Switching Architectures: L2 Protection Features

Campus Network Design & Operations Workshop

UNIVERSITY OF OREGON

Last updated 1st October 2024

NSRC
Network Startup Resource Center

Other Layer 2 Features

- Link Aggregation
- Network Protection
- Switch Configuration Advice:
  - Network Management
  - Documentation

Link Aggregation

- Also known as *port bundling, link bundling*
- You can use multiple links in parallel as a single, logical link
  - For increased capacity
  - For redundancy (fault tolerance)

Link Aggregation

- Also known as *port bundling, link bundling*
- You can use multiple links in parallel as a single, logical link
  - For increased capacity
  - For redundancy (fault tolerance)
- LACP (Link Aggregation Control Protocol) is a standardized method (802.1AX) of negotiating these bundled links between switches

Link Aggregation

- Also known as *port bundling, link bundling*
- You can use multiple links in parallel as a single, logical link
  - For increased capacity
  - For redundancy (fault tolerance)
- LACP (Link Aggregation Control Protocol) is a standardized method (802.1AX) of negotiating these bundled links between switches
- Proprietary methods exist too (Cisco's PAgP, EtherChannel; Juniper's Aggregated Ethernet, etc)

LACP Operation

- Two switches connected via multiple links will send LACPDU packets, identifying themselves and the port capabilities

LACP Operation

- Two switches connected via multiple links will send LACPDU packets, identifying themselves and the port capabilities
- They will then automatically build the logical aggregated links, and then pass traffic.

LACP Operation

- Two switches connected via multiple links will send LACPDU packets, identifying themselves and the port capabilities
- They will then automatically build the logical aggregated links, and then pass traffic.
- Switch ports can be configured as active or passive

LACP Operation

- Two switches connected via multiple links will send LACPDU packets, identifying themselves and the port capabilities
- They will then automatically build the logical aggregated links, and then pass traffic.
- Switch ports can be configured as active or passive
- Software implementations of LACP exist allowing Linux, BSD servers etc. to combine ports as well. Useful for things like NAS devices.

LACP Operation

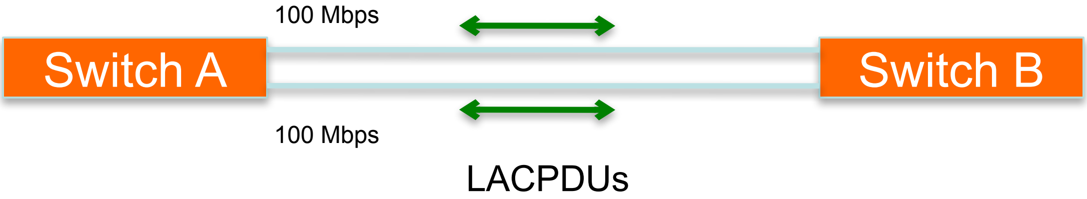

- Switches A and B are connected to each other using two sets of Fast Ethernet ports
- LACP is enabled and the ports are turned on
- Switches start sending LACPDUs, then negotiate how to set up the aggregation

LACP Operation

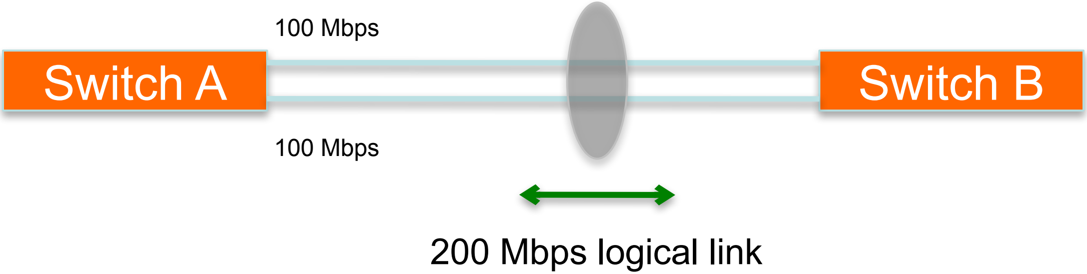

- The result is an aggregated 200 Mbps logical link
- The link is also fault tolerant: If one of the member links fail, LACP will automatically take that link off the bundle, and keep sending traffic over the remaining link

Distributing Traffic in Bundled Links

- Bundled links distribute frames using a hashing algorithm, based on:
  - Source and/or Destination MAC address
  - Source and/or Destination IP address
  - Source and/or Destination Port numbers
- This can lead to unbalanced use of the links, depending on the nature of the traffic
- Always choose the load-balancing method that provides the most distribution

Questions?

