External Layer 2 Switching Domain Connected to an SD-Access Edge Node

Best Practices Guide

June 2021

First Publish: 28 June 2021

Contents

Document Organization

This document is organized into the following sections:

Document Conventions

How to Read Deployment Commands

The guide uses the following conventions for commands that are entered in the command-line interface (CLI).

Commands to enter at a CLI prompt:

```
configure terminal
```

Commands that specify a value for a variable (variable is in italics):

```
ntp server 10.4.0.1
```

Commands with variables that you must define (definition is in braces):

```
router bgp {autonomous-system-number}
```

Commands at a CLI prompt (entered commands are in bold):

```
Router# enable
```

Long command lines that wrap on a printed page (underlined text is entered as one command):

```
monitor capture CAP interface
GigabitEthernet1/0/1 both limit pps 10000
```

When discussing physical or logical network interfaces, this document uses the words *interface* and *port* interchangeably.

Many verification commands are provided with two outputs. When multiple outputs are provided, one output shows the undesired result, and the other output shows the desired results. This convention is used to help explain explicitly what to look for as a result of the configuration to allow for a comparison element before and after.

Verification commands are shown using the command line interface (CLI) although Cisco DNA Center Command Runner can also be used. Configurations are primarily shown using the CLI. Some configurations must be performed in the Cisco DNA Center User Interface (UI) using particular workflows or sequences of steps. When this is applicable, these configurations are shown using the UI. Screen captures from UI were taken using a Cisco DNA Center cluster running software versions 2.2.2.0 through 2.2.2.3.

About This Document

To accelerate customers’ Cisco SD-Access journey, connecting an External Layer 2 Switching Domain to a Fabric Edge Node is now supported. This External Layer 2 Switching Domain is external to the Fabric which means it is not managed by the SD-Access package in Cisco DNA Center. Other Cisco DNA Center applications such as Cisco Plug and Play (PnP), Assurance, and Software Image Management (SWIM) can still be used to monitor and manage the External Layer 2 Switching Domain infrastructure if the devices are supported by the application. Please see the Cisco DNA Center Supported Devices Matrix Compatibility Information for details.

This document captures the recommended practices that should be implemented on the Fabric Edge Nodes when connecting an External Layer 2 Switching Domain to it. These recommended practices are designed to minimize the common issues that arise in the Layer 2 part of the network while protecting the fabric network against Layer 2 loops and Layer 2 misconfiguration.

Tech tip

An External Layer 2 Switching Domain can also be connected to a Border Node with a Layer 2 handoff. However, the scope of this guide is limited exclusively to connecting the External Layer 2 Switching Domain to a Fabric Edge Node.

The recommended practices captured in this document can be configured on the Edge Nodes via templates in Cisco DNA Center or manually using the devices’ CLI.

Supported Topology Details

The supported hardware and software platforms for Edge Nodes are scoped by a given Cisco DNA Center release. Please see the Cisco Software-Defined Access Compatibility Matrix for the supported hardware and software per release.

The Edge Node can be a standalone switch (single switch), a switch stack (hardware stacking), or operate in StackWise Virtual. The Edge Node becomes the first-hop Layer 3 gateway for the clients located in the External Layer 2 Switching Domain. Except as noted below, this makes the SD-Access fabric largely agnostic to the External Layer 2 Switching Domain which can have many potential variations.

Supported connectivity between the Edge Node and an External Layer 2 Switching Domain is as follows:

- The External Layer 2 Switching Domain can be connected to a single Edge Node or to a single, logical Edge Node such as a hardware stack or StackWise Virtual.
- The External Layer 2 Switching Domain **cannot** be connected to two devices operating as independent Edge Nodes.
- The External Layer 2 Switching Domain **cannot** be connected to or through Classic or Policy Extended Nodes.
- The physical interconnect between the Edge Node and the External Layer 2 Switching Domain can be a trunk port or an access port. Further recommendations on when to use a trunk or access port are provided later in this guide.
- This physical interconnect can be a single interface or a Layer 2 EtherChannel.
- The Layer 2 EtherChannel is supported as a Multichassis EtherChannel (MEC). This means it can terminate on multiple, physical Edge Nodes operating as a single logical unit such as a hardware stack or StackWise Virtual.

Figure 1. **External Layer 2 Switching Domain Topology Examples**

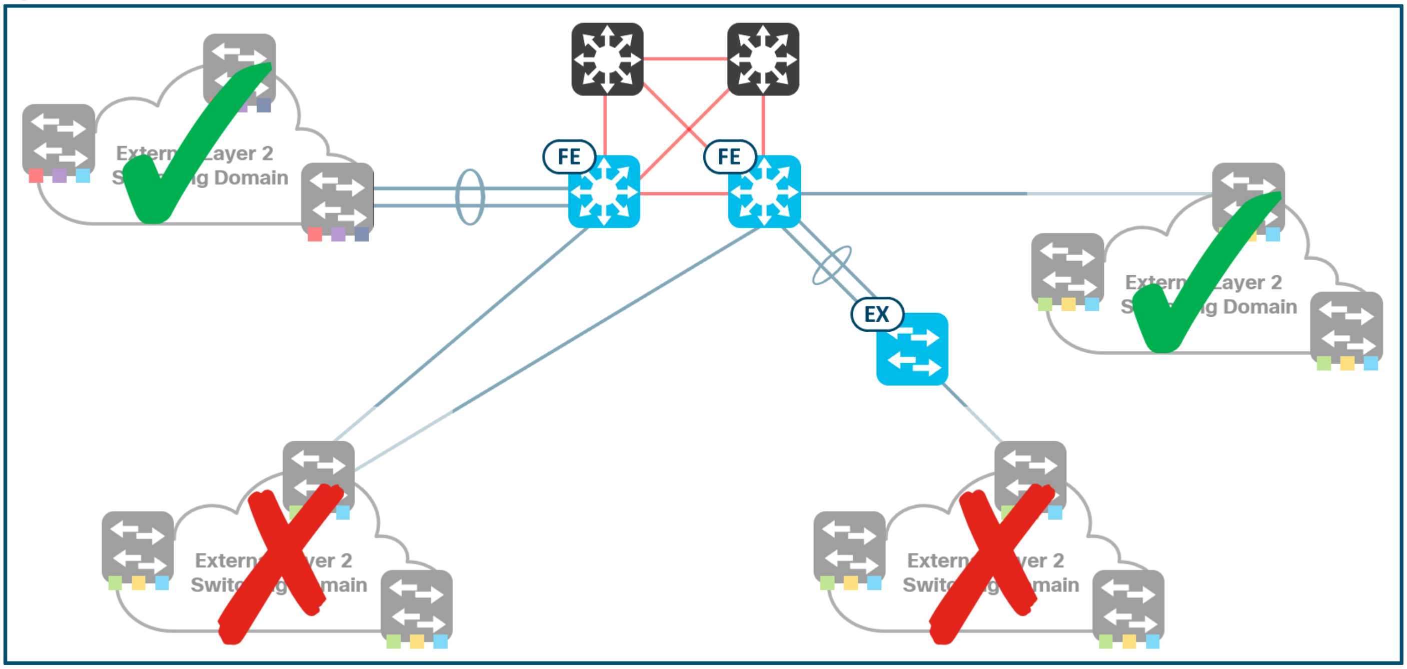
FE
FE
EX
External Layer 2 Switching Domain
External Layer 2 Switching Domain
External Layer 2 Switching Domain
External Layer 2 Switching Domain

Fabric Edge Node High-Availability

The Fabric Edge Node functionality is not limited to the access layer in the network hierarchy. With the ability to connect External Layer 2 Switching Domains along with features such as Catalyst 9000 Series Policy Extended Nodes, the Edge Node function may be deployed in the distribution layer of a classic three-tier hierarchical network architecture model.

A common convention and long-standing recommendation in physical layer planning is to dual home or dual multi-home access layer devices to the distribution block switches. In migration scenarios, the boundary switch of the External Layer 2 Switching Domain can also be considered as an *access layer switch*. However, as of this writing, an External Layer 2 Switching Domain cannot be connected to more than one Edge Node as noted above. If the Edge Node is at the distribution layer, this recommendation about dual-homing the access switch presents a challenge.

There are two methods to address this challenge: hardware stacking and StackWise Virtual (SVL). Both methods take multiple physical devices and allow them to operate as one logical device. If there are deep crosslinks or non-uniformity in the External Layer 2 Switching Domain and redundancy is required at the Edge Node, deploy either a hardware stack or SVL, depending on the platform.

However, there are further caveats for StackWise Virtual in Cisco DNA Center versions earlier than 2.2.2.3. In these versions of software, a Fabric Access Point is not supported downstream from an SVL Edge Node, whether being directly connected or connected through another device such as an Extended Node or third-party switch.

