Interfaces User Guide

Version 23.8.1.2, 06 November 2023

Registered Address	Support	Sales
26, Kingston Terrace, Princeton, New Jersey 08540, United States		
		+91 80 4850 5445
http://www.rtbrick.com	support@rtbrick.com	sales@rtbrick.com

Table of Contents

1. Introduction

RtBrick Full Stack (RBFS) supports various types of interfaces, including physical and logical interfaces. On hardware platforms, RBFS physical interfaces represent the ports of a switch. This guide describes how to configure and verify RBFS interfaces. Features like routing protocols or access services will typically run on top of the interfaces.

1.1. Interface Types

Physical Interfaces

In RBFS, physical interfaces (IFP) typically represent the physical ports of a hardware switch. For example, ifp-0/0/1 represents switch port 1. On the physical interface level, you can configure various parameters associated with Layer 1 of the ISO/OSI reference model.

Logical Interfaces

For each physical interface, you can create one or multiple interface units also referred to as logical interfaces (IFL) in RBFS. A logical interface is associated with the Layer 2 operation. In addition, you can configure Layer 3 parameters like IP addresses on interface units, and assign interface units to routing instances.

Loopback Interfaces

A loopback interface is typically used to represent and identify a device itself. Loopback interfaces are preferred because they do not depend on the status of a physical port, and will always be up. Loopback interfaces are virtual interfaces, they are also represented as physical interfaces and interface units in RBFS, reflecting Layer 1 and Layer 2/3 operation.

Host Interfaces

Linux virtual ethernet (veth) interfaces connect an LXC container with the Linux host OS. In RBFS, a veth interface to the Linux bridge lxcbr0 is created by default. In virtual topologies, you can create additional veth interfaces and Linux bridges. RBFS host interfaces represent veth interfaces in RBFS.
For example, if the container interface eth1 connects to the host interface vethXYZ123, ifp-0/0/1 can be bound to eth1 to represent it in RBFS. Host interfaces can be used like any other physical interface.

Memory Interfaces

Memory interfaces (memif) are virtual interfaces used for creating virtual topologies. They connect multiple containers running RBFS to each other. When configuring memif interfaces:

- Endpoints match on the memif ID, i.e. the memif ID needs to be the same on both ends.

- memif IDs need to be unique on the host.
- The memif interface name is locally significant only.
- One endpoint needs to be configured as a master, while the other one is configured as a slave.

1.2. Interface Numbering

RBFS interface numbers match the port numbers on the switch faceplate. An interface is named in the ifp-<chassis-ID>/<front-panel-block-number>/<port> format. For example, ifp-0/0/1.

- Chassis ID—always 0 for the currently supported platforms
- Front Panel Block—represents group of ports on the faceplate
- Port—matches the port number on switch faceplate

Virtual interfaces follow the same structure, for example, lo-0/0/1 or memif-0/0/1.

Logical interfaces are numbered: ifl-<Node ID>/<Chip ID>/<Port ID>/<Unit ID>, for example ifl-0/0/1/1.

1.3. Community Support for Interfaces

You can tag an interface address with a community or extended community. RBFS will create a direct route for each interface address. If a community or extended community is configured for an interface address, RBFS will add it to the direct route. Communities can be used in policies. For example, when redistributing direct routes, you can match these communities and define desired policy rules.

1.4. Unnumbered Interfaces

An unnumbered interface is a point-to-point interface that is not explicitly configured with a dedicated IP address and subnet. Instead, it borrows (or links to) from a loopback interface, and uses it as the source IP address for packets originating from the interface. The unnumbered interface can "borrow" the IP address from another interface that is already configured on the switch, thereby conserving network and address space.

When configuring IP unnumbered on an interface:

- The IP address of the unnumbered interface cannot be borrowed.
- A logical interface can borrow IP address from a loopback interface, not vice versa.

1.5. Interface States

RBFS supports various interfaces such as physical interface, logical interface, and LAG interface and it uses various indicators to show the various states of interfaces. All interfaces have different states such as the following:

- **Admin state**: Indicates whether the interface is enabled (Up) or disabled (Down).
- **Link state**: Indicates whether the interface is linked (Up) or not linked (Down).
- **Oper state**: Indicates whether the interface is operational (Up) or not functional (Down).
- **IPv4 state**: Indicates that the interface is configured with the IPv4 address. The state can be Up or Down based on configuration.
- **IPv6 state**: Indicates that the interface is configured with the IPv6 address. The state can be Up or Down based on configuration.
- **MPLS state**: Indicates that the interface is an MPLS-enabled interface. The state can be Up or Down based on configuration.

Physical Interface States:

The physical interface states include:

- admin
- oper (operational)
- link

Logical Interface States:

The logical interface states include:

- admin
- oper (operational)
- link
- IPv4
- IPv6
- MPLS

LAG Interface States

The LAG interface states include:

- admin

- oper(operational)

The interface state can either be Up or Down. Up shows that it is ready to pass packets, and Down shows that it is not ready to transmit packets. The operational state shows that the interface is operational and is ready to transmit packets. Oper state is Up only if both the Admin state and Link state are Up.

1.6. Path MTU Discovery

Path MTU Discovery (PMTUD) is a mechanism for discovering the MTU (Maximum Transmission Unit) size of all routers on a network path. This mechanism finds the MTU size allowed for each device for a TCP/UDP connection path on the network between two hosts.

For the Path MTU Discovery to work, you should set a Don’t Fragment (DF) flag (for IPv4 packets) in the IP headers of outgoing packets. Any device whose MTU is smaller than the packet along the path cannot forward the packet to the destination, instead that sends an ICMP 'Fragmentation Needed' (Type 3, Code 4) message to the source router.

The IPv6 packets do not allow fragmentation in transit, so the IPv6 headers do not have the ‘Don’t Fragment’ option. So when a router on the network path receives IPv6 packet that exceeds the MTU, the routers drop the packets and send back the corresponding ICMP message to the source router.

The advantages of PMTUD include:

- Packets, which conform to the lowest MTU along the path are sent.
- Avoid fragmentation which adds latency.
- Avoid packet drops by intermediate devices due to the 'Don’t Fragment' flag while the host device continues sending packets.

In RBFS, the Path MTU Discovery feature is supported only on Q2C platform. It is not supported on Q2A platform.

To enable Path MTU Discovery, you must configure the MTU profile with the action redirect-to-cpu. So the CPU sends an ICMP message (that contains the allowed MTU of the router) to the source router. In this way, routers learn the MTU size of all routers on the path, so that they can determine the packet size in the network to avoid packet fragmentation.

For disabling the Path MTU Discovery, you can configure MTU profile with the action drop.

When an RBFS device receives MTU-violated packets, it responds with the ICMP error messages. ICMP error messages contain error codes that are sent to the source if fragmentation is required. The following are the types of ICMP error

messages for the IPv4 and IPv6 packets.