Network Protection Features

- Vendors have introduced features which can protect against certain problems which can occur in your network
  - These are not standardized
  - Vendors often have similar features but with different names
  - We’ll show the Cisco names here

Problem 1: Rogue switches

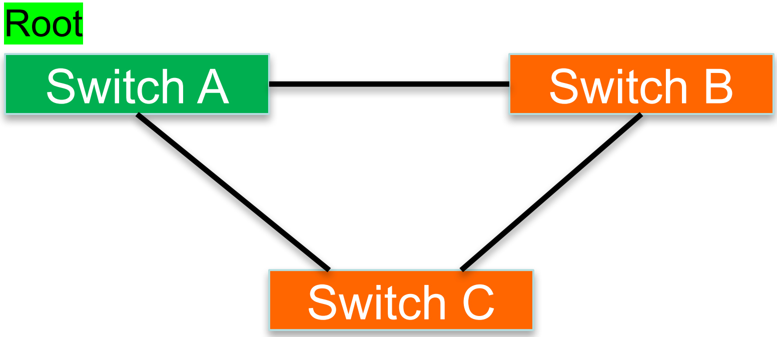

Problem 1: Rogue switches

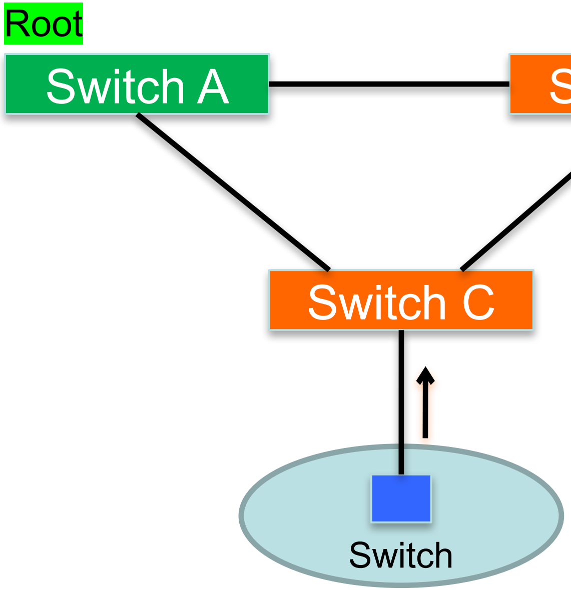

What happens if someone plugs in their own switch into one of your edge ports – and this switch has a lower root bridge priority?

Solution 1a: “Root Guard”

- Enable “Root Guard” on edge ports
- Switch can still be plugged in, and can participate in STP
- However, if it ever tries to become root, the port is shut down
  - Error condition must be cleared manually, unless you configure automatic recovery (*errdisable-timeout*)

Solution 1b: “BPDU Guard”

- Enable “BPDU Guard” on edge ports
  - A more brutal solution!
- If *any* spanning tree BPDU at all is received on this port, the port is immediately shut down
  - Prevents users plugging in their own switches, mostly
  - Does not detect the dumbest, non-STP switches or hubs
  - Does not prevent “connection sharing” at layer 3 (NAT)

