CHAPTER 7

Switching Infrastructure

Baseline switching security is concerned with ensuring the availability of the Layer 2 switching network. This section highlights the key steps to securing and preserving the switching infrastructure, including:

- Restrict Broadcast Domains
- Spanning Tree Protocol (STP) Security
- VLAN Best Common Practices

CSF Methodology Assessment

The results of applying the CSF methodology for baseline switching security are presented in Table 7-1 and Table 7-2. The tables highlight the technologies and features identified for baseline switching security and which are integrated in Network Security Baseline.

Total Visibility

Table 7-1 Switching Infrastructure—Total Visibility

Identify	Monitor	Correlate
	• Logging – Syslog – SNMP	

Complete Control

Table 7-2 Switching Infrastructure—Complete Control

Harden	Isolate	Enforce
	• Restrict Broadcast Domains • VLAN • L3 hierarchical design	• STP Security – Disable dynamic trunking – PVST – BPDU Guard – Root guard • VLAN Best Common Practices

Restrict Broadcast Domains

By definition, LAN switches are responsible for forwarding unknown frames, multicast frames and broadcast frames throughout the LAN segment, forming a broadcast domain. While broadcast domains facilitate Layer 2 connectivity between systems on a LAN segment, designing networks with unnecessarily large broadcast domains has potential drawbacks.

First, in large networks, the flooding of unknown, multicast and broadcast frames may degrade performance, even to the point of breaking connectivity. In addition, a broadcast domain defines a failure domain, whereby typically all systems and switches on the same LAN segment suffer during a failure. Therefore, the larger the broadcast domain, the bigger the impact of a failure. Finally, larger broadcast domains increase the chances of security incidents.

To avoid the challenges described above, it is a good practice to segment broadcast domains into multiple IP subnets or VLANs using a hierarchical design.The use of hierarchical design principles provides the foundation for implementing scalable and reliable LANs.

Figure 7-1 shows a recommended hierarchical design. This design uses a building block approach leveraging a high-speed routed core network layer to which are attached multiple independent distribution blocks. The distribution blocks comprise two layers of switches: the actual distribution nodes that act as aggregators, and wiring closet access switches. The hierarchical design segregates the functions of the network into these separate building blocks to provide for availability, flexibility, scalability, and fault isolation.

A hierarchical design like the one proposed here helps restrict the size of broadcast domains, improving convergence, easing deployments, and reducing the scope of failure domains. This is done by isolating a VLAN to a single wiring closet or single switch. As a result, better convergence and load-balancing upstream can be achieved through the use of L3 protocols, with no need for STP and redundancy protocols such as HSRP and VRRP. In addition, L3 designs are not subject to the same bandwidth and cable plant constraints as L2 designs; and failures are typically confined to a neighbor or route loss, instead of impacting entire broadcast domains like in L2 designs. In cases where L3 to the edge is not viable, broadcast domains should still be restricted to have no loops, with no blocked links, and each access switch having its own, unique VLANs

