



Configuring Basic Performance Routing

Performance Routing (PfR) provides additional intelligence to classic routing technologies to track the performance of, or verify the quality of, a path between two devices over a Wide Area Networking (WAN) infrastructure to determine the best egress or ingress path for application traffic.

Cisco Performance Routing complements classic IP routing technologies by adding intelligence to select 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.

This document contains an introduction to the basic concepts and tasks required to implement Performance Routing using Cisco IOS XE Software.

- [Restrictions for Configuring Basic Performance Routing, on page 1](#)
- [Migrating to Cisco-SDWAN from PfR, on page 2](#)
- [Information About Performance Routing, on page 2](#)
- [How to Configure Basic Performance Routing, on page 9](#)
- [Configuration Examples for Configuring Basic Performance Routing, on page 17](#)
- [Additional References, on page 18](#)
- [Feature Information for Configuring Basic Performance Routing, on page 19](#)

Restrictions for Configuring Basic 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.



Note In Cisco IOS XE Release 3.3S, and later releases, master controller configuration is supported.

Migrating to Cisco-SDWAN from PfR

From Cisco IOS XE 17.4, PfRv3 is not supported. PfRv3 is integrated with IWAN and support for IWAN is discontinued from Cisco IOS XE 17.4 therefore ending the support for PfRv3.

Customers are encouraged to migrate to Cisco's SD-WAN solution, that has Application-aware Routing and a centralized controller to enable SLA-based routing along with performance measurement and monitoring.

The SD-WAN Team has IWAN-migration offerings to ease the transition to Cisco's SD-WAN solution, with particular focus on current IWAN customers, including exclusive commercial offers and technical resources. For more details, reach out to your Cisco representative.

Related Documents:

- <https://www.cisco.com/c/dam/en/us/td/docs/routers/sdwan/migration-guide/iwan-to-sdwan-migration-guide.pdf>

Information About Performance Routing

Performance Routing Overview

Performance Routing (PfR) is an advanced Cisco technology to allow businesses to complement classic routing technologies with additional serviceability parameters to select the best egress or ingress path. It complements these classic routing technologies with additional intelligence. PfR can select an egress or ingress WAN interface based upon parameters like reachability, delay, cost, jitter, MOS score, or it can use interface parameters like load, throughput and monetary cost. Classic routing (for example, EIGRP, OSPF, RIPv2, and BGP) generally focuses upon creating a loop-free topology based upon the shortest or least cost path.

PfR gains additional intelligence using measurement instrumentation. It uses interface statistics, Cisco IP SLA for active monitoring, and NetFlow for passive monitoring. No prior knowledge or experience of IP SLA or NetFlow is required, PfR automatically enables these technologies without any manual configuration.

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 technology can reduce network costs by facilitating more efficient load balancing and by increasing application performance without WAN upgrades.

PfR is an integrated Cisco IOS solution that allows you to monitor IP traffic flows and then define policies and rules based on traffic class performance, link load distribution, link bandwidth monetary cost, and traffic type. PfR provides active and passive monitoring systems, dynamic failure detection, and automatic path correction. Deploying PfR enables intelligent load distribution and optimal route selection in an enterprise network.

Performance Routing Versus Optimized Edge Routing

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 Optimized Edge Routing (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.

Performance Routing Versus Classic Routing Technologies

PfR was developed to identify and control network performance issues that traditional IP routing cannot address. In traditional IP routing, each peer device communicates its view of reachability to a prefix destination with some concept of a cost related to reaching the metric. The best path route to a prefix destination is usually determined using the least cost metric, and this route is entered into the routing information base (RIB) for the device. As a result, any route introduced into the RIB is treated as the best path to control traffic destined for the prefix destination. The cost metric is configured to reflect a statically engineered view of the network, for example, the cost metric is a reflection of either a user preference for a path or a preference for a higher bandwidth interface (inferred from the type of interface). The cost metric does not reflect the state of the network or the state of the performance of traffic traveling on that network at that time. Traditional IP routed networks are therefore adaptive to physical state changes in the network (for example, interfaces going down) but not to performance changes (degradation or improvement) in the network. Occasionally, degradation in traffic can be inferred from either the degradation in performance of the routing device or the loss of session connectivity, but these traffic degradation symptoms are not a direct measure of the performance of the traffic and cannot be used to influence decisions about best-path routing.