Problem 2: Unidirectional links

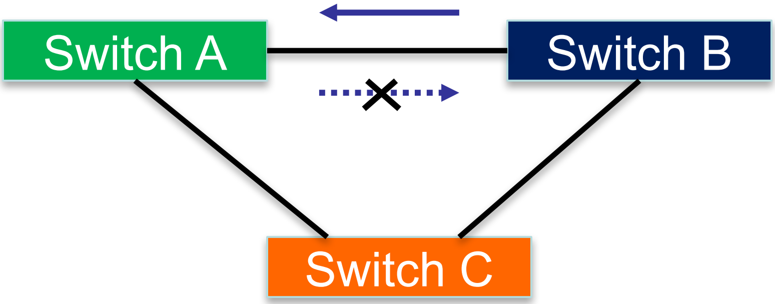

Problem 2: Unidirectional links

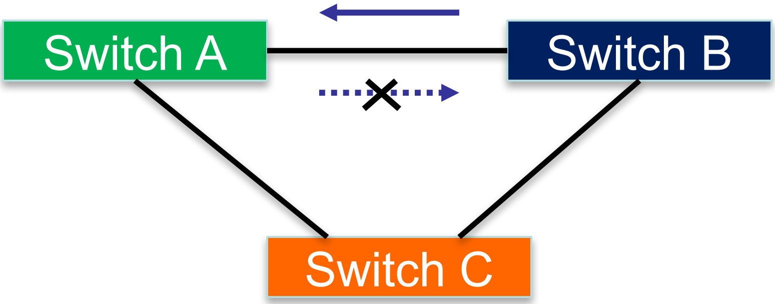

- Switch A can see BPDUs from switch B – but not vice versa
    - Typically, due to faulty leg on a bidirectional fiber link, or mis-patching

Problem 2: Unidirectional links

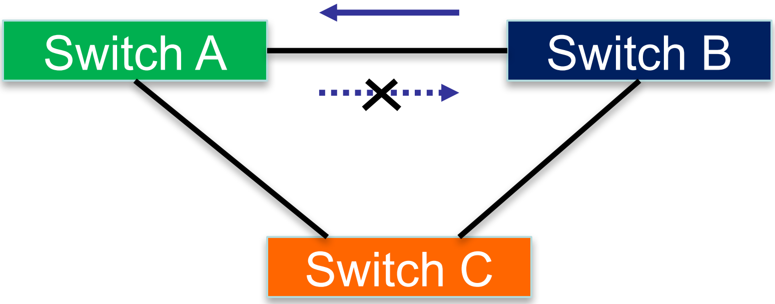

- Switch A can see BPDUs from switch B – but not vice versa
  - Typically, due to faulty leg on a bidirectional fiber link, or mis-patching
- Major STP problems, e.g. two simultaneous roots!
  - Hence loops, broadcast storms etc

Solution 2: “UDLD”

- “Unidirectional Link Detection”
  - Configure at both ends of fiber switch-to-switch trunks
  - Cisco protocol, but some other vendors implement and interoperate
- Sends periodic echo/response packets
- Shuts down link if not working bidirectionally
- “Aggressive mode” gives best protection, but will shut down link if far end doesn’t have UDLD enabled

Problem 3: Failed control plane

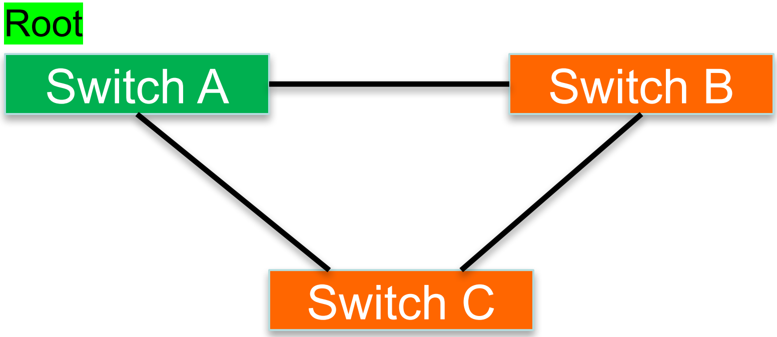

- Switch C is forwarding packets in hardware
  - Hardware forwards STP BPDUs to CPU for processing

Problem 3: Failed control plane

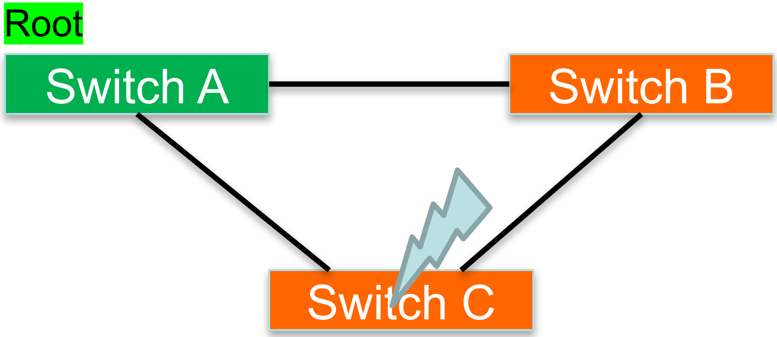

- Switch C is forwarding packets in hardware
  - Hardware forwards STP BPDUs to CPU for processing
- At some point Switch C's CPU locks up
  - Switches A and B no longer see BPDUs from C
  - But data frames are still being forwarded (inc. broadcasts)

