



Technical Documentation

RBFS Overview and Platform

Version 25.4.1.4, 20 March 2026

Table of Contents

1. RBFS Overview	1
1.1. RBFS At-a-Glance	1
1.1.1. Modular and Resilient Architecture	1
1.1.2. Automation and Management Interfaces	1
1.1.3. Comprehensive Feature Set	2
1.1.4. Hardware Flexibility and Ecosystem	2
1.2. Why RBFS	3
1.3. Architecture and the Key Functional Components	3
1.3.1. Brick Data Store	4
1.3.2. Brick Microservices	4
1.4. Supported Topologies	5
1.5. Interfaces to Operate and Manage RBFS	6
1.6. Features and Components	6
1.6.1. Routing	6
1.6.2. Layer 2 Services	7
1.6.3. Multicast	8
1.6.4. Access and Subscriber Management	9
1.6.5. RBFS (Hierarchical) Quality of Service	10
1.6.6. RBFS Carrier-Grade Network Address Translation	10
1.6.7. Ethernet VPN - Virtual Private Wire Service	11
1.6.8. RBFS Redundancy	11
1.6.9. Zero Touch Provisioning	11
1.6.10. Scalability in RBFS	11
1.6.11. Security in RBFS	12
1.6.12. Logging and Observability in RBFS	13
1.6.13. RBFS Software Installation	13
2. Platform	15
2.1. Platform Overview	15
2.1.1. Supported Platforms	16
2.1.2. End-of-Life Policy	16
2.1.3. Guidelines and Limitations	16
2.1.4. Brick Daemon (BD)	17
2.1.5. Setting Up System Parameters	17
2.1.6. CPU Watchdog Timer Utility for Hardware Platforms	18
2.1.7. Displaying Platform Information	19
2.1.8. Displaying RBFS Version Information	20
2.1.9. Configuring Platform Profiles	20
2.2. Platform Hardware Information	26

2.2.1. RBFS Multiservice Edge Image on Edgecore CSR440 (AS7535-28XB)	26
2.2.2. RBFS Multiservice Edge Image on UfiSpace S9510-28DC	27
2.2.3. RBFS Access Leaf & Multiservice Edge Images on Edgecore AGR420 (AS7946-74XKSB) . .	28
2.2.4. RBFS Access Leaf and Multiservice Edge Images on UfiSpace S9600-72XC	29
2.2.5. RBFS Spine Image on Edgecore AGR400 (AS7946-30XB)	30
2.2.6. RBFS Spine Image on UfiSpace S9600-32X	31
2.2.7. RBFS Access Leaf and Multiservice Edge Images on UfiSpace S9600-102XC	31
2.2.8. RBFS L2 Wholesale (L2BSA) Image on UfiSpace S9500-22XST	32
2.3. Firmware Versions	33
2.4. Feature Support Matrix	35
2.4.1. Overview	35
2.4.2. Access-Leaf Image	36
2.4.3. Multiservice Edge Image	40
2.4.4. Spine Image	46
2.4.5. L2 Wholesale (L2BSA) Image	49
2.5. Resource Limits/Feature Support	51
2.5.1. Access-Leaf Image	53
2.5.2. Multiservice Edge Image	53
2.5.3. Spine Image	54
2.5.4. L2 Wholesale (L2BSA) Image	55
2.5.5. Show CLI for Displaying ACL Table Information	57
2.5.6. Static, System, and Dynamic Label Ranges	59

1. RBFS Overview

1.1. RBFS At-a-Glance

RtBrick Full Stack (RBFS) is a carrier-grade network operating system developed specifically for disaggregated telco networks. It supports deployments ranging from small local carriers to Tier 1 to Tier 4 operators, meeting stringent scalability, reliability, and performance requirements. By disaggregating software from hardware, RBFS helps operators modernize their infrastructure without being locked into proprietary systems. It is deployed on bare-metal switches and transforms them into fully functional and programmable routing platforms. These platforms are capable of implementing a broad set of use cases across broadband access, edge, metro, CGNAT, and peering network segments.

RBFS can be operated in a variety of roles within the network. These include its use as a Broadband Network Gateway (BNG) for subscriber termination, as a peering router in the operator's core, or as a system supporting integrated Carrier-Grade NAT (CGNAT) functionality. This role flexibility makes it suitable for both centralized and distributed deployments in modern telecom architectures.

1.1.1. Modular and Resilient Architecture

RBFS has been built around a modular, microservice-oriented architecture. Each protocol or service runs independently as an isolated process, providing high fault isolation and resilience. This architecture simplifies lifecycle operations such as updates, restarts, and diagnostics. For example, individual components can be restarted without requiring a system-wide reboot or service interruption.

A key element of this design is the distributed datastore known as the Brick Datastore (BDS). This datastore ensures that configuration and operational state are consistently available to all modules. Deterministic behavior is achieved through this shared state, enabling reliable operation and support for advanced functions such as transactional configuration, rollback, and state persistence across restarts.

1.1.2. Automation and Management Interfaces

RBFS integrates with industry-standard management interfaces and protocols, including CLI, NETCONF, RESTCONF/YANG, gNMI, RESTful APIs, and telemetry streaming. All configuration and operational data are modeled using YANG or OpenAPI, enabling full automation and lifecycle management through external systems. This allows seamless integration into orchestration platforms, provisioning systems, or in-house tooling, supporting both push and pull models.

Support is provided for both synchronous and asynchronous workflows, allowing operators to manage thousands of devices consistently at scale. Built-in support for Prometheus streaming further enhances observability and troubleshooting capabilities. RBFS provides fine-grained access controls and audit mechanisms to support secure operations in multi-team environments.

1.1.3. Comprehensive Feature Set

RBFS supports a wide range of access and aggregation technologies. These include traditional technologies such as PPPoE and L2TPv2 (LAC), as well as IPoE over DHCPv4 and DHCPv6 for modern broadband deployments. The platform offers complete BNG functionality, including DHCP proxy and relay modes, dynamic subscriber session handling, and seamless IP address management.

Advanced features such as hierarchical quality of service (HQoS), multicast replication for IPTV, full CGNAT with high bandwidth scalability, and support for Lawful Interception are included in the base system. The system is also equipped to support dual-stack IPv4 and IPv6 services, ensuring readiness for evolving network requirements.

RBFS also includes mechanisms for stateful redundancy, both Active/Active and Active/Standby, and fast failover, enabling uninterrupted service delivery in large-scale environments. Combined with flexible traffic steering and policy enforcement, RBFS provides a robust foundation for broadband and edge services.

1.1.4. Hardware Flexibility and Ecosystem

RBFS has been validated on a wide range of white-box platforms powered by merchant silicon, including devices from UfiSpace and Edgecore. The operating system runs on OCP-compliant switches, making it compatible with various hardware vendors.

This disaggregation of hardware and software enables operators to optimize both capital and operational expenditures. Devices can be selected based on specific performance, port density, and energy efficiency needs, while the RBFS software stack provides a consistent set of features and interfaces across the network. This approach aligns with the industry's shift toward cloud-native, software-defined infrastructure.

RBFS delivers the flexibility, scalability, and feature richness required to support the transformation of traditional telecom networks into fully disaggregated, automated, and programmable infrastructures.




Figure 1. RBFS Overview

1.2. Why RBFS

RBFS has been designed to address the operational, economic, and architectural challenges of modern service provider networks. Its open and disaggregated design enables the rapid introduction of new features and services, while fostering a collaborative and competitive ecosystem of hardware, silicon, and software vendors.

By disaggregating software from hardware, RBFS gives operators the freedom to choose white-box switches based on technical or commercial criteria, avoiding vendor lock-in and reducing hardware dependency. This disaggregated approach significantly lowers both capital and operational expenditures. It supports cloud-native deployment models, making it particularly suitable for environments that require elasticity, programmability, and seamless integration into DevOps pipelines.

RBFS is built on a cloud-native microservices architecture, offering major advantages over traditional monolithic systems. Each application is isolated into its own process, which allows for independent scaling and fault isolation. This architecture provides greater agility and simplifies the deployment of software updates and new features without affecting overall system stability.

Operational efficiency is further enhanced by native support for automation and observability. This accelerates innovation and reduces operational overhead.

With its ability to meet the technical requirements for diverse networks, RBFS represents a future-proof platform that combines performance, openness, and automation in a single solution.

1.3. Architecture and the Key Functional Components

RBFS has been designed based on a microservices architecture to cater a rapidly growing broadband traffic. An RBFS container contains multiple microservices, known as daemons. These microservices are the building blocks of the RBFS ecosystem and they can communicate with each other through a centralized in-memory datastore called Brick Data Store (BDS).




Figure 2. RBFS Microservices

1.3.1. Brick Data Store

RBFS has a schema-driven and in-memory database called BDS (Brick Data Store). As an in-memory data store, BDS relies mainly on the main memory for the storage of data which is contrary to the databases that store data on disks. BDS has architecturally been designed to enable very minimal response time by removing the time to access data stored in disks. BDS acts as a control plane and provides all required data and instructions to the daemons for their functioning.

1.3.2. Brick Microservices

RBFS's microservices architecture enables it to perform various complementary functions and deliver services.

For example, the subscriber microservice (`subscriberd`) manages the current subscriber state and is responsible for authentication, authorization, and accounting. The `ribd` microservice is responsible for route selection, next-hop resolution, tunnel selection and recursion.

There are microservices such as `CtrlD` (Controller) and `ApiGwD` (API Gateway) which are part of the RBFS ecosystem. These daemons sit in the middle (on the Host System) and manage all the communication between the client and backend services running in the container. The API Gateway (`ApiGwD`) microservice provides a single point access to expose services running inside of the RBFS container.

RBFS microservice and other dependencies are packaged as a Debian LXC container. The RBFS container is hosted on the RtBrick Host, an open-source operating system, which can be run on white box switches.

RBFS can perform various roles such as Spine, Leaf, and Multiservice Edge which serve different use-cases. The software images of these various roles contain daemons that are required to serve these roles for their different functions. The RBFS Multiservice Edge software image contains all the RBFS daemons packaged in a container, other roles such as Spine and Leaf include only the daemons which are required to carry out their respective functions. For example, the Spine image includes (in addition to other daemons) the interior gateway protocol daemons such as `isis.appd`, `isis.iod`, `ospf.appd`, and `ospf.iod` which are not required in the Leaf image.

The daemons such as `alertmanager`, `confd`, `etcd`, `fibd`, `hostconfd`, `ifmd` and so on are present in the images of both the Spine and Leaf roles as these daemons are required in both of these roles.

Containerization of Daemons

RBFS daemons and other dependencies are packaged as a Debian LXC container. This containerization is a logical layer that helps to make the applications secure, flexible, and portable by providing isolation. This RBFS container is hosted on the RtBrick Host, Debian-based operating system, which can be run on white box switches.

RBFS can perform various roles such as Spine, Leaf, and Multiservice Edge which have different functions to serve. The software images of these various roles contain daemons that are required to serve these roles for their different functions. Though, the RBFS Multiservice Edge software image contains all the RBFS daemons packaged in a container, other roles such as Spine and Leaf include only the daemons which are required to carry out their respective functions.

Launching Microservices Dynamically

When the RBFS container starts up, it installs different sets of microservices depending on the image role and platform. This is done to minimize resource consumption. In RBFS, the microservices are divided into two categories: base microservices and on-demand microservices. RBFS containers will have all microservices installed according to the platform and image role, but not all will be enabled on bootup. Only the base microservices will be enabled and started on bootup. On-demand microservices will only be started when their respective configurations are configured and will stop once all dependent configurations are deleted.

For instance, when the user configures BGP with the CLI command `set instance <instance> protocol bgp`, the `rtbrick-bgp.appd.1` and `rtbrick-bgp.iod.1` services will start. And, once the BGP configuration is deleted, "rtbrick-bgp.appd.1" and "rtbrick-bgp.iod.1" will be stopped after 5 minutes (graceful shutdown time).