ICMPv4 Message Types

The type field identifies the type of the message sent by the host. The type field contains more specific information about the error condition.

The table below lists the ICMPv4 message types.

Type	Description
3	Destination Unreachable. This alerts a source host of delivery problems encountered while trying to reach the destination.

Destination Unreachable uses the following code values to further describe the function of the ICMP message being sent.

Code	Description
4	Fragmentation Needed and Don't-Fragment (DF) was Set. This message occurs when a router receives a packet that requires fragmentation, but the router has the DF flag turned on.

ICMPv6 Message Types

The type field identifies the type of the message sent by the host. The type field contains more specific information about the error condition.

The table below lists the ICMPv6 message type.

Type	Description
2	Packet Too Big. A Packet Too Big MUST be sent by a router in response to a packet that it cannot forward because the packet is larger than the MTU of the outgoing link.

Code	Description
0	No code

You can change this behavior by enabling fragmentation. For more information about enabling host path fragmentation, see the section "Enabling Hostpath Fragmentation".

All outgoing packets are validated against the defined MTU on the egress path.

- If MTU is violated and the MTU-profile action is drop, then packets are dropped.
- If MTU is violated and MTU-profile action is redirect-to-cpu, then packets are sent to the CPU port. A 20MB policer is used to protect the CPU port from overwhelming MTU-violated traffic. If fragmentation is enabled, one of the following operations takes place:
  - If the 'Don't Fragment' (DF) bit is not set in the received packet (only for IPv4), the packets are fragmented and sent to the outgoing port.
  - If the DF bit is set in a packet, it drops the packet and sends an ICMP error message back to the source.
- When fragmentation is disabled, packets are dropped and ICMP error messages are sent to the source.

For information about configuring the MTU profile, see MTU Profile Configuration.

1.7. IP Fragmentation

If the maximum transmission unit (MTU) of an outgoing interface is less than the original packet that needs to be routed, the packet needs to be fragmented. RBFS supports IP fragmentation on the Q2C platforms. The packets are sent to the CPU. CPU manages the fragmentation by sending an ICMP message (requesting packet fragmentation) to the source router.

If the packet needs to be fragmented and the Do-Not-Fragment (DF) flag is specified, then the device is going to send an ICMP Error code "Fragmentation Needed and DF set" to the source.

By default, IPv6 fragmentation is always handled at source as the IPv6 packets are not fragmented in transit. When the transit device receives an MTU-violated packet, it sends a "Packet Too Big" ICMPv6 message that it cannot forward the packet because it is larger than the MTU of the outgoing link.

1.7.1. Guidelines and Limitations of IP Fragmentation

The following guidelines and limitations apply to IP Fragmentation:

- You can control the fragmentation on the Q2C platform by configuring the set forwarding-options fragmentation ipv4 state CPU command. For more information about configuring fragmentation, see "2.2.2. Enabling Hostpath Fragmentation."
- There will be no ICMP error message sent in response to MTU-violated multicast packets.

1.8. MTU Profile

The Maximum Transmission Unit (MTU) is the size of the packet that is allowed in the network. MTU is configured as a MTU profile in the hardware. MTU profile configuration enables you to specify MTU size, MTU attachment point type, and MTU action. These profiles are attached to the interface entities (points) such as IFP, IFL, PPP, and L3 interfaces.

In the new generation silicon like Broadcom Qumran2C (Q2C), resources are conserved by creating profiles of the resources, and multiple entities like IFP, IFL and L3 interfaces utilize these profiles. To better manage MTU resources and platform capabilities, RBFS supports configuring MTU profiles and attaching these profiles to the attachment points.

Attachment Points: The MTU profiles are attached to the interface entities such as physical (IFP), logical (IFL), PPP, and L3 interfaces. RBFS supports the following attachment points for the MTU profiles:

- Port-level
- L3 interface level (IPv4 and IPv6)
- PPPoE subscriber level (L2 IFL)

MTU Size

A user-configured MTU size can range from 64 to 9216 in RBFS.

For MTU profiles of type "pppoe", users should provide L3 MTU size (IPv4/IPv6 headers).

MTU Type

An MTU type specifies the attachment point of the MTU profile. The MTU types supported are as follows:

- **physical**: When checking MTU, the entire packet size is considered.
- **ipv4**: MTU check is for IPv4 packets only.
- **ipv6**: MTU check is for IPv6 packets only.
- **ip**: MTU profile of type IP check for both IPv4 and IPv6 packets.
- **pppoe**: The MTU profile is applied to the PPPoE subscriber interface and the user is required to provide the L3 MTU size. Based on its best match algorithm, the Subscriber Management service associates these profiles with PPPoE subscribers.

- MTU profiles for L3 logical interfaces must be explicitly configured, and if not configured, no default MTU size is set for IPv4 and IPv6.
- On the Q2C platform, L3 interfaces can only be configured with IPv4 MTU profile or IPv6 MTU profile, but not both. However, with the type "ip" MTU profile, you can configure MTU for both IPv4 and IPv6 traffic with a common MTU size.

MTU Action

The MTU action is the action that a router takes when the MTU check fails. The redirect-to-cpu and drop are the two supported actions in RBFS supports.

Drop: If a router on the network path finds that the MTU size of a packet is larger than the configured MTU size of that router, it simply drops that packet.

redirect-to-cpu: If a router on the network path finds that the MTU size of a packet is larger than the configured MTU size of that router, it redirects that packet to CPU. CPU sends an ICMP error message requesting a packet fragmentation to the source router. This ICMP message also contains the MTU size of the router, so that the source router can learn the MTU size of the router on the path.

1.8.1. MTU Profile Limitations

The following limitations apply to the MTU profile:

- There is a limit to how many MTUs can be used by each hardware.
  - On the Q2C platform, the limit is as follows:
    - Maximum number of MTU profiles: 8
    - Maximum number of L3 MTU profiles: 3 (MTU type: IP/IPv4/IPv6)
    - Maximum number of PPPoE MTU profiles: 6 (including the default PPPoE profile)
    - Maximum number of physical MTU profiles: 7

1.9. Supported Platforms

Not all features are necessarily supported on each hardware platform. Refer to the Platform Guide for the features and the sub-features that are or are not supported by each platform.

1.10. Guidelines & Limitations

QAX-based Platforms

An additional restriction applies to ports on QAX-based platforms: because of hardware design, physical ports are grouped into quads (groups of 4, also known as port groups). Each quad must have the same physical parameters: speed, link-training, duplex.

