RBFS Overview

Version 23.8.1, 12 September 2023

Table of Contents

1. RtBrick Full Stack (RBFS) Overview

RBFS At-a-Glance

RtBrick Full Stack (RBFS) is a disaggregated and open network operating system that is presently productized and available as a Broadband Network Gateway (BNG). RBFS acts as an access software for establishing and managing subscriber sessions for broadband subscribers. It aggregates traffic from various subscriber sessions and routes the traffic to the network of the service provider.

RBFS establishes and maintains a connection with the Customer Premise Equipment (CPE), so that subscribers can access and use the network services from a network service provider.

RBFS runs as an Ubuntu container on the Open Network Linux operating system on white boxes which can perform Layer 2 and Layer 3 switching.

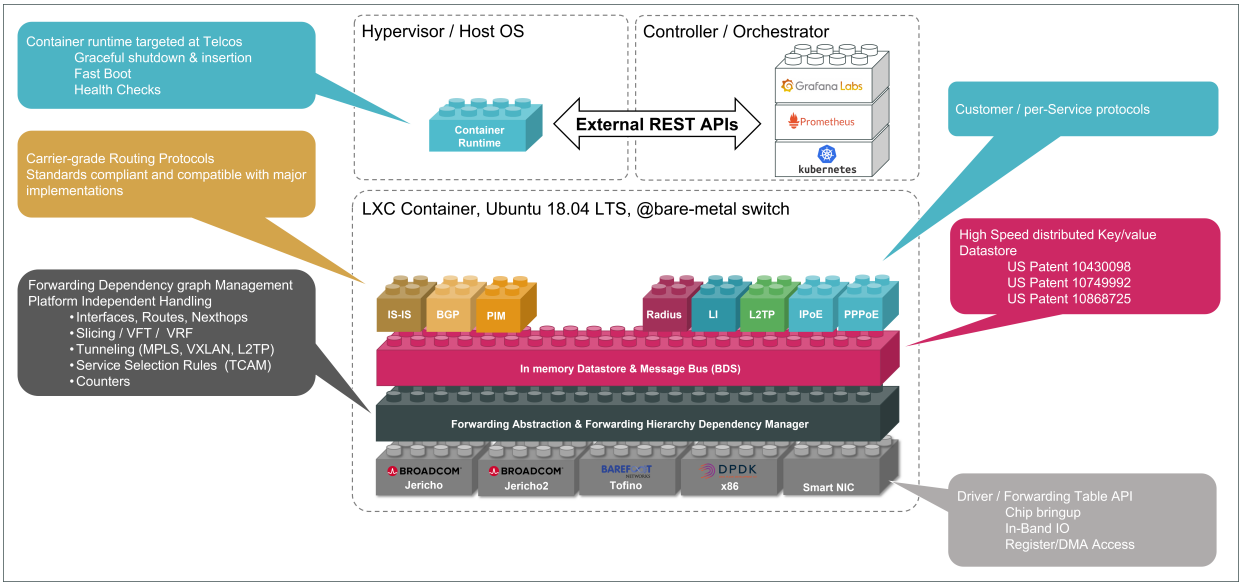

Why RBFS

RBFS' open and disaggregated architectural design fosters a faster deployment of new features and services within a short period and it promotes a collaborative ecosystem of hardware and other component vendors. By separating the hardware from the software, RBFS enables you to choose the white box switches of your choice without any vendor lock-in. It helps to reduce the deployment and

operational costs significantly by promoting disaggregated BNG that is suitable for cloud-native ecosystems.

RBFS, built on the microservices architecture, offers some key benefits compared to traditional monolithic systems. It offers greater agility and provides a higher degree of automation that reduces operational overheads. RBFS works well with continuous integration (CI) and continuous delivery (CD) practices and tools.

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).