For deployments that need high availability for Edge Nodes and have deployed SD-Access Wireless on Cisco DNA Center versions earlier than 2.2.2.3, use hardware stacking for Edge Node deployment, regardless of if they are at the access or distribution layer.

Tech tip
An Access Point connected to the External Layer 2 Switching Domain can operate as a Fabric AP. Fabric APs build a VXLAN tunnel to their first-hop Layer 3 switch which must be the Edge Node. For Fabric AP support downstream from an SVL Edge Node, a minimum version of Cisco DNA Center 2.2.2.3 and IOS XE 17.3.3 is required. Please consult the Cisco DNA Center Release Notes and Cisco Software-Defined Access Compatibility Matrix for further details.

Recommended Practices – General

This chapter is organized into the following sections:

Chapter	Section
Recommended Practices - General	Virtual Trunking Protocol Spanning Tree Mode Spanning Tree Protocol Primer Spanning Tree Root Bridge BPDU Filter EtherChannel (Link Aggregation) Error Disable Detect Error Disable Recovery Control Plane Policing (CoPP) Summary

Virtual Trunking Protocol (VTP)

VTP is a Layer 2 messaging protocol that maintains VLAN configuration consistency by managing the addition, deletion, and renaming of VLANs within a VTP domain. VTP has three different versions (Version 1-3), and a device can operate in one of four different VTP modes (Server, Client, Transparent, and Off).

Tech tip

By default, Cisco switches supported as Edge Nodes run VTP Version 1 and operate in VTP Server mode. When a device is onboarded through LAN Automation, Cisco DNA Center provisions the device in VTP Transparent mode.

The SD-Access fabric must not participate in VTP, and the fabric cannot be used as a transport for VTP. To facilitate this, the VTP mode should be set to `off`.

Step 1. Disable VTP on all fabric devices that are switching platforms.

```
Fabric Device(config)# vtp mode off
```

Step 2. Confirm the VTP operating mode is Off.

```
Fabric Device# show vtp status
VTP Version capable                : 1 to 3
VTP version running                : 1
VTP Domain Name                    :
VTP Pruning Mode                   : Disabled
VTP Traps Generation               : Enabled
Device ID                          : xxxx.xxxx.xxxx

Feature VLAN:
--------------
VTP Operating Mode                   : Server[1]
Maximum VLANs supported locally      : 1005
Number of existing VLANs             : 5
Configuration Revision               : 0
! Output omitted for brevity
```

[1] VTP Operating Mode is Server.

```
Fabric Device# show vtp status
VTP Version capable                : 1 to 3
VTP version running                : 1
VTP Domain Name                    :
VTP Pruning Mode                   : Disabled
VTP Traps Generation               : Enabled
Device ID                          : xxxx.xxxx.xxxx

Feature VLAN:
--------------
VTP Operating Mode                   : Off[1]
Maximum VLANs supported locally      : 1005
Number of existing VLANs             : 5
Configuration Revision               : 0
! Output omitted for brevity
```

[1] VTP Operating Mode is **Off**. This is the desired result.

Tech tip

In certain migration scenarios, VTP may currently be in use in the network. When deploying Edge Nodes at the distribution layer and using Catalyst 9000 Series Policy Extended Node or simply a Layer 2 Switched Access network, the existing VTP will not propagate across the fabric.

If VTP is currently in use in a migration environment, please be aware that the addition and removal of VLANs within the legacy domain will need to be addressed manually either switch by switch or by creating a VTP server on each portion of the network connected to the same Edge Node.

Spanning Tree Mode

By default, Cisco switches supported as Edge Nodes are configured to use Rapid Per-VLAN Spanning Tree (Rapid PVST+). This mode should not be changed. Ensure the Edge Node is configured to use Rapid PVST+.

Step 1. Configure the Edge Node to operate in Rapid Per-VLAN Spanning Tree.

```
Edge Node(config)# spanning-tree mode rapid-pvst
```

Step 2. Confirm the Edge Node is operating in Rapid Per-VLAN Spanning Tree.

```
Edge Node# show spanning-tree summary
Switch is in rapid-pvst mode

! Output omitted for brevity
```

Tech tip

Running the same version of Spanning Tree Protocol on devices participating in the STP domain provides the greatest degree of compatibility and interoperability. If the External Layer 2 Switching Domain uses Multiple Spanning Tree Protocol (MSTP), please consult with your Cisco Subject Matter Expert for the details of running MSTP on the Fabric Edge Node.

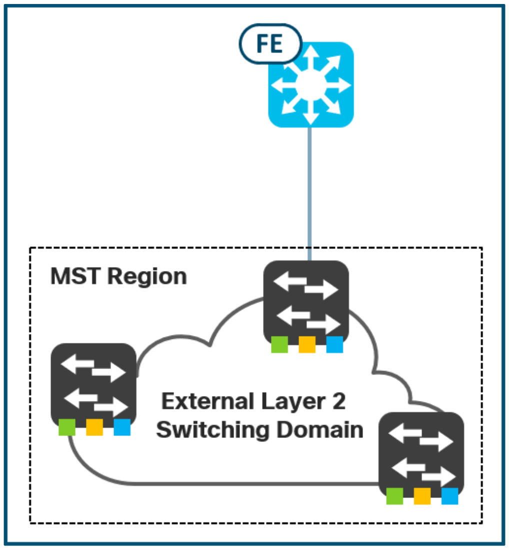

Spanning Tree Protocol Primer

When creating fault-tolerant internetworks, a loop-free path is needed between nodes in the network. The spanning tree algorithm calculates a loop-free topology throughout a switched Layer 2 network. In order to provide path redundancy, as well as to avoid loop conditions, Spanning Tree Protocol (STP) forces certain redundant connected ports into the blocked state and leaves other connected ports in the forwarding state. If a port in the forwarding state becomes unavailable, STP recomputes the loop-free topology in the network and restores the connectivity. As a result, one or more of the previously blocked ports will have moved to the forwarding state.

Classic STP operation can be broadly summarized in three steps:

- Elect the root switch (root bridge).
- On every non-root switch, select the port closest to the root bridge called the Root Port and put it in the Forwarding state. There is exactly one Root Port on every non-root switch.
- On every switch, for every non-Root port, decide whether the port is eligible to be a Designated port; if so, put it in the Forwarding, otherwise put it in the Blocking (Discarding) state. There is exactly one Designated port for every link in the topology.

A loop-free topology in the network is formed by a set of links interconnecting the switches where on each link, one of the ports is Designated and the other port is Root.

With Per-VLAN Spanning Tree and Rapid Per-VLAN Spanning Tree, every VLAN has its own independent loop-free topology, including its own root switch. A single switch can act as the root switch for multiple and even all VLANs created in the network.

Spanning Tree Root Bridge

The selection of the root switch for a particular VLAN is very important. The root switch elections are based on the Bridge ID which is a unique 8-byte value for each switch. The Bridge ID consists of a configurable priority, VLAN number, and the base

MAC address of the switch. Within a VLAN, the root switch will be the switch whose Bridge ID – taken simply as an unsigned 8-byte integer – is the lowest.

If the STP priorities of the switches are left at their default values, the placement of the root switch and the resulting spanning tree might not be optimal with respect to link speeds, creating suboptimal paths in the network.