The following tables are provided for easy identification of ports that need to have the same physical settings:

Edgecore 7316-26XB Port Groups:

Port	Speed	Duplex	Port Group
ifp-0/0/0	100G	Full	0
ifp-0/0/1	100G	Full	1
ifp-0/1/0	10G	Full	2
ifp-0/1/1	10G	Full	
ifp-0/1/2	10G	Full	
ifp-0/1/3	10G	Full	
ifp-0/1/4	10G	Full	3
ifp-0/1/5	10G	Full	
ifp-0/1/6	10G	Full	
ifp-0/1/7	10G	Full	
ifp-0/1/8	10G	Full	4
ifp-0/1/9	10G	Full	
ifp-0/1/10	10G	Full	
ifp-0/1/11	10G	Full	
ifp-0/1/12	10G	Full	5
ifp-0/1/13	10G	Full	
ifp-0/1/14	10G	Full	
ifp-0/1/15	10G	Full	
ifp-0/1/16	25G	Full	6
ifp-0/1/17	25G	Full	
ifp-0/1/18	25G	Full	
ifp-0/1/19	25G	Full	

Port	Speed	Duplex	Port Group
ifp-0/1/20	25G	Full	7
ifp-0/1/21	25G	Full	
ifp-0/1/22	25G	Full	
ifp-0/1/23	25G	Full	

UfiSpace S9500-22XST Port Groups:

Port	Speed	Duplex	Port Group
ifp-0/0/0	10G	Full	8
ifp-0/0/1	10G	Full	
ifp-0/0/2	10G	Full	
ifp-0/0/3	10G	Full	
ifp-0/0/4	10G	Full	4
ifp-0/0/5	10G	Full	
ifp-0/0/6	10G	Full	
ifp-0/0/7	10G	Full	
ifp-0/0/8	10G	Full	11
ifp-0/0/9	10G	Full	
ifp-0/0/10	10G	Full	
ifp-0/0/11	10G	Full	
ifp-0/0/12	25G	Full	3
ifp-0/0/13	25G	Full	
ifp-0/0/14	25G	Full	
ifp-0/0/15	25G	Full	
ifp-0/0/16	25G	Full	2
ifp-0/0/17	25G	Full	
ifp-0/0/18	25G	Full	
ifp-0/0/19	25G	Full	
ifp-0/0/20	100G	Full	0
ifp-0/0/21	100G	Full	1

A PHY Quad can be associated with network interface (NIF) ports of identical type only. For example, a quad cannot be a mix of XLGE and XE ports. An exception is GE and XE ports which can coexist in the same quad. This means all the ports in a port group should have the same physical interface configuration (that is, speed/duplex/link-training). Ports in a port group are only allowed to support 1G and 10G speeds; any other combination is not allowed. If a port within a port group is misconfigured, then it would require changing the speeds/interface type of all ports within the port group to a different type and then back into the original type.

2. Configuring Interfaces

2.1. Configuration Hierarchy

The diagram illustrates the interface configuration hierarchy.

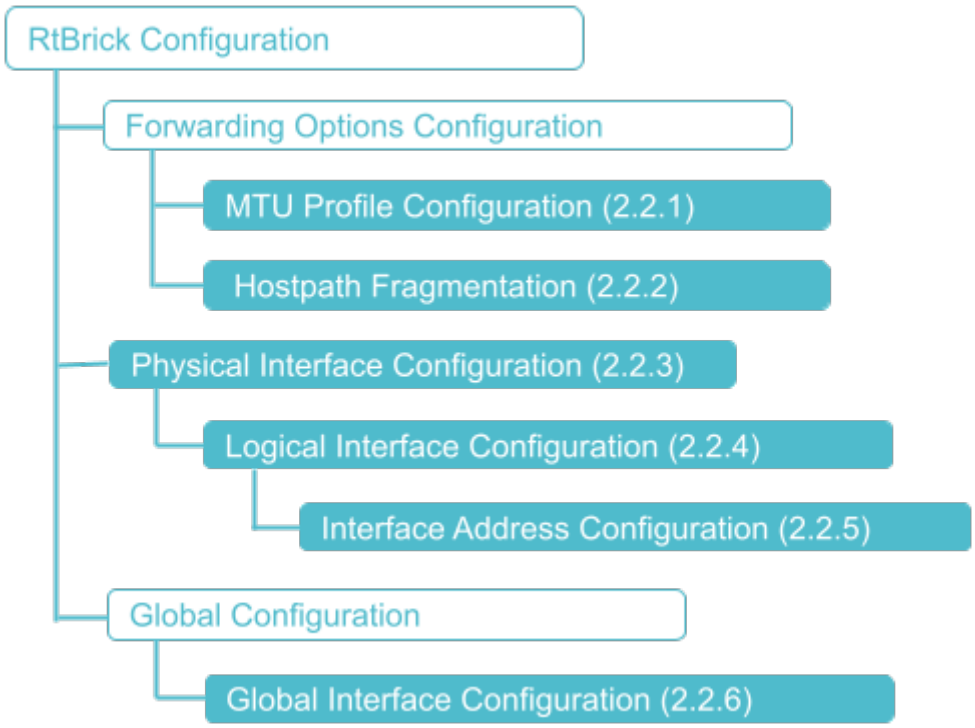

2.2. Configuration Syntax and Commands

The following sections describe the interface configuration syntax and commands.

2.2.1. MTU Profile Configuration

This section describes how to configure MTU profiles.

Syntax:

set forwarding-options mtu-profile <attribute> <value>

Attribute	Description
mtu-profile <mtu-profile-name>	MTU profile name
mtu-size <mtu-size>	MTU size. Range: 64 to 9216 bytes

Attribute	Description
mtu-type <mtu-type>	Specify the MTU type: • **physical**: Port based MTU profile • **pppoe**: subscriber IFL-based MTU profile for L2TP and PPPoE. This MTU profile is used by PPPoE subscribers to set the default MTU size of 1492. A configured size of 1492 bytes limits the size of the IPv4 or IPv6 header plus payload. • **ipv4**: MTU profile of type IPv4. Only IPv4 traffic on the logical interface will be impacted. • **ipv6**: MTU profile of type IPv6. Only IPv6 traffic on the logical interface will be impacted. • **ip**: MTU profile of type IP. Both IPv4 and IPv6 traffic on the logical interface will be impacted.
action <mtu-action>	Specify the MTU action. The following options are supported: **drop**: This indicates that when the MTU check fails, the action "drop" is performed. **redirect-to-cpu**: This is an action of redirecting packets to the CPU in a traffic behavior. A redirect-to-cpu action must be configured for fragmentation to occur.

