

H3C S6800 Switch Series

TRILL Configuration Guide

New H3C Technologies Co., Ltd.
<http://www.h3c.com>

Software version: Release 2609 and later
Document version: 6W106-20200419

Copyright © 2020, New H3C Technologies Co., Ltd. and its licensors

All rights reserved

No part of this manual may be reproduced or transmitted in any form or by any means without prior written consent of New H3C Technologies Co., Ltd.

Trademarks

Except for the trademarks of New H3C Technologies Co., Ltd., any trademarks that may be mentioned in this document are the property of their respective owners.

Notice

The information in this document is subject to change without notice. All contents in this document, including statements, information, and recommendations, are believed to be accurate, but they are presented without warranty of any kind, express or implied. H3C shall not be liable for technical or editorial errors or omissions contained herein.

Preface

This configuration guide describes the background, working principles, and configuration methods of TRILL.

This preface includes the following topics about the documentation:

- [Audience](#).
- [Conventions](#).
- [Documentation feedback](#).

Audience

This documentation is intended for:

- Network planners.
- Field technical support and servicing engineers.
- Network administrators working with the S6800 switch series.

Conventions

The following information describes the conventions used in the documentation.

Command conventions

Convention	Description
Boldface	Bold 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 a minimum of 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
Boldface	Window names, button names, field names, and menu items are in Boldface. For example, the New User window opens; click OK .
>	Multi-level menus are separated by angle brackets. For example, File > Create > Folder .

Symbols

Convention	Description
 WARNING!	An alert that calls attention to important information that if not understood or followed can result in personal injury.
 CAUTION:	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.
 IMPORTANT:	An alert that calls attention to essential information.
NOTE:	An alert that contains additional or supplementary information.
 TIP:	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 module, such as a firewall, load balancing, NetStream, SSL VPN, IPS, or ACG module.

Examples provided in this document

Examples in this document might use devices that differ from your device in hardware model, configuration, or software version. It is normal that the port numbers, sample output, screenshots, and other information in the examples differ from what you have on your device.

Documentation feedback

You can e-mail your comments about product documentation to info@h3c.com.

We appreciate your comments.

Contents

Configuring TRILL.....	1
Overview	1
Basic concepts	1
TRILL frame formats	1
How TRILL works.....	2
TRILL forwarding mechanisms	3
Ping TRILL and tracert TRILL	5
Protocols and standards	6
Licensing requirements	7
Configuration restrictions and guidelines	7
TRILL configuration task list.....	7
Enabling TRILL	8
Configuration restrictions and guidelines	8
Configuration procedure.....	9
Configuring the system ID and nickname for an RB	9
Configuring the link type of a TRILL port.....	10
Configuring the DRB priority of a TRILL port	10
Setting the link cost for a TRILL port.....	11
Configuring announcing VLANs and the designated VLAN.....	11
Configuring TRILL timers	12
Configuring TRILL LSP parameters and features.....	13
Setting TRILL LSP parameters	13
Enabling TRILL LSP fast advertisement	14
Enabling TRILL pseudonode bypass	15
Setting the SPF algorithm parameters	15
Configuring TRILL distribution trees.....	15
Setting basic distribution tree parameters.....	15
Enabling TRILL distribution tree multithread calculation	16
Enabling load balancing over TRILL distribution trees.....	16
Configuring TRILL equal-cost routes	17
Enabling incremental flush for TRILL multicast routing entries	18
Enabling TRILL to forward traffic from EVB S-channels	18
Enabling logging of TRILL neighbor changes	19
Configuring SNMP for TRILL	19
Configuring TRILL GR.....	20
Associating a TRILL port with a track entry	20
Using ping TRILL and tracert TRILL to test network connectivity	21
Displaying and maintaining TRILL	21
TRILL configuration example	22
Network requirements	22
Configuration procedure.....	22
Verifying the configuration.....	24

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 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 protocol frames and data frames.

