

Cisco Performance Routing

As enterprise organizations grow their businesses, the demand for real-time application performance and a better application experience for users increases. For example, voice and TelePresence applications are becoming integral parts of corporate networks, and performance of these applications is crucial. In order to improve application performance, companies have typically deployed two common solutions: provide additional bandwidth by deploying more network connections or offer 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. Cisco Performance Routing (PfR) addresses network performance problems by allowing the network to intelligently choose a path that meets the current requirements of application performance. In addition, this technology allows the network to choose resources appropriately to reduce operational costs incurred by enterprises.

Cisco PfR complements classic routing technologies by adding intelligence to select best paths to meet performance requirements of applications. The first phase of Cisco PfR intelligently optimizes performance of applications over WANs and intelligently load balances traffic to the Internet. Later phases of PfR will enhance application intelligence and extend this technology across the enterprise network.

Cisco PfR selects an egress or ingress WAN path based on parameters that affect application performance, including reachability, delay, cost, jitter, and Mean Opinion Score (MOS). The technology can also select an egress or ingress WAN path to intelligently load balance traffic based on usage or circuit cost to reduce costs incurred by enterprises. To achieve this balance, PfR selects a WAN path based on interface parameters such as reachability, load, throughput, and link cost of using a path. Classic routing (Enhanced IGRP [EIGRP], Open Shortest Path First [OSPF], Routing Information Protocol Version 2 [RIPv2], Border Gateway Protocol [BGP], etc..) protocols generally focus on providing reachability by creating a loop-free topology based on shortest or least-cost path. Cisco PfR focuses on providing application performance by understanding application requirements and current network performance characteristics.

Cisco PfR gains this additional intelligence from the following:

- Allows network administrators to provide business policy or application requirements to IP routing
- Monitors the network performance by taking advantage of embedded Cisco IOS® Software intelligence in switches and routers

The following types of businesses typically use Cisco PfR technology:

- Large, medium, and small enterprises with mission-critical Internet presence requirements
- Enterprises with multiple diverse WANs for availability requirements
- Enterprises with remote offices with a primary and backup WAN service
- Small and home offices with dual Internet connections

Cisco PfR: An Overview

Many businesses have multiple WAN and Internet connections to enhance the network and its applications. These businesses face complex implementation challenges when trying to split inbound and outbound traffic across multiple links while ensuring continuous network availability. This complexity increases further when these connections are from different service providers or have different characteristics such as bandwidth, quality, or cost. Cisco Performance Routing removes this complexity from network operators and provides needed application performance when the network is operational but parts of the network are experiencing performance-degradation problems.

Cisco PfR in its truest form is “application routing based on network performance”. It automatically detects path degradation and responds to avoid continued degradation. In many cases for a multihomed enterprise, the traffic on an initial path is routed through another egress path that can meet the application performance requirements. This routing is different from classic routing, because classic routing looks only at reachability and does not look into the required traffic service needs, such as low loss or low delay. In addition, Cisco PfR allows a multihomed enterprise to use all available WAN or Internet links. It can track throughput, link usage, and link cost, and automatically determine the best load balancing to optimize throughput, load, and cost—and network operators define the Cisco PfR policies that implement these adaptive routing techniques.

Cisco PfR policies can be based on the following parameters:

- WAN out-bound performance (traffic exiting from an enterprise): Delay, loss, reachability, throughput, jitter, and MOS
- WAN in-bound performance (traffic arriving into an enterprise): Delay, loss, reachability, and throughput
- WAN and Internet path parameters: Reachability, throughput, load, and link usage cost

Cisco Performance Routing consists of two distinct elements: border routers and a primary controller. The border routers connect enterprises to the WAN; the primary controller is a software entity supported by Cisco IOS Software on a router platform. Border routers gather traffic and path information and send this information to a primary controller, which places all received information into a database. The primary controller is configured with the requested service policies, so it is aware of everything that happens at the network edge and can automatically detect and take action when certain parameters are out-of-policy (OOP).

No licensing is required to enable Cisco Performance Routing.

