# Configure VRF Route Leak on Nexus Switches

## Contents


**Introduction**

**Prerequisites**

- Requirements
- Components Used
- Limitations

**Configure**

- Default VRF to VRF
- VRF to VRF
- VRF to Default VRF

**Verify**

**Related Information**


## Introduction

This document describes how to configure Route Leak on Cisco Nexus NX-OS based Switches.

## Prerequisites

### Requirements

Cisco recommends that you have knowledge of these topics:

- Nexus NX-OS Software
- Routing Protocols like Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), Border Gateway Protocol (BGP), and so on.

### Components Used

The information in this document is based on Cisco Nexus 7000 with NXOS version 7.3(0)D1(1).

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

### Limitations

You must leak a route directly from the source VRF into the target VRF. You cannot leak a route that is currently leaked from another VRF.

Consider that a BGP session from the Nexus cannot be established to a peer IP when routed through a different VRF on the Nexus.

---

**Note**: When leaking routes between VRFs with BGP, you might experience that route changes are not reflected in the destination VRF until several minutes later. This can happen due to BGP dampening mechanisms which ignore route changes based on metric changes to avoid cases in which an IGP is flapping (link issues, for example) and you do not want these flaps to be propagated beyond. If you suffer from this condition and need the routes to be immediately updated on the destination VRF as well, please take a look at the guide, Configure dampen-igp-metric in BGP on Cisco Nexus Switches, in the Related Information section.

---

## Configure

The leak between VRFs is performed at BGP process level. Because of this, it is necessary to add the routes to the BGP process first, specifically in the BGP table.

---

**Note**: The terms Default VRF and Global Routing Table are used interchangeably in this document.

---

### Default VRF to VRF

In this case, Nexus has received two routes in its Default VRF via EIGRP. The configuration leaks the routes in VRF BLUE.

For the purpose of this example, only the route 192.168.2.0/24 is leaked.

**Global Routing Table Output**

```
<#root>

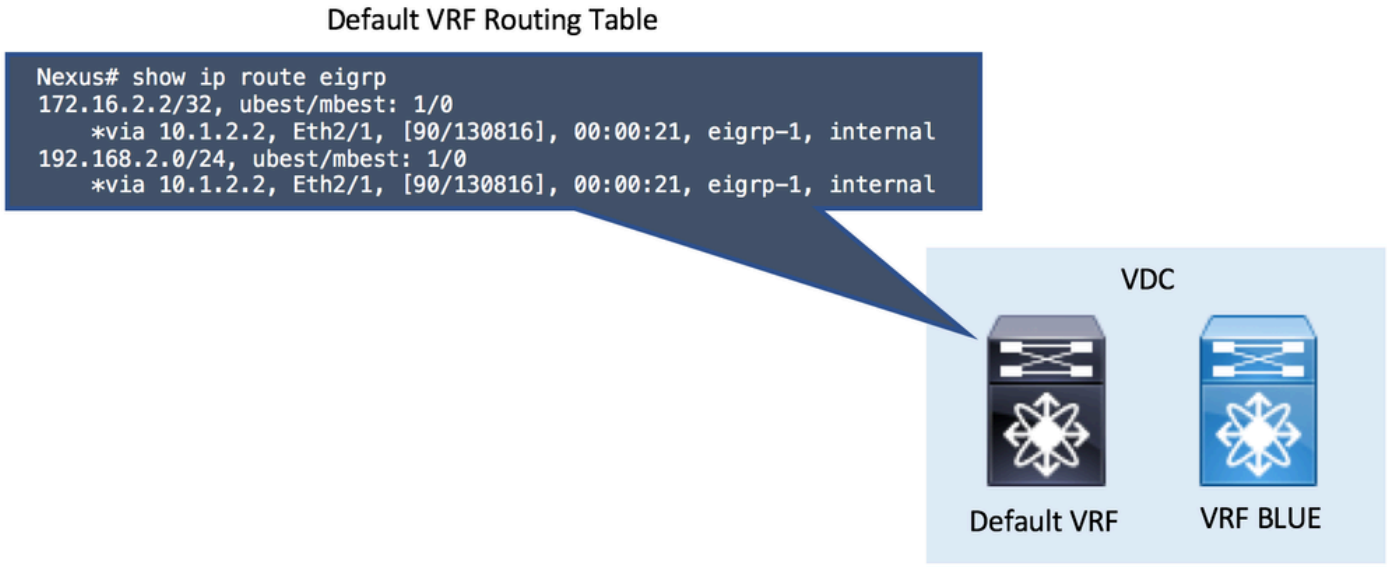
Nexus#

show ip route eigrp

IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

172.16.2.2/32, ubest/mbest: 1/0
```

```
    *via 10.1.2.2, Eth2/1, [90/130816], 00:00:21, eigrp-1, internal
192.168.2.0/24, ubest/mbest: 1/0
    *via 10.1.2.2, Eth2/1, [90/130816], 00:00:21, eigrp-1, internal
Nexus#
```




**Note**: In NX-OS, enabling the features in Global Configuration Mode is required. In order to enable BGP, the command is **feature bgp**.

