



Hewlett Packard
Enterprise

HPE FlexNetwork 10500 Switch Series

TRILL Configuration Guide

Part number: 5998-7136R
Software version: 10500-CMW710-R7178
Document version: 6W100-20160129

© Copyright 2016 Hewlett Packard Enterprise Development LP

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

Intel®, Itanium®, Pentium®, Intel Inside®, and the Intel Inside logo are trademarks of Intel Corporation in the United States and other countries.

Microsoft® and Windows® are trademarks of the Microsoft group of companies.

Adobe® and Acrobat® are trademarks of Adobe Systems Incorporated.

Java and Oracle are registered trademarks of Oracle and/or its affiliates.

UNIX® is a registered trademark of The Open Group.

Contents

Configuring TRILL	1
Overview	1
Basic concepts	1
TRILL frame formats	1
How TRILL works	2
TRILL forwarding mechanisms	3
Protocols and standards	5
Configuration restrictions and guidelines	5
Hardware requirements	5
Feature compatibility requirements	6
TRILL configuration task list	6
Enabling TRILL	7
Configuration restrictions and guidelines	7
Configuration procedure	7
Configuring the system ID and nickname for an RB	7
Configuring the link type of a TRILL port	8
Configuring the DRB priority of a TRILL port	8
Configuring the link cost for a TRILL port	9
Configuring announcing VLANs and the designated VLAN	9
Configuring TRILL timers	10
Configuring TRILL LSP parameters and features	11
Setting LSP parameters	11
Enabling TRILL LSP fast advertisement	12
Enabling TRILL pseudonode bypass	13
Setting the SPF algorithm parameters	13
Configuring TRILL distribution trees	13
Setting basic distribution tree parameters	13
Enabling TRILL distribution tree multithread calculation	14
Enabling load balancing over TRILL distribution trees	14
Configuring TRILL ECMP routes	15
Enabling incremental flush for TRILL multicast routing entries	15
Enabling logging of TRILL neighbor changes	16
Configuring SNMP for TRILL	16
Configuring TRILL GR	17
Associating a TRILL port with a track entry	17
Displaying and maintaining TRILL	18
TRILL configuration example	19
Network requirements	19
Configuration procedure	19
Verifying the configuration	21
Document conventions and icons	23
Conventions	23
Network topology icons	24
Support and other resources	25
Accessing Hewlett Packard Enterprise Support	25
Accessing updates	25
Websites	26
Customer self repair	26
Remote support	26
Documentation feedback	26
Index	28

Configuring TRILL

TRansparent Interconnection of Lots of Links (TRILL) uses IS-IS to provide transparent Layer 2 forwarding.

Overview

TRILL combines the simplicity and flexibility of Layer 2 switching with the stability, scalability, and rapid convergence capability of Layer 3 routing. All these advantages make TRILL very suitable for large flat Layer 2 networks in data centers.

Basic concepts

- **RBridge**—Routing bridge (RB) that runs TRILL. RBs are classified into ingress RBs, transit RBs, and egress RBs, depending on their positions in the TRILL network. A frame enters the TRILL network through an ingress RB, travels along transit RBs, and leaves the TRILL network through an egress RB, as shown in [Figure 2](#).
- **TRILL network**—A Layer 2 network that contains RBs, as shown in [Figure 3](#).
- **System ID**—Unique identifier of an RB in the TRILL network. The system ID is 6 bytes in length.
- **Nickname**—Address of an RB in the TRILL network. The nickname is 2 bytes in length.
- **Link State Database**—The LSDB contains all link state information in the TRILL network.
- **Link State Protocol Data Unit**—An LSP describes local link state information and is advertised between neighbor devices.
- **Designated Routing Bridge (DRB)**—Similar to the designated IS (DIS) in IS-IS, a DRB exists in a broadcast network. It helps simplify network topology, and assigns AVFs and appointed ports for the VLANs on each RB in the broadcast network.
- **Appointed VLAN-x Forwarder (AVF) and appointed port**—To avoid loops, TRILL requires all traffic of a VLAN on a broadcast network to enter and leave the TRILL network through the same port of an RB. The RB is the VLAN's AVF, and the port is the VLAN's appointed port.

For more information about LSDB, LSPDU, and DIS, see *Layer 3—IP Routing Configuration Guide*.

TRILL frame formats

TRILL frames include control frames and data frames.

TRILL control frames include TRILL Hello, LSP, CSNP, PSNP, MTU-probe, and MTU-ack. These control frames use 802.1Q encapsulation and have a fixed destination multicast address 0180-C200-0041.

TRILL data frames have a specific format, as shown in [Figure 1](#). A TRILL header and an outer Ethernet header are added to the original Ethernet frame.

Figure 1 TRILL data frame format

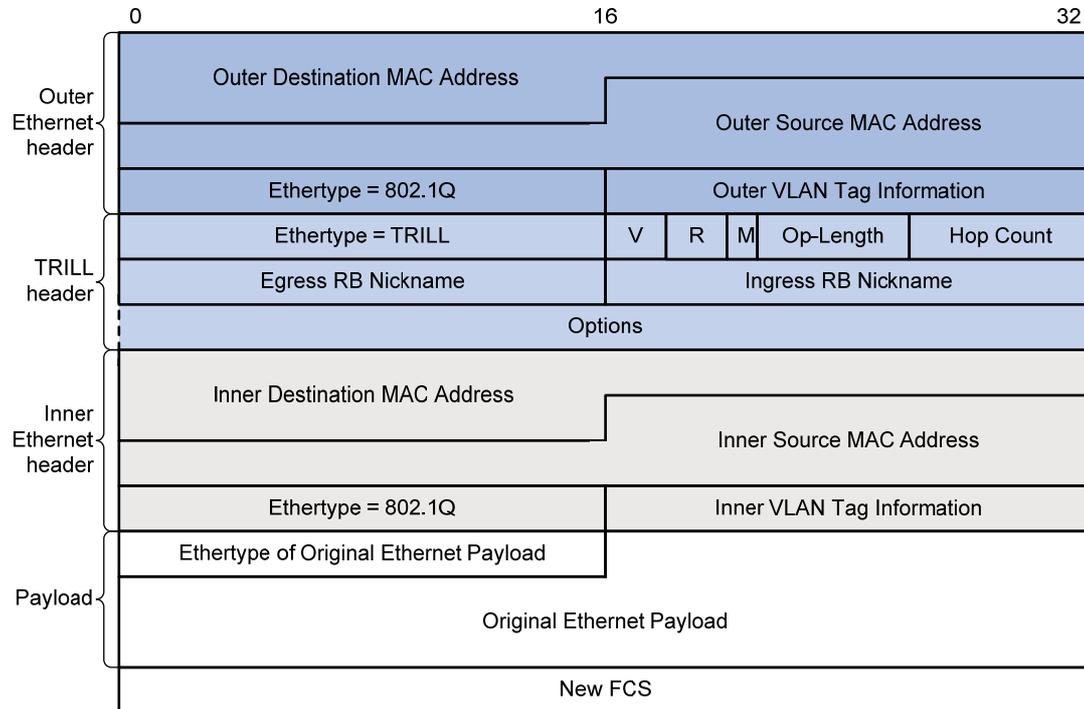


Table 1 describes the fields in the TRILL header.

Table 1 TRILL header fields

Field	Description
Ethertype	The Ethertype is fixed to TRILL.
V	Version number, which is 0. When an RB receives a TRILL frame, it checks the V field and drops the frame if the V field is not 0.
R	Reserved for future extension. An ingress RB sets the R field to 0 when adding a TRILL header. Transit RBs and egress RBs ignore the field.
M	Multidestination attribute: <ul style="list-style-type: none"> 0—Known unicast frame. 1—Multidestination frame (multicast, broadcast, or unknown unicast frame).
Op-Length	Length of the Options field. 0 indicates that the Options field does not exist.
Hop Count	Hop count, which is used to avoid loops. An RB drops a TRILL frame whose hop count is decremented to 0.
Egress RB Nickname	Nickname of the egress RB.
Ingress RB Nickname	Nickname of the ingress RB.
Options	Options field. This field exists when the Op-Length field is non-zero.

How TRILL works

TRILL establishes and maintains adjacencies between RBs by periodically advertising Hello frames, distributes LSPs among RB neighbors, and generates an LSDB for all RBs in the network. Based on the LSDB, each RB uses the SPF algorithm to calculate forwarding entries destined to other RBs.