TRILL protocol frames include TRILL Hello, LSP, CSNP, PSNP, MTU-probe, and MTU-ack. These protocol 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

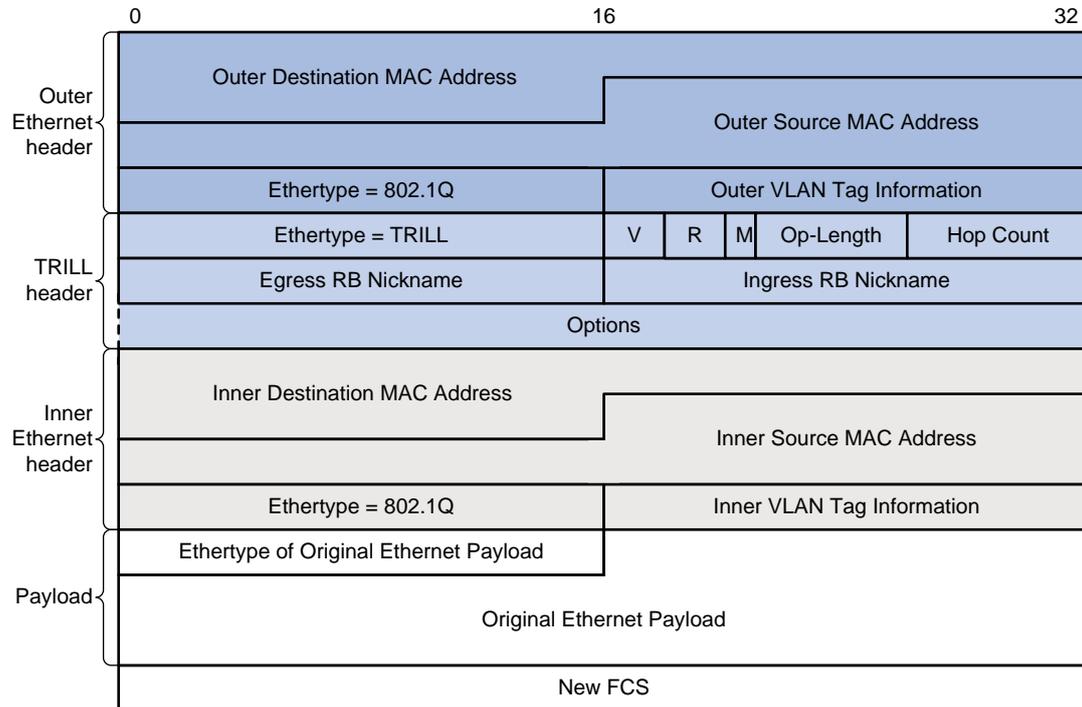


Table 1 describes the fields in the TRILL header.

Table 1 TRILL header fields

Field	Description
Ethertype	The Ether type is fixed to 0x22F3.
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.

