



PfRv3 Remote Prefix Tracking

Performance Routing Version 3 (PfRv3) is an intelligent-path control mechanism for improving application delivery and WAN efficiency. The PfRv3 Remote Prefix Tracking feature enhances networks running Performance Routing Version 3 (PfRv3) to learn the prefix of a remote device from the Routing Information Base (RIB) table.

- [Feature Information for PfRv3 Remote Prefix Tracking, on page 1](#)
- [Information About PfRv3 Remote Prefix Tracking, on page 1](#)
- [How to Display Site Prefixes, on page 5](#)
- [Additional References for PfRv3 Remote Prefix Tracking, on page 11](#)

Feature Information for PfRv3 Remote Prefix Tracking

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 PfRv3 Remote Prefix Tracking

Feature Name	Releases	Feature Information
PfRv3 Remote Prefix Tracking	Cisco IOS release 3.16.6, 15.6M2, 15.5.3M6, 15.7M, 16.3.5, and Cisco IOS XE Everest 16.6.1.	The following command was modified: show domain default vrf .

Information About PfRv3 Remote Prefix Tracking

Site Prefixes Database

Site Prefixes are LAN side prefixes owned by each site. The site prefix database is central to the site concept in PfRv3. Site prefix database reside on the master controller.

- The master site learns the remote site prefix through SAF advertised by remote MC. Master site learns the local site prefix from the local borders. The border learns the prefix from RIB and sends the prefix learned to the local master
- The border site prefix database is populated by SAF messages published by all the remote site master and local site master.
- By default, MCs and BRs delete site prefixes every 24 hours.


Learning Local Site Prefixes

Border routers collect the prefix from the RIB table and send it to the local master controller. After receiving prefixes from a border router, the local master controller filters prefixes as per the following criteria.

1. If a prefix is learned on a tunnel interface, the prefix is marked remote and not added to local LAN list.
2. If a prefix is learned from NHRP, the prefix is not added to LAN list.
3. If a prefix is learned on a physical interface of the tunnel interface, the prefix is not added to LAN list.
4. If an enterprise prefix is configured on the hub and the prefix is part of the enterprise prefix list configured on hub, the branch master adds the prefix from the RIB table to the LAN list.

The prefixes in the LAN list are added to the site prefix database as local site prefix list.

Figure 1: Learning Local Site Prefixes




Learning Remote Site Prefixes

In order to learn from advertisements via the peering infrastructure from remote peers, every MC and BR subscribes to the peering service for the subservice of site prefix. MCs publish and receive site prefixes. BRs only receive site prefixes. MC learns prefixes from the border and filters the prefixes as explained in the previous section and publishes the prefix to all sites. This message is received by all MCs and BRs that

subscribe to the peering service. The message is decoded and added to the site prefix databases at those MCs and BRs.

Figure 2: Pfrv3-discovery-site-prefix.png



PfRv3 Remote Prefix Tracking via Egress Flow

Prior to Cisco IOS XE Everest 16.6.1, the site prefix was learnt via the egress flow on the WAN interface. The prefix thus, learnt is published to all remote sites in the network using the EIGRP SAF message. If a remote site does not receive a new SAF message within 24 hours, the prefix is removed from the local-prefix database. If the routing is updated within 24 hours, corresponding prefix table will not be updated. Since, the prefix is learned from the egress traffic, sometimes-wrong prefixes are learnt due to redirected traffic. These wrongly learnt prefixes are not cleaned up until the 24 hour age out time.

Additionally, the prefix reachability is not tracked per channel. For example, if the prefix belongs to a specific site, it is assumed that prefix is reachable through all the channels available for that site. This results in a traffic blackhole when the prefix is not reachable through the selected channel.

PfRv3 Remote Prefix Tracking via RIB table


The PfRv3 Remote Prefix Tracking feature prevents the above scenarios by learning the local site prefixes from the RIB table instead of the egress flow. The prefixes are advertised to the remote sites. Changes to RIB table are tracked and are accordingly notified to all remote sites. Therefore, all sites are updated automatically with the precise site prefix information. Remote site tracks the prefix learnt via the WAN interface. While controlling the traffic, remote sites validate the reachability of the prefix on all channels available for a site.

There is no specific configuration required for this feature. You only need to configure the WAN interfaces.

How Site Prefix is Learnt?

The following workflow illustrates the process of how site prefix is learnt.