Figure 2. Suboptimal Forwarding Based on STP Root Bridge Selection

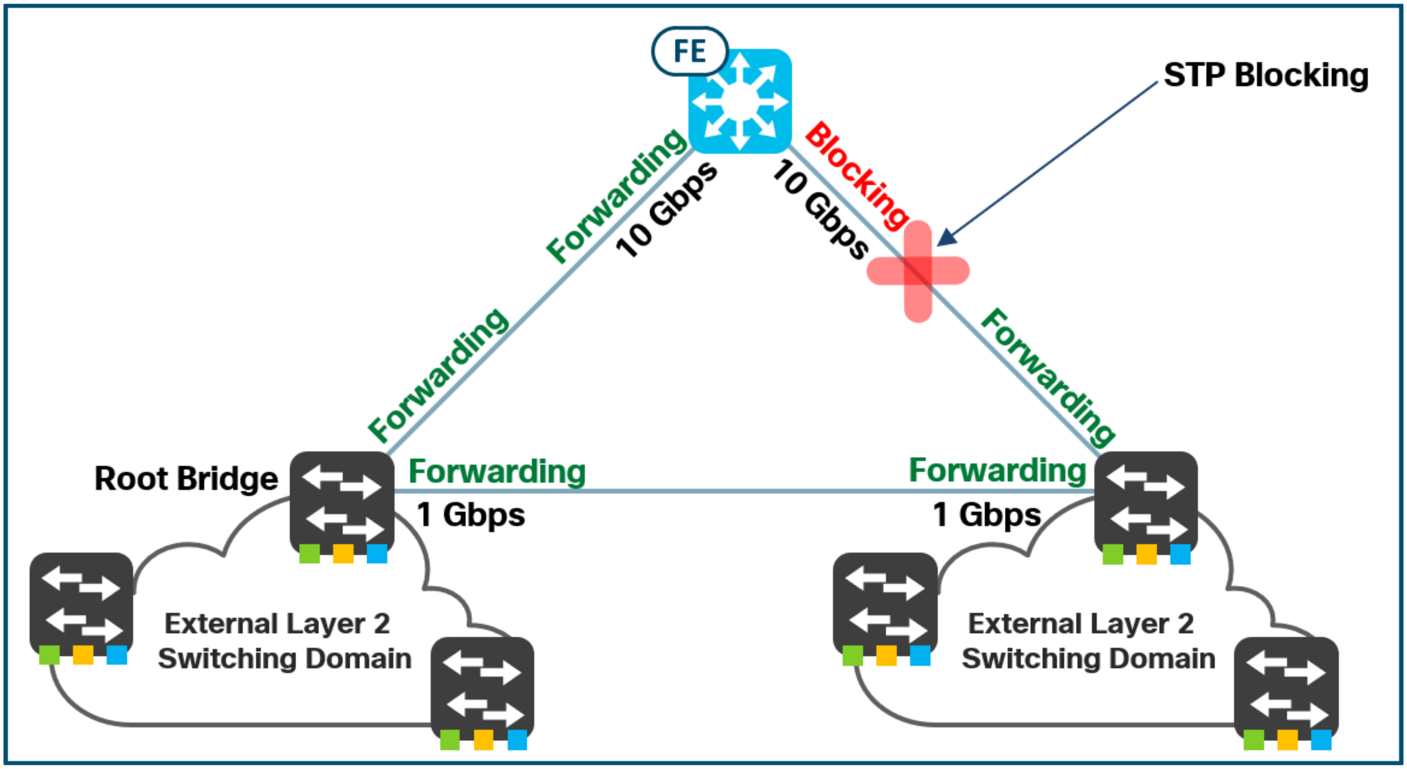

In the example network above, the root bridge election has resulted in a suboptimal forwarding path. While STP has created a loop-free topology by placing the redundant link into the blocking state, traffic destined from the External Layer 2 Switching Domain on the right to the Edge Node does not take the most direct nor highest speed and highest bandwidth path.

In SD-Access, the interconnections between Edge Nodes and Border Nodes are required to be **fully routed**. VLANs are terminated on the Edge Nodes where they are attached, and VLANs are only stretched across the routed fabric using an overlay VXLAN encapsulation.

Since the fabric is routed internally, STP running in VLANs on Edge Nodes is not extended into the fabric. STP Bridge Protocol Data Units (BPDUs) are not tunneled across the SD-Access fabric. As a result, the fabric terminates individual STP domains at the Edge Node.

When connecting External Layer 2 Switching Domains to an Edge Node, the placement of the STP root bridge at the Edge Node provides the highest degree of predictable network convergence and end-user experience. Unless the External Layer 2 Switching Domain requires a particular STP design, configuring the Edge Node as the root bridge establishes a very predictable and proven STP design. Using this design, the root switch is colocated with the switch that terminates the VLAN and the switch that implements the default gateway functions.

Step 1. Configure the Edge Node to be the STP root bridge for all potential VLANs.

```
Edge Node(config)# spanning-tree vlan 1-4094 priority 0
```

Step 2. Confirm the Edge Node is the STP root bridge.

```
Edge Node# show spanning-tree
! Output omitted for brevity

VLAN1021
  Spanning tree enabled protocol rstp
  Root ID     Priority    33789
              Address     f86b.d92c.3560[1]
              Cost        2000
              Port        2097 (Port-channel9)
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    33789  (priority 32768 sys-id-ext 1021)
              Address     f86b.d92c.3ba0[2]
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time  300 sec

Interface           Role  Sts  Cost      Prio.Nbr Type
------------------- ---- ---  --------- -------- --------------------------------
Po9                 Root[3] FWD[4] 2000      128.2097 P2p
```

[1] This is the root bridge ID. Compare this to **2**.

[2] The is the bridge ID of the local switch. Compare this to **1**.

[3] This interface is a root port. A root bridge does not have root ports.

[4] This interface is Forwarding state.

```
Edge Node# show spanning-tree
! Output omitted for brevity

VLAN1021
  Spanning tree enabled protocol rstp
  Root ID     Priority    1021
              Address     f86b.d92c.3ba0[1]
              This bridge is the root[2]
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    1021   (priority 0 sys-id-ext 1021)
              Address     f86b.d92c.3ba0[3]
              Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
              Aging Time  300 sec

Interface           Role  Sts  Cost      Prio.Nbr Type
------------------- ---- ---  --------- -------- --------------------------------
Po9                 Desg[4] FWD[5] 2000      128.2097 P2p
```

[1] This is the root bridge ID. Compare this to **3**.

[2] This device is the root bridge for this VLAN. This is the desired result.

[3] The is the bridge ID of the local switch. Compare this to **1**.

[4] This interface is Designated role.

[5] This interface is Forwarding state.

Regarding PortFast

The PortFast feature enables a port to enter the spanning-tree Forwarding state immediately, bypassing the Listening and Learning states.

When an Authentication Template is selected at the fabric-site level or through **Host Onboarding** > **Port Assignment**, Cisco DNA Center enables PortFast on all access ports.

BPDU Filter

BPDU Filter can effectively disable STP on a port. By default, BPDU Filter is disabled and should remain in this state. Do not configure BPDU Filtering on the Edge Node either on an interface or in global configuration mode.

Step 1. Confirm BPDU Filtering is **disabled** on Catalyst 3650, 3850, 4500, 9200/L, 9300/L, 9400, 9500 Series Switches. Do not enable BPDU Filtering.

```
C9300# show spanning-tree summary
Switch is in rapid-pvst mode
Root bridge for: none
EtherChannel misconfig guard                is enabled
Extended system ID                          is enabled
Portfast Default                            is disabled
PortFast BPDU Guard Default                 is enabled
Portfast BPDU Filter Default                is disabled[1]
Loopguard Default                           is disabled
UplinkFast                                  is disabled
BackboneFast                                is disabled
```

[1] BPDU Filtering is **disabled**. This is the desired result.

Step 2. Confirm BPDU Filtering is **disabled** on 9500H Series Switches. Do not enable BPDU Filtering.

```
C9500H# show spanning-tree summary
Switch is in rapid-pvst mode
Root bridge for: none
EtherChannel misconfig guard                is enabled
Extended system ID                          is enabled
Portfast Default                            is disabled
Portfast Edge BPDU Guard Default            is enabled
Portfast Edge BPDU Filter Default           is disabled[1]
Loopguard Default                           is disabled
PVST Simulation Default                     is enabled but inactive in rapid-pvst mode
Bridge Assurance                            is enabled
UplinkFast                                  is disabled
BackboneFast                                is disabled
```

[1] BPDU Filtering is **disabled**. This is the desired result.

EtherChannel (Link Aggregation)

When bundling multiple physical links into a single logical one, use Link Aggregation Control Protocol (LACP) or Port Aggregation Protocol (PAgP) to negotiate the link aggregation (LAG).

LACP and PAgP use a negotiation and keep-alive mechanism. The negotiation mechanism is used to determine which ports are eligible and safe to become aggregated in the channel. The keepalive mechanism continuously monitors whether the momentary conditions allow the member ports to remain aggregated. The main advantage of using LACP or PAgP is the prevention of switching loops caused by EtherChannel misconfiguration or miswiring.

The **`channel-group`** ***`group_number`*** **`mode on`** command configures a physical interface to be unconditionally aggregated without the negotiation protocols. Using **`mode on`**, the switch will force all compatible ports configured with this mode to become active in the EtherChannel. **`Mode on`** should only be used in the rare situations where link aggregation is needed, and the switch in the External Layer 2 Switching Domain supports LAG but does not support LACP or PAgP. If **`mode on`** is used, it must be configured on both ends of the channel. **`Mode on`** should be used with significant caution and **only** when absolutely necessary.

Tech tip

As a recommended practice with EtherChannels, any additional interface-level configuration should be performed on the Port-channel interface. Once individual, physical interfaces have been bundled in an EtherChannel, additional configuration should be added to them only if the Port-channel does not support the commands that need to be added. As an example, UDLD, MACsec, and so one.