Example 1: Configuration of the MTU Profile for the Physical Port

```
{
  "ietf-restconf:data": {
    "rtbrick-config:forwarding-options": {
      "mtu-profile": [
        {
          "mtu-profile-name": "portMtu",
          "size": 5000,
          "type": "physical",
          "action": "redirect-to-cpu"
        }
      ]
    }
  }
}
```

Example 2: MTU Profile Configuration of Type IPv4

```
{
  "ietf-restconf:data": {
    "rtbrick-config:forwarding-options": {
      "mtu-profile": [
        {
          "mtu-profile-name": "ipv4Mtu",
          "size": 1300,
          "type": "ipv4",
          "action": "redirect-to-cpu"
        }
      ]
    }
  }
}
```

Example 3: MTU Profile Configuration of Type IPv6

```
{
  "ietf-restconf:data": {
    "rtbrick-config:forwarding-options": {
      "mtu-profile": [
        {
          "mtu-profile-name": "ipv6Mtu",
          "size": 1400,
          "type": "ipv6",
          "action": "redirect-to-cpu"
        }
      ]
    }
  }
}
```

Example 4: Configuration of the MTU Profile for PPPoE

```
{
  "ietf-restconf:data": {
    "rtbrick-config:forwarding-options": {
      "mtu-profile": [
        {
          "mtu-profile-name": "pppoeMtu",
          "size": 1492,
          "type": "pppoe",
          "action": "redirect-to-cpu"
        }
      ]
    }
  }
}
```

2.2.2. Enabling Hostpath Fragmentation

This section describes how to enable or disable fragmentation by CPU. It is necessary to configure MTU profile action "redirect-to-cpu" so that fragmentation

takes place. By default, fragmentation is disabled.

To know about the guidelines and limitations before configuring, see the section '1.6.1. Guidelines and Limitations of IP Fragmentation".

Syntax:

set forwarding-options fragmentation ipv4 state <value>

Attribute	Description
disabled | cpu	Enables fragmentation of IPv4 packets. There are two options: disabled—Fragmentation is disabled. It is the default setting. cpu—Fragmentation is performed by CPU.

Example: Configuration of Hostpath Fragmentation

```
{
  "ietf-restconf:data": {
    "rtbrick-config:forwarding-options": {
      "fragmentation": {
        "ipv4": {
          "state": "cpu"
        }
      }
    }
  }
}
```

2.2.3. Physical Interface Configuration

This section describes configuration options at the physical interface (IFP) level.

Syntax:

set interface <interface-name> <attribute> <value>

Attribute	Description
<interface-name>	Name of the interface. Example: ifp-0/0/1.
admin-status <down|up>	Administrative state of the interface.
auto-negotiation <true|false>	Enable or disable auto-negotiation.
class-of-service <profile-name>	Apply class-of-service profile name.
description	Configure physical interface description.

Attribute	Description
host-if <container-interface>	Configure a host interface. For example, if the container interface eth1 connects to the host interface vethXYZ123, use this command option to bound hostif-0/0/1 to eth1. Please note the Linux virtual ethernet (veth) interface needs to be created separately. It cannot be created via RBFS configuration.
forward-error-correction <fec-type>	Configure Forward Error Correction (FEC) on the physical interface. FEC allows you to send the required information to correct errors through the link along with the payload data. A benefit of "forward" in FEC is that errors detected at the receiver do not need to be retransmitted. Currently, the supported FEC types are: base-r, rsfec, none. NOTE: rsfec is the only FEC supported for 100G on the QAX platform.
link-training <true|false>	Enable or disable link training.
master <true|false>	Memif role, master or slave, applicable only to memif interface. One end needs to be configured as master, and the other one as slave.
memif-id <id>	Configure memif ID , applicable only to memif interface. Needs to match on both ends.
mtu-profile <mtu-profile-name>	Attach MTU profile to a physical interface. This is a mandatory attribute.
mru <size>	Maximum receive unit size on the physical interface.
speed <speed>	Configure speed mode for the interface. Port speed refers to the maximum amount of data transmitted. The speed value is specified in Gigabits per second (Gbps). Currently, RBFS supports 10G and 100G ports, and you can make the following changes: • 100G port speed can be changed to 40G • 10G port speed can be changed to 1G

Example 1: Physical Interface Configuration

```
{
    "rtbrick-config:interface": [
      {
        "name": "ifp-0/0/1",
        "description": "Link to leaf1",
        "speed": "10G",
        "mtu-profile": "portMtu",
        "mru": 5000
      }
    ]
}
```

Example 2: Memory Interface Configuration

A End:

```
{
    "rtbrick-config:interface": [
      {
        "name": "ifp-0/0/1",
        "description": "Master",
        "memif-id": 11,
        "master": "true",
      }
    ]
}
```

B End:

```
{
    "rtbrick-config:interface": [
      {
        "name": "ifp-0/0/1",
        "description": "Slave",
        "memif-id": 11,
        "master": "false",
      }
    ]
  }
```

Example 3: Host Interface Configuration

```
{
    "rtbrick-config:interface": [
        "name": "ifp-0/0/1",
        "description": "Represents eth1 as ifp-0/0/1 in RBFS",
        "host-if": "eth1",
    ]
}
```

Example 4: MRU Configuration for Physical Interface

```
{
    "rtbrick-config:interface": [
      {
        "name": "ifp-0/0/7",
        "mru": 5000
      }
    ]
  }
```

Example 5: FEC Configuration for Physical Interface

```
{
    "rtbrick-config:interface": [
      {
        "name": "ifp-0/0/40",
        "forward-error-correction": "base-r"
      }
    ]
  }
```

2.2.4. Logical Interface Configuration

This section describes configuration options at the logical interface (IFL) level.

Syntax:

set interface <interface-name> **unit** <unit-id> <attribute> <value>

Attribute	Description
unit <unit-id>	Create a logical interface (also referred to as a sub-interface) under the physical interface.
admin-status <down|up>	Administrative state of the logical interface.
class-of-service <profile-name>	Apply class-of-service profile name.
description <description>	Description of the logical interface.
inner-vlan <inner-vlan-id>	Inner VLAN ID.
instance <instance>	Assign the logical interface to an instance.
ipv4-admin-status <down|up>	Enable or disable IPv4.
ipv4-mtu-profile <ipv4-mtu-profile>	Attach IPv4 MTU profile to an L3 interface.
ipv6-admin-status <down|up>	Enable or disable IPv6.

Attribute	Description
ipv6-mtu-profile <ipv6-mtu-profile>	Attach IPv6 MTU profile to an L3 interface.
ip-mtu-profile <ip-mtu-profile>	Attach IP MTU profile to an L3 interface.
mpls-admin-status <down|up>	Enable or disable MPLS.
mpls-mtu <mpls-mtu-size>	MPLS maximum transmission unit size.
neighbor <ipv4|ipv6> <ip-address> mac <mac-address>	Configure a static IPv4 or IPv6 neighbor.
unnumbered interface <loopback-interface-name>	Configure an un-numbered interface.
vlan <outer-vlan-id>	Outer VLAN ID.

