# IMPLEMENTING A LAYER 2 ENTERPRISE INFRASTRUCTURE WITH MSTP

## Table of Contents



## Table of Figures

## Introduction

The focus of this document is on implementing a Layer 2 design for enterprise networks using Multiple Spanning Tree Protocol (MSTP). Both Juniper Networks® EX Series Ethernet Switches and MX Series 3D Universal Edge Routers run the Juniper Networks Junos® operating system and support a similar set of Layer 2 features. By supporting Layer 2, Layer 3, or a combination of both, these devices provide flexibility in designing and deploying networks. For example, Layer 2 can be implemented initially, and Layer 3 functionality can be enabled later on the same devices at no additional costs.

This document provides an overview of the design considerations and protocols used, highlighting some of the differences between EX Series and MX Series supported features. It concludes with an implementation example, configuration guidelines, and verification and troubleshooting procedures.

## Scope

This document provides guidelines and an implementation example for Layer 2 enterprise environments using EX Series Ethernet Switches in the access layer and MX Series 3D Universal Edge Routers for core and aggregation. After a brief features and protocols overview, we present the topology implemented and provide configuration examples, verification, and troubleshooting procedures. The design presented in this document uses MSTP to prevent loops. Other design options using RSTP, RTG, and Virtual Chassis are presented in separate implementation guides.

This document is intended for network design and operation engineers or other technical audiences who are supporting enterprise customers with Layer 2 deployments using the EX Series and MX Series.

## Design Considerations

The network architecture presented in this document is based on a collapsed campus model where EX Series switches are used in the access layer while the core and aggregation layers are combined using MX Series routers. The access switches are dual-homed to two aggregation/core switches. This is needed in order to provide physical redundancy while allowing for load balancing. Device redundancy at Layer 2 introduces the potential of broadcast storms with packets traveling endlessly and crippling the network. A mechanism is therefore needed to prevent Layer 2 loops. The EX Series switches support three standard versions: 802.1D Spanning Tree Protocol (STP), 802.1w Rapid Spanning Tree Protocol (RSTP), and 802.1s MSTP. The MX Series routers add support for Virtual Spanning Tree Protocol (VSTP) which is compatible with Per-VLAN Spanning Tree Plus (PVST+) and Rapid-PVST+ protocols supported on Cisco Systems routers and switches. The version of STP that is implemented in this solution is MSTP.

Layer 3 will be enabled only at the aggregation/core layer. MX Series routers are configured with integrated routing and bridging (IRB) and Virtual Router Redundancy Protocol (VRRP). They will also act as DHCP relay agents to allow clients in the access layer to obtain IP addresses dynamically.

### VLANs and 802.1Q

A VLAN is a logical grouping of end devices allowing communication as if they were on the same LAN.

On the EX Series, ports that are assigned to a VLAN can be configured as either access or trunk ports. A port in access mode connects to a network device such as a desktop computer, an IP telephone, a printer, a file server, or a security camera. The interface itself belongs to a single VLAN. Trunk interfaces are generally used to interconnect switches to one another. The frames on a trunk port are tagged as defined in 802.1Q standard.Juniper Networks EX3200 Ethernet Switch and EX4200 Ethernet Switch both support a maximum of 4096 VLANs. VLANs 0 and 4095 are reserved by Junos OS.

### Multiple Spanning Tree Protocol

Redundancy is an important design consideration for high availability and resiliency in an enterprise network. Although RSTP provides faster convergence than STP, it still does not make good use of all available paths within a redundant Layer 2 network. With RSTP, all traffic from all VLANs follows the same path as determined by the spanning tree; therefore, redundant paths are not utilized. MSTP overcomes this limitation and allows load sharing through the use of multiple spanning-tree instances (MSTI).

MSTP allows for building multiple spanning trees over trunks by grouping and associating a set of VLANs to each spanning-tree instance. MSTP also provides the capability to logically divide a Layer 2 network into regions. Every region has a unique identifier and can contain multiple instances of spanning trees. All regions are bound together using a Common Instance Spanning Tree (CIST), which is responsible for creating a loop-free topology across regions, whereas the MSTI controls topology inside regions. MSTP uses RSTP as a converging algorithm and is interoperable with earlier versions of STP.