By default, the following base microservices will be running in the container.

- rtbrick-confd
- rtbrick-etcd
- rtbrick-fibd
- rtbrick-hostconfd
- rtbrick-ifmd
- rtbrick-lldpd
- rtbrick-mribd
- rtbrick-opsd
- rtbrick-pool
- rtbrick-resmond
- rtbrick-resmond-agent
- rtbrick-restconfd
- rtbrick-ribd
- rtbrick-staticd

When you make other RBFS configurations, the required on-demand microservices will be automatically enabled.

1.4. Supported Topologies

RBFS can be deployed in a spine-leaf architecture and can also be deployed standalone in a single switch which consolidates all the features in one switch.

A Multiservice Edge architecture offers all the functionalities of a spine-leaf BNG architecture on a single bare-metal switch.

A spine-leaf architecture is a two-tier network topology that consists of two switching layers — a

spine and a leaf. In this topology, two layers of switches interconnect. The leaf layer consists of access switches that aggregate traffic and connect directly to the spine which is the core network.

The RBFS spine-leaf topology provides better scalability. This topology enables horizontal scaling by adding additional switches to add more capacity to handle increased traffic. This topology is also useful for higher bandwidth.

1.5. Interfaces to Operate and Manage RBFS

RBFS provides a CLI and a rich set of commands that you can use to operate, configure, monitor, and manage the system and its various components. Using the RBFS CLI, you can configure static IPv4, IPv6, MPLS, and multicast routes.

In addition to the CLI, RBFS also offers industry-standard tools and utilities such as RESTCONF/YANG.

RBFS supports REST-based industry-standard tools such as RESTCONF/YANG and Operational State API to enable communication with the software and underlying devices. RESTCONF is a programmatic interface that enables you to programmatically access RBFS devices and manage configurations.

The Operational State API microservice (opsd) provides the operation state of the system. It forms a stable contract between RBFS and network management systems and inspects the operational state of the device to diagnose and troubleshoot problems.

RBFS APIs allow to access and consume RBFS data simply and securely.




Figure 3. RBFS API Overview

1.6. Features and Components

1.6.1. Routing

RBFS, at its core, is a routing software that supports both IP routing and MPLS routing. In dynamic IP routing, RBFS supports all major routing protocols that include OSPFv2 and IS-IS (interior gateway protocols) and BGP (exterior gateway protocol).

RBFS also supports Protocol Independent Multicast (PIM), a multicast routing protocol that runs

over existing unicast infrastructure. PIM-SSM uses a subset of PIM sparse mode and IGMP to permit a client to receive multicast traffic directly from the source.

BGP

BGP is a standard exterior gateway protocol (EGP) supported by RtBrick. BGP is considered a “Path Vector” routing protocol and maintains a separate routing table based on the shortest Autonomous System (AS) path and various other route attributes.

IS-IS

IS-IS, or Intermediate System to Intermediate System, is an open standard routing protocol. ISO published the standard as a way to route datagrams as part of their OSI stack. IETF later republished the standard, and added IP route support.

OSPFv2/v3

OSPF (Open Shortest Path First) is an Interior Gateway Protocol that distributes routing information within a single Autonomous System (AS) in an IP network. OSPF is a link-state routing protocol that uses link-state information to form a routing table and exchange the routing information with the neighbors.

RtBrick FullStack (RBFS) supports OSPF version 2 (OSPFv2) and OSPF version 3 (OSPFv3), including authentication, LDP-IGP sync, and redistribution policy. RBFS does not support OSPFv3 Virtual Link.

LDP

Label distribution protocol (LDP) is the most commonly used protocol in the MPLS network. It generates and distributes labels and thus helps in MPLS packet switching and forwarding. By using LDP, label-switching routers in an MPLS network can exchange label mapping information to create label-switched paths (LSPs) for switching data packets. RtBrick FullStack (RBFS) supports Dual-stack, which means LDP can exchange FEC-label bindings over either IPv4 or IPv6 networks.

Static Routing

RBFS supports static routing that allows you to configure routes manually.

Segment Routing

RBFS supports segment routing using the IS-IS and OSPF protocols. In segment routing, the source router decides the path (throughout the network) to the destination and encodes the path details in the packet header as an ordered list of instructions. The routers on the path do not take any forwarding decisions but just execute the forwarding instructions.

1.6.2. Layer 2 Services

L2X

Layer 2 Cross-Connect (L2X) is a data plane feature that connects two physical ports (IFPs) using Layer 2 switching. L2X can switch the traffic between two IFPs to provide the trunk service for an

Ethernet switch.

EVPN-VPWS

Ethernet Virtual Private Network (EVPN) is a Layer 2 internetworking technology similar to BGP/MPLS IP VPN. EVPN uses extended BGP reachability information and advertisements between different Layer 2 networks at various sites in the control plane.

The EVPN Virtual Private Wire Service (VPWS) is a point-to-point (P2P) service that is built on the EVPN service architecture. EVPN-VPWS uses MPLS tunnels to traverse the backbone network. It offers a Layer 2 packet forwarding mode that connects access circuits (ACs) as per the specifications of RFC 8214.

BGP-signaled L2VPN

The BGP-signaled L2VPN uses BGP for signaling and auto-discovery to establish multipoint Layer 2 VPN over the MPLS backbone network. The remote cross-connect is a point-to-point (P2P) service that connects two locations using the MPLS core network and MP-BGP. The remote cross-connect uses MPLS tunnels to traverse the backbone network. It offers a Layer 2 packet forwarding mode that connects access circuits (ACs). The RBFS implementation of BGP-signaled L2VPN is in accordance with RFC-6624 which supports L2VPN using BGP for auto-discovery and Signaling.

1.6.3. Multicast

IGMP

Internet Group Management (IGMP) protocol allows a host to advertise its multicast group membership to neighboring switches and routers. IGMP is a standard protocol used by the TCP/IP protocol suite to achieve dynamic multicasting.

PIM

PIM SSM builds shortest-path trees (SPTs) rooted at the source immediately because in SSM, the router closest to the interested receiver host is informed of the unicast IP address of the source for the multicast traffic. That is, PIM SSM bypasses the RP connection stage through shared distribution trees, as in PIM sparse mode, and goes directly to the source-based distribution tree.

Multicast VPN

The Multicast VPN (MVPN) feature provides the ability to support multicast over a Layer 3 VPN. Multicast allows the efficient distribution of information between a single multicast source and multiple receivers. IP multicast is used to stream video, voice, and data to an MPLS VPN network core. The RBFS MVPN implementation is based on RFC 6513 “Multicast in MPLS/BGP IP VPNs” and RFC 6514 “BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs”.

Routing Policy

RBFS routing policies allow to control and modify the behavior of routing protocols such as IS-IS, OSPF, and BGP. RBFS has a generic routing policy framework that serves multiple purposes and applications. In RBFS, the routing policy implementation is performed by four major components:

1.6.4. Access and Subscriber Management

RtBrick's modular and scalable subscriber management offers the next-generation access infrastructure (ng-access) that supports protocols such as PPPoE, IPoE, L2TPv2, DHCPv4 and DHCPv6 and RADIUS. It provides subscriber authentication, access, service creation, activation, and deactivation. It collects accounting statistics for the subscriber sessions. RBFS enables you to address the challenges such as interoperability with numerous client devices from various vendors which requires a well-implemented and industry-proven access protocol stack, including support for all relevant RFCs. RBFS subscriber management infrastructure provides the next generation of internet access protocols designed for carrier-grade services.

Support for PPPoE, IPoE, and L2TPv2

RBFS supports subscriber session management protocols such as Point-to-Point Protocol over Ethernet (PPPoE), Layer Two Tunneling Protocol (L2TPv2), and IP over Ethernet (IPoE) to deliver network access services to broadband subscribers.

PPPoE establishes a PPP connection over the ethernet. In RBFS, the PPPoE microservice (`pppoed`) manages PPPoE and PPP sessions.

IP-over-Ethernet (IPoE) is an alternative to PPPoE to deliver network access services to broadband subscribers. IPoE does not require client dial-in software and is easy to use when accessing the network. In RBFS, the IPoE microservice (`ipoed`) manages IPoE services using DHCPv4 and DHCPv6.

The L2TPv2 microservice (`l2tpd`) is used for the L2TPv2 tunnel and session handling. L2TP is a Layer-3 tunneling protocol that initiates a tunnel between an L2TP access concentrator (LAC) and an L2TP network server (LNS). This enables Point-to-Point Protocol (PPP) link layer to be encapsulated and transferred across the internet.

Accounting

RBFS accounting is the process of tracking subscriber activities and network usage in a subscriber session for auditing and billing. Accounting tracks information such as subscriber identity, the number of packets and bytes transferred from and to the network, start and stop times of the sessions and so on. The accounting keeps track of resources used by the subscriber during the sessions. This includes the session time called time accounting and the number of packets and bytes transmitted during the session called volume accounting. In RBFS, accounting can be performed based on classes or types of services such as video, VoIP, and data.

Support for Lawful Interception

RBFS supports Lawful Interception (LI) to allow legal authorities to obtain communications network data for analysis or evidence. LI is a technique of intercepting certain user data streams tunneling the intercepted traffic to a mediation device with the data and only the users with appropriate credentials can access the intercepted data.

HTTP Redirect Service

RBFS HTTP Redirect service allows network service providers to intercept and redirect HTTP request traffic from subscribers to a designated captive portal instead of the original destination. This powerful service has a multitude of use cases, ranging from subscriber re-authentication to enforcing acceptance of network usage policies. It allows network service providers to re-authenticate subscribers when necessary and ensure that users explicitly accept network usage policies before accessing services. By implementing the RBFS HTTP Redirect Service, network service providers can efficiently manage user access and enforce compliance with network regulations and policies, ultimately enhancing the overall security and user experience within their network environment.

1.6.5. RBFS (Hierarchical) Quality of Service

RBFS Quality of Service (QoS) is a method of prioritizing network traffic for mission-critical applications and high-priority network services such as voice and video. It provides control over a variety of traffic types and ensures that critical data traffic gets sufficient network resources such as bandwidth.

RBFS can perform priority forwarding of data packets throughout the network. For this preferential forwarding, it identifies and classifies the network traffic. So that the critical network packets get sufficient resources. RBFS QoS ensures the required level of service and provides cost benefits to network providers by enabling them to use network resources efficiently.

RBFS also supports Hierarchical Quality of Service (HQoS), a mechanism that allows you to specify Quality of Service (QoS) behavior for different traffic classes. QoS allows classifying services such as voice and video, but using HQoS, you can apply QoS policies to different users, VLANs, logical interfaces, and so on. RBFS employs HQoS by using the mechanisms such as classifier, queuing, scheduler, policer, shaper, and remarking. HQoS provides a higher degree of granularity in traffic management.

1.6.6. RBFS Carrier-Grade Network Address Translation

RBFS is multi-service edge routing software with which you can deliver both CGNAT and BNG functionalities on a single open switch to reduce costs and increase efficiency.

The RBFS CGNAT or NAT444 solution supports Network Address Port Translation (NAPT), which has the potential to conserve IPv4 addresses for service providers. NAPT is an effective method for allowing multiple devices to connect to the Internet using a single public IPv4 address.

The solution can address the IPv4 depletion challenge of service providers. Using the RBFS CGNAT, service providers can serve a large number of subscribers using a limited number of public IPv4 addresses.

RBFS CGNAT solution has some unique characteristics. Both BNG functionalities and CGNAT functionalities can coexist in a single RBFS device. RBFS CGNAT implements NAT in the chipset that allows for the delivery of CGNAT functionality in-line, fully integrated into the packet processing pipeline alongside other functions in the data plane itself, without requiring any additional chipset resources.

RBFS CGNAT supports deterministic NAT mode of address translation, which provides a consistent mapping of private IPv4 addresses with public IPv4 addresses and port ranges. This mode ensures a one-to-one mapping of private IPv4 addresses with public IPv4 addresses, allowing you to specify the private address and its matching public address and port range. The given private IPv4 address is always translated to the same public address.