Cisco IOS Software PfR Support

Note: The initial implementation of Cisco PfR uses Cisco IOS Optimized Edge Routing (OER) techniques. Cisco PfR functions are enabled with OER command-line interface (CLI) commands. Search for PfR features in Cisco Feature Navigator (<http://www.cisco.com/go/fn>) through the “Optimized Edge Routing” keyword. Additional PfR technologies will supersede current OER functions in future Cisco IOS Software releases.

Cisco Performance Routing takes advantage of the vast intelligence imbedded in Cisco IOS Software to determine the optimal path based upon network and application policies. It is an evolution of the Cisco IOS OER technology with a much broader scope. The application intelligence and end-to-end network strategy of Cisco PfR are significantly broader than the scope of OER. The initial phase of PfR uses OER technology extensively to meet emerging application demands on the enterprise network.

Available starting with Cisco IOS Software Release 12.3(8)T, the Cisco OER solution supports both the border router and the primary controller functions. The following functions have been added to OER since its inception:

- Cisco IOS Software Release 12.3(8)T: Initial OER support

- Cisco IOS Software Release 12.3(11)T: Added VPN IP Security (IPsec) and generic routing encapsulation (GRE) tunnel optimization
- Cisco IOS Software Release 12.3(11)T: Added port- and protocol-based prefix learning
- Cisco IOS Software Release 12.3(11)T: Added support for policy-rules configuration
- Cisco IOS Software Release 12.3(14)T: Added support for cost-based optimization and traceroute reporting
- Cisco IOS Software Release 12.4(2)T: Added application-aware routing: Policy-Based Routing (PBR)
- Cisco IOS Software Release 12.4(2)T: Added active probe source address
- Cisco IOS Software Release 12.4(6)T: Added voice traffic optimization
- Cisco IOS Software Release 12.4(9)T: Added differentiated services code point (DSCP) monitoring
- Cisco IOS Software Release 12.4(9)T: Added BGP inbound optimization
- Cisco IOS Software Release 12.2(33)SRB: Added support on Cisco7600
- Cisco IOS Software Release 12.2(33)SXH: Added border router support on Cisco Catalyst® 6500

Additional features can be found on the Cisco Feature Navigator at: <http://www.cisco.com/go/fn>.

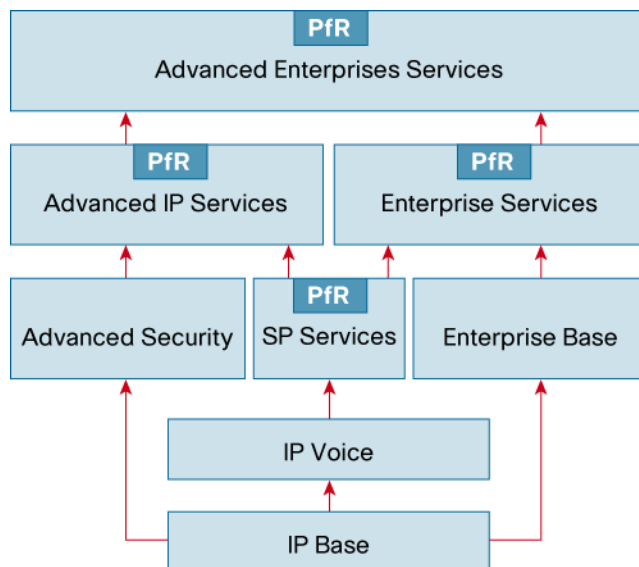
Requirements to run Cisco PfR follow:

- Border routers must run Cisco Express Forwarding
- The primary controller must be able to reach and communicate with border routers for operations.

Supported WAN interfaces include Ethernet, serial, tunnel, High-Speed Serial Interface (HSSI), ISDN Basic Rate Interface (BRI), ATM, Packet over SONET/SDH (PoS), VLAN, Dialer, Multilink, Frame Relay, and port channel.

The required Cisco IOS Software feature sets to run Cisco Performance Routing include SP Services, Advanced IP Services, Enterprise Services, and Advanced Enterprise Services (refer to Figure 1).

