Cisco live!

6-9 March 2018 · Melbourne, Australia

Designing Layer 2 Networks – Avoiding Loops, Drops, Flooding

Roland Salinas - Technical Marketing Engineer

BRKCRS-2661

Cisco live!

Cisco Spark

Questions?

Use Cisco Spark to communicate with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install Spark or go directly to the space
4. Enter messages/questions in the space

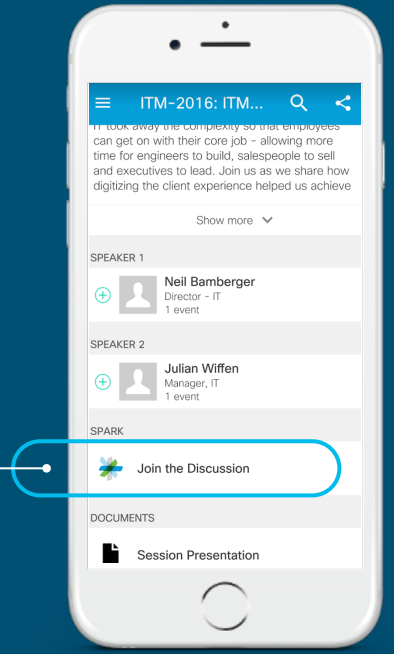

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Agenda

- L2 Network Design Challenges
- Layer 1 and Layer 2 Best Practices
- Spanning Tree Toolkit
- Integrated Security Toolkit
- Control Plane Protection
- Alternative Designs

Presentation Legend

Key Points

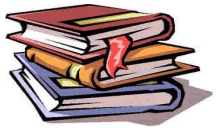

Reference Material

Standalone Multilayer Switch

Virtual Switching System

Layer 2 Link

Layer 3 Link

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

L2 Network Design Challenges

Traditional Multi-Layer Design – No L2 Loops

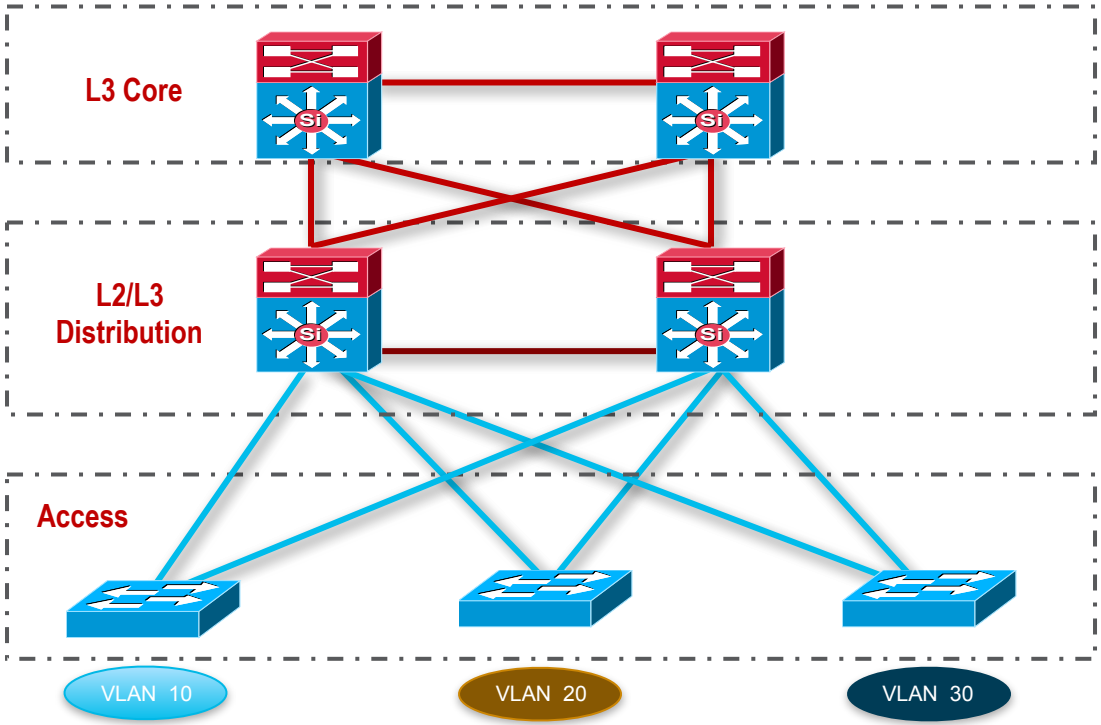

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Traditional Multi-Layer Design – No L2 Loops

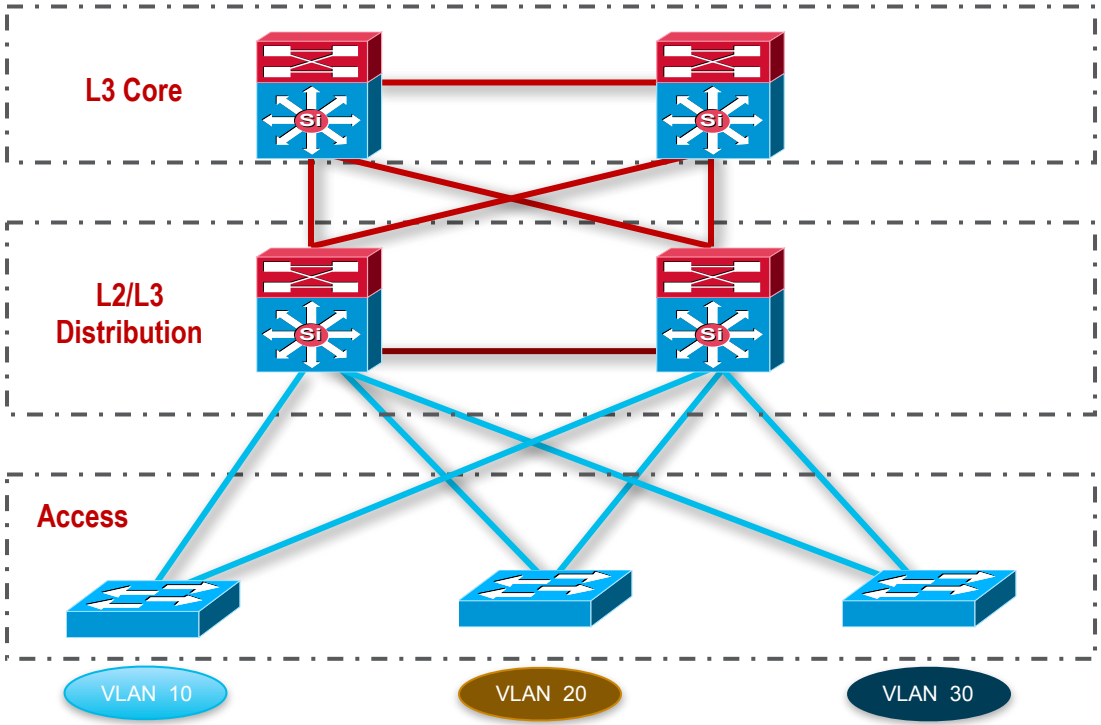

One switch per subnet, per VLAN

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Traditional Multi-Layer Design – No L2 Loops

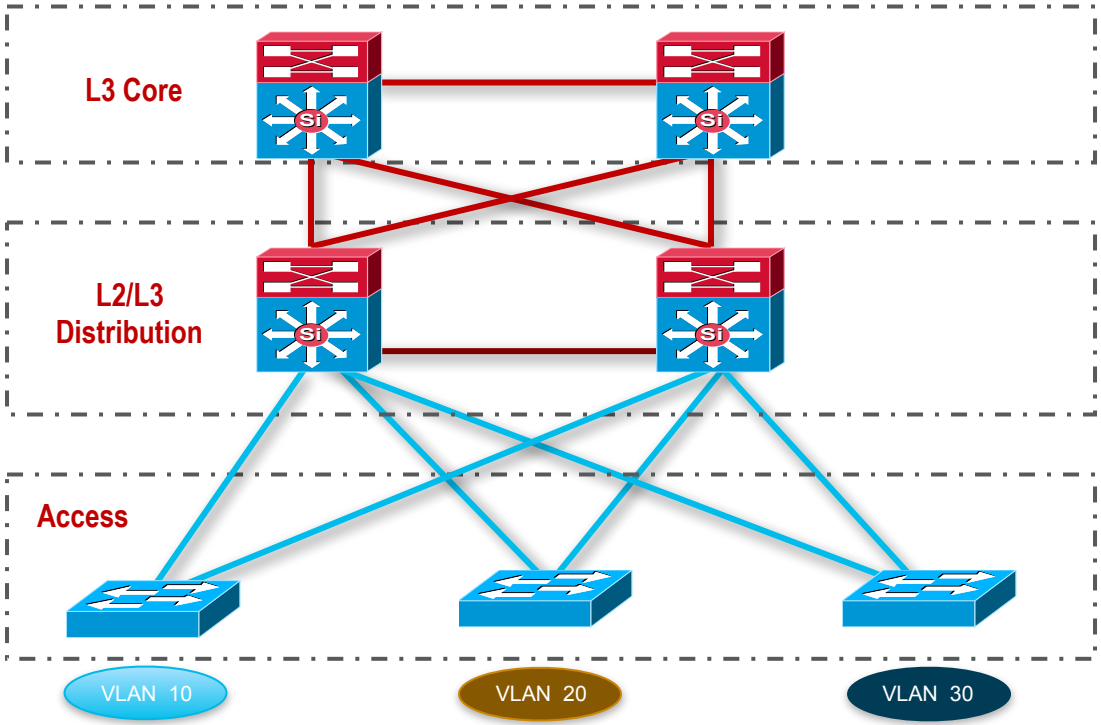

One switch per subnet, per VLAN

Simple design ✔

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Traditional Multi-Layer Design – No L2 Loops

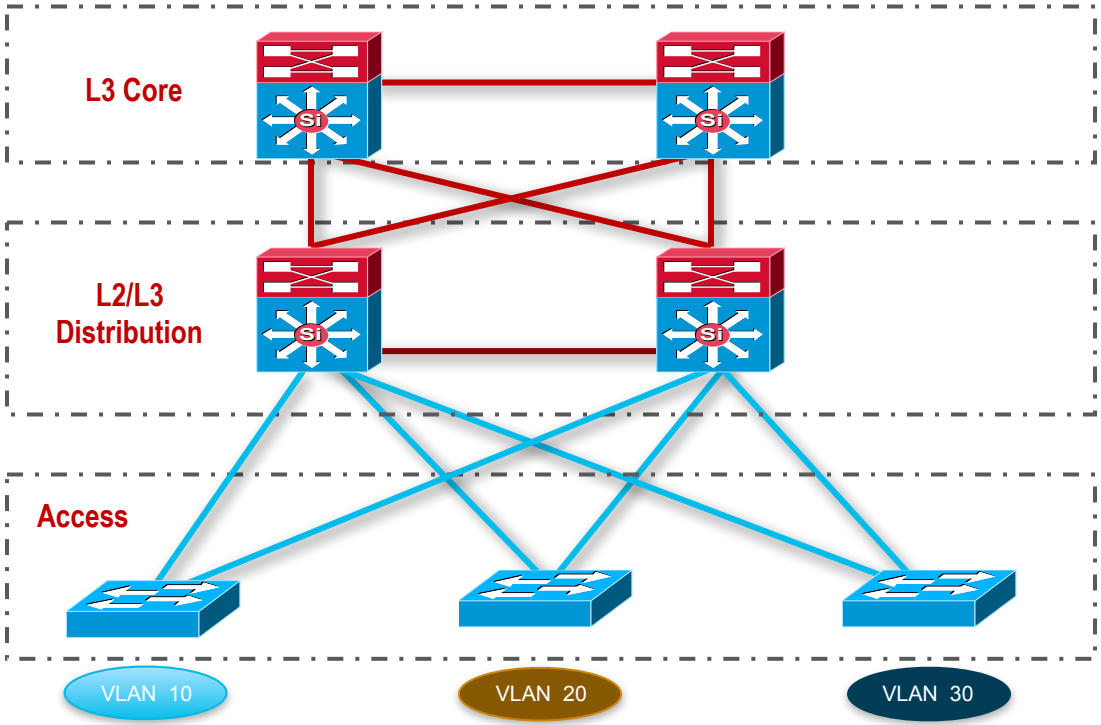

One switch per subnet, per VLAN

Simple design ✔

Limits L2 broadcast domain size to the size of the access switch

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Traditional Multi-Layer Design – With L2 Loops

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Traditional Multi-Layer Design – With L2 Loops

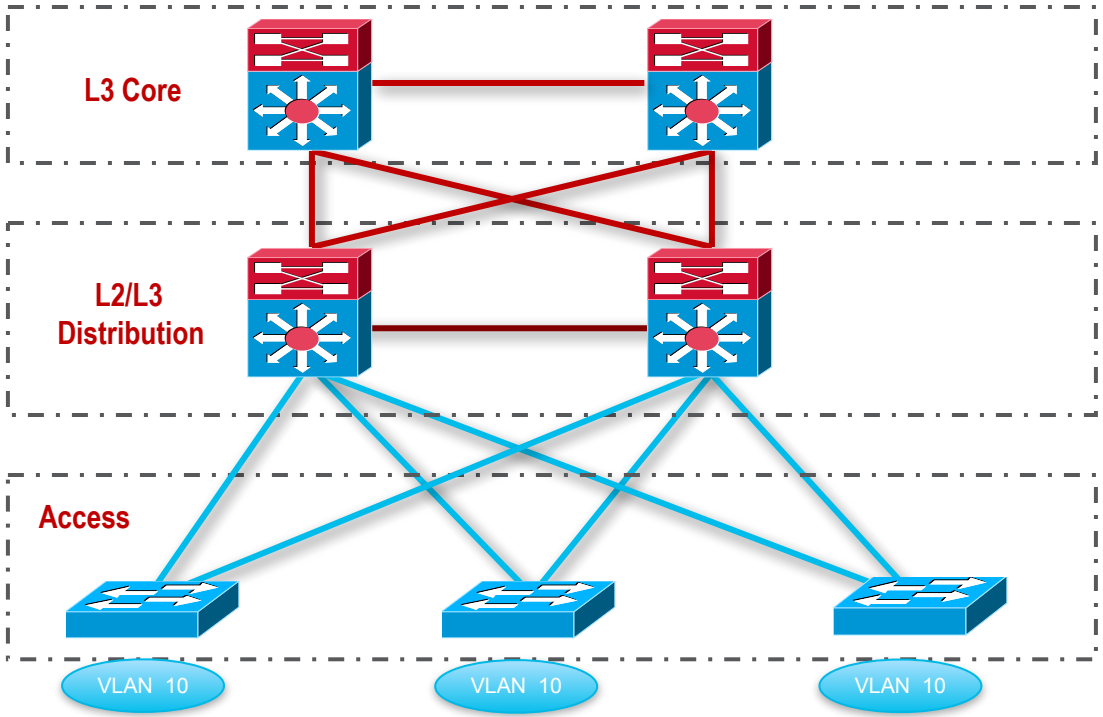

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Traditional Multi-Layer Design – With L2 Loops

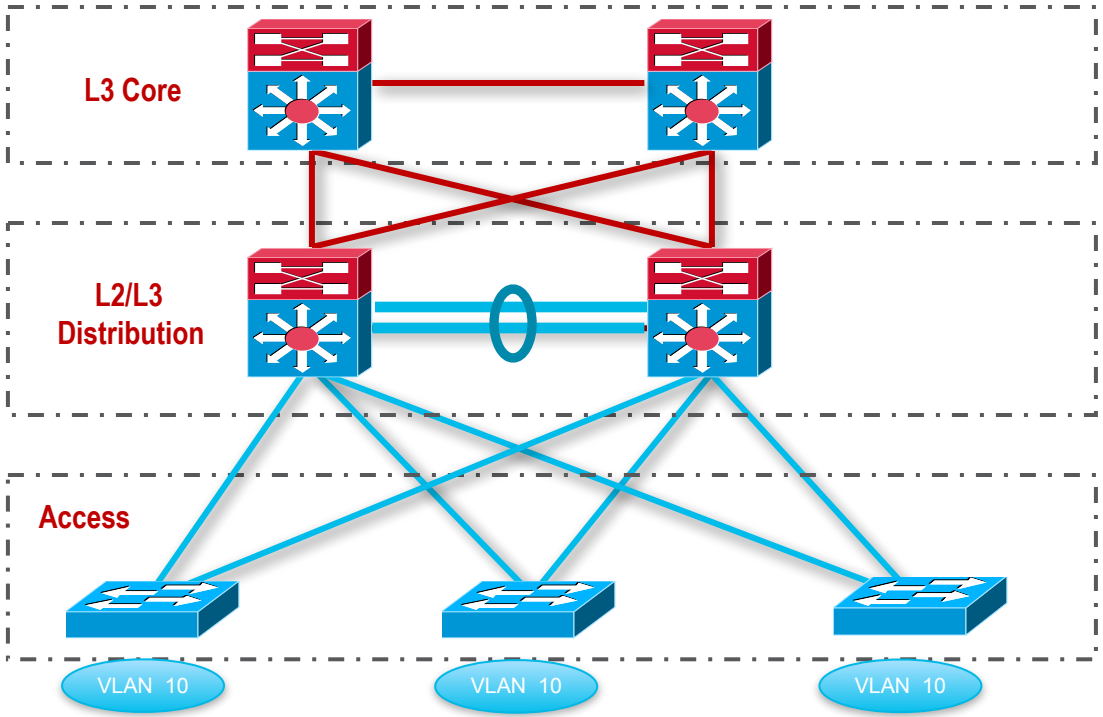

Extending the L2 domain beyond the single switch

Best practice says

- Distribution link must be an L2 link
- Redundant Links

Now we have the loop

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

L2 Loop – What's the Problem?

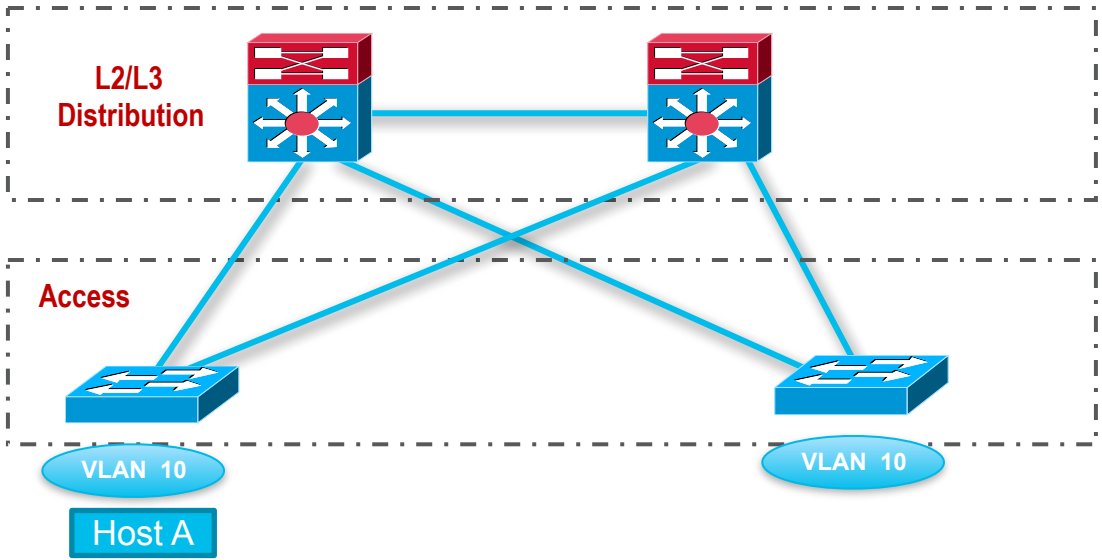

- Broadcast and multicast storm
- Source MAC address appear to be moving around as the MAC gets learned on different ports
- Frames are replicated repeatedly

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

L2 Loop – What's the Problem?

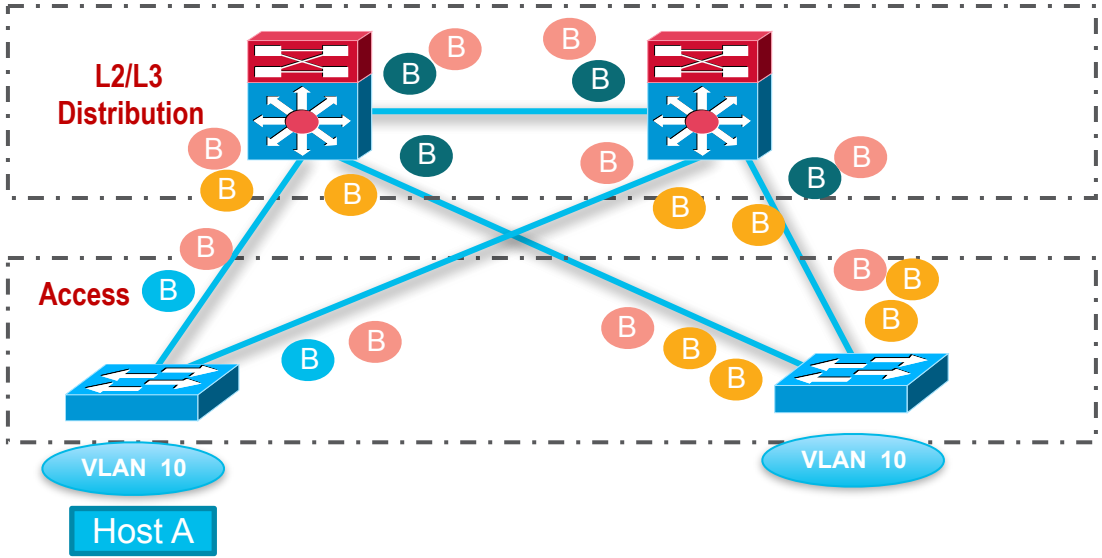

- Broadcast and multicast storm
- Source MAC address appear to be moving around as the MAC gets learned on different ports
- Frames are replicated repeatedly

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

L2 Loops and the Effect of Frame Replication on Interface Bandwidth

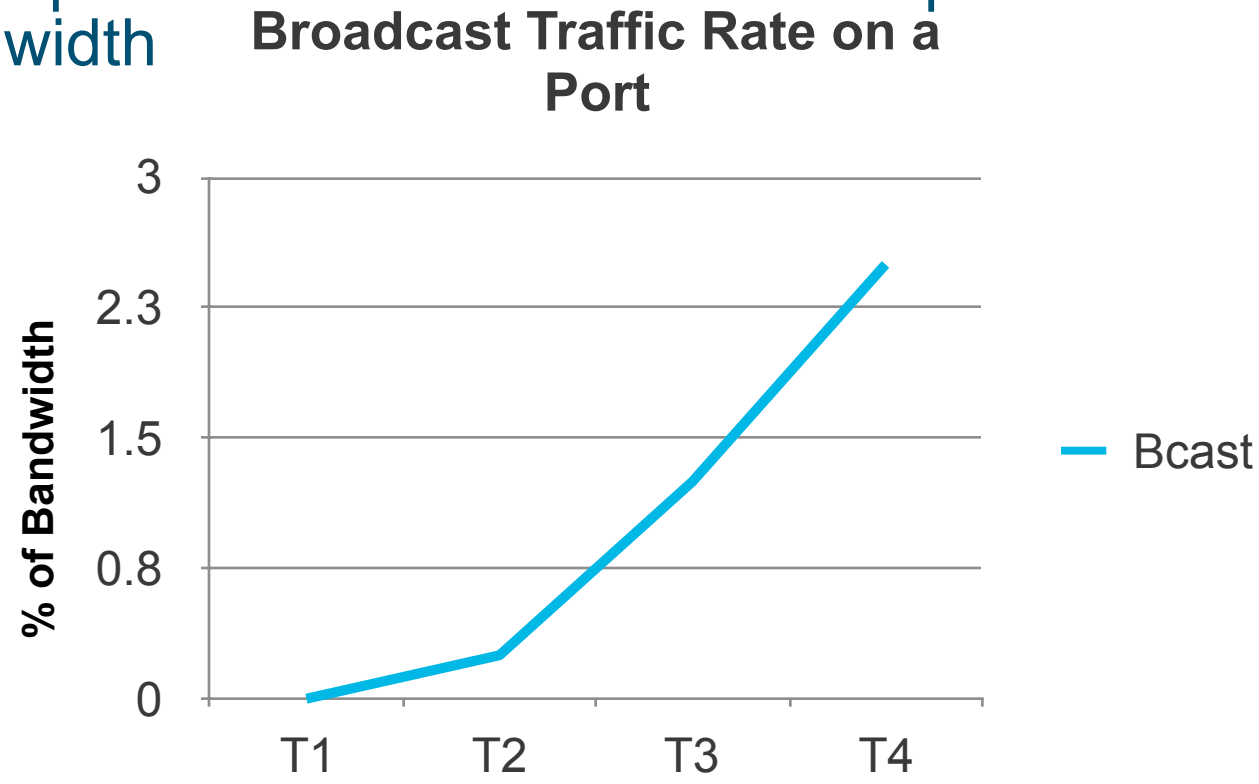

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Effects of a Broadcast Storm

