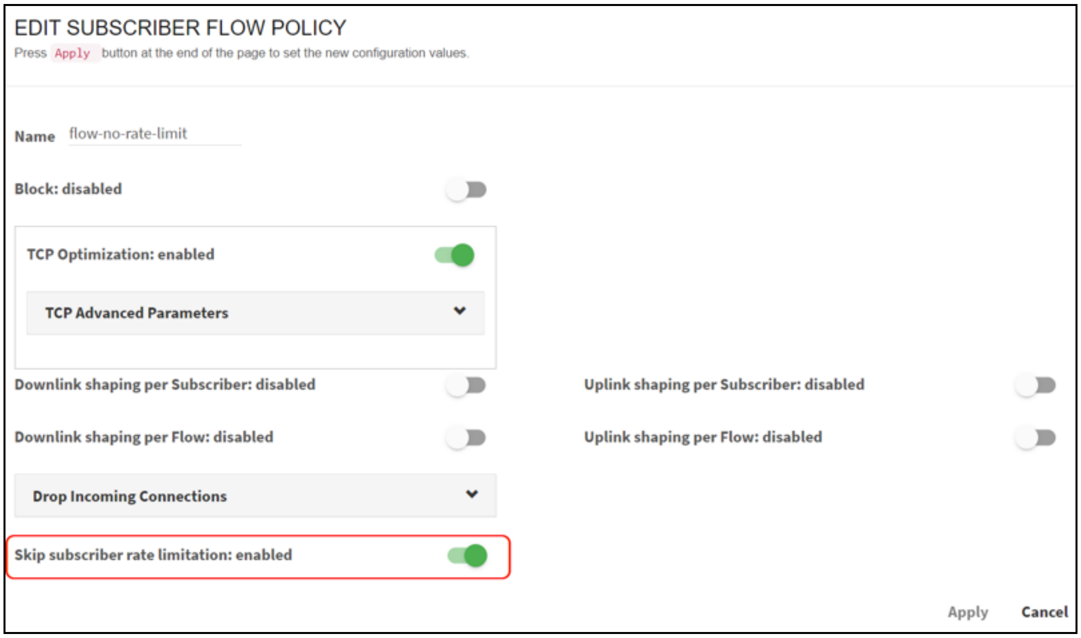

EDIT SUBSCRIBER FLOW POLICY

Press Apply button at the end of the page to set the new configuration values.

Name flow-no-rate-limit

Block: disabled

TCP Optimization: enabled

TCP Advanced Parameters

Downlink shaping per Subscriber: disabled

Uplink shaping per Subscriber: disabled

Downlink shaping per Flow: disabled

Uplink shaping per Flow: disabled

Drop Incoming Connections

Skip subscriber rate limitation: enabled

Figure 77: *The edit subscriber flow policy*

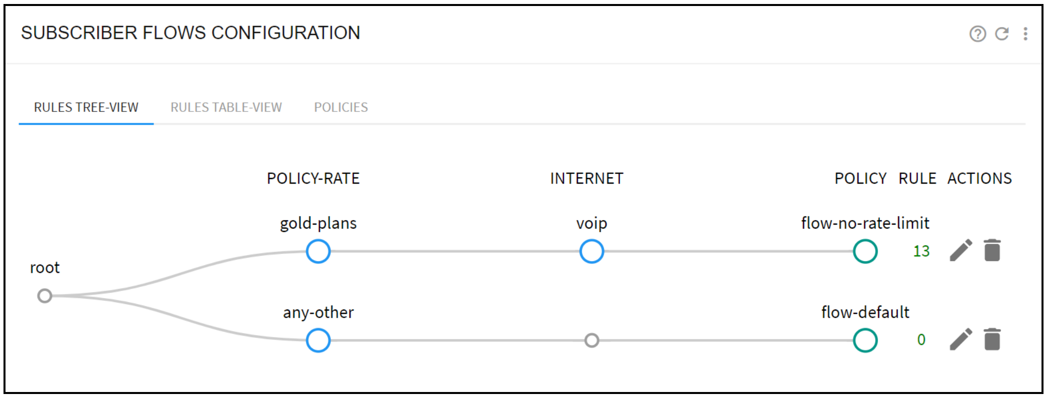

SUBSCRIBER FLOWS CONFIGURATION

Figure 78: *The subscriber flows configuration page*

Blocking applications

In this scenario, some applications require blocking, for example, servers that are sources of attacks. To perform this, an Internet profile is created (**malicious-apps** in the example) to identify the IP addresses to block. Next, a subscriber flow policy is defined with block action (with name **flow-block** in this example) and, finally, the Internet profile and the subscriber flow policy are combined in a subscriber flow rule.

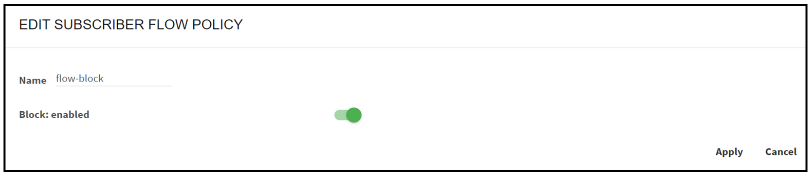
EDIT SUBSCRIBER FLOW POLICY

Name flow-block

Block: enabled

Figure 79: *The edit subscriber flow policy page*

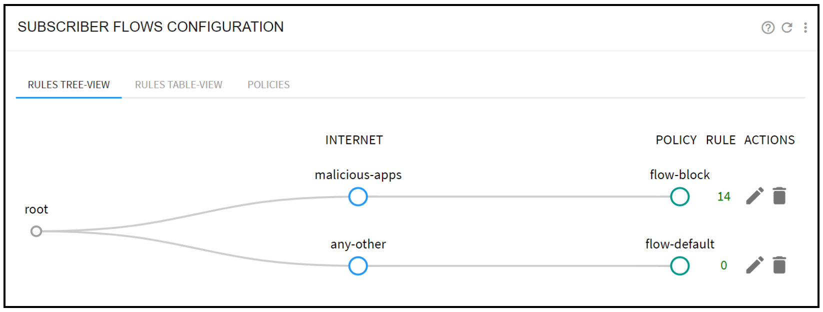
SUBSCRIBER FLOWS CONFIGURATION

Figure 80: *The subscriber flows configuration page*

Exclude traffic from TCP optimization

QoE does not optimize some traffic. For example, some subscribers. To that end, an access profile is defined (**subs-no-tcpo** in the example), with the subscriber IP addresses to exclude. A subscriber flow policy is defined with optimization set to off (**flow-no-tcpo** in the example) and, then the access profile and the subscriber flow policy are combined in a subscriber flow rule.

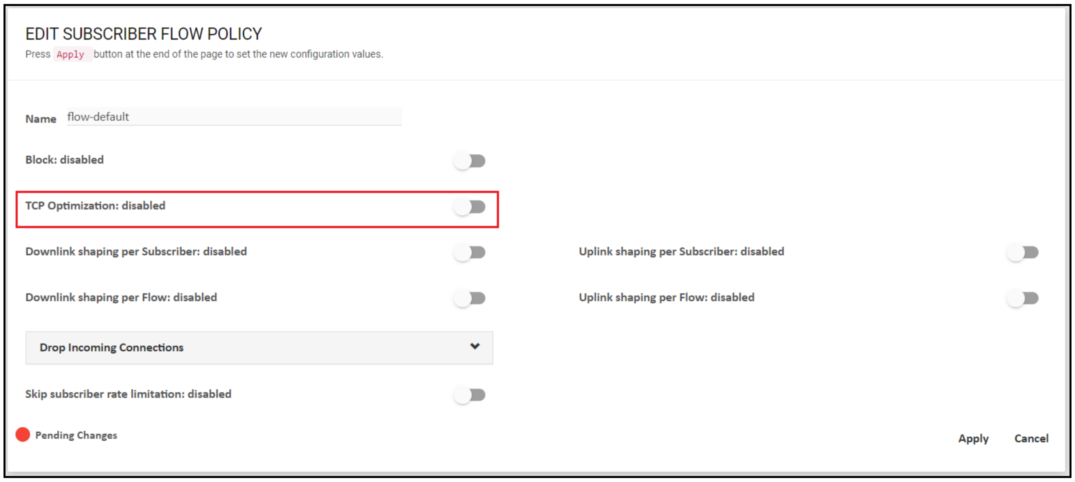
EDIT SUBSCRIBER FLOW POLICY

Press Apply button at the end of the page to set the new configuration values.

Name flow-default

Block: disabled

TCP Optimization: disabled

Downlink shaping per Subscriber: disabled

Uplink shaping per Subscriber: disabled

Downlink shaping per Flow: disabled

Uplink shaping per Flow: disabled

Drop Incoming Connections

Skip subscriber rate limitation: disabled

● Pending Changes

Figure 81: *The edit subscriber flow policy page*

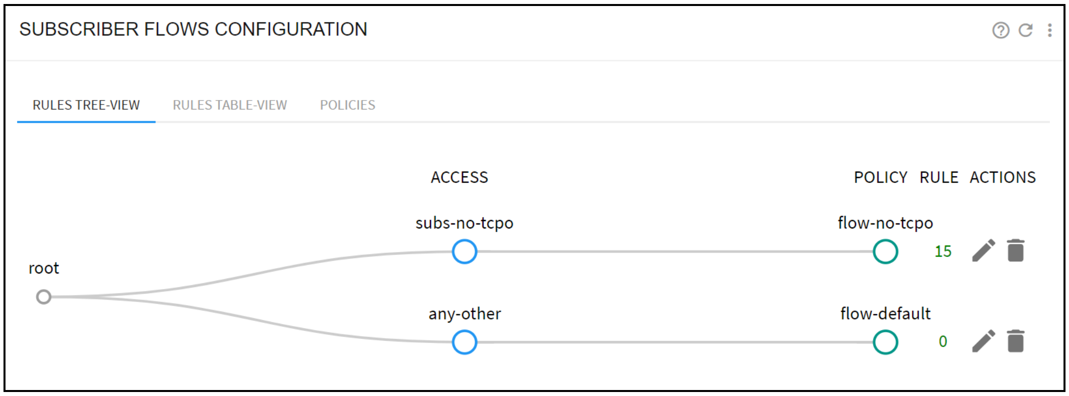

Figure 82: *The subscriber flows configuration page*

This setup is equivalent to the IP address blacklist in QoE Release 3.0 software. Another example is to use an Internet profile to exclude some applications per TCP port.

Chapter 10: Subscriber Quotas

The time and volume quotas can be associated to a subscriber IP address. Once a quota is exhausted, the subscriber IP address is restricted (by default, the traffic is blocked). Both a time and a volume quota can be associated together to an IP address, where the restriction happens if any of the quotas is exhausted.

> **Note**
>
> Quotas are assigned to the IP addresses. If a subscriber changes to the new the IP address, then that new IP address does not have a quota associated until one is provisioned through the REST API.

A time quota grants access for a period. There are two ways to define a time quota:

- As an absolute time. For example, 05/23/2023 (23rd of May, 2023).
- As an extension of current date. For example, 15 days from now.

A volume quota grants access for a volume of traffic. There are two ways to define a time quota:

- As an absolute amount. For example, 10 GB.
- As an extension of amount. For example, 5 GB on top of existing 10.

This chapter contains the following sections:

Quota general configuration

To configure general options of quota, navigate to **Status > Subscribers > Subscriber Quotas** and extend **Advanced quotas parameters**. Figure 83 shows the Quotas page.

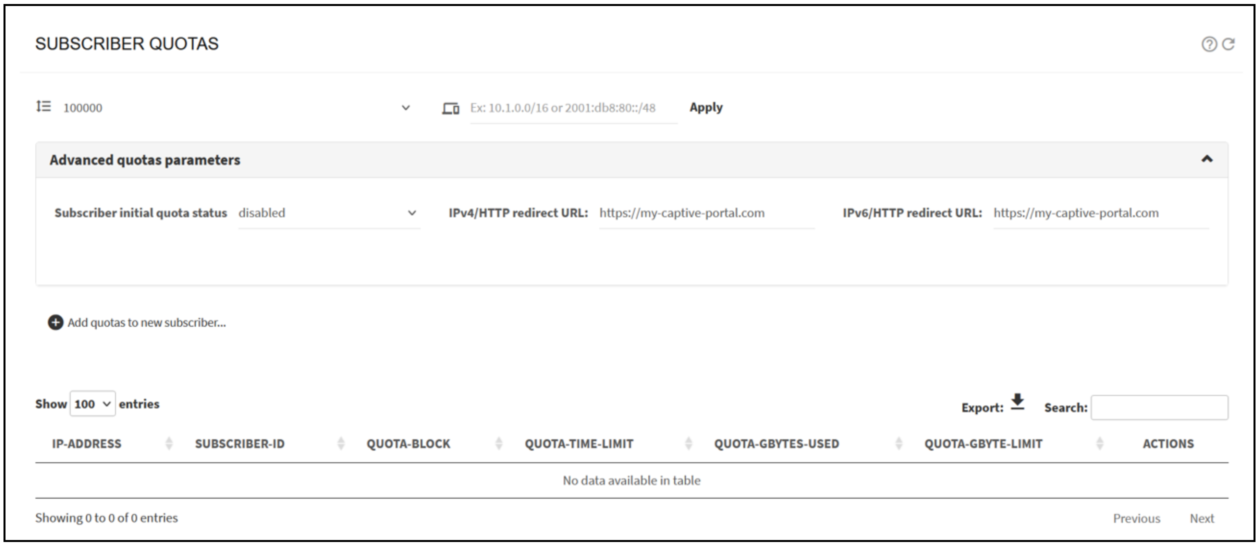

SUBSCRIBER QUOTAS

100000

Ex: 10.1.0.0/16 or 2001:db8:80::/48 Apply

Advanced quotas parameters

Subscriber initial quota status: disabled

IPv4/HTTP redirect URL: https://my-captive-portal.com

IPv6/HTTP redirect URL: https://my-captive-portal.com

Add quotas to new subscriber...

Show 100 entries

Export: Search:

IP-ADDRESS	SUBSCRIBER-ID	QUOTA-BLOCK	QUOTA-TIME-LIMIT	QUOTA-GBYTES-USED	QUOTA-GBYTE-LIMIT	ACTIONS
No data available in table						

Showing 0 to 0 of 0 entries

Previous Next

Figure 83: *The Quotas page*

The subscriber initial quota status defines the options to do with IP addresses without an assigned quota:

- When set to disabled (the default), traffic is allowed, without restrictions.
- When set to blocked, traffic is blocked until a valid quota is assigned.

The redirect URL fields specify the sites to redirect HTTP traffic when an IP address is blocked (captive portal) because of quota exhaustion. There is one field to redirect IPv4 traffic and another for IPv6 traffic. The two fields can have the same URL if the same captive portal is used for both IPv4 and IPv6.

- If the field is empty, no redirection is attempted.
- If a URL is specified, a redirection is attempted to that URL for the corresponding IP version of the HTTP traffic.

Note

Only HTTP redirections are supported, the site to redirect the traffic can be HTTPS. The URL used is https://my-captive-portal.com. HTTPS redirections are not supported, the modern browsers are protected against redirection attempts for security reasons.

If HTTP re-directions are used, a policy is required to allow the traffic to the redirection sites (and associated DNS queries). For more information, see Captive portal policies section.

Associating quotas to subscriber IPs

To associate a quota to a subscriber IP address, navigate to **Status** > **Subscribers** > **Subscriber Quotas**. It can also be accessed from **Configuration** > **Subscriber Quotas**.

Click **Add Quota to new Subscriber**...

Figure 84 defines a time quota as an absolute time.

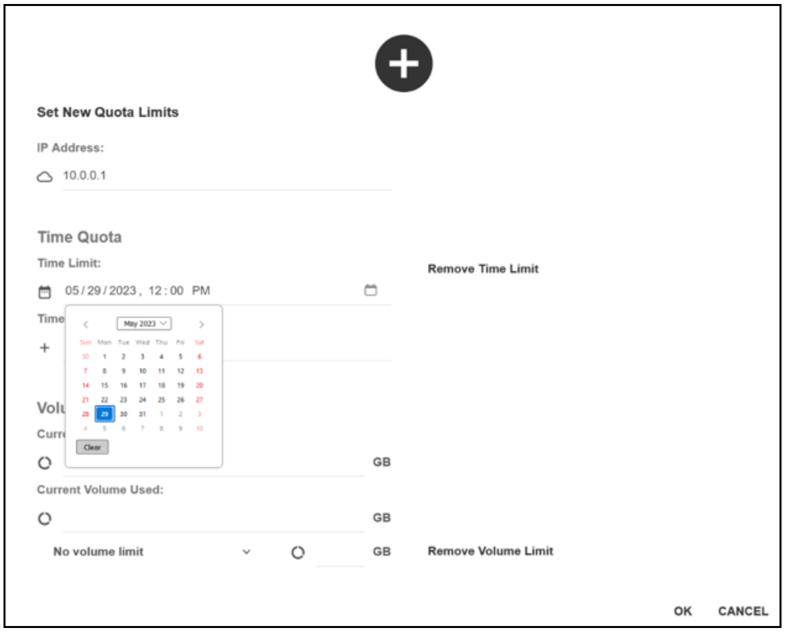

Set New Quota Limits

IP Address:

10.0.0.1

Time Quota

Time Limit:

05 / 29 / 2023 , 12 : 00 PM

Remove Time Limit

Time

Vol

Curr

GB

Current Volume Used:

GB

No volume limit GB

Remove Volume Limit

Figure 84: *Associating Quotas to Subscriber IPs*

It is also possible to define the time quota relative to current date and time (+1 month in the example).

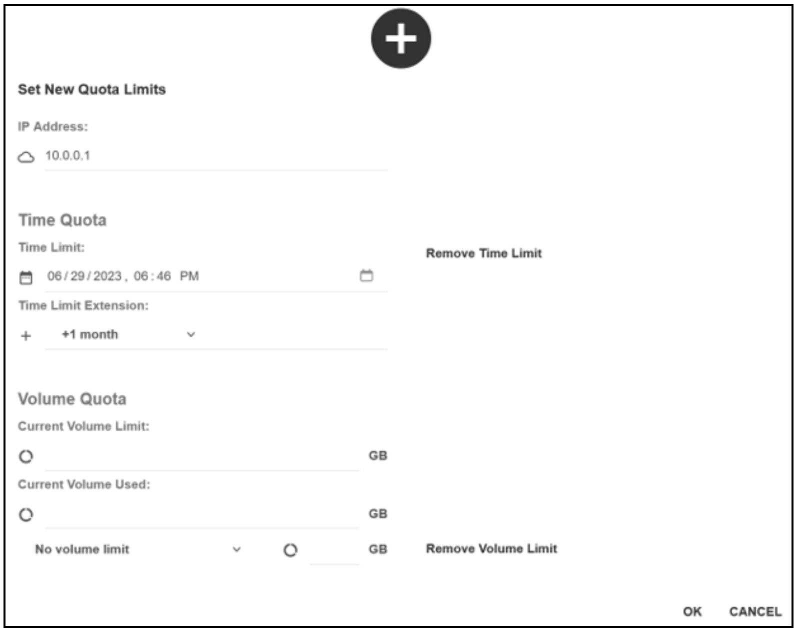

Set New Quota Limits

IP Address:

10.0.0.1

Time Quota

Time Limit:

06 / 29 / 2023 , 06 : 46 PM

Remove Time Limit

Time Limit Extension:

+1 month

Volume Quota

Current Volume Limit:

GB

Current Volume Used:

GB

No volume limit GB

Remove Volume Limit

To define a volume quota, first define an absolute value (20 GB in the following example).

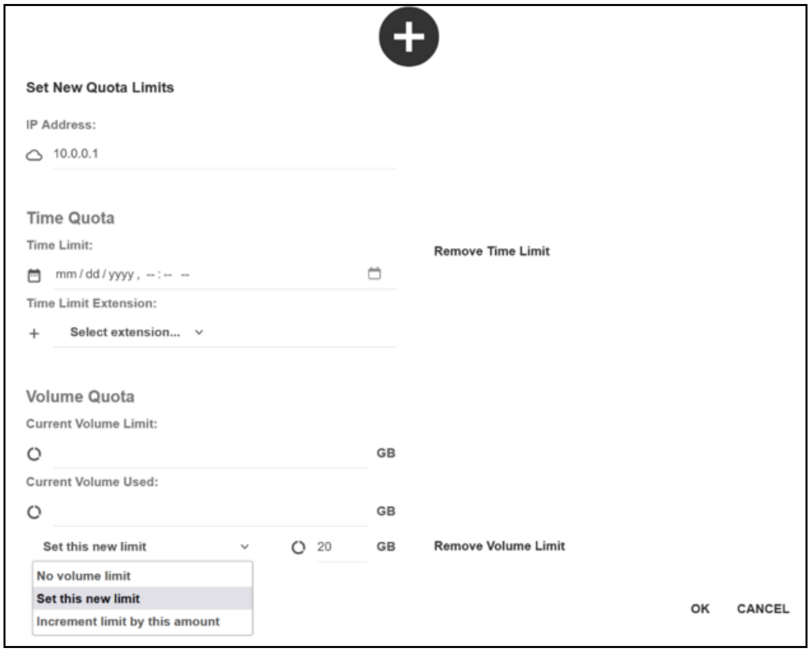

Set New Quota Limits

IP Address:

10.0.0.1

Time Quota

Time Limit:

mm / dd / yyyy , -- : -- --

Time Limit Extension:

Select extension...

