



BGP Support for IP Prefix Import from Global Table into a VRF Table

The BGP Support for IP Prefix Import from Global Table into a VRF Table feature introduces the capability to import IPv4 unicast prefixes from the global routing table into a Virtual Private Network (VPN) routing/forwarding (VRF) instance table using an import route map.

- [Prerequisites for BGP Support for IP Prefix Import from Global Table into a VRF Table, on page 1](#)
- [Restrictions for BGP Support for IP Prefix Import from Global Table into a VRF Table, on page 1](#)
- [Information About BGP Support for IP Prefix Import from Global Table into a VRF Table, on page 2](#)
- [How to Import IP Prefixes from Global Table into a VRF Table, on page 3](#)
- [Configuration Examples for BGP Support for IP Prefix Import from Global Table into a VRF Table, on page 9](#)
- [Additional References for Internal BGP Features, on page 10](#)
- [Feature Information for BGP Support for IP Prefix Import from Global Table into a VRF Table, on page 11](#)

Prerequisites for BGP Support for IP Prefix Import from Global Table into a VRF Table

- Border Gateway Protocol (BGP) peering sessions are established.
- CEF or dCEF (for distributed platforms) is enabled on all participating routers.

Restrictions for BGP Support for IP Prefix Import from Global Table into a VRF Table

- Only IPv4 unicast and multicast prefixes can be imported into a VRF with this feature.
- IPv4 prefixes imported into a VRF using this feature cannot be imported into a VPNv4 VRF.
- The global prefixes should be in the BGP table, so that this feature can import them into the BGP VRF table.

- IPv4 prefixes imported into a VRF using this feature cannot be imported into a second VPNv4 VRF.

Information About BGP Support for IP Prefix Import from Global Table into a VRF Table

Importing IPv4 Prefixes into a VRF

The BGP Support for IP Prefix Import from Global Table into a VRF Table feature introduces the capability to import IPv4 unicast prefixes from the global routing table into a Virtual Private Network (VPN) routing/forwarding instance (VRF) table using an import route map. This feature extends the functionality of VRF import-map configuration to allow IPv4 prefixes to be imported into a VRF based on a standard community. Both IPv4 unicast and multicast prefixes are supported. No Multiprotocol Label Switching (MPLS) or route target (import/export) configuration is required.

IP prefixes are defined as match criteria for the import map through standard Cisco filtering mechanisms. For example, an IP access-list, an IP prefix-list, or an IP as-path filter is created to define an IP prefix or IP prefix range, and then the prefix or prefixes are processed through a match clause in a route map. Prefixes that pass through the route map are imported into the specified VRF per the import map configuration.

Black Hole Routing

The BGP Support for IP Prefix Import from Global Table into a VRF Table feature can be configured to support Black Hole Routing (BHR). BHR is a method that allows the administrator to block undesirable traffic, such as traffic from illegal sources or traffic generated by a Denial of Service (DoS) attack, by dynamically routing the traffic to a dead interface or to a host designed to collect information for investigation, mitigating the impact of the attack on the network. Prefixes are looked up, and packets that come from unauthorized sources are rendered as null routes by the ASIC at line rate.

Classifying Global Traffic

The BGP Support for IP Prefix Import from Global Table into a VRF Table feature can be used to classify global IP traffic based on physical location or class of service. Traffic is classified based on administration policy and then imported into different VRFs. On a college campus, for example, network traffic could be divided into an academic network and residence network traffic, a student network and faculty network, or a dedicated network for multicast traffic. After the traffic is divided along administration policy, routing decisions can be configured with the MPLS VPN--VRF Selection Using Policy Based Routing feature or the MPLS VPN--VRF Selection Based on Source IP Address feature.

Unicast Reverse Path Forwarding

Unicast Reverse Path Forwarding (Unicast RPF) can be optionally configured with the BGP Support for IP Prefix Import from Global Table into a VRF Table feature. Unicast RPF is used to verify that the source address is in the Forwarding Information Base (FIB). The **ip verify unicast vrf** command is configured in interface configuration mode and is enabled for each VRF. This command has **permit** and **deny** keywords that are used to determine if the traffic is forwarded or dropped after Unicast RPF verification.