Control Plane Saturation

- Bandwidth gets consumed due to frame replication
- CPU utilisation of network attached devices can start to reach high levels due to processing the broadcast traffic

- Network Instability
  - MAC addresses move from one port another
  - Traffic drops due to bandwidth oversubscription
  - This can occur with broadcast , multicast and unknown unicast traffic

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Causes of Loops and Broadcast Storms

- Malfunctioning NIC's
- OS /software problems
- Oversubscribed network devices
- Additional devices added unexpectedly (switches, APs)
- Human error
- Malicious activity
- Others…

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

CLOSED

Solution: Mitigate Vulnerabilities & Harden the Network

- Layer 1 Best Practices
- Layer 2 Best Practices
- Spanning Tree Protocol Best Practices
- Harden the control plane

“Make the network fail closed”

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Layer 1 and Layer 2 Best Practices

- Fibre links versus copper
- Point to point links
- Storm control
- Vlan Trunking Protocol (VTP)
- Dynamic Trunking Protocol (DTP)
- Unidirectional Link Detection (UDLD)

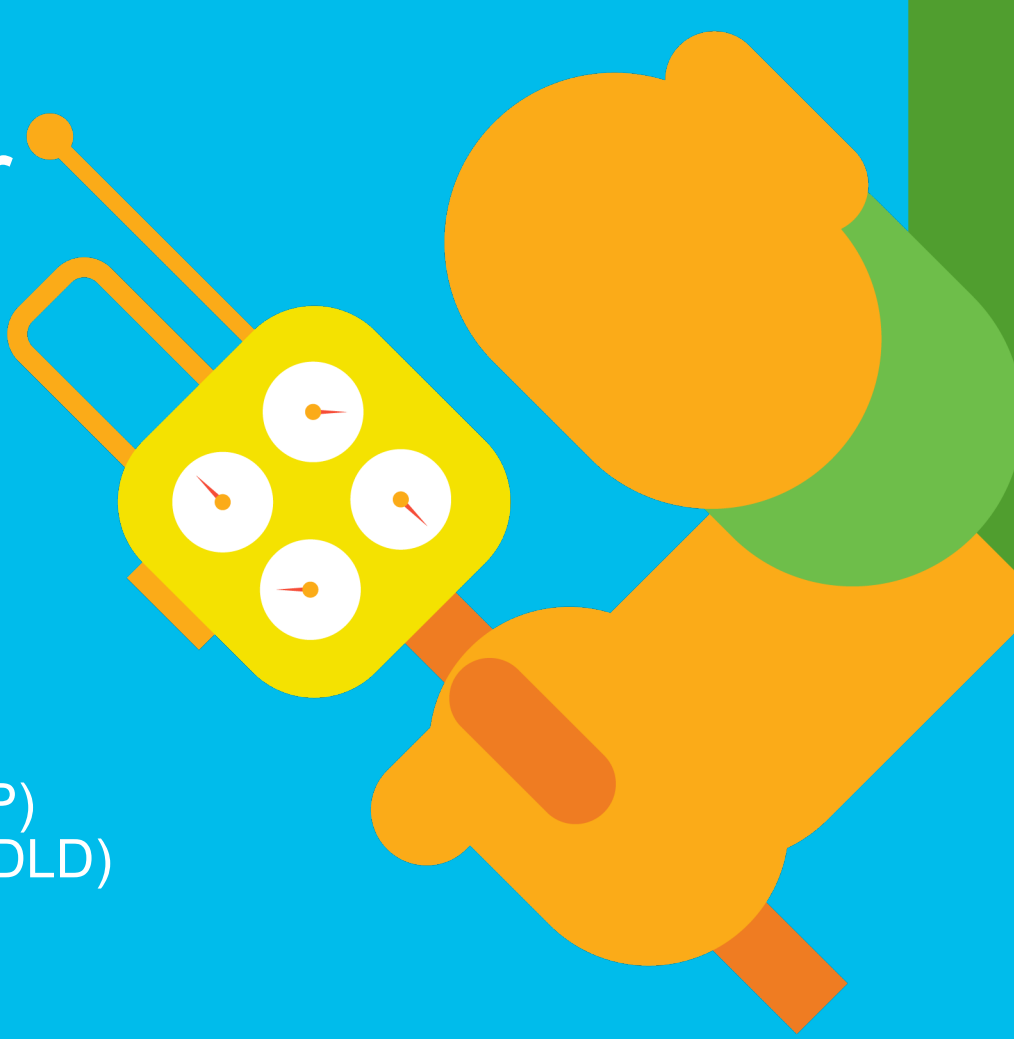

Reduce Link Down Detection Time

Use Fibre Links Versus Copper Links Whenever Possible

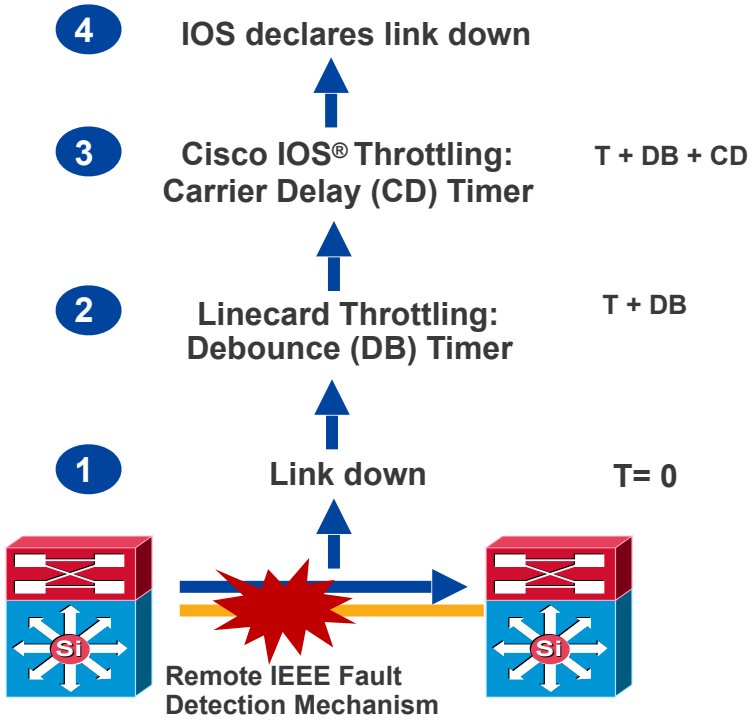

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Reduce Link Down Detection Time

Use Fibre Links Versus Copper Links Whenever Possible

Direct point-to-point fibre provides for fast failure detection

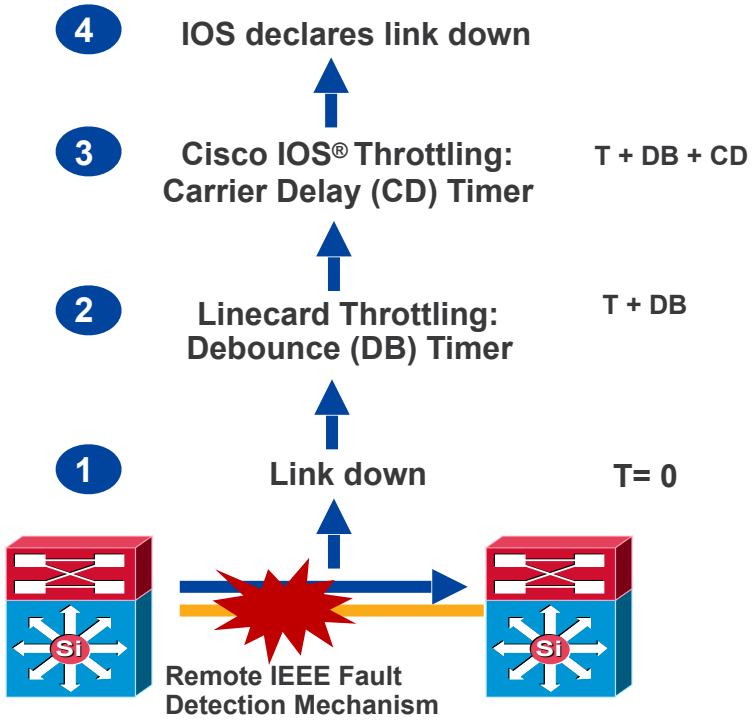

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Reduce Link Down Detection Time

Use Fibre Links Versus Copper Links Whenever Possible

Direct point-to-point fibre provides for fast failure detection

IEEE 802.3z and 802.3ae link negotiation define the use of remote fault indicator and link fault signalling mechanisms

- Bit D13 in the Fast Link Pulse (FLP) can be set to indicate a physical fault to the remote side

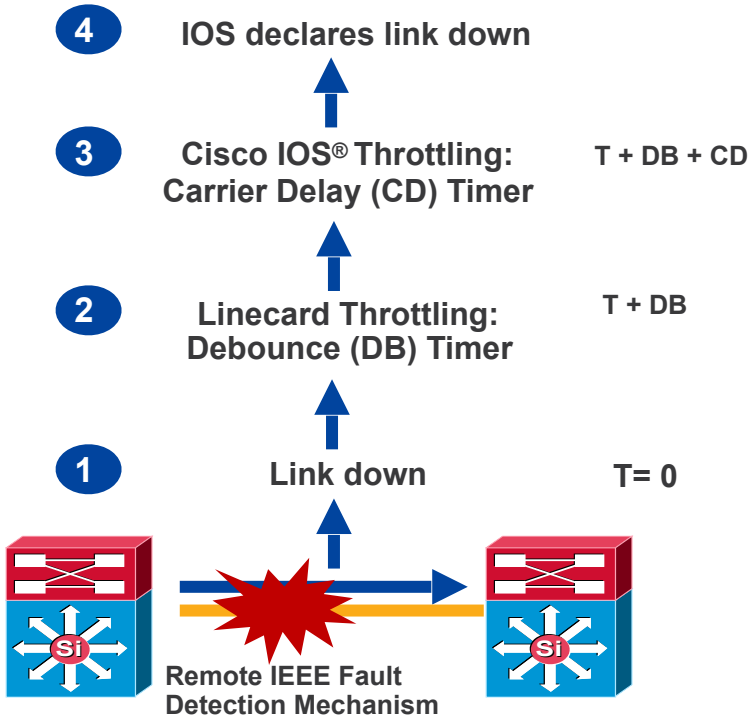

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Reduce Link Down Detection Time

Use Fibre Links Versus Copper Links Whenever Possible

Direct point-to-point fibre provides for fast failure detection

IEEE 802.3z and 802.3ae link negotiation define the use of remote fault indicator and link fault signalling mechanisms

- Bit D13 in the Fast Link Pulse (FLP) can be set to indicate a physical fault to the remote side

The default (disabled) debounce timer on GigE and 10GigE fibre linecard is 10 msec

- copper linecard is 300 msec

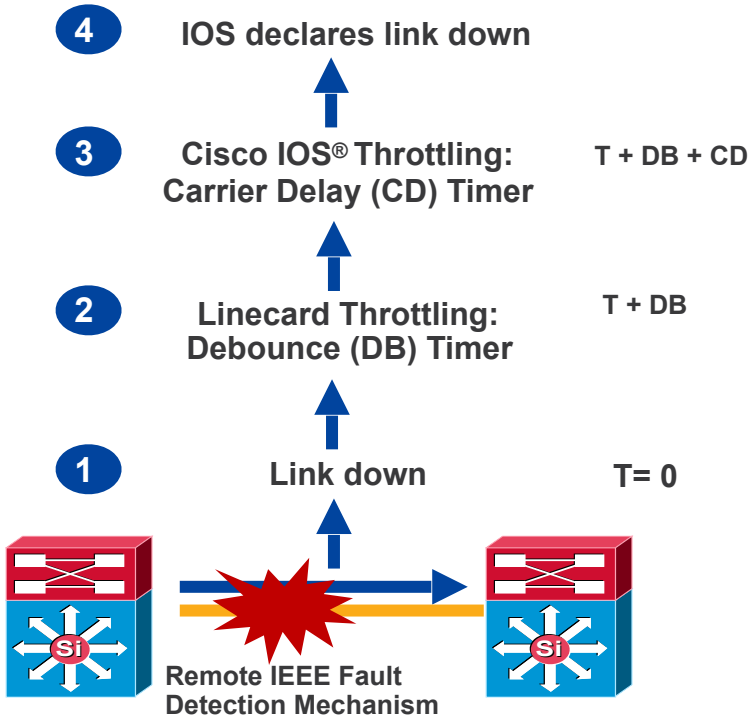

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Reduce Link Down Detection Time

Use Fibre Links Versus Copper Links Whenever Possible

Direct point-to-point fibre provides for fast failure detection

IEEE 802.3z and 802.3ae link negotiation define the use of remote fault indicator and link fault signalling mechanisms

- Bit D13 in the Fast Link Pulse (FLP) can be set to indicate a physical fault to the remote side

The default (disabled) debounce timer on GigE and 10GigE fibre linecard is 10 msec

- copper linecard is 300 msec

Carrier-delay

- Default 10 microseconds is adequate for most applications

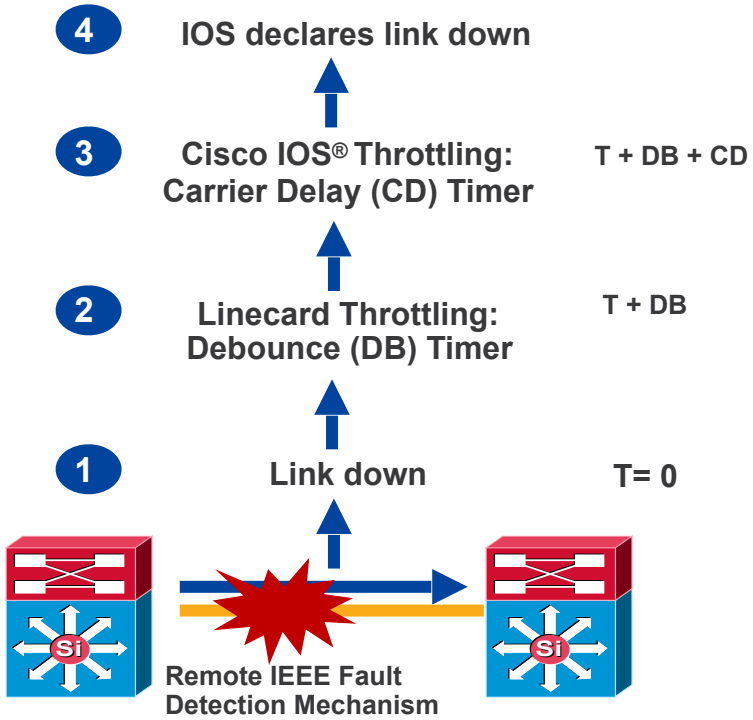

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Use Point-to-Point links Whenever Possible

- Eliminate or avoid at all cost intermediate L1 devices
- Can Mask link failures
- Use point-to-point link only

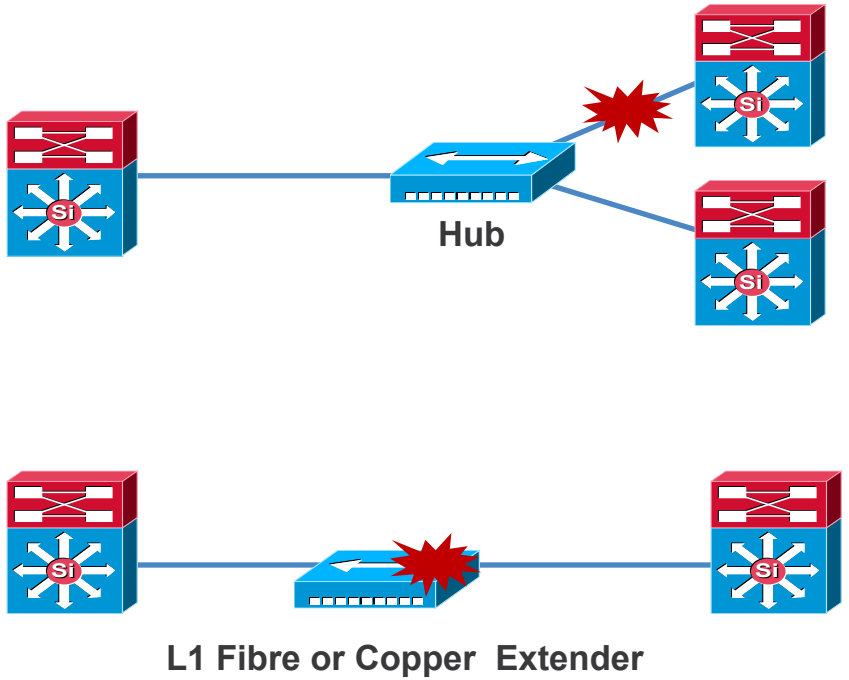

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Use Point-to-Point links Whenever Possible

- Eliminate or avoid at all cost intermediate L1 devices
- Can Mask link failures
- Use point-to-point link only

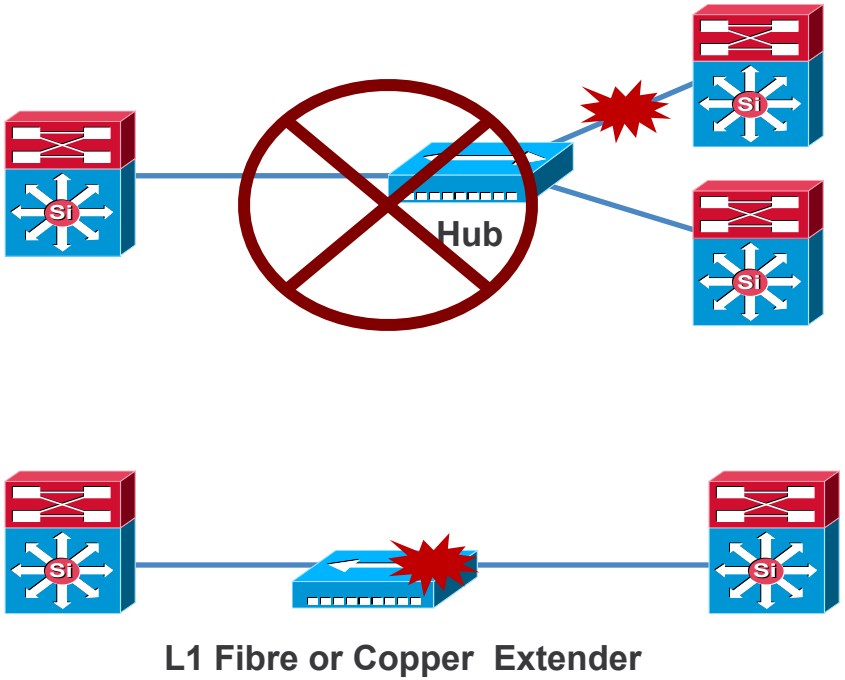

L1 Fibre or Copper Extender

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Use Point-to-Point links Whenever Possible

- Eliminate or avoid at all cost intermediate L1 devices
- Can Mask link failures
- Use point-to-point link only

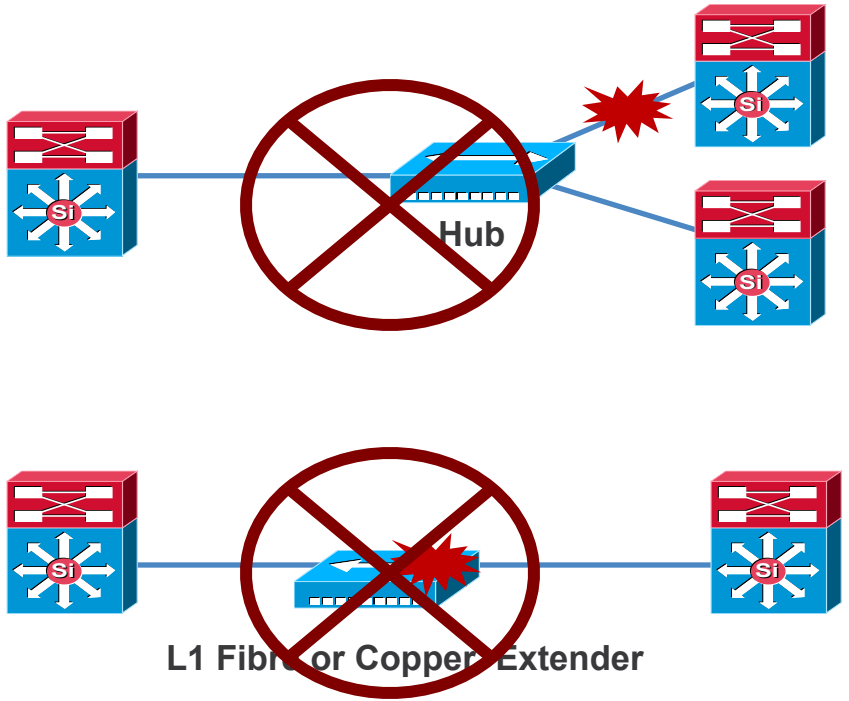

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Storm Control / Traffic Suppression

Use hardware-based rate-limiting to protect and harden the broadcast domain

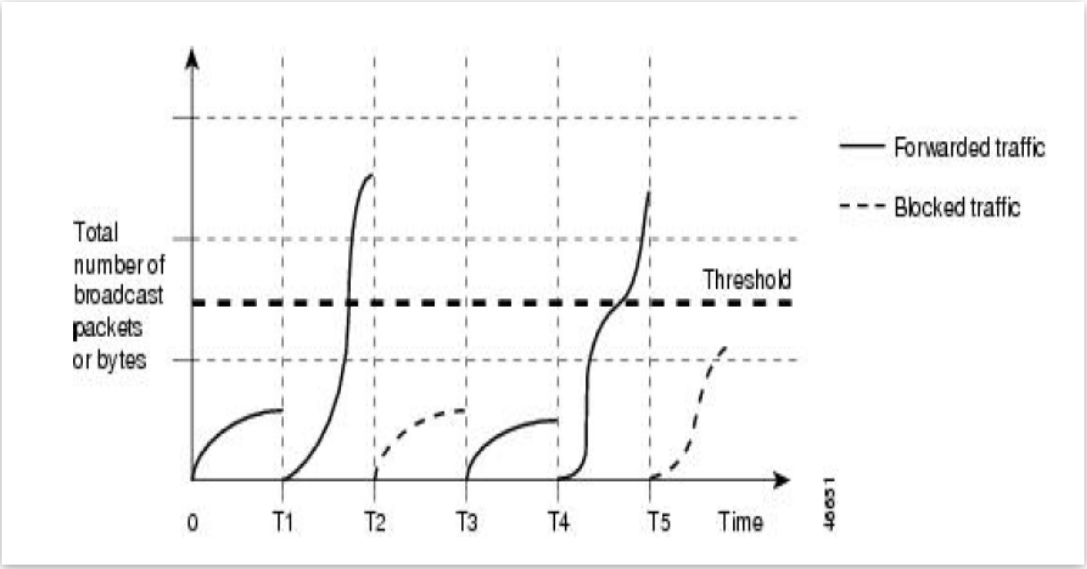

Limit broadcast, multicast and unknown unicast to a specific rate

- 1 sec sample interval

The rate-limiter will drop all profiled traffic above the threshold rate during the sample period

- All profile traffic including legitimate traffic

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Determining Typical Storm Control Rates

- Take some time to monitor bcast, mcast rates under normal conditions
  - Use Top N tools (if available on the platform)
  - Netlfow monitoring
  - Interface counters and SNMP tools
  - Wire Shark monitoring

Example using Top N reports on Catalyst 6500

```
VSS01#collect top counters interface ten sort-by broadcast interval 30
TopN collection started.
VSS01#
*Mar  6 16:58:18.735: %TOPN_COUNTERS-SW2-5-STARTED: TopN collection for report 1 started by console
```


© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Show Top Reports Example

```
VSS01#collect top counters interface tenGigabitEthernet sort-by broadcast
interval 30
VSS01# show top counters interface report 1
Started By         : console
Start Time         : 16:58:18 UTC Thu Mar 6 2014
End Time           : 16:58:48 UTC Thu Mar 6 2014
Port Type          : TenGigEthernet
Sort By            : broadcast
Interval           : 30 seconds

Port          Band  Util Bytes                Packets    Broadcast  Multicast  In-  Buf-
              width      (Tx + Rx)             (Tx + Rx)  (Tx + Rx)  (Tx + Rx)  err  ovfl
------------ ----- ---- ------------------ ---------- ---------- ---------- ---- ----

Te1/2/4       10000 0    21559                151        9          142        0    0
Te2/2/4       10000 0    21559                151        9          142        0    0
Te1/2/5       10000 0    18256                140        1          137        0    0
Te2/1/5       10000 0    18160                139        1          137        0    0
Te2/4/1       10000 0    168                    2          0          0          0    0
Te1/4/4       10000 9    7072412223   17148485  0          3          0    0
Te2/4/5       10000 0    168                    2          0          0          0    0
Te1/4/3       10000 9    7072658842   17148739  0          3          0    0
Te1/4/2       10000 0    6496                  82          0         82          0    0
Te1/4/8       10000 0    168                    2          0          0          0    0
```


© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

RMON History Example

```
C3850#show rmon history
Entry 30 is active, and owned by
 Monitors ifIndex.27 every 15 second(s)
 Requested # of time intervals, ie buckets, is 30,
  Sample # 8 began measuring at 3w4d
   Received 128 octets, 0 packets,
   0 broadcast and 0 multicast packets,
   0 undersized and 0 oversized packets,
   0 fragments and 0 jabbers,
   0 CRC alignment errors and 0 collisions.
   # of dropped packet events is 0
   Network utilization is estimated at 0
  Sample # 9 began measuring at 3w4d
   Received 460 octets, 0 packets,
   0 broadcast and 4 multicast packets,
   0 undersized and 0 oversized packets,
   0 fragments and 0 jabbers,
   0 CRC alignment errors and 0 collisions.
   # of dropped packet events is 0
   Network utilization is estimated at 0
  Sample # 10 began measuring at 3w4d
   Received 949 octets, 0 packets,
   0 broadcast and 2 multicast packets,
   0 undersized and 0 oversized packets,
   0 fragments and 0 jabbers,
```

RMON History Example on Catalyst 3850

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Storm Control Recommendations

Based on your network metrics, choose a value that will allow for peak Broadcast, multicast, unknown unicast plus 50%

- 1% bcast rate is common for 1GbE interfaces
- 0.5% is a common rate for 10GbE interfaces

The higher the interface speed the less the percentage needs to be

Be very cautions when configuring storm control for multicast frames, as this can limit BPDUs

Verify platform specific support and caveats

- Some platforms treat traffic types individually, some will group bcast and mcast together
- Some legacy hardware do not support storm-control or implement software-based mechanisms

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

IEEE802.1Q Trunk Interfaces

- Typically deployed on interconnection between access and distribution layers
- Used to extended multiple VLANs between L2 switches
- Should I use or how to use the Vlan Trunking Protocol?
- Should I use or how to use the Dynamic Trunking Protocol?

