PfR Simplification Phase 1

Performance Routing (PfR) is an advanced Cisco technology to allow businesses to complement traditional IP routing technologies such as Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), Routing Information Protocol Version 2 (RIPv2), and Border Gateway Protocol (BGP) with additional serviceability parameters to select the best egress or ingress path. It complements these traditional IP routing technologies with additional intelligence. PfR can select an egress or ingress WAN interface based upon parameters like reachability, delay, cost, jitter, Mean Opinion Score (MOS) score, or it can use interface parameters like load, throughput, and monetary cost. Traditional IP routing technologies generally focus on creating a loop-free topology based upon the shortest or least cost path.

Although PfR automatically enables IP SLA or NetFlow technologies, the initial configuration of PfR is more complicated than for traditional IP routing technologies due to PfR policy definition and the setting of many performance parameters. Cisco used feedback from customers to reduce the complexity of PfR configuration and align default values to match customer requirements. Phase 1 of the PfR simplification project introduces dynamic tunnels between PfR border routers, revised default values, removal of some CLI, and changes to default behavior. The changes result in fewer configuration steps before PfR is implemented in your network.

Information About PfR Simplification Phase 1

CLI and Default Value Changes to Simplify PfR

With CSCtr26978 a series of CLI and default value changes designed to make configuration of PfR simpler were introduced. Some commands and keywords were removed, and defaults changed to reflect customer environments.

Enforce Route Control by Default

In response to customer feedback, with CSCtr26978 the **mode route control** command is now the default behavior instead of the **mode route observe** command. In control mode, the master controller coordinates information from the border routers and makes policy decisions. The master controller monitors prefixes and exits based on default and user-defined policies, and implements changes to optimize prefixes and to select the best exit.

If you want to passively monitor and report without making any changes, you can still configure PfR to use the observe mode. In observe mode, the master controller monitors prefixes and exit links based on default and user-defined policies and then reports the status of the network and the decisions that should be made, but it does not implement any changes.

Default Change for Mode Verify Bidirectional CLI

In response to customer feedback, with CSCtr26978 the default behavior changed to disable the verification of bidirectional traffic. If you need to verify bidirectional traffic, configure the **mode verify bidirectional** command in master controller configuration mode.

CLI Default Value Changes to Simplify PfR

Command	Default Before CSCtr26978	Default After CSCtr26978
backoff	300, 3000, 300 seconds	90, 900, 90 seconds
holddown	300 seconds	90 seconds
max-xmit-utilization	75 percent	90 percent
monitor-period	5 minutes	1 minute
periodic-interval	120 minutes	0 minutes

Removal of PfR API and Proxy CLI

All CLI commands and functionality involved with the PfR application programming interface (API) and proxy process were removed to simplify PfR. With CSCtr26978, the following CLI commands were removed:

- **api provider (PfR)**
- **debug pfr api**
- **host-address (PfR)**
- **show api provider (PfR)**
- **show pfr proxy**

Removal of OER CLI

Although the Optimized Edge Routing (OER) syntax was replaced in most images with the PfR syntax, the OER syntax is still recognized. When you enter OER syntax the software changes the syntax to the new PfR syntax in the running configuration. With CSCtr26978, the OER syntax was removed.

Removal of Mode Select-Exit CLI

For most customer deployments we do not recommend using the passive monitoring mode with the exit selection of select-exit best because the statistics may change by the time all the links have been examined and the decision may not be accurate. To simplify the PfR configuration, with CSCtr26978 the default behavior is now select-exit good where the first in-policy link is selected. The **mode select-exit** command and **best** and **good** keywords have been removed.

Load Balancing With Link Groups and Resolver Changes

With CSCtr33991 changes were introduced to the PfR link group and resolver behaviors to simplify the configuration and understanding of PfR. The limitation of configuring a range resolver and link grouping at the same time was removed. Without any awareness of link group configuration, load balancing was performed across all the links. Link groups provide the ability to define a group of exit links as a preferred set of links, or a fallback set of links for PfR to use when optimizing traffic classes specified in a PfR policy.