Remove Time Limit

Volume Quota

Current Volume Limit:

GB

Current Volume Used:

GB

Set this new limit 20 GB

Remove Volume Limit

No volume limit
Set this new limit
Increment limit by this amount

After the volume quota is created, it can be extended editing the quota and using the option **Increment limit by this amount** (5 GB in the following example).

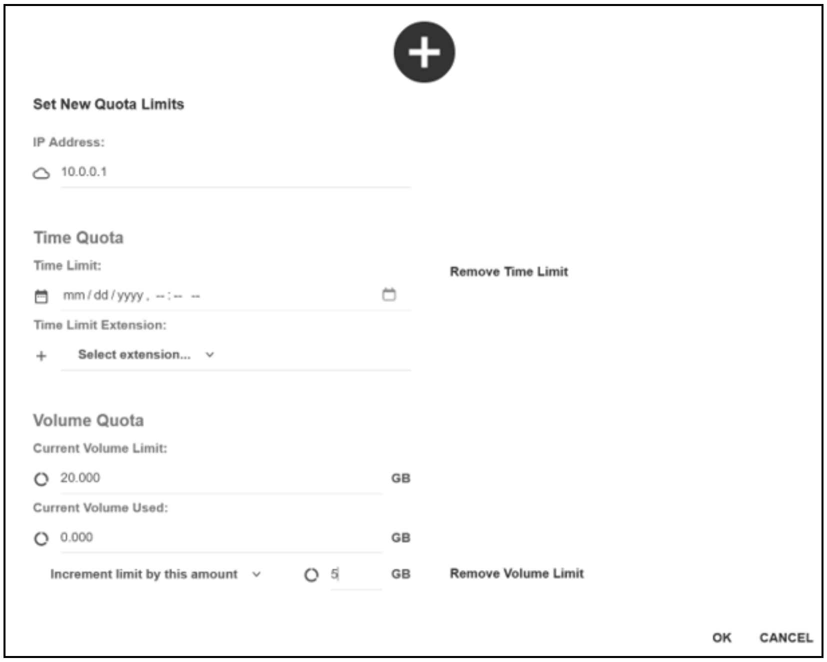

Set New Quota Limits

IP Address:

10.0.0.1

Time Quota

Time Limit:

mm / dd / yyyy , -- : -- --

Time Limit Extension:

Select extension...

Remove Time Limit

Volume Quota

Current Volume Limit:

20.000 GB

Current Volume Used:

0.000 GB

Increment limit by this amount 5 GB

Remove Volume Limit

Time and volume quotas can co-exist, and the subscriber traffic is restricted if anyone becomes exhausted. You can edit the quota and remove the time or volume component by **Remove Time Limit** or **Remove Volume Limit** options respectively.

Checking the state of quota

To view the quota, navigate to **Status** > **Subscribers** > **Subscriber Quotas**. The volume quotas displays the consumption. Figure 85 shows the Subscriber quotas.

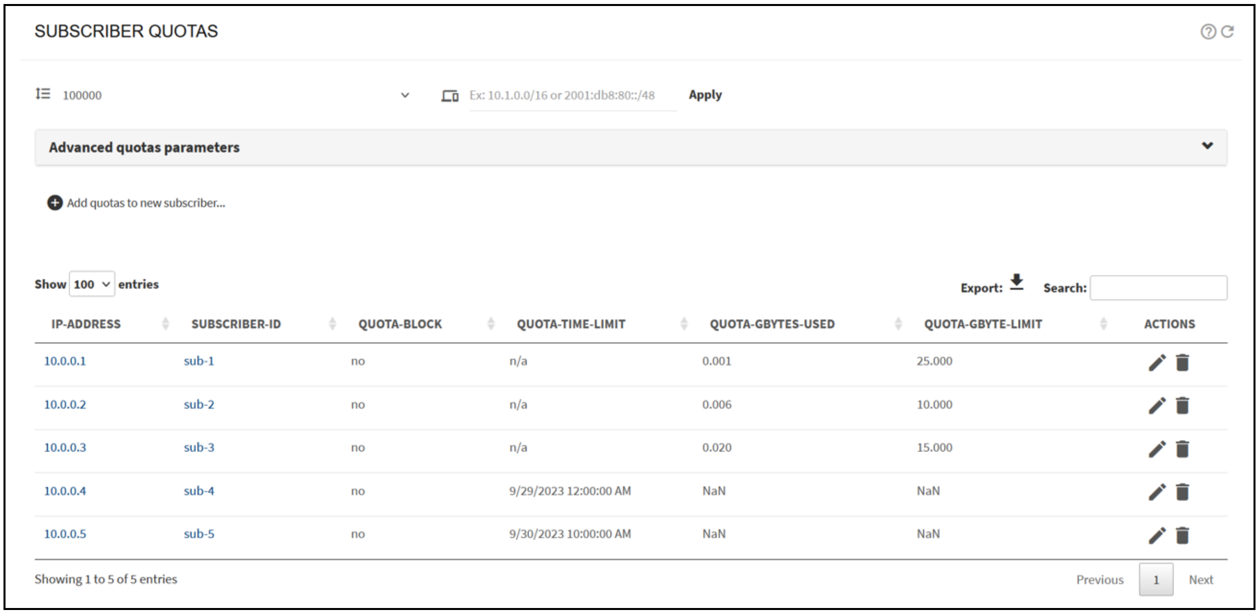
SUBSCRIBER QUOTAS

100000 | Ex: 10.1.0.0/16 or 2001:db8:80::/48 | Apply

Advanced quotas parameters

Add quotas to new subscriber...

Show 100 entries | Export: | Search:

IP-ADDRESS	SUBSCRIBER-ID	QUOTA-BLOCK	QUOTA-TIME-LIMIT	QUOTA-GBYTES-USED	QUOTA-GBYTE-LIMIT	ACTIONS
10.0.0.1	sub-1	no	n/a	0.001	25.000	
10.0.0.2	sub-2	no	n/a	0.006	10.000	
10.0.0.3	sub-3	no	n/a	0.020	15.000	
10.0.0.4	sub-4	no	9/29/2023 12:00:00 AM	NaN	NaN	
10.0.0.5	sub-5	no	9/30/2023 10:00:00 AM	NaN	NaN	

Showing 1 to 5 of 5 entries

Figure 85: *Subscriber quotas*

In Figure 85 , there are three volume quotas, and two time quotas (note that month is given before day, so 9/29/2023 is 29th of September 2023). For volume quotas, the volume already consumed is also shown (for example, 10.0.0.3 has a quota of 15 GB and it has consumed 20 MB).

Slow down when quota exhausted

By default, the traffic is completely blocked when the quota is exhausted, but it is possible to limit the traffic to a slow speed while the quota is not topped up again.

Figure 86 changes the flow-default policy, so it slows down traffic when the quota is exhausted.

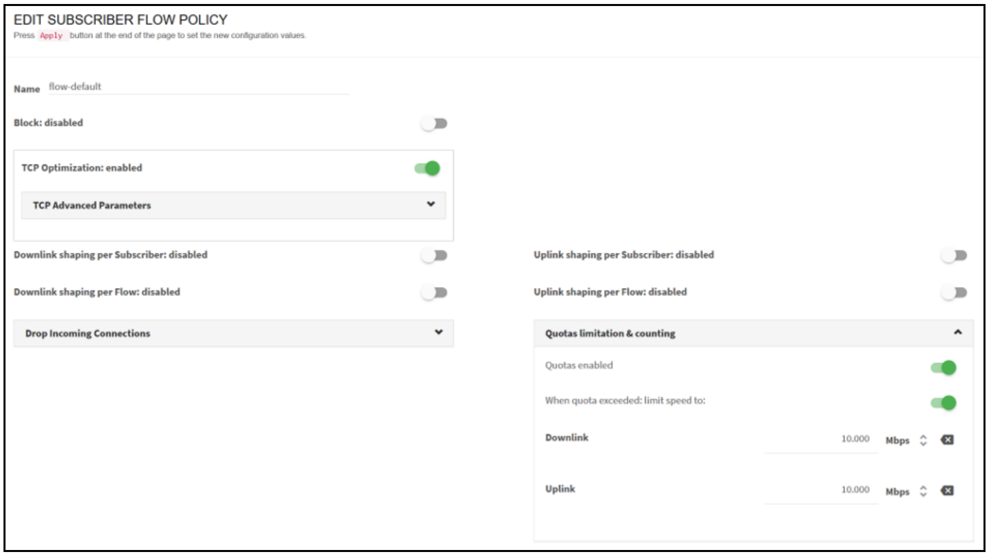
EDIT SUBSCRIBER FLOW POLICY

Press `Apply` button at the end of the page to set the new configuration values.

Name: flow-default

Block: disabled

TCP Optimization: enabled

TCP Advanced Parameters

Downlink shaping per Subscriber: disabled

Downlink shaping per Flow: disabled

Drop Incoming Connections

Uplink shaping per Subscriber: disabled

Uplink shaping per Flow: disabled

Quotas limitation & counting

Quotas enabled

When quota exceeded: limit speed to:

Downlink 10.000 Mbps

Uplink 10.000 Mbps

Figure 86: *The Edit subscriber flow policy page*

Captive portal policy

The quota general configuration section describes the redirection definition to a captive portal if the quota is exhausted. The captive portal implementation requires that traffic to it is not subject to the quota. This is implemented using flow policies. Figure 87 shows the subscriber flows configuration page.

In the following example, two traffic categories are to be out of the quota control (policy flow-no-quota):

- Traffic going to the captive portal.
- Traffic to some specific DNS servers (use to resolve the captive portal URL).

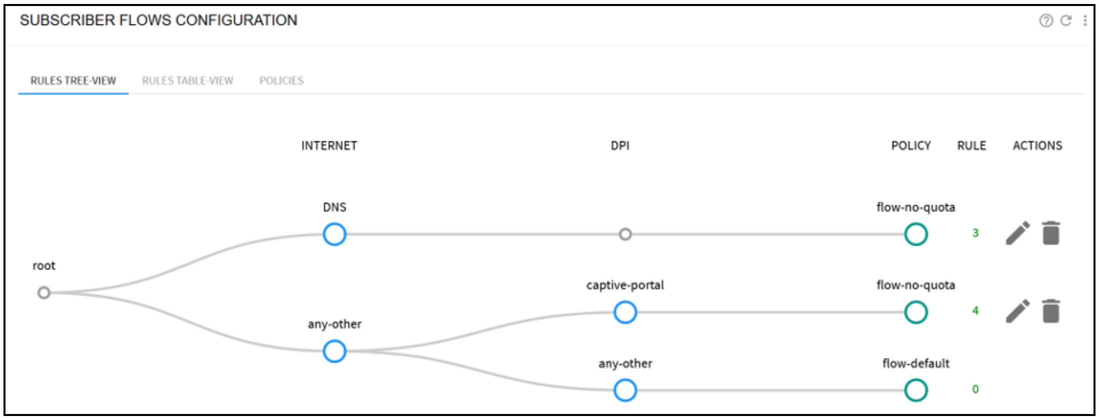

SUBSCRIBER FLOWS CONFIGURATION

INTERNET	DPI	POLICY	RULE	ACTIONS
DNS		flow-no-quota	3	
any-other	captive-portal	flow-no-quota	4	
	any-other	flow-default	0	

Figure 87: *The subscriber flows configuration page*

The policy that is not affected by quota exhaustion has the quota switch set to off: Figure 88 shows the Edit subscriber flow policy page.

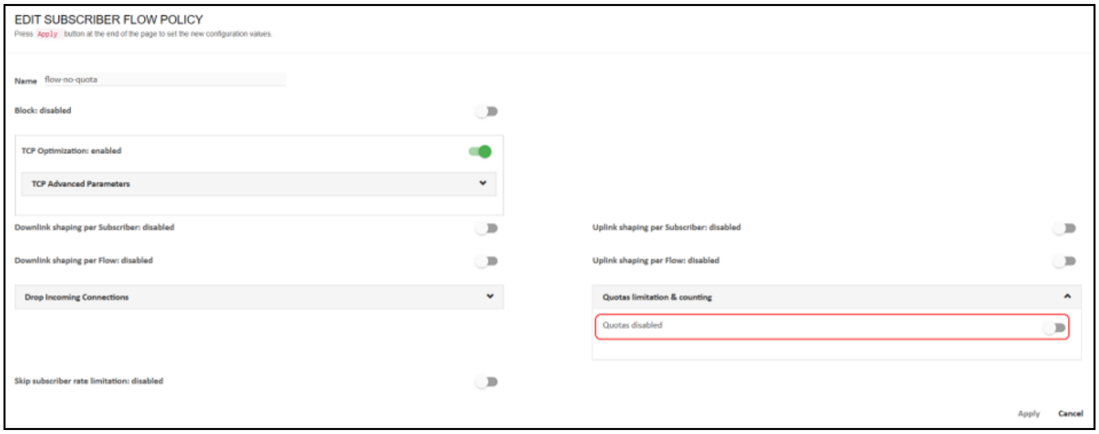

EDIT SUBSCRIBER FLOW POLICY

Press `Apply` button at the end of the page to set the new configuration values.

Name: flow-no-quota

Block: disabled

TCP Optimization: enabled

TCP Advanced Parameters

Downlink shaping per Subscriber: disabled

Downlink shaping per Flow: disabled

Drop Incoming Connections

Skip subscriber rate limitation: disabled

Uplink shaping per Subscriber: disabled

Uplink shaping per Flow: disabled

Quotas limitation & counting

Quotas disabled

Figure 88: *The Edit subscriber flow policy page*

Quotas managed by REST API

The QoE REST API can be also used to manage time and volume quotas. For more information, refer to *QoE Applicance REST API Guide.*

Time quota

There are two ways to define a time quota:

- **As an absolute time**: as POSIX time, defined as the number of seconds elapsed since midnight Coordinated Universal Time (UTC) of January 1, 1970. For example, 1672531200 is UTC Sunday, 1 January 2023 0:00:00. Absolute time is UTC, so convert your local time to UTC when setting the quota.

- **As seconds relative to current time**: for example, a 3600 second quota will be exhausted an hour from now.

To enable a time quota of one hour, the following command is used:

```
curl -k -u myuser:mypassword -X POST
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"quota": {"timeRemaining": 3600}}'
```

To extend the quota to two hours from now, the following command is used:

```
curl -k -u myuser:mypassword -X PUT
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"quota": {"timeRemaining": 7200}}'
```

To remove the quota, so the subscriber is no longer subject to a time quota, the following command is used:

```
curl -k -u myuser:mypassword -X PUT
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"quota": {"time": null}}'
```

Volume quota

A volume quota grants access for a number of Kilo bytes of the traffic. The QoE convention is that 1 Kilo byte is 1000 bytes.

To enable a 1GB volume quota, the following command is used:

```
curl -k -u myuser:mypassword -X POST
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"quota": {"volume": 1000000}}'
```

To extend the quota adding 500 MB, the following command is used:

```
curl -k -u myuser:mypassword -X PUT
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"quota": {"volumeIncrement": 500000}}'
```

To remove the quota, so the subscriber is no longer subject to a volume quota, the following command is used:

```
curl -k -u myuser:mypassword -X PUT
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"quota": {"volume": null}}'
```

Volume and time quotas

To enable a 1GB volume quota and 1 month, the following command is used:

```
curl -k -u myuser:mypassword -X POST
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"quota": {"volume": 1000000, "timeRemaining":
2678400}}'
```

The following command is used to extend the volume quota in 500 MB that is keeping the time quota unchanged:

```
curl -k -u myuser:mypassword -X PUT
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"quota": {"volumeIncrement": 500000}}'
```

The following command is used to remove both quotas, so the subscriber is no longer subject to them:

```
curl -k -u myuser:mypassword -X PUT
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"quota": {"volume": null, "time": null}}'
```

Checking the state of quota

The following command is used to check the quota state through the REST API:

```
curl -k -u myuser:mypassword -X GET
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35

{

    "subscriberIp": "10.0.0.35",

    "quota" : {

    "volume" : 1000000000,

    "volumeConsumed" : 647474875

    "time" : 1676628377,

    "timeRemaining" : 5364849

},

"policyRate" : ""

}
```

Chapter 11: Subscriber Groups

A subscriber group is a set of subscribers. A subscriber can be associated to a group by an IP address or a subscriber ID. A subscriber can belong to up to eight different subscriber groups.

The QoE can display the network metrics per subscriber group. Depending on the groups definition, the metrics allow analysis per access points, network subnets or any customer categorization that can be mapped into groups.

This chapter contains the following sections:

- Creating the subscriber groups
- Viewing subscriber group metrics
- Managing the subscriber groups from REST API

Creating the subscriber groups

To create a subscriber group, navigate go to **Status** > **Subscribers** > **Subscriber Groups.** It can also be accessed from **Configuration > Subscriber Groups**.

Click **Add Subscriber Group ...**.

A subscriber group name must be without spaces.

Subscriber members can be defined by:

- IP address, clicking on **Add IP address...**
- Subscriber ID, clicking on **Add Subscriber ID...**

IP addresses and/or IDs can be added loading a textual file, with one entry per line. File loading options are in the menu in the upper right.

Examples:

- The menu option **Merge IP Address from File...** will add the IP addresses contained in the file to those already existing in the group.
- The menu option **Replace Subscriber IDs with File...** will replace current set of subscriber IDs by those contained in the loaded file.

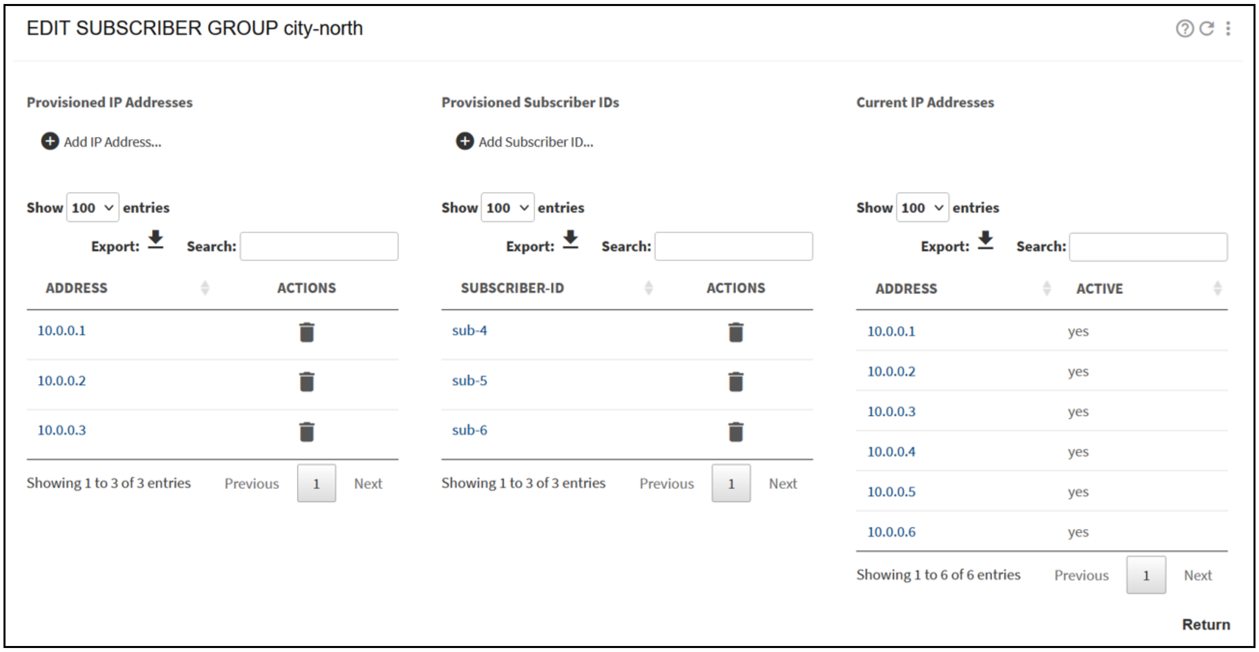

EDIT SUBSCRIBER GROUP city-north

Provisioned IP Addresses

ADDRESS	ACTIONS
10.0.0.1	🗑
10.0.0.2	🗑
10.0.0.3	🗑

Showing 1 to 3 of 3 entries

Provisioned Subscriber IDs

SUBSCRIBER-ID	ACTIONS
sub-4	🗑
sub-5	🗑
sub-6	🗑

Showing 1 to 3 of 3 entries

Current IP Addresses

ADDRESS	ACTIVE
10.0.0.1	yes
10.0.0.2	yes
10.0.0.3	yes
10.0.0.4	yes
10.0.0.5	yes
10.0.0.6	yes

Showing 1 to 6 of 6 entries

Figure 89: *The Edit subscriber group page*

The **Current IP Addresses** table on the right shows all the IP addresses part of the group, either because they are directly added or because they are associated to a subscriber ID part of the group. The table also indicates which IP addresses have traffic (Active yes).