Solution 3: “Loop Guard”

- When you *stop* receiving BPDUs on a link, Spanning Tree thinks that there's no longer a switch connected at the far end
  - and therefore, it’s safe to use, not part of a loop
  - but in this case, it *is* part of a loop, so you get a broadcast storm
- Solution: enable “Loop Guard” on switch-to-switch links
- If you *had been* receiving STP BPDUs on a port, but then they stop, it marks the port in a loop-inconsistent state and blocks

Problem 4: Rogue DHCP servers

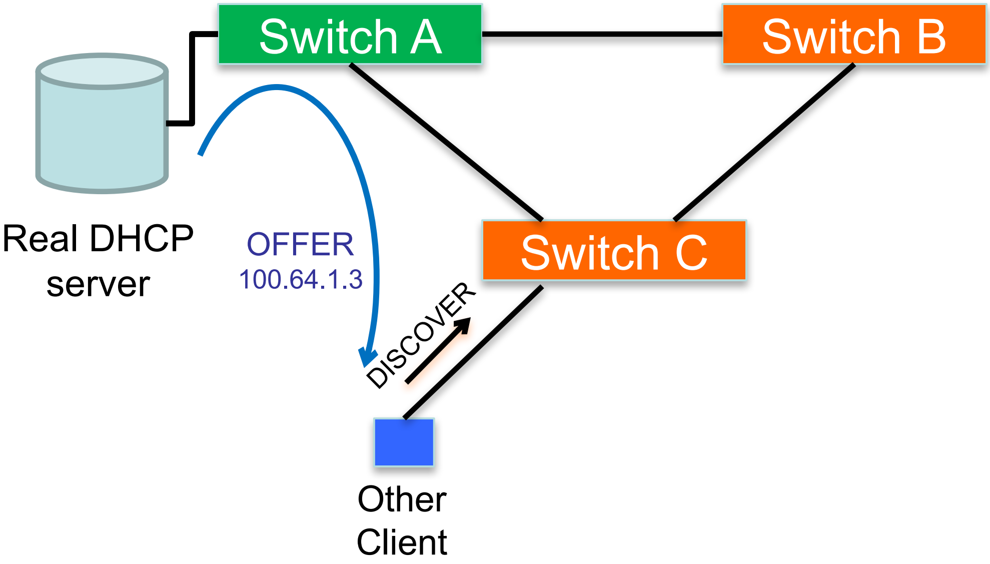

Problem 4: Rogue DHCP servers

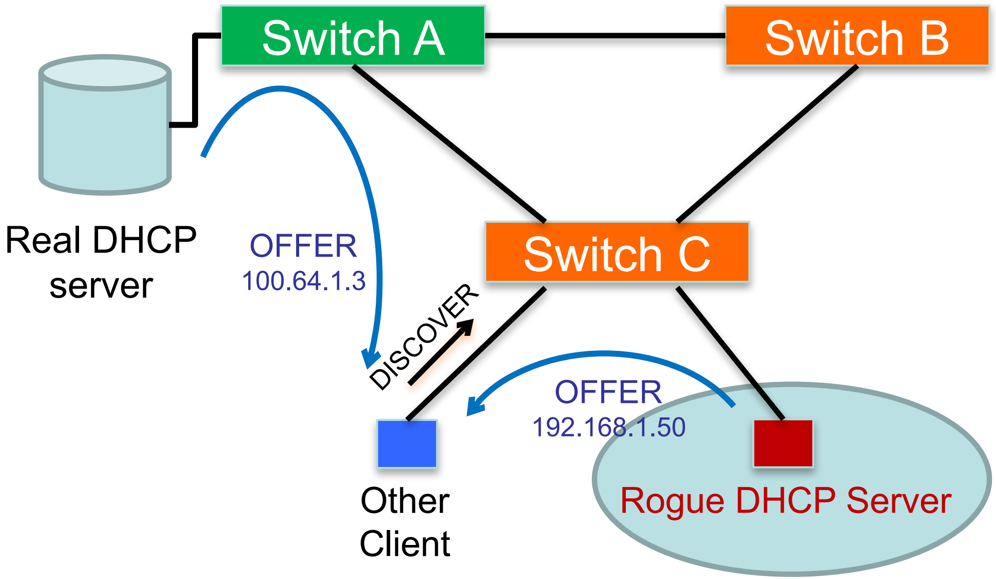

- Someone plugs in a device which acts as a DHCP server

Solution 4: “DHCP snooping”

- This is a common problem, often accidental not malicious
- User plugs in a consumer router/wireless access point using one of its “LAN” ports
  - but these devices contain a DHCP server!
  - client gets two offers and accepts the first one it receives
  - wrong IP = lost connectivity. Affects everyone on the same VLAN
- “DHCP Snooping” blocks DHCP responses except from trusted ports (those which connect to the upstream/core network)
- “RA Guard” is a similar feature for IPv6 router advertisements

ARP and NDP spoofing?

- Some vendors tell you to lock down ARP/NDP and bridge tables
  - dynamic ARP inspection?
  - lock IP addresses to specific MAC addresses?
  - lock MAC addresses to specific physical ports?
- Hard to manage, doesn’t scale, and doesn’t enhance security much
- Recommendation:
  - keep your servers on a different subnet to your users, and your infrastructure management IPs on a different subnet again

Questions?

Network Management

- Enable syslog and/or SNMP traps
  - Collect and process in centralized log server
    - Spanning Tree Changes
    - Duplex mismatches
    - Wiring problems

Network Management

- Enable syslog and/or SNMP traps
  - Collect and process in centralized log server
    - Spanning Tree Changes
    - Duplex mismatches
    - Wiring problems

- Monitor configurations
  - Use RANCID or Oxidized to report any changes in the switch configuration

Network Management

- Collect forwarding tables with SNMP
  - Allows you to find a MAC address in your network quickly
