



Understanding Performance Routing

This module describes how Performance Routing (PfR) operates to help you understand how to implement the technology in your network. After configuration, the PfR technology runs through a series of phases that start with profiling traffic classes, measuring the traffic classes, apply policies to the traffic classes, controlling the traffic classes to meet the policy conditions, and finally verifying the result of the traffic class optimization.



Note

The PfR configuration modules refer to the PfR syntax introduced in Cisco IOS Release 15.1(2)T. If you are running Cisco IOS Release 15.1(1)T, or an earlier release, or any 12.2SR or 12.2SX image, you need to consult the [Optimized Edge Routing Configuration Guide](#) to help you locate all the Optimized Edge Routing documentation.

- [Finding Feature Information, page 1](#)
- [Prerequisites for Understanding Performance Routing, page 2](#)
- [Information About Understanding Performance Routing, page 2](#)
- [Where To Go Next, page 28](#)
- [Additional References, page 28](#)
- [Feature Information for Understanding Performance Routing, page 29](#)

Finding Feature Information

Your software release may not support all the features documented in this module. For the latest feature information and caveats, see the release notes for your platform and software release. To find information about the features documented in this module, and to see a list of the releases in which each feature is supported, see the Feature Information Table at the end of this document.

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.

Prerequisites for Understanding Performance Routing

- Only border router functionality is included in the Cisco IOS XE Release 3.1S and 3.2S images; no master controller configuration is available. The master controller that communicates with the Cisco ASR 1000 series router being used as a border router in the Cisco IOS XE Release 3.1S and 3.2S images must be a router running Cisco IOS Release 15.0(1)M, or a later 15.0M release. In Cisco IOS XE Release 3.3S, and later releases, master controller configuration is supported.
- Before understanding the PfR phases, you need to understand an overview of how PfR works and how to set up basic PfR network components. See the "[Configuring Basic Performance Routing](#)" module for more details.
- Cisco Express Forwarding (CEF) must be enabled on all participating devices. No other switching path is supported, even if otherwise supported by policy-based routing (PBR).

Information About Understanding Performance Routing

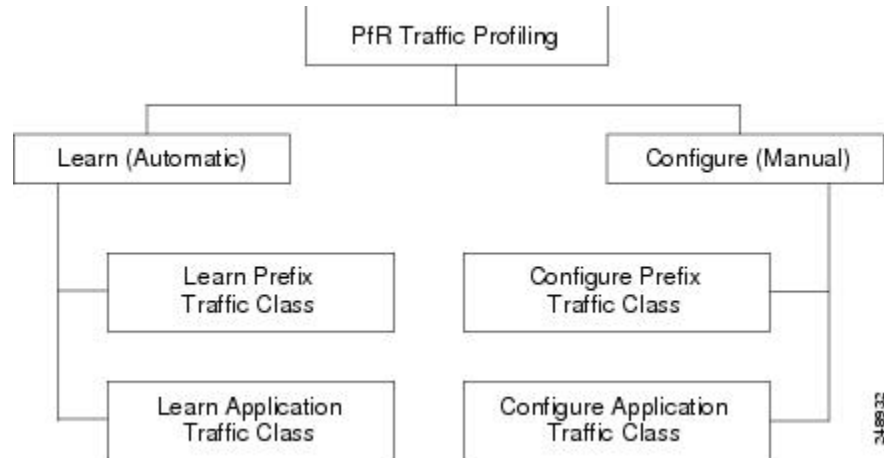
Profile Phase Concepts

Traffic Class Profiling Overview

Before optimizing traffic, PfR has to determine the traffic classes from the traffic flowing through the border routers. To optimize traffic routing, subsets of the total traffic must be identified, and these traffic subsets are named traffic classes. The list of traffic classes entries is named a Monitored Traffic Class (MTC) list. The entries in the MTC list can be profiled either by automatically learning the traffic flowing through the device or by manually configuring the traffic classes. Learned and configured traffic classes can both exist in the MTC list at the same time. The PfR profile phase includes both the learn mechanism and the configure

mechanism. The overall structure of the PfR traffic class profile process and its component parts can be seen in the figure below.

Figure 1: PfR Traffic Class Profiling Process



The ultimate objective of this phase is to select a subset of traffic flowing through the network. This subset of traffic--the traffic classes in the MTC list--represents the classes of traffic that need to be routed based on the best performance path available.

Automatic Traffic Class Learning

PfR can automatically learn the traffic classes while monitoring the traffic flow through border routers. Although the goal is to optimize a subset of the traffic, you may not know all the exact parameters of this traffic and PfR provides a method to automatically learn the traffic and create traffic classes by populating the MTC list. Several features have been added to PfR since the original release to add functionality to the automatic traffic class learning process.

Within the automatic traffic class learning process there are now three components. One component describes the automatic learning of prefix-based traffic classes, the second component describes automatic learning of application-based traffic classes, and the third component describes the use of learn lists to categorize both prefix-based and application-based traffic classes. These three components are described in the following sections:

Prefix Traffic Class Learning Using PfR

The PfR master controller can be configured, using NetFlow Top Talker functionality, to automatically learn prefixes based on the highest outbound throughput or the highest delay time. Throughput learning measures prefixes that generate the highest outbound traffic volume. Throughput prefixes are sorted from highest to lowest. Delay learning measures prefixes with the highest round-trip response time (RTT) to optimize these highest delay prefixes to try to reduce the RTT for these prefixes. Delay prefixes are sorted from the highest to the lowest delay time.

PfR can automatically learn two types of prefixes:

- outside prefix--An outside prefix is defined as a public IP prefix assigned outside the company. Outside prefixes are received from other networks.
- inside prefix--An inside prefix is defined as a public IP prefix assigned to a company. An inside prefix is a prefix configured within the company network.

In the BGP Inbound Optimization feature the ability to learn inside prefixes was introduced. Using BGP, PfR can select inside prefixes to support best entrance selection for traffic that originates from prefixes outside an autonomous system destined for prefixes inside the autonomous system. In prior releases, only outside prefixes were supported. For more details about inside prefix PfR support, see the BGP Inbound Optimization Using Performance Routing module.

Automatic prefix learning is configured in PfR Top Talker and Top Delay learning configuration mode. The **learn** (PfR) command is used to enter this mode from PfR master controller configuration mode. When automatic prefix learning is enabled, prefixes and their delay or throughput characteristics are measured on the border routers. Performance measurements for the prefix-based traffic classes are reported to the master controller where the learned prefixes are stored in the MTC list.

Prefixes are learned on the border routers through monitoring the traffic flow using the embedded NetFlow capability. All incoming and outgoing traffic flows are monitored. The top 100 flows are learned by default, but the master controller can be configured to learn up to 2500 flows for each learn cycle.

The master controller can be configured to aggregate learned prefixes based on type; BGP or non-BGP (static). Prefixes can be aggregated based on the prefix length. Traffic flows are aggregated using a /24 prefix length by default. Prefix aggregation can be configured to include any subset or superset of the network, from single host route (/32) to a major network address range. For each aggregated prefix, up to five host addresses are selected to use as active probe targets. Prefix aggregation is configured with the **aggregation-type**(PfR) command in PfR Top Talker and Delay learning configuration mode.

Application Traffic Class Learning Using PfR

PfR can learn Layer 3 prefixes, and Layer 4 options such as protocol or port numbers can be added as filters to the prefix-based traffic class. The protocol and port numbers can be used to identify specific application traffic classes; protocol and port number parameters are monitored only within the context of a prefix and are not sent to the master controller database (MTC list). The prefix that carries the specific traffic is then monitored by the master controller. PfR application traffic class learning also supports Differentiated Services Code Point (DSCP) values in addition to protocol and port numbers, and these Layer 4 options are entered in the MTC list.

DSCP Value, Port, and Protocol Learning by PfR

PfR has the ability to filter and aggregate application traffic by DSCP value, port number or protocol. Traffic classes can be defined by a combination of keys comprising of protocol, port numbers, and DSCP values. The ability to filter out traffic that is not required, and the ability to aggregate the traffic in which you are interested, was introduced. Information such as protocol, port number, and DSCP value is now sent to the master controller database in addition to the prefix information. The new functionality allows PfR to both actively and passively monitor application traffic. Using new CLI and access lists, PfR can be configured to automatically learn application traffic classes.

Learn List Configuration Mode

PfR supports a learn list configuration mode to simplify the learning of traffic classes. Learn lists are a way to categorize learned traffic classes. In each learn list, different criteria including prefixes, application definitions, filters, and aggregation parameters for learning traffic classes can be configured. A traffic class is automatically learned by PfR based on each learn list criteria, and each learn list is configured with a sequence number. The sequence number determines the order in which learn list criteria are applied. Learn lists allow different PfR policies to be applied to each learn list; in previous releases, the traffic classes could not be divided, and an PfR policy was applied to all the learned traffic classes.