To be part of a common MSTP region, a group of switches must share the same configuration attributes which consist of:

- Configuration name
- Revision level
- VLANs mapping to MSTI instance

If one of these attributes differs between two switches, they are considered part of different regions. In order for regions to communicate, a Common Spanning Tree (CST) instance runs across all regions. The CST also forwards traffic for VLANs which are not covered by any MSTI. Up to 64 MSTIs are supported in each region on the MX Series and EX Series.

### Integrated Routing and Bridging

Integrated routing and bridging (IRB) interfaces on the MX Series support both Layer 2 bridging and Layer 3 routing on the same interface. Frames are bridged if they are not sent to the router's media access control (MAC) address. Frames sent to the router's MAC address are routed to other interfaces configured for Layer 3 routing.

The EX Series Ethernet Switches also support routed interfaces called Routed VLAN Interfaces (RVIs). These are not implemented in this guide. As opposed to IRBs which route bridge domains, RVIs route VLANs. A port of a switch VLAN is identified by an interface and a VLAN-id which is globally significant across the switch.

### Virtual Router Redundancy Protocol

EX Series switches and MX Series routers support Virtual Router Redundancy Protocol (VRRP). With VRRP, routers viewed as a redundancy group share the responsibility for forwarding packets as if they owned the IP address corresponding to the default gateway configured on the hosts. At any time, one of the VRRP routers acts as the master, while other VRRP routers act as backup routers. If the master router fails, a backup router becomes the new master. Using this approach, router redundancy is always provided, allowing traffic on the LAN to be routed without relying on a single router.

### DHCP Relay

DHCP requests sent from a client to a server are normally restricted to the same physical segment, LAN, or VLAN on which the client resides. In the event that the server and client are on different LANs or VLANs, a relay agent is needed. The main advantage of this feature is that a single DHCP server can serve clients on remote LANs or VLANs, eliminating the need for a dedicated DHCP server in each LAN or VLAN environment. Both EX Series and MX Series devices can be configured to relay requests to a DHCP/BOOTP server and use the DHCP Relay Agent option (option 82) in the relayed messages. Since Layer 3 is not implemented on the access switches in this guide, the MX Series routers will act as the DHCP relay agents.

## Implementation

### Configuration Guidelines

#### Interface and VLAN Configuration

Switch ports can be configured with either access mode or trunk mode. Access ports typically belong to a single VLAN and transmit and receive untagged Ethernet frames. A trunk port typically connects to another switch or to a customer's edge router. Interfaces configured for trunk mode handle traffic for multiple VLANs, multiplexing the traffic for all configured VLANs over the same physical connection, and separating the traffic by tagging it with the appropriate VLAN-id.

Below are sample interface and VLAN configurations for both the EX Series and MX Series:

EX-VC-1:

```
vlans {
    ENG {
        vlan-id 200;
    }
    HR {
        vlan-id 100;
    }
    SALES {
        vlan-id 300;
    }
    SERVER {
        vlan-id 600;
    }
}

interfaces {
    /* Access port examples */
    ge-0/0/0 {
            unit 0 {
                family ethernet-switching {
                    port-mode access;
                    vlan {
                        members HR;
                    }
                }
            }
        }
    ge-0/0/1 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members ENG;
                }
            }
        }
    }
   /*Trunk port example */
   ge-0/1/0 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [ HR ENG SERVER ];
                }
            }
        }
    }
}
```

The example above shows a port-based assignment of VLANs on the EX Series switch where VLANs are added under the interface stanza. It is also possible to use a VLAN-based assignment where interfaces are added under the VLAN stanza (as shown below). Some users may be more familiar with the first form of configuration. Others may prefer the second method in that functionalities are grouped by feature rather than per interface. Both methods result in the same configuration from a software perspective.

**EX Series:**

```
vlans {
    HR {
        vlan-id 100;
        interface {
            ge-0/0/0.0;
            ge-0/0/1.0;
        }
    }
}
```

### MSTP Configuration

The root bridge selection can be influenced by changing the bridge priorities from their default value of 32K for each MSTI instance. Two MSTP instances are configured in the example below: MSTI 1 for VLANs 100, 200, and 600, and MSTI 2 for VLAN 300. MX-A is configured to act as the root bridge for MSTI 1 by setting the bridge priority to 4K. In the same fashion, MX-B will be configured to act as the root bridge for MSTI 2.