TRILL forwarding mechanisms

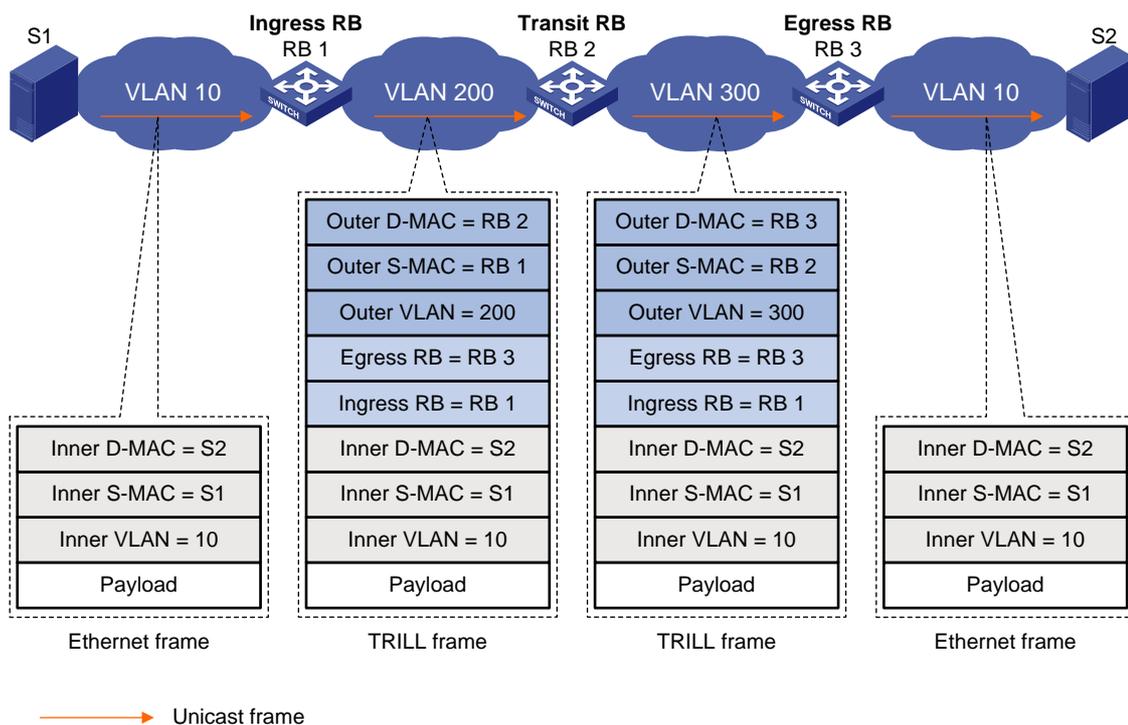
Different types of frames are forwarded by using different forwarding mechanisms. The following sections describe these mechanisms.

Unicast frame forwarding mechanism

As shown in Figure 2, a unicast frame is forwarded as follows:

1. When a unicast frame enters the TRILL network, the ingress RB encapsulates the original Ethernet frame with the following headers:
 - A TRILL header (similar to an IP header).
 - An outer Ethernet header (similar to the Ethernet header of a regular Ethernet frame).
2. RBs forward the frame hop by hop according to the egress RB nickname in the TRILL header in the same way routers forward IP packets. Each hop replaces the outer Ethernet header with an appropriate outer Ethernet header, and decrements the hop count in the TRILL header.
3. Upon receiving the TRILL frame, the egress RB de-encapsulates it to obtain the original Ethernet frame, and sends the frame to the target device.

Figure 2 Unicast frame forwarding flow



The outer Ethernet header enables traditional Ethernet switches to forward TRILL frames and connect RBs.

Multidestination frame forwarding mechanism

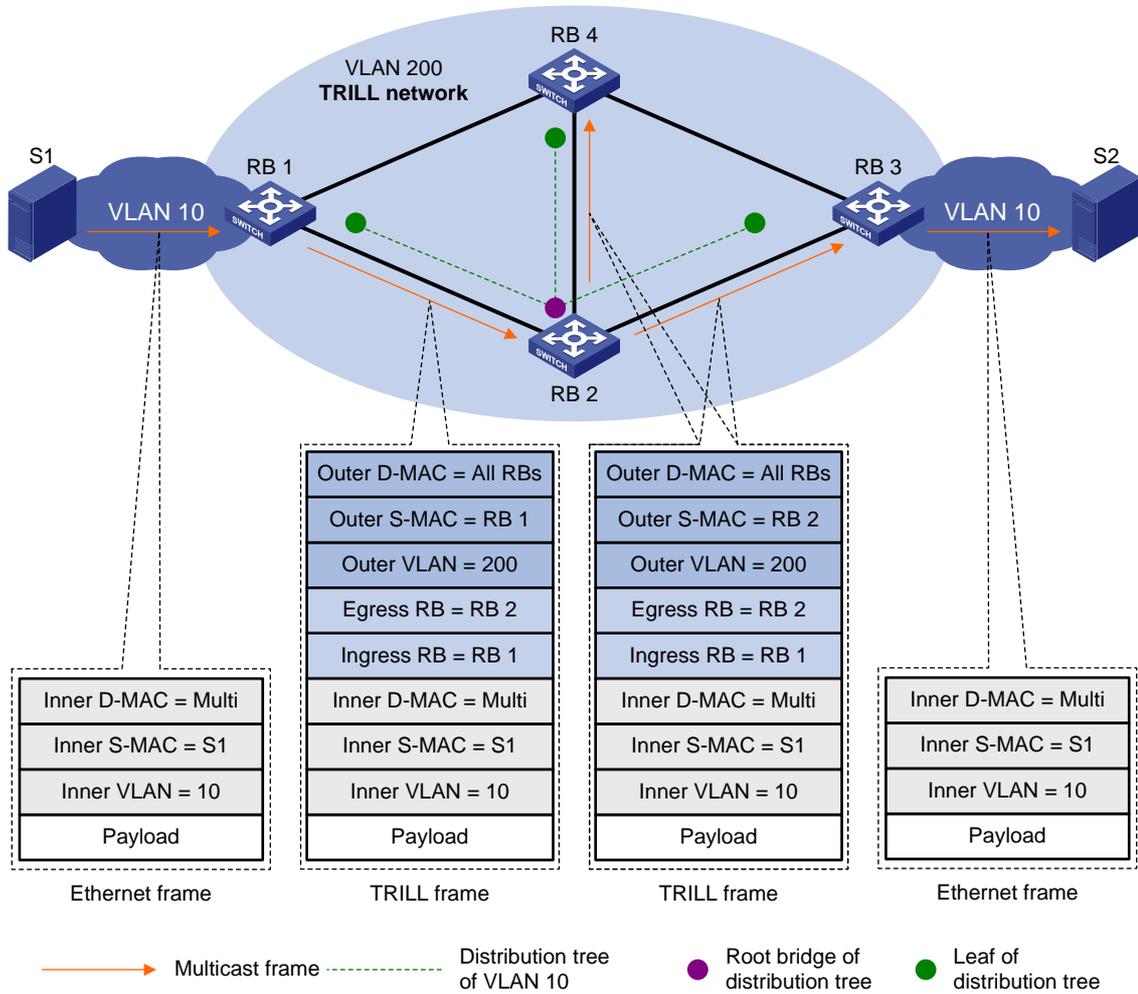
In a TRILL network, RBs perform the following tasks:

- Compute a TRILL distribution tree for each VLAN based on the LSDB.
- Guide the forwarding of multidestination frames in each VLAN by using its TRILL distribution tree.

As shown in Figure 3, when a multicast frame from VLAN 10 enters the TRILL network, RB 1, which is an ingress RB, encapsulates the multicast frame into a TRILL frame. In the frame, the egress RB is RB 2, the root bridge of the TRILL distribution tree for VLAN 10. When the frame arrives at the root

bridge, it is distributed throughout the TRILL distribution tree. Then, the TRILL frame is de-encapsulated by RB 3 and sent to the destination station S2. Because the network segment where RB 4 resides does not have a receiver of this frame, RB 4 drops the frame.

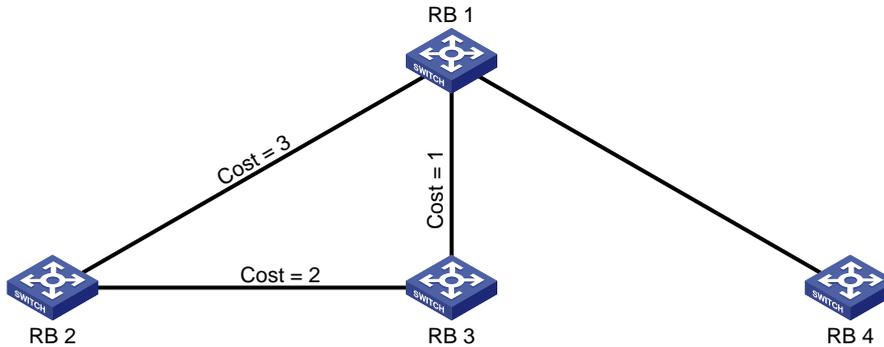
Figure 3 Multicast frame forwarding flow



TRILL selects distribution trees for forwarding multidestination frames based on the VLANs to which the frames belong. Because the topologies of TRILL distribution trees are different, traffic can be load shared. However, equal-cost links are not used for load sharing.

