Chapter 3
Campus Networks: VLANs, DTP, VTP, LAG
NET3011 – 17W

(Section on PVLANs in separate slide deck)

VLANs: Avoiding a flat network

- VLANs isolate traffic and thus provide a logical broadcast domain
- Ports in the same VLAN share broadcasts; ports in different VLANs do not
- VLANs may exist on a single switch ("Local") or may span multiple physical LAN segments and switches ("End-to-end")
- VLANs are a L2 domain and thus every VLAN must terminate at (at least) one L3 routed port in order to go beyond the subnet

VLAN Operation

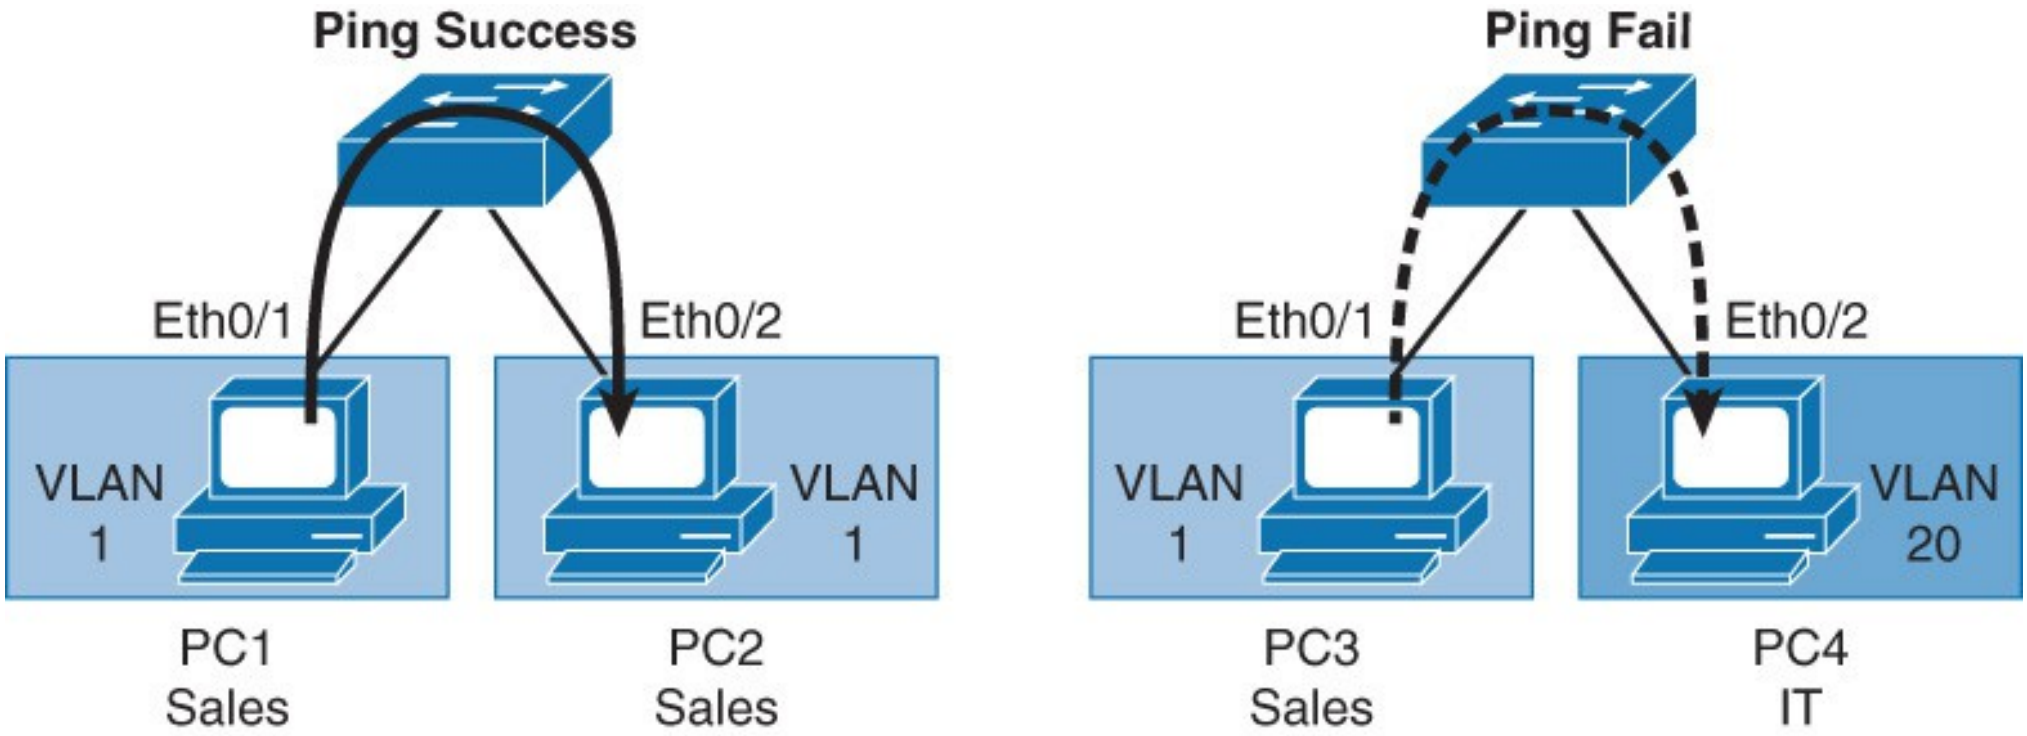

- VLAN is an independent LAN network.
- VLAN = broadcast domain.
- VLAN maps to logical network (subnet).
- VLANs provide segmentation, security, and network flexibility.