Step 1. EtherChannel configuration for fabric devices should be performed using the Cisco DNA Center EtherChannel workflow by navigating to:

Provision > **Fabric** > **Fabric Domain** > **Fabric Site** > **Fabric Infrastructure** > **Device** > **Port Channel** > **Create Port Channel**

Figure 3. **SD-Access Fabric Infrastructure – Create Port Channel**

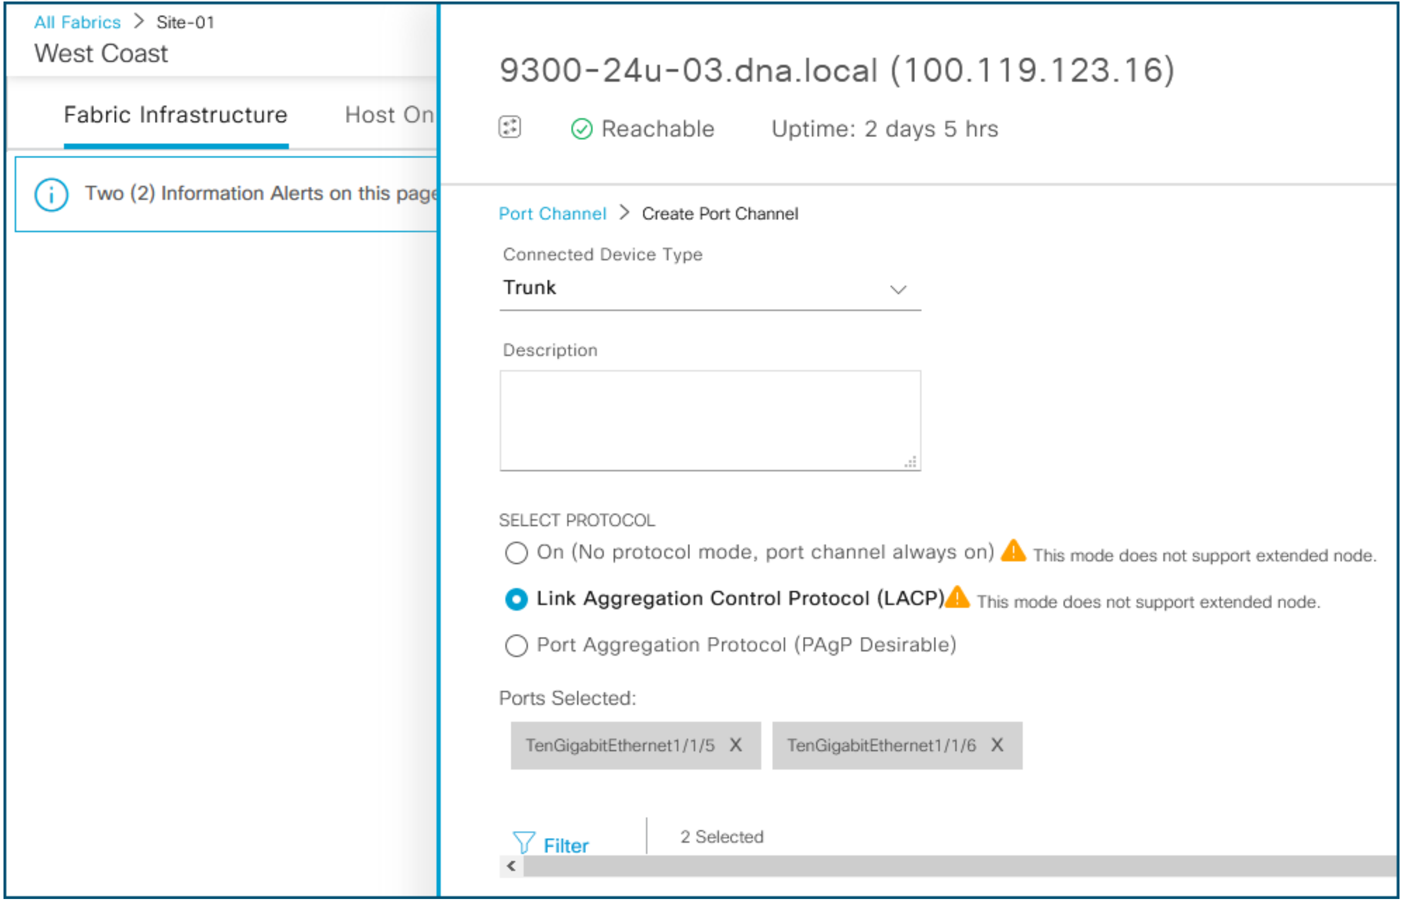

All Fabrics > Site-01

West Coast

Fabric Infrastructure   Host On

Two (2) Information Alerts on this page

9300-24u-03.dna.local (100.119.123.16)

Reachable   Uptime: 2 days 5 hrs

Port Channel > Create Port Channel

Connected Device Type

Trunk

Description

SELECT PROTOCOL

○ On (No protocol mode, port channel always on) ⚠ This mode does not support extended node.

● Link Aggregation Control Protocol (LACP) ⚠ This mode does not support extended node.

○ Port Aggregation Protocol (PAgP Desirable)

Ports Selected:

TenGigabitEthernet1/1/5 X   TenGigabitEthernet1/1/6 X

Filter   2 Selected

Step 2. Confirm the EtherChannel protocol mode.

```
Edge Node# show etherchannel summary
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator

        M - not in use, minimum links not met
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

        A - formed by Auto LAG


Number of channel-groups in use: 2
Number of aggregators:           2

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------
33      Po33(SU)         PAgP[1]      Te1/1/3(P)     Te1/1/4(P)
44      Po44(SU)         LACP[2]      Te1/1/5(P)     Te1/1/6(P)
```

[1] Port Channel 33 is configured using PAgP.

[2] Port Channel 44 is configured using LACP.

Error Disable Detect

Error Disable (errdisable) Detect is a feature that disables the port of a switch when certain error situations are detected. In the errdisable state, the port is effectively shut down, and no traffic is sent or received on the port.

The Error Disable Detect feature is enabled by default and should be left at the default setting.

Do not disable Error Disable Detect.

Step 1. Confirm Error Disable Detect is enabled.

```
Edge Node C9300# show errdisable detect[1]
ErrDisable Reason            Detection        Mode
-----------------            ---------        ----
arp-inspection               Enabled          port
bpduguard                    Enabled          port
channel-misconfig            Enabled          port
community-limit              Enabled          port
dhcp-rate-limit              Enabled          port
dtp-flap                     Enabled          port
evpn-mh-core-isolation       Enabled          port
gbic-invalid                 Enabled          port
iif-reg-failure              Enabled          port
inline-power                 Enabled          port
invalid-policy               Enabled          port
l2ptguard                    Enabled          port
link-flap                    Enabled          port
loopback                     Enabled          port
loopdetect                   Enabled          port
lsgroup                      Enabled          port
mac-limit                    Enabled          port
pagp-flap                    Enabled          port
port-mode-failure            Enabled          port
pppoe-ia-rate-limit          Enabled          port
psecure-violation            Enabled          port/vlan
security-violation           Enabled          port
sfp-config-mismatch          Enabled          port
sgacl_limitation:enforcem    Enabled          port
sgacl_limitation:multiple    Enabled          port
storm-control                Enabled          port
udld                         Enabled          port
psp                          Enabled          port
dual-active-recovery         Enabled          port
evc-lite input mapping fa    Enabled          port
vsl-and-non-vsl-port-pair    Enabled          port
fasthello-and-non-fasthel    Enabled          port
mvrp                         Enabled          port
mrp-miscabling               Enabled          port
```

[1] Error Disable Detect items may vary based on platform and code version. This example is from a Catalyst 9300 Series Switch running IOS XE 17.4.1.

Error Disable Recovery

To recover a port that is in an errdisable state, manual intervention is generally required. The administrator must access the switch and configure the specific port with **`shutdown`** followed by the **`no shutdown`** command.

Error Disable Recovery is a feature that automatically brings the port out of the errdisable state without manual intervention. The port is reenabled after a configured timer interval, and the switch retries the operation to bring the port up. Error Disable Recovery is disabled by default.

When Error Disable Recovery is enabled, the default timer interval is 300 seconds. The timer interval applies to all recovery causes and cannot be configured per cause.

As of this writing (June 2021 and Version 2.2.2.3), Cisco DNA Center provisions the following commands when a device is configured as an Edge Node. Do not modify the interval configuration.

```
switch(config)# errdisable recovery interval 300
switch(config)# errdisable recovery cause all
```

Control Plane Policing (CoPP)

CoPP is an IOS-wide CPU protection mechanism. CoPP applies to all packets that are punted to the device CPU for handling. The default CoPP policy values vary by platform and is either 600 pps (packets per second) or 750 pps for broadcast traffic. This value is defined in the **`system-cpp-police-data`**, and these values should not be changed.