Figure 3: Site Prefix Learning Workflow



WAN Interfaces Configuration

You must configure the WAN interfaces on a border router in a branch using the **domain domain-name dynamic-path** command. For more information, see “Configuring Branch Border Router” in the *Performance Routing Version 3* chapter.

Prefix Learning on Border Router

On initialization, the border device learns the entire prefix from the RIB table and stores in the local prefix database, where the information is classified per VRF. Any changes in the RIB database, such as addition or deletion of prefixes, are accounted in the prefix database as appropriate. Prefixes learned from the RIB on the local border are forwarded to the local master controller. The prefix information in the border device can be viewed using the **show domain default vrf vrf name border route-import** command.

Forwarding the Prefix to Master Controller

Master controller learns about a new prefix added or removed in the RIB table from the border device.

On a branch site, when the WAN interfaces are configured using the **domain domain-name dynamic-path** command., the wan interface details are shared with the master controller by all border routers in a site. The master controller classifies this prefix information as WAN or LAN prefix, as appropriate.

On a hub site, The prefixes are learnt and classified similar to a branch site. The only difference is the command used to configure the WAN interface, which is **domain path service-provider-name path-id number** command.



Note It is mandatory to configure prefixes on the hub and the transit hub. It is also mandatory to configure the **domain domain-name dynamic-path** in branch tunnel interface.

Prefix Classification by Master Controller

Master controller filters the prefix using the criteria described in the *Learning Local Site Prefixes* section and updates the local prefix database. The local prefix database is published to all the subscribers using the EIGRP SAF message. The prefix information in the border device can be viewed using the following commands:

- **show domain {domain-name | default} vrf vrf-name master route-import local all**
- **show domain {domain-name | default} vrf vrf-name master route-import border border-ip**

- **show domain** {*domain-name* | **default**} **vrf** *vrf-name* **master route-import local**
- **show domain** {*domain-name* | **default**} **vrf** *vrf-name* **master route-import remote**
- **show domain** {*domain-name* | *default*} **vrf** *vrf-name* **border route-import**
- **show domain** {*domain-name* | *default*} **vrf** *vrf-name* **border local-prefix interface** *interface-name*

Path Preference

When a master controller receives prefixes from a border router, the master controller evaluates the traffic classes to a device, whose prefixes are listed in the RIB table and performs a policy decision to select a channel.

A channel is added to a channel list of a traffic class when a device associated with a prefix is reachable. The master controller decides on a path to a device based on the reachability of device (with a prefix in the RIB) on a channel. Prefixes are validated as follows:

- The list of interfaces on which prefixes are reachable is obtained from the prefix database and the prefix is verified for reachability via the same interface as the channel interface.
- A list of routes is obtained for a prefix that is reachable via an interface.

The channel is verified for the next hop address and if the next hop matches the appropriate prefix route. If the parent route of a device pertaining to a prefix matches the channel next hop, it indicates that the device with the prefix is reachable through a channel. If prefixes cannot be reached on a channel, a syslog message is displayed.



Note Maximum secondary paths must be configured on the border devices using the maximum-paths command so that prefixes are reachable. This command are enabled in the EIGRP or BGP router configuration mode.

How to Display Site Prefixes

Displaying Site Prefixes Learnt By a Border Router

SUMMARY STEPS

1. **show domain** *domain-name* **vrf** *vrf-name* **border site-prefix**
2. **show domain default vrf** *vrf name* **border route-import**
3. **show domain default vrf** *vrf name* **border route-import interface**
4. **show monitor event-trace pfrv3 all**

DETAILED STEPS

Step 1 **show domain** *domain-name* **vrf** *vrf-name* **border site-prefix**

Use this command to verify the reachability of the prefix on all channels.

Step 2 **show domain default vrf *vrf name* border route-import**

Use this command to view the prefix information learnt by a border device from the RIB table.