Learn list configuration mode uses **traffic-class** commands to simplify the learning of traffic classes. Four types of traffic classes--to be automatically learned--can be profiled:

- Traffic classes based on destination prefixes
- Traffic classes representing custom application definitions using access lists
- Traffic classes based on a static application mapping name with optional prefix lists to define destination prefixes
- Traffic classes based on a NBAR application mapping name with optional prefix lists to define destination prefixes

Only one type of **traffic-class** command can be specified per learn list, and the **throughput** (PfR) and **delay** (PfR) commands are also mutually exclusive within a learn list.

Static Application Mapping Using PfR

The static application mapping feature introduced the ability to define an application using a keyword to simplify the configuration of application-based traffic classes. PfR uses well-known applications with fixed ports, and more than one application may be configured at the same time. For more details about static application mapping, see the Static Application Mapping Using Performance Routing feature.

PfR Application Mapping Using NBAR

PfR supports the ability to profile an application-based traffic class using NBAR. Network-Based Application Recognition (NBAR) is a classification engine that recognizes and classifies a wide variety of protocols and applications, including web-based and other difficult-to-classify applications and protocols that use dynamic TCP/UDP port assignments. PfR uses NBAR to recognize and classify a protocol or application, and the resulting traffic classes are added to the PfR application database to be passively and actively monitored. For more details about PfR application mapping using NBAR, see the Performance Routing with NBAR/CCE Application Recognition feature.

Manual Traffic Class Configuration

PfR can be manually configured to create traffic classes for monitoring and subsequent optimizing. Automatic learning generally uses a default prefix length of /24 but manual configuration allows exact prefixes to be defined. Within the manual traffic class configuration process there are two components-- manually configuring prefix-based traffic classes and manually configuring application-based traffic classes, both of which are described in the following sections:

Prefix Traffic Class Configuration Using PfR

A prefix or range of prefixes can be selected for PfR monitoring by configuring an IP prefix list. The IP prefix list is then imported into the MTC list by configuring a match clause in a PfR map. A PfR map is similar to an IP route map. IP prefix lists are configured with the **ip prefix-list** command and PfR maps are configured with the **pfr-map** command in global configuration mode.

The prefix list syntax operates in a slightly different way with PfR than in regular routing. The **ge** keyword is not used and the **le** keyword is used by PfR to specify only an inclusive prefix. A prefix list can also be used to specify an exact prefix.

A master controller can monitor and control an exact prefix of any length including the default route. If an exact prefix is specified, PfR monitors only the exact prefix.

A master controller can monitor and control an inclusive prefix using the **le** keyword and the *le-value* argument set to 32. PfR monitors the configured prefix and any more specific prefixes (for example, configuring the 10.0.0.0/8 le 32 prefix would include the 10.1.0.0/16 and the 10.1.1.0/24 prefixes) over the same exit and records the information in the routing information base (RIB).



Note

Use the inclusive prefix option with caution in a typical PfR deployment because of the potential increase in the amount of prefixes being monitored and recorded.

An IP prefix list with a deny statement can be used to configure the master controller to exclude a prefix or prefix length for learned traffic classes. Deny prefix list sequences should be applied in the lowest PfR map sequences for best performance. The master controller can also be configured to tell border routers to filter out uninteresting traffic using an access list.



Note

IP prefix lists with deny statements can be applied only to learned traffic classes.

Two types of prefix can be manually configured for PfR monitoring using an IP prefix list:

- outside prefix--An outside prefix is defined as a public IP prefix assigned outside the company. Outside prefixes are received from other networks.
- inside prefix--An inside prefix is defined as a public IP prefix assigned to a company. An inside prefix is a prefix configured within the company network.

In the BGP Inbound Optimization feature the ability to manually configure inside prefixes was introduced. Using BGP, PfR can be configured to select inside prefixes to support best entrance selection for traffic that originates from prefixes outside an autonomous system destined for prefixes inside the autonomous system. In prior releases, only outside prefixes were supported.

For more details about inside prefix PfR support, see the BGP Inbound Optimization Using Performance Routing module.

Application Traffic Class Configuration Using PfR

PfR supports the manual configuration of Layer 3 prefixes during the PfR profile phase. Application-aware routing for policy-based routing (PBR) is also supported. Application-aware routing allows the selection of

traffic for specific applications based on values in the IP packet header, other than the Layer 3 destination address through a named extended IP access control list (ACL). Only named extended ACLs are supported. The extended ACL is configured with a permit statement and then referenced in a PfR map. The protocol and port numbers can be used to identify specific application traffic classes, but protocol and port number parameters are monitored only within the context of a prefix, and are not sent to the MTC list. Only the prefix that carries the specific application traffic is profiled by the master controller. With application-aware routing support, active monitoring of application traffic was supported. Passive monitoring of application traffic is also supported. Application traffic classes can be defined using DSCP values as well as protocol and port numbers. DSCP values, port numbers, and protocols in addition to prefixes, are all now stored in the MTC list.

Learn list configuration mode uses **match traffic-class** commands under PfR map configuration mode to simplify the configuration of traffic classes. Four types of traffic classes--to be manually configured--can be profiled:

- Traffic classes based on destination prefixes
- Traffic classes representing custom application definitions using access lists
- Traffic classes based on a static application mapping name and a prefix list to define destination prefixes
- Traffic classes based on NBAR application mapping name and a prefix list to define destination prefixes

Only one type of **match traffic-class** command can be specified per PfR map.

For a series of well-known applications, static ports have been defined and each application can be defined by entering a keyword. For more details about static application mapping, see the Static Application Mapping Using Performance Routing feature.

PfR supports the ability to profile an application-based traffic class using NBAR. NBAR is a classification engine that recognizes and classifies a wide variety of protocols and applications, including web-based and other difficult-to-classify applications and protocols that use dynamic TCP/UDP port assignments. PfR uses NBAR to recognize and classify a protocol or application, and the resulting traffic classes are added to the PfR application database to be passively and actively monitored. For more details about PfR application mapping using NBAR, see the Performance Routing with NBAR/CCE Application Recognition feature.

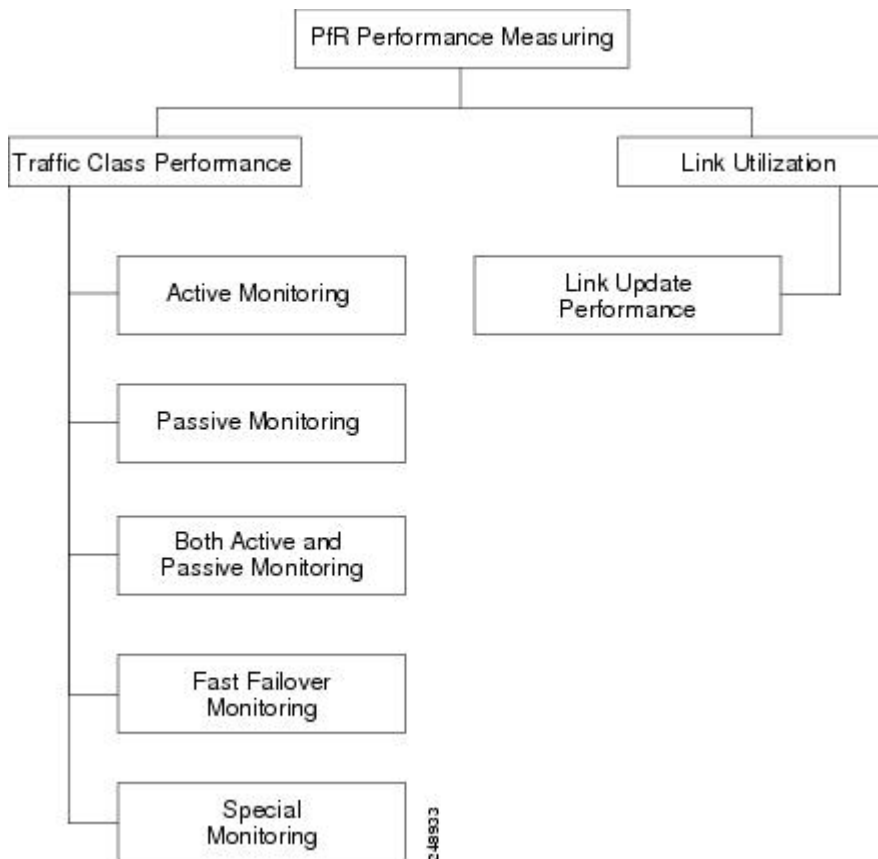
Measure Phase Concepts

Traffic Class Performance Measurement Overview

The PfR measure phase is the second step in the PfR performance loop and it follows the PfR profile phase where the traffic class entries fill the Monitored Traffic Class (MTC) list. The MTC list is now full of traffic class entries and PfR must measure the performance metrics of these traffic class entries. Monitoring is defined here as the act of measurement performed periodically over a set interval of time where the measurements are compared against a threshold. PfR measures the performance of traffic classes using active and passive monitoring techniques but it also measures, by default, the utilization of links. The master controller can be configured to monitor learned and configured traffic classes. The border routers collect passive monitoring and active monitoring statistics and then transmit this information to the master controller. The PfR measure phase is complete when each traffic class entry in the MTC list has associated performance metric measurements.