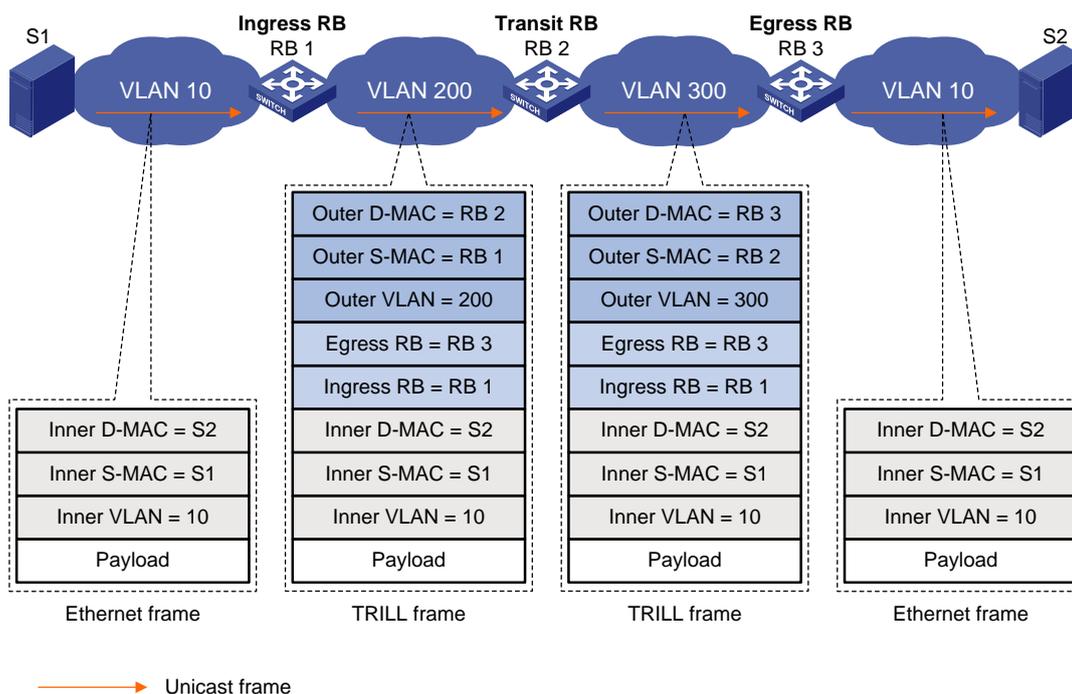
TRILL forwarding mechanisms

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:
 - o A TRILL header (similar to an IP header).
 - o 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 Layer 2 unicast frame forwarding flow



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

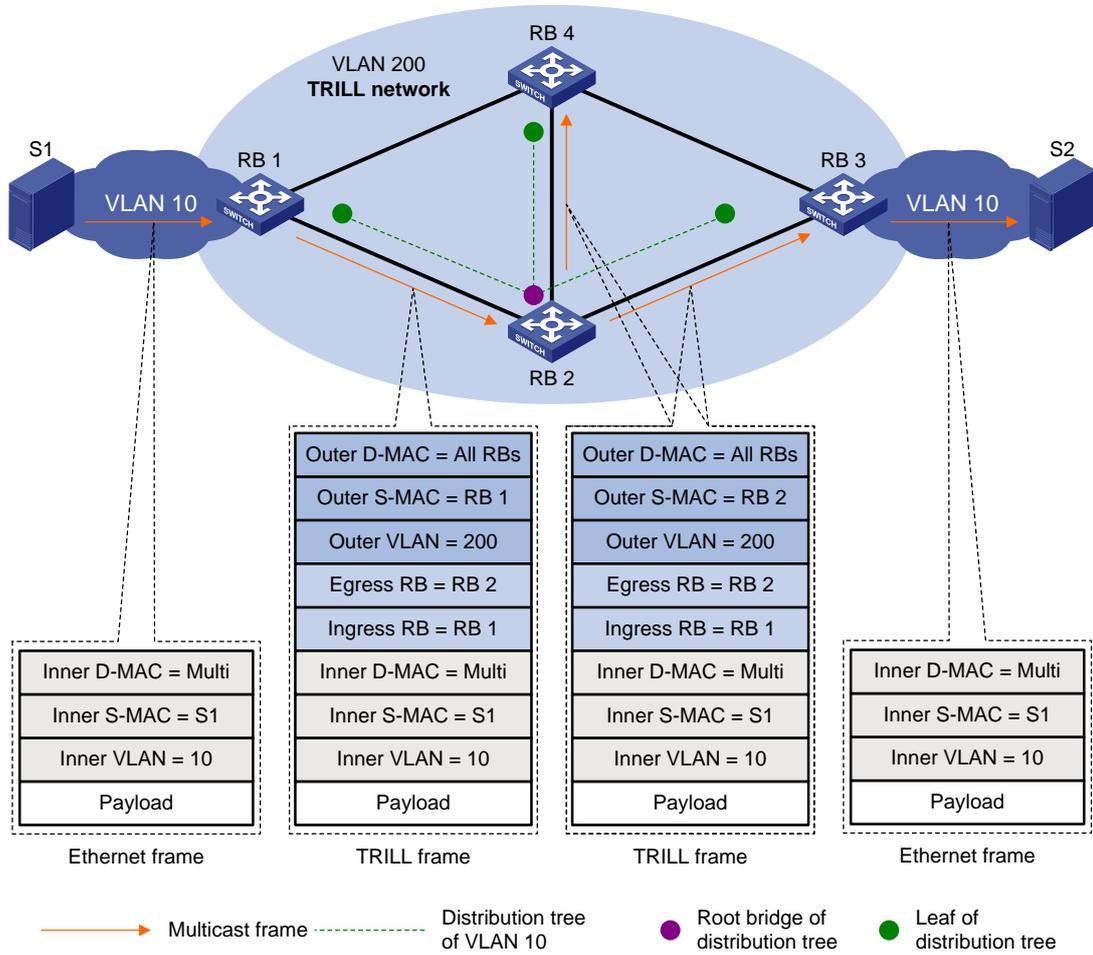
Multidestination frame forwarding mechanism

In a TRILL network, RBs perform the following operations:

- 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, and the destination MAC address is multicast address 0180-C200-0040. 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.