Example:

```
B1MCBR# show domain default vrf green border route-import
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
```

Proto	Prefix	Location	Next-Hop	Index	Interface	In-RIB
L	10.20.0.1/32	Local	0.0.0.0	29	Ethernet0/2.30	YES
C	10.20.0.0/24	Local	0.0.0.0	29	Ethernet0/2.30	YES
L	10.20.1.1/32	Local	0.0.0.0	25	Ethernet0/1.30	YES
C	10.20.1.0/24	Local	0.0.0.0	25	Ethernet0/1.30	YES
D	10.20.2.0/24	Local	10.20.0.2	29	Ethernet0/2.30	YES
L	51.1.0.4/32	Local	0.0.0.0	24	Tunnel10	YES
C	51.1.0.0/16	Local	0.0.0.0	24	Tunnel10	YES
D	52.1.0.0/16	Local	10.20.0.2	29	Ethernet0/2.30	YES
C	100.20.1.1/32	Local	0.0.0.0	22	Loopback1	YES
D	100.20.2.1/32	Local	10.20.0.2	29	Ethernet0/2.30	YES
S	100.20.3.1/32	Local	10.20.0.3	29	Ethernet0/2.30	YES

Step 3 **show domain default vrf *vrf name* border route-import interface**

Use this command to view the prefix information associated with an interface.

Example:

```
B1MCBR# show domain default vrf green border route-import interface Loopback1
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override, p - overrides from PfR
```

Proto	Prefix	Location	Next-Hop	Index	Interface	In-RIB
C	100.20.1.1/32	Local	0.0.0.0	22	Loopback1	YES

Step 4 **show monitor event-trace pfrv3 all**

Enables debugging by collecting trace.

Displaying Site Prefixes Learnt By a Master Controller

SUMMARY STEPS

1. `show domain default vrf vrf name master route-import`
2. `show domain default vrf vrf name master route-import interface`
3. `show domain default vrf vrf name master local-prefix`

DETAILED STEPS

Step 1 `show domain default vrf vrf name master route-import`

Use this command to view the prefix information learnt by a master controller.

Example:

```
B1MCBR# show domain default vrf green master route-import all
```

```
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
a - application route
+ - replicated route, % - next hop override, p - overrides from PfR
```

Enterprise Prefix List:

Prefix: 100.20.0.0, Mask: 16

Prefix: 100.30.0.0, Mask: 16

Prefix: 100.0.0.0, Mask: 8

Proto	Prefix	Location	BR-IP	Next-Hop	Index	Interface
	IF-Role	In-RIB				
B	10.10.1.0/24	Remote	100.20.1.1	10.20.1.2	25	Ethernet0/1.30
	WAN	YES				
B	10.10.2.0/24	Local	100.20.2.1	10.20.2.2	25	Ethernet0/1.32
	LAN	YES				
B	10.10.3.0/24	Remote	100.20.1.1	10.20.1.2	25	Ethernet0/1.30
	WAN	YES				
B	10.10.4.0/24	Local	100.20.2.1	10.20.2.2	25	Ethernet0/1.32
	LAN	YES				
B	10.15.1.0/24	Remote	100.20.1.1	10.20.1.2	25	Ethernet0/1.30
	WAN	YES				
B	10.15.2.0/24	Local	100.20.2.1	10.20.2.2	25	Ethernet0/1.32
	LAN	YES				
L	10.20.0.1/32	Local	100.20.1.1	0.0.0.0	29	Ethernet0/2.30
	LAN	YES				
L	10.20.0.2/32	Local	100.20.2.1	0.0.0.0	28	Ethernet0/2.30
	LAN	YES				
C	10.20.0.0/24	Local	100.20.1.1	0.0.0.0	29	Ethernet0/2.30
	LAN	YES				
C	10.20.0.0/24	Local	100.20.2.1	0.0.0.0	28	Ethernet0/2.30
	LAN	YES				
L	10.20.1.1/32	Remote	100.20.1.1	0.0.0.0	25	Ethernet0/1.30
	WAN	YES				
C	10.20.1.0/24	Remote	100.20.1.1	0.0.0.0	25	Ethernet0/1.30
	WAN	YES				

Displaying Site Prefixes Learnt By a Master Controller

D	10.20.1.0/24	Remote	100.20.2.1	10.20.0.1	28	Ethernet0/2.30
	LAN	YES				
L	10.20.2.1/32	Local	100.20.2.1	0.0.0.0	25	Ethernet0/1.32
	LAN	YES				
D	10.20.2.0/24	Local	100.20.1.1	10.20.0.2	29	Ethernet0/2.30
	LAN	YES				
C	10.20.2.0/24	Local	100.20.2.1	0.0.0.0	25	Ethernet0/1.32
	LAN	YES				
B	10.30.1.0/24	Remote	100.20.1.1	10.20.1.2	25	Ethernet0/1.30
	WAN	YES				
B	10.30.2.0/24	Local	100.20.2.1	10.20.2.2	25	Ethernet0/1.32
	LAN	YES				
L	51.1.0.4/32	Remote	100.20.1.1	0.0.0.0	24	Tunnel10
	WAN	YES				
C	51.1.0.0/16	Remote	100.20.1.1	0.0.0.0	24	Tunnel10
	WAN	YES				

