APPLICATION NOTE

J SERIES AND BRANCH SRX SERIES ETHERNET SWITCHING CONFIGURATION GUIDE

Table of Contents

Table of Figures

Introduction

Juniper Networks® SRX Series Services Gateways for the branch and J Series Services Routers enable the enterprise to provide services without boundaries. The SRX Series products provide a comprehensive suite of Ethernet switching functionality. Ethernet switching features eliminate the need for Layer 2 switches in small branch offices and act as an aggregate switch in medium-sized branch offices.

Juniper Networks Junos® operating system Release 9.2 for J Series routers introduces Ethernet switching features, integrated routing and bridging, and support for several Layer 2 protocols. These features have been present in branch SRX Series Services Gateways since their release.

Scope

This application note provides an overview of the Junos OS Layer 2 features for J Series and branch SRX Series Services Gateways. It describes common deployment scenarios, with detailed configurations. SRX Series data center products (SRX1400, SRX3000 line and SRX5000 line) do not support Ethernet switching functionality. All features discussed in this document reference SRX Series Services Gateways for the branch (Juniper Networks SRX100 Series Services Gateways, SRX200 Series Services Gateway, , and SRX650 Services Gateway). All features and configurations discussed in this document are based on standalone deployment of J Series and branch SRX Series Service Gateways. Please refer SRX technical documentation for Ethernet Switching features in SRX chassis cluster environment.

The Ethernet switching features are limited by both hardware and software. The scope is defined in the following section.

Table 1: Hardware Scope

PLATFORMS	ON-BOARD	UPIM	MPIM	XPIM
J2320	✖	✓	✖	✖
J2350	✖	✓	✖	✖
J4350	✖	✓	✖	✖
J6350	✖	✓	✖	✖
SRX100	✓	✖	✖	✖
SRX110	✓	✖	✖	✖
SRX210	✓	✖	✖*	✖
SRX220	✓	✖	✖*	✖
SRX240	✓	✖	✖*	✖
SRX650	✖	✖	✖	✓**

* Ethernet switching support is planned for future release for 1 Gigabit Ethernet SFP MPIM on the SRX210 and SRX240.

** As of Junos OS Release 10.2, Ethernet switching is not supported on 10GbE XPIM.

Software Scope

Ethernet switching on the J Series and branch SRX Series is based on Juniper Networks EX Series Ethernet Switches functionality. As of Junos OS Release 11.2, the J Series and branch SRX Series support the following:

- Layer 2 switching of traffic, including support for both trunk and access ports
- Routed VLAN interface (or integrated routing and bridging)
- Spanning Tree Protocol (STP)
- Rapid Spanning Tree Protocol (RSTP)
- Multiple Spanning Tree Protocol (MSTP)
- Link aggregation, both static and using Link Aggregation Control Protocol (LACP)

- GARP VLAN Registration Protocol (GVRP)
- IEEE 802.1x authentication
  - Single/single-secure/multiple supplicant modes
  - Dynamic VLAN assignment
  - Guest VLAN and server-reject VLANs
  - RADIUS server failure conditions
  - MAC authentication
  - Authentication bypass
  - VoIP VLAN
- IGMP snooping
- IEEE 802.1ad dot1q tunneling (Q-in-Q)
- Link Layer Discovery Protocol (LLDP)

Limitations in Ethernet Switching Implementation

- As of Junos OS Release 11.2, the following EX Series functionality is not supported on the J Series and branch SRX Series. Additionally, future features added to EX Series platforms are not expected to be automatically ported to the J Series and branch SRX Series.
  - Layer 2 access control lists (ACLs)
  - Quality of service (QoS) for switching ports
  - SNMP MIB support (for the new Layer 2 features)
  - Virtual chassis
  - Port security
  - L2 CoS functionality
- On J Series platforms, Ethernet switching is supported on only one universal PIM (uPIM) per J Series chassis.
- MSTP is not supported on the SRX210.
- The IGMP snooping and Q-in-Q feature is not available for the SRX100.
- The J Series and SRX100 do not support advanced 802.1x features such as dynamic VLAN, guest VLAN, server-
- reject VLAN, server fail operations, and VoIP VLAN. But RADIUS accounting and MAC authentication are available for the SRX100.
- Advanced Q-in-Q features such as push, customer bundling, etc. are only supported on the SRX650.

Only SRX Series Services Gateways for branch support Ethernet switching features in chassis cluster environment. This document discussion Ethernet Switching features on standalone deployments. For Ethernet switching in chassis cluster environment please refer SRX technical documentation.

Most of the limitations discussed in this section are expected to be fixed in later Junos OS releases. Please refer to **Future Support Reference** for more information.

Life of Packet in Ethernet Switching

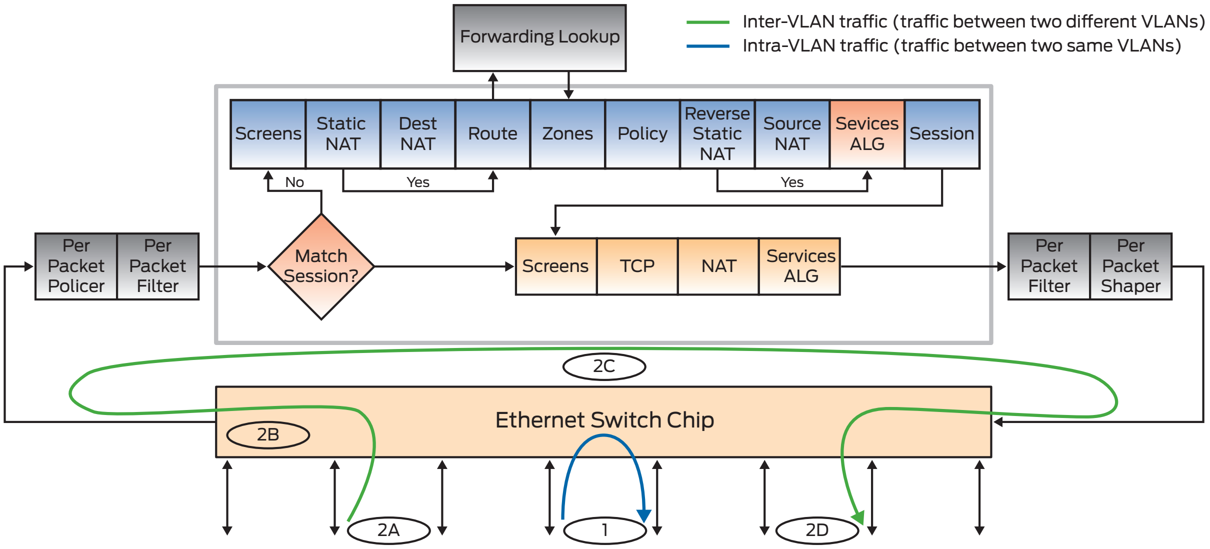

Figure 1: Life of packet in Ethernet switching

1. Intra-VLAN traffic—Once interfaces are configured in the same VLAN through CLI/Juniper Networks J-Web Software, the "Ethernet switch chip" is programmed accordingly, MAC learning, and STP states are maintained at chip. Packets in the same VLAN are switched internally at the Ethernet switch chip. They do not go through a flow architecture, and none of the security features are applied to this traffic.

2. Inter-VLAN traffic—Packets for different VLANs are routed/forwarded through a flow architecture.

2A. Incoming traffic is classified according to port based VLAN.

2B. The destination MAC address of inter-VLAN traffic is matched with the routed VLAN interface at the Ethernet switch chip, and all these packets are sent to a flow module for further processing.

2C. In the flow module, inter-VLAN traffic goes through all security checks and is routed to a different VLAN.

2D. Routed traffic is sent back to the Ethernet switch chip, which further sends out packets through the interface of the destination VLAN.

Junos OS Release 11.2 Ethernet Switching Configuration Scenarios

This section discusses several deployment scenarios and their associated configurations.

Enabling Ethernet Switching on the J Series

The J Series platform supports two different modes of switching. Plain "switching" is legacy bridge mode operation wherein a uPIM is treated as a bridge and all its Ethernet ports are part of this bridge. None of the features discussed in this document are supported in this mode. And details of this mode are beyond the scope of this document. "Enhanced switching" mode converts uPIM on the J Series to a modern L2 switch. All protocols and features discussed in this document are applicable to this mode. Enhanced switching is configured under the `[chassis fpc pic ethernet]` level of the configuration hierarchy. For example, the following configuration enables a PIM in slot 6:

```
fpc 6 {
    pic 0 {
        ethernet {
            pic-mode enhanced-switching;
        }
    }
}
```

Note: In the current implementation, only one universal PIM per chassis can be configured with enhanced switching.

Enabling Ethernet Switching on Branch SRX Series

The Ethernet switching feature is enabled by default on branch SRX Series platforms. There are no explicit configurations required to enable it.

Configuring Layer 2 Switching

Physical interfaces can operate in several modes. When an interface is configured with a Layer 3 address (such as an IPv4, IPv6, or ISO address), the interface routes traffic based on the destination address of each packet. If an interface is not given a Layer 3 address but is configured as part of the Ethernet switching protocol family, the interface forwards traffic based on the link layer destination address. The following configuration defines an interface as a switching port (note that the Layer 2 configuration is limited to unit 0 of an interface):

```
interface {
                    ge-<slot number>/0/<port number> {
                       unit 0 {
                          family ethernet-switching;
                       }
                    }
}
```

Configuring VLAN

As in most modern switches, broadcast domains can be segmented using virtual LANs or VLANs, an approach that allows device segmentation by assigning ports to different broadcast domains. Traffic can be forwarded between member interfaces of the same VLAN, but not between interfaces that belong to different VLANs, effectively allowing the same physical device to be shared between different non-connected networks (a later section of this document describes how to forward traffic between different VLANs).

