

Use of ACI Operation with L2 Switches and Spanning Tree Link Types

Contents

[Introduction](#)

[Operation](#)

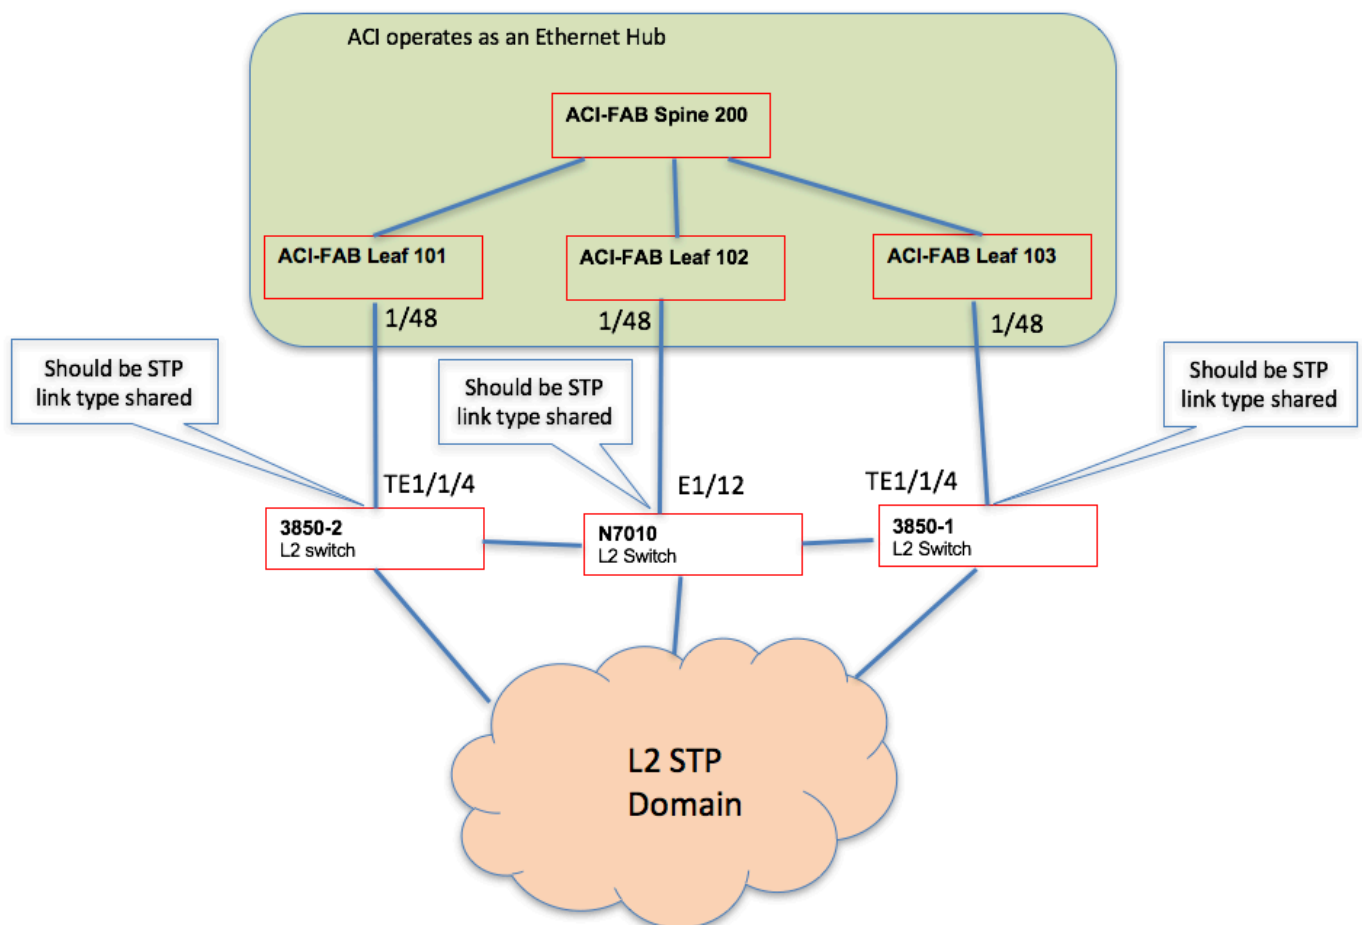
[Conclusion](#)

Introduction

This document describes Application Centric Infrastructure (ACI) link configuration for use with Layer 2 (L2) switches and spanning tree link types.