1.6.7. Ethernet VPN - Virtual Private Wire Service

RBFS Ethernet VPN - Virtual Private Wire Service (EVPN-VPWS) technology provides point-to-point Layer 2 services over an IP or MPLS network. It is based on the EVPN (Ethernet VPN) technology, which extends the BGP (Border Gateway Protocol) to handle MAC addresses and Ethernet segments in addition to IP prefixes.

EVPN-VPWS allows service providers to offer Layer 2 services with better scalability, flexibility, and ease of operation compared to traditional Layer 2 technologies like VPLS (Virtual Private LAN Service). It uses BGP as the control plane protocol to distribute MAC reachability information across the network, enabling efficient MAC learning and forwarding.

1.6.8. RBFS Redundancy

RBFS supports deployment in redundancy mode that protects from link and node failures. Node and link outages that may occur on an RBFS access network can bring down the subscriber services. RBFS Redundancy helps to minimize the impact of these events and to reduce interruptions and downtime by providing a resilient system.

RBFS Redundancy protects subscriber services from various software and hardware outages. It provides mechanisms to enhance network resiliency that enables subscriber workloads to remain functional by ensuring a reliable switchover in the event of a node or link outage. With RBFS Redundancy, if one node goes down, another node can automatically take over the services.

RBFS Redundancy protects subscriber groups using an active standby node cluster model. RBFS Redundancy architecture consists of an active-standby node cluster and one node is active that runs workloads at a time. The peer node, which is identical to the first node, mirrors the concurrent subscriber state data from the peer and takes over workloads in the event of a node or link failure.

1.6.9. Zero Touch Provisioning

By leveraging the Zero Touch Provisioning (ZTP) feature, you can automate many of the RBFS deployment and setup tasks. ZTP allows you to set up and configure the platforms automatically by eliminating the repetitive manual tasks in a large-scale environment. This feature significantly reduces human touch points and errors prone by manual interventions and makes the deployment easier.

1.6.10. Scalability in RBFS

RBFS allows horizontal scaling to enhance system capacity. You can add additional switches to the spine and leaf layers to enhance capacity to handle increased subscriber traffic.

RBFS offers subscriber management capacity in a scale-out architecture called the Point-of-

Deployment (PoD), also known as a SEBA PoD (SDN-enabled PoD). A large-scale PoD consists of access leaf routers aggregated by a layer of spine routers in an auto-provisioned CLOS topology. The access leaf routers provide subscriber management functionality. For even greater scalability, a layer of border leaf routers can be added to the core of the network provider network to provide more connectivity.

The leaf routers can be scaled out horizontally to increase the number of subscribers supported on the PoD, providing a pay-as-you-grow model. PPPoE subscribers can be terminated on the access leaf routers or tunneled to an LNS over L2TPv2. L2 Cross Connect (L2X) allows subscriber traffic to be tunneled out of the PoD at Layer 2, providing connectivity.

1.6.11. Security in RBFS

In RBFS, security is integrated into the foundation of the network. RBFS implements several techniques and methods to safeguard the entire network infrastructure. RBFS has a comprehensive set of security capabilities that deploy multiple security controls to protect different areas of the system and network.

Security features for RBFS Control Plane

RBFS Control Plane security feature enables filtering and rate-limiting the traffic transmitted from the forwarding plane to the control plane. RBFS uses Access Control Lists (ACLs) and policers to secure the router's control plane.

All routing protocols, management protocols, and service protocols run in the control plane. The output of these protocols results in databases such as routing tables, MAC tables, ARP tables, and so on, which eventually get programmed in the forwarding plane.

ACLs are the building blocks of control-plane security. RBFS employs fundamental mechanisms - Protocol ACLs and Route Lookup - for redirecting control plane traffic to the CPU and policers for controlling CP traffic to the CPU.

All routing protocols (BGP, OSPF, and ISIS), Management Protocols (SSH, RESTCONF/YANG, and so on), Service Protocols (RADIUS, NTP, and TACACS+), and Access Protocols (PPPoE, DHCP, L2TP, and PPP) automatically create Access Control Lists (ACLs) required to punt the protocol traffic to the CPU Control Plane.

The RBFS Control Plane Security feature adds policers to all protocol ACLs. This feature creates a set of default policers and applies them to the protocol ACLs to secure the control plane from DDoS attacks.

Security features for RBFS Management Plane

RBFS provides the capability to restrict access to the management plane only to authenticated and authorized entities. The authentication identifies the entity and the authorization validates if the entity is allowed to execute the action.

RBFS supports the security protocol, TACACS (Terminal Access Controller Access Control System). RBFS provides a Pluggable Authentication Module (PAM) that enables it to work with TACACS for centralized authentication for users who try to access a router.

For management plane security, RBFS implements token-based authentication that provides access to the management plane through APIs only to the authenticated entities.

RBFS uses JSON web token, an open standard token, that defines a compact and self-contained way for securely transmitting information between parties as a JSON object. The ApiGwD microservice validates the access token against a JSON web key set (JWS).

1.6.12. Logging and Observability in RBFS

RBFS logging is the process of writing log messages during the execution of an event. Logging provides reports about the events in the entire RBFS ecosystem at different functional areas. You can configure logging based on the different severity levels available. RBFS also allows you to send logs to third-party log management servers such as Graylog where you can view and analyze the real-time data. It provides you the ability to trace out the errors of the applications in real-time.

Operational state visibility is crucial for troubleshooting, testing, monitoring, and capacity management. To enable operational visibility, it is required to collect router metrics periodically. RBFS allows the ingestion of time-series data allows to send operational queries.

RBFS uses Prometheus, an open-source system monitoring and alerting tool, for monitoring and metric collection. Prometheus collects time-stamped data for events, network data, application performance, and so on. The tool allows analyzing metrics with the **PromQL** query language. Additionally, RBFS provides an optional alert management tool. You can use both of these tools together with its own services to integrate them into the RBFS ecosystem.

Observability Using SNMP

RBFS SNMP (Simple Network Management Protocol) provides a network monitoring mechanism that collects state information from various network devices and components. With SNMP, you can monitor interfaces, CPU usage, temperature of the device, bandwidth usage, and so on. For example, if an interface goes down on one of the devices, SNMP can quickly alert this. The RBFS SNMP implementation allows retrieving system state information using the Protocol Data Unit (PDU) from various network components.

SNMP allows performing various operations that include GET for retrieving data, SET for modifying data, TRAP for notifying an event and so on. These operations provide management access to the MIB hierarchy. RBFS supports the SNMP version 2c and SNMP version 3.

Resource Monitoring

Monitoring the device and its various components is very crucial to analyze the health of devices. RBFS provides resource monitoring capabilities to keep track of various components of the devices. RBFS has a dedicated microservice called **resmond** to discover and monitor the device resources. With RBFS Resource Monitoring, you can continuously observe the health of the system resources such as CPU, Memory, Processes, Disks, Sensor, and Optics.

1.6.13. RBFS Software Installation

RBFS software is available at RtBrick Image Store (<https://releases.rtbrick.com/>) where you can

download the latest version. For more information on RBFS software installation, see [RBFS Installation](#).

2. Platform

2.1. Platform Overview

RBFS allows customers to select their preferred hardware platforms. This document provides information about RBFS-supported hardware platforms and the features available on each platform. It also includes information on firmware versions, the feature support matrix, and resource limits. This guide helps you understand the various supported platforms and their capabilities and help you evaluate and choose the right platform for your needs.

The RBFS Software runs as an LXC container on the RtBrick Host system, deployed on bare-metal switches. The switch can function independently as standalone unit leveraging all functionalities into a single platform or can be combined into a leaf-spine architecture for scalable deployments. The software supports Zero-Touch Provisioning (ZTP) installation and provides REST-based APIs for configurations and operations.

The guide is organized into the following sections:

- **Platform Overview:** A high-level summary of the platforms.
- **Platform Hardware Information:** Detailed specifications of the hardware, including CPU, ASIC, Form Factors, memory, storage, and network interfaces.
- **Firmware Versions:** Information on supported firmware versions and their compatibility with RBFS.
- **Feature Support Matrix:** Information about RBFS features supported on various platforms and for various roles.
- **Resource Limits/Feature Support:** Details on resource allocation, such as maximum routes and interfaces.

The following image shows a high-level overview of the platform architecture.




Figure 4. Platform Overview

Platform hardware consists of forwarding ASICs and an RtBrick Host operating system. A RBFS container that resides on top of this software includes all necessary packages to deliver access and

routing protocols.

2.1.1. Supported Platforms

RtBrick's software has been validated on the following hardware platforms.

Hardware Platform	Switch ASIC	Role
Edgecore CSR440 (AS7535-28XB)	Q2A	Multiservice Edge
UfiSpace S9510-28DC	Q2A	Multiservice Edge
Edgecore AGR420 (AS7946-74XKSB)	Q2C	Access Leaf, Multiservice Edge
UfiSpace S9600-72XC	Q2C	Access Leaf, Multiservice Edge
Edgecore AGR400 (AS7946-30XB)	Q2C	Spine
UfiSpace S9600-32X	Q2C	Spine
UfiSpace S9600-102XC	Q2C	Access Leaf, Multiservice Edge
UfiSpace S9500-22XST	QAX	L2 Wholesale (L2BSA)

For a list of features and sub-features supported by each platform, see [Feature Support Matrix](#).

2.1.2. End-of-Life Policy

RtBrick periodically introduces software support for new hardware platforms and use cases. Likewise, support for older software is discontinued to ensure that RtBrick can provide appropriate attention to software critical to drive business functions. The [End-of-Life Policy page](#) details the platforms that are no longer supported or have limited support from RtBrick.

2.1.3. Guidelines and Limitations

QAX-based Platforms

- An additional restriction applies to ports belonging to a port group on QAX-based platforms. For more information, see section [Guidelines and Limitations](#) of the *Interfaces User Guide*.

- MPLS forwarding is not supported on double-tagged interfaces. It is supported only on untagged and single-tagged interfaces.

2.1.4. Brick Daemon (BD)

RBFS runs multiple Brick Daemons (BD). Every application that runs within RBFS is fundamentally a brick daemon. For example, forwarding daemon (fibd), configuration daemon (confd), BGP (bgp.iod or bgp.appd), or interface management daemon (ifmd).

Brick Daemon (BD) Restartability

If a brick daemon fails, RBFS will restart it automatically. All microservices can restart automatically for an unlimited number of times. It improves system resilience in the event of recoverable failures. If the automatic restart does not succeed, you can use the Ubuntu system control to start a daemon.

For more information about troubleshooting the Brick Daemons, see section "2.2. Brick Daemons" of the *RBFS NOC Troubleshooting Guide*.

Limited Core Files to Save Disk Space

RBFS can prevent disk space exhaustion caused by repeated process crashes. Each process is allowed to retain up to three core files. When a new core file is generated, the oldest file will be automatically deleted. It improves disk space management while retaining recent crash data for debugging.

2.1.5. Setting Up System Parameters

You can configure basic host system parameters such as 'element name' and 'pod name' using the `set system host` command.

Syntax:

set system host <attribute> <value>

Attribute	Description
element-name <element-name>	Specify the name of the element (container). A pod can contain a group of elements.
pod-name <pod-name>	Specifies the name of the Pod. Pod stands for point (zone) of deployment.
nameserver <nameserver>	Specify the IP address of the DNS server. It is allowed to configure a maximum of three servers.
role	Specify the role of the element. Roles include Spine, Access Leaf, Border Leaf, and so on.



The `hostname` is defined in the DHCP server. The maximum length of the `hostname` is restricted to 64 characters.

