Cambium Networks™

USER GUIDE

QoE Appliance

Release 4.20

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

Overview

This chapter contains the following sections:

- Audience
- Conventions
- Functionality

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 > for the menu navigation. For example, **Status** > **System** > **System Information**.

Command line interactions are displayed inside a box, in fixed-width font and using bold for the commands entered by the user:

```
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
```

Functionality

The QoE is an advanced network optimization product to improve speed, reduce latency and congestion, and provide full traffic visibility. This user manual describes in detail all the product functionalities and how to make the best use of them. In this section, we provide an overview with links to some interesting sections where to start.

Access the QoE GUI

Once the system has been installed and deployed, see here how to access its user interface.

Traffic optimization: TCP Optimization

The TCP Optimization functionality makes sure that TCP traffic flowing through it reaches optimal speeds, resulting in faster downloads and uploads.

This feature is enabled by default. The benefits in the traffic can be found here. You will have to wait a few hours for the system to gather enough information.

Traffic optimization: Automatic Congestion Management (ACM)

The ACM continuously monitors all network users, their speed, latency and packet losses, and detect if they are reaching a congestion. When congestion is detected, it will take action to mitigate it and reduce latency and losses, providing a much better Quality-of-Experience.

This feature is enabled by default. The benefits in the traffic can be found here. You will have to wait a few hours for the system to gather enough information.

Bandwidth management policies

The QoE platform includes the most advanced queuing technology in the market, which delivers the best possible Quality of Service (QoS) and Quality of Experience (QoE). You can enforce rate plans or per-application bandwidth limitations in an extremely flexible way.

You have some examples of policies here.

The QoE has several ways to get the bandwidth policies from external systems, such as RADIUS, using our REST API or some of the supported billing connectors.

Network visibility

The QoE platform provides you the key metrics to monitor the quality you are delivering to your users and the quality you are getting from the service providers connecting you to the Internet.

Explore what is available here.

Hardware requirements

This chapter contains the following sections:

- General requirements
- Hardware dimensioning
- Virtual platforms

General requirements

The QoE uses dedicated commercial off-the-shelf servers or virtual machines, configured according to the network capacity and connectivity requirements.

See rest of sections for more details

Supported CPUs

- Intel Xeon CPUs (Nehalem or later)
- AMD Epyc CPUs

Dual-CPU servers are supported. Following normal NUMA requirements, RAM must be balanced between the two CPUs (for example, 16GB + 16 GB).

Currently, the server maximum number of CPU cores supported is 256. It requires bqn R4.18 or more, with bqnkernel-R3.0.13 or later. For up to 128 cores, previous bqnkernel and bqn releases can be used.

Supported hard disks

Solid State Drives (SSDs) are recommended for performance and reliability reasons.

The following disk types are supported:

- SATA
- SAS
- NMVe

Supported network interfaces

A QoE server needs at least three network ports: one for management and another two for packet processing.

Ports for packet processing should be Intel-based, with one of the following controllers:

- 1 Gbps
  - Intel I210
  - Intel I350
- 2.5 Gbps
    - Intel i226-V (requires bqnkernel-R3.0.13 or later).

- 10 Gbps
  - Intel X520
  - Intel X540
  - Intel X550
  - Inter X553
  - Intel X710
  - Intel XL710
- 25 Gbps
  - Intel XXV710
  - Intel E810
- 40 Gbps
  - Intel XL710
- 100 Gbps
  - Intel E810

Other network interface models can be supported using pcap mode, but with much lower throughput capacity (up to 1 Gbps).

Supported network interface transceivers

For optical interfaces, the transceivers must be Intel-compatible of one of the following types:

- 1G SFP
  - 1000BASE-SX
  - 1000BASE-LX
- 10G SFP+
  - 10GBASE-SR/1000BASE-SX
  - 10GBASE-LR/1000BASE-LX
- 25G SFP28
- 40G QSFP+
  - 40GBASE-SR4
  - 40GBASE-LR4
- 100G QSFP28

Hardware dimensioning

The minimum hardware configuration is for 1Gbps to 150 Gbps. The following table summarizes the CPU, RAM and disk required depending on the network capacity. The processors shown are examples of verified systems. Processors with similar performance and characteristics are also recommended.

	CPU	RAM	Disk
1 Gbps	i5/ i7/E3-1220/Atom C3558	8 GB	60 GB
10 Gbps	E3-1240/E-2334	32 GB	120 GB
20 Gbps	Xeon 4214	64 GB	240 GB
40 Gbps	2x Xeon E4214	128 GB	240 GB
100 Gbps	2xAMD 7532 Rome	256 GB	480 GB
150 Gbps	2xAMD 75F3 Milan	512 GB	480 GB

Virtual platforms

QoE supports the following virtual platforms:

- VMware
- KVM
    - Linux kernel of the host machine should be version 4.11 or later.
    - QEMU should be version 2.9 or later.

All resources are dedicated (pinned) to the virtual machine (no over subscription). Depending on the traffic load, check with Cambium Networks for the required resources. As a general guideline, use the following resources:

	vCPU*	RAM	Disk
1 Gbps	2	8 GB	60 GB
10 Gbps	14	32 GB	120 GB

* Each vCPU is equivalent to one core of an Intel Xeon E5-2630 v4 @ 2.20GHz CPUs, with hyper threading enabled.

For the data plane interfaces, the supported configuration is using Intel network cards with **PCI passthrough**, because of performance and reliability reasons.

Software installation

Before starting with the installation, check that the server meets the hardware requirements. After that, follow the steps described in this section.

For more information about the installation, refer to *QoE Quick Start Guide*.

This chapter contains the following sections:

- Generate bootable drive
- BIOS settings
- Installing the software from bootable drive
- Post installation steps
- Installation in KVM
- Installation in VMware

Generate bootable drive

A new installation requires the server to boot an ISO image (QoE Installation Image). The ISO image contains the installation programs, operating system, and product software.

Get the latest iso file from Cambium Networks Support Site.

The QoE ISO is already bootable. To burn the ISO file into a USB drive:

- In **Linux**, use the dd command:

  ```
  # dd if=path-to-bqn-iso of=/dev/usb-device
  ```

- Example (USB driver in /dev/sdb and displaying progress information):

  ```
  # dd status=progress if=./bqnos-R4.9.2-20230120.iso of=/dev/sdb
  ```

- In **Windows**, use Rufus (https://rufus.ie). Select MBR partition scheme as shown in Figure 1 and select DD image mode as shown in Figure 2.

Figure 1: *Select MBR partition scheme*

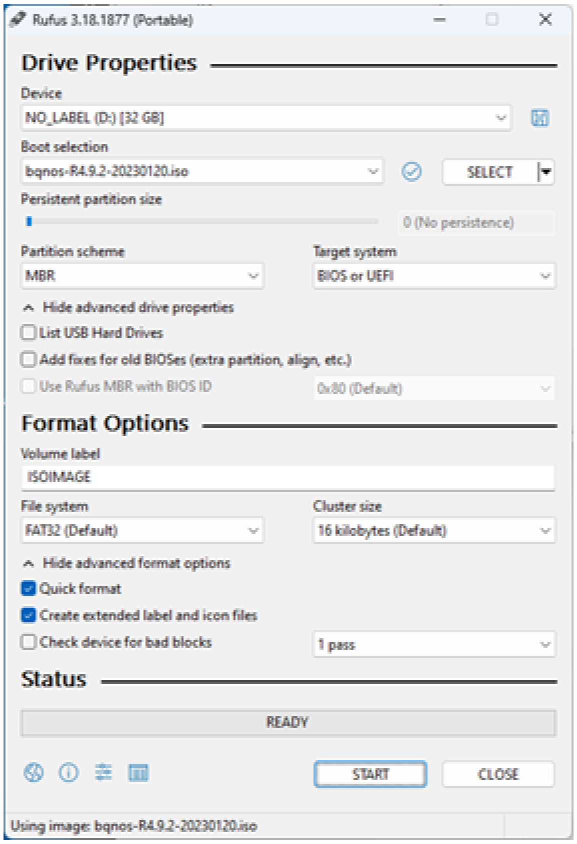

Rufus 3.18.1877 (Portable)

Drive Properties

Device: NO_LABEL (D:) [32 GB]

Boot selection: bqnos-R4.9.2-20230120.iso

Persistent partition size: 0 (No persistence)

Partition scheme: MBR

Target system: BIOS or UEFI

Hide advanced drive properties

□ List USB Hard Drives

□ Add fixes for old BIOSes (extra partition, align, etc.)

□ Use Rufus MBR with BIOS ID: 0x80 (Default)

Format Options

Volume label: ISOIMAGE

File system: FAT32 (Default)

Cluster size: 16 kilobytes (Default)

Hide advanced format options

☑ Quick format

☑ Create extended label and icon files

□ Check device for bad blocks: 1 pass

Status

READY

Using image: bqnos-R4.9.2-20230120.iso

Figure 2: *Select DD image mode*

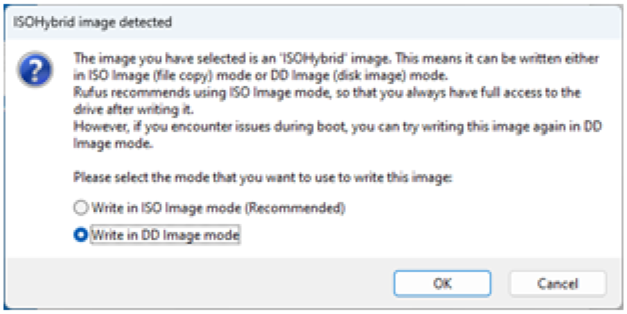

ISOHybrid image detected

The image you have selected is an 'ISOHybrid' image. This means it can be written either in ISO Image (file copy) mode or DD Image (disk image) mode.
Rufus recommends using ISO Image mode, so that you always have full access to the drive after writing it.
However, if you encounter issues during boot, you can try writing this image again in DD Image mode.

Please select the mode that you want to use to write this image:

○ Write in ISO Image mode (Recommended)

● Write in DD Image mode

Alternatively, in servers with a Lights-Out Management (LOM) interface (such as HPE iLO, Dell iDRAC or Supermicro IPMI), the ISO file can be mounted directly as a virtual drive.

BIOS settings

To access the BIOS/System settings, press a special key during the Power-On Self Test (POST) sequence (F2, F9, DEL or some other key displayed on the screen).

This section describes the basic settings to check in the server BIOS.

For more information, refer to *QoE Quick Start Guide*.

Boot mode

The server should have its factory default boot mode (for example, DUAL).

ISO supports both legacy/MBR/BIOS mode or UEFI mode. Some servers only support UEFI boot mode, in some servers, either mode is available.

> Note
>
> If UEFI mode is selected as the boot mode, ensure that **Secure Boot** is **disabled** in the BIOS/System settings.

Server clock

The server clock must be set to UTC (coordinated universal time) in the BIOS. The current UTC time can be obtained at: https://www.timeanddate.com/worldclock/timezone/utc. Verify both date and time of date settings.

> Note
>
> Configure the time zone later during QoE ISO installation.

Disk setup

The RAID controller requires available disks to be defined in the RAID controller configuration, even if there is only one disk or if no RAID redundancy scheme is used. Therefore, prior to the QoE installation process, it is necessary to create a logical or virtual drive inside the RAID controller. In addition, the logical or virtual drive must be marked as a bootable disk.

The RAID controller configuration is available during the server Power-On Self Test (POST) sequence.

> Note
>
> If no RAID controller is available, skip this section.

The RAID type depends on the number of physical disks available:

- If only one physical disk: RAID 0.
- If two physical disks: RAID 1
- If four or more physical disks: RAID 1+0 (stripe of mirrors)

> Note
>
> If there is an odd number of physical disks, the remaining disk is used as a spare.

For more information, refer to *QoE Quick Start Guide*.

Installing the software from bootable drive

After the BIOS Settings setup steps, insert the USB drive. Start or reboot the system, and then select the one-time boot menu and choose the USB drive.

Once the QoE Installation image is loaded, the following screen is displayed:

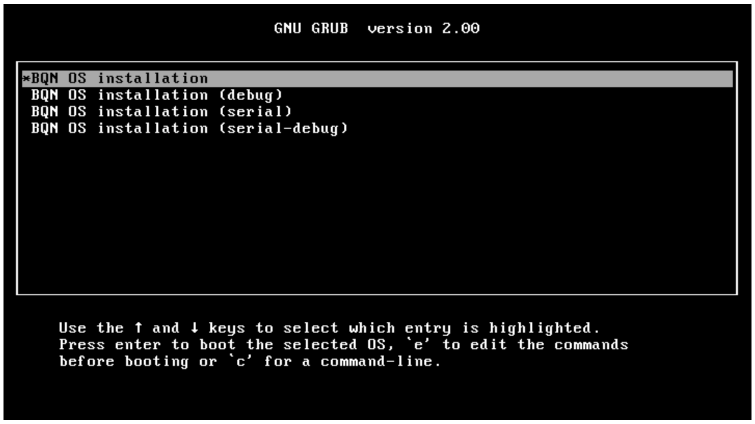

```
                    GNU GRUB  version 2.00

*BQN OS installation
 BQN OS installation (debug)
 BQN OS installation (serial)
 BQN OS installation (serial-debug)

    Use the ↑ and ↓ keys to select which entry is highlighted.
    Press enter to boot the selected OS, `e' to edit the commands
    before booting or `c' for a command-line.
```

Select the first option and press *ENTER*, QoE OS installation is started (if none of the option is selected, then the first option is selected automatically after ten seconds).

The following are some notes on navigating the ISO installation screens:

- Use the right and left arrows or the TAB key to move between buttons.
- Press *ENTER* to select a button.
- Use the up and down arrow keys to move across a list.
- Press a letter to place the list on the first element starting by that letter.
- Press the space bar to select a list item.

Disk selection

If more than one disk is available, a menu with the available options is shown (else the installation navigates to the partition window, see next section). Figure 3 shows the disk selection screen.

Figure 3: *Disk selection screen*

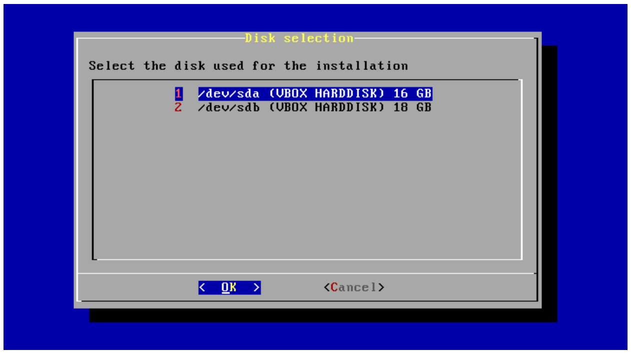

```
              Disk selection
 Select the disk used for the installation

          1  /dev/sda (VBOX HARDDISK) 16 GB
          2  /dev/sdb (VBOX HARDDISK) 18 GB

          <  OK  >          <Cancel>
```

Select a disk to install the QoE software (normally /dev/sda) and press *ENTER*.

Note

Do not select the USB drive for the installation if it is listed (the USB drive is the source of the software, not the installation target).

Disk partitions

Select the **Default** partition scheme in all cases, except if the server is not new and it has a previous installation that did not use DOS or MBR disk label format. Figure 4 shows the disk partition screen.

Figure 4: *Disk partition screen*

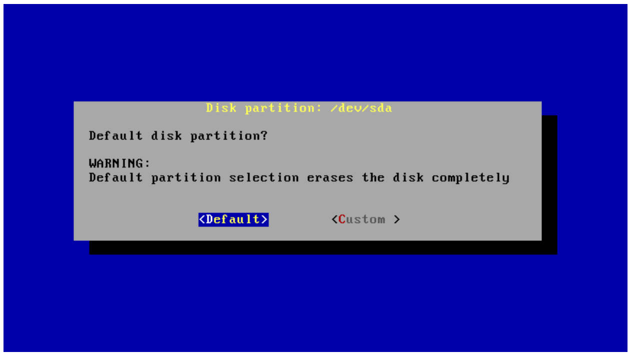

```
              Disk partition: /dev/sda

Default disk partition?

WARNING:
Default partition selection erases the disk completely

               <Default>        <Custom >
```

The default partition scheme uses the whole disk and erases all the existing data. This partition scheme creates the following two partitions:

- A swap partition, with a size of half the amount of physical RAM memory. The default swap size is not bigger than 64 GB or 25% of the total disk space.

- A QOE OS Linux partition, sized using all the remaining disk space.

Custom partition

If a custom partition is needed (a server with a previous non-DOS/MBR installation), refer to *QoE Quick Start Guide*. It is recommended to use the default partition option.

Keyboard map selection

US keyboard is selected by default. To select a different keymap, move in the list (for example, press **e** plus three arrow downs places the list on Spanish **es** keymap). Select the list item with the spacebar and press **ENTER**. Figure 5 shows the keymap selection screen

Figure 5: *Keymap selection screen*

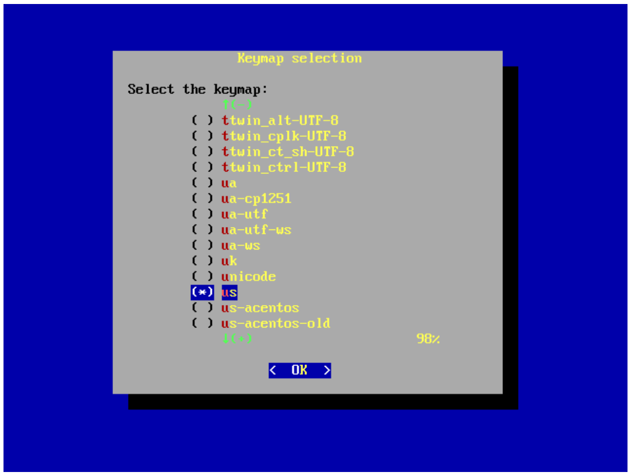

```
              Keymap selection

Select the keymap:
             ↑(-)
        ( ) ttwin_alt-UTF-8
        ( ) ttwin_cplk-UTF-8
        ( ) ttwin_ct_sh-UTF-8
        ( ) ttwin_ctrl-UTF-8
        ( ) ua
        ( ) ua-cp1251
        ( ) ua-utf
        ( ) ua-utf-ws
        ( ) ua-ws
        ( ) uk
        ( ) unicode
        (*) us
        ( ) us-acentos
        ( ) us-acentos-old
             ↓(+)                         98%

                    <  OK  >
```

Host name

Enter the hostname of the QoE server (this can be changed later during the post installation procedure or, after that, using QoE UI). Figure 6 shows the hostname selection screen.

Figure 6: *Hostname selection screen*

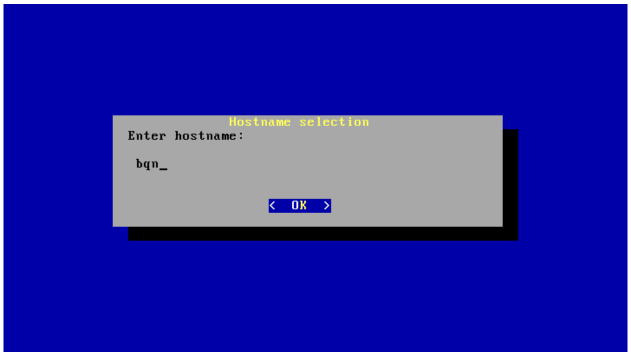

```
                 Hostname selection
 Enter hostname:

  bqn_

                    <  OK  >
```

Time zone selection

Central European Time (CET) time zone is selected by default. To select a different time zone, move in the list, select with the spacebar and press **ENTER** (this can be changed later during the post installation procedure or, after that, using QoE UI). Figure 7 shows the timezone selection screen.

Figure 7: *Timezone selection screen*

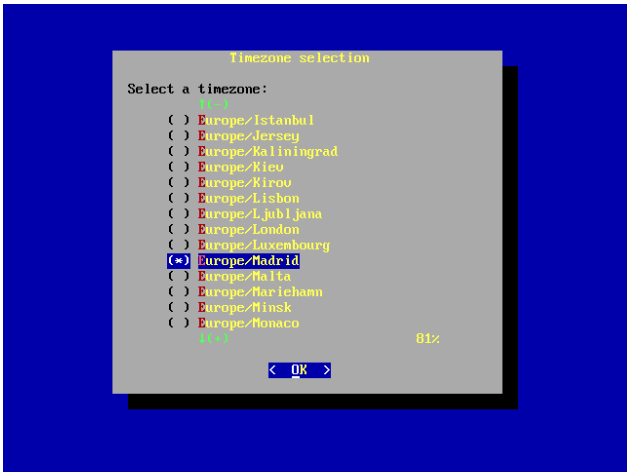

```
                 Timezone selection

 Select a timezone:
          ↑(-)
      ( ) Europe/Istanbul
      ( ) Europe/Jersey
      ( ) Europe/Kaliningrad
      ( ) Europe/Kiev
      ( ) Europe/Kirov
      ( ) Europe/Lisbon
      ( ) Europe/Ljubljana
      ( ) Europe/London
      ( ) Europe/Luxembourg
      (*) Europe/Madrid
      ( ) Europe/Malta
      ( ) Europe/Mariehamn
      ( ) Europe/Minsk
      ( ) Europe/Monaco
          ↓(+)                      81%

                    <  OK  >
```

Root password

It is important that keep the root password in a safe place to log into the QoE server during the post installation. The root password can be changed then.

Enter the root password as shown in Figure 8.

Figure 8: *Enter the root password*

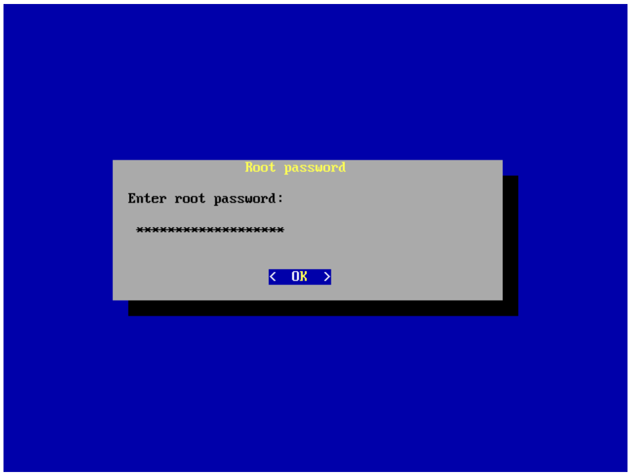

```
Root password

Enter root password:

 *******************
```

Reenter the root password as shown in Figure 9.

Figure 9: *Reenter the root password*

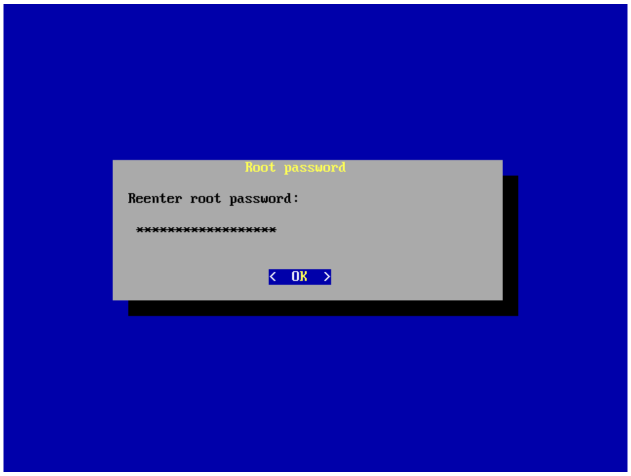

```
Root password

Reenter root password:

 ******************
```

Install boot loader

Select **Install** to install the boot loader at Master Boot Record (MBR) as shown in Figure 10.

Figure 10: *Install boot loader at MBR*

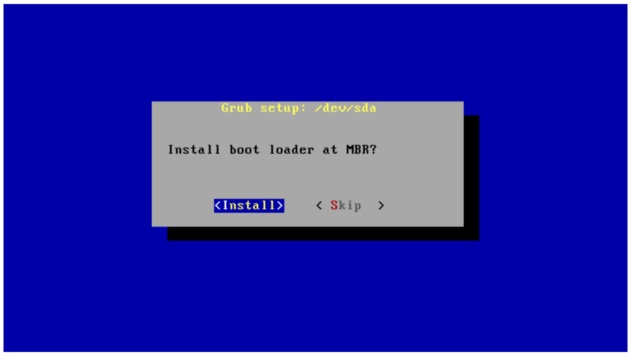

```
Grub setup: /dev/sda

Install boot loader at MBR?
```

Configuration of the QoE management interface

Follow the steps to configure the QoE management interface:

1. Select **Yes** from the QoE setup screen to start configuring the QoE management interface.

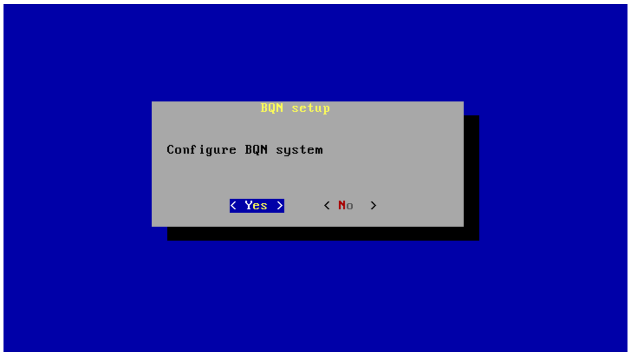

```
BQN setup

Configure BQN system
```

2. Select the **QoE management network interface**. The first port in the server motherboard is preferred (typically named en0o1) for the network interface. This can be changed later during the post installation procedure or after that using QoE UI.

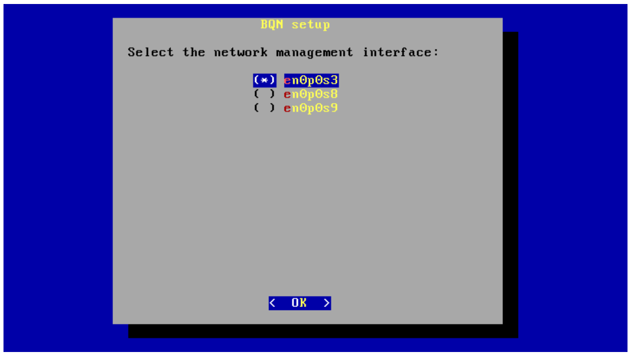

```
BQN setup

Select the network management interface:

(*) en0p0s3
( ) en0p0s8
( ) en0p0s9
```

3. Enter the QoE management IP address with its subnet mask (by default, 192.168.0.121/24):

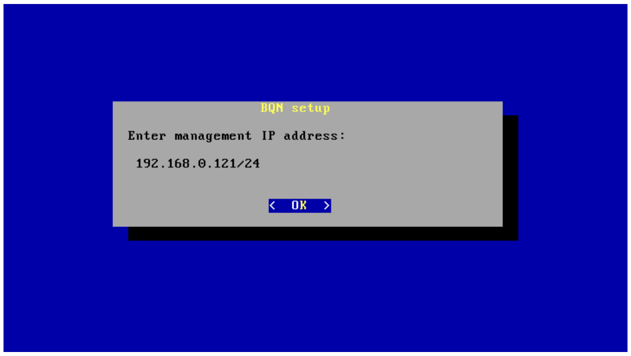

```
BQN setup

Enter management IP address:

 192.168.0.121/24
```

4. Select **Yes** to configure the default gateway and specify its address in the screen that follows (by default, 192.168.0.1):

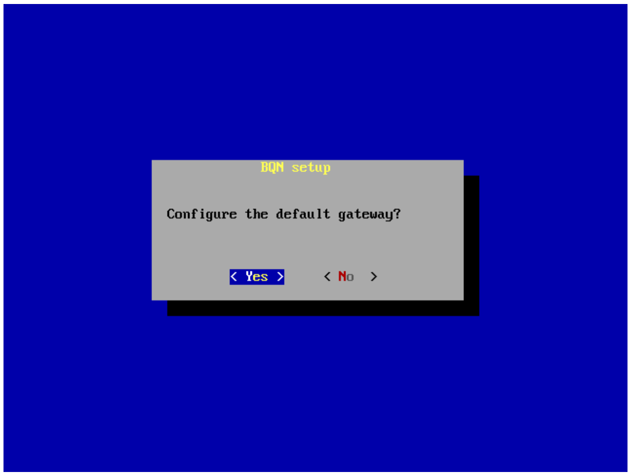

BQN setup

Configure the default gateway?

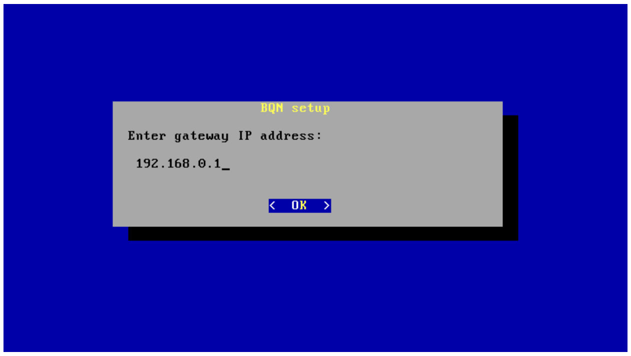

BQN setup

Enter gateway IP address:

192.168.0.1_

A dialog box appears after the successful configuration.

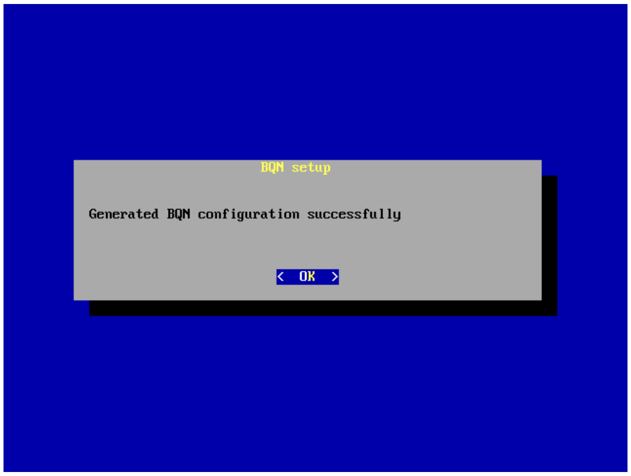

BQN setup

Generated BQN configuration successfully

The management settings can be changed later during the post installation procedure or, after that, using QoE GUI.

Reboot server

Select **Yes** to initiate the server reboot.

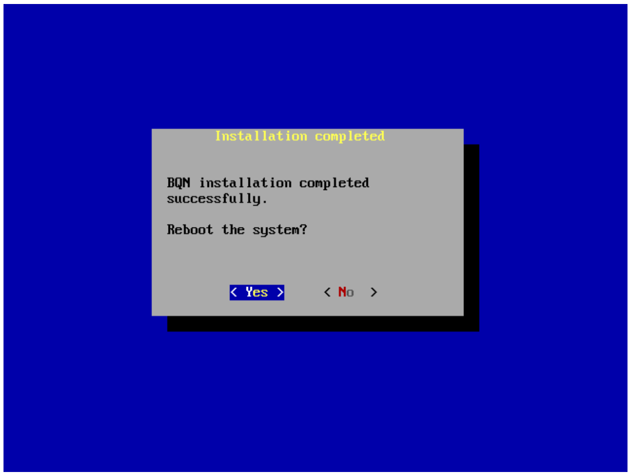

Installation completed

BQN installation completed successfully.

Reboot the system?

Remove the USB driver and select **Reboot**.

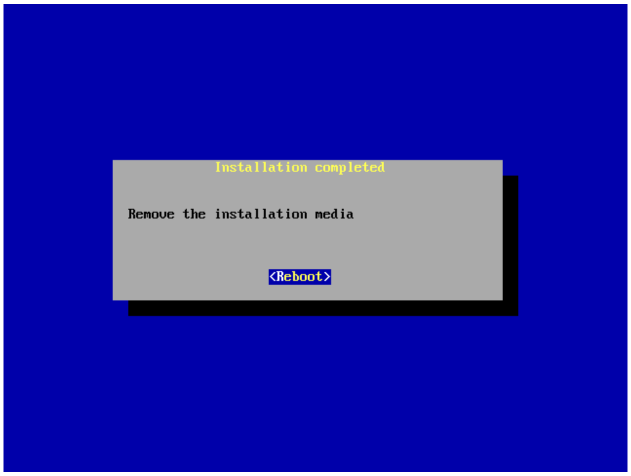

Installation completed

Remove the installation media

Post installation steps

After the software installation, reboot the system and then perform the following post installation procedures.

Configuration wizard - basic case

Log into the system as root, start the QoE shell with the bqnsh command, and run **wizard bta**, a wizard command generates an initial configuration, create the default administration users and enable the UI.

```
bqn01:~ # bqnsh

Copyright (c) 2009-2015 Bequant S.L.

root@bqn01# wizard bta

System vendor: Dell Inc.

System name: PowerEdge R230

System serial: ABCDEFG

CPU model: 12th Gen Intel(R) Core(TM) i7-12700H
CPU cores: 4
Management interface: en0p0s3
Management IP: 192.168.0.121/24
Management gateway: 192.168.0.1
Wire 1: en0p0s8(access)-en0p0s9(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]: yes
Set "bqnadm" user password to "ABCDEFG"
Set "bqnop" user password to "ABCDEFG"
root@bqn01#
```

Note

If the hardware model is supported by the wizard, the system serial number (ABCDEFG in this example) is used to set the passwords of the default administrator (**bqnadm**) and operator (**bqnop**) automatically. These passwords can be changed later (see Optional steps section).

Configuration wizard - interactive case

The wizard can also run in interactive mode to decide at each step of the configuration. You can change the management interface or IP address that is introduced during the installation process. *ENTER* key takes the value suggested by the wizard between square brackets.

In this example, the following changes are made to defaults:

- Register a DNS server (8.8.8.8).
- The management address is 10.0.0.23/24 on VLAN 12 with default gateway 10.0.0.1.

Note

After changing the management IP, log in again from the new network and to the new IP address.

```
1:~ # bqnsh
Copyright (c) 2009-2015 Bequant S.L.
root@bqn01# wizard bta interactive
Available network interfaces:
en0p0s3
en0p0s8
en0p0s9
```

```
Enter management interface [en0p0s3]:
Enable VLAN on management interface? (yes/no) [no]: yes
VLAN ID (1-4095): 12
Enter management IP address and prefix [192.168.0.121/24]: 10.0.0.23/24
Enter default gateway IP address [192.168.56.1]: 10.0.0.1
Configure a nameserver? (yes/no) [no]: 8.8.8.8
Available network interfaces:
en0p0s8
en0p0s9
Select access-side interface for wire: en0p0s8
Select internet-side interface for wire: en0p0s9
Enable SDR generation? (yes/no) [yes]:
Enter random optimization percentage [99]:
Enter random udr generation percentage [2]:
```

Update to the latest QoE package

Check for the latest QoE bpkg package from Cambium Networks and install it at this point. Follow the steps described in Updates within the same major release section.

Final verifications

UI Access

Verify that the QoE server UI is accessible through HTTPS using the management IP address and the **bqnadm** user just created (you cannot log into the UI as root) in the previous example, https://10.0.0.23.

License activation

To enable the different functionalities in QoE, a license is acquired. The first step to acquire a license is, connect QoE to the cloud license server. The management port must have an access to the Internet until the license is acquired. The dashboard on the UI displays the status of the license server connection.

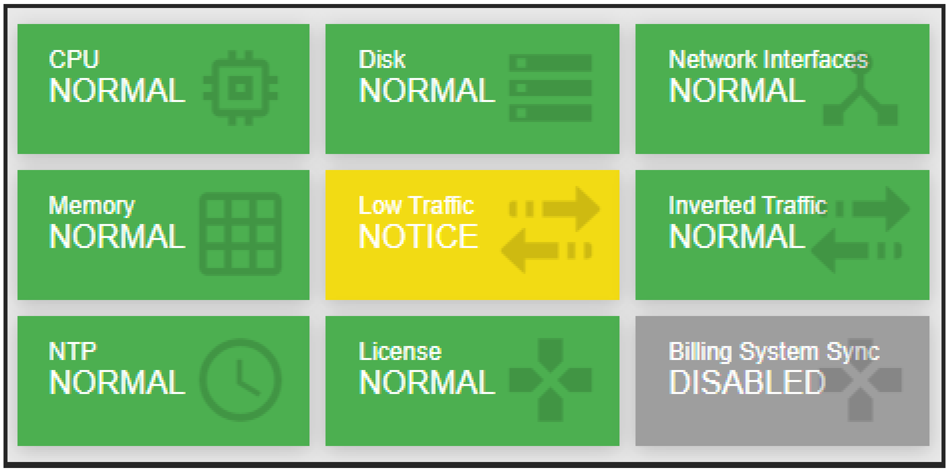

CPU NORMAL

Disk NORMAL

Network Interfaces NORMAL

Memory NORMAL

Low Traffic NOTICE

Inverted Traffic NORMAL

NTP NORMAL

License NORMAL

Billing System Sync DISABLED

- If the color is not in **Green**, then the connection is not established, and QoE does not register with the license server to acquire the license.
- If there is a firewall, then open TCP port 13152 for the IP addresses **146.59.206.4** (primary) and **46.26.190.166** (backup).

To check the connectivity with the cloud license server, perform the following steps:

1. ssh to the QoE as admin user (default bqnadm)
2. system shell
3. nc -vz 146.59.206.4 13152 (or for the backup server nc -vz 46.26.190.166 13152), this will check the TCP connectivity with the cloud license server.
4. exit (return to QoE shell)

NTP synchronization

QoE system clock must be correct at all times. Verify that the QoE server can access NTP servers. The NTP icon in the dashboard should be in normal state (it can take a few seconds to synchronize).

Additional information

See Troubleshooting section for a description of the state of the dashboard icons and procedure to get them to the normal state. For more information refer to *QoE Quick Start Guide*.

Installation in KVM

To run QoE with a KVM hypervisor, the following are the KVM software requirements:

- Linux kernel of the host machine should be version 4.11 or later.
- QEMU should be version 2.9 or later.

The following sections describe the configuration of the KVM. Use **virt-manager** the Virtual Machine Manager part of the **libvirt** package to interact with QEMU, but equivalent commands must be available in other popular managers.

For more information, refer to *QoE Quick Start Guide*.

Virtual Machine creation

To create a virtual machine, perform the following steps:

1. Copy the **QoE ISO** file to the host machine where the VM (virtual machine) to be created.
2. Select **New Virtual Machine**.
3. Under **Connection**, choose the host where the VM to be created (local or remote to the virt-manager).
4. Select **Choose ISO** or **CDROM** install media as the installation method.
5. Type the path of the QoE ISO file that is copied previously to the host machine.
6. Under **Choose the operating system** where you install, enter **openSUSE**.

7. Under **Choose memory and CPU setting**, select the required memory and the number of CPUs (see Virtual platforms section).

8. Select the required disk size.

9. Check the option **Customize configuration** before install.

Virtual disk

Select either **VirtIO** or **SATA** disk bus types.

Network interfaces

In the VM, a minimum of three interfaces must be created. The first interface is used for management and it is normally added during the VM creation. If it is not created, select **Add hardware** and choose **Network**.

The physical ports used for the data plane of the virtual machine must be Intel and configured in PCI-passthrough mode:

1. Verify that the CPU of the host server has the VT-d instruction set enabled. This is set under the server BIOS menu.

2. Select **intel_iommu enabled** for host OS kernel parameter.

3. Select **Add Hardware**.

4. Choose **PCI Host Device**.

5. Under **Host Device** select the network interface to be used in PCI passthrough.

Installation in VMware

To run QoE with a VMware hypervisor, create minimum three interfaces in the VM, one for management and two for data plane traffic.

- The management interface can be connected to a pre-existing virtual switch and used for the management of other VMs.

- The physical ports used for the data plane of the virtual machine must be Intel and configured in PCI-passthrough mode. Under **Host** > **Manage** > **Hardware** > **PCI** devices, select **Toggle Passthrough** in the two data plane ports.

- Create the virtual machine with Guest OS family equal to Linux and Guest OS version equal to SUSE Linux Enterprise 12.

- Under **Memory**, mark **Reserve all guest memory**.

- Navigate to **Add other device** > **PCI device** page, verify that the two physical ports previously placed in passthrough mode are available. Select the both physical ports.

- Select the required memory, CPUs and disk space (For more information, see Virtual platforms section).

- Configure **CD/DVD Media** of **Datastore ISO** file and select the QoE ISO file to install.

For more information, refer to *QoE Quick Start Guide*.

Deploying QoE in the network

This chapter contains the following sections:

- Choosing the right deployment location
- Network requirements
- Setting up a bypass
- Connecting the QoE to the network

Choosing the right deployment location

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

It is crucial to make sure that traffic through the QoE is symmetric: if 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 too. Otherwise, the QoE will not be able to limit the rate or do shaping for that subscriber (because it would not see all the traffic), and the TCP optimization could even block downloads towards that subscriber (because it would not see that some traffic has been 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.

Place the QoE closer or further from the subscribers, we 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 will then help us over that difficult hop. QoE can be installed on both ends of a very challenging transmission hop (like a satellite link), but we should also look at deployment options that result in a minimum number of QoE nodes.

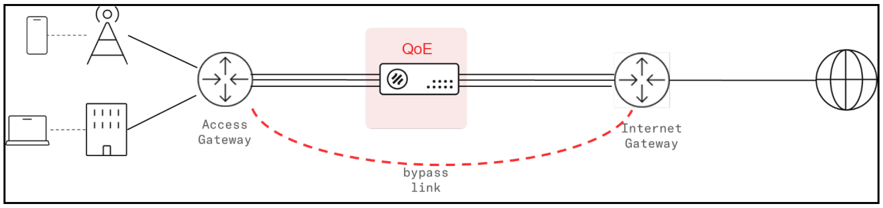

It is recommended that a bypass path is established between the neighboring nodes of the QoE (Access and Internet gateways in the diagram above), so if there is a failure in the active link or the QoE, the traffic is automatically steered through the bypass. See the section Setting up a bypass.

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 Cambium Networks license manager (For more information on IP addresses and ports, see Cambium Networks Support Site).
- Access from the QoE management IP address to the NTP servers (UDP port 123) configured.

Data path requirements

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

Traffic standards supported:

- 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.

Note

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

Aggregate links through the QoE

If the two ports of each wire are bridged transparently through the QoE, aggregate links can go through the QoE transparently over several wires. 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 messages (for example, ARP or fast LACP).

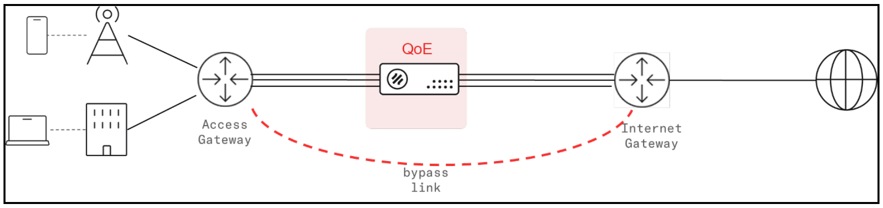

External bypass device

An external bypass device is connected to the external links and to the QoE. The device will trigger 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 by 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 11 shows the external bypass device.

Note

Currently, only Niagara devices are supported.

Figure 11: *External bypass device*

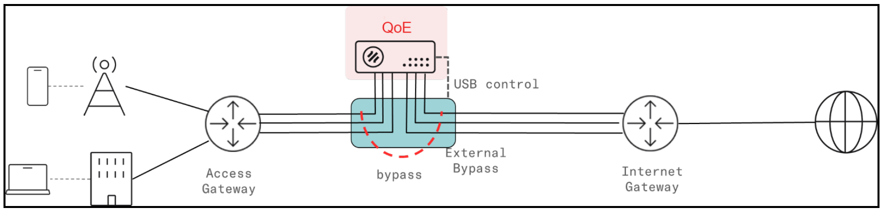

Network card with internal bypass

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.