When N equal-cost links exist in the network, each TRILL distribution tree selects the link with the largest pseudonode ID for forwarding packets. As shown in Figure 4, two equal-cost links exist between RB 1 and RB 3. Assume the link directly connecting RB 1 to RB 3 has the largest pseudonode ID. Both the TRILL distribution tree rooted at RB 3 and the TRILL distribution tree rooted at RB 4 select the link. For more information about pseudonode IDs, see *Layer 3—IP Routing Configuration Guide*.

Figure 4 Multicast ECMP



TRILL distribution trees support Equal Cost Multiple Path (ECMP), also known as multicast ECMP. When multicast ECMP is enabled, TRILL assigns equal-cost links to multiple TRILL distributions trees. This improves the load sharing performance.

When N equal-cost links exist in the network, each TRILL distribution tree selects an equal-cost link for forwarding packets through $J \bmod N$ in root bridge priority order. J is the priority sequence number of a TRILL distribution tree and starts from 0.

As shown in [Figure 4](#):

- The link directly connecting RB 1 to RB 3 is assigned to the TRILL distribution tree rooted at RB 3.
- The link RB 1-RB 2-RB 3 is assigned to the TRILL distribution tree rooted at RB 4.

TRILL distribution trees support fast root switching. When an RB detects that the root of a distribution tree is unreachable, the RB deletes the LSP of the root from its LSDB. This triggers recalculation of all distribution trees in the TRILL network. Multidestination traffic is switched to new distribution trees.

Protocols and standards

- RFC 6325, *Routing Bridges (RBridges): Base Protocol Specification*
- RFC 6326, *Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS*
- RFC 6327, *Routing Bridges (RBridges): Adjacency*
- RFC 1195, *Use of OSI IS-IS for Routing in TCP/IP and Dual Environments*

Configuration restrictions and guidelines

When you configure TRILL, follow the restrictions and guidelines in this section.

Hardware requirements

TRILL is supported only on the following interface modules:

- SE interface modules:
 - LSU1CGC2SE0 (JG916A).
 - LSUM2GP44TSSE0 (JH191A, JH199A).
 - LSUM2GT48SE0 (JH192A, JH200A).
- SF interface modules.
- SG interface modules.
- EC interface modules.

Feature compatibility requirements

When you use TRILL with other features, follow these restrictions and guidelines:

- Do not enable both TRILL and EVB on a port. The allowed VLAN list of a TRILL-enabled port cannot overlap with that of an EVB-enabled port. For more information about EVB, see *EVB Configuration Guide*.
- Do not configure EVI and VXLAN in a TRILL network. For information about EVI, see *EVI Configuration Guide*. For information about VXLAN, see *VXLAN Configuration Guide*.
- Do not assign a TRILL port to a VLAN that is associated with an OpenFlow instance. If the TRILL port is assigned to the VLAN, traffic might be forwarded incorrectly because forwarding in the VLAN is based on flow table entries.
- Configuration in Layer 2 Ethernet interface view takes effect only on the current port. Configuration in Layer 2 aggregate interface view takes effect on the current interface and its member ports. Configuration on the member port of an aggregate interface takes effect after the member port leaves the aggregation group. On a Layer 2 aggregate interface, TRILL distribution tree calculation is done on the aggregate interface rather than its member ports.
- To connect a spanning tree network to a TRILL network, make sure the following requirements are met:
 - The spanning tree protocol is disabled on the TRILL network.
 - An edge port is used to connect the spanning tree network to the TRILL network. The edge port can transit to the forwarding state before DRB election is finished. This prevents multiple DRBs from being elected.

For more information about spanning tree protocols, see *Layer 2—LAN Switching Configuration Guide*.

- As a best practice, do not enable loop detection on TRILL ports, because TRILL avoids loops. For more information about loopback detection, see *Layer 2—LAN Switching Configuration Guide*.
- If IRF is used, retain the IRF bridge MAC address permanently. Otherwise, traffic interruption might occur after an IRF split. For more information about IRF configuration, see *IRF Configuration Guide*.
- Do not configure the TPID for VLAN tags on RBs. For more information about TPIDs, see *Layer 2—LAN Switching Configuration Guide*.
- To avoid loops, do not connect multiple TRILL ports of an RB to a broadcast network, because TRILL ports cannot detect each other when they are on a broadcast network. If there are multiple TRILL ports, they might be elected as the appointed ports of a VLAN at the same time and result in loops.

TRILL configuration task list

Tasks at a glance
(Required.) Enabling TRILL
(Optional.) Configuring the system ID and nickname for an RB
(Optional.) Configuring the link type of a TRILL port
(Optional.) Configuring the DRB priority of a TRILL port
(Optional.) Configuring the link cost for a TRILL port
(Optional.) Configuring announcing VLANs and the designated VLAN
(Optional.) Configuring TRILL timers

(Optional.) Configuring TRILL LSP parameters and features
(Optional.) Setting the SPF algorithm parameters
(Optional.) Configuring TRILL distribution trees
(Optional.) Configuring TRILL ECMP routes
(Optional.) Enabling incremental flush for TRILL multicast routing entries
(Optional.) Enabling logging of TRILL neighbor changes
(Optional.) Configuring SNMP for TRILL
(Optional.) Configuring TRILL GR
(Optional.) Associating a TRILL port with a track entry

Enabling TRILL

After you enable TRILL on a port, TRILL can operate correctly by using default settings. A port with TRILL enabled is called a TRILL port.

Configuration restrictions and guidelines

To enable TRILL on a port, first enable TRILL globally.

Configuration procedure

To enable TRILL:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable TRILL globally and enter TRILL view.	trill	By default, TRILL is disabled globally.
3. Return to system view.	quit	N/A
4. Enter Layer 2 Ethernet or aggregate interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
5. Enable TRILL on the port.	trill enable	By default, TRILL is disabled on a port.

Configuring the system ID and nickname for an RB

The system ID and nickname of an RB are identifiers of the RB in the TRILL network.

- **System ID**—Unique identifier of an RB in the TRILL network. The system ID can be automatically assigned or manually configured.
- **Nickname**—Address of an RB in the TRILL network. The address can be automatically assigned or manually configured. When multiple RBs in the TRILL network have the same nickname, the RB with the highest priority uses the nickname. When the RBs also have the

same priority, the RB with the highest system ID uses the nickname. The system automatically assigns new nicknames to the other RBs.

The system resets the TRILL process when the RB's system ID changes.

To configure the system ID and nickname for the RB:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Configure the system ID for the RB.	system-id <i>system-id</i>	By default, the RB automatically generates a system ID based on its MAC address upon startup.
4. Configure the nickname for the RB.	nickname <i>nickname</i> [priority <i>priority</i>]	By default, TRILL automatically assigns nicknames to RBs, and the priority for a RB to hold a nickname is 64.

Configuring the link type of a TRILL port

The following link types are available for a TRILL port:

- **Access**—Access ports include access ports without the alone attribute and access ports with the alone attribute. Access ports with the alone attribute do not send or receive Hello frames and do not participate in DRB election or AVF negotiation. Access ports without the alone attribute can process only local data frames and Hello frames.
- **Hybrid**—A hybrid port combines the attributes of an access port and a trunk port, and can process local data frames and passing data frames.
- **Trunk**—A trunk port can process passing data frames and some of Layer 2 protocol frames (for example, LLDP frames), but it cannot process local data frames.

To configure the link type of a TRILL port:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Layer 2 Ethernet or aggregate interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Configure the link type of the TRILL port.	trill link-type { access [alone] hybrid trunk }	By default, the link type of a TRILL port is access without the alone attribute.

Configuring the DRB priority of a TRILL port

On a broadcast network, TRILL must elect a DRB. An RB with a higher DRB priority is preferred in DRB election. When two RBs have the same DRB priority, the RB with a higher MAC address takes precedence.

To configure the DRB priority of a TRILL port:

Step	Command	Remarks
1. Enter system view.	system-view	N/A

2. Enter Layer 2 Ethernet or aggregate interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Configure the DRB priority of the TRILL port.	trill drb-priority <i>priority</i>	By default, the DRB priority of a TRILL port is 64.

Configuring the link cost for a TRILL port

The link cost of a TRILL port can be automatically calculated by the system or manually configured.

- A manually configured link cost takes precedence over a calculated link cost.
- If no configured link cost exists and automatic link cost calculation is enabled, the calculated link cost takes effect.
- If no configured link cost exists and automatic link cost calculation is disabled, the default link cost of 2000 is used.

The system automatically calculates the link cost of a TRILL port by using the following formula: link cost = 20000000000000/interface baud rate.