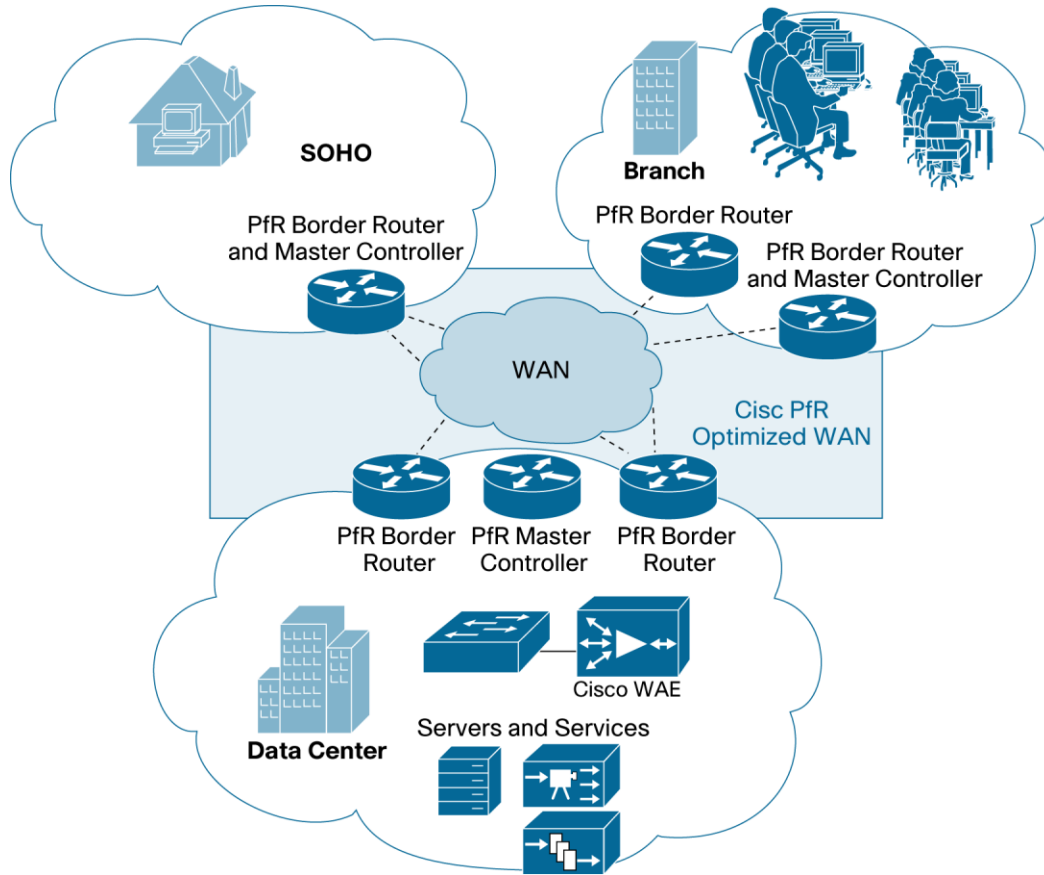
Figure 1. Cisco IOS Software Feature Sets Required to Run Cisco PfR



Cisco PfR Topologies

Cisco PfR can enhance performance of applications in many types of topologies. Figure 2 shows the topologies of most common enterprise implementations, going from a small edge router for the small office or home office (SOHO) to central-site data center deployments.

Figure 2. Enterprise Network Topologies



Typical business services enhanced by Cisco PfR technology include:

- Internet edge: Internet content services where there is more than one WAN access link to the data center can benefit from Cisco PfR. These main sites have two or more links to the Internet with border routers. Remote sites are completely unknown Internet entities (home users, business users, etc...).
- WAN or metropolitan-area network (MAN) hub: Centralized content services in headquarters where the remote sites are known and have static address assignments can benefit from Cisco PfR.
- WAN or MAN branch: Branch or SOHO sites trying to access Internet resources over the best link if more than one access link is available can benefit from Cisco PfR. Based on the site scale requirements, the primary controller and border router functions may be co-located on the same router.
- Cisco Wide Area Application Services (WAAS) and Cisco PfR: Cisco WAAS optimizes the application sessions, whereas Cisco PfR optimizes the network routing topology. Combining Cisco WAAS with Cisco PfR provides advanced application performance. (For more information about Cisco WAAS, visit <http://www.cisco.com/go/waas>.)
- Any to any: Distributed organizations with no real data center or hub site but whose communications flow between all locations can benefit from Cisco PfR. Cisco PfR can optimize path selection when the router is hosting IPSec VPN technologies such as GET-VPN and IPsec with GRE