B1MCCR# show domain default vrf green master route-import local

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
 E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
 ia - IS-IS inter area, * - candidate default, U - per-user static route
 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
 a - application route
 + - replicated route, % - next hop override, p - overrides from PfR

Enterprise Prefix List:

Prefix: 100.20.0.0, Mask: 16

Prefix: 100.30.0.0, Mask: 16

Prefix: 100.0.0.0, Mask: 8

Proto	Prefix	IF-Role	Location	BR-IP	Next-Hop	Index	Interface
			In-RIB				
L	10.20.0.1/32		Local	100.20.1.1	0.0.0.0	29	Ethernet0/2.30
	LAN	YES					
C	10.20.0.0/24		Local	100.20.1.1	0.0.0.0	29	Ethernet0/2.30
	LAN	YES					
C	10.20.0.0/24		Local	100.20.2.1	0.0.0.0	28	Ethernet0/2.30
	LAN	YES					
D	10.20.2.0/24		Local	100.20.1.1	10.20.0.2	29	Ethernet0/2.30
	LAN	YES					
C	10.20.2.0/24		Local	100.20.2.1	0.0.0.0	25	Ethernet0/1.32
	LAN	YES					
C	100.20.1.1/32		Local	100.20.1.1	0.0.0.0	22	Loopback1
	LAN	YES					
D	100.20.1.1/32		Local	100.20.2.1	10.20.0.1	28	Ethernet0/2.30
	LAN	YES					
D	100.20.2.1/32		Local	100.20.1.1	10.20.0.2	29	Ethernet0/2.30
	LAN	YES					
C	100.20.2.1/32		Local	100.20.2.1	0.0.0.0	23	Loopback1
	LAN	YES					
S	100.20.3.1/32		Local	100.20.1.1	10.20.0.3	29	Ethernet0/2.30
	LAN	YES					

B1MCCR# show domain default vrf green master route-import remote

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
 E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2

ia - IS-IS inter area, * - candidate default, U - per-user static route
 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
 a - application route
 + - replicated route, % - next hop override, p - overrides from PfR

Enterprise Prefix List:

Prefix: 100.20.0.0, Mask: 16

Prefix: 100.30.0.0, Mask: 16

Prefix: 100.0.0.0, Mask: 8

Proto	Prefix	Location	BR-IP	Next-Hop	Index	Interface
	IF-Role	In-RIB				
L	10.20.1.1/32	Remote	100.20.1.1	0.0.0.0	25	Ethernet0/1.30
	WAN	YES				
C	10.20.1.0/24	Remote	100.20.1.1	0.0.0.0	25	Ethernet0/1.30
	WAN	YES				
L	51.1.0.4/32	Remote	100.20.1.1	0.0.0.0	24	Tunnel10
	WAN	YES				
C	51.1.0.0/16	Remote	100.20.1.1	0.0.0.0	24	Tunnel10
	WAN	YES				
D	52.1.0.0/16	Remote	100.20.1.1	10.20.0.2	29	Ethernet0/2.30
	LAN	NO				
D	52.1.0.0/16	Remote	100.20.1.1	51.1.0.3	24	Tunnel10
	WAN	YES				
B	10.10.1.0/24	Remote	100.20.1.1	10.20.1.2	25	Ethernet0/1.30
	WAN	YES				
B	10.10.3.0/24	Remote	100.20.1.1	10.20.1.2	25	Ethernet0/1.30
	WAN	YES				
B	10.15.1.0/24	Remote	100.20.1.1	10.20.1.2	25	Ethernet0/1.30
	WAN	YES				
B	10.30.1.0/24	Remote	100.20.1.1	10.20.1.2	25	Ethernet0/1.30
	WAN	YES				
D	100.10.0.0/16	Remote	100.20.1.1	10.20.0.2	29	Ethernet0/2.30
	LAN	NO				
D	100.10.0.0/16	Remote	100.20.1.1	51.1.0.2	24	Tunnel10
	WAN	YES				