Example: System Parameters Configuration

```
supervisor@rtbrick>LEAF01: op> show config system host
{
  "rtbrick-config:host": {
    "element-name": "ufi10.q2c.u9.r4.nbg.rtbrick.net",
    "pod-name": "nbg4"
  }
}
```

2.1.6. CPU Watchdog Timer Utility for Hardware Platforms

The CPU watchdog timer utility is located in BMC and helps to detect any CPU failure. It also enables the CPU to recover from faults. By default, Watchdog Timer functionality is enabled, and it does not require any configuration from users. However, you can configure it to change the default settings.

Enable CPU Watchdog Timer in Hardware

Use the following command to enable watchdog timer on the hardware.

Syntax:

set system platform-management watchdog CPU <attribute> <value>

Attribute	Description
action	Specifies the possible timeout actions: • hard-reset: A reset will take place based on the configured settings. • no-action: No action will be taken. • power-cycle: The watchdog timer will power cycle the device. • shutdown: The device will shut down per the configured setting. Default value: power-cycle.
interval	Specifies the watchdog timer interval in seconds. Default value: 1800 seconds.

Example: Enable CPU Watchdog Timer

```
supervisor@rtbrick>LEAF01: op> show config system platform-management watchdog
{
  "rtbrick-config:watchdog": [
    {
      "type": "CPU",
      "action": "hard-reset",
      "interval": 3600
    }
  ]
}
```

```
}  
]  
}
```

You can use the **ipmitool** utility to view the status of the CPU watchdog timer.

Log into the RtBrick Host and enter the command as shown below to validate CPU watchdog timer functionality.

```
supervisor@onl>rtbrick:~ $ sudo ipmitool mc watchdog get  
Watchdog Timer Use:      OS Load (0x43)  
Watchdog Timer Is:       Started/Running  
Watchdog Timer Actions:  Power Cycle (0x03)  
Pre-timeout interval:    0 seconds  
Timer Expiration Flags:  0x00  
Initial Countdown:       1800 sec  
Present Countdown:       1795 sec
```

2.1.7. Displaying Platform Information

To display platform information, use the **show platform** command, as shown in the example below.

```
supervisor@rtbrick>LEAF01: op> show platform  
x86_64-ufispace_s9600_72xc-r0  
Role           : multiservice-edge  
Platform       : q2c(BCM88820_A1)  
External Processor : OP2(model: X register-value: 0x69A)  
Vendor         : Ufi Space  
Manufacturer    : Ufi Space  
Manufacture date : 06/28/2021 10:51:29  
MAC address     : E8:C5:7A:8F:78:0D  
Part number     : S9600-72XC-RB6B  
Serial number   : WJ91B67T00009B3  
Product name    : S9600-72XC-R  
ONIE version    : 2022.02v04  
Label revision  : N/A  
Diag version    : 0.1.4  
Country code    : CN  
Device version  : 1
```



- Information about external processors is displayed only for Q2C platforms and non-spine image roles. "N/A" displayed otherwise.
- When using virtual platforms, the "show platform" CLI command does not provide any output.

2.1.8. Displaying RBFS Version Information

To display RBFS version information, use the `show version` command, as shown in the example below.

```
supervisor@rtbrick>LEAF01: op> show version
UUID       : 4e83e9a4-eb65-4ab0-a756-3ae3536dd6ad
Version    : 25.4.0-g6daily.20251126013758+Bdevelopment.C0313b7f9
Role       : multiservice-edge
Platform   : q2c
Format     : lxd
Build date  : 2025-11-26 01:37:58 UTC
```

To display detailed version information for RBFS along with library versions, use the `show version detail` command, as shown in the example below.

```
supervisor@rtbrick>LEAF01: op> show version detail
UUID       : 4e83e9a4-eb65-4ab0-a756-3ae3536dd6ad
Version    : 25.4.0-g6daily.20251126013758+Bdevelopment.C0313b7f9
Role       : multiservice-edge
Platform   : q2c
Format     : lxd
Build date  : 2025-11-26 01:37:58 UTC
```

Component	Version	Timestamp	Branch
alertmanager	0-g6daily.2025112122211+Bdevelopment.Cfd733666	2025-11-11 10:41:02	development
cligen	0.1.0-g6daily.20251105111321+Bdevelopment.Cd8964728	2025-11-05 10:12:50	development
clixon	4.3.1-g6daily.20251125140131+Bdevelopment.Cfc95084c	2025-11-13 03:40:34	development
ems-service-event	0.2.0-g6daily.20251104173513+Bdevelopment.Cd923496f	2025-10-28 09:33:01	development
etcd	0-g6daily.20251105113559+Bdevelopment.C8987db4e	2025-10-28 09:32:14	development
fwd-plugin-bcm-q2c-s96...	4.2.0-g6daily.20251126010001+Bdevelopment.Cca388617	2025-11-26 00:40:28	development
hostconfd	0-g6daily.20251121153022+Bdevelopment.C4b49f4ba	2025-11-21 08:26:47	development
hostnetconfd	0.7.0-g6daily.20251119153025+Bdevelopment.C0a756d06	2025-11-19 08:13:05	development
json-builder	0.1.0-g6daily.20251105112151+Bdevelopment.C7bf78ef6	2025-10-28 09:26:24	development
json-parser	1.1.0-g6daily.20251104173526+Bdevelopment.C95257c44	2025-10-28 09:25:00	development
libbgp	1.2.0-g6daily.20251126005928+Bdevelopment.Cb86c1199	2025-11-21 11:52:25	development
libcjson	1.0.0-g6daily.20251104173453+Bdevelopment.C1bfcae2f	2025-10-28 09:24:36	development
libconfd	1.1.1-g6daily.20251125140948+Bdevelopment.C6f603a5b	2025-11-04 12:37:53	development
libdict	1.0.1-g6daily.20251104173534+Bdevelopment.C3698deb7	2025-10-28 09:28:14	development
libegrep	1.1.0-g6daily.20251104173611+Bdevelopment.Cd8731a66	2025-10-28 09:28:29	development
libforwarding	4.2.0-g6daily.20251126005549+Bdevelopment.Cca388617	2025-11-26 00:40:28	development
libh2o	2.2.6-g6daily.20251105112618+Bdevelopment.Ccabf806c	2025-10-28 09:28:27	development
libigmp	1.1.0-g6daily.20251126005259+Bdevelopment.C63aa1a88	2025-10-28 09:24:52	development
libisis	1.2.0-g6daily.20251126005938+Bdevelopment.Cae8678be	2025-11-26 00:38:37	development
liblaccp	0.2.0-g6daily.20251126010938+Bdevelopment.C6b68d457	2025-11-04 12:37:48	development
libldp	0.3.0-g6daily.20251126005023+Bdevelopment.C01e28f15	2025-11-05 06:11:14	development
liblldp	1.1.1-g6daily.20251126010933+Bdevelopment.C81ae1f31	2025-10-28 09:24:48	development
liblwip	1.2.0-g6daily.20251105112452+Bdevelopment.C7c22f157	2025-11-04 08:51:45	development
libnat	1.4.0-g6daily.20251126005258+Bdevelopment.C8f3e30ca	2025-11-05 06:15:54	development
libng-access	1.2.0-g6daily.20251126005302+Bdevelopment.C470442a6	2025-11-19 17:15:01	development

```
<...>
```

2.1.9. Configuring Platform Profiles

Network operators can determine ASIC scale profiles and features to meet their specific business requirements. Currently, the platform profile configuration is supported on the Multiservice Edge image on the Q2C and Q2A platforms. The profiles can be viewed using the `show platform-profile`

command.

Guidelines & Limitations

- The system only allows the configuration of the profiles and features it supports.
- This system allows to configure any one of the supported profiles.
- You can configure multiple features for one profile.
- You can make changes to the features, such as removing or modifying them.
- A system reboot is required for the changes to take effect.

Platform Profile Support

Platform	Supported Profile	Supported Features	Default
Q2C Multiservice Edge	1q, 4q, nat_standalone, nat_1q, nat_4q	IPoE N:1, sFlow, and so on.	4q
Q2A Multiservice Edge	1q, 4q, nat_standalone, nat_1q, nat_4q	IPoE N:1, sFlow, access-multifield-classifier, and so on.	4q (ipoe-n-1)

To configure profile for a specific platform, use the command below.

Syntax:

set system platform profile <profile_name>

Attribute	Description
<profile-name>	Specifies the profile that you want to activate. Supported platform profiles include 1q, 4q, nat_1q, nat_standalone, and nat_4q. The new platform profile will be activated only after a system reboot
feature <feature>	Specifies the feature that enabled along with the profile.
<profile_name>	Specifies the profile that you want to activate. Supported platform profiles include 1q, 4q, nat_1q, nat_standalone, and nat_4q. The new platform profile will be activated only after a system reboot. If no profile is configured, the system defaults to the profile 4q.
<feature>	Specifies the feature that is enabled along with the profile.  To enable support for 400G transceivers, you must specify 400Gx1 (one 400G port mode) or 400Gx2 (two 400G ports mode) as the feature name.
ipoe mode	Choose the desired redundancy mode: standby-forwarding or standby-non-forwarding.

Example: Configuring 1q platform profile

```
set system platform
set system platform profile 1q
```

```
supervisor@rtbrick.net: cfg> show config system platform
{
  "rtbrick-config:platform": {
    "profile": [
      {
        "profile_name": "1q"
      }
    ]
  }
}
```

Example: Enabling the feature ipoe-n-1 for the profile nat_1q

```
supervisor@switch: cfg> set system platform profile nat_1q feature ipoe-n-1
```

```
"rtbrick-config:system": {
  "platform": {
    "profile": [
      {
        "profile_name": "nat_1q"
      }
    ]
  }
}

"rtbrick-config:system": {
  "platform": {
    "profile": [
      {
        "profile_name": "nat_1q",
        "feature": [
          "ipoe-n-1"
        ]
      }
    ]
  }
}
```

Example: Enabling the redundancy mode for the profile nat_1q

```
supervisor@switch: cfg>set system platform profile nat_1q ipoe-mode standby-non-forwarding
```

```
platform": {
  "profile": [
    {
      "profile_name": "nat_1q",
      "feature": [
        "ipoe-n-1"
      ],
      "ipoe-mode": "standby-non-forwarding"
    }
  ]
}
```

To view the profiles configured for a platform after reboot, use the `show platform-profile` command as shown in the example below

Example: Viewing Platform Profile on the Q2A Platform

```
supervisor@rtbrick>LEAF01: op>show platform-profile
Platform-profile: 4q
  Role: multiservice-edge
  MDB Profile: custom_q2a_scale_4q
  Counter Profile: multiservice_edge
  Features: -|ipoe-n-1|-|-
```

```
supervisor@rtbrick>LEAF01: op>show platform-profile
Platform-profile: nat_4q
  Address Translation Service: True
  Role: multiservice-edge
  MDB Profile: custom_q2a_scale_nat
  Counter Profile: multiservice_edge_nat_extend
  Features: -|ipoe-n-1|-|sflow
```

Example: Viewing Platform Profile on the Q2C Platform

```
supervisor@rtbrick>LEAF01: op> show platform-profile
Platform-profile: nat_4q
  External Processor Variant: OP2_M
  Address Translation Service: True
  Role: multiservice-edge
  MDB Profile: custom_rtb_multiservice_edge_nat
  Counter Profile: multiservice_edge_nat_extend
  Forwarding Routes on External Processor: False
```

Ingress Accounting is External: False
Egress Accounting is External: True

The following example configures the platform profile named **4q** to enable the **400Gx1** feature, which activates support for one 400G transceiver on the system.

```
set system platform profile 4q feature 400Gx1
```

```
"rtbrick-config:system": {  
  "platform": {  
    "profile": [  
      {  
        "profile_name": "4q",  
        "feature": [  
          "400Gx1"  
        ]  
      }  
    ]  
  }  
}
```

For more information about the 400G ZR+ Optics, see the [Support for 400G ZR+ Optics on UfiSpace S9510-28DC Platform](#) section.