To see the status of the card:

```
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 will trigger the bypass if there is a power outage (POWER-CFG on and, when there is no bypass, it will steer the traffic through the QoE Software (STEER-CFG on). POWER-LIVE and STER-LIVE report the current status on the network card and should be equal to the configured state. The interface pair should match a wire.

To activate the bypass manually (for example to perform a planned maintenance task on the QoE server), the following configuration changes must be done on 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 example, once 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#
```

It is also possible to disable the triggering of the bypass (for example, because an external bypass path based on dynamic routing is going to be used instead):

```
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, whose name is usually en0o1).

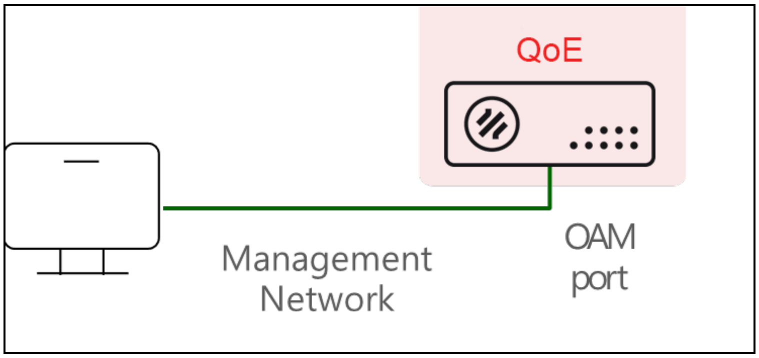

Switch the QoE server on and access the management GUI following the steps described in Logging in section.

The QoE factory management IP address is **192.168.0.121/24**, with default gateway **192.168.0.1**. The management user is **bqnadm**. 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, see the section Changing the management IP address.

Verify date and time

Verify whether 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. For more information on license manager IP addresses and ports, see Cambium Networks Support Site.

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 take some time to synchronize, so they may not be synchronized initially, but they should after a few hours.

Network interface mapping

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

- 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).
- Check in **Status** > **Interfaces** > **Link State** which interface changes the link state to up.
- Make sure you reload the page every time you change the physical connection, because the page does not reload automatically.

Connect subscriber traffic ports

Subscriber traffic interfaces are paired in wires, 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 is affected).

1. Use the network interface mapping to locate the ports of each wire defined in **Configuration** > **Interfaces** > **Data Wires**. Figure 12 shows the network interface mapping.

Figure 12: *Network Interface mapping*

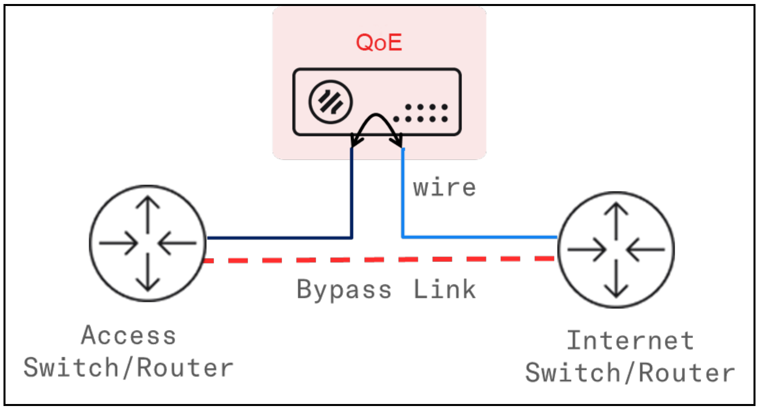

2. The network interfaces in use must be up and with the link detected. Both can be monitored in the GUI **Status** > **Interfaces** > **Link State**. Figure 13 shows the network interfaces page.

Figure 13: *Network Interfaces page*

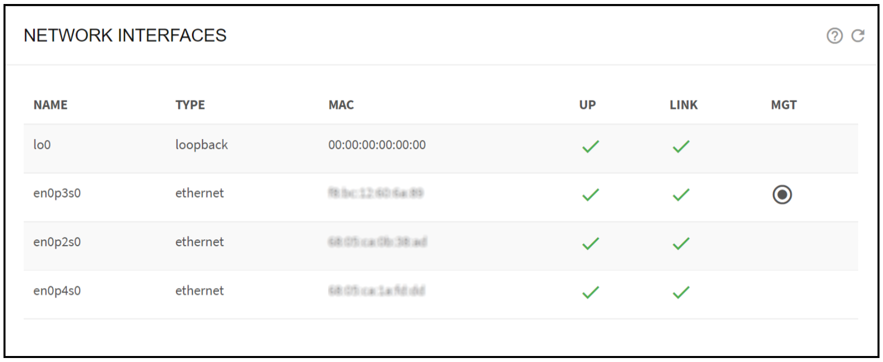

NETWORK INTERFACES

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

3. Once the traffic is steered, you can see its instant value in **Status** > **Interfaces** > **Throughput**. Figure 14 shows the current network interface throughput page.

Figure 14: *Current Network Interface Throughput page*

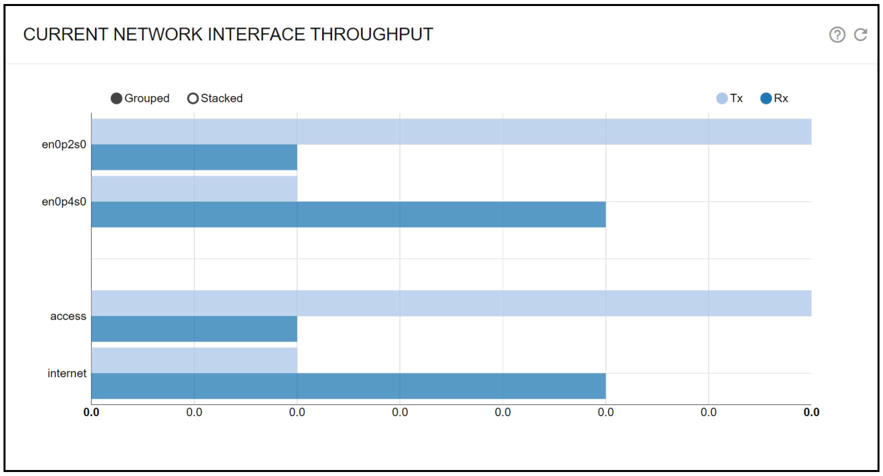

CURRENT NETWORK INTERFACE THROUGHPUT

The installation is complete.

Initial steps

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

The GUI has contextual help: press the ? help icon on the page for which help is needed.

This chapter contains the following sections:

- Logging in
- Changing the time
- Changing the management IP address
- Wire configuration
- IPMI configuration
- Remote access using port forwarding

Logging in

Login page

To access the GUI, open a browser and 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 will signal it as unsecure. Ignore the warning and go to the web page.

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

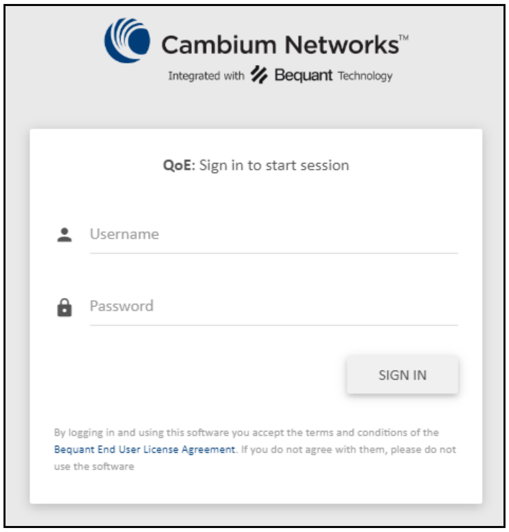

Cambium Networks™
Integrated with Bequant Technology

QoE: Sign in to start session

Username

Password

SIGN IN

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

Note

You cannot use the **root** username to log in to the UI.

Initial dashboard page

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

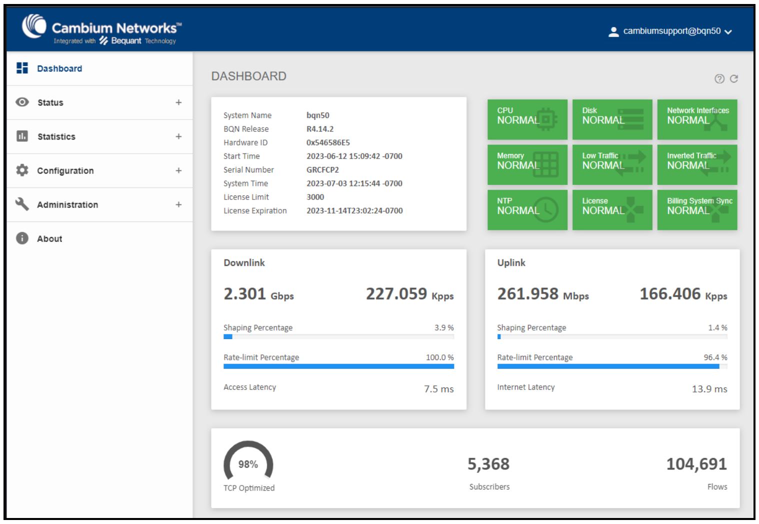

System Name	bqn50
BQN Release	R4.14.2
Hardware ID	0x546586E5
Start Time	2023-06-12 15:09:42 -0700
Serial Number	GRCFCP2
System Time	2023-07-03 12:15:44 -0700
License Limit	3000
License Expiration	2023-11-14T23:02:24-0700

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 dashboard must show all icons in green. The **Network Interfaces** icon will not be green until all the configured wires are connected (if there are interfaces not being used in any of the configured wires, it will remain in orange) and the icon **Low Traffic** will not be in green until the traffic flows through the QoE. Click on the icon for more information about the QoE status. See Billing systems for steps to take when an icon is not in normal state.

Tables

The data table displays the billing information of the subscribers. Figure 15 shows the data table.

Figure 15: *Data table*

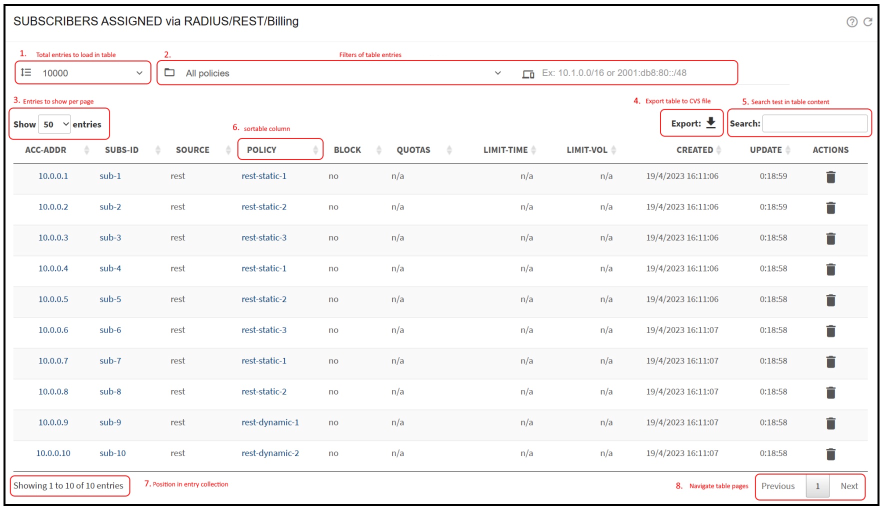
SUBSCRIBERS ASSIGNED via RADIUS/REST/Billing

1. Total entries to load in table
2. Filters of table entries
3. Entries to show per page
4. Export table to CVS file
5. Search test in table content
6. sortable column
7. Position in entry collection
8. Navigate table pages

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

The most important features are highlighted in Figure 15:

1. The first selector on the upper left defines how many entries to load in the table (10,000 in the example). The default value of each table is chosen to load all entries in normal circumstances. If the default is too low, a bigger number can be selected here.

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

3. The **show entries** define the page size.

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

5. The **search** field shows the table entries containing the substring specified in the text field.

6. Columns with arrows are sortable. To sort by a column, click on that column label. Clicking again will use reversed ordering.

7. A label on the lower left informs about where we are in the list of entries.

8. The buttons on the lower right navigate along the table pages.

When exporting tables to CVS, because a comma has special meaning, commands are scaped 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 is present, both rules are applied. For example, “a,b” is exported as “””a,b”””.

Changing the time

If it is necessary to change the system time, navigate to **Administration** > **System Date** > **Set Date & Time**. Figure 16 shows the system date and time page.

Figure 16: *The System date and time page*

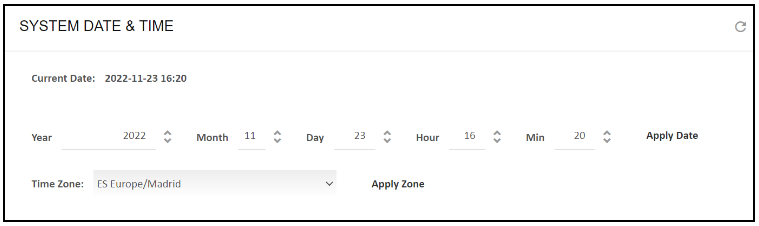

SYSTEM DATE & TIME

Current Date: 2022-11-23 16:20

Year 2022 Month 11 Day 23 Hour 16 Min 20

Time Zone: ES Europe/Madrid

Apply Date changes the local date and time, and **Apply Zone** changes 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, select on the lateral menu **Configuration** > **Interfaces** > **Management**. IP settings include the IP address and mask, the default gateway, and the VLAN identifier (if any). Figure 17 shows the management interface settings page.

Figure 17: *The Management Interface Settings page*

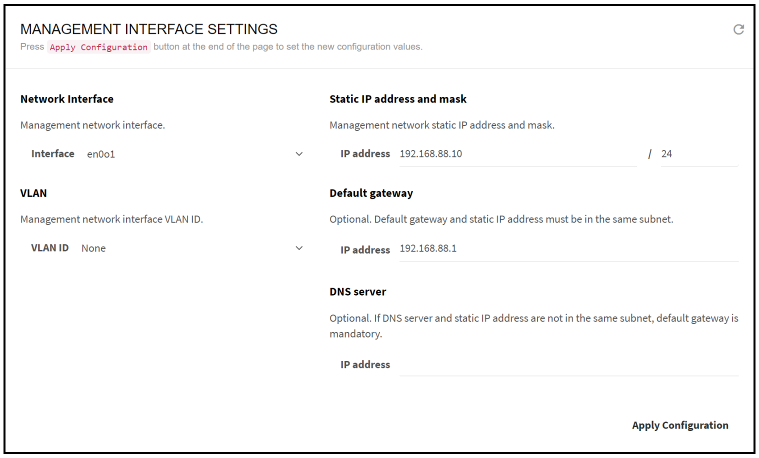

MANAGEMENT INTERFACE SETTINGS

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:

An optional DNS server IP address can be configured (needed if server names are used in the integration with a billing system).

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 into the GUI.

Wire configuration

A wire is a network interface pair processing subscriber traffic.

To configure wires, select **Configuration** > **Interfaces** > **Data Wires**.

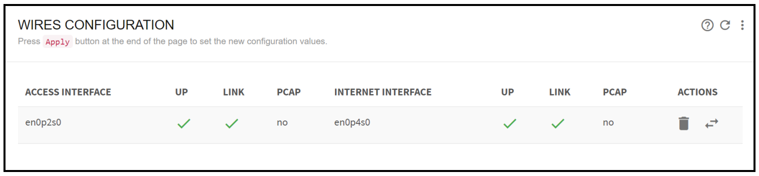

WIRES CONFIGURATION

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

ACCESS INTERFACE	UP	LINK	PCAP	INTERNET INTERFACE	UP	LINK	PCAP	ACTIONS
en0p2s0	✓	✓	no	en0p4s0	✓	✓	no	🗑 ⇆

Wires are directional, with the first network interface connected to the access towards the subscribers and the second interface on the Internet side. If a mistake has been made in connecting the ports, they can be swapped by clicking on the ⇆ arrows icon.

To add a wire, click on the ⋮ menu icon and select **Add Wire....**. A form allows you to select the access and Internet interfaces (the form lists the interfaces available). The **pcap** option is selected only for cards other than Intel (pcap allows compatibility at the price of reducing server performance).

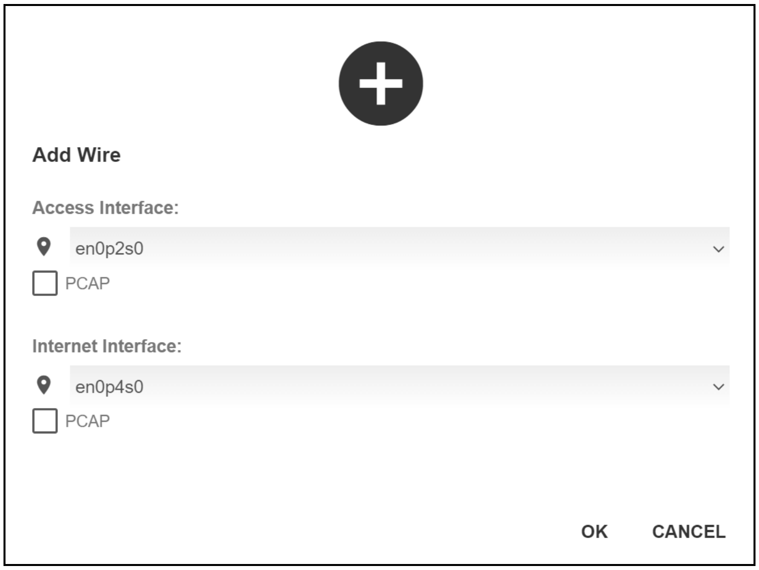

Add Wire

Access Interface:

en0p2s0

□ PCAP

Internet Interface:

en0p4s0

□ PCAP

To remove a wire, click the 🗑 delete icon. To modify a wire (for example, to enable the **pcap** option), you have to remove the existing wire and then, before applying the changes, add the wire with the change.

Note

Do not delete the wires unless indicated by the Cambium support personnel, as misconfigurations may lead to service loss.

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

IPMI configuration

Some servers have a lights-out module for general management (power on/off, hardware monitoring, and so on). Use 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). Similarly, follow the steps for SSH (port 22).

Navigate to: **IP > Firewall > Address Lists** and create an address list with all the source IP addresses that is allowed access to the QoE server (for example, qoe-allowed-src).

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

1. In the **General** tab, **Chain** is set to **dstnat**.
2. In the **General** tab, **Dst. Address** is set to the public IP address used to access the QoE from the Internet.
3. In the **General** tab, **Protocol** is set to TCP.
4. In the **General** tab, **Dst. Port** is set to the public IP address used to access the QoE from the Internet.
5. In the **Advanced** tab, **Src. Address List** is set to the name of the previously created list (qoe-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.

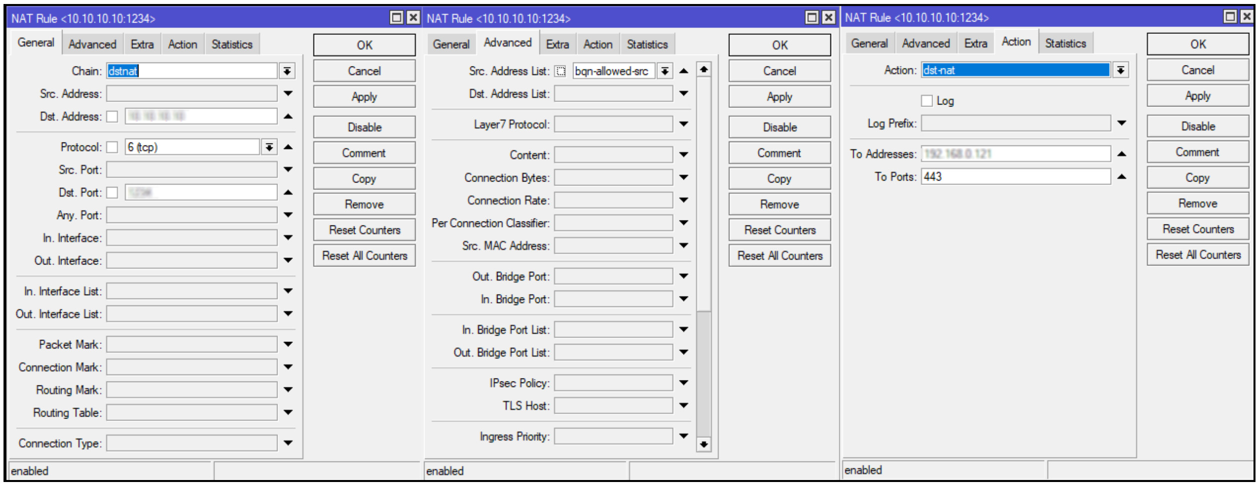

NAT Rule <10.10.10.10:1234>

General tab:

Chain: dstnat

Src. Address:

Dst. Address: □ [redacted]

Protocol: □ 6 (tcp)

Src. Port:

Dst. Port: □ [redacted]

Any. Port:

In. Interface:

Out. Interface:

In. Interface List:

Out. Interface List:

Packet Mark:

Connection Mark:

Routing Mark:

Routing Table:

Connection Type:

enabled

NAT Rule <10.10.10.10:1234>

Advanced tab:

Src. Address List: □ bqn-allowed-src

Dst. Address List:

Layer7 Protocol:

Content:

Connection Bytes:

Connection Rate:

Per Connection Classifier:

Src. MAC Address:

Out. Bridge Port:

In. Bridge Port:

In. Bridge Port List:

Out. Bridge Port List:

IPsec Policy:

TLS Host:

Ingress Priority:

enabled

NAT Rule <10.10.10.10:1234>

Action tab:

Action: dst-nat

□ Log

Log Prefix:

To Addresses: [redacted]

To Ports: 443

enabled

Management tasks

This chapter contains the following sections:

Updating the software

Updates within the same major release

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

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

> **Note**
>
> The software activation involves traffic interruption for some seconds. Therefore, it is advised to activate the software when the throughput is low.

The installation is performed in **Administration** > **Software**, by clicking on the ⋮ menu icon and selecting **Install...** A file selector pops up to choose the package, that it is transferred to the QoE server and installed.

The activation is done in **Administration** > **Software** by clicking on the ↻ cycle icon of the package to activate (highlighted in red as shown in the Figure 18).

This operation forces you to log back into the QoE after a few seconds, during which time, the traffic flow is interrupted.

Figure 18: *The Software Status page*

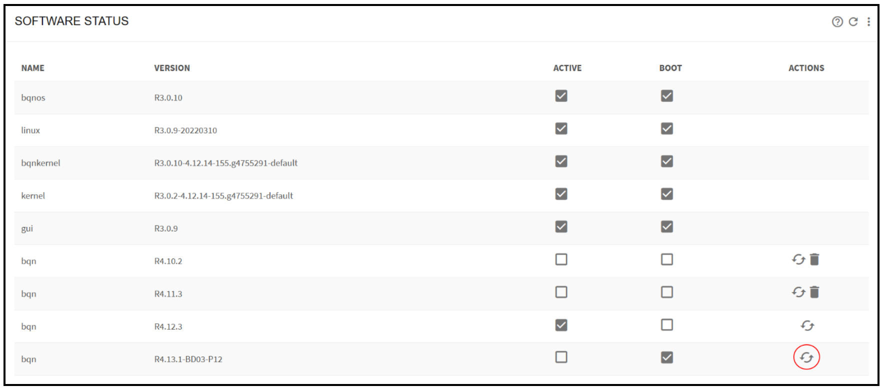

SOFTWARE STATUS

NAME	VERSION	ACTIVE	BOOT	ACTIONS
bqnos	R3.0.10	☑	☑	
linux	R3.0.9-20220310	☑	☑	
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	R4.10.2	□	□	⟳ 🗑
bqn	R4.11.3	□	□	⟳ 🗑
bqn	R4.12.3	☑	□	⟳
bqn	R4.13.1-BD03-P12	□	☑	⟳

Packages no longer in use (non-active and not selected for boot) can be removed with the dust bin icon.

Updates across major releases

For detailed instructions before performing an upgrade to a new major release, see Cambium Networks Support Site.

To migrate to a new major release (for example, from R3 to R4), the platform packages will need updates (bqnos, kernel, bqnkernel, linux and gui), in addition to the qoe package. Platform packages require a reboot to activate. The following is the process for reboot:

1. Install new qoeos, wait for one minute and reboot.
2. Install kernel and qoekernel, wait for one minute and reboot.
3. Install linux, wait for one minute 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. 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 by Cambium Support from **Administration** > **Diagnostic**.

The file is placed in the download folder of the browser being used.

Backup/restoring the configuration

You can save the server configuration to a local file in **Administration** > **Backup** > **Save**.

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

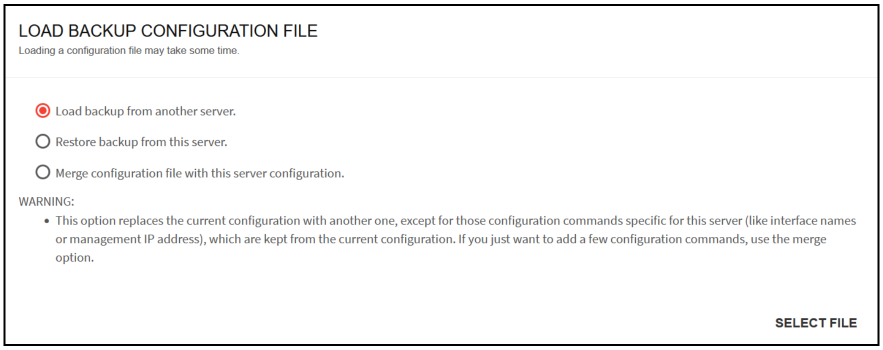

LOAD BACKUP CONFIGURATION FILE

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.

SELECT FILE

The following are the backup/restoring options:

- **Load a 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 a 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 a 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 dialog box appears with 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-poo**l: 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, see Troubleshooting section.

To configure the SNMP agent, navigate to: **Administration** > **SNMP**. Figure 19 shows the SNMP page.

Figure 19: *The SNMP page*

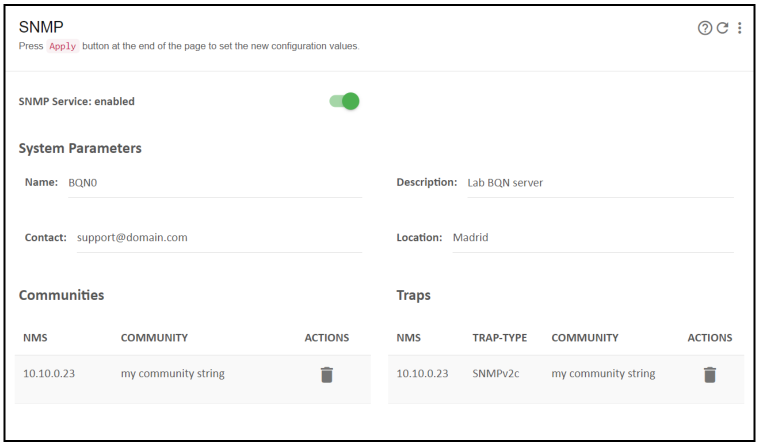

SNMP

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	🗑

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 can be 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 magnifying glass icon next to the interface name. It is available in the following pages:

- **Status** > **Interface**s > **Link State**
- **Status** > **Interfaces** > **Data Wires**

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

Figure 20: *Traffic captures*

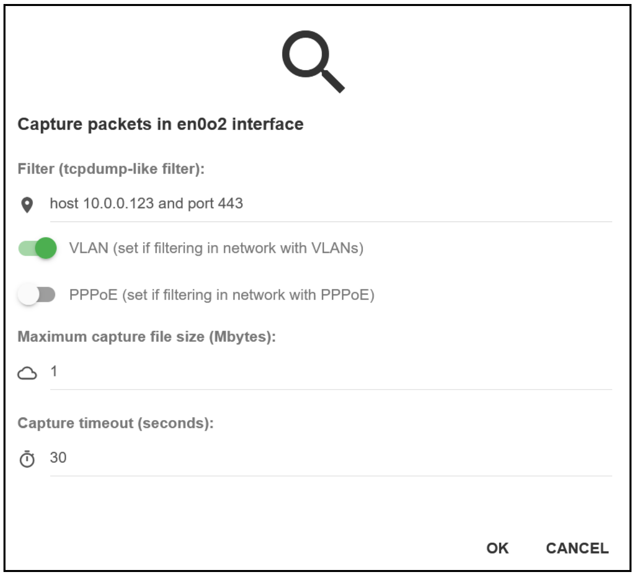

Capture packets in en0o2 interface

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

The **filter** field accepts the format of `tcpdump` filters. If it is optional, and if empty, all traffic will be 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 Figure 20, the network has VLANs.

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

Capture timeout (600 seconds maximum). The capture stops if the maximum time is elapsed. It can also be stopped before that by pressing the **CANCEL** button.

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 your requirements.

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

Logs

In order to help in the debugging of complex issues, the GUI displays two types of logs:

- 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.

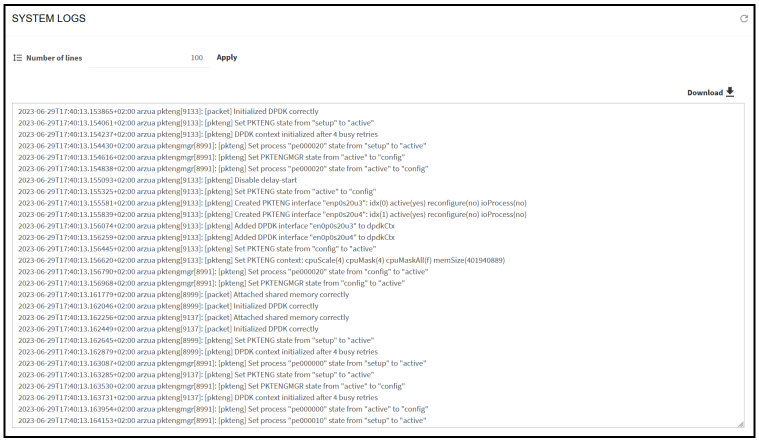

```
2023-06-29T17:40:13.153865+02:00 arzua pkteng[9133]: [packet] Initialized DPDK correctly
2023-06-29T17:40:13.154061+02:00 arzua pkteng[9133]: [pkteng] Set PKTENG state from "setup" to "active"
2023-06-29T17:40:13.154237+02:00 arzua pkteng[9133]: [pkteng] DPDK context initialized after 4 busy retries
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"
```

Software bypass

It is possible to make some traffic to go through the QoE transparently, without being processed by QoE. Such traffic will be captured in one of the network interfaces and relayed transparently to its peer interface in the same wire. That way, the QoE software will have no impact on such traffic. Figure 21 shows the software bypass options.

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).
- Some IPv4 and/or IPv6 addresses or address ranges.

Figure 21: *Software bypass options*

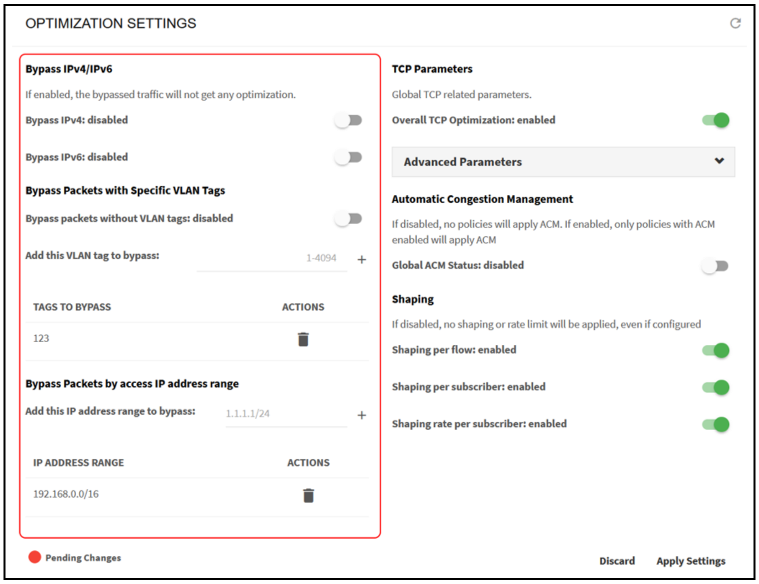

OPTIMIZATION SETTINGS

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
123	🗑

Bypass Packets by access IP address range

Add this IP address range to bypass: 1.1.1.1/24 +

IP ADDRESS RANGE	ACTIONS
192.168.0.0/16	🗑

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: disabled

Shaping

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

Shaping per flow: enabled

Shaping per subscriber: enabled

Shaping rate per subscriber: enabled

● Pending Changes

To bypass some traffic, navigate to **Configuration** > **Optimization Settings** and enable the corresponding toggle.

For VLANs and IP ranges, type in the value, press + to add it and apply the settings.

> **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 in **Administration** > **Users**.

Figure 22: *The Available Users page*

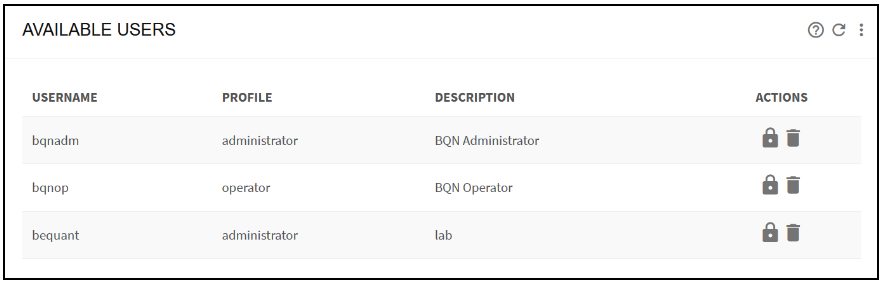

AVAILABLE USERS

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

Secure setup

Session timeout

A configurable timeout will disconnect 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** as shown in Figure 23.

Figure 23: *The General Settings page*

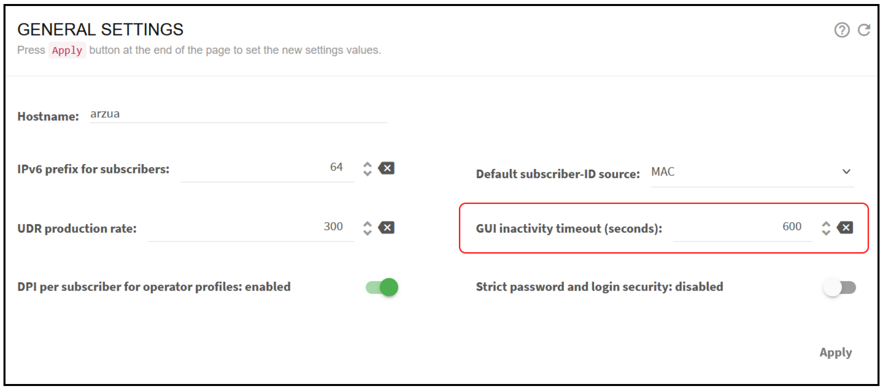

GENERAL SETTINGS

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

The inactivity timeout is applied to new sessions.

Strong user passwords

By default, when setting a user password, any value is valid. It is possible to strengthen the system security by forcing some minimum complexity to user passwords. To set the password, navigate to **Administration** > **General Settings** and enable **Strict password and login security** as shown inFigure 24.

Figure 24: *The General Settings page*

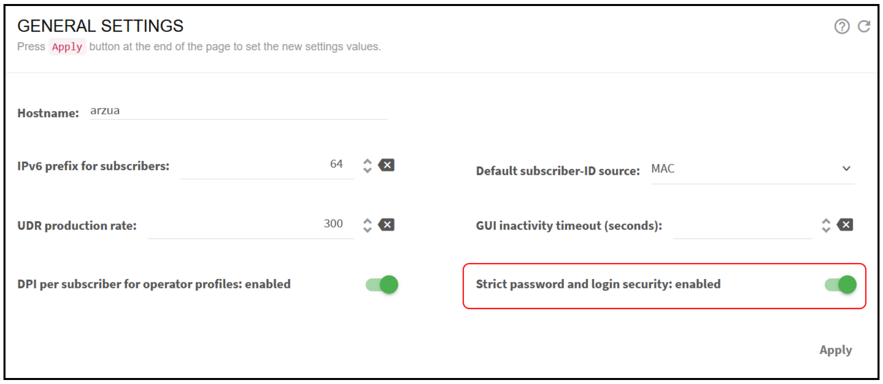

GENERAL SETTINGS

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):

DPI per subscriber for operator profiles: enabled

Strict password and login security: enabled

When the strict password switch is set to ON, the following minimum password complexity is enforced (and password change rejected if not met):

- Length of at least 8 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 **pam_cracklib** simplicity test. This test rejects poor passwords such as:

- 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 applies only to the management interface (not to the interfaces configured in wires), select on the lateral menu **Configuration** > **Interfaces** > **Management Firewal**l. This will show the IP address ranges allowed to access the management interface. By default, no IP address ranges are configured and allowed.

To add an allowed IP address range, click ⋮ and press on **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 GUI and also the subnet of the management IP address (the GUI will include 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 once the **Apply** button. It is important not to press **Apply** before all entries have been deleted, because a premature **Apply** would keep 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, navigate to: **Administration** > **General Settings** and disable the switch **DPI per subscriber for operator profiles** as shown in Figure 25.

Figure 25: *The General Settings page*

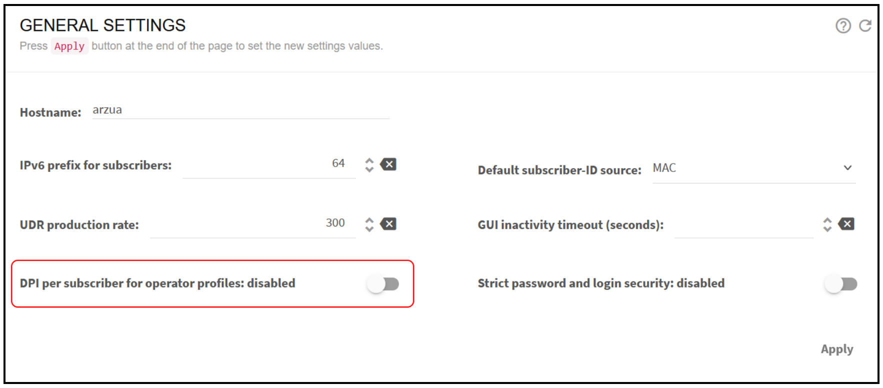

GENERAL SETTINGS

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):

DPI per subscriber for operator profiles: disabled

Strict password and login security: disabled

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 will be 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 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.

Sending system logs to a syslog server

It is possible to configure the QoE to send its system logs to an external syslog server. The syslog server must support BSD syslog protocol (IETF RFC 3164) or syslog protocol (IETF RFC 5454).

To do the setup, log as root through SSH and follow these steps:

1. Create the following directory, so the changes are persistent:

```
mkdir -p /bqn/root/etc/rsyslog.d/
```

2. Copy the original configuration file into that directory:

```
cp /etc/rsyslog.d/remote.conf /bqn/root/etc/rsyslog.d/
```

3. Edit the configuration file:

```
vim /bqn/root/etc/rsyslog.d/remote.conf
```

4. If the BSD syslog protocol is used (IETF RFC 3164), add this line to the configuration file:

```
*.* action(type="omfwd" target="server-ip" port="server-port"
protocol="tcp")
```

5. Replace **server-IP** by the syslog server IP address and **server-port** by its port (example, 514).

If the format needed is syslog protocol (IETF RFC 5454):

```
*.* action(type="omfwd" target="server-ip" port="server-port"
protocol="tcp" template="RSYSLOG_SyslogProtocol23Format")
```

6. Reboot the system:

```
reboot
```

The QoE system logs in /var/log/messages and the messages are sent to the syslog server.

If later on, you do changes to the configuration file, you need to restart the rsyslog service for the changes to be applied:

```
systemctl restart rsyslog.service
```

7. Check that the service is OK by requesting its status:

```
bqn:~ # systemctl status rsyslog.service

rsyslog.service - System Logging Service

Loaded: loaded (/bqn/img/linux/usr/lib/systemd/system/rsyslog.service;
enabled)

Active: active (running) since Thu 2023-10-05 10:37:17 CEST; 1 months 10
days ago

Main PID: 6992 (rsyslogd)

CGroup: /system.slice/rsyslog.service

`-6992 /usr/sbin/rsyslogd -n

. . .
```

Configuration using TLS

To encrypt the communication with the syslog server, the QoE needs the package linux-R3.0.13-20231130 or later.

Transfer the server certification authority certificate to the QoE server rsyslog directory:

```
scp ca.pem root@bqn:/bqn/root/etc/rsyslog.d
```

Edit the configuration file:

```
vim /bqn/root/etc/rsyslog.d/remote.conf
```

Add the following lines before the *.* action line (chosen port is normally 6514, review it in action line):

```
$DefaultNetstreamDriverCAFile /bqn/root/etc/rsyslog.d/ca.pem

$DefaultNetstreamDriver gtls # use gtls netstream driver

$ActionSendStreamDriverMode 1 # require TLS for the connection

$ActionSendStreamDriverAuthMode anon # server is NOT authenticated
```

To apply the changes, restart the rsyslog service:

```
systemctl restart rsyslog.service
```

TCP Optimization

This chapter contains the following sections:

- Metrics
- Configuration

Metrics

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

Status > **TCP Optimization** > **Speed & Acceleration** shows the average speed values of non-accelerated TCP connections (no TCPO) and accelerated ones (TCPO) and the speed increase in percentage during the last 24 hours (n/s if no statistically significant difference found because of variability or not enough samples). Figure 26 shows the TCP speed and acceleration page.

The following metrics are shown:

- 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).

Figure 26: *TCP speed and acceleration page*

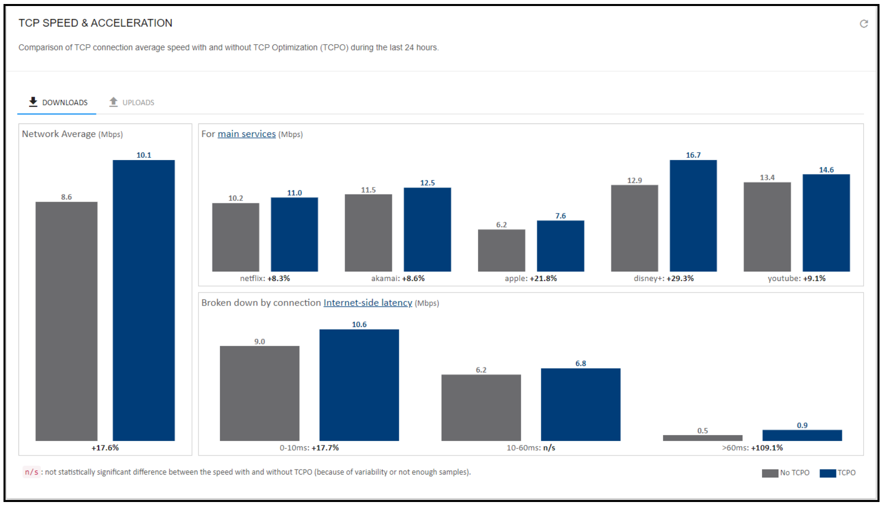

TCP SPEED & ACCELERATION

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).

Sometimes, due to no or little traffic, there might be categories or graph without information. Extend the period of calculation (1 day by default).

The menu option **Statistics** > **TCP Optimization** > **TCP Speed** shows the speed evolution over time. Figure 27 shows the speed evolution over time.

Figure 27: *Speed evolution over time*

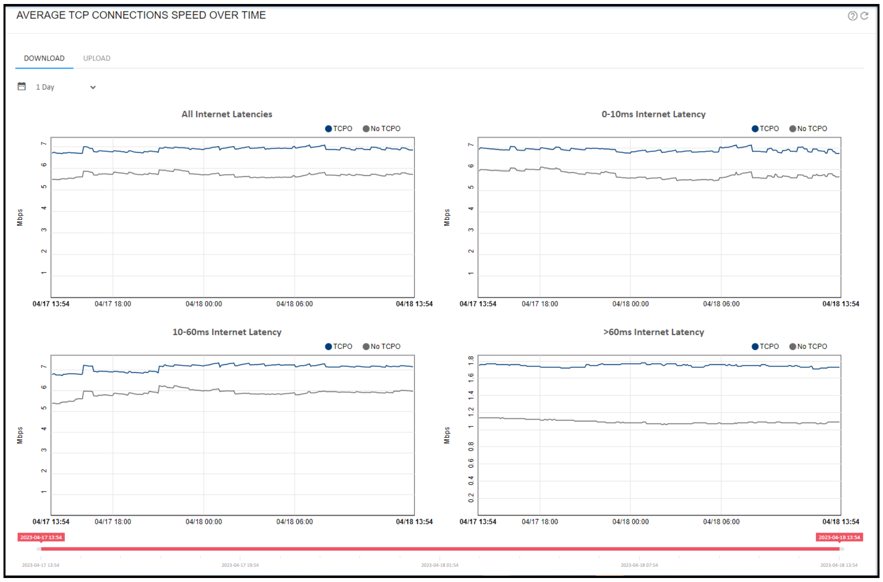

Sometimes, due to no or little traffic, there might be categories or graph periods without information.

Navigate to: **Statistics** > **TCP Optimization** > **TCP Acceleration** to view the acceleration evolution over time. Figure 28 shows the acceleration evolution over time.

Figure 28: *Acceleration evolution over time*

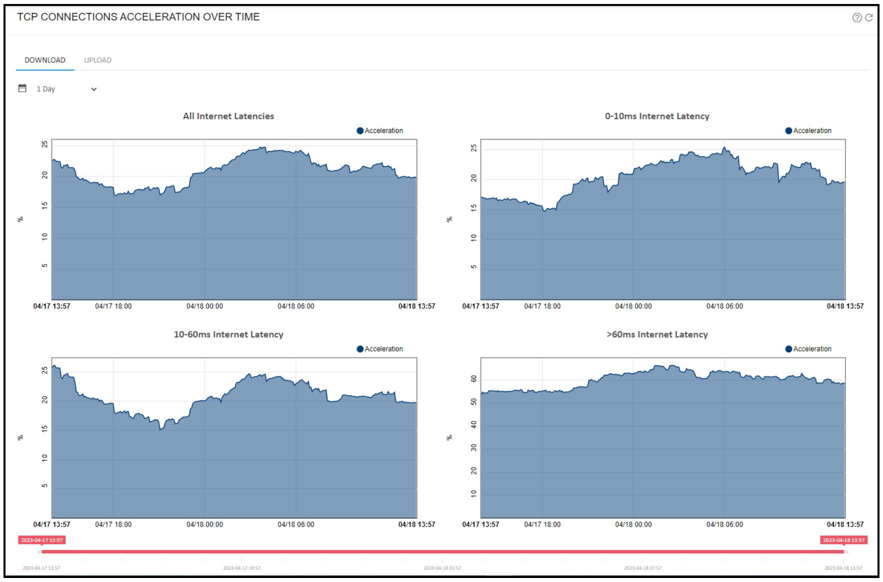

Configuration

To configure TCP Optimization parameters, navigate to: **Configuration** > **Optimization Settings**.

The latency thresholds can be changed in RTTi-small and RTTi-large. You can see the Internet latencies of your network in **Statistics** > **System** > **Latency and Statistics** > **DPI Service Analysis** > **Latency per Service**.

Latency thresholds affect only the way the metrics are classified and displayed, not the TCP optimization, which adjusts continuously to each individual connection characteristics, latencies included. Figure 29 shows the TCPO/ACM settings page.

Figure 29: *TCPO/ACM settings page*

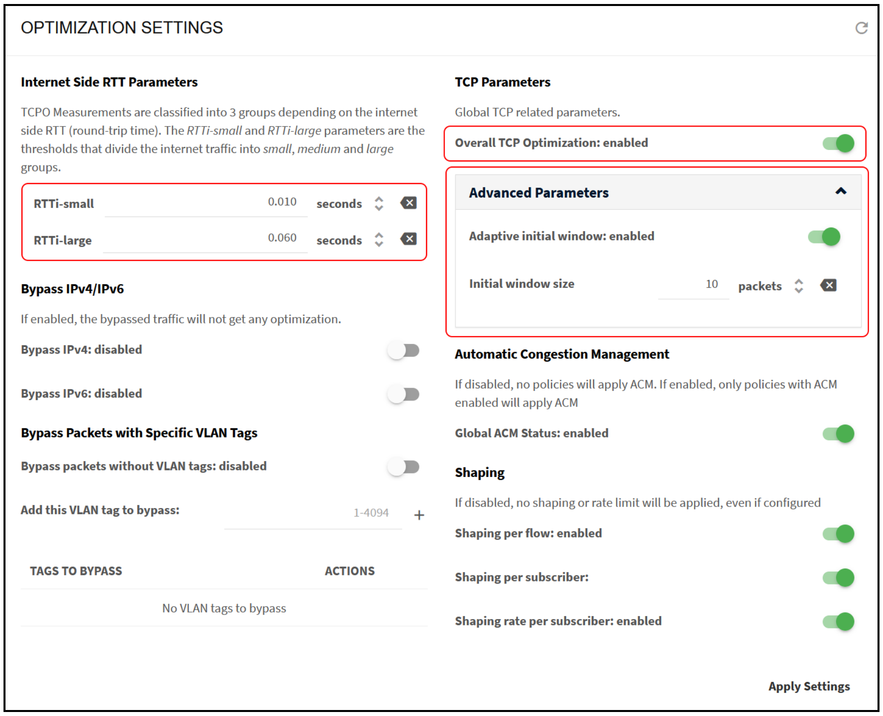

OPTIMIZATION SETTINGS

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

To disable all TCP optimization in the QoE, switch off the toggle **Overall TCP Optimization**. This is done to temporally disable TCPO 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.

Automatic Congestion Management (ACM) optimization

When the subscriber rate limits are unknown, QoE can automatically detect them using machine learning. So 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, when rate limits are known.

This chapter contains the following sections:

- ACM metrics
- Configuration

ACM 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 30 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 due to ACM.
- **Retransmission Reduction with ACM** depicts the reduction in the packet retransmission percentage obtained by the ACM.

Figure 30: *ACM metrics*

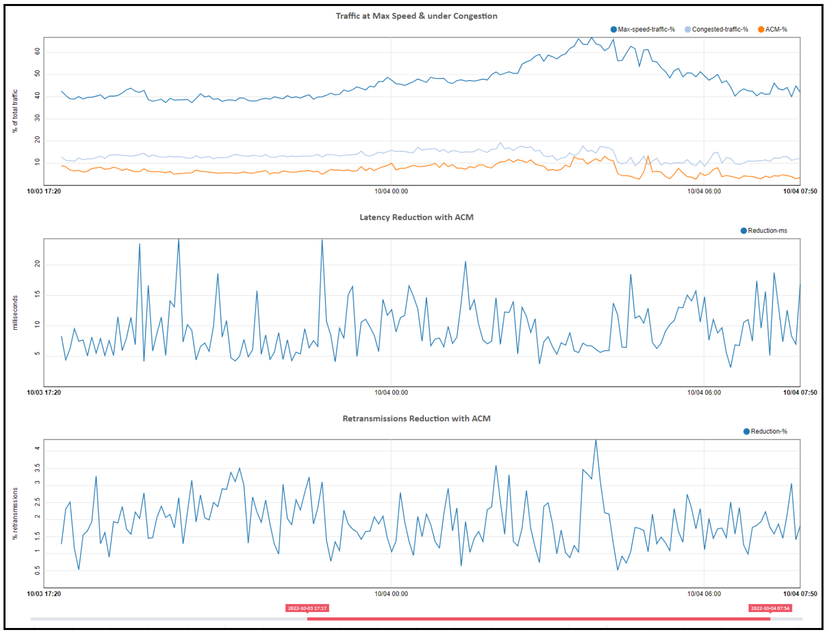

Also, these charts are available per Subscriber. Navigate to: **Status** > **Subscribers**, click **Subscriber ID** or IP address to view the subscriber dashboard.

Configuration

ACM improves the network quality with no need of configuration fine tuning, so it is recommended to keep it enabled at all times. Nevertheless, it is possible to disable ACM for all policies going to **Configuration > Optimization Settings** and disable **Global ACM Status**. Figure 31 shows the Optimization Settings page.

Figure 31: *The Optimization Settings page*

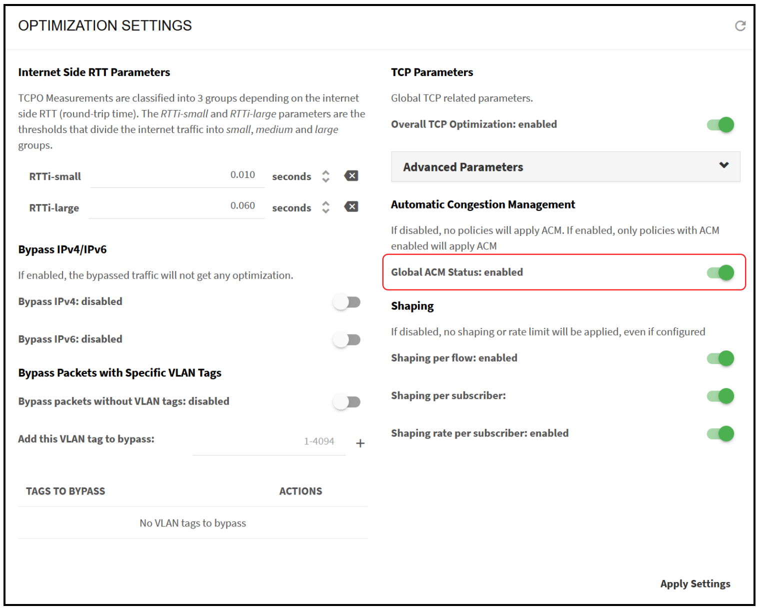

OPTIMIZATION SETTINGS

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

Network visibility

This chapter contains the following sections:

Subscriber dashboard

The subscriber dashboard includes a comprehensive set of useful information about the subscriber current and past performance. It is the ideal place where to analyze and diagnose issues reported in the 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**

The subscriber dashboard contains a lot of information, so use the scroll bar on the right-hand side to browse through it. Figure 32 shows the **Subsrcriber Dashboard** page.

Figure 32: *The Subscriber Dashboard page*

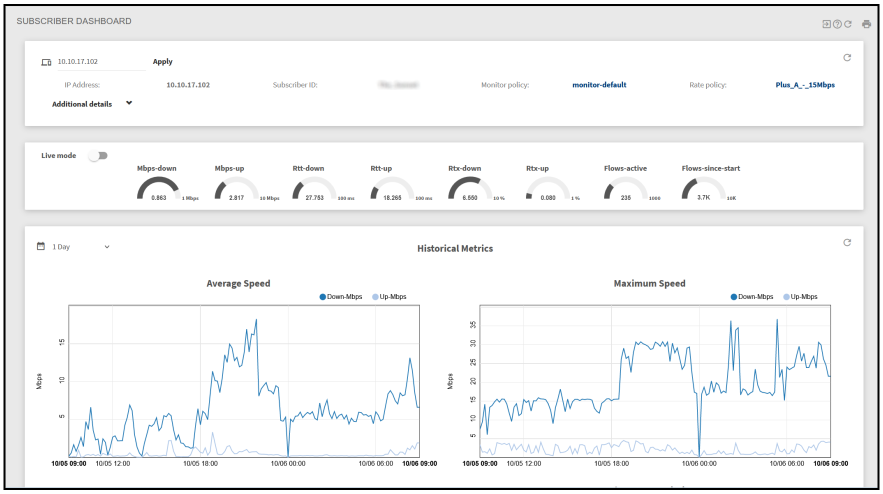

SUBSCRIBER DASHBOARD

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

1 Day

Historical Metrics

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 **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 service usage over time.
- Internet latencies experienced by this subscriber for the most used services. Clicking on the bar on the chart will take you to the histogram with the latency distribution and from there, to the histogram evolution over time.
- Table with the active flows of the subscriber. The flow policy is provided (click on the policy name to go to the policy definition). If available, it also shows the DPI domain (**DOMAIN** column) and the DPI category of the flow (**DPI** column).

Figure 33 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 whether the subscriber is below or above the network average quality.

Figure 33: *Average and maximum speed*

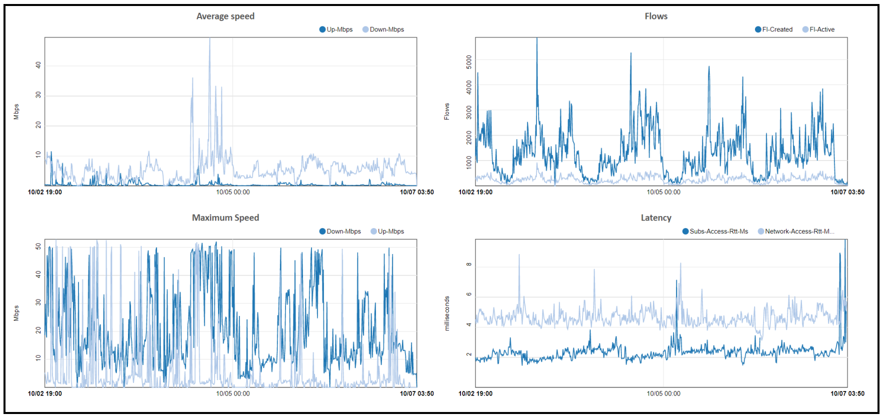

Figure 34 shows the packet retransmission on the access side, also including the network average as a reference. The last three charts contain information about the Automatic Congestion Management (ACM) feature for that subscriber.

Figure 34: *Packet retransmissions*

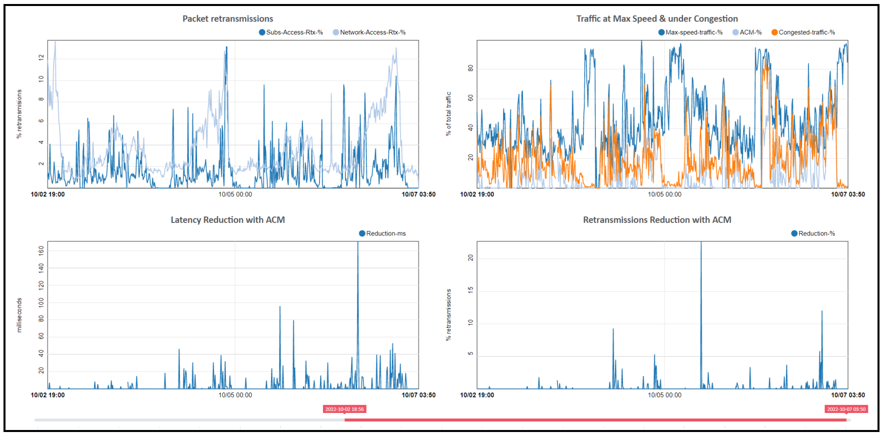

To monitor the subscriber more closely, set the **Live mode** switch to ON. Charts will appear with the current evolution of the main metrics, including DPI composition. A maximum time span can be shown (latest number of minutes). This is configurable to 2, 5 or 10 minutes, using the selector next to the **Live mode** switch. When the time span is reached, for example, after 5 minutes, the charts will continue to update, showing the latest five minutes window.

A flow table will also be included and will be continuously refreshed. Flows are ordered in the table with active flows first, and among them, the faster flows. DPI categories and domains are shown. 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. Figure 35 shows the main active flows.

Figure 35: *Main Active flows*

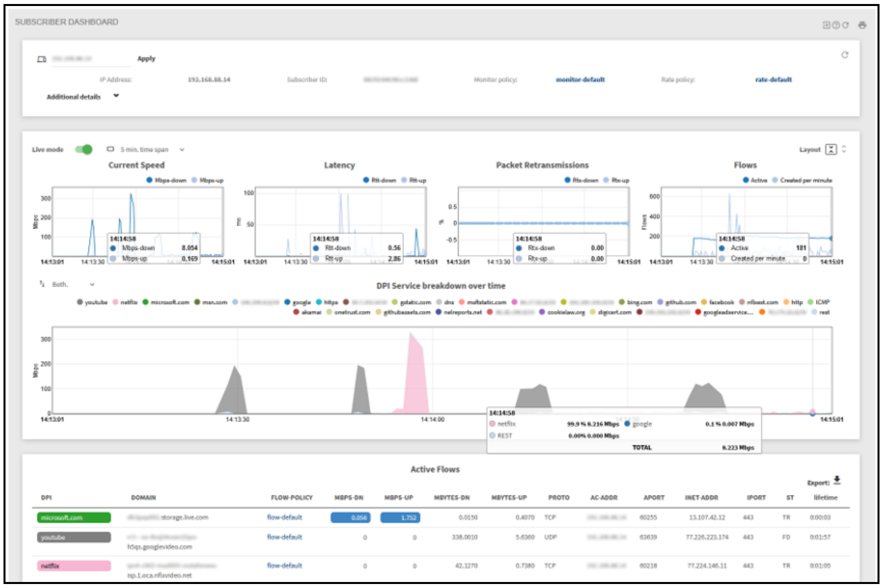

SUBSCRIBER DASHBOARD

[redacted] Apply

IP Address: 192.168.88.14 Subscriber ID: [redacted] Monitor policy: monitor-default Rate policy: rate-default

Additional details

Active Flows

Export:

DPI	DOMAIN	FLOW-POLICY	MBPS-DN	MBPS-UP	MBYTES-DN	MBYTES-UP	PROTO	AC-ADDR	APORT	INET-ADDR	IPORT	ST	lifetime
microsoft.com	[redacted] storage.live.com	flow-default	0.056	1.752	0.0150	0.4070	TCP	[redacted]	60255	13.107.42.12	443	TR	0:00:03
youtube	[redacted] h5qs.googlevideo.com	flow-default	0	0	338.0010	5.6360	UDP	[redacted]	63639	77.226.223.174	443	FD	0:01:57
netflix	[redacted] isp.1.oca.nflxvideo.net	flow-default	0	0	42.1270	0.7380	TCP	[redacted]	60218	77.224.146.11	443	TR	0:01:05

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

Subscriber Identifier (ID)

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

There are several possible sources for the 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**.
- **External system**: An external system can use the QoE REST API, RADIUS or one of our Billing. For more information, see Billing systems.

Subscriber QoE metrics page

Subscriber QoE metrics page provides access to the list of active subscribers along with their metrics. If the subscriber of interest is not listed, enter the IP address on the filter field. To view Subsciber QoE metrics page, navigate to **Status** > **Subscribers** > **QoE Metrics**.

Figure 36: *The Subscriber QoE metrics page*

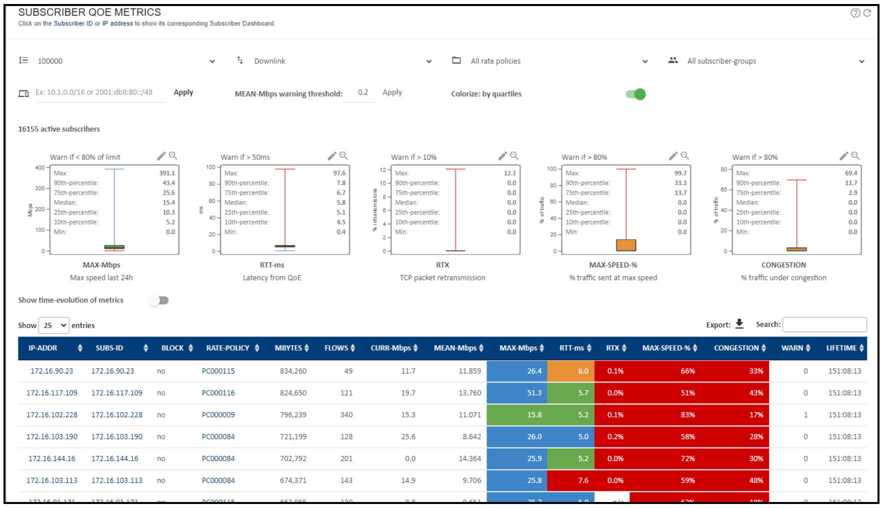

SUBSCRIBER QOE METRICS

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

To facilitate the identification of subscribers with quality of experience issues, the **WARN** column shows a score with the number of metrics above their warning threshold. Sort the table so the highest WARN values are shown first. In the example above, the first three subscribers have access latency issues (values above a threshold set to 10ms 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 perhaps 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**, on the window upper right). If a customer does not reach at least this mean speed, then the customer does not get a warning score and it will not be highlighted. The threshold can be changed, but in general this is not needed.

Status > Subscribers > QoE Metrics contains the information that were in **Status** > **Subscribers** in previous versions.

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 37 shows the maximum and minimum values, the median (value separating the lower half and higher half) and the 1st and 3rd quartiles (values separating the lower 25% and the higher 25%, respectively):

Figure 37: *Box plots*

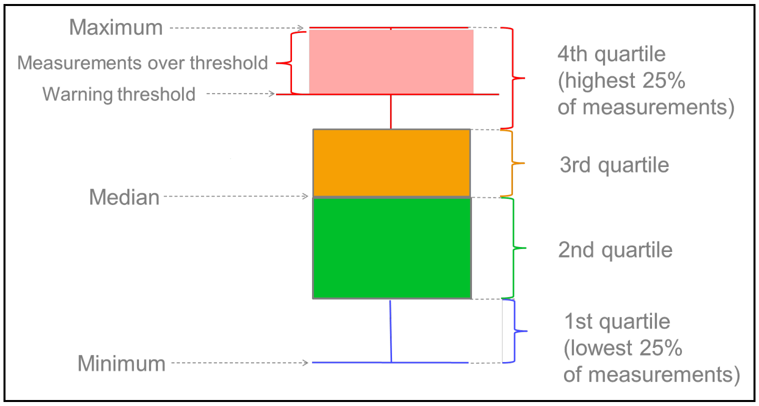

The exact values of the percentiles are shown on the right-hand side of a boxplot.

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 will shade in red all values between 21ms and the maximum. In the table, the RTT values above the threshold will be 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 38: *Warning thresholds*

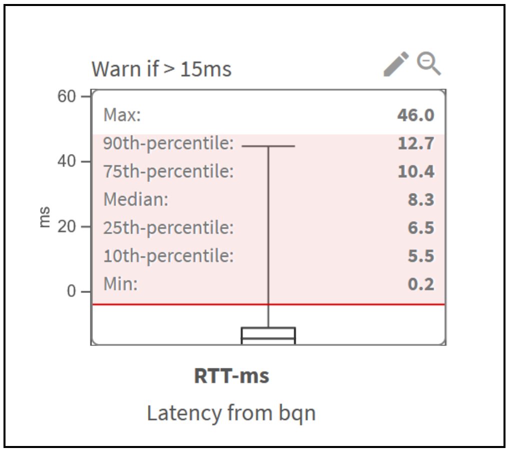

Warn if > 15ms

Max:	46.0
90th-percentile:	12.7
75th-percentile:	10.4
Median:	8.3
25th-percentile:	6.5
10th-percentile:	5.5
Min:	0.2

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

Figure 39: *Threshold settings*

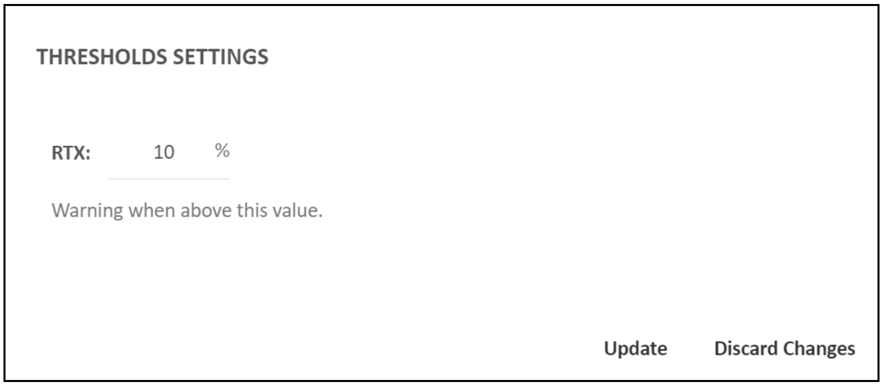

THRESHOLDS SETTINGS

RTX: 10 %

Warning when above this value.

Highlight percentiles

It is also possible to see the quartile of a value by setting the colorize switch in the upper right of the page to ON. Figure 40 shows the Subsriber QoE metrics page.

The quartile codes are:

- First quartile: blue
- Second quartile (up to median): green
- Third quartile: orange
- Fourth quartile (from third up to maximum): red

Figure 40: *The Subscriber QoE Metrics page*

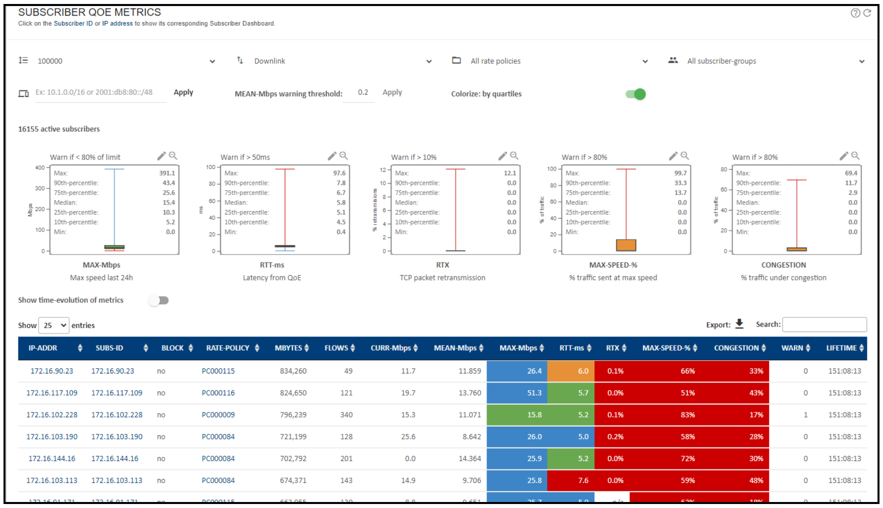

SUBSCRIBER QOE METRICS

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

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

Active subscriber table

Active subscriber table contains the following information:

- **IP-ADDR**: subscriber's IP address.
- **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 of the subscriber session.

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

Metrics over time

The QoE metrics page in Figure 41 shows the temporal evolution of metrics of some selected subscribers.

Figure 41: *The QoE metrics page*

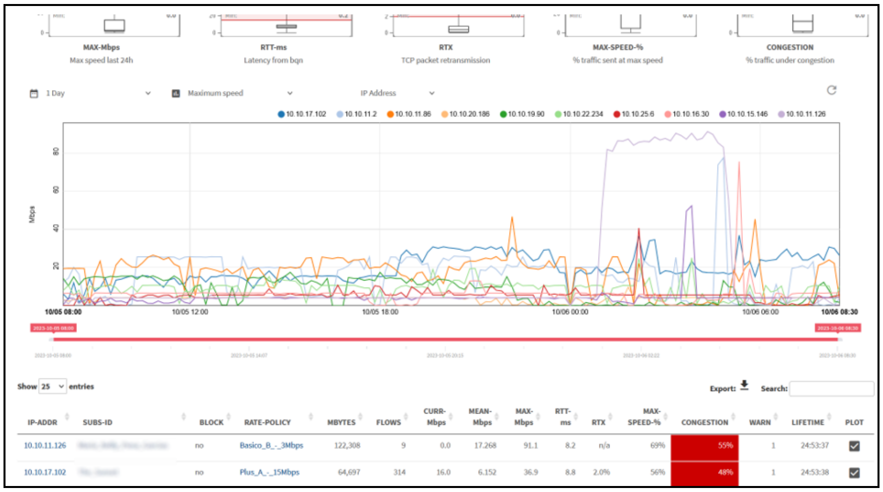

IP-ADDR	SUBS-ID	BLOCK	RATE-POLICY	MBYTES	FLOWS	CURR-Mbps	MEAN-Mbps	MAX-Mbps	RTT-ms	RTX	MAX-SPEED-%	CONGESTION	WARN	LIFETIME	PLOT
10.10.11.126	[redacted]	no	Basico_B_-_3Mbps	122,308	9	0.0	17.268	91.1	8.2	n/a	69%	55%	1	24:53:37	☑
10.10.17.102	[redacted]	no	Plus_A_-_15Mbps	64,697	314	16.0	6.152	36.9	8.8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

By default, the chart shows the first 10 subscribers of the metrics table. The **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 removes 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

The menu entry **Statistics** > **Subscribers Analysis** > **Hourly Volume** displays the subscribers with the largest traffic consumption over time. Usage can be shown per IP address or by Subscriber Id using the appropriate selection. Figure 42 shows the usage per IP address.

Figure 42: *The top subscribers by time page*

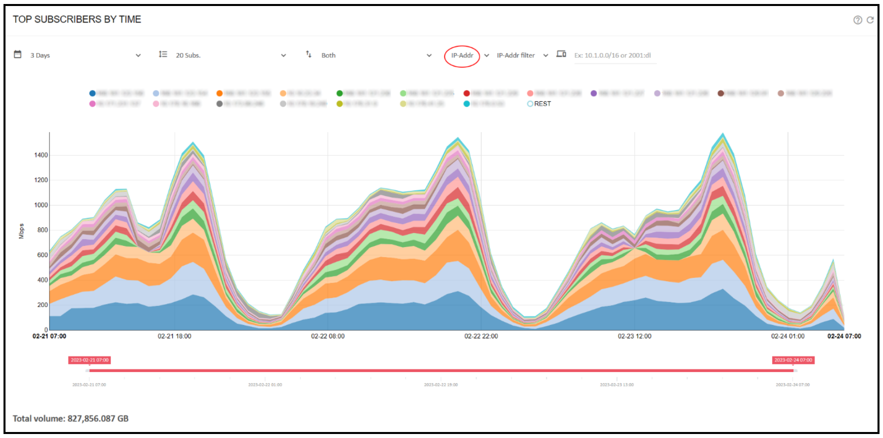

TOP SUBSCRIBERS BY TIME

3 Days | 20 Subs. | Both | IP-Addr | IP-Addr filter | Ex: 10.1.0.0/16 or 2001:dl

Total volume: 827,856.087 GB

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

The **Statistics** > **Subscribers Analysis** > **Total Volume** shows the total in the period being considered. Figure 43 shows usage per Subscriber ID.

Figure 43: *The top subscribers page*

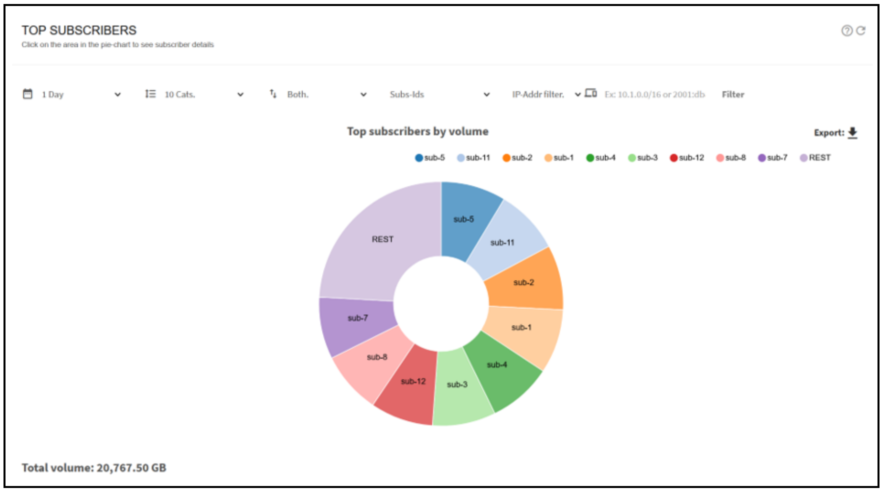

TOP SUBSCRIBERS

Click on the area in the pie-chart to see subscriber details

1 Day | 10 Cats. | Both. | Subs-Ids | IP-Addr filter. | Ex: 10.1.0.0/16 or 2001:db | Filter

Top subscribers by volume

Export:

Total volume: 20,767.50 GB

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

It is also possible to export to a CVS file with the usage of up to 500 top subscribers.

Traffic and subscribers per policy

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

Figure 44: *Split of subscribers into different rate policies*

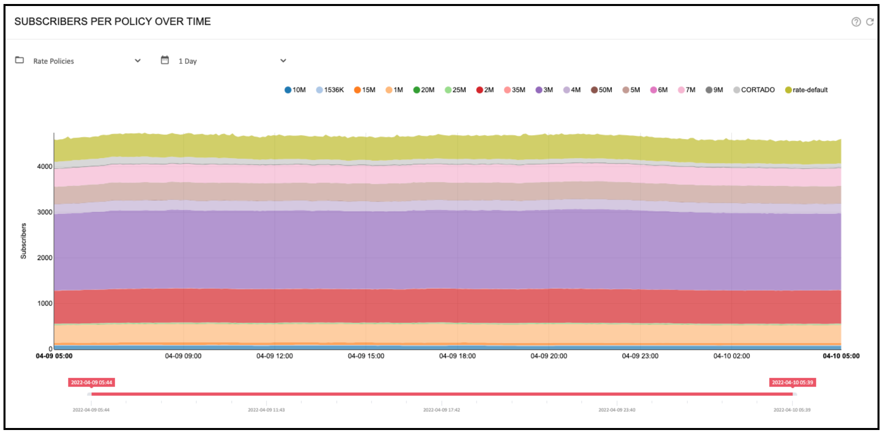

Statistics > Subscribers > Per Policy shows the split of traffic volume into different rate policies.

Figure 45: *Split of traffic volume into different rate policies*

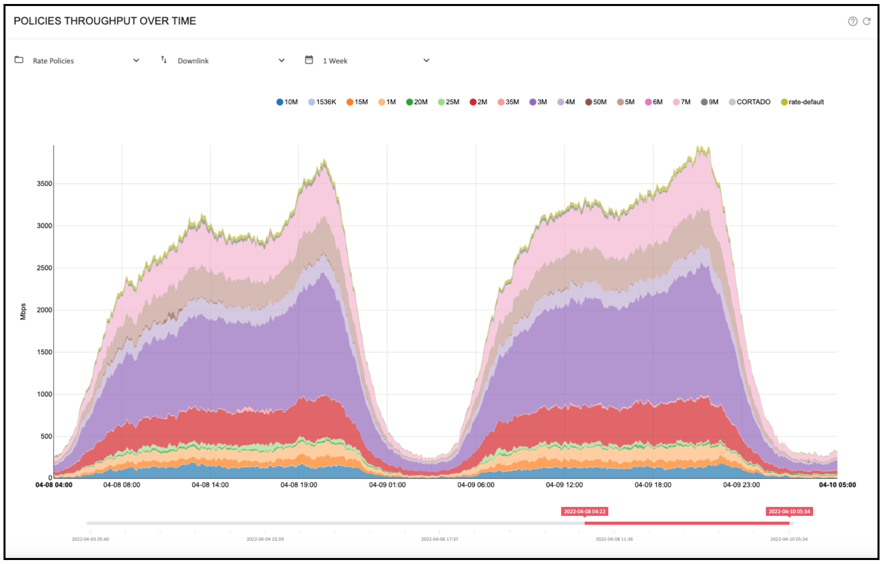

Statistics > Subscribers > Per Policy can also show the split of traffic volume into different flow policies.

Figure 46: *Split of traffic volume into different flow policies*

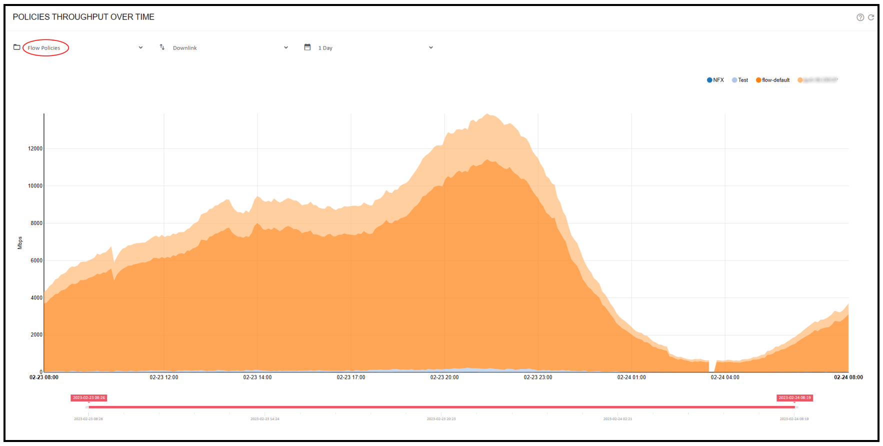

POLICIES THROUGHPUT OVER TIME

Flow Policies | Downlink | 1 Day

Traffic per service

In **Statistics** > **Service Analysis**, there is information about traffic composition.

QoE displays the overall traffic composition per service in **Statistics** > **Service Analysis** > **Total Volume per Service**. shows DPI total statistics. Figure 47 shows the DPI total statistics.

Figure 47: *DPI total statistics*

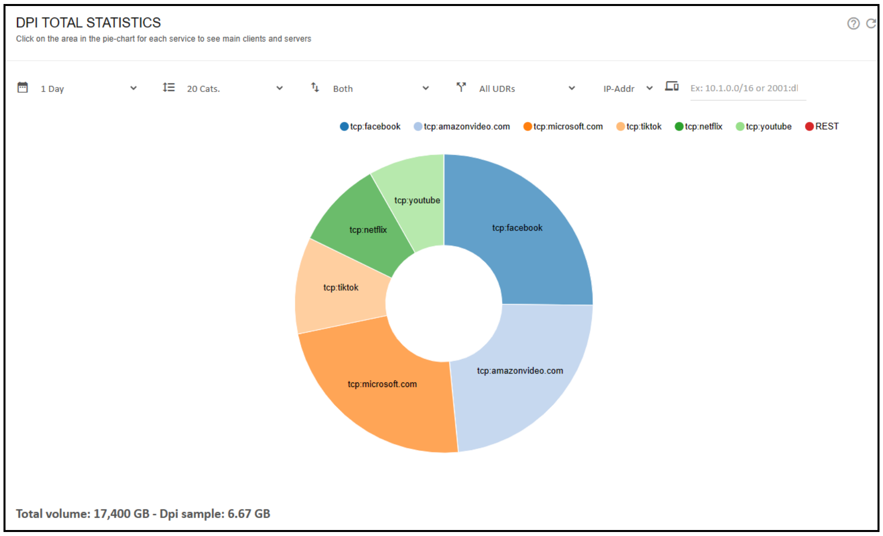

DPI TOTAL STATISTICS

Click on the area in the pie-chart for each service to see main clients and servers

1 Day | 20 Cats. | Both | All UDRs | IP-Addr | Ex: 10.1.0.0/16 or 2001:d

Total volume: 17,400 GB - Dpi sample: 6.67 GB

The evolution over time can be obtained in **Statistics > Service Analysis > Hourly Volume per Service**. Figure 48 shows DPI hourly statistics.

Figure 48: *DPI hourly statistics*

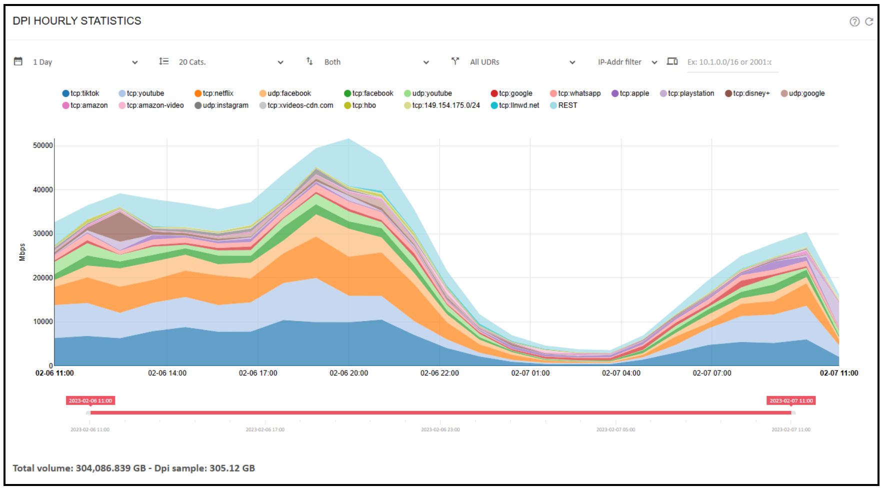

DPI HOURLY STATISTICS

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

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, those samples can be excluded by selecting *Only auto UDRs*. This could be the 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 49 shows the average Internet latency per service page.

Figure 49: *Average Internet latency per service page*

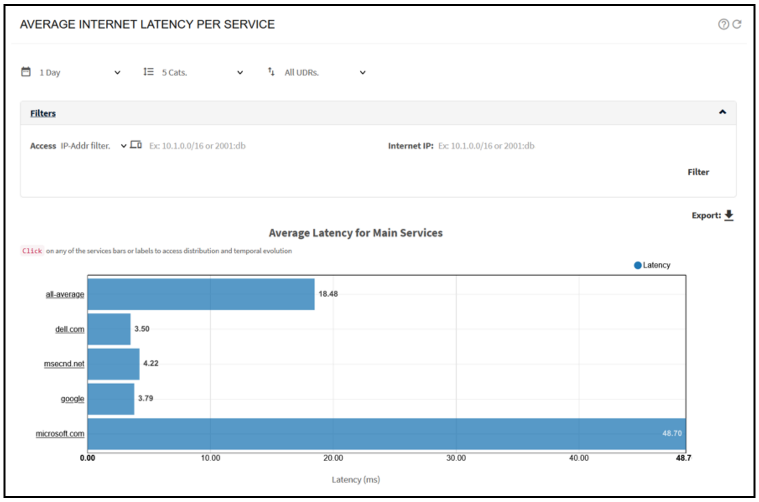

AVERAGE INTERNET LATENCY PER SERVICE

1 Day | 5 Cats. | All UDRs.

Filters

Access IP-Addr filter. Ex: 10.1.0.0/16 or 2001:db

Internet IP: Ex: 10.1.0.0/16 or 2001:db

Filter

Export:

Average Latency for Main Services

Click on any of the services bars or labels to access distribution and temporal evolution

To view the latency distribution of a particular service, click on that service bar on the chart. Figure 50 shows the distribution of Internet latencies page.

Figure 50: *Distribution of Internet latencies page*

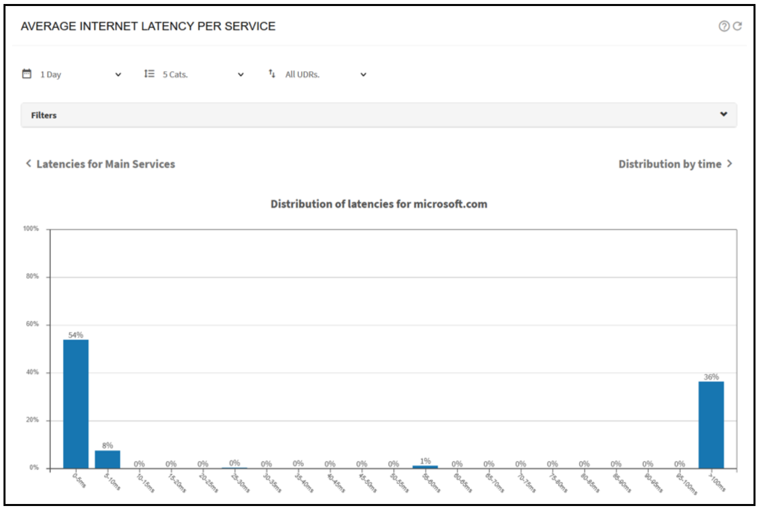

AVERAGE INTERNET LATENCY PER SERVICE

1 Day · 5 Cats. · All UDRs.

Filters

< Latencies for Main Services · Distribution by time >

The distribution shows the percentage of measurements falling into the different intervals of the range of latency values.

To view the evolution over time of those Internet latencies, click on the **Distribution by time** link on the upper right of the chart. Figure 51 shows the distribution of Internet latencies over time page.

Figure 51: *Distribution of Internet latencies over time page*

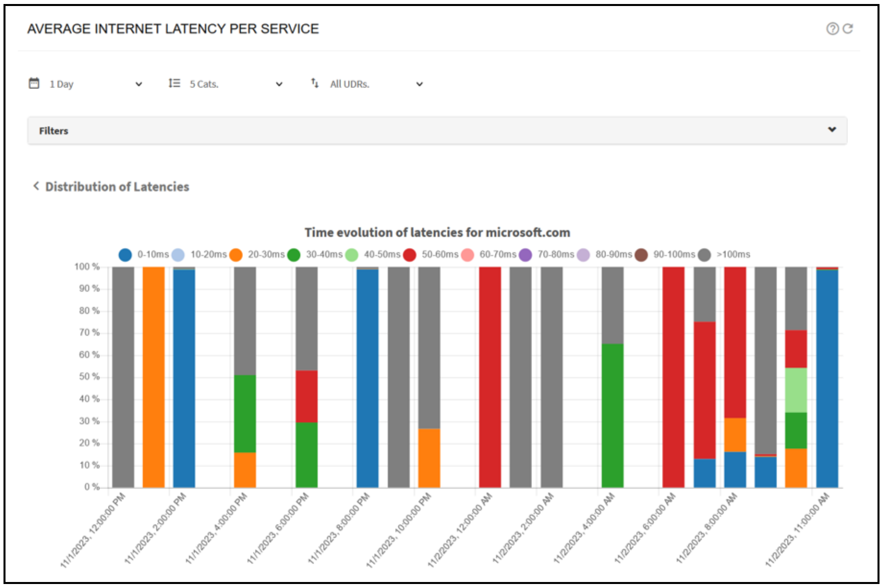

AVERAGE INTERNET LATENCY PER SERVICE

1 Day · 5 Cats. · All UDRs.

Filters

< Distribution of Latencies

For each time interval, the chart shows the relative percentages of each latency range.

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 will let you study the latency of different services coming from different providers, and specifically for certain subscribers. Figure 52 shows the average Internet latency per service.

Figure 52: *Average Internet latency per service*

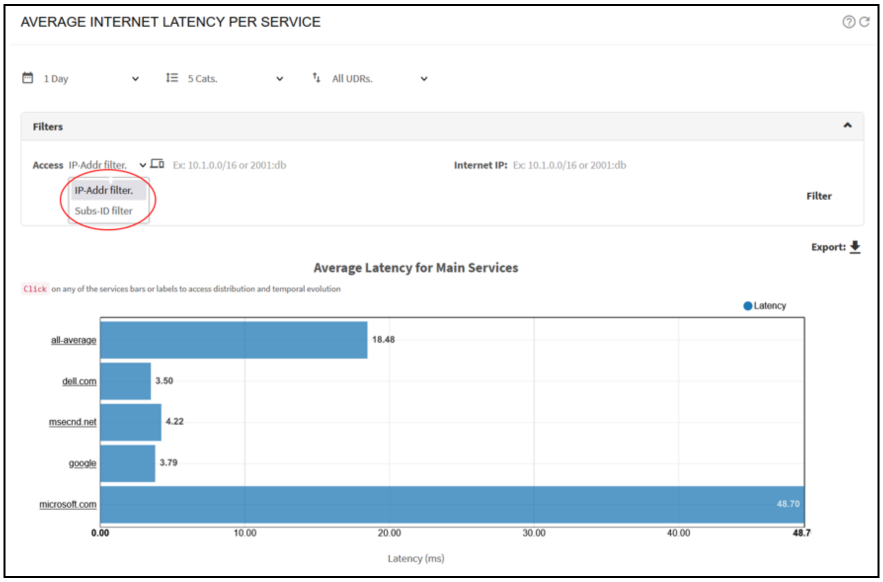

AVERAGE INTERNET LATENCY PER SERVICE

1 Day 5 Cats. All UDRs.

Filters

Access IP-Addr filter. Ex: 10.1.0.0/16 or 2001:db

Internet IP: Ex: 10.1.0.0/16 or 2001:db

Click on any of the services bars or labels to access distribution and temporal evolution

Main subscribers per service

It is possible to identify the main subscribers of a given service, both subscriber IDs 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 will be shown, with server IP addresses at the top and subscriber addresses at the bottom. The histogram shows the percentage of the total service volume of that IP address.

It is also possible to show top subscriber IDs using the selector at the bottom of the page. Figure 53 shows the percentage of traffic page.

Figure 53: *Percentage of traffic*

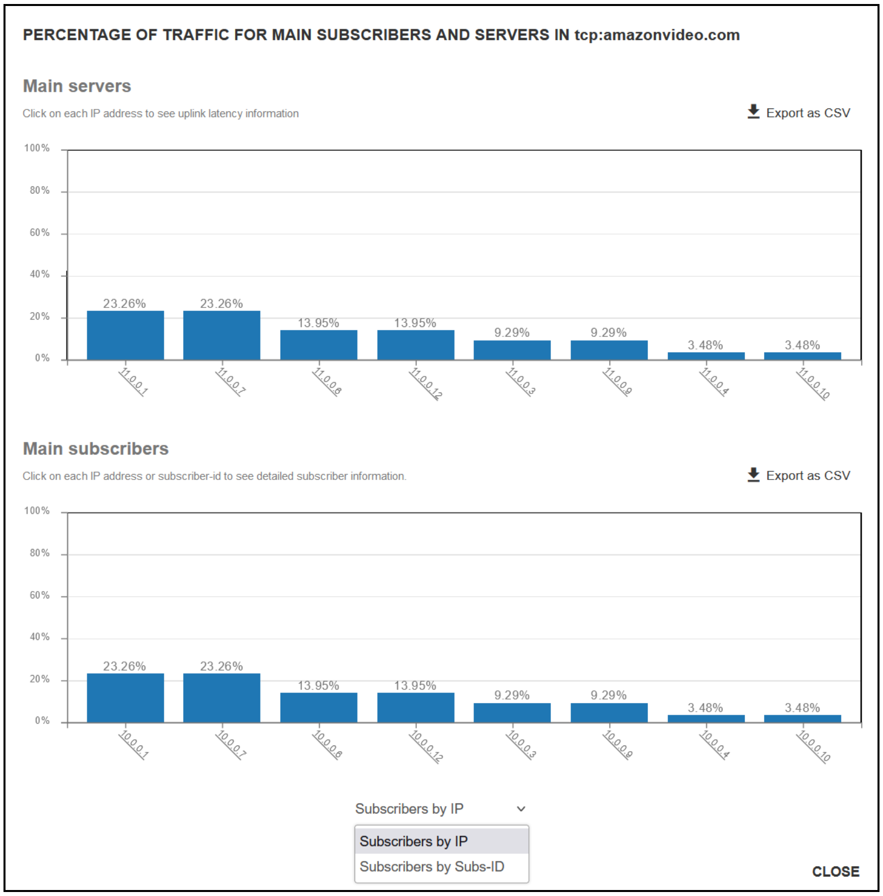

PERCENTAGE OF TRAFFIC FOR MAIN SUBSCRIBERS AND SERVERS IN tcp:amazonvideo.com

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.

Addresses are clickable: server IPs 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 the **Export as CSV** icon.

Overall traffic metrics

Overall traffic metrics shows the temporal evolution of total traffic throughput, adding both directions and all wires. To view Overall traffic metrics, navigate to **Statistics** >**Throughput** > **Overview.** Figure 54 shows the throughput over time page.

Figure 54: *Throughput over time page*

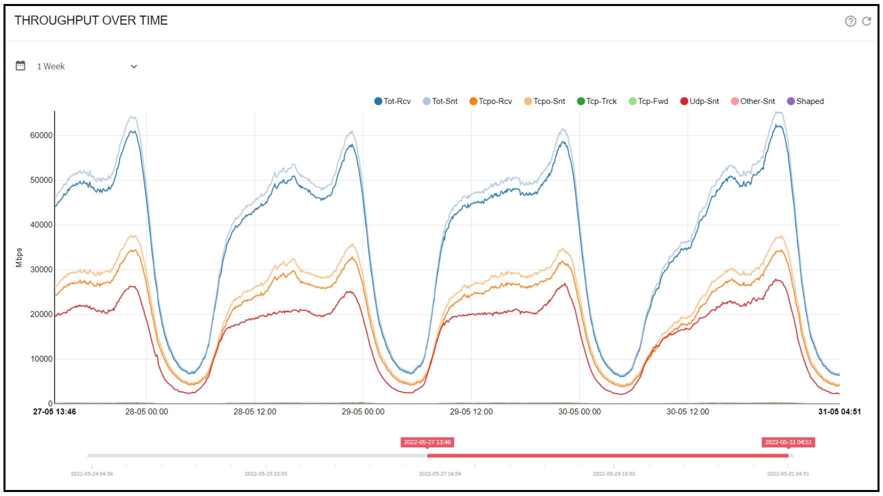

Statistics > **Throughput** > **Protocols** shows the temporal evolution of traffic throughput, broken down by:

- **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 55 shows the throughput overview page.

Figure 55: *Throughput overview page*

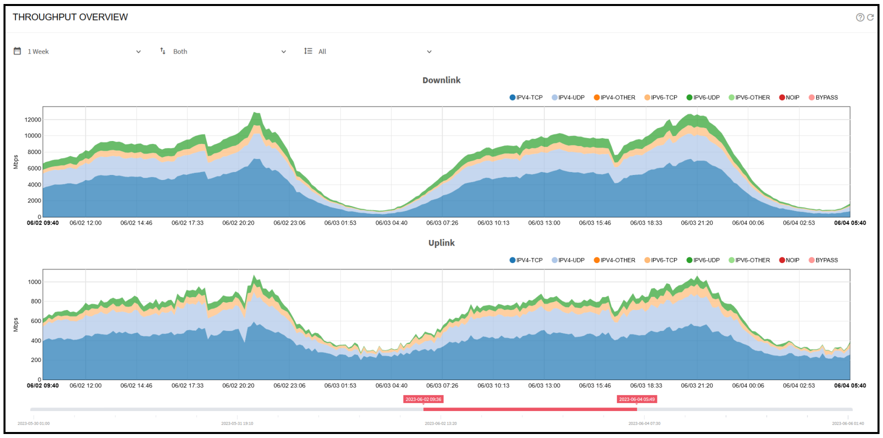

The evolution over time per network interface is available in **Statistics** > **Throughput** > **Interfaces**. Figure 56 shows the network interface throughput over time page.

Figure 56: *Network interface throughput over time page*

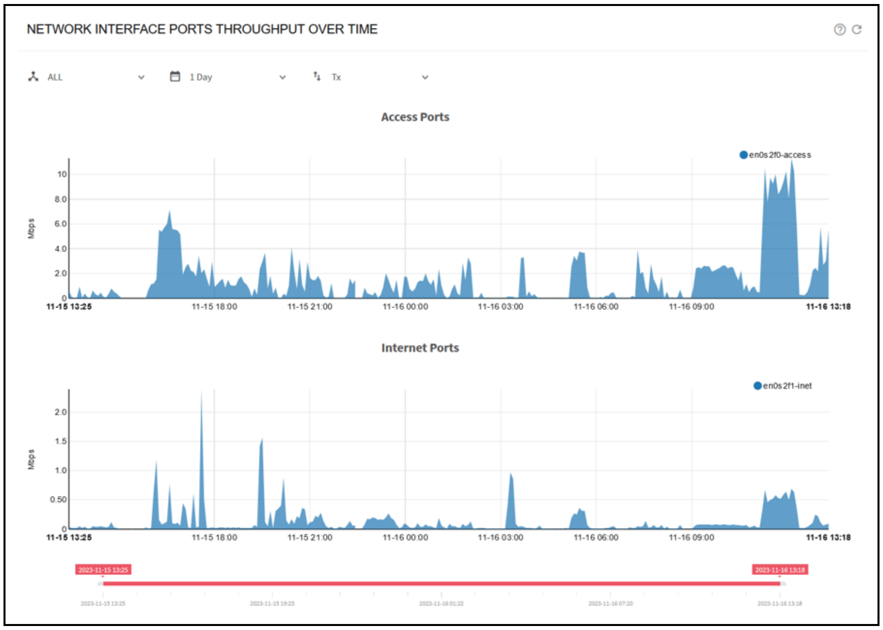

It is possible to check how much traffic is being processed according to each of the configured policies. For Subscriber Flows policies, it can be checked in **Statistics** > **Throughput** > **Subscriber Flows Policies** and similarly for **Subscriber Rate Policies** and **Subscriber Monitoring Policies**.

The chart in **Statistics** > **System** > **Latencies** shows the access RTT (RTT-Access). It is the average across all flows of the minimum value per flow. Figure 57 shows the latency over time page.

Figure 57: *Latency over time page*

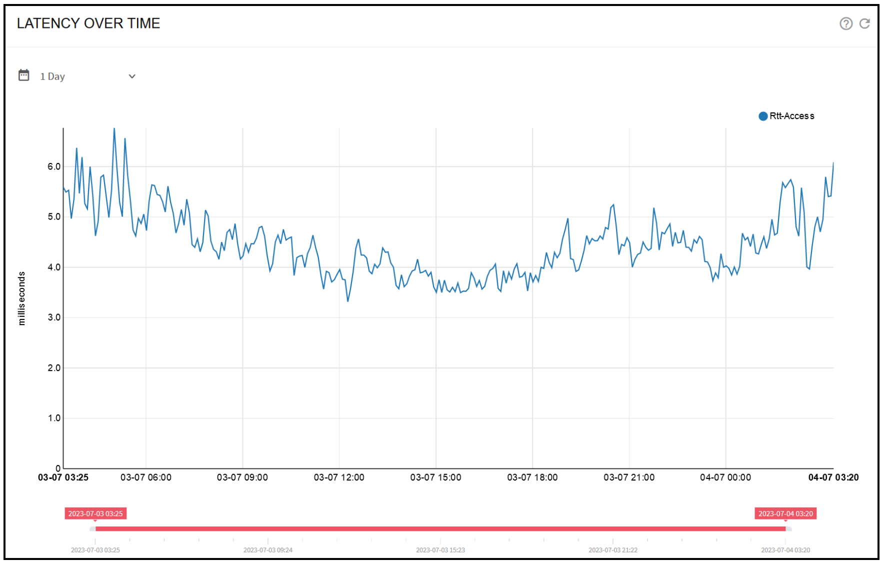

Also, **Statistics** > **System** > **Retransmissions** shows the average retransmission percentages in downlink and uplink directions. Figure 58 shows the average TCP retransmissions over time page.

Figure 58: *Average TCP retransmissions over time*

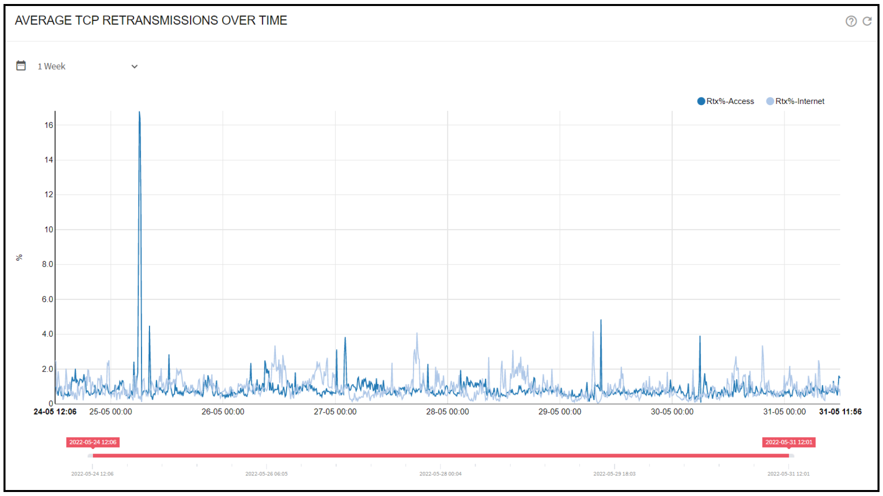

To see the number of flows per policy and per protocol, click on **Statistics** > **Flow** > **Per Policy and Statistics** > **Flow** > **Per Protocol** respectively. Figure 59 shows the active flows over time page.

Figure 59: *Active flows over time page*

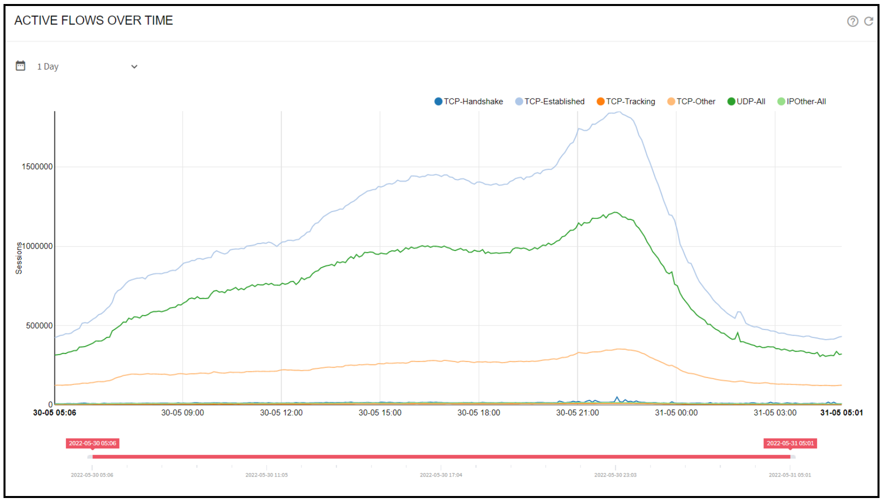

You can also see the instantaneous number of Flows per protocol in **Status** > **Flows** > **Per Protocol** and per subscriber in **Status** > **Flows** > **Per Subscriber**. Figure 60 shows the active flows page.

Figure 60: *Active flows page*

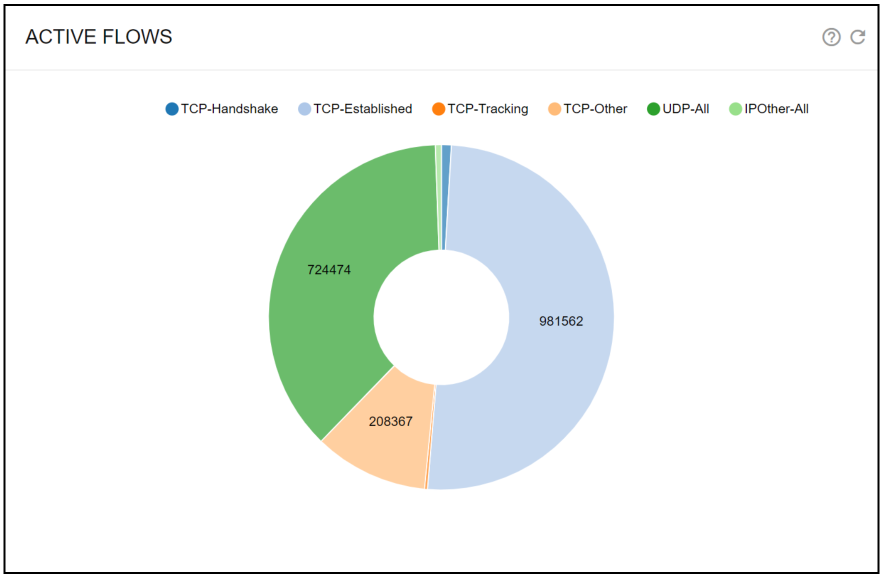

DoS

The QoE detects Denial of Service attacks. To do 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.
- **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 20Mbps plan will have a DoS threshold of 3*20=60Mbps. Figure 61 shows the DoS settings page.

Figure 61: *DoS settings page*

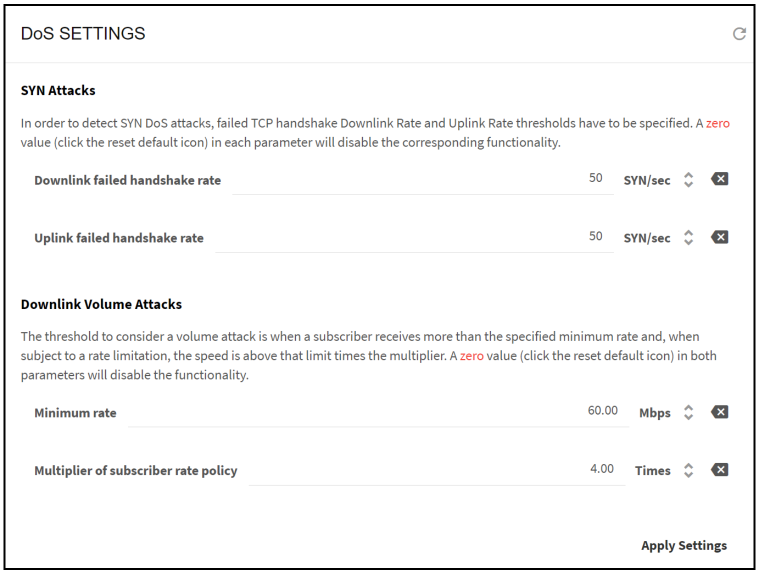

DoS SETTINGS

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**

Apply Settings

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. Figure 62 shows the DoS attack over time page.

Figure 62: *DoS attacks over time*

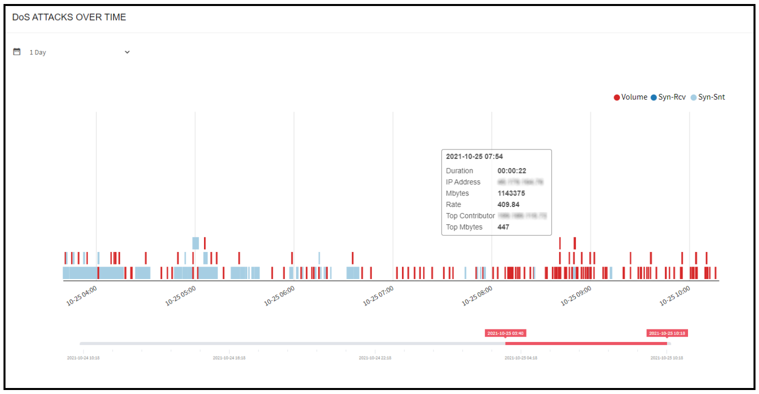

DoS ATTACKS OVER TIME

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

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.

Introduction to policies

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

There are three types 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 details, 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 we 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.
- Blocking only incoming connections from Internet of specific traffic types.
- Quota counting: decide this traffic volume counts when checking the volume quota.
- Quota limitation: what to do when the quota is reached, whether traffic is blocked or slow down to some specified speed.

Through rate policies, we control the following functionalities:

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

Through monitoring policies, we control the following functionalities:

- The amount of sampling when 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 what policy to apply depending on the traffic characteristics.

Additionally, rate policies can managed from an external system, creating them dynamically and assigning them to the subscribers. The QoE supports many 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 those subscribers without a policy assignment from the external system).

This chapter contains the following sections:

- Policy enforcement
- Check a subscriber status
- Check a policy
- Disable policy speed limits

Policy enforcement

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

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 the field **IPv6 prefix for subscribers**.

A new subscriber session is identified when the first packet from an access IP address is received. This is when the subscriber rate and monitoring rules are evaluated, to choose which policies to apply.

Check a subscriber status

You can check the rate policies applied to a subscriber in **Status > Subscribers > Subscriber Attributes**. It lists the subscribers, with the applied policy in **RATE-POLICY** column. Figure 63 shows the Subscriber attributes page.

Figure 63: *The Subscriber attributes page*

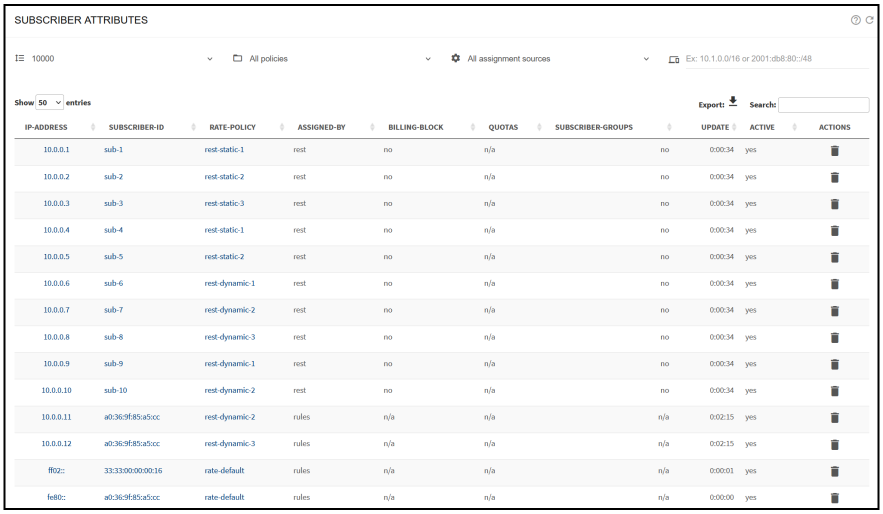

SUBSCRIBER ATTRIBUTES

10000 | All policies | All assignment sources | Ex: 10.1.0.0/16 or 2001:db8:80::/48

Show 50 entries

Export: Search:

IP-ADDRESS	SUBSCRIBER-ID	RATE-POLICY	ASSIGNED-BY	BILLING-BLOCK	QUOTAS	SUBSCRIBER-GROUPS	UPDATE	ACTIVE	ACTIONS
10.0.0.1	sub-1	rest-static-1	rest	no	n/a	no	0:00:34	yes	🗑
10.0.0.2	sub-2	rest-static-2	rest	no	n/a	no	0:00:34	yes	🗑
10.0.0.3	sub-3	rest-static-3	rest	no	n/a	no	0:00:34	yes	🗑
10.0.0.4	sub-4	rest-static-1	rest	no	n/a	no	0:00:34	yes	🗑
10.0.0.5	sub-5	rest-static-2	rest	no	n/a	no	0:00:34	yes	🗑
10.0.0.6	sub-6	rest-dynamic-1	rest	no	n/a	no	0:00:34	yes	🗑
10.0.0.7	sub-7	rest-dynamic-2	rest	no	n/a	no	0:00:34	yes	🗑
10.0.0.8	sub-8	rest-dynamic-3	rest	no	n/a	no	0:00:34	yes	🗑
10.0.0.9	sub-9	rest-dynamic-1	rest	no	n/a	no	0:00:34	yes	🗑
10.0.0.10	sub-10	rest-dynamic-2	rest	no	n/a	no	0:00:34	yes	🗑
10.0.0.11	a0:36:9f:85:a5:cc	rest-dynamic-2	rules	n/a	n/a	n/a	0:02:15	yes	🗑
10.0.0.12	a0:36:9f:85:a5:cc	rest-dynamic-3	rules	n/a	n/a	n/a	0:02:15	yes	🗑
ff02::	33:33:00:00:00:16	rate-default	rules	n/a	n/a	n/a	0:00:01	yes	🗑
fe80::	a0:36:9f:85:a5:cc	rate-default	rules	n/a	n/a	n/a	0:00:00	yes	🗑

The **ASSIGNED-BY** column indicates the origin of the policy: QoE configured rules, radius, QoE rest API or a billing system.

The list of groups this subscriber belongs to are listed in **SUBSCRIBER-GROUPS**.

If the subscriber has a quota, **QUOTA** column will show **enabled**, which is also a link to the quota status page.

Clicking on the subscriber IP address or Subscriber ID leads to the Subscriber dashboard (see Network visibility section for more information).

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

You can dig in the active flows of a subscriber in **Status > Flows > Details**, which shows the policy applied to each flow in **FLOW-POLICY** column, along with other information. Figure 64 shows the Flows per subscriber page.

Figure 64: *The Flows per subscriber page*

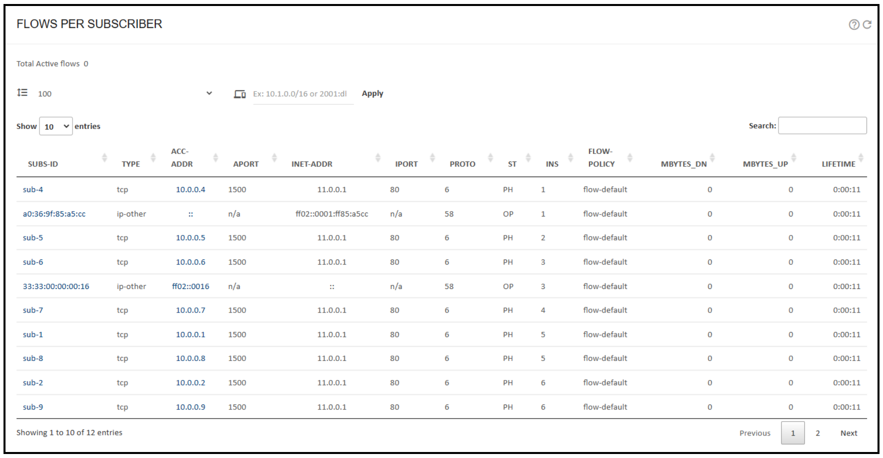

FLOWS PER SUBSCRIBER

Total Active flows 0

100 | Ex: 10.1.0.0/16 or 2001:dl | Apply

Show 10 entries | Search:

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

Status > Subscribers > Subscriber Attributes contains the information that were in **Status > Radius/REST/Billing > Subscribers** in previous versions.

Check a policy

Given a policy, it is possible to see how many subscriber IP addresses are under each policy navigating to **Status > Policies**. To check flow policies, navigate to **Status** > **Policies** > **Flow Policies**. Figure 65 shows the flow policies page.

Figure 65: *The Flow policies page*

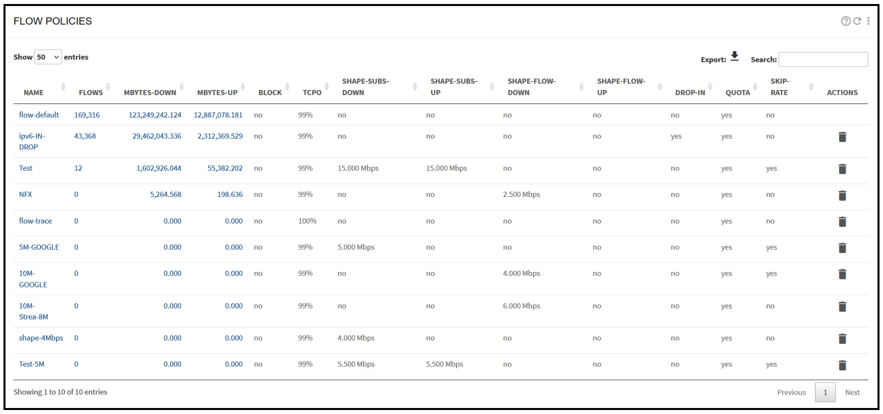

FLOW POLICIES

Show 50 entries	Export:	Search:

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	Previous 1 Next

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

To check rate policies, navigate to **Status > Policies > Rate Policies.** Figure 66 shows the rate policies page.

Figure 66: *The Rate policies page*

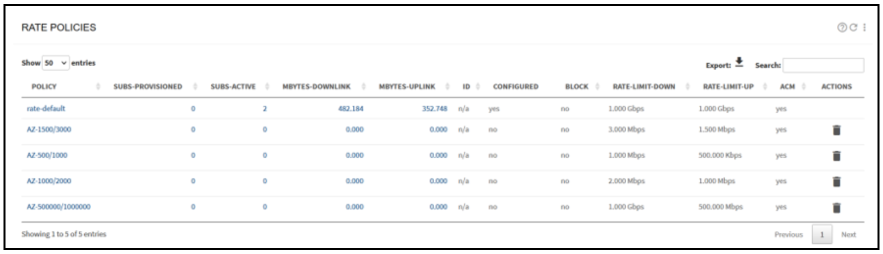

RATE POLICIES

Show 50 entries	Export:	Search:

POLICY	SUBS-PROVISIONED	SUBS-ACTIVE	MBYTES-DOWNLINK	MBYTES-UPLINK	ID	CONFIGURED	BLOCK	RATE-LIMIT-DOWN	RATE-LIMIT-UP	ACM	ACTIONS
rate-default	0	2	482.184	352.748	n/a	yes	no	1.000 Gbps	1.000 Gbps	yes	
AZ-1500/3000	0	0	0.000	0.000	n/a	no	no	3.000 Mbps	1.500 Mbps	yes	🗑
AZ-500/1000	0	0	0.000	0.000	n/a	no	no	1.000 Mbps	500.000 Kbps	yes	🗑
AZ-1000/2000	0	0	0.000	0.000	n/a	no	no	2.000 Mbps	1.000 Mbps	yes	🗑
AZ-500000/1000000	0	0	0.000	0.000	n/a	no	no	1.000 Gbps	500.000 Mbps	yes	🗑

Showing 1 to 5 of 5 entries	Previous 1 Next

SUBS-PROVISIONED says how many subscribers are associated to that policy. **SUBS-ACTIVE** shows how many of them are active (they are running traffic). If the policy definition is configured in the QoE locally, **CONFIGURED** is “yes”. For policies created dynamically via API (RADIUS, REST, Billing) it will show “no”. BLOCK “no” means that the policy does not block traffic and “yes” the opposite. The table also shows the policy rate limits in downlink and uplink directions and whether it has the ACM active or not.