Deployment Considerations

Cisco PfR technology provides significant flexibility for deployment in a network. The following considerations are important when deploying this technology:

- Cisco PfR technology is embedded in Cisco IOS Software: The solution allows border router and primary controller functions to coexist on a single Cisco router or exist on two different Cisco routers. It requires no additional appliance in the network.
- A minimum PfR design requirement is for one primary controller function, one border router function, and two WAN links.
- To benefit from PfR technology, the WAN paths must traverse different paths. The next-hop addresses on WAN links should be on different subnets. Hence, two next-hop upstream routers on a single WAN link are not supported. (Note: This restriction affects the use of PfR at Internet exchange locations where multiple next-hops are reachable from the same peering interface.)
- Each border router is configured to peer with a primary controller (even if both functions are on a single router).
- The primary controller is configured on a router along with the list of border routers that it controls (even if both functions are on a single router).
- The control session between a border router and a primary controller is secured with password authentication.
- The control session between the primary controller and the border router occurs between local interfaces on each router. You can configure the actual source interface, and hence the IP addresses used for this session, by identifying the local interface of the router.
- For PfR to manage and control appropriate WAN paths, WAN links on border routers need to be identified as either internal- or external-facing interfaces.
- You do not need to configure NetFlow, IP Service-Level Agreement (SLA), or other Cisco IOS Software features explicitly on either the primary controller or the border router.
- The border router must be in the traffic forwarding path, but there is no need for the primary controller to be in the traffic forwarding path.
- Currently, a primary controller can manage up to 10 border routers.

Each border router needs the following configuration:

- Cisco Express Forwarding must be enabled.
- The IPv4 address of the primary controller must be included.
- The PfR local interface must be able to communicate with the primary controller.
- The session between the border router and the primary controller must have an authentication password.

A primary controller needs the following configuration:

- IPv4 addresses of its border routers must be included.
- The session between the border router and the primary controller must have an authentication password
- Cisco PfR policies must be configured.

Cisco PfR Operation

The first step to configure Cisco PfR is to identify a primary controller and one or more border routers. Consequently, you must identify the external and internal interfaces per border router, and configure for the sessions between the primary controller and the border router with authentication and IP addresses.

The next step is to decide which **prefixes** or **traffic** classes you want to monitor. Monitoring can occur dynamically (the default) or by human configuration with the help of Cisco IOS Software access control lists (ACLs). The latter option is especially useful if detailed information about the application or prefixes that PfR needs to optimize is available.

Each border router gathers information about the interesting traffic classes or prefixes and sends it to the primary controller, which collects all the information received from its border routers and compares it with the policies you've configured.

If the primary controller determines that everything is within policy, then routing continues, but if the primary controller detects that defined traffic classes, prefixes, or links are experiencing OOP characteristics, it triggers an event, and PfR **may** take action to bring defined traffic classes back into policy. If PfR is configured in **observe** mode, then the OOP event just triggers a syslog message; if in **control** mode, the primary controller directs the border routers to reroute the OOP traffic to another WAN link, with the objective to bring the defined traffic classes back into policy.

The redirection is accomplished by influencing the network routing behavior for the traffic flows. The primary controller sends control messages to each of the border routers to let them know to take action and inject IP routes or PBR entries. Route and PBR entries are injected only on the border routers; no routes are injected on the primary controller.

For **outbound WAN optimization** (traffic going from the local network to the world outside), Cisco PfR uses PBR, static routes, or BGP routes to control the actual outbound path. These injected entries can be redistributed into an Interior Gateway Protocol (IGP; for example, EIGRP, OSPF, or RIP).

Static Route Injection

A Cisco PfR Primary Controller can enforce the use of a particular border router as the preferred exit by injecting temporary static routes. These static routes exist only in the memory of the router, and are intentionally not saved to the permanent configuration.

BGP Control Techniques

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

If the traffic associated with the interesting traffic is defined only by a prefix, the primary controller can instruct a border router to inject a BGP route into the BGP table to influence traffic to use a different link. All PfR injected routes remain local to an autonomous system, and these injected routes are never shared with external BGP peers. As a safeguard to ensure this behavior, when PfR injects a BGP route, it sets the **no-export community** on it. This process is automatic, and requires no user configuration.