Example: VLANs in a building

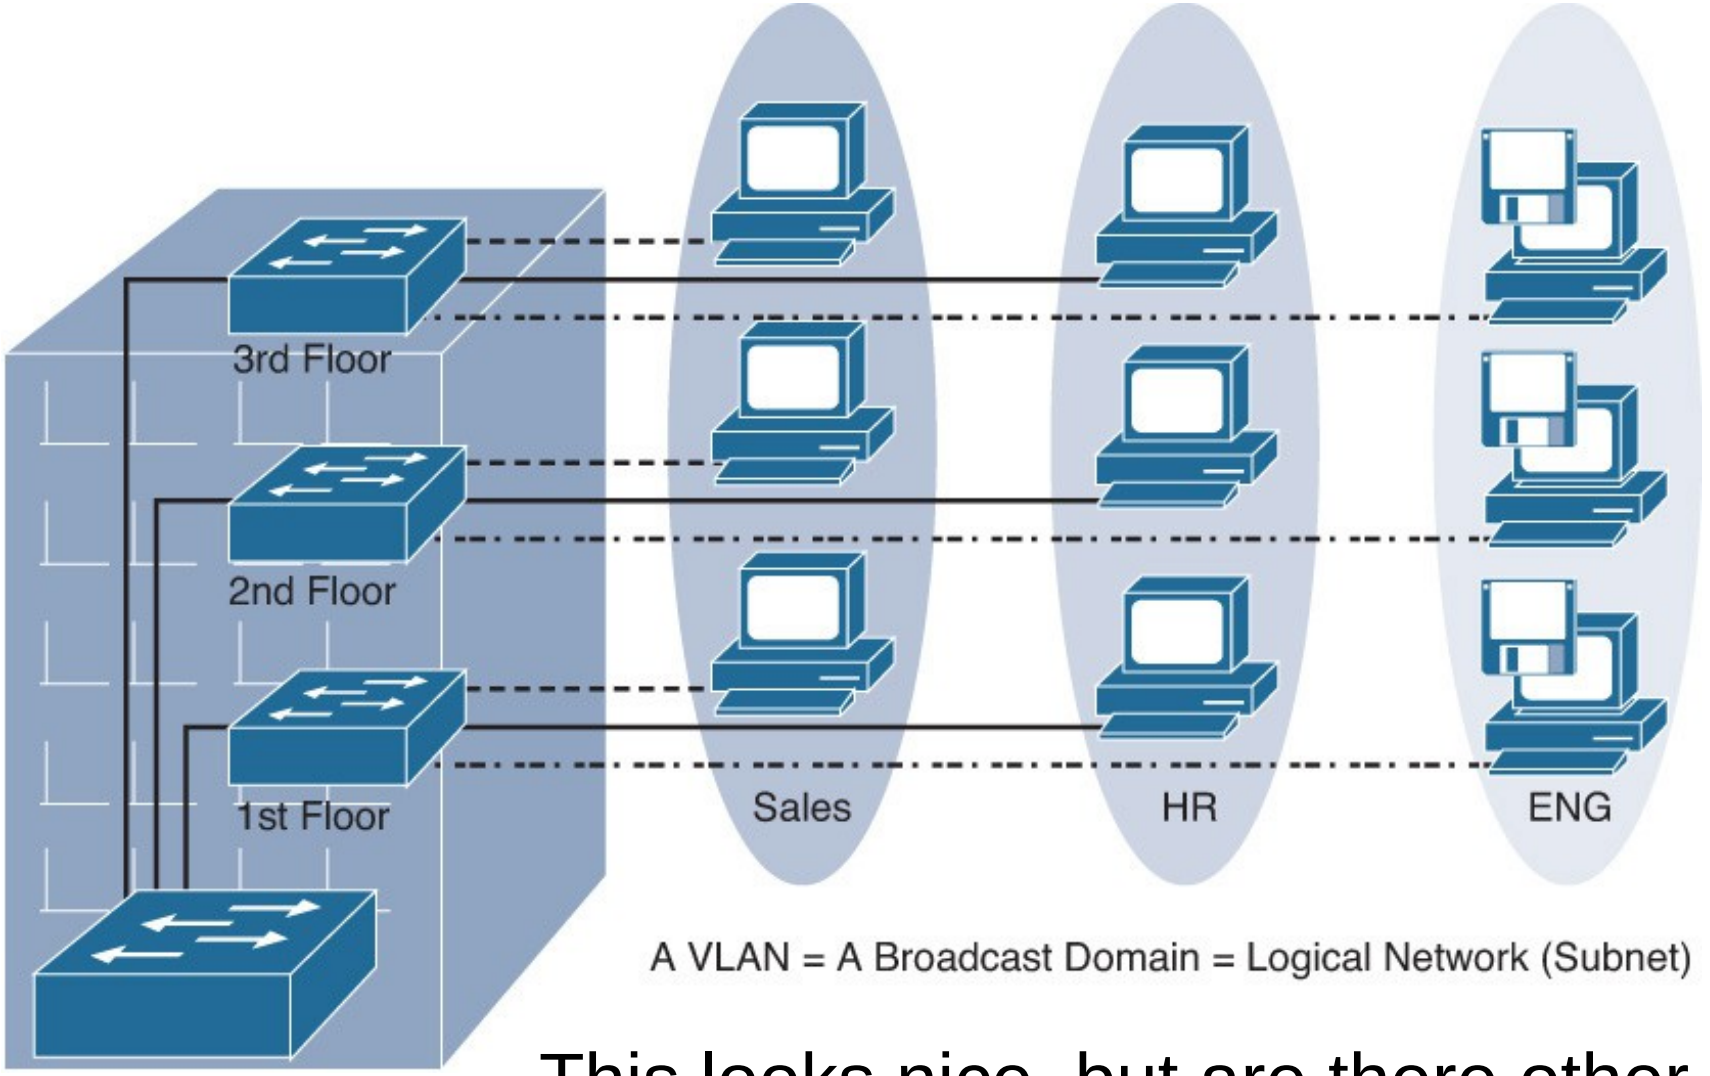

This looks nice, but are there other topologies or ways of organizing VLANs?

