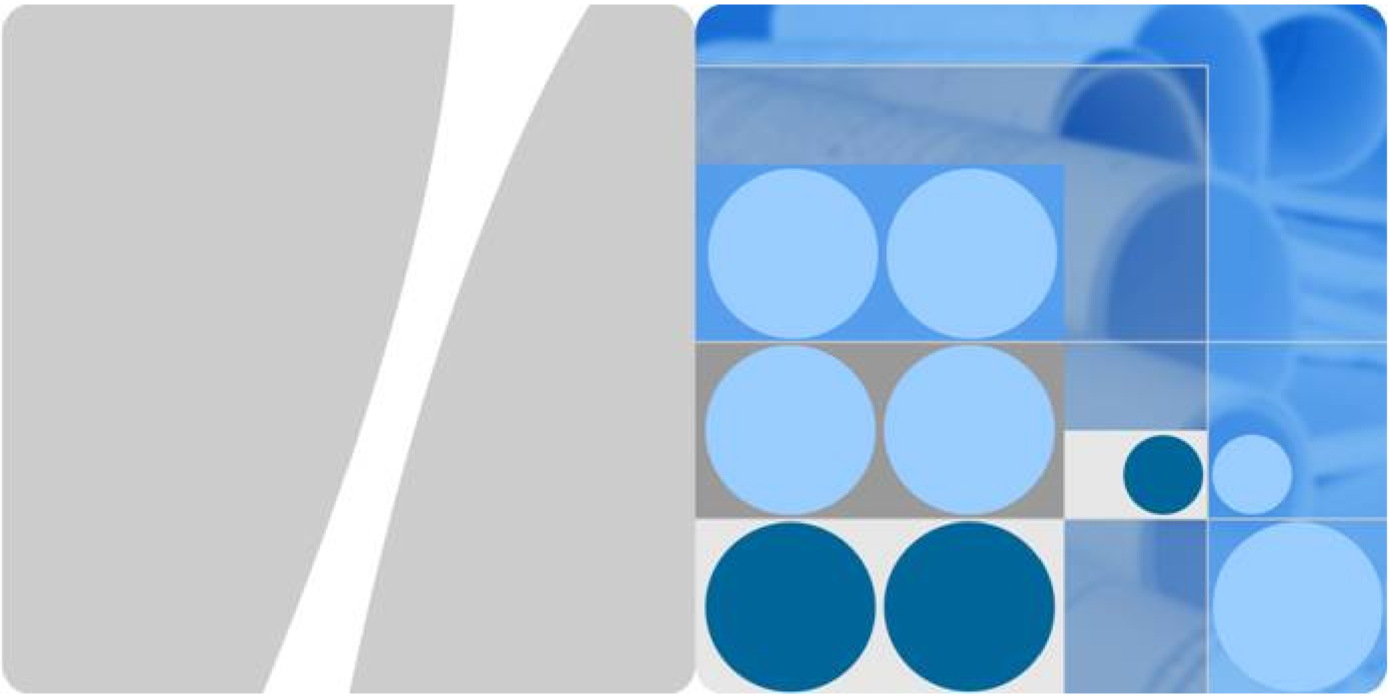

CloudEngine 12800 Series Switches

TRILL Technology White Paper

Issue	**01**
Date	**2016-06-21**

HUAWEI TECHNOLOGIES CO., LTD.

Trademarks and Permissions

HUAWEI and other Huawei trademarks are trademarks of Huawei Technologies Co., Ltd.
All other trademarks and trade names mentioned in this document are the property of their respective holders.

Notice

The purchased products, services and features are stipulated by the contract made between Huawei and the customer. All or part of the products, services and features described in this document may not be within the purchase scope or the usage scope. Unless otherwise specified in the contract, all statements, information, and recommendations in this document are provided "AS IS" without warranties, guarantees or representations of any kind, either express or implied.

The information in this document is subject to change without notice. Every effort has been made in the preparation of this document to ensure accuracy of the contents, but all statements, information, and recommendations in this document do not constitute a warranty of any kind, express or implied.

Huawei Technologies Co., Ltd.

Address: Huawei Industrial Base
Bantian, Longgang
Shenzhen 518129
People's Republic of China

Website: http://e.huawei.com

Contents

1 Introduction to TRILL

This section describes the definition and purpose of TRILL.

Definition

Transparent Interconnection of Lots of Links (TRILL) is a protocol that applies Layer 3 link-state routing technologies to Layer 2 networks. The TRILL protocol extends Intermediate System to Intermediate System (IS-IS) to Layer 2 to build large Layer 2 networks for data centers, providing solutions for data center services.

Purpose

In the cloud computing era, a data center usually uses a distributed architecture for mass data storage, query, and search services. In this architecture, cluster computing between servers generates heavy east-west traffic. As virtualization technologies are widely used in cluster computing, each server needs to compute much more data than before, and therefore throughput of a physical server increases by multiple times. In addition, virtual machines (VMs) must be able to dynamically migrate within a data center, to improve service reliability, reduce costs of IT services and network operation and maintenance, and allow for more flexible service deployment.

Because of these characteristics of cloud-computing data centers, the traditional hierarchical network structure with Layer 2 access (xSTP) and Layer 3 aggregation/core (routing) cannot satisfy requirements of data centers. Currently, a large Layer 2 fat tree architecture is widely used in data centers. TRILL helps build a non-blocking large Layer 2 network that supports smooth VM migration and can adapt to increasing network scales. The following table describes the advantages of a TRILL network over a traditional network using xSTP and Layer 3 routing protocols.

Table 1-1 Comparison between TRILL and xSTP networks

Requirements of Cloud Computing Data Center Networks	**Description**	**TRILL Network**	**xSTP Network**
Smooth VM migration	As one of core cloud computing technologies, server virtualization has been widely used. To maximize service reliability, reduce costs of IT services and network operation and maintenance, and allow flexible service deployment in a data center, VMs must be able to dynamically migrate within the data center but not just on an aggregation or access switch.	Deployed on a large Layer 2 network, TRILL supports dynamic VM migration in the entire data center.	In a traditional network with Layer 2 xSTP access and Layer 3 IP routing, the IP address of a VM will change if the VM migrates to another network segment. Therefore, VMs can only migrate within the same network segment.
Non-blocking, low-delay data forwarding	In a cloud-computing data center, most of traffic is east-west traffic, which is different from the traffic model on traditional carrier networks. Non-blocking, low-delay forwarding is required on data center networks to ensure normal service operation.	On a TRILL network, each device uses the shortest path tree (SPT) algorithm to calculate the shortest paths from itself to all the other nodes. If multiple equal-cost links are available, load balancing can be implemented among the unicast forwarding entries. Load balancing among equal-cost paths fully uses network bandwidth and implements line-speed forwarding on each node.	Redundant links are blocked and traffic is forwarded over a single path, which greatly wastes bandwidth and hinders construction of a non-blocking network.

Requirements of Cloud Computing Data Center Networks	Description	TRILL Network	xSTP Network
Large network scale	In the cloud computing era, a large data center may need to support as many as millions of servers. To implement non-blocking forwarding, hundreds or thousands of switches must be deployed on the data center network, and therefore loop prevention protocols are required. When a network node or link fails, fast network convergence must be triggered to quickly restore services. In addition, network maintenance must be simple enough to facilitate service deployment.	● Network scale: supports about 1000 switches theoretically. ● Loop prevention: uses loop-free IS-IS on the control plane, so no loop exists. ● Convergence rate: uses the IS-IS routing protocol to generate forwarding entries. Moreover, a TRILL header contains the Hop-Count field to allow temporary loops in a short period. These features implement subsecond convergence. ● Network maintenance: requires only simple configuration. Many parameters such as nickname and system ID can be automatically generated, and most protocol parameters can retain their default values. You only need to maintain one protocol (TRILL) instead of managing unicast and multicast protocols separately.	● Network scale: supports only about 100 devices because the network diameter cannot exceed 7 hops. ● Loop prevention: blocks redundant ports to eliminate rings. ● Convergence rate: completes convergence in seconds, due to limitations of the convergence mechanism. ● Network maintenance: requires a heavy workload because multiple routing protocols such as IGP and PIM must be maintained on the network.

Requirements of Cloud Computing Data Center Networks	Description	TRILL Network	xSTP Network
Multitenant	In the cloud computing era, a physical data center is shared by multiple tenants. Each tenant has a virtual data center instance, enabling tenants to exclusively use the server, storage, and network resources in the respective instances, while isolating data traffic of different tenant.	Currently, TRILL uses VLAN IDs to identify tenants and isolates traffic of tenants by VLANs. In the early stage of the cloud computing industry and large Layer 2 network operation, the limit of 4096 VLAN IDs will not become a bottleneck. Later, TRILL will use the FineLabel field to identify tenants. The FineLabel field is 24 bits and can support a maximum of 16M tenants, meeting requirements for future increase of tenants.	Only a maximum of 4096 tenants are supported, and the capacity cannot be expanded.
Scalability	Data center networks must have high scalability to adapt to fast development of data centers.	A traditional xSTP-based Layer 2 network can be seamlessly connected to a TRILL-based large Layer 2 network. TRILL allows large network scales, fast convergence rates, and high scalability.	The network has small scale, slow convergence rate, and low scalability.

Benefits

TRILL brings the following benefits:

- Large Layer 2 data centers support non-blocking VM migration, which facilitates network management.
- TRILL devices can be seamlessly connected to traditional xSTP networks, which reduces network upgrade costs.

2 Principles

About This Chapter

This section describes the principles of TRILL.

2.1 Basic Concepts

This section introduces basic concepts about Transparent Interconnection of Lots of Links (TRILL). **Figure 2-1** shows basic roles in the typical TRILL networking.

Figure 2-1 Large-scale Layer 2 TRILL networking

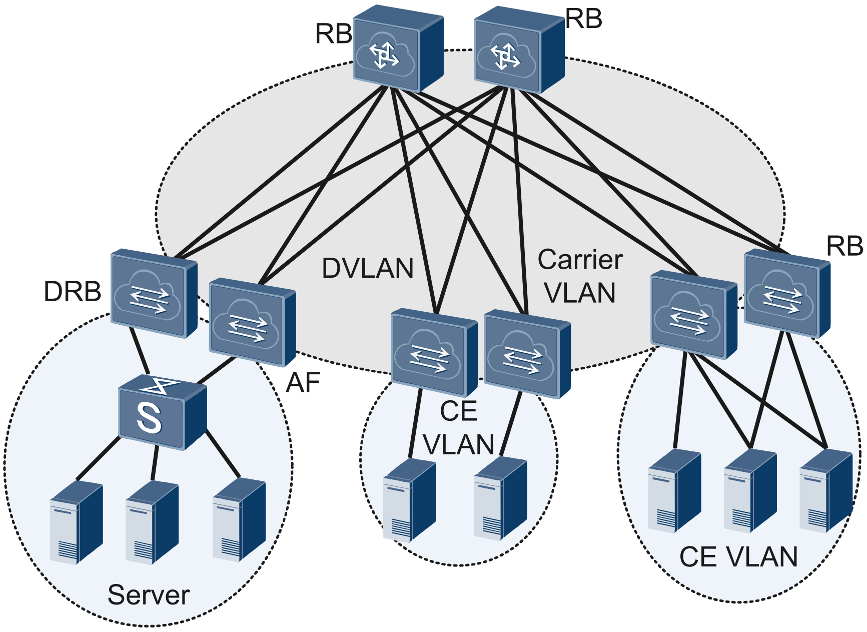

Devices in TRILL Networking

RB

Router bridge (RB) is a Layer 2 switch running TRILL. RBs are classified into ingress RB, transit RB, and egress RB according to their locations on a TRILL network. An ingress RB indicates the ingress from which packets enter the TRILL network. A transit RB indicates the intermediate node through which packets pass on the TRILL network. An egress RB indicates the egress from which packets leave the TRILL network.

DRB

A designated routing bridge (DRB) is an RB that functions as a transit device and performs special tasks on TRILL networks. On a TRILL broadcast network, if two RBs are located on the same virtual local area network (VLAN), the RB whose interface with a higher DRB priority or larger MAC address is selected as the DRB when they are establishing neighbor relationships. The DRB communicates with each device on the network to synchronize all the link state databases (LSDBs) on the VLAN, sparing every two devices from communicating for LSDB synchronization. DRBs perform the following tasks:

- Generate pseudonode link state protocol data units (LSPs) when more than two RBs exist on the network.
- Send complete sequence number protocol data units (CSNPs) to synchronize LSDBs.

- Select an carrier VLAN as the Designated VLAN, the DVLAN will transmit user packets and TRILL control packets.
- Select the appointed forwarder (AF). Only one RB can function as the AF for a customer edge (CE) VLAN.

AF

An AF is an RB elected by the DRB to forward user traffic. Non-AF RBs cannot forward user traffic. As shown in **Figure 2-1**, loops may occur if a server is dual-homed to a TRILL network but does not have double network adapters working in load balancing mode. Therefore, an RB must be elected to forward user traffic.

VLANs on a TRILL Network

Table 2-1 VLANs on a TRILL network

VLAN Name	Description	Packet Supported
CE VLAN	A CE VLAN connects to the TRILL network and is usually configured on the edge devices of a TRILL network to generate multicast routes.	Native Ethernet packets
Admin VLAN	A special CE VLAN transmits TRILL network management traffic.	TRILL network management traffic.
Carrier VLAN	A carrier VLAN transmits TRILL control packets and TRILL data packets. A maximum of three carrier VLANs can be configured on an RB. In the inbound direction, native Ethernet packets are encapsulated into TRILL packets in carrier VLANs. In the outbound direction, TRILL packets are decapsulated and restore to native Ethernet packets.	TRILL control packets and data packets
Designated VLAN	To combine or separate TRILL networks, multiple carrier VLANs are configured on a TRILL network. However, only one carrier VLAN is selected to forward TRILL control and data packets. The selected VLAN is called a designated VLAN.	TRILL control packets and data packets

Nickname

Each RB on a TRILL network has a unique nickname. The nickname is similar to an IP address in terms of function.

A nickname has one priority and one root priority.

- When a nickname conflict occurs on a TRILL network, the priority determines which RB's nickname is to be advertised to other RBs.
  a. The RB with the highest priority advertises its nickname.
  b. If the RBs with the same nickname have the same priority, the RB with the largest system ID advertises its nickname.
- An RB uses its root priority to run for the root of multicast tree. The RBs with the highest and second-highest root priority are selected as the roots of two multicast trees.

Interface Roles

Interfaces of switches on TRILL networks are classified into the following types:

- Trunk interfaces: connect switches and transmit TRILL data packets and protocol packets only.
- Access interfaces: transmit Native Ethernet packets and protocol packets only.
- Hybrid interfaces: transmit both TRILL data and protocol packets and Native Ethernet packets by default.
- P2P port: On a P2P network, the ports between two RBs are P2P ports. P2P ports are special trunk ports, and switches connected using the P2P ports do not participate in DRB election.

By default, the type of TRILL interfaces is p2p.

NET

Similar to IS-IS, TRILL uses network entity titles (NETs) to identify network layer information about switches. A NET includes the following elements:

- Area ID: An area ID identifies an area. An IS-IS network has multiple areas, while a TRILL network has only one area. The TRILL area ID is 00.
- System ID: identifies a host or switch and has a fixed length of 48 bits.

  In actual applications, a system ID can be automatically generated or configured. You can specify the system ID (unique on the entire network) when using the **network-entity (TRILL)** command to configure a NET. If this command is not configured, the system generates a system ID. The generated system ID is the same as the bridge MAC address of RB.
- SEL (also referred to as NSAP Selector or N-SEL): The role of a SEL is similar to that of the protocol identifier of IP. Each transport protocol has one unique SEL. The SEL of TRILL is 00.

2.2 TRILL Packet Formats

TRILL packets include TRILL control packets and TRILL data packets.

2.2.1 TRILL Control Packets

TRILL switches exchange control packets to communicate with each other. This section describes major TRILL control packets used on a TRILL network. TRILL uses IS-IS as the control plane protocol, uses PDUs to process control information, extends IS-IS PDUs.

TRILL PDUs use the same format as IS-IS PDUs, except that TRILL PDUs use extended IS-IS TLVs. For details on fields in a TRILL PDU, see "IS-IS Configuration" in the *CloudEngine 12800 Series switch Configuration Guide - IP Routing*.

TRILL PDU Format

All PDUs used by TRILL can be classified into three types: Hello, LSP, and SNP. The first eight bytes are fixed in all TRILL PDUs, as shown in **Figure 2-2**.

Figure 2-2 TRILL PDU structure

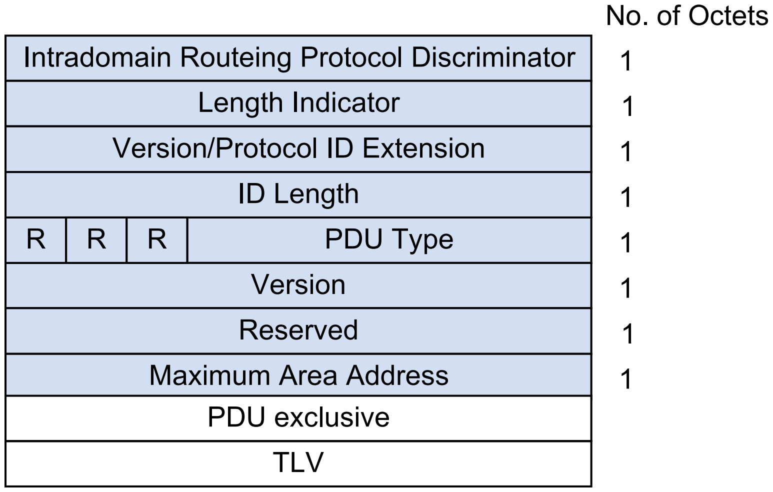

	No. of Octets
Intradomain Routeing Protocol Discriminator	1
Length Indicator	1
Version/Protocol ID Extension	1
ID Length	1
R | R | R | PDU Type	1
Version	1
Reserved	1
Maximum Area Address	1
PDU exclusive	
TLV	

The PDU fields are described as follows:

- Intradomain Routing Protocol Discriminator: identifies a network-layer PDU.
- Length Indicator: indicates the length of the fixed header.
- ID Length: indicates the length of the intra-domain system ID.
- PDU Type: indicates the PDU type.
- Maximum Area Address: indicates the maximum number of area addresses allowed by TRILL. Currently, the TRILL area address can only be 00.
- PDU Exclusive: varies depending on the PDU type and is described in the following PDU formats.
- TLV: indicates the type/length/value, which varies depending on the PDU type.

TRILL Hello PDU Format

Hello PDUs are used to establish and maintain neighbor relationships. LAN Hello PDUs are used on broadcast networks and P2P Hello PDUs are used on non-broadcast networks. LAN Hello PDUs and P2P Hello PDU have different formats.

Figure 2-3 shows the LAN Hello PDU format.

Figure 2-3 LAN Hello PDU format

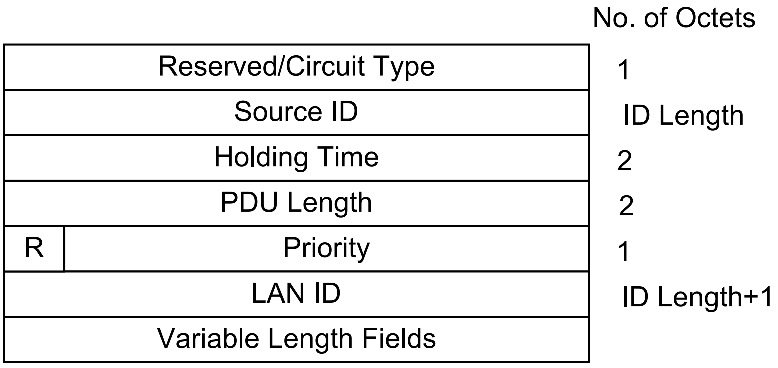

<table>
<tr><th colspan="2"></th><th>No. of Octets</th></tr>
<tr><td colspan="2">Reserved/Circuit Type</td><td>1</td></tr>
<tr><td colspan="2">Source ID</td><td>ID Length</td></tr>
<tr><td colspan="2">Holding Time</td><td>2</td></tr>
<tr><td colspan="2">PDU Length</td><td>2</td></tr>
<tr><td>R</td><td>Priority</td><td>1</td></tr>
<tr><td colspan="2">LAN ID</td><td>ID Length+1</td></tr>
<tr><td colspan="2">Variable Length Fields</td><td></td></tr>
</table>

Figure 2-4 shows the P2P Hello PDU format.

Figure 2-4 P2P Hello PDU format

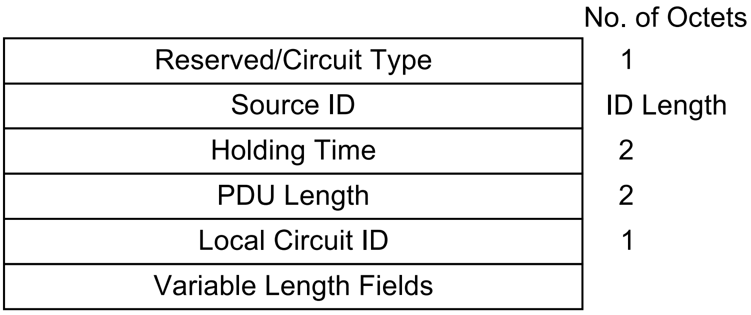

	No. of Octets
Reserved/Circuit Type	1
Source ID	ID Length
Holding Time	2
PDU Length	2
Local Circuit ID	1
Variable Length Fields	

As shown in **Figure 2-4**, most fields in a P2P Hello PDU are the same as those in a LAN Hello PDU. The P2P Hello PDU does not have the priority and LAN ID fields, but has a local circuit ID field indicating the local link ID.

TRILL LSP PDU Format

Link state PDUs (LSPs) are used to exchange link-state information. **Figure 2-5** shows the LSP PDU format.

Figure 2-5 TRILL LSP PDU format

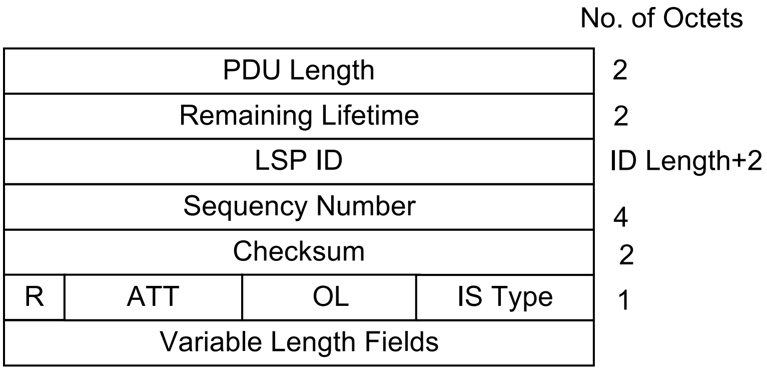

<table>
<tr><th colspan="4"></th><th>No. of Octets</th></tr>
<tr><td colspan="4">PDU Length</td><td>2</td></tr>
<tr><td colspan="4">Remaining Lifetime</td><td>2</td></tr>
<tr><td colspan="4">LSP ID</td><td>ID Length+2</td></tr>
<tr><td colspan="4">Sequncy Number</td><td>4</td></tr>
<tr><td colspan="4">Checksum</td><td>2</td></tr>
<tr><td>R</td><td>ATT</td><td>OL</td><td>IS Type</td><td>1</td></tr>
<tr><td colspan="4">Variable Length Fields</td><td></td></tr>
</table>

TRILL SNP PDU Format

Sequence number PDUs (SNPs) describe the all or some of LSPs to synchronize and maintain all LSDBs. SNPs are classified into the following types:

- Complete SNP (CSNP): carries summary of all LSPs in an LSDB, ensuring LSDB synchronization between neighboring switches. On a broadcast network, the DRB periodically sends CSNPs. The default interval for sending CSNPs is 10 seconds. On a point-to-point link, CSNPs are sent only when the neighbor relationship is established for the first time.

  Figure 2-6 shows the CSNP format.

  Figure 2-6 TRILL CSNP format

  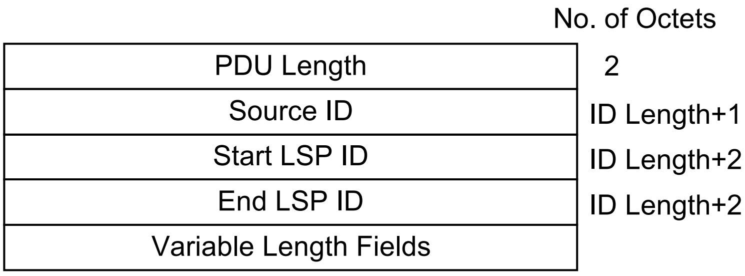

	No. of Octets
PDU Length	2
Source ID	ID Length+1
Start LSP ID	ID Length+2
End LSP ID	ID Length+2
Variable Length Fields	

- Partial SNP (PSNP): lists only the sequence numbers of recently received LSPs. A PSNP can acknowledge multiple LSPs at a time. If a device finds that its LSDB is not updated, it sends a PSNP to request a neighbor to send a new LSP.

  Figure 2-7 shows the PSNP format.

  Figure 2-7 TRILL PSNP format

  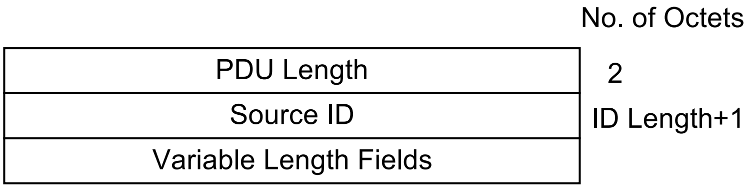

	No. of Octets
PDU Length	2
Source ID	ID Length+1
Variable Length Fields	

2.2.2 TRILL Data Packets

Figure 2-8 shows the TRILL data packet format.