**EX-VC-1**

```
protocols {
    mstp {
        configuration-name REGION1;
        revision-level 1;
        interface ge-0/0/0.0 {
            edge;
        }
        interface ge-0/0/1.0 {
            edge;
        }
        interface ge-0/0/22.0 {
            edge;
        }
        interface ge-0/1/0.0 {
            mode point-to-point;
        }
        interface ge-0/1/1.0 {
            mode point-to-point;
        }
        msti 1 {
            vlan [ HR ENG SERVER ];
        }
        msti 2 {
            vlan SALES;
        }
    }
}
```

**MX-A:**

```
protocols {
    mstp {
        configuration-name REGION1;
        revision-level 1;
        bridge-priority 4k;
        interface ge-8/0/0 {
            mode point-to-point;
        }
        interface ge-8/0/1 {
            mode point-to-point;
        }
        interface ge-8/0/2 {
            mode point-to-point;
        }
        interface ge-8/0/3 {
            mode point-to-point;
        }
        interface xe-2/0/0 {
            mode point-to-point;
        }
        msti 1 {
            bridge-priority 4k;
            vlan [ 100 200 600 ];
        }
        msti 2 {
            bridge-priority 8k;
            vlan 300;
        }
    }
}
```

To be part of a common MSTP region, a group of switches must share the same configuration attributes which consist of:

- Configuration name
- Revision level
- VLANs mapping to MSTI instance

As a result, in the example above, even though EX-VC-1 may not have access ports in VLAN 300, MSTI 2 is still configured on this switch. The purpose of this is to make the VLAN-to-MSTI mapping consistent between all switches so that they can be part of the same MSTP region.

Configuring the revision level for MSTP is not necessary as it defaults to 0. In the example above, the revision level is explicitly set to a value of 1 in order to highlight the need to have a consistent revision number across the MSTP region.

### IRBs and VRRP Configuration

An IRB is configured on the MX Series routers in two steps:

1. Configuring the IRB interface using the irb statement.

2. Referencing the IRB interface at the bridge domain level of the configuration.

VRRP can be configured on the IRB interface so that redundant links can be used to carry traffic between the bridge domain and the router network.

The example below shows IRBs configured with VRRP groups and virtual addresses. The priority is set to 254 for groups 1, 2, and 4 on MX-A. This makes MX-A the master for these groups while MX-B will be left with the default priority of 100 and will act as the backup. The "accept-data" command allows an IRB interface to accept packets destined for a virtual IP address.

**MX-A:**

```
interfaces {
      irb {
        unit 100 {
            family inet {
                address 10.10.10.4/24 {
                    vrrp-group 1 {
                        virtual-address
10.10.10.1;
                        priority 254;
                        accept-data;
                    }
                }
            }
        }
        unit 200 {
            family inet {
                address 20.20.20.4/24 {
                    vrrp-group 2 {
                        virtual-address
20.20.20.1;
                        priority 254;
                        accept-data;
                    }
                }
            }
        }
}
```

```
        unit 300 {
           family inet {
                address 30.30.30.4/24 {
                    vrrp-group 3 {
                        virtual-address
30.30.30.1;
                        accept-data;
                    }
                }
            }
        }
        unit 600 {
            family inet {
                address 60.60.60.4/24 {
                    vrrp-group 4 {
                        virtual-address
60.60.60.1;
                        priority 254;
                        accept-data;
                    }
                }
            }
        }
    }
}
```

 

### Bridge Domains Configuration

Bridge domains limit the scope of MAC learning (and thereby the size of the MAC table). They also determine where the device should propagate frames sent to broadcast, unknown unicast, and multicast MAC addresses. Each interface belonging to a bridge domain needs to be referenced under the corresponding bridge domain stanza. IRB interfaces are also referenced using the "routing-interface" statement as shown below:

### DHCP Relay Configuration

The DHCP server is configured under the "forwarding option" stanza as illustrated in the example below. Here, the MX Series routers are configured to forward BOOTP/DHCP requests coming from IRBs 100 and 200 to the DHCP server address. This is done using the relay-agent-option (option 82) in the relayed messages.