To address performance issues for traffic within a network, PfR manages traffic classes. Traffic classes are defined as subsets of the traffic on the network, and a subset may represent the traffic associated with an application, for example. The performance of each traffic class is measured and compared against configured or default metrics defined in an PfR policy. PfR monitors the traffic class performance and selects the best entrance or exit for the traffic class. If the subsequent traffic class performance does not conform to the policy, PfR selects another entrance or exit for the traffic class.

Basic Performance Routing Deployment

PfR is configured on Cisco routers using Cisco IOS command-line interface (CLI) configurations. Performance Routing comprises two components: the Master Controller (MC) and the Border Router (BR). A PfR deployment requires one MC and one or more BRs. Communication between the MC and the BR is protected by key-chain authentication. 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.

A PfR-managed network must have at least two egress interfaces that can carry outbound traffic and can be configured as external interfaces, see the figure below. These interfaces should connect to an ISP or WAN link (Frame-Relay, ATM) at the network edge. The router must also have one interface (reachable by the internal network) that can be configured as an internal interface for passive monitoring. There are three interface configurations required to deploy PfR: external interfaces, internal interfaces, and local interfaces.

PfR Border Router

The BR component resides within the data plane of the edge router with one or more exit links to an ISP or other participating network. 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. The BR is where all policy decisions and changes to routing in the network are enforced. The BR participates in prefix monitoring and route optimization by reporting prefix and exit link measurements to the master controller and then by enforcing policy changes received from the master controller.

The BR enforces policy changes by injecting a preferred route to alter routing in the network. A BR process can be enabled on the same router as a master controller process.

For more details about the Border router only functionality in Cisco IOS XE Releases 2, 3.1S and 3.2S, see the "Performance Routing Border Router Only Functionality" module. In Cisco IOS XE Release 3.3S, and later releases, master controller configuration is supported.

PfR Primary Controller

The primary controller is a single router that acts as the central processor and database for the Performance Routing system. The primary controller component does not reside in the forwarding plane and, when deployed in a standalone fashion, has no view of routing information contained within the Branch Router (BR). The primary controller maintains communication and authenticates the sessions with the BRs. The role of the primary controller is to gather information from the BR or BRs to determine whether or not traffic classes are in or out of policy, and to instruct the BRs how to ensure that traffic classes remain in policy using route injection or dynamic PBR injection.

In Cisco IOS XE Release 2, 3.1S and 3.2S, PfR supports the ASR 1000 series router as a border router and the primary controller must be running a Cisco IOS Release 15.0(1)M image. In Cisco IOS XE Release 3.3S, and later releases, primary controller configuration is supported.

PfR Component Version

When new PfR functionality is introduced that changes the API between the MC and the BR, the version number for the Performance Routing components, master controller and border router, is incremented. The version number of the master controller must be equal or higher to the version number for the border routers. The version numbers for both the master controller and the border routers are displayed using the **show pfr master** command. In the following partial output, the MC version is shown in the first paragraph and the BR versions are shown in the last column of the information for the border routers.

```
Router# show pfr master
OER state: ENABLED and ACTIVE
Conn Status: SUCCESS, PORT: 7777
Version: 2.0
Number of Border routers: 2
Number of Exits: 2
.
.
.
Border      Status  UP/DOWN      AuthFail  Version
1.1.1.2     ACTIVE  UP           00:18:57  0 2.0
1.1.1.1     ACTIVE  UP           00:18:58  0 2.0
.
.
.
```

The version numbers are not updated at each Cisco IOS XE software release for a specific release train, but if the Cisco IOS XE software image is the same release on the devices configured as a master controller and all the border routers, then the versions will be compatible.

Key Chain Authentication for PfR

Communication between the master controller and the border router is protected by key-chain authentication. The authentication key must be configured on both the master controller and the border router before

communication can be established. The key-chain configuration is defined in global configuration mode on both the master controller and the border router before key-chain authentication is enabled for master controller-to-border router communication. For more information about key management, see the "Managing Authentication Keys" section of the Configuring IP Routing Protocol-Independent Features chapter in the *Cisco IOS IP Routing: Protocol Independent Configuration Guide*.