To configure the link cost for a TRILL port:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Enable automatic link cost calculation for TRILL ports.	auto-cost enable	By default, automatic link cost calculation is enabled for TRILL ports.
4. Return to system view.	quit	N/A
5. Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
6. Configure the link cost for the TRILL port.	trill cost <i>value</i>	The default setting is 2000.

Configuring announcing VLANs and the designated VLAN

The concepts and symbols used to describe a VLAN on a port are as follows:

- **Enabled VLAN**—A VLAN enabled on the port.
- **Forwarding VLAN**—A VLAN for which the port is the appointed port.
- \cap and \cup —Set operation symbols. \cap indicates set-theoretic intersection, and \cup indicates set-theoretic union.

RBs send Hello frames in a set of VLANs. The VLAN set is calculated as follows:

- **DRB**—Enabled VLANs \cap (announcing VLANs \cup designated VLAN).
- **Non-DRB**—Enabled VLANs \cap (designated VLAN \cup (announcing VLANs \cap forwarding VLANs)).

To prevent Hello frames from consuming excessive CPU resources, reduce the number of announcing VLANs.

RBs use the designated VLAN to forward TRILL protocol frames (except Hello frames) and local data frames. For RBs to establish adjacencies and forward TRILL data frames, make sure the designated VLAN is an enabled VLAN.

To configure announcing VLANs and the designated VLAN:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Configure announcing VLANs.	trill announcing-vlan { <i>vlan-list</i> null }	By default, no announcing VLANs are configured, and announcing VLANs are enabled VLANs.
4. Configure the designated VLAN.	trill designated-vlan <i>vlan-id</i>	By default, no designated VLAN is configured.

Configuring TRILL timers

You can configure the following TRILL timers:

- **Hello interval** and **Hello multiplier**—The RB advertises Hello frames at the Hello interval to maintain a TRILL adjacency. The shorter the Hello interval, the faster the network convergence. However, a shorter Hello interval consumes more system resources. The adjacency holding time is obtained by multiplying the Hello interval by the Hello multiplier. The RB advertises the adjacency holding time to neighbors through Hello frames. If a neighbor does not receive any Hello frame from the RB within the adjacency holding time, it removes the TRILL adjacency with the RB.
- **AVF inhibition time**—As the AVF of a VLAN, the RB guarantees that frames of the VLAN enter and leave a broadcast network through the same port. Other RBs on the broadcast network do not process frames from the VLAN.

To avoid loops, the RB suppresses its AVF role during the inhibition time when one of the following conditions exists:

- The RB detects a root bridge change on the broadcast network.
- Other RBs advertise a different AVF for the VLAN.

When the inhibition time expires, the RB restores its AVF role if it is still the AVF of the VLAN.

- **CSNP interval**—On a broadcast network, the RB advertises CSNPs at the CSNP interval to perform network-wide LSDB synchronization if it is elected as the DRB. A CSNP records all LSP digests of the RB's local LSDB. A remote RB compares a received CSNP against its local LSDB to verify whether some LSPs are aged out or missing. If the CSNP has an LSP digest that the local LSDB does not have, the remote RB sends a PSNP packet to request the LSP.

To configure TRILL timers:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Layer 2 Ethernet or aggregate interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Configure the Hello	trill timer hello <i>interval</i>	The default setting is 10 seconds.

interval.		This command sets the Hello interval for an RB. The Hello interval of a DRB is 1/3 of the Hello interval of an RB. This allows for DRB failures to be quickly detected.
4. Configure the Hello multiplier.	trill timer holding-multiplier count	The default setting is 3.
5. Configure the AVF inhibition time.	trill timer avf-inhibited time	The default setting is 30 seconds.
6. Configure the CSNP interval.	trill timer csnp interval	The default setting is 10 seconds.

Configuring TRILL LSP parameters and features

Setting LSP parameters

You can modify the following LSP parameters:

- **LSP maximum age**—The RB uses the LSP maximum age as the remaining lifetime of the LSPs that it originates. When the RB detects that the remaining lifetime of an LSP reaches 0 seconds in the LSDB, the RB performs the following tasks:
 - Removes the LSP's content.
 - Keeps the LSP's digest.
 - Sets the LSP's remaining lifetime to 0 and purges the LSP from the network by advertising the LSP to other RBs.
- **LSP refresh interval**—A locally originated LSP is forcibly refreshed when its remaining lifetime is no greater than n : $n = \text{LSP maximum age} - \text{LSP refresh interval}$. This mechanism avoids frequent LSP aging and ensures network stability.
- **LSP generation timer parameters**—By adjusting the TRILL LSP generation timer parameters, you can prevent frequent network changes from consuming excessive bandwidth and device resources.

When the network is stable, the LSP generation timer is set to the minimum interval for each LSP generation. When the network is unstable, the LSP generation timer is added by the incremental interval for each LSP generation until the maximum interval is reached.

- **Maximum length of originated LSPs**—The RB selects the smallest value from the following values as the actual maximum length of LSPs to be sent to a neighbor:
 - The configured maximum length of originated LSPs.
 - The interface MTU.
 - The maximum originated LSP length carried in the LSPs sent by the neighbor.
- **Maximum length of received LSPs**—When the RB receives an LSP that exceeds the length, the RB drops the LSP.
- **Overload bit of LSPs**—The RB sets the Overload bit in LSPs if the RB fails and cannot correctly perform route selection and packet forwarding. When the RB cannot record the complete LSDB because of insufficient memory, routing calculation errors occur. To make troubleshooting easier, temporarily exclude the RB from the TRILL network by setting the Overload bit for the LSPs sent by the RB.
- **Minimum LSP interval and maximum number of LSPs transmitted per interval**—To avoid frequent LSP aging in the network, RBs periodically advertise LSPs. The actual refresh interval of an LSP is determined by both the minimum LSP interval and the maximum number of LSPs transmitted per interval. To prevent LSPs from being aged out accidentally, set the LSP maximum age and the LSP refresh interval appropriately.

To adjust LSP-related parameters:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Set the LSP maximum age.	timer lsp-max-age <i>time</i>	The default setting is 1200 seconds.
4. Set the LSP refresh interval.	timer lsp-refresh <i>time</i>	The default setting is 900 seconds.
5. Set the TRILL LSP generation timer parameters.	timer lsp-generation <i>maximum-interval</i> [<i>minimum-interval</i> [<i>incremental-interval</i>]]	By default, the maximum LSP generation interval is 2 seconds, the minimum interval is 10 milliseconds, and the incremental interval is 20 milliseconds.
6. Configure the maximum length of originated LSPs.	lsp-length originate <i>size</i>	The default setting is 1458 bytes. To prevent the system from generating error messages, do not set the maximum length of originated LSPs to be greater than the maximum length of received LSPs.
7. Configure the maximum length of received LSPs.	lsp-length receive <i>size</i>	The default setting is 1492 bytes. To prevent the system from generating error messages, do not set the maximum length of originated LSPs to be greater than the maximum length of received LSPs.
8. Set the Overload bit of LSPs and set the lifetime for the set Overload bit.	set overload [<i>timeout</i>]	By default, the Overload bit is not set. Do not perform this task on the root RB of a TRILL distribution tree. The root RB cannot forward traffic when the Overload bit of LSPs is set on the RB.
9. Return to system view.	quit	N/A
10. Enter Layer 2 Ethernet or aggregate interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
11. Configure the minimum LSP interval and the maximum number of LSPs transmitted per interval.	trill timer lsp <i>interval</i> [count <i>count</i>]	By default, the minimum LSP interval is 10 milliseconds, and the maximum number of LSPs transmitted per interval is 5.

Enabling TRILL LSP fast advertisement

LSP fast advertisement enables TRILL to immediately advertise the specified number of LSPs that invoke SPF calculation. This mechanism improves network convergence time.

To enable TRILL LSP fast advertisement:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A

3. Enable TRILL LSP fast advertisement.	flash-flood [flood-count <i>flooding-count</i> max-timer-interval <i>flooding-interval</i>] *	By default, TRILL LSP fast advertisement is disabled.
---	--	---

Enabling TRILL pseudonode bypass

This feature disables a DRB from generating LSPs for the pseudonode when the DRB has only one neighbor on a broadcast network. This reduces the number of LSPs in the network.

To enable the pseudonode bypass feature:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Enable the pseudonode bypass feature.	trill bypass-pseudonode enable	By default, the pseudonode bypass feature is disabled.

Setting the SPF algorithm parameters

The RB uses the SPF algorithm to calculate a shortest path tree with itself as the root based on the LSDB. The RB uses the shortest path tree to determine the next hop to a destination network. By adjusting the SPF calculation intervals, you can prevent resource overconsumption when the network is unstable.

When the network is stable, the SPF calculation interval for continuous calculations is reduced to *minimum-interval*. When the network is unstable, the SPF calculation interval is added by *incremental-interval* $\times 2^{n-2}$ (n is the number of continuous SPF calculation times) for each SPF calculation until the maximum interval is reached.