By default, all switching-enabled ports form part of the same broadcast domain. If an interface is enabled for Layer 2 switching but not associated with any VLAN, it becomes part of the default VLAN (VLAN ID 1 in the J Series and SRX Series). To configure a new domain, a VLAN has to be defined under the [vlans] hierarchy and given a unique identifier (VLAN ID).

```
vlans {
                    <vlan name> {
                       vlan-id <id>;
                    }
}
```

Please note the following limitation in the J Series and branch SRX Series devices for using VLAN IDs.

Table 2: Supported VLAN Range on J Series and branch SRX Series

PLATFORM	SUPPORTED VLAN RANGE
J Series	1-4094
SRX100	1-4094
SRX110	1-4094
SRX210	**1-4094***
SRX220	**1-4094***
SRX240	1-3967
SRX650	1-3967

*VLAN 4093 is reserved for internal purpose in the SRX200 line.

Attaching Switch Ports to VLANs

Additionally, you can specify which interfaces are part of the newly created VLAN. There are two ways to allocate interfaces. (These ways are identical from a functional point of view—it is up to you to choose the method you prefer). The first way, under the `[interface <name> unit 0 family ethernet-switching]` hierarchy, is to declare the VLAN as part of an interface configuration.

```
interface {
                    ge-<slot number>/0/<port number> {
                       unit 0 {
                          family ethernet-switching {
                             vlan members <vlan name or id>
                          }
                       }
                    }
}
```

The second way, under the `[vlan <vlan name> interface]` hierarchy, is to define VLAN member interfaces.

```
vlans {
                    <name> {
                       interfaces {
                          <interface name>;
                          <interface name>;
                          …
                       }
                    }
}
```

Extending Broadcast Domains and Configuring Tagged Interfaces

Modern switching networks can be large enough to require the use of multiple switches (some require a tiered approach, with many switching layers). When multiple bridging domains span more than one switching device, it is convenient to allow traffic from many domains to be forwarded through the same link, while still separating the traffic from different domains. VLAN tagging (IEEE 802.1q) provides this functionality by extending the Ethernet header with a VLAN identifier (a 12-bit value) used to differentiate traffic from different VLANs. VLAN tagging reduces the number of interfaces needed to connect devices because a single interface can then carry traffic from multiple domains. Switching interfaces that carry tagged traffic are referred to as trunk ports. An interface is called an access port when it carries single VLAN untagged traffic. An access port cannot be part of multiple VLANs.

```
interface {
                    ge-*/*/* {
                       unit 0 {
                          family ethernet-switching {
                             port-mode trunk;
                             vlan {
                                members [<vlan name or id> <vlan name or id> …]
                             }
                          }
                       }
                    }
}
```

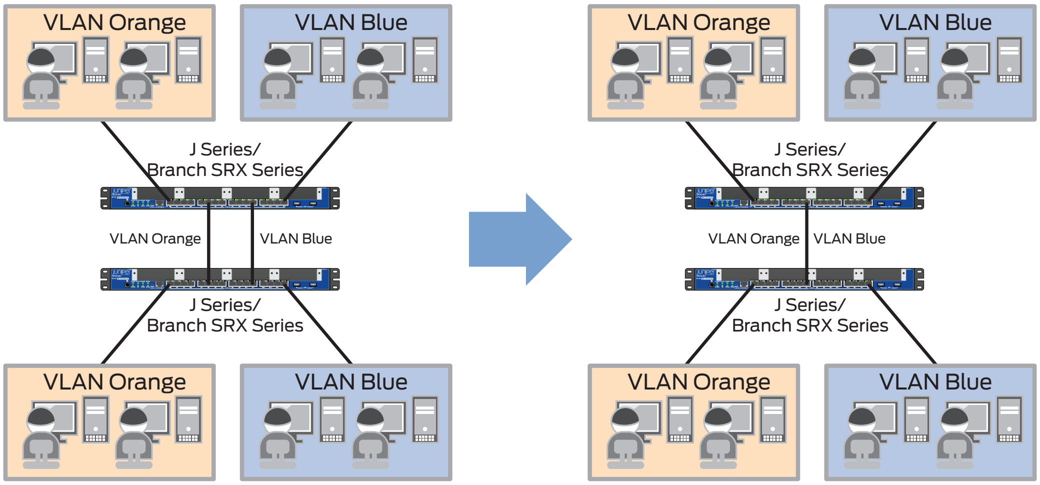

Figure 2: VLAN tagging

By default, all switching interfaces are access ports. An interface can be configured as a trunk port by simply setting the port-mode value to trunk under the `[family ethernet-switching]`. As shown in Figure 1, a trunk port can then be defined as part of multiple VLANs, which allows a switching port defined as a trunk port to be associated with more than one VLAN. Traffic forwarded from a trunk port is tagged using the VLAN ID of the originating VLAN, while received traffic is forwarded to the appropriate VLAN for distribution.

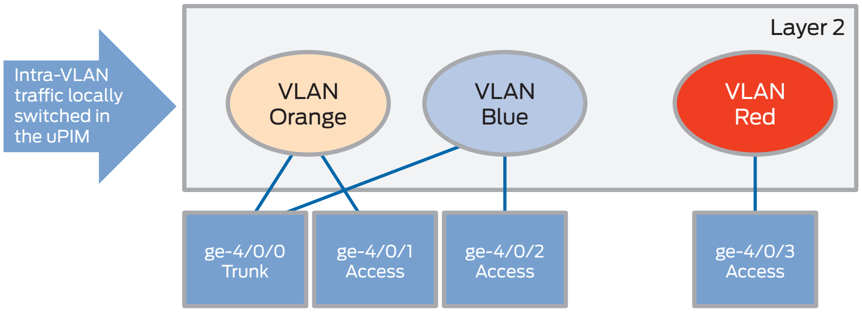

Figure 3: Trunk and access ports

Configuring Routed VLAN Interface (Integrated Routing and Bridging)

As previously discussed, traffic can be forwarded between member interfaces of the same VLAN, but not between interfaces that belong to different VLANs. Traffic inside the same VLAN is switched and traffic across a different VLAN is routed. Hence, Layer 3 device/interfaces are needed to forward traffic from one VLAN to another VLAN. The J Series and SRX Series provide logical Layer 3 interfaces called routed VLAN interfaces (or integrated routing and bridging) for this purpose. Each VLAN domain is tied to one of the logical routed VLAN interfaces. This scenario is equivalent to placing a switch in front of a router. Traffic that is not destined for the router is switched based on the Layer 2 information, and traffic that reaches the router is forwarded based on the Layer 3 information. As different VLAN domains can have unique Layer 3 addresses, traffic between VLAN domains can then be routed by Junos OS software—provided that security policies allow it.

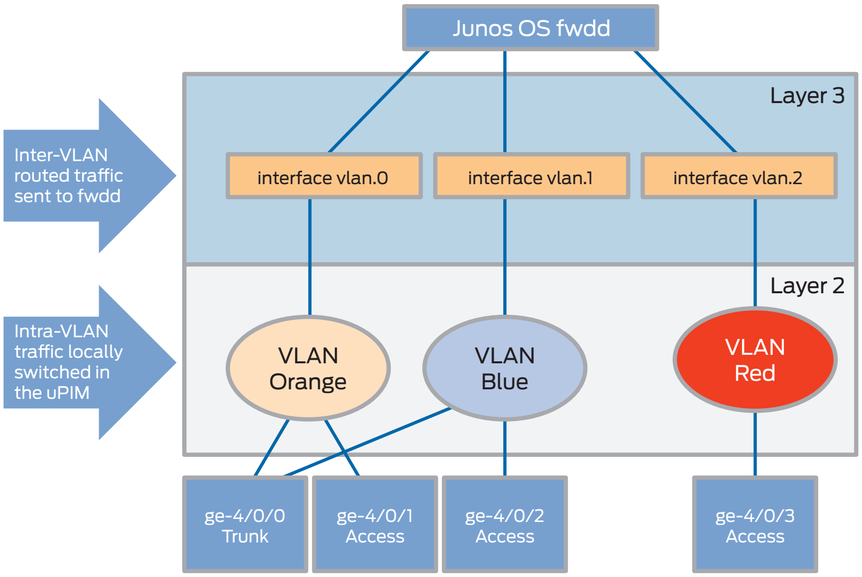

Figure 4: Intra-VLAN and inter-VLAN packet forwarding

A logical Layer 3 interface or routed VLAN interface can be created under the [interfaces vlan] hierarchy. After the logical interface is created, it must be associated with a particular VLAN using the l3-interface keyword.

```
interfaces {
                        vlan {
                            unit <unit number> {
                                family {
                                    inet {
                                        address <ip address>/<netmask>;
                                    }
                                }
                            }
                        }
}
vlans {
                        <vlan name> {
                            l3-interface vlan.<unit of newly created vlan ifl>;
                        }
}
```

Routed VLAN interfaces are no different than any other Layer 3 interfaces in Junos OS and thus require the same configuration. In particular, these interfaces have to be assigned to a security zone, and security policies have to explicitly allow traffic to be forwarded between these interfaces and any other configured Layer 3 interfaces.