PfR-Managed Network Interfaces

A PfR-managed network must have at least two egress interfaces that can carry outbound traffic and that can be configured as external interfaces. These interfaces should connect to an ISP or WAN link at the network edge. The router must also have one interface (reachable by the internal network) that can be configured as an internal interface for passive monitoring. There are three interface configurations required to deploy PfR:

- *External interfaces* are configured as PfR-managed exit links to forward traffic. The physical external interface is enabled on the border router. The external interface is configured as a PfR external interface on the master controller. The master controller actively monitors prefix and exit link performance on these interfaces. Each border router must have at least one external interface, and a minimum of two external interfaces are required in an PfR-managed network.
- *Internal interfaces* are used only for passive performance monitoring with NetFlow. No explicit NetFlow configuration is required. The internal interface is an active border router interface that connects to the internal network. The internal interface is configured as an PfR-internal interface on the master controller. At least one internal interface must be configured on each border router.
- *Local interfaces* are used only for master controller and border router communication. A single interface must be configured as a local interface on each border router. The local interface is identified as the source interface for communication with the master controller.

The following interface types can be configured as external and internal interfaces:

- ATM
- Channelized Interface (T3/STM1 down to T1)
- Fast Ethernet
- Gigabit Ethernet
- 10 Gigabit Ethernet
- Packet-over-SONET (POS)
- Serial
- Tunnel (not supported with NAT in Cisco IOS XE Releases 2, 3.1S, and later releases)
- VLAN (QinQ is not supported)

The following interface types can be configured as local interfaces:

- ATM
- Fast Ethernet
- Gigabit Ethernet
- 10 Gigabit Ethernet

- Packet-over-SONET (POS)
- Serial
- Tunnel (not supported with NAT in Cisco IOS XE Releases 2, 3.1S, and later releases)
- VLAN (QinQ is not supported)

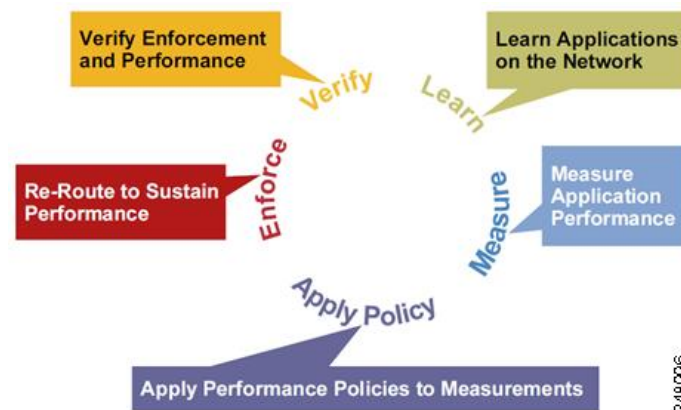
Performance Routing DMVPN mGre Support

- PfR does not support split tunneling.
- PfR supports hub-to-spoke links only. Spoke-to-spoke links are not supported.
- PfR is supported on DMVPN Multipoint GRE (mGRE) deployments. Any multipoint interface deployment that has multiple next hops for the same destination IP address is not supported (for example, Ethernet).

PfR Network Performance Loop

Every traditional routing protocol creates a feedback loop among devices to create a routing topology. Performance Routing infrastructure includes a performance routing protocol that is communicated in a client-server messaging mode. The routing protocol employed by PfR runs between a network controller called a master controller and performance-aware devices called border routers. This performance routing protocol creates a network performance loop in which the network profiles which traffic classes have to be optimized, measures and monitors the performance metrics of the identified traffic classes, applies policies to the traffic classes, and routes the identified traffic classes based on the best performance path. The diagram below shows the five PfR phases: profile, measure, apply policy, enforce, and verify.

Figure 1: PfR Network Performance Loop



To understand how PfR operates in a network, you should understand and implement the five PfR phases. The PfR performance loop starts with the profile phase followed by the measure, apply policy, control, and verify phases. The flow continues after the verify phase back to the profile phase to update the traffic classes and cycle through the process.

Profile Phase

In medium to large networks there are hundreds of thousands of routes in the RIB to which a device is trying to route traffic. Because performance routing is a means of preferring some traffic over another, a subset of

the total routes in the RIB has to be selected to optimize for performance routing. PfR profiles traffic in one of two ways, automatic learning or manual configuration.

- Automatic Learning—The device profiles the traffic that has to be performance routed (optimized) by learning the flows that pass through the device and by selecting those flows that have the highest delay or the highest throughput.
- Manual configuration—In addition to, or instead of learning, you can configure a class of traffic to performance route.

