

[PfR:Home](#)

Cisco Performance Routing (PfR) FAQs

Please let us know if you have any additions you would like added to this via the discussion page at: [Talk:Performance Routing FAQs](#).

Contents

- [1 About Performance Routing](#)
 - ◆ [1.1 Q. What is Cisco® Performance Routing?](#)
 - ◆ [1.2 Q. Are There Any Case Studies or Success Stories with PfR?](#)
 - ◆ [1.3 Q. Where can I find the Performance Routing \(PfR\) documentation?](#)
 - ◆ [1.4 Q. What is the difference between Performance Routing and Cisco IOS Optimized Edge Routing \(OER\)?](#)
 - ◆ [1.5 Q. What are the two components of Performance Routing? What do they do?](#)
 - ◆ [1.6 Q. In PfR terminology, what is an "Internal Interface" and an "External Interface"? Do the Border Routers \(BR's\) use the external interfaces for PfR Peering with each other?](#)
 - ◆ [1.7 Q. What problem does Cisco Performance Routing solve?](#)
 - ◆ [1.8 Q. What are the benefits to customers of Cisco Performance Routing?](#)
 - ◆ [1.9 Q. A lot of cases a customer has MPLS cloud with single provider and a backup cloud. How would you switchover from one to the other?](#)
 - ◆ [1.10 Q. Customer wants to use one provider cloud for one application and other cloud for other application, how would we do it using PfR?](#)
 - ◆ [1.11 Q. What is the challenge with integrating PfR with other routing protocols?](#)
 - ◆ [1.12 Q. What are the newest features in Performance Routing \(PfR\)?](#)
 - ◆ [1.13 Q. What are the list of feature names in Cisco Feature Navigator \(FN\) for Performance Routing \(PfR\)?](#)
 - ◆ [1.14 Q. What Cisco Hardware Platforms is PfR supported on? What is the recommended hardware? Are there any PfR deployment guides?](#)
 - ◆ [1.15 Q. What are the base hardware requirements for Cisco Performance Routing \(PfR\)?](#)
 - ◆ [1.16 Q. What are the Classic Software Image requirements for Cisco Performance Routing \(PfR\)?](#)
 - ◆ [1.17 Q. What are the New Universal Software Image requirements for Cisco Performance Routing \(PfR\)?](#)
 - ◆ [1.18 Q. What are the Software Image requirements for Cisco Performance Routing \(PfR\) on the ASR1k?](#)
 - ◆ [1.19 Q. Is there a feature matrix that lists the features with platform support and correlates them with the Cisco IOS Software releases?](#)
 - ◆ [1.20 Q. Is Cisco Performance Routing supported globally?](#)
 - ◆ [1.21 Q. Can you name some of the PfR customers?](#)
- [2 PfR Active Probing through Cisco \(IP SLAs\)](#)
 - ◆ [2.1 Q. Does Performance Routing require explicit IP SLA's configuration?](#)
 - ◆ [2.2 Q. Do I have to understand IP SLA's comprehensively?](#)
 - ◆ [2.3 Q. How does Performance Routing use IP SLA's active probing?](#)
 - ◆ [2.4 Q. What is a Mean Opinion Score or \(MOS\) metric?](#)
 - ◆ [2.5 Q. What is the overhead for PfR IP SLA's Fast Probing?](#)

Performance_Routing_FAQs

- ◆ [2.6 Q. If a customer wants to use Cisco IP SLAs to test and bring backup path up when the primary path goes down or Out Of Policy \(OOP\), can PfR help automate this?](#)
- ◆ [2.7 Q. In PfR Active Fast Mode, why can't we just change the probe frequency to match fast mode? Why have another mode?](#)
- ◆ [2.8 Q. I found something on the PfR wiki that seems to suggest multiple probes to same destination \(different ports\) is 'better' than a single probe. Would someone be able to elaborate?](#)
- [3 Command-Line Interface](#)
 - ◆ [3.1 Q. Is Performance Routing \(PfR\) configuration complex?](#)
- [4 Inbound Optimization](#)
 - ◆ [4.1 Q. Can inbound traffic be optimized?](#)
- [5 Learn](#)
 - ◆ [5.1 Q. When using the learn delay function, Performance Routing learns only prefixes that generate a three-way handshake during the learning period. Does the learn delay function learn only TCP traffic?](#)
 - ◆ [5.2 Q. Does learn delay learn the top N prefixes with the highest or worst delay?](#)
 - ◆ [5.3 Q. In PfR, is "learning" done using passive or active monitoring?](#)
 - ◆ [5.4 Q. When using learn throughput, Performance Routing identifies, through NetFlow, the top N prefixes that have the most throughput during the learn period. Does a learn throughput configuration allow Performance Routing to learn prefixes for which there is any sort of IP traffic?](#)
 - ◆ [5.5 Q. Why does the learn traffic-class count not properly reflected via "show oer master"?](#)
- [6 Miscellaneous](#)
 - ◆ [6.1 Q. Can a Performance Routing Master Controller and a Border Router be configured on the same physical router?](#)
 - ◆ [6.2 Q. When do I need a dedicated Master Controller?](#)
 - ◆ [6.3 Q. Does PfR have authentication \(key-chains\) between PfR Master Controllers \(MC's\) and Border Routers \(BR's\)?](#)
 - ◆ [6.4 Q. In PfR terminology, what is an "Internal Interface" and an "External Interface"? Do the Border Routers \(BR's\) use the external interfaces for PfR Peering with each other?](#)
 - ◆ [6.5 Q. How do I use Performance Routing for performance monitoring when I only have one external interface?](#)
 - ◆ [6.6 Q. How do I deploy Performance Routing in networks not running internal BGP \(iBGP\) within the enterprise?](#)
 - ◆ [6.7 Q. What Performance Routing limitations exist on Cisco Catalyst® 6500 Switches and Cisco 7600 Series Routers?](#)
 - ◆ [6.8 Q. What does the "Disable hardware switching to enable netflow aggregation export in version 9 format" message mean? Do I have to disable hardware switching?](#)
 - ◆ [6.9 Q. Is Performance Routing VRF-aware?](#)
 - ◆ [6.10 Q. Does Performance Routing enforce symmetric routing?](#)
 - ◆ [6.11 Q. Will Performance Routing cause route oscillation?](#)
 - ◆ [6.12 Q. How are Performance Routing system logging \(syslog\) messages formatted?](#)
- [7 Passive Monitoring \(NetFlow\)](#)
 - ◆ [7.1 Q. Does Performance Routing require explicit NetFlow configuration?](#)
 - ◆ [7.2 Q. How does Performance Routing use passive monitoring?](#)
- [8 Policy](#)
 - ◆ [8.1 Q. Is the PfR Monitor mode specific to a policy? Would it override what's in the PfR global policy?](#)
 - ◆ [8.2 Q. Can you please elaborate on : INPOLICY and OOPOLICY?](#)
 - ◆ [8.3 Q. If the configured policy cannot be met because available paths are not meeting requirement, then what happens next?](#)
 - ◆ [8.4 Q. When does PfR monitoring actually start? Lets say am using external interface-1, and that interface goes OOPOLICY \(OOP\). When does PfR start monitoring all the other](#)