To adjust SPF algorithm parameters:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Set the SPF calculation interval for TRILL.	timer spf <i>maximum-interval</i> [<i>minimum-interval</i> [<i>incremental-interval</i>]]	By default, the maximum SPF calculation interval is 10 seconds, the minimum SPF calculation interval is 10 milliseconds, and the SPF calculation incremental interval is 20 milliseconds.

Configuring TRILL distribution trees

Setting basic distribution tree parameters

In a TRILL network, RBs perform the following tasks:

- Compute TRILL distribution trees according to the LSDB.
- Use the TRILL distribution trees to guide the forwarding of multdestination frames.

An RB with a higher priority is selected as the root bridge of a TRILL distribution tree.

An LSP sent by an RB carries the following TRILL distribution tree information:

- The number of TRILL distribution trees that the RB wants all RBs to compute.
- The maximum number of TRILL distribution trees that the RB can compute (this number is fixed at 15).
- The number of TRILL distribution trees that the RB has computed.

Each RB can compute a maximum of m TRILL distribution trees. An RB determines the number of TRILL distribution trees to compute (n) by selecting the lower value from the following two values:

- The number of TRILL distribution trees that the highest-priority RB wants all RBs to compute.
- The smallest m value across the TRILL network.

The RB selects the first n nicknames from the nickname list advertised by the highest-priority RB. The RB uses the selected nicknames as the root nicknames for computing distribution trees.

To set basic TRILL distribution tree parameters:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Set a priority for the RB.	tree-root priority <i>priority</i>	The default setting is 32768.
4. Configure the number of TRILL distribution trees that the RB wants all RBs to compute.	trees calculate <i>count</i>	The default setting is 1.

Enabling TRILL distribution tree multithread calculation

This feature enables a multicore CPU device to improve TRILL distribution tree calculation efficiency by using each thread to calculate a distribution tree.

To enable TRILL distribution tree multithread calculation:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Enable TRILL distribution tree multithread calculation.	multicast multi-thread enable	By default, this feature is disabled.

Enabling load balancing over TRILL distribution trees

By default, ingress traffic is not load balanced over TRILL distribution trees after a forwarding VLAN is deleted on the RB. To load balance ingress traffic of the remaining forwarding VLANs over the existing distribution trees, you can enable load balancing over TRILL distribution trees.

Ingress traffic is load balanced in any of the following conditions, regardless of whether load balancing is enabled or not:

- A forwarding VLAN is added.
- A distribution tree is added or deleted.

When a distribution tree is added, the RB switches ingress traffic to the new tree to implement load balancing. However, the RB cannot use the new distribution tree to forward traffic before other RBs

are ready to use the new tree. In this case, you can set a delay timer for the RB to switch ingress traffic to the new distribution tree.

If traffic is not evenly distributed over distribution trees when load balancing is disabled, you can also perform one-time load balancing over TRILL distribution trees.

To enable load balancing over TRILL distribution trees:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Enable load balancing over TRILL distribution trees.	ingress assign-rule load-balancing	By default, load balancing over TRILL distribution trees is disabled.
4. Set the delay timer for the RB to switch ingress traffic to a new TRILL distribution tree.	ingress assign-delay <i>seconds</i>	The default delay timer is 300 seconds.
5. Perform one-time load balancing over TRILL distribution trees.	set ingress-load-balancing	Perform this task when load balancing over TRILL distribution trees is disabled.

Configuring TRILL ECMP routes

TRILL unicast ECMP routes share traffic to the same destination. You can configure the maximum number of TRILL unicast ECMP routes.

When TRILL multicast ECMP is disabled, TRILL distribution trees do not use equal-cost routes to share traffic. When TRILL multicast ECMP is enabled, TRILL assigns equal-cost routes to multiple TRILL distribution trees to improve load sharing performance.

To configure TRILL ECMP routes:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Configure the maximum number of TRILL unicast ECMP routes.	max-unicast-load-balancing <i>number</i>	The default maximum number is 8.
4. Enable TRILL multicast ECMP.	multicast-ecmp enable [p2p-ignore]	By default, TRILL multicast ECMP is disabled. Make sure the status of TRILL multicast ECMP is the same across the TRILL network.

Enabling incremental flush for TRILL multicast routing entries

TRILL multicast routing entries are classified into three levels according to the following key combinations:

- **RB**—Root bridge of a TRILL distribution tree.

- **RB+VLAN**—Root bridge and VLAN of a TRILL distribution tree.
- **RB+VLAN+MAC**—Root bridge and VLAN of a TRILL distribution tree and a MAC address.

An entry that is identified by fewer keys is at a higher level.

The incremental flush feature enables the device to compare the outgoing port list and local receiving flag of an entry with its next higher level entry. If the two entries have the same outgoing port list and local receiving flag, the higher level entry is issued to the TRILL FIB. For example, if entry RB 2 and entry RB 2+VLAN 10 have the same outgoing port list and local receiving flag, entry RB 2 is issued.

This feature reduces the number of flushed entries in the scenarios where an entry and its next higher level entry have the same outgoing port list and local receiving flag. Enabling this feature in other scenarios causes the system to issue a large number of entries at the same time and degrades the device performance.

To enable incremental flush for TRILL multicast routing entries:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Enable incremental flush for TRILL multicast routing entries.	flush-policy difference	By default, incremental flush is disabled for TRILL multicast routing entries.

Enabling logging of TRILL neighbor changes

Perform this task to output logs of TRILL neighbor changes to the configuration terminal.

To enable logging of TRILL neighbor changes:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Enable logging of TRILL neighbor changes.	log-peer-change enable	By default, logging of TRILL neighbor changes is enabled.

Configuring SNMP for TRILL

To report critical TRILL events to an NMS, enable SNMP notifications for TRILL. For TRILL event notifications to be sent correctly, you must also configure SNMP as described in *Network Management and Monitoring Configuration Guide*.

TRILL shares the standard IS-IS MIB with IS-IS. The standard IS-IS MIB provides only single-instance MIB objects. For SNMP to correctly identify TRILL's management information in the standard IS-IS MIB, you must configure a unique context for TRILL.

Context is a method introduced to SNMPv3 for multiple-instance management. For SNMPv1/v2c, you must specify a community name as a context name for protocol identification.

To configure SNMP for TRILL:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enable SNMP notifications for TRILL.	snmp-agent trap enable trill [adjacency-state-change area-mismatch buffsize-mismatch id-length-mismatch lsdboverload-state-change lsp-parse-error lsp-size-exceeded max-seq-exceeded maxarea-mismatch new-drb own-lsp-purge protocol-support rejected-adjacency skip-sequence-number topology-change version-skew] *	By default, SNMP notifications are enabled for TRILL.
3. Enter TRILL view.	trill	N/A
4. Configure the context name for TRILL.	snmp context-name <i>context-name</i>	By default, no context name is configured for TRILL.

Configuring TRILL GR

Graceful Restart (GR) ensures the continuity of packet forwarding when a protocol restarts or an active/standby switchover occurs on the RB. The RB advertises the restart status to its neighbors, and allows the neighbors to re-establish connections. GR involves the following roles:

- **GR restarter**—Graceful restarting router. It must be GR capable.
- **GR helper**—A neighbor of the GR restarter. It helps the GR restarter to complete the GR process.

By default, the device acts as the GR helper. Configure TRILL GR on the target GR restarter.

To configure TRILL GR:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Enable GR for TRILL.	graceful-restart	By default, GR is disabled for TRILL.
4. (Optional.) Configure the GR interval for TRILL.	graceful-restart interval <i>interval</i>	The default setting is 300 seconds.
5. (Optional.) Suppress the SA bit during graceful restart.	graceful-restart suppress-sa	By default, the SA bit is set during graceful restart.

Associating a TRILL port with a track entry

Associate a track entry with a TRILL port to fast detect the loss of connectivity to the neighbor on the port. Track can collaborate with Connectivity Fault Detection (CFD) to monitor the link state of the neighbor. For more information about CFD and Track, see *High Availability Configuration Guide*.

To use CFD to detect link failures in a TRILL network, you must configure the outward-facing MEPs. CFD supports only single-hop detection. CFD packets cannot be forwarded by RBs.

To associate a TRILL port with a track entry:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter Layer 2 Ethernet interface view or Layer 2 aggregate interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Associate a track entry with the interface.	trill track <i>track-entry-number</i>	By default, an interface is not associated with any track entries.

Displaying and maintaining TRILL

Execute the **display** commands in any view and the **reset** command in user view.