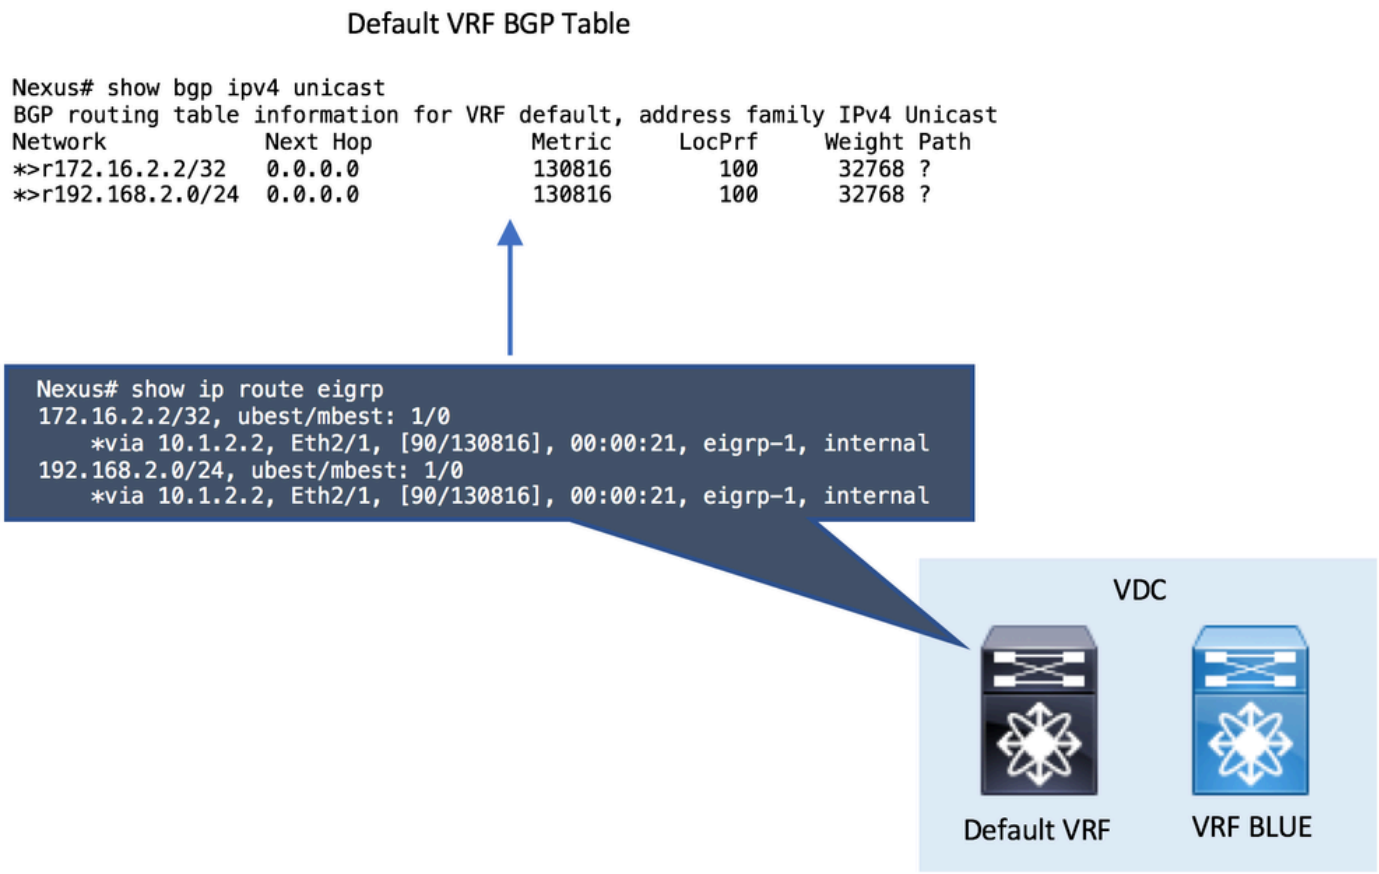
- Step 1. Redistribute into BGP.

Redistribute the routes which exist in the Default VRF Routing Table in BGP. Since the routes are in the Default VRF, the **redistribute** command in BGP goes under the global **address-family ipv4 unicast** section. Use the correct parameter for the **redistribute** command. This depends on how the routes are in the Default VRF (directly connected, eigrp, ospf, and so on).

**Note**: You can skip Step 1 in all scenarios if the routes to be leaked are installed as BGP routes in the origin VRF. In this example, the origin VRF is the Default VRF (Global Routing Table).

| **Redistribute into BGP** |
| --- |
| |

```
route-map ALL permit 10
!
router bgp 65535
  address-family ipv4 unicast
    redistribute eigrp 1 route-map ALL
```




**Note**: In NX-OS a route-map is always required as a parameter in order to selectively redistribute routes. An empty route-map permit statement created is valid in order to match any and all the routes.