Figure 7-1 *Hierarchical Design*

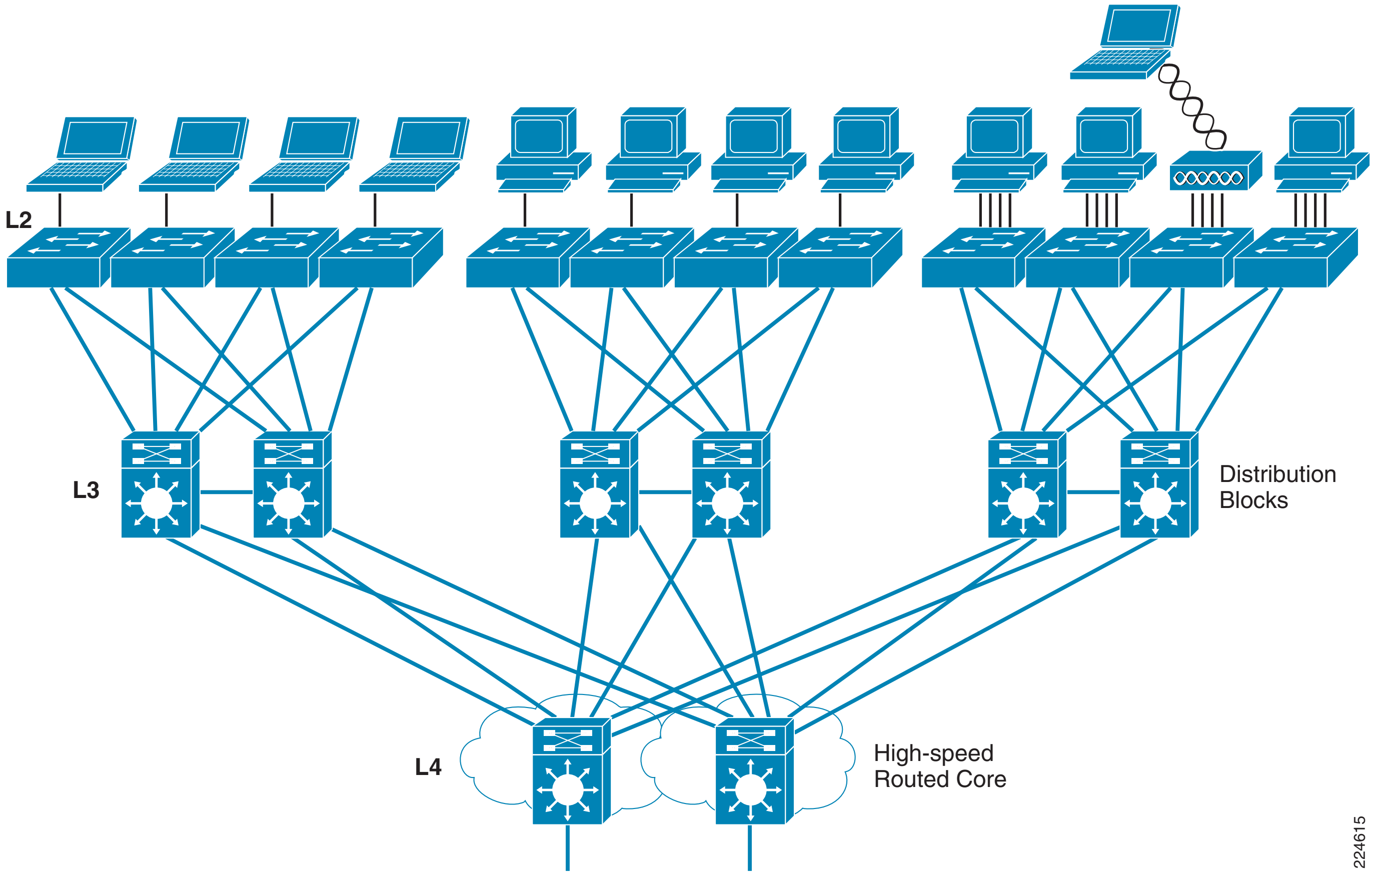

Spanning Tree Protocol Security

Spanning Tree Protocol (STP) is a link management protocol, defined in the IEEE 802.1D, for bridged networks. STP provides path redundancy while preventing undesirable loops in networks consisting of multiple active paths.

Loops occur when multiple active paths exist between hosts, and which could result in an endless loop of traffic in the LAN that could bring the network down. STP implements an algorithm that guarantees a loop-free topology. With STP, all switches and bridges in the LAN exchange BPDU messages containing topology information. The STP algorithm uses the topology information to build a topological tree where only one active path exists at a time between any two hosts. Redundant paths are shutdown and used as backups in case the primary paths fail. Changes to the physical topology normally trigger a recalculation of the topological tree.

A newer version of STP exists, called Rapid-STP, is defined in IEEE 802.1w. Rapid-STP (RSTP) works similarly to STP but provides better convergence after a failure of a switch, switch port, or a LAN. RSTP significantly reduces the time to reconfigure the active topology of the network when changes occur to the physical topology or its configuration parameters. RSTP supersedes STP specified in 802.1d, but remains compatible with STP.

STP is a useful protocol but, unfortunately, both versions of the protocol were conceived with no security in mind and, as a result, are both vulnerable to several types of attacks. STP does not implement any authentication and encryption to protect the exchange of BPDUs. Because of the lack of authentication, anyone can speak to a STP-enabled device. An attacker could very easily inject bogus BPDUs, triggering a topology recalculation. A forced change to the STP topology could lead to a denial of service condition, or leave the attacker as a man-in-the-middle. In addition, because BPDUs are not encrypted, it is fairly simple to intercept BPDUs in transit, revealing important topology information.