Measure Phase

After profiling traffic classes that are to be performance routed, PfR measures the performance metrics of these individual traffic classes. There are two mechanisms—passive monitoring and active monitoring—to measure performance metrics, and one or both could be deployed in the network to accomplish this task. Monitoring is the act of measuring at periodic intervals.

Passive monitoring is the act of measuring the performance metrics of the traffic flow as the flow is traversing the device in the data path. Passive monitoring uses NetFlow functionality and cannot be employed for measuring performance metrics for some traffic classes, and there are some hardware or software limitations.

Active monitoring consists of generating synthetic traffic using IP Service Level Agreements (SLAs) to emulate the traffic class that is being monitored. The synthetic traffic is measured instead of the actual traffic class. The results of the synthetic traffic monitoring are applied to performance route the traffic class represented by the synthetic traffic.

Both passive and active monitoring modes can be applied to the traffic classes. The passive monitoring phase may detect traffic class performance that does not conform to an PfR policy, and then active monitoring can be applied to that traffic class to find the best alternate performance path, if available.

Support for NetFlow or IP SLAs configuration is enabled automatically.

Apply Policy Phase

After collecting the performance metrics of the class of traffic to be optimized, PfR compares the results with a set of configured low and high thresholds for each metric configured as a policy. When a metric, and consequently a policy, goes out of bounds, it is an Out-of-Policy (OOP) event. The results are compared on a relative basis—a deviation from the observed mean—or on a threshold basis—the lower or upper bounds of a value—or a combination of both.

There are two types of policies that can be defined in PfR: traffic class policies and link policies. Traffic class policies are defined for prefixes or for applications. Link policies are defined for exit or entrance links at the network edge. Both types of PfR policies define the criteria for determining an OOP event. The policies are applied on a global basis in which a set of policies is applied to all traffic classes, or on a more targeted basis in which a set of policies is applied to a selected (filtered) list of traffic classes.

With multiple policies, many performance metric parameters, and different ways of assigning these policies to traffic classes, a method of resolving policy conflicts was created. The default arbitration method uses a default priority level given to each performance metric variable and each policy. Different priority levels can be configured to override the default arbitration for all policies, or a selected set of policies.

Enforce Phase

In the PfR enforce phase (also called the control phase) of the performance loop, the traffic is controlled to enhance the performance of the network. The technique used to control the traffic depends on the class of

traffic. For traffic classes that are defined using a prefix only, the prefix reachability information used in traditional routing can be manipulated. Protocols such as Border Gateway Protocol (BGP) or RIP are used to announce or remove the prefix reachability information by introducing or deleting a route and its appropriate cost metrics.

For traffic classes that are defined by an application in which a prefix and additional packet matching criteria are specified, PfR cannot employ traditional routing protocols because routing protocols communicate the reachability of the prefix only and 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.

Verify Phase

During the PfR enforce phase if a traffic class is OOP, then PfR introduces controls to influence (optimize) the flow of the traffic for the traffic class that is OOP. A static route and a BGP route are examples of controls introduced by PfR into the network. After the controls are introduced, PfR will verify that the optimized traffic is flowing through the preferred exit or entrance links at the network edge. If the traffic class remains OOP, PfR will drop the controls that were introduced to optimize the traffic for the OOP traffic class and cycle through the network performance loop.

PfR and the Enterprise Network

Enterprise networks use multiple Internet Service Provider (ISP) or WAN connections at the network edge for reliability and load distribution. Existing reliability mechanisms depend on link state or route removal on the border router to select the best exit link for a prefix or set of prefixes. Multiple connections protect enterprise networks from catastrophic failures but do not protect the network from brownouts, or soft failures, that occur because of network congestion. Existing mechanisms can respond to catastrophic failures at the first indication of a problem. However, blackouts and brownouts can go undetected and often require the network operator to take action to resolve the problem. When a packet is transmitted between external networks (nationally or globally), the packet spends the vast majority of its life cycle on the WAN segments of the network. Optimizing WAN route selection in the enterprise network provides the end-user with the greatest performance improvement, even better than LAN speed improvements in the local network.

Although many of the examples used to describe PfR deployment show ISPs as the network with which the edge devices communicate, there are other solutions. The network edge can be defined as any logical separation in a network: can be another part of the network such as a data center network within the same location, as well as WAN and ISP connections. The network, or part of the network, connected to the original network edge devices must have a separate autonomous system number when communicating using BGP.