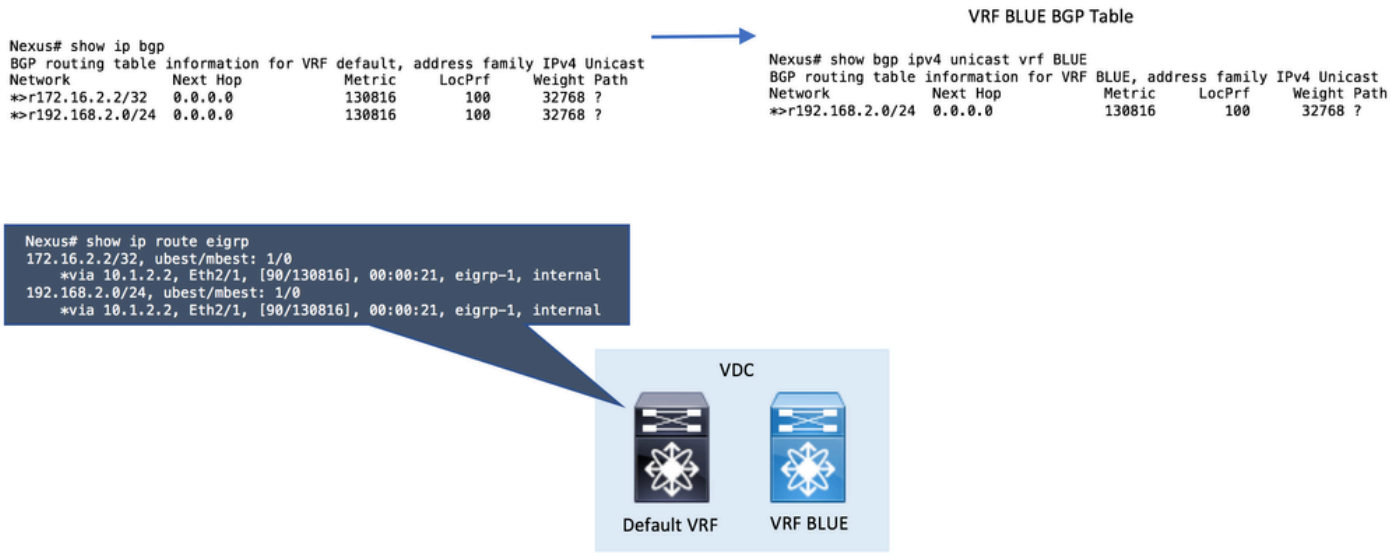
- Step 2. Configure Import VRF default in the destination VRF.

The **import vrf default** command is configured in the destination VRF. The command line requires a route-map as a parameter in order to explicitly define the routes to be imported in the destination VRF, which in this case is the VRF named BLUE.

**Configure Import VRF Default in the Destination VRF**

```
ip prefix-list NETWORK seq 5 permit 192.168.2.0/24
!
route-map GLOBAL-TO-VRF permit 10
  match ip address prefix-list NETWORK
```

```
!
vrf context BLUE
  address-family ipv4 unicast
    import vrf default map GLOBAL-TO-VRF
```




- Step 3. Check the destination VRF Routing Table.

You can confirm in the destination VRF that the routes are now seen via BGP. These BGP routes in the VRF can now be redistributed in any other Routing Protocol that runs in the same VRF.

**Check the Destination VRF Routing Table**

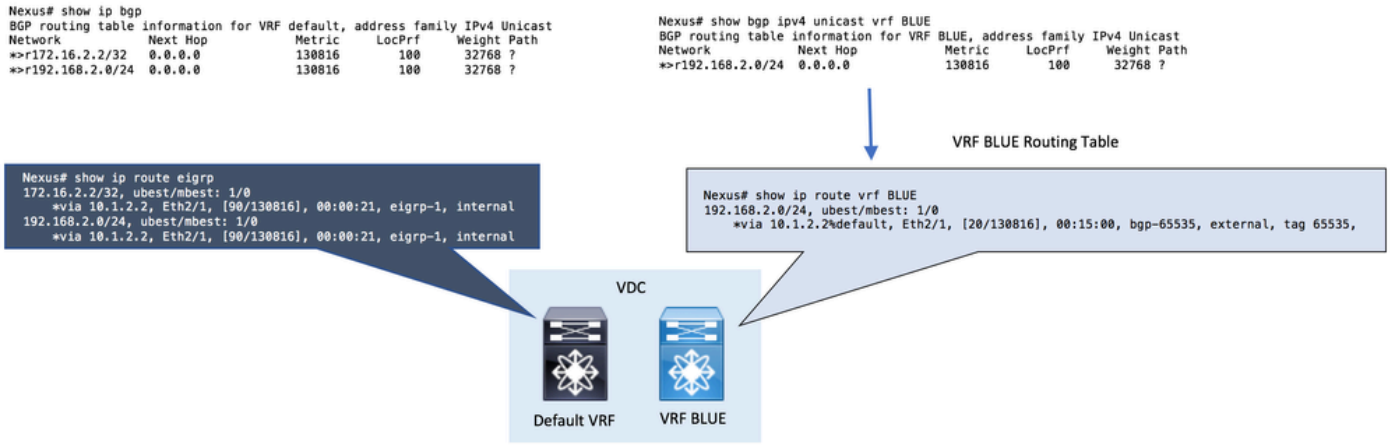
```
<#root>

Nexus#

show ip route vrf BLUE

IP Route Table for VRF "BLUE"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

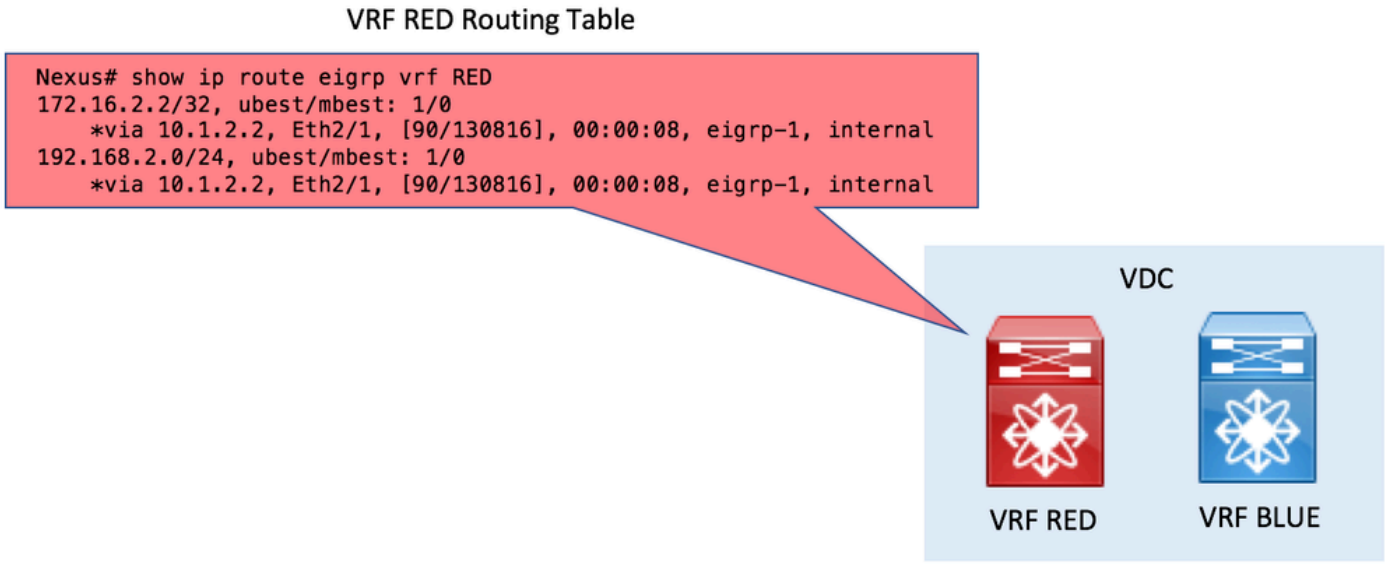
192.168.2.0/24, ubest/mbest: 1/0
    *via 10.1.2.2%default, Eth2/1, [20/130816], 00:15:00, bgp-65535, external, tag 65535,
Nexus#
```