Example 1: Logical Interface Configuration with IPv4 MTU Profile

```
{
  "rtbrick-config:interface": [
     {
       "name": "ifp-0/0/1",
       "unit": [
         {
           "unit-id": 1,
           "description": "VLAN 101",
           "instance": "default",
           "ipv4-mtu-profile": "ipv4Mtu"

         }
       ]
     }
  ]
}
```

Example 2: Logical Interface Configuration with IPv6 MTU Profile

```
{
  "rtbrick-config:interface": [
     {
       "name": "ifp-0/0/1",
       "unit": [
         {
           "unit-id": 1,
           "description": "VLAN 101",
           "instance": "default",
           "ipv6-mtu-profile": "ipv6Mtu"

         }
       ]
     }
  ]
}
```

Example 3: Logical Interface Configuration with IP MTU Profile

```
{
  "rtbrick-config:interface": [
     {
       "name": "ifp-0/0/1",
       "unit": [
         {
           "unit-id": 1,
           "description": "VLAN 101",
           "instance": "default",
           "ip-mtu-profile": "ipMtu"

         }
       ]
     }
  ]
}
```

2.2.5. Interface Address Configuration

This section describes how to configure interface IP addresses.

Syntax:

set interface <interface-name> **unit** <unit-id> **address** <afi> <attribute> <value>

Attribute	Description
<afi>	Address family identifier (AFI). Supported values: ipv4 and ipv6
<prefix4|prefix6>	Assign IPv4 or IPv6 address to the interface unit.
community <community-value>	Configure list of communities associated with the address.

Attribute	Description
extended-community <community-value>	Configure list of extended communities associated with the address.
label <label-value>	Configure label associated with the address. Supported MPLS label values are 0 - 1048575. The reserved MPLS label range is 0 - 15. In RBFS, BGP uses the label range 20000 - 100000. It is recommended to assign label values outside of these reserved ranges to avoid conflicts.
secondary <true|false>	Configure a secondary IPv4 address.

Example: Interface Address Configuration

```
{
    "rtbrick-config:interface": [
      {
        "name": "lo-0/0/1",
        "unit": [
          {
            "unit-id": 1,
            "address": {
              "ipv4": [
                {
                  "prefix4": "198.51.100.103/24",
                  "label": 12346
                }
              ]
            }
          }
        ]
      }
    ]
  }
```

2.2.6. Global Interface Configuration

This section describes a configuration option applied globally to all interfaces.

Syntax:

set global interface all <attribute> <value>

Attribute	Description
admin-status <up|down>	Configure state of the interface.

- The interface level enable/disable command has higher precedence than the global interface enable/disable command.
- You can disable all unused physical interfaces.
- Before executing the global interface disable all command ensure that all physical interfaces are in the link Up state.

Example: Enabling or Disabling all Interfaces

```
{
  "ietf-restconf:data": {
    "rtbrick-config:global": {
      "interface": {
        "all": {
          "admin-status": "down"
        }
      }
    }
  }
}
```

3. Operational Commands

3.1. Interface Show Commands

The interface show commands provide detailed information about the status and parameters of RBFS interfaces.

3.1.1. Interface Summary Commands

Syntax:

show interface <option>

Option	Description
summary	Displays a summary of all interfaces including physical, logical, and address information.
<interface-name>	Displays a summary of an interface including physical, logical, and address information.
physical	Displays all physical interface including loopback, cpu and recycle ports.
logical	Displays all logical interfaces for all instances.
logical <instance-name>	Displays all logical interfaces for the given instance.
address	Displays all IPv4 and IPv6 addresses for all instances.
address <instance-name>	Displays all IPv4 and IPv6 addresses for the given instance.

Example 1: Summary Output for All Interfaces

```
supervisor@rtbrick>LEAF01: op> show interface summary
Interface                Admin    Link     Oper          IPv4 Address             IPv6 Address
ifp-0/0/1                Up       Down     Down
ifp-0/0/2                Up       Down     Down
ifp-0/0/3                Up       Down     Down
ifp-0/0/4                Up       Up       Up
ifp-0/0/5                Up       Down     Down
ifp-0/0/6                Up       Down     Down
ifp-0/0/7                Up       Down     Down
ifp-0/0/8                Up       Down     Down
ifp-0/0/9                Up       Down     Down
ifp-0/0/10               Up       Up       Up
ifl-0/0/10/100           Up       Up       Up            198.51.100.22/24         2001:db8:0:100::/32
ifl-0/0/10/200           Up       Up       Up            198.51.100.32/24         2001:db8:0:10::/32
ifl-0/0/10/300           Up       Up       Up            -                        2001:db8:0:160::/32
ifp-0/0/11               Up       Down     Down
ifp-0/0/12               Up       Down     Down
ifp-0/0/13               Up       Down     Down
ifp-0/0/14               Up       Down     Down
ifp-0/0/15               Up       Down     Down
ifp-0/0/16               Up       Down     Down
ifp-0/0/17               Up       Down     Down
ifp-0/0/18               Up       Down     Down
ifp-0/0/19               Up       Down     Down
ifp-0/0/20               Up       Down     Down
ifp-0/0/21               Up       Down     Down
ifp-0/0/22               Up       Down     Down
ifp-0/0/23               Up       Down     Down
ifp-0/0/24               Up       Down     Down
ifp-0/0/25               Up       Down     Down
ifp-0/0/26               Up       Down     Down
ifp-0/0/27               Up       Up       Up
ifp-0/0/28               Up       Down     Down
ifp-0/0/29               Up       Down     Down
ifp-0/0/30               Up       Down     Down
ifp-0/0/31               Up       Down     Down
ifp-0/0/32               Up       Down     Down
ifp-0/0/33               Up       Down     Down
ifp-0/0/34               Up       Down     Down
ifp-0/0/35               Up       Down     Down
ifp-0/0/36               Up       Down     Down
ifp-0/0/37               Up       Down     Down
ifp-0/0/38               Up       Down     Down
ifp-0/0/39               Up       Down     Down
ifp-0/0/40               Up       Down     Down
ifp-0/0/41               Up       Down     Down
ifp-0/0/42               Up       Down     Down
ifp-0/0/43               Up       Down     Down
ifp-0/0/44               Up       Down     Down
ifp-0/0/45               Up       Down     Down
ifp-0/0/46               Up       Down     Down
ifp-0/0/47               Up       Down     Down
ifp-0/0/48               Up       Down     Down
ifp-0/0/49               Up       Down     Down
ifp-0/0/50               Up       Down     Down
ifp-0/0/51               Up       Down     Down
ifp-0/0/52               Up       Up       Up
ifp-0/0/53               Up       Up       Up
ifp-0/0/54               Up       Down     Down
cpu-0/0/200              Up       Up       Up
cpu-0/0/201              Up       Down     Down
cpu-0/0/202              Up       Down     Down
cpu-0/0/203              Up       Down     Down
recycle-0/0/75           Up       Up       Up
recycle-0/0/75/0         Up       Up       Up
recycle-0/0/76           Up       Up       Up
recycle-0/0/76/0         Up       Up       Up
```