STP introduces some security risks but, in topologies where a loop-free design is not possible, STP should be used along with the Cisco features (see Table 7-3) developed to address its risks. Not using STP would result in a loop becoming another attack vector.

Table 7-3 STP Security Features

STP Attacks and Vulnerabilities	Attack Objectives and Risk	Possible Countermeasures
Illegitimate trunk		• Disable Dynamic Trunking
STP spans VLANs	Attack on one VLAN impacts all other VLANs	• Restrict STP domain using Per-VLAN Spanning Tree (PVST)
Unauthorized spanning tree participation Bogus BPDU packets Superior BPDUs sent to become root bridge	Network instability Attacker sees frames he should not Can be used for MITM, DoS, etc	• BPDU guard • Root Guard

Cisco IOS offers a number of features that help protect bridged networks using STP against the common attacks. The following are the recommended best practices:

- Disable VLAN dynamic trunk negotiation trunking on user ports
- Use Per-VLAN Spanning Tree (PVST)
- Configure BPDU Guard
- Configure STP Root Guard
- Disable unused ports and put them into an unused VLAN
- Implement Port Security

  See Chapter 4, "Device Resiliency and Survivability."
- Enable Traffic Storm Control

  See Chapter 4, "Device Resiliency and Survivability."

Disable Dynamic Trunking

Dynamic trunk negotiation is a feature that facilitates the deployment of switches by an interface automatically configuring itself as a trunk according to the interface type of its neighboring. However, this feature can easily be abused to set up an illegitimate trunk. For this reason, dynamic trunking should be disabled on all ports connecting to end users.

Cisco IOS, by default, sets an interface to dynamic negotiation mode. This can be enabled using the command switchport mode and setting the port mode type to access. An example is shown in the following:

```
Router(config)# interface type slot/port
Router(config-if)# switchport mode access
```

The configuration makes the port a non-trunking, non-tagged single VLAN Layer 2 interface.

Note Catalyst 6500 switches running Cisco IOS software support the macro command switchport host. The switchport host macro command was designed to facilitate the configuration of switch ports that connect to end stations. Entering this command sets the switch port mode to access, enables spanning tree PortFast, and disables channel grouping, all at the same time. The switchport host macro command can be used as an alternative to the switchport mode access command.

For more information on the **switchport mode** command on the Catalyst 6500, refer to the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.1E/native/command/reference/S1.html#wpxref79609

For more information on the **switchport mode** command on the Catalyst 4500, refer to the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/31sg/command/reference/snmp_vtp.html#wp1210450

Per VLAN Spanning Tree (PVST)

Per-VLAN Spanning Tree (PVST) is a feature available on Catalyst 6500 and 4500 Series switches that implements a separate instance of spanning tree for each VLAN configured in the network. Having a separate instance of STP per VLAN makes the network more resilient to attacks against spanning tree. If a problem occurs in one VLAN, the effects are contained in that VLAN, shielding the rest of the network.

There are different versions of PVST, but which all maintain separate spanning tree instances per VLAN, and work in a similar fashion. Per VLAN Spanning Tree (PVST) is the original version, and which uses ISL trunking. Per VLAN Spanning Tree Plus (PVST+) provides the same functionality as PVST using 802.1Q trunking technology rather than ISL. PVST+ is an enhancement to the 802.1Q specification and is not supported on non-Cisco devices. Rapid-Per-VLAN-Spanning Tree (Rapid-PVST+) is another version of PVST that provides faster convergence of the spanning tree by using Rapid Spanning Tree Protocol (RSTP) with the existing configuration for PVST+.

PVST is enabled by default in Cisco IOS and it is recommended that PVST is always enabled.

On a Catalyst 6500 or 4500 running Cisco IOS, the default spanning tree protocol is PVST+. Rapid-PVST+ is also supported on these platforms.