The overall structure of the PfR measure phase and its component parts can be seen in the figure below.

Figure 2: PfR Performance Measuring Process



PfR measures the performance of both traffic classes and links, but before monitoring a traffic class or link PfR checks the state of the traffic class or link. PfR uses a policy decision point (PDP) that operates according to a traffic class state transition diagram.

After determining the state of the traffic class or link, PfR may initiate one of the following performance measuring processes.

Traffic Class Performance Measurement Techniques

PfR uses three methods of traffic class performance measurement:

- Passive monitoring--measuring the performance metrics of traffic class entries while the traffic is flowing through the device using NetFlow functionality.
- Active monitoring--creating a stream of synthetic traffic replicating a traffic class as closely as possible and measuring the performance metrics of the synthetic traffic. The results of the performance metrics of the synthetic traffic are applied to the traffic class in the MTC list. Active monitoring uses integrated IP Service Level Agreements (IP SLAs) functionality.
- Both active and passive monitoring--combining both active and passive monitoring in order to generate a more complete picture of traffic flows within the network.

Fast failover monitoring mode is another variation of the combined active and passive monitoring modes. In fast failover monitoring mode, all exits are continuously probed using active monitoring and passive monitoring. When fast failover monitoring mode is enabled, the probe frequency can be set to a lower frequency than for other monitoring modes, to allow a faster failover capability.

No explicit NetFlow or IP SLAs configuration is required and support for NetFlow and IP SLAs is enabled automatically. You can use both active and passive monitoring methods for a traffic class.

After the master controller is defined and PfR functionality is enabled, the master controller uses both passive and active monitoring by default. All traffic classes are passively monitored using integrated NetFlow functionality. Out-of-policy traffic classes are actively monitored using IP SLA functionality. You can configure the master controller to use only passive monitoring, active monitoring, both passive and active monitoring, or fast failover monitoring. The main differences between the different modes can be seen in the table below.

Table 1: Mode Comparison Table

Comparison Parameter	Active Mode	Passive Mode	Combined Mode	Fast Failover Mode
Active/IP SLA	Yes	No	Yes	Yes
Passive/NetFlow	No	Yes	Yes	Yes
Monitoring of Alternate Paths	On Demand	On Demand	On Demand	Continuous
Best Failover Time	10 seconds	~ 1 minute	~ 1.1 minute	3 seconds
Support for Round Trip Delay	Yes	Yes	Yes	Yes
Support for Loss	Only with Jitter probe	Only for TCP traffic	Only for TCP traffic	Only for TCP traffic and Jitter probe
Support for Reachability	Yes	Only for TCP traffic	Only for TCP traffic	Yes
Support for Jitter	Yes	No	No	Yes
Support for MOS	Yes	No	No	Yes

Passive Monitoring

Cisco IOS PfR uses NetFlow, an integrated technology in Cisco IOS software, to collect and aggregate passive monitoring statistics on a per traffic class basis. Passive monitoring is enabled along with active monitoring by default when an PfR managed network is created. Passive monitoring can also be enabled explicitly using the **mode monitor passive** command. Netflow is a flow-based monitoring and accounting system, and NetFlow support is enabled by default on the border routers when passive monitoring is enabled.

Passive monitoring uses only existing traffic; additional traffic is not generated. Border routers collect and report passive monitoring statistics to the master controller approximately once per minute. If traffic does not

go over an external interface of a border router, no data is reported to the master controller. Threshold comparison is done at the master controller. Passive monitoring supports traffic classes defined by prefix, port, protocol, and DSCP value.

PfR uses passive monitoring to measure the following metrics for all the traffic classes:

- Delay--PfR measures the average delay of TCP flows for a given prefix. Delay is the measurement of the round-trip response time (RTT) between the transmission of a TCP synchronization message and receipt of the TCP acknowledgement.
- Packet loss--PfR measures packet loss by tracking TCP sequence numbers for each TCP flow. PfR estimates packet loss by tracking the highest TCP sequence number. If a subsequent packet is received with a lower sequence number, PfR increments the packet loss counter. Packet loss is measured in packets per million.
- Reachability--PfR measures reachability by tracking TCP synchronization messages that have been sent repeatedly without receiving a TCP acknowledgement.
- Throughput--PfR measures throughput by measuring the total number of bytes and packets for each traffic class for a given interval of time.

**Note**

Although all traffic classes are monitored, delay, loss, and reachability information is captured only for TCP traffic flows. Throughput statistics are captured for all non-TCP traffic flows.

DSCP values, port numbers, and protocols in addition to prefixes, are all sent from border routers to the master controller. Passive monitoring statistics are gathered and stored in a prefix history buffer that can hold a minimum of 60 minutes of information depending on whether the traffic flow is continuous. PfR uses this information to determine if the prefix is in-policy based on the default or user-defined policies. No alternative path analysis is performed as the traffic for a traffic class is flowing through one transit device in the network. If the traffic class goes OOP and only passive monitoring mode is enabled, the traffic class is moved to another point and the measurement repeated until a good or best exit is found. If the traffic class goes OOP and both passive and active monitoring modes are enabled, active probing is executed on all the exits and a best or good exit is selected.

Active Monitoring

If PfR passive monitoring techniques create too much overhead on a network device, or the performance metrics of a traffic class cannot be measured using the PfR passive monitoring mode, then PfR active monitoring techniques are performed. Active monitoring involves creating a stream of synthetic traffic that replicates a traffic class as closely as possible. The performance metrics of the synthetic traffic are measured and the results are applied to the traffic class entry in the MTC list. Active monitoring supports traffic classes defined by prefix, port, protocol, and DSCP value.

PfR uses active monitoring to measure the following metrics for all the traffic classes:

- Delay--PfR measures the average delay of TCP, UDP, and ICMP flows for a given prefix. Delay is the measurement of the round-trip response time (RTT) between the transmission of a TCP synchronization message and receipt of the TCP acknowledgement.
- Reachability--PfR measures reachability by tracking TCP synchronization messages that have been sent repeatedly without receiving a TCP acknowledgement.

- Jitter--Jitter means interpacket delay variance. PfR measures jitter by sending multiple packets to a target address and a specified target port number, and measuring the delay interval between packets arriving at the destination.
- MOS--Mean Opinion Score (MOS) is a standards-based method of measuring voice quality. Standards bodies like the ITU have derived two important recommendations: P.800 (MOS) and P.861 (Perceptual Speech Quality Measurement [PSQM]). P.800 is concerned with defining a method to derive a Mean Opinion Score of voice quality. MOS scores range between 1 representing the worst voice quality, and 5 representing the best voice quality. A MOS of 4 is considered "toll-quality" voice.

The creation of synthetic traffic in Cisco network devices is activated through the use of Cisco IOS IP SLA probes. PfR is integrated with IP SLAs functionality such that PfR will use IP SLA probes to actively monitor a traffic class. When active monitoring is enabled, the master controller commands the border routers to send active probes to set of target IP addresses. The border sends probe packets to no more than five target host addresses per traffic class, and transmits the probe results to the master controller for analysis.

Active probe monitoring periods are defined as short-term which consists of the last 5 probe results, and long-term which consists of the last 60 probe results.

IP SLA Active Probe Types Used by PfR

IP SLAs are an embedded feature set in Cisco IOS software and they allow you to analyze IP service levels for IP applications and services, to increase productivity, to lower operational costs, and to reduce occurrences of network congestion or outages. IP SLAs use active traffic monitoring--the generation of traffic in a continuous, reliable, and predictable manner--for measuring network performance. The accuracy of measured data is enhanced by enabling the IP SLAs Responder, available in Cisco routers, on the destination device. For more details about IP SLAs, see the [IP SLAs Configuration Guide](#).