Performance_Routing_FAQs

- available defined external links?
 - ◆ 8.5 Q. When using link grouping and a link goes out of policy (OOP), can Performance Routing revert back to the primary link?
 - ◆ 8.6 Q. When using select-exit best with mode active or mode both, does Performance Routing also initially probe all exits to find the best exit and then only probe all exits again if the current exit goes OOP?
 - ◆ 8.7 Q. When using select-exit best with mode fast, does Performance Routing always probe all exits and, at each periodic interval, can Performance Routing move the prefix to the best exit if the identified best exit is different from the current exit?
- 9 Route Control
 - ◆ 9.1 Q. What is a parent route?
 - ◆ 9.2 Q. Can you elaborate more on the Parent Route and why it's so important to PfR?
 - ◆ 9.3 Q. Which parent route IP routing protocols are supported?
 - ◆ 9.4 Q. Can Performance Routing control a traffic class that is more granular than a network address? If yes, how?
 - ◆ 9.5 Q. Is prefix splitting supported with Cisco Performance Routing? If so, how is it done?
 - ◆ 9.6 Q. How does the Master Controller inject a route?
 - ◆ 9.7 Q. Does the Master Controller have visibility of the Border Router BGP routing tables that it controls?
 - ◆ 9.8 Q. What happens if the Master Controller fails?
 - ◆ 9.9 Q. Does Performance Routing support the notion of redundant Master Controllers?
 - ◆ 9.10 Q. How does PfR verify that the application is reachable?
- 10 Scale
 - ◆ 10.1 Q. How is CPU utilization affected when Performance Routing is enabled?
 - ◆ 10.2 Q. How many Border Routers are supported, and how many exits in total?
- 11 Traffic Class Optimization
 - ◆ 11.1 Q. How does PfR Profile Traffic Classes?
 - ◆ 11.2 Q. How does PfR "Select the best traffic class" based on variance or priority and what is the elimination process for selecting this path?
 - ◆ 11.3 Q. Do I have to specify the traffic classes that are to be optimized?
 - ◆ 11.4 Q. If a traffic class is configured using an ACL with "any any" instead of specifying a destination subnet, is it supported? For example, "access-list 101 permit tcp any any eq telnet."
 - ◆ 11.5 Q. In a PfR System default configuration with a Aggregation-type prefix-length 24 defined, how is aggregation done for less specific routes in the RIB?
- 12 PfR Route Maps
 - ◆ 12.1 Q. How is an application defined in PfR? How is aggregation done and defined in PfR?
 - ◆ 12.2 Q. Where and how are PfR Dynamic Route-Maps applied? Over which interface, internal or external, will the Dynamic Route-Maps be applied on the BR?
 - ◆ 12.3 Q. In a PfR "PFR-MAP" , what does the associated number specify?

About Performance Routing

Q. What is Cisco® Performance Routing?

A. Cisco Performance Routing complements classic IP routing technologies by adding intelligence to select the best paths to meet application performance requirements. The first phase of Performance Routing technology intelligently optimizes application performance over enterprise WANs and to and from the Internet. This technology will evolve to help enable application performance optimization throughout the enterprise network through an end-to-end, performance-aware network.

Cisco Performance Routing selects an egress or ingress WAN path based on parameters that affect application performance, including reachability, delay, cost, jitter, and Mean Opinion Score (MOS). This

Performance_Routing_FAQs

technology can reduce network costs by facilitating more efficient load balancing and by increasing application performance without WAN upgrades. Traditional IP routing protocols such as Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), Routing Information Protocol (RIP), and Border Gateway Protocol (BGP) generally focus on providing reachability by creating a loop-free topology based on shortest or least-cost path. Cisco Performance Routing adds intelligence to routing decisions based on measured application and network performance.

Q. Are There Any Case Studies or Success Stories with Pfr?

A. Yes. Just recently in June of 2010 Cisco Inc. and SAS Inc. concluded a joint Case Study on their use of PfR. You can find this on the Cisco /go/pfr page under the 'Case Studies' heading. Here is a direct link to the document which is titled: Global Business Analytics Leader Reduces WAN Network Costs. Today **SAS Inc.** uses Performance Routing (PfR) to help ensure top performance for business-critical communications, while controlling costs.

Q. Where can I find the Performance Routing (PfR) documentation?

A. Today, we are designing part of our new PfR documentation strategy around the Cisco Docwiki server. We already have some basic and solution-type documents at PfR:Solutions. The new PfR command syntax can be found at [1], The feature configuration modules are documented for each release train. Cisco IOS Release 15.2M&T can be found at [2], Cisco IOS Release 15.1M&T can be found at [3], Cisco IOS Release 15.2S can be found at [4] and the Cisco IOS XE Release 3S PfR documents can be found at [5]. The old OER syntax documentation for Cisco IOS Release 15.1 and can be found at [6]. The documentation links can also be found at the PfR Marketing page.

Q. What is the difference between Performance Routing and Cisco IOS Optimized Edge Routing (OER)?

A. Cisco Performance Routing takes advantage of the vast intelligence embedded in Cisco IOS Software to determine the optimal path based upon network and application policies. Cisco Performance Routing is an evolution of the Cisco IOS OER technology with a much broader scope. OER was originally designed to provide route control on a per destination prefix basis, but Performance Routing has expanded capabilities that facilitate intelligent route control on a per application basis. The expanded capabilities provide additional flexibility and more granular application optimization than OER.

Optimized Edge Routing was first introduced in Cisco IOS Releases **12.3(8)T**, and **12.2(33)SRB**. Performance Routing is an extension of OER. The leading keyword **PfR** syntax was introduced in Cisco IOS Release **15.1(2)T**.

Q. What are the two components of Performance Routing? What do they do?

A. Performance Routing has two logical components: the Master Controller (MC) and the Border Router (BR).

The BR component resides within the data plane of the edge router. The BR uses NetFlow to passively gather throughput and TCP performance information. The BR also sources all IP service-level agreement (SLA) probes used for explicit application performance monitoring. As the policy enforcement point, the BR receives commands from the MC to actively control routes through route injection or dynamic policy-based routing (PBR) injection.

The MC acts as the central processor and database for the Performance Routing system. The MC component does not reside in the forwarding plane and, when deployed in a standalone fashion, has no view of routing information contained within the BR. The role of the MC is to gather information from the BR or BRs to

Q. What is Cisco® Performance Routing?

Performance_Routing_FAQs

determine whether or not traffic classes are in or out of policy, and to instruct the BRs how to make sure that traffic classes remain in policy using route injection or dynamic PBR injection. The MC also provides an API for reporting capabilities.

Depending on your Performance Routing deployment scenario and scaling requirements, the MC may be deployed on a dedicated router or may be deployed along with the BR on the same physical router.

Q. In PfR terminology, what is an "Internal Interface" and an "External Interface"? Do the Border Routers (BR's) use the external interfaces for PfR Peering with each other?

