



Performance Routing Version 3

Performance Routing Version 3 (PfRv3) is the evolution of Performance Routing (PfR). PfRv3 is an intelligent-path control mechanism for improving application delivery and WAN efficiency. It protects critical applications, increases bandwidth utilization, and serves as an integral part of the Cisco Intelligent WAN (IWAN) solution. PfRv3 uses differentiated services code points (DSCP) and application-based policy framework to provide a multi-site aware bandwidth and path control optimization.

- [Feature Information for PfRv3, on page 1](#)
- [Hardware and Software Support, on page 2](#)
- [Restrictions for Configuring Performance Routing v3, on page 2](#)
- [Information About PfRv3, on page 3](#)

Feature Information for PfRv3

The following table provides release information about the feature or features described in this module. This table lists only the software release that introduced support for a given feature in a given software release train. Unless noted otherwise, subsequent releases of that software release train also support that feature.

Use Cisco Feature Navigator to find information about platform support and Cisco software image support. To access Cisco Feature Navigator, go to www.cisco.com/go/cfn. An account on Cisco.com is not required.

Table 1: Feature Information for Configuring PfRv3

Feature Name	Releases	Feature Information
PfRv3	15.4(3)M	<p>Performance Routing v3 (PfRv3) is the evolution of Performance Routing.</p> <p>PfRv3 is an intelligent-path control mechanism for improving application delivery and WAN efficiency. It protects critical applications, increases bandwidth utilization, and serves as an integral part of the Cisco Intelligent WAN (IWAN) solution.</p> <p>The following commands were modified by this feature: domain default, vrf default, master, source-interface, site-prefixes, password, monitor-interval, route-control, load-balance, enterprise-prefix, advanced, minimum-mask-length, mitigation-mode, threshold-variance, smart-probes, collector, class, match, priority, path-preference, border, domain-path.</p>

Hardware and Software Support

Cisco Performance Routing Version 3 (PfRv3) supports the following Cisco platforms and software releases:

Device	Cisco IOS Software Release	Hub/Remote Site
Cisco ISR 4000 Series Routers	Cisco IOS XE 3.13 or later	Hub or remote site
Cisco ASR 1000 Series Routers	Cisco IOS XE 3.13 or later	Hub site
Cisco CSR 1000v Series Routers	Cisco IOS XE 3.14 or later	Hub site (master controller only)
Cisco ISR-G2 Series Routers	Cisco IOS 15.5(1)T1 or later Cisco IOS 15.4(3)M1 or later	Remote site

Restrictions for Configuring Performance Routing v3

- Asymmetric routing is not supported for application-based policy.

- A new session cannot be established with application-based policy during blackout failure until route converges to backup path. For application-based flows, application ID is not recognized by Network Based Application Recognition (NBAR2) until session gets established and packet exchanges directly. You can configure Differentiated Services Code Point (DSCP) based policy for fast failover with blackout failure.
- PfRv3 does not support High Availability (HA) for both master and border routers. ESP switch over can trigger temporary unreachable event for one to two seconds.
- IPv6 is not supported.
- Network Address Translation (NAT) is not supported.
- Remarking DSCP for traffic flows on WAN interface is not supported.
- On a HUB Master Controller (MC), when a class is configured for matching application within a PfRv3 domain, the list of NBAR application names are limited if there is no active Border Router (BR).



Note Use at least one active BR for the MC to display all possible NBAR application names based on the protocol pack installed in BR.



Note PfRv2 is not supported on Cisco IOS 15.6(3)M and Cisco IOS 15.7(3)M or later releases. Cisco IOS XE 16.3.1 has PfRv2 CLIs, but the functionality is not supported.

Information About PfRv3

Performance Routing v3 Overview

Performance Routing Version 3 (PfRv3) is a one-touch provisioning and multi-site coordination solution that simplifies network provisioning. It enables intelligence of Cisco devices to improve application performance and availability. PfRv3 is an application-based policy driven framework that provides a multi-site aware bandwidth and path control optimization for WAN and cloud-based applications.

PfRv3 monitors network performance and selects best path for each application based on criteria such as reachability, delay, jitter, and loss. It evenly distributes traffic and maintains equivalent link utilization levels and load balances traffic.

It is tightly integrated with existing AVC components such as Performance Monitoring, Quality of Service (QoS), and NBAR2. PfRv3 is useful for enterprise and managed service providers looking for ways to increase their WAN reliability and availability while saving cost.

Benefits of PfRv3


Performance Routing Version 3 provides the following benefits:

- Centralized provisioning — Policies are defined on the hub-master controller and then distributed to all branches. Hence, per-site provisioning is not required in PfRv3.
- Simple provisioning — PfRv3 has simplified policies with pre-existing templates that a user can choose from.
- Enterprise domain — All sites belong to an enterprise domain and are connected with peering.
- Application and DSCP-based policies — Policies are provisioned based on applications. PfRv3 provides application visibility such as bandwidth, performance, and correlation to Quality of Service (QoS) queues by using Unified Monitoring.
- Automatic discovery — PfRv3 sites are discovered using peering. Each site peers with the hub site. The WAN interfaces are automatically discovered on the branch sites.
- Scalable passive monitoring — PfRv3 uses Unified Monitor to monitor traffic going into WAN links and traffic coming from the WAN links. It monitors performance metrics based on per DSCP instead of per flow or per prefix basis.
- Smart probing — PfRv3 uses a probing mechanism that generates traffic only when there is no traffic. It generates real-time transport protocol traffic, which allows measuring jitter and packet loss using performance monitors.
- Scaling — Smart probing and enhanced passive metrics help to attain scale up to 2000 branches.
- VRF awareness — Different policies can be configured for different VRFs.