The following types of active probes can be configured:

- ICMP Echo--A ping is sent to the target address. PfR uses ICMP Echo probes, by default, when an active probe is automatically generated. Configuring an ICMP echo probe does not require knowledgeable cooperation from the target device. However, repeated probing could trigger an Intrusion Detection System (IDS) alarm in the target network. If an IDS is configured in a target network that is not under your control, we recommend that you notify the administrator of this target network.
- Jitter--A jitter probe is sent to the target address. A target port number must be specified. A remote responder must be enabled on the target device, regardless of the configured port number. Loss policy is supported for active monitoring if the jitter probe is used.
- TCP Connection--A TCP connection probe is sent to the target address. A target port number must be specified. A remote responder must be enabled if TCP messages are configured to use a port number other than TCP port number 23, which is well-known.
- UDP Echo--A UDP echo probe is sent to the target address. A target port number must be specified. A remote responder must be enabled on the target device, regardless of which port number is configured.

PfR marks the probe packets with the DSCP value by default if the monitored traffic classes have the DSCP field set to a nonzero value.

Creation of Active Probe for a Traffic Class

To create an active probe for a traffic class, a probe type has to be discovered, and a probe target assigned to the traffic class. To discover a probe type, PfR uses one of the following methods:

- **Learned probe**--Active probes are automatically generated when a traffic class is learned using the NetFlow TopTalker Learn mechanism. Five targets are learned for each traffic class and, by default, the active probe is set as an ICMP echo probe.
- **Configured probe**--Active probes can also be configured on the master controller by specifying the probe type, target address and port if needed. Configured traffic classes can be configured to use any of the IP SLA active probes.

To assign a probe target for a traffic class, PfR uses one of the following methods:

- **Longest match**--By default, PfR assigns a probe target to the traffic class with the longest matching prefix in the MTC list. This is referred to as a default probe assignment.
- **Forced assignment**--An IP SLA probe can be configured using a PfR map and the results of the probe are assigned to specific traffic classes associated with the PfR map. This specific assignment of active probe results is called a forced target probe assignment.

The active probe is sourced from the border router and transmitted through an external interface (the external interface may, or may not, be the preferred route for an optimized prefix). When creating an active probe through an external interface for a specified target, the target should be reachable through the external interface. To test the reachability of the specified target, PfR performs a route lookup in the BGP and static routing tables for the specified target and external interface. Protocol Independent Route Optimization (PIRO) introduced the ability of PfR to search for a parent route--an exact matching route, or a less specific route--in any IP Routing Information Base (RIB). The BGP routing table is searched first, followed by the static routing table, and finally the RIB.

In active monitoring mode, the probes are activated from all the border routers to find the best performance path for the specific traffic class. The active probes for that traffic class are not activated again unless the traffic class goes OOP.

By default, the frequency of an active probe used by PfR is set to 60 seconds. The frequency of an active probe can be increased for each policy by configuring a lower time-interval between two probes. Increased probe frequency can reduce the response time and, for voice traffic, provide a better approximation of the MOS-low count percentage.

PfR Active Probe Source Address

PfR supports the ability to configure an active probe source address. By default, active probes use the source IP address of the PfR external interface that transmits the probe. The active probe source address feature is configured on the border router. When this command is configured, the primary IP address of the specified interface is used as the active probe source. The active probe source interface IP address must be unique to ensure that the probe reply is routed back to the specified source interface. If the interface is not configured with an IP address, the active probe will not be generated. If the IP address is changed after the interface has been configured as an active probe source, active probing is stopped, and then restarted with the new IP address. If the IP address is removed after the interface has been configured as an active probe source, active probing is stopped and not restarted until a valid primary IP address is configured.

PfR Voice Traffic Optimization Using Active Probes

PfR supports outbound optimization of voice traffic using active probes on the basis of voice metrics such as delay, reachability, jitter, and Mean Opinion Score (MOS).

For more details about optimizing voice traffic, see the "[PfR Voice Traffic Optimization Using Active Probes](#)" module.

Combined Monitoring

Cisco IOS PfR can also be configured to combine both active and passive monitoring in order to generate a more complete picture of traffic flows within the network. There are some scenarios in which you may want to combine both PfR monitoring modes.

One example scenario is when you want to learn traffic classes and then monitor them passively, but you also want to determine the alternate path performance metrics in order to control the traffic classes. The alternate path performance metrics, in the absence of the actual traffic flowing through the alternate path in the network, can be measured using the active probes. PfR automates this process by learning traffic classes at five targets and probing through all the alternate paths using active probes.

Fast Failover Monitoring

Fast monitoring sets the active probes to continuously monitor all the exits (probe-all), and passive monitoring is enabled too. Fast failover monitoring can be used with all types of active probes: ICMP echo, Jitter, TCP connection, and UDP echo. When the **mode monitor fast** command is enabled, the probe frequency can be set to a lower frequency than for other monitoring modes, to allow a faster failover ability. Under fast monitoring with a lower probe frequency, route changes can be performed within 3 seconds of an out-of-policy situation. When an exit becomes OOP under fast monitoring, the select best exit is operational and the routes from the OOP exit are moved to the best in-policy exit. Fast monitoring is a very aggressive mode that incurs a lot of overhead with the continuous probing. We recommend that you use fast monitoring only for performance sensitive traffic. For example, a voice call is very sensitive to any performance problems or congested links, but the ability to detect and reroute the call within a few seconds can demonstrate the value of using fast monitoring mode.

**Note**

In fast monitoring mode, probe targets are learned as well as learned prefixes. To avoid triggering large numbers of probes in the network, use fast monitoring mode only for real time applications and critical applications with performance sensitive traffic.

Link Utilization Measurement Techniques

Link Utilization Threshold


After an external interface is configured for a border router, PfR automatically monitors the utilization of the external link (an external link is an interface on a border router that typically links to a WAN). Every 20 seconds, by default, the border router reports the link utilization to the master controller. Both egress (transmitted) and ingress (received) traffic utilization values are reported to the master controller. If the exit or entrance link utilization is above the default threshold of 75-percent, the exit or entrance link is in an OOP state and PfR starts the monitoring process to find an alternative link for the traffic class. The link utilization threshold can be manually configured either as an absolute value in kilobytes per second (kbps) or as a percentage.

Link Utilization Range

PfR can also be configured to calculate the range of utilization over all the links. Both egress (transmitted) and ingress (received) traffic utilization values are reported to the master controller. In the diagram below

there are two border routers with exit links to the Internet through two ISPs. The master controller determines which link on one of the border routers--either BR1 or BR2 in the diagram below--is used by a traffic class.

Figure 3: PfR network diagram



PfR range functionality attempts to keep the exit or entrance links within a utilization range, relative to each other to ensure that the traffic load is distributed. The range is specified as a percentage and is configured on the master controller to apply to all the exit or entrance links on border routers managed by the master controller. For example, if the range is specified as 25-percent, and the utilization of the exit link at BR1 (in the diagram above) is 70-percent, then if the utilization of the exit link at BR2 (in the diagram above) falls to 40-percent, the percentage range between the two exit links will be more than 25-percent and PfR will attempt to move some traffic classes to use the exit link at BR1 to even the traffic load. If BR1 (in the diagram above) is being configured as an entrance link, the link utilization range calculations work in the same way as for an exit link, except that the utilization values are for received traffic, not transmitted traffic.



Note

If you are configuring link grouping, configure the **no max-range-utilization** command because using a link utilization range is not compatible with using a preferred or fallback set of exit links configured for link grouping. With CSCtr33991, this requirement is removed and PfR can perform load balancing within a PfR link group.

Apply Policy Phase Concepts

Apply Policy Phase Overview

The PfR apply policy phase is the third step in the PfR performance loop following after the profile phase that identifies the traffic classes, and the measure phase where each traffic class entry in the MTC list is monitored to determine performance metric measurements. The apply policy phase compares the measured performance metrics against well-known or configured thresholds to determine if the traffic is meeting specified levels of service, or if some action is required. If the performance metric does not conform to the threshold, a decision is made by PfR to move the traffic class or exit into another state.

An PfR policy is a rule that defines an objective and contains the following attributes:


- A scope.
- An action.

- A triggering event or condition.

For example, a policy can be configured to maintain a delay of less than or equal to 100 milliseconds for packets sent to a specific traffic class entry. The scope is the network traffic sent to the specific traffic class entry, the action is a routing table change, and the triggering event is a measured delay of greater than 100 milliseconds for this traffic. The action may not be executed until PfR is configured to control the traffic in the PfR control phase. By default, PfR runs in an observe mode during the profile, measure, and apply policy phases.

In the PfR apply policy phase you can configure and apply policies. Different types of PfR policies can be configured--see the figure below--and specific PfR parameters and options can be included within a policy. In this document, a parameter is a configurable element that can be fine-tuned, and an option is a configurable element that is either enabled or disabled. After an PfR policy is configured, the policy can be applied to learned traffic classes or configured traffic classes. PfR policies can be applied globally--to all the traffic classes--or to just a specific set of traffic classes.