VLAN Design: End-to-End

- End-to-end VLAN refers to individual VLANs that are widely dispersed throughout a campus on multiple switches.
- **Require** trunk links … and thus DTP and VTP (… but see Best Practices, p. 65-66 of FLG)

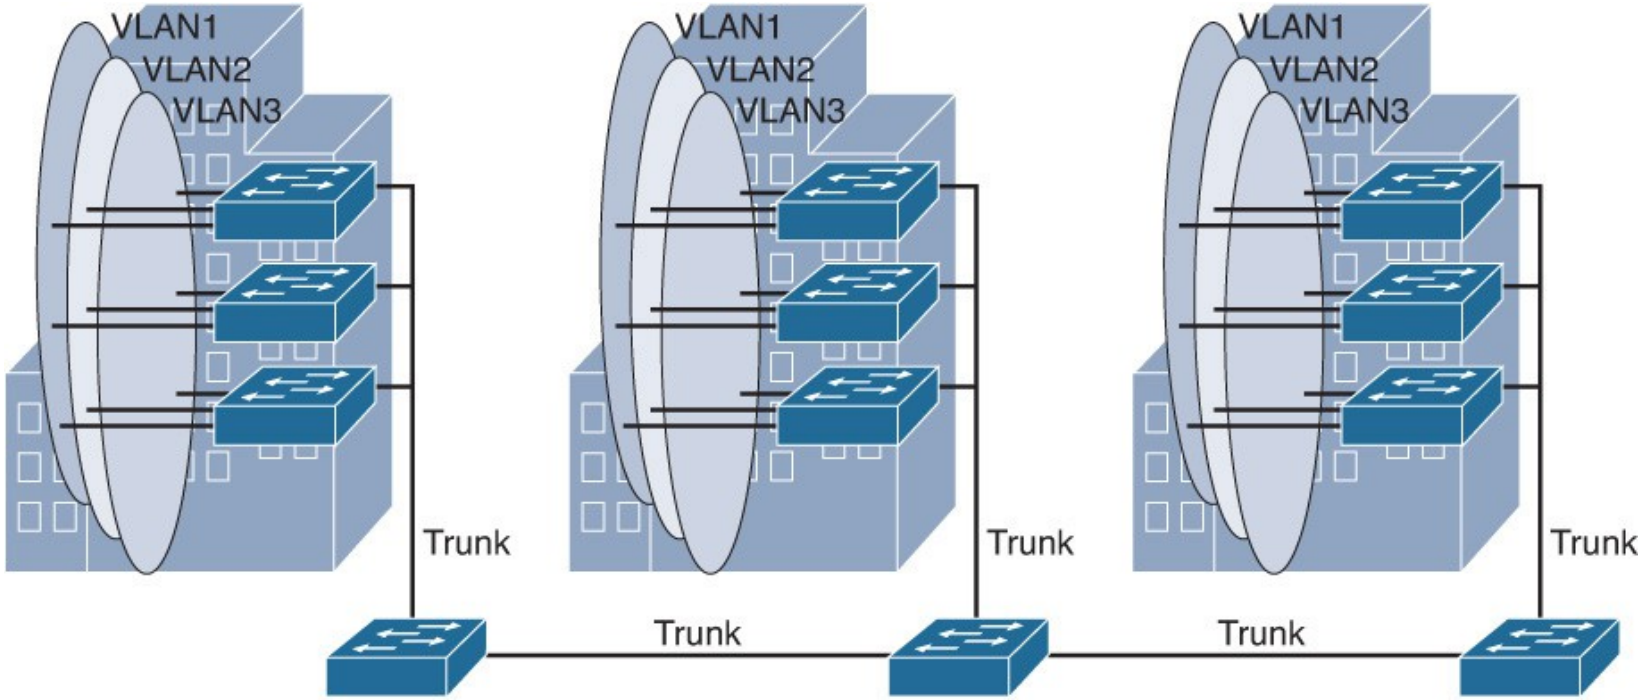

End-to-End VLAN Characteristics

- VLANs are dispersed geographically throughout the network
- Users are grouped into each VLAN regardless of the physical location
- As a user moves throughout a campus, the VLAN membership of that user remains the same, regardless of the physical switch to which this user attaches.
- Users are typically associated with a given VLAN for network management reasons. This is why they are kept in the same VLAN, therefore the same group, as they move through the campus
- All devices on a given VLAN typically have addresses on the same IP subnet
- Switches commonly operate in a server/client VTP mode