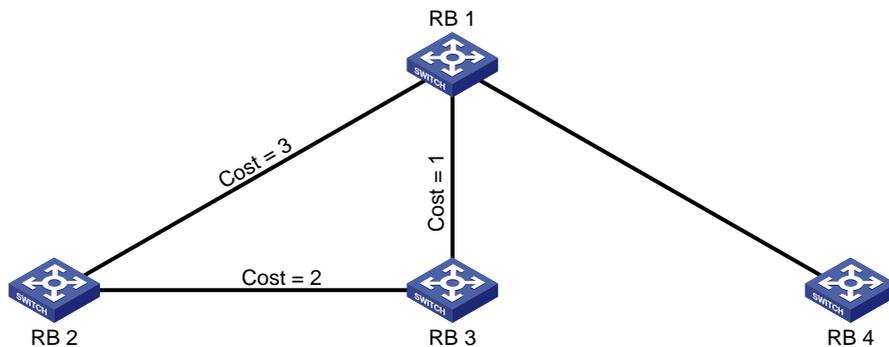
Figure 3 Layer 2 multicast frame forwarding flow



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 2. Assume the link directly connecting RB 1 to RB 2 has the largest pseudonode ID. Both the TRILL distribution tree rooted at RB 1 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 different TRILL distributions trees to improve 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 $J \text{ mod } N$ 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 2 is assigned to the TRILL distribution tree rooted at RB 1.
- The link RB 1-RB 3-RB 2 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.

Ping TRILL and tracet TRILL

You can use ping TRILL and tracet TRILL to test TRILL network connectivity when network failure occurs or new RBs are added to the network.

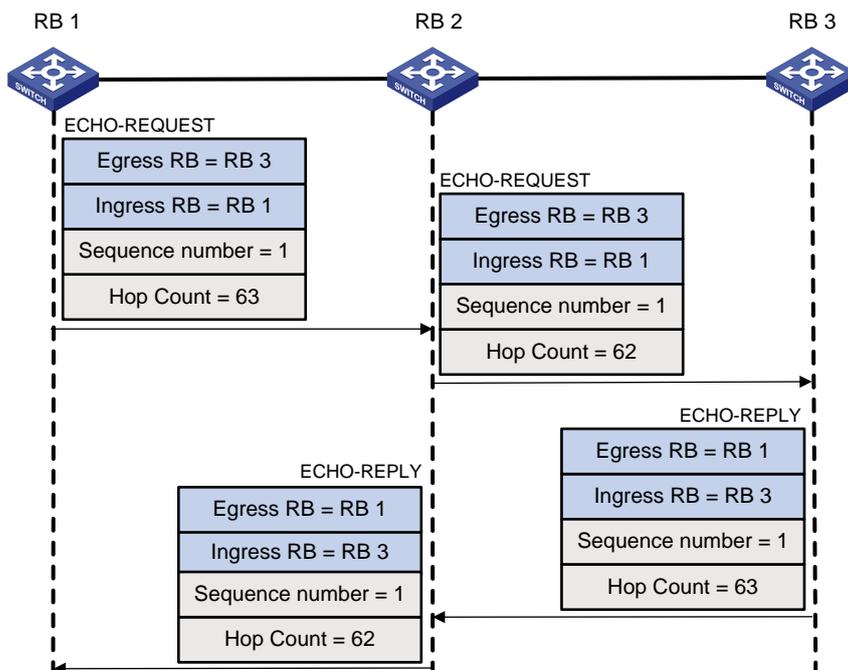
Ping TRILL and tracet TRILL are implemented through the TRILL Operation, Administration, and Maintenance (OAM) protocol.

Ping TRILL

Use ping TRILL to determine if an RB is reachable.

As shown in Figure 5, the source RB sends OAM echo requests to ping the destination RB. Upon receiving the requests, the destination RB responds to the source RB with OAM echo replies. The source RB outputs statistics about the ping TRILL operation, including the number of sent echo requests, the number of received echo replies, and the round-trip time. You can measure the network performance by analyzing the statistics.