PfR is implemented as an integrated part of Cisco core routing functionality. Deploying PfR enables intelligent network traffic load distribution and dynamic failure detection for data paths at the network edge. While other routing mechanisms can provide both load distribution and failure mitigation, only PfR can make routing adjustments based on criteria other than static routing metrics, such as response time, packet loss, path availability, and traffic load distribution. Deploying PfR allows you to optimize network performance and link load utilization while minimizing bandwidth costs and reducing operational expenses.

Typical Topology on Which PfR is Deployed

The figure below shows a typical PfR-managed enterprise network of a content provider. The enterprise network has three exit interfaces that are used to deliver content to customer access networks. The content provider has a separate service level agreement (SLA) with a different ISP for each exit link. The customer

access network has two edge routers that connect to the Internet. Traffic is carried between the enterprise network and the customer access network over six service provider (SP) networks.

Figure 2: A Typical PfR Deployment



PfR monitors and controls outbound traffic on the three border routers (BRs). PfR measures the packet response time and path availability from the egress interfaces on BR1, BR2 and BR3. Changes to exit link performance on the border routers are detected on a per-prefix basis. If the performance of a prefix falls below default or user-defined policy parameters, routing is altered locally in the enterprise network to optimize performance and to route around failure conditions that occur outside of the enterprise network. For example, an interface failure or network misconfiguration in the SP D network can cause outbound traffic that is carried over the BR2 exit interface to become congested or fail to reach the customer access network. Traditional routing mechanisms cannot anticipate or resolve these types of problems without intervention by the network operator. PfR can detect failure conditions and automatically alter routing inside of the network to compensate.



Note In Cisco IOS XE Releases 2, 3.1S and 3.2S, PfR supports the ASR 1000 series router as a border router only and the master controller must be running a Cisco IOS Release 15.0M image for version compatibility. In Cisco IOS XE Release 3.3S, and later releases, master controller configuration is supported.

How to Configure Basic Performance Routing

Setting Up the PfR Master Controller


Perform this task to set up the PfR master controller to manage an PfR-managed network. This task must be performed on the router designated as the PfR master controller. For an example network configuration of a master router and two border routers, see the figure below. Communication is first established between the master controller and the border routers with key-chain authentication being configured to protect the

communication session between the master controller and the border routers. Internal and external border router interfaces are also specified.



Note In Cisco IOS XE Release 3.1S, and later releases, PfR supports the ASR 1000 series router as a border router only and the master controller must be running a Cisco IOS Release 15.0M image. In Cisco IOS XE Release 3.3S, and later releases, master controller configuration is supported.

Figure 3: Master Controller and Border Router Diagram



To disable a master controller and completely remove the process configuration from the running configuration, use the **no pfr master** command in global configuration mode.

To temporarily disable a master controller, use the **shutdown** command in PfR master controller configuration mode. Entering the **shutdown** command stops an active master controller process but does not remove any configuration parameters. The **shutdown** command is displayed in the running configuration file when enabled.

Before you begin

Interfaces must be defined and reachable by the master controller and the border routers before a PfR-managed network can be configured.

To set up a PfR-managed network, you must configure routing protocol peering or redistribution between border routers and peer routers in order for PfR to control routing.



Tip We recommend that the master controller be physically close to the border routers to minimize communication response time in PfR-managed networks. If traffic is to be routed between border routers, the border routers also should be physically close each other to minimize the number of hops.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **key chain** *name-of-chain*
4. **key** *key-id*
5. **key-string** *text*
6. **exit**
7. Repeat Step 3 through Step 7.
8. Repeat Step 3 through Step 7 with appropriate changes to configure key chain authentication for each border router.