Tech tip

The Catalyst 9300 and 9500 Series Switches use a different verification command than the Catalyst 9400 and 9500H Series Switches. Please note the use of **`fed switch active`** and **`fed active`** based on platform.

The following outputs were issued on devices running IOS XE 17.4.1. Policer indexes may vary by release. The intention of the outputs below is to show the default rate for BROADCAST queue for a given platform.

Step 1. Confirm CoPP broadcast traffic policy on Catalyst 3650, 3850, 4500, 9200/L, 9300/L Series Switches.

```
C9300# show platform hardware fed switch active qos queue stats internal cpu policer

                         CPU Queue Statistics
=========================================================================================
                                          (default) (set)     Queue        Queue
QId PlcIdx Queue Name               Enabled   Rate     Rate    Drop(Bytes)  Drop(Frames)
-----------------------------------------------------------------------------------
0    11     DOT1X Auth                 Yes     1000     1000     0           0
1    1      L2 Control                 Yes     2000     2000     0           0
2    14     Forus traffic              Yes     4000     4000     0           0
3    0      ICMP GEN                   Yes     600      600      0           0
4    2      Routing Control            Yes     5400     5400     0           0
5    14     Forus Address resolution   Yes     4000     4000     0           0
6    0      ICMP Redirect              Yes     600      600      0           0
7    16     Inter FED Traffic          Yes     2000     2000     0           0
8    4      L2 LVX Cont Pack           Yes     1000     1000     0           0
9    19     EWLC Control               Yes     13000    13000    0           0
10   16     EWLC Data                  Yes     2000     2000     0           0
11   13     L2 LVX Data Pack           Yes     1000     1000     0           0
12   0      BROADCAST                  Yes     600      600      0           0
```

Step 2. Confirm CoPP broadcast traffic policy on Catalyst 9400 Series Switches.

```
C9400# show platform hardware fed active qos queue stats internal cpu policer

                         CPU Queue Statistics
=========================================================================================
                                          (default)(set)     Queue        Queue
QId PlcIdx Queue Name              Enabled  Rate      Rate    Drop(Bytes) Drop(Frames)
-----------------------------------------------------------------------------------
0    11     DOT1X Auth                Yes    1000      1000     0           0
1    1      L2 Control                Yes    2000      2000     0           0
2    14     Forus traffic             Yes    4000      4000     0           0
3    0      ICMP GEN                  Yes    600       600      0           0
4    2      Routing Control           Yes    5400      5400     0           0
5    14     Forus Address resolution  Yes    4000      4000     0           0
6    0      ICMP Redirect             Yes    600       600      0           0
7    16     Inter FED Traffic         Yes    2000      2000     0           0
8    4      L2 LVX Cont Pack          Yes    1000      1000     0           0
9    19     EWLC Control              Yes    13000     13000    0           0
10   16     EWLC Data                 Yes    2000      2000     0           0
11   13     L2 LVX Data Pack          Yes    1000      1000     0           0
12   0      BROADCAST                 Yes    600       600      0           0
```

Step 3. Confirm CoPP broadcast traffic policy on Catalyst 9500 Series Switches.

```
C9500# show platform hardware fed switch active qos queue stats internal cpu policer

                         CPU Queue Statistics
=========================================================================================
                                          (default)(set)     Queue        Queue
QId PlcIdx Queue Name              Enabled  Rate      Rate    Drop(Bytes) Drop(Frames)
------------------------------------------------------------------------------------
0    11     DOT1X Auth                Yes    1000      1000     0           0
1    1      L2 Control                Yes    2000      2000     0           0
2    14     Forus traffic             Yes    4000      4000     0           0
3    0      ICMP GEN                  Yes    600       600      0           0
4    2      Routing Control           Yes    5400      5400     0           0
5    14     Forus Address resolution  Yes    4000      4000     0           0
6    0      ICMP Redirect             Yes    600       600      0           0
7    16     Inter FED Traffic         Yes    2000      2000     0           0
8    4      L2 LVX Cont Pack          Yes    1000      1000     0           0
9    19     EWLC Control              Yes    13000     13000    0           0
10   16     EWLC Data                 Yes    2000      2000     0           0
11   13     L2 LVX Data Pack          Yes    1000      1000     0           0
12   0      BROADCAST                 Yes    600       600      0           0
```

Step 4. Confirm CoPP broadcast traffic policy on Catalyst 9500H Series Switches.

```
C9500H# show platform hardware fed active qos queue stats internal cpu policer

                         CPU Queue Statistics
=========================================================================================
                                          (default)(set)     Queue        Queue
QId PlcIdx Queue Name              Enabled  Rate      Rate    Drop(Bytes) Drop(Frames)
------------------------------------------------------------------------------------
0    11     DOT1X Auth                Yes      1000    1000     0           0
1    1      L2 Control                Yes      2000    2000     0           0
2    14     Forus traffic             Yes      4000    4000     0           0
3    0      ICMP GEN                  Yes      750     750      0           0
```

```
4    2     Routing Control          Yes    5500   5500    0          0
5    14    Forus Address resolution Yes    4000   4000    0          0
6    0     ICMP Redirect            Yes    750    750     0          0
7    16    Inter FED Traffic        Yes    2000   2000    0          0
8    4     L2 LVX Cont Pack         Yes    1000   1000    0          0
9    19    EWLC Control             Yes    13000  13000   0          0
10   16    EWLC Data                Yes    2000   2000    0          0
11   13    L2 LVX Data Pack         Yes    1000   1000    0          0
12   0     BROADCAST                Yes    750    750     0          0
```

Tech tip

The `show running-config` command does not display information about classes configured under `system-cpp policy` when they are left at default values.

Use `show policy-map system-cpp-policy` or `show policy-map control-plane` commands instead.

General Recommendations Summary

The following table outlines a summary of this section.

Table 1. General Recommendations Summary

Parameter	Recommended Value
VTP Mode	OFF
Spanning Tree Mode	Rapid PVST+
Spanning Tree Root Bridge	Configure the Edge Node as Root Bridge.
BPDU Filter	Disabled (Do not enable BPDU Filtering)
EtherChannel	Use LACP or PAgP wherever possible.
Error Disable Detect	Enabled
Error Disable Recovery	This is enabled with a 300-second timer interval by Cisco DNA Center on Edge Nodes.
Control Plane Policing (CoPP)	Use default settings for broadcast traffic.

Recommend Practices – Edge Node’s Connecting Interface is a Trunk Port

This chapter is organized into the following sections:

Chapter	Section
Recommended Practices – Edge Node Trunk Port	Trunk Encapsulation Mode Trunk Port Interface Mode Dynamic Trunking Protocol Native VLAN Native VLAN Tagging VLAN Pruning Summary

If it is necessary to extend multiple VLANs between the Fabric Edge Node and the External Layer 2 Switching Domain, the connecting port on the Edge Node must be configured as a trunk port.

Trunk Encapsulation Mode

The interface on the boundary switch of the External Layer 2 Switching Domain that connects to a Fabric Edge Node must support the 802.1Q encapsulation. This is the only trunking encapsulation method available to switches supported as an SD-Access Edge Node. However, the External Layer 2 Switching Domain could be of a vintage that supports some other encapsulation such as Inter-Switch Link (ISL).

For legacy Cisco switches in the External Layer 2 Switching Domain that support multiple trunk encapsulations, the following interface level command enforces the 802.1Q encapsulation. This command needs to be configured only on the legacy switch trunk interface connecting to a Fabric Edge Node, not necessarily in the entire External Layer 2 Switching Domain.

```
External Switch(config-if)# switchport trunk encapsulation dot1q
```

Tech tip

Cisco IOS and IOS-XE based switches that only support the 802.1Q encapsulation do not even offer this command in their CLI, as there are no alternatives to choose from.

Trunk Port Interface Mode

The switchport must be hard-coded as a trunk interface.

Step 1. Configure the interface connected to the switch in the External Layer 2 Switching Domain as a trunk.