Figure 5 Ping TRILL packet forwarding



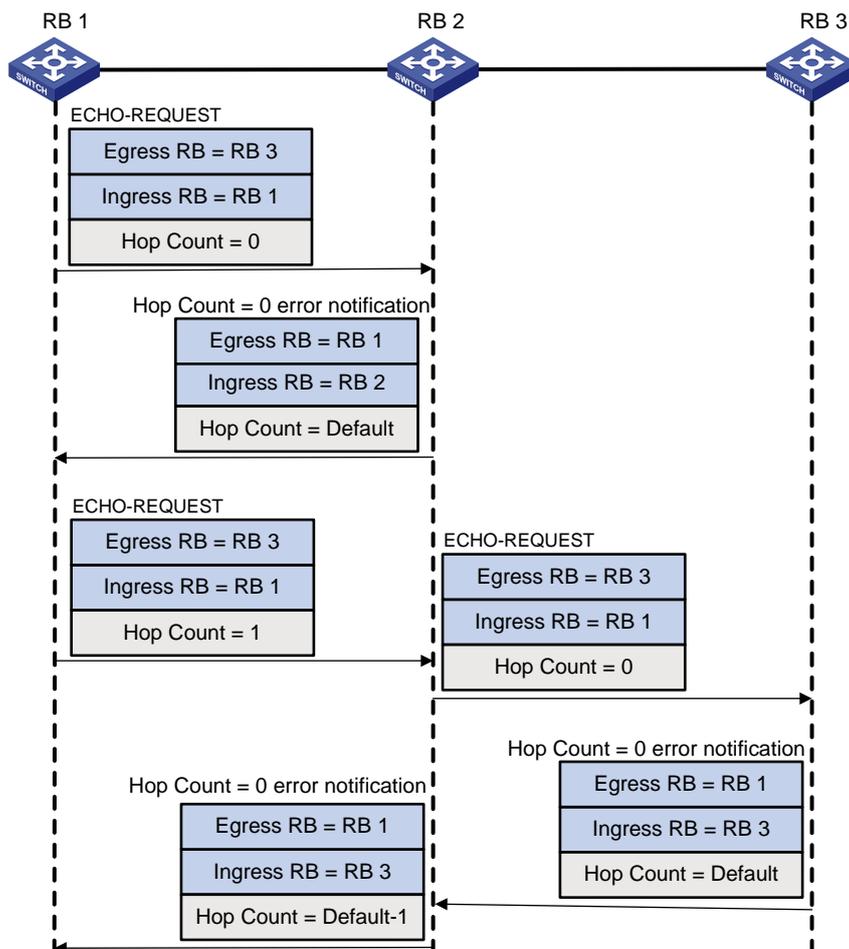
Tracert TRILL

Tracert TRILL enables retrieval of the nicknames of RBs in the path to a destination RB. In the event of network failure, use tracert TRILL to test network connectivity and identify failed nodes.

Tracert TRILL operates as shown in [Figure 6](#).

1. RB 1 sends RB 3 an OAM echo request with a hop count value of 0.
2. The first hop RB 2 responds by sending a hop count error notification to the source RB because the hop count of the request is 0. The notification uses the nickname of RB 2 as the ingress RB. In this way, RB 1 can get the nickname of the first hop RB.
3. RB 1 sends RB 3 an OAM echo request with a hop count value of 1.
4. RB 2 forwards the request to RB 3 according to the TRILL unicast routing table and decrements the hop count value by 1.
5. The second hop RB 3 responds to the source RB with a hop count error notification. The notification uses the nickname of RB 3 as the ingress RB.
6. RB 2 forwards the hop count error notification to RB 1. RB 1 gets the nickname of the second hop RB 3.

Figure 6 Tracert TRILL packet forwarding



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*
- RFC 7978, *Transparent Interconnection of Lots of Links (TRILL): RBridge Channel Header Extension*
- RFC 6905, *Requirements for Operations, Administration, and Maintenance (OAM) in Transparent Interconnection of Lots of Links (TRILL)*

Licensing requirements

TRILL requires a license to run on the device. For information about feature licensing, see *Fundamentals Configuration Guide*.

Configuration restrictions and guidelines

When you configure TRILL, follow these restrictions and guidelines:

- 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 a member port of an aggregate interface takes effect after the member port leaves the aggregation group.
- To connect a spanning tree network to a TRILL network, make sure the following requirements are met:
 - The spanning tree protocol is disabled on TRILL ports.
 - 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 *Virtual Technologies Configuration Guide*.
- Do not configure the TPID for VLAN tags on RBs. For more information about TPIDs, see *Layer 2—LAN Switching Configuration Guide*.