Figure 2-8 TRILL data packet format

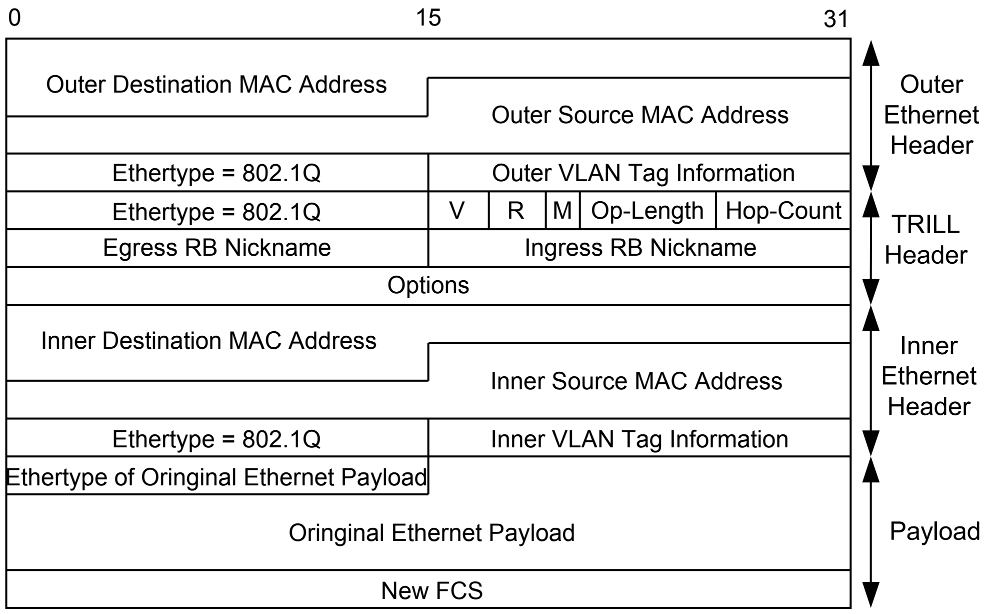

A TRILL data packet is generated by adding a TRILL header and an outer Ethernet header to the original Ethernet packet. The fields in a TRILL header are described as follows:

- Ethertype: fixed as TRILL.
- V: version number, which is 0 currently. Each RB must check the version number when receiving a TRILL packet. If the version is incorrect, the RB discards the packet.
- R: reserved for extension. This field is set to 0 on an ingress RB and ignored on transit and egress RBs.
- M: multi-destination attribute. The value 0 indicates known unicast packets and the value 1 indicates multicast, broadcast, and unknown unicast packets.
- Op-Length: length of the Options field. The value 0 indicates that the Options field is unavailable.
- Hop-Count: used to prevent loops. When the Hop-Count field of a TRILL packet is set to 0, the RB discards the packet.
- Egress RB Nickname: In a unicast packet, the field indicates the nickname of the egress RB. In a multicast packet, the field indicates the nickname of the multicast tree root used for forwarding.
- Ingress RB Nickname: nickname of the ingress RB.
- Options: This field is available only when the value of Op-Length is not 0.

For details on how TRILL data packets are forwarded on a TRILL network, see **TRILL Forwarding Process**.

2.3 TRILL Mechanism

On a TRILL network, RBs must complete the following steps to communicate with each other:

1. **Establishing TRILL Neighbor Relationships**

2. **Synchronizing LSDBs**
3. **Calculating Routes**

Establishing TRILL Neighbor Relationships

TRILL devices send Hello packets (TRILL Hello PDUs) to establish neighbor relationships. Because of different port types, the Hello packets sent on broadcast and P2P links are different; however, the processes of establishing a neighbor relationship over these links are similar. **Figure 2-9** illustrates the process of establishing a neighbor relationship between two RBs.

Figure 2-9 Process of establishing a TRILL neighbor relationship

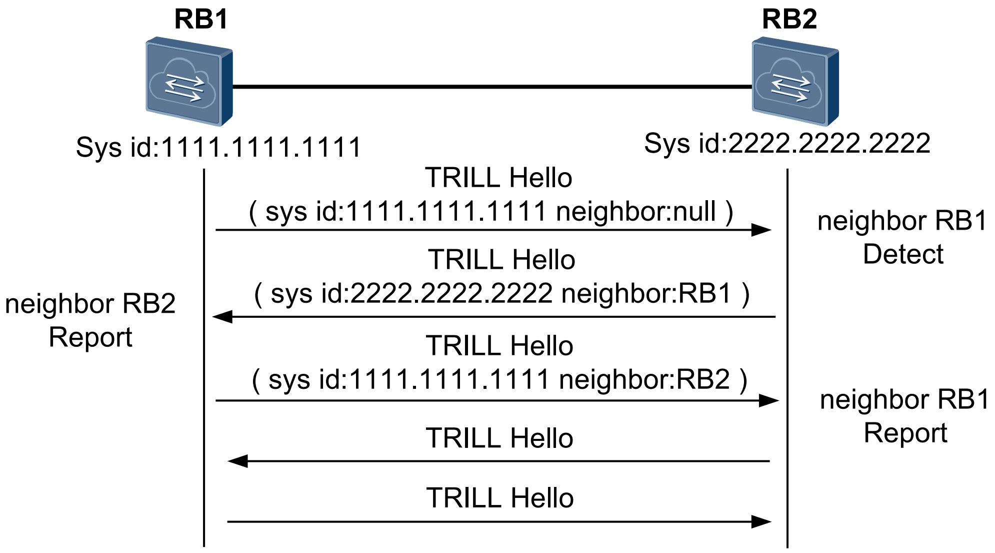

The process of establishing a neighbor relationship between two RBs on a TRILL network is as follows:

1. RB1 sends a TRILL Hello packet. After receiving the packet, RB2 detects that the neighbor field does not contain the local MAC address and sets the status of neighbor RB1 to **Detect**.
2. RB2 replies with a TRILL Hello packet. After receiving the TRILL Hello packet, RB1 detects that the neighbor field contains the local MAC address and sets the status of neighbor RB1 to **Report**.
3. RB1 sends another TRILL Hello packet to RB2. After detecting that the neighbor field contains the local MAC address, RB2 sets the status of neighbor RB1 to **Report**. The neighbor relationship has been set up between the two RBs.
4. The two RBs periodically send Hello packets to each other to maintain the neighbor relationship. If an RB fails to receive a response from the other after sending three Hello packets consecutively, the RB considers the neighbor relationship invalid.

To improve the convergence rate and communication efficiency on broadcast networks, TRILL introduces the following mechanisms:

- Electing a DRB

  On a broadcast network, each two RBs need to exchange information. If there are *n* RBs on the network, *n* x (*n*-1)/2 adjacencies need to be established. Once the status of any RB changes, information must be transmitted many times, resulting in a waste of bandwidth. To address the problem, TRILL defines a DRB. All the RBs send information only to the DRB, which then broadcasts the link status to the network. DRB election begins after the TRILL neighbor state turns to **Detect**.

A DRB is elected according to the following rules:

a. The RB whose interface has a higher DRB priority is elected as the DRB.

b. If the interfaces on the two ends have the same DRB priority, the RB with a larger MAC address is elected as the DRB.

- Specifying an AF

When unknown unicast or multicast traffic passes a TRILL network, a loop may occur because traffic is broadcast in a VLAN. As shown in **Figure 2-10**, multicast traffic from user A is forwarded to the TRILL network by a Layer 2 switch. If RB1 and RB3 belong to the same VLAN, the multicast traffic is forwarded to both the two RBs, and therefore a loop occurs. An AF can be specified to solve this problem. An AF is elected for each CE VLAN by the DRB. Only the AFs can function as ingress and egress RBs, and non-AFs can only be transit RBs. If RB1 in **Figure 2-10** is specified as the AF, user traffic is forwarded by RB1 and does not pass RB3, so loops will not occur.

Figure 2-10 Networking for AF selection

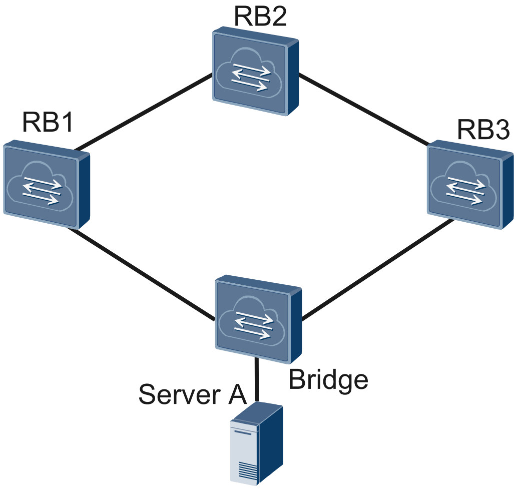

AFs are specified by the DRB. The DRB checks the CE VLANs enabled on the two ingress RBs on the TRILL network. The RB with the same VLAN ID as the user traffic is specified as the AF. If multiple RBs have the same VLAN ID as the user traffic, the AF is elected according to the following rules:

a. The RB with the highest DRB priority is elected as the AF.

b. If DRB priorities are the same, the RB with the largest MAC address is elected as the AF.

c. If the MAC addresses are the same, the RB with the largest port ID is elected as the AF.

d. If the port IDs are the same, the RB with the largest system ID is elected as the AF.

NOTE

- An RB can be specified as an AF only when its connected TRILL ports are access or hybrid ports.
- On a broadcast network, if two RBs have the same nickname, neither of them can be the AF.
- If the DRB is changed, all AF information is deleted and a new AF is elected.

- Specifying a DVLAN

  When multiple carrier VLANs exist on a TRILL network, a DVLAN must be specified on network interfaces of RBs to forward traffic. When sending TRILL protocol packets or forward TRILL data packets, RBs sets the VLAN ID in the Ethernet frame header to the DVLAN on the transmission link. A DVLAN can be manually configured or specified by a DRB.

Synchronizing LSDBs

After a DRB is elected, the LSDBs maintained by all RBs on the network are synchronized. An LSDB is the basis for generating a forwarding table. Therefore, LSDB synchronization is essential to correct data traffic forwarding on the network. The LSDB synchronization process varies depending on the network type.

- Figure 2-11 shows the LSDB update process on a broadcast link.

Figure 2-11 LSDB update on a broadcast link

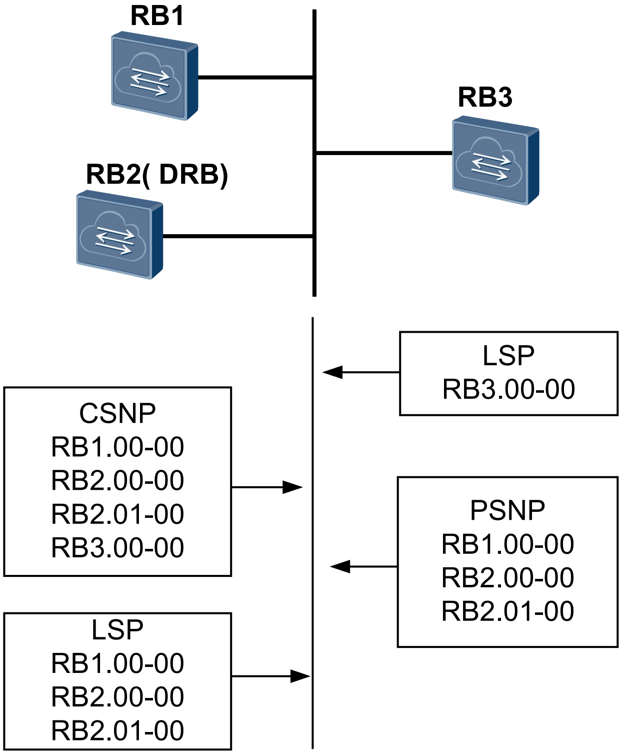

a. A newly added switch RB3 sends Hello packets to establish neighbor relationships with the other switches in the broadcast domain.

b. After neighbor relationships are set up, RB3 sends its own LSP to the multicast address 01-80-C2-00-00-41. All neighbors on the network receive the LSP.

c. The DRB in the network segment adds the LSP received from RB3 to its LSDB. After the **CSNP** timer expires, the DRB sends CSNPs to synchronize the LSDBs on the network. By default, CSNP timer is 10 seconds.

d. After receiving a CSNP from the DRB, RB3 checks its LSDB and sends a **PSNP** to request for the LSPs it does not have.

e. After receiving the PSNP, the DRB sends the requested LSPs to RB3. RB3 then synchronizes its LSDB with the LSPs. During LSDB update, the DRB performs the following operations:

   i. The DRB receives the LSP and searches for the matching record in the LSDB. If no matching record exists, the DRB adds the LSP to the LSDB and broadcasts the new LSDB.

   ii. If the sequence number of the received LSP is greater than the sequence number of the corresponding LSP in the LSDB, the DRB replaces the local LSP with the received LSP, and broadcasts the new LSDB.

   iii. If the sequence number of the received LSP is smaller than the sequence number of the corresponding LSP in the LSDB, the DRB sends the local LSP to the inbound interface.

   iv. If the sequence number of the received LSP is the same as the sequence number of the corresponding LSP in the LSDB, the DRB compares the remaining lifetime of the two LSPs. If the remaining lifetime of the received LSP is smaller than the remaining lifetime of the corresponding LSP in the LSDB, the DRB replaces the local LSP with the received LSP and broadcasts the new LSDB. If the remaining lifetime of the received LSP is larger than the remaining lifetime of the LSP in the LSDB, the DRB sends the local LSP to the inbound interface.

   v. If the sequence number and the remaining lifetime of the received LSP are the same as those of the corresponding LSP in the LSDB, the DRB compares the checksum values of the two LSPs. If the checksum of the received LSP is larger than the checksum of the LSP in the LSDB, the DRB replaces the local LSP with the received LSP and broadcasts the new LSDB. If the checksum of the received LSP is smaller than the checksum of the LSP in the LSDB, the DRB sends the local LSP to the inbound interface.

   vi. If the sequence number, remaining lifetime, and checksum of the received LSP are the same as those of the corresponding LSP in the LSDB, the DRB discards the LSP.

- **Figure 2-12** shows the LSDB update process on a P2P link.

Figure 2-12 LSDB update on a P2P link

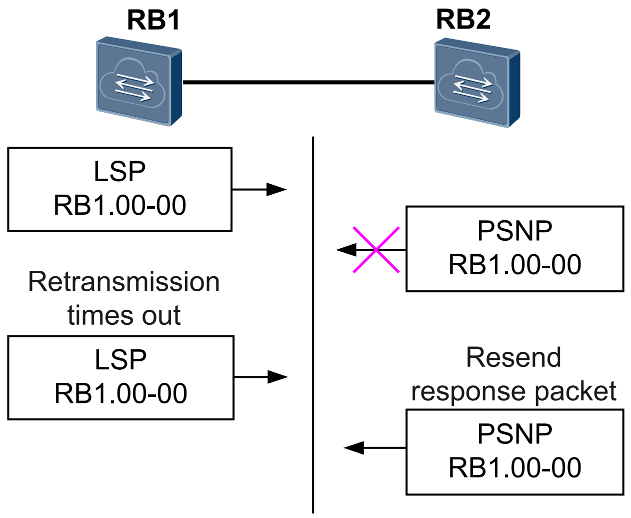

a. After a P2P neighbor relationship is set up, RB1 and RB2 exchange CSNPs to synchronize their LSDBs. In the following example, RB1 sends a CSNP to RB2. If the LSDB on RB2 is not synchronized with the CSNP, RB2 sends a PSNP to request for the corresponding LSP.

b. RB1 sends the required LSP to the neighbor. Meanwhile, it starts the LSP retransmission timer and waits for a PSNP from the neighbor as an acknowledgement of LSP reception. If RB1 does not receive the PSNP from the neighbor when the LSP retransmission timer expires, it resends the LSP.

c. After receiving the PSNP from the neighbor, RB1 performs the following operations:

   i. If the sequence number of the received LSP is greater than the sequence number of the corresponding LSP in the LSDB, RB1 adds the LSP to its LSDB, and then sends a PSNP to acknowledge the received LSP. After that, RB1 sends the LSP to all its neighbors except the neighbor that sends the LSP.

   ii. If the sequence number of the received LSP is smaller than the sequence number of the corresponding LSP in the LSDB, RB1 directly sends its LSP to the neighbor and waits for a PSNP from the neighbor.

   iii. If the sequence number of the received LSP is the same as the sequence number of the corresponding LSP in the LSDB, RB1 compares the remaining lifetime of the two LSPs. If the remaining lifetime of the received LSP is smaller than the remaining lifetime of the LSP in the LSDB, RB1 replaces the local LSP with the received LSP, sends a PSNP, and sends the LSP to all neighbors except the neighbor that sends the LSP. If the remaining lifetime of the received LSP is larger than the remaining lifetime of the LSP in the LSDB, RB1 sends the local LSP to the neighbor and waits for a PSNP.

   iv. If the sequence number and the remaining lifetime of the received LSP are the same as those of the corresponding LSP in the LSDB, RB1 compares the checksums of the two LSPs. If the checksum of the received LSP is larger than the checksum of the LSP in the LSDB, RB1 replaces the local LSP with the received LSP, sends a PSNP, and sends the LSP to all neighbors except the neighbor that sends the LSP. If the checksum of the received LSP is smaller than the checksum of the LSP in the LSDB, RB1 sends the local LSP to the neighbor and waits for a PSNP.

v. If the sequence number, remaining lifetime, and checksum of the received LSP are the same as those of the corresponding LSP in the LSDB, RB1 discards the LSP.

Calculating Routes

When the LSDBs maintained by all the RBs on a TRILL network are synchronized (that is, convergence is implemented), each RB uses the shortest path first (SPF) algorithm to calculate the unicast and multicast forwarding tables based on the LSDB. The calculation process is as follows:

- Generating a unicast routing table: Each RB uses itself as the root to generate the shortest paths to other nodes. Based on neighbor information, the RB obtains the outbound interface and next hop to each neighboring node, and generates a nickname unicast forwarding table according to the nickname information advertised by the neighbors.
- Generating a multicast routing table: To facilitate multicast traffic transmission, more than one multicast distribution tree is generated on a TRILL network. The generation process is as follows:

  a. Electing a root RB: Based on the root priorities of the nicknames advertised by all the devices on the entire network and the number of distribution trees supported, each device obtains the nickname with the highest root priority and the smallest number of distribution trees. The RB whose nickname has the highest root priority is elected as the root RB. If RBs have the same root priority, the RB with a larger system ID is elected as the root RB.

  b. Electing a distribution tree root: The root RB can specify roots of multicast distribution trees. If no root is specified, *N* RBs with the highest nickname root priorities are selected as the roots.

  c. Calculating a shortest path tree (SPT): *N* roots are used as source nodes to calculate the shortest path tree to all the other nodes on the entire network.

  d. Generating a reverse path forwarding (RPF) check table: The RPF check table is created based on the spanning tree information advertised by each ingress RB. The RPF check table is used to prevent loops.

  e. Pruning the SPT: The SPT is pruned based on information advertised by each ingress RB.

NOTE

Other nodes must have reachable unicast routes to the node with the highest nickname root priority. Therefore, unicast route calculation must be completed before multicast route calculation.

2.4 TRILL Forwarding Process

On a TRILL network, RBs send Hello packets to each other to establish neighbor relationships and send LSPs to synchronize LSDBs. After that, all RBs on a network have the same LSDB. Based on LSDB information, each RB then uses the SPF algorithm to calculate the shortest paths, outbound interfaces, and next hops to other RBs on the entire network. The RB uses the advised nickname information in the LSDB to generate a nickname forwarding table.

When user packets reach a TRILL network, they are forwarded in different processes based on the destination MAC addresses:

- Known unicast addresses: **Process of Forwarding Known Unicast Traffic**
- Unknown unicast, multicast, and broadcast addresses: **Process of Forwarding Unknown Unicast, Multicast, and Broadcast Traffic**

Process of Forwarding Known Unicast Traffic

Figure 2-13 illustrates how the known unicast traffic sent from server A to server C is forwarded.

Figure 2-13 Process of forwarding known unicast traffic

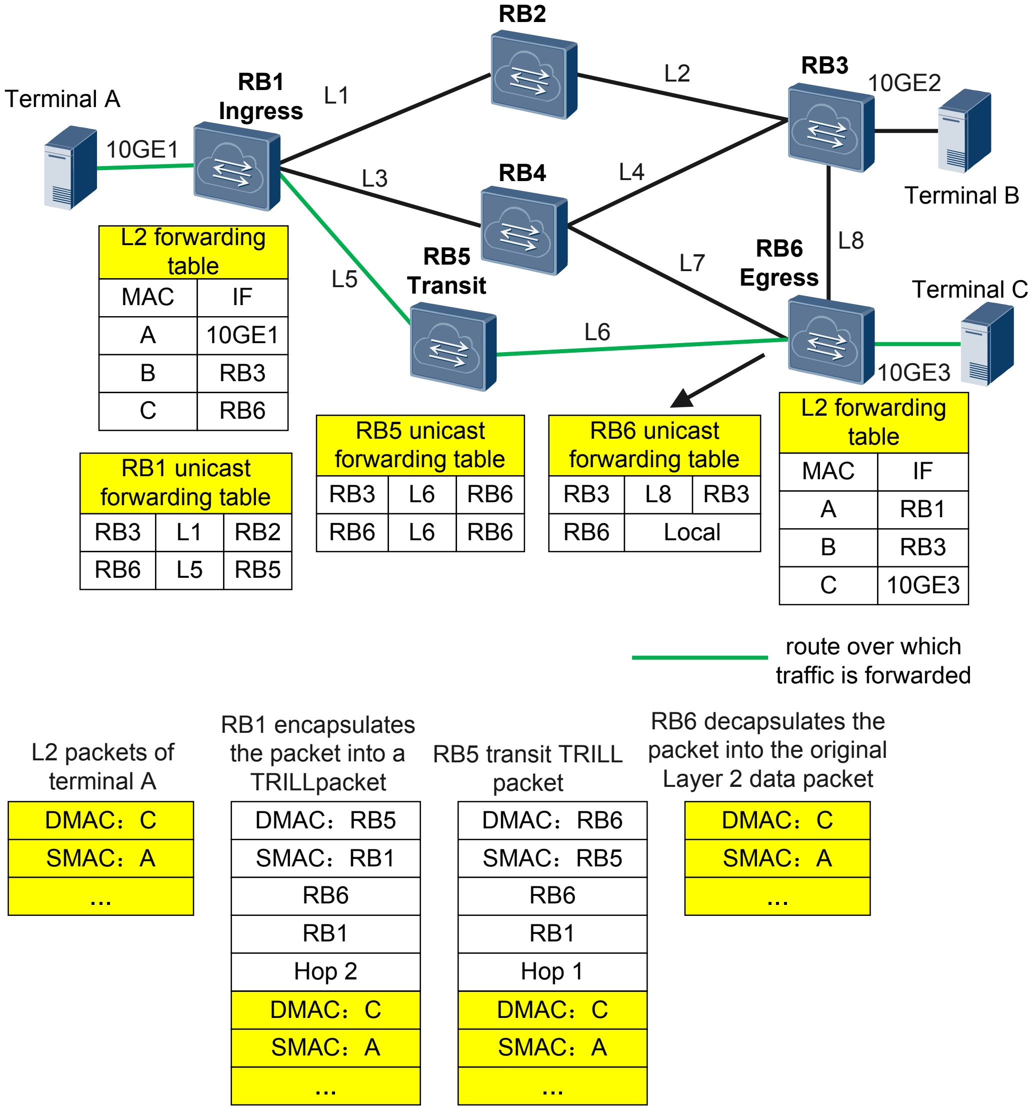

L2 forwarding table	
MAC	IF
A	10GE1
B	RB3
C	RB6

RB1 unicast forwarding table		
RB3	L1	RB2
RB6	L5	RB5

RB5 unicast forwarding table		
RB3	L6	RB6
RB6	L6	RB6

RB6 unicast forwarding table		
RB3	L8	RB3
RB6	Local	

L2 forwarding table	
MAC	IF
A	RB1
B	RB3
C	10GE3

L2 packets of terminal A	RB1 encapsulates the packet into a TRILLpacket	RB5 transit TRILL packet	RB6 decapsulates the packet into the original Layer 2 data packet
DMAC：C	DMAC：RB5	DMAC：RB6	DMAC：C
SMAC：A	SMAC：RB1	SMAC：RB5	SMAC：A
...	RB6	RB6	...
	RB1	RB1	
	Hop 2	Hop 1	
	DMAC：C	DMAC：C	
	SMAC：A	SMAC：A	
	...	...	

1. The ingress RB (RB1) receives a Layer 2 packet from server A, and searches the Layer 2 forwarding table for the egress RB nickname matching the destination MAC address of the packet. After finding the egress RB nickname, RB1 looks up the unicast forwarding table to find the outbound interface L5 and next hop RB5 to the destination RB. RB1 then encapsulates the Layer 2 packet into a TRILL data packet and forwards the packet to the next hop through the outbound interface.

2. When transit RB (RB5) receives the TRILL data packet, it obtains the egress RB nickname from the TRILL header and searches the unicast forwarding table for the egress RB nickname. Finding that the destination RB is RB6, RB5 forwards the TRILL data packet to RB6 through outbound interface L6.
3. The egress RB (RB6) receives the TRILL data packet, and finds that the egress RB nickname in the TRILL header is its own nickname. Then RB6 decapsulates the TRILL packet to obtain the original Layer 2 data packet, and forwards the Layer 2 data packet through the matching outbound interface according to the destination MAC address of the packet.