Configuring Link Aggregation

When connecting two switches together, sometimes it is advantageous to use two or more parallel connections, normally to provide redundancy. It is also desirable to increase bandwidth between switches. The challenge is that Layer 2 networks have to be loop free, and loop avoidance protocols such as Spanning Tree Protocol (and all its variations and extensions such as RSTP and MSTP) deactivate all but one of these parallel connections—allowing parallel connections to solve the redundancy problem, but not the bandwidth limitation.

The solution to this problem is to use link aggregation, which defines how to load-balance traffic across multiple links (while guaranteeing that packets from a given flow are not reordered). The physical interfaces that form part of a link aggregation group can be statically configured or negotiated between endpoints using LACP (specified in IEEE 802.3ad). Endpoints are normally switches, but they can be servers with multiple network interface cards or NICs.

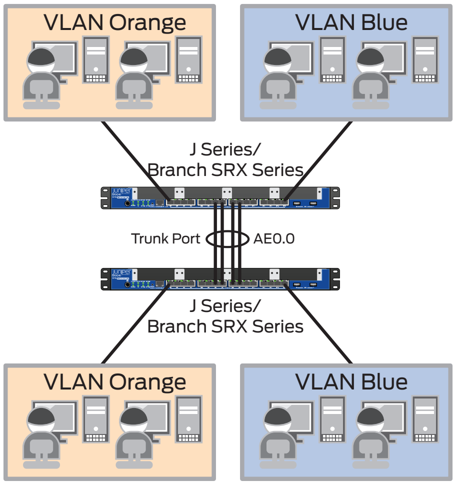

Figure 5: Link aggregation

To configure link aggregation, first create an aggregate interface by defining the number of aggregated interfaces in the system and associate all the physical interfaces that are part of the aggregate bundle to one of the newly created aggregated interfaces.

```
chassis {
aggregated-devices {
    ethernet {
        device-count <number of aggregated interfaces to create>;
    }
}
}
```

Aggregate device count refers to the total number of aggregated interfaces in the system and not the number of physical interfaces per aggregate bundle.

This configuration creates aggregate interfaces named ae0 to ae<device-count -1>. After these interfaces are created, you have to associate physical interfaces with them, which you do under the gigabit-ethernet-options hierarchy.

```
interface {
                    <interface name> {
                       gigabit-ethernet-options {
                          802.3ad {
                             <bndle interface name>;
                          }
                       }
                    }
}
```

LACP is not required between, but if configured, it enables automatic traffic switchover when one or more links fail. It also prevents common misconfiguration errors by confirming that both devices are set up for link aggregation. LACP can be enabled under the aggregated-ethernet-options section of the aggregated interface (make sure that at least one of the endpoints is configured as active, as passive endpoint does not initiate LACP PDU exchange). Link-speed under aggregated-ethernet-options specifies the link speed of each member interface that joins the bundle. And minimum-links keyword specifies the minimum number of active links required for the bundle to be considered "up." The default value of minimum-links is 1 for the J Series and branch SRX Series devices. A maximum of eight links can be bundled in a single AE (LAG) interface.

```
interface {
                    <aggregate interface name> {
        aggregated-ether-options {
                              link-speed [100m|1g];
                    minimum-links <number from 1 to 8>;
                    lacp {
                    active|passive;
              }
         }
         }
}
```

After a bundle interface is created, it can be configured just like any other interface. For example, you can enable switching, add the interface to a VLAN (or a group of VLANs), and enable VLAN tagging.

Configuring Spanning Tree Protocol

Layer 2 switching networks tend to create loops in the network when there are redundant paths available between the source and destination. When such loops are created, a single packet can cause enormous traffic and easily bring down an entire Layer 2 network. J Series Services Routers and SRX Series Services Gateways provide loop prevention in Layer 2 switching networks using STP, RSTP, and MSTP. A loop-free network in spanning-tree topologies is created through the exchange of a special type of frame called bridge protocol data unit (BPDU). Peer STP applications running on the switch interfaces use BPDUs to communicate. Ultimately, the exchange of BPDUs determines which interfaces block traffic (preventing loops) and which interfaces forward traffic.

STP uses the information provided by the BPDUs to elect a root bridge/switch, identify root ports for each switch, identify designated ports for each physical LAN segment, and prune specific redundant links to create a loop-free tree topology. All leaf devices calculate the best path to the root device and place their ports in blocking or forwarding states based on the best path to the root. The resulting tree topology provides a single active Layer 2 data path between any two end stations.

Spanning Tree Protocol (IEEE 802.1D)

STP is a legacy protocol defined in the IEEE 802.1D standard. STP is configured under the [edit protocol] hierarchy.

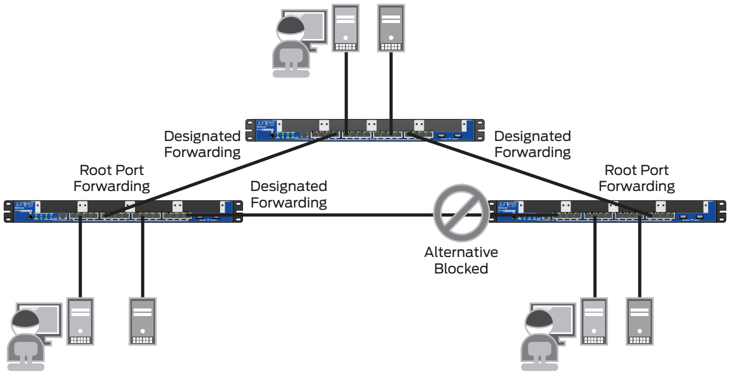

Figure 6: Spanning Tree Protocol

```
protocols {
    stp {
                    bridge-priority <bridge priority>;
                    interface <interface name> {
                    cost <interface cost>;
                    }
    }
}
```

Junos OS provides a number of options to control over the Spanning Tree Protocol. Bridge priority of L2 switches determines which switch to be the root of the network (the switch with the lowest priority is elected as the root of the topology). It also an important parameter in determining root port (the interface that connects to the root of the topology). In Junos OS, bridge priority can be configured under [protocols stp] with a keyword bridge priority with value multiples of 4k, starting with 0 up to 60k. The default bridge priority value is 32k. Another important parameter that controls the Spanning Tree Protocol is link cost. Link costs are dependent upon interface speed. But link costs can be overridden with configuration under [protocols stp interface <interface name]. Junos OS provides other configuration options such as hello-time, forward-delay, and max-age, which control timer mechanisms in protocol state machine.

Rapid Spanning Protocol (IEEE 802.1w)

Legacy Spanning Tree Protocol is very slow in converging loop-free topology. It takes around 30-50 seconds to converge and start forwarding data packets. Also, topology change propagation is largely dependent on root bridge/ switch. Rapid Spanning Tree Protocol or RSTP is a new standard defined by IEEE to overcome these limitations. RSTP uses a messaging mechanism, unlike the timer mechanism in STP, and it is not dependent on root bridge/switch for propagation of topology in the network. It also introduces new port roles, alternative and backup ports as redundant links for root and designated ports, respectively. In the event of link failures, these alternative or backup ports take over immediately. RSTP can be configured as the following:

```
protocols {
    rstp {
                    bridge-priority <bridge priority>;
                    interface <interface name> {
                    cost <interface cost>;
                    }
                    interface <interface name> {
                    edge;
                    }
    }
}
```

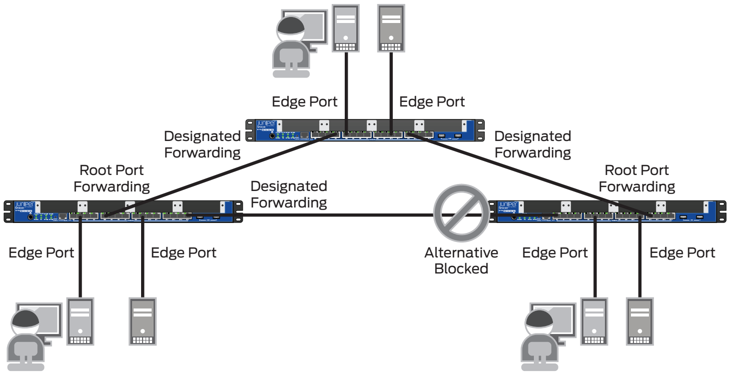

Figure 7: Rapid Spanning Tree Protocol

There is no difference between STP and RSTP in terms of configuration. RSTP also provides configuration options bridge priority and interface cost to control tree topology. An important feature that is available with RSTP is the edge port feature. When an interface is configured as an edge port, it forwards data immediately. And topology changes in the network do not affect the edge port. This configuration is useful when end hosts are connected to interfaces. To avoid the wrong configuration, the edge port starts participating in a spanning-tree state machine when it receives BPDUs. The edge port is configured under the [protocols stp interface <interface name] hierarchy.

Multiple Spanning Tree Protocol

Although RSTP provides faster convergence time than STP, it still does not solve a problem inherent in STP—all VLANs within a LAN must share the same Spanning Tree Protocol. To solve this problem, J Series Services Routers and SRX Series Services Gateways use MSTP to create a loop-free topology in networks with multiple spanning-tree regions.

An MSTP region allows a group of switches to be modeled as a single bridge. Multiple spanning-tree instances (MSTIs) are contained in an MSTP region. MSTIs provide different paths for different VLANs. This functionality facilitates better load sharing across redundant links.