Port Configuration for 400G Optics

The table below provides a summary of the ports NIF for 400G ZR+ optics on the UfiSpace S9510-28DC platform in both modes (**mode_1_400G** and **mode_2_400G**).

400G ZR+ Optics Port Summary

Mode	Resource	Ports	Port Type	Speed	Bandwidth
mode_1_400G	400G	1	NIF	400	400
	100G	3	NIF	100	300
	25G	4	NIF	25	100
	10G	20	NIF	10	200
	OAM	1	Internal	10	10
	OLP	1	Internal	10	10
	Eventor	1	Internal	10	10
	ERP	1	Internal	10	10
	KR Port	4	Internal	10	40
	SAT Interface	1	Internal	100	100

Mode	Resource	Ports	Port Type	Speed	Bandwidth
mode_2_400G	400G	2	NIF	400	800
	100G	2	NIF	100	200
	OAM	1	Internal	10	10
	OLP	1	Internal	10	10
	Eventor	1	Internal	10	10
	ERP	1	Internal	10	10
	KR Port	4	Internal	10	40
	SAT Interface	1	Internal	100	100



The **mode_1_400G** mode supports one 400G port and all twenty-eight ports can be used, whereas in the **mode_2_400G** mode, two 400G ports are supported and only four ports can be used.

The **show interface physical** command output below shows the ports that are available for UfiSpace S9510-28DC platform in the **mode_2_400G** mode.

```
supervisor@rtbrick.net: op> show interface physical
Interface  Admin  Link  Oper  Damp  MAC Address      Speed  Duplex  Uptime
ifp-0/0/0  Up     Up    Up    stable  5c:07:58:78:8f:f9  400G   Full    Tue Dec 02 10:12:59 GMT +0000 2025
ifp-0/0/1  Up     Up    Up    stable  5c:07:58:78:90:01  400G   Full    Tue Dec 02 10:12:59 GMT +0000 2025
ifp-0/0/2  Up     Up    Up    stable  5c:07:58:78:90:09  100G   Full    Tue Dec 02 10:12:53 GMT +0000 2025
ifp-0/0/3  Up     Up    Up    stable  5c:07:58:78:90:0d  100G   Full    Tue Dec 02 10:12:53 GMT +0000 2025
```

Viewing Transceiver Details

The operational commands such as **show transceiver**, **show interface physical**, and **show transceiver <interface>** displays current status, high-level inventory, and detailed diagnostic information for the 400G optics and physical network interfaces on the UfiSpace S9510-28DC platform.

Viewing Transceiver Summary:

The **show transceiver** command provides a list of all installed transceivers across various interfaces on the platform, detailing their module type, name, connector type, power class, vendor, part number, and serial number.

```
supervisor@rtbrick.net: op> show transceiver
Interface  Module Type  Transceiver Name  Connector Type  PC  Vendor  Part Number
Serial Number  Material Number
ifp-0/0/0  QSFP-DD      ZR400-OFEC-16QAM  LC (Lucent Connector)  8(14.0W)  FLEXOPTIX  D.C0164HG.16.yTM
F824HMC    n/a
ifp-0/0/1  QSFP-DD      ZR400-OFEC-16QAM  LC (Lucent Connector)  8(14.0W)  FLEXOPTIX  D.C0164HG.16.yTM
F824GSF    n/a
ifp-0/0/2  QSFP28       100GBASE-LR4      LC (Lucent Connector)  5(4.0W)   LambdaGain  LL1S31B0A
L16C8M0034 T-UNIQSF40907039
ifp-0/0/3  QSFP28       100GBASE-LR4      LC (Lucent Connector)  5(4.0W)   LambdaGain  LL1S31B0A
L16D100019 T-UNIQSF40907039
```

ifp-0/0/4	SFP	10GBASE-CR Passive	Copper Pigtail	n/a	FS	SFP-PC01
S2108213165-1	n/a					

Viewing Physical Interfaces

The `show interface physical` command displays the status of all physical interfaces on the platform, including their administrative and operational state, link status, MAC address, configured speed and duplex, and uptime. This summary confirms that the interfaces (like the 400G and 100G ports) are currently operational and stable.

```
supervisor@rtbrick.net: op> show interface physical
```

Interface	Admin	Link	Oper	Damp	MAC Address	Speed	Duplex	Uptime
ifp-0/0/0	Up	Up	Up	stable	5c:07:58:78:8f:f9	400G	Full	Tue Nov 11 09:48:51 GMT +0000 2025
ifp-0/0/1	Up	Up	Up	stable	5c:07:58:78:90:01	400G	Full	Tue Nov 11 09:48:51 GMT +0000 2025
ifp-0/0/2	Up	Up	Up	stable	5c:07:58:78:90:09	100G	Full	Tue Nov 11 09:48:45 GMT +0000 2025
ifp-0/0/3	Up	Up	Up	stable	5c:07:58:78:90:0d	100G	Full	Tue Nov 11 09:48:45 GMT +0000 2025
cpu-0/0/200	Up	Up	Up	stable	5c:07:58:78:90:c0	100G	Full	Tue Nov 11 09:48:43 GMT +0000 2025
recycle-0/0/211	Up	Up	Up	stable	5c:07:58:78:90:cb	100G	Full	Tue Nov 11 09:48:43 GMT +0000 2025
recycle-0/0/212	Up	Up	Up	stable	5c:07:58:78:90:cc	100G	Full	Tue Nov 11 09:48:43 GMT +0000 2025
recycle-0/0/213	Up	Up	Up	stable	5c:07:58:78:90:cd	100G	Full	Tue Nov 11 09:48:43 GMT +0000 2025
recycle-0/0/214	Up	Up	Up	stable	5c:07:58:78:90:ce	100G	Full	Tue Nov 11 09:48:43 GMT +0000 2025
recycle-0/0/215	Up	Up	Up	stable	5c:07:58:78:90:cf	100G	Full	Tue Nov 11 09:48:43 GMT +0000 2025

Viewing 400G Transceiver Details

This example below provides detailed information for the 400G QSFP-DD optical transceiver named ZR400-OFEC-16QAM, including its vendor, part number, power class, connector type (LC), and capability to support a 600km link.

```
supervisor@rtbrick.net: op> show transceiver ifp-0/0/0
ifp-0/0/0
  Transceiver name: ZR400-OFEC-16QAM, Vendor name: FLEXOPTIX, Module type: QSFP-DD, Serial number: F824HMC
  Part number: D.C0164HG.16.yTM, Material number: n/a
  Power class: 8(14.0W), Power class mode: HIGH
  Connector type: LC (Lucent Connector), Link length support: 600km(sm), Wavelength: 1552.50 nm
  Capabilities: [Tx-Disable, Diagnostic Monitoring, Supply Voltage Monitoring, Temperature Monitoring, Transmitter Power
Measurement]
  Manufactured date: Mon Jul 15 00:00:00 GMT +0000 2024
  Transceiver Diagnostic Monitoring
    Attribute      Value      High Alarm  High Warn  Low Alarm  Low Warn
    Temperature    45.9453 C  False      False     False     False
    Voltage        3.2724 V  False      False     False     False
  Transceiver Lane Diagnostic Monitoring
    Lane: 0
    Attribute      Value      High Alarm  High Warn  Low Alarm  Low Warn
    RxPower        -15.2290 dBm False      False     False     False
    TxPower        -0.1060 dBm False      False     False     False
    TxBiascurrent  53.1600 mA  False      False     False     False
```

2.2. Platform Hardware Information

2.2.1. RBFS Multiservice Edge Image on Edgecore CSR440 (AS7535-28XB)

The RBFS Multiservice Edge is a software image that supports full BNG functionality on a single

image. This image is supported on the Edgecore CSR440 (AS7535-28XB) platform.

Hardware Specification

Edgecore CSR440 (AS7535-28XB) Hardware Specification

Model	Edgecore CSR440 (AS7535-28XB)
Form-factor	1RU, 19 Inch, Rack-Mountable
Switching Capacity	800 Gbps.
Switch ASIC	Broadcom Qumran-2A BCM88483
Co-Processor	—
CPU	Intel Broadwell (8-Core)
Role	Multiservice Edge
System Memory	DDR4 SO-DIMM 2x 8GB SDRAM with ECC support
Storage (SSD)	128 GB
Interfaces	• 24 x SFP28 (each supports 1/10 GbE or 25 GbE) • 2 x 100G QSFP28 (each supports 50/100 GbE) • 2 x 400G QSFP-DD (each supports 50/100/200/400 GbE)
Max. Number of LAG Interfaces Supported Per LAG	28
Number of LAG Members Supported Per LAG	10

For more information, click the link below.

<https://www.edge-core.com/productsInfo.php?cls=291&cls2=342&cls3=343&id=1004>

2.2.2. RBFS Multiservice Edge Image on UfiSpace S9510-28DC

The RBFS Multiservice Edge is a software image that supports full BNG functionality on a single image. This image is supported on the UfiSpace S9510-28DC platform.



RBFS is supported only on the 'Premium' variant of this platform. It is not supported on the 'Standard' variant. For detailed specifications of these variants, refer to the [ODM's Datasheet](#).

Hardware Specification

UfiSpace S9510-28DC Hardware Specification

Model	UfiSpace S9510-28DC
Form-factor	1RU
Switching Capacity	800 Gbps.
Switch ASIC	Broadcom Qumran-2A BCM88483

CPU	Premium: Intel Denverton-NS 8-Core @ 1.7GHz
Role	Multiservice Edge
System Memory	16GB DDR4
Storage (SSD)	Premium: 128GB
Interfaces	• 2 x 400G • 2 x 100G • 24 x 25G
Max. Number of LAG Interfaces Supported Per LAG	28
Number of LAG Members Supported Per LAG	10

For more details on the hardware specifications, refer to the [ODM's Datasheet](#).

Support for 400G ZR+ Optics on UfiSpace S9510-28DC Platform

RBFS supports 400G ZR+ Optics on UfiSpace S9510-28DC Platform. The table below lists the key specifications for the 400G ZR+ Optical Transceiver.

400G ZR+ Optical Transceiver Specifications

Supported Platforms	Product	Module Type Number	Product Type	Cable Type
UfiSpace S9510-28DC	Optical Transceiver	400G QSFP-DD	Optical Transceiver	SMF

For details about configuring 400G ZR+ optics on the UfiSpace S9510-28DC platform, see the [Configuring Platform Profiles](#) section.

2.2.3. RBFS Access Leaf & Multiservice Edge Images on Edgecore AGR420 (AS7946-74XKSB)

The RBFS Access Leaf is a software image that supports subscriber termination functionality on the Leaf Switch in a Spine Leaf deployment for BNG. This image is supported on the Edgecore AGR420 (AS7946-74XKSB) platform.

Hardware Specification

Edgecore AGR420 Hardware Specification

Model	Edgecore AGR420 (AS7946-74XKSB)
Form-factor	2RU, 19 Inch, Rack-Mountable
Switching Capacity	2.4 Tbps. NOTE: 200 Gbps bandwidth is reserved internally, and hence 2.2 Tbps is available for external use.
Switch ASIC	Broadcom Qumran-2C BCM88820

Co-Processor	BCM16K
CPU	Intel Broadwell (8-Core)
Role	Access Leaf, Multiservice Edge
System Memory	2 x 16 GB
Storage (SSD)	128 GB
Interfaces	• 10 x 100G • 64 x 25G
Max. Number of LAG Interfaces Supported Per LAG	24
Number of LAG Members Supported Per LAG	10

For more information, click the link below.

<https://www.edge-core.com/solution-inquiry.php?cls=5&id=129>

2.2.4. RBFS Access Leaf and Multiservice Edge Images on UfiSpace S9600-72XC

The RBFS Access Leaf is a software image that supports subscriber termination functionality on the Leaf Switch in a Spine Leaf deployment for BNG. The RBFS Multiservice Edge is a software image that supports full BNG functionality on a single image. Both these images are supported on the UfiSpace S9600-72XC platform.