## VRF to VRF

In this case, Nexus has received two routes in its VRF called RED via EIGRP. The configuration leaks the routes in VRF BLUE.

| VRF RED Routing Table Output |
| --- |
| <#root><br><br>Nexus#<br><br>**show ip route eigrp vrf RED**<br><br>IP Route Table for VRF "RED"<br>'*' denotes best ucast next-hop<br>'**' denotes best mcast next-hop<br>'[x/y]' denotes [preference/metric]<br>'%<string>' in via output denotes VRF <string><br><br>172.16.2.2/32, ubest/mbest: 1/0<br>*via 10.1.2.2, Eth2/1, [90/130816], 00:00:08, eigrp-1, internal<br>192.168.2.0/24, ubest/mbest: 1/0<br>*via 10.1.2.2, Eth2/1, [90/130816], 00:00:08, eigrp-1, internal<br>Nexus# |

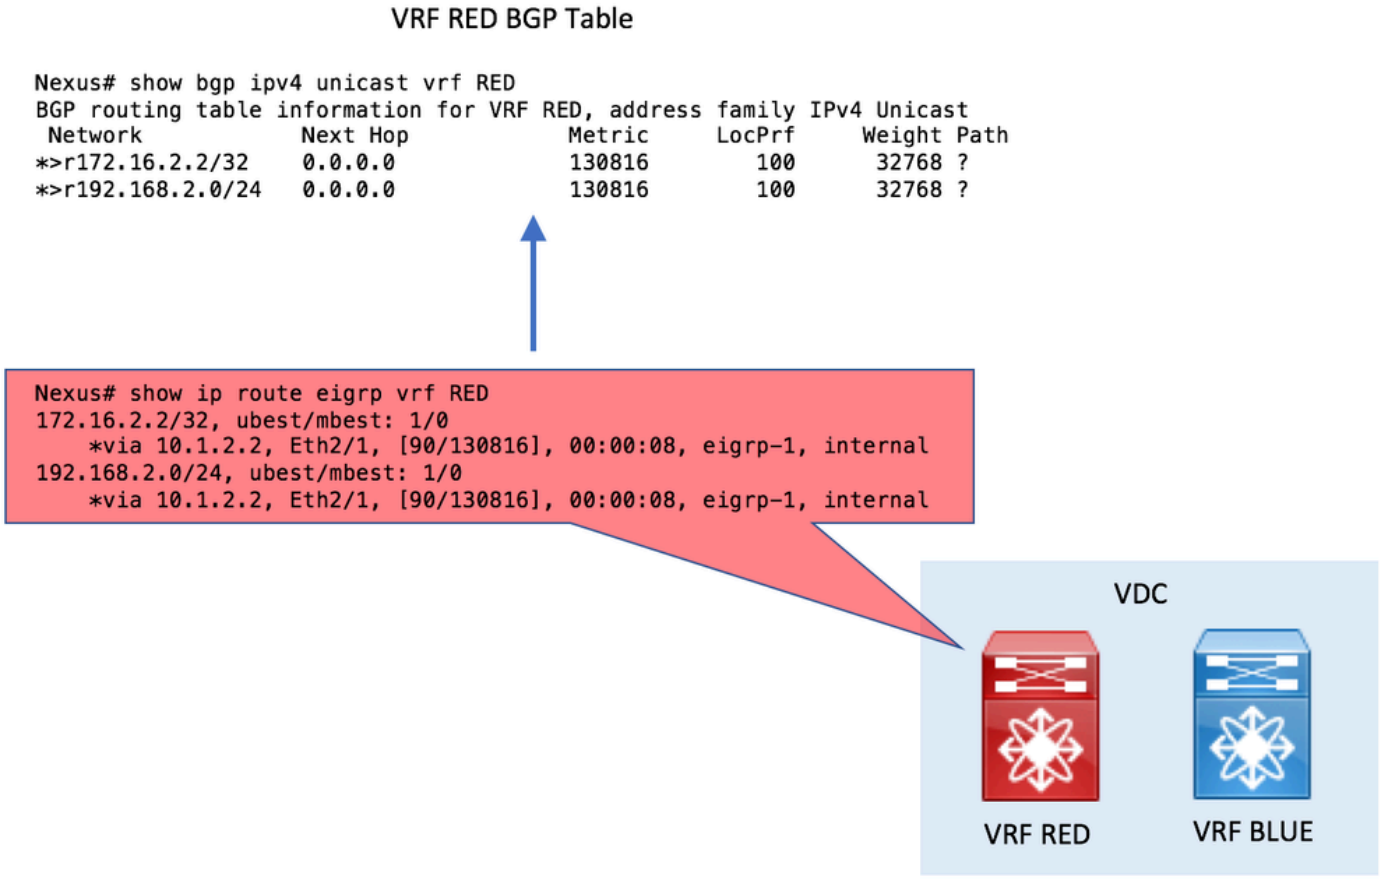



- Step 1. Redistribute into BGP.