**MX-A:**

```
forwarding-options {
    helpers {
        bootp {
            server 60.60.60.2;
            relay-agent-option;
            interface {
                irb.100;
                irb.200;
            }
        }
    }
}
```

### DHCP Relay Configuration

The DHCP server is configured under the "forwarding option" stanza as illustrated in the example below. Here, the MX Series routers are configured to forward BOOTP/DHCP requests coming from IRBs 100 and 200 to the DHCP server address. This is done using the relay-agent-option (option 82) in the relayed messages.

**MX-A:**

```
forwarding-options {
    helpers {
        bootp {
            server 60.60.60.2;
            relay-agent-option;
            interface {
                irb.100;
                irb.200;
            }
        }
    }
}
```

## Verification

Below are some of the commands that can be used to verify the MSTP setup:

EX Series:

- show ethernet-switching interfaces
- show ethernet-switching table

MX Series:

- show bridge mac-table
- show vrrp summary
- show vrrp extensive
- show interface irb terse

Both:

- show spanning-tree interface
- show spanning-tree bridge
- show spanning-tree statistics interface
- show spanning-tree mstp configuration

## Troubleshooting

The following commands can be used for troubleshooting with MSTP:

EX Series:

- clear ethernet-switching table
- restart ethernet-switching

MX Series:

- show bootp statistics

Both:

- monitor traffic interface <name> layer2-headers
- monitor traffic interface <name> size <size> detail
- set protocols mstp traceoptions file <filename>
- set protocols mstp traceoptions flag all
- show system core-dumps

## Implementation Example

### Network Topology

All switches in our network are configured as part of the same region. Two MSTI instances are defined, one regrouping VLANs 100, 200, and 600 and the second one for VLAN 300. MX-A is configured as the root bridge for MSTI 1 while MX-B is the root bridge for MSTI 2. The following diagram shows the resulting logical topology including port states and roles for each path:

