Chapter 9
High Availability
NET3011 – 17W

High Availability – Basics

- High Availability (HA) generally refers to the design of hardware or software (OS) which allow a device to continue operating despite failures (with minimal delay)
- High Availability goes beyond the mechanisms provided by FHRP and redundant topologies
- High Availability typically includes items such as redundant power supplies and fan trays, entire duplicate control planes, and technologies for multi-chassis LAG
- HA also typically provides for hot-swapping components with any requirement to reconfigure or reboot
- Many routing protocols have features to support HA, such as "graceful restart" in OSPF (RFC3623)

High Availability – Some Examples

- The dozen pages in Ch 9 present *examples* of HA, as implemented by Cisco, including:
  - stacking physical switches into 1 virtual switch
    - Stackwise & FlexStack (access layer) and
    - Virtual Switch System (VSS) (distribution layer)
  - redundant Control Plane (aka Supervisor) boards
    - Stateful SwitchOver (SSO)
  - fail-safe Data Plane operation (focuses on L3)
    - Non-Stop Forwarding (NSF)
- We'll look at these as specific examples of general techniques, but will *not* consider the CLI details

HA – Stacking Switches (1)

- Many vendors have a way of stacking switches to make them operate as a single logical unit
- Backplane connectors form a bi-directional ring