Process of Forwarding Multicast, Broadcast, and Unknown Unicast Traffic

On a TRILL network, an RB calculates a distribution tree for each VLAN based on the LSDB to guide the forwarding of multicast, broadcast, and unknown unicast packets. Multicast packets are used as an example to describe the forwarding process, as shown in **Figure 2-14**.

Figure 2-14 Process of forwarding multicast traffic

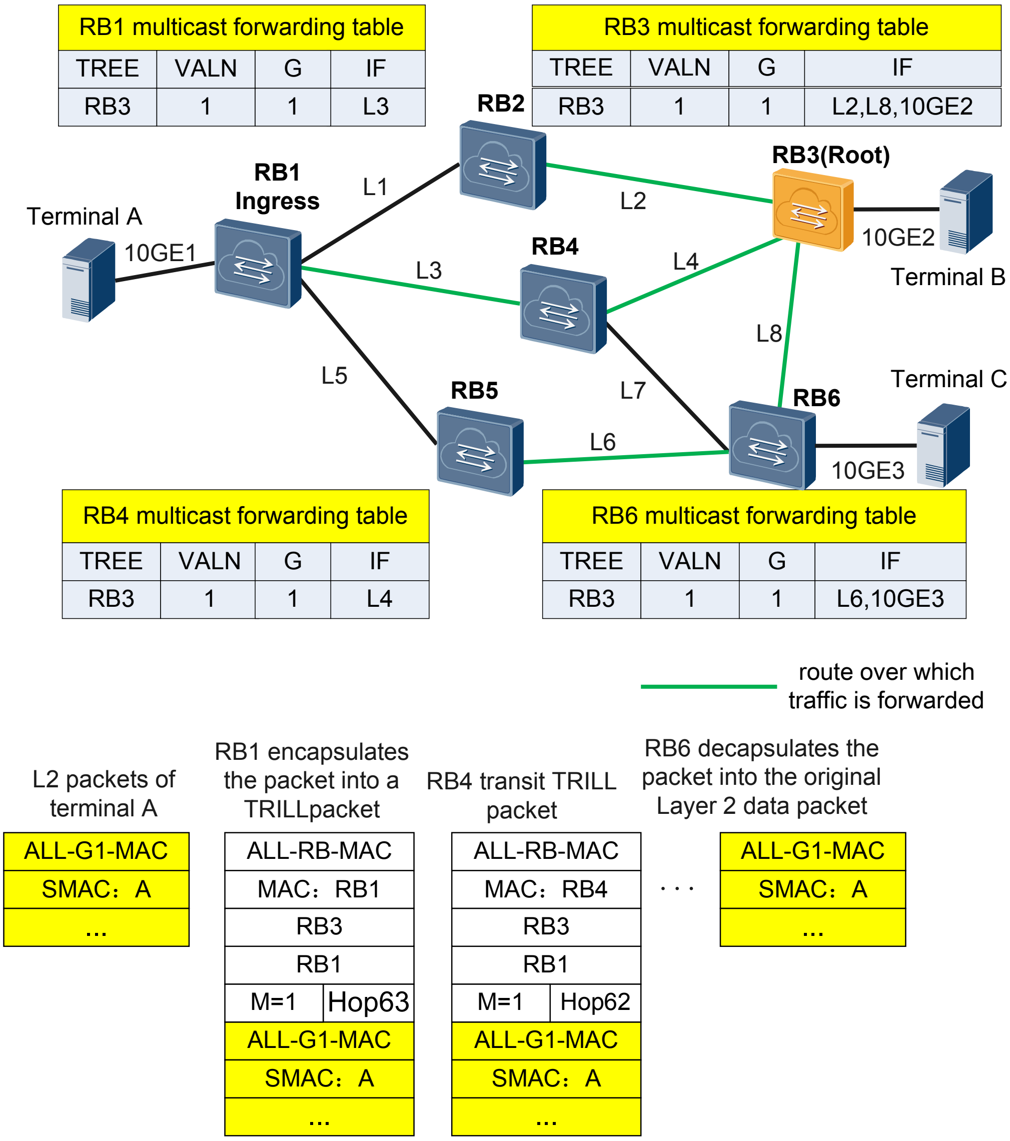

RB1 multicast forwarding table

TREE	VALN	G	IF
RB3	1	1	L3

RB3 multicast forwarding table

TREE	VALN	G	IF
RB3	1	1	L2,L8,10GE2

RB4 multicast forwarding table

TREE	VALN	G	IF
RB3	1	1	L4

RB6 multicast forwarding table

TREE	VALN	G	IF
RB3	1	1	L6,10GE3

L2 packets of terminal A

ALL-G1-MAC
SMAC：A
...

RB1 encapsulates the packet into a TRILLpacket

ALL-RB-MAC	
MAC：RB1	
RB3	
RB1	
M=1	Hop63
ALL-G1-MAC	
SMAC：A	
...	

RB4 transit TRILL packet

ALL-RB-MAC	
MAC：RB4	
RB3	
RB1	
M=1	Hop62
ALL-G1-MAC	
SMAC：A	
...	

. . .

RB6 decapsulates the packet into the original Layer 2 data packet

ALL-G1-MAC
SMAC：A
...

1. The ingress RB (RB1) receives a Layer 2 packet from server A, and finds that the destination MAC address is a multicast MAC address. RB1 selects the multicast distribution tree for the VLAN to which the packet belongs and encapsulates the Layer 2 packet in to a TRILL data packet. In the TRILL header, the M bit is set to 1, identifying a multicast packet. RB1 then looks up the multicast forwarding table according to the root RB nickname to find the outbound interface list, and forwards the TRILL packet to the outbound interface.
2. The transit RB (RB4) receives the TRILL data packet, and checks the TRILL header. As the M bit in the TRILL header is 1, RB4 looks up the multicast forwarding table according to the egress RB nickname in the TRILL header, and forwards the multicast packet to the outbound interfaces in the matching forwarding entry.
3. The root RB (RB3) receives the TRILL data packet, and distributes the packet to all the outbound interfaces.
4. The egress RB (RB6) decapsulates the TRILL data packet to obtain the original Layer 2 data packet and forwards the packet through the matching outbound interface.

2.5 TRILL Multi-Homing Active-Active Access

Background

On a Transparent Interconnection of Lots of Links (TRILL) network, access devices (such as switches and servers) are often dual-homed to TRILL devices to enhance reliability. When one TRILL device fails, services are not interrupted.

In this scenario, if you associate the VLAN appointed forwarder (AF) or MSTP with TRILL to eliminate loops, servers must connect to the TRILL network through Layer 2 access switches. This access mode also requires link redundancy backup, wasting bandwidth. You can use TRILL multi-homing active-active access to enable servers with dual NICs to be directly dual-homed to a TRILL network and forward traffic simultaneously. This access mode ensures reliability and fully utilizes network bandwidth.

Figure 2-15 TRILL multi-homing active-active access networking

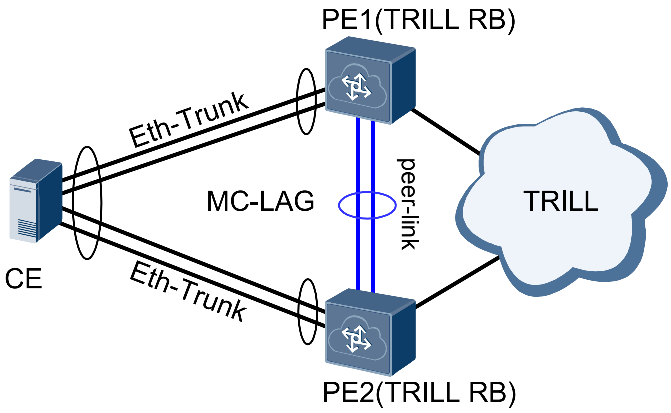

As shown in **Figure 2-15**, CEs are dual homed to a TRILL network. Access-side M-LAG is deployed to ensure device- and link-level reliability. Two dual-homed PEs (RBs on the TRILL network) use the same pseudo nickname. In this way, the peer end considers the PEs a logical device on the TRILL network.

Working Mechanism of Access-Side M-LAG

For details on working mechanism of access-side M-LAG, see Working Mechanism of M-LAG.

Working Mechanism of Network-Side TRILL

In a TRILL multi-homing active-active access solution, two RBs obtain pseudo nicknames through manual configuration or automatic negotiation and finally obtain the same pseudo nickname. RB1 and RB2 encapsulate pseudo nicknames into packets.

In addition, RB1 and RB2 check whether their pseudo nicknames are the same. If not, the E-Trunk on the two RBs does not take effect. If so, the two RBs exchange their actual nicknames and MAC addresses. When the peer-link or one RB becomes faulty, TRILL notifies the fault to the other RB in a timely manner to make the active-active solution ineffective.

Figure 2-16 TRILL multi-homing active-active access networking

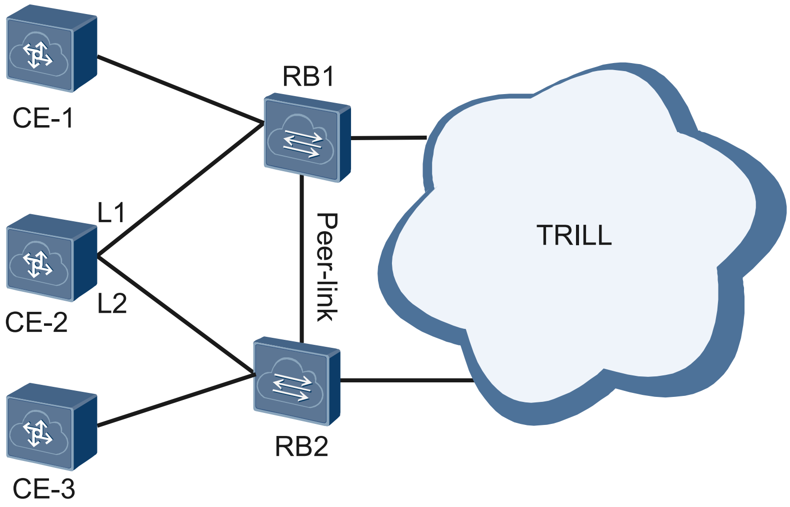

As shown in **Figure 2-16**, a peer-link is deployed between RB1 and RB2, and the two RBs have the same pseudo nickname. CE-2 together with RB1 and RB2 can implement a TRILL multi-homing active-active access solution. TRILL processes traffic of different types and from different directions differently. **Table 2-2** describes how TRILL processes these traffic types.

Table 2-2 Traffic processing in a TRILL multi-homing active-active access scenario

This object indicates traffic type.	Traffic Model	Traffic Processing
Unicast traffic from a non-active-active interface, for example, CE-1	**Figure 2-17** Unicast traffic from a non-active-active interface 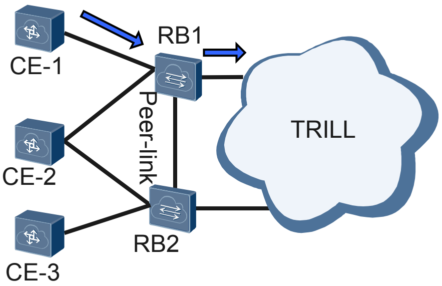	RB1 forwards the traffic in a unicast manner.
Unicast traffic from an active-active interface	**Figure 2-18** Unicast traffic from an active-active interface 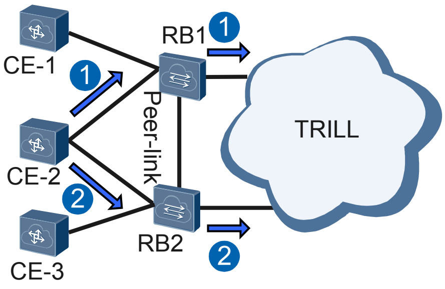 RB1 1 CE-1 1 TRILL Peer-link CE-2 2 RB2 2 CE-3	RB1 and RB2 work in load balancing mode to forward the traffic together. **NOTE** ① and ② represent different data flows.

This object indicates traffic type.	Traffic Model	Traffic Processing
Multicast traffic from a non-active-active interface, for example, CE-1	**Figure 2-19** Multicast traffic from a non-active-active interface 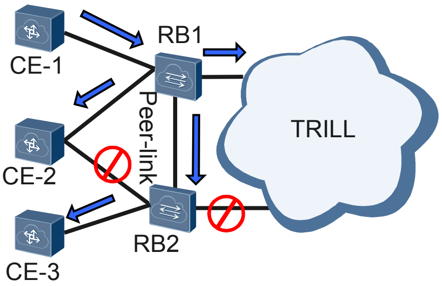	RB1 encapsulates pseudo nickname into the received multicast traffic and then forwards the traffic to each next-hop device. When the traffic arrives at RB2, RB2 forwards the traffic only to CE-3 but not to CE-2 or the TRILL network side according to the unidirectional isolation mechanism.

This object indicates traffic type.	Traffic Model	Traffic Processing
Multicast traffic from an active-active interface	**Figure 2-20** Multicast traffic from an active-active interface 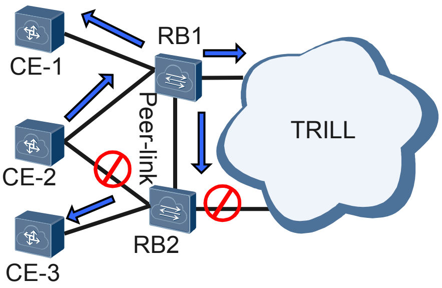	Multicast traffic from CE-2 is load balanced between RB1 and RB2. The following uses the forwarding process on RB1 as an example. RB1 encapsulates the pseudo nickname into the received multicast traffic and then forwards the traffic to each next-hop device. When the traffic arrives at RB2, RB2 forwards the traffic only to CE-3 but not to CE-2 or the TRILL network side according to the unidirectional isolation mechanism.

This object indicates traffic type.	Traffic Model	Traffic Processing
Unicast traffic from the TRILL network side	**Figure 2-21** Unicast traffic from the TRILL network side 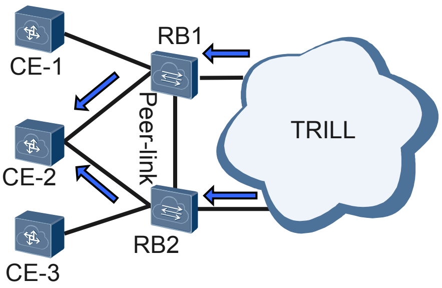 	If the unicast traffic is sent to an active-active interface, traffic is load balanced between RB1 and RB2 and then forwarded to the device that is dual-homed to the two RBs because the traffic is encapsulated with a pseudo nickname. The following uses the traffic sent to CE-1 as an example. Because the traffic is encapsulated with the pseudo nickname, the traffic is load balanced between RB1 and RB2. Traffic arriving at RB2 is sent to RB1 through the peer-link and then from RB1 to CE-1.

This object indicates traffic type.	Traffic Model	Traffic Processing
Multicast traffic from the TRILL network side	**Figure 2-22** Multicast traffic from the TRILL network side 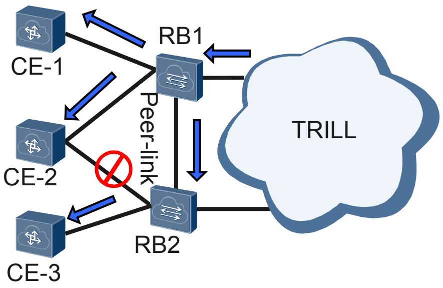 	According to the TRILL multicast forwarding principles, each RB joins a different multicast tree. Therefore, only one RB (RB1 or RB2) processes the multicast traffic. The following uses RB1 as an example. RB1 decapsulates the traffic and then forwards the traffic to each user-side interface. Because the peer-link is isolated from the backup interface, traffic arriving at RB2 is not forwarded to CE-2, avoiding routing loops.

2.6 TRILL NSR

During network operation, if an active/standby switchover is triggered due to a system failure or performed manually by the network administrator for software update or system maintenance, routes are interrupted and therefore traffic is lost. TRILL uses non-stop routing (NSR) to solve this problem.

NSR applies to devices with a backup control plane. This technology ensures that the control plane on a neighbor does not detect faults on the system control plane of the local device, preventing service interruption caused by an active/standby switchover.

NSR Implementation

TRILL NSR implements real-time data synchronization between the active and standby control planes as follows:

- TRILL backs up configuration data and dynamic data (information about interfaces, neighbors, and LSDBs).
- TRILL uses the RawLink socket to send and receive packets and does not back up the socket status.
- TRILL does not back up routes and shortest path trees (SPTs) that can be restored using the source data in the backup process.
- After an active/standby switchover occurs, the new master MPU restores the operation data and takes over services from the original master MPU. The neighbor is unaware of the active/standby process.

For detailed NSR description, see "NSR Configuration" in the *CloudEngine 12800 Series switch Configuration Guide - Reliability*.

3 Applications

About This Chapter

This section describes the application of TRILL in data centers.

3.1 Application of TRILL in Data Centers

Service Overview

Figure 3-1 shows the data center network of an enterprise.

Figure 3-1 Typical networking of a data center

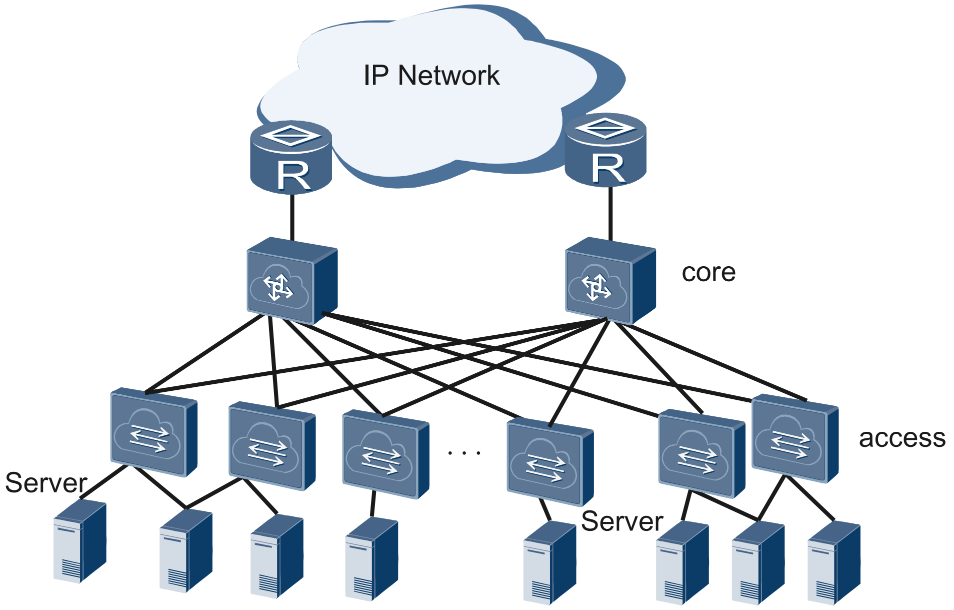

TRILL is used to build a flat Layer 2 network, which implements unblocked forwarding and smooth VM migration on the entire network. When deploying a data center networking using TRILL, **configure basic TRILL functions** on all devices and perform the following operations at different network layers:

- Access layer: Configure a CE VLAN on the user side to connect to the server so that user traffic is transmitted through the TRILL network. If the server is dual-homed to the TRILL network through the access device, on the edge RB connected to the access device to prevent loops. On the network side, you can **adjust TRILL route selection** and **adjust the TRILL network convergence speed** to ensure efficient forwarding.
- Core layer: On the network side, you can **adjust TRILL route selection** and **adjust the TRILL network convergence speed** to ensure efficient forwarding. On the egress side, you can connect to other enterprise networks through the egress router, or configure virtual systems (VSs) on the devices at the core layer and use one VS as the egress gateway to connect to other enterprise networks.

3.2 M-LAG Application in a TRILL Dual-Homing Network

Dual Homing a CE to a TRILL Network

As shown in **Figure 3-2**, M-LAG is deployed between RB1 and RB2 and nicknames of the two RBs are bound to a DFS group so that the CE is dual-homed to the TRILL network. RB1 and RB2 use the same pseudo nickname to form a logical device, and load balance traffic. When one access link or device fails, traffic can be fast switched to the other link or device.

Figure 3-2 Dual homing to a TRILL network through M-LAG

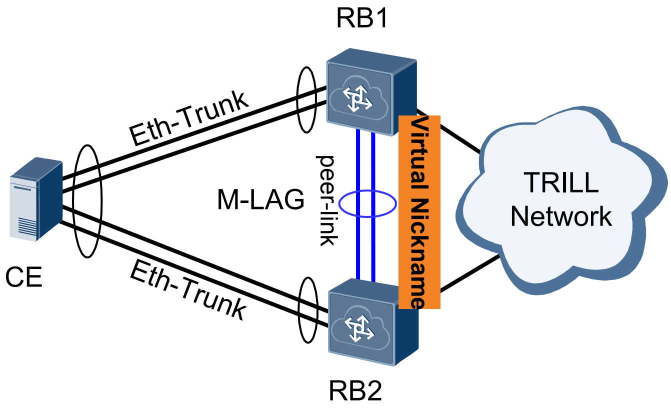

4 Configuration Notes

This section provides the points of attention when configuring TRILL.

Involved Network Elements

Other network elements are required to support Transparent Interconnection of Lots of Links (TRILL).

License Support

The TRILL function is controlled by a license. By default, this function is disabled on new purchased CE12800 series switches. To use the TRILL feature, apply for and purchase the license from the equipment supplier.

Version Support

Table 4-1 Products and minimum version supporting TRILL

Series	Device Model	Minimum Version Required
CE12800	CE12804/CE12808/ CE12812	V100R001C00
	CE12816	V100R003C00
	CE12804S/CE12808S	V100R005C00

Feature Dependencies and Limitations

TRILL conflicts with some other features:

- When VXLAN or EVN is configured on the device, TRILL cannot be configured. In V100R005C00, To configure TRILL, disable VXLAN or EVN first. Then restart the device to make the TRILL configuration take effect.
- In V100R005C00 and earlier versions, the TRILL function cannot be configured simultaneously with any of the FCoE, port security, MAC limit, URPF, DHCP snooping,

or 802.1x functions. In V100R005C10 and later versions, by default, the TRILL and preceding functions cannot be configured simultaneously. To use these functions simultaneously, run the **trill adjacency-check disable** command. The TRILL function takes precedence over the preceding functions. If the TRILL function is configured after the preceding functions are configured, only the TRILL function takes effect.

- When MPLS is configured on a device, TRILL cannot be enabled on the interfaces on the MPLS forwarding path. Otherwise, MPLS will not take effect.
- In V100R005C00, Leaf switches do not support the TRILL function, that is, they cannot function as access or network devices. Parent switches support TRILL.In V100R005C10, Leaf switches can only function as access devices. Parent switches support TRILL.
- TRILL can be configured in the admin-VS if there is no VS in port mode. No VS in port mode can be configured after TRILL is configured in the admin-VS. If VSs in port mode are configured, TRILL cannot be configured in any VS in port mode, including the admin-VS. All VSs in group mode support TRILL configuration.
- NetStream and sFlow cannot sample TRILL packets.
- VBST cannot be associated with TRILL.

Pay attention to the following points when configuring VLAN on a TRILL network:

- A carrier VLAN must be a new VLAN. A CE VLAN and admin VLAN must be the VLANs created using the **vlan** command and different from the carrier VLAN.
- A device can have a maximum of three carrier VLANs configured, and the carrier VLAN with the smallest VLAN ID defaults to be the DVLAN. You can also run the **trill designated-vlan** command to configure a DVLAN.
- When a user network connects to the TRILL network through AFs, ensure that AFs can exchange TRILL packets over the user network. You are advised to configure devices of the user network to allow packets with DVLAN ID to pass through.
- A VPLS network cannot connect to the TRILL network through CE VLAN.
- When the length of packets sent by user-side devices (CE VLAN) is smaller than 512 bytes, interfaces on a TRILL network may be unable to provide line-rate forwarding.
- In V100R005C00 and earlier versions, devices do not support the forwarding of TRILL packets at Layer 3 through routes after terminating these packets.
- When a user network connects to the TRILL network through AFs, admin VLAN must be configured on the TRILL network and an IP address must be configured for the VLANIF interface of the admin VLAN.
- A TRILL CE VLAN or admin VLAN cannot be configured as a super VLAN.

On a TRILL network, the port type defaults to **p2p**. The port type configuration rules are as follows:

- When a port is at the edge of the TRILL network and connects to a user VLAN, the port type is usually set to access.
- When a port is in the middle of the TRILL network and transmits only TRILL packets, the port type is usually set to trunk.
- On a P2P network, the ports between two RBs are generally set as P2P ports.
- If a port needs to connect to a user VLAN and transmit TRILL packets, the port type is usually set to hybrid.

When configuring association between TRILL and MSTP (changing the root bridge), pay attention to the following points:

- When the same bridge MAC address is configured for two devices using the **stp bridge-address** command, to simulate the two devices into one root bridge, ensure that all STP configurations of the two devices are the same, including the device priority and timer parameters.
- Before configuring the **stp tc-notify trill vlan** *vlan-id* command on a device, you must configure the **stp disable** command on the device's interface that has the **trill enable port-mode** { **hybrid** | **p2p** | **trunk** } command configured. Otherwise, TC packets may be looped.
- If association between STP/RSTP/MSTP and TRILL is configured on a VS of a device, you are advised to run the **carrier up-hold-time** *interval* command on all the interfaces that connect a TRILL network to an MSTP network to set the delay in reporting a port Up event to 50s.