A. A basic PfR definition for these is best described as follows. PfR expects traffic to be flowing through the defined Internal to the External interfaces or visa-versa. BR's will only Peer with each other over the Internal Interfaces.[Media:Example.ogg](#)

Q. What problem does Cisco Performance Routing solve?

A. As enterprise organizations grow their businesses, there is a demand for real-time application performance, and better application experiences for users. For example, voice and Cisco TelePresence® applications have become integral parts of business success. Performance of these applications is crucial to their successful implementation. In order to improve application performance, companies have typically deployed two common solutions: providing additional bandwidth by deploying more network connections (for example, through multihoming), and using application optimization technologies (for example, Cisco Wide Area Application Services [WAAS]). Additional WAN bandwidth may improve aggregate throughput, but may not improve delay or loss for critical applications. Application optimization technologies such as Cisco WAAS can improve performance with data-reduction techniques, but fluctuating network performance can still affect applications. Performance Routing addresses network performance problems by helping enable the network to intelligently choose a path that currently meets the performance requirements of a specific application. In addition, using cost-based link policies, Performance Routing allows the network to intelligently choose link resources as needed to reduce operational costs.

Q. What are the benefits to customers of Cisco Performance Routing?

A.

- Bandwidth cost minimization allows enterprises to minimize traffic sent over expensive links or consolidate multiple flat-rate connections to fewer and lower-cost services.
- Automatic performance optimization reduces engineering operating expenses associated with manual network performance analysis and tuning of the routing infrastructure.
- Automatic performance optimization also helps ensure that mission-critical applications perform with the speed, availability, and reliability required for business success.
- Users can experience improved response time because Performance Routing can automatically detect and route around poorly performing paths using an alternate path.
- Due to the ability to actively detect and route around "black hole" conditions in the network, Performance Routing can minimize the effects of network outages.

Q. A lot of cases a customer has MPLS cloud with single provider and a backup cloud. How would you switchover from one to the other?

A. Cisco PfR Link Grouping is designed for this kind of application, with the MPLS groups as the primary link, DMVPN as backup group. See our Cisco [Performance Routing Link Groups](#) document for more information.

Q. What are the two components of Performance Routing? What do they do?

Q. Customer wants to use one provider cloud for one application and other cloud for other application, how would we do it using PFR?

A. By defining an application traffic class in PfR (one for each application), and then map the traffic class to a specific link group. This is the same way you would define paths for specific applications as well.

Q. What is the challenge with integrating PFR with other routing protocols?

A. The challenge is that not all routing protocols have the same level of usable information for PFR to work with. If the routing protocol data base doesn't have sufficient information, the information in the RIB is used (i.e., Protocol Independent Route Optimization (e.g. PIRO)). EIGRP has the concept of "feasible successor" that can be used as an alternate route. OSPF, on the other hand, only has one route in the routing table for a given destination, which forces PFR to use the parent route information in the RIB to satisfy its requirements.

Q. What are the newest features in Performance Routing (PFR)?

A. There are two new features recently released:

1. **Protocol Independent Route Optimization (PIRO)**. You can find this in the Cisco Feature Navigator (FN) listed under "**PFR - Protocol Independent Route Optimization (PIRO)**".

Protocol Independent Route Optimization (**PIRO**) has introduced the ability of Performance Routing (PFR) to search for a parent route, an exact matching route, or a less specific route in the IP Routing Information Base (RIB). This allows PFR to be deployed in any IP-routed environment including Interior Gateway Protocols (IGPs) such as **OSPF** and **IS-IS** since it is performing less direct specific control of the underlying routing protocols.

2. The next new feature allows **PFR** to Control EIGRP Routes with mGRE DMVPN Hub-and-Spoke Support, aptly named PfR EIGRP mGRE DMVPN, Hub-and-Spoke support. In the Cisco Feature Navigator (FN) this is named as "**PFR EIGRP mGRE DMVPN Hub-and-Spoke support**".

This new PFR EIGRP mGRE DMVPN Hub-and-Spoke Support feature has introduced the ability to inject routes into the EIGRP routing table, which allows Performance Routing (PFR) to control prefixes and applications over EIGRP routes. This feature also adds support for multipoint Generic Routing Encapsulation (mGRE) Dynamic Multipoint Virtual Private Network (DMVPN) deployments that follow a hub-and-spoke network design.

Q. What are the list of feature names in Cisco Feature Navigator (FN) for Performance Routing (PFR)?

A. There are three keywords you should look for and those are "**OER**", "**PFR**", and "**Performance Routing**". Here is the list of PFR supported feature names as of July, 2011:

FN Keyword	Full Feature Name Syntax
OER	OER - Application Aware Routing with Static Application
OER	OER - Inbound Optimization thru BGP
OER	OER - Voice Traffic Optimization
OER	OER Border Router Only Functionality
OER	OER Port and Protocol Based Prefix Learning
OER	OER Support for Cost Based Optimization and Traceroute Reporting
OER	OER Support for Policy-Rules Configuration and Port-Based Prefix Learning
OER	OER VPN IPsec/GRE Tunnel Optimization

Q. Customer wants to use one provider cloud for one application and other cloud for other application, how wo

Performance_Routing_FAQs

Performance Routing	Performance Routing - Link Groups
Performance Routing	Performance Routing with NBAR/CCE Application Recognition
PfR	PfR - Protocol Independent Route Optimization (PIRO)
PfR	PfR EIGRP mGRE DMVPN Hub-and-Spoke support
PfR	PfR RSVP Control
PfR	PfR Data Export v1.0 NetFlow v9 Format

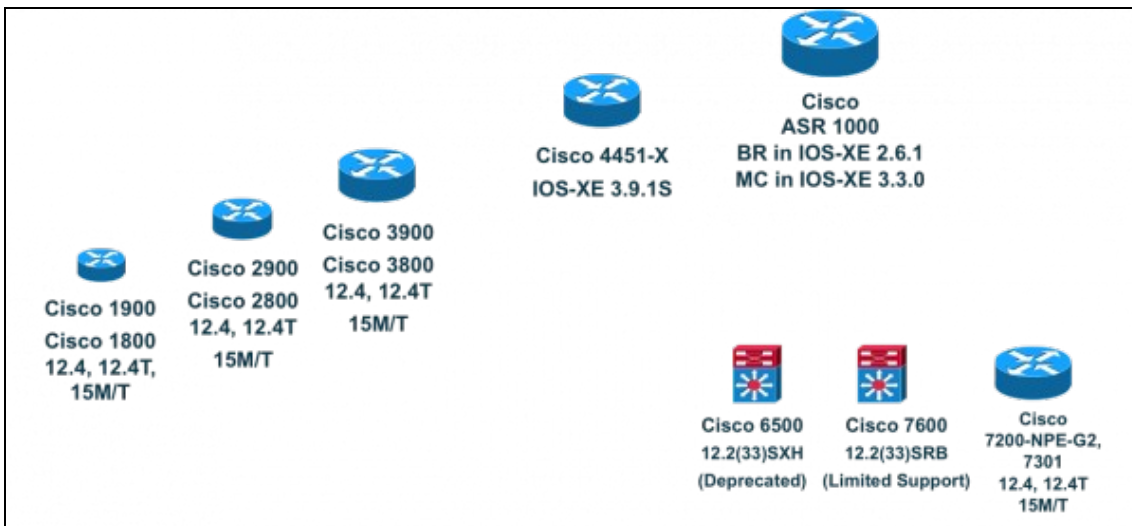
Q. What Cisco Hardware Platforms is PfR supported on? What is the recommended hardware? Are there any PfR deployment guides?