To view the created subscriber groups navigate to **Status** > **Subscribers** > **Subscriber Groups**. Click **Edit** to edit the table.

Viewing subscriber group metrics

To view the subscriber groups, navigate to **Status** > **Subscribers** > **Subscriber Groups**. A table in Figure 90 displays the defined subscriber groups and their metrics.

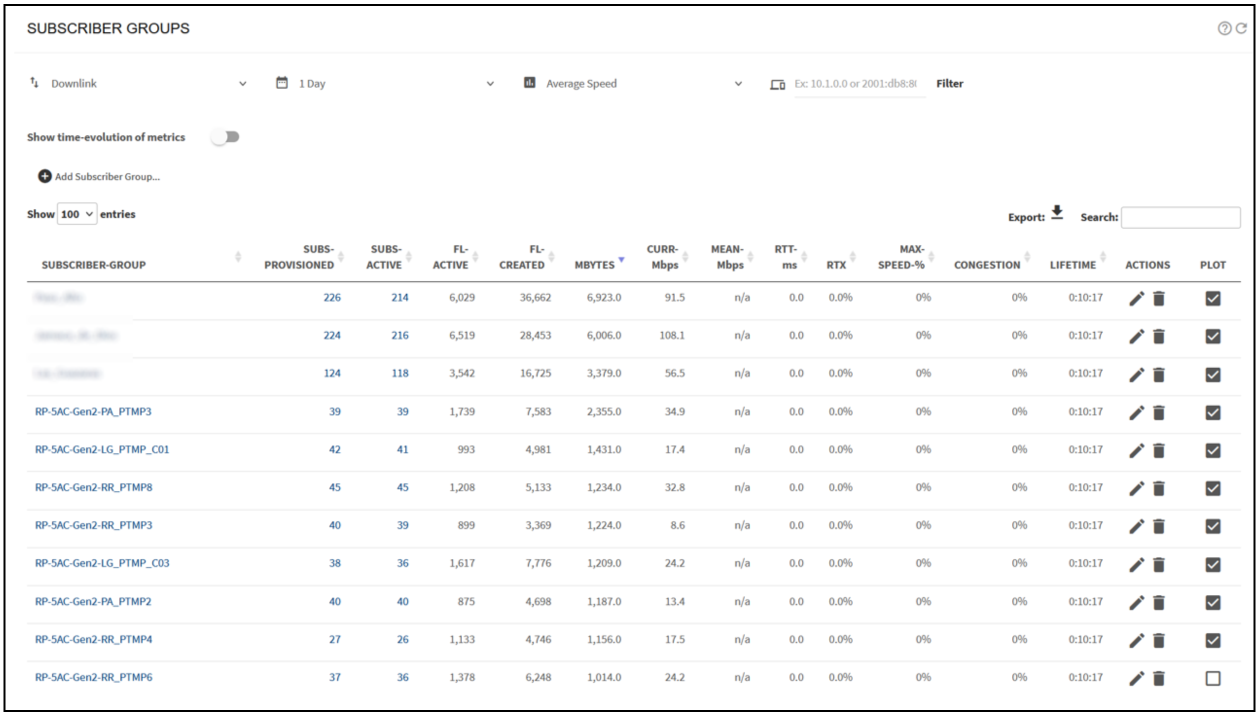

SUBSCRIBER GROUPS

Downlink | 1 Day | Average Speed | Ex: 10.1.0.0 or 2001:db8:8(| Filter

Show time-evolution of metrics

Add Subscriber Group...

Show 100 entries

Export: Search:

SUBSCRIBER-GROUP	SUBS-PROVISIONED	SUBS-ACTIVE	FL-ACTIVE	FL-CREATED	MBYTES	CURR-Mbps	MEAN-Mbps	RTT-ms	RTX	MAX-SPEED-%	CONGESTION	LIFETIME	ACTIONS	PLOT
[redacted]	226	214	6,029	36,662	6,923.0	91.5	n/a	0.0	0.0%	0%	0%	0:10:17	✎ 🗑	☑
[redacted]	224	216	6,519	28,453	6,006.0	108.1	n/a	0.0	0.0%	0%	0%	0:10:17	✎ 🗑	☑
[redacted]	124	118	3,542	16,725	3,379.0	56.5	n/a	0.0	0.0%	0%	0%	0:10:17	✎ 🗑	☑
RP-5AC-Gen2-PA_PTMP3	39	39	1,739	7,583	2,355.0	34.9	n/a	0.0	0.0%	0%	0%	0:10:17	✎ 🗑	☑
RP-5AC-Gen2-LG_PTMP_C01	42	41	993	4,981	1,431.0	17.4	n/a	0.0	0.0%	0%	0%	0:10:17	✎ 🗑	☑
RP-5AC-Gen2-RR_PTMP8	45	45	1,208	5,133	1,234.0	32.8	n/a	0.0	0.0%	0%	0%	0:10:17	✎ 🗑	☑
RP-5AC-Gen2-RR_PTMP3	40	39	899	3,369	1,224.0	8.6	n/a	0.0	0.0%	0%	0%	0:10:17	✎ 🗑	☑
RP-5AC-Gen2-LG_PTMP_C03	38	36	1,617	7,776	1,209.0	24.2	n/a	0.0	0.0%	0%	0%	0:10:17	✎ 🗑	☑
RP-5AC-Gen2-PA_PTMP2	40	40	875	4,698	1,187.0	13.4	n/a	0.0	0.0%	0%	0%	0:10:17	✎ 🗑	☑
RP-5AC-Gen2-RR_PTMP4	27	26	1,133	4,746	1,156.0	17.5	n/a	0.0	0.0%	0%	0%	0:10:17	✎ 🗑	☑
RP-5AC-Gen2-RR_PTMP6	37	36	1,378	6,248	1,014.0	24.2	n/a	0.0	0.0%	0%	0%	0:10:17	✎ 🗑	□

Figure 90: *Subscriber groups and their metrics*

To see the charts with evolution over time of the metrics of the main groups, set the switch **Show time-evolution of metrics** to On as shown in Figure 91.

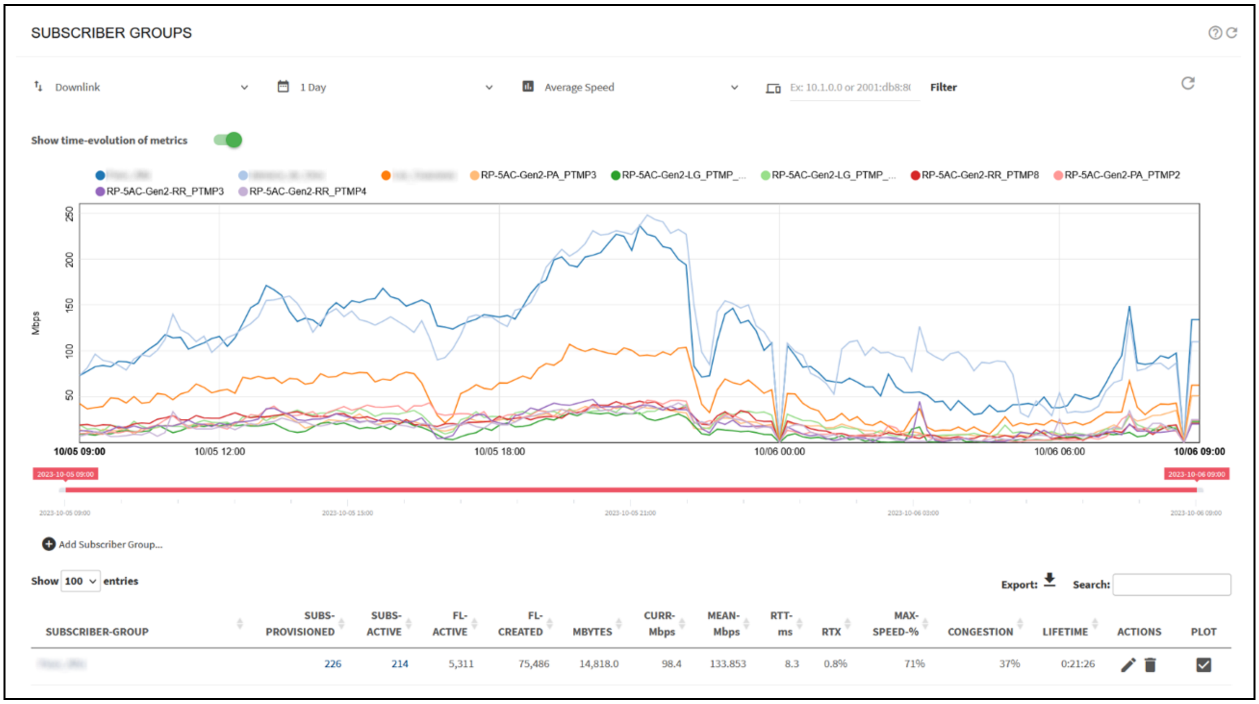

SUBSCRIBER GROUPS

Downlink | 1 Day | Average Speed | Ex: 10.1.0.0 or 2001:db8:8(| Filter

Show time-evolution of metrics

Add Subscriber Group...

Show 100 entries

Export: Search:

SUBSCRIBER-GROUP	SUBS-PROVISIONED	SUBS-ACTIVE	FL-ACTIVE	FL-CREATED	MBYTES	CURR-Mbps	MEAN-Mbps	RTT-ms	RTX	MAX-SPEED-%	CONGESTION	LIFETIME	ACTIONS	PLOT
[redacted]	226	214	5,311	75,486	14,818.0	98.4	133.853	8.3	0.8%	71%	37%	0:21:26	✎ 🗑	☑

Figure 91: *The Subscriber groups page*

The chart will show by default the first 10 subscriber groups. The **PLOT** column in the table indicates the groups included in the chart. Up to 30 subscriber groups can be shown at the same time selecting their plot tick box and refreshing the chart (reload icon in the upper right of the chart). Unselecting the plot tick will remove that group from the chart.

Metrics are for the downlink or uplink directions, and for a time period of three months.

The available metrics are:

- Average speed
- Active flows
- Flows created per minute
- Latency
- Retransmission
- Congestion
- Percentage of traffic going at maximum speed.

Those metrics are similar to the ones shown for a subscriber (for more information, see Subscriber dashboard section).

By default, up to 30 groups are displayed in the chart. You can decide the groups to include in the graph by checking or unchecking the **PLOT** option for each group.

It is possible to use a filter, to select only the groups that have as member an IP address or a subscriber ID:

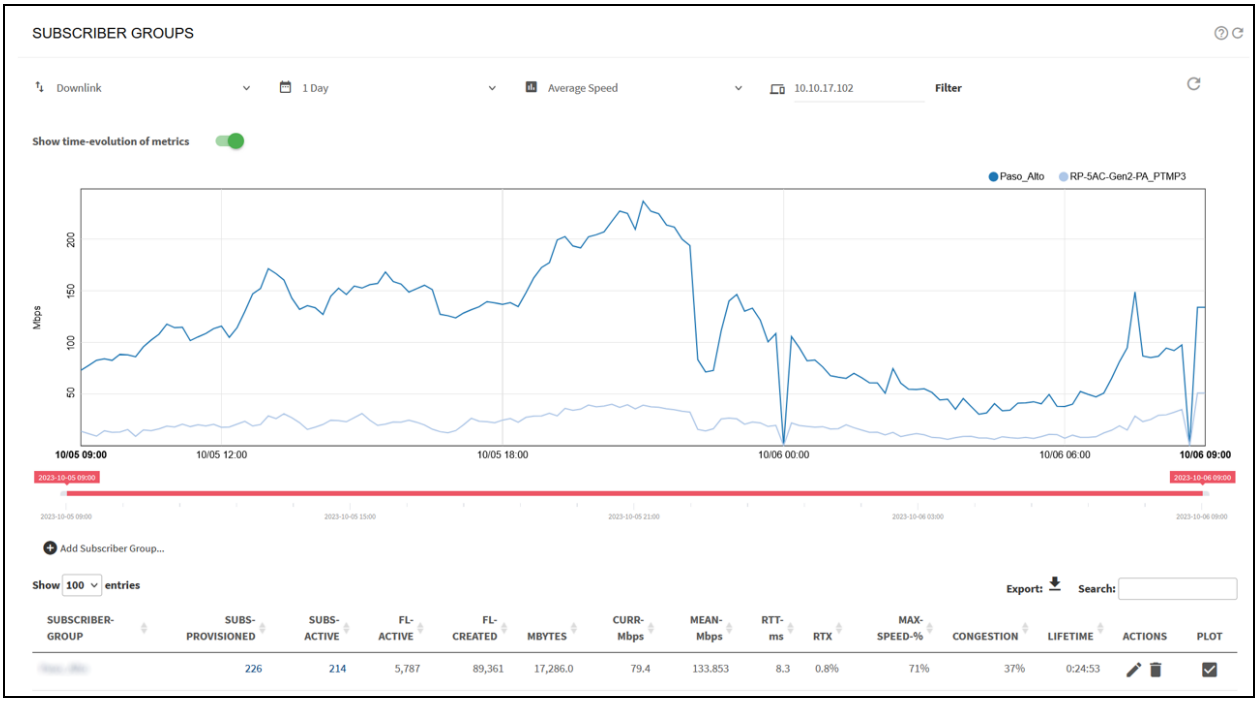

SUBSCRIBER GROUPS

Downlink · 1 Day · Average Speed · 10.10.17.102 · Filter

Show time-evolution of metrics

Show 100 entries

SUBSCRIBER-GROUP	SUBS-PROVISIONED	SUBS-ACTIVE	FL-ACTIVE	FL-CREATED	MBYTES	CURR-Mbps	MEAN-Mbps	RTT-ms	RTX	MAX-SPEED-%	CONGESTION	LIFETIME	ACTIONS	PLOT
[redacted]	226	214	5,787	89,361	17,286.0	79.4	133.853	8.3	0.8%	71%	37%	0:24:53	✎ 🗑	☑

Figure 92: *The Subscriber groups page*

To view the subscriber group dashboard, click subscriber group name with the metrics in separate charts. Figure 93 shows the subscriber group dashboard.

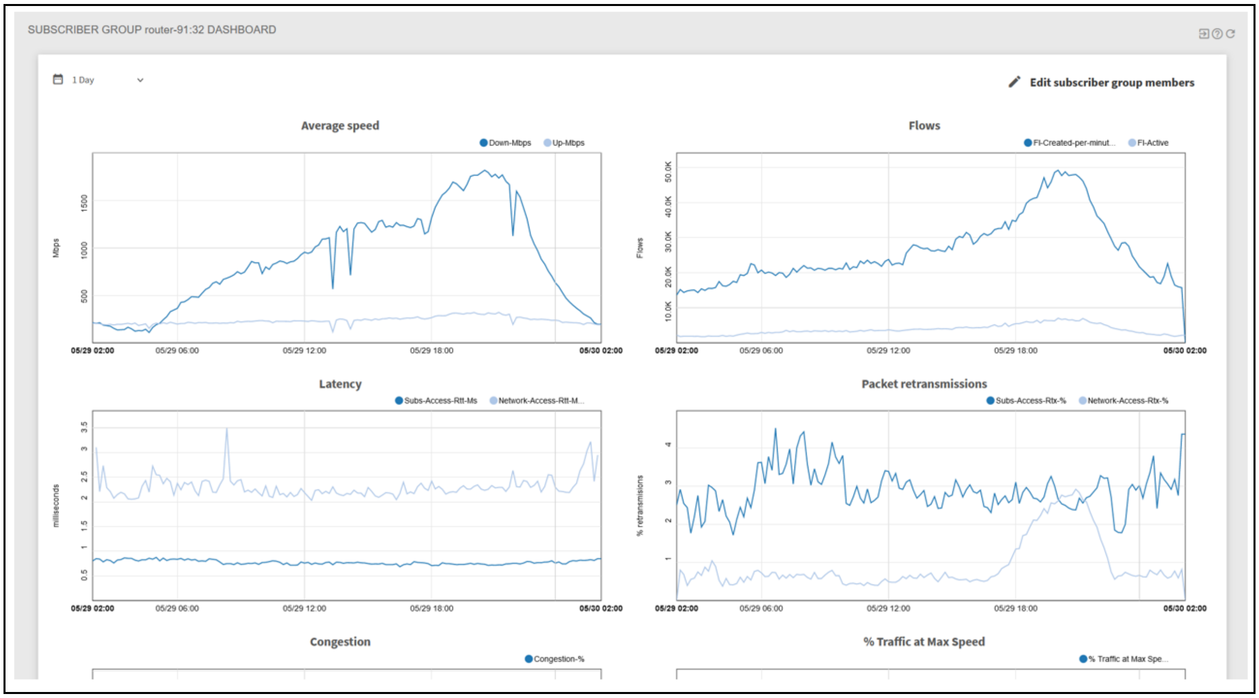

Figure 93: *Subscriber group dashboard*

It is also possible to navigate to this page from **Statistics** > **Subscribers** > **Subscribers Groups**.

Managing the subscriber groups from REST API

QoE REST API is also used to manage subscriber groups.

To add an existing IP 10.0.0.35 to a subscriber group “city-north”, the following command is used:

```
curl -k -u myuser:mypassword -X PUT
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"subscriberGroups": ["city-north”]}'
```

The subscriber group is created automatically if it is not exist. The list of groups should always reflect the full list of membership. For example, to add an existing IP 10.0.0.35 to a subscriber group “vip”, the following command is used:

```
curl -k -u myuser:mypassword -X PUT
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"subscriberGroups": ["city-north”, “vip”]}'
```

The group “city-north” is included, else the IP will be removed from the group. You can remove an IP from a group by omitting the group from the list. Also, to remove the IP from all groups, set an empty list:

```
curl -k -u myuser:mypassword -X PUT
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35 -H "Content-Type:
application/json" --data '{"subscriberGroups": []}'
```

If the IP has group memberships, then they are returned by a GET:

```
curl -k -u myuser:mypassword -X GET
https://192.168.0.121:3443/api/v1/subscribers/10.0.0.35

{

    "subscriberIp" : "10.0.0.35",
    "subscriberId" : "sub-12",

    "subscriberGroups" : [ "city-north", "vips" ],

    "policyRate" : "Plan-200Mbps"

}
```

The membership of a specific subscriber group can be obtained as follows:

```
curl -k -u myuser:mypassword -X GET
https://192.168.0.121:3443/api/v1/subscriberGroups/city-north

{

    "subscriberGroupName" : "city-north",

    "memberSubscriberIps" : [ "10.0.0.35", "10.0.0.15", "10.0.0.25",
    "10.0.0.10" ]

}
```

To get the full list of subscriber groups, the following command is used:

```
curl -k -u myuser:mypassword -X GET
https://192.168.0.121:3443/api/v1/subscriberGroups

{

    "items" : [

      {

         "subscriberGroupName" : "city-north",

         "memberSubscriberIps" : [ "10.0.0.35", "10.0.0.15", "10.0.0.25",
         "10.0.0.10" ]

      },

      {

         "subscriberGroupName" : "vips",

         "memberSubscriberIps" : [ "10.0.0.35", "10.0.0.15", "10.0.0.25",
         "10.0.0.20" ]

      }

    ]

}
```

You can add to a group subscribers by their IDs, sending a PORT with the ID list to the subscriber group end point:

```
curl -k -u myuser:mypassword -X POST
https://192.168.0.121:3443/api/v1/subscriberGroups/city-north -H "Content-Type:
application/json" --data '{memberSubscriberIds: [“sub-1”, “sub-2”, “sub-3”]}’
```

For more information, refer to *QoE Appliance REST API Guide*.

Chapter 12: Billing Systems

In addition to RADIUS and REST, subscriber data can be retrieved from a number of supported billing systems.

> **Notes**
>
> Billing integrations are supported only for IPv4 addresses.

This chapter contains the following sections:

- Azotel
- Gestfy
- ISPSolution
- Microwisp
- Powercode
- Sonar
- Splynx
- UISP
- Visp.net
- WISPControl
- Wisphub
- Wispro

Azotel

QoE retrieves the customers IP, upload rate and download rate from Azotel to enforce speed limits to apply (uploadrate and downloadrate).

To activate Azotel, navigate to **Configuration** > **RADIUS/REST/Billing** > **Billing Systems**, select **Azotel** and enable the switch.