  - You can use simple text files + grep, or a web tool with DB backend *(e.g. Netdisco, LibreNMS)*

Network Management

- Collect forwarding tables with SNMP
  - Allows you to find a MAC address in your network quickly
  - You can use simple text files + grep, or a web tool with DB backend *(e.g. Netdisco, LibreNMS)*

- Enable LLDP (or CDP or similar)
  - Shows how switches are connected to each other and to other network devices
  - LLDP is “Link Layer Discovery Protocol” (IEEE 802.1AB)
  - CDP is “Cisco Discovery Protocol”

Out-of-band (OOB) management

- How to manage devices when the network itself is down?
  - some network devices have a separate management ethernet port with independent IP address and gateway
  - enterprise servers have "integrated lights out management" (ILO/LOM)
- Build a separate out-of-band management network
  - with its own switch (even a dumb one will do)
  - independent of your core network
  - separately firewalled, or add ADSL / 4G LTE etc for full OOB access
  - can carry SNMP monitoring traffic as well as provide SSH access

Console servers

- Using a console server is quicker than going to the location with a laptop and serial cable to access the device console
- Some example devices include:
  - OpenGear’s CM7100 (16-96 serial ports)
  - OpenGear’s ACM7008-2
    - https://opengear.com/products/acm7000-resilience-gateway
  - AirConsole TS
    - https://www.get-console.com/shop/en/24-device-servers

Console servers

- Alternatively, build a serial console server yourself
  - Simple Linux PC (mini-PC is sufficient)
  - Multi-port USB hub
  - USB to serial cables

- Out of band access to network devices is essential to rapidly resolve issues
  - Highly recommended!

Documentation

- Document where your switches are located
  - Name switch after building name
    - E.g. building1-sw1
  - Keep files with physical location
    - Floor, closet number, etc.

Documentation

- Document where your switches are located
  - Name switch after building name
    - E.g. building1-sw1
  - Keep files with physical location
    - Floor, closet number, etc.

- Document your edge port connections
  - Room number, jack number, server name

Questions?

UNIVERSITY OF OREGON

NSRC
Network Startup Resource Center