A. See Below.

Q. What are the base hardware requirements for Cisco Performance Routing (PfR)?

A. Refer to Figure 1.

Figure 1. Cisco Performance Routing Platform Support

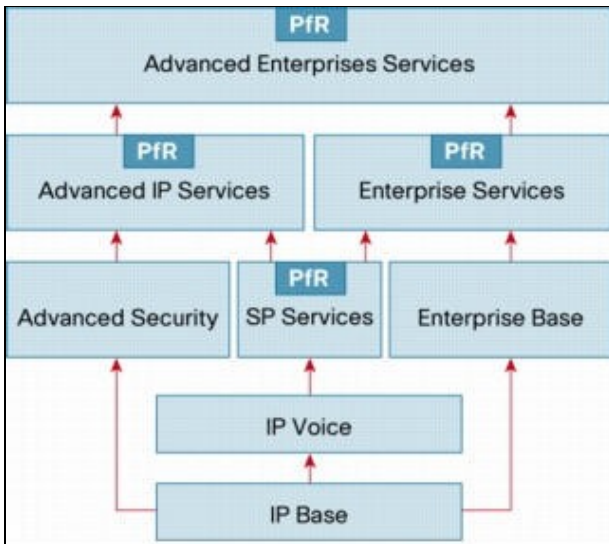


Q. What are the Classic Software Image requirements for Cisco Performance Routing (PfR)?

A. Refer to Figure 2.

Figure 2. Classic Cisco IOS Software Feature Sets

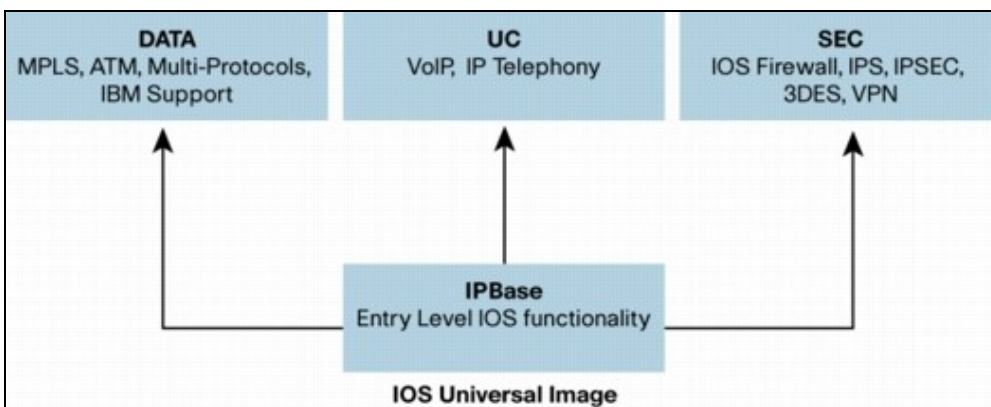
Q. What are the list of feature names in Cisco Feature Navigator (FN) for Performance Routing (PfR)?



Q. What are the New Universal Software Image requirements for Cisco Performance Routing (PfR)?

A. Refer to Figure 3.

Figure 3. New Cisco ISR G2 Simplified Feature Sets (**PfR is within the DATA Package**)



Q. What are the Software Image requirements for Cisco Performance Routing (PfR) on the ASR1k?

- ASR1001: Use Universal Image (U or UK9) with Advanced IP Services (AIS) or Advanced Enterprise Services (AES) technology package license
- All other ASR1000 (ASR1002-F, ASR1002, ASR1004, ASR1006, ASR1013): Use Advanced IP Services (AIS/AISK9) or Advanced Enterprise Services (AES/AESK9) images

Q. Is there a feature matrix that lists the features with platform support and correlates them with the Cisco IOS Software releases?

A. Consult the table below.

Obviously certain release trains correspond to specific platforms as most customers already are aware of. For example, the SR and 15S trains are for the 7600 platform, the 15SY and SX trains are for the **6500** Series

Q. What are the Classic Software Image requirements for Cisco Performance Routing (PfR)?

Performance_Routing_FAQs

platforms, the Cisco XE train is for the **ASR1000** Series platforms, and the mainline and T-trains are cross-platform releases.

Available starting with Cisco IOS Software Release **12.3(8)T**, the Cisco **PfR** solution supports both the border router and the master controller functions. The following list show where new features and functionality have been initially added to **PfR** since its inception:

IOS/XE Software Release	PfR Feature Name
12.3(8)T	Initial OER support
12.3(11)T	Added VPN IP Security (IPSec) and generic routing encapsulation (GRE) tunnel optimization
12.3(11)T	Added port- and protocol-based prefix learning
12.3(11)T	Added support for policy-rules configuration
12.3(14)T	Added support for cost-based optimization and traceroute reporting
12.4(2)T	Added application-aware routing: Policy-Based Routing (PBR)
12.4(2)T	Added active probe source address
12.4(6)T	Added voice traffic optimization
12.4(9)T	Added differentiated services code point (DSCP) monitoring
12.4(9)T	Added BGP inbound optimization
12.2(33)SRB	Added support on Cisco 7600 Series
12.2(33)SXH	Added border router support on Cisco Catalyst® 6500 Series
12.4(24)T	Added PfR - Protocol Independent Route Optimization (PIRO)
15.0(1)M	Added PfR EIGRP mGRE DMVPN Hub-and-Spoke Support
15.1(2)T	Added the new PfR command syntax
15.2(1)T	Added PfR RSVP Control support
15.2(2)T	Added PfR Data Export v1.0 NetFlow v9 Format support
2.6.1	Added support for PfR Border Router features on Cisco ASR 1000 Series
3.3.0S	Added support for PfR Master Controller on Cisco ASR 1000 Series
3.4.0S	Added support for the PfR RSVP Control and PfR Data Export v1.0 NetFlow v9 Format features on Cisco ASR 1000 Series

You can also use Feature Navigator ([FN](#)) to precisely determine which features are included in specific releases on specific platforms of your choice.

Q. Is Cisco Performance Routing supported globally?

A. Yes, the Cisco Technical Assistance Center (TAC) provides worldwide PfR customer support.

Q. Can you name some of the PfR customers?

A. As much as we would like to show how successful PfR is today, we cannot disclose our customers names publicly without their permission. PfR has well over 200 Major known customers today and growing.

PfR Active Probing through Cisco (IP SLAs)

Q. Is there a feature matrix that lists the features with platform support and correlates them with the Cisco IOS S

Q. Does Performance Routing require explicit IP SLA's configuration?

A. No; by default, Performance Routing automatically configures **IP SLA's** Internet Control Message Protocol (**ICMP**) probes for you. Depending on the type of traffic being controlled, you may want to use a different type of **IP SLA's** probe. For example, if you want to control voice traffic through Performance Routing using a jitter-based policy, you will need to explicitly configure an **IP SLA's** jitter probe and specify the target IP address of an appropriate **IP SLA's** Responder.