RBFS is supported on this platform only with the Co-Processor 'Broadcom OP2 BCM16K' (Premium). Other variants of co-processor are not supported.

Hardware Specification

UfiSpace S9600-72XC Hardware Specification

Model	UfiSpace S9600-72XC
Form-factor	2RU, 436W x 87.7H x 609.6D mm (17.17"x3.45"x24")
Switching Capacity	2.4 Tbps. NOTE: 200 Gbps bandwidth is reserved internally, and hence 2.2 Tbps is available for external use.
Switch ASIC	Broadcom Qumran-2C BCM88820
Co-Processor	BCM16K
CPU	Intel Skylake-D D-2145NT 8 Cores @1.9GHz
Role	Access Leaf, Multiservice Edge
Storage (SSD)	128 GB
System Memory	2x 16GB DDR4 R-DIMM with ECC

Interfaces	• 64 x 25GE SFP28 ports • 8 x 100GE QSFP28 ports • 2 x 10GE SFP+ management ports • 1 x RJ45 serial console port
Max. Number of LAG Interfaces Supported Per LAG	72
Number of LAG Members Supported Per LAG	10

For more information, click the link below.

<https://ufispace.com/products/telco/aggregation/s9600-72xc-25g-100g-open-aggregation-router-tcam>

2.2.5. RBFS Spine Image on Edgecore AGR400 (AS7946-30XB)

The RBFS Spine is a software image that supports aggregation functionality across the access leaves in a Spine Leaf deployment for BNG. This image is supported on Edgecore AGR400 (AS7946-30XB) platform.

Hardware Specification

Edgecore AGR400 Hardware Specification

Model	EdgeCore AGR400 (AS7946-30XB)
Form-factor	2RU, 19 Inch, Rack-Mountable
Switching Capacity	2.4 Tbps
Switch ASIC	Broadcom Qumran-2C BCM88823
Co-Processor	-
CPU	Intel Broadwell (8-Core)
Role	Spine
System Memory	SDRAM DDR4 SO-DIMM 32GB (16 GB x 2)
Storage (SSD)	128 GB
Interfaces	• 26 x 100G • 4 x 25G
Max. Number of LAG Interfaces Supported Per LAG	33
Number of LAG Members Supported Per LAG	10

For more information, click the link below.

<https://www.edge-core.com/solution-inquiry.php?cls=5&id=129>

2.2.6. RBFS Spine Image on UfiSpace S9600-32X

The RBFS Spine is a software image that supports aggregation functionality across the access leaves in a Spine Leaf deployment for BNG. This image is supported on UfiSpace S9600-32X platform.

Hardware Specification

UfiSpace S9600-32X Hardware Specification

Model	UfiSpace S9600-32X
Form-factor	2RU, 436W x 87.8H x 762D mm (17.17"x3.46"x30")
Switching Capacity	2.4 Tbps
Switch ASIC	Broadcom Qumran-2C BCM88820
CPU	Intel Skylake-D D-2145NT 8 Core @1.9GHz
Role	Spine
System Memory	1x32GB DDR4 with ECC
Storage (SSD)	128GB
Interfaces	• 31 x 40GE/100GE QSFP28 ports • 4 x 1GE/10GE/25GE SFP28 ports (break out from Port 0) • 1 x RJ45 serial console port
Max. Number of LAG Interfaces Supported Per LAG	32
Number of LAG Members Supported Per LAG	10

For more information on the UfiSpace S9600-32X platform, click [here](#).

2.2.7. RBFS Access Leaf and Multiservice Edge Images on UfiSpace S9600-102XC

The RBFS Access Leaf is a software image that supports subscriber termination functionality on the Leaf Switch in a Spine Leaf deployment for BNG. The RBFS Multiservice Edge is a software image that supports full BNG functionality on a single image. Both these images are supported on the UfiSpace S9600-102XC platform.

Hardware Specification

UfiSpace S9600-102XC Hardware Specification

Model	UfiSpace S9600-102XC
Form-factor	2RU in height with physical dimensions of: Width 436mm (17.16"), Depth 609.6mm (24"), and Height 87.7mm (3.45")
Switching Capacity	2.4 Tbps NOTE: 200 Gbps bandwidth is reserved internally, and hence 2.2 Tbps is available for external use.

Switch ASIC	Broadcom Qumran-2C BCM88820
Co-Processor	BCM16K
CPU	Intel Skylake-D D-2145NT 8 core / 1.9GHz
Role	Access Leaf, Multiservice Edge
Storage (SSD)	128 GB
System Memory	16GB x2 DDR4 ECC RDIMM
Interfaces	• 96x25G • 6x100G
Max. Number of LAG Interfaces Supported Per LAG	102
Number of LAG Members Supported Per LAG	10

For more information, click the link below.

<https://www.ufispace.com/products/telco/aggregation/s9600-102xc-25g-100g-open-aggregation-router>

2.2.8. RBFS L2 Wholesale (L2BSA) Image on UfiSpace S9500-22XST

The RBFS L2BSA is a software image that supports transparent forwarding on the A10-NSP interface received from U Interface and vice versa.

Hardware Specification

UfiSpace S9500-22XST Hardware Specification

Model	UfiSpace S9500-22XST
Form-factor	1RU, 440w x 43.5h x 302d mm (17.32" x 1.713" x 11.89")
Switching Capacity	300 Gbps
Switch ASIC	Broadcom Qumran-AX BCM88470
Co-Processor	—
CPU	Intel Broadwell-DE D1519 4 Cores @1.5GHz
Role	L2 Wholesale (L2BSA)
System Memory	1x8GB DDR4 SO-DIMM with ECC
Storage (SSD)	32GB
Interfaces	• 2 x 100GE QSFP28 port • 8 x 25GE SFP28 ports • 8 x 10GE SFP+ ports • 4 x 1GE RJ45 ports

Max. Number of LAG Interfaces Supported Per LAG	22
Number of LAG Members Supported Per LAG	10

For more information, click the link below.

<https://www.ufispace.com/products/telco/access/s9500-22xst-rj45-disaggregated-cell-site-gateway>

2.3. Firmware Versions

To check the firmware versions for a specific platform, connect to the device, access the CLI prompt, and run the **show firmware** command. Additionally, the **show platform** command can be used to display the hardware platform information, which includes the diagnostic firmware (DIAG) version.

The table below shows the firmware versions that are embedded in the hardware platforms supported by RBFS Release 25.4.1.4.

Firmware Versions

Platform	DIAG Version	Firmware Versions
S9600-102XC UfiSpace Q2C (BCM88820)	0.0.2	<pre> x86_64-ufispace_s9600_102xc-r0 CPU CPLD: 1.17.5 Mainboard CPLD1: 0.0.9 Mainboard CPLD2: 0.0.1 Mainboard CPLD3: 0.0.1 Mainboard CPLD4: 0.0.1 Mainboard CPLD5: 0.0.1 BIOS: SKLD_BIOS_R04.10 BMC: 2.39.0 HW ID: 2 DEPH ID: 0 SKU ID: 28 ID type: 0 Build ID: 0 MU Version: v2.1.0 </pre>

Platform	DIAG Version	Firmware Versions
S9600-72XC UfiSpace Q2C (BCM88820)	0.5.5	x86_64-ufispace_s9600_72xc-r0 CPU CPLD: 1.17.5 Mainboard CPLD1: 1.9.12 Mainboard CPLD2: 0.9.17 Mainboard CPLD3: 0.9.17 Mainboard CPLD4: 0.18.9 BIOS: SKLD_BIOS_R04.12 BMC: 3.5.0 HW ID: 3 DEPH ID: 0 SKU ID: 6 ID type: 0 Build ID: 4 MU Version: v3.3.0
S9600-32X UfiSpace Q2C (BCM88820)	0.1.8	x86_64-ufispace_s9600_32x-r0 CPU CPLD: 1.17.5 Mainboard CPLD1: 0.21.6 Mainboard CPLD2: 0.11.4 Mainboard CPLD3: 0.10.4 BIOS: SKLD_BIOS_R04.12 BMC: 3.37.0 HW ID: 3 Build ID: 7 MU Version: v3.0.2
AGR420 AS7946-74XKSB EdgeCore Q2C (BCM88820)	02.0b.00.01_OP 2_ILKN_202202 24	x86_64-accton_as7946_74xkb-r0 CPU CPLD: 225 System CPLD: 23308 Fan CPLD: 513 BMC: 0.1.0
AGR400 AS7946-30XB EdgeCore Q2C (BCM88823)	02.0a.00.06	x86_64-accton_as7946-30xb-r0 CPU CPLD: 217 System CPLD: 23299 Fan CPLD: 513 BMC: 0.1.0

Platform	DIAG Version	Firmware Versions
S9510-28DC UfiSpace Q2A (BCM88483)	DIAG_DNX_4.1.4	x86_64-ufispace_s9510_28dc-r0 Mainboard CPLD: 0.5.3 BIOS: SIADLite_C3000_SERIES_R02.15 BMC: 5.1.0 HW Rev: 3 HW Build: 1 MU Version: v3.2.1
CSR440 AS7535-28XB EdgeCore Q2A (BCM88483)	02.02.00.07	x86_64-accton_as7535_28xb-r0 CPLD: 2 FPGA: 20997 Fan CPLD: 513 BMC: 0.3.0
S9500-22XST UfiSpace QAX (BCM88470)	4.0.21-1	x86_64-ufispace_s9500_22xst-r0 CPU CPLD: 21 Mainboard CPLD: 2 BIOS: T770994T01_R04.13 BMC: 4.8.0 HW Rev: 3 HW Build: 0 UCD: R0.2 MU Version: v3.0.1

2.4. Feature Support Matrix

2.4.1. Overview

RtBrick supports the following images (also known as roles).

- [Access-Leaf Image](#)
- [Multiservice Edge Image](#)
- [Spine Image](#)
- [L2 Wholesale \(L2BSA\) Image](#)

The following sections provide information about what RtBrick features are supported by respective images for each hardware platform.

2.4.2. Access-Leaf Image

The following table shows the RBFS feature support for access-leaf images.

Access-Leaf Images Feature Support

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	EdgeCore AGR420 (AS7946-74XKSB) (Q2C)
Routing Protocols	BGP	Yes	Yes	Yes
	BGP FlowSpec	Yes	Yes	Yes
	BGP RPKI	Yes	Yes	Yes
	IS-IS	No	No	No
	LDP	No	No	No
	OSPFv2/v3	No	No	No
	Policy	Yes	Yes	Yes
	Segment Routing (MPLS)	Yes	Yes	Yes
Layer 2 Services	L2X (Local & Remote)	Yes	Yes	Yes
	EVPN-VPWS	Yes	Yes	Yes
	BGP-signaled L2VPN	Yes	Yes	Yes
	LDP-signaled L2VPN	Yes	Yes	Yes
Layer 3 Services	L3VPN	Yes	Yes	Yes

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	EdgeCore AGR420 (AS7946-74XKSB) (Q2C)
Forwarding	HQoS	Yes	Yes	Yes
	Multifield (MF) Classifier	Yes	Yes	Yes
	OAM (IP Ping & Traceroute)	Yes	Yes	Yes
	OAM MPLS Ping	No	No	No
	Ethernet Connectivity Fault Management (CFM)	Yes	Yes	Yes
	BFD	Yes	Yes	Yes
	LLDP	Yes	Yes	Yes
	Inband Management	Yes	Yes	Yes
	LAG (Static, LACP)	Yes	Yes	Yes
	Mirroring	Yes	Yes	Yes
	ECMP & Hierarchical FEC	Yes	Yes	Yes
Multicast	IGMPv2/v3	Yes	Yes	Yes
	PIM-SSM	Yes	Yes	Yes