Figure 4: PfR Apply Policy Phase Structure



In the figure above you can see that there are three types of PfR policies plus some operational options and parameters that can be configured. Use the following links to review more information about each policy type, parameter, or option:

After an PfR policy is configured, you can see from the figure above that a policy can be applied to learned traffic classes or configured traffic classes on a global basis for all traffic classes or for a specific set of traffic classes.


When configuring multiple policy parameters for traffic classes, it is possible to have multiple overlapping policies. To resolve the potential conflict of which policy to run, PfR uses its resolve function: a flexible mechanism that allows you to set the priority for most of the policy types.

PfR Policy Decision Point

When running an PfR policy that compares the traffic class performance metrics with default or configured thresholds, a traffic class may change state. PfR uses a policy decision point (PDP) that operates according to the traffic class state transition diagram shown in the figure below. The state transition diagram below contains the following states:

- **Default**--A traffic class is placed in the default state when it is not under PfR control. Traffic classes are placed in the default state when they are initially added to the central policy database, the MTC. A traffic class will transition into and out of the default state depending on performance measurements, timers, and policy configuration.
- **Choose Exit**--This is a temporary state in which the PDP compares the current state of the traffic class against its policy settings and chooses the optimal exit for the traffic class. PfR will try to keep a traffic class flowing through its current exit but, as in the default state, performance measurements, timers, and policy configurations can cause the master controller to place a traffic class in this state for the duration of the exit link selection process. The traffic class remains in the choose exit state until it is moved to the new exit.
- **Holddown**--A traffic class is placed in the holddown state when the master controller requests a border router to forward the traffic class to be monitored using probes. Measurements are collected for the selected traffic class until the holddown timer expires unless the exit used by this traffic class is declared unreachable. If the exit is unreachable, the traffic class transitions back to the choose exit state.

Figure 5: PfR Traffic Class State Transition Diagram



- **In-Policy**--After performance measurements are compared against default or user-defined policy settings and an exit selection is made, the traffic class enters an in-policy state. When a traffic class is in the in-policy state, the traffic class is forwarded through an exit that satisfies the default or user-defined settings. The master controller continues to monitor the traffic class, but no action is taken until the periodic timer expires, or an out-of-policy message is received from a measurement collector, when the traffic class transitions back to the choose exit state.



Note

When observe mode is running, a prefix goes into an in-policy state only if the exit selected for that prefix is the current exit.

- **Out-of-Policy (OOP)**--A traffic class is placed in this state when there are no exits through which to forward the traffic class that conform to default or user-defined policies. While the traffic class is in this state, the backoff timer controls exiting from this state. Each time the traffic class enters this state, the amount of time the traffic class spends in this state increases. The timer is reset for a traffic class when

the traffic class enters an in-policy state. If all exit links are out-of-policy, the master controller may select the best available exit.

Traffic Class Performance Policies

PfR traffic class performance policies are a set of rules that govern performance characteristics for traffic classes that can be network addresses (prefixes) or application criteria such as protocol, port number, or DSCP value. Network addresses can refer to individual endpoints within a network (e.g. 10.1.1.1/32) or to entire subnets (e.g. 10.0.0.0/8). The major performance characteristics that can be managed within an PfR policy are:

With the exception of reachability, none of these performance characteristics can be managed within the constructs of conventional routing protocol metrics. Cisco PfR extends the concept of reachability (beyond ensuring that a particular route exists in the routing table) by automatically verifying that the destination can be reached through the indicated path. Using Cisco PfR provides the network administrator with a new and powerful toolset for managing the flow of traffic.

Reachability

Reachability is specified as the relative percentage or the absolute maximum number of unreachable hosts, based on flows per million (fpm), that PfR will permit from a traffic class entry. If the absolute number or relative percentage of unreachable hosts is greater than the user-defined or the default value, PfR determines that the traffic class entry is out-of-policy and searches for an alternate exit link.

To configure parameters for reachability, use the **unreachable** (PfR) command. This command has two keywords, **relative** and **threshold**. The **relative** keyword is used to configure the relative percentage of unreachable hosts. The relative unreachable host percentage is based on a comparison of short-term and long-term measurements. The short-term measurement reflects the percentage of hosts that are unreachable within a 5-minute period. The long-term measurement reflects the percentage of unreachable hosts within a 60-minute period. The following formula is used to calculate this value:

$$\text{Relative percentage of unreachable hosts} = ((\text{short-term percentage} - \text{long-term percentage}) / \text{long-term percentage}) * 100$$

The master controller measures the difference between these two values as a percentage. If the percentage exceeds the user-defined or default value, the traffic class entry is determined to be out-of-policy. For example, if 10 hosts are unreachable during the long-term measurement and 12 hosts are unreachable during short-term measurement, the relative percentage of unreachable hosts is 20 percent.

The **threshold** keyword is used to configure the absolute maximum number of unreachable hosts. The maximum value is based on the actual number of hosts that are unreachable based on fpm.

Delay

Delay (also referred as latency) is defined as the delay between when the packet was sent from the source device and when it arrived at a destination device. Delay can be measured as one-way delay or round-trip delay. The largest contributor to latency is caused by network transmission delay.

PfR supports defining delay performance characteristics with respect to voice traffic. Round-trip delay affects the dynamics of conversation and is used in Mean Opinion Score (MOS) calculations. One-way delay is used for diagnosing network problems. A caller may notice a delay of 200 milliseconds and try to speak just as the other person is replying because of packet delay. The telephone industry standard specified in ITU-T G.114 recommends the maximum desired one-way delay be no more than 150 milliseconds. Beyond a one-way

delay of 150 milliseconds, voice quality is affected. With a round-trip delay of 300 milliseconds or more, users may experience annoying talk-over effects.

Packet Loss

Packet loss can occur due an interface failing, a packet being routed to the wrong destination, or congestion in the network.

Packet loss for voice traffic leads to the degradation of service in which a caller hears the voice sound with breaks. Although average packet loss is low, voice quality may be affected by a short series of lost packets.

Jitter

PfR supports defining jitter performance characteristics. Jitter means interpacket delay variance. When multiple packets are sent consecutively from source to destination, for example, 10 ms apart, and if the network is behaving ideally, the destination should be receiving them 10 ms apart. But if there are delays in the network (like queuing, arriving through alternate routes, and so on) the arrival delay between packets might be greater than or less than 10 ms. Using this example, a positive jitter value indicates that the packets arrived more than 10 ms apart. If the packets arrive 12 ms apart, then positive jitter is 2 ms; if the packets arrive 8 ms apart, then negative jitter is 2 ms. For delay-sensitive networks like VoIP, both positive and negative jitter values are undesirable; a jitter value of 0 is ideal.

Mean Opinion Score (MOS)

PfR supports defining MOS performance characteristics. With all the factors affecting voice quality, many people ask how voice quality can be measured. Standards bodies like the ITU have derived two important recommendations: P.800 (MOS) and P.861 (Perceptual Speech Quality Measurement [PSQM]). P.800 is concerned with defining a method to derive a Mean Opinion Score of voice quality. MOS scores range between 1 representing the worst voice quality, and 5 representing the best voice quality. A MOS of 4 is considered "toll-quality" voice.

Jitter and MOS performance characteristic can be configured in an PfR policy as well as delay and packet loss to determine the quality of a phone call over an IP network.

PfR Link Policies

PfR link policies are a set of rules that are applied against PfR-managed external link (an external link is an interface on a border router on the network edge). Link policies define the desired performance characteristics of the links. Instead of defining the performance of an individual traffic class entry that uses the link (as in traffic class performance policies), link policies are concerned with the performance of the link as a whole. Link policies can be applied to exit (egress) links and entrance (ingress) links. The following performance characteristics are managed by link policies:

- Traffic Load (Utilization)
- Range
- Cost

Traffic Load

A traffic load (also referred to as utilization) policy consists of an upper threshold on the amount of traffic that a specific link can carry. Cisco IOS PfR supports per traffic class load distribution. Every 20 seconds, by default, the border router reports the link utilization to the master controller, after an external interface is

configured for a border router. Exit link and entrance link traffic load thresholds can be configured as an PfR policy. If the exit or entrance link utilization is above the configured threshold, or the default threshold of 75-percent, the exit or entrance link is in an OOP state and PfR starts the monitoring process to find an alternative link for the traffic class. The link utilization threshold can be manually configured either as an absolute value in kilobits per second (kbps) or as a percentage. A load utilization policy for an individual interface is configured on the master controller under the border router configuration.