Q. Do I have to understand IP SLA's comprehensively?

A. No, by default the activation of IP SLA's and active monitoring using probes occurs without any interaction or awareness of the person configuring Performance Routing (PfR). The Border Router (BR) sends the IP SLA's results to the Master Controller where routing decisions are made based on the policy configuration.

Q. How does Performance Routing use IP SLA's active probing?

A. Performance Routing (PfR) can use passive monitoring, active monitoring, or both types of monitoring to determine the performance of prefixes on all possible paths. You can configure the IP SLA's probes to use system defaults or instead configure them in the following ways:

- If active probing is enabled, but no other specific probe configuration is created (using system defaults), Performance Routing dynamically sends out IP SLA's ICMP Echo probes to a subset of Netflow Top Talker destinations if you previously configured Performance Routing to build a Top Talker list.
- IP SLA's UDP Jitter active probe can calculate the voice over IP (**VoIP**) Mean Opinion Score (**MOS**) by measuring delay, loss, and jitter by generating simulated VoIP codec's as the probe.

Q. What is a Mean Opinion Score or (MOS) metric?

A. **Mean Opinion Score** - MOS is a Voice Call Metric which provides a numerical indication of the perceived quality of received media after compression and transmission. The MOS is expressed as a single digit number in the range from 1 to 5, where 1 is lowest (worst) perceived audio quality, and 5 is the highest (best) perceived audio quality.

MOS tests for voice are specified by ITU-T recommendation P.800. See the Wikipedia's Mean Opinion Score page for more information, or Cisco's Understanding Codecs section on Codec Mean Opinion Scores.

Q. What is the overhead for PfR IP SLA's Fast Probing?

A. PfR set in Fast Mode probes all possible exits all the time, which can add significant overhead in exchange for the faster path error detection provided. This can be useful in cases where critical application traffic paths must be maintained like VoIP or Video.

Q. If a customer wants to use Cisco IP SLAs to test and bring backup path up when the primary path goes down or Out Of Policy (OOP), can PfR help automate this?

A. Yes. Using PfR Mode Monitor Active, Mode Monitor Both, or Mode Monitor Fast will utilize IP SLAs under the covers automatically. Essentially whenever a customer chooses to use any PfR active mode, then yes IP SLA will be used to determine when paths or exit links are up or down.

Q. In PfR Active Fast Mode, why can't we just change the probe frequency to match fast mode? Why have another mode?

A. In PfR Active mode alone, PfR only probes the active exit alone and continues to do so until an out of policy event is detected on that exit. At that moment in Active mode, all the exits are probed to find and identify the next INPOLICY exit. In contrast to that, in PfR Fast Mode, all the defined exits are probed all the time (continually).

Q. I found something on the PfR wiki that seems to suggest multiple probes to same destination (different ports) is ?better? than a single probe. Would someone be able to elaborate?

A. Jitter/UDP probe relies on control packet. If control packet is lost then no measurement is done. So if you only have one target and if control packet is lost then there will be no probe results. Now, if the link is totally unreachable or high loss it wouldn't matter. But if there is low loss then you would be better off having three probes. At least one other probe will succeed. It is okay to lose probe packet. The other thing is with the MOS. MOS is calculated for the entire probe. With one target you will only have one MOS value. With three targets you will have 3 MOS values.

Command-Line Interface

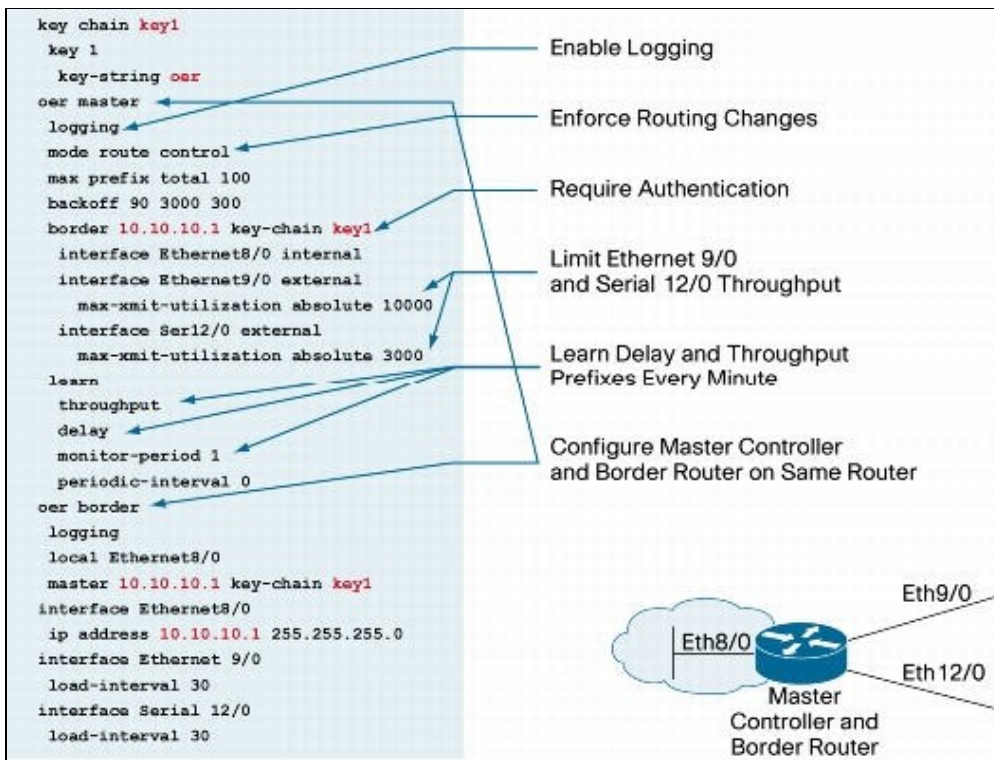
Q. Is Performance Routing (PfR) configuration complex?

A. Basic Performance Routing configuration using the command-line interface (CLI) is very simple, with most of the configuration performed on the Master Controller. For an example of a simple link-utilization-based Performance Routing policy, see Figure 3.

Note: The legacy 'oer' keyword is currently being deprecated and is replaced by 'pfr' in the CLI in Cisco IOS Release 15.1(2)T and Cisco IOS XE Release 3.1S. Optimized Edge Routing (OER) evolved into Performance Routing (PfR). OER represents the older technology that only optimized destination networks (for example, prefixes), PfR represents the newer technology that is application aware and much more flexible and powerful.

Figure 3. **Cisco Performance Routing Configuration**

Q. In PfR Active Fast Mode, why can't we just change the probe frequency to match fast mode? Why have another



Inbound Optimization

Q. Can inbound traffic be optimized?

A. Yes, Performance Routing can optimize inbound traffic. Performance Routing downgrades a non-preferred entrance by pre-pending BGP autonomous system numbers or by adding a community to external route advertisements associated with this entrance.