To further simplify PfR, CSCtr33991 changed the behavior where range resolvers are considered after performance resolvers (such as delay, throughput, or loss).

Note The cost resolver cannot be configured with a performance resolver.

Delay, Range, and Utilization Resolver Changes

Before CSCtr3399	After CSCtr3399
Support of utilization and range resolvers.	With CSCtr33991, the **range** and **utilization** keywords in the **resolve** and **set resolve** commands were removed to disable support for the utilization and range resolvers.
Delay, range, and utilization resolvers are the default resolvers in the default global policy.	PfR automatically performs load balancing; default resolvers were removed from the default global policy.

Performance Resolver and Link Group Load Balancing

Before PfR performs load balancing traffic across available exits, rules to consider configured performance resolvers (such as delay or loss) and any configured link group were introduced with CSCtr33991. The rules are evaluated in the following order:

1. If no performance resolver is configured and no link group is used, PfR automatically performs load balancing across all links.

2. If no performance resolver is configured but link group is used, PfR automatically performs load balancing within the primary link group.

3. If performance resolvers are configured but no link group is used, PfR automatically performs load balancing across qualified links after those performance resolvers.

4. If performance resolvers are configured and a link group is used, PfR automatically performs load balancing across qualified links within the primary link group.

Load Balancing Within a Link Group

With CSCtr33991, the behavior of triggering range out-of-policy (OOP) for an exit by comparing the load of an exit to all other exits, is changed to comparing the load of an exit with all the exits in the same link group.

The maximum utilization (soft limit) of all the PfR-managed exit links is checked before PfR calls a resolver and, if none of the exits is below the soft limit, the exit selection is performed by ignoring the soft limit.

The existing behavior of moving any traffic class to balance the traffic load has been replaced by the ability to move any traffic class in the link group (whether primary or fallback) to balance the traffic load.

When any performance resolver is configured, the following rules apply in the specified order:

1. If only one qualified link is in the primary group, move traffic classes to this link.

2. If more than one qualified link is in the primary group, move traffic classes and perform load balancing across these links.

3. If more than one qualified link is in the fallback group, move traffic classes to one of the fallback group links.

4. If no qualified link is in both the primary and fallback groups, do not move the traffic class.

5. If no links are under the maximum utilization (soft limit) in the primary or fallback link groups, ignore the soft limit and move traffic classes to the best link.

When no performance resolver is configured, the following rules apply in the specified order:

1. If one or more qualified links are under the maximum utilization in the primary group, perform load balancing across these links in the primary group.

2. If more than one qualified link is in the fallback group, move traffic classes to one of the fallback group links.

3. If no links are under the maximum utilization (soft limit) in the primary or fallback link groups, perform load balancing across the primary group links.

Automatic Enable of Throughput Learning

To simplify PfR configuration, CSCtr2697 enabled PfR learn mode using throughput-based learning by default.

After feedback from customers, the default periodic interval of 120 minutes was changed to 90 minutes and the default monitor period was changed from 5 minutes to 1 minute.

The automatic enabling of PfR learn mode can be switched off using the **no learn** command if manual configuration is preferred.

Automatic PBR Route Control When No Parent Route Exists

When a PfR master controller (MC) decides to control a prefix using a protocol BGP, for example, it sends the control request to a selected PfR border router (BR). If the MC receives the successful control notification from the BR, it will notify all the other BRs to exclude the prefix. Some BRs may not have a parent route to this prefix via the same protocol. When no parent route exists for the prefix, this is detected as a RIB mismatch, the prefix is moved into a default state, and the control procedure begins again.

To simplify PfR, CSCtr26978 introduced new behavior when no parent route is detected. In this situation, PfR automatically switches to using dynamic policy-based routing (PBR) instead of trying all the other routing protocols in the following order; BGP, EIGRP, static, and PBR. With CSCtr26978, the existing **mode route protocol pbr** command behavior was enabled by default. Configuration of the **no mode route protocol pbr** command initially sets the traffic classes to be uncontrolled and PfR then uses a single protocol to control the traffic class in the following order: BGP, EIGRP, static, and PBR.