Redistribute the routes which exist in VRF RED Routing Table in BGP. Since the routes are in the VRF RED, the **redistribute** command in BGP goes under the vrf RED address-family ipv4 unicast section.

**Redistribute into BGP**

```
route-map ALL permit 10
!
router bgp 65535
  vrf RED
    address-family ipv4 unicast
      redistribute eigrp 1 route-map ALL
```

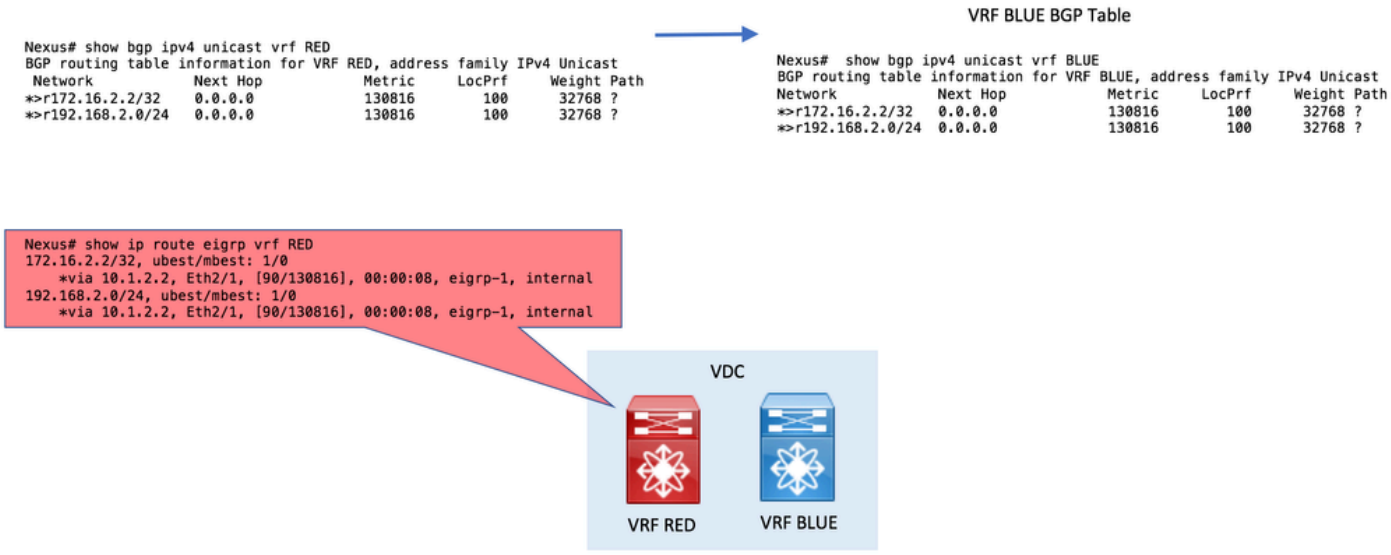



- Step 2. Create export and import Route-Targets.

In order to leak between VRFs, the use of Route-Targets is required. The origin VRF exports a Route-Target value. The destination VRF imports the same Route-Target value.

**Create Export and Import Route-Targets**

```
vrf context RED
  address-family ipv4 unicast
    route-target export 1:1
!
vrf context BLUE
  address-family ipv4 unicast
    route-target import 1:1
```




- Step 3. Check the destination VRF Routing Table.

You can confirm in the destination VRF that the routes are now seen via BGP. These BGP routes in the VRF can now be redistributed in any other Routing Protocol that runs in the same VRF.