A user and password will also be needed for REST/JSON access to Azotel. That user/password must be created in the Azotel system with allowed access from the QoE IP address. For more information, see Azotel documentation.

The QoE uses its management address for Azotel queries. If the QoE reaches Azotel over the Internet, then Azotel verifies a public IP address and an authorization by the Azotel system.

Provide also the Azotel system IP address or server name and port number (443 by default).

Figure 94 shows the billing systems integrations page.

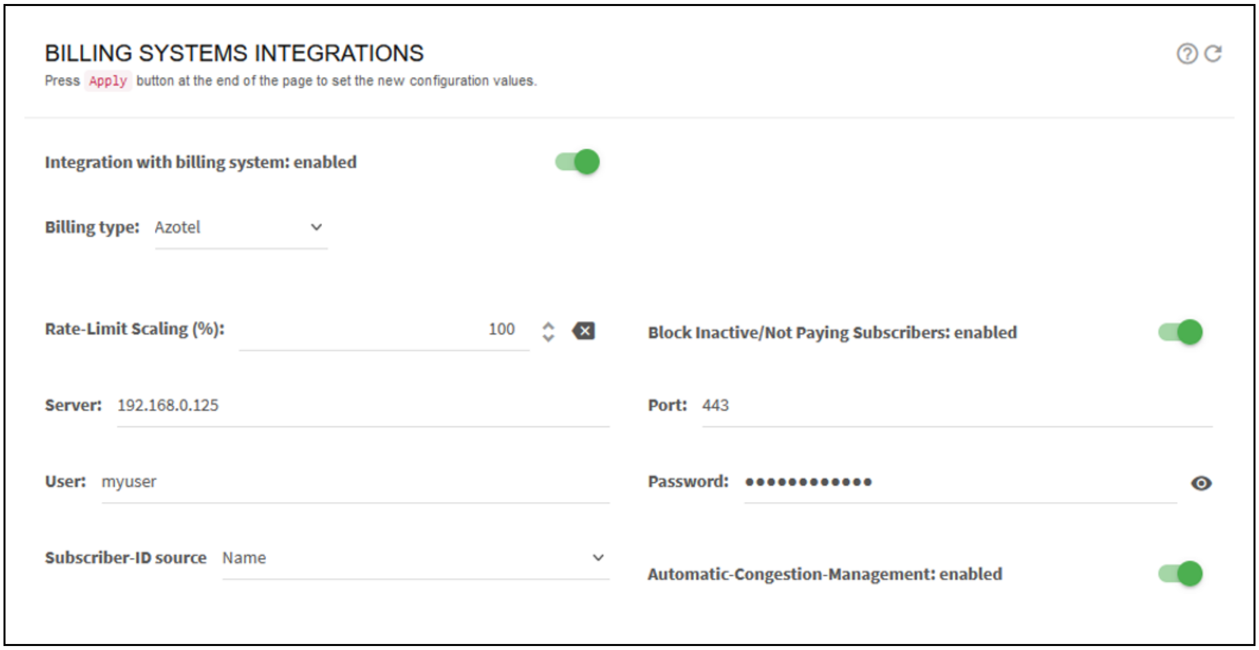

BILLING SYSTEMS INTEGRATIONS

Press Apply button at the end of the page to set the new configuration values.

Integration with billing system: enabled

Billing type: Azotel

Rate-Limit Scaling (%): 100

Block Inactive/Not Paying Subscribers: enabled

Server: 192.168.0.125

Port: 443

User: myuser

Password: ●●●●●●●●●●●●

Subscriber-ID source Name

Automatic-Congestion-Management: enabled

Figure 94: *The billing systems integrations page*

Azotel customers in a status other than **current** is blocked (they are regarded as lacking a valid subscription). You can change this behaviour to non-blocking disabling the switch **Block Inactive/Not Paying Subscribers**.

The following fields in Azotel billing can be used as source of the subscriber ID:

- Customer-ID
- Name
- Nickname

Gestfy

Gestfy uses QoE REST API. For more information on Gestfy, refer to *QoE Appliance REST API Guide*.

ISPSolution

ISPSolution uses QoE REST API. For more information on ISPSolution, refer to *QoE Appliance REST API Guide*.

Microwisp

Mikrowisp uses RADIUS. For more information on Microwisp, refer to *QoE Appliance RADIUS Guide*.

Powercode

QoE retrieves CPE equipment of a certain category (1 by default). For subscribers with that category of equipment, it retrieves the rate limits of their Internet service (InternetInfo).

To activate Powercode, navigate to **Configuration > RADIUS/REST/Billing > Billing Systems**, select **Powercode** and enable the switch. Enter the Powercode system IP address or server name and its SSH port number (22 by default).

The QoE server requires SSH access to the Powercode server using a Unix User/Password. It also requires read access to the MySQL database. If MySQL user/password is different to the Unix user/password, that must be specified under **MySQL Credentials**. The MySQL user must have read access to the following tables of the MySQL database:

- Services
- InternetInfo
- Equipment
- Customer
- CustomerServices
- AddressRange

Figure 95 shows the billing systems integrations page.

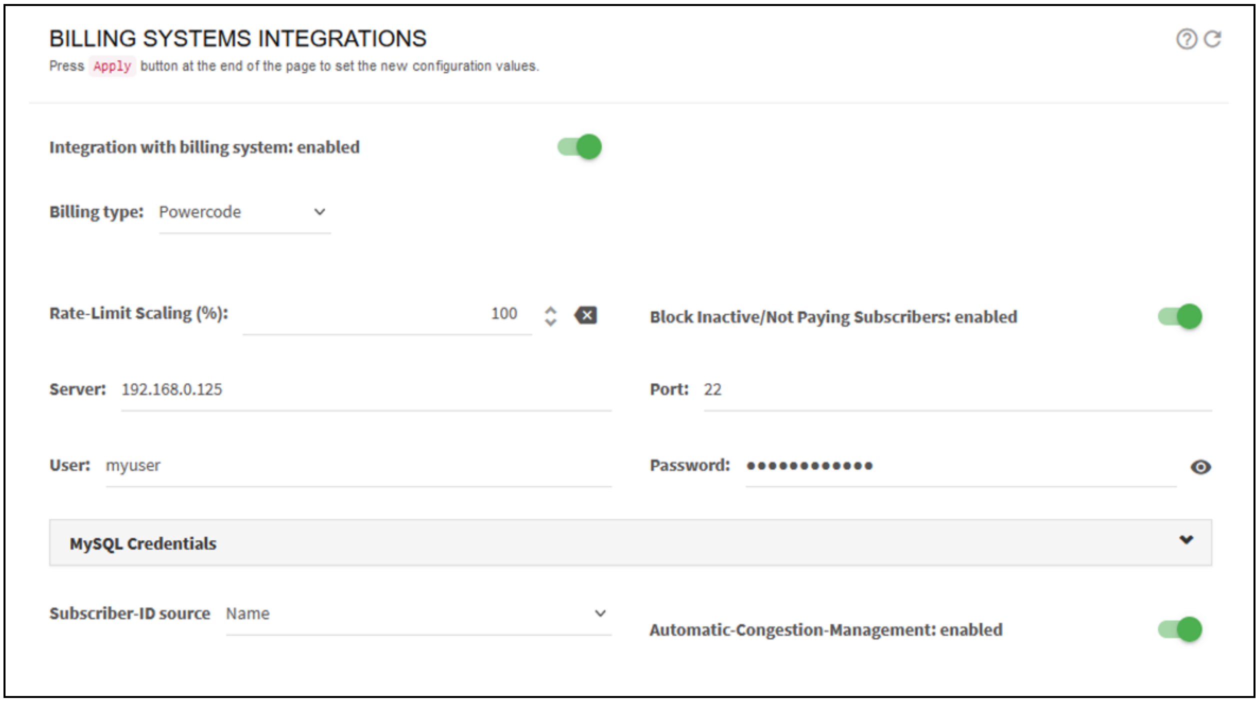

BILLING SYSTEMS INTEGRATIONS

Press Apply button at the end of the page to set the new configuration values.

Integration with billing system: enabled

Billing type: Powercode

Rate-Limit Scaling (%): 100

Block Inactive/Not Paying Subscribers: enabled

Server: 192.168.0.125

Port: 22

User: myuser

Password: ●●●●●●●●●●●●●

MySQL Credentials

Subscriber-ID source Name

Automatic-Congestion-Management: enabled

Figure 95: *The billing systems integrations page*

In status, Powercode customers other than **Active** are blocked (they are regarded as lacking a valid subscription). This behaviour can be changed to non-blocking disabling the switch **Block Inactive/Not Paying Subscribers** (see Blocking of inactive subscribers section).

The following fields in Powercode billing can be used as source of the subscriber ID:

- Customer-ID
- Equipment-ID
- MAC address
- Name

Warning

Do not change the Subscriber-Id source after the first sync with the billing system. If it is changed, then all statistics associated with the old Subscriber-Id are not available on the UI.

REST-API powercode

Powercode billing restricts REST API to three requests per second. So the preferred integration uses the SQL access as described in the previous section. REST-API can be used if SQL is not used and the number of subscribers is low (one thousand or less).

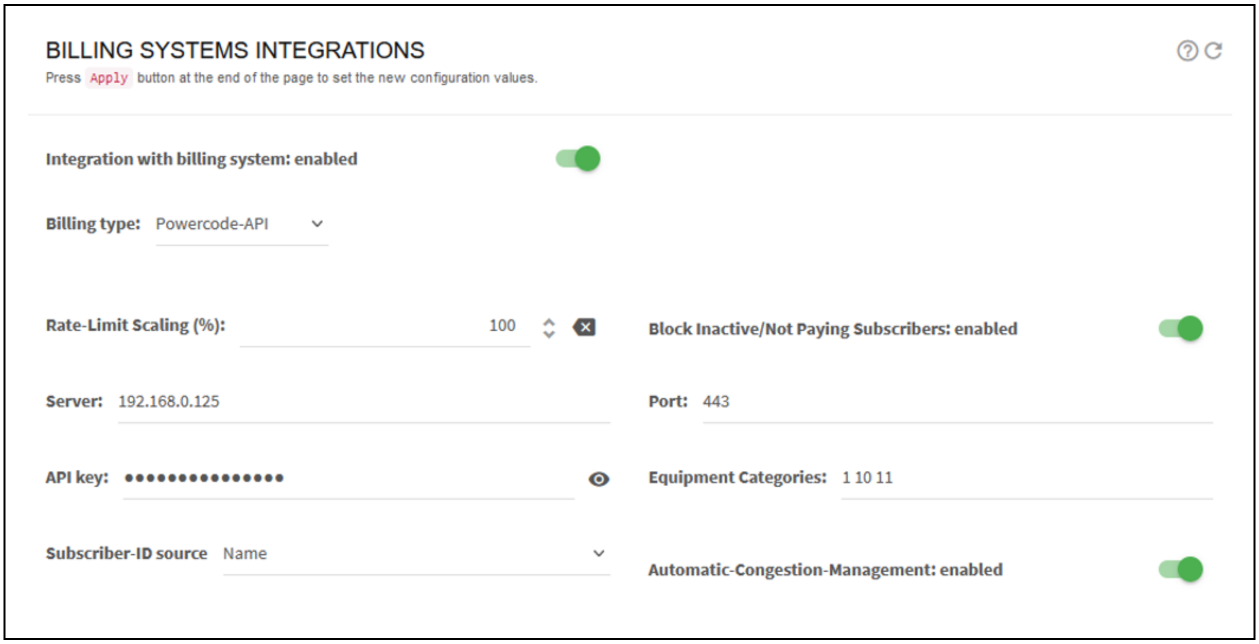

BILLING SYSTEMS INTEGRATIONS

Press `Apply` button at the end of the page to set the new configuration values.

Integration with billing system: enabled

Billing type: Powercode-API

Rate-Limit Scaling (%): 100

Block Inactive/Not Paying Subscribers: enabled

Server: 192.168.0.125

Port: 443

API key: •••••••••••••••

Equipment Categories: 1 10 11

Subscriber-ID source Name

Automatic-Congestion-Management: enabled

Figure 96: *REST-API powercode*

QoE retrieves CPE equipment of a certain category (1 by default). For subscribers with that category of equipment, it retrieves the rate limits of their Internet service (internetInfo). To activate Powercode, navigate to **Configuration** > **RADIUS/REST/Billing** > **Billing Systems**, select **Powercode-API** and enable the switch.

An API Key is required to active the Powercode. The API key must be created in the Powercode system, with allowed access from QoE IP address (QoE uses its management address for Powercode queries). It also provides the Powercode system IP address or server name and port number (444 by default).

If the CPE equipment category in the Powercode database is other than 1, then change it to 1. More than one category can be specified typing the category numbers separated by spaces (for example, **10 11 12** for categories 10, 11 and 12).

Warning

Do not change the Subscriber-Id source after the first sync with the billing system. If it is changed, then all statistics associated with the old Subscriber-Id are not available on the UI.

Sonar

QoE retrieves the customer tariff and get from it the speed limits to apply. To activate Sonar, navigate to **Configuration** > **RADIUS/REST/Billing** > **Billing Systems**, select **Sonar** and enable the switch.

An API Key is required and it must be created in the Sonar system. For creating the API Key, refer to Sonar Knowledge Base article.

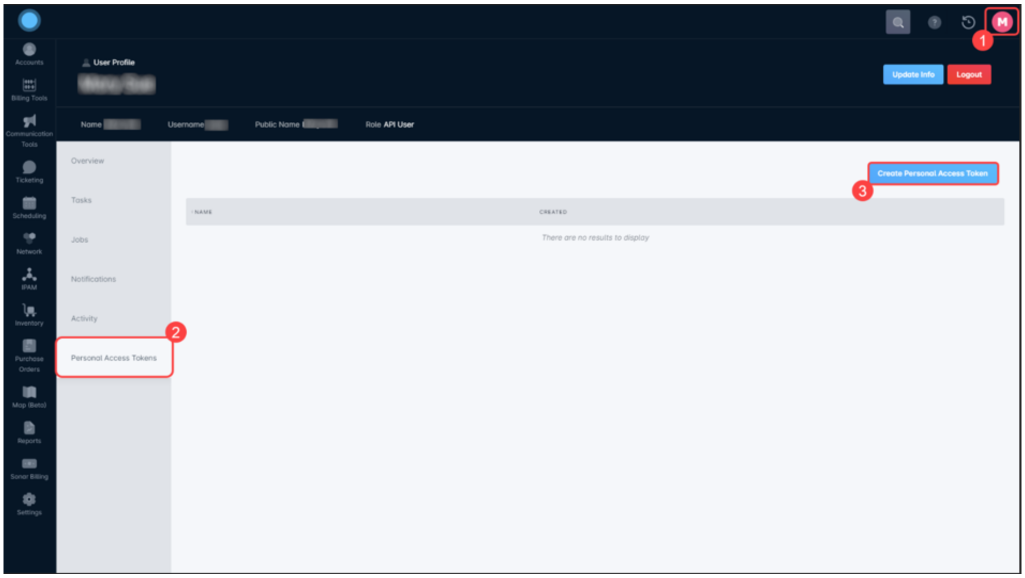

User Profile

[redacted]

Name [redacted] Username [redacted] Public Name [redacted] Role API User

There are no results to display

Figure 97: *Creating an API Key*

Enter the IP address of the Sonar system or server name and port number (443 by default). Figure 98 shows the billing systems integration.

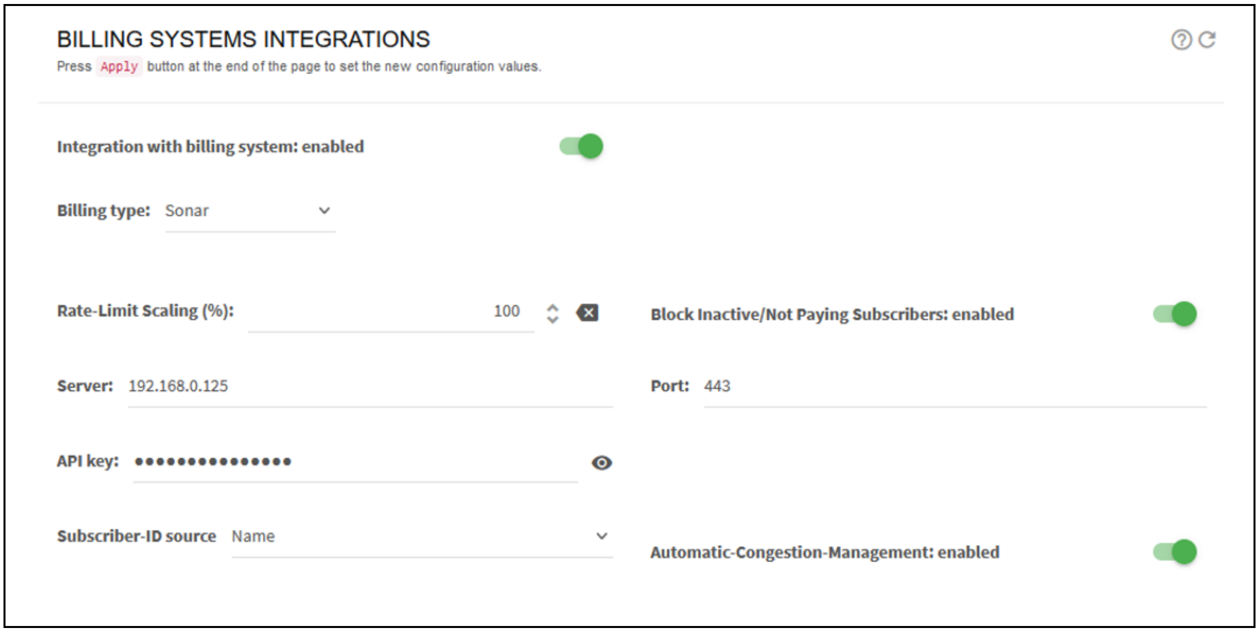

BILLING SYSTEMS INTEGRATIONS

Press `Apply` button at the end of the page to set the new configuration values.

Integration with billing system: enabled

Billing type: Sonar

Rate-Limit Scaling (%): 100

Block Inactive/Not Paying Subscribers: enabled

Server: 192.168.0.125

Port: 443

API key: •••••••••••••••

Subscriber-ID source Name

Automatic-Congestion-Management: enabled

Figure 98: *Billing systems integration*

Sonar customers with **account_status** > **name field** with a value other than **Active** or **Employee**, are blocked (they are regarded as lacking a valid subscription). This behavior can be changed to non-blocking disabling the switch **Block Inactive/Not Paying Subscribers** (see Blocking of inactive subscribers section).

The following fields in Sonar billing can be used as source of the subscriber ID:

- Customer-ID
- Name

Warning

Do not change the Subscriber-Id source after the first sync with the billing system. If it is changed, then all statistics associated with the old Subscriber-Id are not available on the UI.

Splynx

QoE retrieves the customer tariff and get from it the speed limits to apply and the burst rates, thresholds, and duration. To activate Splynx, navigate to **Configuration** > **RADIUS/REST/Billing** > **Billing Systems**, select **Splynx** and enable the switch. shows the Splynx activation.

An API key and secret are required to activate the Splynx. They must be created in the Splynx system with the following settings:

- Enable basic authorization for this key.
- Leave empty **Allowed list for IPs** or include the QoE IP address. QoE uses its management address for Splynx queries if QoE reaches Splynx over the Internet, Splynx notices a public IP address and requires authorization by the Splynx system.
- Add view permissions for database items **Tariff plans** > **Internet and Customers** > **Customers online**.