9. **pfr master**
10. **logging**
11. **border** *ip-address* [**key-chain** *key-chain-name*]
12. **interface** *type number* **external**
13. **exit**
14. **interface** *type number* **internal**
15. **exit**
16. Repeat Step 11 through Step 15 with appropriate changes to establish communication with each border router.
17. **keepalive** *timer*
18. **end**
19. **show running-config**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: <pre>Router> enable</pre>	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: <pre>Router# configure terminal</pre>	Enters global configuration mode.
Step 3	key chain <i>name-of-chain</i> Example: <pre>Router(config)# key chain border1_PFR</pre>	Enables key-chain authentication and enters key-chain configuration mode. <ul style="list-style-type: none"> • Key-chain authentication protects the communication session between the master controller and the border router. The key ID and key string must match in order for communication to be established. • In this example, a key chain is created for use with border router 1.
Step 4	key <i>key-id</i> Example: <pre>Router(config-keychain)# key 1</pre>	Identifies an authentication key on a key chain. <ul style="list-style-type: none"> • The key ID must match the key ID configured on the border router.
Step 5	key-string <i>text</i> Example: <pre>Router(config-keychain-key)# key-string b1</pre>	Specifies the authentication string for the key and enters key-chain key configuration mode. <ul style="list-style-type: none"> • The authentication string must match the authentication string configured on the border router. • Any encryption level can be configured.

	Command or Action	Purpose
		<ul style="list-style-type: none"> In this example, a key string is created for use with border router 1.
Step 6	exit Example: <pre>Router(config-keychain-key)# exit</pre>	Exits key-chain key configuration mode and returns to key-chain configuration mode.
Step 7	Repeat Step 3 through Step 7.	Exits key-chain configuration mode and returns to global configuration mode.
Step 8	Repeat Step 3 through Step 7 with appropriate changes to configure key chain authentication for each border router.	--
Step 9	pfr master Example: <pre>Router(config)# pfr master</pre>	<p>Enters PfR master controller configuration mode to configure a router as a master controller.</p> <ul style="list-style-type: none"> A master controller and border router process can be enabled on the same router (for example, in a network that has a single router with two exit links to different service providers).
Step 10	logging Example: <pre>Router(config-pfr-mc)# logging</pre>	<p>Enables syslog messages for a master controller or border router process.</p> <ul style="list-style-type: none"> The notice level of syslog messages is enabled by default.
Step 11	border ip-address [key-chain key-chain-name] Example: <pre>Router(config-pfr-mc)# border 10.1.1.2 key-chain border1_PFR</pre>	<p>Enters PfR-managed border router configuration mode to establish communication with a border router.</p> <ul style="list-style-type: none"> An IP address is configured to identify the border router. At least one border router must be specified to create an PfR-managed network. A maximum of 20 border routers can be controlled by a single master controller. The value for the <i>key-chain-name</i> argument must match the key-chain name configured in Step 3. <p>Note The key-chain keyword and <i>key-chain-name</i> argument must be entered when a border router is initially configured. However, this keyword is optional when reconfiguring an existing border router.</p>
Step 12	interface type number external Example:	<p>Configures a border router interface as an PfR-managed external interface.</p> <ul style="list-style-type: none"> External interfaces are used to forward traffic and for active monitoring.

	Command or Action	Purpose
	<pre>Router(config-pfr-mc-br)# interface GigabitEthernet 0/0/0 external</pre>	<ul style="list-style-type: none"> A minimum of two external border router interfaces are required in an PfR-managed network. At least one external interface must be configured on each border router. A maximum of 400 external interfaces can be controlled by single master controller. <p>Tip Configuring an interface as an PfR-managed external interface on a router enters PfR border exit interface configuration mode. In this mode, you can configure maximum link utilization or cost-based optimization for the interface.</p> <p>Note Entering the interface command without the external or internal keyword places the router in global configuration mode and not PfR border exit configuration mode. The no form of this command should be applied carefully so that active interfaces are not removed from the router configuration.</p>
Step 13	<p>exit</p> <p>Example:</p> <pre>Router(config-pfr-mc-br-if)# exit</pre>	Exits PfR-managed border exit interface configuration mode and returns to PfR-managed border router configuration mode.
Step 14	<p>interface type number internal</p> <p>Example:</p> <pre>Router(config-pfr-mc-br)# interface GigabitEthernet 1/0/0 internal</pre>	<p>Configures a border router interface as an PfR controlled internal interface.</p> <ul style="list-style-type: none"> Internal interfaces are used for passive monitoring only. Internal interfaces do not forward traffic. At least one internal interface must be configured on each border router.
Step 15	<p>exit</p> <p>Example:</p> <pre>Router(config-pfr-mc-br)# exit</pre>	Exits PfR-managed border router configuration mode and returns to PfR master controller configuration mode.
Step 16	Repeat Step 11 through Step 15 with appropriate changes to establish communication with each border router.	--
Step 17	<p>keepalive timer</p> <p>Example:</p> <pre>Router(config-pfr-mc)# keepalive 10</pre>	<p>(Optional) Configures the length of time that an PfR master controller will maintain connectivity with an PfR border router after no keepalive packets have been received.</p> <ul style="list-style-type: none"> The example sets the keepalive timer to 10 seconds. The default keepalive timer is 60 seconds.