Task	Command
Display TRILL adjacency table information.	display trill adjacent-table [count nickname <i>nickname</i> interface <i>interface-type</i> <i>interface-number</i>]
Display brief TRILL information.	display trill brief
Display TRILL FIB information.	display trill fib [count nickname <i>nickname</i>]
Display TRILL GR status.	display trill graceful-restart status
Display TRILL ingress forwarding information.	display trill ingress-route [vlan <i>vlan-list</i>]
Display TRILL port information.	display trill interface [<i>interface-type</i> <i>interface-number</i>]
Display TRILL LSDB information.	display trill lsdb [local lsp-id <i>lsp-id</i> verbose] *
Display all ingress entries in the TRILL multicast FIB (MFIB).	display trill mfib ingress [vlan <i>vlan-id</i> [local-entry remote-entry]]
Display all egress entries in the TRILL MFIB.	display trill mfib transit [nickname <i>nickname</i> [prune-entry rpf-entry vlan <i>vlan-id</i> [mac-address <i>mac-address</i>]]]
Display information about the TRILL multicast routing table.	display trill multicast-route [tree-root <i>nickname</i> [vlan <i>vlan-list</i> [mac <i>mac-address</i>]]]
Display the TRILL neighbor table.	display trill neighbor-table
Display TRILL neighbor statistics.	display trill peer [interface <i>interface-type</i> <i>interface-number</i>]
Display the TRILL RPF check table information.	display trill rpf-table tree-root <i>nickname</i>
Display TRILL topology information.	display trill topology [verbose]
Display information about the TRILL unicast routing table.	display trill unicast-route [nickname <i>nickname</i>] [verbose]
Clear dynamic running statistics of the TRILL process.	reset trill

TRILL configuration example

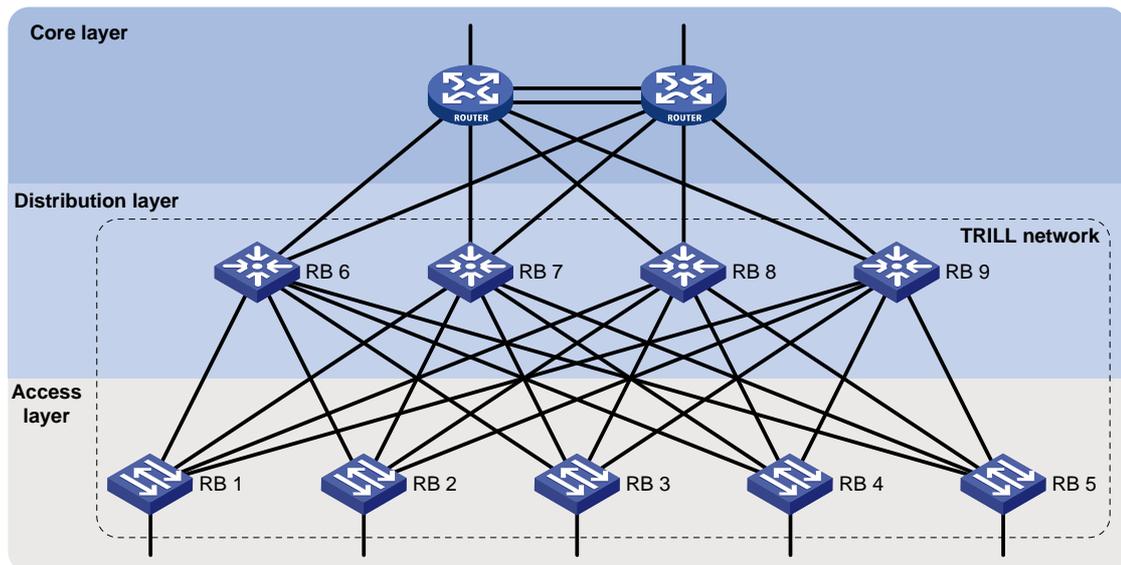
Network requirements

As shown in [Figure 5](#), a Layer 2 data center network has three layers: the core layer, distribution layer, and access layer. A port connected to a higher layer device is an uplink port, and a port connected to a lower layer device is a downlink port.

Configure TRILL in the network as follows:

- Enable TRILL on the downlink ports of access layer devices to connect terminal devices to the TRILL network.
- Enable TRILL on the uplink ports of access layer devices, and configure these uplink ports as trunk ports to pass TRILL frames to the TRILL network.
- Enable TRILL on the downlink ports of distribution layer devices, and configure these downlink ports as trunk ports to forward TRILL data frames.
- Enable TRILL on the uplink ports of the distribution layer devices. These ports send the de-encapsulated TRILL data frames to the core layer.
- In the TRILL network, configure four TRILL distribution trees with RB 6 through RB 9 as the root bridges. RB 6 through RB 9 are in descending priority order.

Figure 5 Network diagram



Configuration procedure

This section provides only TRILL-related configurations.

1. Configure the downlink ports of access layer devices:

Enable TRILL globally on RB 1, and enable TRILL on downlink port Ten-GigabitEthernet 1/0/1 of RB 1.

```
<RB1> system-view
[RB1] trill
[RB1-trill] quit
[RB1] interface ten-gigabitethernet 1/0/1
[RB1-Ten-GigabitEthernet1/0/1] trill enable
```

- ```
[RB1-Ten-GigabitEthernet1/0/1] quit
```
- # Configure RB 2 through RB 5 in the same way RB 1 is configured. (Details not shown.)
- Configure the uplink port of access layer devices:
 

# Enable TRILL on uplink port Ten-GigabitEthernet 1/0/2 through Ten-GigabitEthernet 1/0/5 of RB 1, and configure these ports as trunk ports.

```
[RB1] interface range ten-gigabitethernet 1/0/2 to ten-gigabitethernet 1/0/5
[RB1-if-range] trill enable
[RB1-if-range] trill link-type trunk
[RB1-if-range] quit
```

# Configure RB 2 through RB 5 in the same way RB 1 is configured. (Details not shown.)
  - Configure the downlink ports of distribution layer devices:
 

# Enable TRILL globally on RB 6, enable TRILL on downlink port Ten-GigabitEthernet 1/0/1 through Ten-GigabitEthernet 1/0/5 of RB 6, and configure these ports as trunk ports.

```
<RB6> system-view
[RB6] trill
[RB6-trill] quit
[RB6] interface range ten-gigabitethernet 1/0/1 to ten-gigabitethernet 1/0/5
[RB6-if-range] trill enable
[RB6-if-range] trill link-type trunk
[RB6-if-range] quit
```

# Configure RB 7 through RB 9 in the same way RB 6 is configured. (Details not shown.)
  - Configure the uplink ports of the distribution layer devices:
 

# Enable TRILL on uplink ports Ten-GigabitEthernet 1/0/6 and Ten-GigabitEthernet 1/0/7 of RB 6.

```
[RB6] interface ten-gigabitethernet 1/0/6
[RB6-Ten-GigabitEthernet1/0/6] trill enable
[RB6-Ten-GigabitEthernet1/0/6] quit
[RB6] interface ten-gigabitethernet 1/0/7
[RB6-Ten-GigabitEthernet1/0/7] trill enable
[RB6-Ten-GigabitEthernet1/0/7] quit
```

# Configure RB 7 through RB 9 in the same way RB 6 is configured. (Details not shown.)
  - Configure TRILL distribution trees:
 

# Set the RB 6's priority to 65535, and set the number of TRILL distribution trees that the RB wants all RBs to compute to 4.

```
[RB6] trill
[RB6-trill] tree-root priority 65535
[RB6-trill] trees calculate 4
[RB6-trill] quit
```

# Set the RB 7's priority to 65534, and set the number of TRILL distribution trees that the RB wants all RBs to compute to 4.

```
[RB7] trill
[RB7-trill] tree-root priority 65534
[RB7-trill] trees calculate 4
[RB7-trill] quit
```

# Set the RB 8's priority to 65533, and set the number of TRILL distribution trees that the RB wants all RBs to compute to 4.

```
[RB8] trill
[RB8-trill] tree-root priority 65533
```

```
[RB8-trill] trees calculate 4
```

```
[RB8-trill] quit
```

# Set the RB 9's priority to 65532, and set the number of TRILL distribution trees that the RB wants all RBs to compute to 4.

```
[RB9] trill
```

```
[RB9-trill] tree-root priority 65532
```

```
[RB9-trill] trees calculate 4
```

```
[RB9-trill] quit
```

## Verifying the configuration

Suppose that the nicknames of RB 1 through RB 9 are 0x5801 through 0x5809.