Cisco PfR also uses BGP local preference to control interesting prefix entries. BGP local preference (Local_Pref) is a discretionary attribute applied to a BGP prefix to specify the degree of preference for that route during route selection. The Local_Pref is a value applied to a BGP prefix, and a higher Local_Pref value causes a route to be preferred over an equivalent route. The primary controller instructs one of the border routers to apply the Local_Pref attribute to a prefix or set of prefixes associated with a traffic-flow entry. The border router then shares the Local_Pref value with all of its internal BGP (iBGP) peers. Local_Pref is a locally significant value within an autonomous system, but it is never

shared with external BGP (eBGP) peers. When the iBGP convergence is complete, the router with the highest Local_Pref for the prefix becomes the exit link from the network.

Policy Route

Cisco PfR can control application traffic using PBR. You can identify application traffic traveling through a particular PfR border router by matching traffic defined in a PfR map as part of a PfR policy.

Cisco PfR offers the following methods to enforce an entrance link selection:

- BGP Autonomous System Number (ASN) prepend: After Cisco PfR selects a best entrance for an inside prefix, extra autonomous system hops (up to a maximum of six) are prepended to the inside prefix BGP advertisement over the other entrances. The extra autonomous system hops on the other entrances increase the probability that the best entrance will be used for the inside prefix. PfR uses this (default) method to control an inside prefix, and no user configuration is required.
- BGP autonomous system number community prepend: After PfR selects a best entrance for an inside prefix, it attaches a 'BGP prepend' community to the inside prefix BGP advertisement from the network to another Autonomous System (AS) such as an Internet service provider (ISP). The BGP prepend community increases the number of autonomous system hops in the advertisement of the inside prefix from the ISP to its peers.

Cisco PfR Traffic Classes and Prefixes

Cisco PfR is all about applying advanced routing control to identified interesting **traffic classes** or **prefixes**. A traffic class can be defined as a combination of prefix, protocol, port numbers, and DSCP values. PfR can be instructed to learn these interesting elements through either human configuration or dynamic learning. To dynamically learn the interesting traffic classes, PfR uses a mechanism based on the Cisco IOS NetFlow technology. You can configure the interesting PfR traffic classes with the **oer-map** and **ip prefix-list** commands. The **oer-map** command is similar to a normal **route-map** command.

Cisco PfR Monitoring

Cisco PfR uses three methods of traffic class performance measurement:

- Passive monitoring: Measuring the performance metrics of interesting prefixes while the traffic is flowing through the device using NetFlow technology
- Active monitoring: Creating a stream of synthetic traffic replicating the interesting traffic classes as closely as possible to measure the performance metrics of the synthetic traffic; uses Cisco integrated IP SLAs technology
- Both active and passive monitoring: Combining both active and passive monitoring to generate a more complete picture of traffic flows within the network

No explicit NetFlow or IP SLAs configuration is required; support for NetFlow and IP SLAs is enabled and configured automatically by the primary controller. You can use both active and passive monitoring methods for each traffic class. The primary controller uses both passive and active monitoring by default. All traffic classes are passively monitored using integrated NetFlow functions, whereas the OOP traffic classes are actively monitored using IP SLA functions.

The primary controller also receives information from each border router about WAN link usage metrics.

Passive monitoring metrics include the following:

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

Active monitoring metrics include the following:

- Delay: Cisco PfR measures the average delay of TCP, User Datagram Protocol (UDP), and Internet Control Message Protocol (ICMP) flows for a given traffic class or prefix.
- Reachability: Cisco PfR measures reachability by tracking TCP synchronization messages that have been sent repeatedly without a received TCP acknowledgement.
- Jitter: Cisco PfR measures jitter, or interpacket delay variance, by sending multiple packets to a target address and a specified target port number and measuring the delay interval between packets arriving at the destination.
- MOS: MOS is a standards-based method of measuring voice quality. Standards bodies such as the ITU have derived two important recommendations: P.800 (MOS) and P.861 (Perceptual Speech Quality Measurement [PSQM]). P.800 defines a method to derive a MOS of voice quality. MOS scores range from 1, representing the worst voice quality, to 5, representing the best voice quality. A MOS of 4.0 is considered "toll-quality" voice. Cisco PfR measures MOS by using IP SLA infrastructure. The Cisco IP SLA Responder feature must be enabled at the PfR target address.

