Connecting Switches to TRUE-IX

Many members choose to connect their TRUE-IX port to a layer 2 switch and then forward their peering traffic to a router virtual interface hosted elsewhere on their network. While connecting layer 2 switches to the TRUE-IX peering LAN is not actively discouraged, incorrect configuration can cause serious and unexpected connectivity problems.

The primary concern is to ensure that only traffic from the router subinterface is presented to the TRUE-IX port. TRUE-IX implements per port mac address counting: if more than 1 mac address is seen on any switch port at any time, that port will automatically be disabled for a cooling off period, and your connectivity to TRUE-IX will temporarily be lost.

This policy prevents two potential problems: firstly, it ensures that layer 2 traffic loops are prevented and secondly, it ensures that no other traffic escapes to the TRUE-IX peering LAN which shouldn't be seen there.

If you choose to connect your TRUE-IX port or ports to a switch, it is critically important to assign one unique vlan for each TRUE-IX connection. If you share an TRUE-IX facing VLAN between multiple TRUE-IX ports or share a TRUE-IX-facing VLAN with any other network, your connection may automatically be shut down due to the security mechanisms implemented by TRUE-IX.

It is also important to disable all switch-generated link-local traffic on your switch port. Typical link-local traffic will include spanning tree BPDUs, keepalive packets and discovery protocols. This traffic is particular problematic because these packets are typically forwarded on a port when the link is first brought up.

If multiple mac addresses are seen on any particular port, one of two things will happen. Either the switch port will shut down for a cooling-down period of 5 minutes, or else the equiment on the client-side of the switch port will have very poor quality connectivity, where arbitrary packets will appear to be dropped without any apparent reason.

Recommended Cisco Configuration

By default, all Cisco switches will broadcast CDP, Spanning Tree Protocol and keepalive packets on all ports. In addition, higher speed switches can default to using UDLD (unidirectional link detection). On IOS-enabled switches, these packets can be disabled using the following commands:

```
  interface GigabitEthernetx/x

    spanning-tree bpdufilter enable

    no keepalive

    no cdp enable

    udld port disable

    no lldp transmit
```

Some older Cisco switches do not support the "spanning-tree bpdufilter enable" command. On these units, it may be necessary to specify the following command:

```
spanning-tree bpduguard enable
```

Recommended Extreme Configuration

By default, Extreme switches will broadcast EDP on all ports. These packets can be disabled using the following commands:

```
  disable edp ports <portname>
```

If Spanning Tree Protocol is enabled on a particular port, it can be disabled using:

```
  disable stpd <stpd_name> ports <portname>
```

Recommended Brocade / Foundry Configuration

On trunk ports, Foundry switches will broadcast FDP (Foundry Discovery protocol) and by default on all ports, Foundry switches will broadcast Spanning Tree BPDUs. These packets can be disabled on a per-interface basis using the following command:

```
  no link-keepalive ethernet x/y

  interface ethernet x/y

    no fdp enable

    no spanning-tree
```