How to Import IP Prefixes from Global Table into a VRF Table

Defining IPv4 IP Prefixes to Import

IPv4 unicast or multicast prefixes are defined as match criteria for the import route map using standard Cisco filtering mechanisms. This task uses an IP access-list and an IP prefix-list.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **access-list** *access-list-number* {deny | permit} *source* [*source-wildcard*] [log]
4. **ip prefix-list** *prefix-list-name* [seq *seq-value*] {deny *network/length* | permit *network/length*} [ge *ge-value*] [le *le-value*]

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable Example: Device> enable	Enables privileged EXEC mode. <ul style="list-style-type: none"> • Enter your password if prompted.
Step 2	configure terminal Example: Device# configure terminal	Enters global configuration mode.
Step 3	access-list <i>access-list-number</i> {deny permit} <i>source</i> [<i>source-wildcard</i>] [log] Example: Device(config)# access-list 50 permit 10.1.1.0 0.0.0.255	Creates an access list and defines a range of IP prefixes to import into the VRF table. <ul style="list-style-type: none"> • The example creates a standard access list numbered 50. This filter will permit traffic from any host with an IP address in the 10.1.1.0/24 subnet.
Step 4	ip prefix-list <i>prefix-list-name</i> [seq <i>seq-value</i>] {deny <i>network/length</i> permit <i>network/length</i> } [ge <i>ge-value</i>] [le <i>le-value</i>] Example: Device(config)# ip prefix-list COLORADO permit 10.24.240.0/22	Creates a prefix list and defines a range of IP prefixes to import into the VRF table. <ul style="list-style-type: none"> • The example creates an IP prefix list named COLORADO. This filter will permit traffic from any host with an IP address in the 10.24.240.0/22 subnet.

Creating the VRF and the Import Route Map

The IP prefixes that are defined for import are then processed through a match clause in a route map. IP prefixes that pass through the route map are imported into the VRF. A maximum of 5 VRFs per router can be configured to import IPv4 prefixes from the global routing table. By default, a maximum of 1000 prefixes per VRF can be imported. You can change the limit to be from 1 to 2,147,483,647 prefixes for each VRF. We recommend that you use caution if you increase the prefix import limit above 1000. Configuring the router to import too many prefixes can interrupt normal router operation.

No MPLS or route target (import/export) configuration is required.

Import actions are triggered when a new routing update is received or when routes are withdrawn. During the initial BGP update period, the import action is postponed to allow BGP to convergence more quickly. Once BGP converges, incremental BGP updates are evaluated immediately and qualified prefixes are imported as they are received.

The following syslog message is introduced by the BGP Support for IP Prefix Import from Global Table into a VRF Table feature. It will be displayed when more prefixes are available for import than the user-defined limit:

```
00:00:33: %BGP-3-AFIMPORT_EXCEED: IPv4 Multicast prefixes imported to multicast vrf exceed
the limit 2
```

You can either increase the prefix limit or fine-tune the import route map filter to reduce the number of candidate routes.



Note

- Only IPv4 unicast and multicast prefixes can be imported into a VRF with this feature.
- A maximum of five VRF instances per router can be created to import IPv4 prefixes from the global routing table.
- IPv4 prefixes imported into a VRF using this feature cannot be imported into a VPNv4 VRF.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **ip vrf** *vrf-name*
4. **rd** *route-distinguisher*
5. **import ipv4** {unicast | multicast} [*prefix-limit*] **map** *route-map*
6. **exit**
7. **route-map** *map-tag* [**permit** | **deny**] [*sequence-number*]
8. **match ip address** {*acl-number* [*acl-number* | *acl-name*] | *acl-name* [*acl-name* | *acl-number*] | **prefix-list** *prefix-list-name* [*prefix-list-name*]}
9. **end**

DETAILED STEPS

	Command or Action	Purpose
Step 1	enable	Enables privileged EXEC mode.