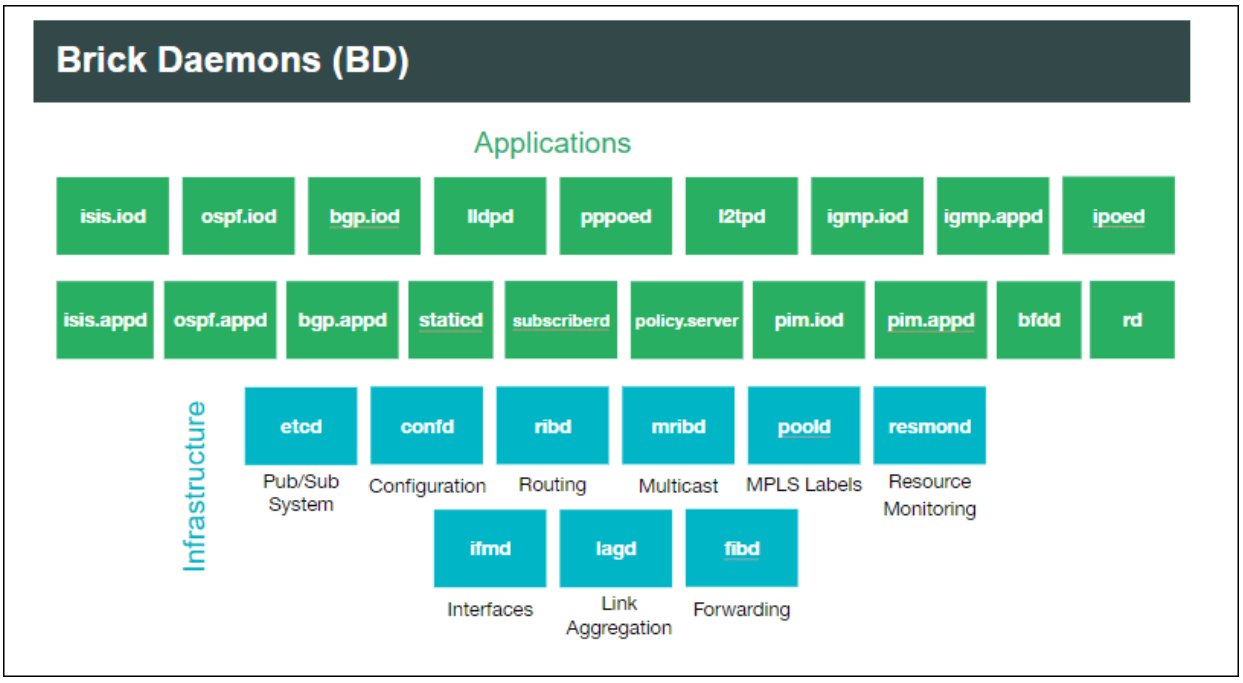

Brick Data Store

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.

Brick Daemons

RBFS microservices architecture allows decoupled daemons to serve various functionalities and services and they have their own realm of responsibilities to serve independently.

For example, the subscriber daemon (subscriberd) manages the current subscriber state and is responsible for authentication, authorization, and accounting. The ribd daemon is responsible for route selection, next-hop resolution, tunnel selection and recursion. The forwarding (fibd) daemon handles packet forwarding, route NH download, VPP and PD layer programming. Daemons such as confd and ifmd take care of various configurations and interface management respectively and together they all compose a comprehensive broadband session.

For the routing protocols such as BGP, there are two daemons - bgp.appd and bgp.iod - available to carry out the various functions of the protocol. The bgp.iod daemon manages sending and receiving of the BGP messages such as open, update, keepalive, and notification and takes care of session management. The best route that is selected by bgp.appd daemon is synced with the bgp.iod daemon so that the routes can further be advertised to other BGP peers.

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

In addition to the RtBrick daemons, you can deploy some other third-party applications in the container to bring additional capabilities to the system. For example, Prometheus is an open-source monitoring and alerting software that you can use for observability and monitoring purposes in the container.

Containerization of Daemons

RBFS daemons and other dependencies are packaged as an Ubuntu 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 Open Network Linux (ONL), an open-source operating system, which can be run on white box switches.

RBFS can perform various roles such as Spine, Leaf, and Consolidated BNG 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 Consolidated BNG 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 core Spine RBFS image must include (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 Access Leaf image.