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

To check monitoring policies, navigate to **Status > Policies > Monitoring Policies.** Figure 67 shows the monitoring policies page.

Figure 67: *The Monitoring policies page*

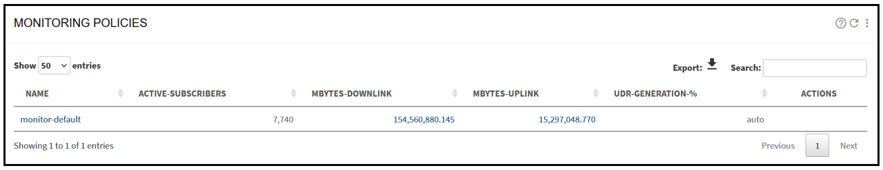

MONITORING POLICIES

Show 50 entries	Export:	Search:

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	Previous 1 Next

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

Status > **Policies** > **Rate Policies** combines the information that were in **Status** > **Radius/REST/Billing** > **Policies** and **Status** > **Policies** in previous versions.

Status > **Policies** > **Flow Policies** and **Status** > **Policies** > **Monitoring Policies** contain the information that were in **Status** > **Policies** in previous versions.

Disable policy speed limits

During testing or for any other reason, it is possible to disable speed limit enforcement globally in the node, regardless of the limits specified by the policies.

Navigate to **Configuration** > **Optimization Settings**. The following are the three speed limit types:

- Individual flow shaping (per flow): limits the speed of one single traffic flow. The limits are defined in flow policies (downlink and uplink shaping per flow).
- Aggregated flow shaping (per subscriber): limits the combined speed of all traffic flows of the same subscriber meeting the policy. The limits are defined in flow policies (downlink and uplink shaping per subscriber).
- Subscriber rate limiting: limits the combined speed of all traffic of the subscriber. The limits are defined in rate policies (maximum subscriber downlink/uplink speeds).
- Subscriber group rate limiting: limits the combined speed of all traffic of all subscribers in the group. The limits are defined in rate policies (maximum subscriber downlink/uplink speeds).

Figure 68 shows the optimization settings page.

Figure 68: *Optimization Settings page*

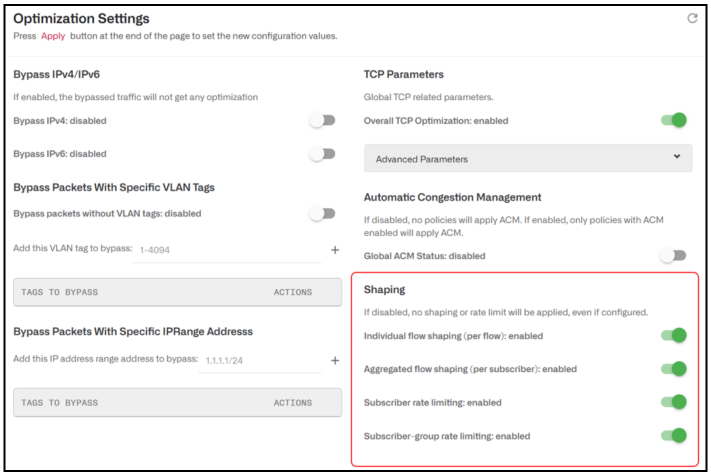

Optimization Settings

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

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

Bypass Packets With Specific IPRange Addresss

Add this IP address range address to bypass: 1.1.1.1/24

TAGS TO BYPASS	ACTIONS

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: disabled

Shaping

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

Individual flow shaping (per flow): enabled

Aggregated flow shaping (per subscriber): enabled

Subscriber rate limiting: enabled

Subscriber-group rate limiting: enabled

Note

While enforcement of speed limits are disabled, speeds may go and will most likely go above those limit. For example, if subscriber rate limit is disabled, subscriber plans will not be enforced.

Configured policies

Locally configured policies are selected using rules that combine profiles with policies:

- Profiles classify the traffic according to some criteria (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 address are in some subnet will have a specific rate limit.

This chapter contains the following sections:

Profiles

Profiles classify the traffic, and, when used by policy rules, determine the policy applied to a subscriber or a flow. There are different profile types, according to the properties being used for traffic classification. The current 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-*” will match subscriber traffic with rate policies named “premium-gold” and “premium-platinum”.
- **Internet Profile:** identifies the traffic involving IP addresses, ports or L4 protocols on the Internet side.
- **Access Profile**: identifies the traffic involving IP addresses, ports or L4 protocols on the Access side.

- **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* will match 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.
- **Deep Packet Inspection (DPI) Profile**: identifies the flows that have a DPI domain that matches one of the domain patterns (signatures) 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). See at the end of this section how to add custom signatures.

Profiles are configured in the menu option **Configuration > Profiles**.

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 the following example, a profile is defined for one of the two wires of the QoE server and another profile for a second wire. Figure 69 shows the interface profiles page.

Figure 69: *The Interface profiles page*

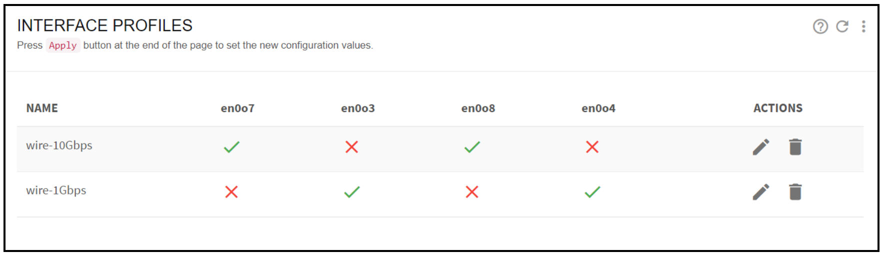

INTERFACE PROFILES

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	✗	✓	✗	✓	✎ 🗑

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.

The following example defines two profiles for two network areas and a third profile for traffic without VLAN tag. Figure 70 shows the VLAN profiles page.

Figure 70: *The VLAN profiles page*

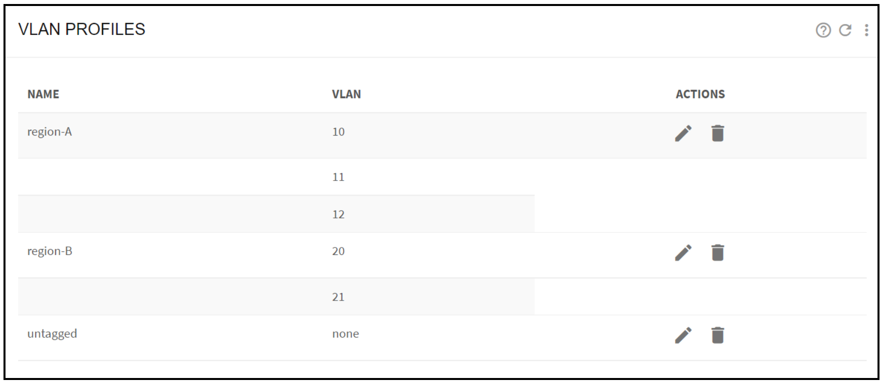

VLAN PROFILES

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

Policy rate profile

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 71 shows the policy rate profiles page.

Figure 71: *The Policy-rate profiles page*

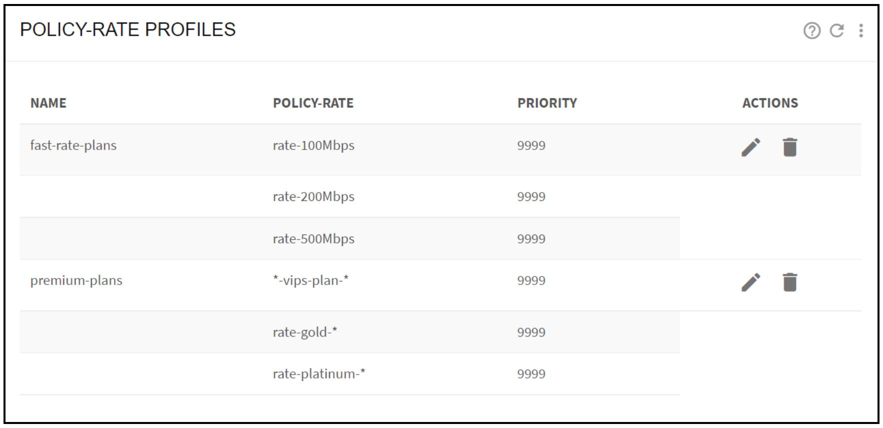

POLICY-RATE PROFILES

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	

Internet and access profiles

Internet profiles check some traffic properties of the end points on the Internet side (content servers). Access profiles do the same for the end points on the access side (subscribers).

The following are the traffic properties:

- IP addresses.
- IP address ranges.
- Layer 4 protocol number (example ICMP, TCP, UDP, OSPF, and so on).
- Port numbers (for TCP and UDP protocols).

Figure 72 shows an example of an access profile containing all private IPv4 ranges.

Figure 72: *The Access profiles page*

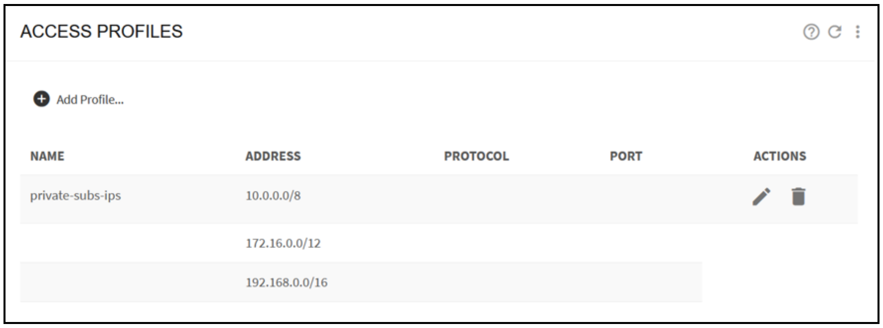

ACCESS PROFILES

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

To add an entry, click on **Add Entry...** menu option on the upper right. Figure 73 shows one IP address entry, matching any protocol.

Figure 73: *Add Access Profile and entry page*

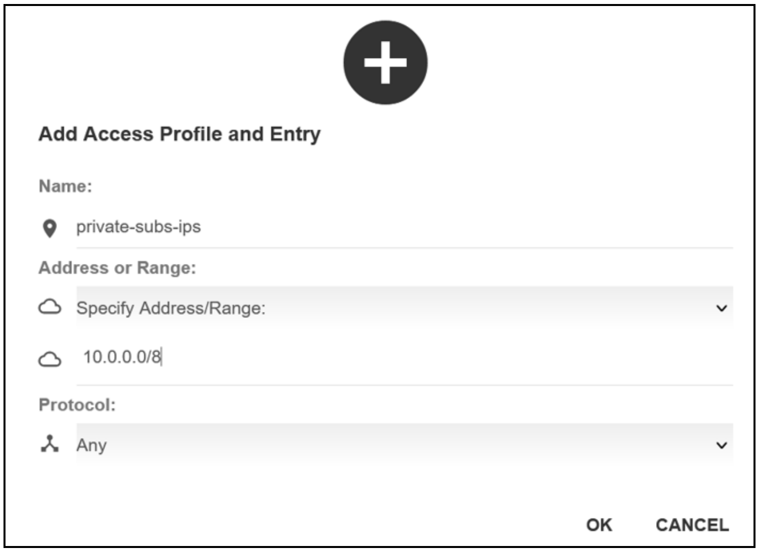

Add Access Profile and Entry

Name: private-subs-ips

Address or Range: Specify Address/Range:

10.0.0.0/8

Protocol: Any

Figure 74 shows some Internet profile examples.

Figure 74: *Internet Profile examples*

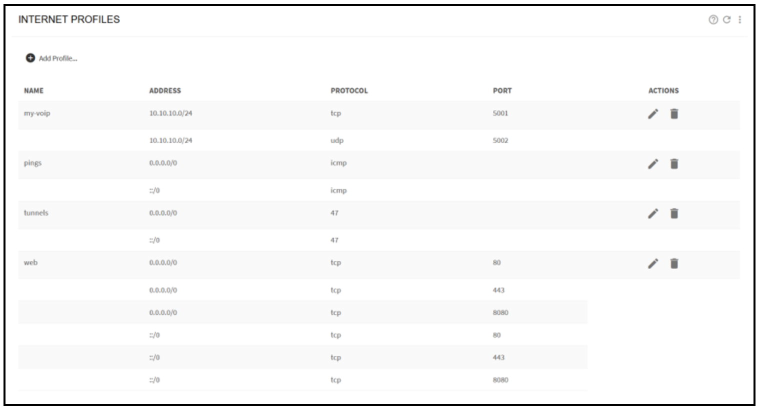

INTERNET PROFILES

NAME	ADDRESS	PROTOCOL	PORT	ACTIONS
my-voip	10.10.10.0/24	tcp	5001	
	10.10.10.0/24	udp	5002	
pings	0.0.0.0/0	icmp		
	::/0	icmp		
tunnels	0.0.0.0/0	47		
	::/0	47		
web	0.0.0.0/0	tcp	80	
	0.0.0.0/0	tcp	443	
	0.0.0.0/0	tcp	8080	
	::/0	tcp	80	
	::/0	tcp	443	
	::/0	tcp	8080	

The profile **my-voip** includes the IP range of a VoIP service that uses TCP and UDP ports 5001 and 5002.

The profile **pings** will match ICMP traffic to any IPv4 or IPv6 destination (0.0.0.0/0 matches any IPv4 address and ::/0 any IPv6 address).

The profile ***web*** matches any IPv4 or IPv6 traffic to TCP ports 80, 8080 and 443.

To add an entry, click on ***Add Entry...*** menu option on the upper right.

Figure 75 shows an entry matching any IPv4 traffic to TCP port 80.

Figure 75: *Add Internet Profile Entry page*

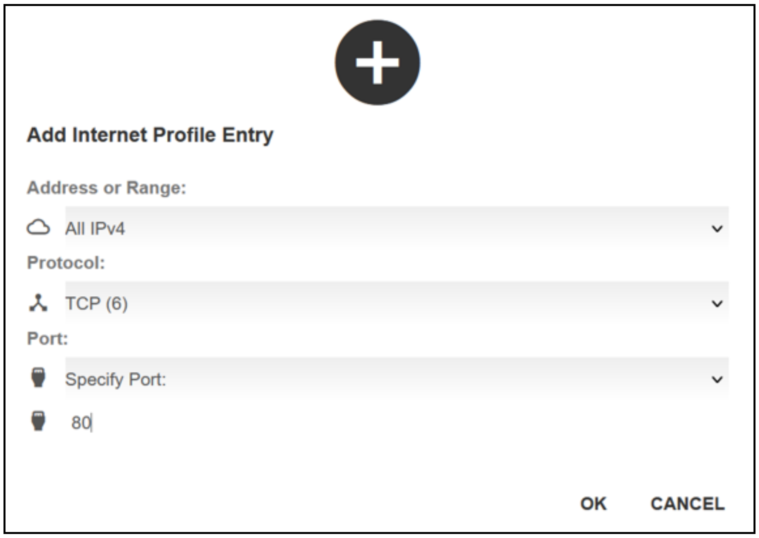

Add Internet Profile Entry

Address or Range:

All IPv4

Protocol:

TCP (6)

Port:

Specify Port:

80

It is possible to load an IP address list from a text file. The text file format has 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 76 shows the subscriber group profiles page.

Figure 76: *The Subscriber group profiles page*

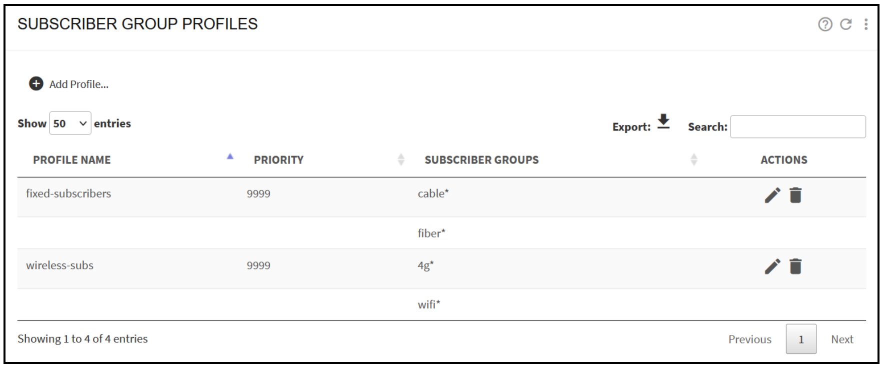

SUBSCRIBER GROUP PROFILES

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

Showing 1 to 4 of 4 entries

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 77 shows the Subscriber ID profiles page.

Figure 77: *The Subscriber ID profiles page*

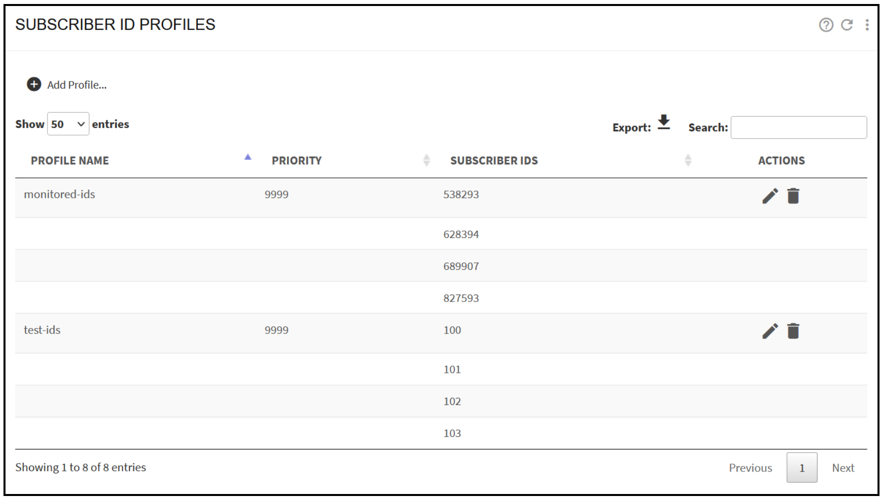

SUBSCRIBER ID PROFILES

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

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 78 shows the time profiles page.

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

Figure 78: *The time profiles page*

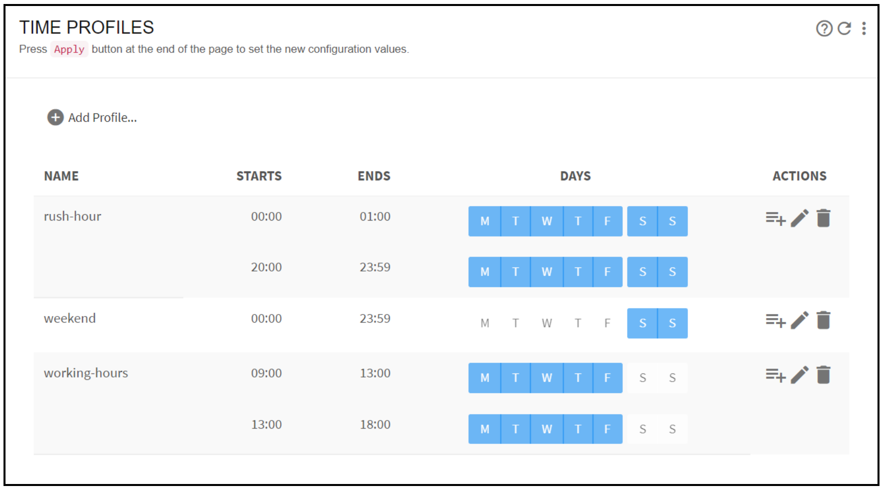

TIME PROFILES

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	

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 79 shows the throughput profiles page.

Figure 79: *The throughput profiles page*

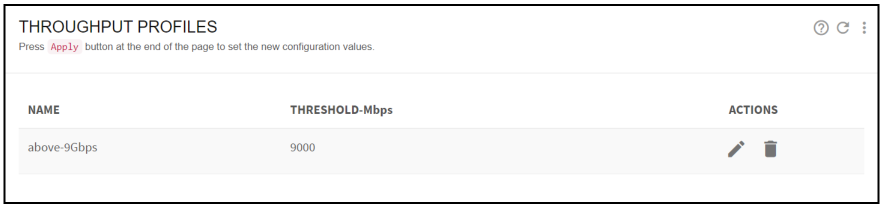

THROUGHPUT PROFILES

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

NAME	THRESHOLD-Mbps	ACTIONS
above-9Gbps	9000	

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...**

Some of the predefine sets are:

- **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 flow policies

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

- **TCP Optimization**: Improves TCP traffic performance. It specifies whether to apply optimization to TCP traffic. It is recommended to enable it (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 will no longer be affected by the rate limitation specified in the rate policy for this subscriber. They will only get the rate limitation specified by this flow policy (if any).

These policies are configured in the menu option **Configuration** > **Subscriber Flows** by selecting the **POLICIES** tab.

Shaping per subscriber

The following example defines a downlink speed limit of 10Mbps, an uplink limit of 8Mbps, and bursts of 3 seconds of double the normal speed. Figure 80 shows the Edit subscriber flow policy page.

Figure 80: *The Edit subscriber flow policy page*

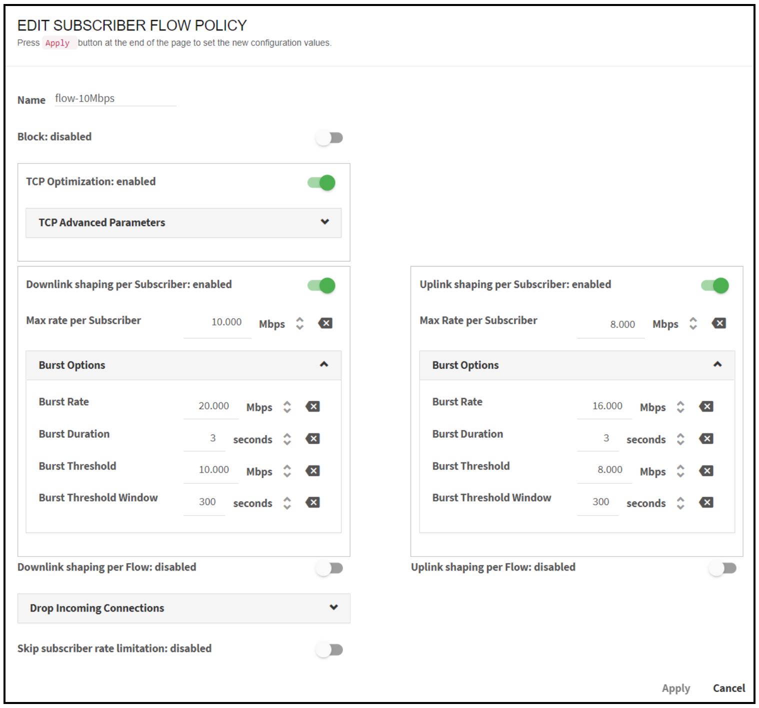

EDIT SUBSCRIBER FLOW POLICY

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

Name flow-10Mbps

Block: disabled

TCP Optimization: enabled

TCP Advanced Parameters

Downlink shaping per Subscriber: enabled

Max rate per Subscriber 10.000 Mbps

Burst Options

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 Options

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

Drop Incoming Connections

Skip subscriber rate limitation: disabled

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 to sustain.
- **Burst Threshold**: An average speed that, if exceeded, prevent a new burst from happening. It is the way to control when a new burst is 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 81: *Burst options*

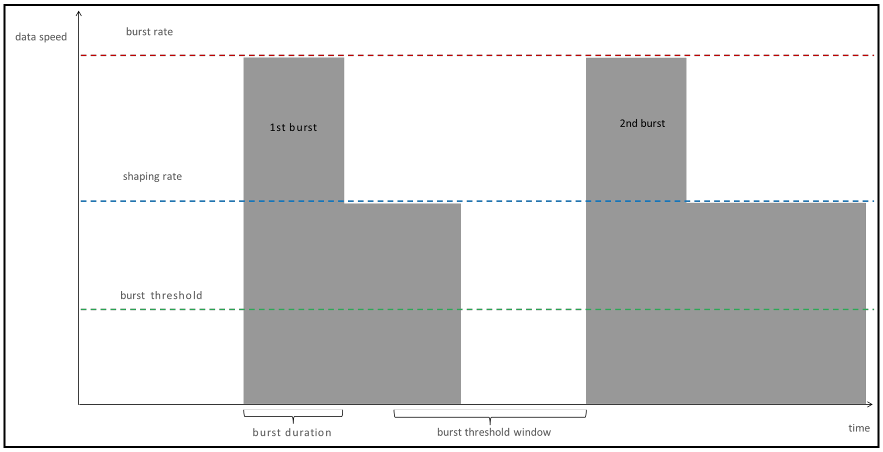

Shaping per flow

The following example is a policy with a limit per flow of 4Mbps in either direction. Figure 82 shows the edit subscriber flow policy.

Figure 82: *The Edit subscriber flow policy*

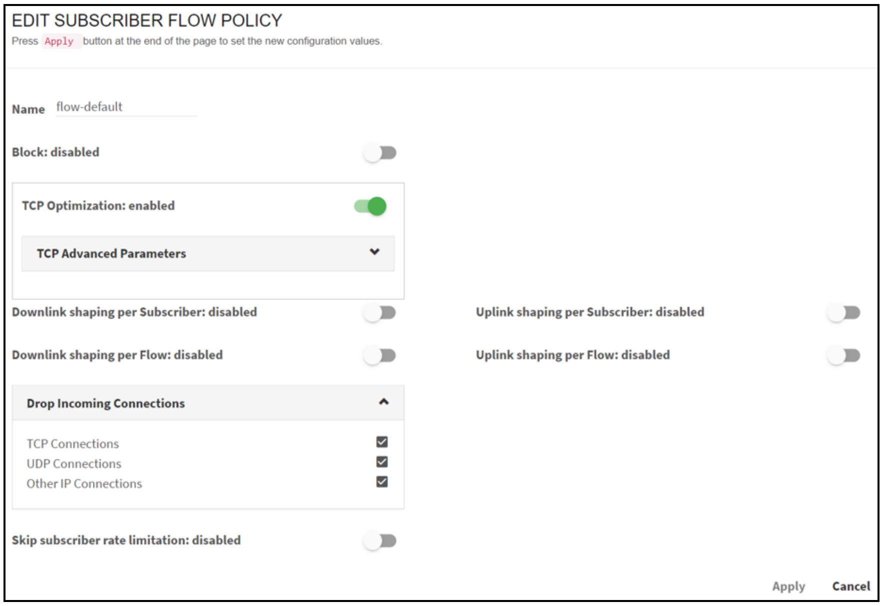

EDIT SUBSCRIBER FLOW POLICY

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

Uplink shaping per Subscriber: disabled

Downlink shaping per Flow: disabled

Uplink shaping per Flow: disabled

Drop Incoming Connections

TCP Connections ☑

UDP Connections ☑

Other IP Connections ☑

Skip subscriber rate limitation: disabled

It is possible to add a shaping per subscriber. Per flow and per subscriber shaping limits will 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 do so, there is a **Drop Incoming Connections** section as part of a **Subscriber Flow** policy. Figure 83 shows the edit subscriber flow policy page.

Figure 83: *The edit subscriber flow policy page*

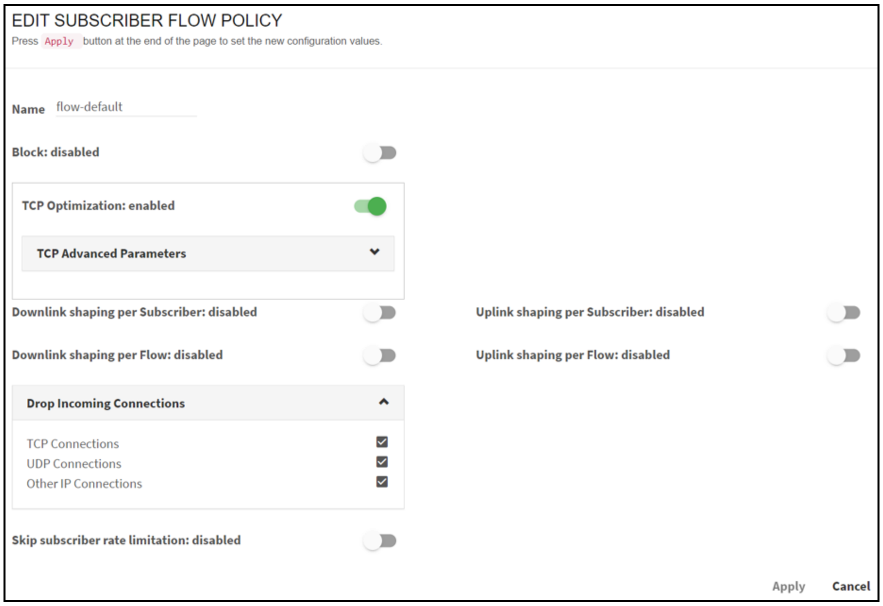

EDIT SUBSCRIBER FLOW POLICY

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

Name flow-default

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

- TCP Connections ☑
- UDP Connections ☑
- Other IP Connections ☑

Skip subscriber rate limitation: disabled

Subscriber rate policies

Subscriber rate policies are applied per subscriber. Figure 84 shows the edit subscriber rate policy.

The following are the 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 will detect congestion and select a rate limit automatically (enabled by default).

Figure 84: *The Edit subscriber rate policy*

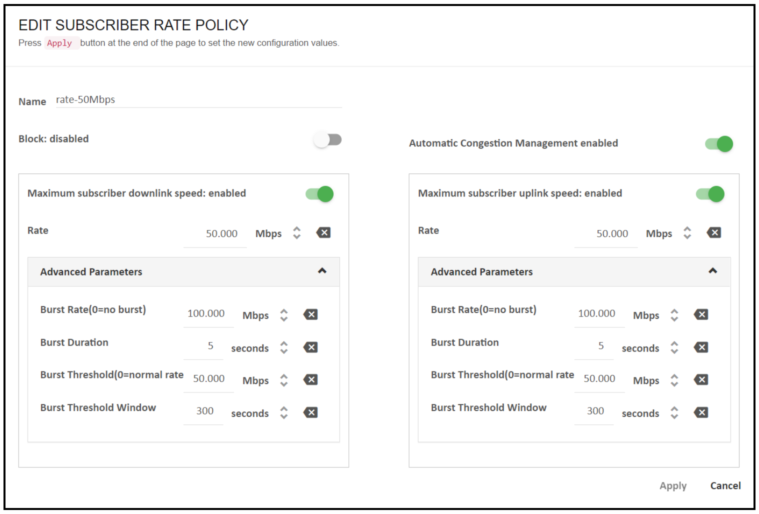

EDIT SUBSCRIBER RATE POLICY

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

To configure the policies, navigate to **Configuration > Subscriber Rates** and select **POLICIES** tab.

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-configured rules (configured rules are a fallback for subscribers without a policy assigned by the external interface).

Enable/disable ACM optimization

The ACM is enabled by default.

To enable or disable ACM in a configured policy, change the **Automatic Congestion Management** field in the Subscriber Rate window. Figure 85 shows the edit subscriber rate policy page.

Figure 85: *The Edit subscriber rate policy page*

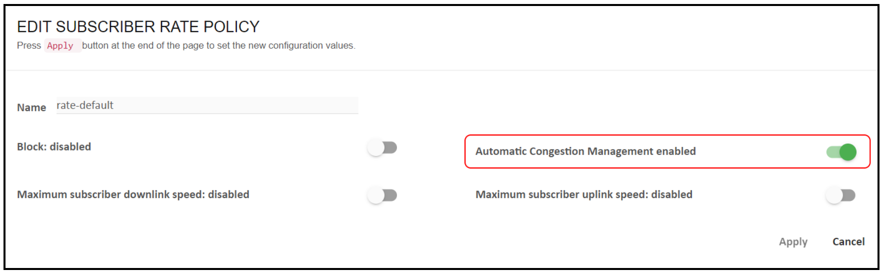

EDIT SUBSCRIBER RATE POLICY

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

The 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 too much CPU and disk to produce. The monitor-default is the recommended mode of operation. Figure 86 shows the edit subscriber monitoring policy page for monitor-default.

Figure 86: *The Edit subscriber monitoring policy page for monitor-default*

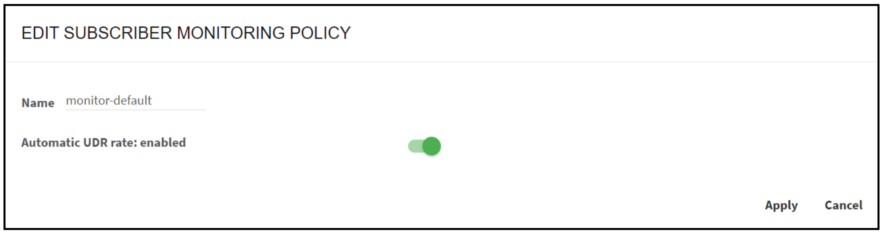

EDIT SUBSCRIBER MONITORING POLICY

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 percent or 100 percent). UDRs for 100 percent of flows will produce and accurate description of the subscriber usage statistics, but for performance reasons, it is recommended to set a 100 percent UDR monitoring policies only on a few subscribers at a time, because these policies can produce a huge number of records that could use all the available disk space and could also make the production of DPI statistics very slow and CPU-consuming. Figure 87 shows the edit subscriber monitoring policy page for monitor-full.

Figure 87: *The Edit subscriber monitoring policy page for monitor-full*

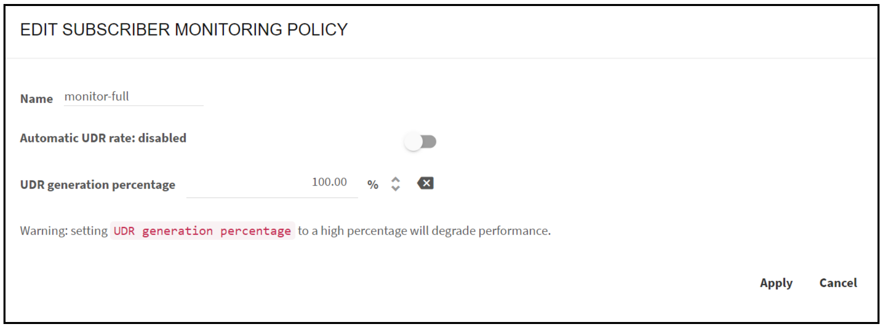

EDIT SUBSCRIBER MONITORING POLICY

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 which configured policies to assign 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 subscriber flow policy for each flow, subscriber rate rules select the appropriate configured 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 will be 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, see 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 in **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 the following 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, would be selected.

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

Figure 88: *The Subscriber flows configuration page*

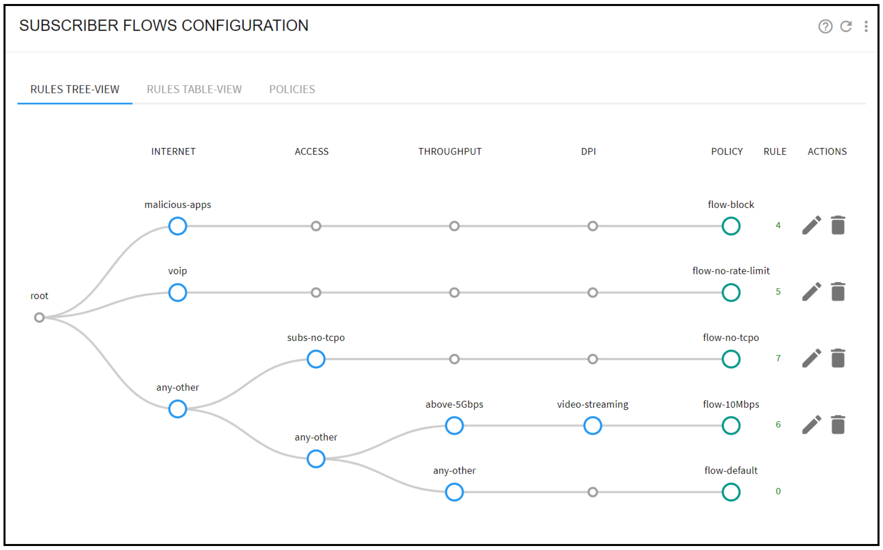

If there are common elements in two profiles of the same type and therefore a rule conflict, the decision tree will signal it so the rules can be reviewed by the operator and the conflict corrected. In the following example, two access profiles have an overlap (they both contain the same IP range). This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

Figure 89: *The Subscriber flows configuration page*

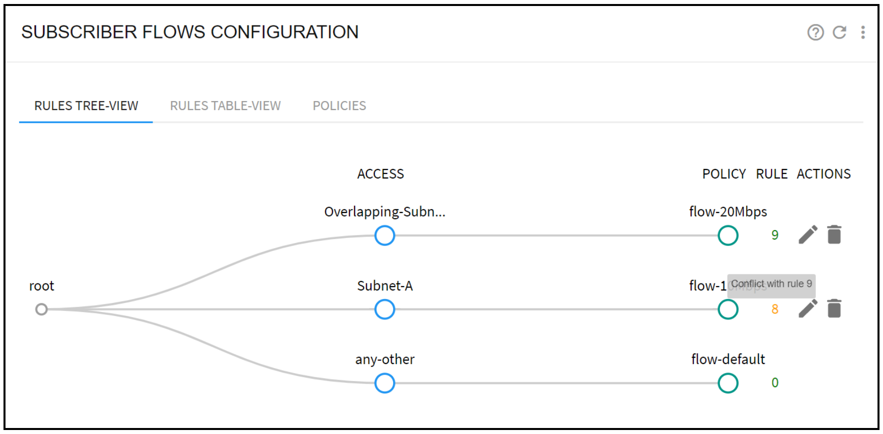

Subscriber rate decision tree

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 will 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. It is available in **Configuration** > **Subscriber Rate** > **Rules** selecting the tab **RULES TREE-VIEW**.

Figure 90: *The Subscriber rate configuration page*

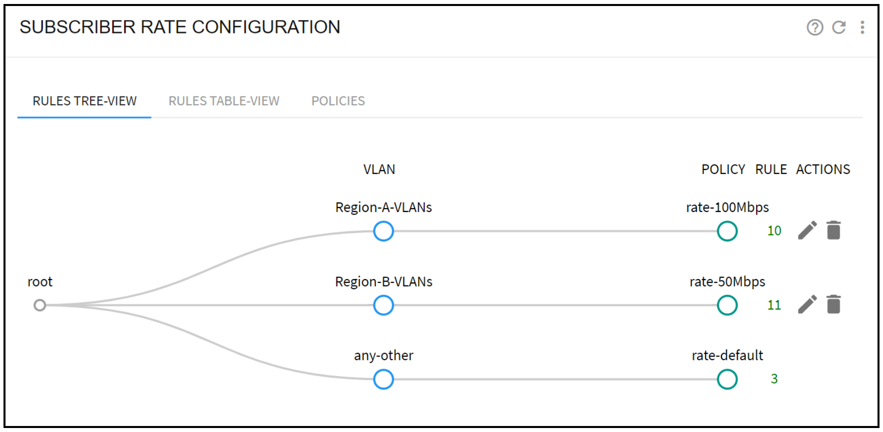

Subscriber monitoring decision tree

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 will specify 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 will 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. It is available in **Configuration** > **Subscriber Monitoring** > **Rules** by selecting the **RULES TREE-VIEW** tab

Profile explicit priority

As explained in previous sections, profiles are evaluated in a priority order based on the profile type (some types take precedence over other types) and, for profiles of the same type, on what profiles are more specific. Most rule sets can be defined relying on this implicit priority order.

However, some profiles use string patterns with wild chars, like ***youtube*** or ***mysubscriberId****, for which it is not easy to decide which profile is more specific. For that reason, those profile types have a priority number if the rules need to decide which one will take precedence over other matching profile of the same type. The lower the number, the higher the profile priority.

For example, a profile with a priority of 10 will take precedence over another profile of the same type with priority 20.

The profile types that have a numerical priority are:

- Policy rate profile
- Subscriber group profile
- Subscriber ID profile
- DPI

Profiles of these types are created with a default value of 9999 (the lowest possible priority, meaning none takes precedence unless the value is changed to a lower number).

Policy examples

Several common examples of policies follow.

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. Furthermore, it uses independent queues per flow and that makes application latencies independent of each other, which greatly improves the experience of interactive applications. Figure 91 shows the queue structure, with a queue per flow and policy control at flow and subscriber levels.

Figure 91: *Queue structure*

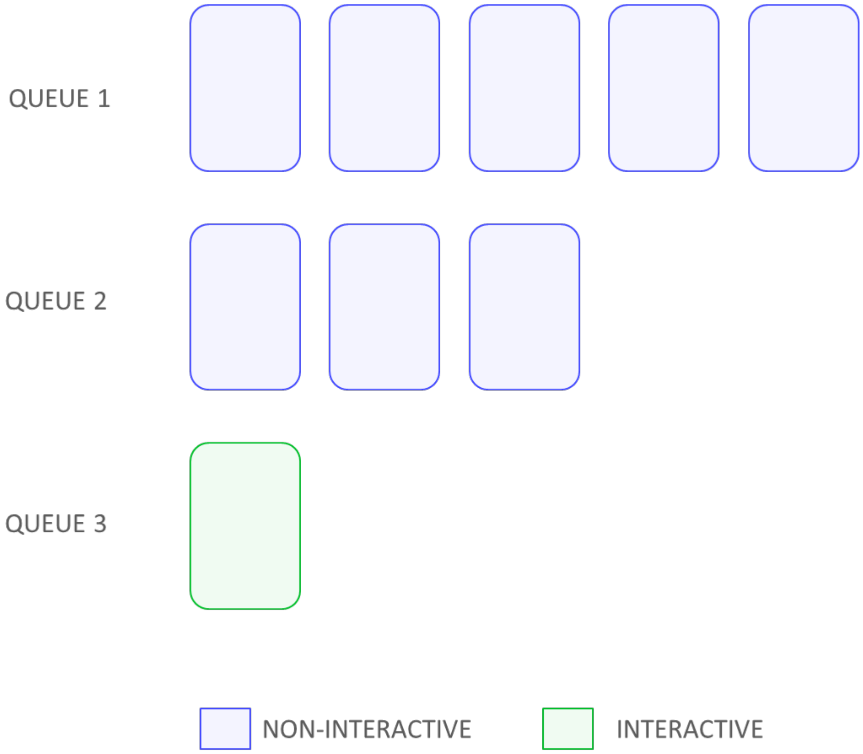

The easiest way to implement Rate Plans is to use the RADIUS, REST or Billing interfaces (see the corresponding sections). Rate policies will then be assigned by an external system for each subscriber. The external system can directly create those policies, or it can assign rate policies configured from the QoE GUI.

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 92 shows an example of the rule tree.

Figure 92: *Rule tree*

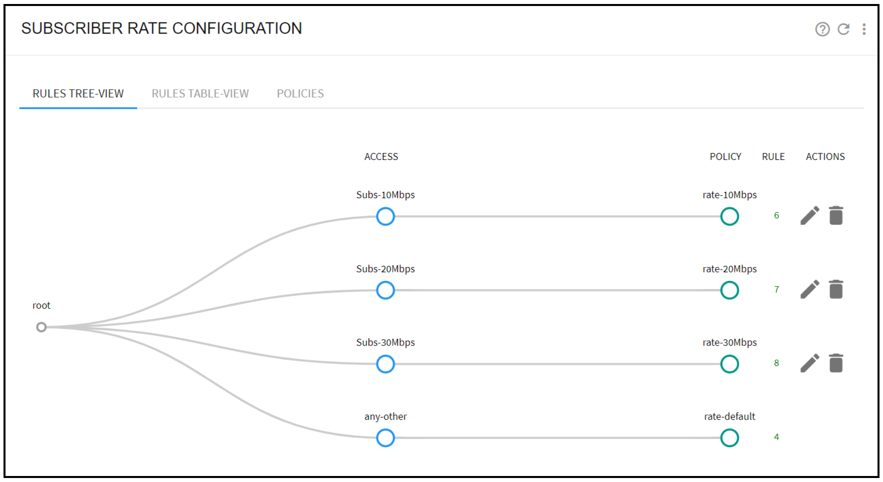

ACCESS	POLICY	RULE	ACTIONS
Subs-10Mbps	rate-10Mbps	6	✎ 🗑
Subs-20Mbps	rate-20Mbps	7	✎ 🗑
Subs-30Mbps	rate-30Mbps	8	✎ 🗑
any-other	rate-default	4	

It is also possible to define such 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 93 shows the subscriber rate configuration page.

Figure 93: *Subscriber rate configuration*

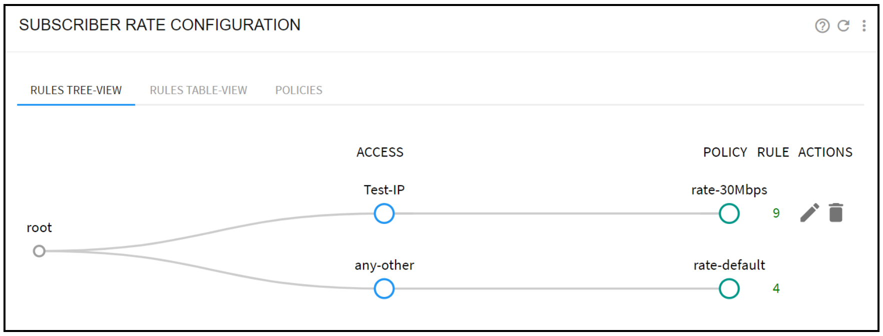

Limiting the speed of some applications

The goal is 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 94 shows the subscriber flows configuration page.

Figure 94: *The Subscriber flows configuration page*

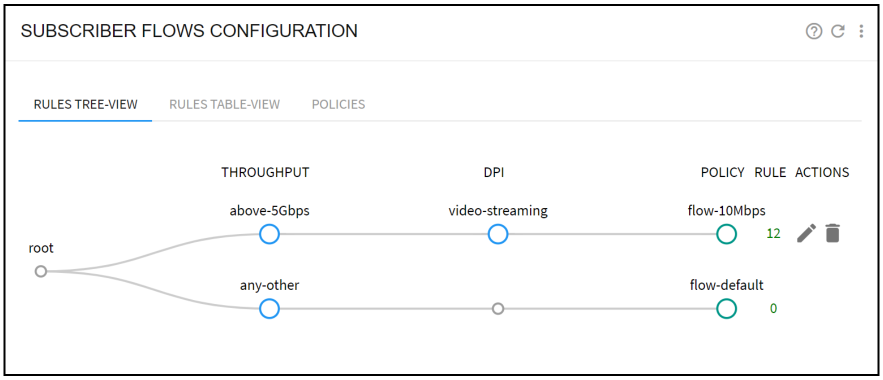

Limiting the speed of some applications with NAT

The goal is 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.

We use the same Throughput and DPI profiles of the previous example, but we use a per-flow shaping policy. Figure 95 shows the edit subscriber flow policy page.

Figure 95: *The Edit subscriber flow policy page*

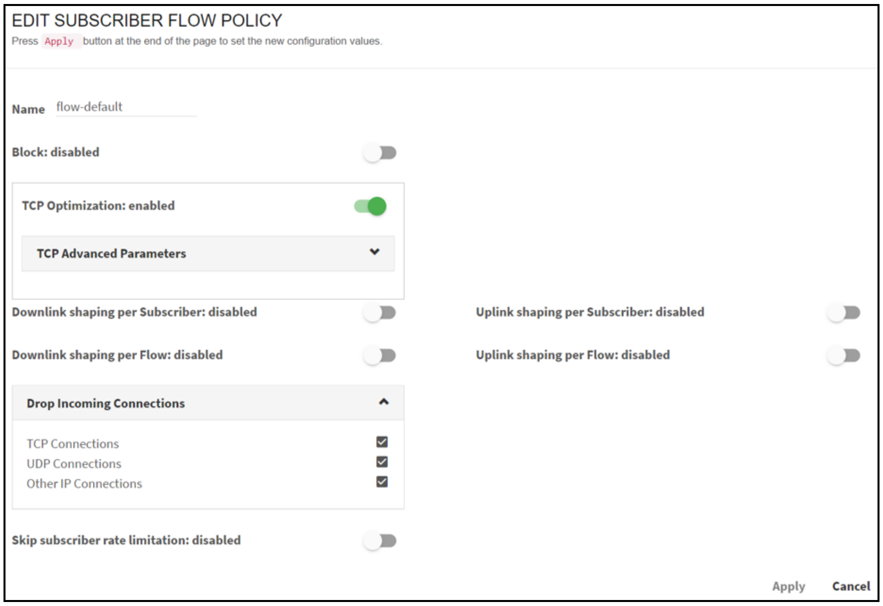

EDIT SUBSCRIBER FLOW POLICY

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

Name flow-default

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

TCP Connections ☑
UDP Connections ☑
Other IP Connections ☑

Skip subscriber rate limitation: disabled

Services not limited by the subscriber rate

The goal is to preserve the quality of experience of some services by granting throughput to them even when the subscriber rate plan is fully used. An example could be a 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.

Figure 96: *The policy-rate profiles page*

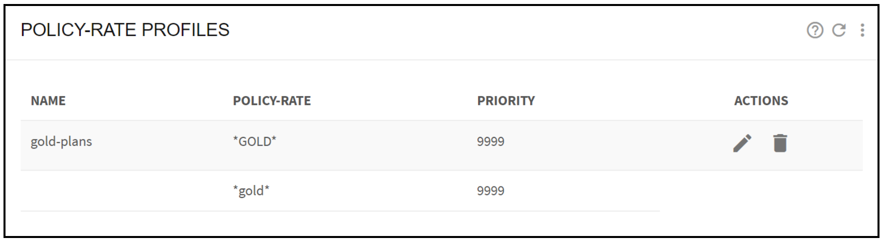

POLICY-RATE PROFILES

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

Figure 97: *The edit subscriber flow policy*

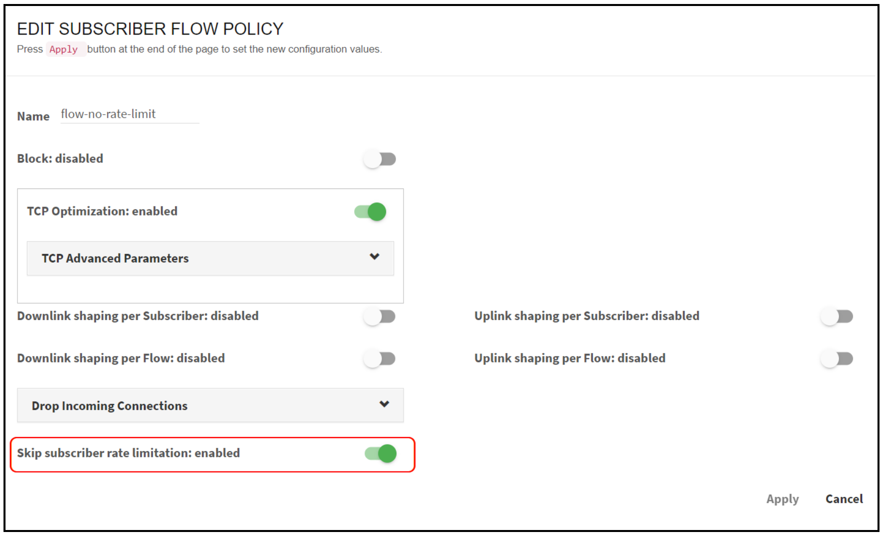

EDIT SUBSCRIBER FLOW POLICY

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 98: *The subscriber flows configuration page*

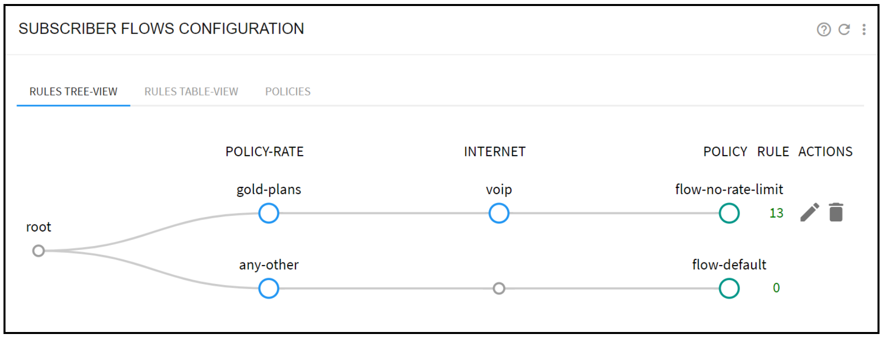

SUBSCRIBER FLOWS CONFIGURATION

Blocking applications

In this scenario, some applications need blocking, for example, servers that are sources of attacks. To do 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.

Figure 99: *The edit subscriber flow policy page*

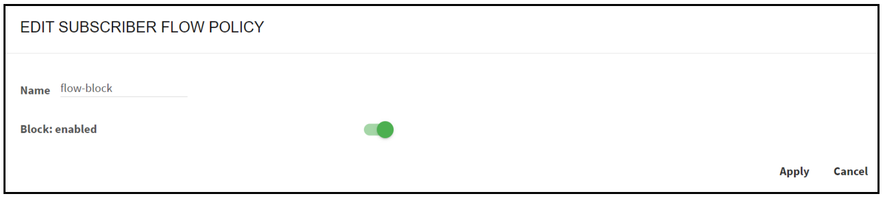

EDIT SUBSCRIBER FLOW POLICY

Name flow-block

Block: enabled

Figure 100: *The subscriber flows configuration page*

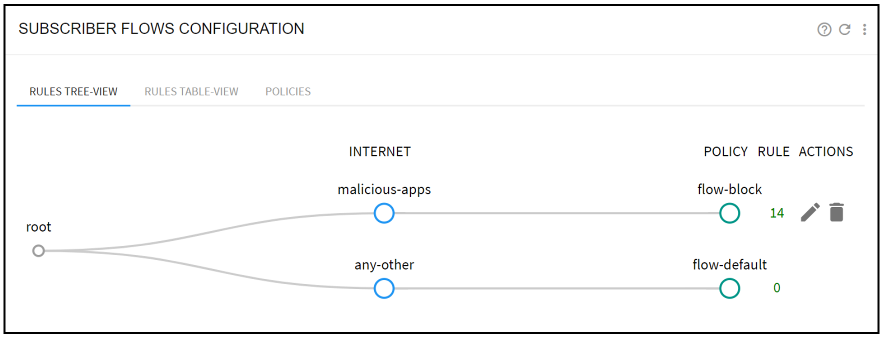

Exclude traffic from TCP optimization

The objective is that the QoE does not optimize certain 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. Next, 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.

Figure 101: *The edit subscriber flow policy page*

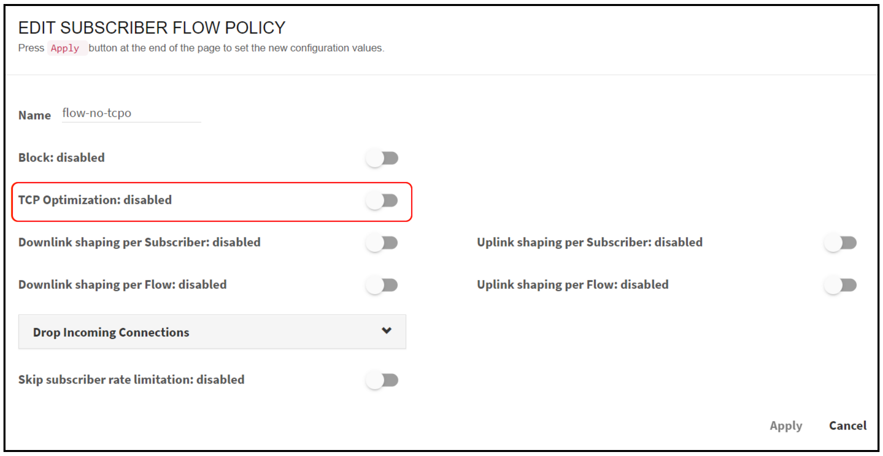
EDIT SUBSCRIBER FLOW POLICY

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

Name flow-no-tcpo

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

Figure 102: *The subscriber flows configuration page*

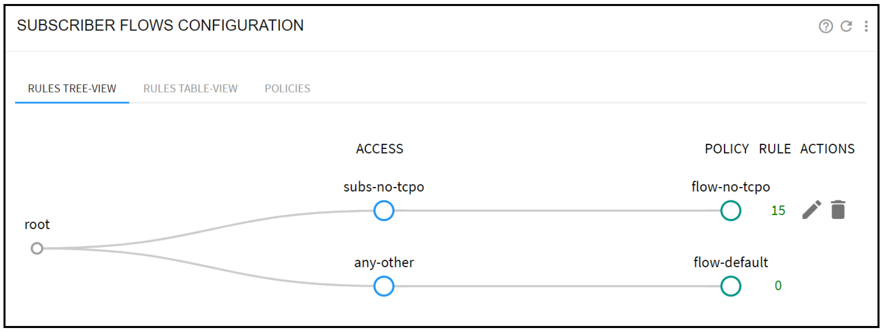

This setup is equivalent to the IP address blacklist in QoE Release 3 software.

Another example is to use an Internet profile to exclude some applications per TCP port.

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 to an IP address at the same time, in which case the restriction happens when any of the quotas is exhausted.

Quotas are assigned to IP addresses. If a subscriber changes the IP address to a new one, 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
- Associating quotas to subscriber IPs
- Checking the quota state
- Slow down when quota is exhausted
- Captive portal policy
- Quotas managed using REST API
- Quotas managed from RADIUS

Quota general configuration

To configure general aspects of quota behavior, navigate to: **Status > Subscribers > Subscriber Quotas** and extend **Advanced Quota Parameters**. Figure 103 shows the subscriber quotas page.

Figure 103: *The Subscriber quotas page*

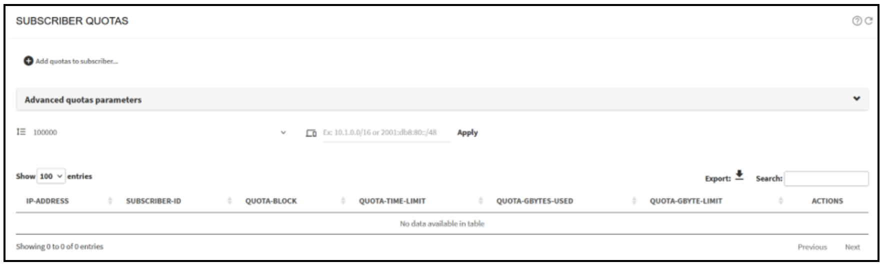

SUBSCRIBER QUOTAS

Add quotas to subscriber...

Advanced quotas parameters

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

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

Subscriber initial quota status defines the oprations with IP addresses without an assigned quota:

- When set to disabled (by default), traffic is allowed, without restrictions.
- When set to blocked, traffic will be 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.

HTTPS redirections are not supported, because modern browsers are protected against redirection attempts for security reasons.

> **Note**
>
> Though only HTTP redirections are supported, the site to redirect the traffic can be HTTPS, where the URL used is https://my-captive-portal.com).

If HTTP redirections are used, a policy is needed to allow the traffic to the redirection sites (and associated DNS queries). For more information, see Captive portal policy 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**.

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

Figure 104 defines a time quota as an absolute time.

Figure 104: *Associating quotas to subscriber IPs*

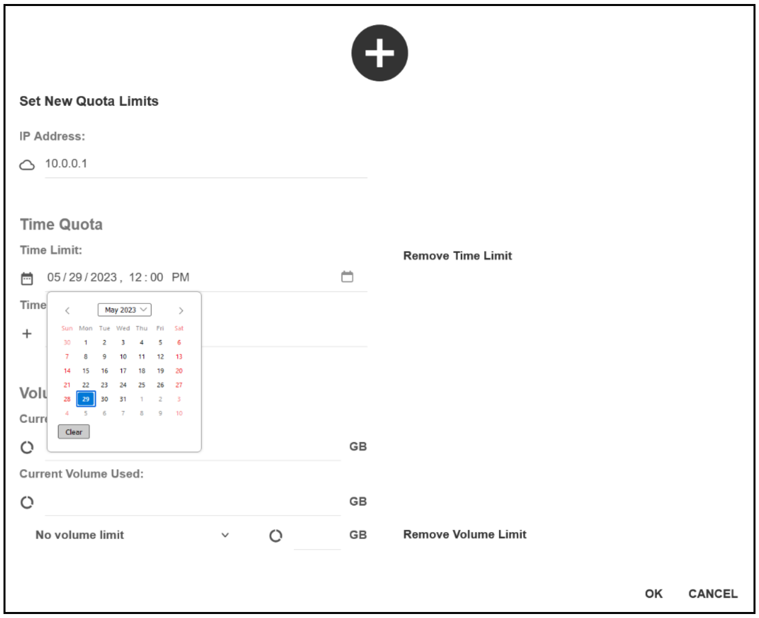

Set New Quota Limits

IP Address: 10.0.0.1

Time Quota

Time Limit: 05 / 29 / 2023 , 12 : 00 PM

Current Volume Used:

No volume limit

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

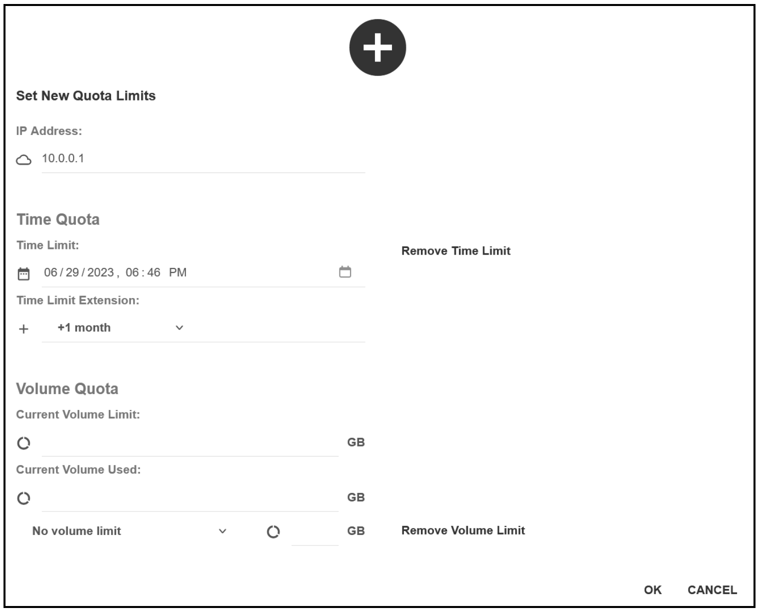

Set New Quota Limits

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

OK CANCEL

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

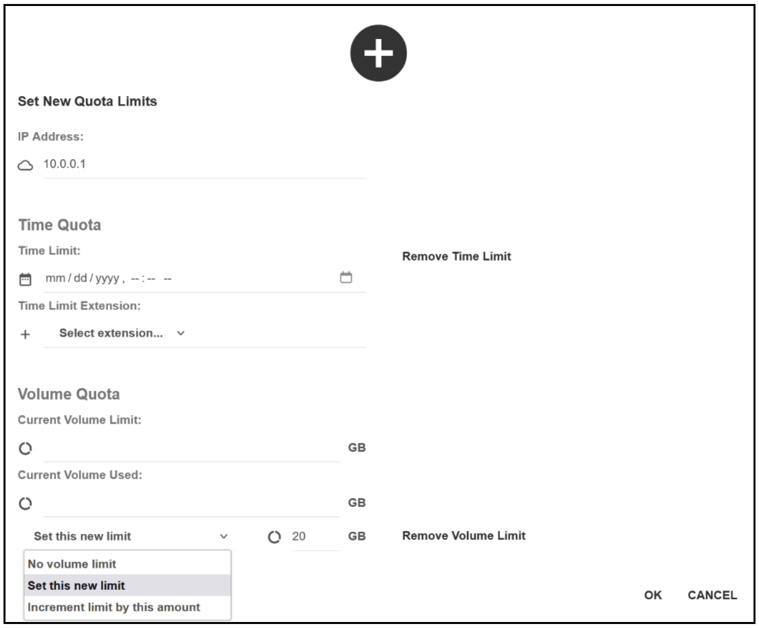

Set New Quota Limits

IP Address: 10.0.0.1

Time Quota

Time Limit: mm / dd / yyyy , -- : -- --

Remove Time Limit

Time Limit Extension: Select extension...

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

OK CANCEL

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

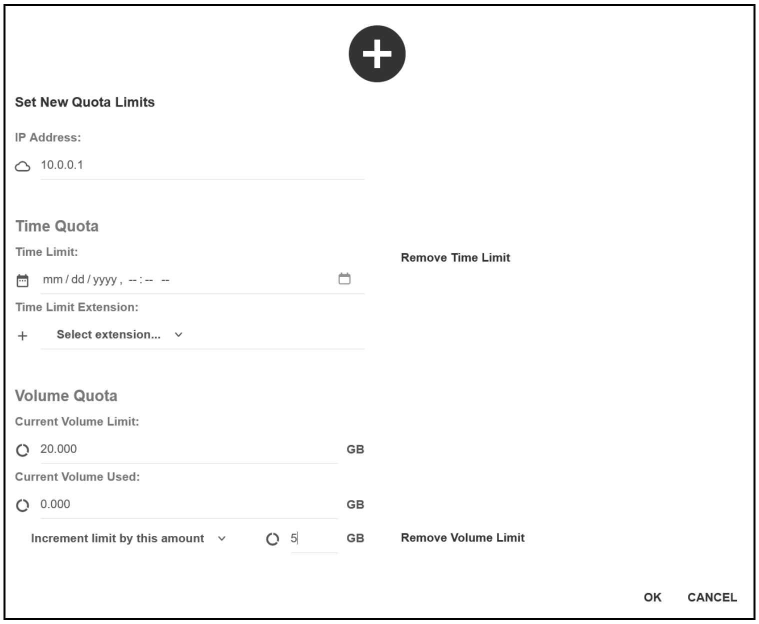

Set New Quota Limits

IP Address: 10.0.0.1

Time Quota

Time Limit: mm / dd / yyyy , -- : -- --

Remove Time Limit

Time Limit Extension: Select extension...

Volume Quota

Current Volume Limit: 20.000 GB

Current Volume Used: 0.000 GB

Increment limit by this amount: 5 GB

Remove Volume Limit

OK CANCEL

The time and volume quotas can coexist, and in this case the subscriber traffic will be restricted when either of the two become exhausted.

It is also possible to edit the quota and remove its time or volume component pressing **Remove Time Limit** or **Remove Volume Limit** respectively.

Checking the quota state

To view the status of subscriber quotas, navigate to: **Status** > **Subscribers** > **Subscriber Quotas**. The volume quotas displays the consumption. Figure 105 shows the subscriber quotas page.

Figure 105: *The Subscriber quotas page*

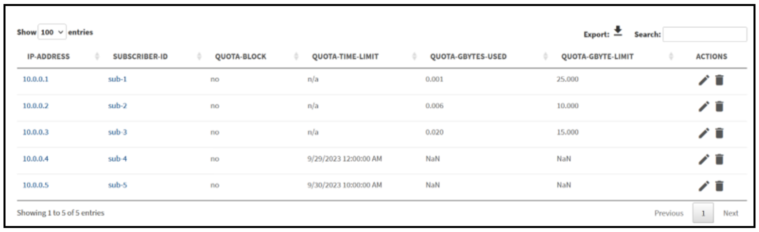

Show 100 entries Export: Search:

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 Previous 1 Next

In Figure 105, 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 is 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 106 changes the flow-default policy, so it slows down traffic when the quota is exhausted.

Figure 106: *The Edit subscriber flow policy*

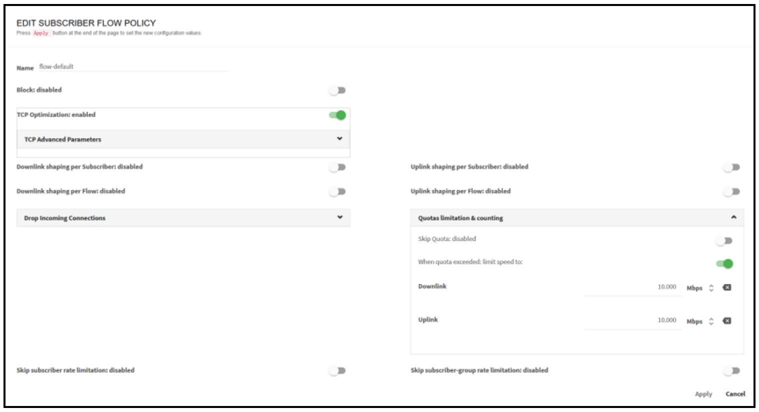

EDIT SUBSCRIBER FLOW POLICY

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

Skip Quota: disabled

When quota exceeded: limit speed to:

Downlink 10.000 Mbps

Uplink 10.000 Mbps

Skip subscriber rate limitation: disabled

Skip subscriber-group rate limitation: disabled

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 107 shows the subscriber flows configuration page.

In the following example, two traffic categories need 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).

Figure 107: *The subscriber flows configuration page*

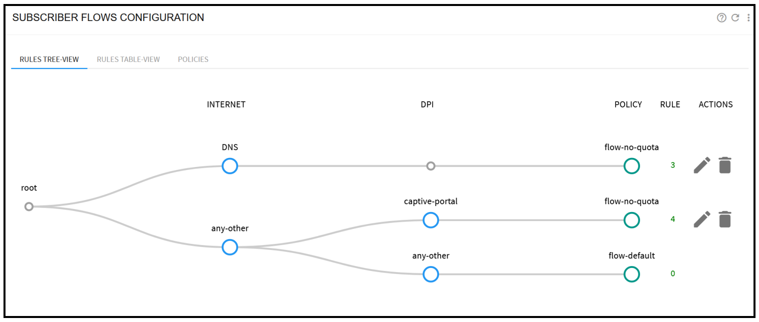

SUBSCRIBER FLOWS CONFIGURATION

The policy that is not affected by quota exhaustion has the quota switch set to OFF. Figure 108 shows the Edit subscriber flow policy page.

Figure 108: *The Edit subscriber flow policy page*

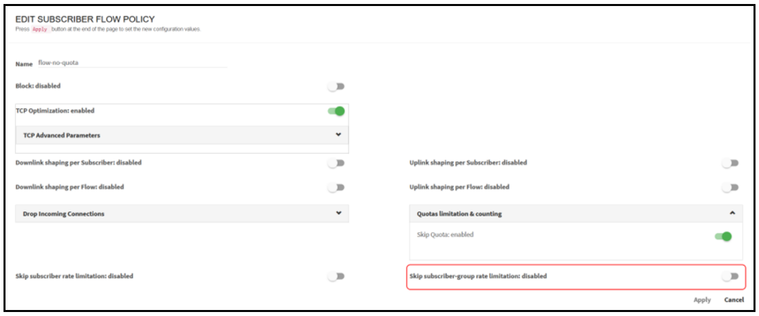

EDIT SUBSCRIBER FLOW POLICY

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

Uplink shaping per Subscriber: disabled

Uplink shaping per Flow: disabled

Quotas limitation & counting

Skip Quota: enabled

Skip subscriber rate limitation: disabled

Skip subscriber-group rate limitation: disabled

Quotas managed using REST API

In addition to the GUI, the QoE REST API can be used to manage time and volume quotas.

For more REST API commands, see REST API reference.

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, use the following command:

```
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:

```
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:

```
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 Kbytes of traffic. The QoE convention is that 1 Kbyte is 1000 bytes.

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

```
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:

```
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:

```
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 at the same time

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

```
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}}'
```

To extend the volume quota in 500 MB, 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}}'
```

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 quota state

To check the quota state through the REST API, use the following command:

```
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" : ""
}
```

Quotas managed from RADIUS

QoE in RADIUS proxy deployment can handle volume and time quotas. For more details, see RADIUS chapter.

Subscriber groups

A subscriber group is a set of subscribers. A subscriber can be associated to a group by the IP address or the subscriber ID. By IP address, a subscriber can belong to up to three different subscriber groups, plus a default group called **all-subscribers** that all IP addresses belong to. Subscribers can belong to up to four different by their subscriber ID.

The QoE can show network metrics per subscriber group. Depending on how the groups are defined, those 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 subscriber groups
- Viewing subscriber group metrics
- Subscriber group rate limiting
- Skipping subscriber group rate limiting
- Managing subscriber groups from REST API
- Managing subscriber groups from RADIUS

Creating subscriber groups

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

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

A subscriber group is given a name, 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...** adds the IP addresses contained in the file to those already existing in the group.
- The menu option **Replace Subscriber IDs with File...**replaces current set of subscriber IDs by those contained in the loaded file.

Figure 109: *The Edit subscriber group page*

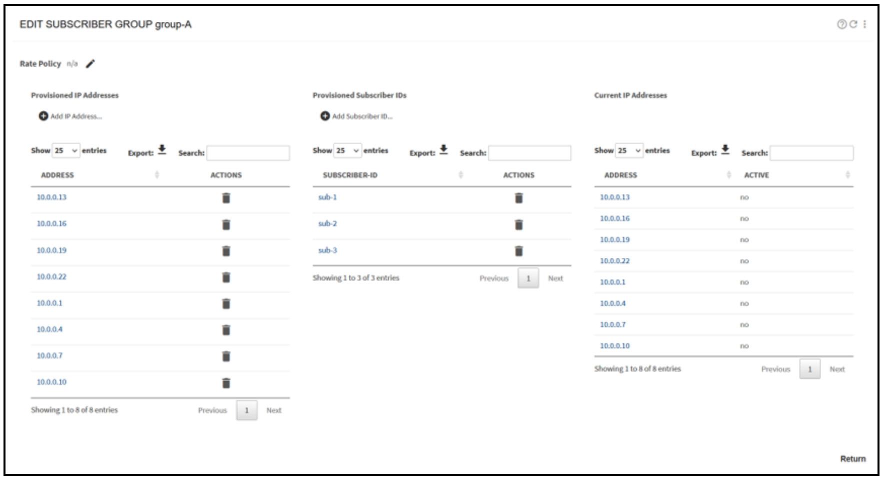

EDIT SUBSCRIBER GROUP group-A

Rate Policy n/a

Provisioned IP Addresses

Add IP Address...

Show 25 entries Export: Search:

ADDRESS	ACTIONS
10.0.0.13	🗑
10.0.0.16	🗑
10.0.0.19	🗑
10.0.0.22	🗑
10.0.0.1	🗑
10.0.0.4	🗑
10.0.0.7	🗑
10.0.0.10	🗑

Showing 1 to 8 of 8 entries Previous 1 Next

Provisioned Subscriber IDs

Add Subscriber ID...

Show 25 entries Export: Search:

SUBSCRIBER-ID	ACTIONS
sub-1	🗑
sub-2	🗑
sub-3	🗑

Showing 1 to 3 of 3 entries Previous 1 Next

Current IP Addresses

Show 25 entries Export: Search:

ADDRESS	ACTIVE
10.0.0.13	no
10.0.0.16	no
10.0.0.19	no
10.0.0.22	no
10.0.0.1	no
10.0.0.4	no
10.0.0.7	no
10.0.0.10	no

Showing 1 to 8 of 8 entries Previous 1 Next

Return

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** or navigate to **Statistics** > **Subscribers** > **Subscriber Groups**. Figure 110 shows the Subscribers Groups page.

Figure 110: *The Subscriber Groups page*

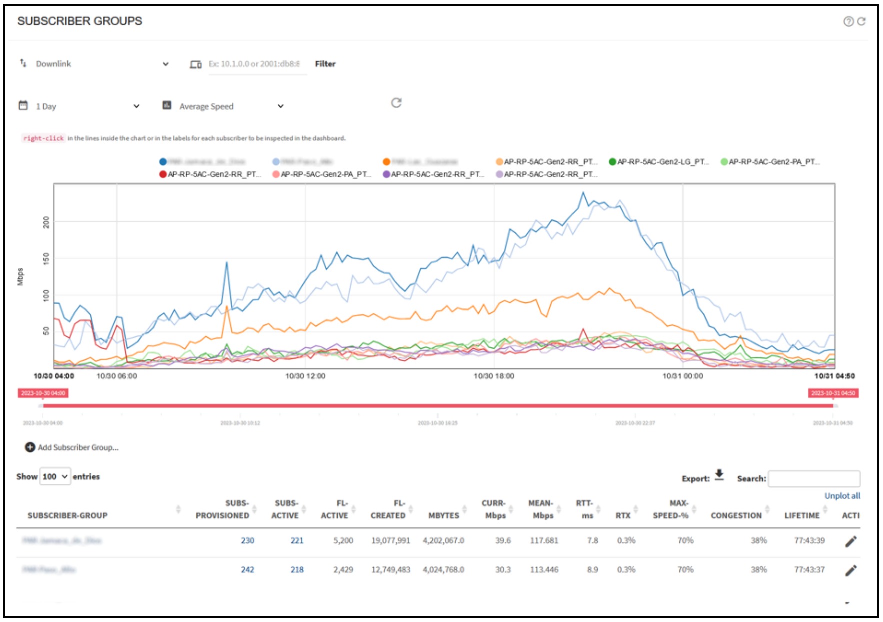

SUBSCRIBER GROUPS

Downlink ⌄ Ex: 10.1.0.0 or 2001:db8:8 Filter

1 Day ⌄ Average Speed ⌄

`right-click` in the lines inside the chart or in the labels for each subscriber to be inspected in the dashboard.

Add Subscriber Group...

Show 100 entries Export: Search:

Unplot all

SUBSCRIBER-GROUP	SUBS-PROVISIONED	SUBS-ACTIVE	FL-ACTIVE	FL-CREATED	MBYTES	CURR-Mbps	MEAN-Mbps	RTT-ms	RTX	MAX-SPEED-%	CONGESTION	LIFETIME	ACTI
[redacted]	230	221	5,200	19,077,991	4,202,067.0	39.6	117.681	7.8	0.3%	70%	38%	77:43:39	✎
[redacted]	242	218	2,429	12,749,483	4,024,768.0	30.3	113.446	8.9	0.3%	70%	38%	77:43:37	✎

By default, the chart shows the first 10 subscriber groups, and a **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. The chart will be refreshed automatically after a few seconds. Unselecting the plot tick will remove that group from the chart.

Metrics can be for the downlink or uplink directions, and for a time period of up to three months.

The following are the available metrics:

- 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 the Subscriber dashboard section).

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

Figure 111: *The Subscriber groups page*

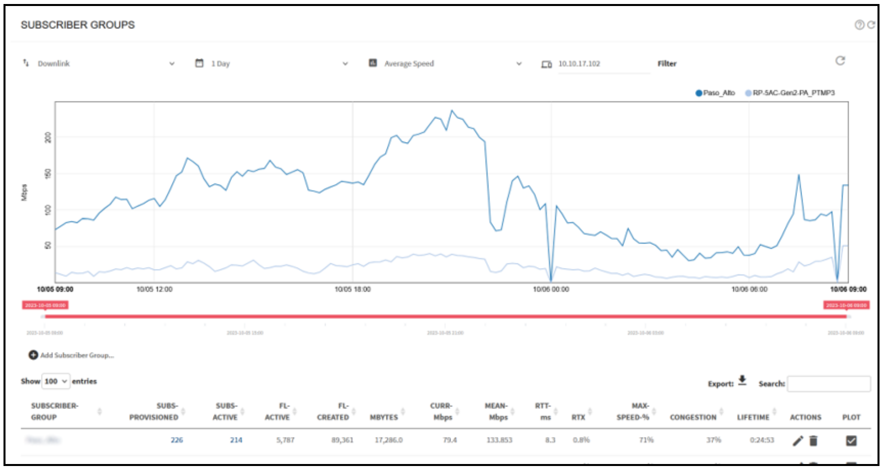

SUBSCRIBER-GROUP	SUBS-PROVISIONED	SUBS-ACTIVE	FL-ACTIVE	FL-CREATED	MBYTES	CURR-Mbps	MEAN-Mbps	RTT-ms	RTX	MAX-SPEED-%	CONGESTION	LIFETIME	ACTIONS	PLOT
[redacted]	226	214	5,787	89,361	17,286.0	79.4	133.853	8.3	0.8%	71%	37%	0:24:53		☑

When filtering by IP address, a **PERMANENT** column tells if the subscriber group has been assigned to this IP address using CLI or REST (permanent) or by RADIUS or Billing (non-permanent). Assignments to a subscriber group by REST or CLI, therefore, takes precedence over RADIUS and Billing, so you can use them to adjust subscriber group membership.

To view the subscriber group dashboard, click subscriber group name with the metrics in separate charts. Figure 112 shows the subscriber group dashboard.

Figure 112: *The Subscriber group dashboard*

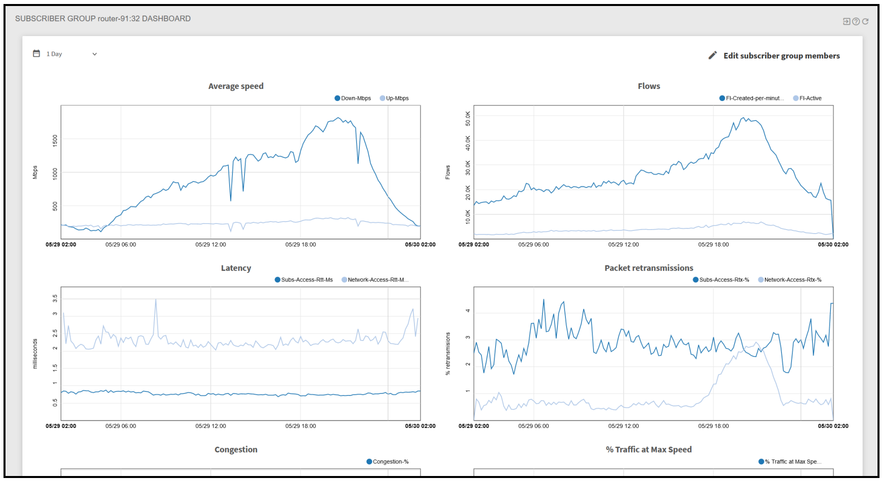

It is also possible to navigate to this page by right-clicking on the subscriber group name in the chart legend or the chart itself.

Subscriber group rate limiting

It is possible to set a rate limit to the overall subscriber group traffic, to be shared by all subscribers of the group. The rate limiting implements max-min fairness, that is, the goal is to choose per-subscriber limits that maximize the lowest ones, so the distribution of the group limits is as equal as possible among the group members.

To apply rate limits per subscriber group, check that the **Subscriber-group rate limiting** feature is enabled in **Configuration** > **Optimization Settings** (it is enabled by default).