	Command or Action	Purpose
	Example: <pre>Router> enable</pre>	<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	ip vrf vrf-name Example: <pre>Router(config)# ip vrf GREEN</pre>	<p>Creates a VRF routing table and specifies the VRF name (or tag).</p> <ul style="list-style-type: none"> The ip vrf vrf-name command creates a VRF routing table and a CEF table, and both are named using the <i>vrf-name</i> argument. Associated with these tables is the default route distinguisher value.
Step 4	rd route-distinguisher Example: <pre>Router(config-vrf)# rd 100:10</pre>	<p>Creates routing and forwarding tables for the VRF instance.</p> <ul style="list-style-type: none"> There are two formats for configuring the route distinguisher argument. It can be configured in the as-number:network number (ASN:nn) format, as shown in the example, or it can be configured in the IP address:network number format (IP-address:nn).
Step 5	import ipv4 {unicast multicast} [prefix-limit] map route-map Example: <pre>Router(config-vrf)# import ipv4 unicast 1000 map UNICAST</pre>	<p>Imports IPv4 prefixes from the global routing table to a VRF table, filtered by the specified route map.</p> <ul style="list-style-type: none"> Unicast or multicast prefixes are specified. Up to a 1000 prefixes will be imported by default. The <i>prefix-limit</i> argument is used to specify a limit from 1 to 2,147,483,647 prefixes. The example references a route map that will import up to 1000 unicast prefixes that pass through the route map.
Step 6	exit Example: <pre>Router(config-vrf)# exit</pre>	Exits VRF configuration mode and enters global configuration mode.
Step 7	route-map map-tag [permit deny] [sequence-number] Example: <pre>Router(config)# route-map UNICAST permit 10</pre>	<p>Defines the conditions for redistributing routes from one routing protocol into another, or enables policy routing.</p> <ul style="list-style-type: none"> The route map name must match the route map specified in Step 5. The example creates a route map named UNICAST.

	Command or Action	Purpose
Step 8	match ip address <i>{acl-number [acl-number acl-name] acl-name [acl-name acl-number] prefix-list prefix-list-name [prefix-list-name]}</i> Example: <pre>Router(config-route-map)# match ip address 50</pre>	Distributes any routes that have a destination network number address that is permitted by a standard or extended access list, and performs policy routing on matched packets. <ul style="list-style-type: none"> • Both IP access lists and IP prefix lists are supported. • The example configures the route map to use standard access list 50 to define match criteria.
Step 9	end Example: <pre>Router(config-route-map)# end</pre>	Exits route-map configuration mode and returns to privileged EXEC mode.

Filtering on the Ingress Interface

The BGP Support for IP Prefix Import from Global Table into a VRF Table feature can be configured globally or on a per-interface basis. We recommend that you apply it to ingress interfaces to maximize performance.

SUMMARY STEPS

1. **enable**
2. **configure terminal**
3. **interface** *type number [name-tag]*
4. **ip policy route-map** *map-tag*
5. **ip verify unicast vrf** *vrf-name {deny | permit}*
6. **end**

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	interface <i>type number [name-tag]</i> Example: <pre>Router(config)# interface Ethernet0/0</pre>	Configures an interface and enters interface configuration mode.

	Command or Action	Purpose
Step 4	ip policy route-map <i>map-tag</i> Example: <pre>Router(config-if)# ip policy route-map UNICAST</pre>	Identifies a route map to use for policy routing on an interface. <ul style="list-style-type: none"> The example attaches the route map named UNICAST to the interface.
Step 5	ip verify unicast vrf <i>vrf-name</i> {deny permit} Example: <pre>Router(config-if)# ip verify unicast vrf GREEN permit</pre>	(Optional) Enables Unicast Reverse Path Forwarding verification for the specified VRF. <ul style="list-style-type: none"> The example enables verification for the VRF named GREEN. Traffic that passes verification will be forwarded.
Step 6	end Example: <pre>Router(config-if)# end</pre>	Exits interface configuration mode and returns to privileged EXEC mode.

Verifying Global IP Prefix Import

Perform the steps in this task to display information about the VRFs that are configured with the BGP Support for IP Prefix Import from Global Table into a VRF Table feature and to verify that global IP prefixes are imported into the specified VRF table.

SUMMARY STEPS

1. **enable**
2. **show ip bgp vpnv4** {all | rd *route-distinguisher* | vrf *vrf-name*}
3. **show ip vrf** [brief | detail | interfaces | id] [*vrf-name*]

DETAILED STEPS

Step 1 enable

Enables privileged EXEC mode. Enter your password if prompted.