This configuration should be performed using the Cisco DNA Center Port Assignment workflow by navigating to the following:

Provision > **Fabric** > **Fabric Domain** > **Fabric Site** > **Host Onboarding** > **Port Assignment**

Figure 4. SD-Access Host Onboarding Port Assignment – Trunk

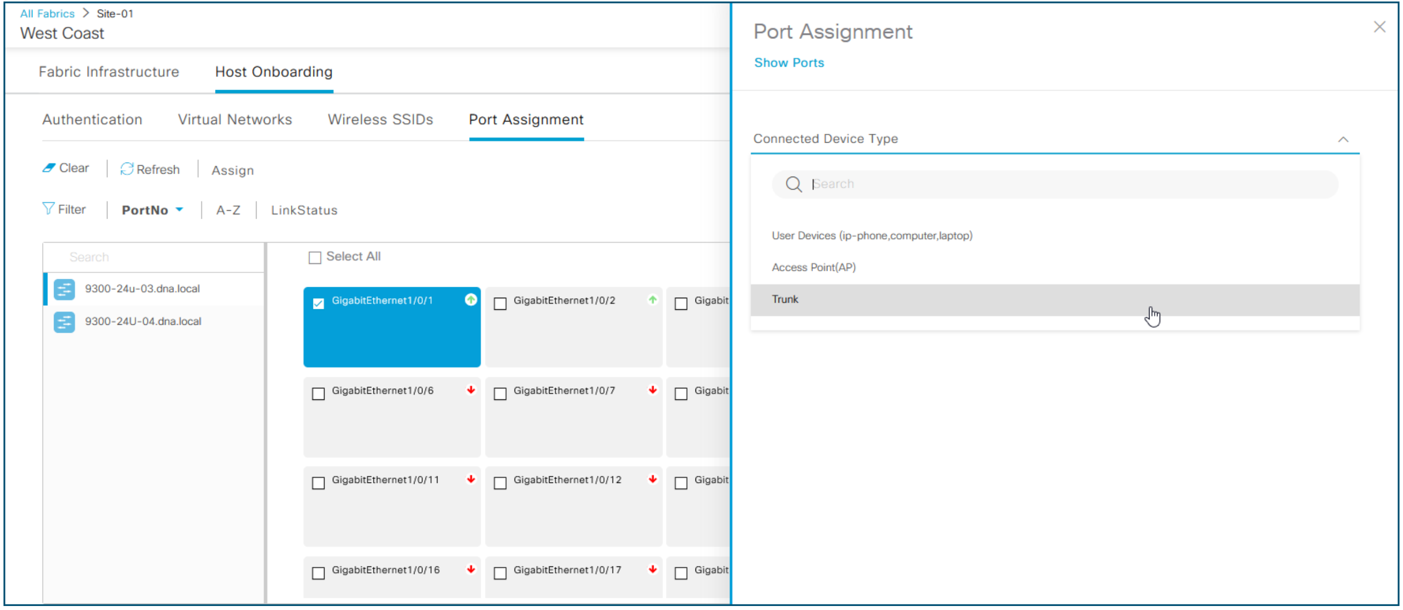

Step 2. Confirm the port is operating as a trunk.

```
Edge Node# show interface trunk

Port         Mode              Encapsulation  Status       Native vlan
Gi1/0/1      on[1]             802.1q[2]      trunking     1
```

[1] The interface is operating as a trunk.

[2] The trunk is using 802.1q encapsulation. This is the desired result.

Dynamic Trunking Protocol (DTP)

Dynamic Trunking Protocol is a feature that is used to negotiate the formation of a trunk between two DTP-capable devices without manual configuration. DTP automatically negotiates whether the port should be put into access mode or trunk mode along with what trunking protocol should be used. DTP is enabled by default on a switchport.

DTP should be disabled on the interface connecting to the External Layer 2 Switching Domain. The command to disable DTP is only necessary for trunk ports, as static access ports do not send DTP packets.

Step 1. Disable DTP on interfaces connecting to the external domain.

```
Edge Node(config-if)# switchport nonegotiate
```

Step 2. Confirm DTP is disabled on the trunk port.

```
Edge Node# show dtp interface TenGigabitEthernet1/1/1
DTP information for TenGigabitEthernet1/1/1:
  TOS/TAS/TNS:                              TRUNK/ON/TRUNK[1]
  TOT/TAT/TNT:                              802.1Q/802.1Q/802.1Q
  Neighbor address 1:                       F86BD92C3BA9
  Neighbor address 2:                       000000000000
  Hello timer expiration (sec/state):       19/RUNNING[2]
  Access timer expiration (sec/state):      never/STOPPED
  Negotiation timer expiration (sec/state): never/STOPPED
  Multidrop timer expiration (sec/state):   never/STOPPED
```

[1] The interface is operating as a trunk.

[2] DTP is **enabled** on the interface.

```
Edge Node# show dtp interface TenGigabitEthernet1/1/1
DTP information for TenGigabitEthernet1/1/1:
  TOS/TAS/TNS:                              TRUNK/NONEGOTIATE/TRUNK[1]
  TOT/TAT/TNT:                              802.1Q/802.1Q/802.1Q
  Neighbor address 1:                       F86BD92C3BA9
  Neighbor address 2:                       000000000000
  Hello timer expiration (sec/state):       never/STOPPED[2]
  Access timer expiration (sec/state):      never/STOPPED
  Negotiation timer expiration (sec/state): never/STOPPED
  Multidrop timer expiration (sec/state):   never/STOPPED
  FSM state:                                S6:TRUNK
  # times multi & trunk                     0
  Enabled:                                  yes
  In STP:                                   no
```

[1] The interface is configured as operating as a trunk.

[2] DTP is **disabled** on the interface. This is the desired result.

Native VLAN

The native VLAN is a concept retaken from IEEE 802.1Q standard that states that each port has an untagged Primary VLAN ID (native VLAN) and may have additional VLAN IDs (tagged VLANs). To be compatible with 802.1Q, each device must implement this concept. The result is the native VLAN on trunk ports. On Cisco switches, the native VLAN is set to VLAN 1 by default.

Some inter-switch protocols on trunks, such as Cisco Discovery Protocol (CDP), operate explicitly in VLAN 1. Their frames will still be tagged with VLAN ID 1 if the native VLAN on the trunk is changed to a different value. Frames of other inter-switch protocols on the trunk, such as LACP, are always untagged regardless of the native VLAN setting. These protocols are considered to always operate in the native VLAN on the trunk. PVST+ and Rapid-PVST+ are two special-case protocols operating across all existing VLANs allowed on the trunk, including VLAN 1.

Carrying user traffic alongside the crucial inter-switch protocols is both a security risk and network-stability risk. Therefore, the long-standing recommended security practice is to change the native VLAN to a VLAN different from VLAN 1. This different VLAN should be distinct from all user VLANs. This is sometimes referred to as setting the native VLAN to an *unused VLAN*.

Despite these long-standing recommendations, practical experience has shown that sometimes this approach can have unintended consequences with various switch-to-switch protocols. Therefore, this guide suggests a slightly modified approach that, if diligently adhered to, provides the same level of security while mitigating the possible contingencies:

- Leave the native VLAN on trunks set to VLAN 1.
- Do not use VLAN 1 to carry any user traffic.
- Do not prune VLAN 1 on trunks.

Native VLAN ID numbers must match on both ends of the trunk. If the native VLANs are mismatched, PVST+ and Rapid-PVST+ may put the two mismatched VLANs into the blocking state.

```
*May 14 08:46:54.548: %SPANTREE-2-BLOCK_PVID_LOCAL: Blocking TenGigabitEthernet1/1/1 on
VLAN0888. Inconsistent local vlan.

*May 14 08:46:55.359: %CDP-4-NATIVE_VLAN_MISMATCH: Native VLAN mismatch discovered on
TenGigabitEthernet1/1/1 (888), with External_Switch TenGigabitEthernet1/1/1 (1).

Edge Node# show spanning-tree vlan 888
```

```
VLAN0888
Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Te1/1/1             Desg BKN[1] *100     128.1    P2p *PVID_Inc
```

[1] The interface is in the Blocking state.

If the connected External Layer 2 Switching Domain uses a different native VLAN on the trunk, the Edge Node's interface must be configured to match this. Only change the native VLAN on the Edge Node if required due to the configuration present in the external domain.

In the example below, the External Layer 2 Switching Domain uses VLAN 99 as the native VLAN. The Edge Node is configured with this VLAN ID, and the connecting interface to the external domain is configured with the same native VLAN.

Step 1. If the External Layer 2 Switching Domain currently has a defined native VLAN, define this same VLAN ID on the Edge Node. Otherwise, skip these steps and do not change the native VLAN.

```
Edge Node(config)# vlan 99
Edge Node(config-vlan)# name NATIVE_VLAN_99
Edge Node(config-vlan)# exit
```

Step 2. Define the native VLAN on the trunk port connecting to the External Layer 2 Switching Domain.

```
Edge Node(config-if)# switchport trunk native vlan 99
```

Step 3. Confirm the native VLAN on the trunk port.

```
Edge Node# show interfaces trunk

Port        Mode             Encapsulation  Status        Native vlan
Te1/1/1     on               802.1q         trunking      99[1]
```

[1] The interface is configured to use VLAN 99 as the native VLAN.