Figure 99 displays the API KEY and the access permits:

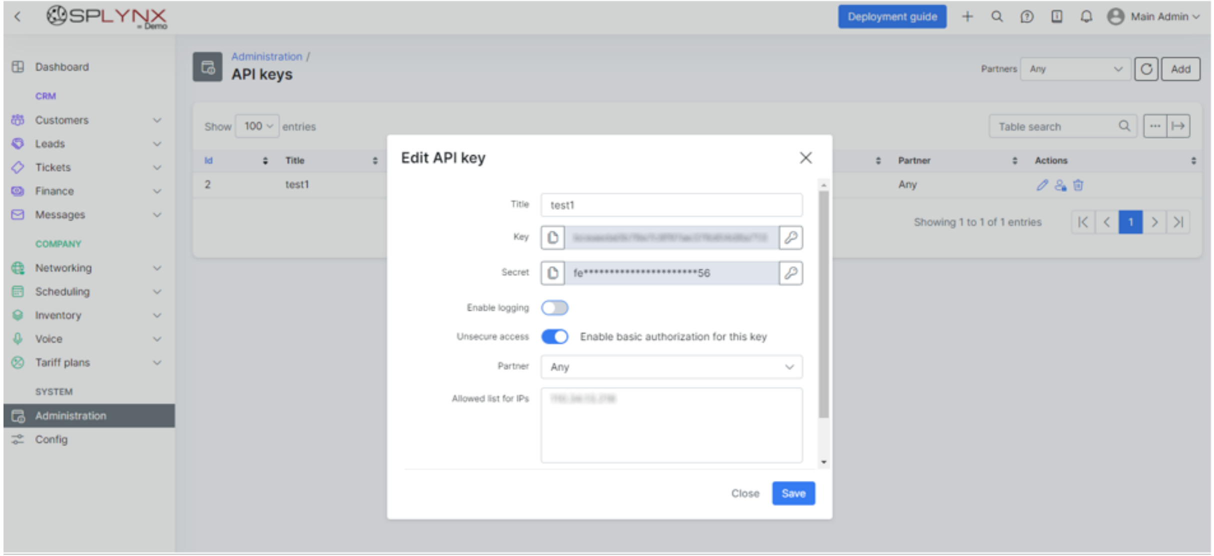

API keys

Show 100 entries

Id	Title	Partner	Actions
2	test1	Any	

Showing 1 to 1 of 1 entries

Edit API key

Title: test1

Key: [redacted]

Secret: fe***********************56

Enable logging: ○

Unsecure access: ● Enable basic authorization for this key

Partner: Any

Allowed list for IPs: [redacted]

Close Save

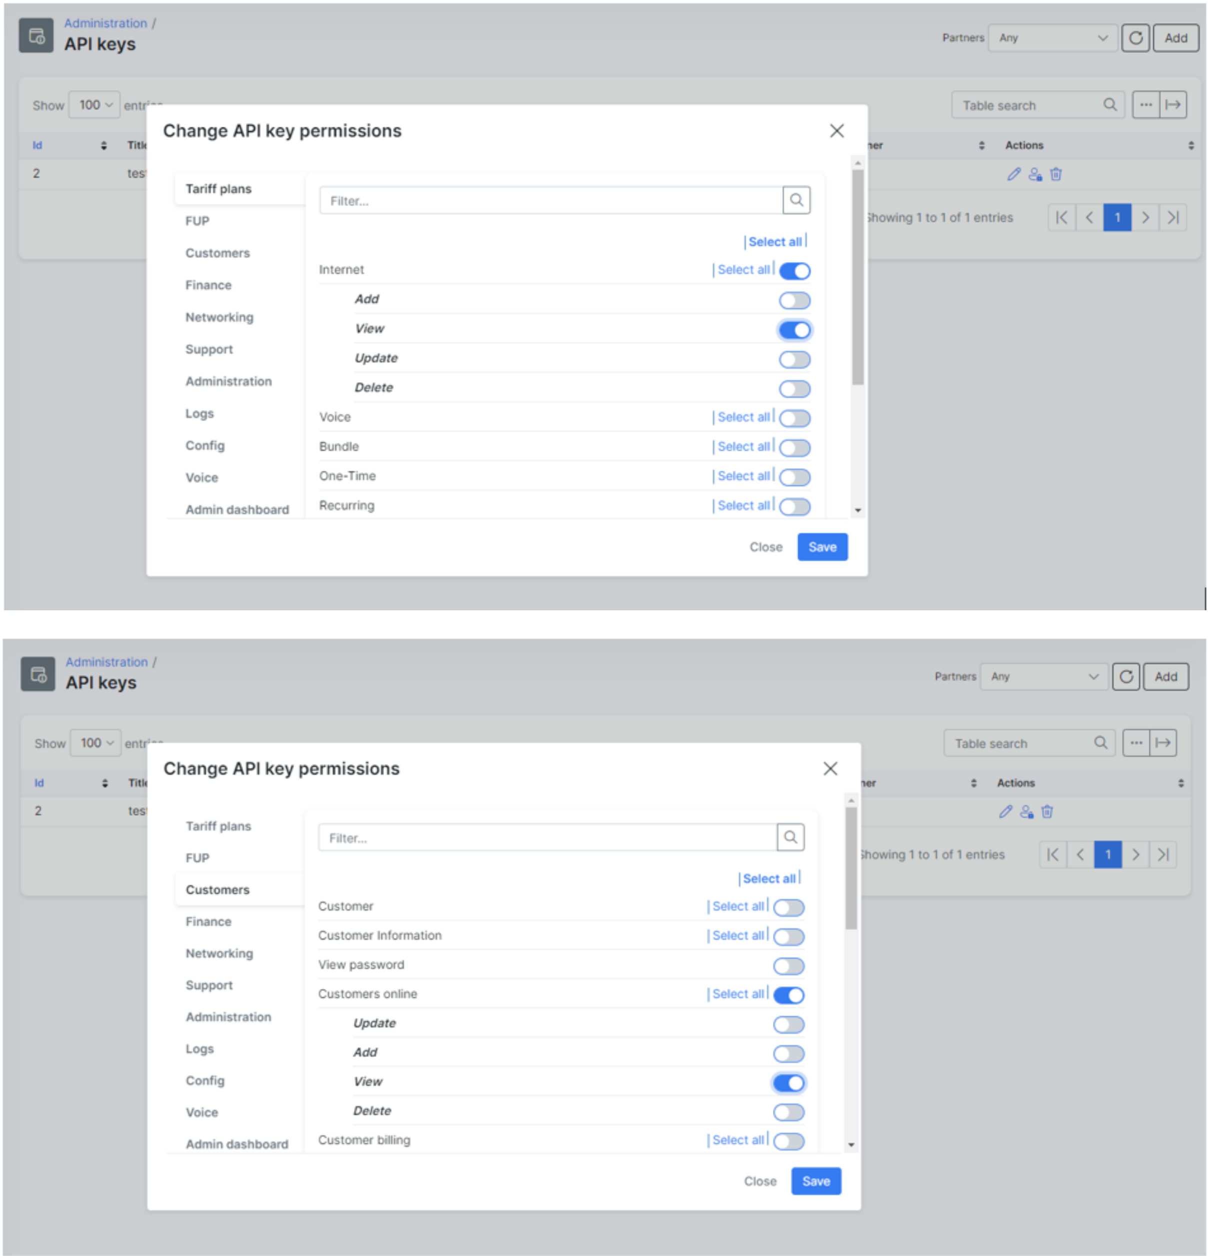

Change API key permissions

Tariff plans

- Filter...
- Select all
- Internet | Select all | ●
  - *Add* ○
  - *View* ●
  - *Update* ○
  - *Delete* ○
- Voice | Select all | ○
- Bundle | Select all | ○
- One-Time | Select all | ○
- Recurring | Select all | ○

Change API key permissions

Customers

- Filter...
- Select all
- Customer | Select all | ○
- Customer Information | Select all | ○
- View password ○
- Customers online | Select all | ●
  - *Update* ○
  - *Add* ○
  - *View* ●
  - *Delete* ○
- Customer billing | Select all | ○

Figure 99: *API KEY and the access permits*

It also provides the Splynx system IP address or server name and port number (443 by default).

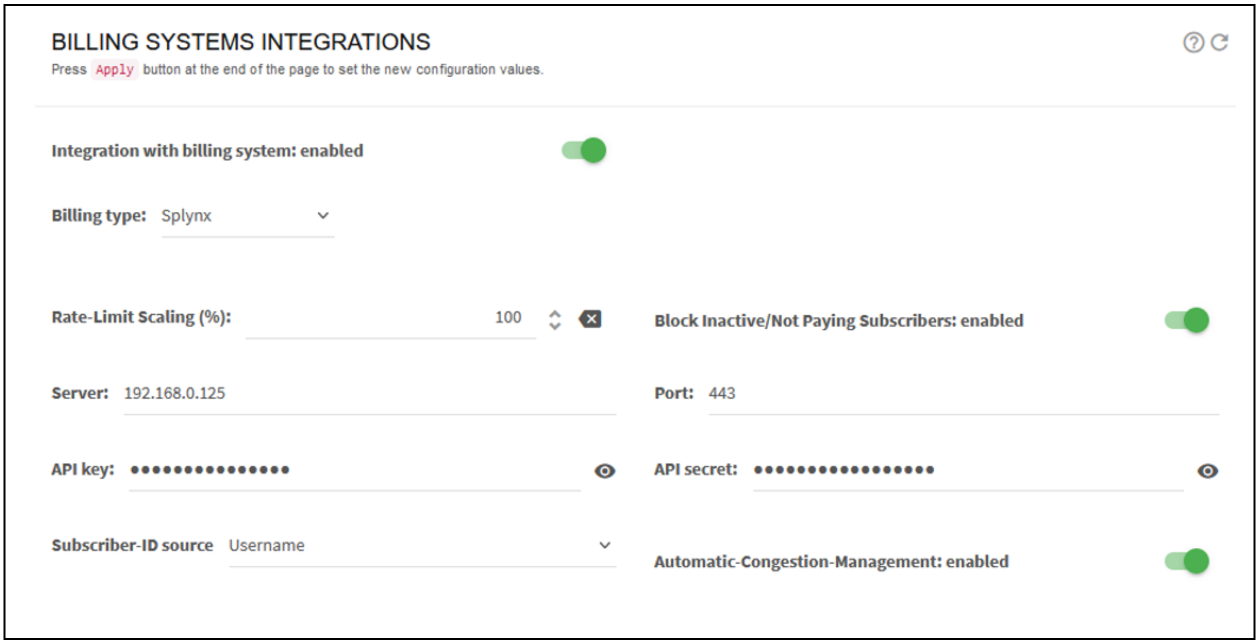

BILLING SYSTEMS INTEGRATIONS

Press Apply button at the end of the page to set the new configuration values.

Integration with billing system: enabled

Billing type: Splynx

Rate-Limit Scaling (%): 100

Block Inactive/Not Paying Subscribers: enabled

Server: 192.168.0.125

Port: 443

API key: •••••••••••••••

API secret: •••••••••••••••••

Subscriber-ID source Username

Automatic-Congestion-Management: enabled

Figure 100: *Splynx*

In a blocked field, Splynx customers with “1” are blocked (they are regarded as lacking a valid subscription). This behavior can be changed to non-blocking disabling the switch with **Block Inactive/Not Paying Subscribers** option (see Blocking of inactive subscribers section).

The following fields in Splynx billing can be used as source of the subscriber ID:

- Customer-ID
- Username
- Login

Warning

Do not change the Subscriber-Id source after the first sync with the billing system. If it is changed, then all statistics associated with the old Subscriber-Id are not available on the UI.

UISP

QoE integrates with UISP using python tool developed by Cambium Networks and is available at the QoE support site for download.

QoE integrates with UISP using python tool developed by Cambium Networks and is available at the Cambium Networks support site for download. For more information, refer to *QoE Appliance REST API Guide*.

Visp.net

QoE retrieves the customer tariff and get it from the speed limits to apply and the burst rates, thresholds and duration.

To activate Visp.net, navigate to **Configuration** > **RADIUS/REST/Billing** > **Billing Systems**, select **Visp** and enable the switch. Figure 101 shows enabling the Visp.net.

A valid client id and secret, user name and password must be provided to QoE to request the temporal API tokens. A client id is unique per Visp installation. A user can be any one of the valid user accounts from the client to access the system.

The provided IP address or server name is used along with the port (443 by default) to request API tokens (*https://<server>:<port>/token*) and also to send API queries (*https://<server>:<port>/graphql*).

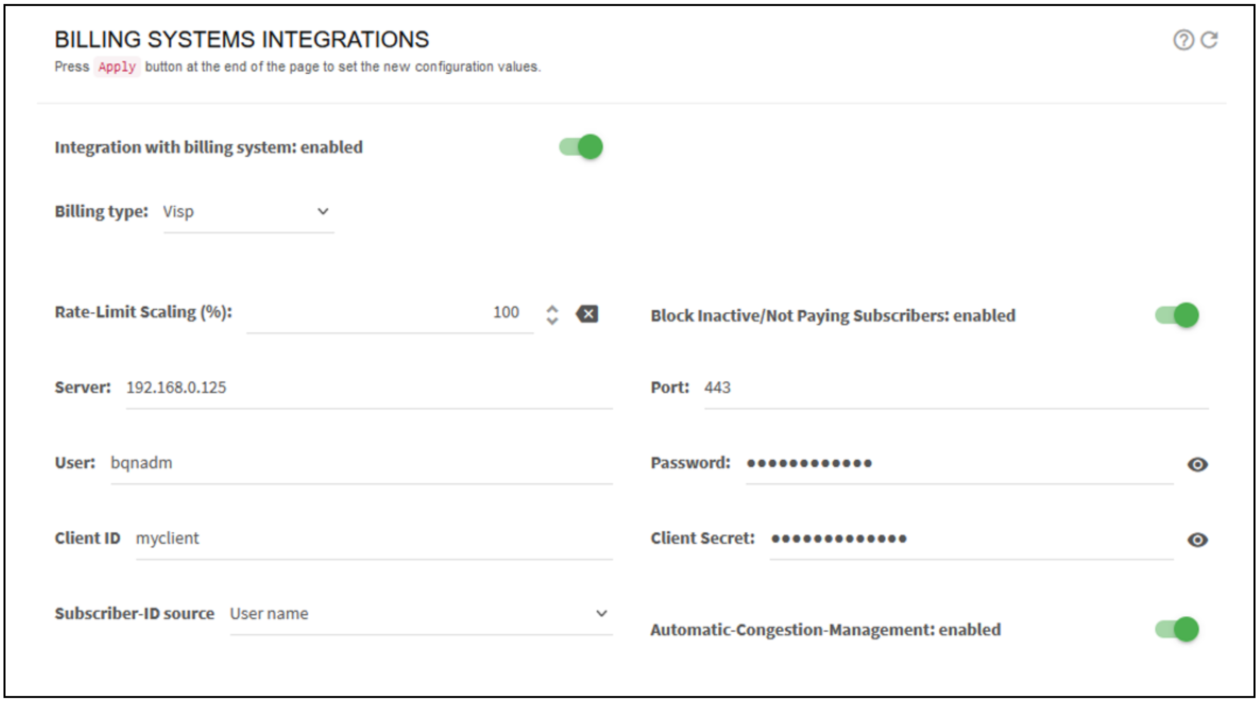

BILLING SYSTEMS INTEGRATIONS

Press Apply button at the end of the page to set the new configuration values.

Integration with billing system: enabled

Billing type: Visp

Rate-Limit Scaling (%): 100

Block Inactive/Not Paying Subscribers: enabled

Server: 192.168.0.125

Port: 443

User: bqnadm

Password: ••••••••••••

Client ID myclient

Client Secret: •••••••••••••

Subscriber-ID source User name

Automatic-Congestion-Management: enabled

Figure 101: *Enabling the Visp.net*

Visp customers with a package status and service instance status other than **ACTIVE** are blocked. You can change this behavior to non-blocking disabling the switch **Block Inactive/Not Paying Subscribers** (see Blocking of inactive subscribers section).

The following fields in Visp billing can be used as source of the subscriber ID:

- Customer-ID
- First + Last name
- Username

Warning

Do not change the Subscriber-Id source after the first sync with the billing system. If it is changed, then all statistics associated with the old Subscriber-Id are not available on the UI.

WISPControl

WISPControl uses RADIUS. For more information on WISPControl, refer to *QoE Appliance RADIUS Guide*.

Wisphub

Wisphub has developed an integration with QoE using QoE REST API. For more information, refer to Wisphub product documentation.

Wispro

The QoE retrieves clients, contracts and plans to get the speed limits to apply.

To activate Visp, navigate to **Configuration** >**RADIUS/REST/Billing** >**Billing Systems**, select Wispro and enable the switch.

A valid API key must be provided. The API key must be generated in the Wispro system. See instructions in the *QoE Appliance REST API Guide*.

The provided IP address or server name will be used, along with the port (443 by default), to send API queries to Wispro (https://<server>:<port>/api/v1).

The following screen shows an example of Wispro configuration:

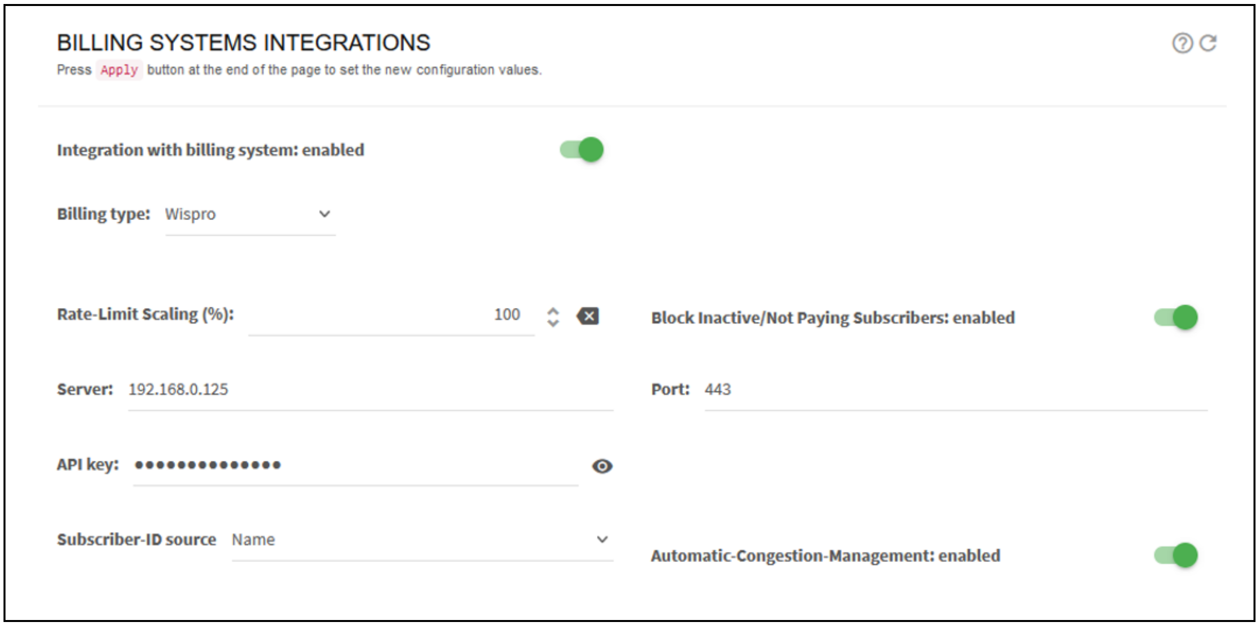

BILLING SYSTEMS INTEGRATIONS

Press Apply button at the end of the page to set the new configuration values.

Integration with billing system: enabled

Billing type: Wispro

Rate-Limit Scaling (%): 100

Block Inactive/Not Paying Subscribers: enabled

Server: 192.168.0.125

Port: 443

API key: ••••••••••••••

Subscriber-ID source Name

Automatic-Congestion-Management: enabled

Only Wispro clients with contracts in **disabled** state will be blocked. You can change this behaviour to non-blocking disabling the switch **Block Inactive/Not Paying Subscribers**.

The following fields in Wispro billing can be used as source of the subscriber ID:

- Subscriber ID (“public_id” in Wispro client).
- Name.
- MAC address.
- Login (“email” field in Wispro client).

If a Wispro client has more than one contract, each IP address will be assigned the speed limits of its contract. If an IP address is repeated in two client contracts (this is an inconsistency in the billing database that should not happen), the speed limits of the last obtained contract will be selected.

General Billing Considerations

Subscriber ID

The billing system can be the source of QoE subscriber IDs. The choices of ID sources depend on the billing system (see each billing section for details).

For billing systems integrated using REST API, the billing system is in full control of the subscriber ID, that can be explicity defined when creating or editing the subscriber.

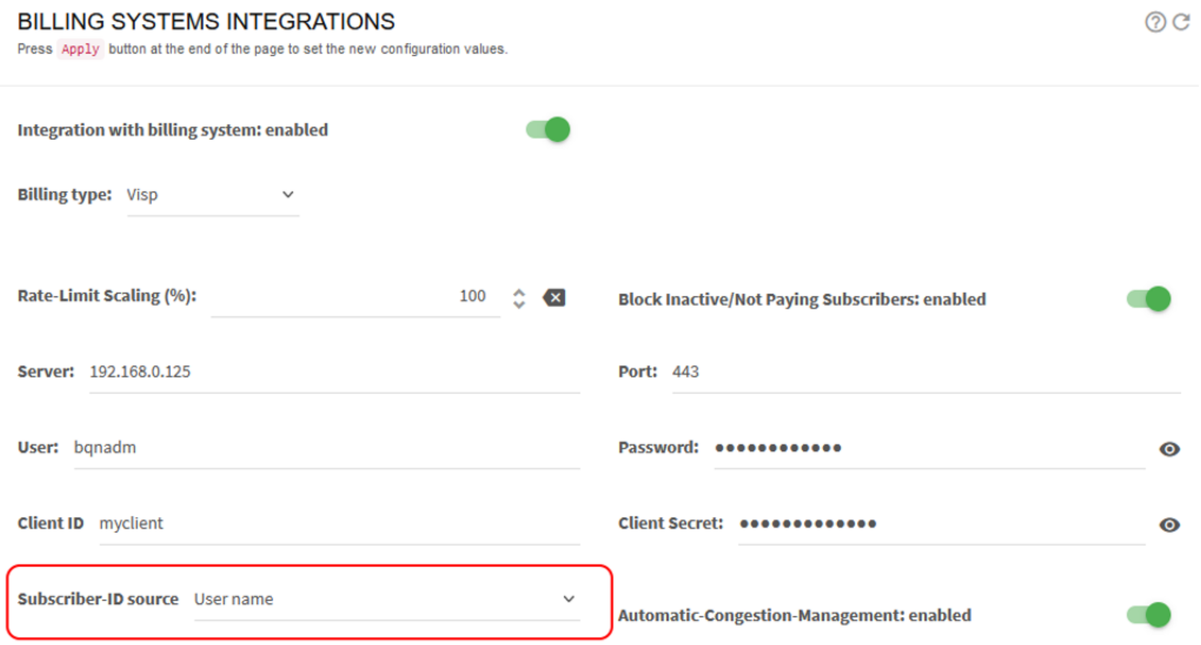

BILLING SYSTEMS INTEGRATIONS

Press Apply button at the end of the page to set the new configuration values.