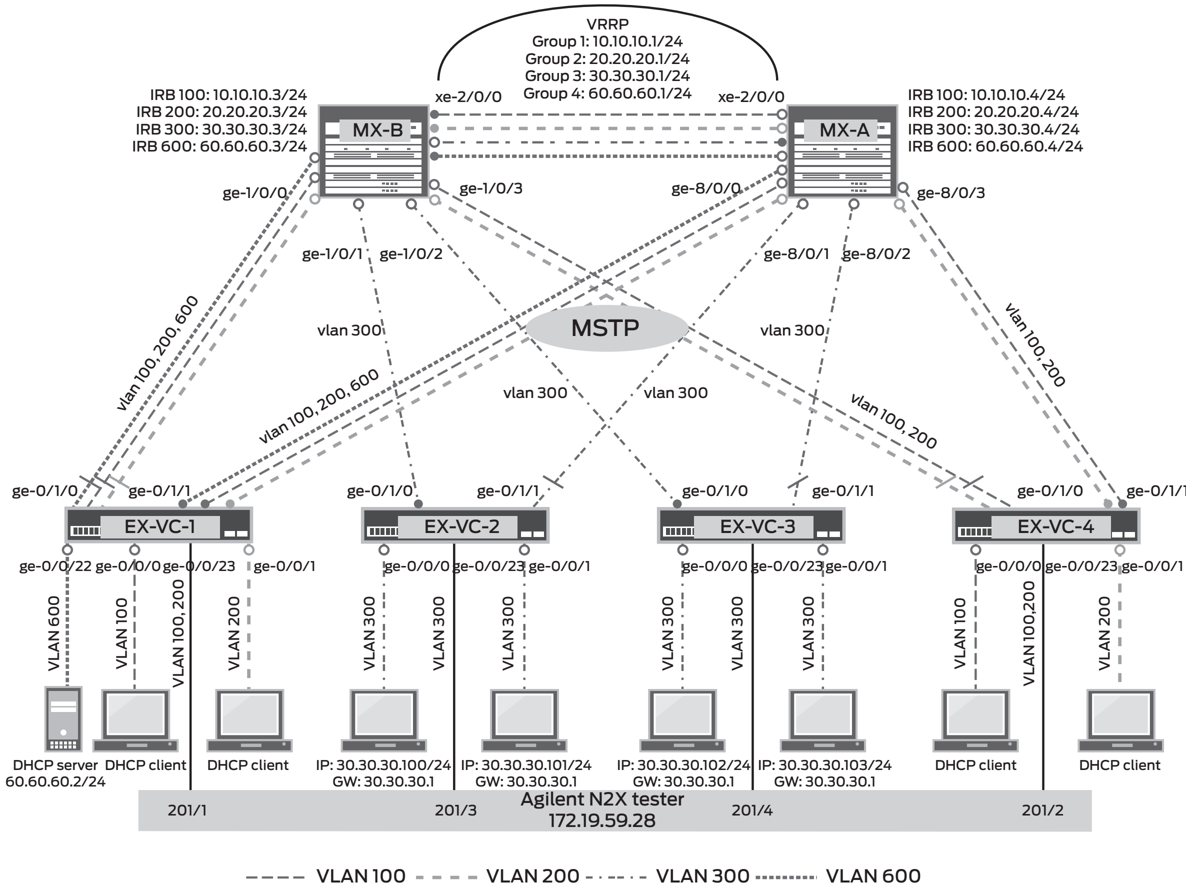


**Figure 1: Logical diagram with MSTP**

### Hardware Used for Validation

The following device platforms are required to implement the topology described:

- Four EX4200 line Ethernet switches
- Two MX Series 3D Universal Edge Routers: Juniper Networks MX240 3D Universal Edge Router, MX480 3D Universal Edge Router, or MX960 3D Universal Edge Router. We have used one MX480 and one MX960 for the aggregation devices.

**Table 1: Hardware**

| EQUIPMENT | COMPONENTS |
|---|---|
| 4 x EX4200 | 4 x 4-port uplink Gigabit Ethernet module (EX-UM-4SFP) |
| | 21 small form-factor pluggable transceivers (SFPs) |
| 1 x MX480 | 2 10-gigabit small form-factor pluggable transceivers (XFPs) |
| 1 x MX960 | |

## Testing Used for Validation

Table 2: Testing Hardware

| EQUIPMENT | COMPONENTS |
|---|---|
| Agilent N2X tester | 4 x 10/100/1000 Mb ports |
| Linux DHCP server | |

## Software Used for Validation

Table 3: Software

| EQUIPMENT | COMPONENTS |
|---|---|
| EX Series and MX Series | Junos OS 9.0 |

## Detailed Configurations

The detailed configurations for EX-VC-1 and MX-A are listed below. For the rest of the configurations, please refer to *Appendix B: Detailed Configurations.*

### EX-VC-1

**...truncated**

```
interfaces {
    ge-0/0/0 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members HR;
                }
            }
        }
    }
    ge-0/0/1 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members ENG;
                }
            }
        }
    }
    ge-0/0/2 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/0/22 {
        ether-options {
            link-mode full-duplex;
            speed {
                100m;
            }
        }
        unit 0 {
            family ethernet-switching {
                port-mode access;
```

### MX-A:

**...truncated**

```
interfaces {
    ge-1/0/0 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 100 {
            vlan-id 100;
        }
        unit 200 {
            vlan-id 200;
        }
        unit 600 {
            vlan-id 600;
        }
    }
    ge-1/0/1 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 300 {
            vlan-id 300;
        }
    }
    ge-1/0/2 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 300 {
            vlan-id 300;
        }
    }
    ge-1/0/3 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 100 {
            vlan-id 100;
        }
        unit 200 {
            vlan-id 200;
```