When configuring association between TRILL and MSTP (retaining the root bridge), pay attention to the following points to ensure that STP packets are transparently transmitted over a TRILL network:

- An admin VLAN must be configured on the TRILL network and an IP address must be configured for the VLANIF interface of the admin VLAN.
- The same PVID must be configured for the interfaces that have the **stp tc-snooping notify trill** command configured, and the PVID is the admin VLAN ID of the TRILL network.
- On an STP/RSTP/MSTP network, the devices that connect the MSTP network to the TRILL network must be configured as root bridges.
- On an STP/RSTP/MSTP network, ensure that the interface that connects the STP/RSTP/MSTP network to the TRILL network is not blocked. To prevent the interface from being blocked, set the path cost of the interface to the smallest value among those of other interfaces.
- You are advised to run the **stp edged-port disable** command on the interface that connects an STP/RSTP/MSTP network to the TRILL network to configure the interface as a non-edge interface.

When configuring TRILL dual-homing access through an E-Trunk/M-LAG, pay attention to the following points:

- In V100R003C00 and V100R003C10, only the following cards support TRILL dual-homing access through an E-Trunk: CE-L48GT-EA, CE-L48GT-EC, CE-L48GS-EA, CE-L48GS-EC, CE-L24XS-BA, CE-L24XS-EA, CE-L48XS-BA, CE-L48XS-EA and CE-L24LQ-EA.
- In V100R005C00 and later versions, only the following cards support TRILL dual-homing access through an M-LAG: CE-L48GT-ED, CE-L48GS-ED, CE-L12XS-ED, CE-L24XS-EC, CE-L24XS-ED, CE-L48XS-EC, CE-L48XS-ED, CE-L48XS-EF, CE-L12LQ-EF, CE-L24LQ-EC, CE-L04CF-EF, CE-L48XT-EC, CE-L02LQ-EC, CE-L24LQ-EC1, CE-L36LQ-EG, CE-L08CC-EC, CE-L12CF-EG and CE-L06LQ-EC.
- To deploy TRILL dual-homing access through an E-Trunk, ensure that all devices on the TRILL networks run the software version of V100R003C00 or V100R003C10.
- To deploy TRILL dual-homing access through an M-LAG, ensure that all devices on the TRILL networks run the software version of V100R005C00 or later.
- In V100R005C00 and later versions, TRILL dual-homing access through an M-LAG conflicts with association between STP/RSTP/MSTP and TRILL, so the two functions cannot be configured simultaneously.

- In a TRILL active-active scenario, peer RBs cannot be upgraded using ISSU.

When configuring a TRILL gateway, pay attention to the following points:

- The TRILL gateway in internal loopback mode cannot forward multicast traffic at Layer 3.
- After the TRILL gateway function is configured, the administrative VLAN does not support data traffic transmission.
- In V100R005C10, after the TRILL gateway in internal loopback mode is configured, the VLANIF interface of a CE VLAN cannot establish neighbor relationships with other devices' interfaces through a TRILL network. In V100R006C00 and later versions, the VLANIF interface of a CE VLAN can establish only OSPF neighbor relationships but not neighbor relationships of other routing protocols with other devices' interfaces through a TRILL network.
- After the TRILL gateway function is configured on the device, the device cannot be upgraded from V100R005C10SPC100 to V100R005C10SPC200 using ISSU in process-switchover mode.
- After the TRILL gateway function is configured, a traffic policy configured on the VLANIF interface of a CE VLAN matches only the inner IPv4 packets decapsulated from TRILL packets.
- A TRILL gateway cannot forward QinQ packets that enter a TRILL domain at Layer 3.
- During an upgrade from V100R005C10SPC100 to V100R005C10SPC200, if a VS in group mode has an M-LAG or the TRILL gateway function configured, the system does not prompt that ISSU is not supported in the ISSU check phase.
- After the TRILL gateway in internal loopback mode is configured on a device, the TRILL gateway in VLAN mapping mode cannot be configured on the device.
- A board that has the TRILL gateway function configured can process only a limited volume of Layer 3 forwarded traffic. Ports on each board can be divided into multiple port groups, and each port group can process at most 100G Layer 3 forwarded traffic. Port group division on each board type is as follows.

Table 4-2 Port groups on cards and ports in each port group

Port Rate on a Card	Card Model	Port Group and Port Assignment (Port Group: Port Number)
GE	● CE-L48GS-EA ● CE-L48GT-EA ● CE-L48GS-EC ● CE-L48GT-EC ● CE-L48GS-ED ● CE-L48GT-ED	Group 0: 0-47

Port Rate on a Card	Card Model	Port Group and Port Assignment (Port Group: Port Number)
10GE	● CE-L48XS-EA ● CE-L48XS-SA ● CE-L48XS-BA ● CE-L48XS-EC ● CE-L48XS-ED ● CE-L48XS-EF ● CE-L48XT-EC	Group 0: 24-47 Group 1: 0-23
	● CE-L24XS-EA ● CE-L24XS-BA ● CE-L24XS-EC ● CE-L24XS-ED	Group 0: 0-23
	CE-L12XS-ED	Group 0: 0-11
	● CE-FWA ● CE-IPSA	Group 0: 0-3
	CE-L48XS-FDA	Group 0: 0-51 (including 10GE, 40GE, and 100GE ports on cards)
40GE	● CE-L24LQ-EA ● CE-L24LQ-EC ● CE-L24LQ-EC1	Group 0: 0-5 Group 1: 6-11 Group 2: 12-17 Group 3: 18-23
	CE-L02LQ-EC	Group 0: 0-1
	CE-L06LQ-EC	Group 0: 0-5
	CE-L12LQ-EF	Group 0: 6-11 Group 1: 0-5
	CE-L36LQ-EG	Group 0: 0-5 Group 1: 6-11 Group 2: 12-17 Group 3: 18-23 Group 4: 24-29 Group 5: 30-35
	CE-L36LQ-FD	Group 0: 0-11 Group 1: 12-23 Group 2: 24-35

Port Rate on a Card	Card Model	Port Group and Port Assignment (Port Group: Port Number)
100GE	● CE-L04CF-EF ● CE-L04CF-EC	Group 0: 2-3 Group 1: 0-1
	CE-L08CC-EC	Group 0: 0-1 Group 1: 2-3 Group 2: 4-5 Group 3: 6-7
	CE-L12CF-EG	Group 0: 0-1 Group 1: 2-3 Group 2: 4-5 Group 3: 6-7 Group 4: 8-9 Group 5: 10-11
	CE-L36CQ-FD	Group 0: 0-5 Group 1: 6-11 Group 2: 12-17 Group 3: 18-23 Group 4: 24-29 Group 5: 30-35
Contiguous ports on a card are added to each port group in sequence. For example, a CE-L48XS-EA card has two port groups, and ports 10GE1/0/0 to 10GE1/0/23 are added to one port group, and ports 10GE1/0/24 to 10GE1/0/47 are added to the other port group.		

When configuring TRILL OAM, pay attention to the following points:

The input parameter must be consistent with the actual hash factor so that a correct outbound interface can be obtained.

In a stack, to view the TRILL unicast forwarding path or use the TRILL unicast trace function, specify the source interface.

There are other configuration restrictions on a TRILL network:

- NETs must be unique on the TRILL network. If a NET is not unique, route flapping may occur.
- Manually changing the nickname of a device will interrupt services on the TRILL network temporarily. Therefore, confirm your operation before changing the nickname. Nicknames of devices on the TRILL network must be unique.
- After devices on a TRILL network are configured to perform pruning based on multicast groups, multicast data within the TRILL network can only be forwarded based on MAC addresses. Multiple IPv4 multicast addresses may be mapped to the same IPv4 multicast MAC address according to the multicast IP-and-MAC address mapping mechanism. When multicast data is forwarded based on MAC addresses and a group IP address for

receivers and the multicast IP address reserved for a protocol are mapped to the same IP multicast MAC address, the protocol cannot run normally. For example, IP multicast address 224.0.0.5 is reserved for the OSPF protocol. If a multicast group uses IP multicast address 225.0.0.5, the two IP multicast addresses are both mapped to IP multicast MAC address 01-00-5E-00-00-05. In this case, the OSPF protocol cannot run normally. Therefore, a proper IP multicast address plan must be made to prevent this problem.

- During an upgrade from V100R005C00SPC300 to V100R005C10SPC200 using ISSU in process-switchover mode, multicast traffic on the TRILL network will be lost.

5 Configuration Task Summary

After the basic TRILL functions are configured, a TRILL network can be constructed. If other TRILL functions are required, configure them according to reference sections.

Table 5-1 describes the RIP configuration tasks.

Table 5-1 TRILL configuration tasks

Scenario	Description	Task
Configuring basic TRILL functions	The following functions can be configured only when the basic TRILL functions are enabled. **NOTE** On the current CE series switches, Layer 3 route forwarding cannot be performed after TRILL packets are terminated.	**Configuring Basic TRILL Functions**

Scenario	Description	Task
Adjusting TRILL route selection	After the basic TRILL functions are configured, each node in the network can communicate with each other using TRILL. The unicast and multicast forwarding tables are generated through TRILL based on the LSDBs to guide the unicast and multicast traffic forwarding. However, on a large network, only the protocol mechanism cannot meet the network planning and traffic management requirements. As TRILL uses the SPF algorithm to calculate unicast and multicast routing tables, some links may be set idle due to high costs. Meanwhile, some links with low costs are too busy to load traffic and load balancing cannot be performed. This results in the network resource waste and affects the network transmission quality. Therefore, to optimize TRILL networks, route selection must be adjusted for accurate network control.	**Adjusting TRILL Route Selection**
Adjusting the TRILL network convergence speed	The network convergence speed determines the network quality. Although TRILL supports fast convergence, it always applies to large data center networks. A complex network slows down the convergence speed. In this situation, the network convergence speed can be manually increased.	**Adjusting the TRILL Network Convergence Speed**

Scenario	Description	Task
Configuring the association between STP and TRILL	You are advised to configure the association between STP/RSTP/MSTP and TRILL on edge devices connecting TRILL networks to STP/RSTP/MSTP networks.	**6.4 Configuring the Association Between STP/ RSTP/MSTP and TRILL**
Configuring TRILL network dual-homing through an M-LAG **NOTICE** ● To deploy the TRILL network dual-homing through an M-LAG, ensure that all devices on the TRILL networks run the software version of DCV100R005C00 or a later version. ● TRILL network dual-homing through an M-LAG and association between STP/RSTP/MSTP and TRILL are mutually exclusive. Do not configure the two functions at the same time.	In a dual-homing access scenario, if the VLAN appointed forwarder (AF) or MSTP is associated with TRILL to eliminate loops, servers must connect to the TRILL network through Layer 2 access switches. This access mode also requires link redundancy backup, causing a waste of bandwidth. You can configure servers to be dual-homed to the TRILL network through M-LAG. The servers then forward traffic simultaneously. This access mode ensures reliability and fully utilizes network bandwidth.	**Configuring TRILL Network Dual-Homing Through an M-LAG**

Scenario	Description	Task
Improving TRILL network security	With development of the Internet, more and more data, voice, and video information is exchanged over networks, and most of these services require high security. TRILL authentication is an encryption method based on network security requirements. It encrypts TRILL packets by adding the authentication field to the packets. When the local RB receives TRILL packets sent from a remote RB, if the authentication passwords are different from the local configuration, the local RB discards the packets to implement self-protection. TRILL supports the following authentication modes: ● simple: supports plain-text authentication, requires simple configuration, and applies to networks with lower security requirements. ● MD5: supports plain-text or cipher-text authentication, requires simple configuration, and applies to networks that require short-time encryption. A single password is generated after this mode is configured, and the password can be changed only manually. ● Keychain: provides an enhanced encryption algorithm and allows users to define a group of passwords as a password string. An encryption/	**Improving TRILL Network Security**

Scenario	Description	Task
	decryption algorithm and a validity period are defined for each password. The keychain algorithm is complex to configure. Keychain authentication allows automatically change of a password based on the configuration. Therefore, keychain authentication is applicable to the network requiring high security. ● hmac-sha256: uses the hmac-sha256 algorithm.	

6 TRILL Configuration

About This Chapter

Transparent Interconnection of Lots of Links (TRILL) uses Intermediate System to Intermediate System (IS-IS) to enable communication over Layer 2 networks. TRILL boasts easy deployment, highly-efficient forwarding, high-speed convergence, and loop prevention. Therefore, it meets data centers' requirements for large-scale Layer 2 networks.

6.1 Configuring Basic TRILL Functions

6.1.1 Enabling TRILL Globally

Context

When configuring the TRILL protocol, enable TRILL globally and then configure other TRILL features. To facilitate maintenance and management, you can configure the following additional TRILL functions:

- Admin VLAN: allows administrators to manage routing bridges (RBs) using remote login.
- Network entity title (NET) or dynamic host name: TRILL uses the MAC address of each RB as its system ID by default, and the MAC address is difficult to memorize.
- Nickname: identifies an RB and is automatically generated by default. To prevent nickname conflicts and facilitate check and management, you can manually set a nickname for an RB.
- Port mode: There are four port types for RBs. Based on the RB roles in the network, you can configure specified types for all RB ports to improve the network service efficiency.

NOTE

- After TRILL is enabled globally, the port security and URPF functions configured on the device do not take effect.
- When a VLAN is created on the ingress or egress node, the transmit node creates the corresponding VLAN and learns the MAC address to ensure normal packet forwarding. You do not need to pay attention to this process.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
trill
```

TRILL is enabled globally and the TRILL view is displayed.

Step 3 Run:

```
carrier-vlan carrier-vlanid
```

The carrier VLAN ID is specified.

NOTE

A created VLAN cannot be configured as the carrier VLAN.

When a user network accesses a TRILL network through AFs, you must ensure that AFs can exchange TRILL packets through the user network. You are advised to configure the related devices to allow packets from a DVLAN to pass through. If only one carrier VLAN is configured, the carrier VLAN is the DVLAN by default. If multiple carrier VLANs are configured, the carrier VLAN with the smallest VLAN ID is the DVLAN by default. You can also run the **trill designated-vlan** command to configure a DVLAN.

Step 4 (Optional) Run:

```
network-entity net
```

The NET is configured.

NOTE

The NET must be unique on the network. If the NETs conflict, route flapping may occur. Therefore, the parameter should be planned before you perform the operations.

Step 5 (Optional) Run:

```
nickname nicknamevalue [ priority priorityvalue ] [ root-priority
rootpriorityvalue ]
```

The nickname, priority, and root priority is configured.

⚠ **NOTICE**

On the TRILL network, manually changing the nickname of a device will interrupt running services temporarily. Therefore, confirm your action before changing the nickname.

NOTE

The nickname must be unique on the network. If nicknames conflict, the TRILL protocol performs the following operations:

- If a configured nickname conflicts with a generated nickname, the RB with a lower priority generates a new nickname.
- If a configured nickname conflicts with another configured nickname, the nickname of the RB with a lower priority is suppressed, and is not advertised to the network.

Step 6 (Optional) Run:

```
port-mode { access | hybrid | p2p | trunk }
```

The port mode is configured.

NOTE

By default, the port mode is **p2p**. The port mode configuration rules are as follows:

- When a port is located at the edge of a TRILL network to connect to a user VLAN, the port mode is **access**.
- When a port is located in the middle of a TRILL network to transmit TRILL packets, the port mode is **trunk**.
- On a P2P network, the mode of the port between two RBs is **p2p**.
- If you require that a port should connect to a user VLAN and transmit TRILL packets, the port mode is **hybrid**.

Step 7 (Optional) If an RB is located at the edge of a TRILL network to connect to a server and the TRILL network, configure a CE VLAN as follows:

1. Run the **quit** command to exit from the TRILL view.
2. Run the **vlan** *vlan-id* command to create a VLAN and enter the VLAN view.
3. Run the **quit** command to exit from the VLAN view.
4. Run the **trill** command to enter the TRILL view.
5. Run the **ce-vlan** { *vlan-id1* [**to** *vlan-id2*] } & <1-10> command to specify a CE VLAN for the TRILL process.

NOTE

- The CE VLAN must have been created using the **vlan** command and must be different from the carrier VLAN.
- The CE VLAN cannot be connected to a VPLS network.
- When the length of the packets sent from the user side is less than 512 bytes, the interfaces on the TRILL network may fail to work at the line speed.

Step 8 (Optional) Run:

```
trill-name symbolic-name
```

The TRILL dynamic host name mapping is configured and the host name is set for the local device.

After this command is executed, the host name *symbolic-name* is advertised to other devices in the domain in LSP packets. When you check TRILL information using the related display commands on other devices, if the dynamic host name mapping is also configured on these devices, the system ID is replaced with *symbolic-name*.

Step 9 (Optional) To facilitate management on the TRILL network, configure the admin VLAN for an RB as follows:

1. Run the **quit** command to exit from the TRILL view.
2. Run the **vlan** *vlan-id* command to create a VLAN and enter the VLAN view.
3. Run the **quit** command to exit from the VLAN view.
4. Run the **trill** command to enter the TRILL view.
5. Run the **admin-vlan** *vlan-id* command to specify the admin VLAN ID for the TRILL process.

   NOTE

   The admin VLAN must have been created using the **vlan** command and must be different from the carrier VLAN.

   When a user network accesses a TRILL network through AFs, the admin VLAN must be configured on the TRILL network and an IP address must be configured for the VLANIF interface of the admin VLAN.
6. Run the **quit** command to exit from the TRILL view.
7. Run the **interface vlanif** *vlan-id* command to create a VLANIF interface and enter the interface view.
8. Run the **ip address** *ip-address* { *mask* | *mask-length* } [**sub**] command to configure an IP address for the VLANIF interface.

   NOTE

   The IP address of the VLANIF interface must be on the same network segment as the IP address of the network edge device to ensure that the nickname is reachable.

Step 10 Run:

```
commit
```

The configuration is committed.

----End

6.1.2 Enabling TRILL on an Interface

Context

After TRILL is enabled globally, TRILL neighbor relationships are not established between the RBs, and TRILL must be enabled on interfaces.

TRILL can be enabled only on trunk or hybrid interfaces. An access or dot1q-tunnel interface can only function as a user access interface and cannot run TRILL.

There are four port modes of TRILL, and the default port mode is **p2p**. Different port modes can be configured for interfaces of different roles based on their network locations, which reduces the number of packets processed by the interfaces and saves system resources.

In a broadcast network, the DRB communicates with each device on the network. Therefore, the DRB must have high performance. You can configure proper DRB priorities for interfaces to enable a high-performance device to be elected as the DRB.

NOTE

An interface can establish neighbor relationships with at most eight interfaces.

If MPLS functions have been configured on the device, interfaces on the MPLS forwarding path cannot have TRILL functions enabled. Otherwise, MPLS functions do not take effect.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
interface interface-type interface-number
```

The interface view is displayed.

Step 3 Run:

```
port link-type { hybrid | trunk }
```

The link type of the interface is set to hybrid or trunk.

Step 4 Run:

```
trill enable [ port-mode { access | hybrid | p2p | trunk } ]
```

The TRILL function is enabled and the port mode is configured on the interface.

NOTE

By default, the port mode is **p2p**. The port mode configuration rules are as follows:

- When a port is located at the edge of a TRILL network to connect to a user VLAN, the port mode is **access**.
- When a port is located in the middle of a TRILL network to transmit TRILL packets, the port mode is **trunk**.
- On a P2P network, the mode of the port between two RBs is **p2p**.
- If you require that a port should connect to a user VLAN and transmit TRILL packets, the port mode is **hybrid**.

Step 5 (Optional) Run:

```
trill drb-priority priority
```

The DRB priority of the interface is configured.

By default, the DRB priority of an interface is 64. If you want to configure the local RB as a DRB, set the priority of the local RB to a large value.

Step 6 Run:

```
commit
```

The configuration is committed.

----End

6.1.3 (Optional) Configuring a Link Cost for a TRILL Interface

Context

Because TRILL uses the Shortest Path First (SPF) algorithm to calculate routing tables, the link cost is crucial to TRILL route selection. Adjusting link costs of TRILL interfaces directly affects route selection. The following link costs take effect on a TRILL interface in ascending order of priority:

- **Automatically calculated link cost**: The default calculation formula is as follows: Link cost of a TRILL interface = Bandwidth reference value/Interface bandwidth. You can change the bandwidth reference value to adjust the link cost of an interface.
- **Global link cost**: Configure a link cost for all TRILL interfaces on a specified RB.
- **Interface link cost**: Configure the link cost for a specified interface.

Procedure

- Adjusting the automatically calculated link cost
  a. Run:
  ```
  system-view
  ```
  The system view is displayed.
  b. Run:
  ```
  trill
  ```
  The TRILL view is displayed.
  c. Run:
  ```
  bandwidth-reference value
  ```
  The bandwidth reference value is adjusted.
  d. Run:
  ```
  commit
  ```
  The configuration is committed.
- Configuring a global link cost
  a. Run:
  ```
  system-view
  ```
  The system view is displayed.
  b. Run:
  ```
  trill
  ```

The TRILL view is displayed.

c. Run:

```
circuit-cost { cost | maximum }
```

A global link cost is configured.

d. Run:

```
commit
```

The configuration is committed.

- Configuring an interface link cost

a. Run:

```
system-view
```

The system view is displayed.

b. Run:

```
interface interface-type interface-number
```

The interface view is displayed.

c. Run:

```
trill cost { cost | maximum }
```

The link cost for the specified interface is configured.

d. Run:

```
commit
```

The configuration is committed.

----End

6.1.4 Checking the Configuration

Procedure

- Run the **display trill interface** [**verbose**] command to view information about the TRILL-enabled interfaces.
- Run the **display trill lsdb** [**verbose**] command to view the LSDB.
- Run the **display trill peer** [**verbose**] command to view information about TRILL neighbors.
- Run the **display trill route** [*nickname*] command to view information about TRILL unicast routes.
- Run the **display trill name-table** command to view information about TRILL dynamic host names.
- Run the **display trill mroute** [*vlan-id*] command to view information about TRILL multicast routes.

----End

6.2 Adjusting TRILL Route Selection

Pre-configuration Task

Before adjusting TRILL route selection, complete the following task:

- **Configuring Basic TRILL Functions**

Configuration Procedure

You can choose one or more configuration tasks as required.

6.2.1 Specifying a DVLAN

Context

To combine or separate TRILL networks, configure multiple carrier VLANs in a TRILL network. However, only one carrier VLAN is responsible for forwarding TRILL data packets. Therefore, a VLAN must be selected from these carrier VLANs to forward TRILL data packets. The DRB can specify a carrier VLAN as the designated VLAN (DVLAN), but the specified VLAN may not meet the network planning requirement. In this case, you can specify the DVLAN based on the network requirement. Only the DVLAN is responsible for forwarding TRILL data.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
interface interface-type interface-number
```

The interface view is displayed.

Step 3 Run:

```
trill designated-vlan vlan-id
```

A DVLAN is specified.

NOTE

Make sure that the specified DVLAN is a TRILL carrier VLAN.

Step 4 Run:

```
commit
```

The configuration is committed.

----End

6.2.2 Configuring TRILL Load Balancing

Context

On large-scale TRILL networks, multiple equal-cost routes destined for the same node may exist. In this situation, traffic is forwarded randomly, which may result in traffic imbalance and poor traffic management. To resolve this issue, enable the equal-cost routes to load-balance the traffic, which can improve link utilization and prevent congestion.

Load balancing can be configured for a TRILL process or interface, and the load balancing configured for an interface takes precedence over that configured for a process.

- To configure load balancing for a TRILL process, you need to specify the maximum number of equal-cost routes for load balancing.

  NOTE

  For the load balancing modes of TRILL packets in Eth-Trunk, see "Link Aggregation Configuration" in the *CloudEngine 12800 Series switch Configuration Guide - Ethernet Switching*.

  For TRILL ECMP load balancing modes, see "IP Routing Table Management" in the *CloudEngine 12800 Series switch Configuration Guide - IP Unicast Routing*.

- To configure load balancing for a TRILL interface, you can designate an RB as an AF for a VLAN or VLANs. After an RB is designated as an AF for a VLAN or VLANs, the RB is always used as an AF by the VLAN or VLANs, and all the RBs in other VLANs are designated as AFs based on the AF priority in order for load balancing.

Procedure

- Configure load balancing for a TRILL process.

  a. Run:

  ```
  system-view
  ```

  The system view is displayed.

  b. Run:

  ```
  trill
  ```

  The TRILL view is displayed.

  c. Run:

  ```
  maximum load-balance number
  ```

  The maximum number of equal-cost routes for load balancing is configured.

  After the command is run, traffic is load-balanced among a maximum of the configured number of equal-cost routes.

  NOTE

  By default, TRILL supports load balancing, and the maximum number of equal-cost routes for load balancing is 32.

  If equal-cost routes available outnumber *number* specified in the command, TRILL selects *number* equal-cost routes in the following sequence:

  - Routes with smaller outbound interface indexes
  - Routes whose next hop RBs have smaller system IDs

  d. Run:

  ```
  commit
  ```

  The configuration is committed.

- Configure load balancing for a TRILL interface.

a. Run:

```
system-view
```

The system view is displayed.

b. Run:

```
interface interface-type interface-number
```

The interface view is displayed.

c. Run:

```
trill load-balance systemid vlan { vlan-id1 [ to vlan-id2 ] }&<1-10>
```

The RB whose system ID is the specified *systemid* is designated as an AF for the specified VLAN or VLANs.

NOTE

The value of *vlan-id2* must be greater than that of *vlan-id1*.

d. Run:

```
commit
```

The configuration is committed.

----End

6.2.3 Configuring an Overload Bit for an RB

Context

If an RB on a TRILL network fails or needs to be upgraded or maintained, you can configure an overload bit for it. After an overload bit is configured for the RB, the RB is isolated from the TRILL network temporarily so that other devices no longer forward traffic to the RB, which prevents traffic interruptions.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
trill
```

The TRILL view is displayed.

Step 3 Run:

```
set-overload [ on-startup [ timeout1 ] [ send-sa-bit [ timeout2 ] ] ] max-reachable-cost
```

The RB is enabled to send LSPs whose overload bit is 1 and routes whose cost is the maximum link cost allowed (16777214).