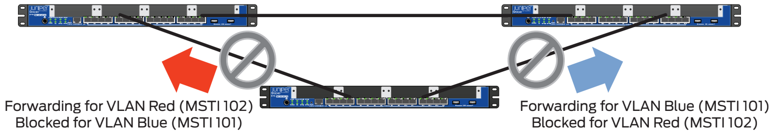
Forwarding for VLAN Red (MSTI 102)
Blocked for VLAN Blue (MSTI 101)
Forwarding for VLAN Blue (MSTI 101)
Blocked for VLAN Red (MSTI 102)

Figure 8: Multiple Spanning Tree Protocol

The MSTP region can support up to 64 MSTIs, and each instance can support anywhere from 1 through 4094 VLANs.

```
protocols {
    mstp {
configuration-name <region name);
                    bridge-priority <bridge priority>;
                    interface <interface name> {
                    cost <interface cost>;
                    }
                    interface <interface name> {
                    edge;
                    }
       msti <msti id> {
             bridge-priority <bridge priority>;
             vlan <vlan id or vlan name list>;
             interface <interface name> {
                 cost <interface cost>;
             }
       }
    }
}
```

MSTP configurations are a bit different compared to STP and RSTP. Similar to OSPF areas, an MSTP network is split into a number of regions. Inside each region different spanning-tree instances (known as MSTIs) run for different groups of VLANs. There is another spanning-tree instance running at a global level—that is, between regions (known as CIST). Junos OS provides configuration options to control CIST and MSTI parameters. As you can see in the previous configuration example, the MSTP region name and CIST bridge parameters are configured under `[protocol mstp]`. MSTI parameters and VLAN list association with MSTI are configured under `[protocols mstp msti <msti id>]`.

Another advantage of splitting a network into regions and running different MSTIs is that topology changes in one MSTI do not affect another MSTI or CIST and are local to that spanning-tree instance only.

Configuring IEEE 802.1x Authentication

IEEE 802.1x, which provides an authentication and authorization mechanism in wireless networks, is gaining popularity in wired networks. It provides network edge security, protecting Ethernet LANs from unauthorized access. An 802.1x-enabled switch (known as an authenticator) blocks all traffic from users (known as supplicants) connected to the switch until user credentials are verified in an authentication server (RADIUS server).

The J Series and SRX Series support three 802.1x modes for supplicants:

- Single—Only the first user is authenticated and the remaining users are tailgated.
- Single secure—Only one user is allowed.
- Multiple—More than one user is allowed and all users need to get authenticated.

As stated earlier, Ethernet switching features including 802.1x are inherited from the EX Series product line. But not all EX Series 802.1x features are available in the J Series and branch SRX Series. These platforms support the following:

- **Dynamic VLAN Assignment**—After successful authentication, it enables the supplicant to be a member of a particular VLAN dynamically. Please note that the VLAN ID needs to be configured in a RADIUS server for the user.
- **Guest VLAN**—This provides limited access to a LAN for 802.1x unsupported supplicants (supplicants that do not understand 802.1x).
- **Server-reject VLAN**—When an 802.1x-compliant supplicant fails to authenticate (because of wrong credentials), then the supplicant is assigned to a configured server-reject VLAN.
- **RADIUS accounting**—Accounting information is sent to the RADIUS accounting server. The information is sent to the server whenever a user (supplicant) logs in or logs out. Accounting information includes the amount of traffic, login and logout time, etc.
- **MAC RADIUS or MAC Authentication**—802.1x unsupported supplicants can be authenticated via a MAC RADIUS feature. Please note that guest VLAN and MAC RADIUS features are mutually exclusive.
- **Support for VoIP**—IP telephones are supported. If the phone is 802.1x enabled, it is authenticated like any other supplicant. If the phone is not 802.1x enabled, but has another 802.1x compatible device connected to its data port, that device is authenticated and then VoIP traffic can flow to and from the phone (providing that the interface is configured in single mode and not in single-secure mode). After successful authentication, RADIUS server communicates VLAN ID to device so that all voice traffic is classified under this VLAN also called VoIP VLAN.
- **Server failure cases**—When the RADIUS server becomes unreachable, the J Series and SRX Series take actions such as the following:
    - Permit—Allow all authentication requests without authentication until the RADIUS server is reachable.
    - Deny—Until the RADIUS server becomes reachable, all authentication requests are blocked.
    - VLAN—Enable authentication requested users to be members of a VLAN.
    - Cache—Imitate the previous authentication result for an authentication requested user.
- **Static MAC bypass**—A list of MAC addresses can be configured on the J Series and branch SRX Series for which 802.1 x authentications are bypassed.

Table 3: Supported 802.1x Features on J Series and Branch SRX Series Platforms

FEATURE	SRX100	SRX110	SRX210	SRX220	SRX240	SRX650	J SERIES
Dynamic VLAN assignment	✖	✖	✓	✓	✓	✓	✖
Authentication bypass	✓	✓	✓	✓	✓	✓	✓
Bypass with VLAN assignment	✖	✖	✓	✓	✓	✓	✖
Guest VLAN	✖	✖	✓	✓	✓	✓	✖
Server-reject VLAN	✖	✖	✓	✓	✓	✓	✖
Server failure fallback	✖	✖	✓	✓	✓	✓	✖
VoIP VLAN	✖	✖	✓	✓	✓	✓	✖
RADIUS accounting	✓	✓	✓	✓	✓	✓	✖
MAC RADIUS or MAC authentication	✓	✓	✓	✓	✓	✓	✖

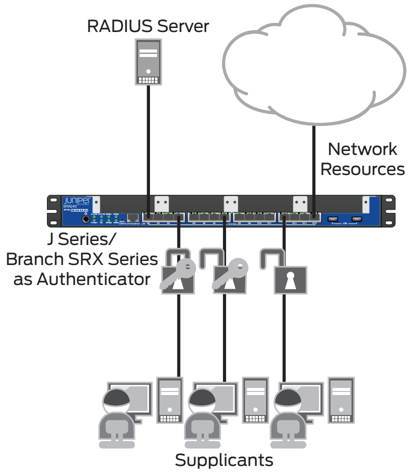

Figure 9: IEEE 802.1x authentication

```
protocols {
    dot1x {
        authenticator {
                          authentication-profile-name abc;
            static {
                <mac radius>/mask;
            }
            interface {
                <interface name> {
                    supplicant (single | single-secure| multiple);
                    guest-vlan <vlan name>;
                    server-reject-vlan <vlan name>;
                    server-fail (permit| deny| vlan-name <vlan name> |cache);
                }
            }
        }
    }
}
access {
    radius-server {
        <RADIUS server IP> secret <RADIUS share secret>
    }
    profile <profile name> {
        authentication-order radius;
        radius {
            authentication-server <RADIUS sever IP>;
        }
    }
}
```

802.1x is enabled on an interface under `[protocols dot1x authenticator]`. Although a supported supplicant type is configured under `[protocols dot1x authenticator interface <interface name> supplicant mode]`, it can be any of three modes—that is, single, single-secure, and multiple. Guest VLAN, server-reject VLAN, server fail conditions, and MAC authentication options are configured under `[protocols dot1x authenticator interface <interface name>]`. The authentication bypass list is configured under `[protocols dot1x authenticator static]`.

The RADIUS server configuration is a must for proper working of the 802.1x protocol. The RADIUS server needs to be defined under [edit access profile]. Also, it is mandatory that an access profile be created for the RADIUS server, and this access profile should be configured under [protocols dot1x authenticator authentication-profile-name].

Configuring IGMP Snooping

At Layer 2 all multicast traffic is treated as broadcast and is flooded to all ports of a switch of the same broadcast domain or VLAN domain. Due to this, a lot of bandwidth is wasted when only a few multicast receivers are connected to this switch. To overcome this limitation on J Series and branch SRX Series platforms, Junos OS provides a feature called IGMP snooping. Internet Group Management Protocol (IGMP) snooping regulates multicast traffic in a switched network. With IGMP snooping enabled, a LAN switch monitors the IGMP transmissions between a host (a network device) and a multicast router, keeping track of the multicast groups and associated member interfaces. The switch uses that information to make intelligent multicast-forwarding decisions and forward traffic to the intended destination interfaces.

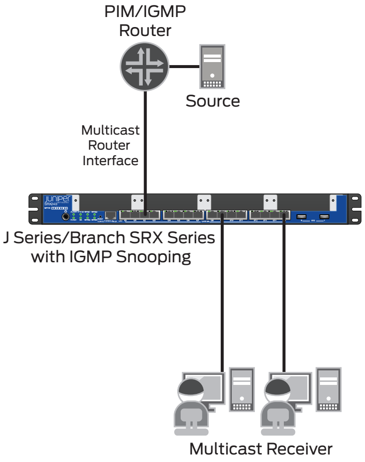

Figure 10: IGMP snooping

```
protocols {
    igmp-snooping {
        vlan vlan10;
    }
}
```

IGMP snooping is configured per VLAN under `[protocols]`. Once it is configured, the switch starts inspecting IGMP communication between multicast receivers (hosts) and IGMP or the PIM router. The interface where IGMP queries are received is identified as the multicast router interface. A binding between a multicast group and an interface is created when join/report messages are received on that interface. When actual multicast data traffic for a particular group is received on a router-connected interface, it is forwarded to only those interfaces for which binding is present for that multicast group. And it continues to forward traffic until it receives IGMP leave or time-out mechanisms in IGMPv1 hosts. All these operations are transparent to the IGMP/PIM router and multicast receiver. Junos OS also provides options for manual configuration of multicast router interfaces and static binding between multicast groups and interfaces. Please note this feature is not available in SRX100.