B1MCBR# show domain default vrf green master route-import border 100.20.1.1

Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
 D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
 N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
 E1 - OSPF external type 1, E2 - OSPF external type 2
 i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
 ia - IS-IS inter area, * - candidate default, U - per-user static route
 o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
 a - application route
 + - replicated route, % - next hop override, p - overrides from PfR

Enterprise Prefix List:

Prefix: 100.20.0.0, Mask: 16

Prefix: 100.30.0.0, Mask: 16

Prefix: 100.0.0.0, Mask: 8

Proto	Prefix	Location	BR-IP	Next-Hop	Index	Interface
	IF-Role	In-RIB				
L	10.20.0.1/32	Local	100.20.1.1	0.0.0.0	29	Ethernet0/2.30
	LAN	YES				
C	10.20.0.0/24	Local	100.20.1.1	0.0.0.0	29	Ethernet0/2.30
	LAN	YES				
L	10.20.1.1/32	Remote	100.20.1.1	0.0.0.0	25	Ethernet0/1.30
	WAN	YES				
C	10.20.1.0/24	Remote	100.20.1.1	0.0.0.0	25	Ethernet0/1.30

Displaying Site Prefixes Learnt By a Master Controller

		WAN	YES					
D	10.20.2.0/24	LAN	YES	Local	100.20.1.1	10.20.0.2	29	Ethernet0/2.30
L	51.1.0.4/32	WAN	YES	Remote	100.20.1.1	0.0.0.0	24	Tunnel10
C	51.1.0.0/16	WAN	YES	Remote	100.20.1.1	0.0.0.0	24	Tunnel10
D	52.1.0.0/16	LAN	NO	Remote	100.20.1.1	10.20.0.2	29	Ethernet0/2.30
D	52.1.0.0/16	WAN	YES	Local	100.20.1.1	0.0.0.0	22	Loopback1
C	100.20.1.1/32	LAN	YES	Local	100.20.1.1	10.20.0.2	29	Ethernet0/2.30
D	100.20.2.1/32	LAN	YES	Local	100.20.1.1	10.20.0.2	29	Ethernet0/2.30

Step 2 show domain default vrf *vrf name* master route-import interface

Use this command to view the prefix information associated with an interface.

Example:

```
Router# show domain default vrf green border local-prefix interface Ethernet0/0.10
Codes: L - local, C - connected, S - static, R - RIP, M - mobile,
       B-BGP D - EIGRP, EX - EIGRP external, O - OSPF,
       IA - OSPF inter area N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2 E1 - OSPF external type 1,
       E2 - OSPF external type 2 i - IS-IS, su - IS-IS summary,
       L1 - IS-IS level-1, L2 - IS-IS level-2 ia - IS-IS inter area,
       * - candidate default, H- NHRP
Local  -- Prefix learned over LAN.
Remote - Prefix learned over WAN.
Prefix      Interface      BR IP      Index  Prefix-site  Proto Next-Hop      Status
-----
100.10.4.1/32 Ethernet0/0.10 100.20.1.1 12     Local        C          -----
Up
```

Step 3 show domain default vrf *vrf name* master local-prefix

Use this command to view the prefix information associated with an border router.

Example:

```
Router# show domain default vrf green master local-prefix border-ip 100.20.1.1
Codes: L - local, C - connected, S - static, R - RIP, M - mobile,
       B-BGP D - EIGRP, EX - EIGRP external, O - OSPF,
       IA - OSPF inter area N1 - OSPF NSSA external type 1,
       N2 - OSPF NSSA external type 2 E1 - OSPF external type 1,
       E2 - OSPF external type 2 i - IS-IS, su - IS-IS summary,
       L1 - IS-IS level-1, L2 - IS-IS level-2 ia - IS-IS inter area,
       * - candidate default, H- NHRP
Local  -- Prefix learned over LAN.
Remote - Prefix learned over WAN.
Prefix      Interface      BR IP      Index  Prefix-site  Proto  Next-Hop      Status
-----
100.10.4.1/32 Ethernet0/0.10 100.20.1.1 12     Local        C
```

Additional References for PfRv3 Remote Prefix Tracking

Related Documents

Related Topic	Document Title
PfRv3commands	Cisco IOS Performance Routing Version 3 Command Reference
Site Prefix Splitting	Site Prefix Splitting