NOTE

timeout1 specifies the period during which the RB stays in the overload state. The default value is 600s.

If **send-sa-bit** is specified in the command, the SA bit carried in Hello packets sent by the RB is 1 so that neighbors of the RB will not advertise the neighbor relationship with the RB to others after receiving the Hello packets. *timeout2* specifies the period during which the SA bit carried in Hello packets sent by the RB stays at 1. The default value is 30s.

Step 4 Run:

```
commit
```

The configuration is committed.

----End

6.2.4 Enabling TRILL Multicast Group-based Pruning

Context

On a TRILL network, an RB calculates a multicast distribution tree (MDT) to forward multicast or broadcast traffic. If the MDT has more than one next hop, the RB replicates the traffic and forwards one copy to each next hop. As a result, each next hop needs to process the traffic on receiving it, wasting bandwidth and forwarding resources.

To address this issue, TRILL provides multicast group-based pruning to ensure that an RB forwards traffic only to the next hop in the same multicast group. This function improves bandwidth efficiency.

NOTE

After devices on a TRILL network are configured to perform pruning based on multicast groups, multicast data within the TRILL network can only be forwarded based on MAC addresses. Multiple IPv4 multicast addresses may be mapped to the same IPv4 multicast MAC address according to the multicast IP-and-MAC address mapping mechanism. When multicast data is forwarded based on MAC addresses and a group IP address for receivers and the multicast IP address reserved for a protocol are mapped to the same IP multicast MAC address, the protocol cannot run normally. For example, IP multicast address 224.0.0.5 is reserved for the OSPF protocol. If a multicast group uses IP multicast address 225.0.0.5, the two IP multicast addresses are both mapped to IP multicast MAC address 01-00-5E-00-00-05. In this case, the OSPF protocol cannot run normally. Therefore, a proper IP multicast address plan must be made to prevent this problem.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
trill
```

The TRILL view is displayed.

Step 3 Run:

```
multicast-group prune enable
```

TRILL multicast group-based pruning is enabled.

Step 4 Run:

```
commit
```

The configuration is committed.

----End

6.2.5 Adjusting Multicast Distribution Trees

Multicast distribution trees (MDTs) can be adjusted for flexible TRILL configuration.

Context

On a TRILL network, either of the following ways can be used to adjust MDTs:

- **Configure a specific number of MDTs**: By default, an RB calculates two MDTs for load balancing. However, configuring an RB to calculate only one MDT in the following scenarios is recommended:
  - One of the two RBs that function as Spine nodes fails, but you do not want other non-Spine-node RBs to be selected as MDT roots.
  - The network is simple or forwarding resources are limited.
- **Configure TRILL to calculate MDTs in compliance with RFC 7180**: By default, TRILL calculates MDTs in compliance with RFC 6325. To allow devices following RFC 6325 to interoperate with devices following RFC 7180, run the **rfc7180 compatible** command to configure TRILL to calculate MDTs in compliance with RFC 7180.

Procedure

- Configure a specific number of MDTs.
  a. Run:
     ```
     system-view
     ```
     The system view is displayed.
  b. Run:
     ```
     trill
     ```
     TRILL is enabled globally, and the TRILL view is displayed.
  c. Run:
     ```
     tree-number compute compute-number
     ```
     The number of MDTs to be calculated by an RB is configured.
  d. Run:
     ```
     commit
     ```
     The configuration is committed.
- Configure TRILL to calculate MDTs in compliance with RFC 7180.
  a. Run:
     ```
     system-view
     ```
     The system view is displayed.
  b. Run:
     ```
     trill
     ```
     TRILL is enabled globally, and the TRILL view is displayed.

c. Run:

```
rfc7180 compatible
```

TRILL is configured to calculate MDTs in compliance with RFC 7180.

d. Run:

```
commit
```

The configuration is committed.

----End

6.2.6 Checking the Configuration

Procedure

- Run the **display trill interface** [*interface-type interface-number* | **verbose**] command to view information about TRILL interfaces.
- Run the **display trill route** [*nickname*] command to view information about TRILL unicast routes.

----End

6.3 Adjusting the TRILL Network Convergence Speed

Pre-configuration Task

Before adjusting the TRILL network convergence speed, complete the following task:

- **Configuring Basic TRILL Functions**

Configuration Procedure

You can choose one or more configuration tasks as required.

6.3.1 Configuring the Interval for Detecting Neighboring Device Faults

Context

Connection status between a TRILL device and its neighboring devices can be monitored by exchanging Hello packets at intervals. A neighboring device is considered Down if the TRILL device does not receive any Hello packets from the neighboring device within the specified period (called the holding time). The device fault then triggers routing table recalculation, and the TRILL network reconverges. To speed up fault detection, use the following methods to accelerate the speed of detecting TRILL neighboring device faults:

- **Setting an interval at which Hello packets are sent**
- **Setting the holding time for TRILL neighboring devices**

Procedure

- Setting an interval at which Hello packets are sent

a. Run:

```
system-view
```

The system view is displayed.

b. Run:

```
interface interface-type interface-number
```

The interface view is displayed.

c. Run:

```
trill timer hello hello-interval
```

The interval for sending Hello packets is configured on the interface.

By default, Hello packets are sent on the interface every 10 seconds.

d. Run:

```
commit
```

The configuration is committed.

- Setting the holding time for TRILL neighboring devices

a. Run:

```
system-view
```

The system view is displayed.

b. Run:

```
interface interface-type interface-number
```

The interface view is displayed.

c. Run:

```
trill timer holding-multiplier number
```

The multiplier between the expiration time of TRILL neighboring devices and the Hello packet sending interval is configured to determine the holding time for the TRILL neighboring devices.

By default, a TRILL neighboring device is considered Down after failing to receive three Hello packets.

d. Run:

```
commit
```

The configuration is committed.

----End

6.3.2 Adjusting SNP and LSP Attributes

Context

When the network status changes, LSPs are sent to advertise these changes on TRILL networks, and SNPs are responsible for synchronizing the LSDB of each device on the network. You can adjust SNP and LSP attributes based on the actual network status to improve the network performance and reliability. The following table describes the SNP and LSP attributes.

Table 6-1 SNP and LSP attributes

Attribute	Description	Default Setting
Interval for sending CSNPs	All devices on a TRILL network periodically send CSNPs through the DRB to synchronize the LSDBs. After the interval for sending CSNPs is adjusted based on the network scale, the LSDBs can be synchronized in real time and CSNPs are not frequently sent to occupy a large amount of the device memory.	10 seconds

Attribute	Description	Default Setting
LSP intelligent timer	In the TRILL network, if the local routing information changes, the device generates a new LSP to notify this change. Many new LSPs caused by frequent routing information changes will occupy a large number of system resources. The LSP intelligent timer can automatically adjust the delay based on the change frequency of routing information, which accelerates the network convergence speed and does not affect the system performance. The intelligent timer provides three parameters. The parameter functions are as follows: ● If only *max-interval* is configured, the intelligent timer functions as an ordinary one-time triggering timer and TRILL generates LSPs at a specified interval. ● If *max-interval* and *init-interval* are configured, *init-interval* determines the delay in LSP generation for the first time, and from the second time on, *max-interval* determines the delay in LSP generation. After the delay remains at the value specified by *max-interval* for three times or the TRILL process is restarted, the delay restores to the value specified by *init-interval*. ● If *max-interval*, *init-interval*, and *incr-*	● *max-interval*: 2 seconds ● *init-interval*: 0 milliseconds ● *incr-interval*: 0 milliseconds

Attribute	Description	Default Setting
	interval are configured, *init-interval* determines the delay in LSP generation for the first time, and *incr-interval* determines the delay in generating the LSPs with the same ID for the second time. From the third time on, the delay in generating an LSP increases twice every time until the delay reaches the value specified by *max-interval*. After the delay remains at the value specified by *max-interval* for three times or the TRILL process is restarted, the delay restores to the value specified by *init-interval*.	
LSP refresh interval	When the network status changes, LSPs are sent to advertise these changes to other devices on the network. Based on the network scale, a long LSP refresh interval may delay the real-time change synchronization and a short interval may result in frequent LSP update and much memory occupation. After the LSP refresh interval is adjusted, the network can converge in real time and a proper number of LSPs are sent.	By default, the LSP refresh interval is 900s, and the maximum lifetime of an LSP is 1200s. Ensure that the LSP refresh interval is more than 300s shorter than the maximum LSP lifetime. This allows new LSPs to reach all routers in an area before existing LSPs expire.

Attribute	Description	Default Setting
Maximum LSP lifetime	The maximum LSP lifetime determines how long LSPs exist on the device. An LSP is deleted if the device does not receive the updated one after the maximum LSP lifetime expires. If the maximum LSP lifetime is set to an excessively short period, the device may discard the original LSPs when receiving new ones, which results in the failure of LSDB synchronization on the network. If the maximum LSP lifetime is set to an excessively long period, the LSDB cannot be updated in real time when the network status changes, which decreases the network convergence speed.	1200 seconds
Minimum interval at which LSPs are sent	When there are a large number of LSPs on the device, multiple LSPs are sent each time at a specified interval. This attribute defines the minimum interval at which LSPs are sent. If the minimum interval is set to an excessively short period or there are too many LSPs to be sent each time, the network resources are occupied. Therefore, the device must be configured to send all LSPs equally based on the actual network load capability.	50 milliseconds

Attribute	Description	Default Setting
Interval at which LSPs are retransmitted over a P2P link	On a P2P network, the devices on the two ends of the link synchronize the LSDBs through LSP flooding. The device on one end sends an LSP; the device on the other end receives the LSP and replies with a PSNP. If the sending device does not receive the PSNP from the peer within a certain period, it resends the LSP. If the retransmission interval is set to an excessively short period, LSPs are retransmitted improperly, which results in high CPU, memory and bandwidth usage.	5 seconds

Procedure

- Setting an interval at which CSNPs are sent
  a. Run:
  ```
  system-view
  ```
  The system view is displayed.
  b. Run:
  ```
  interface interface-type interface-number
  ```
  The interface view is displayed.
  c. Run:
  ```
  trill timer csnp csnp-interval
  ```
  The interval for sending CSNPs is set on the interface.
  d. Run:
  ```
  commit
  ```
  The configuration is committed.
- Configuring the LSP intelligent timer
  a. Run:
  ```
  system-view
  ```
  The system view is displayed.
  b. Run:
  ```
  trill
  ```
  The TRILL view is displayed.

c. Run:

   ```
   timer lsp-generation max-interval [ init-interval [ incr-interval ] ]
   ```

   The intelligent timer used for generating LSPs is configured.

d. Run:

   ```
   commit
   ```

   The configuration is committed.

- Setting the LSP refresh interval

  a. Run:

     ```
     system-view
     ```

     The system view is displayed.

  b. Run:

     ```
     trill
     ```

     The TRILL view is displayed.

  c. Run:

     ```
     timer lsp-refresh refresh-time
     ```

     The LSP refresh interval is set.

  d. Run:

     ```
     commit
     ```

     The configuration is committed.

- Setting the maximum LSP lifetime

  a. Run:

     ```
     system-view
     ```

     The system view is displayed.

  b. Run:

     ```
     trill
     ```

     The TRILL view is displayed.

  c. Run:

     ```
     timer lsp-max-age age-time
     ```

     The maximum lifetime is set for LSPs.

  d. Run:

     ```
     commit
     ```

     The configuration is committed.

- Setting the minimum interval at which LSPs are sent

  a. Run:

     ```
     system-view
     ```

     The system view is displayed.

  b. Run:

     ```
     interface interface-type interface-number
     ```

     The interface view is displayed.

c. Run:

```
trill timer lsp-throttle throttleinterval [ count countnumber ]
```

The minimum interval for sending LSPs is set.

d. Run:

```
commit
```

The configuration is committed.

- Setting an interval at which LSPs are retransmitted over a P2P link

a. Run:

```
system-view
```

The system view is displayed.

b. Run:

```
interface interface-type interface-number
```

The interface view is displayed.

c. Run:

```
trill timer lsp-retransmit retransmit-interval
```

The interval for retransmitting LSPs over a P2P link is set.

d. Run:

```
commit
```

The configuration is committed.

----End

6.3.3 Setting the SPF Calculation Interval

Context

When the network status changes, TRILL calculates routes using the SPF algorithm to ensure accuracy. However, when the network status is unstable, frequent SPF calculation consumes excessive CPU resources, affecting services.

To solve this problem, configure an intelligent timer to control the SPF calculation interval. That is, set the SPF calculation interval to a small value to speed up TRILL route convergence, and set the interval to a large value after the TRILL network becomes stable.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
trill
```

The TRILL view is displayed.

Step 3 Run:

```
timer spf max-interval [ init-interval [ incr-interval ] ]
```

The SPF intelligent timer is configured.

Specify the parameters based on the following rules:

- If you only specify *max-interval*, the intelligent timer functions as an ordinary timer, and TRILL performs SPF calculation after routes are converged and *max-interval* expires.
- If you specify both *max-interval* and *init-interval*, *init-interval* determines the delay in SPF calculation for the first time, and from the second time on, *max-interval* determines the delay in SPF calculation. After the delay remains at the value specified by *max-interval* for three times or the TRILL process is restarted, the delay restores to the value specified by *init-interval*.
- If you specify *max-interval*, *init-interval*, and *incr-interval*, *init-interval* determines the delay in SPF calculation for the first time, and *incr-interval* determines the delay in SPF calculation for the second time. From the third time on, the delay in SPF calculation increases twice every time until the delay reaches the value specified by *max-interval*. After the delay remains at the value specified by *max-interval* for three times or the TRILL process is restarted, the delay restores to the value specified by *init-interval*.

Step 4 Run:

```
commit
```

The configuration is committed.

----End

6.3.4 Checking the Configuration

Procedure

- Run the **display trill interface verbose** command to view detailed information about TRILL interfaces.

----End

6.4 Configuring the Association Between STP/RSTP/MSTP and TRILL

Context

The association between STP/RSTP/MSTP and TRILL enables an MSTP network to connect to a TRILL network through two uplinks, implementing seamless expansion of the original data center network.

As shown in **Figure 6-1**, TRILL is deployed on RB1, RB2, RB3, RB4, RB5, and RB6, and servers communicate over the TRILL network. MSTP is deployed on MS1, MS2, and MS3 to prevent traffic loops. Traffic loops will also occur among RB1, RB3, RB4, MS1, and MS2. However, MSTP cannot be deployed on RB1 to prevent traffic loops.

Figure 6-1 Networking diagram for configuring the association between STP/RSTP/MSTP and TRILL

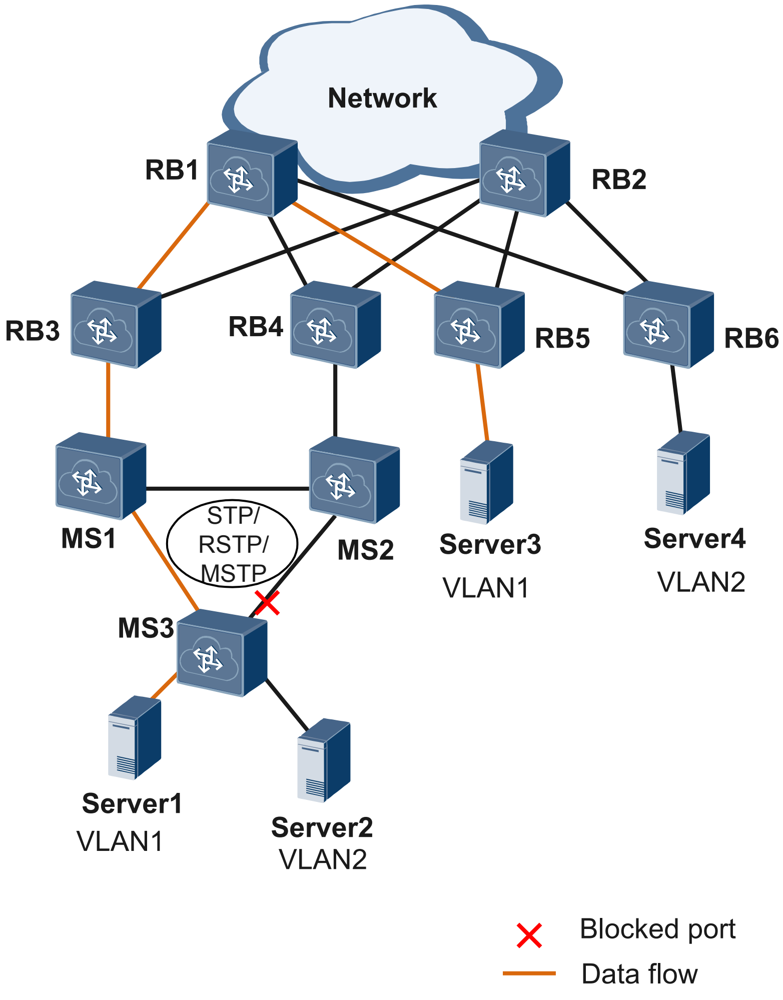

To prevent traffic loops among RB1, RB3, RB4, MS1, and MS2, two methods are available. You can choose one to perform configurations.

- Change the root bridge.

  Run MSTP on RB3 and RB4 and simulate RB3 and RB4 as one root bridge. That is, ensure that BPDUs sent by RB3 and RB4 carry the same bridge ID. MSTP then blocks one port after calculating the spanning tree, preventing traffic loops among RB1, RB3, RB4, MS1, and MS2. In most cases, however, each device has a different bridge MAC address and so has a different bridge ID. To ensure that RB3 and RB4 have the same bridge ID, you can configure the bridge MAC address used by the current device to participate in calculating the spanning tree.

  In **Figure 6-2**, MSTP blocks the MS3 interface connecting to MS2 and the MS2 interface connecting to MS1, and the transmission path for traffic from Server1 to Server3 is Server1 -> MS3 -> MS1 -> RB3 -> RB1 -> RB5 -> Server3.

Figure 6-2 Networking of the association between STP/RSTP/MSTP and TRILL (changing the root bridge)

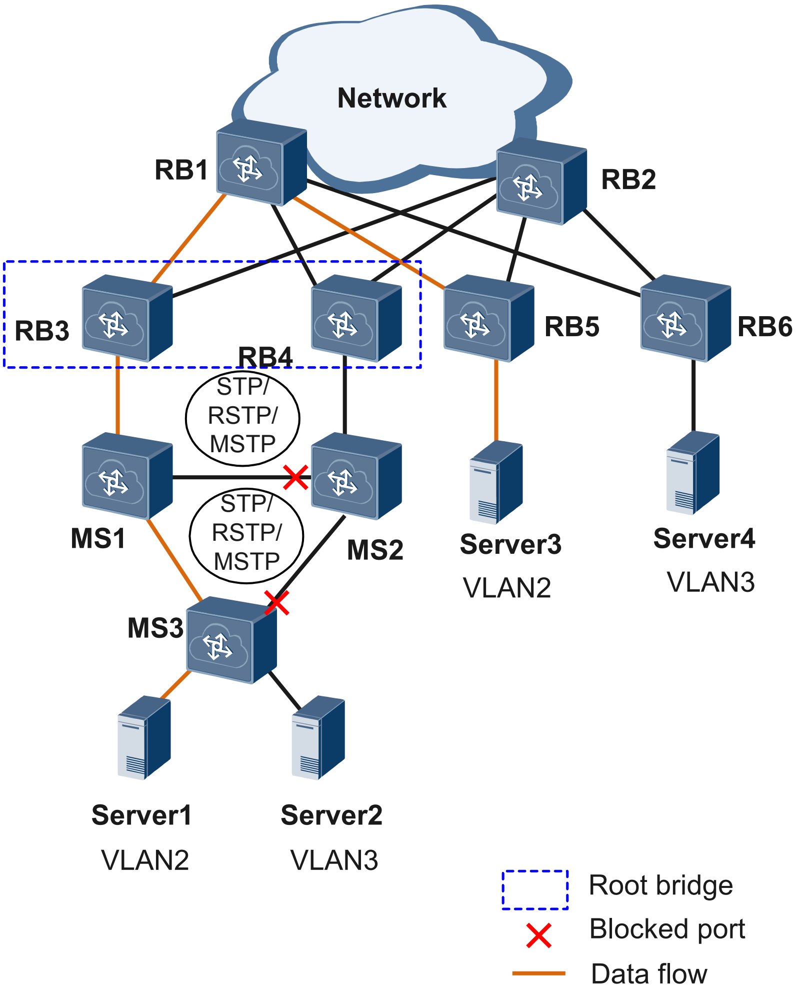

In **Figure 6-3**, if the link between MS1 and MS3 fails, MSTP recalculates the spanning tree, blocks the MS3 interface connecting to MS1, and unblocks the MS3 interface connecting to MS2. All devices, including those on the TRILL network and connected networks, must be notified of the topology change and update their MAC address entries and ARP entries accordingly. However, devices on the TRILL network cannot process topology change (TC) packets generated by MSTP. To allow TRILL devices to process TC packets and ensure uninterrupted traffic forwarding, configure the association between MSTP and TRILL. In **Figure 6-3**, MSTP blocks the MS3 interface connecting to MS1, and the transmission path for traffic from Server1 to Server3 is Server1 -> MS3 -> MS2 -> RB4 -> RB1 -> RB5 -> Server3.

Figure 6-3 Networking of the association between STP/RSTP/MSTP and TRILL (changing the root bridge in the case of a link fault)

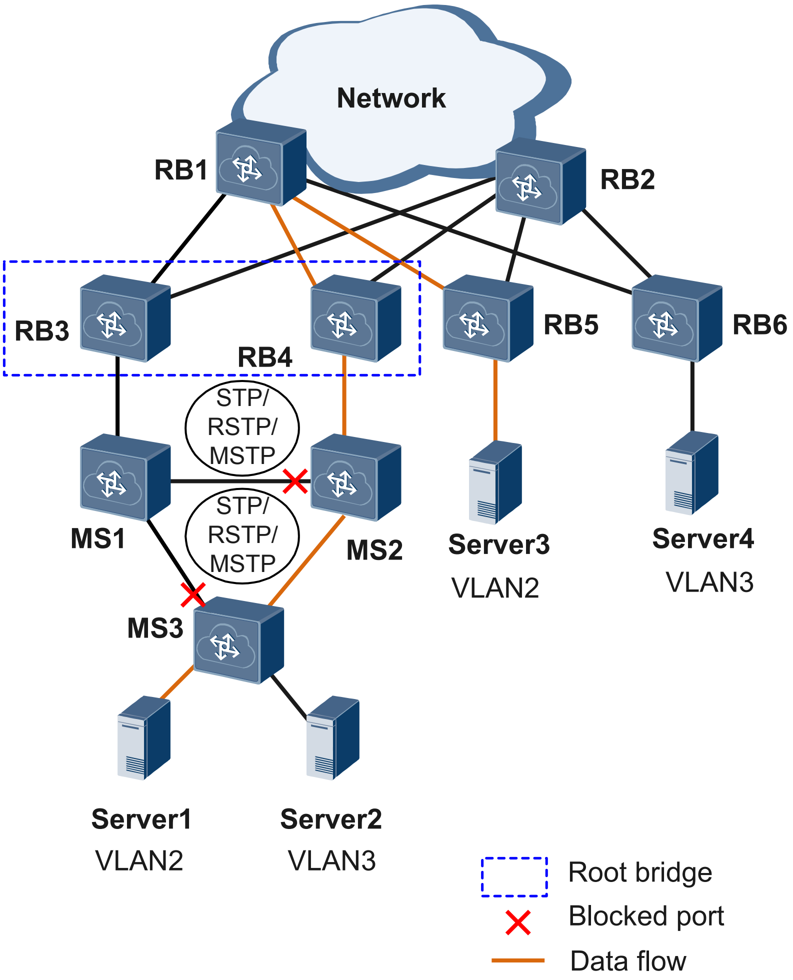

- Retain the root bridge.

  Configure the RB3 and RB4 interfaces connecting to the MSTP network to transparently transmit spanning tree packets in the TRILL network. Spanning tree packets are then transparently transmitted in the TRILL network. If the TRILL network receives topology change (TC) packets, the TRILL network updates MAC address entries in a timely manner. This configuration does not affect the existing configurations in the original MSTP network, and MS1 is still the root bridge.

  As shown in **Figure 6-4**, MSTP blocks the MS3 interface connecting to MS2, and the transmission path for traffic between Server1 and Server3 is Server1->MS3->MS1->RB3->RB1->RB5->Server3.