Configuring 802.1q Tunneling

Q-in-Q tunneling allows service providers on Ethernet access networks to extend a Layer 2 Ethernet connection between two customer sites. This feature is very useful when J Series and branch SRX Series devices are deployed in a service provider network as a provider edge (PE) device. A PE device sends and encapsulates incoming VLAN tagged packets from customers into a provider VLAN, and the receiving PE device de-encapsulates the provider VLAN and forwards packets to receiving customers. In this way the customer Layer 2 information (VLAN, priority) is intact when it is received at the other end.

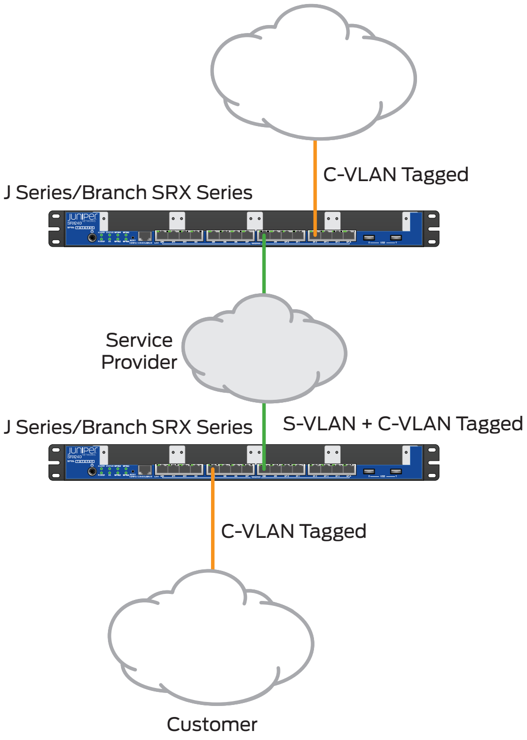

Figure 11: Q-in-Q tunneling

In Q-in-Q tunneling, as a packet travels from a customer VLAN (C-VLAN) to a service provider's VLAN, a customer-specific 802.1Q tag is added to the packet. This additional tag is used to segregate traffic into service-provider-defined service VLANs (S-VLANs). The original customer 802.1Q tag of the packet remains and is transmitted transparently, passing through the service provider's network. As the packet leaves the S-VLAN in the downstream direction, the extra 802.1Q tag is removed.

In a Q-in-Q deployment, customer packets from downstream interfaces are transported without any changes to source and destination MAC addresses. You can disable MAC address learning at both the interface level and the VLAN level. Disabling MAC address learning on an interface disables learning for all the VLANs of which that interface is a member.

```
vlans {
    <vlan name> {
        vlan-id <vlan id>;
        dot1q-tunneling {
            customer-vlans (native | <vlan id range>);
        }
        interface {
            <interface name> {
                mapping {
                    (native | <vlan id>) {
                        push;
                    }
                }
            }
        }
        no-mac-learning;
    }
}
ethernet-switching-options {
    interfaces {
        <interface name> {
            no-mac-learning;
        }
    }
}
```

When Q-in-Q tunneling is enabled on J Series and branch SRX Series platforms, it is assumed that trunk interfaces are to be part of the service provider network and access interfaces are to be customer facing. An access interface can receive both tagged and untagged frames in this case. There are three ways to map C-VLANs to an S-VLAN:

- All-in-one bundling—Use the dot1q-tunneling statement at the `[vlan <vlan name>]` hierarchy to map without specifying customer VLANs. All packets from a specific access interface are mapped to the S-VLAN.
- Many-to-one bundling—Use the customer-vlans statement at the `[vlan <vlan name>]` hierarchy to specify which C-VLANs are mapped to the S-VLAN.
- Mapping C-VLAN on a specific interface—Use the mapping statement at the `[vlan <vlan name>]` hierarchy to map a specific C-VLAN on a specified access interface to the S-VLAN.
- Please note that only the SRX650 supports all types—all-in-one, many-to-one, and C-VLAN mapping. The rest of the SRX Series platforms (except the SRX100) and J Series support only all-in-one bundling. To disable MAC learning on VLAN, configure `no-mac-learning` under `[vlan <vlan name>]`. And to disable at the interface level, add the same keyword under `[ethernet-switching-options interface <interface name>]`. Please note this feature is not available in SRX100.

Configuring Link Layer Discover Protocol (LLDP) and LLDP-MED

Discovery Protocol–Media Endpoint Discovery (LLDP-MED) to learn and distribute device information on network links. The information allows the switch to quickly identify a variety of devices, resulting in a LAN that interoperates smoothly and efficiently.

LLDP-capable devices transmit information in type, length, and value (TLV) messages to neighbor devices. Device information can include specifics, such as chassis and port identification and system name and system capabilities. The TLVs leverage this information from parameters that have already been configured in the Junos OS.

LLDP-MED goes one step further, exchanging IP-telephony messages between the switch and the IP telephone. These TLV messages provide detailed information on Power over Ethernet (PoE) policy. The PoE Management TLVs let the switch ports advertise the power level and power priority needed. For example, the switch can compare the power needed by an IP telephone running on a PoE interface with available resources. If the switch cannot meet the resources required by the IP telephone, the switch could negotiate with the telephone until a compromise on power is reached.

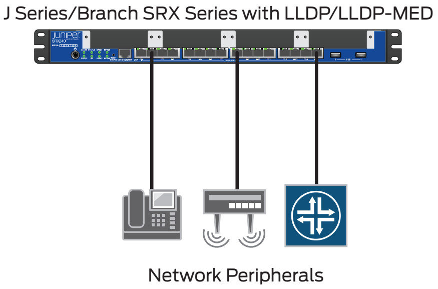

Figure 12: LLDP and LLDP-MED

The switch also uses these protocols to ensure that voice traffic gets tagged and prioritized with the correct values at the source itself. For example, 802.1p CoS and 802.1Q tag information can be sent to the IP telephone.

```
Protocols {
    lldp {
        interface <interface name>;
    }
    lldp-med {
        interface <interface name>;
    }
}
```

LLDP and LLDP-MED are enabled on an interface by configuring under [`protocols lldp`] and [`protocols lldp-med`], respectively.

The following basic LLDP TLVs are supported:

- Chassis identifier—This is the MAC address associated with the local system.
- Port identifier—This is the port identification for the specified port in the local system.
- Port description—This is the user-configured port description. The port description can be a maximum of 256 characters.
- System name—This is the user-configured name of the local system. The system name can be a maximum of 256 characters.
- System description—The system description contains information about the software and current image running on the system. This information is not configurable, but it is taken from the software.
- System capabilities—These pertain to the primary function performed by the system. Capabilities that the system supports are bridge or router. This information is not configurable, but it is based on the model of the product.
- Management address—This is the IP management address of the local system.

The following LLDP-MED TLVs are supported:

- LLDP-MED capabilities—A TLV advertises the primary function of the port. The capabilities values range from 0 through 15:
    - 0 – Capabilities
    - 1 – Network policy
    - 2 – Location identification
    - 3 – Extended power via medium-dependent interface power-sourcing equipment (MDI-PSE)
    - 4 – Inventory
    - 5 – 15 – Reserved
- LLDP-MED device class values:
    - 0 – Class not defined
    - 1 – Class 1 device
    - 2 – Class 2 device
    - 3 – Class 3 device
    - 4 – Network connectivity device
    - 5 – 255 - Reserved
- Network policy—A TLV advertises the port VLAN configuration and associated Layer 2 and Layer 3 attributes. Attributes include the policy identifier; application types, such as voice or streaming video; 802.1Q VLAN tagging; and 802.1p priority bits and DiffServ code points.
- Endpoint location—A TLV advertises the physical location of the endpoint.
- Extended power via MDI—A TLV advertises the power type, power source, power priority, and power value of the port. It is the responsibility of the PSE device (network connectivity device) to advertise the power priority on a port.

J Series and Branch SRX Series Ethernet Switching Configuration Examples

Simple Ethernet Switching

This example details the configuration needed to use a J Series device and a branch SRX Series device as simple Layer 2 switches. The topology is illustrated in Figure 13.

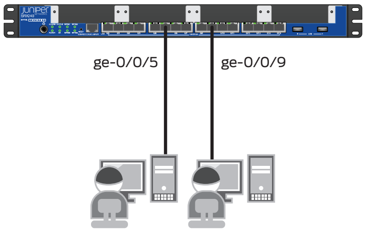

Figure 13: Simple Ethernet switching

The associated configurations are as follows:

```
set interfaces ge-0/0/5 unit 0 family ethernet-switching
set interfaces ge-0/0/9 unit 0 family ethernet-switching
```

Troubleshooting

Both interfaces, ge-0/0/5 and ge-0/0/9, should be part of the default VLAN.

```
regress@SRX-1> show vlans
Name           Tag     Interfaces
default        1
                       ge-0/0/5.0*, ge-0/0/9.0*
```

Adding VLANs

Now suppose that this small branch office has two departments—SALES and OPERATIONS. To isolate the departments and prevent traffic from leaking between domains, VLANS are added to the design—resulting in a new topology, as illustrated in Figure 14.