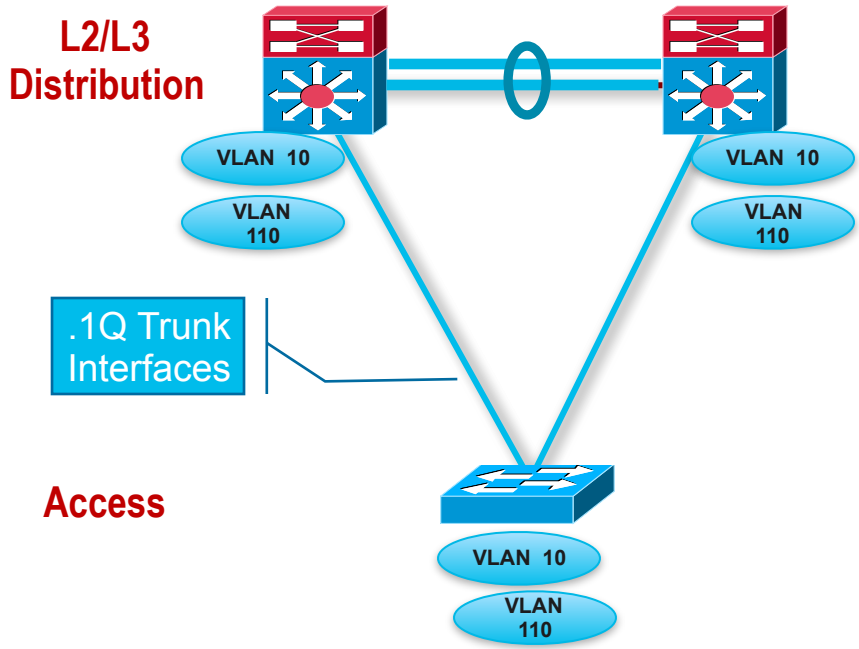

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Vlan Trunking Protocol (VTP)

Vlan Trunking Protocol (VTP) - used to maintain Vlan configuration consistency across a given domain of switches

- VLAN configurations are made at the VTP server
- VLAN configuration is propagated via VTP protocols to client switches
- Intended to minimise misconfigurations and inconsistencies

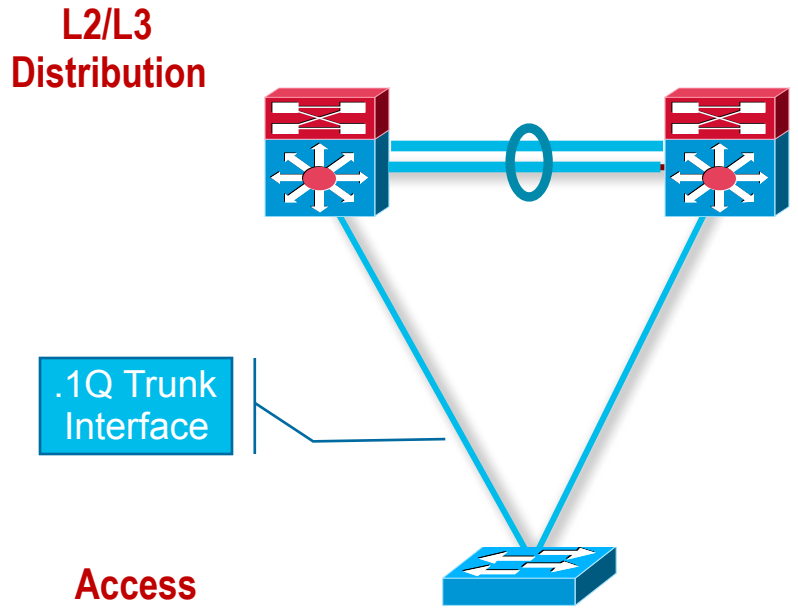

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Vlan Trunking Protocol (VTP)

Vlan Trunking Protocol (VTP) - used to maintain Vlan configuration consistency across a given domain of switches

- VLAN configurations are made at the VTP server
- VLAN configuration is propagated via VTP protocols to client switches
- Intended to minimise misconfigurations and inconsistencies

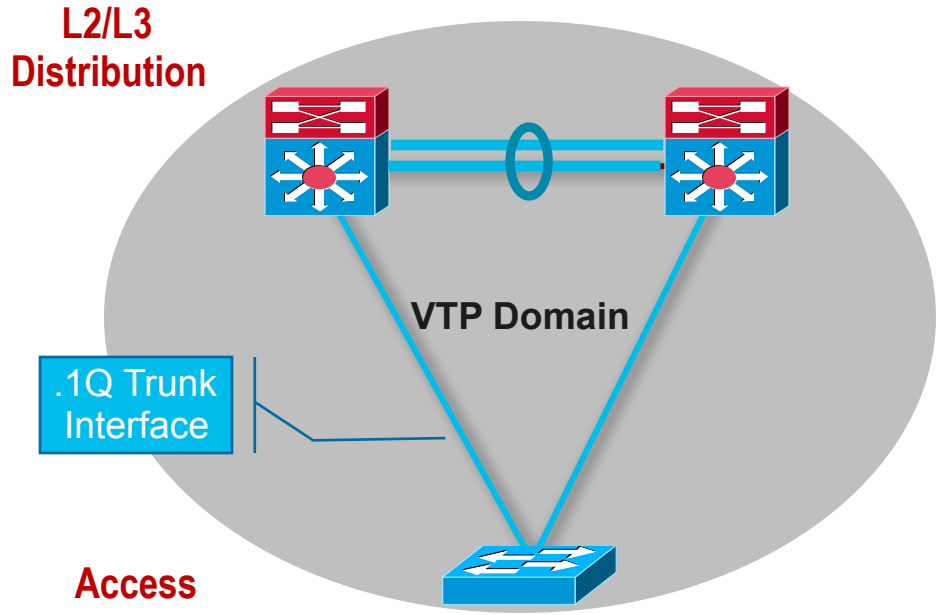

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Vlan Trunking Protocol (VTP)

Vlan Trunking Protocol (VTP) - used to maintain Vlan configuration consistency across a given domain of switches

- VLAN configurations are made at the VTP server
- VLAN configuration is propagated via VTP protocols to client switches
- Intended to minimise misconfigurations and inconsistencies

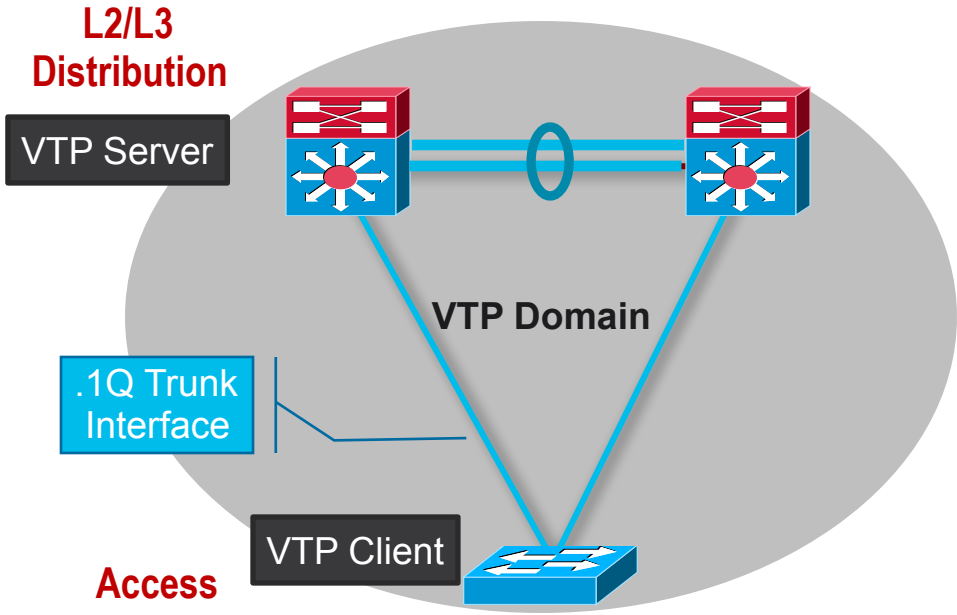

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Vlan Trunking Protocol (VTP)

Vlan Trunking Protocol (VTP) - used to maintain Vlan configuration consistency across a given domain of switches

- VLAN configurations are made at the VTP server
- VLAN configuration is propagated via VTP protocols to client switches
- Intended to minimise misconfigurations and inconsistencies

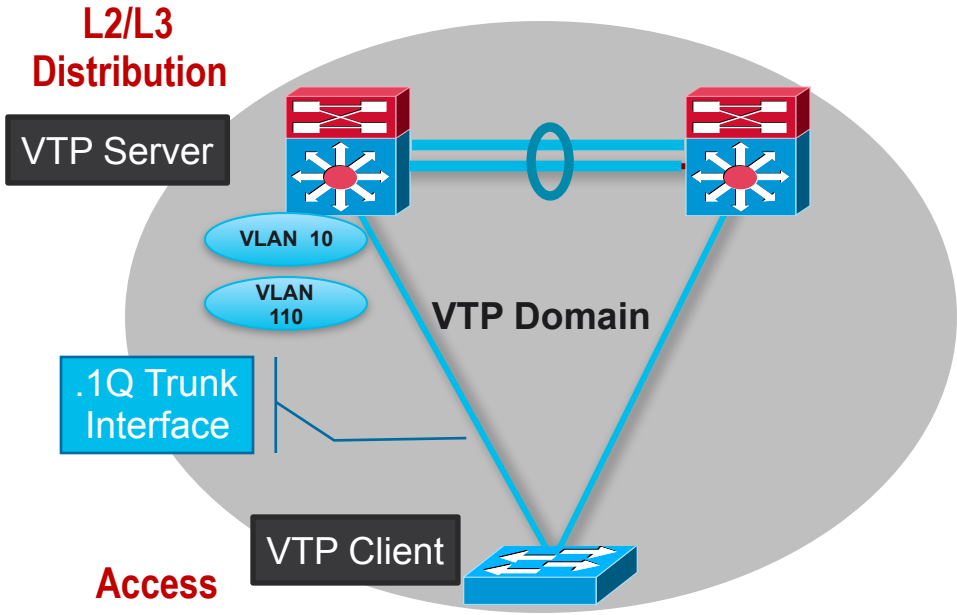

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Vlan Trunking Protocol (VTP)

Vlan Trunking Protocol (VTP) - used to maintain Vlan configuration consistency across a given domain of switches

- VLAN configurations are made at the VTP server
- VLAN configuration is propagated via VTP protocols to client switches
- Intended to minimise misconfigurations and inconsistencies

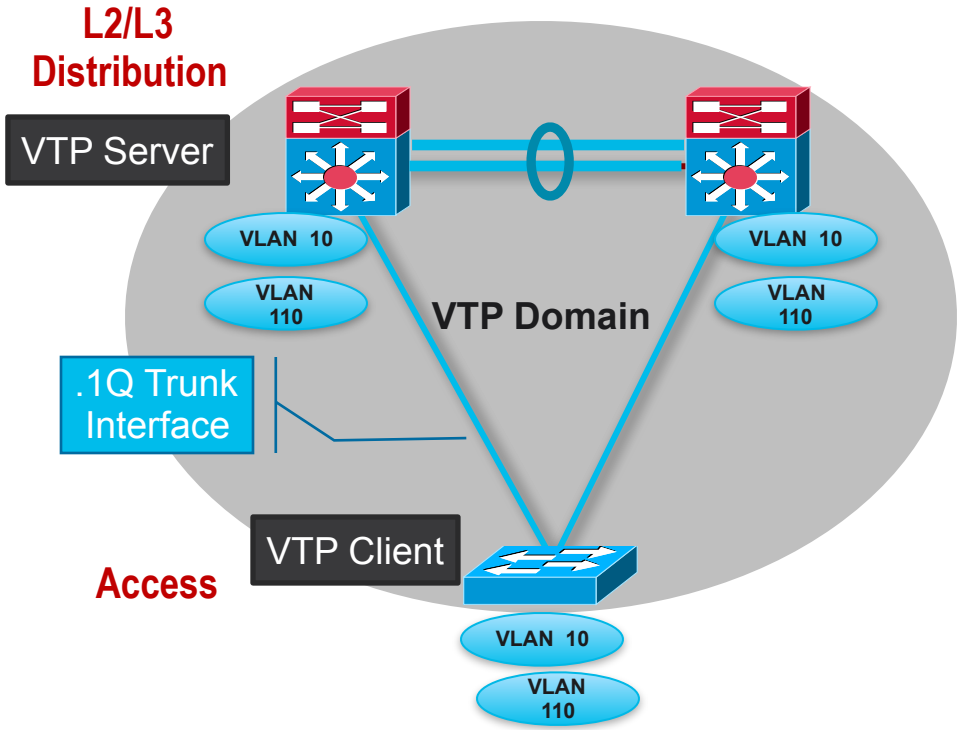

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Vlan Trunking Protocol (VTP) Recommendations

- Weigh the benefits and drawbacks of deploying VTP for your organisation
- VTP Transparent mode is suitable for majority of current campus deployments
  - Promotes use of stringent change control processes
  - Eliminates/creates some possibilities for misconfigurations
  - Coupled with manual VLAN pruning can minimises extensive use of virtual port resources

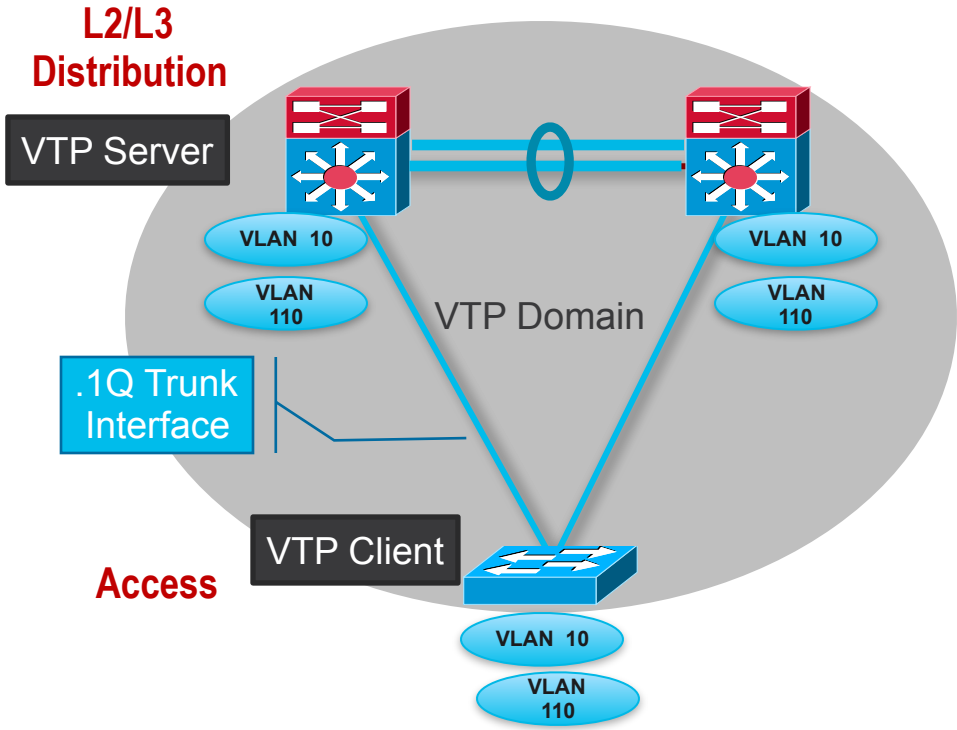

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Vlan Trunking Protocol (VTP) Recommendations

- Configure VTP mode transparent
  - Router(config)# vtp mode transparent
- Change the native VLAN to something unused to avoid VLAN hopping
  - Router (config) # switchport trunk native vlan 999
- Manually prune all VLANS except those needed to the downstream switch

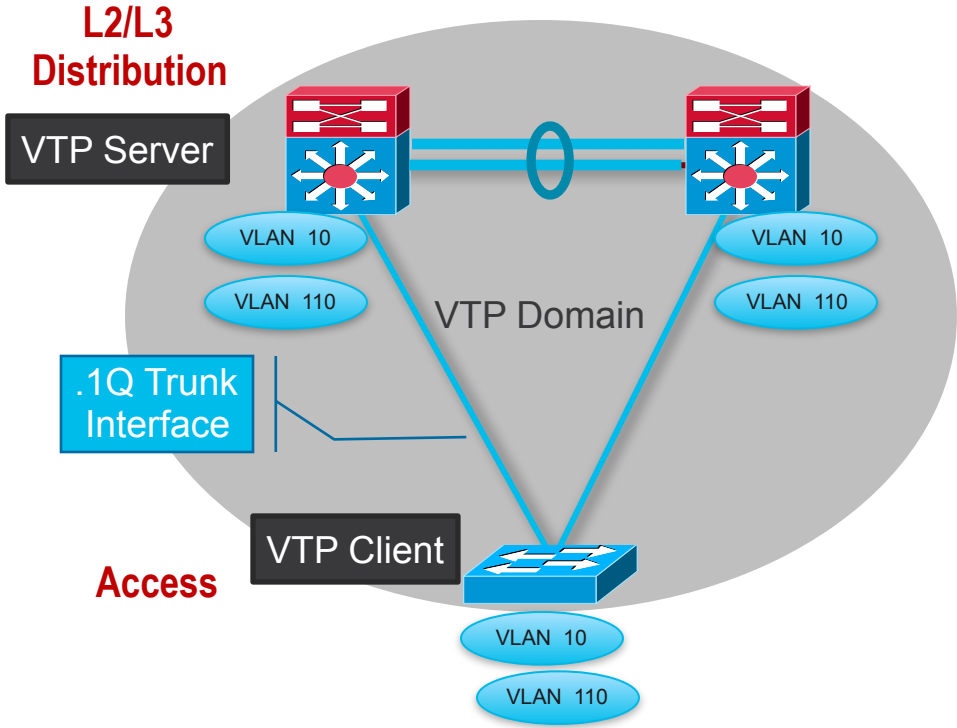

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Dynamic Trunking Protocol

DTP is an Layer 2 protocol that negotiates configuration parameters between a switch port and its neighbour.

- Encapsulation type
- Native VLAN
- Hardware capability

DTP ensures that ports and their neighbours are in consistent states.

- also helps protect against the flooding of tagged frames by non-trunk ports
- potentially a serious security risk

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

DTP Dynamic Trunk Protocol

- Automatic formation of trunked switch-to-switch interconnection
  - Trunk: always be a trunk
  - Desirable: ask if the other side can/will
  - Auto: if the other sides asks I will
  - Access: don't become a trunk
- Negotiation of 802.1Q or ISL encapsulation
  - ISL: try to use ISL trunk encapsulation
  - 802.1q: try to use 802.1q encapsulation
  - Negotiate: negotiate ISL or 802.1q encapsulation with peer
  - Non-negotiate: always use encapsulation that is hard set

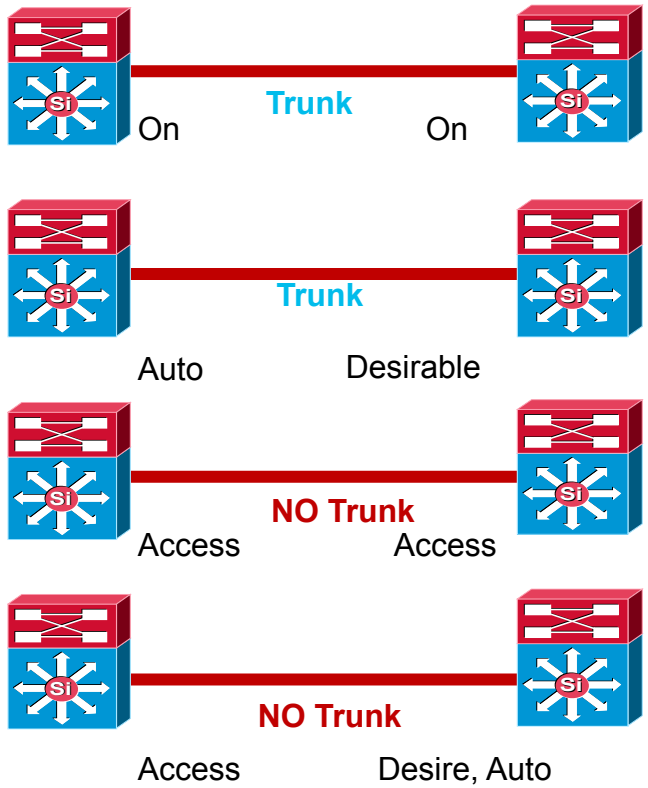

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Switchport Mode Command Options

Mode	DTP Frames Transmitted ?	Final State ? (local port)
Dynamic Auto	Yes, periodic	Trunking, if far side requests
Trunk	Yes, periodic	Trunking, unconditionally
Nonegotiate	No	Trunking, unconditionally
Dynamic desirable	Yes, periodic	It ends up in trunking state only if the remote mode is on, auto, or desirable.
Access	No, in steady state, but transmits informs in order to speed up remote-end detection after a change from on.	Non-trunking

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Optimising Convergence: Trunk Tuning

Trunk Auto/Desirable Takes Some Time

Tuning DTP with “nonegotiate” improves link up convergence time

```
IOS(config-if)# switchport mode trunk
IOS(config-if)# switchport nonegotiate
```

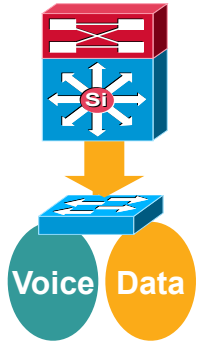

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Trunking/VTP/DTP - Quick Summary

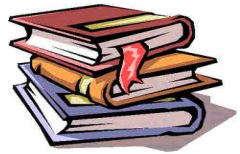

- VTP transparent should be used; there is a trade off between administrative overhead and the temptation to span existing VLANS across multiple access layer switches
- One can consider a configuration that uses DTP ON/ON and NO NEGOTIATE; there is a trade off between performance/HA impact and maintenance and operations implications
- An ON/ON and NO NEGOTIATE configuration is faster from a link up (restoration) perspective than a desirable/desirable alternative. However, in this configuration DTP is not actively monitoring the state of the trunk and a misconfigured trunk is not easily identified
- It's really a balance between fast convergence and your ability to manage configuration and change control …

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Unidirectional Links

- Example topology where a uni-directional can cause an L2 loop
- If SW3 stops receiving BPDUs form SW2, SW3 will change its Blocking Port to Forwarding after STP timeout

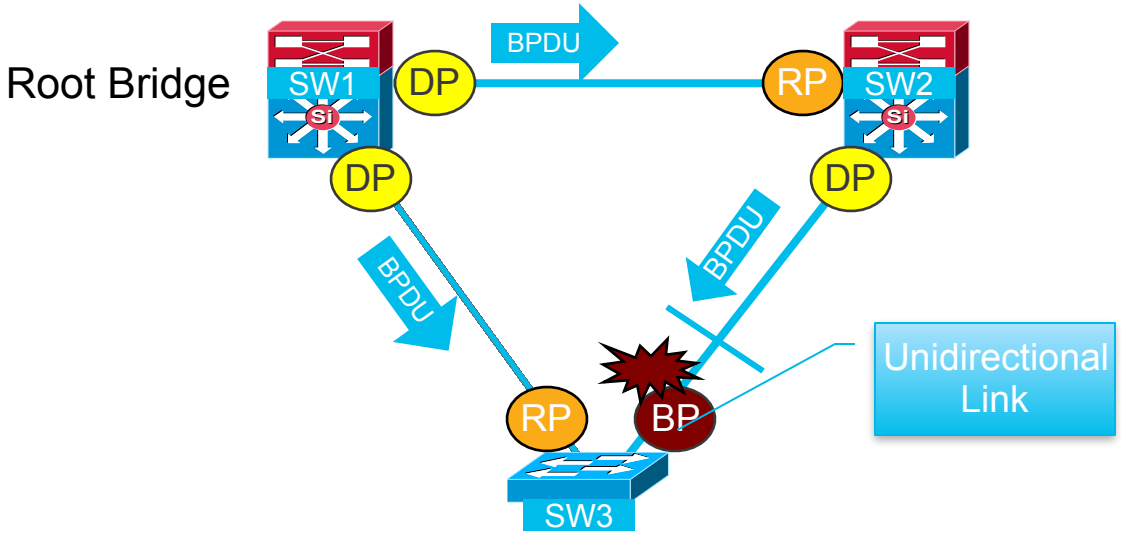

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Unidirectional Link Detection Protocol (UDLD)

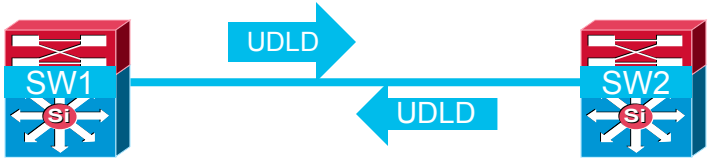

- UDLD works by exchanging protocol packets between the neighbouring devices.
- Both devices on the link must support UDLD and have it enabled on respective ports.
- UDLD protocol packets contain the port's own device/port ID, and the neighbour's device/port IDs seen by UDLD on that port. Neighbouring ports should see their own device/port ID (echo) in the packets received from the other side.