Tech tip

The native VLAN is locally significant to the interconnection between the switches. If per-port *Native VLAN Tagging* is supported by the hardware platform, every trunk interface of the same switch can be configured with a different native VLAN. This scenario is applicable to a multi-tenant environment that has multiple External Layer 2 Switching Domains connected to an Edge Node.

In a migration scenario, where the SD-Access Fabric and External Layer 2 Switching Domains are under a single administrative entity, a single native VLAN is generally used.

Native VLAN Tagging

By definition, the native VLAN is not tagged on the trunk port. The Double Tagging (Double Encapsulation) Attack can leverage this operation to perform VLAN hopping. While this does require the attacker to be in the same VLAN as the native VLAN, there may be deployment scenarios where user or management traffic is placed in the native VLAN despite recommendations otherwise.

To mitigate this VLAN hopping attack, the trunk port can be configured to tag all traffic in all VLANs including the VLAN configured as native (effectively suppressing the concept of a native VLAN). This feature is somewhat imprecisely called the *Native VLAN Tagging*.

Tech tip

> On Cisco Catalyst switches, Native VLAN Tagging causes the switch to tag all VLANs on the trunk port. It also changes the way the switch processes untagged frames. Untagged frames arriving on the trunk port are dropped and not further forwarded.
>
> Native VLAN Tagging is configured globally and applies to all interfaces operating as trunk ports.

If the External Layer 2 Switching Domain connected to the Edge Node is configured to tag the native VLAN, then the Edge Node must be configured to match this behavior. Only tag the native VLAN on the Edge Node if required due to the configuration present in the external domain.

> **Tech tip**
>
> On Cisco Catalyst switches, when native VLAN tagging is enabled, untagged frames are silently dropped when received on a trunk port. Switches from other vendors may act differently. Please check their corresponding documentation for details.
>
> In a multi-tenant environment that is connected to the same Edge Node, one domain may tag the native VLAN and the other may not. To support this environment, ensure that the Edge Node platform can tag the native VLAN on a per-port basis.
>
> If the platform supports Native VLAN Tagging, it can be configured on a per-port basis after being enabled globally. Once enabled globally, Native VLAN Tagging is enabled on all interfaces operating as trunk ports by default.
>
> Support for per-port Native VLAN Tagging varies by hardware platform and software release. Please see the corresponding release notes for a given hardware and software combination for details.

Step 1. Confirm the platform's support of per-port Native VLAN Tagging.

These outputs are from a Catalyst 9200 and 9300 Series switches running IOS XE 17.4.1.

```
C9200# show vlan dot1q tag native
dot1q native vlan tagging is disabled[1]
C9200#
```

[1] Platforms that do not support per-port Native VLAN Tagging have this single line of output for this command.

```
C9300# show vlan dot1q tag native
dot1q native vlan tagging is disabled globally

Per Port Native Vlan Tagging State[1]
----------------------------------

Port         Operational        Native VLAN
                Mode            Tagging State
-------------------------------------------
```

[1] Platforms that support per-port Native VLAN Tagging have additional output in this command which show the per-port status.

Step 2. If the External Layer 2 Switching Domain is currently tagging the native VLAN, configure the Edge Node to tag of all frames the trunk port, including the native VLAN. Otherwise, skip these steps and do not configure Native VLAN Tagging.

```
Edge Node(config)# vlan dot1q tag native
```

Step 3. Confirm Native VLAN Tagging is enabled globally.

```
Edge Node# show vlan dot1q tag native
dot1q native vlan tagging is disabled globally[1]
```

[1] Native VLAN Tagging is globally **disabled**.

```
Edge Node# show vlan dot1q tag native
```

```
dot1q native vlan tagging is enabled globally[1]
```

[1] Native VLAN Tagging is globally **enabled**. For this specific step, this is the desired result.

Step 4. If required in a multi-tenant environment, Native VLAN Tagging can be enabled on a per-port basis with the following command.

```
Edge Node(config-if)# switchport trunk native vlan tag
```

Step 5. Confirm Native VLAN Tagging is enabled on the applicable ports.

```
Edge Node# show vlan dot1q tag native
dot1q native vlan tagging is enabled globally[1]

Per Port Native Vlan Tagging State
----------------------------------

Port         Operational        Native VLAN
                Mode            Tagging State
-------------------------------------------

! Output omitted for brevity.
Po9          trunk                 enabled[2]
Po10         trunk                 disabled[3]
```

[1] Native VLAN Tagging is globally **enabled**.

[2] Native VLAN Tagging is **enabled** on interface Port-Channel 9.

[3] Native VLAN Tagging is **disabled** on interface Port-Channel 10.

Tech tip

Native VLAN Tagging applies to trunk ports only. Access ports are, in general, immune to VLAN hopping attacks through double tagging. In the current generation of switches supported as Edge Nodes, tagged frames received on an access port configured without a voice VLAN are discarded, effectively invalidating the Double Tagging Attack.

Access ports configured with a voice VLAN will accept tagged frames in the voice VLAN as well as tagged frames in the access VLAN (due to the possibility of the end host using the 802.1p CoS marking).

VLAN Pruning

VLAN Pruning (not to be confused with VTP pruning) refers to the ability to define a set of allowed VLANs on a trunk. Traffic in VLANs that are not allowed on a trunk will be blackholed on the port. No traffic in these VLANs will be sent out from the trunk, and traffic received in these VLANs will be silently discarded.

To avoid disruption of certain switch-to-switch protocols, VLAN 1 should not be pruned off the trunk.

If the port on the connecting switch in the External Layer 2 Switching Domain is configured with VLAN pruning, the allowed VLANs on the Edge Node interface must also match. The allowed VLANs on the interface should be consistent on both devices forming the trunk.

Step 1. Prune VLANs on the trunk port connected to the External Layer 2 Switching Domain.

```
Edge Node(config-if)# switchport trunk allowed vlan vlan-list
```

Step 2. Confirm VLAN pruning is enabled and VLANs are pruned.

```
Edge Node# show interfaces trunk

Port        Mode             Encapsulation  Status        Native vlan
Te1/1/1     on               802.1q         trunking      1
```

```
Port        Vlans allowed on trunk
Te1/1/1     1-4094[1]

Port        Vlans allowed and active in management domain
Te1/1/1     1,1021-1030

Port        Vlans in spanning tree forwarding state and not pruned
Te1/1/1     1,1021-1030[2]
```

[1] All VLANs are allowed on the trunk.

[2] VLANs 1 and 1021-1030 are configured on the switch and forwarding on the interface.

```
Edge Node# show interfaces trunk

Port         Mode              Encapsulation  Status        Native vlan
Te1/1/1      on                802.1q         trunking      1

Port        Vlans allowed on trunk
Te1/1/1     1,1021,1023,1025[1]

Port        Vlans allowed and active in management domain
Te1/1/1     1,1021,1023,1025

Port        Vlans in spanning tree forwarding state and not pruned
Te1/1/1     1,1021,1023,1025[2]
```

[1] VLANs 1, 1021, 1023, and 1025 are allowed on the trunk.

[2] VLANs 1, 1021, 1023, and 1025 are forwarding and not pruned. This is an example desired result is VLAN Pruning is used.

Trunk Port Recommendations Summary

The following table outlines a summary of this section.

Table 2. General Recommendations Summary

Parameter	Recommended Value
Trunk Encapsulation Mode	802.1Q
Trunk Port	Used when the External Layer 2 Switching Domain utilizes VLANs
DTP	Disabled
Native VLAN	**Keep user traffic off VLAN 1**. Do not change the Native VLAN unless required due to the configuration in the External Layer 2 Switching Domain.
Native VLAN Tagging	Do not tag the Native VLAN unless required due to the configuration in the External Layer 2 Switching Domain.
VLAN Pruning	Enabled Do not prune VLAN 1 unless required due to the configuration in the External Layer 2 Switching Domain.

Recommend Practices – Edge Node’s Connecting Interface is an Access Port

This chapter is organized into the following sections:

Chapter	Section
Recommended Practices – Edge Node Access Port	Access Port Interface Mode Summary

An access port should be used on the Edge Node to attach the External Layer 2 Switching Domain if the domain does not and will not utilize multiple VLANs.

Access Port Interface Mode

The switchport must be hard-coded as an access interface. This should be done using the Cisco DNA Center User Interface.

Step 1. In Cisco DNA Center, navigate to:

Provision > **Fabric** > **Fabric Domain** > **Fabric Site** > **Host Onboarding** > **Port Assignment**.

Step 2. Select the applicable Edge Node.