Source: p. 45 of FLG

VLAN Design: Local

- Devices connect to their geographically nearest switch, regardless of function or organizational role
- VLANs are confined to a wiring closet, e.g. to a single access switch, then connect to Distribution (… and see Best Practices on p. 65-66 of FLG)

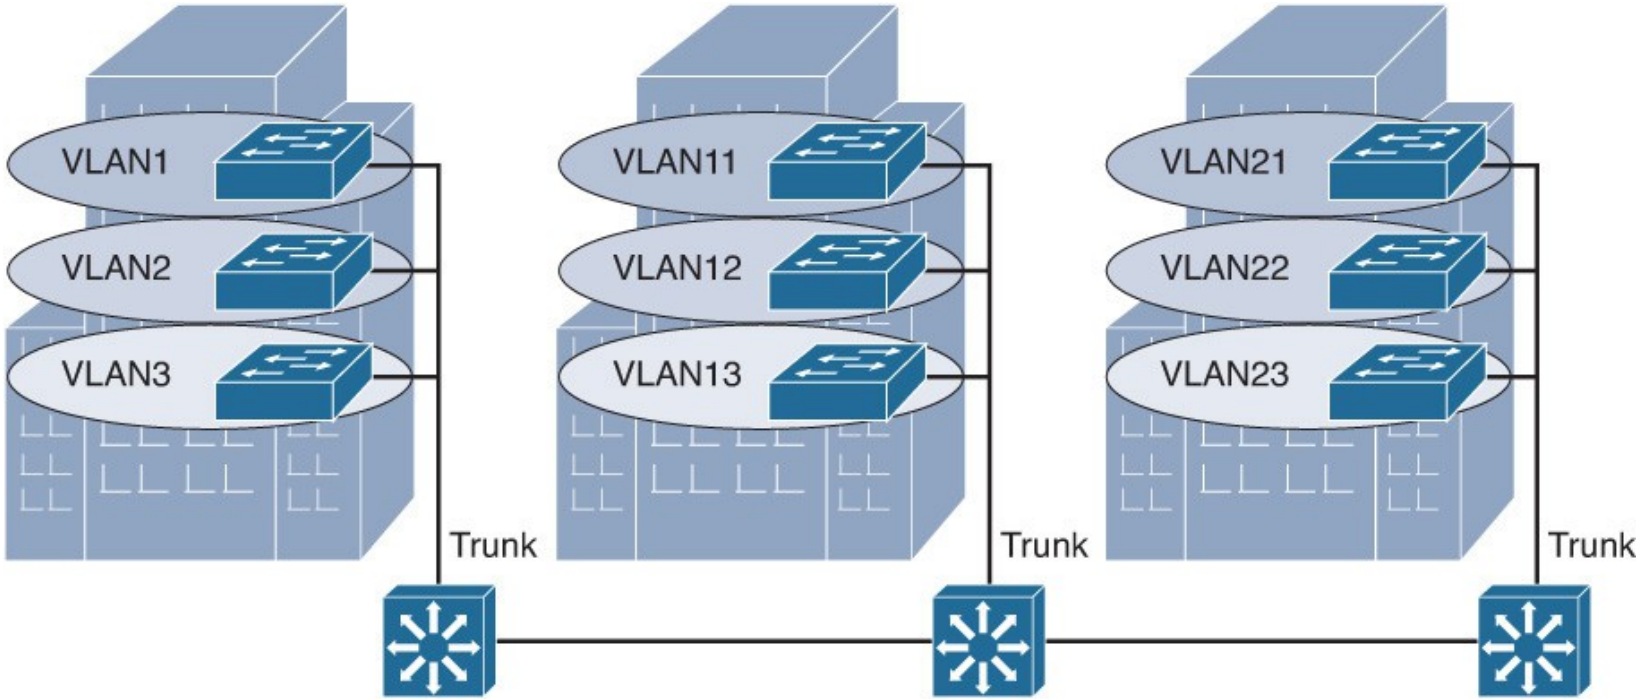

Local VLAN Characteristics

- The network administrator should create local VLANs with physical boundaries in mind rather than the job functions of the users on the end devices
- Generally, VLANs exist between the access and distribution layers
- Traffic from a local VLAN is routed at the distribution and core levels to reach destinations on other networks
- Configure the VTP mode in transparent mode because VLANs on a given access switch should not be advertised to all other switches in the network, nor do they need to be manually created in any other switch VLAN databases
- A network that consists entirely of local VLANs can benefit from *increased convergence times offered via routing protocols*, instead of a spanning tree for Layer 2 networks. It is usually recommended to have **one to three VLANs per access layer switch**.

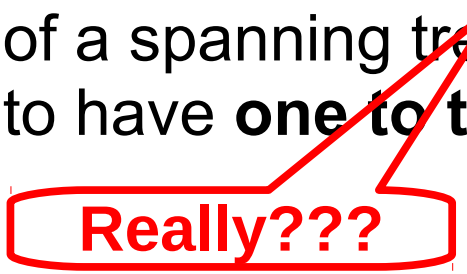

Source: p. 46 of FLG

VLAN Design: End-to-end or Local?

- BEWARE of contradictions in the FLG:

"there are many reasons to implement end-to-end VLANs. The main reason to implement local VLANs is simplicity." [1]

… and yet …

- "L3 … in the access layer has become more common over typical L2 switching" [2]
- "L3 switching in the access layer VLANs scales better" [3]
- "Layer 3 in the access layer are becoming more popular" [4]