Learn

Q. When using the learn delay function, Performance Routing learns only prefixes that generate a three-way handshake during the learning period. Does the learn delay function learn only TCP traffic?

A. Yes, bidirectional TCP traffic is required to use the learn delay function of Performance Routing. The learn delay function is limited to TCP traffic, but throughput-based learning can be used for non-TCP traffic. Using the learn delay function, Performance Routing learns the top N prefixes with the highest delay values.

Q. Does learn delay learn the top N prefixes with the highest or worst delay?

A. Yes.

Q. In PfR, is "learning" done using passive or active monitoring?

A. When learning is enabled, PfR learns traffic-classes passively from data gathered through Netflow.

Q. When using learn throughput, Performance Routing identifies, through NetFlow, the top N prefixes that have the most throughput during the learn period. Does a learn throughput configuration allow Performance Routing to learn prefixes for which there is any sort of IP traffic?

A. Yes, unless learn filter is configured.

Q. Is Performance Routing (PfR) configuration complex?

Q. Why does the learn traffic-class count not properly reflected via "show oer master"?

A. An access-list under a pfr learn list is not used for a filter. It is used to define the application signature. Applications can be identified either using port, protocol, dscp or any combination of it. Each entry in an access-list serves as a different application.

Here is the example. Let say all my voice traffic are marked with dscp ef and I have following range of prefixes at the branch office:

10.1.1.0/24 To 10.1.15.0/24

I want to learn and optimize voice traffic to this branch office. In this case I would configure the following items:

1. A Prefix list for filtering destination networks.

```
Prefix-list BRANCH seq 10 permit 10.1.16.0/20
```

2. Next an Access-list to define VoIP traffic

```
Ip access-list extended VOIP  
Permit ip any any dscp ef
```

3. Then Learn the VoIP traffic destined to Branch office.

```
Learn  
List seq 10 refname BRANCH_VOIP  
Traffic-class access-list VOIP filter BRANCH  
Aggregation prefix-length 24
```

With above config it will learn following traffic-classes:

```
10.1.1.0/24 dscp ef  
10.1.2.0/24 dscp ef  
10.1.3.0/24 dscp ef  
10.1.4.0/24 dscp ef  
10.1.5.0/24 dscp ef  
10.1.6.0/24 dscp ef  
10.1.7.0/24 dscp ef  
10.1.8.0/24 dscp ef  
10.1.9.0/24 dscp ef  
10.1.10.0/24 dscp ef  
10.1.11.0/24 dscp ef  
10.1.12.0/24 dscp ef  
10.1.13.0/24 dscp ef  
10.1.14.0/24 dscp ef  
10.1.15.0/24 dscp ef
```

Miscellaneous

Q. Can a Performance Routing Master Controller and a Border Router be configured on the same physical router?

A. Yes, although scaling issues in your network may require that the MC and BR be deployed on separate physical routers.

Q. When do I need a dedicated Master Controller?

A. If you need to control more than 1000 traffic classes, it is recommended that you configure a dedicated Master Controller. If you are deploying Performance Routing for a data center, we recommend that you configure a dedicated Master Controller. For small office or home office (SOHO) setups, it is acceptable if the Master Controller and the Border Router coexist on the same physical router.

Q. Does PfR have authentication (key-chains) between PfR Master Controllers (MC's) and Border Routers (BR's)?

A. Yes. See the PfR Configuration Guide that covers this concept quite well. Cisco IOS Performance Routing Configuration Guide, [Release 15.2M&T](#) for more information.

Q. In PfR terminology, what is an "Internal Interface" and an "External Interface"? Do the Border Routers (BR's) use the external interfaces for PfR Peering with each other?

A. A basic PfR definition for these is best described as follows. PfR expects traffic to be flowing through the defined Internal to the External interfaces or visa-versa. BR's will only Peer with each other over the Internal Interfaces.

Q. How do I use Performance Routing for performance monitoring when I only have one external interface?

A. Performance Routing is designed for multihoming, which requires a minimum of two external interfaces. However, Performance Routing can be configured to monitor application performance over a single interface by creating a second dummy interface, which will not forward packets. Use a Dialer interface for the second interface but do not attach any physical interfaces to the Dialer. The Dialer interface needs an IP address because Performance Routing checks that the interfaces are routable. The IP address assigned to the Dialer does not need to be advertised.

The following example configuration shows how to add a Dialer interface as the second dummy external interface to an existing configuration:

On PfR Border Router

```
configure terminal
interface Dialer 1
ip address 10.9.9.9 255.255.255.254
end
```

On PfR Master Controller

```
configure terminal
pfr master
border 10.1.1.1
interface Dialer 1 external
end
```

Performance_Routing_FAQs

The following is an example Master Controller configuration after adding the dummy Dialer interface:

On PfR Master Router

```
pfr master
border 10.1.1.1 key-chain key1
interface Dialer1 external
interface Ethernet2/0 external
interface Ethernet0/0 internal
!
learn
throughput
periodic-interval 0
monitor-period 1
```

Q. How do I deploy Performance Routing in networks not running internal BGP (iBGP) within the enterprise?

A. To synchronize routing within the enterprise and take advantage of optimal Performance Routing routes, route redistribution into the local Interior Gateway Protocol (IGP; for example, EIGRP, OSPF, or RIP) needs to occur. Static routes injected by Performance Routing are tagged by an identifier that can be specifically redistributed. If you are running iBGP, BGP routes are usually not redistributed into an IGP. For more configuration details about route redistribution into IGPs, see the Setting Up OER Network Components module.

Q. What Performance Routing limitations exist on Cisco Catalyst® 6500 Switches and Cisco 7600 Series Routers?

A. In these platforms, due to architectural limitations with NetFlow on the EARL7 chip, Performance Routing cannot determine performance (delay/loss/reachability) characteristics from passive monitoring of TCP flows; Performance Routing can only measure passive throughput. Note that these platforms support active IP SLA measurements and throughput-based load balancing. In addition, Cisco Catalyst 6500 Switches support only Border Router functionality and cannot be used as a Performance Routing Master Controller.

Q. What does the "Disable hardware switching to enable netflow aggregation export in version 9 format" message mean? Do I have to disable hardware switching?

A. This is just an informational message. You do not have to disable hardware switching, and you can safely ignore this message. Performance Routing uses NetFlow Version 9 format internally, but the set of hardware you are running does not support NetFlow Version 9 export format. All this message is saying is that the platform cannot export version 9 when hardware switching. Note that while Performance Routing uses NetFlow for its operation, Performance Routing performance information is not exported in NetFlow export messages.

Q. Is Performance Routing VRF-aware?

A. No.

Q. Does Performance Routing enforce symmetric routing?

A. No, Performance Routing selects the best path based on real-time performance data. Symmetric routing is not always optimal. In an enterprise deployment, if Performance Routing is deployed on both ends of the network, it will select the best path in both directions. However, these bidirectional decisions are made independently and do not guarantee path symmetry.