 

```
                vlan {
                    members SERVER;
                }
            }
        }
    }
    ge-0/0/23 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [ HR ENG ];
                }
            }
        }
    }
    ge-0/1/0 {
        enable;
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [ HR ENG SERVER
];
                }
            }
        }
    }
    ge-0/1/1 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [ HR ENG SERVER
];
                }
            }
        }
    }
    ge-0/1/2 {
        unit 0 {
            family ethernet-switching;
        }
    }
    ge-0/1/3 {
        unit 0 {
            family ethernet-switching;
        }
    }
    vme {
        unit 0 {
            family inet {
                address 172.19.59.190/24;
            }
        }
    }
}
routing-options {
protocols {
    mstp {
```

```
        }
    }
    xe-2/0/0 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 100 {
            vlan-id 100;
        }
        unit 200 {
            vlan-id 200;
        }
        unit 300 {
            vlan-id 300;
        }
        unit 600 {
            vlan-id 600;
        }
    }
    irb {
        unit 100 {
            family inet {
                address 10.10.10.3/24 {
                    vrrp-group 1 {
                        virtual-address
10.10.10.1;
                        accept-data;
                    }
                }
            }
        }
        unit 200 {
            family inet {
                address 20.20.20.3/24 {
                    vrrp-group 2 {
                        virtual-address
20.20.20.1;
                        accept-data;
                    }
                }
            }
        }
        unit 300 {
            family inet {
                address 30.30.30.3/24 {
                    vrrp-group 3 {
                        virtual-address
30.30.30.1;
                        priority 254;
                        accept-data;
                    }
                }
            }
        }
        unit 600 {
            family inet {
                address 60.60.60.3/24 {
                    vrrp-group 4 {
                        virtual-address
60.60.60.1;
                        accept-data;
```

```
        configuration-name REGION1;
        revision-level 1;
        interface ge-0/0/0.0 {
            edge;
        }
        interface ge-0/0/1.0 {
            edge;
        }
        interface ge-0/0/22.0 {
            edge;
        }
        interface ge-0/1/0.0 {
            mode point-to-point;
        }
        interface ge-0/1/1.0 {
            mode point-to-point;
        }
        msti 1 {
            vlan [ HR ENG SERVER ];
        }
        msti 2 {
            vlan SALES;
        }
    }
}
vlans {
    ENG {
        vlan-id 200;
    }
    HR {
        vlan-id 100;
    }
    SALES {
        vlan-id 300;
    }
    SERVER {
        vlan-id 600;
    }
}
```

```
                    }
                }
            }
        }
    }
    lo0 {
        unit 0 {
            family inet {
                address 127.0.0.1/32;
            }
        }
    }
}
forwarding-options {
    helpers {
        bootp {
            server 60.60.60.2;
            relay-agent-option;
            interface {
                irb.100;
                irb.200;
            }
        }
    }
}
protocols {
    mstp {
        configuration-name REGION1;
        revision-level 1;
        bridge-priority 8k;
        interface ge-1/0/0 {
            mode point-to-point;
        }
        interface ge-1/0/1 {
            mode point-to-point;
        }
        interface ge-1/0/2 {
            mode point-to-point;
        }
        interface ge-1/0/3 {
            mode point-to-point;
        }
        interface xe-2/0/0 {
            mode point-to-point;
        }
        msti 1 {
            bridge-priority 8k;
            vlan [ 100 200 600 ];
        }
        msti 2 {
            bridge-priority 4k;
            vlan 300;
        }
    }
}
bridge-domains {
    ENG {
        domain-type bridge;
        vlan-id 200;
        interface ge-1/0/0.200;
```

```
            interface ge-1/0/3.200;
            interface xe-2/0/0.200;
            routing-interface irb.200;
        }
        HR {
            domain-type bridge;
            vlan-id 100;
            interface ge-1/0/0.100;
            interface ge-1/0/3.100;
            interface xe-2/0/0.100;
            routing-interface irb.100;
        }
        SALES {
            domain-type bridge;
            vlan-id 300;
            interface ge-1/0/1.300;
            interface ge-1/0/2.300;
            interface xe-2/0/0.300;
            routing-interface irb.300;
        }
        SERVER {
            domain-type bridge;
            vlan-id 600;
            interface ge-1/0/0.600;
            interface xe-2/0/0.600;
            routing-interface irb.600;
        }

}
```

## Summary

With the EX Series Ethernet Switches and MX Series 3D Universal Edge Routers, Juniper Networks offers its enterprise customers compelling end-to-end solutions that can meet the requirements of either Layer 2 or Layer 3 deployments. In Layer 2 environments, network administrators are faced with the task of preventing and possibly troubleshooting loops. This document has shown how to implement a layer design using MSTP. Three other designs using RSTP, RTGs, and Virtual Chassis technology are described in separate implementation guides.