Integration with billing system: enabled

Billing type: Visp

Rate-Limit Scaling (%): 100

Block Inactive/Not Paying Subscribers: enabled

Server: 192.168.0.125

Port: 443

User: bqnadm

Password: ●●●●●●●●●●●●

Client ID myclient

Client Secret: ●●●●●●●●●●●●●

Subscriber-ID source User name

Automatic-Congestion-Management: enabled

Block Inactive / Not Paying Subscribers

By default, QoE will block non-paying subscribers. What is a non-paying subscriber depends on the billing system (see each specific billing section for details). To prevent QoE from blocking non-paying subscribers, disable the switch Block Inactive/Not Paying Subscribers.

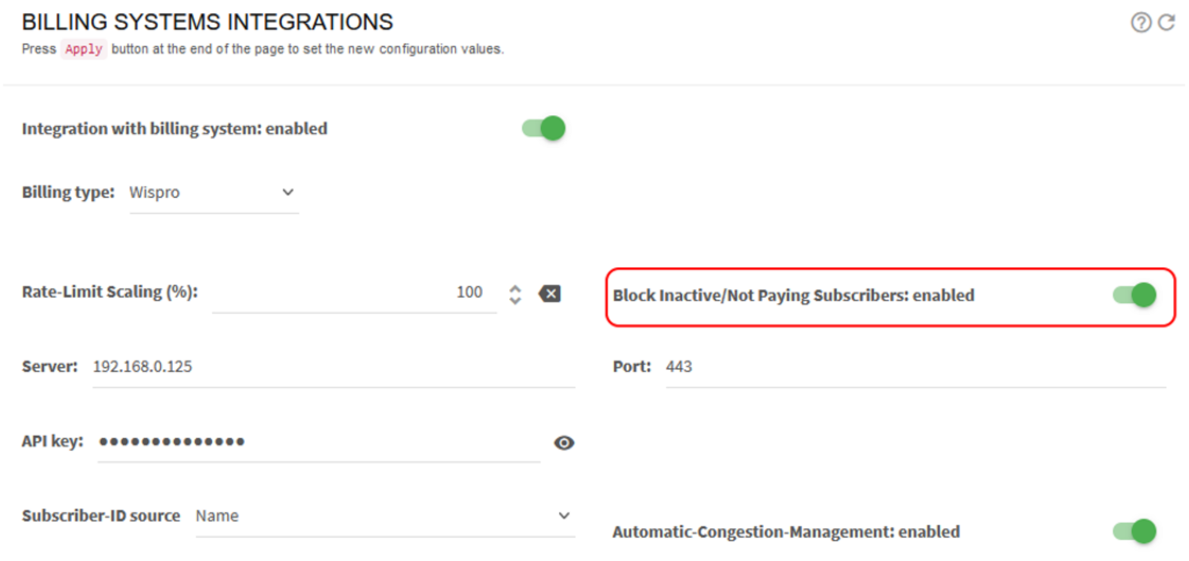

BILLING SYSTEMS INTEGRATIONS

Press Apply button at the end of the page to set the new configuration values.

Integration with billing system: enabled

Billing type: Wispro

Rate-Limit Scaling (%): 100

Block Inactive/Not Paying Subscribers: enabled

Server: 192.168.0.125

Port: 443

API key: ••••••••••••••

Subscriber-ID source Name

Automatic-Congestion-Management: enabled

For billing systems integrated using REST API, the billing system is in full control of the rate limit definition and it can block a subscriber by assigning it to a policy with 0 rate limit speed.

Rate Scaling Factor

By default, the QoE applies the rate limits as specified by the billing system. It is possible to apply a scaling factor to those limits using the **Rate-Limit Scaling** field.

To enforce a speed limit lower than the one in the billing, use a factor less than 100%. For example, a limit in the billing of 200 Mbps with a factor of 90% will be 180 Mbps).

To enforce a speed higher than in the billing, use a factor bigger 100% and up to 200% (maximum factor possible). For example, a rate of 200 Mbps in the billing with a factor of 150% will be 300 Mbps.

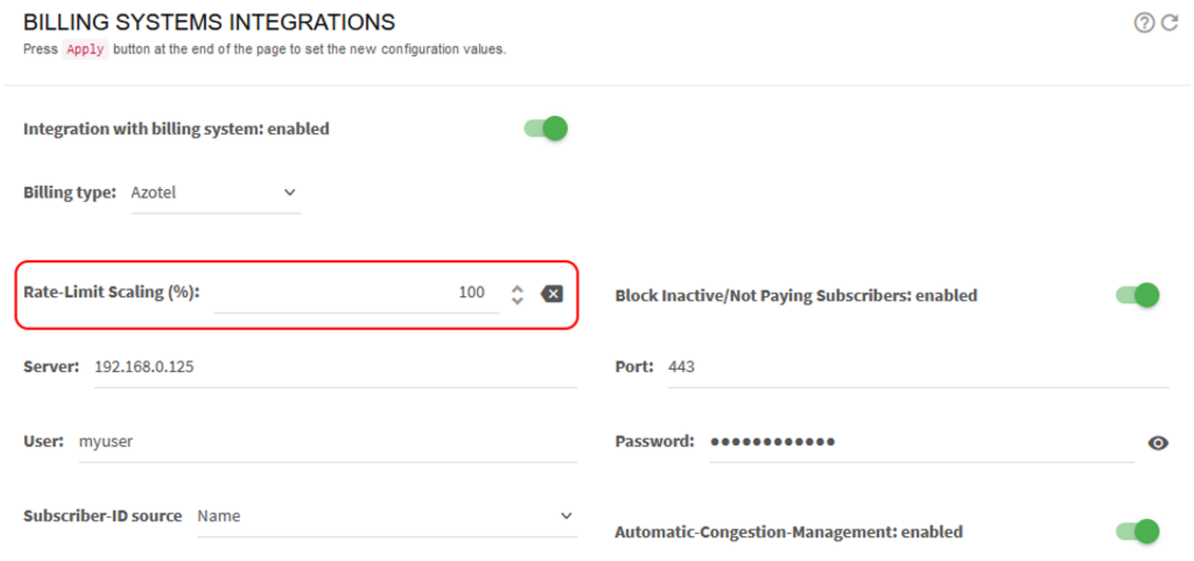

BILLING SYSTEMS INTEGRATIONS

Press Apply button at the end of the page to set the new configuration values.

Integration with billing system: enabled

Billing type: Azotel

Rate-Limit Scaling (%): 100

Block Inactive/Not Paying Subscribers: enabled

Server: 192.168.0.125

Port: 443

User: myuser

Password: ••••••••••••

Subscriber-ID source Name

Automatic-Congestion-Management: enabled

For billing systems integrated using REST API, the billing system is in full control of the rate limit definition and can decide which factor if any apply to the limits sent to the QoE.

Chapter 13: Troubleshooting

This chapter contains the following sections:

Installation related issues

For the installation related issues, verify the following:

- Copy the ISO to a USB drive with MBR partitions and DD mode.
- The server BIOS must have its factory default boot mode (for example, DUAL).
- If there is a RAID controller, a logical drive must be configured and marked as **Bootable**.
- Verify if the server used meets the QoE hardware requirements for installation.
- Ensure that the installation is done in the hard disk and not overwritten in the USB drive.

Installation issues

For the installation related issues, check the following:

- The ISO should be copied to a USB drive with MBR partitions and DD mode.
- The server BIOS should have its factory default boot mode (for example, DUAL).

- If there is a RAID controller, a logical drive must be configured and marked as bootable.
- Check if server used in the installation meets QoE hardware requirements.
- Check that the installation was done in the hard disk and not overwritten in the USB drive.

No access to the management IP address

QoE uses a dedicated network interface for the management. The management interface supports both the SSH and Web (HTTPs) services. If any problem in accessing the configured management IP, then verify the following:

- Ensure that the management network interface port is connected to the appropriate network.
- Verify that the link state of the management network interface is working. If the management interface is connected to a network switch, verify that the port in the switch is up and its attributes match the properties shown by the **show interface** command.
- If you are accessing the management IP address from a different network, ensure that static routing is configured to the access network, as explained in the **Network Interface** section in the *QoE Appliance User Interface Guide*.
- If there are firewalls in the management network, then permit the access to TCP port 22 for the SSH service, and TCP port 443 for the WEB service.
- Verify using the system console that the management IP address and network prefix are correct. Connect a monitor and a keyboard to the server and login as **root**:

  ```
  bqn0:# bqnsh

  root@bqn0# show interface management detail

  Interface: en0o1

  IP address: 192.168.0.121/24

  Default gateway: 192.168.0.1

  Nameserver: n/a
  ```

- If you suspect the OAM IP settings are incorrect or unknown, connect a monitor and a keyboard to the server and login as **root** to change it using the bta wizard in interactive mode. For example, to change the management interface to en0o1 with IP address 10.10.10.12/24 (press *Enter* to accept suggested response):

  ```
  bqn0:# bqnsh

  root@bqn0# wizard bta interactive

  Available network interfaces:

  en0o1

  en0p0s0

  en0p0s1

  Enter management interface [en0p0s0]: en0o1

  Enable VLAN on management interface? (yes/no) [no]:

  Enter management IP address and prefix [192.168.0.120/24]: 10.10.10.12/24

  Enter default gateway IP address [192.168.0.1]: 10.10.10.1
  ```

```
Configure a nameserver? (yes/no) [no]:

Available network interfaces:

en0p0s0

en0p0s1

Select access-side interface for wire: en0p0s0

Select internet-side interface for wire: en0p0s1

Enable SDR generation? (yes/no) [yes]:

Enter random optimization percentage [99]:

Enter random udr generation percentage [2]:

System vendor: Dell

System name: bqn

System serial: 0

System supported: yes

CPU model: 12th Gen Intel(R) Core(TM) i7-12700H

CPU cores: 4

Management interface: en0o1

Management IP: 10.10.10.12/24

Management gateway: 10.10.10.1

Wire 1: en0p0s0(access)-en0p0s1(internet)

SDR generation: enabled

BTA random optimization: 99%

UDR random generation: 2%

If the proposed configuration is not valid execute the command

wizard bta interactive

to manually enter the configuration.

Proceed with configuration? (yes/no) [yes]:

root@bqn0# show interface management detail

Interface: en0o1

IP address: 10.10.10.12/24

Default gateway: 10.10.10.1

Nameserver: n/a
```

If the interface is not available at the time of the change (for example, participated in a wire), a reboot message appears. After the reboot, QoE has the new IP and management network interface.

```
Management interface en0o1 seems not to be set. A process reboot may fix the
problem.

Proceed with the process reboot? (yes/now) [yes]: yes
```

- The QoE management interface may be protected by its own firewall. The problem is that the source IP address is not included in that firewall whitelist. This happens even for addresses from the same subnet of the QoE management IP if the subnet is not part of the firewall rules. The firewall can be disabled temporarily until the connection to the management port is restored. Connect a monitor and a keyboard to the server and login as **root**:

  ```
  bqn0:# bqnsh

  root@bqn0# show interface firewall

  IFACE CHAIN RANGE

  en0o1 input 10.0.0.0/8

  en0o1 input 172.16.0.0/12

  en0o1 input 192.168.0.0/16

  root@bqn0# clear interface en0o1 firewall input root@QoE0# show interface
  firewall

  IFACE CHAIN RANGE

  root@bqn0#
  ```

Once the management IP is reachable, define the new whitelist of allowed source IP ranges.

Web is not accessible

If any problem in accessing the Web, then verify the followings:

- Check that the management IP address is accessible using SSH.
- Check that HTTPS is used to access (HTTP is not supported) (for example, https://192.168.0.121).
- Check that **admin** account is used (cannot use root for accessing UI).
- When installing from scratch, ensure that the **wizard bta** command is executed (else, the UI web service is not active and no **admin** user is created).
- Verify that the SSH port of the QoE server is not modified. To access the QoE using a port other than 22, you can define port forwarding rules in router on the access path, but cannot change the QoE SSH. Login into the server as root and verify the SSH port is 22 as follows:

  ```
  bqn:~ # grep Port /etc/ssh/sshd_config

  #Port 22

  #GatewayPorts no
  ```

If required, comment the line that specifies a port other than 22.

- Check that the browser (such as Edge, Firefox, Chrome) is supported. **MS Explorer** is not supported.

Network interfaces down

If the **Network Interfaces** icon in the dashboard is not in green, then navigate to **Configuration > Interfaces > Data Wires** and update the network interfaces.

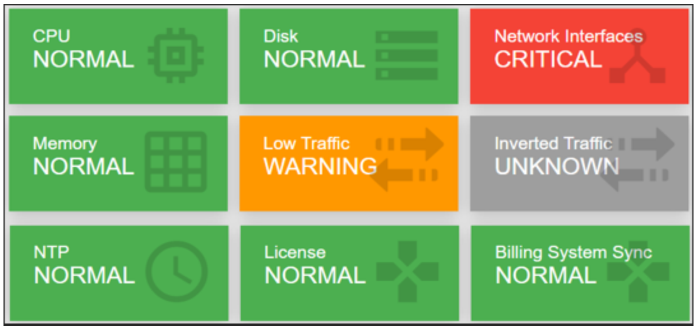

Figure 102: *The QoE dashboard*

Red in color (Critical)

- If there is no wire configured, then create and configure a new wire.
- If there are wires configured but their interfaces are not in UP state, this most likely indicates that the interfaces are not Intel compatible: remove the wire and create a new one with both interfaces in **pcap** mode. This should place the interfaces in UP state, but with much lower throughput capacity (less than 1 Gbps).
- If there are wires configured, with interfaces in UP state but with the LINK down, there is a problem in the connection with the other equipment. Connect both interface ports to one another in a loop using a suitable cable/fiber.
  - If both interfaces are in up state, then the problem is on the other equipment.
  - If the link is still down and optic ports are used, then verify the following:
    - transceivers are Intel-compatible.
    - transceivers are supported (for supported hard disks, see Supported Network Interfaces).
    - transceivers of the type required by the installation (for example, SFP+-LR in an installation with monomode fiber and SFP+-LR on the other side).

Yellow in color (Notice)

- If the wires that appear as down are supposed to be with traffic, follow the steps in Red in color (Critical) section.
- If the wires that appear as down, are not in use and you want to remove the notice signal, then delete the unused wires. Consider that the changes in wire configuration stop the traffic for some seconds.

Inverted traffic

If the **Inverted Traffic** icon in the dashboard is in orange color (Warning), then it indicates that the traffic throughput in the uplink direction is bigger than the downlink direction. This can be normal in small deployments (less than a hundred subscribers, like QoE in a lab) but in a network deployment, it indicates that some of the wires are connected incorrectly, with the access port connected to the Internet side and vice versa.

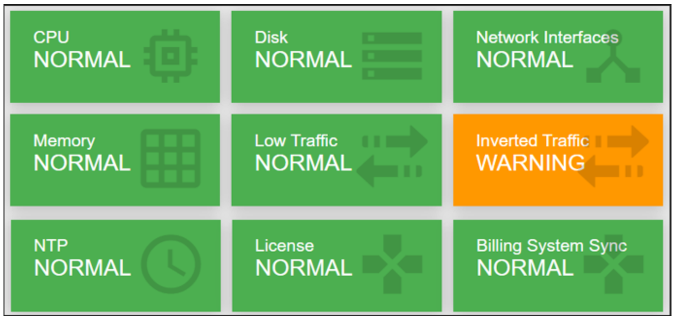

CPU NORMAL
Disk NORMAL
Network Interfaces NORMAL
Memory NORMAL
Low Traffic NORMAL
Inverted Traffic WARNING
NTP NORMAL
License NORMAL
Billing System Sync NORMAL

Figure 103: *The QoE dashboard*

To verify the inverted traffic, navigate to **Statistics > Throughput > Interfaces** and select the wire interface configured in the access side. If it shows more received throughput than the sent throughput, then the wire is inverted. It can be confirmed by selecting the throughput of the Internet-side interface to see that its sent traffic is bigger than the received traffic.

To fix this issue, navigate to **Configuration > Interfaces > Data Wires**, click on **Swap interfaces** (⇄) icon from the inverted wire.

Low traffic

If the **Low Traffic** icon in the dashboard is not in green color, then it indicates the high traffic.

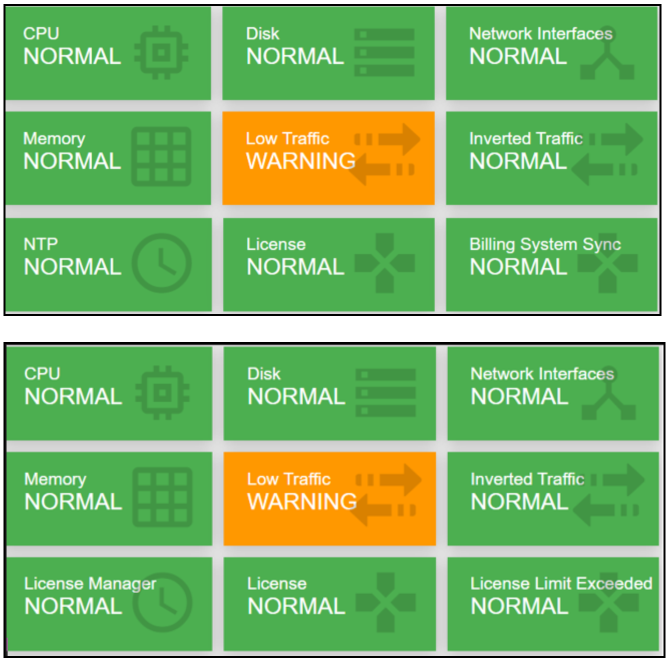

CPU NORMAL
Disk NORMAL
Network Interfaces NORMAL
Memory NORMAL
Low Traffic WARNING
Inverted Traffic NORMAL
NTP NORMAL
License NORMAL
Billing System Sync NORMAL

CPU NORMAL
Disk NORMAL
Network Interfaces NORMAL
Memory NORMAL
Low Traffic WARNING
Inverted Traffic NORMAL
License Manager NORMAL
License NORMAL
License Limit Exceeded NORMAL

Figure 104: *The QoE dashboard*

Yellow in color (Notice)

Hover the mouse over icon to confirm that the **Traffic-low notice** is displayed. It indicates that, very little traffic going through the QoE server. This is normal, if the system is waiting for traffic routed through it. However, in a system which is in production can be an indication for some failure, elsewhere in the network is preventing the traffic to reach the QoE server.

Orange in color (Warning)

Hover the mouse over icon. Either **Traffic-uplink** or **Traffic-downlink** should be in warning. This is because there is no traffic going through the QoE in that direction and therefore the traffic is asymmetric. Navigate to **Configuration > TCPO/ACM Settings** and set the **Overall TCP Optimization** to **OFF** for any issue.

Fix the traffic routing, so the QoE gets both the directions (uplink and downlink). After it is complete and the icon returns to green in color, navigate to **Configuration > TCPO/ACM Settings** to enable overall TCP optimization.

License manager

If the **License Manager** icon in the dashboard is in Yellow and the text says **license-mgr-connection: notice**, then the QoE server cannot reach the license manager. The license manager is responsible of validating QoE software licenses and also helps the Cambium Networks to provide a support by reporting the server problems.