Operation

The L2 interface connected to ACI leafs can be configured with link-type shared, not point-to-point (P2P).



ACI Topology

Spanning Tree Protocol (STP) assumes that when the link type is P2P, there can be only one other STP

device connected on that interface. When a topology change occurs, the process is as follows:

1. Switch 3850-1 sends the proposal.
2. Switch 3850-2 sends the agreement (since the link type is P2P, it thinks there are no other switches that see the proposal).
3. Switch 3850-1 can transition immediately from blocking to forwarding upon receipt of the agreement.
4. Switch N7010 can wish to send a proposal, but 3850-1 already forwards it.

If the link types are configured as shared, this is the flow:

1. Switch 3850-1 sends the proposal.
2. Switch 3850-2 does not send the agreement (because the line type is shared and another switch wants to send the proposal).
3. Switch N7010 does not send the agreement or proposal.
4. Switch 3850-1 sends the proposal.
5. Switch 3850-1 transitions from blocking to learning, and then to forwarding, based on STP timers.

This leads to a slower convergence, but also assures that no loops are formed. In shared mode, the switches do not send the agreement because there can be more than one other switch on the segment, and they cannot all be in agreement. So in shared mode, the lack of agreement leads the proposal sender to wait longer before it transitions to forwarding.

Any L2 switch interfaces that are attached to ACI leaf can be configured as shown here:

```
<#root>
```

```
RTP-AGG1(config-if)#
```

```
spanning-tree link-type shared
```

The default behavior is STP link-type P2P. This is based on the link operating in full-duplex, which is the default behavior for switch-to-switch links. If the link comes up in half-duplex, the STP link type defaults to shared.

The IEEE 802.1d specification says:

17.12 RSTP and point-to-point links

The rapid transition of a Designated Port to Forwarding depends on the Port being directly connected to at most one other Bridge [It is an Edge Port (17.3, 17.19.17), or is attached to a point-to-point LAN, rather than a shared medium]. The `adminPointToPointMAC` and `operPointToPointMAC` parameters (6.4.3) provide management and signalling of the point-to-point status to RSTP state machines. A newly selected Root Port can transition to Forwarding rapidly, even if attached to shared media.

17.20.4 EdgeDelay

Returns the value of `MigrateTime` if `operPointToPointMAC` is `TRUE`, and the value of `MaxAge` otherwise.

17.21.9 recordAgreement()

If `rstpVersion` is `TRUE`, `operPointToPointMAC` (6.4.3) is `TRUE`, and the received Configuration Message has the Agreement flag set, the agreed flag is set and the proposing flag is cleared. Otherwise, the agreed flag is cleared.

ACI floods STP Bridge Protocol Data Units (BPDUs) to the VXLAN network identifier (VNID) assigned to the FD VLAN (VNID is assigned through the VLAN pool so encapsulation has to be part of same VLAN pool to be in part of same STP domain). So in effect, it operates as an Ethernet hub when it comes to Spanning Tree. When external L2 switches are connected to an ACI leaf, they default to full-duplex operation. In the STP world, this equates to a P2P link type.

In order to confirm if two Endpoint Groups (EPGs) are part of the same STP Domain, enter this command:

```
<#root>
module-1#
show system internal eltmc info vlan 49 | grep fabric_encap
    fabric_encap_type:          VXLAN   :::   fabric_encap:
11196
```

 **Note:** VLAN 49 is the FD/ EPG VLAN. All BDPUs can be flooded through the fabric in VNID 11196.

Conclusion

In ACI, it operates as an Ethernet hub. In P2P if a switch sees a proposal, it sends the agreement immediately and a transition from blocking to forwarding can occur. In shared mode, the interface on the switch cannot send the agreement immediately when it sees a proposal. This causes the sender of the proposal to wait longer (this depends on the timers configured) before it transitions from blocking for forwarding.

This applies to all versions of ACI firmware.