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

UDLD Aggressive and UDLD Normal

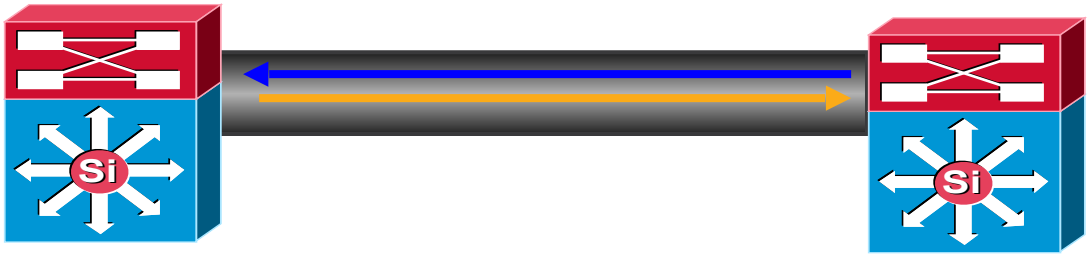

- Timers are the same—15-second hellos by default

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

UDLD Aggressive and UDLD Normal

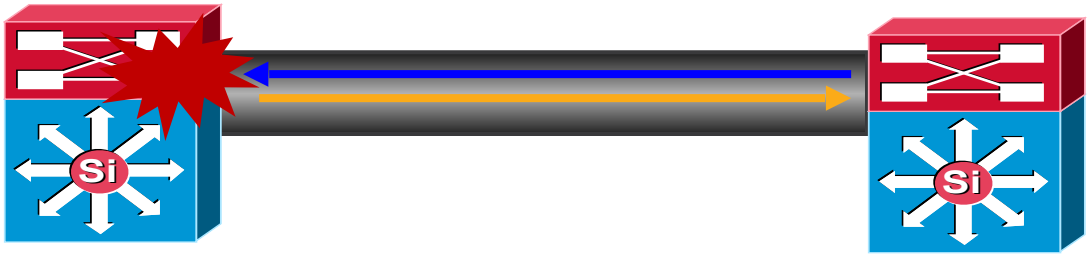

- Timers are the same—15-second hellos by default
- UDLD—Normal Mode—will transition a port to “undetermined” state – will not prevent an L2 loop from occurring

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

UDLD Aggressive and UDLD Normal

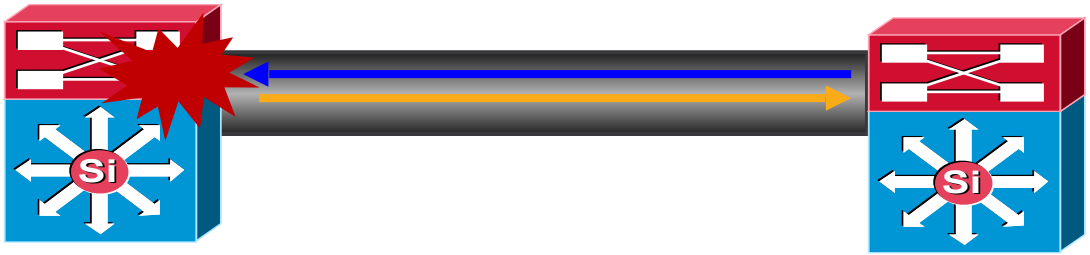

- Timers are the same—15-second hellos by default
- UDLD—Normal Mode—will transition a port to “undetermined” state – will not prevent an L2 loop from occurring
- UDLD—Aggressive—err-disable **both** ends of the connection
  due to err-disable when aging and re-establishment of UDLD communication fails

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

UDLD Aggressive and UDLD Normal

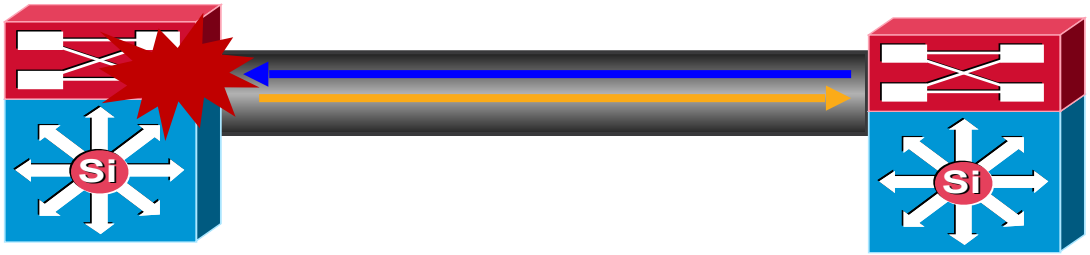

- Timers are the same—15-second hellos by default
- UDLD—Normal Mode—will transition a port to “undetermined” state – will not prevent an L2 loop from occurring
- UDLD—Aggressive—err-disable **both** ends of the connection
  due to err-disable when aging and re-establishment of UDLD communication fails
  - Aggressive Mode—after aging on a previously bi-directional link—tries eight times (once per second) to reestablish connection then err-disables port

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

UDLD Aggressive and UDLD Normal

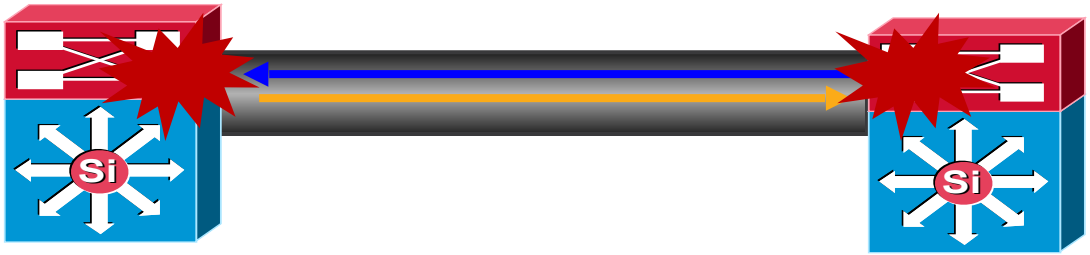

- Timers are the same—15-second hellos by default
- UDLD—Normal Mode—will transition a port to “undetermined” state – will not prevent an L2 loop from occurring
- UDLD—Aggressive—err-disable **both** ends of the connection
  due to err-disable when aging and re-establishment of UDLD communication fails
  - Aggressive Mode—after aging on a previously bi-directional link—tries eight times (once per second) to reestablish connection then err-disables port
  - Use `errdisable recovery` configuration to automate recovery

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Best Practices - UDLD Configuration

- Typically deployed on any fibre optic interconnection
- Use UDLD aggressive mode with caution
  - Ensure an out-of-band console connection to the device in then event that an err-disabled port will cut-off in-band management
  - Err-disable recovery method
    - Use Automated
- Turn on in global configuration to avoid operational error/ misses
- Config example
  - Cisco IOS:
    `udld enable`
  - Enables UDLD globally on fibre-optic LAN ports.
- Make sure UDLD timers will expire before the STP convergence = Maxage+ 2 X Forward Delay time

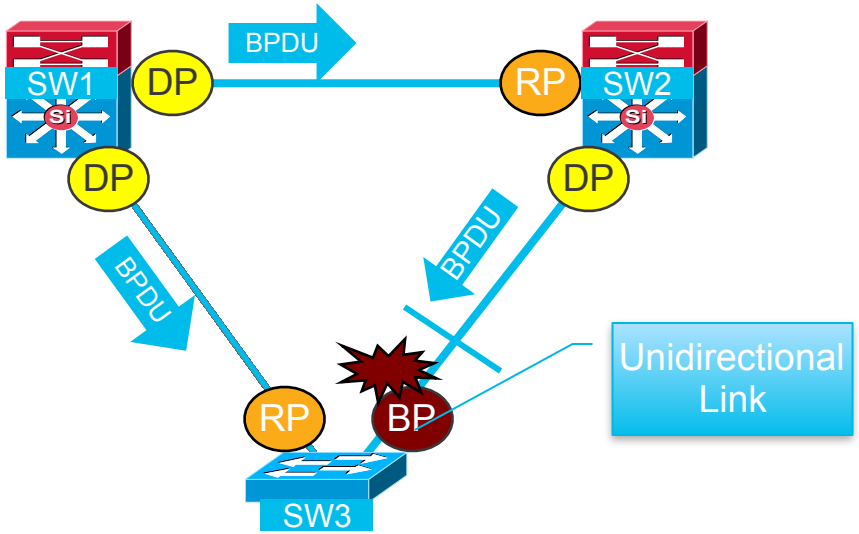

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Spanning Tree Toolkit

Multilayer Network Design

Layer 2 Access with Layer 3 Distribution

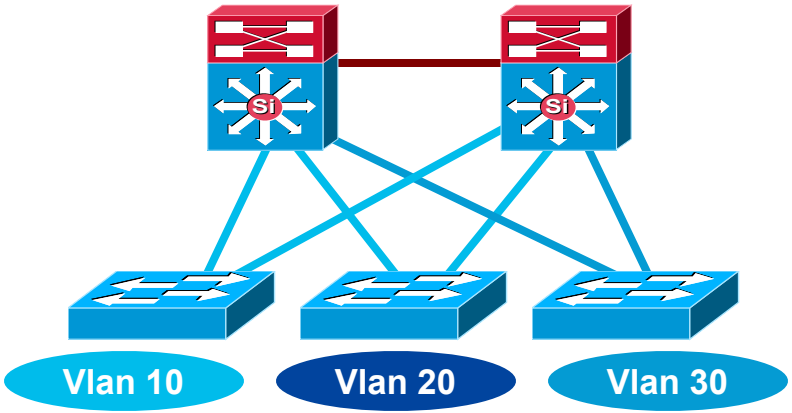

- Each access switch has unique VLANs
- No Layer 2 loops
- Layer 3 link between distribution
- No blocked links

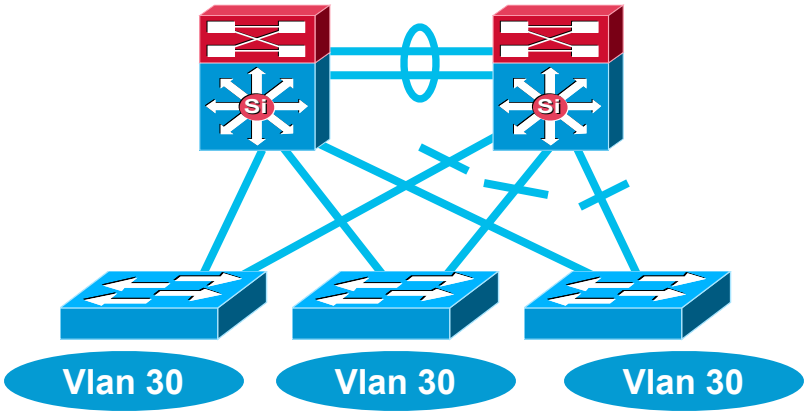

- At least some VLANs span multiple access switches
- Layer 2 loops
- Layer 2 and 3 running over link between distribution
- Blocked links

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Spanning Tree Protocol Options

- Rapid-PVST+ greatly improves the restoration times for any VLAN that requires a topology convergence due to link UP
- PVST+ (IEEE 802.1d)
  - Traditional spanning tree implementation
- Rapid PVST+ (IEEE 802.1w)
  - Scales to large size (~10,000 logical ports)
  - **Easy to implement, proven, scales**
  - **Default STP mode in current IOS releases**
- MST (IEEE 802.1s)
  - Permits very large scale STP implementations (~30,000 logical ports)
  - **Not as flexible as rapid PVST+**

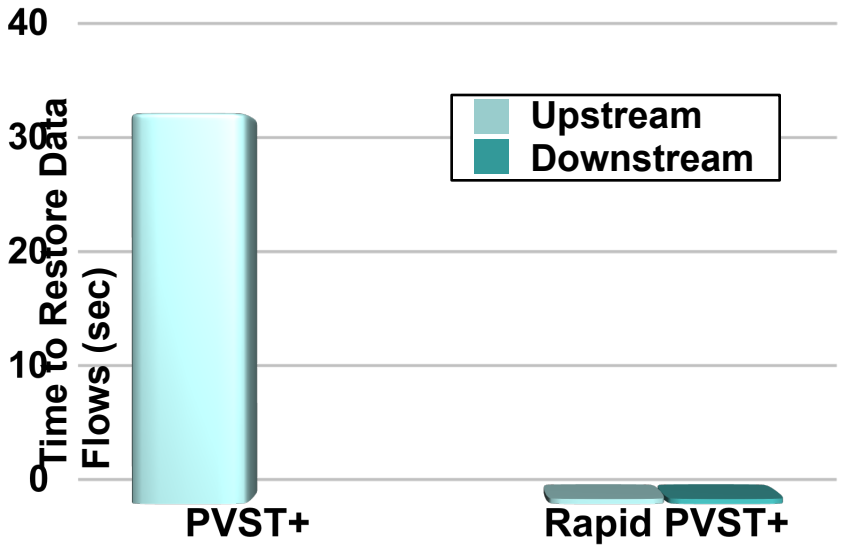

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Optimising the Layer 2 Design

STP Toolkit – PortFast and BPDU Guard

- PortFast is configured on edge ports to allow them to quickly move to forwarding bypassing listening and learning and avoids TCN (Topology Change Notification) messages
- BPDU Guard can prevent loops by moving PortFast configured interfaces that receive BPDUs to errdisable state
- BPDU Guard prevents ports configured with PortFast from being incorrectly connected to another switch
- When enabled globally, BPDU Guard applies to all interfaces that are in an operational PortFast state

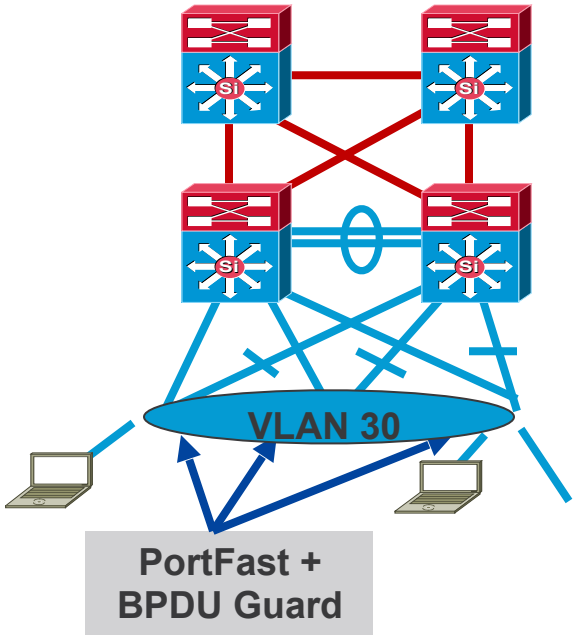

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Optimising the Layer 2 Design

STP Toolkit – PortFast and BPDU Guard

- PortFast is configured on edge ports to allow them to quickly move to forwarding bypassing listening and learning and avoids TCN (Topology Change Notification) messages
- BPDU Guard can prevent loops by moving PortFast configured interfaces that receive BPDUs to errdisable state
- BPDU Guard prevents ports configured with PortFast from being incorrectly connected to another switch
- When enabled globally, BPDU Guard applies to all interfaces that are in an operational PortFast state

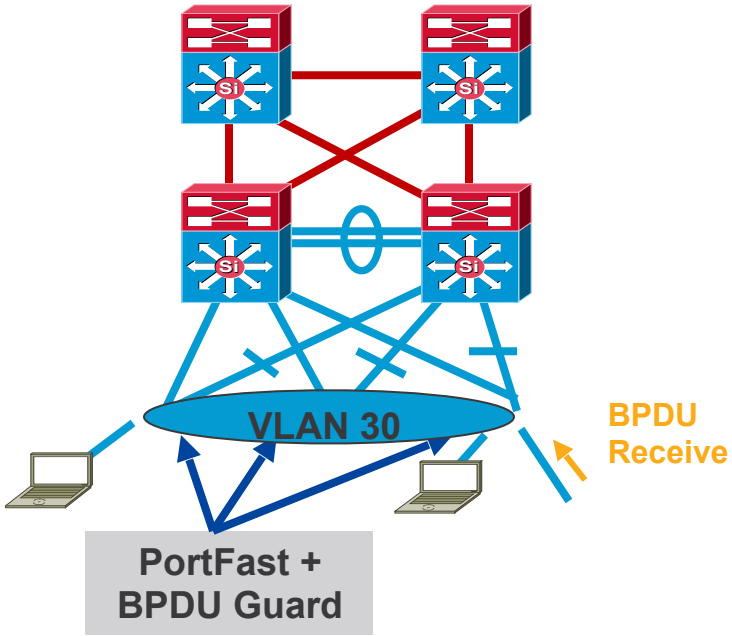

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Optimising the Layer 2 Design

STP Toolkit – PortFast and BPDU Guard

- PortFast is configured on edge ports to allow them to quickly move to forwarding bypassing listening and learning and avoids TCN (Topology Change Notification) messages
- BPDU Guard can prevent loops by moving PortFast configured interfaces that receive BPDUs to errdisable state
- BPDU Guard prevents ports configured with PortFast from being incorrectly connected to another switch
- When enabled globally, BPDU Guard applies to all interfaces that are in an operational PortFast state

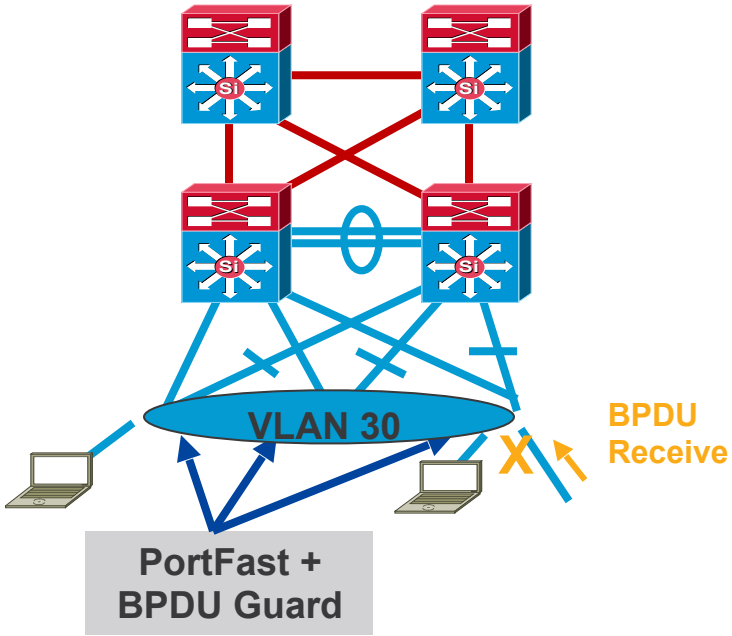

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Force Spanning Tree to Perform as Expected

- Place the root where you want it
  - Root primary/secondary macro
  - spanning-tree vlan 1-4094 root primary
- The root bridge should stay where you put it
  - RootGuard
  - LoopGuard
  - UplinkFast (PVST+)
  - UDLD
- Only end-station traffic should be seen on an edge port
  - PortFast
  - BPDU Guard
  - RootGuard
  - Port-security

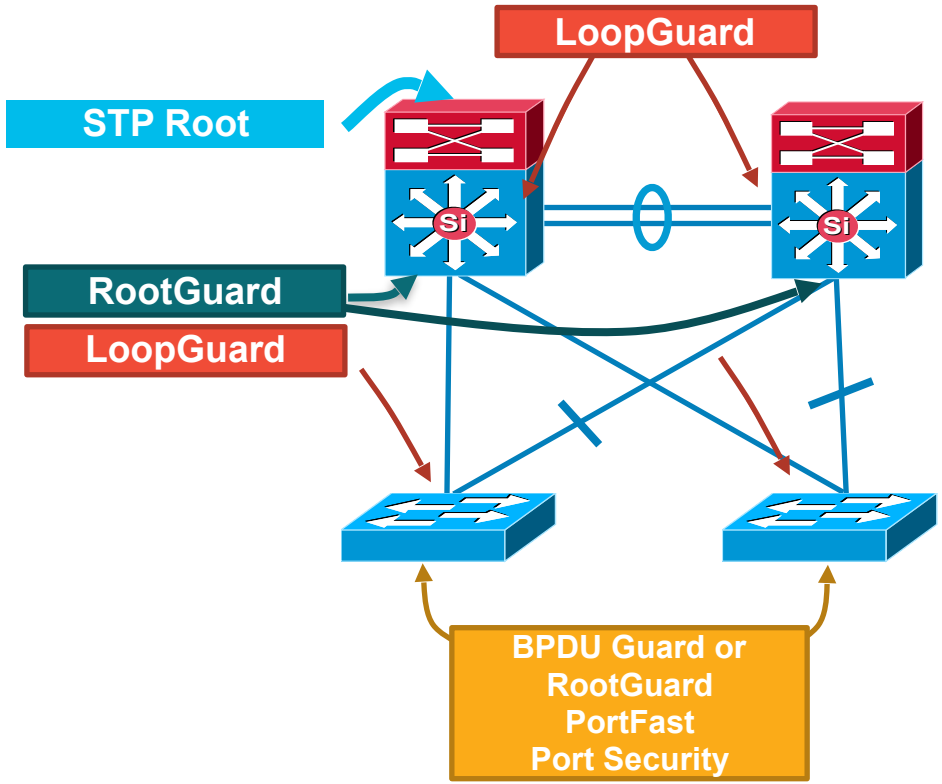

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Avoid Daisy Chaining Access Layer Switches