- "local VLANs have become more efficient" [5]
- "there are few benefits to extending a VLAN throughout an enterprise" [6]
- "It is recommended to have two to three VLANs per access block rather than span all the VLANs across all access blocks." [7]

[1]p. 47 [2]p. 17 [3]p. 18 [4]p. 19 [5]p. 47 [6]p. 48 [7]p. 48

VLAN Design: Pros and Cons

- See FLG p. 46-48, 49 for pros / cons of each type **Learn** them!
- E-to-E: +Grouping, +virtualization, +security, +QoS, +avoiding needless routing, +specialized VLANs, -management, -broadcast propagation, -troubleshooting
- Local: +Simpler, +deterministic/predictable, +better link utilization(?), +better HA due to matched L2/L3 topology, +finite failure domain, +scalable design

VLAN Organization

- ~~Five~~ ~~Six~~ Seven generic types of VLANs:
  - **default** = VLAN 1 (also Native VLAN by default)
  - **native** = traffic on this VLAN isn't tagged
  - **data** = regular user traffic; may consist of many VLANs
  - **voice** = needed in modern campus networks
  - **wireless** = particularly for controller-based WLANs
  - **management** = for SSH, HTTPS, SNMP access
  - **garbage** or **black-hole** = for security purposes, a VLAN that is suspended and also never allowed on trunks; all un-used interfaces are placed in this VLAN.
- We'll dig into voice VLANs at the end of the section (FLG p. 67-68); NET3900 handled WLANs

Trunking Basics

- 2 methods of tagging: ISL and 802.1Q
- ISL is **dead**; you could forget about it *except* that many Cisco switches that (still) support it may need to have 802.1Q explicitly chosen in the configuration

Memorize the 802.1Q tag format!