Link usage metrics include the following:

- Link usage threshold: If you configure an external interface for a border router, Cisco PfR automatically monitors the use of that external link (an external link is an interface on a border router that typically links to a WAN).
- Link usage range: You can configure Cisco PfR to calculate the range of usage over all the links.
- Cisco PfR range function: Cisco PfR attempts to keep the exit or entrance links within a usage range, relative to each other to ensure that the traffic load is distributed. The range is specified as a percentage and is configured on the primary controller to apply to all exit or entrance links on border routers managed by the primary controller.

Cisco PfR Timers

A variety of timers can affect the behavior of Cisco PfR. First, PfR has to learn prefixes. When learned, the router periodically checks to see if the actual traffic statistics match the configured policy expectations. The frequency and the speed for this checking rely upon another set of timers. If it finds an OOP traffic class, it may reroute that traffic to another path, and then place it in a hold-down state during which the rerouted flow remains unchanged for a configurable amount of time.

Three timers influence prefix learning:

- Periodic interval: Timer to let the primary controller know how often it needs to learn new prefixes from the border routers
- Monitor period: Timer that tells how long top-talker and top-delay measurements are observed
- Expire after: Timer to let the primary controller time out previously learned top-talker or top-delay prefixes

The primary controller uses the following timers to verify the configured policy with the actual traffic:

- Monitoring statistics: Cisco PfR statistics are based upon two time periods (5- and 60-minute periods).
- Hold down: This timer sets the inertia or dampening of the PfR route manipulation; the higher this timer is configured, the slower PfR can react on prefixes that were controlled and rerouted to a better path (default is 5 minutes).
- Back-off: Each time the back-off timer expires and PfR fails to discover an in-policy exit, the back-off interval is increased by a specified step (if configured) or a minimum number of seconds that increases up to a maximum number of seconds.
- Periodic: This command tells how fast the primary controller checks to see if the actual policy is compliant with the configured policy (default is 300 seconds).

Cisco PfR Policy

The Cisco PfR policy, which you configure on the Cisco PfR Primary Controller, defines if the interesting traffic classes should be learned, how fast they are learned and updated, how frequently the actual received service level is checked with the configured expected service level, etc.

A global PfR policy defines the PfR default behavior for timers and measurement thresholds. You can overwrite this policy per traffic class by means of a more specific definition. You configure WAN link policies when identifying the border routers; these policies are declared under each PfR-controlled WAN path.

You can have multiple policy criteria for a single interesting Cisco PfR prefix; you also can have multiple overlapping policies. To resolve the potential conflict of which policy to run, PfR uses its resolve function: a flexible mechanism that allows you to set the priority for a PfR policy. You assign each policy a unique value, and the policy with the lowest value is selected as the highest-priority policy. By default, PfR assigns the highest priority to delay policies, followed by usage policies. Assigning a priority value to any policy overrides the default settings.

Conclusion

Cisco Performance Routing enhances application performance in a network in order to meet business requirements by intelligently choosing a best path that reduces the cost, distributes load among all possible WAN paths, and improves overall network performance. Cisco Performance Routing is a Cisco IOS Software technology that takes advantage of embedded Cisco IOS Software intelligence to measure path performance information, and the solution requires no additional equipment.

Easily deployed with minimal configuration, Cisco PfR offers significant flexibility to tune and optimize the solution for your network to support current and future applications.

References

- Cisco Performance Routing on Cisco.com: <http://www.cisco.com/go/pfr>
- Cisco Performance Routing configuration guide: http://www.cisco.com/en/US/docs/ios/pfr/configuration/guide/15_1/pfr_15_1_book.html
- Cisco Performance Routing command reference: http://www.cisco.com/en/US/docs/ios/pfr/command/reference/pfr_book.html



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco Logo are trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and other countries. A listing of Cisco's trademarks can be found at www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1005R)

Printed in USA

C78-424310-02 07/10