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.) Setting the link cost for a TRILL port
(Optional.) Configuring announcing VLANs and the designated VLAN
(Optional.) Configuring TRILL timers

Tasks at a glance
(Optional.) Configuring TRILL LSP parameters and features
(Optional.) Setting the SPF algorithm parameters
(Optional.) Configuring TRILL distribution trees
(Optional.) Configuring TRILL equal-cost routes
(Optional.) Enabling incremental flush for TRILL multicast routing entries
(Optional.) Enabling TRILL to forward traffic from EVB S-channels
(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
(Optional.) Using ping TRILL and tracer TRILL to test network connectivity

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

When you enable TRILL, follow these restrictions and guidelines:

- Before you enable TRILL on a port, first enable TRILL globally.
- Enable or disable TRILL on all ports in a VLAN, so that the ports in a VLAN have the same TRILL status (enabled or disabled).
- Do not enable both TRILL and EVB on a port. The allowed VLANs of a TRILL-enabled port cannot overlap with those of an EVB-enabled port if the **trill evb-support** command is not configured on the EVB-enabled port. For more information about EVB, see *EVB Configuration Guide*.
- When you set up a TRILL network, avoid the case that multiple TRILL neighbors are established for one RB port, for example, the networks in [Figure 7](#) and [Figure 8](#).

[Figure 7](#) shows a typical network where two TRILL neighbors are established for the same port of an RB. In [Figure 8](#), because TRILL is disabled on the port connecting RB 2 to RB 3, the port will transparently transmit the TRILL Hello frames from RB 3. As a result, two TRILL neighbors are established for the port connecting RB 1 to RB 2.

Figure 7 Two TRILL neighbors are established for a port (1)

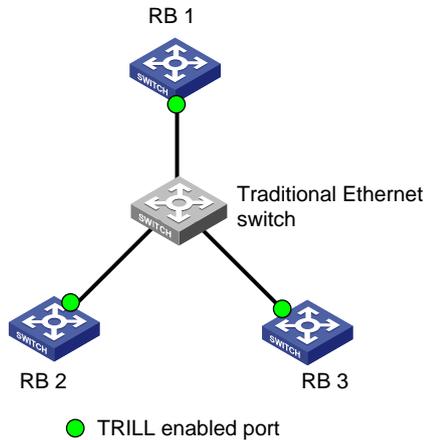
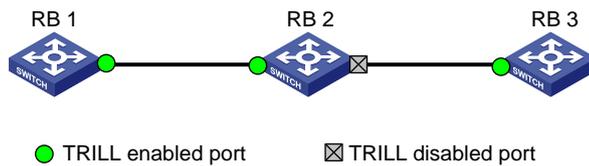


Figure 8 Two TRILL neighbors are established for a port (2)



Configuration procedure

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. If a TRILL port is on a PEX, the link type of the TRILL port can only be access.
- **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	interface <i>interface-type</i>	N/A

Step	Command	Remarks
or aggregate interface view.	<i>interface-number</i>	
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.

Setting the link cost for a TRILL port

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

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

If you manually set the link cost for a TRILL port, make sure its peer TRILL port uses the same link cost.

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

To set 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. Set the link cost for the TRILL port.	trill cost <i>cost-value</i>	The default setting is 2000.

Configuring announcing VLANs and the designated VLAN

The concepts and symbols used to describe a VLAN on 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 VLAN is 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. The system automatically selects an enabled VLAN as the designated VLAN.

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

Step	Command	Remarks
2. Enter Layer 2 Ethernet or aggregate interface view.	interface <i>interface-type</i> <i>interface-number</i>	N/A
3. Configure the Hello interval.	trill timer hello <i>interval</i>	The default setting is 10 seconds. 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 <i>count</i>	The default setting is 3.
5. Configure the AVF inhibition time.	trill timer avf-inhibited <i>time</i>	The default setting is 30 seconds.
6. Configure the CSNP interval.	trill timer csnp <i>interval</i>	The default setting is 10 seconds.