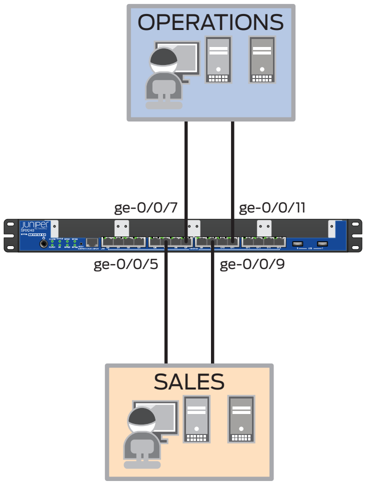

Figure 14: Ethernet switching with multiple VLANs

```
set vlans OPERATIONS vlan-id 20
set vlans SALES vlan-id 10
set interfaces ge-0/0/5 unit 0 family ethernet-switching vlan members SALES
set interfaces ge-0/0/7 unit 0 family ethernet-switching vlan members OPERATIONS
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members SALES
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members OPERATIONS
```

Troubleshooting

The following command shows interfaces and VLAN association:

```
regress@SRX-1> show vlans
Name           Tag     Interfaces
OPERATIONS     20
                       ge-0/0/7.0*, ge-0/0/11.0*
SALES          10
                       ge-0/0/5.0*, ge-0/0/9.0*
default        1
                       None
```

Routing Traffic Between VLANs

Now assume that this small branch needs to provide connectivity between the different business units, but that the connectivity must be controlled by assigning each business unit its own Layer 3 segment. Consequently, traffic between units is routed and inspected by the firewall module, where traffic policies can be enforced, as illustrated in Figure 15. The following configuration adds two Layer 3 interfaces, one for each VLAN, which serve as default gateways for the respective network segments. These new VLAN interfaces are then added to security zones, and security policies are defined to allow traffic between the zones. In this example, two security zones—SALES and OPERATIONS—are created, and HTTP traffic is allowed between them (bidirectional).

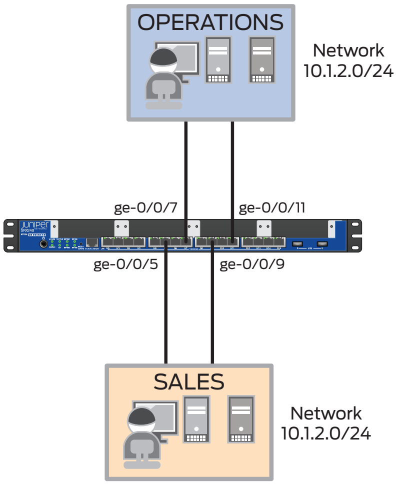

Figure 15: Inter-VLAN forwarding

```
set vlans OPERATIONS vlan-id 20
set vlans OPERATIONS l3-interface vlan.20
set vlans SALES vlan-id 10
set vlans SALES l3-interface vlan.10
set interfaces ge-0/0/5 unit 0 family  ethernet-switching vlan members SALES
set interfaces ge-0/0/7 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces ge-0/0/9 unit 0 family  ethernet-switching vlan members SALES
set interfaces ge-0/0/11 unit 0 family  ethernet-switching vlan members
OPERATIONS
set interfaces vlan unit 10 family inet address 10.1.1.1/24
set interfaces vlan unit 20 family inet address 10.1.2.1/24
set security zones security-zone SALES interfaces vlan.10
set security zones security-zone OPERATIONS interfaces vlan.20
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
source-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
```

```
destination-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
application junos-http
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP then
permit
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
source-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
destination-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
application junos-http
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP then
permit
```

Troubleshooting

Along with VLAN associations, routed VLAN interfaces should be linked up to forward traffic between VLANs.

```
regress@SRX-1> show vlans
Name           Tag     Interfaces
OPERATIONS     20
                       ge-0/0/7.0*, ge-0/0/11.0*
SALES          10
                       ge-0/0/5.0*, ge-0/0/9.0*
default        1
                       None
regress@SRX-1> show interfaces vlan terse
Interface               Admin Link Proto    Local                Remote
vlan                    up    up
vlan.10                 up    up   inet     10.1.1.1/24
vlan.20                 up    up   inet     10.1.2.1/24
```

Adding Tagged Interface

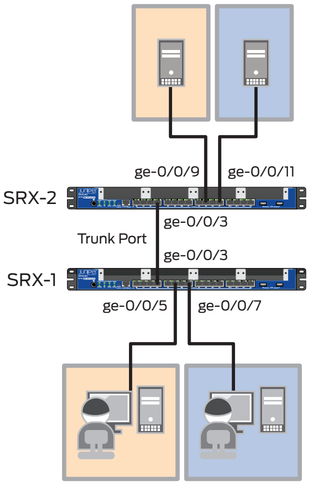

Figure 16: Trunk port (or adding tagged interface)

Now assume that the J Series and SRX Series are connected to another SRX Series device. SALES and OPERATIONS users belonging to one switch want to access their respective servers in another switch, keeping their VLAN domain separately as shown in Figure 16. As you can see, interfaces ge-0/0/3 on both devices are connected to each other and configured as a trunk port to carry SALES and OPERATIONS VLAN traffic.

SRX-1 Configurations

```
set vlans OPERATIONS vlan-id 20
set vlans OPERATIONS l3-interface vlan.20
set vlans SALES vlan-id 10
set vlans SALES l3-interface vlan.10
set interfaces ge-0/0/3 unit 0 family  ethernet-switching port-mode trunk
set interfaces ge-0/0/3 unit 0 family  ethernet-switching vlan members SALES
set interfaces ge-0/0/3 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces ge-0/0/5 unit 0 family  ethernet-switching vlan members SALES
set interfaces ge-0/0/7 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces vlan unit 10 family inet address 10.1.1.1/24
set interfaces vlan unit 20 family inet address 10.1.2.1/24
set security zones security-zone SALES interfaces vlan.10
set security zones security-zone OPERATIONS interfaces vlan.20
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
source-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
destination-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
application junos-http
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP then
permit
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
source-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
destination-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
application junos-http
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP then
permit
```

SRX-2 Configurations

```
set vlans OPERATIONS vlan-id 20
set vlans OPERATIONS l3-interface vlan.20
set vlans SALES vlan-id 10
set vlans SALES l3-interface vlan.10
set interfaces ge-0/0/3 unit 0 family  ethernet-switching port-mode trunk
set interfaces ge-0/0/3 unit 0 family  ethernet-switching vlan members SALES
set interfaces ge-0/0/3 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces ge-0/0/9 unit 0 family  ethernet-switching vlan members SALES
set interfaces ge-0/0/11 unit 0 family  ethernet-switching vlan members
OPERATIONS
set interfaces vlan unit 10 family inet address 10.1.1.1/24
set interfaces vlan unit 20 family inet address 10.1.2.1/24
set security zones security-zone SALES interfaces vlan.10
set security zones security-zone OPERATIONS interfaces vlan.20
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
source-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
destination-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
application junos-http
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP then
permit
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
source-address any
```

```
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
destination-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
application junos-http
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP then
permit
```

Troubleshooting

Access ports should be untagged members of VLANs, and trunk ports are tagged members of VLANs. A trunk port is part of a multiple VLAN.

```
regress@SRX-1> show  ethernet-switching interfaces
Interface     State  VLAN members         Tag   Tagging  Blocking
ge-0/0/3.0    up     OPERATIONS           20    tagged   unblocked
                     SALES                10    tagged   unblocked
ge-0/0/5.0    up     SALES                10    untagged unblocked
ge-0/0/7.0    up     OPERATIONS           20    untagged unblocked
regress@SRX-2> show  ethernet-switching interfaces
Interface     State  VLAN members         Tag   Tagging  Blocking
ge-0/0/3.0    up     OPERATIONS           20    tagged   unblocked
                     SALES                10    tagged   unblocked
ge-0/0/9.0    up     SALES                10    untagged unblocked
ge-0/0/11.0   up     OPERATIONS           20    untagged unblocked
```

Increasing Capacity with Link Aggregation

As the small branch office grows, with increasing numbers of applications requiring additional bandwidth, a bottleneck is created between the router and the switch. To alleviate this problem, link aggregation is configured, and a new link between the devices is added.

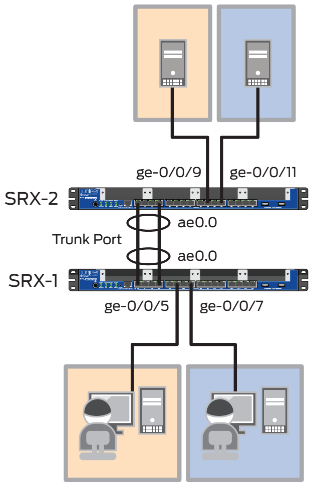

Figure 17: Link aggregation

Interfaces ge-0/0/1 and ge-0/0/3 are bundles to aggregated Ethernet interface ae0 on both switches. And this ae0.0 is configured as a trunk port to carry SALES and OPERATIONS VLAN traffic.