Figure 6-4 Networking diagram for configuring the association between STP/RSTP/MSTP and TRILL (retaining the root bridge)

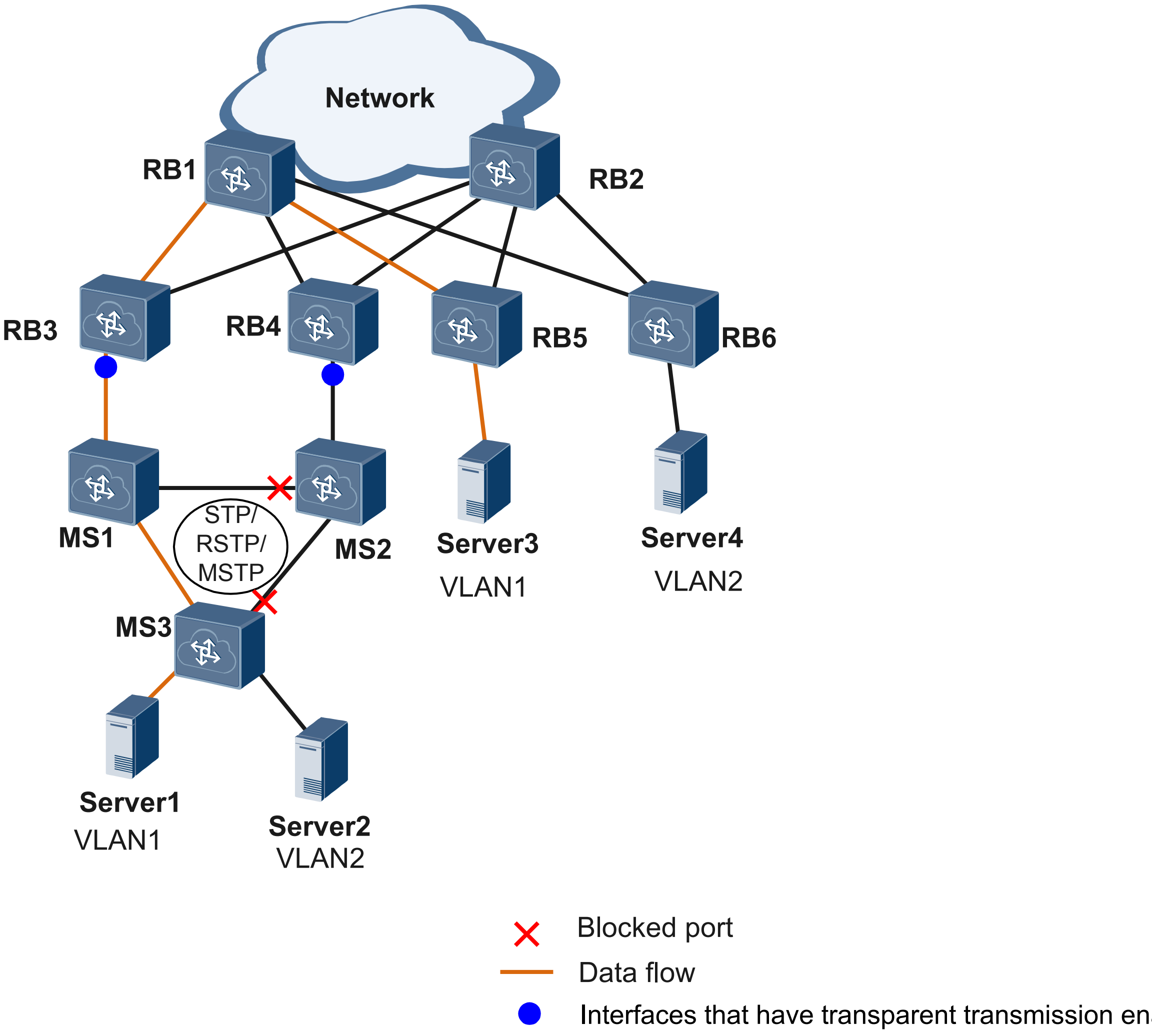

As shown in **Figure 6-5**, if a fault occurs on the link between MS1 and MS3, MSTP recalculates the spanning tree, blocks the MS3 interface connecting to MS1, and unblocks the MS3 interface connecting to MS2. This topology change needs to be notified to all the devices using TC packets, including devices on the TRILL network and connected networks so that these devices can detect the topology change and update MAC address entries in a timely manner. The TRILL network, however, does not process the TC packets. To ensure uninterrupted service traffic forwarding in this situation, configure the transparent transmission of spanning tree packets in the TRILL network. Spanning tree packets are then transparently transmitted in the TRILL network. If the TRILL network receives TC packets, the TRILL network updates MAC address entries in a timely manner. As shown in **Figure 6-5**, MSTP blocks the MS3 interface connecting to MS1, and the transmission path for traffic between Server1 and Server3 becomes Server1->MS3->MS2->RB4->RB1->RB5->Server3.

Figure 6-5 Networking diagram for configuring the association between STP/RSTP/MSTP and TRILL (retaining the root bridge in the case of a link fault)

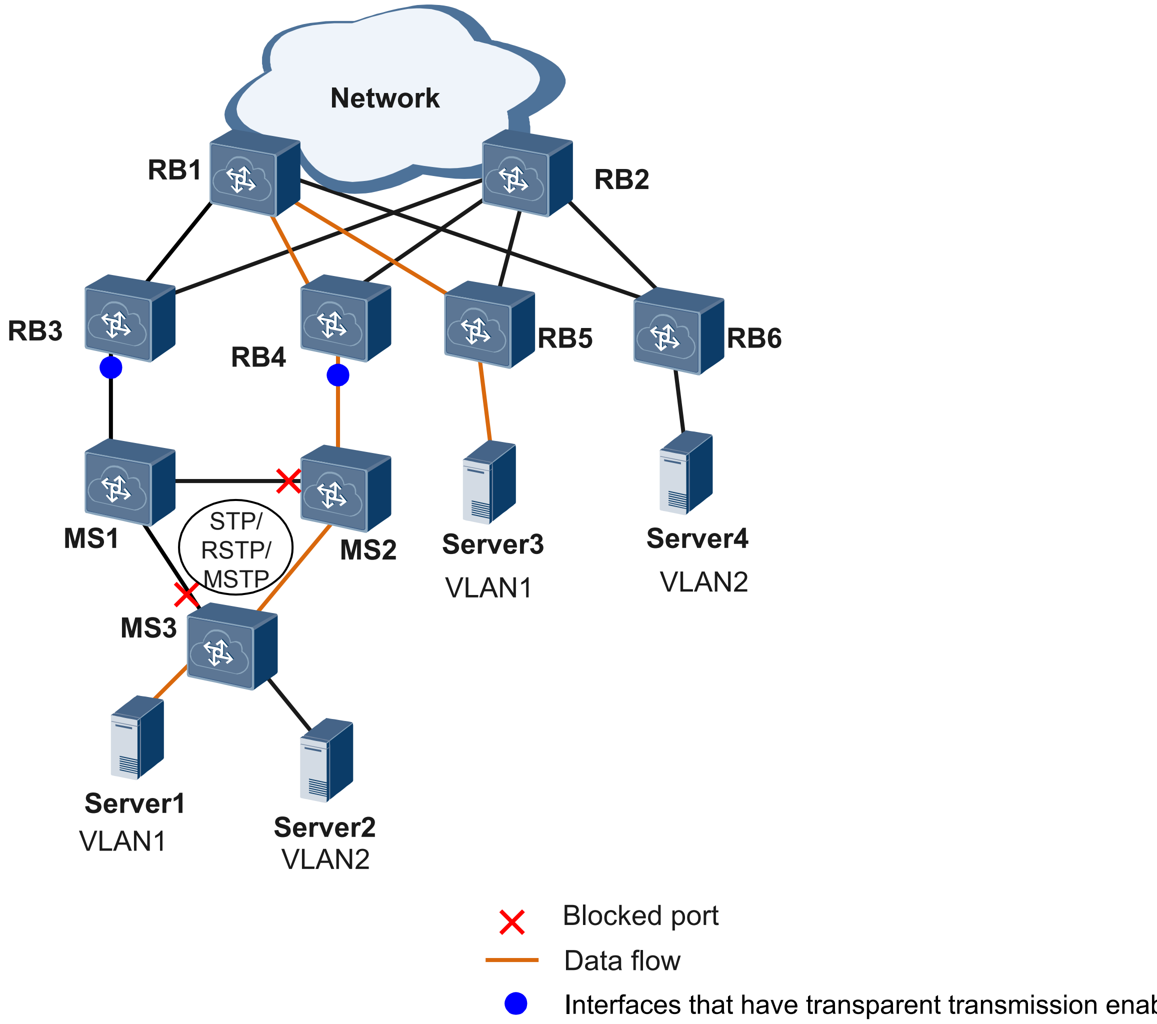

Pre-configuration Tasks

Before configuring the association between STP/RSTP/MSTP and TRILL, complete the following tasks:

- Configuring the STP/RSTP Function or Configuring the MSTP Function
- **Configuring Basic TRILL Functions**

Procedure

- Configure the association between STP/RSTP/MSTP and TRILL (changing the root bridge).

  a. Run:

  ```
  system-view
  ```

  The system view is displayed.

b. Run:

```
stp bridge-address mac-address
```

The bridge MAC address for spanning tree calculation is configured on RB3 and RB4.

By default, a device uses its MAC address as the bridge MAC address to calculate the spanning tree.

NOTE

- A bridge ID identifies a device. If two devices send packets with the same bridge ID to another device, the packets are considered to be sent by one device. Exercise caution when you run the **stp bridge-address** command.
- After you run the **stp bridge-address** command to configure the same bridge MAC address for two devices, to allow the devices to simulate the same root bridge, ensure that the devices have the same spanning tree protocol configurations, such as the device priority and timer settings.

c. Run:

```
stp tc-notify trill vlan vlan-id
```

The association between STP/RSTP/MSTP and TRILL is enabled on RB3 and RB4.

By default, the association between STP/RSTP/MSTP and TRILL is disabled.

vlan-id must be the ID of an admin VLAN on the TRILL network. A VLANIF interface must be configured for the admin VLAN.

NOTE

Before the **stp tc-notify trill vlan** *vlan-id* command is configured on the device, the **stp disable** command must be configured on the interface configured with the **trill enable port-mode** { **hybrid** | **p2p** | **trunk** } command.

If the association between STP/RSTP/MSTP and TRILL is configured on a VS, you are advised to run the **carrier up-hold-time** *interval* command to set the delay in reporting Up events to 50s on all the interfaces connected to the MSTP network on the TRILL network.

d. Run:

```
commit
```

The configuration is committed.

- Configure the association between STP/RSTP/MSTP and TRILL (retaining the root bridge).

a. Run:

```
system-view
```

The system view is displayed.

b. Run:

```
interface interface-type interface-number
```

The interface view is displayed.

c. Run:

```
stp disable
```

STP/RSTP/MSTP is disabled on the interface.

By default, STP/RSTP/MSTP is enabled on an interface.

d. Run:

```
stp tc-snooping notify trill
```

Transparent transmission of spanning tree packets is enabled in the TRILL network.

By default, spanning tree packets are not transparently transmitted in a TRILL network.

NOTE

To ensure that spanning tree packets are transparently transmitted in the TRILL network, ensure that:

- An admin VLAN has been configured in the TRILL network and an IP address has been configured for the VLANIF interface of the admin VLAN.
- The same PVID has been configured for the interfaces on which the **stp tc-snooping notify trill** command is executed, and the PVID is the ID of the admin VLAN of the TRILL network.
- You are advised to run the **stp edged-port disable** command on the interfaces connecting the STP/RSTP/MSTP network to the TRILL network to configure the interfaces as non-edge ports.

NOTICE

If the interface that connects the TRILL network to the MSTP network is faulty, MSTP cannot fast detect the fault, causing services to be interrupted for at least 48 seconds.

e. Run:

```
commit
```

The configuration is committed.

----End

Checking the Configuration

Run the **display current-configuration** command to check the association between STP/RSTP/MSTP and TRILL.

6.5 Configuring TRILL Network Dual-Homing Through a M-LAG

Eth-Trunk provides board-level reliability, whereas M-LAG provides device-level reliability.

Pre-configuration Tasks

Before configuring TRILL network dual-homing through a M-LAG, **6.1 Configuring Basic TRILL Functions**.

6.5.1 Configuring a DFS Group

Context

A DFS group is used to pair devices. To exchange heartbeat packets, a DFS group must be bound to a nickname, which is used to communicate with the peer end.

In a TRILL active-active scenario, configure a pseudo nickname for a DFS group. Two devices dual-homed to the TRILL network must have the same pseudo nickname. In this way, the peer end considers the two devices a logical device on the TRILL network and does not need to reconstruct the network topology.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
dfs-group dfs-group-id
```

A DFS group is created and its view is displayed, or the view of an existing DFS group is displayed.

Step 3 Run:

```
source nickname nickname-value
```

The DFS group is bound to a nickname.

Step 4 Run:

```
pseudo-nickname nickname-value1 [ priority priority ]
```

A pseudo nickname is configured for the DFS group.

Step 5 (Optional) Run:

```
priority priority
```

The priority of the DFS group is set.

The priority of a DFS group is used for master/backup negotiation between two devices. The device with a higher priority is the master, and a larger priority value indicates a higher priority.

If the priorities of the two devices are the same, the device with a smaller system MAC address is the master.

By default, the priority of a DFS group is 100.

Step 6 Run:

```
commit
```

The configuration is committed.

----End

6.5.2 Configuring an Interface as a Peer-link Interface

Context

A peer-link is a direct aggregated link between two devices configured with M-LAG. The peer-link is used to transmit negotiation packets and some traffic. The two devices exchange protocol packets over the peer-link to ensure normal running of M-LAG.

Prerequisites

The direct link between two devices configured with M-LAG has been configured as an aggregated link.

NOTE

When member interfaces of a peer-link are deployed on the same card, a fault of the card causes the peer-link fault. To improve reliability, it is recommended that member interfaces of the peer-link be deployed on different cards.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
interface eth-trunk trunk-id
```

The Eth-Trunk interface view is displayed.

Step 3 Run:

```
mode { lacp-static | lacp-dynamic }
```

The Eth-Trunk is configured to work in LACP mode.

By default, an Eth-Trunk works in manual load balancing mode. To ensure M-LAG reliability, you must configure the Eth-Trunk to work in LACP mode.

Step 4 Run:

```
undo stp enable
```

STP is disabled on the interface.

NOTE

STP needs to be disabled because two devices need to be simulated into one STP root bridge and the directly connected interface cannot be blocked.

When V-STP needs to be configured, you do not need to perform this operation.

Step 5 Run:

```
peer-link peer-link-id
```

The interface is configured as a peer-link interface.

- An interface configured as a peer-link interface joins all VLANs by default.
- If the ERPS control VLAN, TRILL Carrier VLAN, or FCoE VLAN needs to be configured, perform **Step 6** to remove the peer-link interface from the control VLAN,

Carrier VLAN, or FCoE VLAN. Otherwise, the control VLAN, Carrier VLAN, or FCoE VLAN cannot be configured.

- If the network-side VLANIF interface is configured, you are advised to perform **Step 6** to remove the peer-link interface from the VLAN. Otherwise, exceptions may occur. For example, heartbeat detection becomes ineffective.

Step 6 (Optional) Run:

```
port vlan exclude { { vlan-id1 [ to vlan-id2 ] } &<1-10> }
```

The VLANs not allowed by the peer-link interface are specified.

Step 7 Run:

```
commit
```

The configuration is committed.

----End

6.5.3 Binding an Interface to a DFS Group

Prerequisites

The links between a server and two upstream devices configured with M-LAG have been configured as aggregated links. To improve reliability, configure link aggregation in LACP mode.

If the two devices that constitute an M-LAG are configured in an SVF and access devices are connected through left switches, run the **extend enable** command in the Eth-Trunk interface view first.

Procedure

- When the Eth-Trunk works in manual load balancing mode, perform the following operations.
  a. Run:
     ```
     system-view
     ```
     The system view is displayed.
  b. Run:
     ```
     interface eth-trunk trunk-id
     ```
     The Eth-Trunk interface view is displayed.
  c. Run:
     ```
     dfs-group dfs-group-id m-lag m-lag-id
     ```
     The Eth-Trunk is bound to a DFS group.

     NOTE

     The two devices configured with M-LAG must use the same M-LAG ID.
  d. Run:
     ```
     commit
     ```
     The configuration is committed.
- When the Eth-Trunk works in LACP mode, perform the following operations.

a. Run:

```
system-view
```

The system view is displayed.

b. Run:

```
interface eth-trunk trunk-id
```

The Eth-Trunk interface view is displayed.

c. Run:

```
mode { lacp-static | lacp-dynamic }
```

The Eth-Trunk is configured to work in LACP mode.

d. Run:

```
dfs-group dfs-group-id m-lag m-lag-id
```

The Eth-Trunk is bound to a DFS group.

NOTE

The two devices configured with M-LAG must use the same M-LAG ID.

e. Run:

```
lacp m-lag priority priority
```

The LACP M-LAG system priority is configured.

- Eth-Trunks on both devices configured with M-LAG must use the same LACP M-LAG system priority.
- The LACP M-LAG system priority configured in the system view takes effect for all Eth-Trunks. The LACP M-LAG system priority configured in the Eth-Trunk interface view takes effect for only the specified Eth-Trunk. If the **lacp m-lag priority** command is configured in the system view and Eth-Trunk interface view, the LACP M-LAG system priority configured in the Eth-Trunk interface view takes effect.
- When multiple M-LAGs are configured on the device, different Eth-Trunks can use different system priorities. In this case, you need to set the LACP M-LAG system priority in the Eth-Trunk interface view.
- The LACP M-LAG system priority is valid for the M-LAG composed of an Eth-Trunk in LACP mode, whereas the LACP system priority configured by the **lacp priority** command is valid for an Eth-Trunk in LACP mode.

  If both the LACP M-LAG system priority and LACP system priority are configured, the Eth-Trunk in LACP mode that joins an M-LAG uses the LACP M-LAG system priority but not the LACP system priority.

f. Run:

```
lacp m-lag system-id mac-address
```

The LACP M-LAG system ID is configured.

- Eth-Trunks on both devices configured with M-LAG must use the same LACP M-LAG system ID.
- The LACP M-LAG system ID configured in the system view takes effect on all Eth-Trunks. The LACP M-LAG system ID configured in the Eth-Trunk interface view takes effect for only the specified Eth-Trunk. If the **lacp m-lag system-id** command is configured in the system view and Eth-Trunk interface view, the LACP M-LAG system ID configured in the Eth-Trunk interface view takes effect.

- When multiple M-LAGs are configured on the device, different Eth-Trunks can use different system IDs. In this case, you need to set the system ID of an LACP M-LAG in the Eth-Trunk interface view.
- The LACP M-LAG system ID is valid for the M-LAG composed of an Eth-Trunk in LACP mode, whereas the LACP system ID is valid for an Eth-Trunk in LACP mode. The LACP system ID is the MAC address of an Ethernet interface on the MPU and cannot be changed.

g. Run:

```
commit
```

The configuration is committed.

----End

6.5.4 Checking the Configuration

Procedure

- Run the **display dfs-group** *dfs-group-id* [**node** *node-id* **m-lag** [**brief**] | **peer-link**] command to check M-LAG information.
- Run the **display stp v-stp** command to check the V-STP status and statistics.
- Run the **display dfs-group m-lag check stp** command to check whether the STP configurations on both ends of the M-LAG are consistent.

----End

Follow-up Procedure

After M-LAG is configured, if the peer-link is faulty but the heartbeat status is normal, some interfaces on the standby devices will enter the Error-Down state. The device records the status of an interface as Error-Down when it detects that a fault occurs. The interface in Error-Down state cannot receive or send packets and the interface indicator is off. You can run the **display error-down recovery** command to check information about all interfaces in Error-Down state on the device.

- When M-LAG is used for dual-homing to a TRILL network and the peer-link fails but the heartbeat is normal, the M-LAG interface on the backup device will enter the error-down state. When the peer-link recovers, the physical interfaces change from error-down to Up after 2 minutes by default (the physical interface on a leaf switch in an SVF of modular and fixed switches changes from error-down to Up after 5 minutes by default).
- When M-LAG is used for dual-homing to an Ethernet, VXLAN network, or IP network and the peer-link fails but the heartbeat is normal, all physical interfaces except the management interface, peer-link interface, and stack interface on the backup device will enter the error-down state. When the peer-link recovers, the M-LAG interface change from error-down to Up after 2 minutes by default (the physical interface on a leaf switch in an SVF of modular and fixed switches changes from error-down to Up after 5 minutes by default). Other physical interfaces go Up immediately by default.

When the interface enters the Error-Down state, locate the cause. You are not advised to manually restore the interface or run the **error-down auto-recovery cause m-lag interval** *interval-value* command in the system view to enable the interface to go Up automatically. Otherwise, extra packets, packet loss or forwarding failure may occur. Exercise caution when you perform the preceding operation.

6.6 Configuring TRILL Gateway

Context

It is costly and complex to deploy independent gateway devices or configure virtual systems (VSs) as gateways on a TRILL network. You can deploy the TRILL gateway function to reduce deployment cost and complexity.

NOTE

When configuring a TRILL gateway, pay attention to the following points:

- The TRILL gateway in internal loopback mode cannot forward multicast traffic at Layer 3.
- After the TRILL gateway function is configured, the administrative VLAN does not support data traffic transmission.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
trill gateway enable
```

The TRILL gateway function in internal loopback mode is enabled

Step 3 Run:

```
commit
```

The configuration is committed.

----End

6.7 Improving TRILL Network Security

Pre-configuration Task

Before configuring TRILL authentication, complete the following task:

- **Configuring Basic TRILL Functions**

Configuration Procedure

You can choose one or more configuration tasks as required.

6.7.1 Configuring TRILL Packet Authentication

Context

In most cases, RBs do not encapsulate authentication information into TRILL packets before sending them or authenticate received TRILL packets. Therefore, networks are open to attacks. To improve network security, configure TRILL authentication.

In TRILL packet authentication, LSPs and SNPs carry authentication information. After receiving the packets, the remote RB authenticates them and discards those that fail the authentication.

RBs in the same area must share the same authentication mode and password so that TRILL packets can be properly flooded. Whether packets pass the authentication does not affect the establishment of neighbor relationships.

NOTICE

If **plain** is selected during the configuration of the TRILL packet authentication mode, the password is saved in the configuration file in plain text. This brings security risks. It is recommended that you select **cipher** to save the password in cipher text.

Simple and MD5 authentication has potential risks. HMAC-SHA256 cipher text authentication is recommended.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
trill
```

The TRILL view is displayed.

Step 3 Run:

```
area-authentication-mode { { simple | md5 | hmac-sha256 key-id key-id }
{ [ cipher ] password-key | plain password } | keychain keychain-name } [ snp-
packet { authentication-avoid | send-only } | all-send-only ]
```

The LSP authentication is configured.

The authentication involves the following situations:

- The RB encapsulates the authentication information into LSPs and SNPs to be sent and checks whether the received packets pass authentication. The RB then discards the packets that do not pass the authentication. In this case, the parameter **snp-packet all-send-only** does not need to be configured.
- The RB encapsulates authentication information into LSPs to be sent and checks whether the received LSPs pass the authentication. The RB neither encapsulates the SNPs to be sent with authentication information nor checks whether the received SNPs pass the authentication. In this case, the parameter **snp-packet authentication-avoid** needs to be specified.
- The RB encapsulates the LSPs and SNPs to be sent with authentication information; however, it checks the authentication result of only the received LSPs, not SNPs. In this case, the parameter **snp-packet send-only** needs to be configured.
- The RB encapsulates the LSPs and SNPs to be sent with authentication information, but does not check whether the received LSPs or SNPs pass the authentication. In this case, the parameter **all-send-only** needs to be specified.

NOTE

If keychain authentication is used, the encryption algorithm must be configured to HMAC-MD5 algorithm.

Step 4 Run:

```
commit
```

The configuration is committed.

----End

6.7.2 Configuring TRILL Interface Authentication

Context

In TRILL interface authentication, authentication information is configured on a TRILL interface, and Hello packets sent through the interface are encapsulated with the information. Only the authenticated Hello packets can be received.

If TRILL interface authentication is configured on both ends, they must share the same authentication mode and password so that neighbor relationships can be established between them.

NOTICE

If **plain** is selected during the configuration of the TRILL interface authentication mode, the password is saved in the configuration file in plain text. This brings security risks. It is recommended that you select **cipher** to save the password in cipher text.

Simple and MD5 authentication has potential risks. HMAC-SHA256 cipher text authentication is recommended.

Procedure

Step 1 Run:

```
system-view
```

The system view is displayed.

Step 2 Run:

```
interface interface-type interface-number
```

The interface view is displayed.

Step 3 Run:

```
trill authentication-mode { { simple | md5 | hmac-sha256 key-id key-id }
{ [ cipher ] password-key | plain password } | keychain keychain-name } [ send-
only ]
```

The authentication mode and password are configured on the interface.

- Configure **send-only** if the TRILL interface needs to encapsulate authentication information into Hello packets to be sent and does not need to check whether the received packets pass the authentication.

- Do not configure **send-only** if the TRILL interface needs to encapsulate authentication information into Hello packets to be sent and check whether the received packets pass the authentication. In addition, configure the same authentication information for all TRILL interfaces in the same VLAN to ensure normal communication.

NOTE

If keychain authentication is used, the encryption algorithm must be configured to HMAC-MD5 algorithm.

Step 4 Run:

```
commit
```

The configuration is committed.

----End

6.7.3 Checking the Configuration

Procedure

- Run the **display trill lsdb verbose** command to view the verbose LSDB information.

----End

7 Maintaining TRILL

About This Chapter

Resetting TRILL facilitates fault location of TRILL networks.

7.1 Resetting TRILL

Background

You can reset TRILL to clear all TRILL data and re-establish neighbor relationships.

NOTICE

Resetting TRILL may interrupt services. Therefore, confirm the action before running the **reset trill all** command.

Procedure

- Run the **reset trill all** command in the user view to reset TRILL.

----End

7.2 Configuring TRILL OAM

Context

After TRILL is configured, unicast trace and ping functions can be used to detect TRILL network connectivity. To check the TRILL unicast forwarding path, use the **display trill forwarding-path unicast** command.

NOTE

The input parameter must be consistent with the actual hash factor so that a correct outbound interface can be obtained.

In a stack, to view the TRILL unicast forwarding path or use the TRILL unicast trace function, specify the source interface.

Procedure

- Run the following **trace trill** commands to locate TRILL network faults.

  # Perform TRILL unicast trace.

  trace trill [**-h** *hop-count-value* | **-t** *timeout*] * *nickname* [**interface** *interface-type interface-number*]

  # Perform flow-based unicast trace.

  trace trill [**-h** *hop-count-value* | **-t** *timeout*] * *nickname* [**source-mac** *mac-address* | **destination-mac** *mac-address* | **source-ip** *ip-address* | **destination-ip** *ip-address* | **source-port** *port-number* | **destination-port** *port-number* | **ce-vlan** *ce-vlan-id* | **eth-type** *eth-type* | **protocol** *protocol-type* | **source-interface** *interface-type interface-number*] *

- Run the **ping trill** [**-c** *count* | **-h** *ttl-value* | **-m** *time* | **-t** *timeout*] * *nickname* command to locate TRILL network faults.

