



Configuring Tenant Routed Multicast, Release 12.1.3

Table of Contents

New and Changed Information	1
Overview of Tenant Routed Multicast	2
Guidelines and Limitations	3
Overview of Tenant Routed Multicast with VXLAN EVPN Multi-Site	4
Tenant Routed Multicast with VXLAN EVPN Multi-Site Operations	5
Configuring TRM for Single Site Using Cisco Nexus Dashboard Fabric Controller	6
Configuring TRM for Multi-Site Using Cisco Nexus Dashboard Fabric Controller	8
Copyright	10

New and Changed Information

The following table provides an overview of the significant changes up to this current release. The table does not provide an exhaustive list of all changes or of the new features up to this release.

Release Version	Feature	Description
NDFC release 12.1.3	Reorganized content	Content within this document was originally provided in the <i>Cisco NDFC-Fabric Controller Configuration Guide</i> or the <i>Cisco NDFC-SAN Controller Configuration Guide</i> . Beginning with release 12.1.3, this content is now provided solely in this document and is no longer provided in those documents.

Overview of Tenant Routed Multicast

Tenant Routed Multicast (TRM) enables multicast forwarding on the VXLAN fabric that uses a BGP-based EVPN control plane. TRM provides multi-tenancy aware multicast forwarding between senders and receivers within the same or different subnet local or across VTEPs.

With TRM enabled, multicast forwarding in the underlay is leveraged to replicate VXLAN encapsulated routed multicast traffic. A Default Multicast Distribution Tree (Default-MDT) is built per-VRF. This is an addition to the existing multicast groups for Layer-2 VNI Broadcast, Unknown Unicast, and Layer-2 multicast replication group. The individual multicast group addresses in the overlay are mapped to the respective underlay multicast address for replication and transport. The advantage of using a BGP-based approach allows the VXLAN BGP EVPN fabric with TRM to operate as fully distributed Overlay Rendezvous-Point (RP), with the RP presence on every edge-device (VTEP).

A multicast-enabled data center fabric is typically part of an overall multicast network. Multicast sources, receivers, and multicast rendezvous points might reside inside the data center but also might be inside the campus or externally reachable via the WAN. TRM allows a seamless integration with existing multicast networks. It can leverage multicast rendezvous points external to the fabric. Furthermore, TRM allows for tenant-aware external connectivity using Layer-3 physical interfaces or subinterfaces.

Guidelines and Limitations

Refer to the following documents for switch-level guidelines and limitations for Tenant Routed Multicast:

- [Guidelines and Limitations for Tenant Routed Multicast](#)
- [Guidelines and Limitations for Layer 3 Tenant Routed Multicast](#)

Following are additional guidelines and limitations at the NDFC level:

- When you perform a brownfield import on a fabric where the VRFs and the networks are deployed with the **Enable Tenant Routed Multicast** option enabled (without using a configuration profile), if the **Overlay Mode** option for the fabric is set to **cli**, the VRF level **TRM Enable** option will not be enabled for VRFs imported into this fabric (in addition to the VRF level **RP Address**, **RP Loopback ID**, **Underlay Mcast Address**, and **Overlay Mcast Groups** options, which will also not be enabled in this case).

In addition, if this VRF is deployed on a new leaf switch, TRM will not be enabled on that leaf switch for that VRF.

Overview of Tenant Routed Multicast with VXLAN EVPN Multi-Site

Tenant Routed Multicast with Multi-Site enables multicast forwarding across multiple VXLAN EVPN fabrics connected via Multi-Site.

The following two use cases are supported:

- Use Case 1: TRM provides Layer 2 and Layer 3 multicast services across sites for sources and receivers across different sites.
- Use Case 2: Extending TRM functionality from VXLAN fabric to sources receivers external to the fabric.