Daisy Chain without a backup link can cause traffic blackhole

Return Path Traffic Has a 50/50 Chance of Being 'Black Holed'

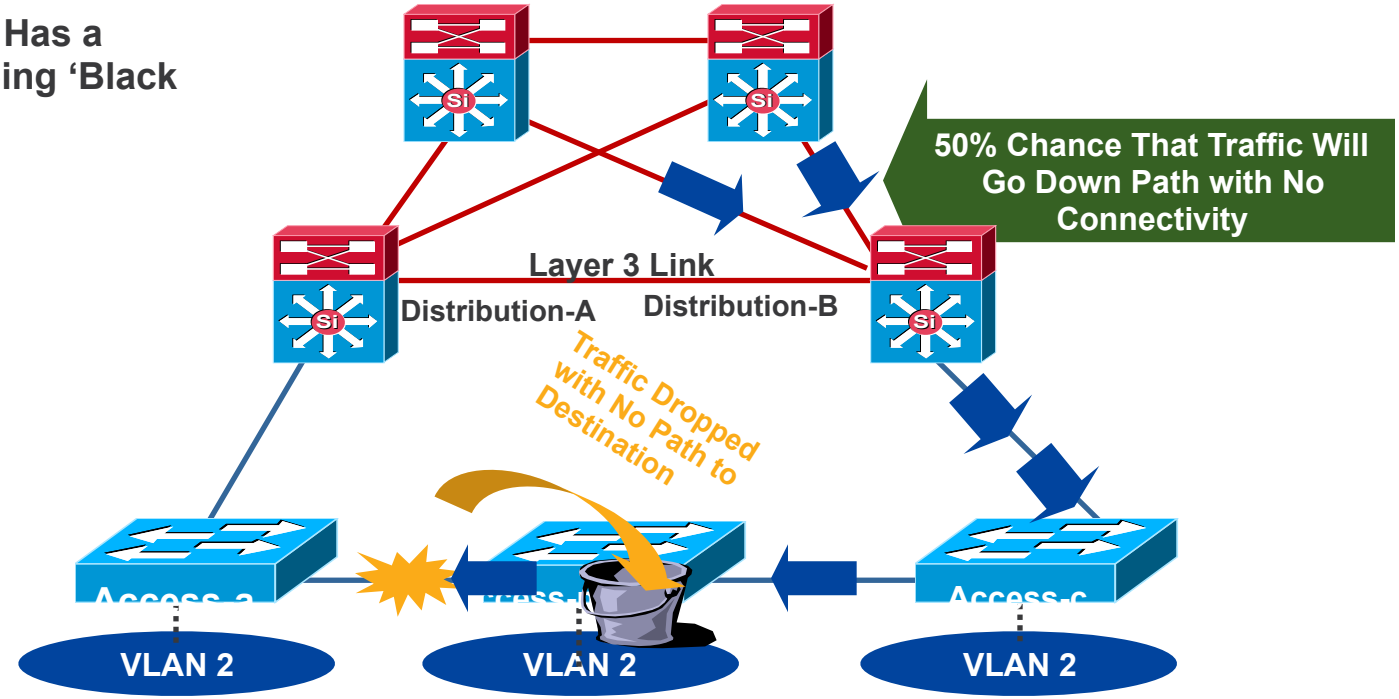

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Avoid Daisy Chaining Access Layer Switches

Daisy Chain without a backup link can cause traffic blackhole

Return Path Traffic Has a 50/50 Chance of Being 'Black Holed'

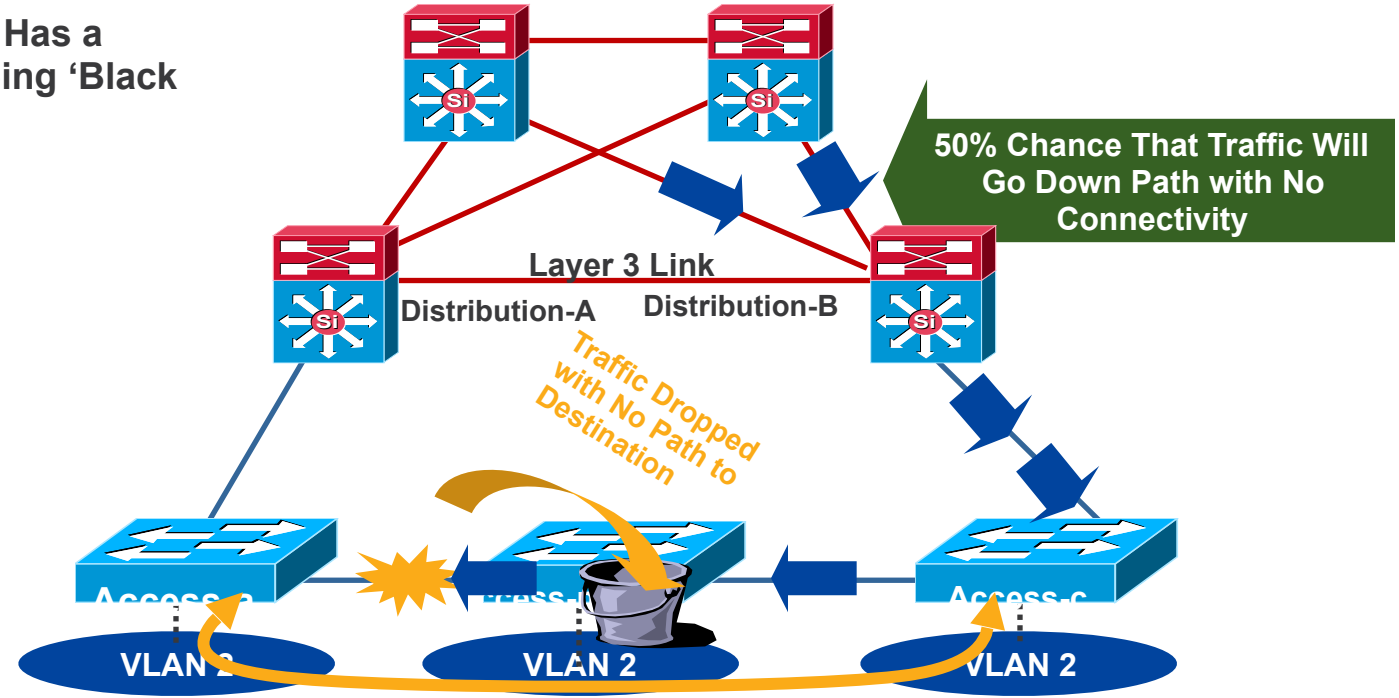

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Use Stacking Technologies for Access Layer Switches

Stacking Technology Addresses Old Problems

- **Stackwise/Stackwise-Plus** technology eliminates the concern
  - Loopback links not required
  - No longer forced to have L2 link in distribution
- If you use modular (chassis-based) switches, these problems are not a concern

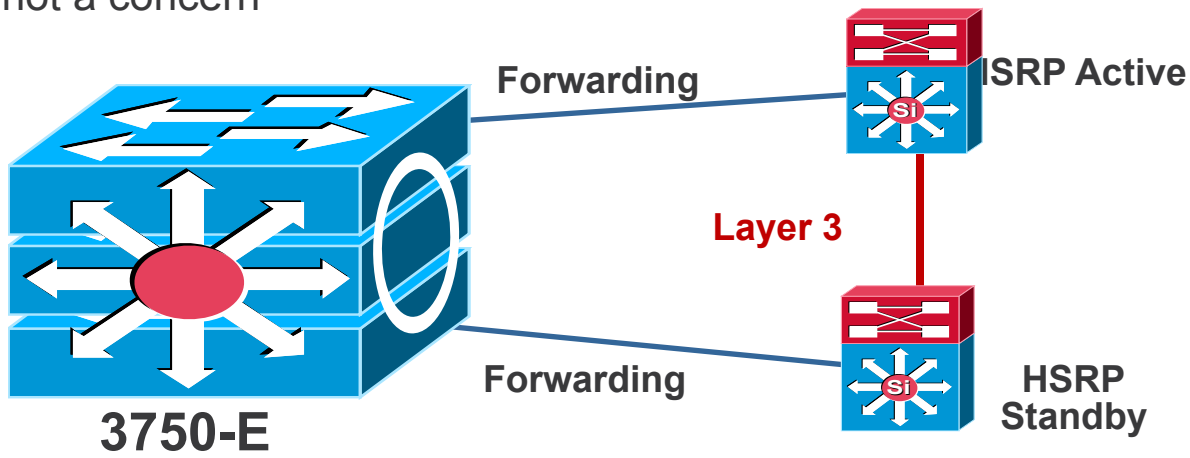

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Asymmetric Routing (Unicast Flooding)

- Affects redundant topologies with shared L2 access
- One path upstream and two paths downstream
- CAM table entry ages out on standby HSRP
- Without a CAM entry packet is flooded to all ports in the VLAN

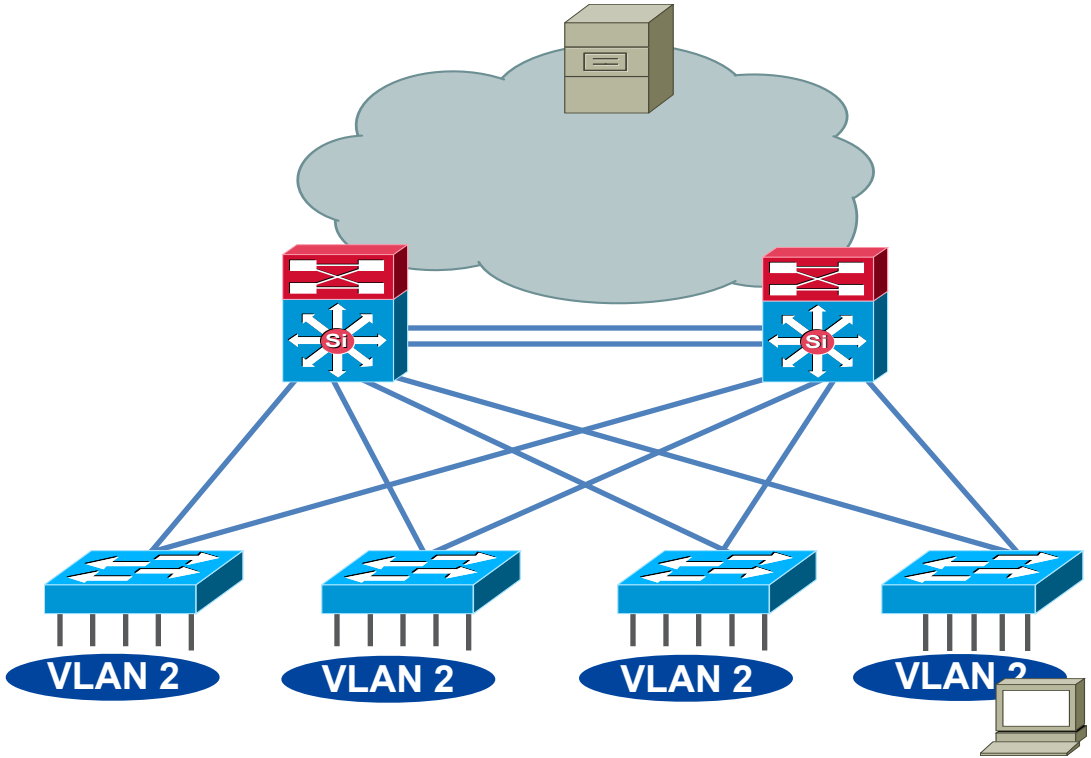

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Asymmetric Routing (Unicast Flooding)

- Affects redundant topologies with shared L2 access
- One path upstream and two paths downstream
- CAM table entry ages out on standby HSRP
- Without a CAM entry packet is flooded to all ports in the VLAN

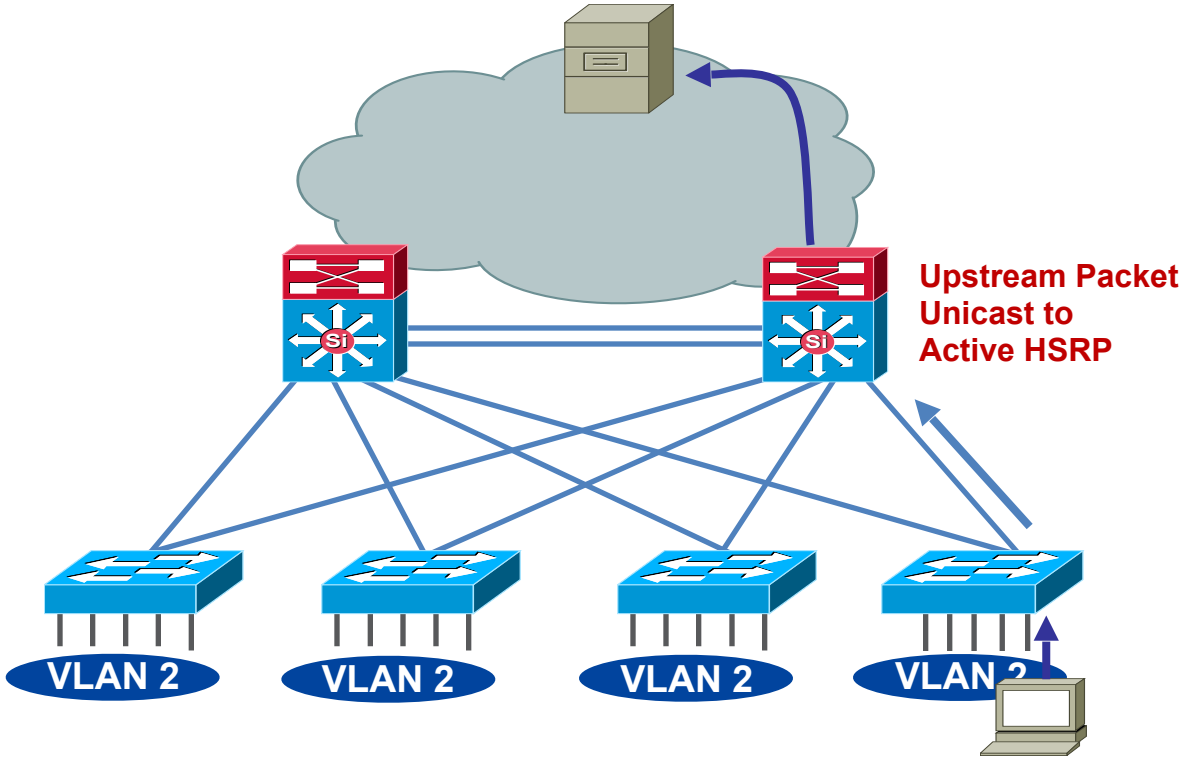

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Asymmetric Routing (Unicast Flooding)

- Affects redundant topologies with shared L2 access
- One path upstream and two paths downstream
- CAM table entry ages out on standby HSRP
- Without a CAM entry packet is flooded to all ports in the VLAN

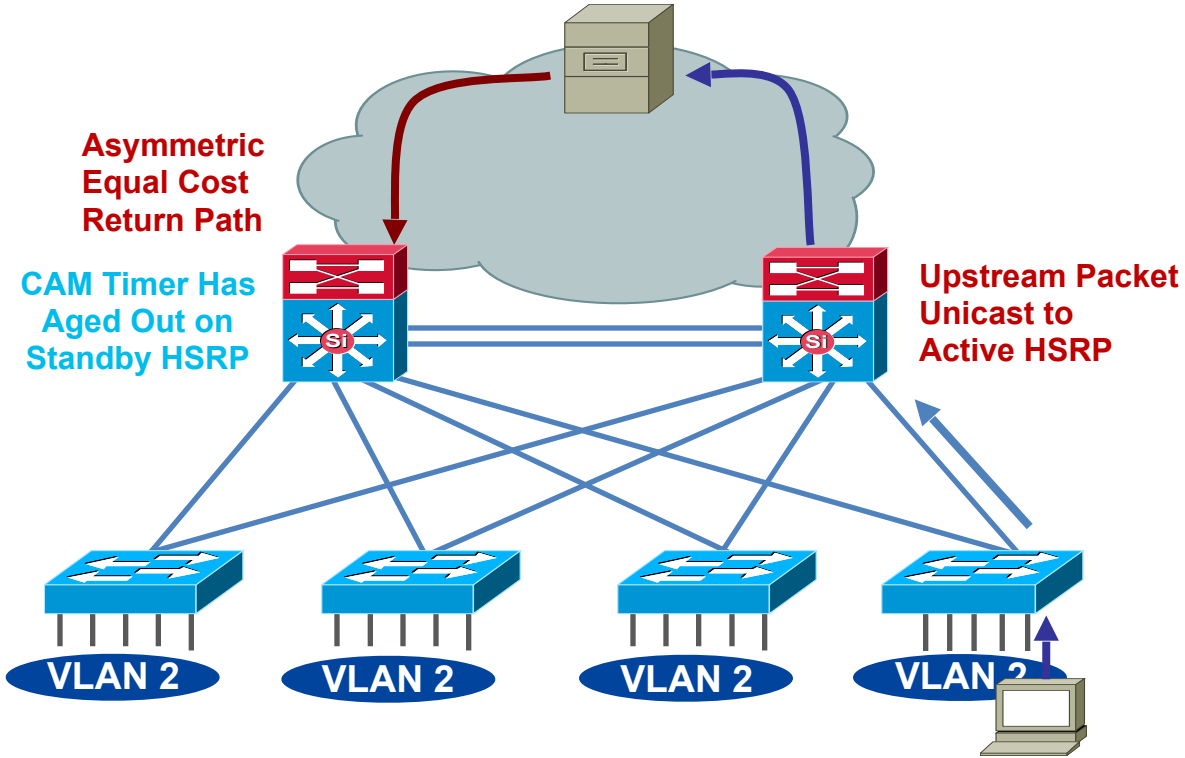

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Asymmetric Routing (Unicast Flooding)

- Affects redundant topologies with shared L2 access
- One path upstream and two paths downstream
- CAM table entry ages out on standby HSRP
- Without a CAM entry packet is flooded to all ports in the VLAN

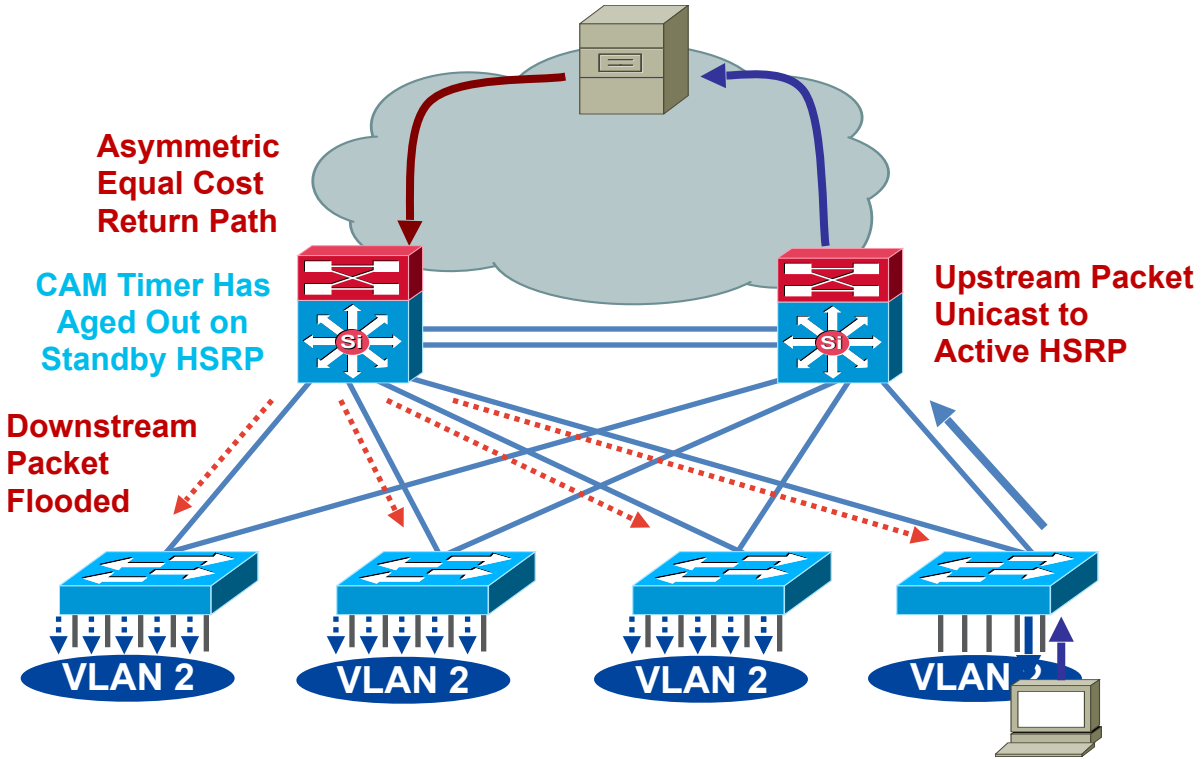

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Best Practices Prevent Unknown Unicast Flooding

- Assign one unique data and voice VLAN to each access switch
- Traffic is now only flooded down one trunk
- Access switch unicasts correctly; no flooding to all ports
- If you have to:
  - Tune ARP and CAM aging timers; CAM timer exceeds ARP timer
  - Bias routing metrics to remove equal cost routes

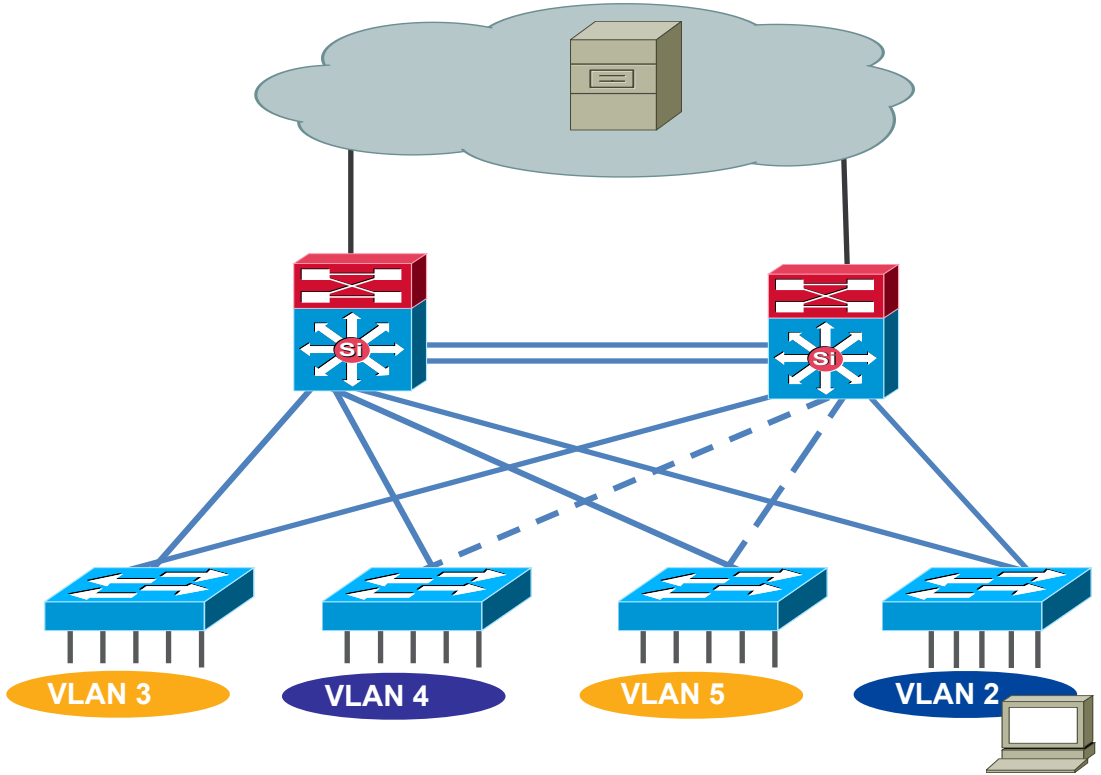

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Best Practices Prevent Unknown Unicast Flooding

- Assign one unique data and voice VLAN to each access switch
- Traffic is now only flooded down one trunk
- Access switch unicasts correctly; no flooding to all ports
- If you have to:
  - Tune ARP and CAM aging timers; CAM timer exceeds ARP timer
  - Bias routing metrics to remove equal cost routes