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	EdgeCore AGR420 (AS7946-74XKSB) (Q2C)
Subscriber Management	PPPoE	Yes	Yes	Yes
	L3 Wholesale (L2TPv2 LAC)	Yes	Yes	Yes
	IPoE	Yes	Yes	Yes
	AAA (RADIUS)	Yes	Yes	Yes
	Dual Stack	Yes	Yes	Yes
	Multicast for IPTV	Yes	Yes	Yes
	L2 Wholesale (L2BSA)	Yes	Yes	Yes
	Lawful Intercept	Yes	Yes	Yes
	Accounting	Yes	Yes	Yes
	PPPoE Single-/double-tagged interfaces	Yes	Yes	Yes
	IPoE Single-/double-tagged interfaces	Yes	Yes	Yes
	PPPoE Untagged Interfaces	Yes	Yes	Yes
	IPoE Untagged Interfaces	Yes	Yes	Yes
	IPoE Hot Standby Redundancy	No	No	No
	Subscriber Filters	Yes	Yes	Yes
	HTTP Redirect Service	Yes	Yes	Yes

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	EdgeCore AGR420 (AS7946-74XKSB) (Q2C)
Infrastructure	Logging	Yes	Yes	Yes
	NTP	Yes	Yes	Yes
	Clock Synchronization G.8275.1	Yes	Yes	No
	LED Control	Yes	Yes	Yes
	IPMI	Yes	Yes	Yes
	Watchdog Timer	Yes	Yes	Yes
Security	Securing the Management Plane	Yes	Yes	Yes
	Securing the Control Plane	Yes	Yes	Yes
	Local User Management	Yes	Yes	Yes
	BGP FlowSpec	Yes	Yes	Yes
	BGP RPKI	Yes	Yes	Yes
	BGP TCP-AO / MD5	Yes	Yes	Yes
	BGP GTSM	Yes	Yes	Yes
	LDP TCP-AO / MD5	Yes	Yes	Yes
	sFlow	Yes	Yes	Yes

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	EdgeCore AGR420 (AS7946-74XKSB) (Q2C)
Telemetry	Resmon	Yes	Yes	Yes
	ASIC Resource Monitoring	Yes	Yes	Yes
	Prometheus TSDB	Yes	Yes	Yes
	SNMPv2c/SNMPv3	Yes	Yes	Yes

2.4.3. Multiservice Edge Image

The following table shows the RBFS feature support for Multiservice Edge images.

Multiservice Edge Images Feature Support

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	Edgecore AGR420 (AS7946-74XKSB) (Q2C)	UfiSpace S9510-28DC (Q2A)	Edgecore CSR440 (AS7535-28XB) (Q2A)
Routing Protocols	BGP	Yes	Yes	Yes	Yes	Yes
	BGP FlowSpec	Yes	Yes	Yes	Yes	Yes
	BGP RPKI	Yes	Yes	Yes	Yes	Yes
	IS-IS	Yes	Yes	Yes	Yes	Yes
	LDP	Yes	Yes	Yes	Yes	Yes
	OSPFv2/v3	Yes	Yes	Yes	Yes	Yes
	Policy	Yes	Yes	Yes	Yes	Yes
	Segment Routing (MPLS)	Yes	Yes	Yes	Yes	Yes

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	Edgecore AGR420 (AS7946-74XKSB) (Q2C)	UfiSpace S9510-28DC (Q2A)	Edgecore CSR440 (AS7535-28XB) (Q2A)
Layer 2 Services	L2X (Local & Remote)	Yes	Yes	Yes	Yes	Yes
	EVPN-VPWS	Yes	Yes	Yes	Yes	Yes
	BGP-signaled L2VPN	Yes	Yes	Yes	Yes	Yes
	LDP-signaled L2VPN	Yes	Yes	Yes	Yes	Yes

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	Edgecore AGR420 (AS7946-74XKSB) (Q2C)	UfiSpace S9510-28DC (Q2A)	Edgecore CSR440 (AS7535-28XB) (Q2A)
Forwarding	HQoS	Yes	Yes	Yes	Yes	Yes
	Multifield (MF) Classifier	Yes	Yes	Yes	Yes	Yes
	OAM (IP Ping & Traceroute)	Yes	Yes	Yes	Yes	Yes
	OAM MPLS Ping	Yes	Yes	Yes	Yes	Yes
	Ethernet Connectivity Fault Management (CFM)	Yes	Yes	Yes	Yes	Yes
	BFD	Yes	Yes	Yes	Yes	Yes
	LLDP	Yes	Yes	Yes	Yes	Yes
	Inband Management	Yes	Yes	Yes	Yes	Yes
	LAG (Static, LACP)	Yes	Yes	Yes	Yes	Yes
	Mirroring	Yes	Yes	Yes	Yes	Yes
	ECMP & Hierarchical FEC	Yes	Yes	Yes	Yes	Yes
Multicast	IGMPv2/v3	Yes	Yes	Yes	Yes	Yes
	PIM-SSM	Yes	Yes	Yes	Yes	Yes

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	Edgecore AGR420 (AS7946-74XKSB) (Q2C)	UfiSpace S9510-28DC (Q2A)	Edgecore CSR440 (AS7535-28XB) (Q2A)
Subscriber Management	PPPoE	Yes	Yes	Yes	Yes	Yes
	L3 Wholesale (L2TPv2 LAC)	Yes	Yes	Yes	Yes	Yes
	IPoE	Yes	Yes	Yes	Yes	Yes
	AAA (RADIUS)	Yes	Yes	Yes	Yes	Yes
	Dual Stack	Yes	Yes	Yes	Yes	Yes
	Multicast for IPTV	Yes	Yes	Yes	Yes	Yes
	L2BSA (L2 Wholesale)	Yes	Yes	Yes	No	No
	Lawful Intercept	Yes	Yes	Yes	Yes	Yes
	Accounting	Yes	Yes	Yes	Yes	Yes

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	Edgecore AGR420 (AS7946-74XKSB) (Q2C)	UfiSpace S9510-28DC (Q2A)	Edgecore CSR440 (AS7535-28XB) (Q2A)
Subscriber Management (Cont'd)	PPPoE Single- /double-tagged interfaces	Yes	Yes	Yes	Yes	Yes
	IPoE Single-/double- tagged interfaces	Yes	Yes	Yes	Yes	Yes
	PPPoE Untagged Interfaces	Yes	Yes	Yes	No	No
	IPoE Untagged Interfaces	Yes	Yes	Yes	No	No
	IPoE Hot Active- Active Redundancy	Yes	Yes	Yes	No	No
	IPoE Hot Active- Standby Redundancy	Yes	Yes	Yes	Yes	Yes
Subscriber Management (Cont'd)	Subscriber Filters	Yes	Yes	Yes	Yes	Yes
	HTTP Redirect Service	Yes	Yes	Yes	Yes	Yes
	Carrier-Grade NAT	Yes	Yes	Yes	Yes	Yes
	Carrier-Grade NAT Appliance	Yes	Yes	No	Yes	Yes

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	Edgecore AGR420 (AS7946-74XKSB) (Q2C)	UfiSpace S9510-28DC (Q2A)	Edgecore CSR440 (AS7535-28XB) (Q2A)
Infrastructure	Logging	Yes	Yes	Yes	Yes	Yes
	NTP	Yes	Yes	Yes	Yes	Yes
	Clock Synchronization G.8275.1	No	No	No	No	No
	LED Control	Yes	Yes	Yes	Yes	Yes
	IPMI	Yes	Yes	Yes	Yes	No
	Watchdog Timer	Yes	Yes	Yes	Yes	Yes
Security	Securing the Management Plane	Yes	Yes	Yes	Yes	Yes
	Securing the Control Plane	Yes	Yes	Yes	Yes	Yes
	Local User Management	Yes	Yes	Yes	Yes	Yes
	BGP FlowSpec	Yes	Yes	Yes	Yes	Yes
	BGP TCP-AO / MD5	Yes	Yes	Yes	Yes	Yes
	BGP RPKI	Yes	Yes	Yes	Yes	Yes
	BGP GTSM	Yes	Yes	Yes	Yes	Yes
	LDP TCP-AO / MD5	Yes	Yes	Yes	Yes	Yes
	sFlow	Yes	Yes	Yes	Yes	Yes

Component	Feature	UfiSpace S9600-72XC (Q2C)	UfiSpace S9600-102XC (Q2C)	Edgecore AGR420 (AS7946-74XKSB) (Q2C)	UfiSpace S9510-28DC (Q2A)	Edgecore CSR440 (AS7535-28XB) (Q2A)
Telemetry	Resmon	Yes	Yes	Yes	Yes	Yes
	ASIC Resource Monitoring	Yes	Yes	Yes	Yes	Yes
	Prometheus TSDB	Yes	Yes	Yes	Yes	Yes
	SNMPv2c/SNMPv3	Yes	Yes	Yes	Yes	Yes

2.4.4. Spine Image

The following table shows the RBFS feature support for spine images.

Spine Images Feature Support

Component	Feature	UfiSpace S9600-32X (Q2C)	EdgeCore AGR400 (AS7946-30XB) (Q2C)
Routing Protocols	BGP	Yes	Yes
	BGP FlowSpec	Yes	Yes
	BGP RPKI	Yes	Yes
	IS-IS	Yes	Yes
	LDP	Yes	Yes
	OSPFv2/v3	Yes	Yes
	Policy	Yes	Yes
	Segment Routing (MPLS)	Yes	Yes

Component	Feature	UfiSpace S9600-32X (Q2C)	EdgeCore AGR400 (AS7946-30XB) (Q2C)
Layer 2 Services	L2X (Local & Remote)	Yes	Yes
	EVPN-VPWS	Yes	Yes
	BGP-signaled L2VPN	Yes	Yes
	LDP-signaled L2VPN	Yes	Yes
Forwarding	HQoS	Yes	Yes
	Multifield (MF) Classifier	Yes	Yes
	OAM (IP Ping & Traceroute)	Yes	Yes
	OAM MPLS Ping	Yes	Yes
	Ethernet Connectivity Fault Management (CFM)	Yes	Yes
	LLDP	Yes	Yes
	Inband Management	Yes	Yes
	LAG (Static, LACP)	Yes	Yes
	Mirroring	Yes	Yes
	Breakout Interfaces	Yes	Yes
	ECMP & Hierarchical FEC	Yes	Yes
Multicast	IGMPv2/v3	No	No
	PIM-SSM	Yes	Yes

Component	Feature	UfiSpace S9600-32X (Q2C)	EdgeCore AGR400 (AS7946-30XB) (Q2C)
Infrastructure	Logging	Yes	Yes
	NTP	Yes	Yes
	Clock Synchronization G.8275.1	No	No
	LED Control	Yes	Yes
	IPMI	Yes	No
	Watchdog Timer	Yes	Yes
Security	Securing the Management Plane	Yes	Yes
	Securing the Control Plane	Yes	Yes
	Local User Management	Yes	Yes
	BGP FlowSpec	Yes	Yes
	BGP RPKI	Yes	Yes
	BGP TCP-AO / MD5	Yes	Yes
	BGP GTSM	Yes	Yes
	LDP TCP-AO / MD5	Yes	Yes
	sFlow	Yes	Yes
Telemetry	Resmon	Yes	Yes
	ASIC Resource Monitoring	Yes	Yes
	Prometheus TSDB	Yes	Yes
	SNMPv2c/SNMPv3	Yes	Yes

2.4.5. L2 Wholesale (L2BSA) Image

The features listed below are tested using the switch with the L2BSA image as a L2SBA wholesale appliance.

The following table shows the RBFS feature support for L2 Wholesale (L2BSA) images.