Similarly, the Access Leaf image should include daemons (in addition to other daemons) such as subscriberd, l2tpd, pppoed, and ipoed which are not present in the Spine image.

You can see 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.

Supported Topologies

RBFS can be deployed in a spine-leaf architecture and can also be deployed standalone in a single switch by consolidating all the features in one 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 advantage of RBFS spine-leaf topology includes higher performance and better scalability. It is inherently scalable by providing many paths between any two points. This topology is easier for horizontal scaling by adding additional switches to add more capacity to handle increased traffic. This topology is also useful for low latency and higher bandwidth.

A consolidated BNG architecture offers all the functionalities of a spine-leaf BNG architecture on a single bare-metal switch. However, this architecture is recommended when there is a small concentration of broadband subscribers.

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.

RBFS supports REST-based industry-standard tools such as RESTCONF 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 daemon (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.

RBMS (RtBrick's Management System) is a GUI-based application that acts as a single pane of glass and allows interactions with RBFS for all operations, from provisioning and management to monitoring and debugging.

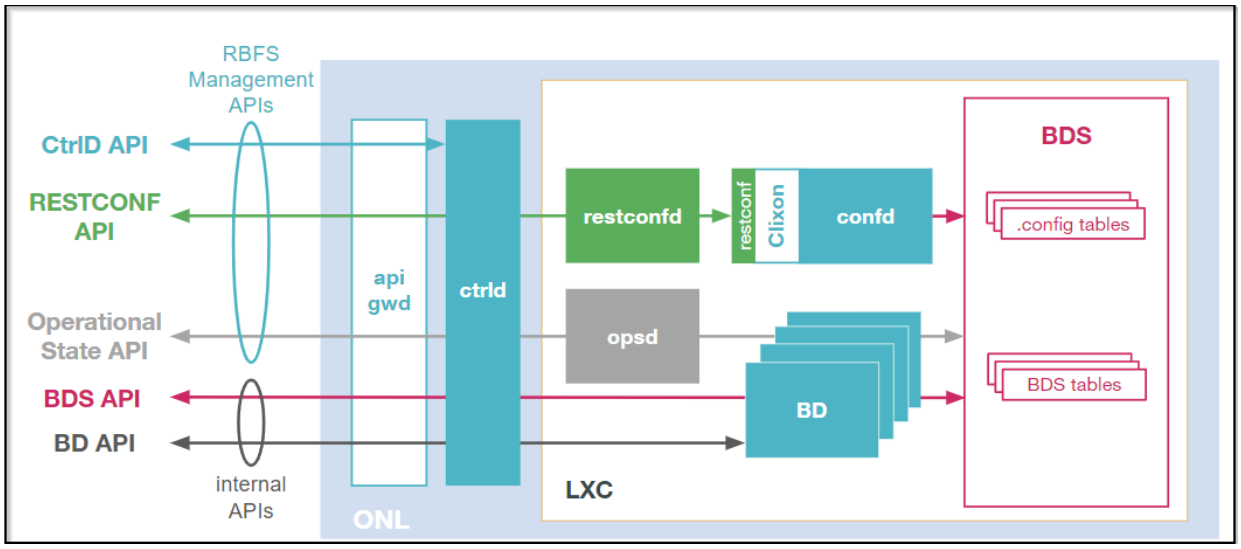

Features and Components

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.

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.

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: Policy Repository, Command Processing Module, Policy Server, and Policy Client.

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 daemon (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 daemon (ipoed) manages IPoE services using DHCPv4 and DHCPv6.

The L2TPv2 daemon (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.

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.

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.

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.

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.

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, 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 daemon validates the access token against a JSON web key set (JKWS).

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.

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 daemon 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.

Port Mirroring

RBFS supports port mirroring, a monitoring technique that can be implemented on network switches. Port mirroring allows to copy and send data packets from one port to another port for monitoring purposes. Port mirroring enables network administrators to troubleshoot the system with a protocol analyzer on the port that has the mirrored data.

RBFS Software Licensing

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 licensing and installation, see RBFS Software Licensing and Installation.

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