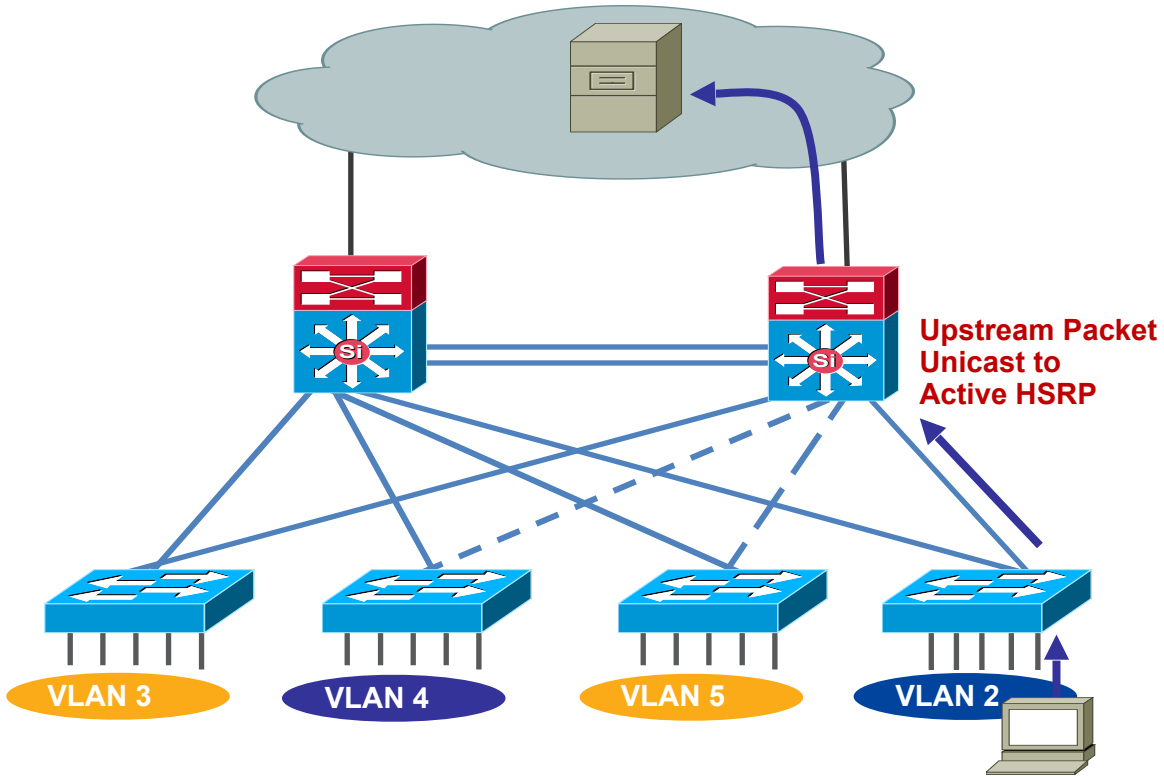

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Best Practices Prevent Unknown Unicast Flooding

- Assign one unique data and voice VLAN to each access switch
- Traffic is now only flooded down one trunk
- Access switch unicasts correctly; no flooding to all ports
- If you have to:
  - Tune ARP and CAM aging timers; CAM timer exceeds ARP timer
  - Bias routing metrics to remove equal cost routes

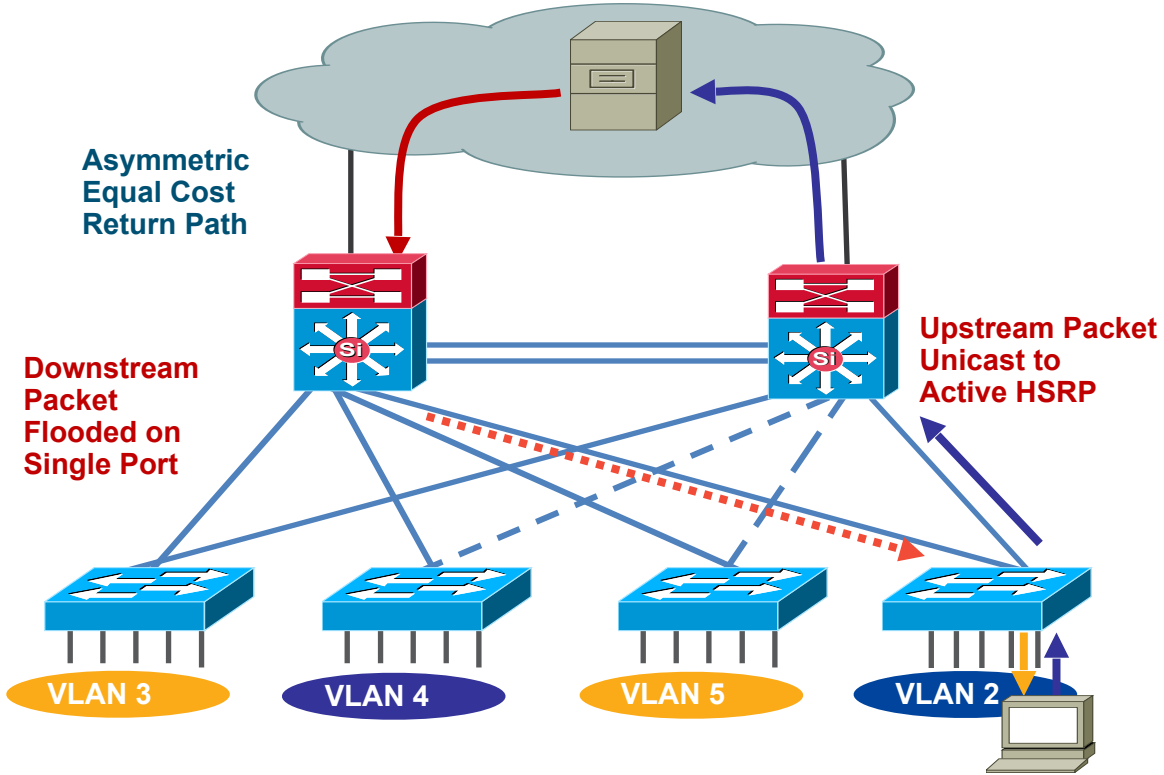

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Integrated Security Toolkit

Securing Layer 2 From Surveillance Attacks

Cutting Off MAC-Based Attacks

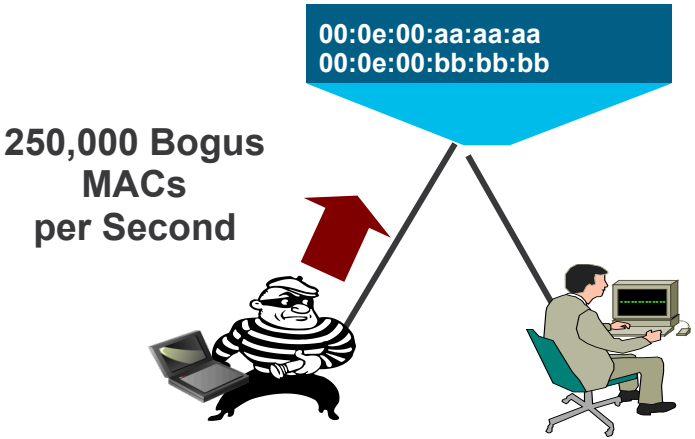

Problem –

Hacking Tools Enable Attackers to Flood Switch CAM Tables with Bogus MACs – Turning the VLAN into a Hub

Switch CAM Table Limit Is Finite Number of MAC Addresses

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Securing Layer 2 From Surveillance Attacks

Cutting Off MAC-Based Attacks

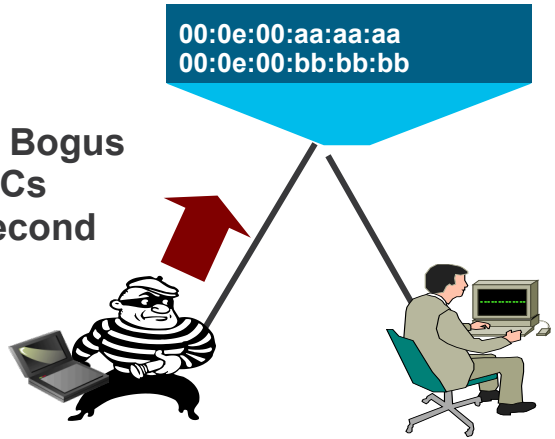

Problem –

Hacking Tools Enable Attackers to Flood Switch CAM Tables with Bogus MACs – Turning the VLAN into a Hub

Switch CAM Table Limit Is Finite Number of MAC Addresses

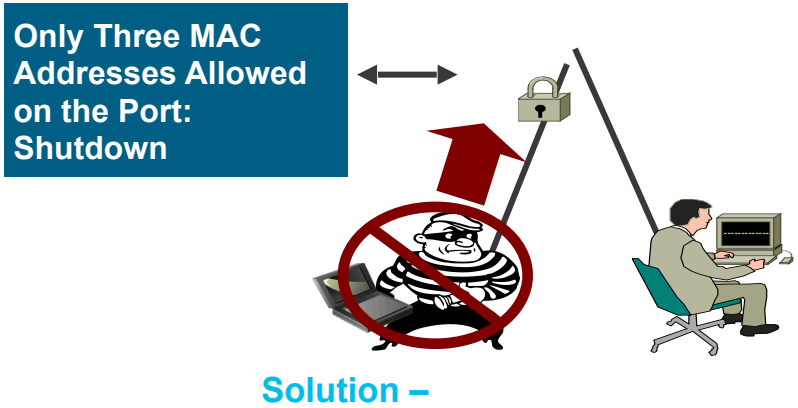

Solution –

Port Security Limits MAC Flooding Attack – Locks Down Port and Sends an SNMP Trap

```
switchport port-security
switchport port-security maximum 10
switchport port-security violation restrict
switchport port-security aging time 2
switchport port-security aging type inactivity
```


© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

DHCP Snooping

Protection Against Rogue / Malicious DHCP Server

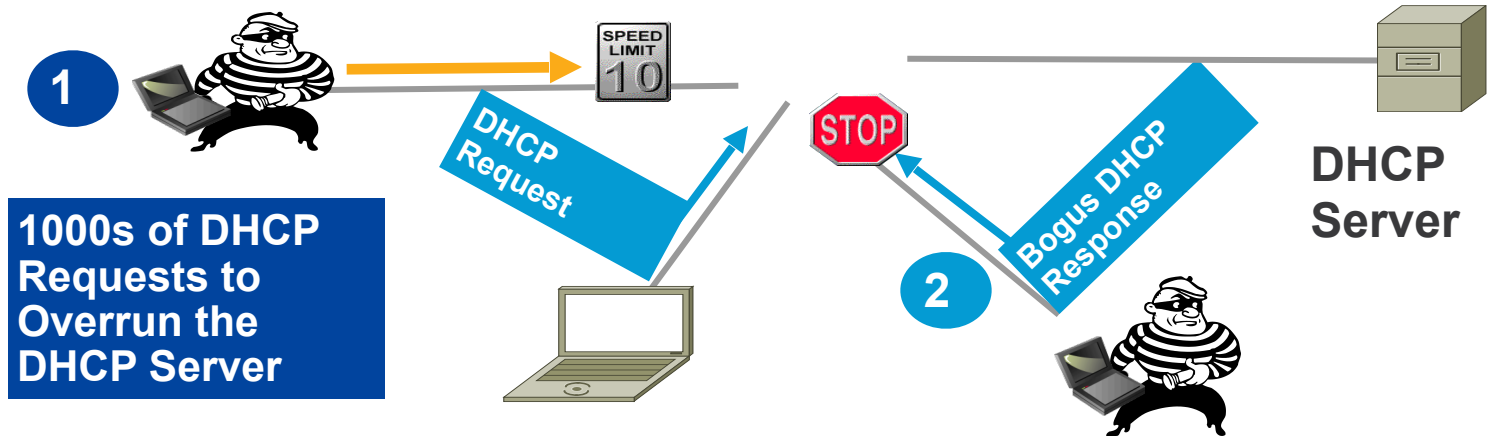

- DHCP requests (discover) and responses (offer) tracked
- Rate-limit requests on trusted interfaces –limits DoS attacks on DHCP server
- Deny responses (offers) on non-trusted interfaces – stop malicious or errant DHCP server

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

DHCP Snooping

Protection Against Rogue / Malicious DHCP Server

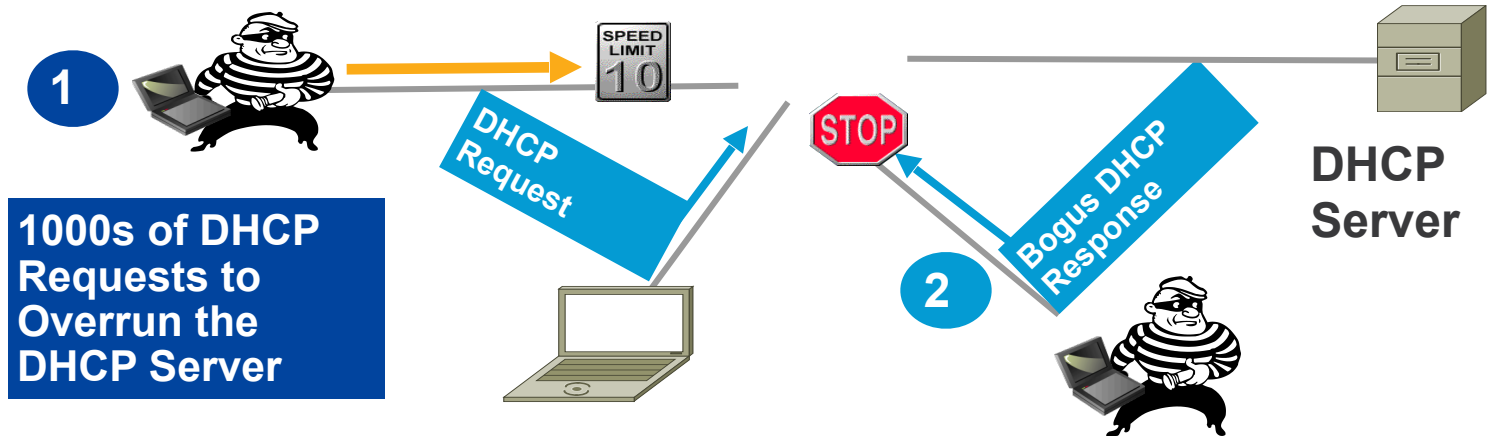

- DHCP requests (discover) and responses (offer) tracked
- Rate-limit requests on trusted interfaces –limits DoS attacks on DHCP server
- Deny responses (offers) on non-trusted interfaces – stop malicious or errant DHCP server

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Securing Layer 2
From Surveillance Attacks

Protection Against **ARP Poisoning –**

- **Dynamic ARP inspection** protects against ARP poisoning (ettercap, dsnif, arpspoof)
- Uses the DHCP Snooping binding table
- Tracks MAC to IP from DHCP transactions
- Rate-limits ARP requests from client ports
- Drop **bogus** gratuitous ARPs – stops ARP poisoning / “Man In The Middle” attacks

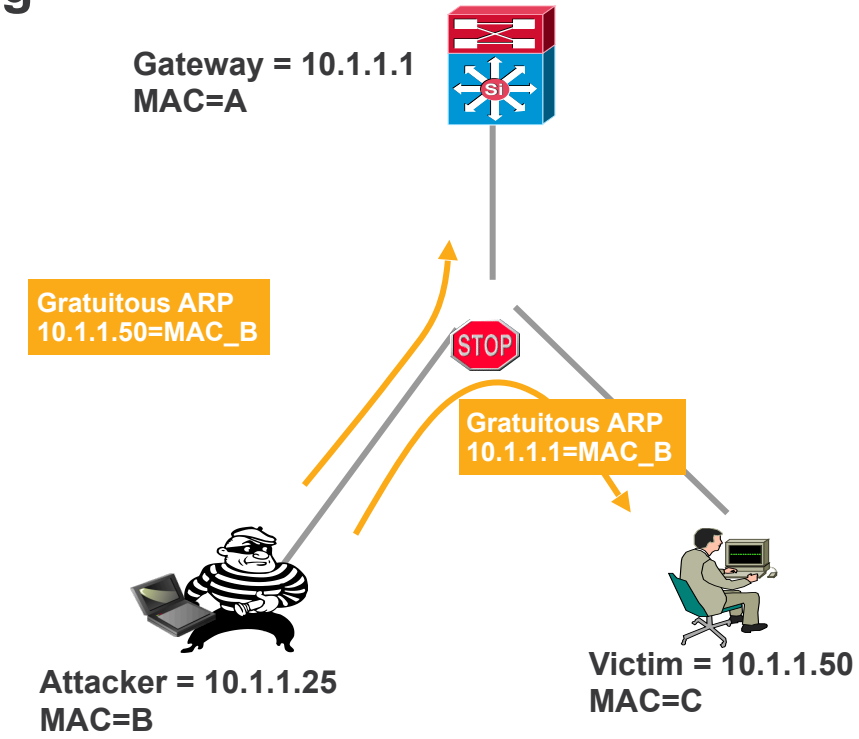

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Securing Layer 2
From Surveillance Attacks

Protection Against **ARP Poisoning –**

- **Dynamic ARP inspection** protects against ARP poisoning (ettercap, dsnif, arpspoof)
- Uses the DHCP Snooping binding table
- Tracks MAC to IP from DHCP transactions
- Rate-limits ARP requests from client ports
- Drop **bogus** gratuitous ARPs – stops ARP poisoning / “Man In The Middle” attacks

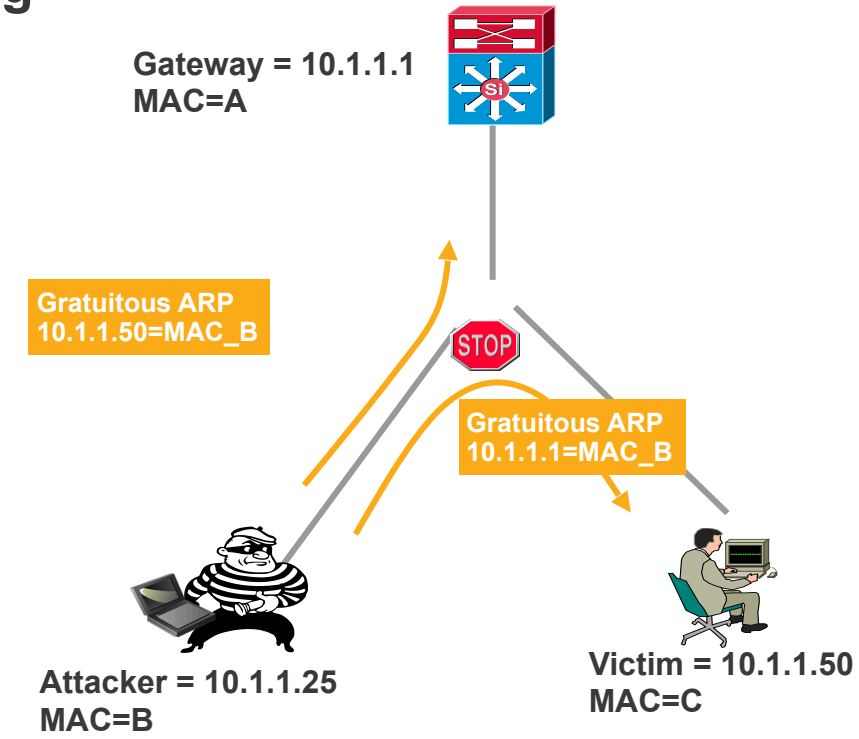

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

IP Source Guard
Protection Against Spoofed IP Addresses

- **IP Source Guard** protects against spoofed IP addresses
- Uses the DHCP Snooping binding table
- Tracks IP address to port associations
- Dynamically programs port ACL to drop traffic not originating from IP address assigned via DHCP

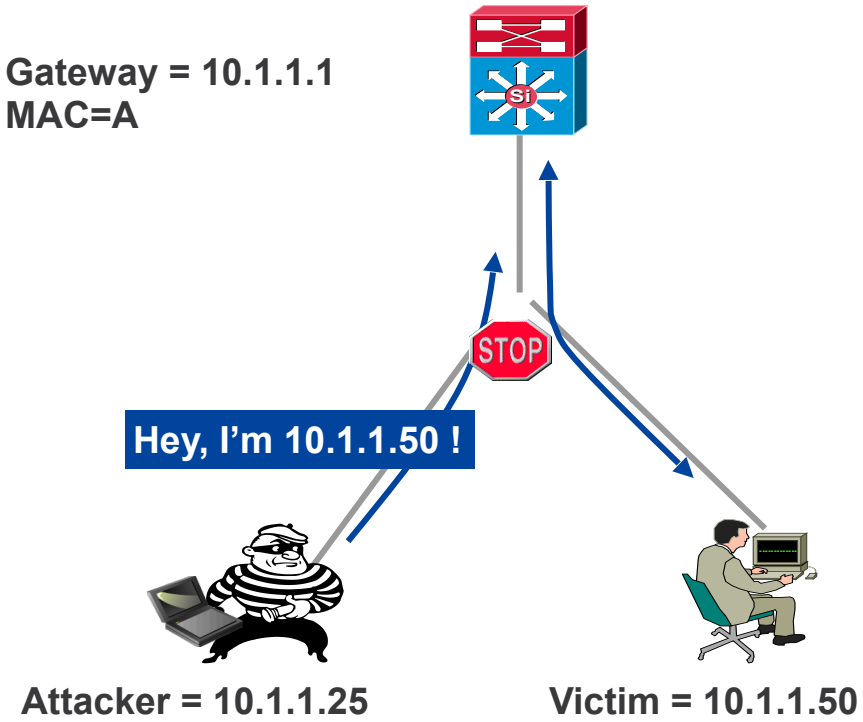

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

IP Source Guard
Protection Against Spoofed IP Addresses

- **IP Source Guard** protects against spoofed IP addresses
- Uses the DHCP Snooping binding table
- Tracks IP address to port associations
- Dynamically programs port ACL to drop traffic not originating from IP address assigned via DHCP

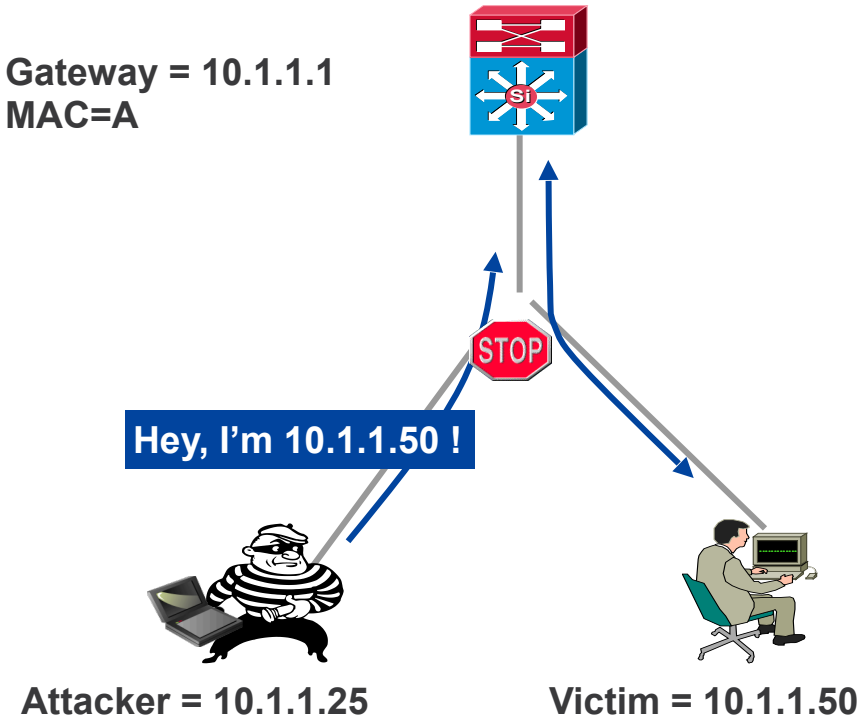

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst Integrated Security Features
Summary

- **Port security** prevents MAC flooding attacks
- **DHCP Snooping** prevents client attack on the switch and server
- **Dynamic ARP Inspection** adds security to ARP using the DHCP snooping table
- **IP Source Guard** adds security to IP source addresses, using the DHCP snooping table

```
ip dhcp snooping
ip dhcp snooping vlan 2-10
ip arp inspection vlan 2-10
!
interface fa3/1
switchport port-security
switchport port-security max 11
switchport port-security violation restrict
switchport port-security aging time 2
switchport port-security aging type inactivity
ip arp inspection limit rate 100
ip dhcp snooping limit rate 100
ip verify source vlandhcp-snooping
!
Interface gigabit1/1
ip dhcp snooping trust
ip arp inspection trust
```


© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Control Plane Protection CoPP