Example 2: Summary Output for One Physical Interface

```
supervisor@rtbrick>LEAF01: op> show interface ifp-0/0/10
Interface                Admin    Link     Oper          IPv4 Address             IPv6 Address
ifp-0/0/10               Up       Up       Up
  ifl-0/0/10/100         Up       Up       Up            198.51.100.22/24               2001:db8:0:100::/32
  ifl-0/0/10/200         Up       Up       Up            198.51.100.32/24               2001:db8:0:10::/32
  ifl-0/0/10/300         Up       Up       Up            -                        2001:db8:0:160::/32
  ifl-0/0/10/1000        Up       Up       Up            -                        2001:db8:0:33::/32
```

Example 3: List of All Physical Interfaces

```
supervisor@rtbrick>LEAF01: op> show interface physical
Interface        Admin   Link    Oper    MAC Address          Speed  Duplex  Uptime
lo-0/0/1         Up      Up      Up      80:a2:35:a0:00:01    -      -       Thu Nov 19 10:41:06 GMT +0000
2020
ifp-0/0/1        Up      Down    Down    80:a2:35:ee:a8:01    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/2        Up      Down    Down    80:a2:35:ee:a8:02    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/3        Up      Down    Down    80:a2:35:ee:a8:03    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/4        Up      Up      Up      80:a2:35:ee:a8:04    10G    Full    Thu Nov 19 10:05:02 GMT +0000
2020
ifp-0/0/5        Up      Down    Down    80:a2:35:ee:a8:05    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/6        Up      Down    Down    80:a2:35:ee:a8:06    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/7        Up      Down    Down    80:a2:35:ee:a8:07    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/8        Up      Down    Down    80:a2:35:ee:a8:08    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/9        Up      Down    Down    80:a2:35:ee:a8:09    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/10       Up      Up      Up      80:a2:35:ee:a8:0a    10G    Full    Fri Nov 20 00:59:12 GMT +0000
2020
ifp-0/0/11       Up      Down    Down    80:a2:35:ee:a8:0b    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/12       Up      Down    Down    80:a2:35:ee:a8:0c    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/13       Up      Down    Down    80:a2:35:ee:a8:0d    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/14       Up      Down    Down    80:a2:35:ee:a8:0e    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/15       Up      Down    Down    80:a2:35:ee:a8:0f    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/16       Up      Down    Down    80:a2:35:ee:a8:10    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/17       Up      Down    Down    80:a2:35:ee:a8:11    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/18       Up      Down    Down    80:a2:35:ee:a8:12    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/19       Up      Down    Down    80:a2:35:ee:a8:13    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/20       Up      Down    Down    80:a2:35:ee:a8:14    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/21       Up      Down    Down    80:a2:35:ee:a8:15    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/22       Up      Down    Down    80:a2:35:ee:a8:16    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/23       Up      Down    Down    80:a2:35:ee:a8:17    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/24       Up      Down    Down    80:a2:35:ee:a8:18    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/25       Up      Down    Down    80:a2:35:ee:a8:19    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/26       Up      Down    Down    80:a2:35:ee:a8:1a    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/27       Up      Up      Up      80:a2:35:ee:a8:1b    10G    Full    Fri Nov 20 00:59:11 GMT +0000
2020
ifp-0/0/28       Up      Down    Down    80:a2:35:ee:a8:1c    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/29       Up      Down    Down    80:a2:35:ee:a8:1d    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
```

```
ifp-0/0/30      Up      Down    Down    80:a2:35:ee:a8:1e    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/31      Up      Down    Down    80:a2:35:ee:a8:1f    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/32      Up      Down    Down    80:a2:35:ee:a8:20    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/33      Up      Down    Down    80:a2:35:ee:a8:21    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/34      Up      Down    Down    80:a2:35:ee:a8:22    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/35      Up      Down    Down    80:a2:35:ee:a8:23    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/36      Up      Down    Down    80:a2:35:ee:a8:24    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/37      Up      Down    Down    80:a2:35:ee:a8:25    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/38      Up      Down    Down    80:a2:35:ee:a8:26    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/39      Up      Down    Down    80:a2:35:ee:a8:27    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/40      Up      Down    Down    80:a2:35:ee:a8:28    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/41      Up      Down    Down    80:a2:35:ee:a8:29    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/42      Up      Down    Down    80:a2:35:ee:a8:2a    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/43      Up      Down    Down    80:a2:35:ee:a8:2b    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/44      Up      Down    Down    80:a2:35:ee:a8:2c    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/45      Up      Down    Down    80:a2:35:ee:a8:2d    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/46      Up      Down    Down    80:a2:35:ee:a8:2e    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/47      Up      Down    Down    80:a2:35:ee:a8:2f    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/48      Up      Down    Down    80:a2:35:ee:a8:30    10G    Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/49      Up      Down    Down    80:a2:35:ee:a8:31    100G   Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/50      Up      Down    Down    80:a2:35:ee:a8:35    100G   Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/51      Up      Down    Down    80:a2:35:ee:a8:39    100G   Full    Mon Nov 16 11:24:09 GMT +0000
2020
ifp-0/0/52      Up      Up      Up      80:a2:35:ee:a8:3d    100G   Full    Tue Nov 17 14:10:46 GMT +0000
2020
ifp-0/0/53      Up      Up      Up      80:a2:35:ee:a8:41    100G   Full    Fri Nov 20 00:59:12 GMT +0000
2020
ifp-0/0/54      Up      Down    Down    80:a2:35:ee:a8:45    100G   Full    Mon Nov 16 11:24:09 GMT +0000
2020
cpu-0/0/200     Up      Up      Up      80:a2:35:ee:a8:c8    100G   Full    Mon Nov 16 11:24:11 GMT +0000
2020
cpu-0/0/201     Up      Down    Down    80:a2:35:ee:a8:c9    100G   Full    Mon Nov 16 11:24:09 GMT +0000
2020
cpu-0/0/202     Up      Down    Down    80:a2:35:ee:a8:ca    100G   Full    Mon Nov 16 11:24:09 GMT +0000
2020
cpu-0/0/203     Up      Down    Down    80:a2:35:ee:a8:cb    100G   Full    Mon Nov 16 11:24:09 GMT +0000
2020
recycle-0/0/75  Up      Up      Up      80:a2:35:ee:a8:4b    100G   Full    Mon Nov 16 11:24:11 GMT +0000
2020
recycle-0/0/76  Up      Up      Up      80:a2:35:ee:a8:4c    100G   Full    Mon Nov 16 11:24:11 GMT +0000
2020
```