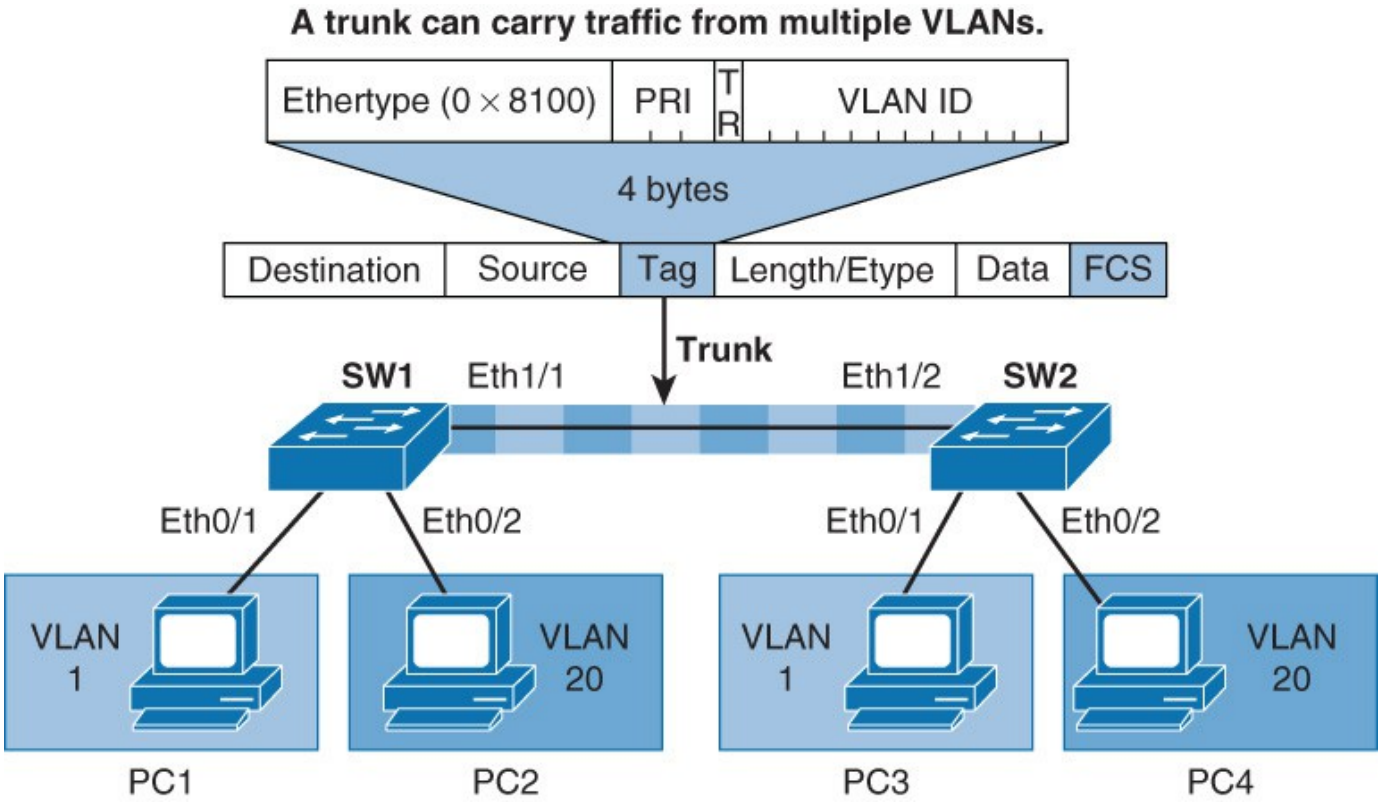

Support for VLANs on Trunks

- Compare with Lotto 649 and Lotto MAX
- Now consider that different categories of switches (access, distribution) may be limited in the number of VLANs that can be "picked" simultaneously by the configuration.
- Examples:
  - 2960 can support 255 VLANs (250 + 5 std)
  - 3560 can support 1005 VLANs (1000 + 5 std)
  - 4000/6000 series supports 4094 VLANs (all)
- As a **separate** issue, number of simultaneous VLANs may be limited by STP ("Per VLAN STP")