TRM Multi-Site is an extension of BGP-based TRM solution that enables multiple TRM sites with multiple VTEPs to connect to each other to provide multicast services across sites in most efficient possible way. Each TRM site is operating independently and border gateway on each site allows stitching across each site. There can be multiple Border Gateways for each site. In a given site, the BGW peers with Route Server or BGWs of other sites to exchange EVPN and MVPN routes. On the BGW, BGP will import routes into the local VRF/L3VNI/L2VNI and then advertise those imported routes into the Fabric or WAN depending on where the routes were learnt from.

Tenant Routed Multicast with VXLAN EVPN Multi-Site Operations

The operations for TRM with VXLAN EVPN Multi-Site are as follows:

- Each Site is represented by Anycast VTEP BGWs. DF election across BGWs ensures no packet duplication.
- Traffic between Border Gateways uses ingress replication mechanism. Traffic is encapsulated with VXLAN header followed by IP header.
- Each Site will only receive one copy of the packet.
- Multicast source and receiver information across sites is propagated by BGP protocol on the Border Gateways configured with TRM.
- BGW on each site receives the multicast packet and re-encapsulate the packet before sending it to the local site.

For information about guidelines and limitations for TRM with VXLAN EVPN Multi-Site, see [Configuring Tenant Routed Multicast](#).

Configuring TRM for Single Site Using Cisco Nexus Dashboard Fabric Controller

This section assumes that a VXLAN EVPN fabric has already been provisioned using Cisco Nexus Dashboard Fabric Controller.

Perform the following steps to enable TRM for a single site.

1. Enable TRM for the selected Easy Fabric as follows:

- a. If the fabric template is **Data Center VXLAN EVPN**, from the Fabric Overview **Actions** drop-down list, choose **Edit Fabric**.
- b. Click the **Replication** tab and configure the fields on the tab as follows:
 - **Enable Tenant Routed Multicast**: Select the check box to enable Tenant Routed Multicast (TRM) that allows overlay multicast traffic to be supported over EVPN/MVPN in the VXLAN BGP EVPN fabric.
 - **Default MDT Address for TRM VRFs**: When you select the **Enable Tenant Routed Multicast (TRM)** check box, the multicast address for Tenant Routed Multicast traffic is auto populated. By default, this address is from the IP prefix specified in the **Multicast Group Subnet** field. When you update either field, ensure that the TRM address is chosen from the IP prefix specified in **Multicast Group Subnet**.
- c. Click **Save** to save the fabric settings.

At this point, all the switches turn "Blue" as it will be in the pending state.

- d. From the Fabric Overview **Actions** drop-down list, choose **Recalculate Config** and then choose **Deploy Config** to enable the following:
 - Enable feature ngmvpn: Enables the Next-Generation Multicast VPN (ngMVPN) control plane for BGP peering.
 - Configure ip multicast multipath s-g-hash next-hop-based: Multipath hashing algorithm for the TRM enabled VRFs.
 - Configure ip igmp snooping vxlan: Enables IGMP Snooping for VXLAN VLANs.
 - Configure ip multicast overlay-spt-only: Enables the MVPN Route-Type 5 on all MPVN enabled Cisco Nexus 9000 switches.
 - Configure and Establish MVPN BGP AFI Peering: This is necessary for the peering between BGP RR and the leaf nodes.

For VXLAN EVPN fabric created using **BGP Fabric** template, **Enable Tenant Routed Multicast (TRM)** and **Default MDT Address for TRM VRFs** fields can be found on the **EVPN** tab

2. Enable TRM for the VRF as follows:

- a. Navigate to **Fabric Overview > VRFs** and edit the selected VRF.
- b. Navigate to the **Advanced** tab and edit the following TRM settings:
 - **TRM Enable** - Select the check box to enable TRM. If you enable TRM, then the RP address and the underlay multicast address must be entered.
 - **Is RP External** - Enable this check box if the RP is external to the fabric. If this field is

unchecked, RP is distributed in every VTEP.



If the RP is external, then select the appropriate option. If the RP is external, then RP loopback ID is greyed out.