Q. How do I use Performance Routing for performance monitoring when I only have one external interface?

Q. Will Performance Routing cause route oscillation?

A. Performance Routing has a hold-down period that prevents route oscillation. After the traffic class is rerouted through a new or better link, by default, Performance Routing does not change the route of a traffic class for a period of time known as the hold-down period. The hold-down period can be configured, and the default is 300 seconds. The only exception to the rule occurs if the traffic class becomes unreachable prior to the completion of the hold-down period.

Q. How are Performance Routing system logging (syslog) messages formatted?

A.

Sep 22 11:27:14.990: %OER_MC-5-NOTICE: Route changed Prefix 100.1.1.0/26, BR 1.1.1.5, i/f Et2/0, Reason Non-OER, OOP Reason None

Sep 22 11:27:32.042: %OER_MC-5-NOTICE: Route changed Prefix 100.1.1.0/26, BR 1.1.1.6, i/f Et2/0, Reason Delay, OOP Reason Timer Expired

Sep 22 11:27:52.242: %OER_MC-5-NOTICE: Prefix Learning WRITING DATA

Sep 22 11:29:55.350: %OER_MC-5-NOTICE: Prefix Learning STARTED

Sep 22 11:30:56.226: %OER_MC-5-NOTICE: Prefix Learning WRITING DATA

Sep 22 11:32:59.434: %OER_MC-5-NOTICE: Prefix Learning STARTED

Sep 22 11:33:04.974: %OER_MC-5-NOTICE: Passive ABS Delay OOP Prefix 100.1.1.0/26, delay 451, BR 1.1.1.6, i/f Et2/0

Sep 22 11:33:04.974: %OER_MC-5-NOTICE: Uncontrol Prefix 100.1.1.0/26, OOP, mode select-exit good

Sep 22 11:33:59.018: %OER_MC-5-NOTICE: Discovered Exit for Prefix 100.1.1.0/26, BR 1.1.1.6, i/f Et2/0

Passive Monitoring (NetFlow)

Q. Does Performance Routing require explicit NetFlow configuration?

A. No, when Performance Routing is enabled, it automatically configures NetFlow for you.

Q. How does Performance Routing use passive monitoring?

A. For passive monitoring, Performance Routing uses many of the services built into the Cisco IOS Software including:

- Cisco Performance Routing builds a list of top throughput and top latency talkers.
- Based on the list of learned and configured prefixes, Performance Routing passively monitors TCP flags for traffic on every flow (of the current exit) to measure latency, packet loss, and reachability.

Policy

Q. Is the PfR Monitor mode specific to a policy? Would it override what's in the PfR global policy?

A. Anything defined in a pfr-map overrides the value in the global policy. If the metric is not specified in the pfr-map, then it inherits from the Global Configuration Policy that metric in its place. So it is possible to use a pfr-map to contain both specific policy and global policy items at the same time, making it more simple to only enter what you need to override from the global policy, and no more.

Q. Can you please elaborate on : INPOLICY and OOPOLICY?

A. The INPOLICY and OOPOLICY (OOP) conditions in PfR are related to whether a traffic-class is within the parameters of its current policy (INPOLICY) or not (OOP). In addition to traffic-classes, links can be INPOLICY or OOP. Policy parameters defined in a pfr-map override the global policy. If a parameter is not specified in the pfr-map, it inherits the value from the global policy.

Q. If the configured policy cannot be met because available paths are not meeting requirement, then what happens next?

A. Under such conditions, PfR is designed to un-control the prefix at that point. This will send the prefix, or traffic class, back to normal routing control. This is PfR's fail-safe default action.

Q. When does PfR monitoring actually start? Lets say am using external interface-1 , and that interface goes OOPOLICY (OOP). When does PfR start monitoring all the other available defined external links?

A. This depends on how PfR has been configured. If your system is configured for PfR Mode Monitor Active, PfR will start probing all of the available exits as soon as the OOP event occurs. Active mode includes mode monitor active, both, active-throughput, and fast. Once PfR locates a new exit that is in policy, the traffic is switched to that exit. At that point only the new current exit is monitored on a regular basis. In PfR Fast Mode, the system will probe all exits all of the time (continually), which allows route changes to be made very quickly (~3 seconds).

Q. When using link grouping and a link goes out of policy (OOP), can Performance Routing revert back to the primary link?

A. Yes, if appropriate link grouping is configured, Performance Routing can revert back to the primary link when the primary link returns to an in-policy condition.

Q. When using select-exit best with mode active or mode both, does Performance Routing also initially probe all exits to find the best exit and then only probe all exits again if the current exit goes OOP?

A. The configuration of select-exit best is independent of how often Performance Routing probes exits. The probing is driven by the timer, the state of the prefix, and an OOP condition. In active, fast, or both modes Performance Routing always selects the best exit. Select-exit best configures Performance Routing to control or un-control a traffic class if the traffic class is OOP.

Q. When using select-exit best with mode fast, does Performance Routing always probe all exits and, at each periodic interval, can Performance Routing move the prefix to the best exit if the identified best exit is different from the current exit?

A. In fast mode, Performance Routing always probes all exits (continually) to allow Performance Routing to reroute traffic quickly.

Route Control

Q. What is a parent route?

A. A parent route is a route that is equal to, or less specific than, the destination prefix of the traffic class being optimized by Performance Routing. The parent route should have a route through the Performance Routing external interfaces. All routes for the parent prefix are called parent routes. For Performance Routing to control a traffic class on a Performance Routing external interface, the parent route must exist on the Performance Routing external interface. BGP and Static routes qualify as Performance Routing parent routes. In Cisco IOS Release 12.4(24)T and later releases, any route in RIB, with an equal or less specific mask than the traffic class, will qualify as a parent route.

Q. Can you elaborate more on the Parent Route and why it's so important to PFR?

A. Yes. For any route that PFR modifies or controls (BGP, Static, PIRO, EIGRP, PBR), having a Parent prefix in the routing table eliminates the possibility of a routing loop occurring. This is naturally a good thing to prevent in routed networks.

Q. Which parent route IP routing protocols are supported?

A. BGP and static routes are supported in all current releases. With the introduction of Performance Routing Protocol Independent Route Optimization (PIRO) in Cisco IOS Software Release 12.4(24)T and later releases, any route in the RIB, with an equal or less specific mask than the traffic class, will qualify as a parent route. PIRO allows Performance Routing to be used with any IP routing protocol. EIGRP support was introduced in Cisco IOS Release 15.0(1)M.

Q. Can Performance Routing control a traffic class that is more granular than a network address? If yes, how?

A. Yes, Performance Routing can control applications that are defined using access control lists (ACLs), network-based application recognition (NBAR) application IDs, or fixed static signatures (Performance Routing has a predefined list of signatures). Performance Routing controls applications using route maps as part of policy-based routing. Route maps are automatically generated, and no user configuration is required.

Q. Is prefix splitting supported with Cisco Performance Routing? If so, how is it done?

A. Yes. Prefix splitting means that a more specific route is derived from a prefix (for example, splitting a host route [32] from a /24 prefix). If you configure a /24 prefix to be optimized but Cisco Performance Routing finds only /16 route in the routing table, it injects a /24 route with the new exit link using the attributes from the /16 route. This injection affects only the internal routes, and this new information is not propagated outside the autonomous system. If you configure a prefix and Performance Routing finds only a default route for it (from any Border Router), the following occurs:

- If BGP is not running, a new static route is injected.
- If BGP is running, a new BGP route is injected.