**Check the Destination VRF Routing Table**

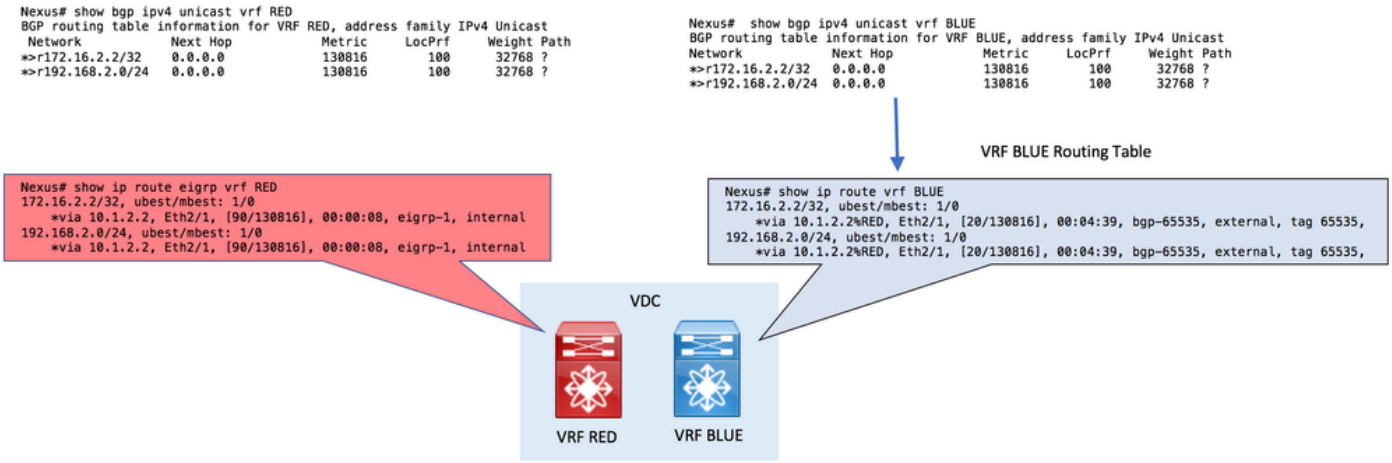
```
<#root>

Nexus#

show ip route vrf BLUE

IP Route Table for VRF "BLUE"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

172.16.2.2/32, ubest/mbest: 1/0
    *via 10.1.2.2%RED, Eth2/1, [20/130816], 00:01:58, bgp-65535, external, tag 65535,
192.168.2.0/24, ubest/mbest: 1/0
    *via 10.1.2.2%RED, Eth2/1, [20/130816], 00:01:58, bgp-65535, external, tag 65535,
Nexus#
```




- Step 4 (Optional). Assign Route-Target to specific routes.

You can optionally use the **export map** command under the origin VRF in order to assign Route-Targets to specific routes to be exported. Use the set extcommunity rt parameter in the route-map in order to assign the Route-Target.

In this example, only network 192.168.2.0/24 is exported with Route-Target 1:1 which is later imported in VRF BLUE.

The result is that only the specified network is leaked.

**Assign Route-Target to Specific Routes**

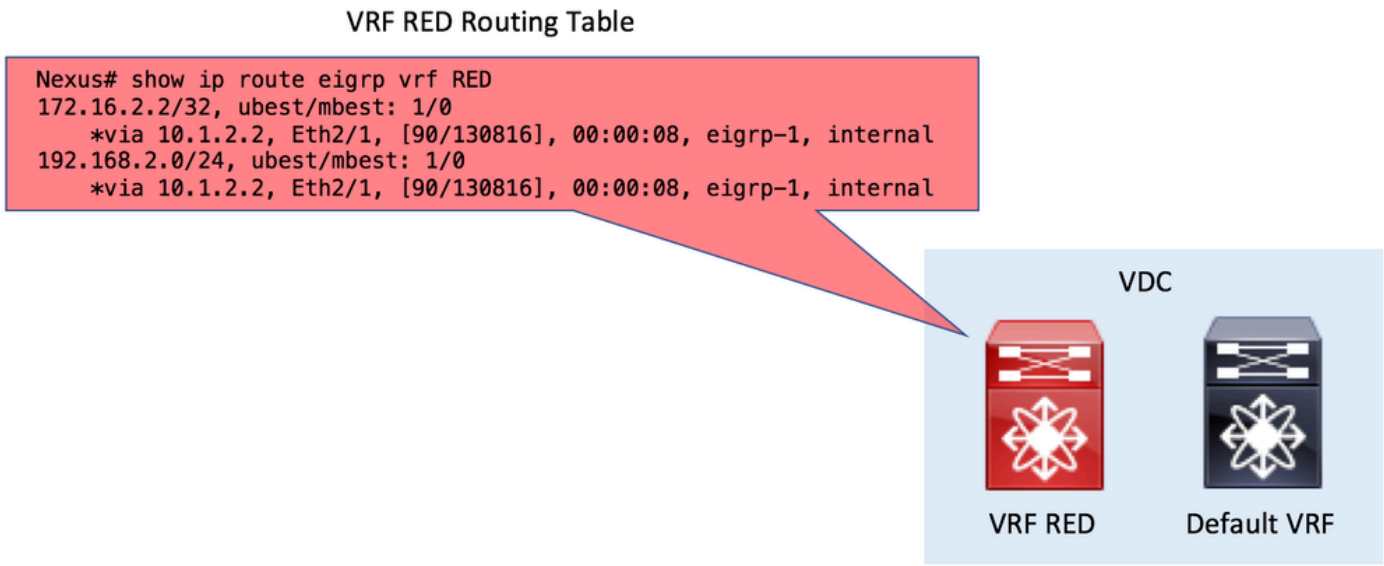
```
ip prefix-list NETWORK seq 5 permit 192.168.2.0/24
!
route-map ADD-RT permit 10
  match ip address prefix-list NETWORK
  set extcommunity rt 1:1
!
vrf context RED
  address-family ipv4 unicast
    export map ADD-RT
!
vrf context BLUE
  address-family ipv4 unicast
    route-target import 1:1
```

## VRF to Default VRF

**Note**: On Nexus 7000 and Nexus 7700 family of Switches, this feature that supports the export of IP prefixes to the global routing table (the default VRF) from any other VRF with the use of the **export vrf default map** command was introduced in NX-OS release 7.3(0)D1(1).

Nexus has received two routes in its VRF called RED via EIGRP. The configuration leaks the routes in the Default VRF.

For the purpose of this example, only the route 192.168.2.0/24 is leaked.