Example:

```
Device# enable
```

Step 2 show ip bgp vpnv4 {all | rd *route-distinguisher* | vrf *vrf-name*}

Displays VPN address information from the BGP table. The output displays the import route map, the traffic type (unicast or multicast), the default or user-defined prefix import limit, the actual number of prefixes that are imported, and individual import prefix entries.

Example:

```
Device# show ip bgp vpnv4 all
```

```

BGP table version is 15, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network        Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 100:1 (default for vrf academic)
Import Map: ACADEMIC, Address-Family: IPv4 Unicast, Pfx Count/Limit: 6/1000
*> 10.50.1.0/24    172.17.2.2                0 2 3 ?
*> 10.50.2.0/24    172.17.2.2                0 2 3 ?
*> 10.50.3.0/24    172.17.2.2                0 2 3 ?
*> 10.60.1.0/24    172.17.2.2                0 2 3 ?
*> 10.60.2.0/24    172.17.2.2                0 2 3 ?
*> 10.60.3.0/24    172.17.2.2                0 2 3 ?
Route Distinguisher: 200:1 (default for vrf residence)
Import Map: RESIDENCE, Address-Family: IPv4 Unicast, Pfx Count/Limit: 3/1000
*> 10.30.1.0/24    172.17.2.2                0      0 2 i
*> 10.30.2.0/24    172.17.2.2                0      0 2 i
*> 10.30.3.0/24    172.17.2.2                0      0 2 i
Route Distinguisher: 300:1 (default for vrf BLACKHOLE)
Import Map: BLACKHOLE, Address-Family: IPv4 Unicast, Pfx Count/Limit: 3/1000
*> 10.40.1.0/24    172.17.2.2                0      0 2 i
*> 10.40.2.0/24    172.17.2.2                0      0 2 i
*> 10.40.3.0/24    172.17.2.2                0      0 2 i
Route Distinguisher: 400:1 (default for vrf multicast)
Import Map: MCAST, Address-Family: IPv4 Multicast, Pfx Count/Limit: 2/2
*> 10.70.1.0/24    172.17.2.2                0      0 2 i
*> 10.70.2.0/24    172.17.2.2                0      0 2 i

```

Step 3 **show ip vrf [brief | detail | interfaces | id] [vrf-name]**

Displays defined VRFs and their associated interfaces. The output displays the import route map, the traffic type (unicast or multicast), and the default or user-defined prefix import limit. The following example output shows that the import route map named UNICAST is importing IPv4 unicast prefixes and that the prefix import limit is 1000.

Example:

```

Device# show ip vrf detail

VRF academic; default RD 100:10; default VPNID <not set>
VRF Table ID = 1
  No interfaces
  Connected addresses are not in global routing table
  Export VPN route-target communities
    RT:100:10
  Import VPN route-target communities
    RT:100:10
  Import route-map for ipv4 unicast: UNICAST (prefix limit: 1000)
  No export route-map

```


Configuration Examples for BGP Support for IP Prefix Import from Global Table into a VRF Table

Example: Importing IP Prefixes from Global Table into a VRF Table

The following example imports unicast prefixes into the VRF named *green* by using an IP prefix list and a route map:

This example starts in global configuration mode:

```
!  
ip prefix-list COLORADO seq 5 permit 10.131.64.0/19  
ip prefix-list COLORADO seq 10 permit 172.31.2.0/30  
ip prefix-list COLORADO seq 15 permit 172.31.1.1/32  
!  
ip vrf green  
  rd 200:1  
  import ipv4 unicast map UNICAST  
  route-target export 200:10  
  route-target import 200:10  
!  
exit  
!  
route-map UNICAST permit 10  
  match ip address prefix-list COLORADO  
!  
exit
```

Example: Verifying IP Prefix Import to a VRF Table

The **show ip vrf** command or the **show ip bgp vpnv4** command can be used to verify that prefixes are imported from the global routing table to the VRF table.