Use **display trill unicast-route** to display the TRILL unicast routing table. For example:

# Display brief information about all entries in the TRILL unicast routing table on RB 1.

```
[RB1] display trill unicast-route
```

| Destination | Interface | NextHop |
|-------------|-----------|---------|
| 0x5801      | N/A       | N/A     |
| 0x5802      | XGE1/0/2  | 0x5806  |
|             | XGE1/0/3  | 0x5807  |
|             | XGE1/0/4  | 0x5808  |
|             | XGE1/0/5  | 0x5809  |
| 0x5803      | XGE1/0/2  | 0x5806  |
|             | XGE1/0/3  | 0x5807  |
|             | XGE1/0/4  | 0x5808  |
|             | XGE1/0/5  | 0x5809  |
| 0x5804      | XGE1/0/2  | 0x5806  |
|             | XGE1/0/3  | 0x5808  |
|             | XGE1/0/4  | 0x5808  |
|             | XGE1/0/5  | 0x5809  |
| 0x5805      | XGE1/0/2  | 0x5806  |
|             | XGE1/0/3  | 0x5807  |
|             | XGE1/0/4  | 0x5808  |
|             | XGE1/0/5  | 0x5809  |
| 0x5806      | XGE1/0/2  | Direct  |
| 0x5807      | XGE1/0/3  | Direct  |
| 0x5808      | XGE1/0/4  | Direct  |
| 0x5809      | XGE1/0/5  | Direct  |

Use **display trill multicast-route** to display the TRILL multicast routing table. For example:

# Display the TRILL multicast routing table on RB 1.

```
[RB1] display trill multicast-route
```

| Root   | Flag  |
|--------|-------|
| 0x5806 | Valid |
| 0x5807 | Valid |
| 0x5808 | Valid |
| 0x5809 | Valid |

# Display the TRILL multicast routing table information for the TRILL distribution tree with RB 6 as the root bridge on RB 1.

```
[RB1] display trill multicast-route tree-root 5806
```

```
Root: 0x5806
```

```
LocalRcvFlag: True
```

```
List of VLANs:
```

```
 1
```

```
List of outgoing ports (1 in total):
```

```
 XGE1/0/2
```

# Document conventions and icons

## Conventions

This section describes the conventions used in the documentation.

### Port numbering in examples

The port numbers in this document are for illustration only and might be unavailable on your device.

### Command conventions

| Convention        | Description                                                                                                                                              |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Boldface</b>   | <b>Bold</b> text represents commands and keywords that you enter literally as shown.                                                                     |
| <i>Italic</i>     | <i>Italic</i> text represents arguments that you replace with actual values.                                                                             |
| [ ]               | Square brackets enclose syntax choices (keywords or arguments) that are optional.                                                                        |
| { x   y   ... }   | Braces enclose a set of required syntax choices separated by vertical bars, from which you select one.                                                   |
| [ x   y   ... ]   | Square brackets enclose a set of optional syntax choices separated by vertical bars, from which you select one or none.                                  |
| { x   y   ... } * | Asterisk marked braces enclose a set of required syntax choices separated by vertical bars, from which you select at least one.                          |
| [ x   y   ... ] * | Asterisk marked square brackets enclose optional syntax choices separated by vertical bars, from which you select one choice, multiple choices, or none. |
| &<1-n>            | The argument or keyword and argument combination before the ampersand (&) sign can be entered 1 to n times.                                              |
| #                 | A line that starts with a pound (#) sign is comments.                                                                                                    |

### GUI conventions

| Convention      | Description                                                                                                                                 |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Boldface</b> | Window names, button names, field names, and menu items are in Boldface. For example, the <b>New User</b> window appears; click <b>OK</b> . |
| >               | Multi-level menus are separated by angle brackets. For example, <b>File &gt; Create &gt; Folder</b> .                                       |

### Symbols

| Convention                                                                                            | Description                                                                                                                                                            |
|-------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  <b>WARNING!</b>   | An alert that calls attention to important information that if not understood or followed can result in personal injury.                                               |
|  <b>CAUTION:</b>   | An alert that calls attention to important information that if not understood or followed can result in data loss, data corruption, or damage to hardware or software. |
|  <b>IMPORTANT:</b> | An alert that calls attention to essential information.                                                                                                                |
| <b>NOTE:</b>                                                                                          | An alert that contains additional or supplementary information.                                                                                                        |
|  <b>TIP:</b>       | An alert that provides helpful information.                                                                                                                            |

# Network topology icons

| Convention                                                                          | Description                                                                                                                                |
|-------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------|
|    | Represents a generic network device, such as a router, switch, or firewall.                                                                |
|    | Represents a routing-capable device, such as a router or Layer 3 switch.                                                                   |
|    | Represents a generic switch, such as a Layer 2 or Layer 3 switch, or a router that supports Layer 2 forwarding and other Layer 2 features. |
|    | Represents an access controller, a unified wired-WLAN module, or the access controller engine on a unified wired-WLAN switch.              |
|    | Represents an access point.                                                                                                                |
|    | Represents a wireless terminator unit.                                                                                                     |
|    | Represents a wireless terminator.                                                                                                          |
|   | Represents a mesh access point.                                                                                                            |
|  | Represents omnidirectional signals.                                                                                                        |
|  | Represents directional signals.                                                                                                            |
|  | Represents a security product, such as a firewall, UTM, multiservice security gateway, or load balancing device.                           |
|  | Represents a security card, such as a firewall, load balancing, NetStream, SSL VPN, IPS, or ACG card.                                      |

# Support and other resources

## Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:  
[www.hpe.com/assistance](http://www.hpe.com/assistance)
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:  
[www.hpe.com/support/hpesc](http://www.hpe.com/support/hpesc)

### Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

## Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.
- To download product updates, go to either of the following:
  - Hewlett Packard Enterprise Support Center **Get connected with updates** page:  
[www.hpe.com/support/e-updates](http://www.hpe.com/support/e-updates)
  - Software Depot website:  
[www.hpe.com/support/softwaredepot](http://www.hpe.com/support/softwaredepot)
- To view and update your entitlements, and to link your contracts, Care Packs, and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:  
[www.hpe.com/support/AccessToSupportMaterials](http://www.hpe.com/support/AccessToSupportMaterials)

---

### ⓘ **IMPORTANT:**

Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HP Passport set up with relevant entitlements.

---

## Websites

| Website                                                       | Link                                                                                                       |
|---------------------------------------------------------------|------------------------------------------------------------------------------------------------------------|
| <b>Networking websites</b>                                    |                                                                                                            |
| Hewlett Packard Enterprise Information Library for Networking | <a href="http://www.hpe.com/networking/resourcefinder">www.hpe.com/networking/resourcefinder</a>           |
| Hewlett Packard Enterprise Networking website                 | <a href="http://www.hpe.com/info/networking">www.hpe.com/info/networking</a>                               |
| Hewlett Packard Enterprise My Networking website              | <a href="http://www.hpe.com/networking/support">www.hpe.com/networking/support</a>                         |
| Hewlett Packard Enterprise My Networking Portal               | <a href="http://www.hpe.com/networking/mynetworking">www.hpe.com/networking/mynetworking</a>               |
| Hewlett Packard Enterprise Networking Warranty                | <a href="http://www.hpe.com/networking/warranty">www.hpe.com/networking/warranty</a>                       |
| <b>General websites</b>                                       |                                                                                                            |
| Hewlett Packard Enterprise Information Library                | <a href="http://www.hpe.com/info/enterprise/docs">www.hpe.com/info/enterprise/docs</a>                     |
| Hewlett Packard Enterprise Support Center                     | <a href="http://www.hpe.com/support/hpesc">www.hpe.com/support/hpesc</a>                                   |
| Hewlett Packard Enterprise Support Services Central           | <a href="http://ssc.hpe.com/portal/site/ssc/">ssc.hpe.com/portal/site/ssc/</a>                             |
| Contact Hewlett Packard Enterprise Worldwide                  | <a href="http://www.hpe.com/assistance">www.hpe.com/assistance</a>                                         |
| Subscription Service/Support Alerts                           | <a href="http://www.hpe.com/support/e-updates">www.hpe.com/support/e-updates</a>                           |
| Software Depot                                                | <a href="http://www.hpe.com/support/softwaredepot">www.hpe.com/support/softwaredepot</a>                   |
| Customer Self Repair (not applicable to all devices)          | <a href="http://www.hpe.com/support/selfrepair">www.hpe.com/support/selfrepair</a>                         |
| Insight Remote Support (not applicable to all devices)        | <a href="http://www.hpe.com/info/insightremotesupport/docs">www.hpe.com/info/insightremotesupport/docs</a> |

## Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

[www.hpe.com/support/selfrepair](http://www.hpe.com/support/selfrepair)

## Remote support

Remote support is available with supported devices as part of your warranty, Care Pack Service, or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

For more information and device support details, go to the following website:

[www.hpe.com/info/insightremotesupport/docs](http://www.hpe.com/info/insightremotesupport/docs)

## Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback ([docsfeedback@hpe.com](mailto:docsfeedback@hpe.com)). When submitting your feedback, include the document title,

part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.