L2 Wholesale (L2BSA) Images Feature Support

Component	Feature	UfiSpace S9500-22XST (QAX)
Routing Protocols	BGP	Yes
	BGP FlowSpec	No
	BGP RPKI	No
	IS-IS	Yes
	LDP	No
	OSPFv2/v3	No
	Policy	Yes
	Segment Routing (MPLS)	Yes
Layer 2 Services	L2X (Local & Remote)	Yes
	EVPN-VPWS	No
	BGP-signaled L2VPN	No
	LDP-signaled L2VPN	No

Component	Feature	UfiSpace S9500-22XST (QAX)
Forwarding	HQoS	Yes
	Multifield (MF) Classifier	No
	OAM (IP Ping & Traceroute)	Yes
	OAM MPLS Ping	No
	Ethernet Connectivity Fault Management (CFM)	No
	LLDP	Yes
	Inband Management	Yes
	LAG (Static, LACP)	Yes
	Mirroring	Yes
	ECMP & Hierarchical FEC	No
Infrastructure	Logging	Yes
	NTP	Yes
	Clock Synchronization G.8275.1	No
	LED Control	Yes
	IPMI	Yes
	Watchdog Timer	Yes

Component	Feature	UfiSpace S9500-22XST (QAX)
Security	Securing the Management Plane	Yes
	Securing the Control Plane	Yes
	Local User Management	Yes
	BGP FlowSpec	No
	BGP TCP-AO / MD5	No
	BGP GTSM	No
	LDP TCP-AO / MD5	No
	sFlow	No
Telemetry	Resmon	Yes
	ASIC Resource Monitoring	No
	Prometheus TSDB	Yes
	SNMPv2c/SNMPv3	Yes

2.5. Resource Limits/Feature Support

Limiting the resource usage or consumption (wherever applicable) helps to improve the system stability and also restricts over utilization of system capacity. In RBFS, the usage limits for the following resources are pre-defined:

- IPv4 Route Count
- IPv6 Route Count
- NAT44 Rule Count
- MTU Profile

- L3 MTU-Profile
- Subscriber MTU Profile
- Physical MTU Profile

In addition, you can track the resource usage of the following features.

- 6PE label
- High Precision QoS

You can specify resource limits on the following images (also known as roles).

- [Access-Leaf Image](#)
- [Multiservice Edge Image](#)
- [Spine Image](#)
- [L2 Wholesale \(L2BSA\) Image](#)

2.5.1. Access-Leaf Image

The following table provides the limits defined for the resources for the access-leaf images that RBFS supports.

Access-Leaf Images - Feature/Resource Usage

Component	Feature/Resource	UfiSpace S9600-102XC (Q2C)	UfiSpace S9600-72XC (Q2C)	EdgeCore AGR420 (AS7946-74XKSB) (Q2C)
FIB	IPv4 Route Count	1,200,000	1,200,000	1,200,000
	IPv6 Route Count	250,000	250,000	250,000
	Low Rate Shaping Enabled (<1000 Kbps)	Yes	Yes	Yes
	ECMP Maximum Next-hop Count	4	4	4
CONFD	MTU-Profile Count	8	8	8
	L3 MTU-Profile Count	3	3	3
	Subscriber MTU-Profile Count	6	6	6
	Physical MTU Profile Count	8	8	8
BGP	6PE label value	2	2	2
	BGP FlowSpec ACL	1,000	1,000	1,000
MPLS	MPLS Stack Depth	8	8	8

2.5.2. Multiservice Edge Image

The following table provides the limits defined for the resources for the Multiservice Edge images that RBFS supports.

Multiservice Edge Images - Feature/Resource Usage

Component	Feature	UfiSpace S9600-102XC (Q2C)	UfiSpace S9600-72XC (Q2C)	Edgecore AGR420 (AS7946-74XKSB) (Q2C)	UfiSpace S9510-28DC (Q2A)	EdgeCore CSR440 (AS7535-28XB) (Q2A)
FIB	IPv4 Route Count	1,200,000	1,200,000	1,200,000	1,200,000	1,200,000
	IPv6 Route Count	250,000	250,000	250,000	500,000	500,000
	NAT44 Rule Count	3,800,000	3,800,000	3,800,000	150,000	150,000
	Low Rate Shaping Enabled (<1,000 Kbps)	Yes	Yes	Yes	Yes	Yes
	ECMP Maximum Next-hop Count	4	4	4	2	2
CONFD	MTU-Profile Count	8	8	8	8	8
	L3 MTU-Profile Count	3	3	3	3	3
	Subscriber MTU-Profile Count	6	6	6	5	5
	Physical MTU Profile Count	8	8	8	8	8
BGP	6PE label value	2	2	2	2	2
	BGP FlowSpec ACL	1,000	1,000	1,000	1,000	1,000
MPLS	MPLS Stack Depth	8	8	8	8	8

2.5.3. Spine Image

The following table provides the limits defined for the resources for the spine images that RBFS supports.

Spine Images - Feature/Resource Usage

Component	Feature	UfiSpace S9600-32X (Q2C)	EdgeCore AGR400 (AS7946-30XB) (Q2C)
FIB	IPv4 Route Count	1,200,000	1,200,000
	IPv6 Route Count	250,000	250,000
	Low Rate Shaping Enabled	No	No
	ECMP Maximum Next-hop Count	4	4
CONFD	MTU-Profile Count	8	8
	L3 MTU-Profile Count	3	3
	Subscriber MTU-Profile Count	6	6
	Physical MTU Profile Count	8	8
BGP	6PE label value	2	2
	BGP FlowSpec ACL	1,000	1,000
MPLS	MPLS Stack Depth	8	8

2.5.4. L2 Wholesale (L2BSA) Image

The following table provides the limits defined for the resources for the L2BSA images that RBFS supports.

L2 Wholesale (L2BSA) Images - Feature/Resource Usage

Component	Feature	UfiSpace S9500-22XST (QAX)
FIB	IPv4 Route Count	—
	IPv6 Route Count	—
	ECMP Maximum Next-hop Count	—

Component	Feature	UfiSpace S9500-22XST (QAX)
CONFID	MTU-Profile Count	8
	L3 MTU-Profile Count	3
	Subscriber MTU-Profile Count	5
	Physical MTU Profile Count	8
BGP	6PE label value	—
	BGP FlowSpec ACL	—
MPLS	MPLS Stack Depth	2

2.5.5. Show CLI for Displaying ACL Table Information

The `show acl table` command displays the access control list table information. The options for the `show acl table` command varies depending on the platforms and their roles.

Syntax:

show acl table <options>

The following are the ACL table types supported across different platforms and their respective roles:

Supported CLI options on the Q2A Multiservice Edge platform:

Option	Description
l2	Displays the information for Layer 2 ACL table.
li	Displays the information for Lawful Interception (LI) ACL table.
l3v4	Displays the information for IPv4 Layer 3 ACL table.
l3v6	Displays the information for IPv6 Layer 3 ACL table.
policer	Displays information related to traffic policers configured in the system.

Supported CLI Option on the Q2C Multiservice Edge Platform:

Option	Description
l2	Displays the information for Layer 2 ACL table.
li	Displays the information for Lawful Interception (LI) ACL table.
pmf_mf_de_act	Displays the information for packet matching function with match fields for deactivation rules.
pmf_l3_de_act	Displays the information for packet matching function for Layer 3 deactivation rules.
ext_l3v4_mf	Displays the information for extended IPv4 Layer 3 ACL table.
ext_l3v6_mf	Displays the information for extended IPv6 Layer 3 ACL table.
ext_l3v4	Displays the information for the extended IPv4 Layer 3 ACL table.
ext_l3v6	Displays the information for extended IPv6 Layer 3 ACL table.
l3v4	Displays the information for the extended IPv4 Layer 3 ACL table.
ext_policer	Displays the information for ACL table type for applying policers on matched traffic.
policer_level	Displays the information for the policer levels.
pmf_policer_de_act	Displays the information for packet matching function for policer deactivation rules.

Example:

```
supervisor@rtbrick>rtbrick.net: cfg> show acl table
<cr>
ext_l3v4          ACL table type
ext_l3v4_mf       ACL table type
ext_l3v6          ACL table type
ext_l3v6_mf       ACL table type
ext_policer       ACL table type
l2                ACL table type
l3v4              ACL table type
li                ACL table type
pmf_l3_de_act     ACL table type
pmf_mf_de_act     ACL table type
pmf_policer_de_act ACL table type
policer_level     ACL table type
```

Supported CLI options on the Q2C Spine (UfiSpace) Platforms:

Option	Description
l2	Displays the information for Layer 2 ACL table.
l3v4	Displays the information for standard IPv4 Layer 3 ACL table.
l3v6	Displays the information for standard IPv6 Layer 3 ACL table.
policer	Displays information related to traffic policers configured in the system.
l2_cp	Displays information about Layer 2 control protocol packets.
sampling_egress	Displays egress packet sampling configuration information.

Supported CLI options on the Q2C Access Leaf (UfiSpace) Platforms:

Option	Description
l2	Displays the information for Layer 2 ACL table.
li	Displays the information for Lawful Interception (LI) ACL table.
pmf_mf_de_act	Displays the information for packet matching function with match fields for deactivation rules.
pmf_l3_de_act	Displays the information for packet matching function for Layer 3 deactivation rules.
ext_l3v4_mf	Displays the information for extended IPv4 Layer 3 ACL table.
ext_l3v6_mf	Displays the information for extended IPv6 Layer 3 ACL table.
ext_l3v4	Displays the information for the extended IPv4 Layer 3 ACL table.
ext_l3v6	Displays the information for extended IPv6 Layer 3 ACL table.
l3v4	Displays the information for standard IPv4 Layer 3 ACL table.

Option	Description
ext_policer	Displays the information for ACL table type for applying policers on matched traffic.
policer_level	Displays the information for the policer levels.
pmf_policer_de_act	Displays the information for packet matching function for policer deactivation rules.

Example:

```

supervisor@rtbrick.net: cfg> show acl table l3v4 direction ingress
Table name: L3V4_INGRESS
Profile id: default
Table stage: ipmf1
Table priority: 5
Handle: 40
Matches supported:
  source_ipv4_prefix
  destination_ipv4_prefix
  ipv4_protocol
  ipv4_tos
  ip_has_options
  ip_tunnel_type
  source_l4_port
  destination_l4_port
  forward_class
  in_vlan_port
  tcp_ctrl_flags
Actions supported:
  drop
  trap_id
  redirect_to_cpu
  policer_and_counter
  forward_class
  remark_codepoint
  fwd_type_ipv4
Contexts supported:
  fwd_type_ipv4_ipmf1

```

2.5.6. Static, System, and Dynamic Label Ranges

The static, system, and dynamic label ranges provide greater flexibility in label management. The total allocatable label range is from 0 to 1,048,575.

- **Static Label Range:** Labels in this range are manually configured by administrators and assigned to specific network service protocols.
- **System Label Range:** Labels in this range include the labels reserved for internal system functions. These labels are not available for user configuration.

- **Dynamic Label Range:** Labels in this range are automatically allocated and managed by network protocols, such as LDP. They are assigned as needed and may change dynamically.

Allocation of Label Ranges

Static, System, and Dynamic Label Allocation Ranges		
Static	System	Dynamic
0 - 499,968 and 999,936 - 1,048,575	499,969 - 519,935	519,936 - 999,935



- In the static range 0 - 499,968, only the range 256 - 499,968 can be used for allocation.
- Static and dynamic ranges can be viewed using the show pool block command.

The output of the show pool block command is displayed as follows:

```
supervisor@rtbrick>LEAF01: op> show pool block
Resource Type      Module Name      Block Size      Block Range
LABEL              poold            519936          0 - 519935
LABEL              poold            48640          999936 - 1048575
```

Registered Address	Support	Sales
40268, Dolerita Avenue Fremont CA 94539		
+1-650-351-2251		+91 80 4850 5445
http://www.rtbrick.com	support@rtbrick.com	sales@rtbrick.com

©Copyright 2024 RtBrick, Inc. All rights reserved. The information contained herein is subject to change without notice. The trademarks, logos and service marks ("Marks") displayed in this documentation are the property of RtBrick in the United States and other countries. Use of the Marks are subject to RtBrick's Term of Use Policy, available at <https://www.rtbrick.com/privacy>. Use of marks belonging to other parties is for informational purposes only.