SRX-1 Configuration

```
set vlans OPERATIONS vlan-id 20
set vlans OPERATIONS l3-interface vlan.20
set vlans SALES vlan-id 10
set vlans SALES l3-interface vlan.10
set chassis aggregated-devices  thernet device-count 2
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/3 gigether-options 802.3ad ae0
set interfaces ge-0/0/5 unit 0 family  ethernet-switching vlan members SALES
set interfaces ge-0/0/7 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family  ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family  ethernet-switching vlan members SALES
set interfaces ae0 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces vlan unit 10 family inet address 10.1.1.1/24
set interfaces vlan unit 20 family inet address 10.1.2.1/24
set security zones security-zone SALES interfaces vlan.10
set security zones security-zone OPERATIONS interfaces vlan.20
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
source-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
destination-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
application junos-http
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP then
permit
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
source-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
destination-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
application junos-http
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP then
permit
```

SRX-2 Configuration

```
set vlans OPERATIONS vlan-id 20
set vlans OPERATIONS l3-interface vlan.20
set vlans SALES vlan-id 10
set vlans SALES l3-interface vlan.10
set chassis aggregated-devices  ethernet device-count 2
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/3 gigether-options 802.3ad ae0
set interfaces ge-0/0/9 unit 0 family  ethernet-switching vlan members SALES
set interfaces ge-0/0/11 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces ae0 aggregated-ether-options lacp active
set interfaces ae0 unit 0 family  ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family  ethernet-switching vlan members SALES
set interfaces ae0 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces vlan unit 10 family inet address 10.1.1.1/24
set interfaces vlan unit 20 family inet address 10.1.2.1/24
set security zones security-zone SALES interfaces vlan.10
set security zones security-zone OPERATIONS interfaces vlan.20
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
source-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
destination-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
application junos-http
```

```
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP then
permit
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
source-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
destination-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
application junos-http
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP then
permit
```

Troubleshooting

The multiplexer state of member interfaces of LAG should be collecting and distributing. L2 switching functionality is configured on an aggregated interface (in this example ae0 is made the trunk port).

```
regress@SRX-1> show lacp interfaces
Aggregated interface: ae0
    LACP state:       Role   Exp   Def  Dist  Col  Syn  Aggr  Timeout  Activity
      ge-0/0/5       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/5     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/7       Actor    No    No   Yes  Yes  Yes   Yes     Fast    Active
      ge-0/0/7     Partner    No    No   Yes  Yes  Yes   Yes     Fast    Active
    LACP protocol:        Receive State  Transmit State          Mux State
      ge-0/0/5                  Current   Fast periodic Collecting distributing
      ge-0/0/7                  Current   Fast periodic Collecting distributing

regress@SRX-1> show  ethernet-switching interfaces
Interface     State  VLAN members        Tag   Tagging  Blocking
ae0.0         up     OPERATIONS          20    tagged   unblocked
                     SALES               10    tagged   unblocked
ge-0/0/5.0    up     SALES               10    untagged unblocked
ge-0/0/7.0    up     OPERATIONS          20    untagged unblocked
```

Loop Avoidance with RSTP

Another J Series and SRX Series device, SRX-3, is connected to both SRX-1 and SRX-2 as shown in Figure 18. To avoid loops in the network, RSTP is configured.

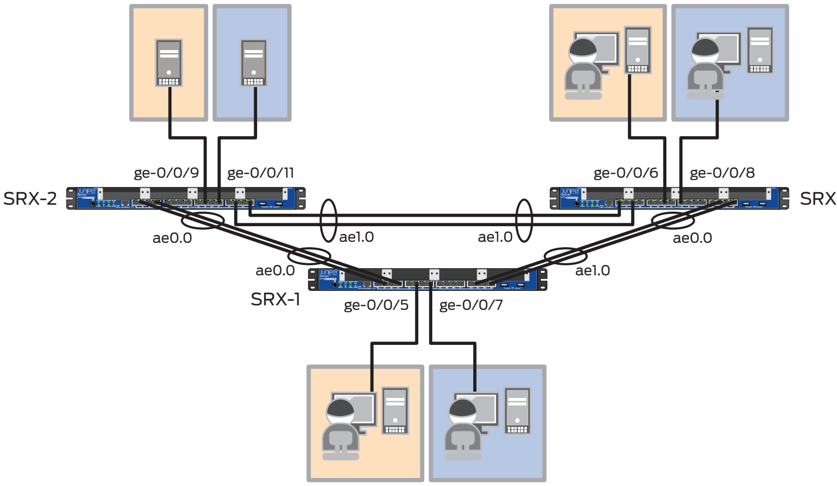

Figure 18: Loop avoidance with RSTP

Rapid Spanning Tree Protocol is enabled on all devices and SRX-2 is made as the root switch. Interfaces connected to end hosts, such as user workstations or servers, are configured as edge ports.

SRX-1 Configurations

```
set vlans OPERATIONS vlan-id 20
set vlans OPERATIONS l3-interface vlan.20
set vlans SALES vlan-id 10
set vlans SALES l3-interface vlan.10
set chassis aggregated-devices  ethernet device-count 2
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/3 gigether-options 802.3ad ae0
set interfaces ge-0/0/15 gigether-options 802.3ad ae1
set interfaces ge-0/0/13 gigether-options 802.3ad ae1
set interfaces ge-0/0/5 unit 0 family  ethernet-switching vlan members SALES
set interfaces ge-0/0/7 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces ae0 unit 0 family  ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family  ethernet-switching vlan members SALES
set interfaces ae0 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces ae1 unit 0 family  ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family  ethernet-switching vlan members SALES
set interfaces ae1 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces vlan unit 10 family inet address 10.1.1.1/24
set interfaces vlan unit 20 family inet address 10.1.2.1/24
set protocols rstp
set protocols rstp interface ge-0/0/5.0 edge
set protocols rstp interface ge-0/0/7.0 edge
set security zones security-zone SALES interfaces vlan.10
set security zones security-zone OPERATIONS interfaces vlan.20
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
source-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
destination-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
application junos-http
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP then
permit
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
source-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
destination-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
application junos-http
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP then
permit
```

SRX-2 Configurations

```
set vlans OPERATIONS vlan-id 20
set vlans OPERATIONS l3-interface vlan.20
set vlans SALES vlan-id 10
set vlans SALES l3-interface vlan.10
set chassis aggregated-devices  ethernet device-count 2
set interfaces ge-0/0/1 gigether-options 802.3ad ae0
set interfaces ge-0/0/3 gigether-options 802.3ad ae0
set interfaces ge-0/0/15 gigether-options 802.3ad ae1
set interfaces ge-0/0/13 gigether-options 802.3ad ae1
set interfaces ge-0/0/9 unit 0 family  ethernet-switching vlan members SALES
set interfaces ge-0/0/11 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces ae0 unit 0 family  ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family  ethernet-switching vlan members SALES
set interfaces ae0 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces ae1 unit 0 family  ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family  ethernet-switching vlan members SALES
set interfaces ae1 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces vlan unit 10 family inet address 10.1.1.2/24
set interfaces vlan unit 20 family inet address 10.1.2.2/24
set protocols rstp bridge-priority 4k
set protocols rstp interface ge-0/0/9.0 edge
set protocols rstp interface ge-0/0/11.0 edge
set security zones security-zone SALES interfaces vlan.10
set security zones security-zone OPERATIONS interfaces vlan.20
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
source-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
destination-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
application junos-http
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP then
permit
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
source-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
destination-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
application junos-http
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP then
permit
```

SRX-3 Configurations

```
set vlans OPERATIONS vlan-id 20
set vlans OPERATIONS l3-interface vlan.20
set vlans SALES vlan-id 10
set vlans SALES l3-interface vlan.10
set chassis aggregated-devices  ethernet device-count 2
set interfaces ge-0/0/13 gigether-options 802.3ad ae0
set interfaces ge-0/0/15 gigether-options 802.3ad ae0
set interfaces ge-0/0/0 gigether-options 802.3ad ae1
set interfaces ge-0/0/2 gigether-options 802.3ad ae1
set interfaces ge-0/0/6 unit 0 family  ethernet-switching vlan members SALES
set interfaces ge-0/0/8 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces ae0 unit 0 family  ethernet-switching port-mode trunk
set interfaces ae0 unit 0 family  ethernet-switching vlan members SALES
set interfaces ae0 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces ae1 unit 0 family  ethernet-switching port-mode trunk
set interfaces ae1 unit 0 family  ethernet-switching vlan members SALES
set interfaces ae1 unit 0 family  ethernet-switching vlan members OPERATIONS
set interfaces vlan unit 10 family inet address 10.1.1.3/24
set interfaces vlan unit 20 family inet address 10.1.2.3/24
set protocols rstp
set protocols rstp interface ge-0/0/6.0 edge
set protocols rstp interface ge-0/0/8.0 edge
set security zones security-zone SALES interfaces vlan.10
set security zones security-zone OPERATIONS interfaces vlan.20
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
source-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
destination-address any
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP match
application junos-http
set security policies from-zone SALES to-zone OPERATIONS policy Allow_HTTP then
permit
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
source-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
destination-address any
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP match
application junos-http
set security policies from-zone OPERATIONS to-zone SALES policy Allow_HTTP then
permit
```

Troubleshooting

`SRX-2` is the root switch. All interfaces on the root switch are in forwarding state.

```
regress@SRX-2> show spanning-tree bridge
STP bridge parameters
Context ID                          : 0
Enabled protocol                    : RSTP
  Root ID                           : 4096.00:22:83:99:b0:50
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Message age                       : 0
  Number of topology changes        : 2
  Time since last topology change   : 458 seconds
  Topology change initiator         : ae1.0
  Topology change last recvd. From  : 80:71:1f:a4:2b:01
  Local parameters
    Bridge ID                       : 4096.00:22:83:99:b0:50
    Extended system ID              : 0
    Internal instance ID            : 0

regress@elanta> show spanning-tree interface
Spanning tree interface parameters for instance 0
Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae0.0            128:1        128:1   4096.00228399b050     20000  FWD    DESG
ae1.0            128:2        128:2   4096.00228399b050     10000  FWD    DESG
ge-0/0/9.0     128:522      128:522   4096.00228399b050     20000  FWD    DESG
ge-0/0/11.0    128:524      128:524   4096.00228399b050     20000  FWD    DESG
```