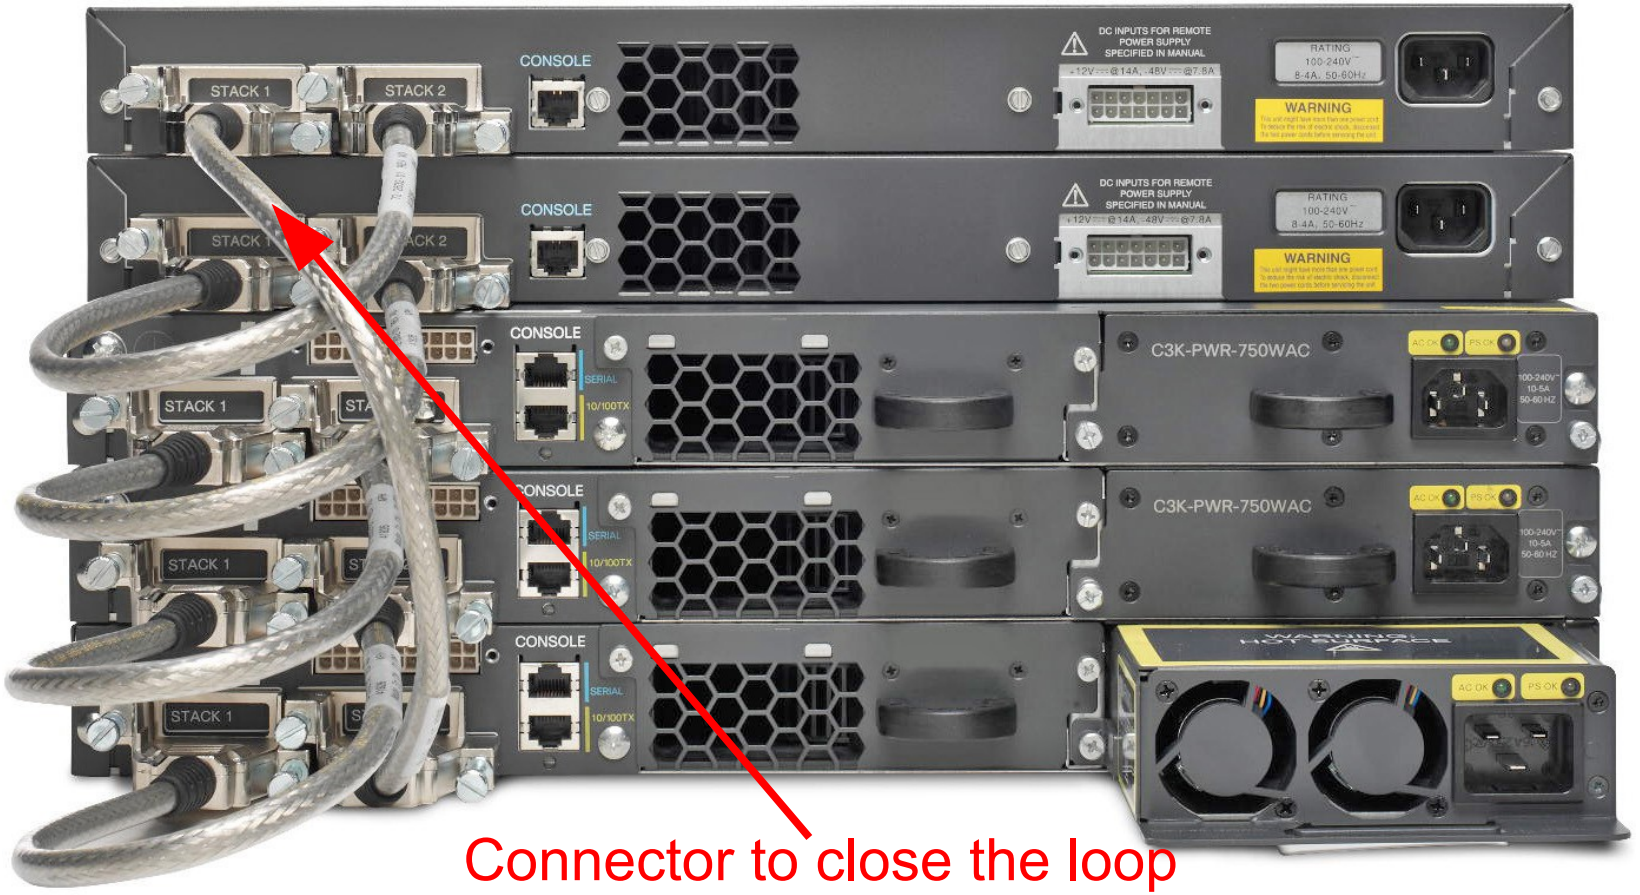

HA – Stacking Switches (2)