Configuring TRILL LSP parameters and features

Setting TRILL LSP parameters

You can set 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 operations:
 - 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 set TRILL LSP 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 time.	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 time 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 set the 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 operations:

- 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 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, 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 equal-cost routes

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

The maximum number of TRILL unicast equal-cost routes is restricted by the maximum number of ECMP routes set by using the **max-ecmp-num** command. The value range for the maximum number of ECMP routes is 1 to 128, and the default value is 8. If you set the maximum number of ECMP routes to *n* ($1 \leq n \leq 128$), the value range for the *number* argument of the **max-unicast-load-balancing** command is 1 to *n*. For more information about the **max-ecmp-num** command, see *Layer 3—IP Routing Command Reference*.

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

To configure TRILL equal-cost routes:

Step	Command	Remarks
1. Enter system view.	system-view	N/A
2. Enter TRILL view.	trill	N/A
3. Set the maximum number of TRILL unicast equal-cost routes.	max-unicast-load-balancing <i>number</i>	The default maximum number of TRILL unicast equal-cost routes equals the maximum number of ECMP routes.
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

Step	Command	Remarks
		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 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 TRILL to forward traffic from EVB S-channels

Perform this task on the interface that is directly connected to a server. TRILL will forward the packets received from the EVB S-channels on the interface to the TRILL network.

Before you perform this task, you must enable EVB on the interface.

To enable TRILL to forward traffic from EVB S-channels:

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

Step	Command	Remarks
3. Configure the link type of the port as trunk.	port link-type trunk	The default link type of a port is access.
4. Enable EVB.	evb enable	By default, EVB is disabled on a port.
5. Enable TRILL to forward traffic from EVB S-channels.	trill evb-support	By default, TRILL does not support forwarding traffic from EVB S-channels. The trill evb-support command and other TRILL commands (including trill enable) are mutually exclusive.

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 on the device. For more information about SNMP configuration, see the network management and monitoring configuration guide for the device.

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 lsdoverflow-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	snmp context-name <i>context-name</i>	By default, no context name

Step	Command	Remarks
context name for TRILL.		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.

Using ping TRILL and tracet TRILL to test network connectivity

To use ping TRILL to test the network connectivity:

Task	Command	Remarks
Determine if an RB with the specified nickname is reachable.	ping trill [-c <i>count</i> -h <i>tth</i> -i <i>interface-type interface-number</i> -m <i>interval</i> -priority <i>priority</i> -t <i>timeout</i>] * <i>nickname</i>	Execute this command in any view. If multiple routes destined for the RB exist, the RB is reachable if any of the routes is reachable.

To use tracet TRILL to test the network connectivity:

Task	Command	Remarks
Display the route to an RB with the specified nickname.	tracet trill [-f <i>first-ttl</i> -i <i>interface-type interface-number</i> -m <i>max-ttl</i> -priority <i>priority</i> -q <i>packet-number</i> -t <i>timeout</i> -v [-name]] * <i>nickname</i>	Execute this command in any view.

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 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 interface-number</i> verbose]
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 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]

Task	Command
Display information about the TRILL unicast routing table.	display trill unicast-route [<i>nickname nickname</i>] [<i>verbose</i>]
Clear dynamic running statistics of the TRILL process.	reset trill