Note that the root bridge ID is populated on all non-root switches. Also note that the root port is connected to the root switch.

```
regress@SRX-1> show spanning-tree bridge
STP bridge parameters
Context ID                          : 0
Enabled protocol                    : RSTP
  Root ID                           : 4096.00:22:83:99:b0:50
  Root cost                         : 10000
  Root port                         : ae0.0
  Hello time                        : 2 seconds
  Maximum age                       : 20 seconds
  Forward delay                     : 15 seconds
  Message age                       : 1
  Number of topology changes        : 4
  Time since last topology change   : 95 seconds
  Topology change initiator         : ae1.0
  Topology change last recvd. From  : 00:22:83:99:b0:c0
  Local parameters
    Bridge ID                       : 32768.00:1b:c0:53:69:88
    Extended system ID              : 0
    Internal instance ID            : 0
```

When there are two redundant links to the root switch, one of them is the root port and another is the alternate port.

```
regress@SRX-3> show spanning-tree interface

Spanning tree interface parameters for instance 0
Interface    Port ID    Designated      Designated         Port    State  Role
                         port ID        bridge ID          Cost
ae0.0            128:1        128:2  32768.001bc0536988     10000  BLK    ALT
ae1.0            128:2        128:2   4096.00228399b050     10000  FWD    ROOT
ge-0/0/6.0     128:519      128:519  32768.80711fa42a90     20000  FWD    DESG
ge-0/0/8.0     128:521      128:521  32768.80711fa42a90     20000  FWD    DESG
```

IEEE 802.1x Authentication

In this example, 802.1x is enabled on interface ge-0/0/5. Unless the credentials of the user connected to the interface are verified, the user is unable to access any of the network resources connected to this device.

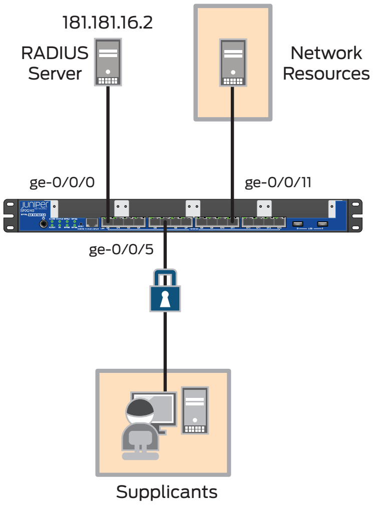

Figure 19: IEEE 802.1x authentication

The RADIUS server must be reachable from the switch. And it must be configured with the supplicant's username and password. No authentication is performed on the configured static MAC address under [`edit protocols dot1x static`].

```
set interfaces ge-0/0/0 unit 0 family inet address 181.181.16.1/24
set interfaces ge-0/0/11 unit 0 family ethernet-switching vlan members SALES
set interfaces ge-0/0/12 unit 0 family ethernet-switching vlan members SALES
set protocols dot1x authenticator authentication-profile-name test
set protocols dot1x authenticator static 00:11:22:33:55:66/48
set protocols dot1x authenticator interface ge-0/0/12.0 supplicant multiple
set access radius-server 181.181.16.2 secret “$9$K76WX-YgJHqfVwqfTzCAvWL”
set access profile test authentication-order radius
```

Troubleshooting

```
regress@SRX-1# run show dot1x interface
802.1X Information:
Interface     Role           State          MAC address        User
ge-0/0/12.0   Authenticator  Connecting
regress@SRX-1# run show dot1x interface
802.1X Information:
Interface     Role           State          MAC address        User
ge-0/0/12.0   Authenticator  Authenticated  00:00:00:80:00:01  user1
regress@SRX-1# run show dot1x authentication-bypassed-users
MAC address        Interface        VLAN
00:11:22:33:55:66  ge-0/0/12.0      configured/default
```

Multicast Snooping with IGMP Snooping Protocol

This example configures IGMP snooping on J Series and SRX Series devices to regulate multicast traffic on a device. A multicast receiver is connected to interface ge-0/0/9, and interface ge-0/0/2 is connected to the PIM/IGMP router from where multicast data packets are sent.

```
set vlans SALES vlan-id 10
set interfaces ge-0/0/2 unit 0 family ethernet-switching vlan members SALES
set interfaces ge-0/0/9 unit 0 family ethernet-switching vlan members SALES
set protocols igmp-snooping vlan SALES
```

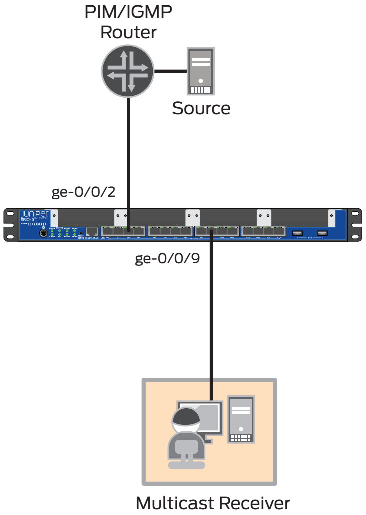

Figure 20: Multicast snooping with IGMP snooping

Troubleshooting

Make sure that the uplink interface (ge-0/0/2) is identified as a multicast router interface. Otherwise, the received join message cannot be forwarded to the PIM/IGMP router.

```
regress@SRX-1# run show igmp-snooping membership detail
VLAN: SALES Tag: 10 (Index: 2)
    Router interfaces:
        ge-0/0/2.0 dynamic Uptime: 00:04:48 timeout: 219
   Group: 230.5.5.5
     ge-0/0/9.0 timeout: 233 Last reporter: 23.23.23.2 Receiver count: 1, Flags:
<V2-hosts>
```

802.1q Tunneling (Q-in-Q Tunneling)

This example shows that the 802.1q tunneling feature in the J Series and branch SRX Series devices can be used as a provider edge (PE) feature in service provider networks.

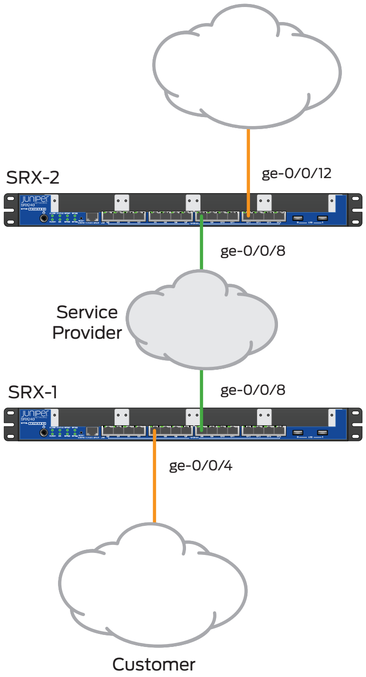

Figure 21: 802.1q tunneling

Interfaces ge-0/0/4 and ge-0/0/12 of SRX-1 and SRX-2 are connected to end customer devices, respectively. And ge-0/0/8 on both devices is connected to a service provider network.

SRX-1 Configurations

```
set vlans SERVICE_PROVIDER vlan-id 100
set vlans SERVICE_PROVIDER dot1q-tunneling
set interfaces ge-0/0/4 unit 0 family ethernet-switching vlan members SERVICE_
PROVIDER
set interfaces ge-0/0/8 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/8 unit 0 family ethernet-switching vlan members SERVICE_
PROVIDER
```

SRX-2 Configurations

```
set vlans SERVICE_PROVIDER vlan-id 100
set vlans SERVICE_PROVIDER dot1q-tunneling
set interfaces ge-0/0/8 unit 0 family ethernet-switching port-mode trunk
set interfaces ge-0/0/8 unit 0 family ethernet-switching vlan members SERVICE_
PROVIDER
set interfaces ge-0/0/12 unit 0 family ethernet-switching vlan members SERVICE_
PROVIDER
```

Troubleshooting

```
regress@SRX-1# run show vlans detail
VLAN: SERVICE_PROVIDER, 802.1Q Tag: 100, Admin State: Enabled
Dot1q Tunneling status: Enabled
Number of interfaces: 2 (Active = 2)
  Untagged interfaces: ge-0/0/4.0*
  Tagged interfaces: ge-0/0/8.0*
VLAN: default, 802.1Q Tag: 1, Admin State: Enabled
```

About Juniper Networks

Juniper Networks is in the business of network innovation. From devices to data centers, from consumers to cloud providers, Juniper Networks delivers the software, silicon and systems that transform the experience and economics of networking. The company serves customers and partners worldwide. Additional information can be found at **www.juniper.net**.

Corporate and Sales Headquarters
Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
Phone: 888.JUNIPER (888.586.4737)
or 408.745.2000
Fax: 408.745.2100
www.juniper.net

APAC Headquarters
Juniper Networks (Hong Kong)
26/F, Cityplaza One
1111 King's Road
Taikoo Shing, Hong Kong
Phone: 852.2332.3636
Fax: 852.2574.7803

EMEA Headquarters
Juniper Networks Ireland
Airside Business Park
Swords, County Dublin, Ireland
Phone: 35.31.8903.600
EMEA Sales: 00800.4586.4737
Fax: 35.31.8903.601

To purchase Juniper Networks solutions, please contact your Juniper Networks representative at 1-866-298-6428 or authorized reseller.

3500196-002-EN Dec 2011

Printed on recycled paper