**Tip**

When configuring load distribution, we recommend that you set the interface load calculation on external interfaces to 30-second intervals with the **load-interval** (PfR) interface configuration command. The default calculation interval is 300 seconds. The load calculation is configured under interface configuration mode on the border router. This configuration is not required, but it is recommended to allow Cisco IOS PfR to respond as quickly as possible to load distribution issues.

Range

A range policy is defined to maintain all links within a certain utilization range, relative to each other in order to ensure that the traffic load is distributed. For example, if a network has multiple exit links, and there is no financial reason to choose one link over another, the optimal choice is to provide an even load distribution across all links. The load-sharing provided by traditional routing protocols is not always evenly distributed, because the load-sharing is flow-based rather than performance- or policy-based. Cisco PfR range functionality allows you to configure PfR to maintain the traffic utilization on a set of links within a certain percentage range of each other. If the difference between the links becomes too great, PfR will attempt to bring the link back to an in-policy state by distributing traffic classes among the available links. The master controller sets the maximum range utilization to 20 percent for all PfR-managed links by default, but the utilization range can be configured using a maximum percentage value. Exit link and entrance link utilization ranges can be configured as a PfR policy.

**Note**

If you are configuring link grouping, configure the **no max-range-utilization** command because using a link utilization range is not compatible with using a preferred or fallback set of exit links configured for link grouping. With CSCtr33991, this requirement is removed and PfR can perform load balancing within a PfR link group.

Cost

Cost-based optimization allow you to configure policies based on the monetary cost (ISP service level agreements [SLAs]) of each exit link in your network. To implement PfR cost-based optimization the PfR master controller is configured to send traffic over exit links that provide the most cost-effective bandwidth utilization, while still maintaining the desired performance characteristics.

Cost Based Optimization can be applied to links that are billed using a fixed or tiered billing method and load balancing based on cost can also be achieved. For more configuration details, see the “Configuring Performance Routing Cost Policies” module.

PfR Link Grouping

In the Performance Routing - Link Groups feature, the ability to define a group of exit links as a preferred set of links, or a fallback set of links for PfR to use when optimizing traffic classes specified in an PfR policy,

was introduced. PfR currently selects the best link for a traffic class based on the preferences specified in a policy and the traffic class performance--using parameters such as reachability, delay, loss, jitter or MOS--on a path out of the specified link. Bandwidth utilization, cost, and the range of links can also be considered in selecting the best link. Link grouping introduces a method of specifying preferred links for one or more traffic classes in an PfR policy so that the traffic classes are routed through the best link from a list of preferred links, referred to as the primary link group. A fallback link group can also be specified in case there are no links in the primary group that satisfy the specified policy and performance requirements. If no primary group links are available, the traffic classes are routed through the best link from the fallback group. To identify the best exit, PfR probes links from both the primary and fallback groups.

**Note**

If you are configuring link grouping, configure the **no max-range-utilization** command because using a link utilization range is not compatible with using a preferred or fallback set of exit links configured for link grouping. With CSCtr33991, this requirement is removed and PfR can perform load balancing within a PfR link group.

For more details about PfR link grouping, see the "Performance Routing Link Groups" module.

PfR Network Security Policies

The ability to configure network security policies either to prevent unauthorized use of the network or to mitigate attacks inside and outside the network is provided by PfR. You can configure PfR to use black hole or sinkhole routing techniques to limit the impact of attacks against your network. Black hole routing refers to the process of forwarding packets to a null interface, meaning that the packets are dropped into a "black hole." Sinkhole routing directs packets to a next hop where the packets can be stored, analyzed, or dropped. Another term for sinkhole routing is honey-pot routing.

PfR Policy Operational Options and Parameters

In addition to the specific types of PfR policies, there are some PfR policy operational parameters or options that can be configured. The operational parameters are timers and the operational options consist of different operational modes. For more details, see the following sections:

PfR Timers Parameters

Three types of timers can be configured as PfR policy operational parameters:

Backoff Timer

The backoff timer is used to adjust the transition period that the master controller holds an out-of-policy traffic class entry. The master controller waits for the transition period before making an attempt to find an in-policy exit. A minimum, a maximum, and an optional step timer value can be configured.

Holddown Timer

The holddown timer is used to configure the traffic class entry route dampening timer to set the minimum period of time that a new exit must be used before an alternate exit can be selected. To prevent the traffic class entry from flapping because of rapid state changes, the master controller does not move the traffic class entry to a different exit even if it goes out-of-policy during the holddown timer period. PfR does not implement policy changes while a traffic class entry is in the holddown state. A traffic class entry will remain in a

holddown state for the default or configured time period. When the holddown timer expires, PfR will select the best exit based on performance and policy configuration. However, an immediate route change will be triggered if the current exit for a traffic class entry becomes unreachable.

Periodic Timer

The periodic timer is used to find a better path for a traffic class entry, even if the traffic class entry is in-policy on the current exit. When the periodic timer expires, the master controller evaluates current exit links for the traffic class entry and, if a better exit exists based on the current measurements and priorities, the traffic class entry is moved to a new in-policy exit link.

When adjusting PfR timers note that a newly configured timer setting will immediately replace the existing setting if the value of the new setting is less than the time remaining. If the value is greater than the time remaining, the new setting will be applied when the existing timer expires or is reset.



Note

Overly aggressive timer settings can keep an exit link or traffic class entry in an out-of-policy state.

PfR Mode Options

Three types of mode options can be configured as PfR policy operational options:

Mode Monitor

The mode monitor option enables the configuration of PfR monitoring settings. Monitoring is defined here as the act of measurement performed periodically over a set interval of time where the measurements are compared against a threshold. PfR measures the performance of traffic classes using active and passive monitoring techniques but it also measures, by default, the utilization of exit links.

Mode Route

The mode route option specifies one of three PfR route control policy settings. Mode route control enables PfR to control routes automatically, mode route metric specifies PfR route protocol-related settings, and mode route observe offers route control advice, but does not take any action. Observe mode monitoring is enabled by default when PfR is enabled. In observe mode, the master controller monitors traffic classes and exit links based on default and user-defined policies and then reports the status of the network and the decisions that should be made but does not implement any changes. Observe mode is used to verify the effect of PfR features before PfR is actively deployed on your network.

If you have different routing protocols operating on your PfR border routers (for example, BGP on one border router and EIGRP on another) you must configure the **protocol** and **pbr** keywords with the mode route command to allow destination-only traffic classes to be controlled using dynamic PBR. Entering the **no mode route protocol pbr** command will initially set the destination-only traffic classes to be uncontrolled and PfR then reverts to the default behavior using a single protocol to control the traffic class in the following order; BGP, EIGRP, static, and PBR.

Mode Select-Exit

The mode select-exit option enables the exit selection settings. The definition of an in-policy traffic class entry is that the measured performance metrics do not exceed a default or configured threshold while the traffic class entry is on the current path. In this situation, PfR does not search for an alternate exit link because the current network path keeps the traffic class entry in-policy. This type of configuration would be activated

by using the **mode select-exit good** command which is the default if the **mode** (PfR) command is not specified. There are other deployment scenarios, where PfR selects the best performance path. This type of configuration can be activated by using the **mode select-exit best** command. In this type of situation, PfR measures alternate path performance metrics while the traffic class entry is in-policy on the current path. PfR moves the current path if a better performance path is found. After the first selection of the best path, however, PfR does not initiate another search unless the periodic timer is configured. When the periodic timer expires, the master controller evaluates current exit links for the traffic class entry and, if a better exit exists based on the current measurements and priorities, the traffic class entry is moved to a new in-policy exit link. Use the periodic timer with the **mode select-exit best** command if you have a deployment scenario where you need PfR to select the best performance path at any given time.

There is one further use of the mode select-exit option. If PfR does not find an in-policy exit for a traffic class entry when the **mode select-exit good** command is operational, PfR transitions the traffic class entry to an uncontrolled state. If PfR does not find an in-policy exit for a traffic class entry when the **mode select-exit best** command is operational, PfR selects the best of the OOP exit links for the traffic class entry.

PfR Policy Application

PfR policies can be applied to learned or configured traffic classes. PfR policies can be applied on a global basis when the policy is configured directly under PfR master controller configuration mode. All traffic classes inherit global policies. If, however, you want to apply a policy to a subset of the traffic classes, then a specific policy can be configured. A specific PfR policy applies only to the specific traffic classes that match a prefix list or access list. Specific policies inherit global policies unless the same policy is overwritten by the specific policy. PfR policies can apply to prefixes alone, or PfR policies can apply to traffic classes that define an application traffic class and may include prefixes, protocols, port numbers, and DSCP values. To apply specific policies to learned or configured traffic classes, PfR map configuration is used.