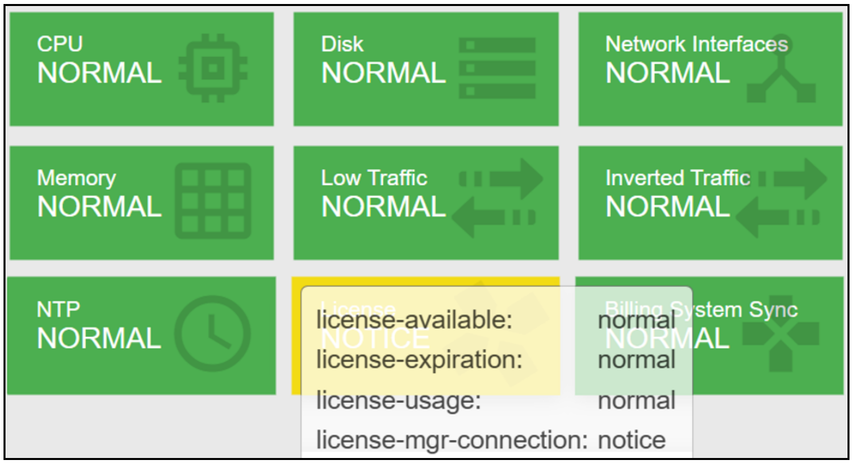

```
license-available:       normal
license-expiration:      normal
license-usage:           normal
license-mgr-connection:  notice
```

Figure 105: *The QoE dashboard*

Ensure that the QoE server can initiate the outgoing connections to the License Manager IP (For details, contact Cambium Network support).

To verify the outgoing connection is possible, log in as **root**. A telnet to the provided IP and the port should work.

```
bqn0:# telnet <ip> <port>

Trying <ip>...

Connected to <ip>.

Escape character is '^]'.
```

License issues

If the **License** icon in the dashboard is in red color, and the text says **license-available: critical**, then there is no valid license.

This may be due to the following reasons:

- There is no license defined in the node
- The license is invalid

- The license is no longer valid (date is expired).

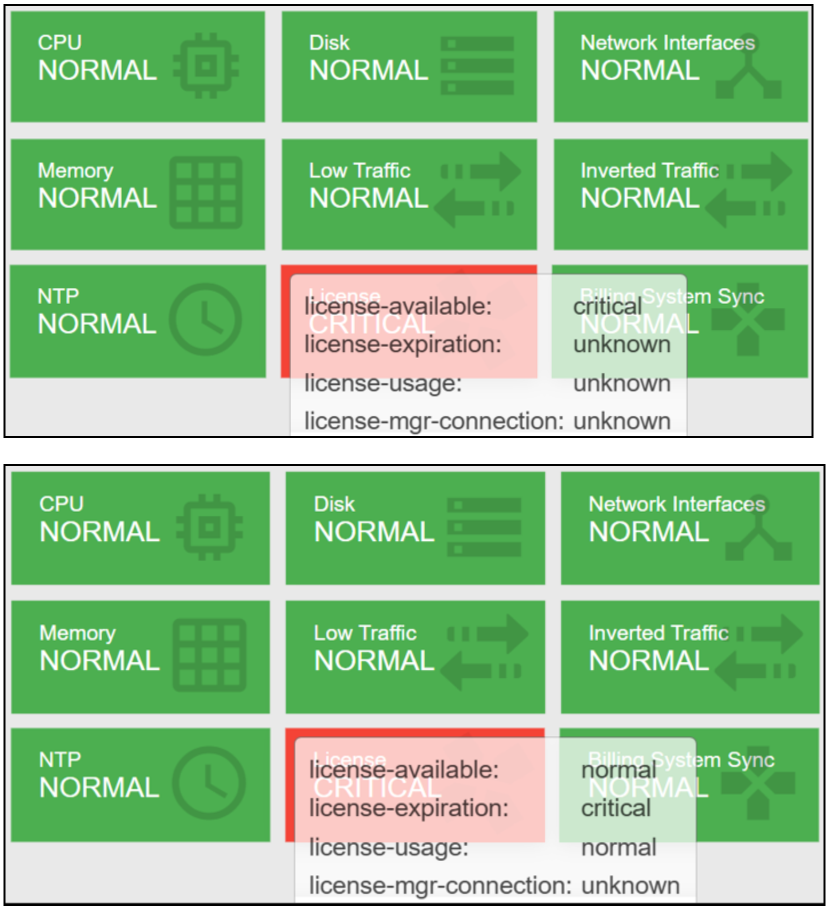

```
license-available:       critical
license-expiration:      unknown
license-usage:           unknown
license-mgr-connection:  unknown
```

```
license-available:       normal
license-expiration:      critical
license-usage:           normal
license-mgr-connection:  unknown
```

Figure 106: *The QoE dashboard*

If the **License** icon in the dashboard is in Red, and the text says **license-available: critical**, then there is no valid license. This can be due to the following reasons:

- There is no license defined in the node
- The license is not valid
- The license is no longer valid (final date is expired).

Contact your distributor for a valid license. You can check the license state in **Administration > License**.

When there is no valid license, the QoE forwards all traffic transparently. That is, the service does not affected, but none of the QoE advanced processing is applied to the traffic.

License limit exceeded

If the **License** icon in the dashboard is in orange and the text says **license-usage: warning**, then the maximum capacity of the license is exceeded (the traffic throughout in the QoE server is above the license limit).

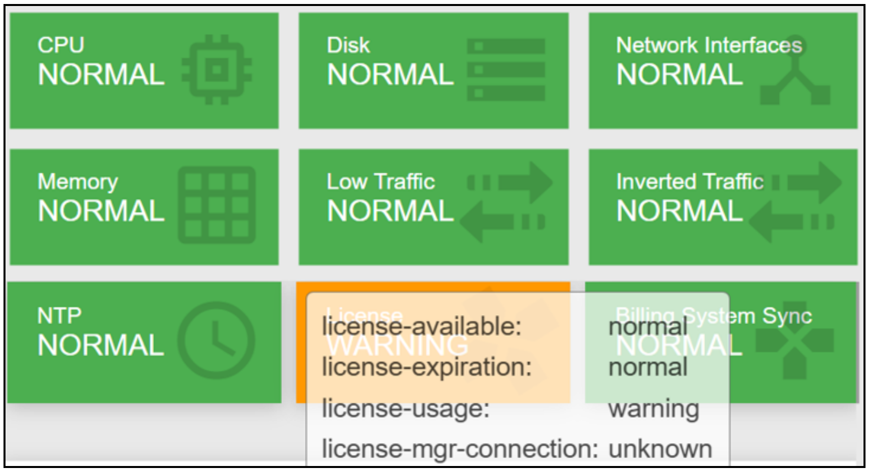

license-available: normal
license-expiration: normal
license-usage: warning
license-mgr-connection: unknown

Figure 107: *The QoE dashboard*

Contact your distributor for a license upgrade. You can check the license capacity in **Administration > License**.

In **Statistics > Throughput > Overview**, a red line shows the license limit along with latest throughput levels.

> **Note**
>
> Exceeding license limit should be temporal, while the license is upgraded to the right capacity.

When the license limit is exceeded, the QoE does **not** drop any packets; it simply bridge them.

The effect on the traffic is different at flow and at subscriber session level.

The following are the effects at flow level:

- While the license limit is being exceeded, new flows will have no license.
- Existing flows prior to the license being exceeded are unaffected: if they were being optimized remain so and likewise if without a license.
- Once the optimized traffic falls below the limit, new flows are optimized again.
- Existing flows prior to the license going below the limit are unaffected: if they had no license remain so and likewise if they were being optimized.
- A flow without a license has:
  - No TCPO
  - No shaping
  - It does not generate metrics (retransmissions, latencies, DPI, etc.).

The following are the effects at subscriber session level:

- While the license limit is being exceeded, new subscriber sessions has no license.
- Existing subscriber sessions prior to the license is exceeded remains initially in their current state. If they were optimized, then remains so and likewise if without a license.

- If the license remains above the limit, subscriber sessions is optimized and may transit to no license if they accumulate enough flows.

- Once the optimized traffic falls below the limit, new subscriber sessions are optimized again.

- Exiting subscriber sessions prior to the license going below the limit remains initially in their current state. If they had no license remain so and likewise if they were optimized.

- If the license remains below the limit, subscriber sessions without a license, then it may transit to optimized if they accumulate enough flows.

- A subscriber session without a license has the following:

    - No ACM

    - No rate limiting.

    - It continues generating traffic volume totals.

    - The other metrics such as retransmissions, latencies are reduced to those from the flows being optimized.

- A subscriber session being optimized is affected by the fact that some of his flows may not have a license and therefore the subscriber may have a reduced set of measurements in comparison with normal operation. This affects the ACM and also the metrics generated for that subscriber.

The result is that more and more traffic no longer gets QoE functionality, and conversely, the amount of traffic getting QoE functionality gets lower, and soon below the license limitation. At that point, new flows get TCPO optimization and subscribers that got their rate limitation de-activated gets activated again. So, with these oscillations, QoE functionality is provided to an amount of traffic equivalent to the license limit. License exceeded events are meant to be temporal, while the license is upgraded to the right capacity.

High CPU load

If the **CPU** icon in the dashboard is not in green color, some CPUs are running at abnormally high levels. This is normally due to unbalanced traffic (concentrated in a few subscriber IPs) or very high traffic is proceeded by the QoE server.

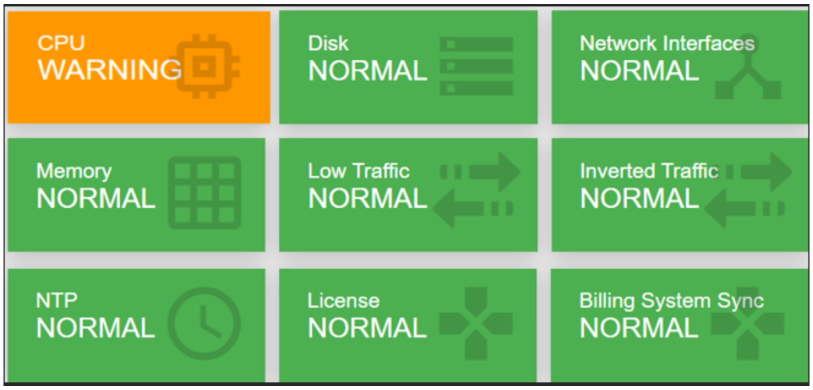

CPU
WARNING

Disk
NORMAL

Network Interfaces
NORMAL

Memory
NORMAL

Low Traffic
NORMAL

Inverted Traffic
NORMAL

NTP
NORMAL

License
NORMAL

Billing System Sync
NORMAL

Figure 108: *The QoE dashboard*

The QoE has internal mechanisms to simplify this situation, and try to prevent the traffic losses by reducing the amount of optimized traffic.

The throughput level can be verified by **Statistics > Throughput > Overview** and the CPU levels in **Statistics > System > CPU**.

There are two alarm types, depending on the CPU load levels:

- **Orange**: If some CPU cores are at high load (above 80 percent usage).
- **Red**: If some CPU cores are at very high load (above 90 percent usage).

Perform the following steps to reduce the CPU load:

- If you are using NAT between the QoE and the end subscribers, increment the number of IPs used by the NAT, so the QoE can distribute the traffic among more addresses. You can do the traffic distribution among IPs under **Statistics > Subscribers > Top By Time**.
- Enable the bypass path or if not possible, reduce the amount of traffic which are routed through the QoE.
- Disable TCP optimization from **Configuration > TCPO/ACM Settings**.
- A hardware upgrade is required. For more information, visit Cambium Networks Support Site.

If the QoE server is used only for TCPO and/or per-flow speed limits, the CPU load distribution can be improved enabling per-flow steeting. Per-flow steering distributes the traffic load across CPU cores per individual flow, instead of the default per-subscriber traffic distribution. This improves CPU load balance, but because per-subscriber control is done per core, does not allow subscriber-level control, such as ACM, policy rates or per subscriber flow limits.

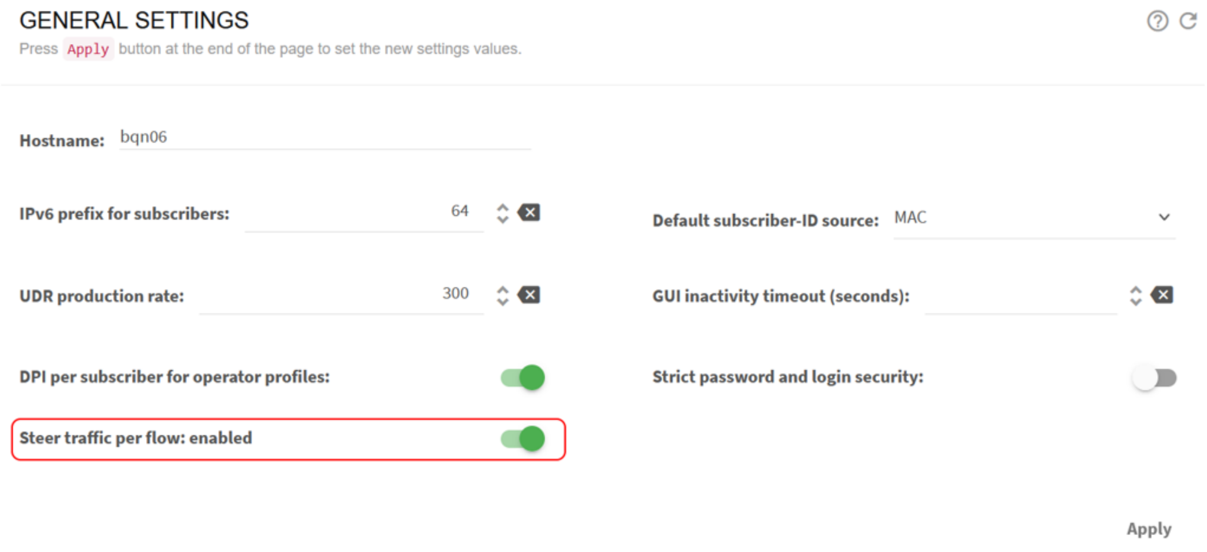

GENERAL SETTINGS

Press Apply button at the end of the page to set the new settings values.

Hostname: bqn06

IPv6 prefix for subscribers: 64

Default subscriber-ID source: MAC

UDR production rate: 300

GUI inactivity timeout (seconds):

DPI per subscriber for operator profiles: ●

Strict password and login security: ○

Steer traffic per flow: enabled ●

Apply

High memory load

If the **Memory** icon in the dashboard is not in green color, some processes are running out of memory. This is normally due to unbalanced traffic (traffic concentrated in a few subscriber IPs) or very high traffic is processed by the QoE server.

The QoE has internal mechanisms to mitigate this situation, and try to prevent traffic losses by reducing the amount of optimized traffic.

The throughput level can be verified by **Statistics > Throughput > Overview** and the memory levels in **Statistics > System > Memory**.

There are two alarm types, depending on the memory load levels:

- **Orange**: If some processes reach high usage (above 90 percent usage).
- **Red**: If some processes reach very high usage (above 95 percent usage).

Perform the following steps to reduce the memory load:

- If you are using NAT between the QoE and the end subscribers, increment the number of IPs used by the NAT, so the QoE can distribute the traffic among more addresses. You can see how traffic is distributed among IPs in **Statistics > Subscribers > Top By Time**.
- Enable the bypass path or if not possible, reduce the amount of traffic being routed through the QoE.
- Disable TCP optimization from **Configuration > TCPO/ACM Settings**.
- A hardware upgrade is required. For more information, visit Cambium Networks Support Site.

No RADIUS messages are received

If no RADIUS information is displayed in the QoE UI and to check if RADIUS messages are incoming, login to QoE shell and run the following command:

```
$ ssh bqnadm@192.168.0.121

bqnadm@bqn# system interface en0o1 capture filter 'udp and port 1813'

listening on eno1, link-type EN10MB (Ethernet), capture size 65535 bytes
```

In this example, the management interface is **eno1**. Check one of your QoE server in **Configuration > Interfaces > Management** (the 0, added by QoE configuration must be removed, for example, **en0s1f0** in QoE UI is **ens1f0** in Linux).

If the QoE Firewall is configured (**Configuration > Interfaces > Management Firewall**), all RADIUS client IPs must be added (in this example, 10.10.10.10 y 10.10.10.11).

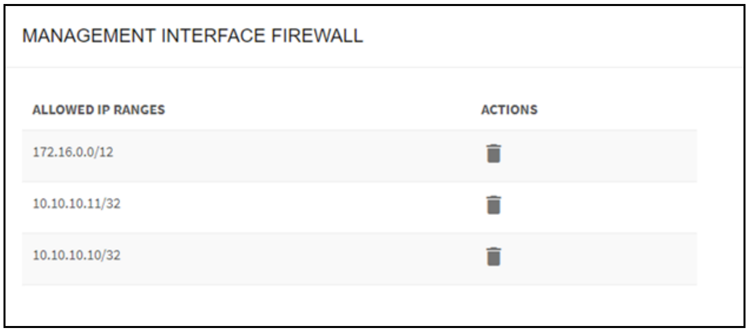

MANAGEMENT INTERFACE FIREWALL

ALLOWED IP RANGES	ACTIONS
172.16.0.0/12	🗑
10.10.10.11/32	🗑
10.10.10.10/32	🗑

Figure 109: *The management interface firewall page*

And now the RADIUS messages are received:

```
#  tcpdump -i eno1 'udp and port 1813’

tcpdump: verbose output suppressed, use -v or -vv for full protocol decode

listening on eno1, link-type EN10MB (Ethernet), capture size 65535 bytes

14:21:20.177347 IP 10.10.10.10.60072 > 192.168.0.121.radius-acct: RADIUS,
Accounting Request (4), id: 0xf0 length: 222

14:21:20.177424 IP 192.168.0.121.radius-acct > 10.10.10.10.60072: RADIUS,
Accounting Response (5), id: 0xf0 length: 20

. . .
```

If RADIUS messages are not received, then the rest of the traffic jumps requires verification. In this example, RADIUS clients are in subnet 10.10.10.0/24 and QoE in subnet 192.168.0.0/24. Verify that there are valid routes between the two subnets and no intermediate firewall is blocking UDP port 1813 (RADIUS Accounting).

If the RADIUS messages are received, but the information is not reflected as expected, enable the event log entering the QoE server shell through ssh:

```
bqnadm@bqn# configure

bqnadm@bqn(config)# api common

bqnadm@bqn(config-api)# event level general

bqnadm@bqn(config-api)# event level radius

bqnadm@bqn(config-api)# event level policy

bqnadm@bqn(config-api)# event level subscriber

bqnadm@bqn(config-api)# root

bqnadm@bqn(config)# commit

bqnadm@bqn(config)# end

bqnadm@bqn#
```

Now, the event log of the received RADIUS requests:

```
bqnadm@bqn# show api event log

2023-05-24T18:16:32.138 [radius] Sent Accounting-Response message to
172.27.1.194:42043: id(134)

2023-05-24T18:16:32.138 [radius] Received Accounting-Request message from
172.27.1.194:42043: id(134) statusType(start) framedIpAddress(10.0.0.11)
mikrotikAddressList() mikrotikRateLimit(45M/90M 90M/100M 45M/90M 5/5)

2023-05-24T18:16:32.138 [policy] Created "RA-45M/90M-90M/100M-45M/90M-5"
policy: rate(45000/90000) burstRate(90000/100000) burstDurationMs
(5000/5000) burstThreshold(45000/90000) burstThresholdWindow(-1/-1)
burstTransitionDurationMs(-1/-1) autoCongestionManagement(yes/yes)

2023-05-24T18:16:32.138 [subscriber] Updated "10.0.0.11" subscriber: policy
(RA-45M/90M-90M/100M-45M/90M-5) sessionId(1234) userName(Sub-102)
callingStationId(+34100100102) nasId() nasPort(4294967295)

2023-05-24T18:16:32.157 [radius] Sent Accounting-Response message to
172.27.1.194:22090: id(75)

2023-05-24T18:16:32.157 [radius] Received Accounting-Request message from
172.27.1.194:22090: id(75) statusType(start) framedIpAddress(10.0.0.12)
mikrotikAddressList() mikrotikRateLimit(10M/20M 0K/0K 0K/0K 0/0)

2023-05-24T18:16:32.157 [subscriber] Updated "10.0.0.12" subscriber: policy
(RA-10M/20M-0K-0K-0) sessionId(1234) userName(Sub-103) callingStationId
(+34100100103) nasId() nasPort(4294967295)

. . .
```

No REST messages are received

If the REST information is not displayed in the QoE UI, then check the reception of REST messages in the QoE. Login to the QoE shell as **root** and execute the following command:

In this example, **eno1** is the management interface.

```
#  tcpdump -i eno1 tcp and port 3443

tcpdump: verbose output suppressed, use -v or -vv for full protocol decode

listening on eno1, link-type EN10MB (Ethernet), capture size 65535 bytes
```

Navigate to **Configuration > Interfaces > Management** and Check in your server (remove the 0 added by QoE configuration, for example, **en0s1f0** in the QoE UI is **ens1f0** in UNIX).

If the QoE Firewall is configured (**Configuration > Interfaces > Management Firewall**), the IPs of all the REST clients must be added (in our example, 10.10.10.10 y 10.10.10.11).