- Run the **display trill forwarding-path unicast role** { **ingress dst-nickname** *dst-nickname* | **transit dst-nickname** *dst-nickname* | **egress out-interface** *eth-trunk interface-number* } { **eth-type** { **ip** | **l2** | **dhcp** | **arp** } | **src-mac** *src-mac-address* | **dst-mac** *dst-mac-address* | **src-ip** *src-ip-address* | **dst-ip** *dst-ip-address* | **ce-vlan** *vlan-id* | **src-interface** *interface-type interface-number* | **protocol** { *protocol* | **gre** | **icmp** | **igmp** | **ip** | **ipinip** | **ospf** | **tcp** [**l4-src-port** *l4-src-port* | **l4-dst-port** *l4-dst-port*]* | **udp** [**l4-src-port** *l4-src-port* | **l4-dst-port** *l4-dst-port*]* } }* command to check the TRILL unicast forwarding path.

----End

8 Configuration Examples

About This Chapter

This section provides configuration examples of TRILL. Refer to the networking diagrams to help familiarize yourself with the configuration. The configuration examples include networking requirements and configuration roadmap.

8.1 Example for Configuring Basic TRILL Functions

Networking Requirements

Figure 8-1 shows a data center network. It is required that an unblocked Layer network be constructed, multi-path forwarding be implemented, and each network node can implement line-speed forwarding.

NOTE

In the figure, 10GE1 indicates 10GE1/0/1, 10GE2 indicates 10GE1/0/2, and 10GE3 indicates 10GE1/0/3.

Figure 8-1 Networking for configuring basic TRILL functions

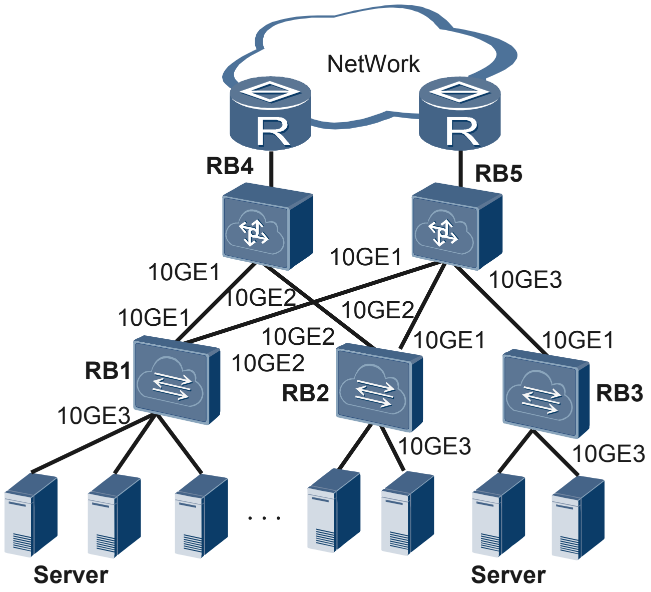

Configuration Roadmap

Using TRILL to construct a Layer 2 network can meet the requirements. The configuration roadmap is as follows:

1. Enable TRILL globally on five RBs so that the RBs can process TRILL packets.
2. Configure mandatory parameters such as the VLAN, NET, and nickname.
3. Enable TRILL on the interfaces connecting the five RBs so that the interfaces can send and receive TRILL packets.
4. Configure an admin VLAN on each of five RBs so that the administrators can manage devices using the NMS.

Procedure

Step 1 Enable TRILL on each RB, and configure their VLANs, NETs, and nicknames. Core devices have high performance. Therefore, you need to set the root priority of the core device RB4 to 65535 to ensure that RB4 becomes the multicast tree root of the TRILL network.

# Configure RB1.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB1
[*HUAWEI] commit
[~RB1] vlan 100
[*RB1-vlan100] commit
[~RB1-vlan100] quit
[~RB1] trill
[*RB1-trill] carrier-vlan 10
[*RB1-trill] ce-vlan 100
[*RB1-trill] network-entity 00.0000.0000.1111.00
[*RB1-trill] nickname 100
[*RB1-trill] commit
[~RB1-trill] quit
```

# Configure RB2.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB2
[*HUAWEI] commit
[~RB2] vlan 100
[*RB2-vlan100] commit
[~RB2-vlan100] quit
[~RB2] trill
[*RB2-trill] carrier-vlan 10
[*RB2-trill] ce-vlan 100
[*RB2-trill] network-entity 00.0000.0000.2222.00
[*RB2-trill] nickname 200
[*RB2-trill] commit
[~RB2-trill] quit
```

# Configure RB3.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB3
[*HUAWEI] commit
[~RB3] vlan 100
[*RB3-vlan100] commit
[~RB3-vlan100] quit
[~RB3] trill
[*RB3-trill] carrier-vlan 10
[*RB3-trill] ce-vlan 100
[*RB3-trill] network-entity 00.0000.0000.3333.00
[*RB3-trill] nickname 300
[*RB3-trill] commit
[~RB3-trill] quit
```

# Configure RB4.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB4
[*HUAWEI] commit
[~RB4] trill
[*RB4-trill] carrier-vlan 10
[*RB4-trill] network-entity 00.0000.0000.4444.00
[*RB4-trill] nickname 400 root-priority 65535
[*RB4-trill] commit
[~RB4-trill] quit
```

# Configure RB5.

```
<HUAWEI> system-view
```

```
[~HUAWEI] sysname RB5
[*HUAWEI] commit
[~RB5] trill
[*RB5-trill] carrier-vlan 10
[*RB5-trill] network-entity 00.0000.0000.5555.00
[*RB5-trill] nickname 500 root-priority 65535
[*RB5-trill] commit
[~RB5-trill] quit
```

Step 2 Configure TRILL on each interface.

# Enable TRILL on the interface of RB1.

```
[~RB1] interface 10ge 1/0/1
[~RB1-10GE1/0/1] undo shutdown
[~RB1-10GE1/0/1] port link-type trunk
[*RB1-10GE1/0/1] trill enable
[*RB1-10GE1/0/1] quit
[*RB1] interface 10ge 1/0/2
[*RB1-10GE1/0/2] undo shutdown
[*RB1-10GE1/0/2] port link-type trunk
[*RB1-10GE1/0/2] trill enable
[*RB1-10GE1/0/2] quit
[*RB1] interface 10ge 1/0/3
[*RB1-10GE1/0/3] undo shutdown
[*RB1-10GE1/0/3] port default vlan 100
[*RB1-10GE1/0/3] commit
[~RB1-10GE1/0/3] quit
```

The configurations on other RBs are similar to that on RB1. For detailed configurations, see **Configuration Files**.

Step 3 Configure admin VLAN

# Configure an admin VLAN on each device of the TRILL network, configure VLANIF interfaces for the admin VLANs and configure an IP address for each VLANIF interface. RB1 is used as an example here.

```
[~RB1] vlan 50
[*RB1-vlan50] quit
[*RB1] interface vlanif 50
[*RB1-Vlanif50] ip address 192.168.10.1 24
[*RB1-Vlanif50] quit
[*RB1] trill
[*RB1-trill] admin-vlan 50
[*RB1-trill] commit
[~RB1-trill] quit
```

The configurations on other RBs are similar to that on RB1. For detailed configurations, see **Configuration Files**.

Step 4 Check the TRILL database and unicast routing table of each RB.

# Run the **display trill interface** command to view information about the TRILL-enabled interface on each RB. RB1 is used as an example here.

```
[~RB1] display trill interface

Interface information for TRILL
----------------------------------------------------------------------------
Total Interface(s): 2

Interface Circuit-ID   State            MTU Type  DRB-State DVLAN  Port-Type
----------------------------------------------------------------------------
10GE1/0/1         001  UP              1497 L1    Non-DRB     --   p2p
10GE1/0/2         002  UP              1497 L1    Non-DRB     --   p2p
```

# Run the **display trill nickname** command to view the nickname of each RB. RB1 is used as an example here.

```
[~RB1] display trill nickname

Nickname information for TRILL
------------------------------------------------------------------
*-Local Nickname, A-Advertised, S-Suppressed / S-Static, D-Dynamic

Total Nickname(s): 5

  Nickname Source ID/Trill Name   State Priority RootPri
--------------------------------------------------------
*      100 0000.0000.1111         A/S        192   32768
       200 0000.0000.2222         A/S        192   32768
       300 0000.0000.3333         A/S        192   32768
       400 0000.0000.4444         A/S        192   32768
       500 0000.0000.5555         A/S        192   32768
```

# Run the **display trill route** command to view the unicast routing table of each RB. RB1 is used as an example here.

```
[~RB1] display trill route

TRILL Unicast Routing Table
--------------------------------------------------------------------
Flags: D-Download To Fib

Total Route(s): 4

Nickname       Cost Flag OutInterface  OuterVlan NextHop            Hop
-----------------------------------------------------------------------
     200       4000 D    10GE1/0/1            10 400/3609-b654-1220   2
                         10GE1/0/2            10 500/3609-b655-1220   2
     300       4000 D    10GE1/0/2            10 500/3609-b655-1220   2
     400       2000 D    10GE1/0/1            10 400/3609-b654-1220   1
     500       2000 D    10GE1/0/2            10 500/3609-b655-1220   1
```

Step 5 Disable load balancing on RB1 by setting the number of equal-cost routes for load balancing to 1.

```
[~RB1] trill
[*RB1-trill] maximum load-balance 1
[*RB1-trill] commit
```

# Check the unicast routing table of RB1.

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

TRILL Unicast Routing Table
--------------------------------------------------------------------
Flags: D-Download To Fib

Total Route(s): 4

Nickname       Cost Flag OutInterface  OuterVlan NextHop            Hop
-----------------------------------------------------------------------
     200       4000 D    10GE1/0/1            10 400/3609-b654-1220   2
     300       4000 D    10GE1/0/2            10 500/3609-b655-1220   2
     400       2000 D    10GE1/0/1            10 400/3609-b654-1220   1
     500       2000 D    10GE1/0/2            10 500/3609-b655-1220   1
```

The preceding table shows that the outbound interface of the route to Nickname 200 is 10GE1/0/1. After the maximum number of equal-cost routes for load balancing is configured to 1, the route with the next hop RB4 is selected as the optimal route because RB4 has a smaller interface index.

Step 6 Restore the default number of equal-cost routes for load balancing on RB1.

```
[*RB1-trill] undo maximum load-balance
[*RB1-trill] commit
```

# Check the TRILL unicast routing table of RB1.

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

TRILL Unicast Routing Table
--------------------------------------------------------------------
Flags: D-Download To Fib

Total Route(s): 4

Nickname       Cost Flag OutInterface  OuterVlan NextHop            Hop
-----------------------------------------------------------------------
     200       4000 D    10GE1/0/1            10 400/3609-b654-1220   2
                         10GE1/0/2            10 500/3609-b655-1220   2
     300       4000 D    10GE1/0/2            10 500/3609-b655-1220   2
     400       2000 D    10GE1/0/1            10 400/3609-b654-1220   1
     500       2000 D    10GE1/0/2            10 500/3609-b655-1220   1
```

The preceding table shows the valid routes with the next hops RB4 (3609-b654-1220) and RB5 (3609-b655-1220) of RB1. By default, the maximum number of equal-cost routes for load balancing is 32.

----End

Configuration Files

- Configurations files of RB1

```
#
sysname RB1
#
vlan batch 50 100
#
trill
 network-entity 00.0000.0000.1111.00
 nickname 100
 carrier-vlan 10
 admin-vlan 50
 ce-vlan 100
#
interface Vlanif50
 ip address 192.168.10.1 255.255.255.0
#
interface 10GE1/0/1
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/3
 port default vlan 100
#
return
```

- Configurations files of RB2

```
#
sysname RB2
#
vlan batch 50 100
#
trill
```

```
 network-entity 00.0000.0000.2222.00
 nickname 200
 carrier-vlan 10
 admin-vlan 50
 ce-vlan 100
#
interface Vlanif50
 ip address 192.168.10.2 255.255.255.0
#
interface 10GE1/0/1
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/3
 port default vlan 100
#
return
```

- Configurations files of RB3

```
#
sysname RB3
#
vlan batch 50 100
#
trill
 network-entity 00.0000.0000.3333.00
 nickname 300
 carrier-vlan 10
 admin-vlan 50
 ce-vlan 100
#
interface Vlanif50
 ip address 192.168.10.3 255.255.255.0
#
interface 10GE1/0/1
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/3
 port default vlan 100
#
return
```

- Configurations files of RB4

```
#
sysname RB4
#
vlan batch 50
#
trill
 network-entity 00.0000.0000.4444.00
 nickname 400 root-priority 65535
 carrier-vlan 10
 admin-vlan 50
#
interface Vlanif50
 ip address 192.168.10.4 255.255.255.0
#
interface 10GE1/0/1
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
```

```
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
return
```

- Configurations files of RB5

```
#
sysname RB5
#
vlan batch 50
#
trill
 network-entity 00.0000.0000.5555.00
 nickname 500 root-priority 65535
 carrier-vlan 10
 admin-vlan 50
#
interface Vlanif50
 ip address 192.168.10.5 255.255.255.0
#
interface 10GE1/0/1
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/3
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
return
```

8.2 Example for Configuring the Association Between MSTP and TRILL

Networking Requirements

As shown in **Figure 8-2**, an enterprise's data center network originally runs MSTP to avoid loops. The customer plans to expand the network capacity to construct a large Layer 2 network using TRILL. It is required that the original MSTP network be dual-homed to the TRILL network to implement seamless expansion.

NOTE

RB5 is a core device. The configuration of RB5 is similar to that of RB1 and so is not mentioned here.

Figure 8-2 Networking for configuring the association between MSTP and TRILL

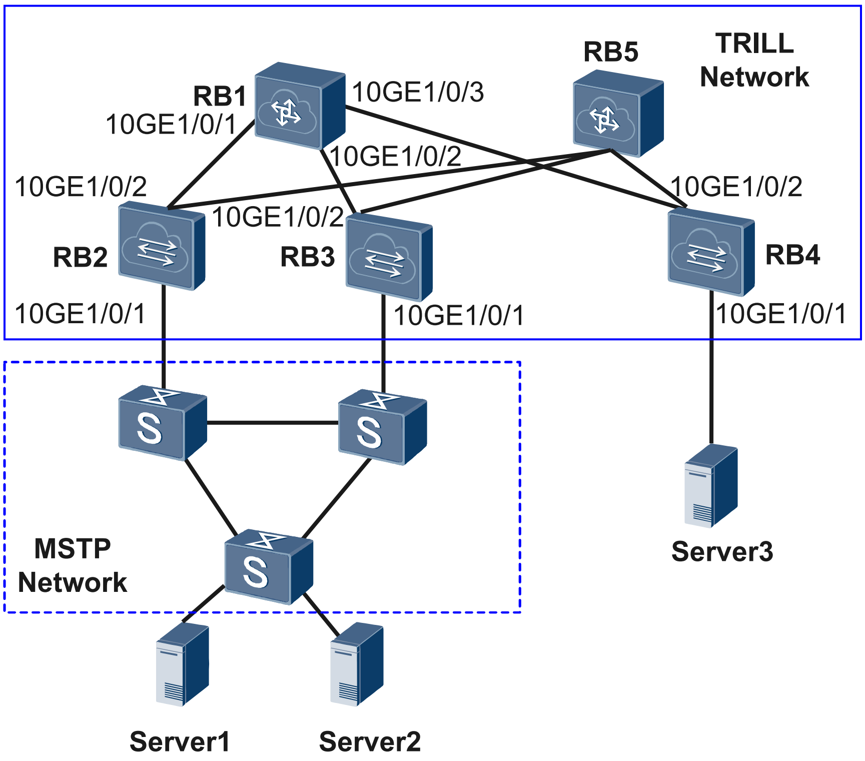

Configuration Roadmap

As the MSTP network needs to be dual-homed to the TRILL network, you are advised to configure the association between MSTP and TRILL to avoid loops. The configuration roadmap is as follows:

1. Configure basic TRILL functions on devices in the new network so that the devices can communicate with each other using TRILL.
2. Configure the association between MSTP and TRILL on the edge device connecting the TRILL network to the MSTP network so that loops are avoided and the TC packets of the MSTP network can be sent through the TRILL network.

Procedure

Step 1 Enable TRILL on each RB, and configure their VLANs, NETs, and nicknames. Core devices have high performance. Therefore, you need to set the root priority of the core device RB1 to 65535 to ensure that RB1 becomes the multicast tree root of the TRILL network.

# Configure RB1.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB1
[*HUAWEI] commit
[~RB1] trill
[*RB1-trill] carrier-vlan 10
[*RB1-trill] network-entity 00.0000.0000.1111.00
[*RB1-trill] nickname 100 root-priority 65535
[*RB1-trill] commit
```

```
[~RB1-trill] quit
```

# Configure RB2.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB2
[*HUAWEI] commit
[~RB2] vlan 100
[*RB2-vlan100] commit
[~RB2-vlan100] quit
[~RB2] trill
[*RB2-trill] carrier-vlan 10
[*RB2-trill] ce-vlan 100
[*RB2-trill] network-entity 00.0000.0000.2222.00
[*RB2-trill] nickname 200
[*RB2-trill] commit
[~RB2-trill] quit
```

# Configure RB3.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB3
[*HUAWEI] commit
[~RB3] vlan 100
[*RB3-vlan100] commit
[~RB3-vlan100] quit
[~RB3] trill
[*RB3-trill] carrier-vlan 10
[*RB3-trill] ce-vlan 100
[*RB3-trill] network-entity 00.0000.0000.3333.00
[*RB3-trill] nickname 300
[*RB3-trill] commit
[~RB3-trill] quit
```

# Configure RB4.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB4
[*HUAWEI] commit
[~RB4] vlan 100
[*RB4-vlan100] commit
[~RB4-vlan100] quit
[~RB4] trill
[*RB4-trill] carrier-vlan 10
[*RB4-trill] ce-vlan 100
[*RB4-trill] network-entity 00.0000.0000.4444.00
[*RB4-trill] nickname 400
[*RB4-trill] commit
[~RB4-trill] quit
```

Step 2 Configure TRILL on each interface.

# Configure RB1.

```
[~RB1] interface 10ge 1/0/1
[~RB1-10GE1/0/1] port link-type trunk
[*RB1-10GE1/0/1] trill enable
[*RB1-10GE1/0/1] quit
[*RB1] interface 10ge 1/0/2
[*RB1-10GE1/0/2] port link-type trunk
[*RB1-10GE1/0/2] trill enable
[*RB1-10GE1/0/2] quit
[*RB1] interface 10ge 1/0/3
[*RB1-10GE1/0/3] port link-type trunk
[*RB1-10GE1/0/3] trill enable
[*RB1-10GE1/0/3] commit
[~RB1-10GE1/0/3] quit
```

# Configure RB2.

```
[~RB2] interface 10ge 1/0/2
[~RB2-10GE1/0/2] port link-type trunk
[*RB2-10GE1/0/2] trill enable
[*RB2-10GE1/0/2] undo stp enable
[*RB2-10GE1/0/2] commit
[~RB2-10GE1/0/2] quit
```

# Configure RB3.

```
[~RB3] interface 10ge 1/0/2
[~RB3-10GE1/0/2] port link-type trunk
[*RB3-10GE1/0/2] trill enable
[*RB3-10GE1/0/2] undo stp enable
[*RB3-10GE1/0/2] commit
[~RB3-10GE1/0/2] quit
```

# Configure RB4.

```
[~RB4] interface 10ge 1/0/1
[*RB4-10GE1/0/1] port default vlan 100
[*RB4-10GE1/0/1] quit
[*RB4] interface 10ge 1/0/2
[*RB4-10GE1/0/2] port link-type trunk
[*RB4-10GE1/0/2] trill enable
[*RB4-10GE1/0/2] commit
[~RB4-10GE1/0/2] quit
```

Step 3 Configure the association between MSTP and TRILL.

# Configure an admin VLAN on each device of the TRILL network, and configure VLANIF interfaces for the admin VLANs. RB1 is used as an example here.

```
[~RB1] vlan 50
[*RB1-vlan50] quit
[*RB1] interface vlanif 50
[*RB1-Vlanif50] ip address 10.1.1.1 24
[*RB1-Vlanif50] quit
[*RB1] trill
[*RB1-trill] admin-vlan 50
[*RB1-trill] commit
[~RB1-trill] quit
```

The configurations on RB2, RB3, and RB4 are the same as that on RB1.

NOTICE

The VLANIF interface is configured for the admin VLAN; however, the admin VLAN still carries TRILL traffic, not Layer 3 traffic.

# Add RB2 and RB3 to the MSTP region of the original network. Configure RB2 and RB3 as the root bridge of the MSTP instance, and configure the same bridge MAC address for them. In addition, configure 10GE1/0/1 to allow the traffic of the CE VLAN and admin VLAN to pass through. RB2 is used as an example here.

```
[~RB2] stp region-configuration
[*RB2-mst-region] region-name RG1
[*RB2-mst-region] instance 1 vlan 100
[*RB2-mst-region] quit
[*RB2] stp instance 1 root primary
[*RB2] stp bridge-address 39-39-39
[*RB2] interface 10ge 1/0/1
[~RB2-10GE1/0/1] port link-type trunk
[*RB2-10GE1/0/1] port trunk allow-pass vlan 100
```

```
[*RB2-10GE1/0/1] quit
[*RB2] commit
```

The configuration on RB3 is the same as that on RB2.

# Enable the association between MSTP and TRILL on RB2 and RB3, and specify the VLAN that carries the association packets as the admin VLAN. RB2 is used as an example here.The configuration on RB3 is the same as that on RB2.

```
[~RB2] stp tc-notify trill vlan 50
[*RB2] commit
```

----End

Configuration Files

- Configurations files of RB1

```
#
sysname RB1
#
vlan batch 50
#
trill
 network-entity 00.0000.0000.1111.00
 nickname 100 root-priority 65535
 carrier-vlan 10
 admin-vlan 50
#
interface Vlanif50
 ip address 10.1.1.1 255.255.255.0
#
interface 10GE1/0/1
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/3
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
return
```

- Configurations files of RB2

```
#
sysname RB2
#
vlan batch 50 100
#
stp bridge-address 0039-0039-0039
stp tc-notify trill vlan 50
stp instance 1 root primary
#
stp region-configuration
 region-name RG1
 instance 1 vlan 100
#
trill
 network-entity 00.0000.0000.2222.00
 nickname 200
 carrier-vlan 10
 admin-vlan 50
 ce-vlan 100
#
interface Vlanif50
 ip address 10.1.1.2 255.255.255.0
#
interface 10GE1/0/1
 port link-type trunk
 port trunk allow-pass vlan 100
#
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 stp disable
 trill enable
#
return
```

- Configurations files of RB3

```
#
sysname RB3
#
vlan batch 50 100
#
stp bridge-address 0039-0039-0039
stp tc-notify trill vlan 50
stp instance 1 root primary
#
stp region-configuration
 region-name RG1
 instance 1 vlan 100
#
trill
 network-entity 00.0000.0000.3333.00
 nickname 300
 carrier-vlan 10
 admin-vlan 50
 ce-vlan 100
#
interface Vlanif50
 ip address 10.1.1.3 255.255.255.0
#
interface 10GE1/0/1
 port link-type trunk
 port trunk allow-pass vlan 100
#
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 stp disable
 trill enable
#
return
```

- Configurations files of RB4

```
#
sysname RB4
#
vlan batch 50 100
#
trill
 network-entity 00.0000.0000.4444.00
 nickname 400
 carrier-vlan 10
 admin-vlan 50
 ce-vlan 100
#
interface Vlanif50
 ip address 10.1.1.4 255.255.255.0
#
interface 10GE1/0/1
```

```
 port default vlan 100
#
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
return
```

8.3 Example for Deploying M-LAG on a Dual-Homing TRILL Network Through V-STP

Networking Requirements

As shown in **Figure 8-3**, the switch connects to a TRILL network. The requirements are as follows:

- When one access link fails, traffic can be fast switched to the other link to ensure reliability.
- The load balancing mode can be used to forward traffic to make full use of bandwidth and ensure that two links are in active state.

Figure 8-3 Dual homing to a TRILL network through M-LAG

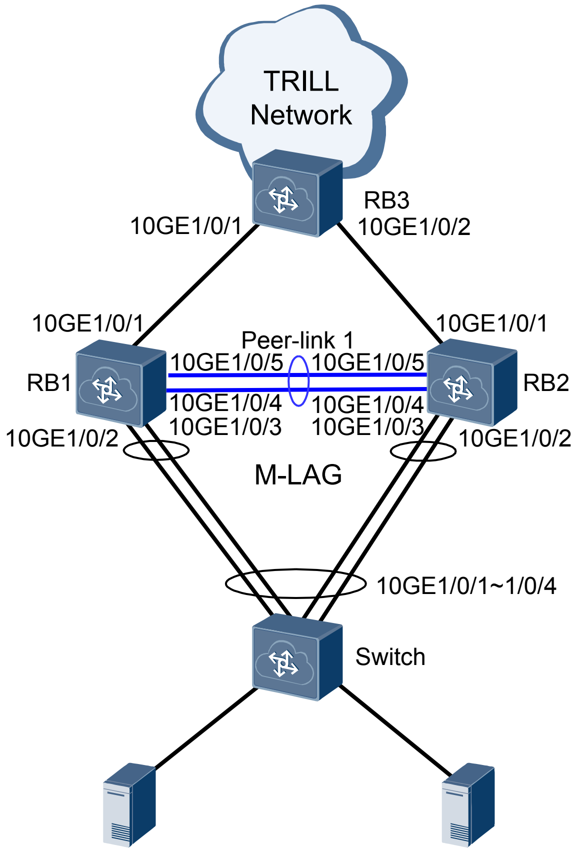

You can configure M-LAG and connect the device to the TRILL network through two RBs. M-LAG is configured between the device and RBs to enhance device-level and link-level reliability. RB1 and RB2 use the same pseudo nickname to form a logical device. On RBs, the DFS group is associated with M-LAG and TRILL. M-LAG and TRILL ensure that service packets are correctly forwarded.