# Index

## A

### access

TRILL access port link, 8

### advertisement

TRILL LSP fast advertisement, 12

### algorithm

TRILL SPF algorithm parameter, 13

appointed VLAN-x forwarder. See [AVF](#)

### associating

TRILL port+Track entry, 17

### AVF

TRILL AVF inhibition time, 10

AVF (TRILL), 1

## C

### calculating

TRILL link cost configuration, 9

### CFD

TRILL port+Track entry association, 17

### changing

TRILL neighbor change logging, 16

### configuring

TRILL, 1, 6, 19

TRILL announcing VLAN, 9

TRILL designated VLAN, 9

TRILL distribution tree, 13

TRILL ECMP routes, 15

TRILL GR, 17

TRILL LSP parameter+feature configuration, 11

TRILL port DRB priority, 8

TRILL port link cost, 9

TRILL port link type, 8

TRILL RB nickname, 7

TRILL RB system ID, 7

TRILL SNMP, 16

TRILL timers, 10

### cost

TRILL ECMP route configuration, 15

TRILL link cost configuration, 9

### CSNP

TRILL CSNP interval, 10

## D

### designated

TRILL designated routing bridge, 1

### displaying

TRILL, 18

### distributing

load balancing over TRILL distribution trees, 14

TRILL distribution tree, 13

TRILL distribution tree basics, 13

TRILL distribution tree multithread calculation, 14

### DRB

TRILL port priority, 8

## E

### enabling

incremental flush for TRILL multicast routing entries, 15

load balancing over TRILL distribution trees, 14

TRILL, 7

TRILL distribution tree multithread calculation, 14

TRILL LSP fast advertisement, 12

TRILL neighbor change logging, 16

TRILL pseudonode bypass, 13

## F

### feature compatibility

TRILL, 6

### flush

incremental flush for TRILL multicast routing entries, 15

### format

TRILL control frame format, 1

TRILL data frame format, 1

### forwarding

TRILL configuration, 1, 6, 19

TRILL frame formats, 1

TRILL frame forwarding (multi-destination), 3

TRILL frame forwarding (unicast), 3

TRILL GR configuration, 17

TRILL mechanisms, 3

TRILL SNMP configuration, 16

### frame

TRILL control frame format, 1

TRILL data frame format, 1

TRILL multi-destination forwarding, 3

TRILL unicast forwarding, 3

## G

### Graceful Restart (GR)

TRILL configuration, 17

TRILL GR helper, 17

TRILL GR restarter, 17

## H

### hardware

TRILL, 5

### hello

TRILL hello interval, 10

TRILL hello multiplier, 10

### hybrid

TRILL hybrid port link, 8

## I

### ID

TRILL RB system ID, 7

### interval

TRILL CSNP interval, 10

TRILL hello interval, 10

### IS-IS

TRILL configuration, 1, 6, 19

## L

### Layer 2

TRILL configuration, 1, 6, 19

### link

TRILL configuration, 1, 6, 19

TRILL link cost configuration, 9

TRILL link state database, 1

TRILL link state protocol data unit, 1

TRILL port link type, 8

### load balancing

load balancing over TRILL distribution trees, 14

### logging

TRILL neighbor change logging, 16

### LSDB

TRILL LSP parameter, 11

TRILL LSP parameter set, 11

### LSP

TRILL LSP fast advertisement, 12

TRILL LSP parameter set, 11

TRILL parameter+feature configuration, 11

## M

### maintaining

TRILL, 18

## N

### naming

TRILL RB nickname, 7

### neighbor

TRILL neighbor change logging, 16

### network

incremental flush for TRILL multicast routing entries, 15

load balancing over TRILL distribution trees, 14

pseudonode bypass, 13

TRILL announcing VLAN, 9

TRILL designated VLAN, 9

TRILL distribution tree, 13

TRILL distribution tree basics, 13

TRILL distribution tree multithread calculation, 14

TRILL ECMP route configuration, 15

TRILL enable, 7

TRILL forwarding mechanisms, 3

TRILL frame formats, 1

TRILL GR configuration, 17

TRILL LSP fast advertisement, 12

TRILL LSP parameter set, 11

TRILL LSP parameter+feature configuration, 11

TRILL neighbor change logging, 16

TRILL port DRB priority, 8

TRILL port link type, 8

TRILL RB nickname, 7

TRILL RB system ID, 7

TRILL routing bridge network, 1

TRILL SNMP configuration, 16

TRILL SPF algorithm parameter, 13

TRILL timer configuration, 10

### network management

TRILL basic concepts, 1

TRILL configuration, 1, 6, 19

nickname (TRILL), 1, 7

### NMS

TRILL SNMP configuration, 16

### notifying

TRILL SNMP configuration, 16

## P

### parameter

TRILL LSP max age parameter, 11

TRILL LSP max number transmitted, 11

TRILL LSP min interval, 11

TRILL LSP refresh interval, 11

### port

TRILL link cost configuration, 9

TRILL port DRB priority, 8

TRILL port link type, 8

TRILL port+Track entry association, 17

### priority

TRILL distribution tree, 13, 13

TRILL port DRB priority, 8

### procedure

associating TRILL port+Track entry, 17

configuring TRILL, 6, 19

configuring TRILL announcing VLAN, 9

- configuring TRILL designated VLAN, 9
- configuring TRILL distribution tree, 13
- configuring TRILL distribution tree basics, 13
- configuring TRILL ECMP routes, 15
- configuring TRILL GR, 17
- configuring TRILL LSP parameter+feature configuration, 11
- configuring TRILL port DRB priority, 8
- configuring TRILL port link cost, 9
- configuring TRILL port link type, 8
- configuring TRILL RB nickname, 7
- configuring TRILL RB system ID, 7
- configuring TRILL SNMP, 16
- configuring TRILL timers, 10
- displaying TRILL, 18
- enabling incremental flush for TRILL multicast routing entries, 15
- enabling load balancing over TRILL distribution trees, 14
- enabling TRILL, 7
- enabling TRILL distribution tree multithread calculation, 14
- enabling TRILL LSP fast advertisement, 12
- enabling TRILL neighbor change logging, 16
- enabling TRILL pseudonode bypass, 13
- maintaining TRILL, 18
- setting TRILL LSP parameters, 11
- setting TRILL SPF algorithm parameters, 13
- protocols and standards
  - TRILL, 5
- pseudonode
  - TRILL pseudonode bypass, 13
- R**
- RB
  - TRILL nickname, 7
  - TRILL system ID, 7
- restrictions
  - TRILL configuration, 5
- route
  - TRILL ECMP route configuration, 15
- routing bridge
  - TRILL, 1
  - TRILL distribution tree, 13
  - TRILL distribution tree basics, 13
- S**
- setting
  - TRILL distribution tree basics, 13
  - TRILL LSP parameters, 11
  - TRILL SPF algorithm parameters, 13
- SNMP

- TRILL SNMP configuration, 16
- SPF
  - TRILL SPF algorithm parameter, 13
- system
  - TRILL RB nickname, 7
  - TRILL RB system ID, 7
- T**
- timer
  - TRILL AVF inhibition time, 10
  - TRILL configuration, 10
  - TRILL CSNP interval, 10
  - TRILL hello interval, 10
  - TRILL hello multiplier, 10
- Track
  - TRILL port+Track entry association, 17
- TRansparent Interconnection of Lots of Links. *Use TRILL*
- TRILL
  - basic concepts, 1
  - basic distribution tree configuration, 13
  - configuration, 1, 6, 19
  - configuration restrictions, 5
  - displaying, 18
  - distribution tree configuration, 13
  - distribution tree multithread calculation, 14
  - ECMP route configuration, 15
  - enable, 7
  - feature compatibility, 6
  - forwarding mechanisms, 3
  - frame formats, 1
  - frame forwarding (multi-destination), 3
  - frame forwarding (unicast), 3
  - GR configuration, 17
  - hardware, 5
  - how it works, 2
  - incremental flush for TRILL multicast routing entries, 15
  - load balancing over TRILL distribution trees, 14
  - LSP fast advertisement, 12
  - LSP parameter set, 11
  - LSP parameter+feature configuration, 11
  - maintaining, 18
  - neighbor change logging, 16
  - port DRB priority configuration, 8
  - port link cost configuration, 9
  - port link type configuration, 8
  - port+Track entry association, 17
  - protocols and standards, 5
  - pseudonode bypass, 13
  - RB nickname configuration, 7

- RB system ID configuration, [7](#)
- SNMP configuration, [16](#)
- SPF algorithm parameter set, [13](#)
- timer configuration, [10](#)
- VLAN configuration, [9](#)

trunk

- TRILL trunk port link, [8](#)

## **U**

unicast

- TRILL multi-destination frame forwarding, [3](#)

- TRILL unicast frame forwarding, [3](#)

## **V**

VLAN

- TRILL announcing VLAN, [9](#)

- TRILL appointed VLAN-x forwarder, [1](#)

- TRILL appointed VLAN-x port, [1](#)

- TRILL designated VLAN, [9](#)