With the guidelines presented in these documents, Juniper Networks customers can integrate the EX Series and MX Series into their Layer 2 networks. They can later enable Layer 3 on the same devices at no additional cost to leverage Junos OS's rich feature set while minimizing capital and operational expenses.

## Appendix A: Conventions/Glossary

| | |
|---|---|
| BOOTP | Bootstrap Protocol |
| BPDU | Bridge protocol data unit |
| CST | Common Spanning Tree |
| DPC | Dense Port Concentrator |
| ESR | Ethernet Services Router |
| IRB | Integrated routing and bridging |
| MSTI | MSTP Instance |
| MSTP | Multiple Spanning Tree Protocol |
| PVST | Per-VLAN Spanning Tree |
| RSTP | Rapid Spanning Tree |
| RTG | Redundant Trunk Group |
| RVI | Routed VLAN Interface |
| STP | Spanning Tree Protocol |
| SFP | Small form-factor pluggable transceiver |
| VLAN | Virtual LAN |
| VRRP | Virtual Router Redundancy Protocol |
| VSTP | Virtual Spanning Tree Protocol |
| XFP | 10-gigabit small form-factor pluggable transceiver |

## Appendix B: Detailed Configurations

### MX-B

…truncated

```
interfaces {
    xe-2/0/0 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 100 {
            vlan-id 100;
        }
        unit 200 {
            vlan-id 200;
        }
        unit 300 {
            vlan-id 300;
        }
        unit 600 {
            vlan-id 600;
        }
    }
    /* Trunk example */
    ge-8/0/0 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 100 {
            vlan-id 100;
        }
        unit 200 {
            vlan-id 200;
```

```
        }
        unit 600 {
            vlan-id 600;
        }
    }
    ge-8/0/1 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 300 {
            vlan-id 300;
        }
    }
    ge-8/0/2 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 300 {
            vlan-id 300;
        }
    }
    ge-8/0/3 {
        vlan-tagging;
        encapsulation extended-vlan-bridge;
        unit 100 {
            vlan-id 100;
        }
        unit 200 {
            vlan-id 200;
        }
    }
    irb {
        unit 100 {
            family inet {
                address 10.10.10.4/24 {
                    vrrp-group 1 {
                        virtual-address 10.10.10.1;
                        priority 254;
                        accept-data;
                    }
                }
            }
        }
        unit 200 {
            family inet {
                address 20.20.20.4/24 {
                    vrrp-group 2 {
                        virtual-address 20.20.20.1;
                        priority 254;
                        accept-data;
                    }
                }
            }
        }
        unit 300 {
            family inet {
                address 30.30.30.4/24 {
                    vrrp-group 3 {
                        virtual-address 30.30.30.1;
                        accept-data;
                    }
                }
```

```
                }
            }
            unit 600 {
                family inet {
                    address 60.60.60.4/24 {
                        vrrp-group 4 {
                            virtual-address 60.60.60.1;
                            priority 254;
                            accept-data;
                        }
                    }
                }
            }
        }
    }
    forwarding-options {
        helpers {
            bootp {
                server 60.60.60.2;
                relay-agent-option;
                interface {
                    irb.100;
                    irb.200;
                }
            }
        }
    }
    protocols {
        mstp {
            configuration-name REGION1;
            revision-level 1;
            bridge-priority 4k;
            interface xe-2/0/0 {
                mode point-to-point;
            }
            interface ge-8/0/0 {
                mode point-to-point;
            }
            interface ge-8/0/1 {
                mode point-to-point;
            }
            interface ge-8/0/2 {
                mode point-to-point;
            }
            interface ge-8/0/3 {
                mode point-to-point;
            }
            msti 1 {
                bridge-priority 4k;
                vlan [ 100 200 600 ];
            }
            msti 2 {
                bridge-priority 8k;
                vlan 300;
            }
        }
    }
    bridge-domains {
        ENG {
            domain-type bridge;
```