Figure 113: *Subscriber group rate limiting*

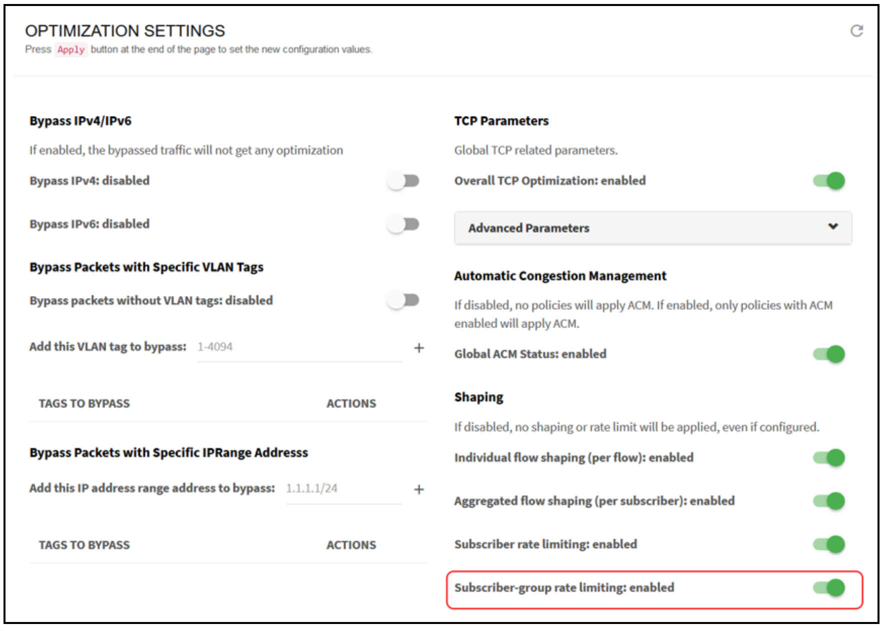

OPTIMIZATION SETTINGS

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

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

Bypass Packets with Specific IPRange Addresss

Add this IP address range address to bypass: 1.1.1.1/24 +

TAGS TO BYPASS	ACTIONS

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.

Individual flow shaping (per flow): enabled

Aggregated flow shaping (per subscriber): enabled

Subscriber rate limiting: enabled

Subscriber-group rate limiting: enabled

The next step is to create a rate policy in **Configuration** > **Subscriber Rates** with the uplink and downlink rate limits. In the example shown in Figure 114, a limit of 500 Mbps is set in the downlink direction and no limit is set in uplink (n/a).

Figure 114: *Limit for downlink and uplink traffic*

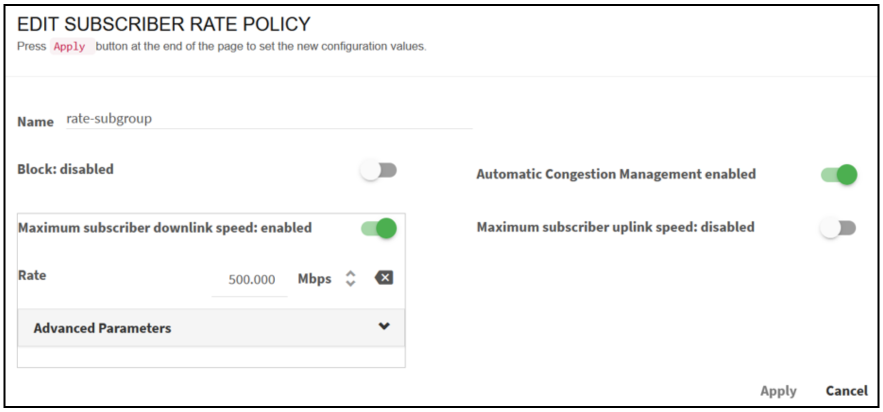

EDIT SUBSCRIBER RATE POLICY

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

Name: rate-subgroup

Block: disabled

Automatic Congestion Management enabled

Maximum subscriber downlink speed: enabled

Maximum subscriber uplink speed: disabled

Rate: 500.000 Mbps

Note

For subscriber groups, only the rate limits are applied (other policy parameters such as burst attributes are ignored).

Finally, navigate to the subscriber group edit page and select the policy rate.

Figure 115: *Edit Subscriber Group*

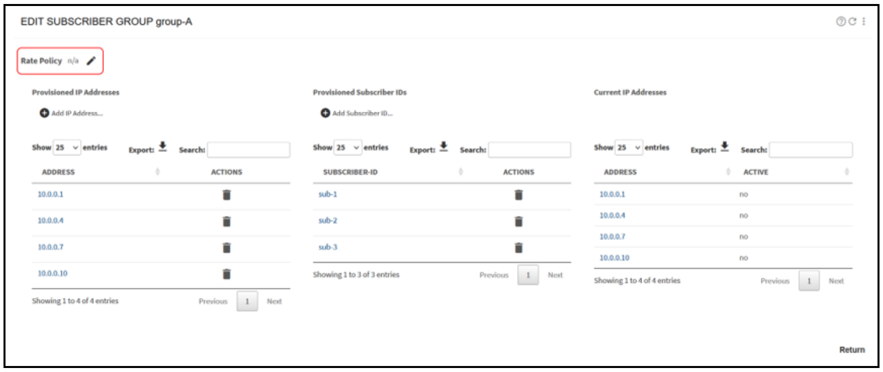

EDIT SUBSCRIBER GROUP group-A

Rate Policy: n/a

Provisioned IP Addresses

Show 25 entries

ADDRESS	ACTIONS
10.0.0.1	
10.0.0.4	
10.0.0.7	
10.0.0.10	

Showing 1 to 4 of 4 entries

Provisioned Subscriber IDs

Show 25 entries

SUBSCRIBER-ID	ACTIONS
sub-1	
sub-2	
sub-3	

Showing 1 to 3 of 3 entries

Current IP Addresses

Show 25 entries

ADDRESS	ACTIVE
10.0.0.1	no
10.0.0.4	no
10.0.0.7	no
10.0.0.10	no

Showing 1 to 4 of 4 entries

A dialogue will list the available rate policies as shown in Figure 116. Select the one to be associated to the subscriber group.

Figure 116: *Rate Policy selection page*

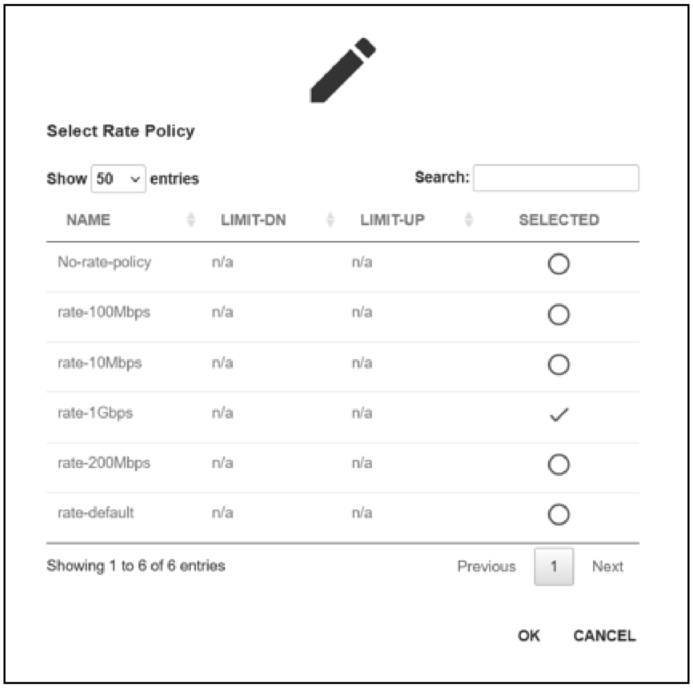

Select Rate Policy

Show 50 entries Search:

NAME	LIMIT-DN	LIMIT-UP	SELECTED
No-rate-policy	n/a	n/a	○
rate-100Mbps	n/a	n/a	○
rate-10Mbps	n/a	n/a	○
rate-1Gbps	n/a	n/a	✓
rate-200Mbps	n/a	n/a	○
rate-default	n/a	n/a	○

Showing 1 to 6 of 6 entries

Skipping subscriber group rate limiting

It is possible to liberate some traffic from subscriber group rate limits. This is done by defining a flow policy which has enabled the switch **Skip subscriber-group rate limitations**, as in the example shown in Figure 117.

Figure 117: *The Edit Subscriber Flow Policy page*

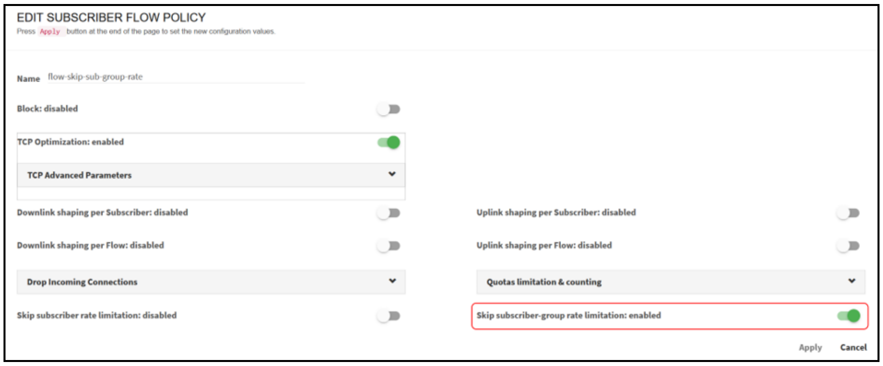

EDIT SUBSCRIBER FLOW POLICY

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

Name flow-skip-sub-group-rate

Block: disabled

TCP Optimization: enabled

Downlink shaping per Subscriber: disabled

Uplink shaping per Subscriber: disabled

Downlink shaping per Flow: disabled

Uplink shaping per Flow: disabled

Skip subscriber rate limitation: disabled

Skip subscriber-group rate limitation: enabled

You can decide what traffic liberate from subscriber group limits by defining rules ate evaluate to this flow policy. In the example shown in Figure 118, a VLAN is exempt from subscriber group limit.

Figure 118: *The Subscriber flows configuration page*

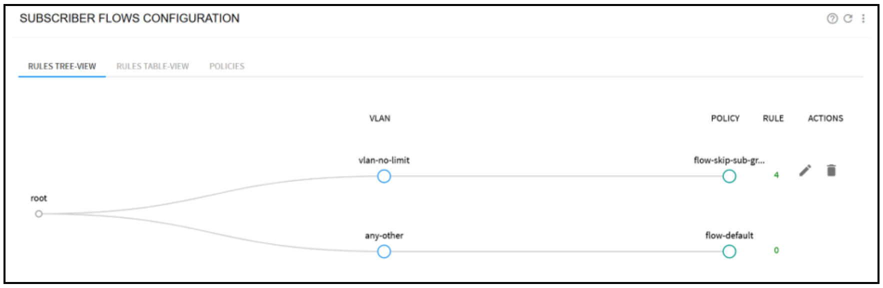

Managing subscriber groups from REST API

The QoE REST API can be used to manage subscriber groups.

To add an existing IP 10.0.0.35 to a subscriber group "city-north", use the following command:

```
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 will be automatically created if it did not exist.

The list of groups should reflect always the full list of membership. For example, to add an existing IP 10.0.0.35 to a subscriber group “vip”:

```
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 because otherwise the IP would be removed from that group. If this is in fact what is desired, 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, they will be 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" ]

}
```

And to get the full list of subscriber groups:

```
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 members to a group by the subscriber 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”]}’
```

You can also set the subscriber group rate limits as follows (1 Gbps downlink and 500 Mbps uplink in the example):

```
curl -k -u myuser:mypassword -X PUT
https://192.168.0.121:3443/api/v1/subscriberGroups/city-north -H "Content-
Type: application/json" --data '{“shapingDownlink”: {“rate”: 1000000},
{“shapingUplink”: {“rate”: 500000}} }’
```

For more REST API commands, see REST API reference.

Managing subscriber groups from RADIUS

QoE in RADIUS proxy deployment can handle subscriber groups. For more details, see RADIUS chapter.

RADIUS interface

This chapter contains the following sections:

Introduction

The QoE can be configured to receive RADIUS messages to support the following functionality:

- Assignment of rate policies to subscribers.
- Dynamic definitions of those rate policies.
- Management of time and/or volume quotas.
- Assignment of RADIUS attribute to QoE subscriber ID.
- Assignment of a RADIUS attribute to QoE subscriber group.

The enforcement of quotas and rate policies is applied per subscriber IP address, so the QoE must have visibility of the subscriber IP addresses, that is, there cannot be a NAT between the subscribers and the QoE because the rate limits are applied per subscriber IP address. It is also important that the wires are connected in the right way (access ports connected on the side of the subscribers).

Note

Currently, only IPv4 address subscribers are supported through RADIUS.

There are two ways for QoE to receive RADIUS messages:

- As a RADIUS accounting server.
- As a RADIUS proxy between RADIUS clients and the RADIUS server (both for Authentication and Accounting).

Figure 119 summarizes a QoE deployment as a RADIUS accounting server.

Figure 119: *QoE deployment as a RADIUS accounting server*

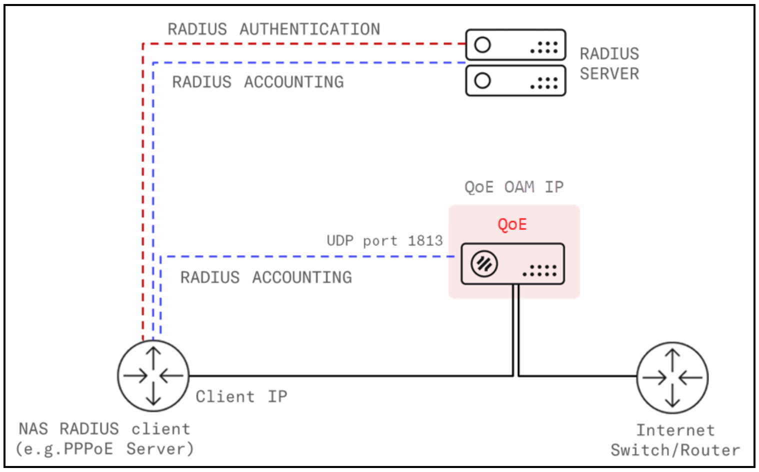

The diagram details the IP and ports used and the flow of the RADIUS signaling.

Similarly, Figure 120 shows a QoE deployment as a RADIUS proxy:

Figure 120: *QoE deployment as a RADIUS proxy*

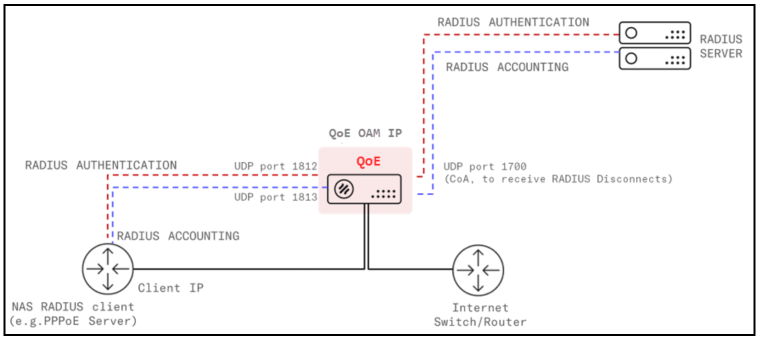

The QoE uses its management IP address (same address as GUI) to receive and send RADIUS messages. It uses the following ports:

- UDP port 1812 for RADIUS Authentication/Authorization
- UDP port 1813 for RADIUS accounting
- UDP port to receive Change of Authorization Disconnect message (1700 by default, though others like 3799 port can be configured).

RADIUS accounting server and RADIUS proxy deployments have pros and cons. The main advantage of RADIUS accounting deployment is that the QoE is not in the RADIUS authentication path.

Note

RADIUS accounting server deployment is the preferred option.

However, Quota management requires a RADIUS proxy deployment. Also, RADIUS proxy deployment is a valid option when RADIUS accounting server is not possible (for example, with Mikrotik routers in DHCP mode).

Accounting server deployment

To integrate the QoE to the RADIUS in Accounting server mode, it is necessary to make configuration changes in the QoE and NAS RADIUS clients. Figure 121 shows a summary.

Figure 121: *Integrating QoE to the RADIUS in Accounting server mode*

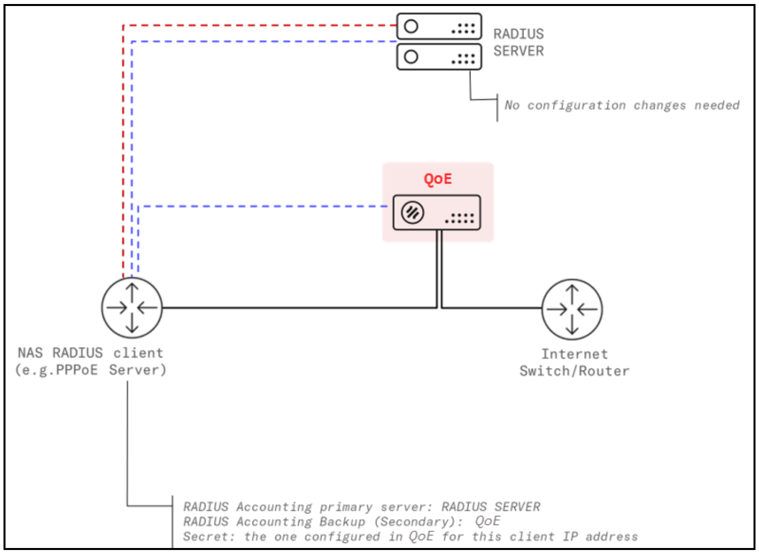
RADIUS
SERVER
No configuration changes needed
QoE
NAS RADIUS client
(e.g.PPPoE Server)
Internet
Switch/Router
RADIUS Accounting primary server: RADIUS SERVER
RADIUS Accounting Backup (Secondary): QoE
Secret: the one configured in QoE for this client IP address

Configuration of the QoE

Follow the steps to configure QoE:

1. Navigate to **Configuration** > **RADIUS/REST/Billing** > **RADIUS**
2. Set the **RADIUS service** switch to ON.
3. Select **RADIUS proxy** option in **RADIUS proxy/Accounting server** field.
4. Configure the RADIUS server IP address and secret in fields **Server IP address** and **Server secret** respectively.
5. Add the IP addresses of each NAS RADIUS client, along with its secret. To do so, in the QoE GUI, navigate to **Configuration** > **RADIUS/REST/Billing** > **RADIUS** and click on **Add Client...**

Figure 122 shows an example.

Figure 122: *Configuration of QoE*

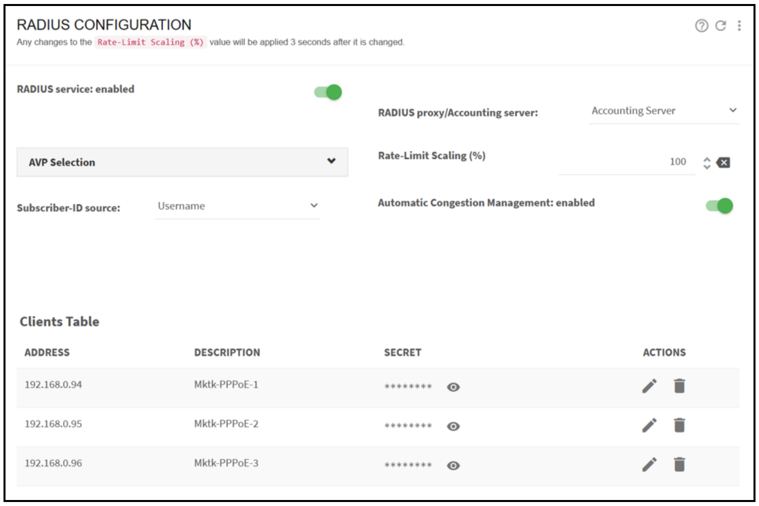

RADIUS CONFIGURATION

Any changes to the `Rate-Limit Scaling (%)` value will be applied 3 seconds after it is changed.

RADIUS service: enabled

RADIUS proxy/Accounting server: Accounting Server

AVP Selection

Rate-Limit Scaling (%) 100

Subscriber-ID source: Username

Automatic Congestion Management: enabled

Clients Table

ADDRESS	DESCRIPTION	SECRET	ACTIONS
192.168.0.94	Mktk-PPPoE-1	********	
192.168.0.95	Mktk-PPPoE-2	********	
192.168.0.96	Mktk-PPPoE-3	********	

An optional description for each NAS RADIUS client can be added (the description cannot contain spaces).

Configuration of the NAS RADIUS clients

The goal is to configure the NAS RADIUS client to send copies of RADIUS accounting messages to the QoE, as if the QoE server is a RADIUS server only for accounting. The steps for each NAS RADIUS client are:

- Configure a RADIUS backup server with QoE IP address and the secret configured previously in QoE for this NAS.
- Make sure the port used is UDP 1813.
- Set Accounting Interim Updates, with an interim period no longer than 5 minutes.

The following instructions are related to a Mikrotik PPPoE server, but similar steps can be followed for other vendors:

- First, the existing RADIUS configuration should not be changed.
- A new RADIUS server will be configured with the QoE as an Accounting Backup server. To create a new RADIUS server, navigate to the RADIUS section and click on **Add New**. The following screen is displayed:

Figure 123: *Creation of new RADIUS server*

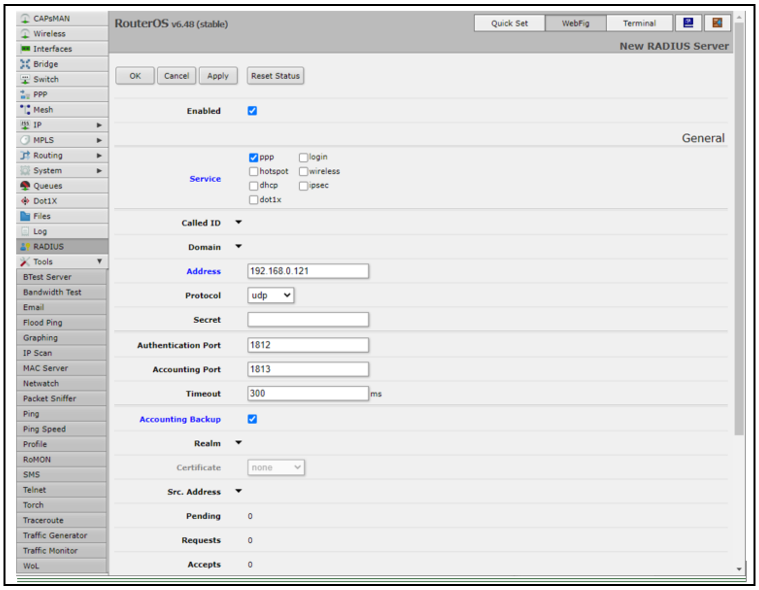

RouterOS v6.48 (stable)

New RADIUS Server

Enabled ☑

General

Service ☑ ppp ☐ login ☐ hotspot ☐ wireless ☐ dhcp ☐ ipsec ☐ dot1x

Called ID

Domain

Address 192.168.0.121

Protocol udp

Secret

Authentication Port 1812

Accounting Port 1813

Timeout 300 ms

Accounting Backup ☑

Realm

Certificate none

Src. Address

Pending 0

Requests 0

Accepts 0

As it can be observed:

- The service of the router/switch should be enabled (usually ppp).
- In **Address** field, the IP address will be the QoE management IP address (you can see which one it is in QoE GUI **Configuration** > **Interfaces** > **Management**).
- The **Accounting Backup** field must be selected (otherwise, the QoE would receive RADIUS Authentication and Authorization messages, which it does not support).
- The accounting port is left in its default value (1813).
- Optionally, a secret can be specified (if used, it must match the one configured in the QoE RADIUS configuration).
- Optionally, a comment can add a description of the RADIUS server (for example, QoE RADIUS).
- After creating the RADIUS server, the list should be as follows:

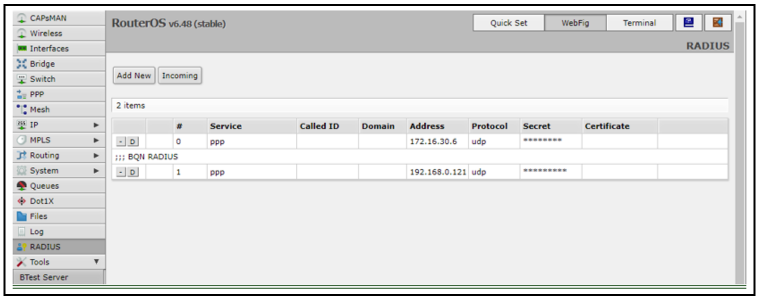

RouterOS v6.48 (stable)

RADIUS

2 items

	#	Service	Called ID	Domain	Address	Protocol	Secret	Certificate
- D	0	ppp			172.16.30.6	udp	********	
;;; BQN RADIUS								
- D	1	ppp			192.168.0.121	udp	*********	

- Make sure the service has RADIUS Interim Updates enabled and with a reasonable period (1-5 minutes). For example, for PPP, in option **PPP** > **Secrets** > **PPP Authentication & Accounting**.

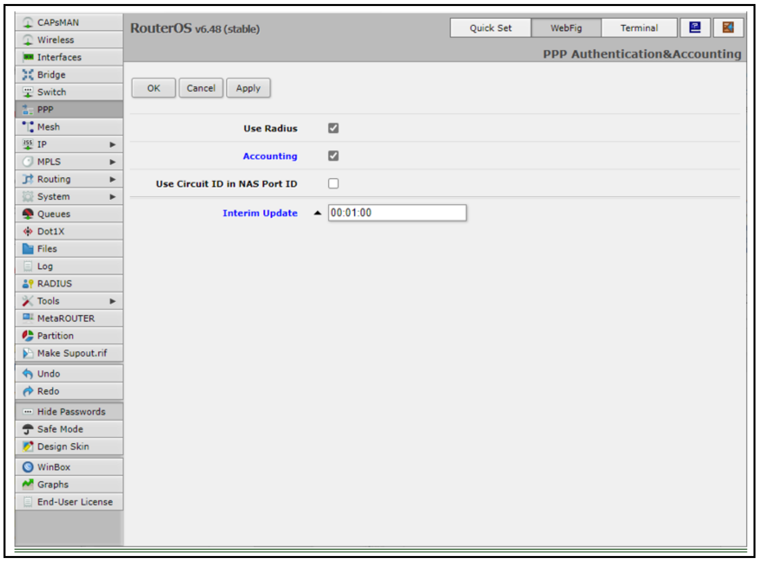

PPP Authentication&Accounting

Use Radius ☑

Accounting ☑

Use Circuit ID in NAS Port ID □

Interim Update 00:01:00

Note

The process must be repeated in all the nodes whose RADIUS is to be sent to the QoE.

Proxy Deployment

To integrate the QoE to the RADIUS in proxy mode, it is necessary to make configuration changes in the QoE, the NAS RADIUS clients and the RADIUS Server. Figure 124 shows a summary.

Figure 124: *Proxy deployment*

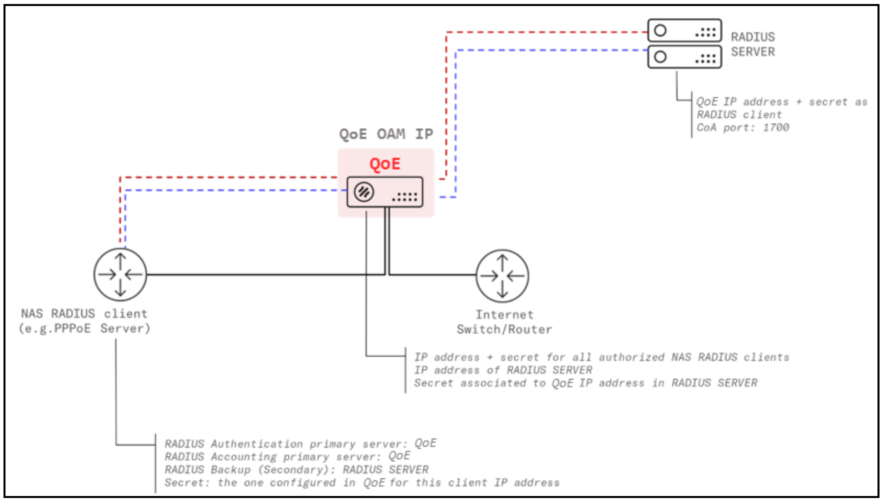
RADIUS SERVER
QoE IP address + secret as RADIUS client
CoA port: 1700
QoE OAM IP
QoE
NAS RADIUS client (e.g.PPPoE Server)
Internet Switch/Router
IP address + secret for all authorized NAS RADIUS clients
IP address of RADIUS SERVER
Secret associated to QoE IP address in RADIUS SERVER
RADIUS Authentication primary server: QoE
RADIUS Accounting primary server: QoE
RADIUS Backup (Secondary): RADIUS SERVER
Secret: the one configured in QoE for this client IP address

Configuration of the QoE

Follow the steps to configure QoE:

1. Navigate to **Configuration** > **RADIUS/REST/Billing** > **RADIUS**
2. Set the **RADIUS service** switch to ON.
3. Select **RADIUS proxy** option in **RADIUS proxy/Accounting server** field.

4. Configure the RADIUS server IP address and secret in the **Server IP address** and **Server secret** fields respectively.

5. Configure the Disconnect port used by the RADIUS server (1700 by default).

6. Add the IP addresses of each NAS RADIUS client, along with its secret. To do so, in the QoE GUI, navigate to **Configuration** > **RADIUS/REST/Billing** > **RADIUS** and click on **Add Client...**

Figure 125 shows an example.

Figure 125: *RADIUS proxy*

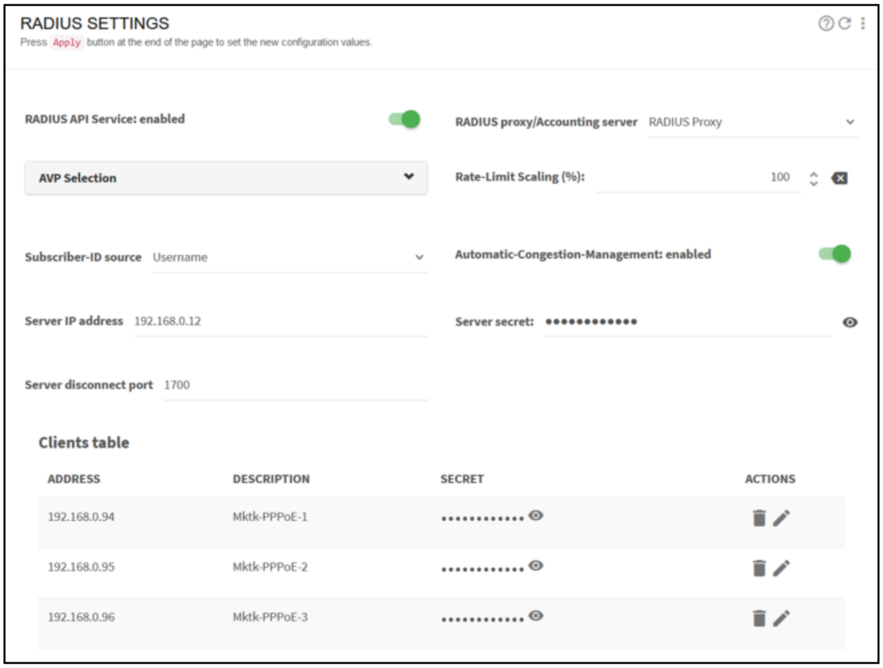

RADIUS SETTINGS

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

RADIUS API Service: enabled

RADIUS proxy/Accounting server RADIUS Proxy

AVP Selection

Rate-Limit Scaling (%): 100

Subscriber-ID source Username

Automatic-Congestion-Management: enabled

Server IP address 192.168.0.12

Server secret: ••••••••••••

Server disconnect port 1700

Clients table

ADDRESS	DESCRIPTION	SECRET	ACTIONS
192.168.0.94	Mktk-PPPoE-1	••••••••••••	
192.168.0.95	Mktk-PPPoE-2	••••••••••••	
192.168.0.96	Mktk-PPPoE-3	••••••••••••	

Configuration of the NAS RADIUS clients

The steps for each NAS RADIUS client are:

- Configure the QoE IP address as primary RADIUS server, using the secret configured previously in the QoE.

- Make sure the ports used are UDP 1812 for authentication, UDP 1813 for accounting and UDP 1700 to receive Disconnects.

- Set the RADIUS server as backup RADIUS, if possible.

Configuration of the RADIUS server

Follow this step:

- Configure the QoE IP address as a valid NAS RADIUS client, with a secret equal to the one previously configured in the QoE.

RADIUS AVP handling

The QoE processes some RADIUS AVPs (Attribute-Value Pairs) to implement QoE own functionality, such as rate policy enforcement, quotas or information for subscriber identification.

By default, the QoE reacts to all the supported AVPs (of which more later on in this chapter), but it is configurable to ignore some of them. When in proxy deployment mode, the QoE will relay ignored AVPs but still will not perform any other action.

One possible reason to ignore AVPs is when others with lower priority are preferred. The policy rate section provides a concrete example.

When in proxy QoE deployment mode, the QoE can also remove some AVPs received in an Accept-Accept message from the RADIUS server before relaying it to the NAS RADIUS client. One possible reason to remove AVPs is to prevent the NAS client to try to act on them, interfering with QoE actions. The quota section provides a concrete example.

The configuration of AVP handling is in **Configuration** > **RADIUS/REST/Billing** > **RADIUS**, under **AVP Selection** field.

Rate policy control

The RADIUS Accounting Start and Interim messages link a subscriber IP address with a subscriber rate policy. The subscriber IP address is received in the **Framed-IP-Address** field.

There are two ways to specify the subscriber rate policy:

- Specifying the subscriber rate policy parameters (like the rate limit), where the RADIUS attribute provides the policy definition and the QoE creates a policy based on that information.
- Specifying the subscriber rate policy name, where the RADIUS attribute contains the name of the policy to choose from the policies that are part of the QoE configuration.

Both accounting server and proxy deployments can be used for rate policy control.

Supported RADIUS AVPs for Rate Policies

The following RADIUS AVPs are supported for rate policy management and are listed in order of priority (parameters evaluated first will take precedence):

Priority	Name	Vendor	ID	Description	Example
1	Mikrotik Rate-Limit	14988	8	Contains the policy rate limits, including optional burst parameters: rx-rate[/tx-rate] [rx-burst-rate [/tx-burst-rate] [rx-burst-threshold[/tx-burst-threshold] [rx-burst-time[/tx-burst-time] [priority] [rx-rate-min[/tx-rate-min]]]]	15M/20M 30M/40M 15M/20M 5/5
2	Ascend Data-Rate	529	197	The first instance defines the downlink rate limit in Mbps. The second instance defines the	50

Priority	Name	Vendor	ID	Description	Example
				uplink rate limit in Mbps. If there is only one instance, it specifies the downlink limit, unless Ascend-Xmit-Rate is present, in which it specifies the uplink limit.	
3	Ascend Xmit-Rate	529	255	The downlink rate limit in Mbps.	100
4	Mikrotik Address-List	14988	19	Defines the policy parameters if meets the following format: [rate-limit][K or M] Downlink-[rate-limit][K or M] Uplink	50M Downlink-50M Uplink
5	CISCO sub-qos-policy	9	1	Defines the rate limits with the following format: sub-qos-policy-out=[Prefix]_[rate-limit]Mbps_Downlink sub-qos-policy-in=[Prefix]_[rate-limit]Mbps_Uplink	sub-qos-policy-out=PPPoE_100Mbps_Downlink" "sub-qos-policy-in=PPPoE_50Mbps_Uplink
6	CISCO rate limits	9	37, 38	Defines the rate limits with the following format: ID=37 (downlink rate limit): [rate-limit][K or M] Downlink. ID=38 (uplink rate limit): [rate-limit][K or M] Uplink.	ID 37: "10M Downlink" ID 38: "5M Uplink"
7	Mikrotik Address-List	14988	19	If Mikrotik Address-List format does not match the format described in priority 4, its content is interpreted as a policy name configured in the QoE.	Gold_Users
8	CISCO sub-qos-policy	9	1	If Cisco rate limits format does not match with the format described in priority 5, its content are interpreted as a policy name configured in the QoE. Both AVPs are combined in the following way: [qos-in AVP]-[qos-out AVP] and if a AVP is absent, it is replaced by an 'u'.	"Plan_50-Plan_100" "u-Plan_100"
9	CISCO rate limits	9	37, 38	If CISCO rate limits format does not match the format described in priority 6, its content is interpreted as a policy name configured in the QoE.	Basic_Plan

Priority	Name	Vendor	ID	Description	Example
10	Class	NA	25	Generic RADIUS parameter. A dynamic policy is created if Class contains the substring: policy=[rate-limit][K|M|G] Uplink-[rate-limit][K|M|G] Downlink	Dynamic policy: “policy=5M Uplink-10M Downlink, other-param=foo” Resulting policy name: “5M_ Uplink-10M_Downlink”

11	Class	NA	25	If Class content does not match the previous pattern, a reference to a configured policy is created with the substring following the = up to a comma:
				policy=[policy.name],
				Otherwise, the parameter is ignored.
					“policy=mypolicy, other-param=foo” Resulting policy name (configured in QoE):
					“mypolicy” “policy=anotherpolicy” Resulting policy name (configured in QoE):
					“anotherpolicy”

12	Connect-Info	NA	77	The speeds are specified as follows: [downlink-rate][K|M|G][/[uplink-rate][K|M|G]] If uplink-rate is absent, it will take the same value as downlink-rate. If units K,M,G are absent, rate will be in bits per second (bps).	50000K/5000K 5M 5000000/5000000 10M/5M
13	NetElastic QoS Profile Name	54268	31	Defines the policy parameters if meets the following format: [rate-limit][K|M|G] Uplink-[rate-limit] [K|M|G] Downlink	50M Uplink-50M Downlink
14	NetElastic QoS Profile Name	54268	31	If NetElastic QoS Profile Name format does not match the format described in priority 13, its content will be interpreted as a policy name configured in the QoE.	Gold Users
15	Huawei Input-Peak-Rate Output-Peak-Rate	2021	121 & 124	Huawei-Input-Peak-Rate is the uplink rate limit in bits per second. Huawei-Output-Peak-Rate is the downlink rate limit in bits per second.	50000000 100000000 (50 Mbps uplink, 100 Mbps downlink).

When more than one AVP is present, the policy is managed according to the priority order. For example, if both **Mikrotik-Rate-Limit** and **Ascend-Data-Rate** are present, **Mikrotik-Rate-Limi**t will take precedence. Also, if both **Ascend-Data-Rate** and **Mikrotik-Address-List** are in the Radius message, **Mikrotik-Address-**

List will be ignored. In any case, it is possible to use any of those information elements, since the QoE can be configured to ignore the ones with more precedence, as can be seen in the following section.

If a RADIUS message contains none of the supported AVPs, the subscriber rate policy previously assigned to the subscriber IP address in this RADIUS message, if any, will be removed and a new one will be chosen based on QoE configured subscriber rate policy rules.

RADIUS provides the rate policy definition

This requires no specific configuration. Make sure that in **Configuration** > **RADIUS/REST/Billing** > **RADIUS**, the RADIUS parameters which will specify the policy are not ignored (they are not ignored by default). Once the QoE starts receiving the RADIUS messages, it will assign each subscriber (for which a RADIUS message is received) a subscriber rate policy with the rate limits defined in the RADIUS message.

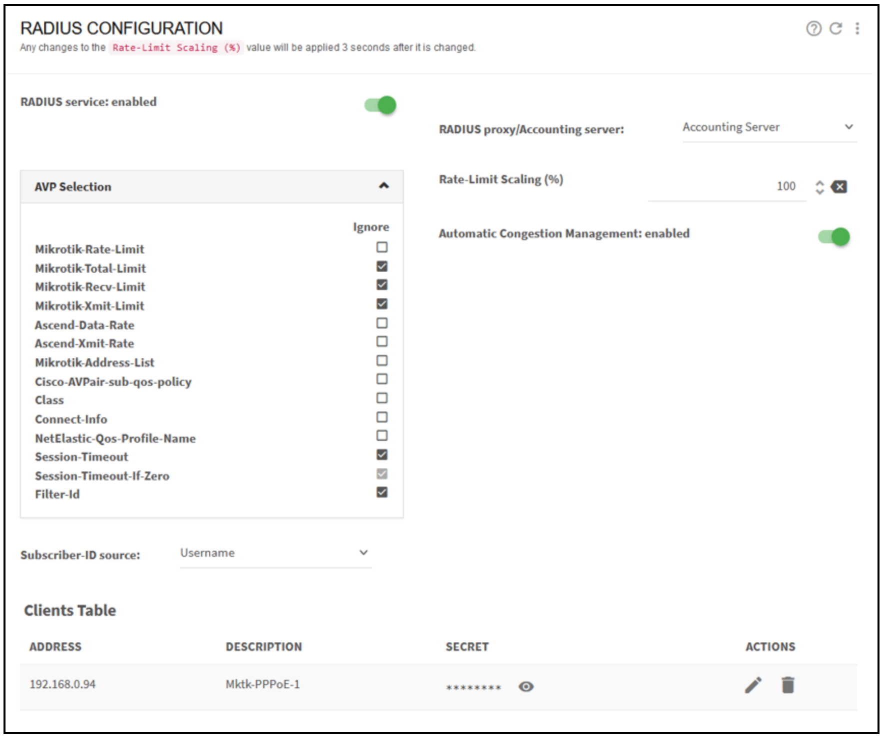

RADIUS CONFIGURATION

Any changes to the `Rate-Limit Scaling (%)` value will be applied 3 seconds after it is changed.

RADIUS service: enabled

RADIUS proxy/Accounting server: Accounting Server

Rate-Limit Scaling (%) 100

Automatic Congestion Management: enabled

AVP Selection

	Ignore
Mikrotik-Rate-Limit	☐
Mikrotik-Total-Limit	☑
Mikrotik-Recv-Limit	☑
Mikrotik-Xmit-Limit	☑
Ascend-Data-Rate	☐
Ascend-Xmit-Rate	☐
Mikrotik-Address-List	☐
Cisco-AVPair-sub-qos-policy	☐
Class	☐
Connect-Info	☐
NetElastic-Qos-Profile-Name	☐
Session-Timeout	☑
Session-Timeout-If-Zero	☑
Filter-Id	☑

Subscriber-ID source: Username

Clients Table

ADDRESS	DESCRIPTION	SECRET	ACTIONS
192.168.0.94	Mktk-PPPoE-1	********	

The name of the policy created dynamically is composed based on the AVP content, with the following format:

```
RA-rx-rate[/tx-rate]-[rx-burst-rate[/tx-burst-rate]-[rx-burst-threshold
[/tx-burst-threshold]-[rx-burst-time[/tx-burst-time]
```

where:

- **RA**: a prefix indicating that it is a policy created from RADIUS.
- **rx-rate**: rate limit in uplink.
- **tx-rate**: rate limit in downlink.
- **rx-burst-rate**: burst rate in uplink.
- **tx-burst-rate**: burst rate in downlink.

- **rx-burst-threshold**: burst threshold in uplink (average speed not to be exceeded before granting a new burst).
- **tx-burst-threshold**: burst threshold in downlink (average speed not to be exceeded before granting a new burst).
- **rx-burst-time**: uplink burst duration, in seconds.
- **tx-burst-time**: downlink burst duration, in seconds.

Rates and thresholds include their units (K for Kbps, M for Mbps, and G for Gbps).

Once the QoE has the rate policies, it will enforce them on a subscriber IP address basis. Regarding the enforcement in the NAS (for example, the PPPoE server), there are three options:

- Remove the rate limits in the NAS (for example, remove Mikrotik queues).
- Make the QoE apply the policy limits reduced by a configurable factor, so QoE will be the enforcement point, not the NAS. The NAS remains a backup at the full rate limit, if QoE enforcement fails.
- Make the QoE remove the parameter received from the RADIUS server so the NAS does not receive it (available only in RADIUS proxy deployments using Mikrotik Rate-Limit parameter).

The following shows an example of policy limit reduction:

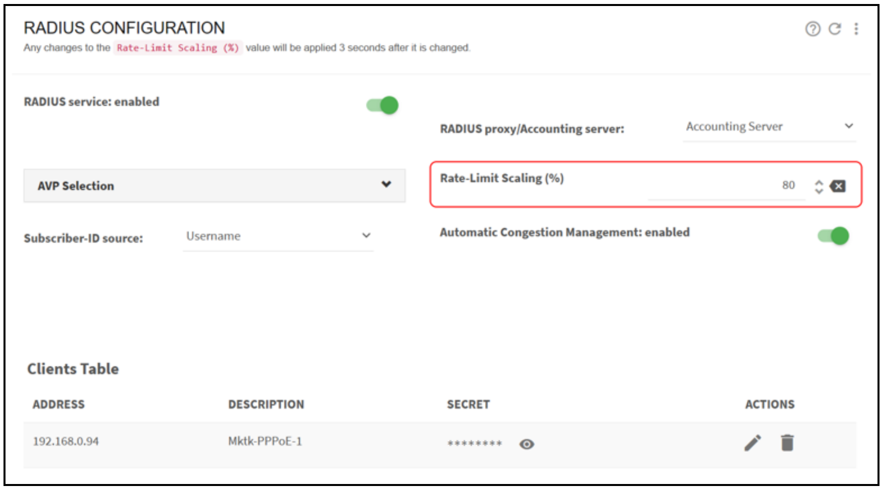

RADIUS CONFIGURATION

Any changes to the `Rate-Limit Scaling (%)` value will be applied 3 seconds after it is changed.

RADIUS service: enabled

RADIUS proxy/Accounting server: Accounting Server

AVP Selection

Rate-Limit Scaling (%) 80

Subscriber-ID source: Username

Automatic Congestion Management: enabled

Clients Table

ADDRESS	DESCRIPTION	SECRET	ACTIONS
192.168.0.94	Mktk-PPPoE-1	********	

For example, with a percentage of 80%, a 125 Mbps limit in RADIUS will be converted to a 100Mbps limit in the QoE (125*0.8).

The percentage is applied to all parameters of the subscriber rate policy (rate limit, burst speed, and burst threshold).

This parameter can also be set to a higher value (for example, 200%) to enforce higher rate limits than the current NAS, so that you will get rate policies assigned to subscribers, but the rates will still be controlled by the NAS. This can be useful during initial testing of the RADIUS interface. Also, it is appropriate if you just want the QoE to get the rate policy without enforcing it (for example to use it to select an appropriate flow policy).

To prevent Mikrotik-Rate-Limit from reaching the NAS in proxy deployments:

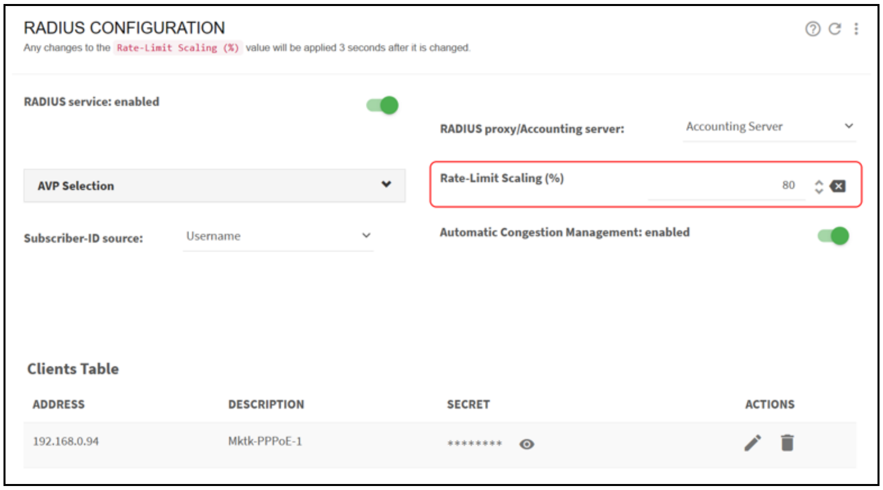

RADIUS CONFIGURATION

Any changes to the `Rate-Limit Scaling (%)` value will be applied 3 seconds after it is changed.

RADIUS service: enabled

RADIUS proxy/Accounting server: Accounting Server

AVP Selection

Rate-Limit Scaling (%) 80

Subscriber-ID source: Username

Automatic Congestion Management: enabled

Clients Table

ADDRESS	DESCRIPTION	SECRET	ACTIONS
192.168.0.94	Mktk-PPPoE-1	********	

Note

QoE is configured in Proxy mode.

RADIUS provides the rate policy name

When a supported RADIUS parameter does not follow the format that allows the QoE to extract the policy definition, the parameter value will be interpreted as the name of one of the subscriber rate policies.

In our example, **Mikrotik-Address-List** is used, but it is similar with other supported RADIUS parameters.

Make sure that the parameter is not ignored and that parameters with higher priority are:

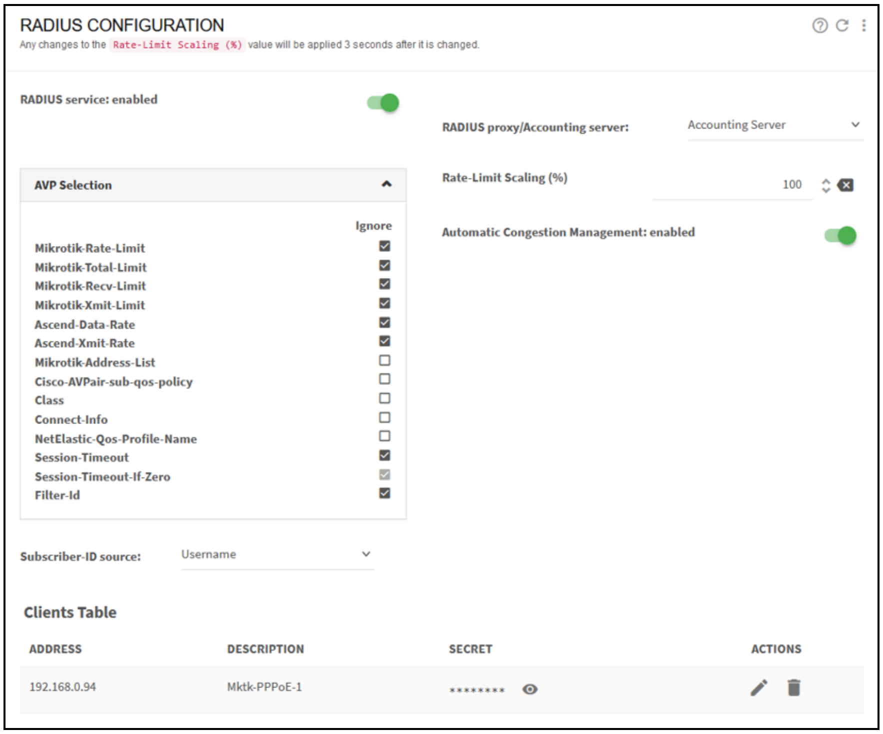

RADIUS CONFIGURATION

Any changes to the `Rate-Limit Scaling (%)` value will be applied 3 seconds after it is changed.

RADIUS service: enabled

RADIUS proxy/Accounting server: Accounting Server

Rate-Limit Scaling (%) 100

Automatic Congestion Management: enabled

AVP Selection

	Ignore
Mikrotik-Rate-Limit	☑
Mikrotik-Total-Limit	☑
Mikrotik-Recv-Limit	☑
Mikrotik-Xmit-Limit	☑
Ascend-Data-Rate	☑
Ascend-Xmit-Rate	☑
Mikrotik-Address-List	☐
Cisco-AVPair-sub-qos-policy	☐
Class	☐
Connect-Info	☐
NetElastic-Qos-Profile-Name	☐
Session-Timeout	☑
Session-Timeout-If-Zero	☑
Filter-Id	☑

Subscriber-ID source: Username

Clients Table

ADDRESS	DESCRIPTION	SECRET	ACTIONS
192.168.0.94	Mktk-PPPoE-1	********	

The name of the policy configured in the QoE is based on the AVP content, with spaces replaced by underscores ("_"). Those are the names that should be used when configuring the policies in the QoE. With **Mikrotik-Address-Lis**t, the AVP content "GOLD PLAN" becomes "GOLD_PLAN".

The policies specified by RADIUS for each subscriber may already be configured in the QoE, in which case they will just be assigned. However, certain policy names specified in RADIUS may not exist in the QoE yet.

In **Status** > **Subscribers** > **Subscriber Attributes**, subscribers associated to an undefined policy are marked in red in **RATE-POLICY** column.

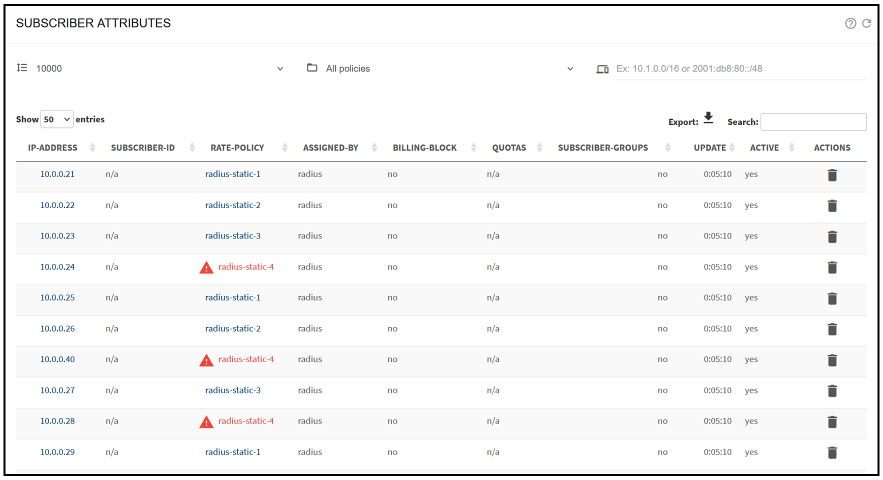

SUBSCRIBER ATTRIBUTES

10000 | All policies | Ex: 10.1.0.0/16 or 2001:db8:80::/48

Show 50 entries — Export: — Search:

IP-ADDRESS	SUBSCRIBER-ID	RATE-POLICY	ASSIGNED-BY	BILLING-BLOCK	QUOTAS	SUBSCRIBER-GROUPS	UPDATE	ACTIVE	ACTIONS
10.0.0.21	n/a	radius-static-1	radius	no	n/a	no	0:05:10	yes	🗑
10.0.0.22	n/a	radius-static-2	radius	no	n/a	no	0:05:10	yes	🗑
10.0.0.23	n/a	radius-static-3	radius	no	n/a	no	0:05:10	yes	🗑
10.0.0.24	n/a	⚠ radius-static-4	radius	no	n/a	no	0:05:10	yes	🗑
10.0.0.25	n/a	radius-static-1	radius	no	n/a	no	0:05:10	yes	🗑
10.0.0.26	n/a	radius-static-2	radius	no	n/a	no	0:05:10	yes	🗑
10.0.0.40	n/a	⚠ radius-static-4	radius	no	n/a	no	0:05:10	yes	🗑
10.0.0.27	n/a	radius-static-3	radius	no	n/a	no	0:05:10	yes	🗑
10.0.0.28	n/a	⚠ radius-static-4	radius	no	n/a	no	0:05:10	yes	🗑
10.0.0.29	n/a	radius-static-1	radius	no	n/a	no	0:05:10	yes	🗑

To configure policy, click on the undefined policy name, with the right name already filled in.

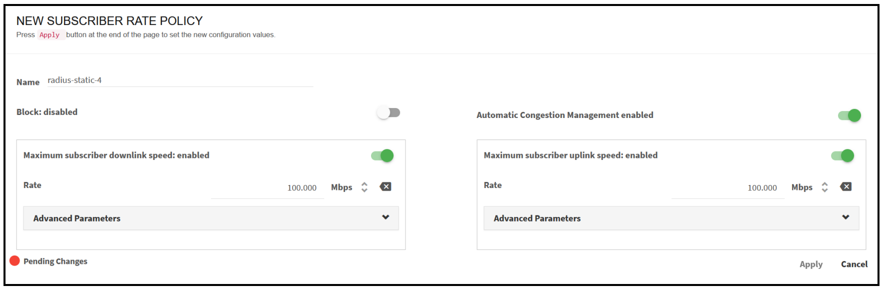

NEW SUBSCRIBER RATE POLICY

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

Name radius-static-4

Block: disabled

Automatic Congestion Management enabled

Maximum subscriber downlink speed: enabled

Rate 100.000 Mbps

Maximum subscriber uplink speed: enabled

Rate 100.000 Mbps

The process is repeated for each subscriber rate policy pending configuration.

As RADIUS assignments are received, the **SUBS_PROVISIONED** counter grows and, as subscriber traffic is received, the **SUBS-ACTIVE** counter will increase.

To view the list of subscribers associated to a subscriber rate policy through RADIUS, navigate to: **Status** > **Subscribers** > **Subscriber Attributes**.

Checking the status of RADIUS rate policies

Status > **Subscribers** > **Subscriber Attributes** shows a table with all subscribers where you can see the assigned rate policy (**RATE-POLICY** column) and if coming from RADIUS (**ASSIGNED-BY** showing **radius**).

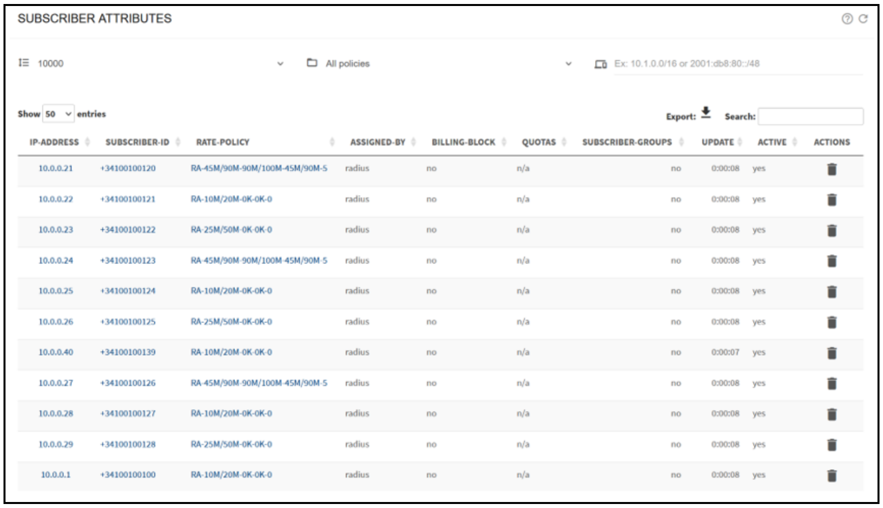

SUBSCRIBER ATTRIBUTES

10000

All policies

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

Show 50 entries

Export:

Search:

IP-ADDRESS	SUBSCRIBER-ID	RATE-POLICY	ASSIGNED-BY	BILLING-BLOCK	QUOTAS	SUBSCRIBER-GROUPS	UPDATE	ACTIVE	ACTIONS
10.0.0.21	+34100100120	RA-45M/90M-90M/100M-45M/90M-5	radius	no	n/a	no	0:00:08	yes	🗑
10.0.0.22	+34100100121	RA-10M/20M-0K-0K-0	radius	no	n/a	no	0:00:08	yes	🗑
10.0.0.23	+34100100122	RA-25M/50M-0K-0K-0	radius	no	n/a	no	0:00:08	yes	🗑
10.0.0.24	+34100100123	RA-45M/90M-90M/100M-45M/90M-5	radius	no	n/a	no	0:00:08	yes	🗑
10.0.0.25	+34100100124	RA-10M/20M-0K-0K-0	radius	no	n/a	no	0:00:08	yes	🗑
10.0.0.26	+34100100125	RA-25M/50M-0K-0K-0	radius	no	n/a	no	0:00:08	yes	🗑
10.0.0.40	+34100100139	RA-10M/20M-0K-0K-0	radius	no	n/a	no	0:00:07	yes	🗑
10.0.0.27	+34100100126	RA-45M/90M-90M/100M-45M/90M-5	radius	no	n/a	no	0:00:08	yes	🗑
10.0.0.28	+34100100127	RA-10M/20M-0K-0K-0	radius	no	n/a	no	0:00:08	yes	🗑
10.0.0.29	+34100100128	RA-25M/50M-0K-0K-0	radius	no	n/a	no	0:00:08	yes	🗑
10.0.0.1	+34100100100	RA-10M/20M-0K-0K-0	radius	no	n/a	no	0:00:08	yes	🗑

Status > Policies > Rate Policies shows a table with all rate policies. Those created from RADIUS will have a **no** in **CONFIGURED** column. **SUBS-PROVISIONED** indicates how many subscribers are associated to this policy and **SUBS-ACTIVE** how many of them are currently active (with traffic).

Time and Volume quota management

Time and Volume quotas are received from the RADIUS server Access Accept message and associated to the subscriber IP address received in the **Framed-IP-Address** field.

RADIUS proxy deployment must be used for quota management.

A table with the RADIUS AVPs supported for Quota management follows (received in AccessAccept from the RADIUS Server). They are ignored by default, so uncheck the ignore tick box to configure the QoE to use them:

Name	Vendor	ID	Description	Example
Mikrotik Total-Limit	14988	17	This is the size of the volume quota in bytes, uplink and downlink traffic combined	25000000 (25MB volume quota)

Mikrotik Total-Limit-Gigawords	14988	18	This AVP gives the number of 4 GBs in the volume quota. It allows to define very big volume quotas in combination with Mikrotik Total-Limit.
			If the volume-quota is > 4 GB:
			Mikrotik-Total-Limit-Gigawords = volume-quota / 4GB (upper 32 bits of volume-quota)
			Mikrotik-Total-Limit-Gigawords = volume-quota % 4GB (lower 32 bits of volume-quota)
				3
				Combined with Mikrotik Total-Limit = 25000000
				Volume quota of 12.025 GBytes.

Session-Timeout	n/a	27	Duration of the time quota in seconds. If zero, the QoE can either ignore it or take it as zero duration and block the session (configurable).	7776000 (90 days)

Name	Vendor	ID	Description	Example
Mikrotik-Recv-Limit	14988	1	This AVP is not supported by QoE to define quota volume limits, but the QoE can be set to remove it in relayed messages, so it does not interfere with Mikrotik Total-Limit.	
Mikrotik-Xmit-Limit	14988	2	This AVP is not supported by QoE to define quota volume limits, but the QoE can be set to remove it in relayed messages, so it does not interfere with Mikrotik Total-Limit.	

The QoE reports usage in its accounting interim requests, using the following RADIUS parameters:

Name	ID	Description	Example
Acct-Input-Octets	42	This is the number of bytes of traffic volume of that Subscriber Session, uplink and downlink traffic combined.	25000000 (25MB volume quota)
Acct-Input-Octets-Gigawords	52	This is the number of 4 GBs of traffic volume. It is combined with Acct-Input-Octets in the same way as Mikrotik-Total-Limit and Mikrotik-Total-Limit-Gigawords. Acct-Input-Octets: lower 32 bits of total volume consumed so far. Acct-Input-Octets: upper (more significant) 32 bits of total volume consumed so far. Mikrotik-Total-Limit-Gigawords = volume-quota / 4GB (upper 32 bits of volume-quota) Mikrotik-Total-Limit-Gigawords = volume-quota % 4GB (lower 32 bits of volume-quota)	5 Combined with Acct-Input-Octets = 500000000 Volume of 20,5 GBytes.

Note

Uplink and downlink traffic are reported combined using the same pair of AVPs.

Because QoE is in proxy mode, its volume counting will drive volume quota enforcement. For example, it is possible to configure QoE rules based on the subscriber ID or group (see Subscriber Identification section) so some applications are excluded from the count.

In most cases, it is important that the quota is enforced by the QoE alone. To achieve this, configure QoE so that quota-related AVPs are removed from both Access-Accept message relayed to the NAS client and from Accounting-Requests coming from the NAS client.

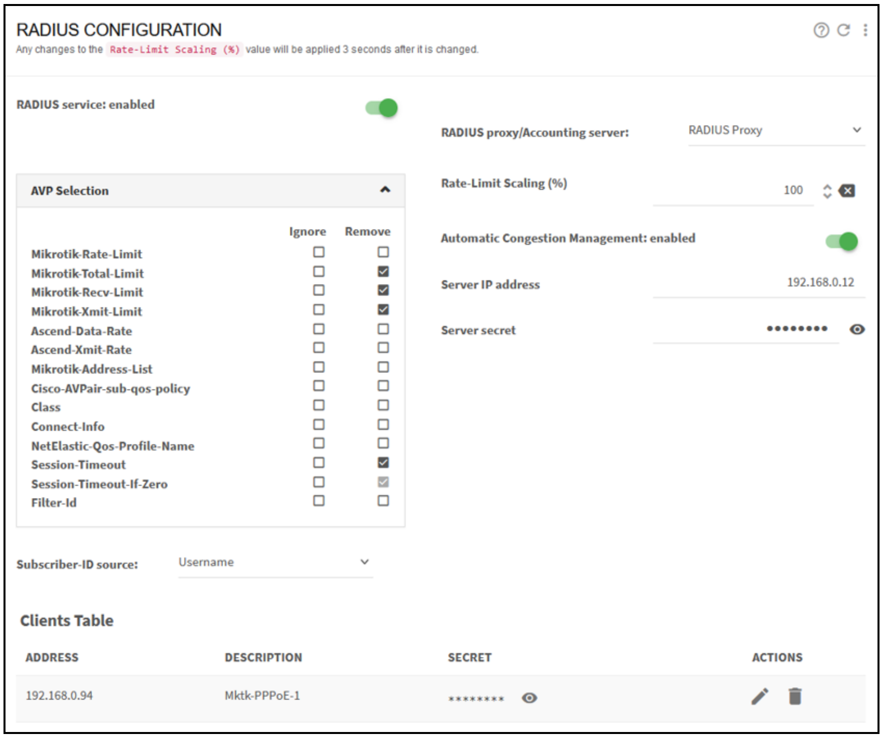

RADIUS CONFIGURATION

Any changes to the `Rate-Limit Scaling (%)` value will be applied 3 seconds after it is changed.

RADIUS service: enabled

RADIUS proxy/Accounting server: RADIUS Proxy

Rate-Limit Scaling (%) 100

Automatic Congestion Management: enabled

Server IP address 192.168.0.12

Server secret ••••••••

AVP Selection

	Ignore	Remove
Mikrotik-Rate-Limit	☐	☐
Mikrotik-Total-Limit	☐	☑
Mikrotik-Recv-Limit	☐	☑
Mikrotik-Xmit-Limit	☐	☑
Ascend-Data-Rate	☐	☐
Ascend-Xmit-Rate	☐	☐
Mikrotik-Address-List	☐	☐
Cisco-AVPair-sub-qos-policy	☐	☐
Class	☐	☐
Connect-Info	☐	☐
NetElastic-Qos-Profile-Name	☐	☐
Session-Timeout	☐	☑
Session-Timeout-If-Zero	☐	☑
Filter-Id	☐	☐

Subscriber-ID source: Username

Clients Table

ADDRESS	DESCRIPTION	SECRET	ACTIONS
192.168.0.94	Mktk-PPPoE-1	********	

Some NAS (for example, Mikrotik) may send a Session-Timeout value of zero, even when no Session-Timeout has been returned by the Access-Accept message. This will be considered as an exhausted time quota and the subscriber to be blocked. To avoid this, you can set the QoE to ignore a zero value in the Session-Timeout AVP selecting **Ignore in Session-Timeout-If-Zero**:

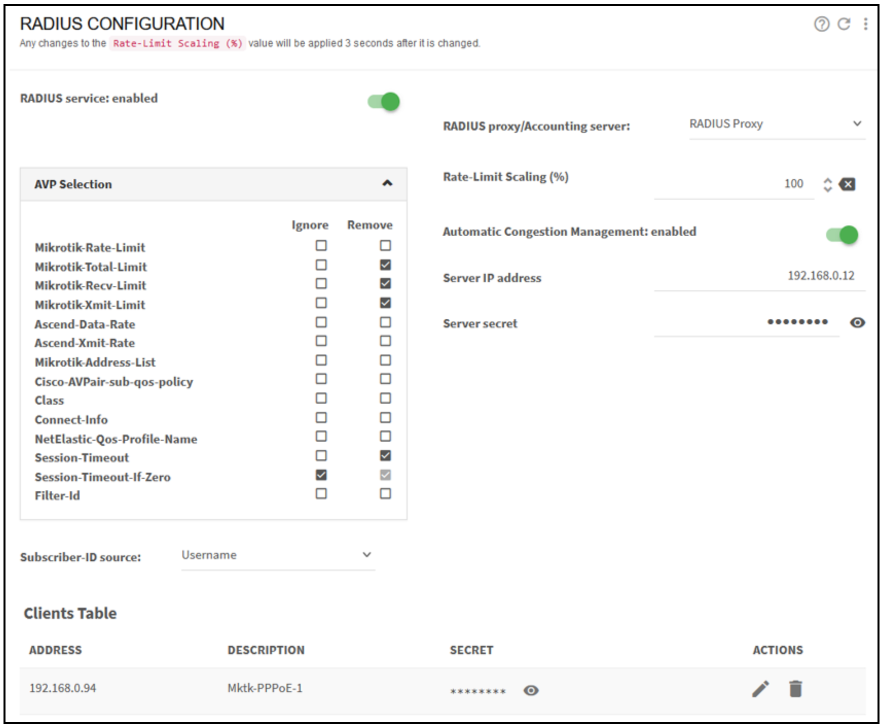

RADIUS CONFIGURATION

Any changes to the `Rate-Limit Scaling (%)` value will be applied 3 seconds after it is changed.

RADIUS service: enabled

RADIUS proxy/Accounting server: RADIUS Proxy

Rate-Limit Scaling (%) 100

Automatic Congestion Management: enabled

Server IP address 192.168.0.12

Server secret ••••••••

AVP Selection

	Ignore	Remove
Mikrotik-Rate-Limit	☐	☐
Mikrotik-Total-Limit	☐	☑
Mikrotik-Recv-Limit	☐	☑
Mikrotik-Xmit-Limit	☐	☑
Ascend-Data-Rate	☐	☐
Ascend-Xmit-Rate	☐	☐
Mikrotik-Address-List	☐	☐
Cisco-AVPair-sub-qos-policy	☐	☐
Class	☐	☐
Connect-Info	☐	☐
NetElastic-Qos-Profile-Name	☐	☐
Session-Timeout	☐	☑
Session-Timeout-If-Zero	☑	☑
Filter-Id	☐	☐

Subscriber-ID source: Username

Clients Table

ADDRESS	DESCRIPTION	SECRET	ACTIONS
192.168.0.94	Mktk-PPPoE-1	********	

Checking the Status of RADIUS Quotas

Status > Subscribers > Subscriber Attributes shows a table with all subscribers where you can see which subscriber have a Quota (QUOTA column with enabled). Clicking on enabled will lead you to the quota details.

Status > Subscribers > Subscriber Quotas shows a table with all the quota details. It will tell you if the quota is exhausted (QUOTA-BLOCK yes), the quota time limit, the quota volume limit in GB and how much of the volume quota has been consumed.

Subscriber Identification

RADIUS information can be the source of subscriber identification information.

Both accounting server and proxy deployments can be used for subscriber identication.

Subscriber ID

The QoE keeps a subscriber ID to track subscriber across IP address changes. The source of this ID can be taken from RADIUS. The list of supported RADIUS parameters are:

Name	ID	Description	Example
Calling-Station-ID	31	A descriptive string of calling number (for example, its MAC address)	“21:EF:3A:21:A4:0D”
User-Name	1	A string with the name of the user to be authenticated.	“name-surname”

In **Configuration** > **RADIUS/REST/Billling** > **RADIUS**, navigate to the **Subscriber-ID source** field and select the appropriate value (Leave unchanged by default).

To check the subscriber ID information, navigate to **Status > Subscribers > Subscriber Attributes** and see **SUBSCRIBER-ID** column.

Subscriber Group

You can assign subscribers to subscriber groups based on RADIUS information. You can see metrics broken down per subscriber groups or define policy rules that treat those groups differently (for example, make a subscribers in a group with a service not accounted as part of those subscribers volume quotas).

The RADIUS parameters that can be source of subscriber groups are listed in the following table.

Name	ID	Description	Example
Filter-Id	11	A string with the filter identifier. Only one instance of this parameter should be included in the RADIUS server Access Accept.	“my-filter-id”

They are ignored by default. Uncheck the ignore tick box to configure the QoE to use them as shown in Figure 126.

Figure 126: *The RADIUS settings page*

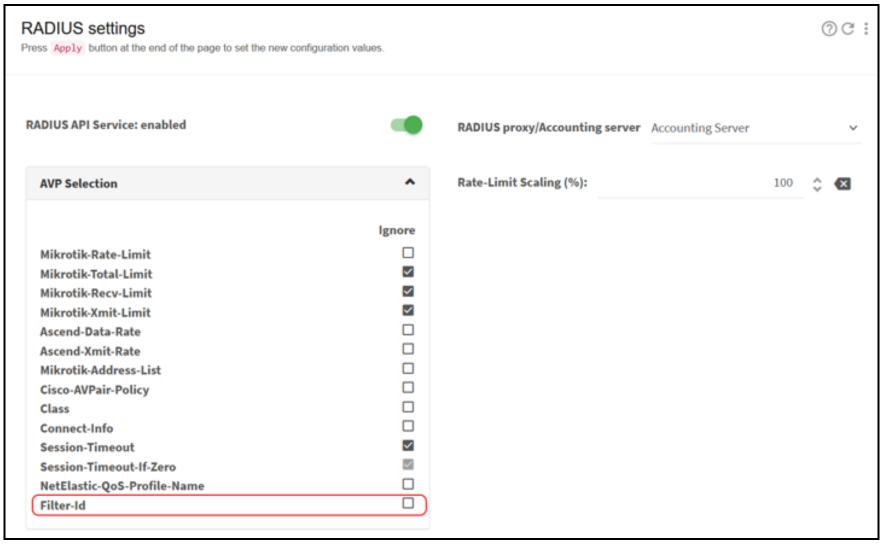

RADIUS settings

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

RADIUS API Service: enabled

RADIUS proxy/Accounting server: Accounting Server

Rate-Limit Scaling (%): 100

AVP Selection

	Ignore
Mikrotik-Rate-Limit	□
Mikrotik-Total-Limit	☑
Mikrotik-Recv-Limit	☑
Mikrotik-Xmit-Limit	☑
Ascend-Data-Rate	□
Ascend-Xmit-Rate	□
Mikrotik-Address-List	□
Cisco-AVPair-Policy	□
Class	□
Connect-Info	□
Session-Timeout	☑
Session-Timeout-If-Zero	☑
NetElastic-QoS-Profile-Name	□
Filter-Id	□

Once enabled, subscribers will be added to a group named after their Filter-Id value. RADIUS sends several instances of Filter-Id AVP, the subscriber will be added to more than one group, one for each Filter-Id value, to a maximum of 3 groups.

To check the subscriber group information, navigate to **Status** > **Subscribers** > **Subscriber Attributes** and see **SUBSCRIBER-GROUPS** column.

Status > **Subscribers** > **Subscriber Groups** shows the list of subscriber groups and the number of subscribers assigned (provisioned) and the number of them active.

If you do not want subscribers to be grouped using RADIUS, set the AVP to be ignored:

Figure 127: *RADIUS settings page*

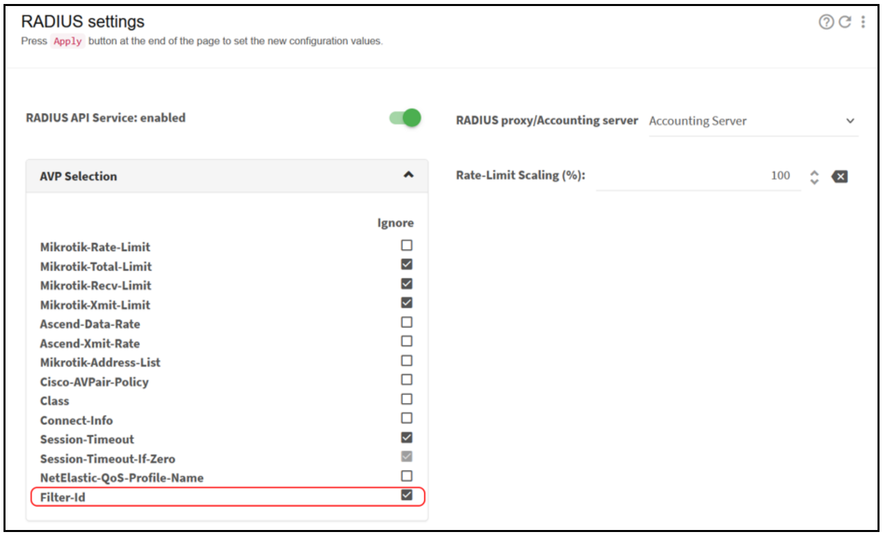

RADIUS settings

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

RADIUS API Service: enabled

RADIUS proxy/Accounting server: Accounting Server

Rate-Limit Scaling (%): 100

AVP Selection

	Ignore
Mikrotik-Rate-Limit	□
Mikrotik-Total-Limit	☑
Mikrotik-Recv-Limit	☑
Mikrotik-Xmit-Limit	☑
Ascend-Data-Rate	□
Ascend-Xmit-Rate	□
Mikrotik-Address-List	□
Cisco-AVPair-Policy	□
Class	□
Connect-Info	□
Session-Timeout	☑
Session-Timeout-If-Zero	☑
NetElastic-QoS-Profile-Name	□
Filter-Id	☑

Enable/Disable ACM Optimization

The ACM is enabled by default for all RADIUS dynamic policies. To enable or disable ACM, change the **Automatic Congestion Management** field in the configuration of the **RADIUS/REST/Billing** > **RADIUS**.

Figure 128 shows the RADIUS configuration page.

Figure 128: *The RADIUS configuration page*

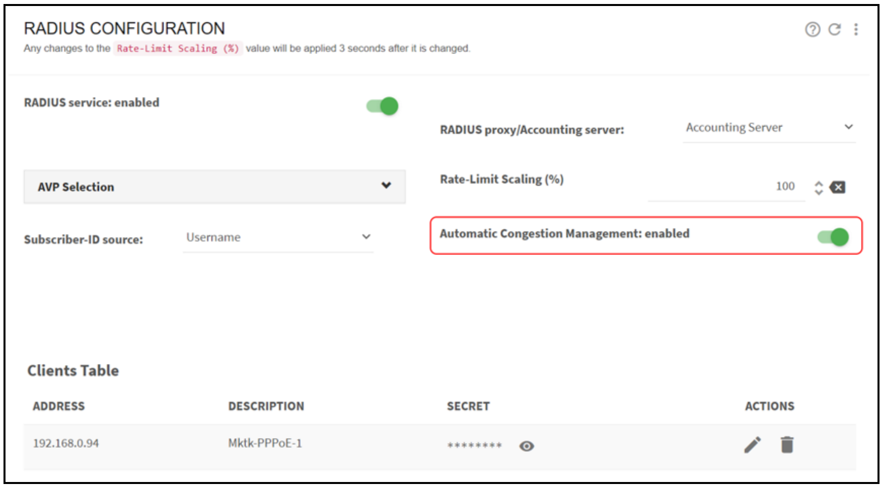

RADIUS CONFIGURATION

Any changes to the `Rate-Limit Scaling (%)` value will be applied 3 seconds after it is changed.

RADIUS service: enabled

RADIUS proxy/Accounting server: Accounting Server

AVP Selection

Rate-Limit Scaling (%) 100

Subscriber-ID source: Username

Automatic Congestion Management: enabled

Clients Table

ADDRESS	DESCRIPTION	SECRET	ACTIONS
192.168.0.94	Mktk-PPPoE-1	********	

REST API

This chapter contains the following sections:

- QoE deployment
- Definition and selection of policies
- Configuring REST in the QoE
- Digital certificate for QoE

For more REST API commands, see REST API reference.

QoE deployment

The QoE rate policies and the assignment of subscribers to rate policies can be performed through a REST API, which can also provide subscriber metrics. To do this, of course, the QoE must view the IP addresses of the subscriber data traffic, that is, there cannot be a NAT between the subscribers and the QoE, because the plan limits will be applied for each IP address. It is also important for the wires to be connected correctly (that is, for the network ports configured as access ports to be connected on the side of the subscribers).

Note

Both IPv4 and IPv6 addresses are supported. For IPv6 addresses, the REST API applies a subnet mask (/64 by default) as all the addresses of same subnet are regarded as it belongs to the same subscriber.

The integration is between a system that controls the plan definition (for example, a billing system) and the QoE, through the QoE management interface. The QoE acts as the REST API server. The QoE uses the management IP address (the same used by the GUI) to receive the REST messages, listening on TCP port 3443 (configurable). Figure 129 shows the QoE deployment.

Figure 129: *QoE deployment*

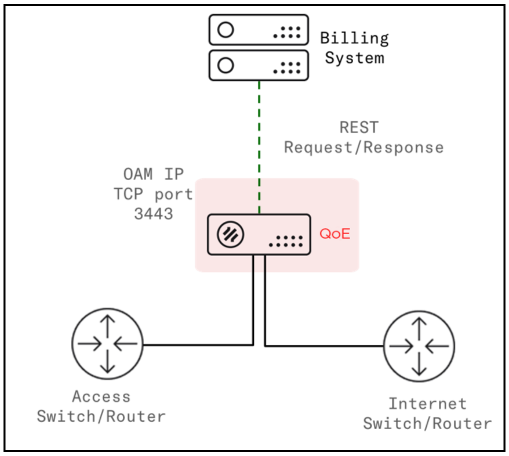

Definition and selection of policies

The following policy operations are possible through the REST API:

- Create rate control policies (the subscription plans).
- Assign a plan to a subscriber.
- Obtain subscriber metrics for up to 3 months: volume used, maximum speed, latency, retransmissions, active/created flows, traffic with congestion and at maximum speed, DPI information, and so on.
- Manage time and volume quotas.

Figure 130 shows an example in which several policies are created and later one each subscriber has one of those policies assigned.

Figure 130: *Policy creation and subscriber assignation*

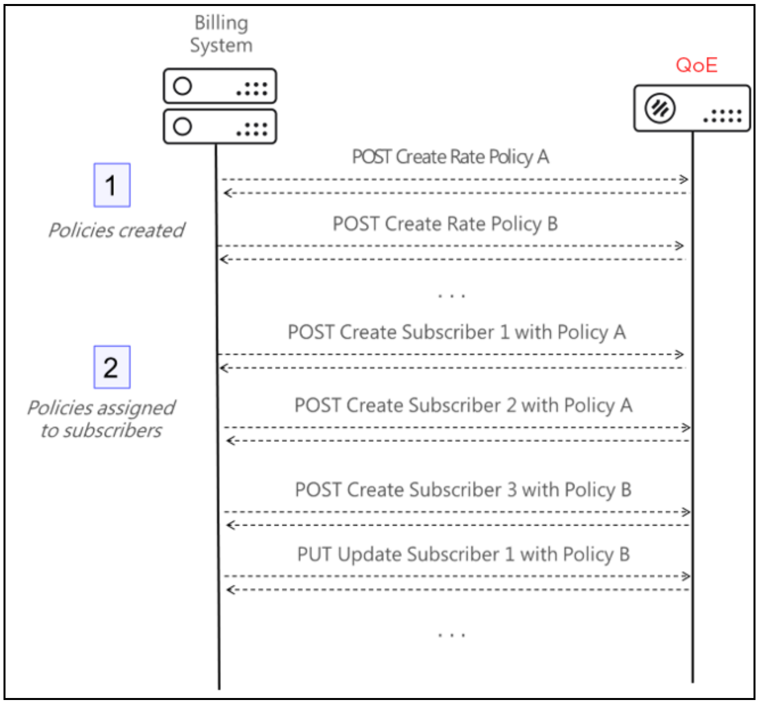
Billing System
QoE
POST Create Rate Policy A
1
Policies created
POST Create Rate Policy B
. . .
POST Create Subscriber 1 with Policy A
2
Policies assigned to subscribers
POST Create Subscriber 2 with Policy A
POST Create Subscriber 3 with Policy B
PUT Update Subscriber 1 with Policy B
. . .

There are operations to create, list, modify and delete both the policies and the subscriber associations.

The REST API is described in more detail in the REST API Guide, that includes the definition of all the possible requests.

Configuring REST in the QoE

To configure REST in the QoE, perform the following steps:

1. To activate the REST API in the QoE, navigate to **Configuration** > **RADIUS/REST/Billing** > **REST API** and switch **Rest API** to ON.
2. Add the IP addresses of all REST clients to the QoE configuration. In the GUI, navigate to **Configuration** > **RADIUS/REST/Billing** > **REST API** and click on **Add Client...**
3. Add the REST users, along with their passwords, to the QoE configuration. From the GUI, navigate to **Configuration** > **RADIUS/REST/Billing** > **REST API** and click on **Add User...**

Figure 131: *REST API configuration*

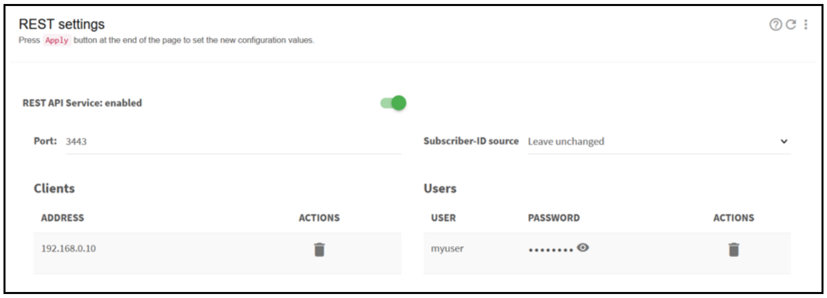

REST settings

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

REST API Service: enabled

Port: 3443

Subscriber-ID source: Leave unchanged

Clients

ADDRESS	ACTIONS
192.168.0.10	🗑

Users

USER	PASSWORD	ACTIONS
myuser	••••••••	🗑

Check REST clients

The REST API integration requires programming the API in the element interacting with the QoE, for example a billing system. To test that the QoE REST configuration is correct, REST messages can be sent manually using curl Linux utility.

As an example, to create a policy named my_rest_policy_1 and asssign it to a subscriber with IP 10.10.1.232:

```
curl -i -k -m 20 -u myuser:mypassword -X POST
https://192.168.0.121:3443/api/v1/policies/rate/my_rest_policy_1 -H
"Content-Type: application/json" --data '{"rateLimitDownlink": {"rate":
1000, "burstRate": 2000}, "rateLimitUplink": {"rate": 500, "burstRate":
1000}}'
```

Check the REST API

Once the policies and subscriber associations are created, the table **Status** > **Policies** > **Rate Policies** displays the policies coming from REST. The table in Figure 132 shows three policies, with provisioned subscribers that are also active (with traffic).

Figure 132: *The Rate policies page*

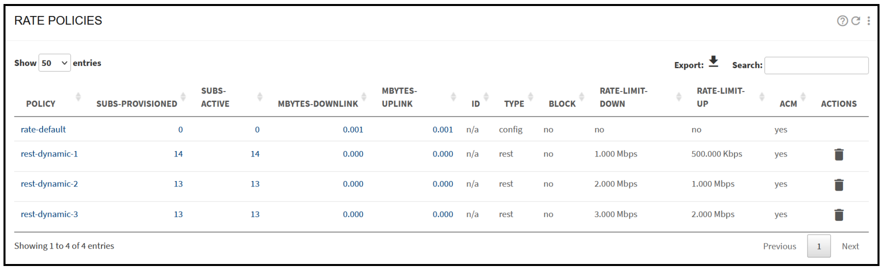

RATE POLICIES

Show 50 entries Export: Search:

POLICY	SUBS-PROVISIONED	SUBS-ACTIVE	MBYTES-DOWNLINK	MBYTES-UPLINK	ID	TYPE	BLOCK	RATE-LIMIT-DOWN	RATE-LIMIT-UP	ACM	ACTIONS
rate-default	0	0	0.001	0.001	n/a	config	no	no	no	yes	
rest-dynamic-1	14	14	0.000	0.000	n/a	rest	no	1.000 Mbps	500.000 Kbps	yes	🗑
rest-dynamic-2	13	13	0.000	0.000	n/a	rest	no	2.000 Mbps	1.000 Mbps	yes	🗑
rest-dynamic-3	13	13	0.000	0.000	n/a	rest	no	3.000 Mbps	2.000 Mbps	yes	🗑

Showing 1 to 4 of 4 entries

Previous 1 Next

To see the list of subscribers with REST policies, select **Status** > **Subscribers** > **Subscriber Attributes**. The table in Figure 133 shows ten subscribers and their associated policies.

Figure 133: *The Subscriber attributes page*

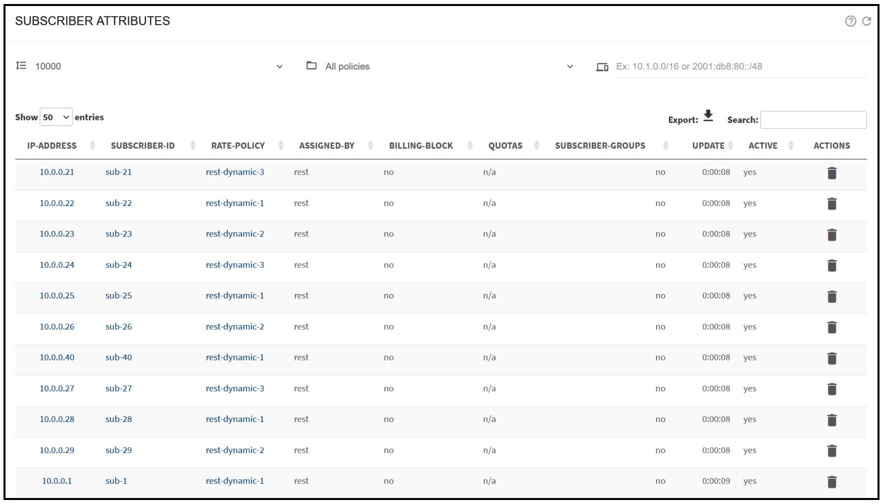

SUBSCRIBER ATTRIBUTES

10000

All policies

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

Show 50 entries

Export: Search:

IP-ADDRESS	SUBSCRIBER-ID	RATE-POLICY	ASSIGNED-BY	BILLING-BLOCK	QUOTAS	SUBSCRIBER-GROUPS	UPDATE	ACTIVE	ACTIONS
10.0.0.21	sub-21	rest-dynamic-3	rest	no	n/a	no	0:00:08	yes	🗑
10.0.0.22	sub-22	rest-dynamic-1	rest	no	n/a	no	0:00:08	yes	🗑
10.0.0.23	sub-23	rest-dynamic-2	rest	no	n/a	no	0:00:08	yes	🗑
10.0.0.24	sub-24	rest-dynamic-3	rest	no	n/a	no	0:00:08	yes	🗑
10.0.0.25	sub-25	rest-dynamic-1	rest	no	n/a	no	0:00:08	yes	🗑
10.0.0.26	sub-26	rest-dynamic-2	rest	no	n/a	no	0:00:08	yes	🗑
10.0.0.40	sub-40	rest-dynamic-1	rest	no	n/a	no	0:00:08	yes	🗑
10.0.0.27	sub-27	rest-dynamic-3	rest	no	n/a	no	0:00:08	yes	🗑
10.0.0.28	sub-28	rest-dynamic-1	rest	no	n/a	no	0:00:08	yes	🗑
10.0.0.29	sub-29	rest-dynamic-2	rest	no	n/a	no	0:00:08	yes	🗑
10.0.0.1	sub-1	rest-dynamic-1	rest	no	n/a	no	0:00:09	yes	🗑

Digital certificate for QoE

The REST clients need a specific digital certificate to validate the access.

Log in as root to QoE Unix shell.

Look for the following directories:

```
/bqn/root/etc/ssl/certs
```

```
/bqn/root/etc/ssl/private
```

If the directories do not exit, create them:

```
mkdir -p /bqn/root/etc/ssl/certs/
mkdir -p /bqn/root/etc/ssl/private/
chmod 700 /bqn/root/etc/ssl/private/
```

Copy the certificate and its key in the following files:

```
/bqn/root/etc/ssl/certs/bqn.crt
```

```
/bqn/root/etc/ssl/private/bqn.key
```

Note

Use scp to transfer previously obtained certificate files to the QoE server.

To apply the change, reboot the QoE server in **Administration** > **Reboot**.

Billing systems

Subscriber data can be retrieved from a number of supported billing systems in addition to RADIUS and REST.

Billing integrations support only IPv4 addresses.

This chapter contains the following sections:

Aradial

Aradial uses RADIUS. For more information, see RADIUS Interface chapter.

Azotel

The QoE retrieves the customer bucket data and get from it the 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 username and password are required 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, but bear in mind that if the QoE reaches Azotel over the Internet, Azotel will see a public IP address and this is the one that will need authorization by the Azotel system.

Provide also the Azotel system IP address or server name and port number (443 by default). Figure 134 shows the billing systems integrations page.

Figure 134: *The Billing systems integrations page*

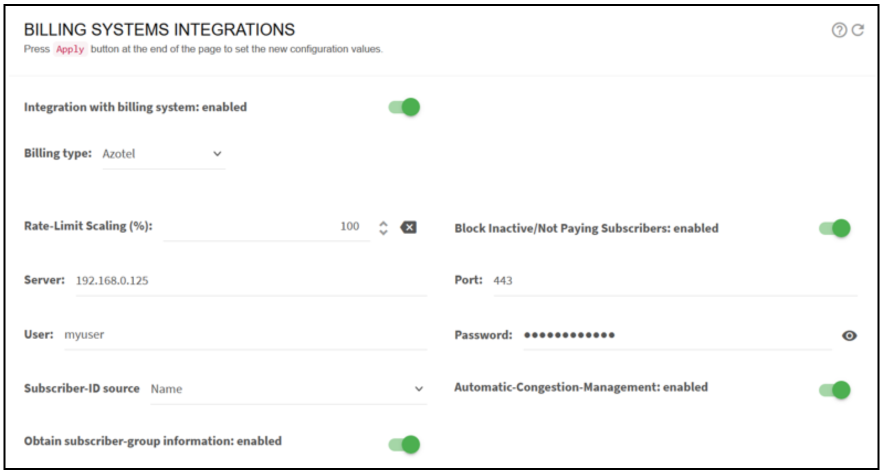

BILLING SYSTEMS INTEGRATIONS

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 ●

Obtain subscriber-group information: enabled ●

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

Policy rates are taken from customer's bucket fields **uploadrate** and **downloadrate**.

With subscriber group information enabled, and if the information is available in Azotel database, the nickname of the parent of the customer CPE (for example, an AP) will be used as a subscriber group, prefixed with **L1-**. Also, the site name where the parent is located will be a subscriber group, with **L2-** prefix. There are two levels: a first level made up of AP subscriber groups (L1) and a second level with sites (L2), that can contain one or more APs.

Gestfy

Gestfy uses QoE REST API. For more information, see REST API section.

ISPSolution

ISPSolution uses QoE REST API. For more information, see REST API section.

Microwisp

Mikrowisp uses RADIUS. For more information, see RADIUS interface section.

Powercode

To activate Powercode, navigate to **Configuration** > **RADIUS/REST/Billing** > **Billing Systems**, select **Powercode** and enable the switch. Figure 135 shows the billing systems integration page.

Provide the Powercode system IP address or server name and its SSH port number (22 by default).

The QoE server needs SSH access to the Powercode server using a Unix User/Password.

If MySQL user/password is different to the Unix user/password, that should be specified in MySQL Credentials. The MySQL user must have read access to the following tables of the MySQL database:

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

Figure 135: *The billing systems integrations page*

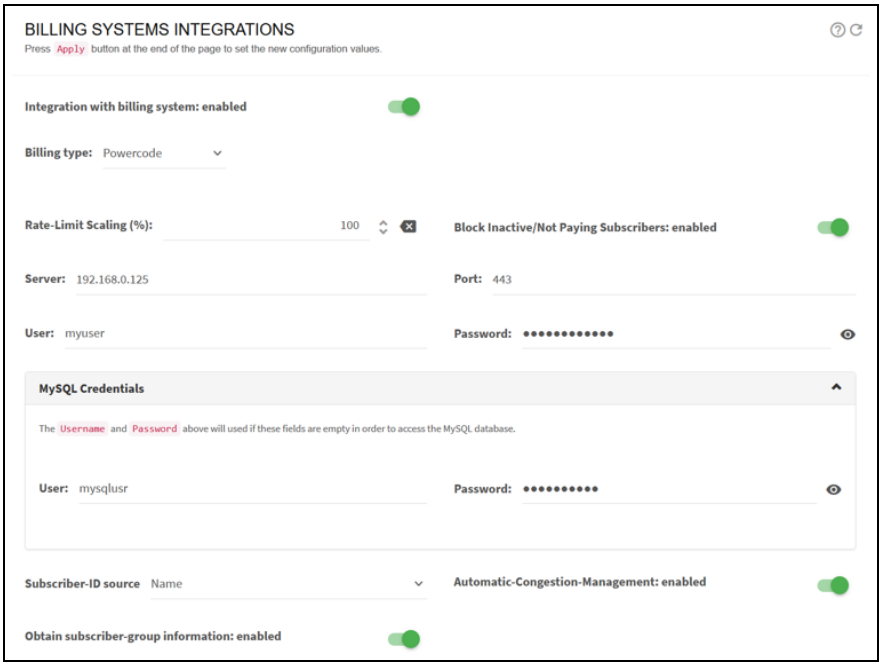

BILLING SYSTEMS INTEGRATIONS

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: 443

User: myuser

Password: ••••••••••••

MySQL Credentials

The Username and Password above will used if these fields are empty in order to access the MySQL database.

User: mysqlusr

Password: ••••••••••

Subscriber-ID source Name

Automatic-Congestion-Management: enabled

Obtain subscriber-group information: enabled

Powercode customers in a status other than **Active** are 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 Powercode billing can be used as source of the subscriber ID:

- Customer-ID
- Equipment-ID

- MAC address
- Name

Policy rates are taken from customer service Internet Info (MaxIn, MaxOut, BurstIn, BurstOut and BurstBucktTime).

With subscriber group information enabled, and if the information is available in Powercode database, the **LocationID** of the equipment of the customer CPE will be used as a subscriber group, prefixed with **L1-**.

REST-API Powercode

Powercode billing restricts REST API to three requests per second. For that reason, the preferred integration is using the SQL access described in the Microwisp section. REST-API can be used when SQL cannot be used and the number of subscribers is low (one thousand or less). Figure 135 shows the billing systems integrations page.

Figure 136: *The Billing systems integrations page*

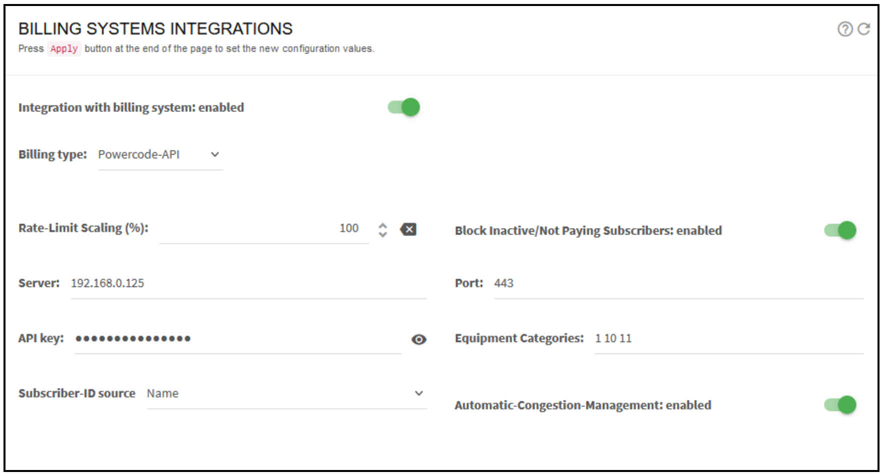

BILLING SYSTEMS INTEGRATIONS

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

Integration with billing system: enabled

Billing type: Powercode-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

The QoE retrieves CPE equipment of a certain category (1 by default). For subscribers with that category of equipment, it will retrieve 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.

The API key must be created in the Powercode system, with allowed access from QoE IP address (the QoE uses its management address for Powercode queries).

Provide 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, change it. 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).

Note

Retrieval of subscriber group information is not possible. Use the SQL-based option instead.

Sonar

Sonar v2 is supported (the one with GraphQL API).

The QoE retrieves the customer tariff and gets the speed limits from it 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, see Sonar Knowledge Base article.

Figure 137: *Creating an API key*

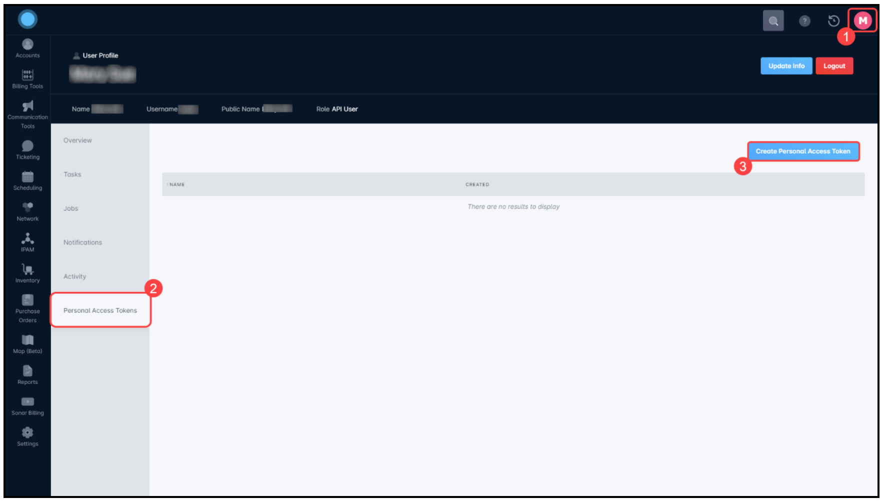

[redacted]

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

NAME	CREATED
There are no results to display	

Enter the IP address of the Sonar system or server name and port number (443 by default). Figure 138 shows the billing systems integration page.

Figure 138: *The billing systems integration page*

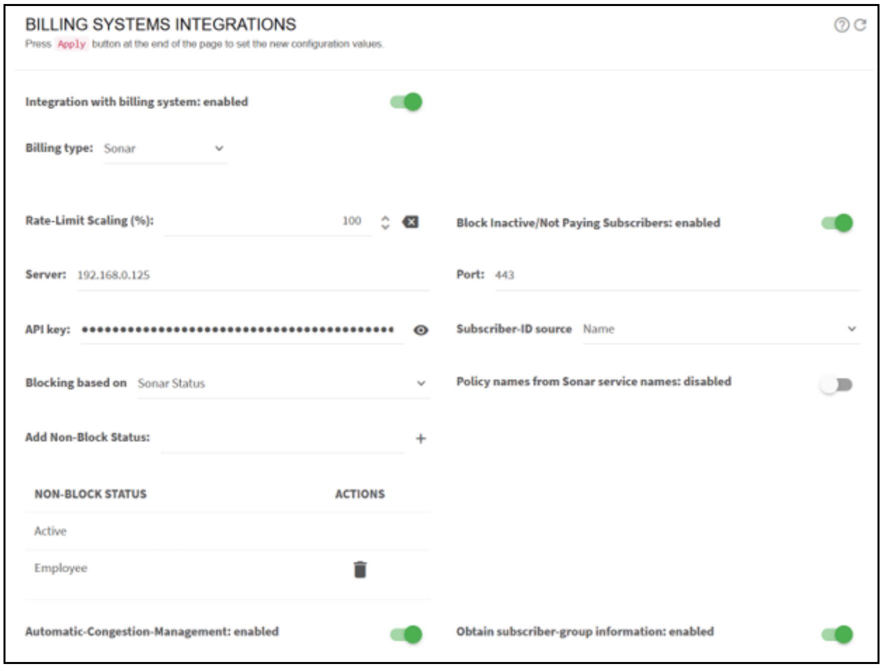

BILLING SYSTEMS INTEGRATIONS

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

Integration with billing system: enabled

Billing type: Sonar

Rate-Limit Scaling (%): 100

Block Inactive/Not Paying Subscribers: enabled

Server: 192.168.0.125

Port: 443

API key: ••

Subscriber-ID source Name

Blocking based on Sonar Status

Policy names from Sonar service names: disabled

Add Non-Block Status:

NON-BLOCK STATUS	ACTIONS
Active	
Employee	

Automatic-Congestion-Management: enabled

Obtain subscriber-group information: enabled

By default, Sonar customers with **account_status** > **name** field with a value other than **Active**, are blocked (they are regarded as lacking a valid subscription). You can add more status in addition to **Active**, so they are not blocked either (see example shown in Figure 138, where status **Employee** is added). You can change the QoE behaviour to non-blocking regardless of the account status by disabling the switch **Block Inactive/Not Paying Subscribers**.

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

- Customer-ID (account ID field in Sonar)
- Name (account name in Sonar)

The speed limits are taken from data service detail fields:

- download_speed_kilobits_per_second
- upload_speed_kilobits_per_second

By default, the name of the rate policy is created internally by the QoE. Alternatively, if the switch **Policy names from Sonar service names** is enabled, the policies will be named using Sonar service name plus the service Id appended to it, in order to have a unique identifier (Sonar service names alone are not guaranteed to be unique).

With subscriber group information enabled, and if the information is available in Sonar database, a subscriber group, prefixed with **L1-**, will be created based on the description of the parent inventory item of the subscriber's inventory item.

Splynx

The 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 **Configuration** > **RADIUS/REST/Billing** > **Billing Systems**, select **Splynx** and enable the switch.

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 but bear in mind that if the 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**.

The following screenshots display the API KEY and the access permits:

Figure 139: *API KEY and access permits*

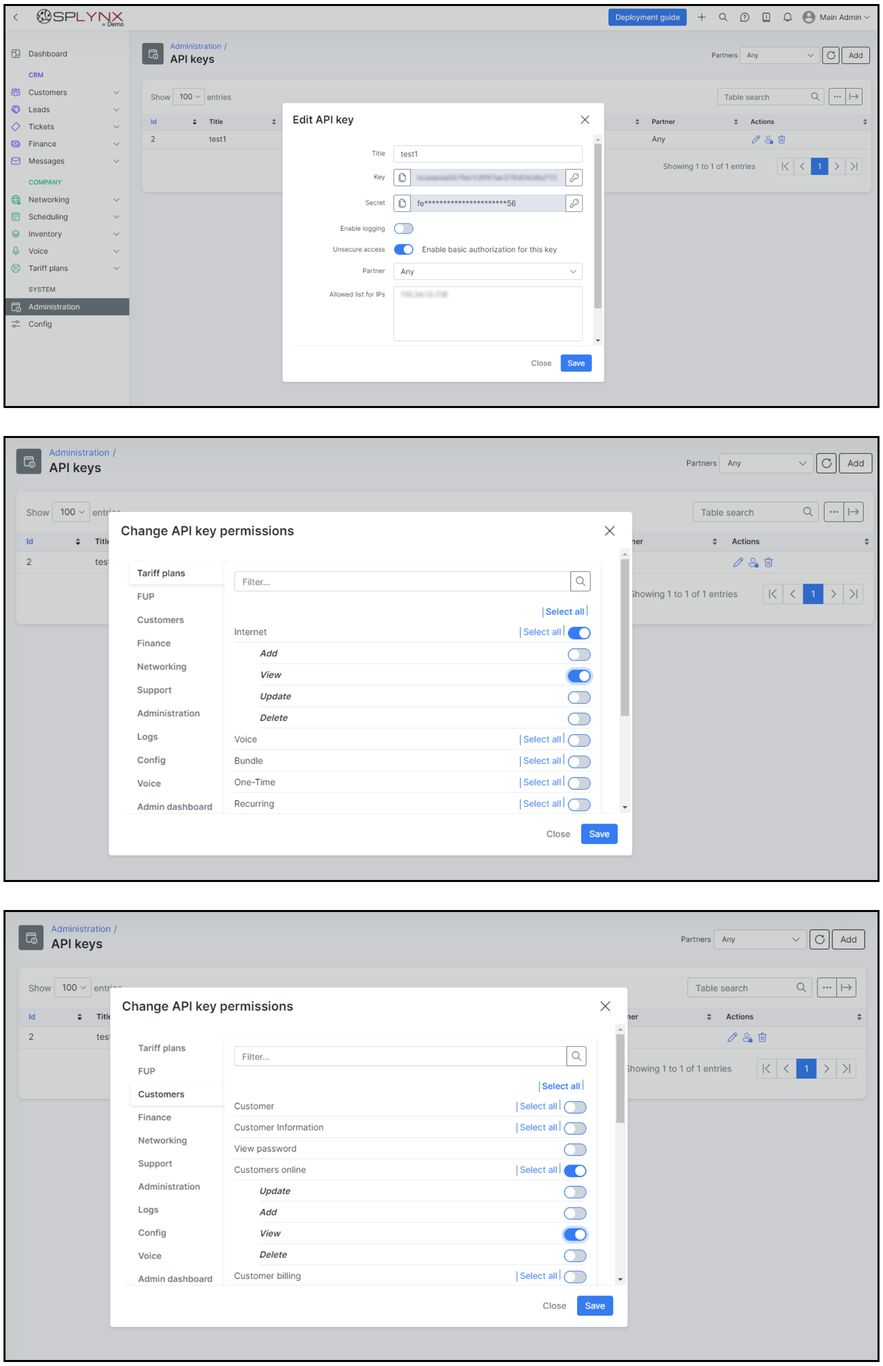

Edit API key

Title: test1

Key: [redacted]

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

Enable logging: ○

Unsecure access: ● Enable basic authorization for this key

Partner: Any

Allowed list for IPs: [redacted]

Change API key permissions — Tariff plans

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

Customer	Select all	○
Customer Information	Select all	○
View password		○
Customers online	Select all	●
Update		○
Add		○
View		●
Delete		○
Customer billing	Select all	○

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

Figure 140: *Splynx*

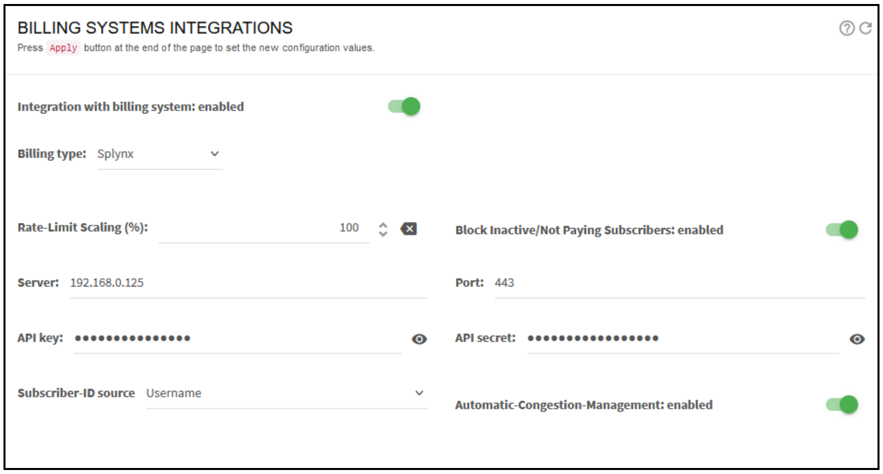

BILLING SYSTEMS INTEGRATIONS

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

Splynx customers with **1** in their blocked field 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**.

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

- Customer-ID
- Username
- Login

The rate policy speed limits are taken from the Internet tariff associate to the online customer:

- speed_download
- speed_upload

> **Note**
>
> Retrieval of subscriber group information is not possible, as there is no standard field in the billing system with that information.

UISP

A tool is available on Cambium QoE support site for integrating UISP (cloud and on-premise) with QoE. The tool includes a user's guide that explains how to set it up and deploy it. The tool uses QoE REST API.

Visp.net

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

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

A valid client id and secret and user name and password must be provided to the QoE in order to request temporal API tokens. A client id is unique per Visp installation. A user is any of the valid user accounts in that 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). Figure 141 shows an example of Visp configuration:

Figure 141: *Visp configuration*

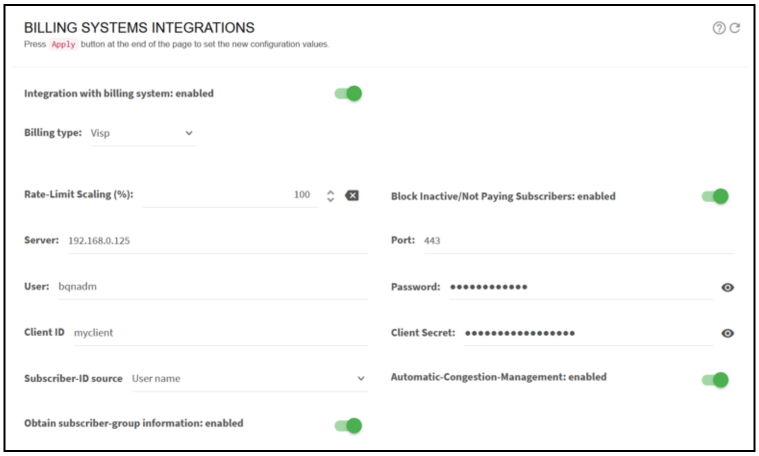

BILLING SYSTEMS INTEGRATIONS

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

Integration with billing system: enabled

Billing type: 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

Obtain subscriber-group information: enabled

With the switch **Block Inactive/Not Paying Subscribers** enabled, the QoE will block subscribers based on two possible criteria:

- Status of service instances and packages (the default). If neither of them are in **ACTIVE** state, the subscriber will be blocked.
- Status of the subscriber account. If this option is chosen, subscribers in states Suspended, Hibernated, and Inactive will be blocked (status codes 8, 9 and 10, respectively).

Visp customers with a package status and service instance status other than **ACTIVE** are blocked. You can change this behaviour to non-blocking disabling the switch **Block Inactive/Not Paying Subscribers**.

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

- Customer-ID
- First + Last name
- Username

Policy rates are obtained from customer's service instance **up_speed** and **down_speed**. The burst rates **up_burst** and **down_burst** are also obtained.

With subscriber group information enabled, the access point identifier associated to the CPE of the subscriber, with an **L1-** prefix, will be the name of the subscriber group, if the information is available in the Visp database.

WISPControl

WISPControl uses RADIUS. For more information, see RADIUS interface section.

Wisphub

Wisphub has developed an integration with QoE using QoE REST API. For more information, see 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. For instructions, see API REST.

The provided IP address or server name is used along with the port (443 by default) to send API queries to Wispro (https://<server>:<port>/api/v1). Figure 142 shows an example of Wispro configuration.

Figure 142: *Wispro configuration*

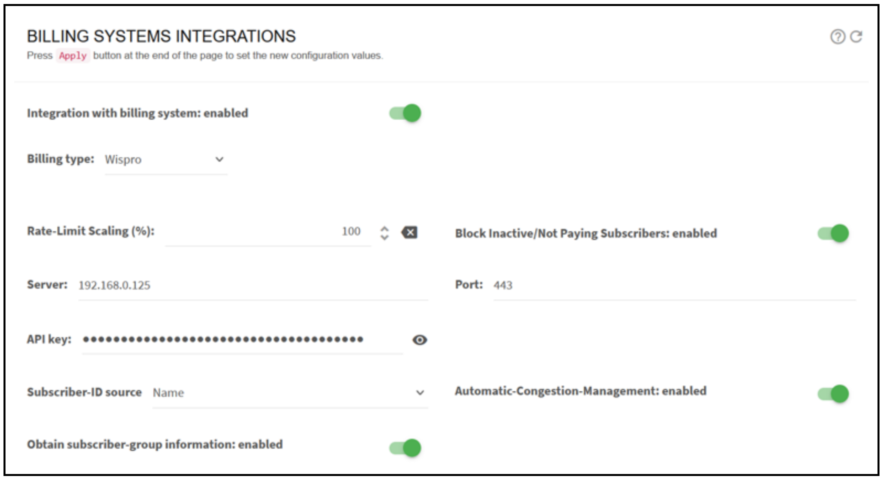

BILLING SYSTEMS INTEGRATIONS

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

Obtain subscriber-group information: enabled

Only Wispro clients with contracts in **disabled** state are blocked. You can change this behaviour to non-blocking by 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).

Policy rates are obtained from the plan of the client contract. The plan **ceil_down_kbps** and **ceil_up_kbps** parameters are used as speed limits.

If a Wispro client has more than one contract, each IP address is 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 are selected.

With subscriber group information enabled, node names will be used as subscriber groups, that is, a subscriber will be associated to a subscriber group after the node name of the contract, if that information is available in Wispro database.

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 explicitly defined when creating or editing the subscriber. Figure 143 shows the billing systems integration page.

Figure 143: *The Billing systems integration page*

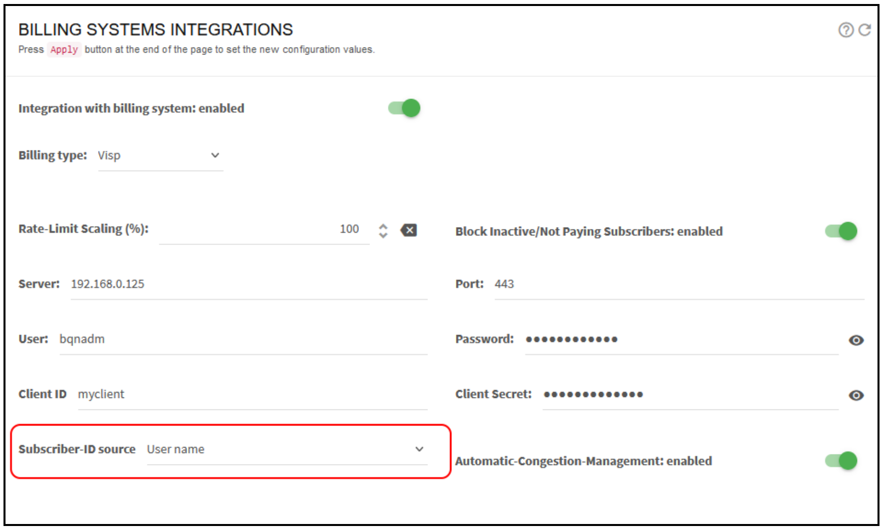

BILLING SYSTEMS INTEGRATIONS

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

Integration with billing system: enabled

Billing type: 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 blocks 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. Figure 144 shows the billing systems integration page.

Figure 144: *The Billing systems integration page*

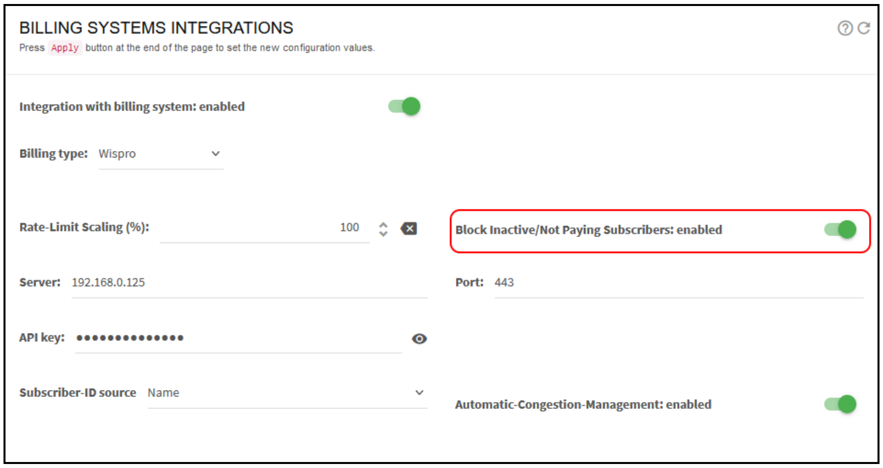

BILLING SYSTEMS INTEGRATIONS

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

Integration with billing system: enabled

Billing type: 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 field **Rate-Limit Scaling**.

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% is 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% is 300 Mbps.

Figure 145: *Rate-Limit Scaling*

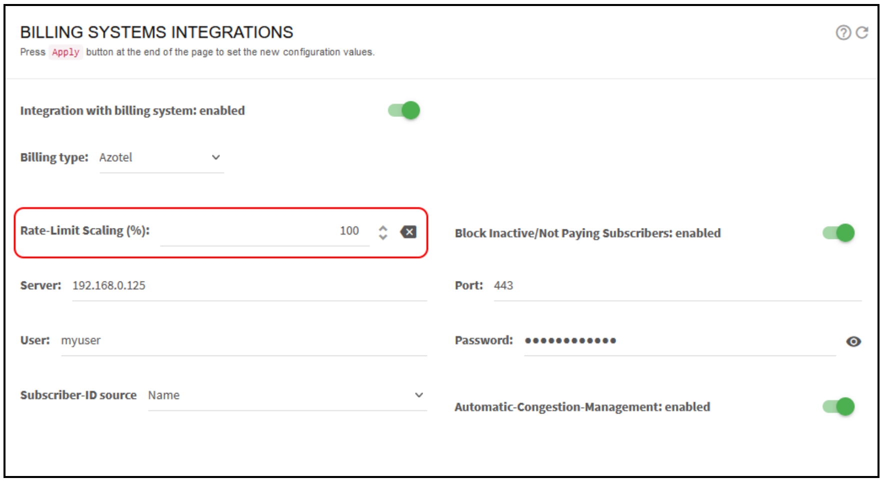

BILLING SYSTEMS INTEGRATIONS

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.

Troubleshooting

This chapter contains the following sections:

Installation issues

For installation-related problems, verify the following:

- The ISO should have been 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 overwrote the USB drive.

No access to the management IP address

The QoE uses a dedicated network interface for management. The management interface supports both the SSH and WEB (HTTPs) services. In case of problems accessing the configured management IP check 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 up. 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 accessing the management IP address from a different network, make sure that static routing is configured to the access network, as explained in the Network Interface section in the User Guide.

- If there are firewalls in the management network, allow 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 message will request a reboot. After the reboot, the QoE should have the new IP and management network interface.

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

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

- The QoE management interface may be protected by its own firewall. The problem could be that your source IP address is not included in that firewall whitelist. This could happen even for addresses from the same subnet of the QoE management IP if the subnet is not part of the firewall rules. You can disable the firewall temporarily until the connection to the management port is restored. Connect a monitor and a keyboard to the server and log in 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@bqn0# show interface
firewall

IFACE CHAIN RANGE

root@bqn0#
```

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

- If the command `wizard bta interactive` fails, connect with a console and a keyboard and try changing the configuration directly. In the following example, the management interface is in en0o2 and we will move it to interface en0o1 and from IP address 192.168.0.121 to IP address 10.0.0.121, with default gateway 10.0.0.1:

```
bqn0:# bqnsh

root@bqn0# show interface management detail

Interface: en0o2

IP address: 192.168.0.121/24

Default gateway: 192.168.0.1

Nameserver: 8.8.8.8

root@bqn0# show interface

IFACE TYPE MAC STATE LINK

lo0 loopback 00:00:00:00:00:00 up yes

en0o1 ethernet ff:2f:6b:72:12:01 up no

en0o2 ethernet ff:2f:6b:72:12:02 up yes

en0o3 ethernet ff:2f:6b:72:12:03 up no

en0o4 ethernet ff:2f:6b:72:12:04 up no

root@bqn0# configure

root@bqn0(config)# interface en0o2

root@bqn0(config-iface)# no ip address 192.168.0.121/24

root@bqn0(config-iface)# no management

root@bqn0(config-iface)# root

root@bqn0(config)# interface en0o1

root@bqn0(config-iface)# ip address 10.0.0.121/24

root@bqn0(config-iface)# management

root@bqn0(config-iface)# root

root@bqn0(config)# router static

root@bqn0(config-static)# route 0.0.0.0/0 via 10.0.0.1

root@bqn0(config-static)# root

root@bqn0(config)# commit

root@bqn0(config)# end

root@bqn0# exit
```

```
bqn0:#
```

Web is not accessible

If there are any problems in accessing the web, verify the following:

- Check that the management IP address is accessible using SSH.
- Check that you are using HTTPS in the access (HTTP is not supported). Example of URL: https://192.168.0.121
- Check that you are using bqnadm user (root cannot be used in GUI access).
- When installing from scratch, make sure the command **wizard bta** was run (otherwise, the GUI web service will not be active a no **bqnadm** user created).
- Verify that the SSH port of the QoE server has not been modified. To access the QoE using a port other than 22, you can define port forwarding rules in router on the access path, but the QoE SSH port cannot be changed. Logging to the server as root, you can verify that the SSH port is 22 as follows:

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

  #Port 22

  #GatewayPorts no
  ```

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

- Check that your browser is supported (Edge, Firefox, Chrome). MS Explorer is not supported.

Network interfaces down

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

Figure 146: *The QoE dashboard*

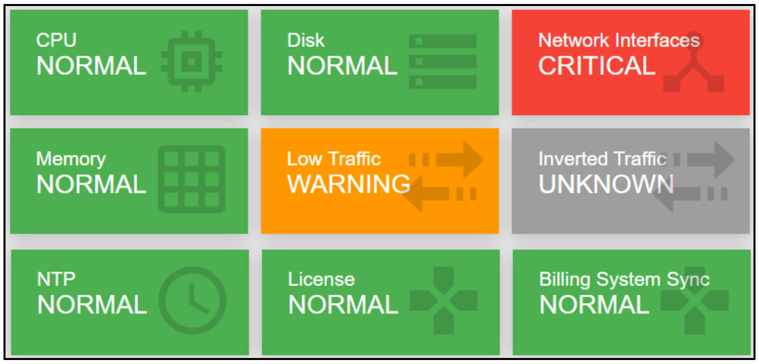

CPU
NORMAL

Disk
NORMAL

Network Interfaces
CRITICAL

Memory
NORMAL

Low Traffic
WARNING

Inverted Traffic
UNKNOWN

NTP
NORMAL

License
NORMAL

Billing System Sync
NORMAL

In Red (Critical)

- If there is no wire configured, create one.
- If there are wires configured but their interfaces are not in UP state, this most likely indicates that the interfaces are not Intel compatible. Navigate to **Configuration > Interfaces > Data Wires**.

- Click on the down interface to see the model (PCI vendor ID) and confirm it is not Intel.
- 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 1Gbps).

- 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, then the problem is on the other equipment.
  - If the link is still down and optic ports are used, verify that the transceivers:
    - are Intel-compatible.
    - are supported (see Supported Network Interfaces for a list).
    - of the type required by the installation (for example, SFP+-LR in an installation with monomode fiber and SFP+-LR on the other side).

In Yellow (Notice)

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

Inverted traffic

If the **Inverted Traffic** icon in the dashboard is in orange (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 in a QoE in a lab) but in a network deployment most likely indicates that some of the wires have been connected incorrectly, with the access port connected to the Internet side and vice versa.

Figure 147: *The QoE dasboard*

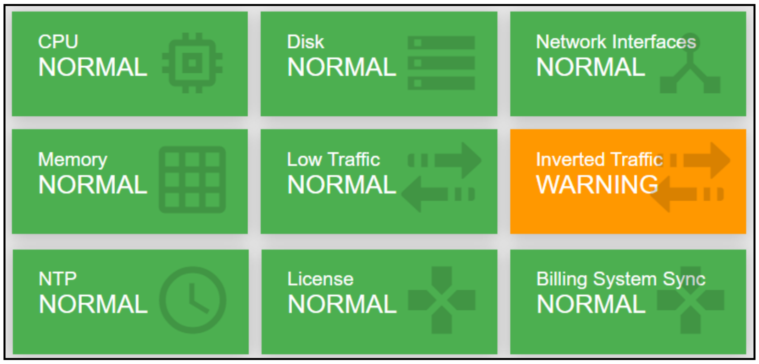

To verify that this is the case, navigate to **Statistics** > **Throughput** > **Interfaces** and select the wire interface configured in the access side. If it shows more received throughput that sent throughput, it is indication that the wire is inverted. This can be confirm selecting the throughput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

To fix the issue, navigate to **Configuration** > **Interfaces** > **Data Wires** and, in the inverted wire, press the **Swap interfaces** icon.

Low traffic

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

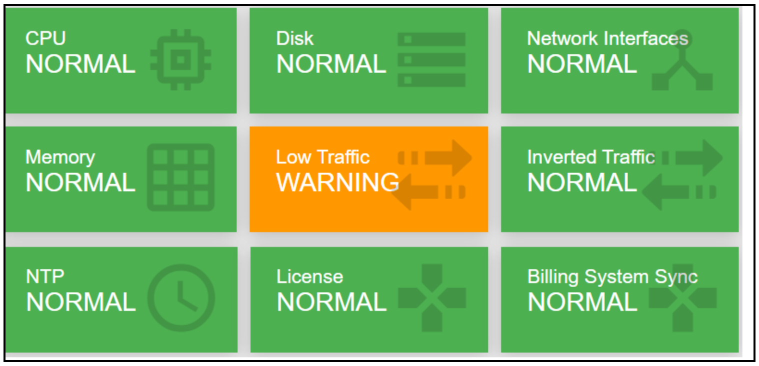

CPU NORMAL

Disk NORMAL

Network Interfaces NORMAL

Memory NORMAL

Low Traffic WARNING

Inverted Traffic NORMAL

NTP NORMAL

License NORMAL

Billing System Sync NORMAL

In Yellow (Notice)

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

In Orange (Warning)

Hover the mouse over the 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. You should navigate to **Configuration** > **TCPO/ACM Settings** and set the **Overall TCP Optimization** to OFF while the issue is being delt with.

Fix the traffic routing so the QoE sees both directions (uplink and downlink). When done and after the icon returns to green, navigate to **Configuration** > **TCPO/ACM Settings** to enable the overall TCP optimization.

License manager

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

Figure 148: *The QoE dashboard*

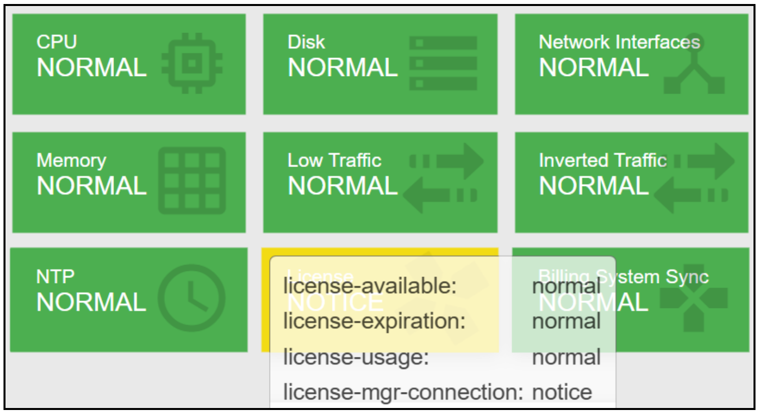

```
license-available:       normal
license-expiration:      normal
license-usage:           normal
license-mgr-connection:  notice
```

Ensure the QoE server can initiate outgoing connections to the License Manager IP (For details, see Cambium Network Support Site).

To verify that an outgoing connection is possible, log in as **root** and a telnet to the provided IP and 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 red, and the text says **license-available: critical**, there is no valid license. This could be due to several reasons:

- There is no license defined in the node
- The license is not valid
- The license is no longer valid (its final date has 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. The service is not affected, but none of the QoE advanced processing is applied to the traffic.

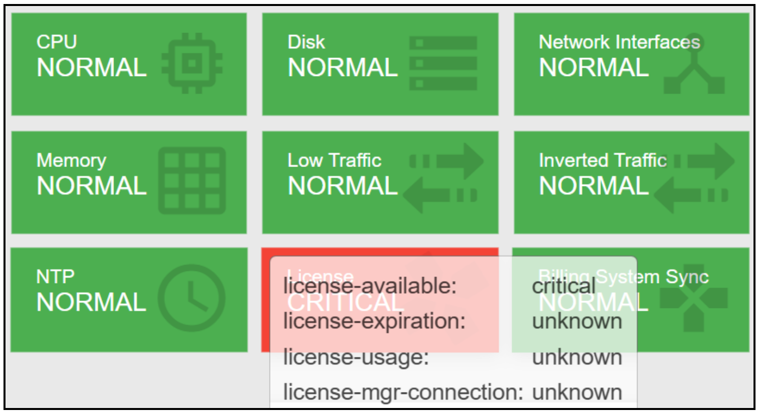

```
license-available:       critical
license-expiration:      unknown
license-usage:           unknown
license-mgr-connection:  unknown
```

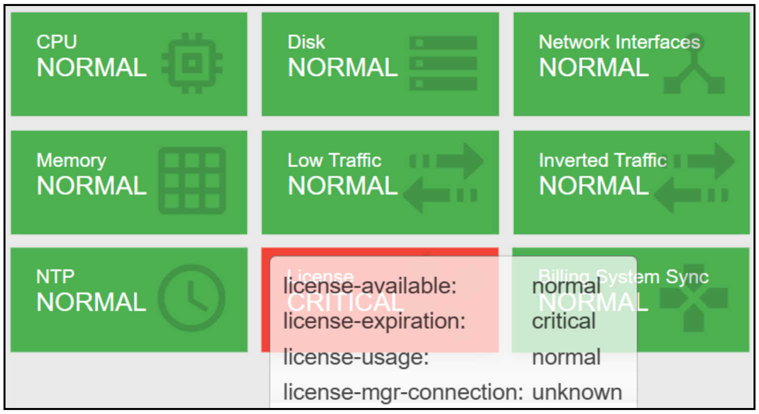

```
license-available:       normal
license-expiration:      critical
license-usage:           normal
license-mgr-connection:  unknown
```

License limit exceeded

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

Figure 149: *The QoE dashboard*

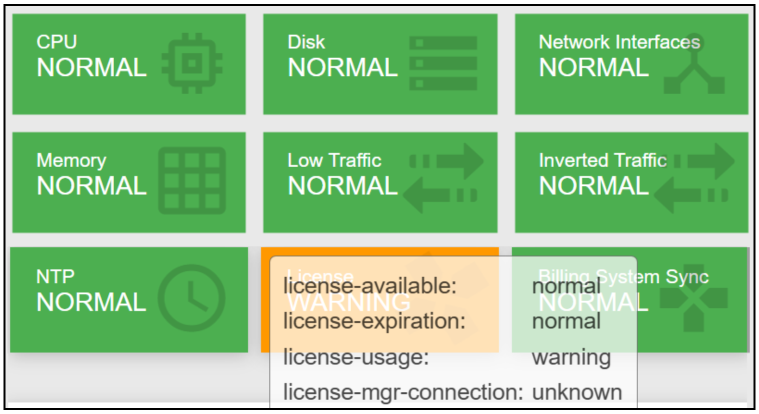

```
license-available:       normal
license-expiration:      normal
license-usage:           warning
license-mgr-connection:  unknown
```

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

In **Statistics** > **Throughput** > **Overview**, a red line will show the license limit along with recent 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 will 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, and so on).

The following are the effects at subscriber session level:

- While the license limit is being exceeded, new subscriber sessions will have no license.
- Existing subscriber sessions prior to the license being exceeded remain initially in their current state: if they were being optimized remain so and likewise if without a license.
- If the license remains above the limit, subscriber sessions being optimized may transition to no license if they accumulate a big number of non-optimized 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 remain initially in their current state: if they had no license remain so and likewise if they were being optimized.
- If the license remains below the limit, subscriber sessions without a license may transition to being optimized if they accumulate a small number of optimized flows.
- A subscriber session without a license has:
  - No ACM
  - No rate limiting.
  - It continues generating traffic volume totals.
  - Quota volume reporting is suspended.
  - Blocking at quota exhaustion will not take place.

- The other metrics, like retransmissions, latencies, and so on are reduced to those from the flows being optimized.

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

The effect of exceeding the licensed limit is that more and more traffic will no longer get QoE functionality, and, conversely, the amount of traffic getting QoE functionality gets lower. Once the amount of traffic getting QoE functionality is below the license limit, new flows and subscriber sessions will regain QoE functionality. This continues until the license limit is exceeded, after which the effect repeats again. With this oscillating behavior, the QoE keeps providing functionality to traffic up to the license limit. The QoE continuously checks the level of throughput versus the license limits (every minute or less), so the adaptation is pretty quick.

High CPU load

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

Figure 150: *The QoE dashboard*

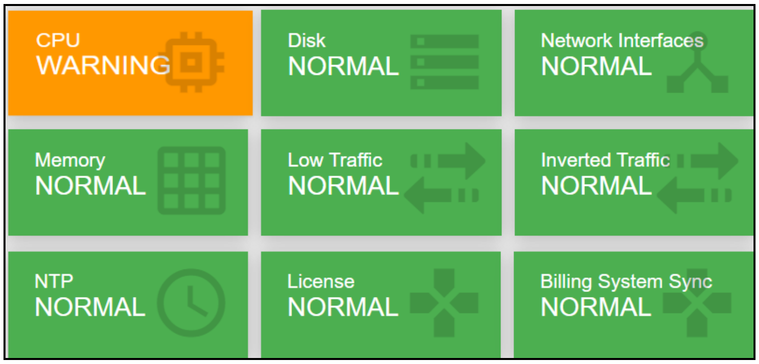

The QoE has internal mechanisms to mitigate this situation, trying to prevent traffic losses by reducing the amount of optimized traffic.

The throughput level can be verified in **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% usage).

- Red if some CPU cores are at very high load (above 90% usage).

Follow these steps:

- 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 in **Configuration** > **TCPO/ACM Settings**.
- A hardware upgrade may be needed. 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.

Figure 151: *The General settings page*

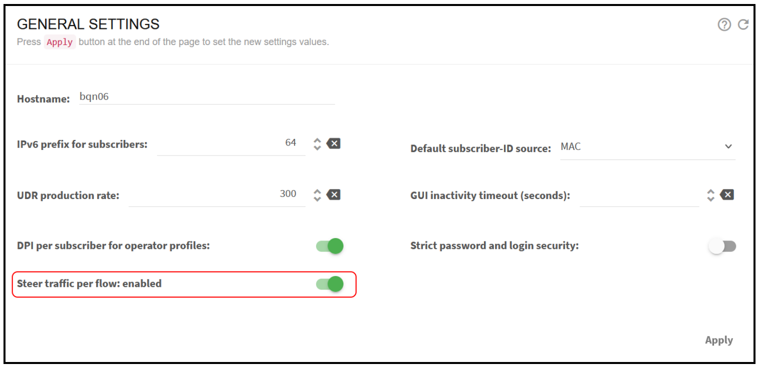

GENERAL SETTINGS

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 ●

High memory load

If the Memory icon in the Dashboard is not in green, some processes are running out of memory. This is normally due to unbalanced traffic (traffic concentrated in a few subscriber IPs) or to too much traffic being processed by the QoE server.

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

The throughput level can be verified in **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% usage).
- **Red**: If some processes reach very high usage (above 95% usage).

Follow these steps:

- 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 in **Configuration** > **TCPO/ACM Setting**s.
- A hardware upgrade may be needed. For more information, visit [Cambium Networks Support Site](#).

No RADIUS messages are received

If no RADIUS information is shown in the QoE GUI, to check if RADIUS messages are incoming, log in 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 servers in **Configuration** > **Interfaces** > **Management** (the 0 added by QoE configuration must be removed, for example, **en0s1f0** in QoEGUI 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).

Figure 152: *The management interface firewall page*

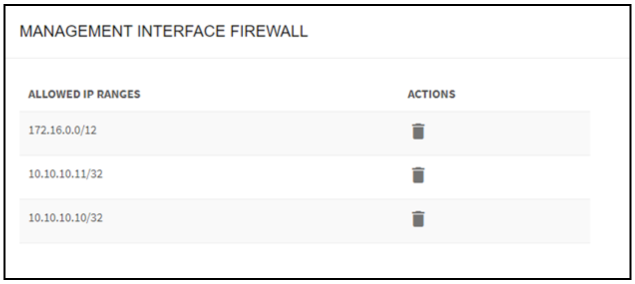

MANAGEMENT INTERFACE FIREWALL

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

And now the RADIUS messages are received:

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

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 still not received, the rest of the traffic jumps need verification. In our example, RADIUS clients are in subnet 10.10.10.0/24 and QoE in subnet 192.168.0.0/24. It has to be verified if there are valid routes between the two subnets and that 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, you can see 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. Log in to the QoE shell as **root** and execute the following command:

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

```
$ ssh bqnadm@192.168.0.121

bqnadm@bqn# system interface en0o1 capture filter ‘tcp and port 3443’

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

Check the one in your server by navigating to **Configuration** > **Interfaces** > **Management** (remove the 0 added by QoE configuration, for example, **en0s1f0** in the QoE GUI 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).

Figure 153: *The management interface firewall page*

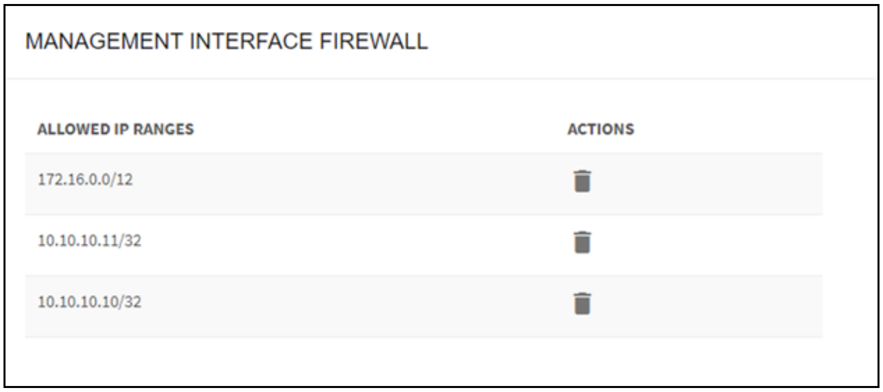

MANAGEMENT INTERFACE FIREWALL

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

And now the messages should be received:

```
bqnadm@bqn# system interface en0o1 capture filter 'tcp and port 3443'

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 yet received, 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. Check 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 through 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, you can see the event log of the received REST requests. 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)

. . .
```

And 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 red, the access to the billing system is failing.

Figure 154: *The QoE dashboard*

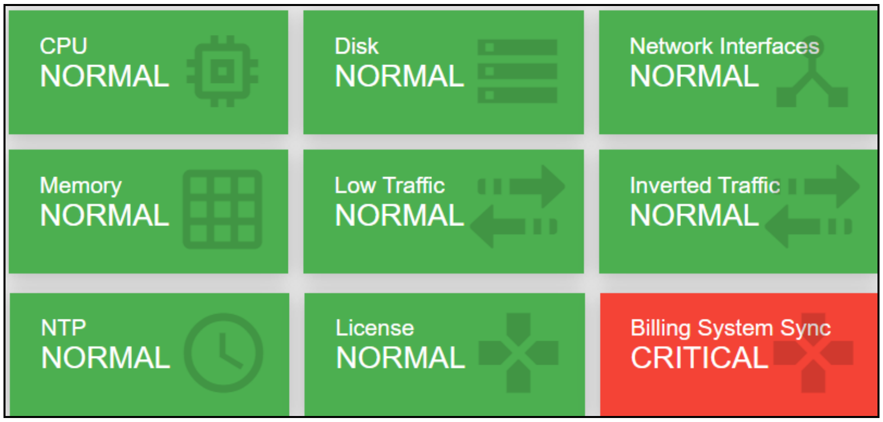

CPU
NORMAL

Disk
NORMAL

Network Interfaces
NORMAL

Memory
NORMAL

Low Traffic
NORMAL

Inverted Traffic
NORMAL

NTP
NORMAL

License
NORMAL

Billing System Sync
CRITICAL

Click on the **Billing** icon to view the billing status page.

Figure 155: *The Billing status page*

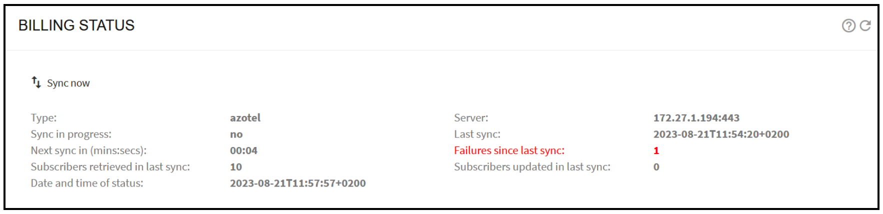

BILLING STATUS

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 pressing the **Sync now** button.

If the failure remains, follow these steps:

1. Check that the billing configuration is correct (direction and credentials).
2. Check 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#
   ```

3. Enable the API logs 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 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, you can see the event log between the QoE and the billing system. In the following example, the QoE connects to an Azotel billing are 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 should give some indication of what is happening.

4. 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#
```

The QoE generates traces of the last five requests 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

. . .
```

5. As a last resort, you can send a request directly from the 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:

```
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

If the NTP icon is yellow, the NTP is not configured. Click on the icon to view the configuration page.

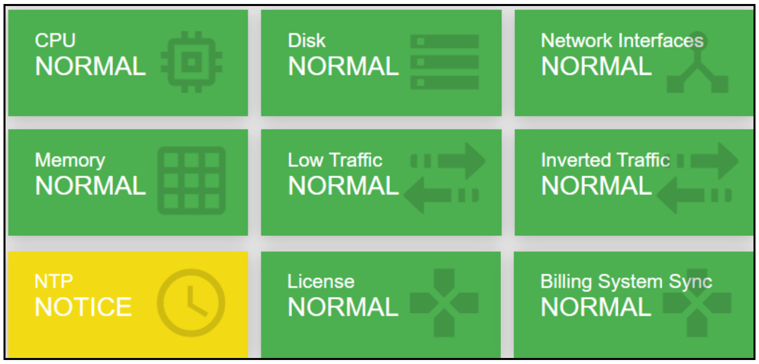

CPU NORMAL	Disk NORMAL	Network Interfaces NORMAL
Memory NORMAL	Low Traffic NORMAL	Inverted Traffic NORMAL
NTP NOTICE	License NORMAL	Billing System Sync NORMAL

If the NTP servers are configured, but the QoE cannot synchronize with any of them, the NTP icon in the dashboard will be in orange:

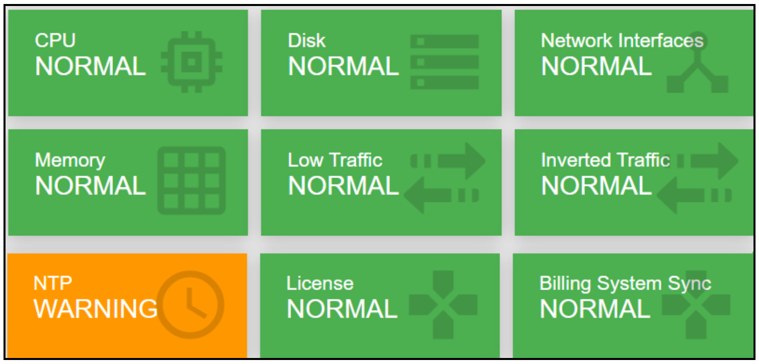

CPU NORMAL	Disk NORMAL	Network Interfaces NORMAL
Memory NORMAL	Low Traffic NORMAL	Inverted Traffic NORMAL
NTP WARNING	License NORMAL	Billing System Sync NORMAL

Note

An abrupt change in the system time because of lack of NTP synchronization. It may lead to brief service losses while the system adjusts. To avoid this from happening, have always at least one NTP server in sync.

To check the list of configured NTP servers, navigate to: **Administration** > **System Date** > **NTP Servers**.

Figure 156: *The NTP servers page*

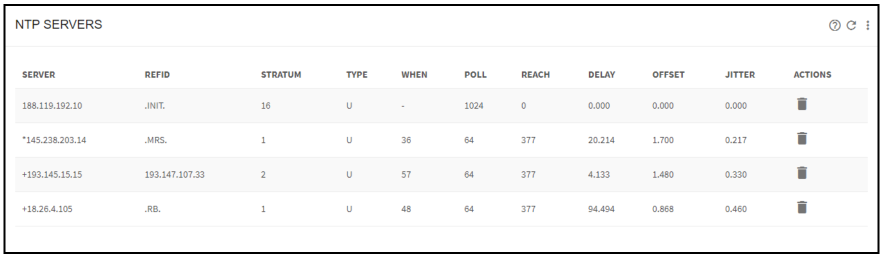

NTP SERVERS

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	🗑

At least one NTP server must be synchronized. In the example shown in Figure 156, the NTP server 145.238.203.14 is chosen for clock synchronization (indicated by the * next to the server IP address) and is contacted 36 seconds ago (column **WHEN**).

To solve the issue, if you have a local NTP server, add it to the list by clicking the ⋮ menu icon and selecting **Add Server...**

If you have no local NTP servers, make sure the UDP port 123 is open from the 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 on the **DISK** icon of the dashboard. Figure 157 shows the disk usage page.

Figure 157: *The disk usage page*

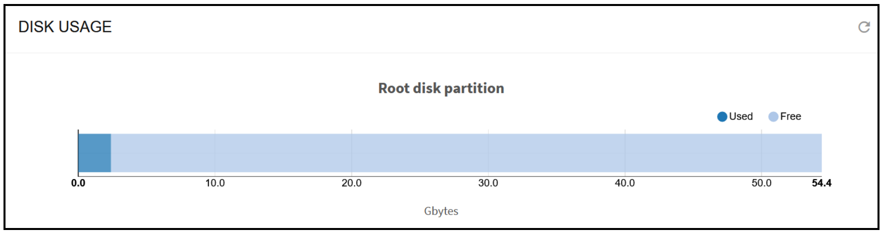

When less than 15% of the disk storage is free, the disk icon is in orange (warning state).

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