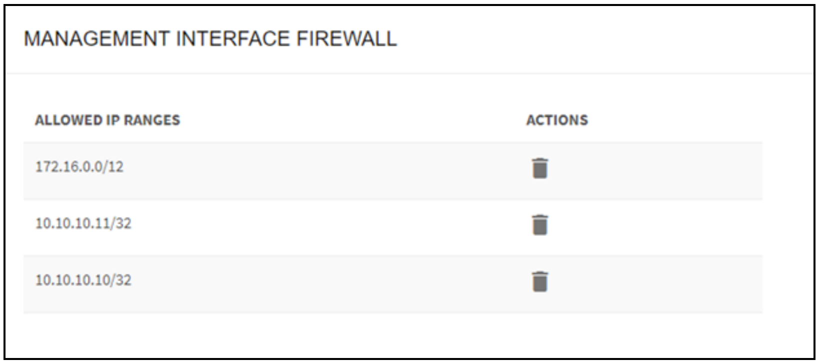

MANAGEMENT INTERFACE FIREWALL

ALLOWED IP RANGES	ACTIONS
172.16.0.0/12	🗑
10.10.10.11/32	🗑
10.10.10.10/32	🗑

Figure 110: *The management interface firewall page*

The following messages are received:

```
#  tcpdump -i eno1 'udp and port 1813'

tcpdump: verbose output suppressed, use -v or -vv for full protocol decode

listening on eno1, link-type EN10MB (Ethernet), capture size 65535 bytes

17:30:30.767149 IP 192.168.88.12.48316 > 192.168.88.13.ov-nnm-websrv: Flags [S],
seq 639501187, win 64240, options [mss 1460,sackOK,TS val 3813494325 ecr
0,nop,wscale 7], length 0

17:30:30.767163 IP 192.168.88.13.ov-nnm-websrv > 192.168.88.12.48316: Flags [S.],
seq 2135448282, ack 639501188, win 28960, options [mss 1460,sackOK,TS val 607264358
ecr 3813494325,nop,wscale 5], length 0

17:30:30.767260 IP 192.168.88.12.48316 > 192.168.88.13.ov-nnm-websrv: Flags [.],
ack 1, win 502, options [nop,nop,TS val 3813494325 ecr 607264358], length 0

. . .
```

If the REST messages are not received, then check the rest of traffic steps. In this example, the REST clients are in 10.10.10.0/24 and the QoE in 192.168.0.0/24. Verify that there are valid routes between both subnets and that no intermediate Firewall is blocking the TCP port 3443.

If the REST messages are received, but the information is not reflected as expected, enable the event log and traces entering the QoE server shell via ssh:

```
bqnadm@bqn# configure

bqnadm@bqn(config)# api common

bqnadm@bqn(config-api)# event level general

bqnadm@bqn(config-api)# event level rest

bqnadm@bqn(config-api)# event level policy

bqnadm@bqn(config-api)# event level subscriber

bqnadm@bqn(config-api)# root

bqnadm@bqn(config)# api rest

bqnadm@bqn(config-rest)# trace request 5
```

```
bqnadm@bqn(config-rest)# trace response 5

bqnadm@bqn(config-rest)# root

bqnadm@bqn(config)# commit

bqnadm@bqn(config)# end

bqnadm@bqn#
```

Now, the event log of the received REST requests is displayed. In the following example, the QoE receives some POSTs assigning subscribers to rate policies:

```
bqnadm@bqn# show api event log

2023-05-24T17:41:40.268 [subscriber] [10.0.0.3] Updated subscriber: policy
(rest-static-3) sessionId() subscriberId(n/a) customerId(1) name() mac()
nasId() nasPort(4294967295) change(0x71)

2023-05-24T17:41:40.268 [rest] [172.27.1.194:49428] Send HTTP response:
code(200) httpLength(70) contentLength(0) hdrExt(0)

2023-05-24T17:41:40.296 [rest] [172.27.1.194:49434] Received HTTP request:
method(POST) hdr(44/248) uri(/api/v1/subscribers/10.0.0.4) authorization
(Basic) contentLength(56) connection(0) transferEncoding(0x0)

2023-05-24T17:41:40.296 [subscriber] [10.0.0.4] Created subscriber: policy
(rest-static-1) sessionId() subscriberId(sub-4) customerId() name() mac()
nasId() nasPort(4294967295) change(0x7fff)

2023-05-24T17:41:40.296 [rest] [172.27.1.194:49434] Send HTTP response:
code(201) httpLength(75) contentLength(0) hdrExt(0)

…
```

The traces of the last requests and responses can be found in the directory /opt/bqn/var/trace:

```
bqn0:~ # ls -al /opt/bqn/var/trace/rest*

-rw-r--r-- 1 root root 318 May 24 15:14 rest-req-0000

-rw-r--r-- 1 root root 364 May 24 11:30 rest-req-0001

-rw-r--r-- 1 root root 358 May 24 11:30 rest-req-0002

. . .
```

No billing messages are received

If the billing system is configured and the **Billing** icon in the dashboard is in Red, the access to the Billing system is failing.

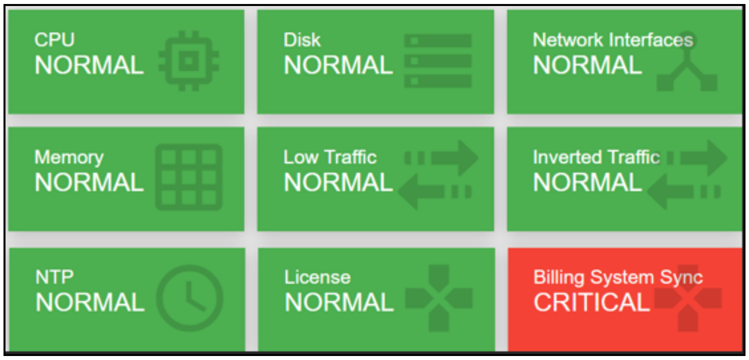

CPU NORMAL

Disk NORMAL

Network Interfaces NORMAL

Memory NORMAL

Low Traffic NORMAL

Inverted Traffic NORMAL

NTP NORMAL

License NORMAL

Billing System Sync CRITICAL

Figure 111: *The QoE dashboard*

Click the **Billing** icon to view the billing status page:

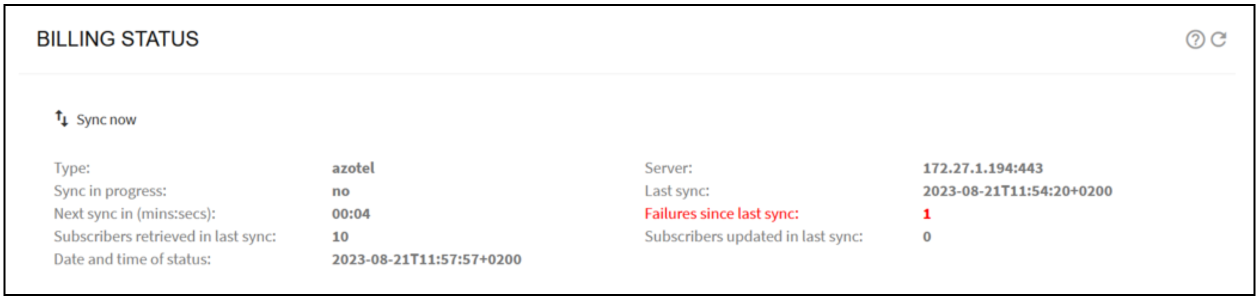

BILLING STATUS

Sync now

Type:	azotel	Server:	172.27.1.194:443
Sync in progress:	no	Last sync:	2023-08-21T11:54:20+0200
Next sync in (mins:secs):	00:04	Failures since last sync:	1
Subscribers retrieved in last sync:	10	Subscribers updated in last sync:	0
Date and time of status:	2023-08-21T11:57:57+0200		

In this example, there was a successful synchronization at 11:54:20 retrieving 10 subscribers, but there was one failed attempt afterwards. You can force a synchronization attempt by pressing the **Sync now** button.

If the failure remains, follow these steps:

- Verify that the billing configuration is correct (direction and credentials).
- Verify that the billing IP address is reachable from the QoE server:

```
bqnadm@bqn# net ping 192.168.0.122

PING 192.168.0.122 (192.168.0.122) 56(84) bytes of data.

64 bytes from 172.27.1.194: icmp_seq=1 ttl=64 time=0.169 ms

64 bytes from 172.27.1.194: icmp_seq=2 ttl=64 time=0.180 ms

64 bytes from 172.27.1.194: icmp_seq=3 ttl=64 time=0.152 ms

^C

bqnadm@bqn#
```

- Enable the API logs entering the QoE server shell via ssh:

```
bqnadm@bqn# configure

bqnadm@bqn(config)# api common

bqnadm@bqn(config-api)# event level general

bqnadm@bqn(config-api)# event level billing

bqnadm@bqn(config-api)# event level policy

bqnadm@bqn(config-api)# event level subscriber

bqnadm@bqn(config-api)# commit

bqnadm@bqn(config-api)# end

bqnadm@bqn#
```

Now, the event log between the QoE and the billing system is displayed. In the following example, the QoE connects to an Azotel billing and retrieves three policies associated to ten subscribers:

```
bqnadm@bqn# show api event log

2022-12-09T10:43:54.480 [billing] Sent HTTP POST request: uri
(/restapi/listCustomerBucketData) host(172.27.1.194:443)
```

```
2022-12-09T10:43:54.482 [policy] Updated "AZ-1000-500" policy: rate
(500/1000)

2022-12-09T10:43:54.482 [subscriber] Created 10.0.0.1 subscriber: policy
(AZ-1000-500) block(no) customerId(10) name(Subscriber_number_10) nickname
()

2022-12-09T10:43:54.482 [policy] Updated "AZ-2000-1000" policy: rate
(1000/2000)

2022-12-09T10:43:54.482 [subscriber] Created 10.0.0.2 subscriber: policy
(AZ-2000-1000) block(no) customerId(11) name(Subscriber_number_11) nickname
()

2022-12-09T10:43:54.482 [policy] Updated "AZ-3000-1500" policy: rate
(1500/3000)

2022-12-09T10:43:54.482 [subscriber] Created 10.0.0.3 subscriber: policy
(AZ-3000-1500) block(no) customerId(12) name(Subscriber_number_12) nickname
()

2022-12-09T10:43:54.482 [subscriber] Created 10.0.0.4 subscriber: policy
(AZ-1000-500) block(no) customerId(13) name(Subscriber_number_13) nickname
()

2022-12-09T10:43:54.482 [subscriber] Created 10.0.0.5 subscriber: policy
(AZ-2000-1000) block(no) customerId(14) name(Subscriber_number_14) nickname
()

2022-12-09T10:43:54.482 [subscriber] Created 10.0.0.6 subscriber: policy
(AZ-3000-1500) block(no) customerId(15) name(Subscriber_number_15) nickname
()

2022-12-09T10:43:54.482 [subscriber] Created 10.0.0.7 subscriber: policy
(AZ-1000-500) block(no) customerId(16) name(Subscriber_number_16) nickname
()

2022-12-09T10:43:54.482 [subscriber] Created 10.0.0.8 subscriber: policy
(AZ-2000-1000) block(no) customerId(17) name(Subscriber_number_17) nickname
()

2022-12-09T10:43:54.482 [subscriber] Created 10.0.0.9 subscriber: policy
(AZ-3000-1500) block(no) customerId(18) name(Subscriber_number_18) nickname
()

2022-12-09T10:43:54.482 [subscriber] Created 10.0.0.10 subscriber: policy
(AZ-1000-500) block(no) customerId(19) name(Subscriber_number_19) nickname
()

2022-12-09T10:43:54.482 [billing] Updated 10/10 billing subscribers

bqnadm@bqn#
```

If there is an error, the log indicates the current status:

- It is also possible to extend the logging to detailed traces of the requests and responses exchanged between the QoE and the billing system:

```
bqnadm@bqn# configure

bqnadm@bqn(config)# api billing

bqnadm@bqn(config-api)# trace request 5

bqnadm@bqn(config-api)# trace response 5
```

```
bqnadm@bqn(config-api)# commit

bqnadm@bqn(config-api)# end

bqnadm@bqn#
```

- QoE generates traces of the last five request and responses between QoE and the billing system. The traces can be found in directory /opt/bqn/var/trace:

```
bqn0:~ # ls -al /opt/bqn/var/trace/billing*

-rw-r--r-- 1 root root 224 Apr 13 18:43 /opt/bqn/var/trace/billing-
req-0000

-rw-r--r-- 1 root root 224 Apr 13 18:43 /opt/bqn/var/trace/billing-
req-0001

-rw-r--r-- 1 root root 1955 Apr 13 18:43 /opt/bqn/var/trace/billing-
rsp-0000

-rw-r--r-- 1 root root 1955 Apr 13 18:43 /opt/bqn/var/trace/billing-
rsp-0001

...
```

- As a last report, send a request directly from QoE to the billing system using UNIX curl command following the billing system API conventions. For example, to send a query to an Azotel billing, execute the following command:

```
PS C:\Users\myuser> ssh root@bqn

bqn:~ # curl -H 'Accept: application/json' -H 'Content-Type:
application/json' -X POST -d '{"api_username": "myuser", "api_
password": "mypassword", "allcustomers": "1"}'
https://demo.azotel.com/restapi/listCustomerBucketData

{"ip":"10.10.0.12","result":0}

bqn:~ #
```

No NTP servers synchronized

To keep the server clock with an accurate time, QoE requires the NTP service. Some public NTP servers are configured by default.

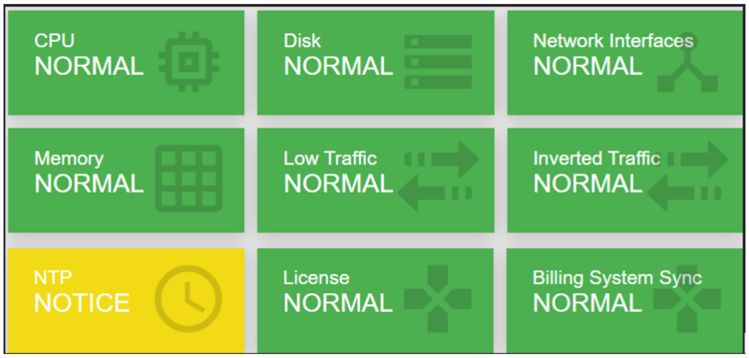

If the NTP servers are configured, but the QoE cannot synchronize with any of them, then the NTP icon in the dashboard is in orange color.

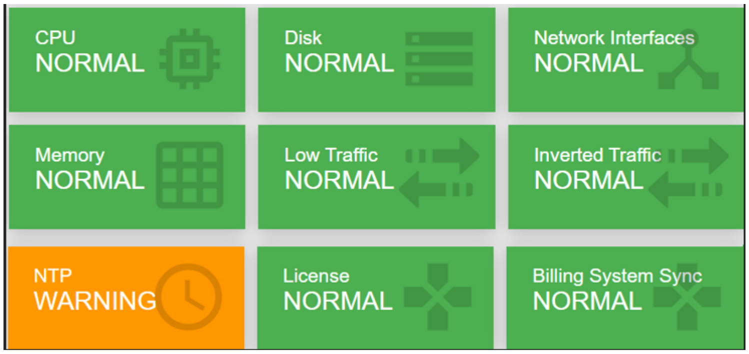

CPU NORMAL

Disk NORMAL

Network Interfaces NORMAL

Memory NORMAL

Low Traffic NORMAL

Inverted Traffic NORMAL

NTP WARNING

License NORMAL

Billing System Sync NORMAL

Note

There is a change in the system time because of lack of NTP synchronization. It may leads to brief service losses while the system adjusts. To avoid this, always have atleast one NTP server in sync.

To check the list of configured NTP servers, navigate to **Administration > System Date > NTP Servers**.

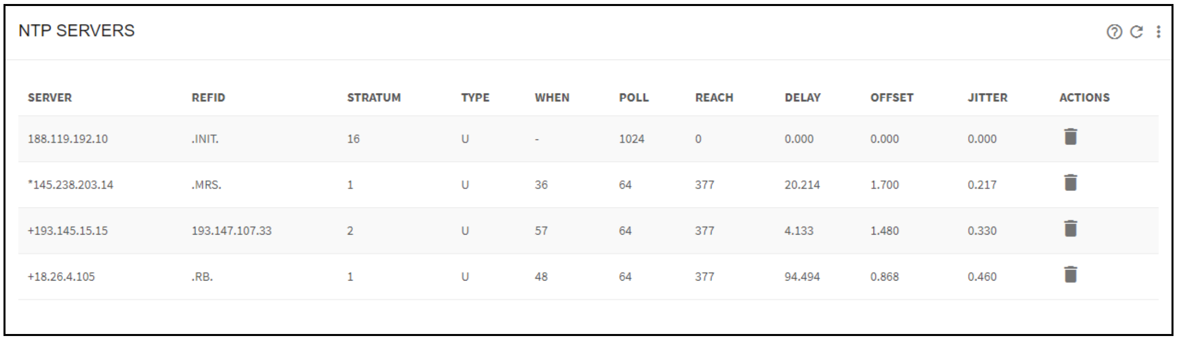

NTP SERVERS

SERVER	REFID	STRATUM	TYPE	WHEN	POLL	REACH	DELAY	OFFSET	JITTER	ACTIONS
188.119.192.10	.INIT.	16	U	-	1024	0	0.000	0.000	0.000	🗑
*145.238.203.14	.MRS.	1	U	36	64	377	20.214	1.700	0.217	🗑
+193.145.15.15	193.147.107.33	2	U	57	64	377	4.133	1.480	0.330	🗑
+18.26.4.105	.RB.	1	U	48	64	377	94.494	0.868	0.460	🗑

Minimum one NTP server must be synchronized. In the example above, the NTP server 145.238.203.14 is chosen for clock synchronization (indicated by the * next to the server IP address) and contacted 36 seconds ago (column WHEN).

To solve this issue, if you have a local NTP server, then click ⋮ menu icon, add it to the list and select **Add Server...**

If you have no local NTP servers, ensure that the UDP port 123 is open from QoE management IP to the Internet and vice-versa, including in the QoE Firewall (if activated).

Disk issues

To check the status of the system disk, click **DISK** icon on the dashboard. Figure 112 shows the disk usage page.

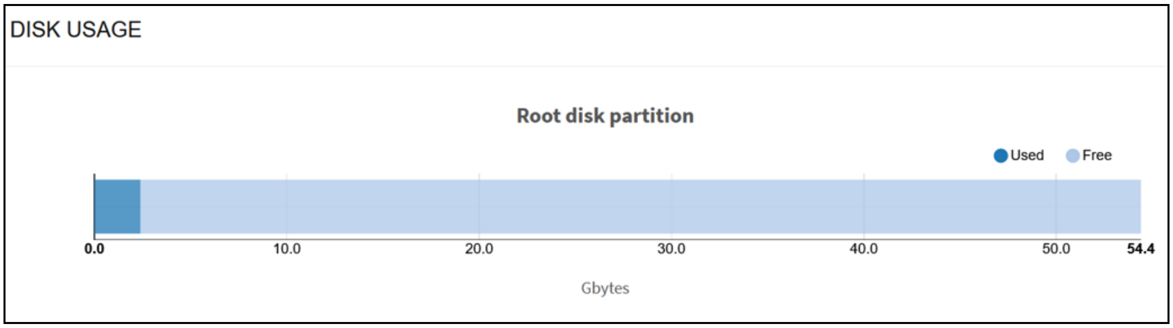

Figure 112: *The disk usage page*

When less than 15% of the disk storage is free, the disk icon is in orange (warning state).

Cambium Networks

Cambium Networks delivers wireless communications that work for businesses, communities, and cities worldwide. Millions of our radios are deployed to connect people, places and things with a unified wireless fabric that spans multiple standards and frequencies of fixed wireless and Wi-Fi, all managed centrally via the cloud. Our multi-gigabit wireless fabric offers a compelling value proposition over traditional fiber and alternative wireless solutions. We work with our Cambium certified ConnectedPartners to deliver purpose-built networks for service provider, enterprise, industrial, and government connectivity solutions in urban, suburban, and rural environments, with wireless that just works.

User Guides	http://www.cambiumnetworks.com/guides
Technical training	https://learning.cambiumnetworks.com/learn
Support website (enquiries)	https://support.cambiumnetworks.com
Main website	http://www.cambiumnetworks.com
Sales enquiries	solutions@cambiumnetworks.com
Warranty	https://www.cambiumnetworks.com/support/standard-warranty/
Telephone number list	http://www.cambiumnetworks.com/contact-us/
Address	Cambium Networks Limited, Unit B2, Linhay Business Park, Eastern Road, Ashburton, Devon, TQ13 7UP United Kingdom

Cambium Networks™ www.cambiumnetworks.com