TRILL configuration example

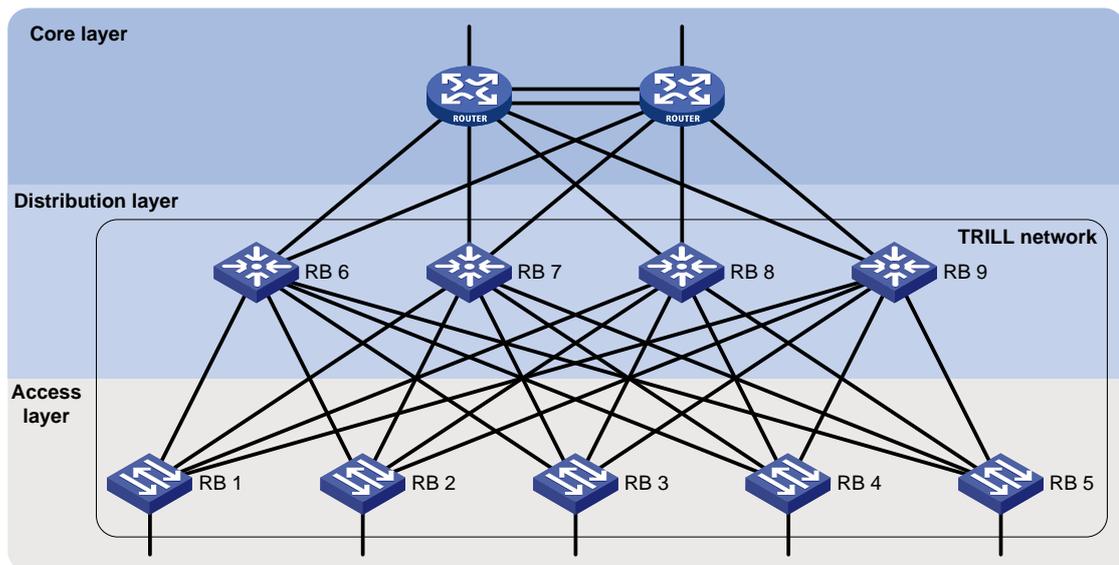
Network requirements

As shown in [Figure 9](#), a Layer 2 data center network has three layers: the core layer, the distribution layer, and the 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 data center 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 9 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.)

2. Configure the uplink ports of access layer devices:

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

```
[RB1] interface ten-gigabitethernet 1/0/2
[RB1-Ten-GigabitEthernet1/0/2] trill enable
[RB1-Ten-GigabitEthernet1/0/2] trill link-type trunk
[RB1-Ten-GigabitEthernet1/0/2] quit
[RB1] interface ten-gigabitethernet 1/0/3
[RB1-Ten-GigabitEthernet1/0/3] trill enable
[RB1-Ten-GigabitEthernet1/0/3] trill link-type trunk
[RB1-Ten-GigabitEthernet1/0/3] quit
[RB1] interface ten-gigabitethernet 1/0/4
[RB1-Ten-GigabitEthernet1/0/4] trill enable
[RB1-Ten-GigabitEthernet1/0/4] trill link-type trunk
[RB1-Ten-GigabitEthernet1/0/4] quit
[RB1] interface ten-gigabitethernet 1/0/5
[RB1-Ten-GigabitEthernet1/0/5] trill enable
[RB1-Ten-GigabitEthernet1/0/5] trill link-type trunk
[RB1-Ten-GigabitEthernet1/0/5] quit
```

Configure RB 2 through RB 5 in the same way RB 1 is configured. (Details not shown.)

3. Configure the downlink ports of distribution layer devices:

Enable TRILL globally on RB 6, enable TRILL on downlink ports 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 ten-gigabitethernet 1/0/1
[RB6-Ten-GigabitEthernet1/0/1] trill enable
[RB6-Ten-GigabitEthernet1/0/1] trill link-type trunk
[RB6-Ten-GigabitEthernet1/0/1] quit
[RB6] interface ten-gigabitethernet 1/0/2
[RB6-Ten-GigabitEthernet1/0/2] trill enable
[RB6-Ten-GigabitEthernet1/0/2] trill link-type trunk
[RB6-Ten-GigabitEthernet1/0/2] quit
[RB6] interface ten-gigabitethernet 1/0/3
[RB6-Ten-GigabitEthernet1/0/3] trill enable
[RB6-Ten-GigabitEthernet1/0/3] trill link-type trunk
[RB6-Ten-GigabitEthernet1/0/3] quit
[RB6] interface ten-gigabitethernet 1/0/4
```

```
[RB6-Ten-GigabitEthernet1/0/4] trill enable
[RB6-Ten-GigabitEthernet1/0/4] trill link-type trunk
[RB6-Ten-GigabitEthernet1/0/4] quit
[RB6] interface ten-gigabitethernet 1/0/5
[RB6-Ten-GigabitEthernet1/0/5] trill enable
[RB6-Ten-GigabitEthernet1/0/5] trill link-type trunk
[RB6-Ten-GigabitEthernet1/0/5] quit
```

Configure RB 7 through RB 9 in the same way RB 6 is configured. (Details not shown.)

4. 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.)

5. Configure TRILL distribution trees:

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

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

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

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

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

```
[RB8] trill
[RB8-trill] tree-root priority 65533
[RB8-trill] trees calculate 4
[RB8-trill] quit
```

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

```
[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.

Display brief information about all entries in the TRILL unicast routing table on an access layer device (for example, 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

Display the TRILL multicast routing table on an access layer device (for example, 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:
```

```
XGE1/0/2
```