PfR Map Configuration for PfR Policies

A PfR map may appear to be similar to a route map but there are significant differences. A PfR map is designed to select learned or configured traffic classes using a match clause and then to apply PfR policy configurations using a set clause. The PfR map can be optionally configured with a sequence number like a route map, but only the PfR map with the lowest sequence number is evaluated. The operation of a PfR map differs from a route map at this point. There are two important distinctions:

- Only a single match clause may be configured for each sequence. An error message will be displayed on the console if you attempt to configure multiple match clauses for a single PfR map sequence.
- A PfR map is not configured with permit or deny statements. However, a permit or deny sequence can be configured for an IP traffic flow by configuring a permit or deny statement in an IP prefix list and then applying the prefix list to the PfR map.



Note Match precedence priority is not supported in PfR maps.

The PfR map applies the configuration of the set clause after a successful match occurs. A PfR set clause can be used to set policy parameters such as the backoff timer, packet delay, holddown timer, packet loss, mode settings, periodic timer, resolve settings, unreachable hosts, and traceroute reporting.

Policies applied by an PfR map take effect immediately. The PfR map configuration can be viewed in the output of the **show running-config** command. PfR policy configuration can be viewed in the output of the

show pfr master policy command. These policies are applied only to traffic classes that match or pass through the PfR map.

Policy Rules Configuration to Apply an PfR Policy

The **policy-rules** (PfR) command allows you to select a PfR map using a sequence number and apply the configuration under PfR master controller configuration mode, providing an improved method to switch between predefined PfR maps. Only one PfR map is used at a time for policy configuration, but many PfR maps can be defined.

Priority Resolution for Multiple PfR Policies

When configuring multiple policy criteria for a single traffic class entry, or a set of traffic classes, it is possible to have multiple overlapping policies. To resolve the potential conflict of which policy to run, PfR uses its resolve function: a flexible mechanism that allows you to set the priority for a PfR policy. Each policy is assigned a unique value, and the policy with the lowest value is selected as the highest priority policy. By default, PfR assigns the highest priority to delay policies, followed by utilization policies. Assigning a priority value to any policy will override the default settings. To configure the policy conflict resolution, use the **resolve** (PfR) command in PfR master controller configuration mode, or the **set resolve** (PfR) command in PfR map configuration mode.

Variance Setting for PfR Policy Conflict Resolution

When configuring PfR resolve settings, you can also set an allowable variance for the defined policy. Variance configures the average delay, as a percentage, that all traffic classes for one exit, or the specific policy traffic classes for an exit, can vary from the defined policy value and still be considered equivalent. For example, if the delay on the best exit link (best exit in terms of delay) for a traffic class entry is 80 milliseconds (ms) and a 10 percent variance is configured, then any other exit links with a delay between 80 and 88 ms for the same traffic class entry are considered equivalent to the best exit link.

To illustrate how variance is used by PfR consider three exit links with the following performance values for delay and jitter for a traffic class entry:

- Exit A--Delay is 80 ms, jitter is 3ms
- Exit B--Delay is 85 ms, jitter is 1ms
- Exit C--Delay is 100 ms, jitter is 5ms

The following PfR policy conflict resolution is configured and applied to the traffic class entry:

```
delay priority 1 variance 10
jitter priority 2 variance 10
```

PfR determines the best exit by looking at the policy with the lowest priority value, which in this example is the delay policy. Exit A has the lowest delay value, but Exit B has a delay value of 85 which is within a 10-percent variance of the delay value at Exit A. Exit A and Exit B can therefore be considered equal in terms of delay values. Exit C is now eliminated because the delay values are too high. The next priority policy is jitter, and Exit B has the lowest jitter value. PfR will select Exit B as the only best exit for the traffic class entry because Exit A has a jitter value that is not within 10-percent variance of the Exit B jitter value.



Note

Variance cannot be configured for cost or range policies.

Enforce Phase Concepts

PfR Enforce Phase Overview

After profiling the traffic classes during the PfR learn phase, measuring the performance metrics of the traffic classes during the measure phase, and using network policies to map the measured performance metrics of traffic class entries in the Monitored Traffic Class (MTC) list against well-known or configured thresholds to determine if the traffic is meeting specified levels of service in the policy phase, the next step in the PfR performance loop is the PfR enforce phase.

PfR, by default, operates in an observation mode and the documentation for the PfR learn, measure, and apply policy phases assumes that PfR is in the observe mode. In observe mode, the master controller monitors traffic classes and exit links based on default and user-defined policies and then reports the status of the network including out-of-policy (OOP) events and the decisions that should be made, but does not implement any changes. The PfR enforce phase operates in control mode, not observe mode, and control mode must be explicitly configured using the **mode route control** command. In control mode, the master controller coordinates information from the border routers in the same way as observe mode, but commands are sent back to the border routers to alter routing in the PfR managed network to implement the policy decisions.

PfR initiates route changes when one of the following occurs:

- A traffic class goes OOP.
- An exit link goes OOP.
- The periodic timer expires and the select exit mode is configured as select best mode.

During the PfR enforce phase, the master controller continues to monitor in-policy traffic classes that conform to the desired performance characteristics, to ensure that they remain in-policy. Changes are only implemented for OOP traffic classes and exits in order to bring them in-policy. To achieve the desired level of performance in your network, you must be aware of the configuration options that can affect the policy decisions made by the master controller.

Another configuration issue to consider when deploying PfR is that if aggressive delay or loss policies are defined, and the exit links are also seriously over-subscribed, it is possible that PfR will find it impossible to bring a traffic class in-policy. In this case, the master controller will either choose the link that most closely conforms to the performance policy, even though the traffic class still remains OOP, or it will remove the prefix from PfR control. PfR is designed to allow you to make the best use of available bandwidth, but it does not solve the problem of over-subscribed bandwidth.

After PfR control mode is enabled, and configuration options are considered, the next step is to review the traffic class control techniques employed by PfR.


PfR Traffic Class Control Techniques

After the PfR master controller has determined that it needs to take some action involving an OOP traffic class or exit link, there are a number of techniques that can be used to alter the routing metrics, alter BGP attributes, or introduce policy-based routing using a route map to influence traffic to use a different link. If the traffic associated with the traffic class is defined only by a prefix then a traditional routing control mechanism such as introducing a BGP route or a static route can be deployed. This control is network wide after redistribution because a prefix introduced into the routing protocol with a better metric will attract traffic for that prefix towards a border router. If the traffic associated with the traffic class is defined by a prefix and

other matching criteria for the packet (application traffic, for example), then traditional routing cannot be employed to control the application traffic. In this situation, the control becomes device specific and not network specific. This device specific control is implemented by PfR using policy-based routing (PBR) functionality. If the traffic in this scenario has to be routed out to a different device, the remote border router should be a single hop away or a tunnel interface that makes the remote border router look like a single hop.

The figure below shows the various traffic class control techniques grouped by exit or entrance link selection.

Figure 6: Controlling Traffic Class Techniques



PfR Exit Link Selection Control Techniques

Before introducing the exit link selection control techniques you need to understand one principle about load balancing with Performance Routing as it applies to exit selection. PfR does not treat a more specific route as a parent route unless you configure the more specific route as a default route.

When searching for a parent route, the software tries to find the most specific route that includes the specified prefix and verifies that it points to the expected exit. If there are two or more static routes that are more specific, each route is inspected for the expected exit. If the expected exit is found, the probe is created.

In the configuration where:

```
ip route 10.4.0.0 255.255.0.0 172.17.40.2
ip route 0.0.0.0 0.0.0.0 serial 6/0
```

Probes for prefix 10.4.1.0/24 and target 10.4.1.1 will not be created over the exit using serial interface 6/0 because the most specific route inclusive of 10.4.1.1 is the exit to 172.17.40.2. If you are looking to load balance the traffic over both exits, the answer is to create a default route of the more specific route. For example:

```
ip route 10.4.0.0 255.255.0.0 172.17.40.2
ip route 10.4.0.0 255.255.0.0 serial 6/0
```

Or

```
ip route 0.0.0.0 0.0.0.0 serial 6/0
ip route 0.0.0.0 0.0.0.0 172.17.40.2
```

In the modified configuration, two probes are created, one for the exit to 172.17.40.2 and one for the exit using serial interface 6/0.

To enforce an exit link selection, PfR offers the following methods:

Static Route Injection

A PfR master controller can enforce the use of a particular border router as the preferred exit link for a traffic class by injecting temporary static routes. These static routes exist only in the memory of the router, and are intentionally not saved to the permanent configuration. There are a few different methods that the master controller can use to inject static routes on the border routers. Existing static routes can be overwritten with new static routes, which have a better routing metric. If a default route, or even a less specific route, exists on the border router, the master controller can add a specific static route for the monitored traffic classes, which will be preferred to the existing default route. Finally, the master controller can also use something known as split prefixes.