Control Plane Performance

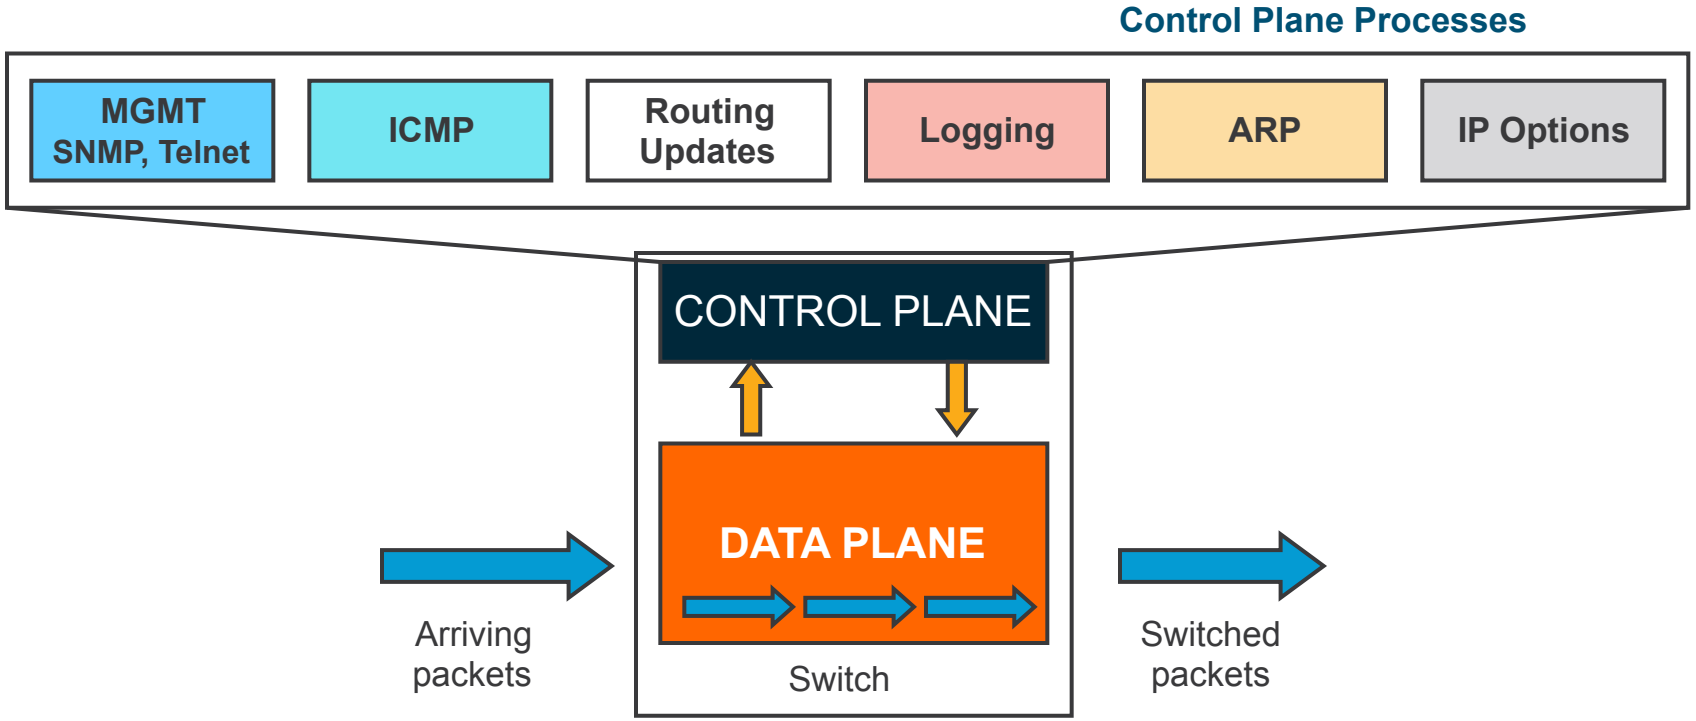

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Control Plane Oversubscription

Control Plane features are processed by the switch CPU so there is a limited amount of processing power available for these tasks. If that CPU is swamped, all other processes stop. This can have a bad impact on the switch if the CPU is overwhelmed for a period of time…

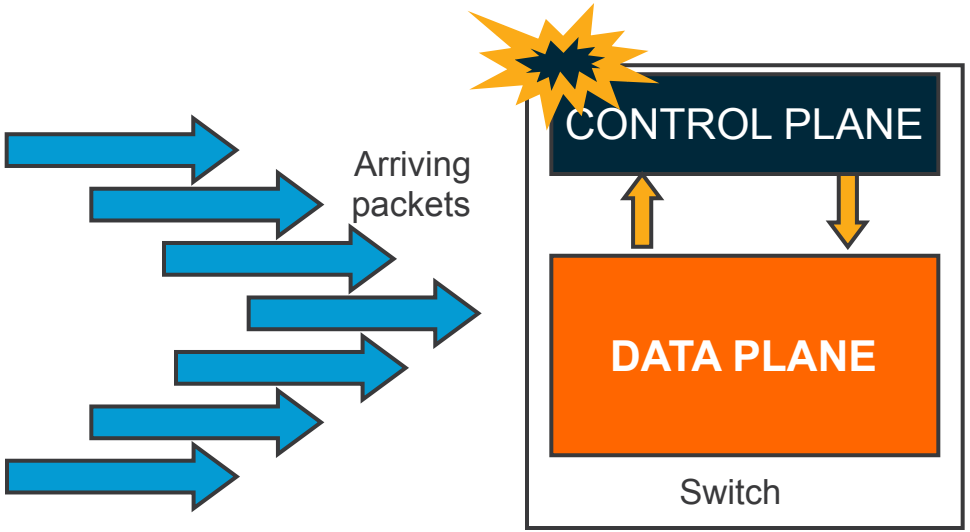

Result of CPU overload

- Dropping Routing Neighbours
- Failure to send Route Updates
- Failure to send STP Updates
- Failure to keep up with Logging requests
- No ARP's processed
- CLI locks up
- Switch locking up
- and more…

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

REFERENCE : Example Protocols and Services Process in Software

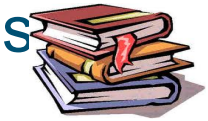

Control Plane Protocols	Control Plane Packet Forwarding
UDLD Protocol	IP Options
PagP Protocol	Fragmentation
LACP Protocol	Select Tunnel Options
SNMP Protocol	ICMP Packets
Syslog Export	MTU failure
Netflow & Netflow Data Export	TTL=1 or TTL=0
Address Resolution Protocol	Packets with Checksum error or error length
HSRP, VRRP, GLBP	RPF Check
Cisco Discovery Protocol	Packets that require ARP resolution
VLAN Trunking Protcol	Non-IP (IPX, Appletalk)
Dynamic Trunking Protocol	ACL logging
Telnet, IP Sec, SSH	Broadcast traffic denied in RACL
BGP, OSPF, EIGRP, RIP, ISIS	Authentication Proxy
Web Cache Control Protocol	PBR traffic for certain “match” or “set” arguments

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Control Plane Protection on Catalyst Switches

	Catalyst 6500 / 6800 Series	Catalyst 4500 Series	Catalyst 3850 series
Hardware Rate-Limiters	Yes	No	Yes
Modular QoS Class Maps	Yes	Yes	Yes (16.x and newer)

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

SUP2T Control Plane Protection

Hardware Rate Limiters Support

Unicast Rate Limiters	
CEF Receive	Traffic Destined to the Router
CEF Receive Secondary	**Traffic destined to an IP address terminated on the C6500**
CEF Glean	Traffic requiring ARP
CEF No Route	Packets with Not Route in the FIB
ICMP Redirect	Packets that Require ICMP Redirects
IP Errors	Packet with IP Checksum or Length Errors
ICMP No Route	ICMP Unreachables for Unroutable Packets
ICMP ACL Drop	ICMP Uncreachables for Admin Deny Packets
RPF Failure	Packets that Fail uRPF Check
L3 Security	CBAC, Auth-Proxy, and IPSEC Traffic
ACL Bridged In	NAT, TCP Int, Reflexive ACLs, Log on ACLs
ACL Bridged Out	NAT, TCP Int, Reflexive ACLs, Log on ACLs
ARP Inspection	Dynamic ARP Inspection Traffic to CPU
DHCP Snoop In	DHCP Snooping Traffic to CPU
IP Features	Security Features (Auth-Proxy,IP Sec, others)
UCAST UNKNOWN FLOOD	L2 unknown unicast traffic
VACL Logging	CLI Notification of VACL Denied Packets
IP Options	Unicast Traffic with IP Options Set
Capture	Used with Optimised ACL Logging

Layer 2 Rate Limiters	
LAYER_2 PT	L2PT Encapsulation/Decapsulation
LAYER_2 PDU	Layer 2 PDUs
MAC PBF In	
IP Admis. on L2 Port	
LAYER_2 PORTSEC	
LAYER_2 SPAN PCAP	

General Rate Limiters	
MTU Failure	Packets Requiring Fragmentation
TTL Failure	Packets with TTL<=1
Capture Pkt	Limits packets punted to the CPU because of Optimised ACL Logging (OAL).
DIAG RESERVED 0	Reserved
DIAG RESERVED 1	Reserved
DIAG RESERVED 2	Reserved
MCAST REPL RESERVED	Reserved

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Control Plane Protection Comparison Sup720 and Sup2T

Feature/Capability	Sup720	Sup2T
Uses Modular QoS CLI for applying policies to the Control Plane Interface	Yes	Yes
Special Case Rate-Limiters for non-IP traffic types (HWRL)	Yes	Yes
Max number of HWRLs supported	8 (L3) 4 (L2)	31 (L3) 26 (L2)
Distributed policing: synchronisation of hardware policers on different linecards	No	Yes
HWRL measured in Packets and Bits per second	No, PPS only	Yes, PPS and BPS
Traffic Counters	CoPP only	Yes, both HWRL and CoPP
Traffic Exceptions configurable in CoPP (allow CoPP based counters and monitoring)	No	Yes
Full suite of Multicast HWRL	No	Yes
Microflow Policers via CoPP	No	Yes
Ability to leak the first packet above the rate (allows packet capture)	No	Yes
CoPP Polices to match ARP and RARP traffic	No	Yes

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

SUP2T Control Plane Protection

Hardware Rate Limiters Support

Multicast Rate Limiters	
MCAST IPV4 FIB MISS	Packets with No mroute in the FIB
MCAST IPv4 IGMP	IGMP Packets
MCAST IPv4 Direct C	Local Multicast on Connected Interface
MCAST IPV4 OPTIONS	Multicast Traffic with IP Options Set
MCAST IPV4 CONTROL PK	IPv4 Multicast Control Traffic to CPU
MCAST IPV6 DIRECTLY C	Packets with No Mroute in the FIB
MCAST IPV6 MLD	IPv6 Multicast Control Traffic to CPU
MCAST IPV6 CONTROL PK	Partial Shortcut Entries
MCAST BRG FLD IP CNTR	Partial Shortcut Entries
MCAST BRG FLD IP	Partial Shortcut Entries
MCAST BRG	Partial Shortcut Entries
MCAST BRG OMF	Partial Shortcut Entries
V6 Route Control	Partial Shortcut Entries
V6 Default Route	Multicast Traffic with IP Options Set
V6 Second Drop	Mulicast Traffic with IP Options Set

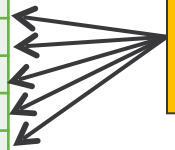

These HWRL were all covered as a single HWRL in the PFC3

Some new HWRL are added and also more granularity to existing HWRLs is provided in the PFC4 versus PFC3

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst 3850 Built in Rate Limiters (IOS-XE 3.x)

```
C3850#show platform qos queue stats internal cpu policer
For Asic 0
 Queue                       Enabled  Rate(default)  Rate(set)     Drop
-----------------------------------------------------------------------
DOT1X Auth                     No       1000          1000          0
L2 Control                     No       500           500           0
Forus traffic                  No       1000          1000          0
ICMP GEN                       Yes      200           200           0
Routing Control                No       500           500           0
Forus Address resolution       No       1000          1000          0
ICMP Redirect                  No       500           500           0
WLESS PRI-5                    No       1000          1000          0
WLESS PRI-1                    No       1000          1000          0
WLESS PRI-2                    No       1000          1000          0
WLESS PRI-3                    No       1000          1000          0
WLESS PRI-4                    No       1000          1000          0
BROADCAST                      Yes      200           200           0
Learning cache ovfl            Yes      100           100           0
Sw forwarding                  Yes      1000          1000          0
Topology Control               No       13000         13000         0
Proto Snooping                 No       500           500           0
DHCP Snooping                  No       1000          1000          0
Transit Traffic                Yes      500           500           0
RPF Failed                     Yes      100           100           0
MCAST END STATION              Yes      2000          2000          0
LOGGING                        Yes      1000          1000          0
Punt Webauth                   No       1000          1000          0
Crypto Control                 Yes      100           100           0
Exception                      Yes      100           100           0
General Punt                   No       500           500           0
NFL SAMPLED DATA               Yes      100           100           0
SGT Cache Full                 Yes      1000          1000
```


© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst 3850 Default Class-maps for Control Plane Policing (IOS-XE 16.x)

```
!
class-map match-any system-cpp-police-topology-control
  description Topology control
class-map match-any system-cpp-police-sw-forward
  description Sw forwarding, SGT Cache Full, LOGGING
class-map match-any system-cpp-default
  description DHCP snooping, show forward and rest of traffic
class-map match-any system-cpp-police-sys-data
  description Learning cache ovfl, Crypto Control, Exception, EGR
Exception, NFL SAMPLED DATA, Gold Pkt, RPF Failed
class-map match-any system-cpp-police-punt-webauth
  description Punt Webauth
class-map match-any system-cpp-police-l2lvx-control
  description L2 LVX control packets
class-map match-any system-cpp-police-forus
  description Forus Address resolution and Forus traffic
class-map match-any system-cpp-police-multicast-end-station
  description MCAST END STATION
class-map match-any system-cpp-police-multicast
  description Transit Traffic and MCAST Data
```

```
class-map match-any system-cpp-police-l2-control
  description L2 control
class-map match-any system-cpp-police-dot1x-auth
  description DOT1X Auth
class-map match-any system-cpp-police-data
  description ICMP_GEN and BROADCAST
class-map match-any system-cpp-police-control-low-priority
  description ICMP redirect and general punt
class-map match-any system-cpp-police-wireless-priority1
  description Wireless priority 1
class-map match-any system-cpp-police-wireless-priority2
  description Wireless priority 2
class-map match-any system-cpp-police-wireless-priority3-4-5
  description Wireless priority 3,4 and 5
class-map match-any non-client-nrt-class
class-map match-any system-cpp-police-routing-control
  description Routing control
class-map match-any system-cpp-police-protocol-snooping
  description Protocol snooping
!
```


© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst 4500 Sup7/Sup8 CoPP Macro and Class-Maps

Switch(config)# qos	Class system-cpp-igmp
Switch(config)# macro global apply system-cpp	Class system-cpp-pim
Switch# show policy-map system-cpp-policy	Class system-cpp-all-systems-on-subnet
Policy Map system-cpp-policy	Class system-cpp-all-routers-on-subnet
Class system-cpp-dot1x	Class system-cpp-ripv2
Class system-cpp-bpdu-range	Class system-cpp-ip-mcast-linklocal
Class system-cpp-cdp	Class system-cpp-dhcp-cs
police 32000 bps 1000 byte conform-action transmit exceed-action drop	Class system-cpp-dhcp-sc
Class system-cpp-sstp	Class system-cpp-dhcp-ss
Class system-cpp-cgmp	* Class telnet-class
Class system-cpp-ospf	police 80000 1000 byte conform-action drop exceed-action drop
Class system-cpp-hsrpv2	

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Sup2T CoPP Default Class-Maps

Default Class-Maps Supported with Sup2T
class-map: class-copp-icmp-redirect-unreachable (match-all)
class-map: class-copp-ucast-rpf-fail (match-all)
class-map: class-copp-vacl-log (match-all)
class-map: class-copp-mcast-punt (match-all)
class-map: class-copp-mcast-copy (match-all)
class-map: class-copp-ip-connected (match-all)
class-map: class-copp-ipv6-connected (match-all)
class-map: class-copp-match-pim-data (match-any)
class-map: class-copp-match-pimv6-data (match-any)
class-map: class-copp-match-mld (match-any)
class-map: class-copp-match-igmp (match-any)
class-map: class-copp-match-ndv6 (match-any)

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

SUP2T Control Plane Protection

Default Class-maps for Control Plane Protection

Using the class-maps allows for better visibility using show commands

```
Router#show policy-map control-plane input class class-copp-options

 Control Plane Interface

  Service-policy input: policy-default-autocopp

  Hardware Counters:

    class-map: class-copp-options (match-all)
      Match: any
      police :
        100 pps 24 limit 24 extended limit
      Earl in slot 1 :
        0 packets
        5 minute offered rate 0 pps
        aggregate-forwarded 0 packets
                          action: transmit
        exceeded 0 packets action: drop
        aggregate-forward 0 pps exceed 0 pps
      Earl in slot 3 :
        997 packets
        5 minute offered rate 196 pps
        aggregate-forwarded 198 packets
                          action: transmit
        exceeded 798 packets action: drop
        aggregate-forward 99 pps exceed 400 pps
      Earl in slot 5 :
        0 packets
        5 minute offered rate 0 pps
```

Default Class map "class-copp-options" classifies traffic with IP Options set

Traffic counters matching the class-map on the DFC in slot3

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst Control Plane Protection

CoPP Deployment - Step 1

Step 1: Identify traffic of interest and classify it into multiple traffic classes:

- BGP
- IGP (EIGRP, OSPF, ISIS)
- Management (telnet, TACACS, ssh, SNM NTP)
- Reporting (SAA)
- Monitoring (ICMP)
- Critical applications (HSRP, DHCP)
- Undesirable
- Default

```
ip access-list extended coppacl-bgp
permit tcp host 192.168.1.1 host 10.1.1.1 eq bgp
permit tcp host 192.168.1.1 eq bgp host 10.1.1.1
!
ip access-list extended coppacl-igp
permit ospf any host 224.0.0.5
permit ospf any host 224.0.0.6
permit ospf any any
!
ip access-list extended coppacl-management
permit tcp host 10.2.1.1 host 10.1.1.1 established
permit tcp 10.2.1.0 0.0.0.255 host 10.1.1.1 eq 22
permit tcp 10.86.183.0 0.0.0.255 any eq telnet
permit udp host 10.2.2.2 host 10.1.1.1 eq snmp
permit udp host 10.2.2.3 host 10.1.1.1 eq ntp
!
```


© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst Control Plane Protection

CoPP Deployment - Step 2

Step 2: Associate the identified traffic with a class, and permit the traffic in each class – Start with a generous amount of allowed traffic

- Must enable QoS globally(Sup720 only), else CoPP will not be applied in hardware
- Always apply a policing action for each class since the switch will ignore a class that does not have a corresponding policing action (for example "police 31500000 conform-action transmit exceed-action drop").
- Alternatively, both conform-action and exceed-action could be set to transmit, but doing so will allocate a default policer as opposed to a dedicated policer with its own hardware counters.
- HW CoPP classes are limited to one match per class-map

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst Control Plane Protection

CoPP Deployment - Step 2

Step 2: Associate the identified traffic with a class, and permit the traffic in each class – Start with a generous amount of allowed traffic

- Must enable QoS globally(Sup720 on

```
class-map match-all copp-bgp
 match access-group name coppacl-bgp
class-map match-all copp-igp
 match access-group name coppacl-igp
class-map match-all copp-management
 match access-group name coppacl-management
class-map match-all copp-reporting
 match access-group name coppacl-reporting
class-map match-all copp-monitoring
 match access-group name coppacl-monitoring
class-map match-all copp-critical-app
 match access-group name coppacl-critical-app
class-map match-all copp-undesirable
 match access-group name coppacl-undesirable
```

```
policy-map copp-policy
 class copp-bgp
  police 30000000 conform-action transmit exceed-action drop
 class copp-igp
  police 30000000 conform-action transmit exceed-action drop
 class copp-management
  police 30000000 conform-action transmit exceed-action drop
 class copp-reporting
  police 30000000 conform-action transmit exceed-action drop
 class copp-monitoring
  police 30000000 conform-action transmit exceed-action drop
 class copp-critical-app
  police 30000000 conform-action transmit exceed-action drop
 class copp-undesirable
  police 30000000 conform-action transmit exceed-action drop
 class class-default
  police 30000000 conform-action transmit exceed-action drop
```

```
control-plane
 service-policy input copp-policy
```


© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst Control Plane Protection

CoPP Deployment - Step 3

Step 3: Adjust classification, and apply liberal CoPP policies for each class of traffic

- show policy-map control-plane displays dynamic information for monitoring control plane policy.
- Statistics include rate information and number of packets/ bytes confirmed or exceeding each traffic class

```
Switch# show policy-map control-plane
 Control Plane Interface
 Service-policy input: copp-policy
<snip>
  Hardware Counters:
  class-map: copp-monitoring (match-all)
   Match: access-group name coppacl-monitoring
   police :
     30000000 bps 937000 limit 937000 extended limit
    Earl in slot 5 :
     0 bytes
     5 minute offered rate 0 bps
     aggregate-forwarded 0 bytes action: transmit
     exceeded 0 bytes action: drop
     aggregate-forward 0 bps exceed 0 bps
    Earl in slot 7 :
     112512 bytes
     5 minute offered rate 3056 bps
     aggregate-forwarded 112512 bytes action: transmit
     exceeded 0 bytes action: drop
     aggregate-forward 90008 bps exceed 0 bps

==========
```


© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst Control Plane Protection

CoPP Deployment—Step 3 (Cont.)