 

```
        vlan-id 200;
        interface ge-8/0/0.200;
        interface ge-8/0/3.200;
        interface xe-2/0/0.200;
        routing-interface irb.200;
    }
    HR {
        domain-type bridge;
        vlan-id 100;
        interface ge-8/0/0.100;
        interface ge-8/0/3.100;
        interface xe-2/0/0.100;
        routing-interface irb.100;
    }
    SALES {
        domain-type bridge;
        vlan-id 300;
        interface ge-8/0/1.300;
        interface ge-8/0/2.300;
        interface xe-2/0/0.300;
        routing-interface irb.300;
    }
    SERVER {
        domain-type bridge;
        vlan-id 600;
        interface ge-8/0/0.600;
        interface xe-2/0/0.600;
        routing-interface irb.600;
    }
}
```

## EX-VC-3

...truncated

```
interfaces {
    ge-0/0/0 {
        unit 0 {
            family ethernet-switching {
                vlan {
                    members SALES;
                }
            }
        }
    }
    ge-0/0/1 {
        unit 0 {
            family ethernet-switching {
                vlan {
                    members SALES;
                }
            }
        }
    }
    ge-0/0/23 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
```

```
                        members SALES;
                    }
                }
            }
        }
        ge-0/1/0 {
            unit 0 {
                family ethernet-switching {
                    port-mode trunk;
                    vlan {
                        members SALES;
                    }
                }
            }
        }
        ge-0/1/1 {
            unit 0 {
                family ethernet-switching {
                    port-mode trunk;
                    vlan {
                        members SALES;
                    }
                }
            }
        }
        ge-0/1/3 {
            unit 0 {
                family ethernet-switching;
            }
        }
        vme {
            unit 0 {
                family inet {
                    address 172.19.59.193/24;
                }
            }
        }
    }
    protocols {
        mstp {
            configuration-name REGION1;
            revision-level 1;
            interface ge-0/0/0.0 {
                edge;
            }
            interface ge-0/0/1.0 {
                edge;
            }
            interface ge-0/1/0.0 {
                mode point-to-point;
            }
            interface ge-0/1/1.0 {
                mode point-to-point;
            }
            msti 1 {
                vlan [ HR ENG SERVER ];
            }
            msti 2 {
                vlan SALES;
            }
```

```
    }
}
vlans {
    ENG {
        vlan-id 200;
    }
    HR {
        vlan-id 100;
    }
    SALES {
        vlan-id 300;
    }
    SERVER {
        vlan-id 600;
    }
}
```

## EX-VC-4

...truncated

```
interfaces {
    ge-0/0/0 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members HR;
                }
            }
        }
    }
    ge-0/0/1 {
        unit 0 {
            family ethernet-switching {
                port-mode access;
                vlan {
                    members ENG;
                }
            }
        }
    }
    ge-0/0/23 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [ HR ENG ];
                }
            }
        }
    }
    ge-0/1/0 {
        unit 0 {
            family ethernet-switching {
                port-mode trunk;
                vlan {
                    members [ HR ENG ];
```

```
                    }
                }
            }
        }
        xe-0/1/0 {
            unit 0 {
                family ethernet-switching;
            }
        }
        ge-0/1/1 {
            unit 0 {
                family ethernet-switching {
                    port-mode trunk;
                    vlan {
                        members [ HR ENG ];
                    }
                }
            }
        }
        xe-0/1/1 {
            unit 0 {
                family ethernet-switching;
            }
        }
        vme {
            unit 0 {
                family inet {
                    address 172.19.59.194/24;
                }
            }
        }
    }
    protocols {
        mstp {
            configuration-name REGION1;
            revision-level 1;
            interface ge-0/0/0.0 {
                edge;
            }
            interface ge-0/0/1.0 {
                edge;
            }
            interface ge-0/1/0.0 {
                mode point-to-point;
            }
            interface ge-0/1/1.0 {
                mode point-to-point;
            }
            msti 1 {
                vlan [ HR ENG SERVER ];
            }
            msti 2 {
                vlan SALES;
            }
        }
    }
    vlans {
        ENG {
            vlan-id 200;
        }
```

```
    HR {
        vlan-id 100;
    }
    SALES {
        vlan-id 300;
    }
    SERVER {
        vlan-id 600;
    }
}
```

## About Juniper Networks

Juniper Networks, Inc. is the leader in high-performance networking. Juniper offers a high-performance network infrastructure that creates a responsive and trusted environment for accelerating the deployment of services and applications over a single network. This fuels high-performance businesses. 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.



8010065-001-EN Dec 2009

Printed on recycled paper