Q. How does the Master Controller inject a route?

A. Master Controller-to-Border Router communication takes place through a Cisco internal mechanism that does not rely on an iBGP session between them. When the Master Controller determines the best path for a specific prefix (whether learned or user-configured), it sends a route control command to a selected Border Router that has the best exit. If an exact route (BGP or static) already exists for this prefix, Cisco Performance Routing either changes the local preference for a BGP route or modifies the next hop of a static

route. If a route does not exist for this prefix but there is a super route or a default route covering this prefix, the Cisco Performance Routing Master Controller injects a route that matches exactly the length of this prefix (called "prefix splitting") by sending a route control command to the selected Border Router. The Border Router then redistributes this route internally with its iBGP peers.

Q. Does the Master Controller have visibility of the Border Router BGP routing tables that it controls?

A. No, the Master Controller does not have a copy of the BGP or static routing tables from any of the Border Routers it controls. The Master Controller does not need to run any routing protocols and is typically not in the data forwarding path.

Q. What happens if the Master Controller fails?

A. If the Performance Routing Master Controller fails, the Border Routers notice that their internal communication to the Performance Routing Master Controller is no longer active. The Cisco IOS Software in the Performance Routing Border Router determines which routes are under Performance Routing control and withdraws these routes, which restores everything back to normal routing. The network experiences the same performance that was occurring before turning on Cisco Performance Routing.

Q. Does Performance Routing support the notion of redundant Master Controllers?

A. Yes, redundant Master Controllers can be deployed using Hot Standby Router Protocol (HSRP). Note that in this scenario, Master Controller failover is not stateful. When the Border Router learns that the primary Master Controller is no longer reachable, it will purge all Performance Routing-learned routes and restore itself to traditional IP routing. After the Master Controller HSRP switchover takes place, the Master Controller-to-Border Router communication is restored, and the system will restart the Performance Routing learning process.

Q. How does PfR verify that the application is reachable?

A. PfR must have a successful probe completion to the destination prefix or network where the application is being hosted (prefix reachability); For Application reachability, the PfR Master Controller (MC) needs to detect that there is bidirectional traffic for a given application before it can be considered reachable.

Scale

Q. How is CPU utilization affected when Performance Routing is enabled?

A. This depends on the existing CPU load and the number of active traffic classes that are controlled.

Q. How many Border Routers are supported, and how many exits in total?

A. A total of 10 Border Routers and 20 exits per Border Router are supported on any one Master Controller for a combined total of up to 200 Interfaces or Exits.

Traffic Class Optimization

Q. How does PfR Profile Traffic Classes?

A. PfR Traffic Class Profiling must be done before optimizing traffic. In order to do this PfR must learn of the traffic classes by monitoring the traffic flowing through the border router(s). To optimize traffic routing, subsets of the total traffic must be identified, and these traffic subsets are named traffic classes. The list of

Q. How does the Master Controller inject a route?

Performance_Routing_FAQs

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.

For more information on this subject, please see our [Using Performance Routing to Profile Traffic Classes document](#) on Cisco.com.

Q. How does PfR "Select the best traffic class" based on variance or priority and what is the elimination process for selecting this path?

A. The concept of variance in PfR is similar in implementation to the variance keyword that enables EIGRP to install multiple unequal cost routes in the local routing table. Variance is a means to specify a range in which two unequal values are considered similar enough to be treated as equal.

In general in the context of PfR, **variance** is a percentage value ranged from 1-100. For example, if the lowest delay on a link is 23ms, and a 10 percent variance is configured, delay values from 23 to 25 (23 + 2.3 rounded down) will be considered equal.

Specifically in case of a delay metric, if an **absolute** threshold value is configured, and if the lowest delay is within threshold, the variance in delay will not be considered. Yet if the lowest delay is more than the threshold the variance will be considered. Thus variance works well with delay if the lowest delay is within the threshold. Variance also performs in this manner using PfR's loss, jitter, or MoS metrics as well.

For example, if the delay threshold is 100ms, and the 10% variance is configured,

```
case-a: delay_of_exit1 = 200ms,  delay_of_exit2 = 210ms,
=> we consider both exit1 and exit2 within the policy,
    because both are within 220ms. Here, 110% of the lowest
    delay is 220ms, as you can see.
```

```
case-b: delay_of_exit1 = 200ms,  delay_of_exit2 = 230ms,
=> we pick up only exit1, because exit2 is not within 220ms.
```

```
case-c: delay_of_exit1 = 10ms,   delay_of_exit2 = 90ms,
=> we consider both exit1 and exit2 within the policy as far as
    delay is concerned. This is because both extis are within the
    100ms threshold. In such a condition the variance isn't part
    of PfR's path consideration.
```

Q. Do I have to specify the traffic classes that are to be optimized?

A. No. The traffic classes to be optimized either can be specified using prefix-list, access-list, NBAR, or static application definition or can be learned in real time.

Q. If a traffic class is configured using an ACL with "any any" instead of specifying a destination subnet, is it supported? For example, "access-list 101 permit tcp any any eq telnet."

A. No, the default prefix cannot be learned, no passive statistics can be gathered for the default prefix, and route control cannot work.

Q. In a PfR System default configuration with a Aggregation-type prefix-length 24 defined, how is aggregation done for less specific routes in the RIB?

A. In PfR you can define what network boundary on which you wish aggregate traffic classes.

Q. How does PfR Profile Traffic Classes?

PfR Route Maps

Q. How is an application defined in PfR? How is aggregation done and defined in PfR?

A. In PfR, an application is defined by a pfr-map. A pfr-map defines the application by using an Access List, Extended Access List, or its Integrated NBAR features.

Q. Where and how are PfR Dynamic Route-Maps applied? Over which interface, internal or external, will the Dynamic Route-Maps be applied on the BR?

A. PfR Route-Maps will be applied onto the Internal Interfaces. For Dynamic Route-Maps to work, PfR requires a direct L3 connection, or a common subnet between the BR's. Without this direct connection PfR's PBR requirement will not be met and dynamic route maps will not function. You can also optionally use a direct L2 connection with a GRE Tunnel configured between the BR's Internal Interfaces to connect them as well in place of a common L3 subnet connection.

Q. In a PfR "PFR-MAP" , what does the associated number specify:

A. In PfR you can define a policy containing multiple traffic-classes. By using different sequence numbers you can control the priority. By default the lower sequence number will be the higher priority under that policy.

```
pfr-map MYMAP 30 <- what is 30
```

Note: The legacy 'oer' and 'oer-map' keywords are currently being deprecated and were replaced by 'pfr' and 'pfr-map' in the CLI in Cisco IOS Release 15.1(2)T and Cisco IOS XE Release 3.1S. Optimized Edge Routing (OER) evolved into Performance Routing (PfR). OER represents the older technology that only optimized destination networks (i.e., prefixes), PfR represents the newer technology that is application aware and much more flexible and powerful.