Configuration Roadmap

The configuration roadmap is as follows:

1. Configure basic TRILL functions on RBs to implement interworking of the TRILL network.
2. Configure M-LAG on RB1 and RB2 so that the device is dual-homed to RB1 and RB2.
3. On RB1 and RB2, associate uplink and downlink interfaces with the Monitor Link group to prevent a user-side traffic forwarding failure and traffic loss due to the uplink fault.

Procedure

Step 1 Configure basic TRILL functions on RBs.

# Configure RB1.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB1
[*HUAWEI] commit
[~RB1] vlan batch 100
[*RB1] trill
[*RB1-trill] network-entity 00.0000.0000.1111.00
[*RB1-trill] nickname 100
[*RB1-trill] carrier-vlan 2
[*RB1-trill] ce-vlan 100
[*RB1-trill] quit
[*RB1] interface 10ge 1/0/1
[*RB1-10GE1/0/1] port link-type trunk
[*RB1-10GE1/0/1] trill enable
[*RB1-10GE1/0/1] quit
[*RB1] commit
```

# Configure RB2.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB2
[*HUAWEI] commit
[~RB2] vlan batch 100
[*RB2] trill
[*RB2-trill] network-entity 00.0000.0000.2222.00
[*RB2-trill] nickname 200
[*RB2-trill] carrier-vlan 2
[*RB2-trill] ce-vlan 100
[*RB2-trill] quit
[*RB2] interface 10ge 1/0/1
[*RB2-10GE1/0/1] port link-type trunk
[*RB2-10GE1/0/1] trill enable
[*RB2-10GE1/0/1] quit
[*RB2] commit
```

# Configure RB3.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB3
[*HUAWEI] commit
[~RB3] trill
[*RB3-trill] network-entity 00.0000.0000.3333.00
[*RB3-trill] nickname 400
[*RB3-trill] carrier-vlan 2
[*RB3-trill] quit
```

```
[*RB3] interface 10ge 1/0/1
[*RB3-10GE1/0/1] port link-type trunk
[*RB3-10GE1/0/1] trill enable
[*RB3-10GE1/0/1] quit
[*RB3] interface 10ge 1/0/2
[*RB3-10GE1/0/2] port link-type trunk
[*RB3-10GE1/0/2] trill enable
[*RB3-10GE1/0/2] quit
[*RB3] commit
```

After the configuration is complete, run the **display trill peer** command to check TRILL peer information.

```
[~RB1] display trill peer

Peer information for TRILL
--------------------------------------------------------------------------------

Total Peer(s): 1

System ID/Trill Name  Interface      Circuit ID        State  HoldTime Type  PRI
--------------------------------------------------------------------------------
0000.0000.3333        10GE1/0/1      0000.0000.3333.02 Report       7s L1     64

```

Step 2 On the switch, bind the uplink interface to an Eth-Trunk.

# Configure the switch.

```
<HUAWEI> system-view
[~HUAWEI] sysname Switch
[*HUAWEI] commit
[~Switch] vlan batch 100
[*Switch] interface eth-trunk 20
[*Switch-Eth-Trunk20] mode lacp-static
[*Switch-Eth-Trunk20] port link-type trunk
[*Switch-Eth-Trunk20] port trunk allow-pass vlan 100
[*Switch-Eth-Trunk20] trunkport 10ge 1/0/1 to 1/0/4
[*Switch-Eth-Trunk20] quit
[*Switch] commit
```

Step 3 Configure the V-STP, DFS group, peer-link, and M-LAG interface on RB1 and RB2.

# Configure RB1.

```
[~RB1] stp mode rstp
[*RB1] stp v-stp enable
[*RB1] dfs-group 1
[*RB1-dfs-group-1] source nickname 100
[*RB1-dfs-group-1] pseudo-nickname 500
[*RB1-dfs-group-1] priority 150
[*RB1-dfs-group-1] quit
[*RB1] interface eth-trunk 1
[*RB1-Eth-Trunk1] trunkport 10ge 1/0/4
[*RB1-Eth-Trunk1] trunkport 10ge 1/0/5
[*RB1-Eth-Trunk1] mode lacp-static
[*RB1-Eth-Trunk1] peer-link 1
[*RB1-Eth-Trunk1] port vlan exclude 2
[*RB1-Eth-Trunk1] quit
[*RB1] interface eth-trunk 10
[*RB1-Eth-Trunk10] mode lacp-static
[*RB1-Eth-Trunk10] port link-type trunk
[*RB1-Eth-Trunk10] port trunk allow-pass vlan 100
[*RB1-Eth-Trunk10] trunkport 10ge 1/0/2
[*RB1-Eth-Trunk10] trunkport 10ge 1/0/3
[*RB1-Eth-Trunk10] dfs-group 1 m-lag 1
[*RB1-Eth-Trunk10] quit
[*RB1] commit
```

# Configure RB2.

```
[~RB2] stp mode rstp
[*RB2] stp v-stp enable
[*RB2] dfs-group 1
[*RB2-dfs-group-1] source nickname 200
[*RB2-dfs-group-1] pseudo-nickname 500
[*RB2-dfs-group-1] priority 120
[*RB2-dfs-group-1] quit
[*RB2] interface eth-trunk 1
[*RB2-Eth-Trunk1] trunkport 10ge 1/0/4
[*RB2-Eth-Trunk1] trunkport 10ge 1/0/5
[*RB2-Eth-Trunk1] mode lacp-static
[*RB2-Eth-Trunk1] peer-link 1
[*RB2-Eth-Trunk1] port vlan exclude 2
[*RB2-Eth-Trunk1] quit
[*RB2] interface eth-trunk 10
[*RB2-Eth-Trunk10] mode lacp-static
[*RB2-Eth-Trunk10] port link-type trunk
[*RB2-Eth-Trunk10] port trunk allow-pass vlan 100
[*RB2-Eth-Trunk10] trunkport 10ge 1/0/2
[*RB2-Eth-Trunk10] trunkport 10ge 1/0/3
[*RB2-Eth-Trunk10] dfs-group 1 m-lag 1
[*RB2-Eth-Trunk10] quit
[*RB2] commit
```

Step 4 Configure the LACP M-LAG system priority and system ID on RB1 and RB2.

# Configure RB1.

```
[~RB1] lacp m-lag priority 10
[*RB1] lacp m-lag system-id 00e0-fc00-0000
[*RB1] commit
```

# Configure RB2.

```
[~RB2] lacp m-lag priority 10
[*RB2] lacp m-lag system-id 00e0-fc00-0000
[*RB2] commit
```

Step 5 On RB1 and RB2, associate uplink and downlink interfaces with the Monitor Link group.

# Configure RB1.

```
[~RB1] monitor-link group 1
[*RB1-mtlk-group1] port 10ge 1/0/1 uplink
[*RB1-mtlk-group1] port eth-trunk 10 downlink 1
[*RB1-mtlk-group1] quit
[*RB1] commit
```

# Configure RB2.

```
[~RB2] monitor-link group 1
[*RB2-mtlk-group1] port 10ge 1/0/1 uplink
[*RB2-mtlk-group1] port eth-trunk 10 downlink 1
[*RB2-mtlk-group1] quit
[*RB2] commit
```

Step 6 Verify the configuration.

Run the **display dfs-group** command to check M-LAG information.

# Check information about the M-LAG with DFS group 1.

```
[~RB1] display dfs-group 1 m-lag
*                : Local node
Heart beat state : OK
Node 1 *
  Dfs-Group ID   : 1
  Priority       : 150
  Address        : nickname 100
  State          : Master
  Causation      : -
  System ID      : 0025-9e95-7c31
  SysName        : RB1
  Version        : V100R006C00
  Device Type    : CE12800
Node 2
  Dfs-Group ID   : 1
  Priority       : 120
  Address        : nickname 200
  State          : Backup
  Causation      : -
  System ID      : 0025-9e95-7c11
  SysName        : RB2
  Version        : V100R006C00
  Device Type    : CE12800
```

# Check M-LAG information on RB1.

```
[~RB1] display dfs-group 1 node 1 m-lag brief
* - Local node

M-Lag ID     Interface      Port State    Status
       1     Eth-Trunk 10   Up            active(*)-active
```

# Check M-LAG information on RB2.

```
[~RB1] display dfs-group 1 node 2 m-lag brief
* - Local node

M-Lag ID     Interface      Port State    Status
       1     Eth-Trunk 10   Up            active-active(*)
```

In the preceding information, the value of **Heart beat state** is **OK**, indicating that the heartbeat is normal. RB1 is used as Node 1, its priority is 150, and its status is **Master**. RB2 is used as Node 2, its priority is 120, and its status is **Backup**. The value of **Causation** is **-**, the values of **Port State** of Node 1 and Node 2 are both **Up**, and the M-LAG status of both Node 1 and Node 2 is **active**, indicating that the M-LAG configuration is correct.

----End

Configuration Files

- RB1 configuration file

```
#
sysname RB1
#
dfs-group 1
 priority 150
 source nickname 100
 pseudo-nickname 500
#
vlan batch 100
#
stp mode rstp
stp v-stp enable
#
lacp m-lag system-id 00e0-fc00-0000
lacp m-lag priority 10
#
trill
 network-entity 00.0000.0000.1111.00
 nickname 100
 carrier-vlan 2
 ce-vlan 100
#
interface Eth-Trunk1
 mode lacp-static
 peer-link 1
 port vlan exclude 2
#
```

```
interface Eth-Trunk10
 port link-type trunk
 port trunk allow-pass vlan 100
 mode lacp-static
 dfs-group 1 m-lag 1
#
interface 10GE1/0/1
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/2
 eth-trunk 10
#
interface 10GE1/0/3
 eth-trunk 10
#
interface 10GE1/0/4
 eth-trunk 1
#
interface 10GE1/0/5
 eth-trunk 1
#
monitor-link group 1
 port 10GE1/0/1 uplink
 port Eth-Trunk10 downlink 1
#
return
```

- RB2 configuration file

```
#
sysname RB2
#
dfs-group 1
 priority 120
 source nickname 200
 pseudo-nickname 500
#
vlan batch 100
#
stp mode rstp
stp v-stp enable
#
lacp m-lag system-id 00e0-fc00-0000
lacp m-lag priority 10
#
trill
 network-entity 00.0000.0000.2222.00
 nickname 200
 carrier-vlan 2
 ce-vlan 100
#
interface Eth-Trunk1
 mode lacp-static
 peer-link 1
 port vlan exclude 2
#
interface Eth-Trunk10
 port link-type trunk
 port trunk allow-pass vlan 100
 mode lacp-static
 dfs-group 1 m-lag 1
#
interface 10GE1/0/1
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/2
 eth-trunk 10
```

```
#
interface 10GE1/0/3
 eth-trunk 10
#
interface 10GE1/0/4
 eth-trunk 1
#
interface 10GE1/0/5
 eth-trunk 1
#
monitor-link group 1
 port 10GE1/0/1 uplink
 port Eth-Trunk10 downlink 1
#
return
```

- RB3 configuration file

```
#
sysname RB3
#
trill
 network-entity 00.0000.0000.3333.00
 nickname 400
 carrier-vlan 2
#
interface 10GE1/0/1
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
return
```

- Switch configuration file

```
#
sysname Switch
#
vlan batch 100
#
interface Eth-Trunk20
 port link-type trunk
 port trunk allow-pass vlan 100
 mode lacp-static
#
interface 10GE1/0/1
 eth-trunk 20
#
interface 10GE1/0/2
 eth-trunk 20
#
interface 10GE1/0/3
 eth-trunk 20
#
interface 10GE1/0/4
 eth-trunk 20
#
return
```

8.4 Example for Configuring the TRILL Gateway Function

Networking Requirements

In **Figure 8-4**, a CE device is single-homed to a TRILL network and needs to access external networks through Layer 3 forwarding. To support this access requirement, a gateway needs to be deployed on the TRILL network. Deploying a gateway device using virtual systems (VSs) or independently requires high costs and complex configurations. You can configure the TRILL gateway function on a core device to enable the core device to integrate the TRILL device functions and gateway functions. To improve reliability, deploy the Virtual Router Redundancy Protocol (VRRP) for gateway redundancy while deploying the TRILL gateway function.

Figure 8-4 Configuring the TRILL gateway function

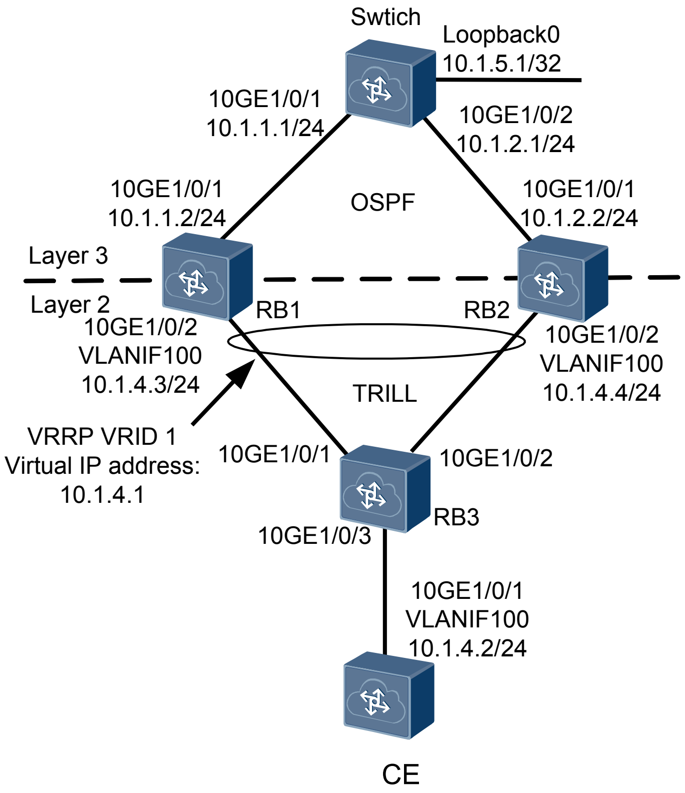

Configuration Roadmap

The configuration roadmap is as follows:

1. Configure basic TRILL functions on router bridges (RBs) for interconnection on the TRILL network.
2. Configure VRRP to implement TRILL gateway redundancy.
3. Configure the TRILL gateway function on the core device to enable the device to forward TRILL packets at Layer 3 after terminating the packets.

Procedure

Step 1 Configure basic TRILL functions on RBs.

# Configure RB1.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB1
[*HUAWEI] commit
[~RB1] vlan batch 50 100
[*RB1] trill
[*RB1-trill] network-entity 00.0000.0000.1111.00
[*RB1-trill] nickname 100
[*RB1-trill] carrier-vlan 2
[*RB1-trill] ce-vlan 100
[*RB1-trill] admin-vlan 50
[*RB1-trill] quit
[*RB1] interface 10ge 1/0/2
[*RB1-10GE1/0/2] port link-type trunk
[*RB1-10GE1/0/2] trill enable
[*RB1-10GE1/0/2] quit
[*RB1] commit
```

# Configure RB2.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB2
[*HUAWEI] commit
[~RB2] vlan batch 50 100
[*RB2] trill
[*RB2-trill] network-entity 00.0000.0000.2222.00
[*RB2-trill] nickname 200
[*RB2-trill] carrier-vlan 2
[*RB2-trill] ce-vlan 100
[*RB2-trill] admin-vlan 50
[*RB2-trill] quit
[*RB2] interface 10ge 1/0/2
[*RB2-10GE1/0/2] port link-type trunk
[*RB2-10GE1/0/2] trill enable
[*RB2-10GE1/0/2] quit
[*RB2] commit
```

# Configure RB3.

```
<HUAWEI> system-view
[~HUAWEI] sysname RB3
[*HUAWEI] commit
[~RB3] vlan batch 50 100
[~RB3] trill
[*RB3-trill] network-entity 00.0000.0000.3333.00
[*RB3-trill] nickname 300
[*RB3-trill] carrier-vlan 2
[*RB3-trill] ce-vlan 100
[*RB3-trill] admin-vlan 50
[*RB3-trill] quit
[*RB3] interface 10ge 1/0/1
[*RB3-10GE1/0/1] port link-type trunk
[*RB3-10GE1/0/1] trill enable
[*RB3-10GE1/0/1] quit
[*RB3] interface 10ge 1/0/2
[*RB3-10GE1/0/2] port link-type trunk
[*RB3-10GE1/0/2] trill enable
[*RB3-10GE1/0/2] quit
[*RB3] commit
```

After the preceding configurations are complete, you can run the **display trill peer** command to check TRILL neighbor information.

```
[~RB3] display trill peer

Peer information for TRILL
-------------------------------------------------------------------------------
--

Total Peer(s): 2

```

```
System ID/Trill Name  Interface          Circuit ID        State  HoldTime Type
PRI
----------------------------------------------------------------------------------
--
0000.0000.1111        10GE1/0/1          0000000053        Report      26s L1
--
0000.0000.2222        10GE1/0/2          0000000055        Report      27s L1
--
```

Step 2 Configure OSPF to implement Layer 3 interconnection between the switch and TRILL gateway.

# Configure the switch.

```
<HUAWEI> system-view
[~HUAWEI] sysname switch
[*HUAWEI] commit
[~switch] interface 10ge 1/0/1
[~switch-10GE1/0/1] undo portswitch
[*switch-10GE1/0/1] ip address 10.1.1.1 24
[*switch-10GE1/0/1] quit
[*switch] interface 10ge 1/0/2
[*switch-10GE1/0/1] undo portswitch
[*switch-10GE1/0/1] ip address 10.1.2.1 24
[*switch-10GE1/0/1] quit
[*switch] interface loopback 0
[*switch-LoopBack0] ip address 10.1.5.1 32
[*switch-LoopBack0] quit
[*switch] ospf
[*switch-ospf-1] area 0
[*switch-ospf-1-area-0.0.0.0] network 10.1.1.0 0.0.0.255
[*switch-ospf-1-area-0.0.0.0] network 10.1.2.0 0.0.0.255
[*switch-ospf-1-area-0.0.0.0] network 10.1.5.1 0.0.0.0
[*switch-ospf-1-area-0.0.0.0] quit
[*switch-ospf-1] quit
[*switch] commit
```

# Configure RB1.

```
[~RB1] interface 10ge 1/0/1
[~RB1-10GE1/0/1] undo portswitch
[*RB1-10GE1/0/1] ip address 10.1.1.2 24
[*RB1-10GE1/0/1] quit
[*RB1] interface vlanif 100
[*RB1-Vlanif100] ip address 10.1.4.3 24
[*RB1-Vlanif100] quit
[*RB1] ospf
[*RB1-ospf-1] area 0
[*RB1-ospf-1-area-0.0.0.0] network 10.1.1.0 0.0.0.255
[*RB1-ospf-1-area-0.0.0.0] network 10.1.4.0 0.0.0.255
[*RB1-ospf-1-area-0.0.0.0] quit
[*RB1-ospf-1] quit
[*RB1] commit
```

# Configure RB2.

```
[~RB2] interface 10ge 1/0/1
[~RB2-10GE1/0/1] undo portswitch
[*RB2-10GE1/0/1] ip address 10.1.2.2 24
[*RB2-10GE1/0/1] quit
[*RB2] interface vlanif 100
[*RB2-Vlanif100] ip address 10.1.4.4 24
[*RB2-Vlanif100] quit
[*RB2] ospf
[*RB2-ospf-1] area 0
[*RB2-ospf-1-area-0.0.0.0] network 10.1.2.0 0.0.0.255
[*RB2-ospf-1-area-0.0.0.0] network 10.1.4.0 0.0.0.255
[*RB2-ospf-1-area-0.0.0.0] quit
```

```
[*RB2-ospf-1] quit
[*RB2] commit
```

Step 3 Configure the CE device to be single-homed to the TRILL network and configure a default route on the CE device.

# Configure the CE device.

```
<HUAWEI> system-view
[~HUAWEI] sysname CE
[*HUAWEI] commit
[~CE] vlan batch 100
[*CE] interface 10ge 1/0/1
[*CE-10GE1/0/1] port link-type trunk
[*CE-10GE1/0/1] port trunk allow-pass vlan 100
[*CE-10GE1/0/1] quit
[*CE] interface vlanif 100
[*CE-Vlanif100] ip address 10.1.4.2 24
[*CE-Vlanif100] quit
[*CE] ip route-static 0.0.0.0 0 10.1.4.1
[*CE] commit
```

# Configure RB3.

```
[~RB3] interface 10ge 1/0/3
[~RB3-10GE1/0/3] port link-type trunk
[*RB3-10GE1/0/3] port trunk allow-pass vlan 100
[*RB3-10GE1/0/3] quit
[*RB3] commit
```

Step 4 Configure VRRP to implement gateway redundancy and improve reliability.

# Configure RB1.

```
[~RB1] interface vlanif 100
[~RB1-Vlanif100] vrrp vrid 1 virtual-ip 10.1.4.1
[*RB1-Vlanif100] vrrp vrid 1 priority 120
[*RB1-Vlanif100] vrrp vrid 1 preempt timer delay 20
[*RB1-Vlanif100] quit
[*RB1] commit
```

# Configure RB2.

```
[~RB2] interface vlanif 100
[~RB2-Vlanif100] vrrp vrid 1 virtual-ip 10.1.4.1
[*RB2-Vlanif100] quit
[*RB2] commit
```

Step 5 Configure the TRILL gateway function.

# Configure RB1.

```
[*RB1] trill gateway enable
[*RB1] commit
```

# Configure RB2.

```
[*RB2] trill gateway enable
[*RB2] commit
```

Step 6 Verify the configuration.

Run the **ping** command on the CE device to ping the switch. The command output shows that the CE device pings the switch successfully.

```
[~CE] ping 10.1.5.1
  PING 10.1.5.1: 56  data bytes, press CTRL_C to break
    Reply from 10.1.5.1: bytes=56 Sequence=1 ttl=254 time=18 ms
    Reply from 10.1.5.1: bytes=56 Sequence=2 ttl=254 time=6 ms
    Reply from 10.1.5.1: bytes=56 Sequence=3 ttl=254 time=5 ms
    Reply from 10.1.5.1: bytes=56 Sequence=4 ttl=254 time=4 ms
    Reply from 10.1.5.1: bytes=56 Sequence=5 ttl=254 time=6 ms

  --- 10.1.5.1 ping statistics ---
    5 packet(s) transmitted
    5 packet(s) received
    0.00% packet loss
    round-trip min/avg/max = 4/7/18 ms
```

----End

Configuration Files

- Switch configuration file

```
#
sysname Switch
#

interface
10GE1/0/1
 undo
portswitch
 ip address 10.1.1.1
255.255.255.0
#

interface
10GE1/0/2
 undo
portswitch
 ip address 10.1.2.1
255.255.255.0
#

interface
LoopBack0
 ip address 10.1.5.1
255.255.255.255
#
ospf
1
 area
0.0.0.0
  network 10.1.1.0
0.0.0.255
  network 10.1.2.0
0.0.0.255
  network 10.1.5.1
0.0.0.0
#
return
```

- RB1 configuration file

```
#
sysname RB1
#
trill gateway enable
#
vlan batch 50 100
#
trill
 network-entity 00.0000.0000.1111.00
 nickname 100
 carrier-vlan 2
 admin-vlan 50
 ce-vlan 100
#
```

```
interface Vlanif100
 ip address 10.1.4.3 255.255.255.0
 vrrp vrid 1 virtual-ip 10.1.4.1
 vrrp vrid 1 priority 120
 vrrp vrid 1 preempt timer delay 20
#
interface 10GE1/0/1
 undo portswitch
 ip address 10.1.1.2 255.255.255.0
#
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
ospf 1
 area 0.0.0.0
  network 10.1.1.0 0.0.0.255
  network 10.1.3.0 0.0.0.255
  network 10.1.4.0 0.0.0.255
#
return
```

- RB2 configuration file

```
#
sysname RB2
#
trill gateway enable
#
vlan batch 50 100
#
trill
 network-entity 00.0000.0000.2222.00
 nickname 200
 carrier-vlan 2
 admin-vlan 50
 ce-vlan 100
#
interface Vlanif100
 ip address 10.1.4.4 255.255.255.0
 vrrp vrid 1 virtual-ip 10.1.4.1
#
interface 10GE1/0/1
 undo portswitch
 ip address 10.1.1.2 255.255.255.0
#
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
ospf 1
 area 0.0.0.0
  network 10.1.2.0 0.0.0.255
  network 10.1.3.0 0.0.0.255
  network 10.1.4.0 0.0.0.255
#
return
```

- RB3 configuration file

```
#
sysname RB3
#
vlan batch 50 100
#
trill
 network-entity 00.0000.0000.3333.00
 nickname 300
 carrier-vlan 2
 admin-vlan 50
 ce-vlan 100
#
interface 10GE1/0/1
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/2
 port link-type trunk
 undo port trunk allow-pass vlan 1
 trill enable
#
interface 10GE1/0/3
 port link-type trunk
 port trunk allow-pass vlan 100
#
return
```

- CE configuration file

```
#
sysname CE
#
vlan batch 100
#
interface Vlanif100
 ip address 10.1.4.2 255.255.255.0
#
interface 10GE1/0/1
 port link-type trunk
 port trunk allow-pass vlan 100
#
ip route-static 0.0.0.0 0.0.0.0 10.1.4.1
#
return
```