In Cisco IOS, the Spanning Tree mode can be modified using the **spanning-tree mode** command. For example, to configure Rapid-PVST+:

```
Router(config)# spanning-tree mode rapid-pvst
```

For more information on the **spanning-tree mode** command on the Catalyst 6500, refer to the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.1E/native/command/reference/S1.html#wp1179548

For more information on the **spanning-tree mode** command on the Catalyst 4500, refer to the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/31sg/command/reference/snmp_vtp.html#wp1144173

BPDU Guard

Since STP does not implement any authentication or encryption to protect the exchange of BPDUs, it is vulnerable to unauthorized participation and attacks, as highlighted earlier. Cisco IOS offers the BPDU Guard feature to restrict participation in spanning tree.

End user ports should not be participating in spanning tree and, by enabling BPDU Guard on these ports, the port is shutdown if a BPDU is received. In this way, BPDU guard helps prevent unauthorized participation in spanning tree and the injection of forged BPDUs.

BPDU can be configured per port or globally. When configured globally, BPDU Guard is only effective on ports in the operational PortFast state.

BPDU Guard requires STP PortFast to be already configured on a port.

In Cisco IOS, BPDU Guard can be enabled on an interface by enabling PortFast and then using the **spanning-tree bpduguard** command as follows:

```
Router(config)# interface fastethernet 3/1
Router(config-if)# spanning-tree portfast
Router(config-if)# spanning-tree bpduguard enable
```

In Cisco IOS, BPDU Guard can be enabled globally by using the **spanning-tree portfast bpduguard default** command as follows:

```
Router(config)# spanning-tree portfast bpduguard default
```

When enabled globally, BPDU Guard applies to all interfaces that are in an operational PortFast state.

For more information on the **spanning-tree bpduguard** command on the Catalyst 6500, refer to the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.1E/native/command/reference/S1.html#wp1179312

For more information on the **spanning-tree bpduguard** command on the Catalyst 4500, refer to the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/31sg/command/reference/snmp_vtp.html#wp1065041

For more information on the **spanning-tree portfast bpduguard default** command on the Catalyst 6500, refer to the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.1E/native/command/reference/S1.html#wp1180500

For more information on the **spanning-tree portfast bpduguard default** command on the Catalyst 4500, refer to the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/31sg/command/reference/snmp_vtp.html#wp1072464

STP Root Guard

As highlighted earlier, since STP does not implement any authentication or encryption to protect the exchange of BPDUs, it is vulnerable to unauthorized participation and attacks. Cisco IOS offers the STP Root Guard feature to enforce the placement of the root bridge.

STP root guard forces a port to become a designated port so that no switch on the other end of the link can become a root switch. If a port configured for root guard receives a superior BPDU, the port it is received on is blocked. In this way, STP root guard blocks other devices from trying to become the root bridge.

STP root guard should be enabled on all ports that will never connect to a root bridge, for example, all end user ports. This ensures that a root bridge will never be negotiated on those ports.

STP root guard requires STP PortFast to be already configured on a port. STP root guard is configured on a per-port basis.

In Cisco IOS, STP Root Guard can be enabled on an interface using the **spanning-tree guard root** command as follows:

```
Router(config)# interface fastethernet 3/1
Router(config-if)# spanning-tree guard root
```

Note Do not enable loop guard and root guard on a port at the same time. Root guard forces a port to always be designated as the root port. Loop guard is effective only if the port is a root port or an alternate port.

For more information on the **spanning-tree guard root** command on the Catalyst 6500, refer to the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst6500/ios/12.1E/native/command/reference/S1.html#wp1179394

For more information on the **spanning-tree guard root** command on the Catalyst 4500, refer to the following URL:

http://www.cisco.com/en/US/docs/switches/lan/catalyst4500/12.2/31sg/command/reference/snmp_vtp.html#wp1031770

VLAN Best Common Practices

VLAN hopping is an attack vector which provides a client with unauthorized access to other VLANs on a switch. This type of attack can be easily mitigated by applying the following best common practices:

- Always use a dedicated VLAN ID for all trunk ports
- Disable all unused ports and put them in an unused VLAN
- Do not use VLAN 1 for anything
- Configure all user-facing ports as non-trunking (DTP off)
- Explicitly configure trunking on infrastructure ports
- Use all tagged mode for the native VLAN on trunks and drop untagged frames

- Set the default port status to "disable"