Step 3. Select the port(s), and click **Assign**.

Step 4. Select **User Devices (ip-phone,computer,laptop)** in the **Connected Device Type** drop down.

Step 5. Select the applicable Data VLAN / Pool from the **VLAN Name / IP Address Pool (Data)** drop down.

Step 6. (Optional) Select the **Scalable Group** from the drop down.

Step 7. (Optional) Select the applicable Voice VLAN / Pool from the **VLAN Name / IP Address Pool (Voice)** drop down.

Step 8. Select **No Authentication** from the **Authentication** drop down.

Step 9. Click **Update**, **Deploy**, and **Apply** to provision the configuration.

Figure 5. SD-Access Host Onboarding Port Assignment – Access

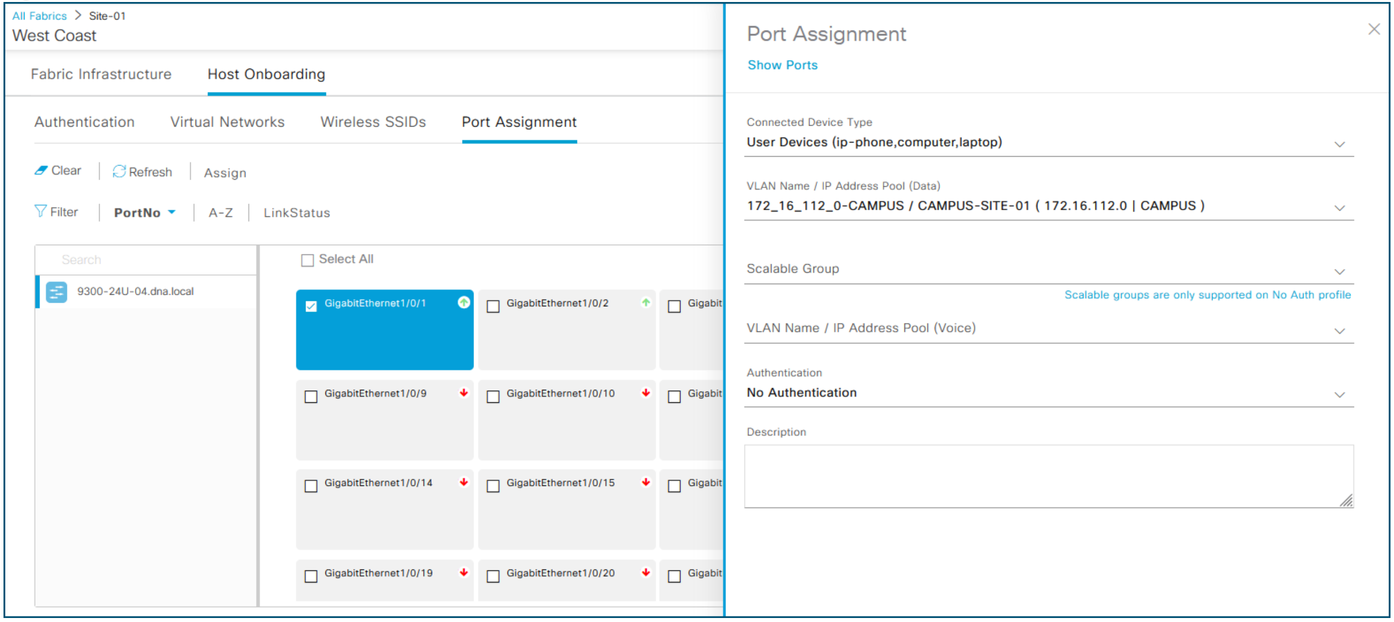

Port Assignment

Connected Device Type
User Devices (ip-phone,computer,laptop)

VLAN Name / IP Address Pool (Data)
172_16_112_0-CAMPUS / CAMPUS-SITE-01 (172.16.112.0 | CAMPUS)

Scalable Group

Scalable groups are only supported on No Auth profile

VLAN Name / IP Address Pool (Voice)

Authentication
No Authentication

Description

Step 10. Confirm the interface is configured as an access port with a VLAN assignment.

```
Edge Node# show interfaces GigabitEthernet 1/1/8 switchport
Name: Gi1/0/1
Switchport: Enabled[1]
Administrative Mode: static access[2]
Operational Mode: static access[3]
Administrative Trunking Encapsulation: dot1q
Negotiation of Trunking: Off
Access Mode VLAN: 1021 (172_16_112_0-CAMPUS)[4]
Trunking Native Mode VLAN: 1 (default)
Administrative Native VLAN tagging: enabled
! Output omitted for brevity
```

[1] The interface is configured as a switchport (operating in Layer 2 mode).

[2] The interface is configured as an access port.

[3] The interface is operating as an access port.

[4] The port is assigned to VLAN 1021.

Tech tip

On the Edge Node, STP BPDUs are generated by all ports operating at Layer 2. BPDUs are sent to a multicast destination MAC address. Switches that do not support STP, regardless of if they are managed or unmanaged, will flood BPDUs like a broadcast, as the multicast destination MAC address is never used as a source MAC address. It can therefore not be learned by the switch.

An *unmanaged* switch, as described in this Tech tip, does not have configuration capabilities and/or makes forwarding decisions based exclusively on MAC addresses and not VLAN IDs. It is also assumed to not support any STP version.

If the External Layer 2 Switching Domain uses *unmanaged* switches and is deployed in a looped topology, this flooding behavior may result in the Edge Node receiving its own BPDU. Cisco DNA Center provisions access ports with both PortFast and BPDU Guard. The reception of a BPDU will result in a port entering *errdisable* state, and the port cannot enter STP Blocking (Discarding) state.

If the External Layer 2 Switching Domain uses *unmanaged* switches and is deployed in a looped topology as shown in the example below, please consult with your Cisco Subject Matter Expert.

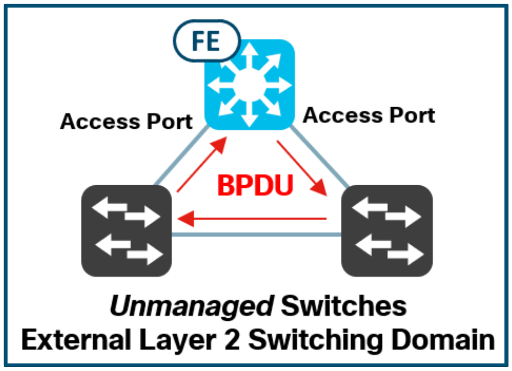

Access Port Recommendations Summary

The following table outlines a summary of this section.

Table 3. General Recommendations Summary

Parameter	Recommended Value
Access Port	Used when the External Layer 2 Switching Domain does not and will not utilize multiple VLANs.
Authentication Mode	No Authentication

Appendix A: Additional References

Cisco Networking Academy's Introduction to VLANs: https://www.ciscopress.com/articles/article.asp?p=2181837&seqNum=11

Cisco Press - Scaling Networks v6 Companion Guide – Spanning Tree Concepts: https://www.ciscopress.com/articles/article.asp?p=2832407&seqNum=4

Cisco Press - Scaling Networks v6 Companion Guide – Spanning Tree Configuration: https://www.ciscopress.com/articles/article.asp?p=2832407&seqNum=6

Configuring Spanning Tree: https://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst4500/12-2/15-02SG/configuration/guide/config/spantree.html

Configuring VLAN Trunks: https://www.cisco.com/en/US/docs/switches/lan/catalyst3850/software/release/3.2_0_se/multibook/configuration_guide/b_consolidated_config_guide_3850_chapter_0100010.html

Errdisable Port State Recovery on the Cisco IOS Platforms:https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/69980-errdisable-recovery.html

PVST Simulation on MST Switches: https://www.cisco.com/c/en/us/support/docs/lan-switching/multiple-instance-stp-mistp-8021s/116464-configure-pvst-00.html

Routing Between VLANs Overview: https://www.cisco.com/en/US/docs/ios/lanswitch/configuration/guide/lsw_rtng_vlan_ovw_ps6350_TSD_Products_Configuration_Guide_Chapter.html

Spanning Tree PortFast BPDU Guard Enhancement: https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/10586-65.html

STP Troubleshooting TechNotes: https://www.cisco.com/c/en/us/tech/lan-switching/spanning-tree-protocol/tsd-technology-support-troubleshooting-technotes-list.html

Understanding and Configuring Spanning Tree Protocol (STP) on Catalyst Switches: https://www.cisco.com/c/en/us/support/docs/lan-switching/spanning-tree-protocol/5234-5.html

Feedback

For comments and suggestions about this guide and related guides, join the discussion on Cisco Community at https://cs.co/en-cvds.

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.