	Command or Action	Purpose
Step 18	end Example: <pre>Router(config-pfr-mc-learn)# end</pre>	Exits PFR Top Talker and Top Delay learning configuration mode and returns to privileged EXEC mode.
Step 19	show running-config Example: <pre>Router# show running-config</pre>	(Optional) Displays the running configuration to verify the configuration entered in this task.

Setting Up a PFR Border Router

Perform this task to set up a PFR border router. This task must be performed at each border router in your PFR-managed network. Communication is first established between the border router and the master controller with key-chain authentication being configured to protect the communication session between the border router and the master controller. A local interface is configured as the source for communication with the master controller, and external interfaces are configured as PFR-managed exit links.

To disable a border router and completely remove the process configuration from the running configuration, use the **no pfr border** command in global configuration mode.

To temporarily disable a border router process, use the **shutdown** command in PFR border router configuration mode. Entering the **shutdown** command stops an active border router process but does not remove any configuration parameters. The **shutdown** command is displayed in the running configuration file when enabled.

Before you begin

- Perform the task, Setting Up the PFR Master Controller, to set up the master controller and define the interfaces and establish communication with the border routers.
- Each border router must have at least one external interface that is either used to connect to an ISP or is used as an external WAN link. A minimum of two external interfaces are required in a PFR-managed network.
- Each border router must have at least one internal interface. Internal interfaces are used for only passive performance monitoring with NetFlow. Internal interfaces are not used to forward traffic.
- Each border router must have at least one local interface. Local interfaces are used only for master controller and border router communication. A single interface must be configured as a local interface on each border router.



Tip For Cisco IOS XE Release 3.1S and 3.2S, PFR supports the ASR 1000 series router as a border router only; the master controller cannot be enabled on an ASR 1000 series router. In Cisco IOS XE Release 3.3S, and later releases, master controller configuration is supported.



Tip We recommend that the border routers be physically close to one another to minimize the number of hops. The master controller also should be physically close to the border routers to minimize communication response time in PFR-managed networks.



Note

- Internet exchange points where a border router can communicate with several service providers over the same broadcast media are not supported.
- When two or more border routers are deployed in a PFR-managed network, the next hop to an external network on each border router, as installed in the RIB, cannot be an IP address from the same subnet.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **key chain** *name-of-chain*
4. **key** *key-id*
5. **key-string** *text*
6. **exit**
7. Repeat Step 6
8. **pfr border**
9. **local** *type number*
10. **master** *ip-address* **key-chain** *key-chain-name*
11. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Router> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Router# configure terminal	Enters global configuration mode.
Step 3	key chain <i>name-of-chain</i> Example: Router(config)# key chain border1_PFR	Enables key-chain authentication and enters key-chain configuration mode. <ul style="list-style-type: none"> • Key-chain authentication protects the communication session between both the master controller and the border router. The key ID and key string must match in order for communication to be established.

	Command or Action	Purpose
Step 4	key <i>key-id</i> Example: <pre>Router(config-keychain)# key 1</pre>	Identifies an authentication key on a key chain and enters key-chain key configuration mode. <ul style="list-style-type: none"> • The key ID must match the key ID configured on the master controller.
Step 5	key-string <i>text</i> Example: <pre>Router(config-keychain-key)# key-string bl</pre>	Specifies the authentication string for the key. <ul style="list-style-type: none"> • The authentication string must match the authentication string configured on the master controller. • Any level of encryption can be configured.
Step 6	exit Example: <pre>Router(config-keychain-key)# exit</pre>	Exits key-chain key configuration mode and returns to key-chain configuration mode.
Step 7	Repeat Step 6 Example: <pre>Router(config-keychain)# exit</pre>	Exits key-chain configuration mode and returns to global configuration mode.
Step 8	pfr border Example: <pre>Router(config)# pfr border</pre>	Enters PFR border router configuration mode to configure a router as a border router. <ul style="list-style-type: none"> • The border router must be in the forwarding path and contain at least one external and internal interface.
Step 9	local <i>type number</i> Example: <pre>Router(config-pfr-br)# local GigabitEthernet 0/0/0</pre>	Identifies a local interface on a PFR border router as the source for communication with an PFR master controller. <ul style="list-style-type: none"> • A local interface must be defined.
Step 10	master <i>ip-address</i> key-chain <i>key-chain-name</i> Example: <pre>Router(config-pfr-br)# master 10.1.1.1 key-chain border1_PFR</pre>	Enters PFR-managed border router configuration mode to establish communication with a master controller. <ul style="list-style-type: none"> • An IP address is used to identify the master controller. • The value for the key-chain-name argument must match the key-chain name configured in Step 3.
Step 11	end Example: <pre>Router(config-pfr-br)# end</pre>	Exits PFR Top Talker and Top Delay learning configuration mode and returns to privileged EXEC mode.