Example 4: List of All Logical Interfaces for All Instances

```
supervisor@rtbrick>LEAF01: op> show interface logical
Interface              Instance            Admin   Link    Oper    Outer VLAN   Inner VLAN  IPv4 Status,MTU
IPv6 Status,MTU
ifl-0/0/10/100         default             Up      Up      Up      -            -           Up,1500
Up,1500
ifl-0/0/10/200         default             Up      Up      Up      200          -           Up,1500
Up,1500
ifl-0/0/10/300         default             Up      Up      Up      300          -           Up,1500
Up,1500
```

Example 5: List of Logical Interfaces for an Instance

```
supervisor@rtbrick: op> show interface logical instance default
Interface              Instance            Admin   Link    Oper    Outer VLAN   Inner VLAN  IPv4 Status,MTU
IPv6 Status,MTU
ifl-0/0/10/100         default             Up      Up      Up      -            -           Up,1500
Up,1500
ifl-0/0/10/200         default             Up      Up      Up      200          -           Up,1500
Up,1500
ifl-0/0/10/300         default             Up      Up      Up      300          -           Up,1500
Up,1500
```

Example 6: List of All Interface Addresses

```
supervisor@rtbrick: op> show interface address
Interface              Instance         IPv4 Address        IPv4 Primary   IPv6 Address
  ifl-0/0/10/100       default        198.51.100.22/24      True           2001:db8:0:100::/32
  ifl-0/0/10/200       default        198.51.100.32/24       True          2001:db8:0:10::/32
  ifl-0/0/10/300       default              -                               2001:db8:0:160::/32
```

3.1.2. Interface Details Commands

Syntax:

show interface <option> **detail**

Option	Description
detail	Without any additional option, displays detailed information for all interfaces.
<interface-name> detail	Displays detailed information for an interface.

Example 7: Detailed Information for a Physical Interface and Its Logical Interfaces

```
supervisor@rtbrick: op> show interface ifp-0/0/10 detail
Interface:ifp-0/0/52
Admin/Link/Operational status: Up/Up/Up
Speed configured: 100G
Speed maximum: 100G
Duplex: Full
Autonegotiation: Disabled
Encapsulation mode: ieee
MRU: 16360
MTU Profile: portMtu
Maximum frame size: 16360
Interface type: ethernet
Interface index: 124929
MAC: 80:a2:35:ee:a8:3d
Uptime: Tue Nov 17 14:10:46 GMT +0000 2020
Flap count: 2
Description: Physical interface #52 from node 0, chip 0
Packet statistics:
Rx packets: 16034 Tx packets: 17295
Rx bytes: 1602124 Tx bytes: 1711264
Interface:ifl-0/0/52/4, Instance:default
Admin/Link/Operational status: Up/Up/Up
IPv4/IPv6/MPLS Status: Up/Up/Up
IPv4/IPv6/MPLS MTU: 1500/1500/1500
Interface type: Logical Sub interface
Interface index: 106497
MAC: 80:a2:35:ee:a8:3d

IPv4 Address IPv6 Address
198.51.100.44/24 2001:db8:0:222::/32

Packet statistics:
Ingress forwarded packets: 16000
Ingress forwarded bytes: 1598656
Ingress drop Packets: 0
Ingress drop bytes: 0
Egress forwarded packets: 0
Egress forwarded bytes: 0
Egress drop packets: 0
Egress drop bytes: 0
```

3.2. MTU Profile Show Command

Syntax:

show mtu profile <option>

Option	Description
-	Without any additional option, displays detailed information for all MTU profiles.
profile-name <mtu-profile-name>	MTU Profile Name

Example 8: Detailed Information About the MTU Profiles

```
supervisor@rtbrick>LEAF01: op> show mtu profile
Profile Name                Type         Size   Action
__default_pppoe__           pppoe        1492   drop
l3IpMtu                     ipv4         1300   drop
l3Ipv6Mtu                   ipv6         1300   drop
portMtu                     physical     1300   drop
portM2                      physical     1400   drop
portM5                      physical     1430   drop
supervisor@rtbrick>LEAF01: op>
```

Example 9: Display Information About the Specified MTU Profile

```
supervisor@rtbrick>LEAF01: op> show mtu profile profile-name l3IpMtu
Profile Name                Type         Size   Action
l3IpMtu                     ipv4         1300   drop
supervisor@rtbrick>LEAF01: op>
```

3.2.1. Interface Statistics Commands

Syntax:

show interface <option> **statistics**

Option	Description
statistics	Without any additional option, displays statistics information for all interfaces.
<interface-name> statistics	Displays statistics information for an interface.