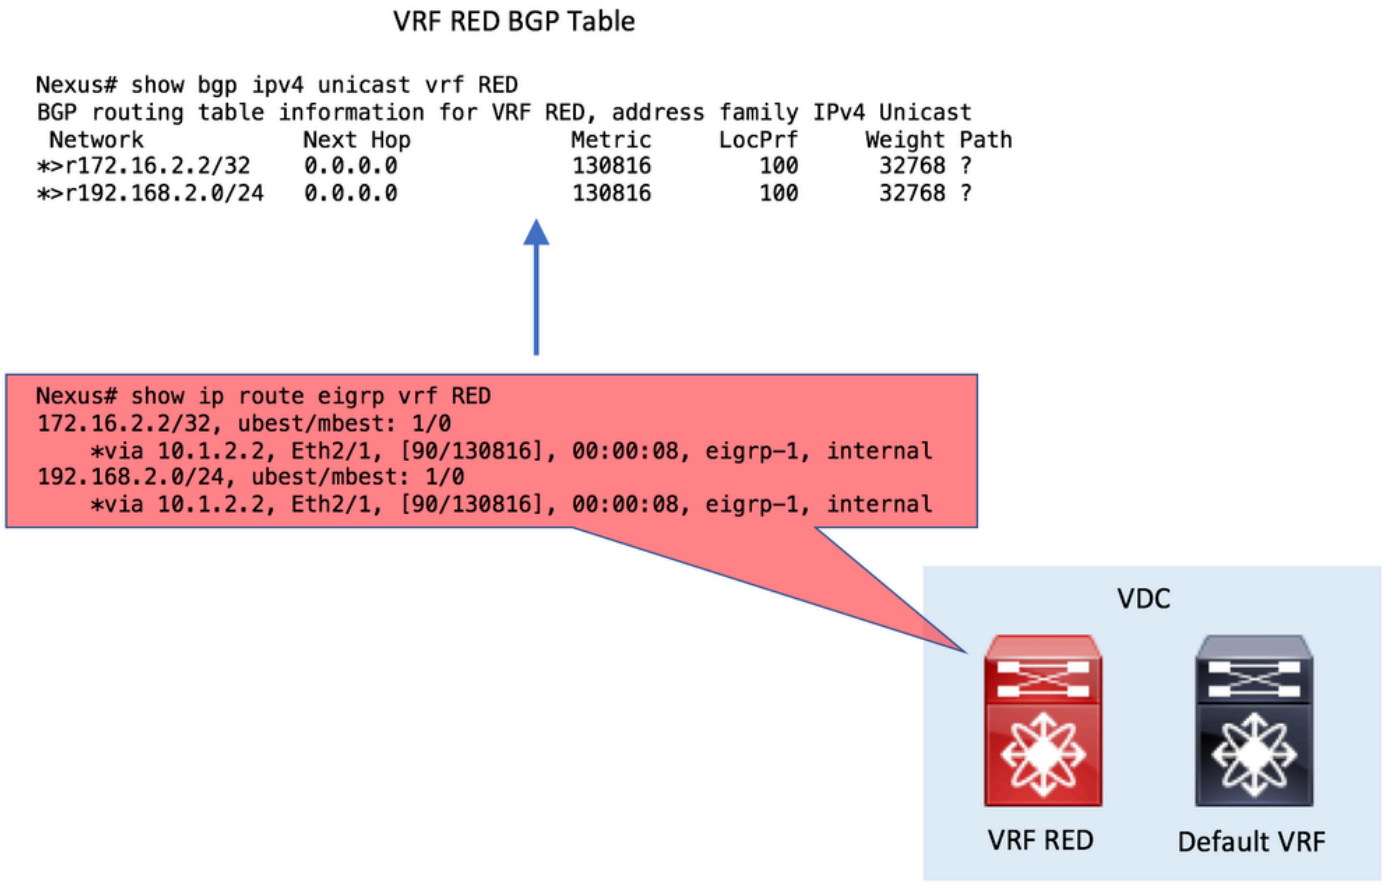



- Step 1. Redistribute into BGP.

Redistribute the routes which exist in VRF RED Routing Table in BGP. Since the routes are in the VRF RED, the **redistribute** command in BGP goes under the vrf RED address-family ipv4 unicast section.

**Redistribute into BGP**

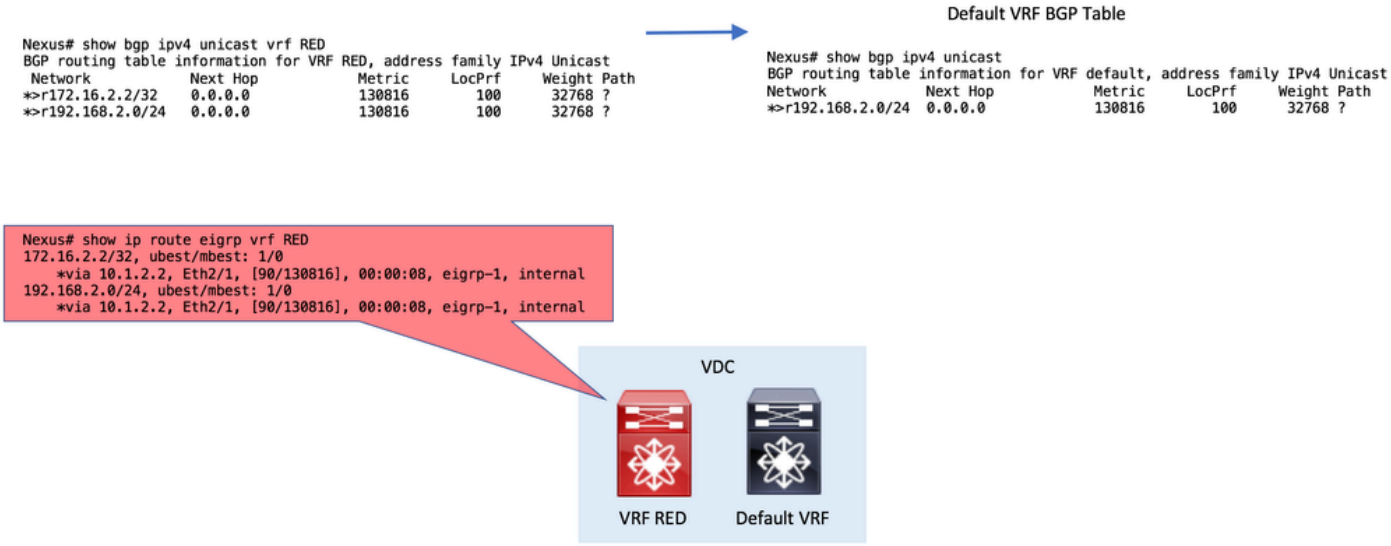
```
route-map ALL permit 10
!
router bgp 65535
  vrf RED
    address-family ipv4 unicast
      redistribute eigrp 1 route-map ALL
```