What to Do Next

If your network is configured to use only static routing, no additional configuration is required. The PfR-managed network should be operational, as long as valid static routes that point to external interfaces on the border routers are configured.

Otherwise, routing protocol peering or static redistribution must be configured between the border routers and other routers in the PfR-managed network.

Configuration Examples for Configuring Basic Performance Routing

Configuring the PfR Master Controller Example

The following configuration example, starting in global configuration mode, shows the minimum configuration required to configure a master controller process to manage the internal network. A key-chain configuration named PFR is defined in global configuration mode.



Note This configuration is performed on a master controller. Only border router functionality is included in Cisco IOS XE Release 3.1S and 3.2S; no master controller configuration is available. The master controller that communicates with the Cisco ASR 1000 series router being used as a border router 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.

```
Router(config)# key chain PFR
Router(config-keychain)# key 1
Router(config-keychain-key)# key-string KEYSTRING2
Router(config-keychain-key)# end
```

The master controller is configured to communicate with the 10.100.1.1 and 10.200.2.2 border routers. The keepalive interval is set to 10 seconds. Route control mode is enabled. Internal and external PfR-controlled border router interfaces are defined.

```
Router(config)# pfr master
Router(config-pfr-mc)# keepalive 10
Router(config-pfr-mc)# logging
Router(config-pfr-mc)# border 10.100.1.1 key-chain PFR
Router(config-pfr-mc-br)# interface GigabitEthernet 0/0/0 external
Router(config-pfr-mc-br)# interface GigabitEthernet 0/0/1 internal
Router(config-pfr-mc-br)# exit
Router(config-pfr-mc)# border 10.200.2.2 key-chain PFR
Router(config-pfr-mc-br)# interface GigabitEthernet 0/0/0 external
Router(config-pfr-mc-br)# interface GigabitEthernet 0/0/1 internal
Router(config-pfr-mc)# exit
```

Configuring a PfR Border Router Example

The following configuration example, starting in global configuration mode, shows the minimum required configuration to enable a border router. The key-chain configuration is defined in global configuration mode.

```
Router(config)# key chain PFR
Router(config-keychain)# key 1
Router(config-keychain-key)# key-string KEYSTRING2
Router(config-keychain-key)# end
```

The key-chain PFR is applied to protect communication. An interface is identified to the master controller as the local interface (source) for PfR communication.

```
Router(config)# pfr border
Router(config-pfr-br)# local GigabitEthernet 1/0/0
Router(config-pfr-br)# master 192.168.1.1 key-chain PFR
Router(config-pfr-br)# end
```

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
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
<p>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.</p>	<p>http://www.cisco.com/cisco/web/support/index.html</p>

Feature Information for Configuring Basic Performance Routing

Table 1: Feature Information for Configuring Basic Performance Routing

Feature Name	Releases	Feature Information
Optimized Edge Routing	Cisco IOS XE Release 2.6.1, Cisco IOS XE Release 3.1S	<p>OER was introduced on the Cisco ASR 1000 series routers. Performance Routing is an extension of OER.</p> <p>PfR syntax was introduced in Cisco IOS XE Release 3.1S.</p> <p>The following commands were introduced or modified: pfr, show pfr master.</p> <p>Note Only border router functionality is included in the Cisco IOS XE Release 2.6.1 and Cisco IOS XE Release 3.1S releases; no master controller configuration is available. 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.</p>
PfR Master Controller support for ASR 1000	Cisco IOS XE Release 3.3S	In Cisco IOS XE Release 3.3S and later releases, master controller functionality is supported.