A split prefix refers to the addition of a more specific route, which will be preferred over a less specific route. For example, if the border router already has a route of 10.10.10.0/24, adding a static route of 10.10.10.128/25 will also cause the addresses 10.10.10.129-10.10.10.254 to be forwarded using the newly injected route. If PfR has been configured to monitor a subset of a larger network, it will add an appropriate route to the existing routing table. PfR can use split prefixes to redirect subsets of an existing prefix to a more optimal exit link, and can use split prefixes for both internal BGP (iBGP) and static routes.

PfR will never inject a route where one does not already exist in the routing protocol table. Before injecting a route of a particular type, PfR will verify that a route exists in the BGP or static table that includes the prefix and points to the exit link. This route may be a default route.

BGP Control Techniques

PfR uses two BGP techniques to enforce the best exit path; injecting a BGP route, or modifying the BGP local preference attribute.

If the traffic associated with the traffic class is defined only by a prefix, the master controller can instruct a border router to inject a BGP route into the BGP table to influence traffic to use a different link. All PfR injected routes remain local to an autonomous system, and these injected routes are never shared with external BGP peers. As a safeguard to ensure this behavior, when PfR injects a BGP route, it will set the no-export community on it. This is done automatically, and does not require any user configuration. However, because these routes now have a special marking, some extra configuration is required to allow the information to be shared with internal BGP peers. For each iBGP peer, the send community configuration must be specified. Although the border routers know about the best exit for the injected route, it may also be necessary to redistribute this information further into the network.

PfR also uses BGP local preference to control traffic classes. BGP local preference (Local_Pref) is a discretionary attribute applied to a BGP prefix to specify the degree of preference for that route during route selection. The Local_Pref is a value applied to a BGP prefix, and a higher Local_Pref value causes a route to be preferred over an equivalent route. The master controller instructs one of the border routers to apply the Local_Pref attribute to a prefix or set of prefixes associated with a traffic class. The border router then shares the Local_Pref value with all of its internal BGP peers. Local_Pref is a locally significant value within an autonomous system, but it is never shared with external BGP peers. Once the iBGP reconvergence is complete, the router with the highest Local_Pref for the prefix will become the exit link from the network.



Note

If a local preference value of 5000 or higher has been configured for default BGP routing, you should configure a higher BGP local preference value in PfR using the **mode** (PfR) command.

EIGRP Route Control

The PfR EIGRP mGRE DMVPN Hub-and-Spoke Support feature introduced PfR route control for EIGRP. When enabled, a parent route check is performed in the EIGRP database for controlling PfR prefixes/routes in addition to the existing BGP and static route databases. For more details, see the ["Using Performance Routing to Control EIGRP Routes with mGRE DMVPN Hub-and-Spoke Support"](#) module.

Policy-Based Routing Control

PfR can control application traffic using policy-based routing. Application traffic traveling through a particular PfR border router can be identified by matching traffic defined in a PfR map as part of a PfR policy. The **match ip address** (PfR) command was enhanced to support extended ACLs. The extended ACL is referenced in a PfR map, and a single match clause can be configured for each PfR map sequence. Set clauses are configured to apply independent PfR policies to the matched traffic, which is a subset of a monitored prefix. The PfR policy is applied to all border routers to enforce policy routing for the application. Matched traffic is policy routed through the PfR external interface that conforms to policy parameters.

DSCP values, as well as prefixes, port numbers, and protocols, can all be used to identify and control application traffic. DSCP values, protocols, and port numbers are sent by the border routers to the master controller for inclusion in the MTC list.

Protocol Independent Route Optimization (PIRO)

PIRO was introduced to extend the ability of PfR to identify and control traffic classes. Prior to PIRO, PfR optimizes paths for traffic classes that have a parent route--an exact matching route, or a less specific route--in BGP or static route databases. PIRO enables PfR to search the IP Routing Information Base (RIB) for a parent route allowing PfR to be deployed in any IP-routed environment including Interior Gateway Protocols (IGPs) such as OSPF and IS-IS.

For more details, see the ["PfR Protocol Independent Route Optimization"](#) module.

PfR Entrance Link Selection Control Techniques

The PfR BGP inbound optimization feature introduced the ability to influence inbound traffic. A network advertises reachability of its inside prefixes to the Internet using eBGP advertisements to its ISPs. If the same prefix is advertised to more than one ISP, then the network is multihoming. PfR BGP inbound optimization works best with multihomed networks, but it can also be used with a network that has multiple connections to the same ISP. To implement BGP inbound optimization, PfR manipulates eBGP advertisements to influence the best entrance selection for traffic bound for inside prefixes. The benefit of implementing the best entrance selection is limited to a network that has more than one ISP connection.

For more details about PfR entrance link selection control techniques, see the ["BGP Inbound Optimization Using Performance Routing"](#) module.

Verify Phase Concepts

Verify Phase Overview

The last phase of the PfR performance loop is to verify that the actions taken during the PfR control phase control actually change the flow of traffic and that the performance of the traffic class or link does move to an in-policy state. PfR uses NetFlow to automatically verify the route control. The master controller expects

a Netflow update for the traffic class from the new link interface and ignores Netflow updates from the previous path. If a Netflow update does not appear after two minutes, the master controller moves the traffic class into the default state. A traffic class is placed in the default state when it is not under PfR control.

In addition to the NetFlow verification used by PfR, there are two other methods you can use to verify that PfR has initiated changes in the network:

- **Syslog report**--The logging command can be configured to notify you of all the main PfR state changes, and a syslog report can be run to confirm that PfR changes have occurred. The master controller is expecting bidirectional traffic, and a syslog report delimited for the specified prefix associated with the traffic class can confirm this.
- **PfR show commands**--PfR show commands can be used to verify that network changes have occurred and that traffic classes are in-policy. Use the **show pfr master prefix** command to display the status of monitored prefixes. The output from this command includes information about the current exit interface, prefix delay, egress and ingress interface bandwidth, and path information sourced from a specified border router. Use the **show pfr border routes** command to display information about PfR controlled routes on a border router. This command can display information about BGP or static routes.

Where To Go Next

To access configuration tasks and configuration examples that implement the concepts contained in this module, see the "Configuring Advanced Performance Routing" module. For details about other Performance Routing modules and features, see the "Related Documents" section.

Additional References

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Command List, All Releases
Cisco IOS PfR commands: complete command syntax, command mode, command history, defaults, usage guidelines, and examples	Cisco IOS Performance Routing Command Reference
Basic PfR configuration for Cisco IOS XE releases	"Configuring Basic Performance Routing" module
Information about configuration for the border router only functionality for Cisco IOS XE Releases 3.1 and 3.2	"Performance Routing Border Router Only Functionality" module
Concepts required to understand the Performance Routing operational phases for Cisco IOS XE releases	"Understanding Performance Routing" module
Advanced PfR configuration for Cisco IOS XE releases	"Configuring Advanced Performance Routing" module

Related Topic	Document Title
IP SLAs overview	“Cisco IOS IP SLAs Overview” module
PfR home page with links to PfR-related content on our DocWiki collaborative environment	PfR:Home

MIBs

MIB	MIBs Link
<ul style="list-style-type: none"> • CISCO-PFR-MIB • CISCO-PFR-TRAPS-MIB 	<p>To locate and download MIBs for selected platforms, Cisco software releases, and feature sets, use Cisco MIB Locator found at the following URL:</p> <p>http://www.cisco.com/go/mibs</p>

Technical Assistance

Description	Link
The Cisco Support and Documentation website provides online resources to download documentation, software, and tools. Use these resources to install and configure the software and to troubleshoot and resolve technical issues with Cisco products and technologies. Access to most tools on the Cisco Support and Documentation website requires a Cisco.com user ID and password.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for Understanding Performance Routing

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 2: Feature Information for Understanding Performance Routing

Feature Name	Releases	Feature Configuration Information
OER Border Router Only Functionality	Cisco IOS XE Release 3.1.S	<p>Border Router Only Functionality was introduced in Cisco IOS XE Release 3.1S. The master controller that communicates with the Cisco ASR 1000 series routers being used as a border router must be a router running Cisco IOS Release 15.0(1)M or a later release.</p> <p>The following command was introduced or modified by this feature: show pfr border passive cache.</p>
PfR Master Controller Support for ASR 1000	Cisco IOS XE Release 3.3.S	<p>PfR Master Controller Support for ASR 1000 introduced master controller support. Previously, only border router support was available. This feature enabled most of the PfR features that are available on other platforms.</p>