- Step 2. Configure Export VRF default in the origin VRF.

The **export vrf default** command is configured in the origin VRF. The command line requires a route-map as a parameter in order to explicitly define the routes to be exported in the Default VRF.

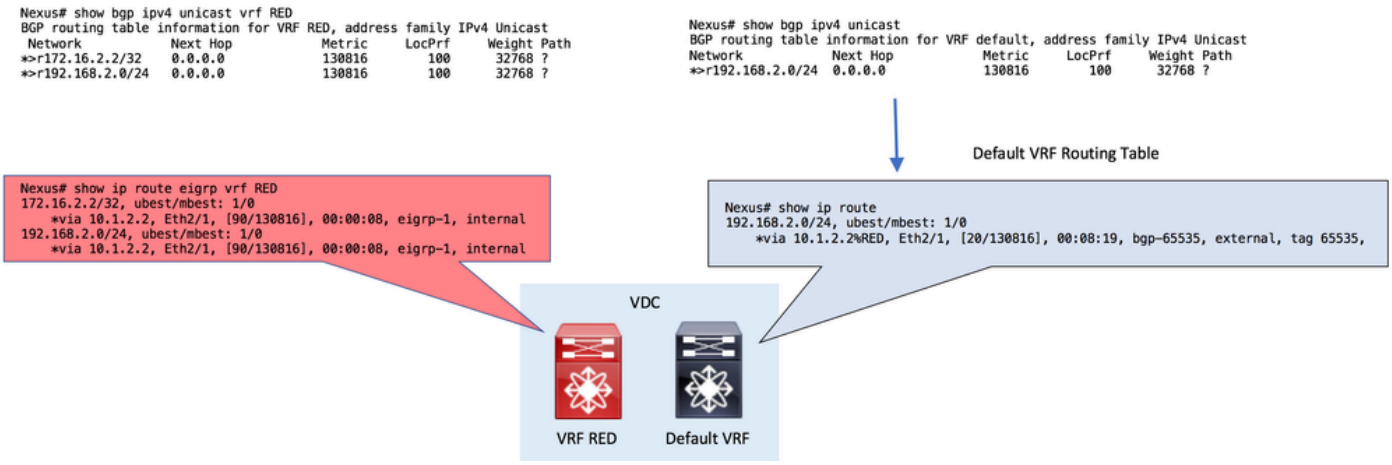
| **Configure Export VRF Default in the Origin VRF** |
| --- |
| ip prefix-list NETWORK seq 5 permit 192.168.2.0/24<br>!<br>route-map GLOBAL-TO-VRF permit 10<br>match ip address prefix-list NETWORK<br>!<br>vrf context RED<br>address-family ipv4 unicast<br>export vrf default map GLOBAL-TO-VRF |




- Step 3. Check the Default VRF Routing Table.

You can confirm in the Default VRF that the routes are now seen via BGP. These BGP routes in the Default VRF can now be redistributed in any other Routing Protocol that also runs in the Default VRF.

| **Check the Default VRF Routing Table** |
| --- |
| <#root><br><br>Nexus#<br><br>**show ip route**<br><br><br>IP Route Table for VRF "default"<br>'*' denotes best ucast next-hop<br>'**' denotes best mcast next-hop<br>'[x/y]' denotes [preference/metric]<br>'%<string>' in via output denotes VRF <string><br><br>192.168.2.0/24, ubest/mbest: 1/0<br>*via 10.1.2.2%RED, Eth2/1, [20/130816], 00:08:19, bgp-65535, external, tag 65535,<br>Nexus# |




# Verify

There are 4 phases in the VRF route leak process. Verification can be done in order:




In order to check the routes are correctly in the Routing Table, the command is:

<#root>

**show ip route [vrf <vrf name>]**

In order to check the routes are correctly in the BGP table, the commands are:

Notice that the second command can be used interchangeably in order to show IPv4 Unicast addresses in the BGP table.

<#root>

**show bgp ipv4 unicast [vrf <vrf name>]**

**show ip bgp [vrf <vrf name>]**

Lastly, the **show forwarding route A.B.C.D/LEN [VRF <vrf name>]** can be used in order to confirm the Layer 3 - route programmed at Line Card level (Hardware Programming).

```
<#root>

Nexus#

show forwarding route 10.1.2.2


slot  1
=======


IPv4 routes for table default/base

'*' denotes recursive route
----------------+-----------------------------------------+----------------------+-----------------
Prefix          | Next-hop                                | Interface            | Labels
----------------+-----------------------------------------+----------------------+-----------------
10.1.2.0/24       Attached                                  Ethernet2/1
Nexus#
```

# Related Information

[Configure dampen-igp-metric in BGP on Cisco Nexus Switches]