The following sample output shows that the import route map named UNICAST is importing IPv4 unicast prefixes and the prefix import limit is 1000:

```
Device# show ip vrf detail  
  
VRF green; default RD 200:1; default VPNID <not set>  
  Interfaces:  
    Se2/0  
VRF Table ID = 1  
  Export VPN route-target communities  
    RT:200:10  
  Import VPN route-target communities  
    RT:200:10  
  Import route-map for ipv4 unicast: UNICAST (prefix limit: 1000)  
  No export route-map  
  VRF label distribution protocol: not configured  
  VRF label allocation mode: per-prefix  
VRF red; default RD 200:2; default VPNID <not set>  
  Interfaces:  
    Se3/0  
VRF Table ID = 2
```

```

Export VPN route-target communities
  RT:200:20
Import VPN route-target communities
  RT:200:20
No import route-map
No export route-map
VRF label distribution protocol: not configured
VRF label allocation mode: per-prefix

```

The following sample output displays the import route map names, the prefix import limit and the actual number of imported prefixes, and the individual import entries:

```

Device# show ip bgp vpnv4 all

BGP table version is 18, local router ID is 10.131.127.252
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
   Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 200:1 (default for vrf green)
Import Map: UNICAST, Address-Family: IPv4 Unicast, Pfx Count/Limit: 1/1000
*>i10.131.64.0/19    10.131.95.252      0      100      0 i
*> 172.16.1.1/32    172.16.2.1         0              32768 i
*> 172.16.2.0/30    0.0.0.0            0              32768 i
*>i172.31.1.1/32    10.131.95.252      0      100      0 i
*>i172.31.2.0/30    10.131.95.252      0      100      0 i
Route Distinguisher: 200:2 (default for vrf red)
*> 172.16.1.1/32    172.16.2.1         0              32768 i
*> 172.16.2.0/30    0.0.0.0            0              32768 i
*>i172.31.1.1/32    10.131.95.252      0      100      0 i
*>i172.31.2.0/30    10.131.95.252      0      100      0 i

```

Additional References for Internal BGP Features

Related Documents

Related Topic	Document Title
Cisco IOS commands	Cisco IOS Master Commands List, All Releases
BGP commands	Cisco IOS IP Routing: BGP Command Reference
BGP overview	“Cisco BGP Overview” module
Basic BGP configuration tasks	“Configuring a Basic BGP Network” module
iBGP multipath load sharing	“iBGP Multipath Load Sharing” module
Connecting to a service provider	“Connecting to a Service Provider Using External BGP” module
Configuring features that apply to multiple IP routing protocols	<i>IP Routing: Protocol-Independent Configuration Guide</i>

RFCs

RFC	Title
RFC 1772	<i>Application of the Border Gateway Protocol in the Internet</i>
RFC 1773	<i>Experience with the BGP Protocol</i>
RFC 1774	<i>BGP-4 Protocol Analysis</i>
RFC 1930	<i>Guidelines for Creation, Selection, and Registration of an Autonomous System (AS)</i>
RFC 2519	<i>A Framework for Inter-Domain Route Aggregation</i>
RFC 2858	<i>Multiprotocol Extensions for BGP-4</i>
RFC 2918	<i>Route Refresh Capability for BGP-4</i>
RFC 3392	<i>Capabilities Advertisement with BGP-4</i>
RFC 4271	<i>A Border Gateway Protocol 4 (BGP-4)</i>
RFC 4893	<i>BGP Support for Four-octet AS Number Space</i>
RFC 5396	<i>Textual Representation of Autonomous system (AS) Numbers</i>
RFC 5398	<i>Autonomous System (AS) Number Reservation for Documentation Use</i>

Technical Assistance

Description	Link
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.	http://www.cisco.com/cisco/web/support/index.html

Feature Information for BGP Support for IP Prefix Import from Global Table into a VRF Table

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 BGP Support for IP Prefix Import from Global Table into a VRF Table

Feature Name	Releases	Feature Information
BGP Support for IP Prefix Import from Global Table into a VRF Table	Cisco IOS XE Release 2.1	<p>The BGP Support for IP Prefix Import from Global Table into a VRF Table feature introduces the capability to import IPv4 unicast prefixes from the global routing table into a Virtual Private Network (VPN) routing/forwarding (VRF) instance table using an import route map.</p> <p>This feature was introduced on the Cisco ASR 1000 Series Aggregation Services Routers.</p> <p>The following commands were introduced or modified by this feature: debug ip bgp import, import ipv4, ip verify unicast vrf.</p>