Example 10: Statistics Information for a Physical Interface and Its Logical Interfaces

```
supervisor@rtbrick>LEAF01: op> show interface ifp-0/0/10 statistics
Interface:  ifp-0/0/10
   Counter              Direction   Unit       Rx         Rx Diff    Rx Rate    Tx         Tx Diff   Tx Rate
   IPv4                 -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   IPv6                 -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   MPLS                 -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Punt                 -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Miss                 RX          Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Drops                -           Packets    4995       -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Error                RX          Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Error                TX          Packets    47         -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   No Buff              RX          Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Traffic Statistics   -           Packets    4995       -          -          68492      -          -
                                    Bytes      489510     -          -          5869876    -          -
   Unicast Statistics   -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
```

```
    Broadcast Statistics  -           Packets    -          -          -          -          -          -
                                      Bytes      -          -          -          -          -          -
    Multicast Statistics  -           Packets    -          -          -          -          -          -
                                      Bytes      -          -          -          -          -          -
    Bcm Statistics:
    inOctets:                                 511632
    inUcastPkts:                              0
    inNonUcastPkts:                           5016
    inErrors:                                 0
    inUnknownProtos:                          0
    outOctets:                                6236484
    outUcastPkts:                             0
    outNonUcastPkts:                          68492
    outErrors:                                0
    etherStatsDropEvents:                     0
    etherStatsMulticastPkts:                  67718
    etherStatsBroadcastPkts:                  5790
    etherStatsUndersizePkts:                  0
    etherStatsFragments:                      0
    etherStatsOversizePkts:                   0
    etherStatsOctets:                         6748116
    etherStatsPkts:                           73508
    etherStatsCollisions:                     0
    etherStatsTXNoErrors:                     68492
    etherStatsRXNoErrors:                     5016
    ifInMulticastPkts:                        5016
    ifOutBroadcastPkts:                       5790
    ifOutMulticastPkts:                       62702
    ifOutBroadcastPkts:                       5790
    bcmReceivedUndersizePkts:                 0
    bcmTransmittedUndersizePkts:              5790
    bcmQmxDot1dBasePortDelayExceededDiscards: 0
    bcmQmxDot1dBasePortMtuExceededDiscards:   0
    bcmQmxDot1dTpPortInFrames:                5016
    bcmQmxDot1dTpPortOutFrames:               68492
    bcmQmxEtherStatsPkts64Octets:             5790
    bcmQmxEtherStatsPkts128to255Octets:       24
    bcmQmxEtherStatsPkts256to511Octets:       0
    bcmQmxEtherStatsPkts512to1023Octets:      0
    bcmQmxEtherStatsPkts1024to1518Octets:     0
    bcmQmxEtherRxOversizePkts:                0
    bcmQmxEtherTxOversizePkts:                0
    bcmQmxEtherStatsJabbers:                  0
    bcmQmxEtherStatsCRCAlignErrors:           0
    bcmQmxDot3StatsFCSErrors:                 0
    bcmQmxDot3StatsSingleCollisionFrames:     0
    bcmQmxDot3StatsMultipleCollisionFrames:   0
    bcmQmxDot3StatsSQETTestErrors:            0
    bcmQmxDot3StatsDeferredTransmissions:     0
    bcmQmxDot3StatsLateCollisions:            0
    bcmQmxDot3StatsExcessiveCollisions:       0
    bcmQmxDot3StatsInternalMacTransmitErrors: 0
    bcmQmxDot3StatsCarrierSenseErrors:        0
    bcmQmxDot3StatsFrameTooLongs:             0
    bcmQmxDot3StatsInternalMacReceiveErrors:  0
    bcmQmxDot3StatsSymbolErrors:              0
    bcmQmxDot3ControlInUnknownOpcodes:        0
    bcmQmxDot3InPauseFrames:                  0
    bcmQmxDot3OutPauseFrames:                 0
    bcmQmxIfHCInOctets:                       511632
    bcmQmxIfHCInUcastPkts:                    0
    bcmQmxIfHCInMulticastPkts:                5016
    bcmQmxIfHCInBroadcastPkts:                0
    bcmQmxIfHCOutOctets:                      6236484
    bcmQmxIfHCOutUcastPkts:                   0
    bcmQmxIfHCOutMulticastPkts:               62702
    bcmQmxIfHCOutBroadcastPckts:              5790
    bcmQmxIeee8021PfcRequests:                0
    bcmQmxIeee8021PfcIndications:             0
    bcmQmxBcmEtherStatsPkts1519to1522Octets:  0
    bcmQmxBcmEtherStatsPkts1522to2047Octets:  0
    bcmQmxBcmReceivedPkts64Octets:            0
    bcmQmxBcmReceivedPkts65to127Octets:       5016
    bcmQmxBcmReceivedPkts128to255Octets:      0
    bcmQmxBcmReceivedPkts256to511Octets:      0
    bcmQmxBcmReceivedPkts512to1023Octets:     0
```

```
   bcmQmxBcmReceivedPkts1024to1518Octets:    0
   bcmQmxBcmReceivedPkts1519to2047Octets:    0
   bcmQmxBcmTransmittedPkts64Octets:         5790
   bcmQmxBcmTransmittedPkts65to127Octets:    62678
   bcmQmxBcmTransmittedPkts128to255Octets:   24
   bcmQmxBcmTransmittedPkts256to511Octets:   0
   bcmQmxBcmTransmittedPkts512to1023Octets:  0
   bcmQmxBcmTransmittedPkts1024to1518Octets: 0
   bcmQmxBcmTransmittedPkts1519to2047Octets: 0
   bcmQmxBcmTransmittedPkts2048to4095Octets: 0
   bcmQmxBcmTransmittedPkts4095to9216Octets: 0

Logical Interface:  ifl-0/0/10/100, Physical Interface:  ifp-0/0/10
   Counter               Direction   Unit       Rx         Rx Diff    Rx Rate    Tx         Tx Diff   Tx Rate
   IPv4                  -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   IPv6                  -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   MPLS                  -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Punt                  -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Miss                  RX          Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Drops                 -           Packets    4995       -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Error                 RX          Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Error                 TX          Packets    47         -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   No Buff               RX          Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Traffic Statistics    -           Packets    4995       -          -          68492      -          -
                                     Bytes      489510     -          -          5869876    -          -
   Unicast Statistics    -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Broadcast Statistics  -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Multicast Statistics  -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Packet Statistics:
     Ingress Forwarded Packets: 1810
     Ingress Forwarded Bytes:   184620
     Ingress Drop Packets:      1
     Ingress Drop Bytes:        102
     Egress Forwarded Packets:  0
     Egress Forwarded Bytes:    0
     Egress Drop Packets:       0
     Egress Drop Bytes:         0
Logical Interface:  ifl-0/0/10/200, Physical Interface:  ifp-0/0/10
   Counter               Direction   Unit       Rx         Rx Diff    Rx Rate    Tx         Tx Diff   Tx Rate
   IPv4                  -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   IPv6                  -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   MPLS                  -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Punt                  -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Miss                  RX          Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Drops                 -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Error                 RX          Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Error                 TX          Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   No Buff               RX          Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Traffic Statistics    -           Packets    -          -          -          6811       -          -
                                     Bytes      -          -          -          573170     -          -
   Unicast Statistics    -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Broadcast Statistics  -           Packets    -          -          -          -          -          -
                                     Bytes      -          -          -          -          -          -
   Multicast Statistics  -           Packets    -          -          -          -          -          -
```

```
                                Bytes      -          -          -          -          -          -
    Packet Statistics:
      Ingress Forwarded Packets: 0
      Ingress Forwarded Bytes:   0
      Ingress Drop Packets:      0
      Ingress Drop Bytes:        0
      Egress Forwarded Packets:  0
      Egress Forwarded Bytes:    0
      Egress Drop Packets:       0
      Egress Drop Bytes:         0
Logical Interface:  ifl-0/0/10/300, Physical Interface:  ifp-0/0/10
   Counter              Direction   Unit       Rx         Rx Diff    Rx Rate    Tx         Tx Diff    Tx Rate
   IPv4                 -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   IPv6                 -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   MPLS                 -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Punt                 -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Miss                 RX          Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Drops                -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Error                RX          Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Error                TX          Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   No Buff              RX          Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Traffic Statistics   -           Packets    -          -          -          5902       -          -
                                    Bytes      -          -          -          531180     -          -
   Unicast Statistics   -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Broadcast Statistics -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Multicast Statistics -           Packets    -          -          -          -          -          -
                                    Bytes      -          -          -          -          -          -
   Packet Statistics:
     Ingress Forwarded Packets: 0
     Ingress Forwarded Bytes:   0
     Ingress Drop Packets:      0
     Ingress Drop Bytes:        0
     Egress Forwarded Packets:  0
     Egress Forwarded Bytes:    0
     Egress Drop Packets:       0
     Egress Drop Bytes:         0
supervisor@rtbrick>LEAF01: op>
```

3.3. Interface Clear Commands

Clear commands allow to reset operational states.

3.3.1. Interface Statistics

This command clears interface counters.

Syntax:

clear interface statistics <option>

Option	Description
-	Without any additional option, the command clears the counters for all interfaces.

Option	Description
<interface-name>	Clears the counters for the given interface.