Step 3: Adjust Classification, and Apply liberal CoPP policies for each class of traffic

- show ip access-lists provides packet count statistics per ACE. Absence of any hits on an entry indicate lack of traffic matching the ACE criteria—the rule might be rewritten
- Hardware ACL hit counters are available in PFC3B/BXL and newer for security ACL TCAM only (not QoS ACL TCAM)

```
Switch#show access-list
Extended IP access list coppacl-bgp
  10 permit tcp host 192.168.1.1 host 10.1.1.1 eq bgp
  20 permit tcp host 192.168.1.1 eq bgp host 10.1.1.1
Extended IP access list coppacl-critical-app
  10 permit ip any host 224.0.0.1
  20 permit udp host 0.0.0.0 host 255.255.255.255 eq bootps
  30 permit udp host 10.2.2.8 eq bootps any eq bootps
Extended IP access list coppacl-igp
  10 permit ospf any host 224.0.0.5 (64062 matches)
  20 permit ospf any host 224.0.0.6
  30 permit ospf any any (17239 matches)
Extended IP access list coppacl-management
  10 permit tcp host 10.2.1.1 host 10.1.1.1 established
  20 permit tcp 10.2.1.0 0.0.0.255 host 10.1.1.1 eq 22
  30 permit tcp 10.86.183.0 0.0.0.255 any eq telnet
  40 permit udp host 10.2.2.2 host 10.1.1.1 eq snmp
  50 permit udp host 10.2.2.3 host 10.1.1.1 eq ntp
Extended IP access list coppacl-monitoring
  10 permit icmp any any ttl-exceeded (120 matches)
  20 permit icmp any any port-unreachable
  30 permit icmp any any echo-reply (17273 matches)
  40 permit icmp any any echo (5 matches)
Extended IP access list coppacl-reporting
  10 permit icmp host 10.2.2.4 host 10.1.1.1 echo
Extended IP access list coppacl-undesirable
  10 permit udp any any eq 1434
```


© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst Control Plane Protection

CoPP Deployment - Step 4

Step 4: Fine tune the control plane policy

- Narrow the ACL permit statements to only allow known authorised source addresses and depending on class defined, apply appropriate policy
  - Routing protocol traffic — no rate limit or very conservative rate limit
  - Management traffic — conservative rate limit
  - Reporting traffic — conservative rate limit
  - Monitoring traffic — conservative rate limit
  - Critical traffic — conservative rate limit
  - Default traffic — low rate limit
  - Undesirable traffic — drop

http://www.cisco.com/en/US/products/hw/switches/ps708/products_white_paper0900aecd802ca5d6.shtml

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst Control Plane Protection

CoPP Deployment - Step 4

Step 4: Fine tune the control plane policy

- Narrow the ACL permit stat… addresses and depending …
  - Routing protocol traffic — no …
  - Management traffic — conse…
  - Reporting traffic — conservat…
  - Monitoring traffic — conserva…
  - Critical traffic — conservative…
  - Default traffic — low rate limit…
  - Undesirable traffic — drop

```
policy-map copp-policy
 class coppclass-bgp
  police 15000000 conform-action transmit exceed-action drop
 class coppclass-igp
  police 15000000 conform-action transmit exceed-action drop
 class coppclass-management
  police 2560000 conform-action transmit exceed-action drop
 class coppclass-reporting
  police 1000000 conform-action transmit exceed-action drop
 class coppclass-monitoring
  police 1000000 conform-action transmit exceed-action drop
 class coppclass-critical-app
  police 7500000 conform-action transmit exceed-action drop
 class coppclass-undesirable
  police 32000 conform-action transmit exceed-action drop
 class class-default
  police 1000000 conform-action transmit exceed-action drop
```

http://www.cisco.com/en/US/products/hw/switches/ps708/products_white_paper0900aecd802ca5d6.shtml

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Catalyst Control Plane Protection

Mitigating attacks with CoPP and CPU RL (Example)

- Multiple concurrent attacks (multicast ttl=1, multicast partial shortcuts, unicast IP options, unicast fragments to receive adjacency, unicast TCP SYN flood to receive adjacency)
- CPU kept within acceptable bounds with no loss of mission critical traffic

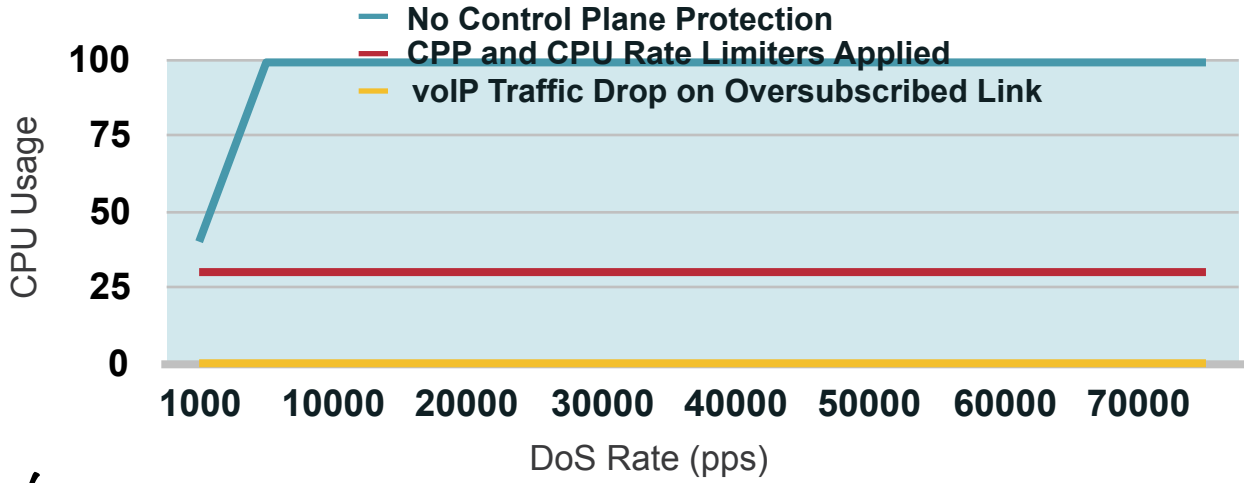

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Alternative Designs

Simplified Layer 2 Access with VSS Distribution

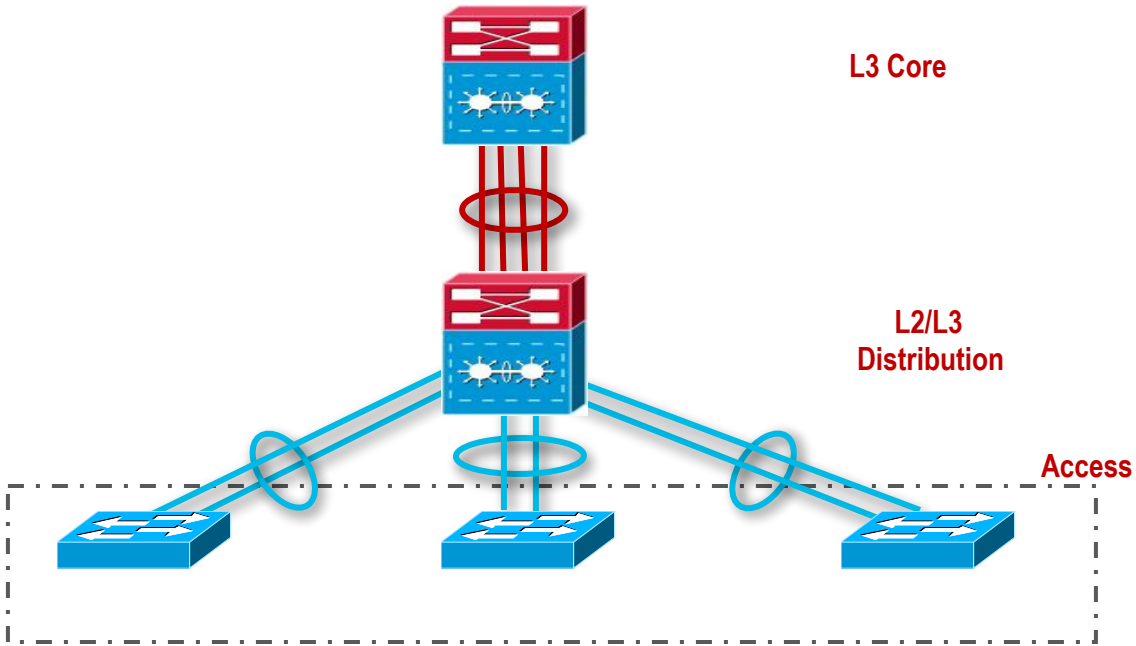

Reduced routing neighbours, Minimal L3 reconvergence

No FHRPs
No Looped topology
Policy Management

Multiple active uplinks per VLAN, No STP convergence

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Simplified Layer 2 Access with VSS Distribution

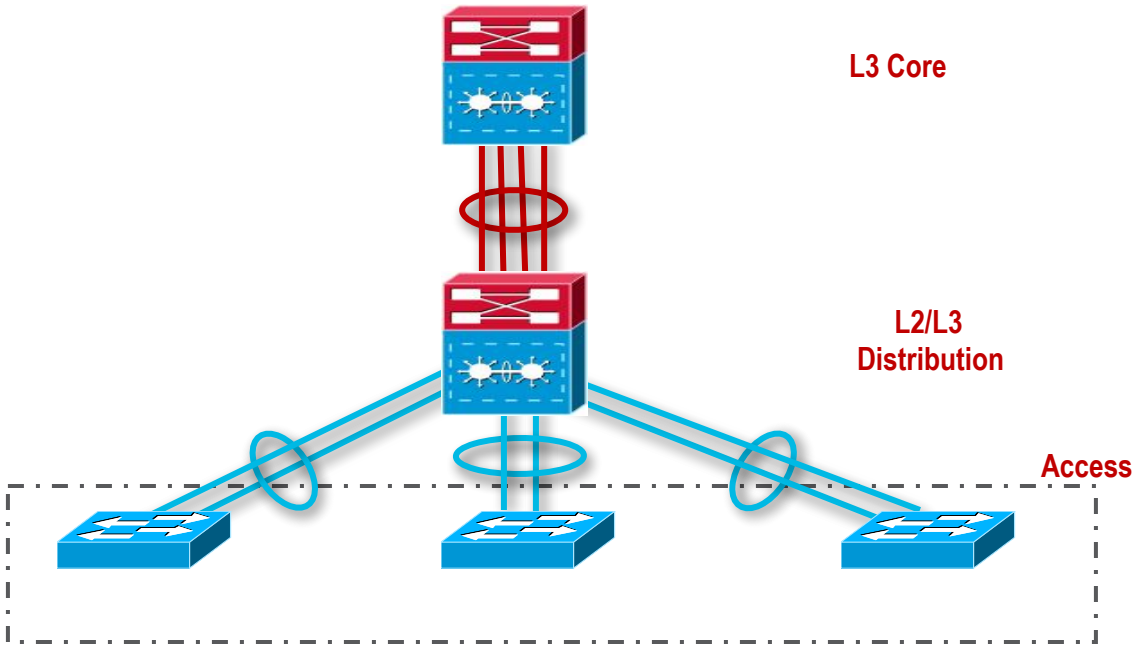

Reduced routing neighbours, Minimal L3 reconvergence

No FHRPs
No Looped topology
Policy Management

Multiple active uplinks per VLAN, No STP convergence

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Simplified Layer 2 Access with VSS Distribution

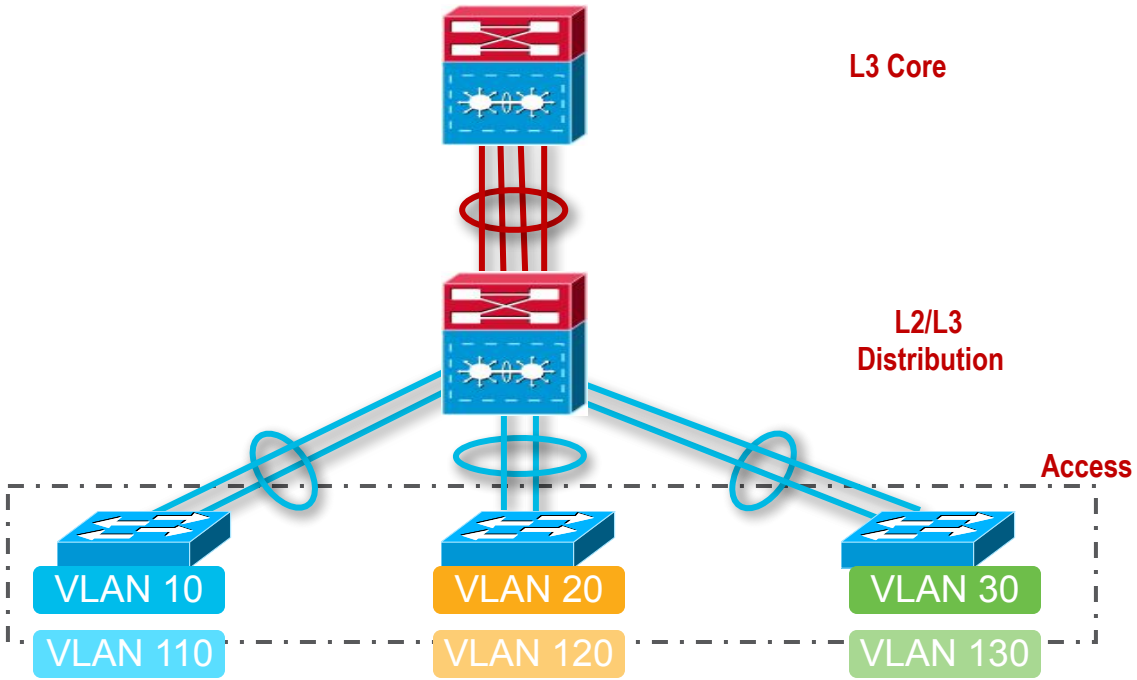

Reduced routing neighbours, Minimal L3 reconvergence

No FHRPs
No Looped topology
Policy Management

Multiple active uplinks per VLAN, No STP convergence

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Simplified Layer 2 Access with VSS Distribution

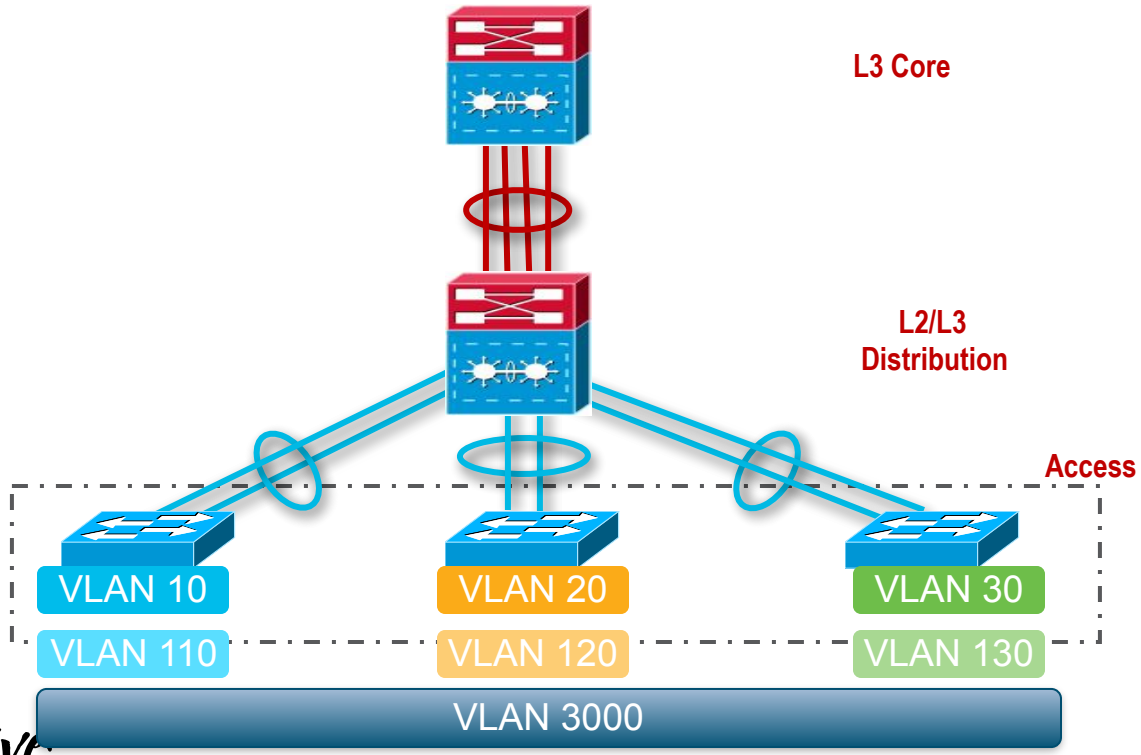

Reduced routing neighbours, Minimal L3 reconvergence

No FHRPs
No Looped topology
Policy Management

Multiple active uplinks per VLAN, No STP convergence

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

VSS Simplifies the Configuration

Standalone Switch 1 (Coordinated Configuration)	Standalone Switch 2 (Coordinated Configuration)	VSS (One simplified configuration)
Spanning Tree Configuration		
! Enable 802.1d per VLAN spanning tree enhancements. spanning-tree mode pvst spanning-tree loopguard default no spanning-tree optimise bpdu transmission spanning-tree extend system-id spanning-tree uplinkfast spanning-tree backbonefast spanning-tree vlan 2,4,6,8,10 priority 24576!	! Enable 802.1d per VLAN spanning tree enhancements. spanning-tree mode pvst spanning-tree loopguard default no spanning-tree optimise bpdu transmission spanning-tree extend system-id spanning-tree uplinkfast spanning-tree backbonefast spanning-tree vlan 3,5,7,9,11 priority 24576!	! Enable 802.1d per VLAN spanning tree enhancements spanning-tree mode rapid-pvst no spanning-tree optimise bpdu transmission spanning-tree extend system-id spanning-tree vlan 2-11 priority 24576
L3 SVI Configuration (sample for 1 VLAN)		
! Define the Layer 3 SVI for each voice and data VLAN interface Vlan4 description Data VLAN ip address 10.120.4.3 255.255.255.0 no ip redirects no ip unreachables ! Reduce PIM query interval to 250 msec ip pim query-interval 250 msec ip pim sparse-mode load-interval 30 ! Define HSRP default gateway with 250/800 msec hello/hold standby 1 ip 10.120.4.1 standby 1 timers msec 250 msec 800 ! Set preempt delay large enough to allow network to stabilise before HSRP ! switches back on power on or link recovery standby 1 preempt delay minimum 180 ! Enable HSRP authentication standby 1 authentication cisco123	! Define the Layer 3 SVI for each voice and data VLAN interface Vlan4 description Data VLAN ip address 10.120.4.3 255.255.255.0 no ip redirects no ip unreachables ! Reduce PIM query interval to 250 msec ip pim query-interval 250 msec ip pim sparse-mode load-interval 30 ! Define HSRP default gateway with 250/800 msec hello/hold standby 1 ip 10.120.4.1 standby 1 timers msec 250 msec 800 ! Set preempt delay large enough to allow network to stabilise before HSRP ! switches back on power on or link recovery standby 1 preempt delay minimum 180 ! Enable HSRP authentication standby 1 authentication cisco123	! Define the Layer 3 SVI for each voice and data VLAN interface Vlan4 description Data VLAN ip address 10.120.2.1 255.255.255.0 no ip redirects no ip unreachables ip pim sparse-mode load-interval 30

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco Digital Network Architecture

Insights and Experiences

Automation and Assurance

Security and Compliance

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Cisco Digital Network Architecture

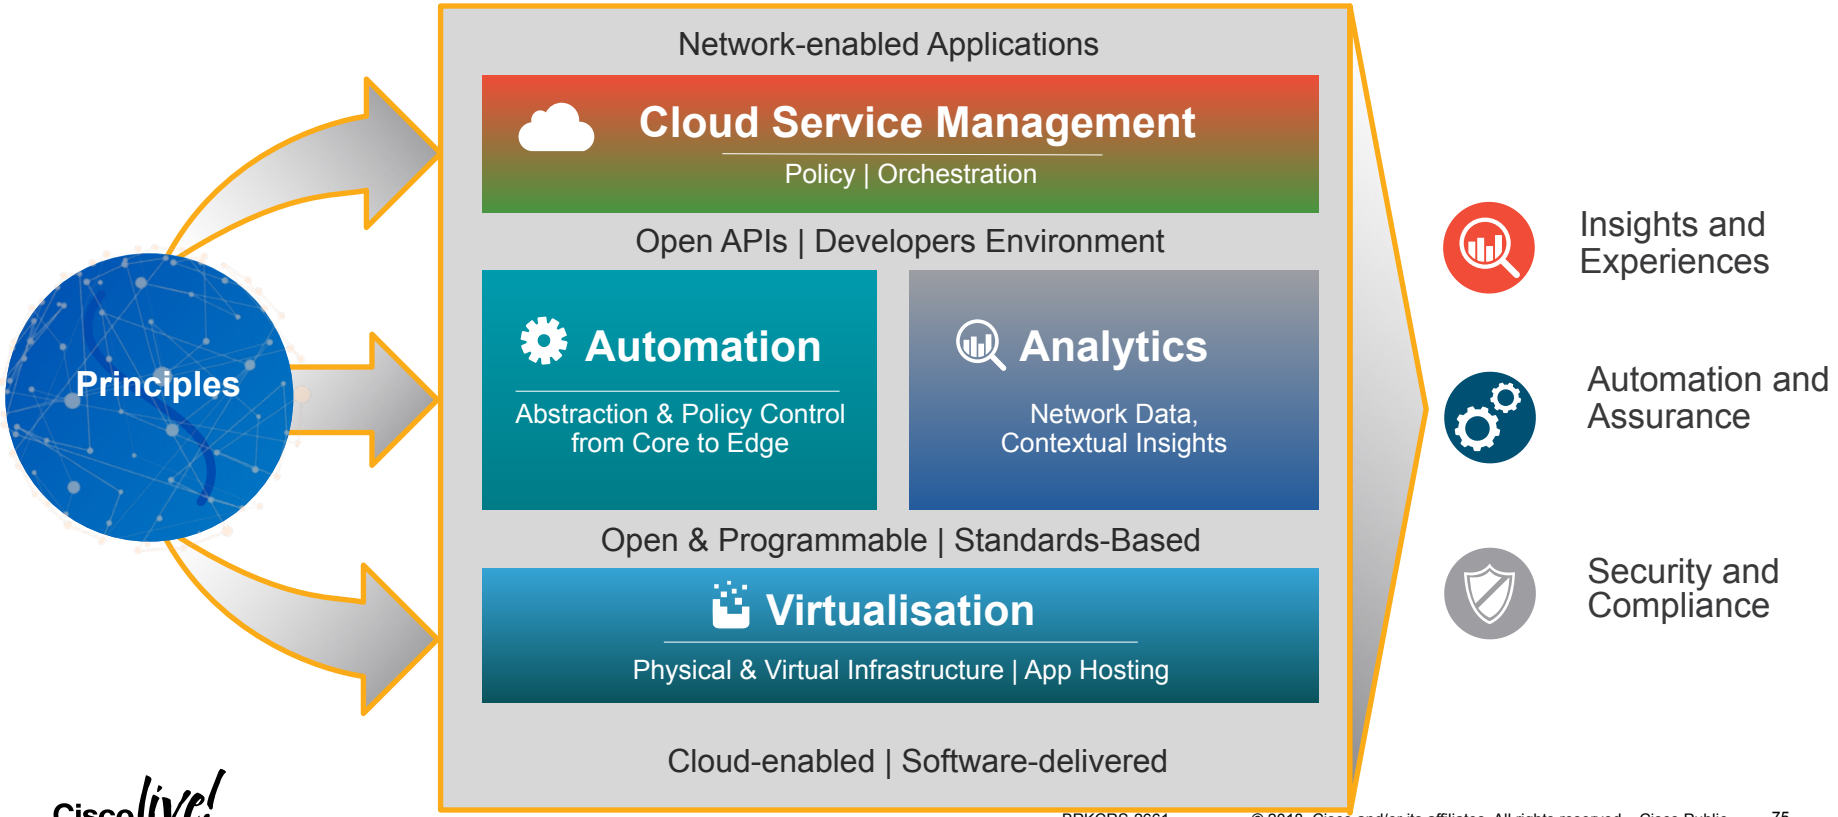

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

What exactly is a Fabric?

A **Fabric** is an **Overlay**

An “Overlay” is a *logical topology* used to *virtually connect* devices, built *on top* of an arbitrary physical “Underlay” topology.

An “Overlay” network often uses *alternate forwarding attributes* to provide *additional services*, not provided by the “Underlay”.

Examples of Network Overlays

- GRE or mGRE
- MPLS or VPLS
- IPSec or DMVPN
- CAPWAP
- LISP
- OTV
- DFA
- ACI

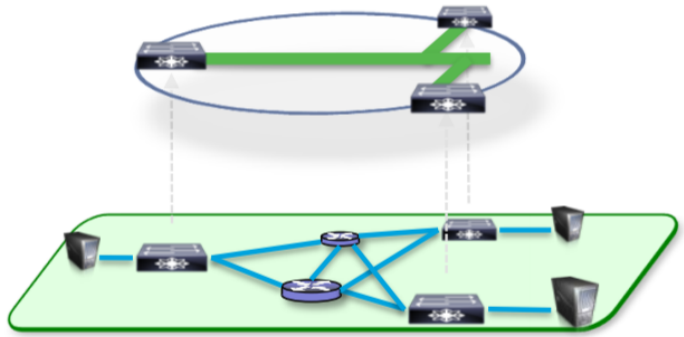

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Where To Get More Information

- Cisco Live Sessions in Melbourne
  - BRKCRS 1800 DNA Campus Fabric – An Introduction – Wednesday 8:30am
  - BRKCRS 3800 DNA Campus Fabric – Deep Dive – Friday 8:45am
  - BRKCRS 2801 DNA Campus Fabric – How to Integrate With Your Existing Network – Thursday 4:30pm
  - BRKCRS 2031 Enterprise Campus Design – Multilayer Architecture and Design Principals – Wednesday 2:30PM
- Cisco Validated Designs
  - Campus LAN and Wireless LAN Design Guide http://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/Oct2016/CVD-Campus-LAN-WLAN-Design-2016OCT.pdf
  - Campus LAN Layer 2 Access with Simplified Distribution Deployment Guide http://www.cisco.com/c/dam/en/us/td/docs/solutions/CVD/Oct2015/CVD-Campus_LAN_L2_Access_Simplified_Dist_Deployment-Oct2015.pdf

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Summary

Designing Layer 2 Networks is Easy!

- Limit the size of the L2 domain as much as possible
- Use L1-L2 best practices to harden the network and eliminate common causes of loops
- Implement Spanning Tree tool kit “Make the network fail closed!”
- Use the Integrated Security tool kit to harden the network form malicious or non-malicious network events
- Harden the control plane of the network devices with Control Plane Protection tools
- Consider alternate designs that minimise L2 loops

© 2018 Cisco and/or its affiliates. All rights reserved. Cisco Public

Q & A

Complete Your Online Session Evaluation

- Give us your feedback and receive a **Cisco Live 2018 Cap** by completing the overall event evaluation and 5 session evaluations.
- All evaluations can be completed via the Cisco Live Mobile App.

Don’t forget: Cisco Live sessions will be available for viewing on demand after the event at www.CiscoLive.com/Global.

Ciscolive!

Thank you

cisco
You're
Cisco live!