- Stacked switches allow multi-chassis LAG

Before

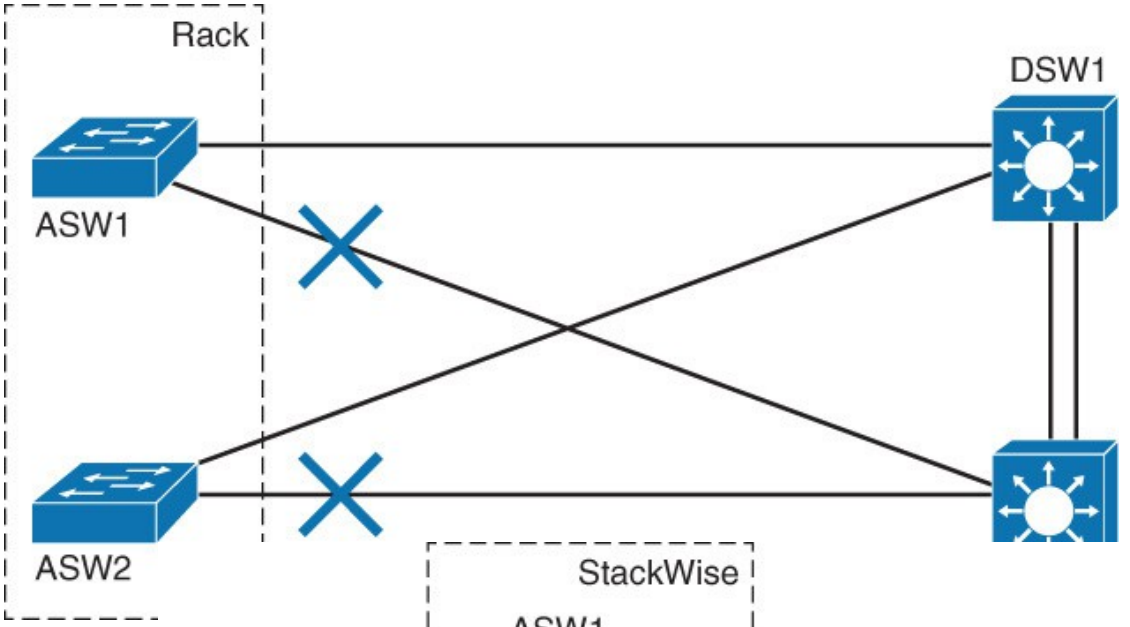

After

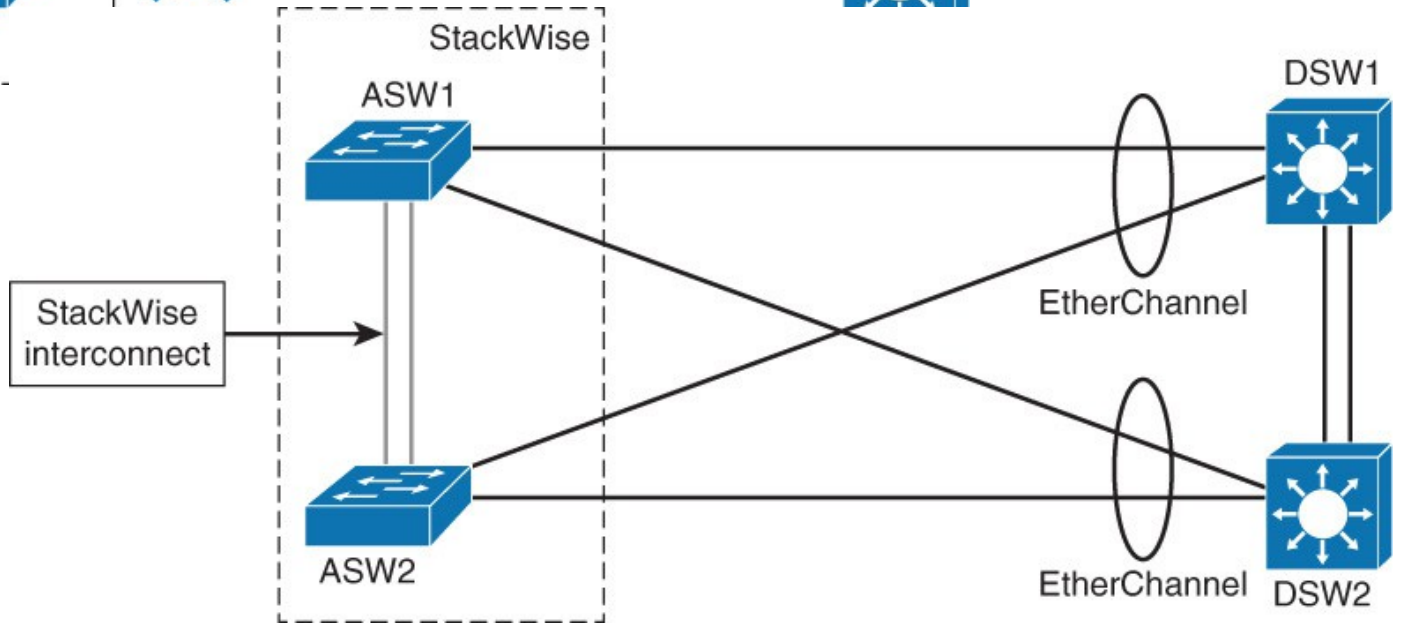

HA – Stacking Switches (3)

- Stacked switches are configured as a single unit
- Switches can be added or removed while the stack is in operation
- Each specific technology has limits on number of switches (Stackwise = 9; FlexStack = 4)
- Backplane connections are typically very fast (Stackwise = 480Gbps; VSS = 2 x 80Gbps)
- Stacked switches being physically many but virtually one may eliminate the need for a number of lower-level protocols (e.g. STP, FHRP)

HA – Redundant Control Plane boards

- Two boards, two roles: active and standby
- Active handles all Control Plane functions: protocols, timers, etc.
- Standby is allowed to boot to a certain level, then monitors Active; if Active fails, Standby completes it's initialization and takes over
- Historically, three levels of booting for Standby:
  - partly booted; must reset/reload all other boards upon taking control; minutes to fail-over (old; outdated)
  - fully booted, but no protocols initialized; no reloading necessary; fail-over within a minute (old; outdated)
  - fully booted, perfectly synchronized with Active, including **state** of all protocols; fail-over within secs
    SSO = Stateful SwitchOver (modern/current)

HA – Redundant Control Plane (2)

- SSO can synchronize most L2 protocols between the two planes:
  - 802.3x (Flow control)
  - 802.3ad (LACP) and PagP
  - 802.1X (Authentication and port security)
  - 802.3af (Inline power)
  - VTP
  - Dynamic ARP inspection / DHCP snooping / IP Source Guard
  - IGMP snooping (Versions 1 and 2)
  - DTP (802.1Q and ISL)
  - MST/PVST+/Rapid-PVST
  - PortFast / UplinkFast / BackboneFast / BPDU Guard & Filtering
  - Voice VLAN
  - Unicast MAC filtering
  - Access control lists (ACLs; VLAN ACLs, port ACLs, router ACLs)
  - QoS (Dynamic buffer limiting [DBL])
  - Multicast storm control / broadcast storm control

HA – Fail-safe Data Plane operation

- NonStop Forwarding (NSF) is a generic term for hardware & protocols capable of continuing to forward traffic using existing FIB info even if/while Control Plane is down
- If your hardware is good enough ($$$) for NSF, it's almost guaranteed to also support SSO, so both Data Plane + Control Plane work together for HA
- Zero packet (L3!) loss or near-zero packet loss is achievable with the combination of NSF + SSO
- NSF also focuses on quickly rebuilding all RIBs:
  - all 4 major routing protocols (BGP, OSPF, ISIS, EIGRP) have extensions (aka "graceful restart") to send special a packet upon Control Plane fail-over
  - it triggers routing updates from NSF-aware neighbors

Reminder

- *Some* details do not appear in these slides
- You are responsible for *reading* the textbook to gain the knowledge (memorization) and understanding (apply the knowledge)