PfRv3 Design Overview

An enterprise organization has a hub and branch site. The hub site consists of a master controller and border router.

Figure 1: PfRv3 Design Topology



- In a network, all the policies are created on the hub-master controller. Policies dictate the desired treatment for a set of specified differentiated service code points (DSCPs) or application IDs (such as telepresence,

WebEx, and so on) in the network. The policies are percolated to all the master controllers on the network via Service Advertisement Framework (SAF). The policies can be modified by the hub-master controller and the modified policies are sent over the SAF framework so that all the nodes in the network are in sync with the hub-master controller. The hub-master controller collects information about flows handled by border routers. This information is exported to the master controller periodically using the performance monitoring instances (PMI) exporter. A domain can be configured on the central location (Hub) and branches. PfRv3 allows only one domain configuration. Virtual Routing and Forwarding (VRF) and roles are defined on a domain.

- PfRv3 is enabled on the WAN interface of the hub-border routers. The border routers give the flow information to the branch-master controller.
- Every branch has a local-master controller. The master controller can be either co-located with a branch router or a separate router. You must configure both local master and branch border on the same domain. Border devices establishes connection with local-master controller only if both are in the same domain. In a scenario where master and border configurations are on different domain, peering rejects all messages from different peers. Border devices are automatically shut down for five minutes. The connection is established only when the domain conflict is resolved.

Based on the flow information provided by the hub-border router, the branch-master (local-master) controller applies appropriate controls on the branch router per flow. It ascertains if a flow is operating within the policy limits or out-of-policy. The master-controller to branch-border communication is done via a TCP connection. This connection is used for tasks such as sending configuration and control information from master controller to branch router and flow information from branch router to master controller.

- The branch router is the enforcer, which classifies and measures metrics and sends them to the local-master controller. It is also responsible for path enforcement.

PfRv3 Configuration Components

PfRv3 comprises of the following configuration components:

- Device setup and role — Identifies devices in the network where PfRv3 should be configured and in what role.
- Policy configurations — Identifies the traffic in the network and determines what policies to apply.

Device Setup and Role

There are four different roles a device can play in PfRv3 configuration:

- Hub-master controller — The master controller at the hub site, which can be either a data center or a head quarter. All policies are configured on hub-master controller. It acts as master controller for the site and makes optimization decision.
- Hub-border router — The border controller at the hub site. PfRv3 is enabled on the WAN interfaces of the hub-border routers. You can configure more than one WAN interface on the same device. You can have multiple hub border devices. On the hub-border router, PfRv3 must be configured with the address of the local hub-master controller, path names, and path-ids of the external interfaces. You can use the global routing table (default VRF) or define specific VRFs for the hub-border routers.

- Branch-master controller — The branch-master controller is the master controller at the branch site. There is no policy configuration on this device. It receives policy from the hub-master controller. This device acts as master controller for the branch site and makes optimization decision.
- Branch- border router — The border device at the branch-site. There is no configuration other than enabling of PfRv3 border-master controller on the device. The WAN interface that terminates on the device is detected automatically.

Domain Policies

Domain policies are defined only on the hub-master controller and then sent over peering infrastructure to all the branch-master controllers. Policies can be defined per application or per differentiated service code point (DSCP). You cannot mix and match DSCP and application-based policies in the same class group. Traffic that does not match any of the classification and match statements falls into a default group, which is load balanced (no performance measurement is done).



Note You can either select an existing template for a policy or customize your policies for a domain type.

The following table lists the existing templates for domain type policy:

Pre-defined Template	Threshold Definition
Voice	Priority 1 one-way-delay threshold 150 threshold 150 (msec) Priority 2 packet-loss-rate threshold 1 (%) Priority 2 byte-loss-rate threshold 1 (%) Priority 3 jitter 30 (msec)
Real-time-video	Priority 1 packet-loss-rate threshold 1 (%) Priority 1 byte-loss-rate threshold 1 (%) Priority 2 one-way-delay threshold 150 (msec) Priority 3 jitter 20 (msec)
Low-latency-data	Priority 1 one-way-delay threshold 100 (msec)) Priority 2 byte-loss-rate threshold 5 (%) Priority 2 packet-loss-rate threshold 5 (%)
Bulk-data	Priority 1 one-way-delay threshold 300 (msec) Priority 2 byte-loss-rate threshold 5 (%) Priority 2 packet-loss-rate threshold 5 (%)
Best-effort	Priority 1 one-way-delay threshold 500 (msec) Priority 2 byte-loss-rate threshold 10 (%) Priority 2 packet-loss-rate threshold 10 (%)

Pre-defined Template	Threshold Definition
Scavenger	Priority 1 one-way-delay threshold 500 (msec) Priority 2 byte-loss-rate threshold 50 (%) Priority 2 packet-loss-rate threshold 50 (%)
Custom	Defines customized user-defined policy values

PfRv3 and Link Group Configuration

PfRv3 allows you to configure the following option for link grouping:

- Allows up to five primary path preferences and four fallback path preferences
- Allows a fallback blackhole configuration
- Allows a fallback routing configuration

During Policy Decision Point (PDP), the exits are first sorted on the available bandwidth and then a second sort algorithm places all primary path preferences in the front of the list followed by fallback preferences. If you have a configuration of primary Internet Service Provider (ISP) 1 and ISP2 and ISP3 as fallback, during policy decision, ISP1 is selected as the primary channel and if ISP2 is equally good it is selected as the fallback. ISP3 is considered only if ISP2 is bad in bandwidth availability.

Routing configuration means that when the traffic is uncontrolled, the routing table takes the responsibility of pushing the flow out of the box.