- **RP Address** - Specifies the IP address of the RP.
 - **RP Loopback ID** - Specifies the loopback ID of the RP, if **Is RP External** is not enabled.
 - **Underlay Mcast Address** - Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.
 - **Overlay Mcast Groups** - Specifies the multicast group subnet for the specified RP. The value is the group range in "ip pim rp-address" command. If the field is empty, 224.0.0.0/24 is used as default.
- c. Click **Save** to save the settings. The switches go into the pending state, that is, blue color. These settings enable the following:
- Enable PIM on L3VNI SVI.
 - Route-Target Import and Export for MVPN AFI.
 - RP and other multicast configuration for the VRF.
 - Loopback interface using the above RP address and RP loopback id for the distributed RP.
3. Enable TRM for the network as follows:
- a. On the **Fabric Overview** window, go to the **Networks** tab.
 - b. Edit the selected network and go to the **Advanced** tab.
 - c. Check the **TRM Enable** check box to enable TRM.
 - d. Click **Save** to save the settings. The switches go into the pending state, that is, the blue color. The TRM settings enable the following:
 - Enable PIM on the L2VNI SVI.
 - Create a PIM policy **none** to avoid PIM neighborship with PIM Routers within a VLAN. The **none** keyword is a configured route map to deny any ipv4 addresses to avoid establishing PIM neighborship policy using anycast IP.

Configuring TRM for Multi-Site Using Cisco Nexus Dashboard Fabric Controller

This section assumes that a Multi-Site Domain (MSD) has already been deployed by Cisco Nexus Dashboard Fabric Controller and TRM needs to be enabled.

Perform the following steps to enable TRM for a multi-site.

1. Enable TRM on the BGWs.

a. Navigate to **Fabric Overview > VRFs**. Make sure that the right DC Fabric is selected under the **Scope** and edit the VRF.

b. Navigate to the **Advanced** tab. Edit the TRM settings.

- **TRM Enable** - Select the check box to enable TRM. If you enable TRM, then the RP address and the underlay multicast address must be entered.
- **Is RP External** - Enable this check box if the RP is external to the fabric. If this field is unchecked, RP is distributed in every VTEP.



If the RP is external, then select the appropriate option. If the RP is external, then RP loopback ID is greyed out.

- **RP Address** - Specifies the IP address of the RP.
- **RP Loopback ID** - Specifies the loopback ID of the RP, if **Is RP External** is not enabled.
- **Underlay Mcast Address** - Specifies the multicast address associated with the VRF. The multicast address is used for transporting multicast traffic in the fabric underlay.
- **Overlay Mcast Groups** - Specifies the multicast group subnet for the specified RP. The value is the group range in "ip pim rp-address" command. If the field is empty, 224.0.0.0/24 is used as default.
- **Enable TRM BGW MSite** - Select the check box to enable TRM on Border Gateway Multi-Site.

c. Click **Save**.

The switches go into the pending state, that is, blue color. These settings enable the following:

- Enable feature ngmvpn: Enables the Next-Generation Multicast VPN (ngMVPN) control plane for BGP peering.
- Enables PIM on L3VNI SVI.
- Configures L3VNI Multicast Address.
- Route-Target Import and Export for MVPN AFI.
- RP and other multicast configuration for the VRF.
- Loopback interface for the distributed RP.
- Enable Multi-Site BUM ingress replication method for extending the Layer 2 VNI

2. Establish MVPN AFI between the BGWs, as follows:

- a. Double-click the MSD fabric to open the **Fabric Overview** window.
- b. Choose **Links**. Filter it by the policy - **Overlays**.
- c. Select and edit each overlay peering to enable TRM by checking the **Enable TRM** check box.
- a. Click **Save** to save the settings.

The switches go into the pending state, that is, the blue color. The TRM settings enable the MVPN peering's between the BGWs, or BGWs and Route Server.

Copyright

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <http://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. (1110R)

© 2017–2024 Cisco Systems, Inc. All rights reserved.