Dynamic PBR Support for PfR

The PfR BR Automatic Adjacencies feature introduces dynamic PBR support. In dynamic route maps, the PBR requirement for both interface and next-hop information is now supplied by PfR in a single set clause. To display the route map or policy information use the **show route-map dynamic** command or the **show ip policy** command.

How to Configure PfR Simplification Phase 1

Enabling PfR Route Observe Mode

With CSCtr26978, the **mode route control** command behavior is the default. Perform this task at the master controller to configure PfR to use route observe mode instead of the default route control mode. In route observe mode, the master controller monitors prefixes and exit links based on default and user-defined policies and then reports the status of the network and the decisions that should be made, but it does not implement any changes. In route control mode, the master controller coordinates information from the borders routers in the same way as route observe mode, but commands are sent back to the border routers to alter routing in the PfR managed network to implement the policy decisions.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **pfr master**
4. **mode route observe**
5. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**pfr master** **Example:** `Router(config)# pfr master`	Enters PfR master controller configuration mode to configure a router as a master controller and to configure global operations and policies.

	Command or Action	Purpose
Step 4	**mode route observe** **Example:** `Router(config-pfr-mc)# mode route observe`	Configures PfR to passively monitor and report without making any changes.
Step 5	**end** **Example:** `Router(config-pfr-mc)# end`	Exits PfR master controller configuration mode and returns to privileged EXEC mode.

Disabling Automatic PBR Route Control

Perform this task at the master controller to disable the default route control behavior when a RIB mismatch is found and allow PfR to use a single protocol to control a traffic class.

Note With CSCtr26978, the **no mode route protocol pbr** command behavior is enabled by default. Perform this task to override the default behavior.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **pfr master**
4. **no mode route protocol pbr**
5. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	**enable** **Example:** `Router> enable`	Enables privileged EXEC mode. • Enter your password if prompted.
Step 2	**configure terminal** **Example:** `Router# configure terminal`	Enters global configuration mode.
Step 3	**pfr master** **Example:** `Router(config)# pfr master`	Enters PfR master controller configuration mode to configure a router as a master controller and to configure global operations and policies.

	Command or Action	Purpose
Step 4	**no mode route protocol pbr** **Example:** `Router(config-pfr-mc)# no mode route protocol pbr`	Disables the automatic PBR route control. • Sets the traffic classes to be uncontrolled and PfR then uses a single protocol to control the traffic class in the following order; BGP, EIGRP, static, and PBR.
Step 5	**end** **Example:** `Router(config-pfr-mc)# end`	Exits PfR master controller configuration mode and returns to privileged EXEC mode.

Configuration Examples for PfR Simplification Phase 1

Example: Verifying PfR Simplification Default Changes

The following example outputs, from privileged EXEC mode, display the new default values and behavior introduced to simplify PfR.

The following partial output shows the new default behavior introduced with CSCtr26978; learn mode is enabled, the monitor period is set to 1 minute, and the periodic interval is set to 0 minutes:

```
.
.
.
Learn Settings:
  current state : ENABLED
  time remaining in current state : 0 seconds
  throughput
  no delay
  no inside bgp
  monitor-period 1
  periodic-interval 0
  aggregation-type prefix-length 24
  prefixes 100 appls 100
  expire after time 720
```

Feature Information for PfR Simplification Phase 1

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

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

Table 1: Feature Information for PfR Simplification Phase 1

Feature Name	Releases	Feature Information
PfR BR Automatic Adjacencies	15.2(2)S 15.2(3)T Cisco IOS XE Release 3.6S	The PfR BR Automatic Adjacencies feature introduces dynamic PBR support. In dynamic route maps, the PBR requirement for both interface and next-hop information is supplied by PfR in a single set clause. No commands were introduced or modified.