Trunking Characteristics

- Trunk links have a *Native* VLAN which is the only VLAN whose traffic is ***not*** tagged;
  - by default = 1
  - best practice: choose & configure a native VLAN number that is unused for any other purpose
- Trunk links have *allowed* and *active* VLANs
  - can configure allowed VLANs
  - active VLANs are a result of network conditions
- Trunking *encapsulation* can be set, but only on Cisco switches that still support ISL; otherwise encapsulation is 802.1Q
- List of *pruning-eligible VLANs* (… relevant only if VTP is enabled)

http://www.cisco.com/c/en/us/td/docs/switches/lan/catalyst2950/software/release/12-1_9_ea1/configuration/guide/scg/swvlan.html#wp1214566

Trunking Subtleties

- Interfaces are configured to strictly *access* mode, configured to strictly *trunk* mode, or configured for dynamic selection of either access or trunk mode
- *Describing* trunk characteristics is ***not*** the same as setting an interface to trunking mode! (Likewise, describing access mode characteristics does **not** configure an interface to access mode)
- *Unmanaged* switches completely ignore VLAN tags, which also means they will pass *all* traffic and thus can be used to join trunk links

MTU Considerations

- Adding one (or more) VLAN tags makes the frame larger by 4 bytes
- MTU commentary on FLG p. 51-52 appears almost completely **incorrect**.

  "To process an 802.1Q tagged frame, a device must enable a maximum MTU of 1522 or higher."

  But have you **ever** had to adjust MTU in lab??

  In 2016:

  max *frame* size = MTU + 22 bytes (always!)

  see: http://www.cisco.com/c/en/us/support/docs/switches/catalyst-4000-series-switches/29805-175.pdf
- Using MPLS will likely require increasing MTU

Trunking and VLAN commands

- You hopefully remember the well-used commands:
  - show vlan [brief]
  - show interfaces trunk
- For additional details on a specific VLAN:
  - show vlan id {#}
  - show vlan name {vlan-name}
- Review all other VLAN and trunk configuration commands presented in the textbook(s)

DTP Basics

- Dynamic Trunking Protocol, a Cisco proprietary protocol (no known open standard equivalents)
- Only one purpose/function: negotiate the creation of trunk links dynamically, *if protocol is enabled*
- Has *nothing* to do with VTP
  … except Cisco mixed them together in one way: DTP won't work unless VTP domain names are identical on both sides of link ("null" can match)
- Exactly and only two options:
  – desirable = negotiate pro-actively
  – auto = negotiate passively ("Don't start nuthin!")

DTP Combinations

	Dynamic Auto	Dynamic Desirable	Trunk	Access
Dynamic Auto	Access	Trunk	Trunk	Access
Dynamic Desirable	Trunk	Trunk	Trunk	Access
Trunk	Trunk	Trunk	Trunk	Problem! ~~Limited Connectivity~~
Access	Access	Access	Problem! ~~Limited Connectivity~~	Access

Source: p. 54 of FLG, except caption is wrong there!

Best Practices for VLANs & Trunking

- FLG p. 65-66 provides a dozen best practices for VLAN and trunk design & configuration (though there is overlap between some of them)
- Note how some recommendations are fluid, depending on how deep L3 goes down, whether the network has reached long-term stability (ever?), and the specific physical topology
- Be sufficiently familiar with these that you can describe at least half of them and recognize the rest on a test

Best Practices

- For local VLAN designs, only one to three VLANs per access module and limit those VLANs to a couple of access switches and the distribution switches.
- For local VLAN designs, avoid VTP; use manually allowed VLANs on trunks.
- Avoid using VLAN 1 as the "blackhole" for all unused ports. Use a dedicated VLAN separate from VLAN 1 to assign all the unused ports.
- Separate the voice VLANs, data VLANs, the management VLAN, the native VLAN, blackhole VLANs, and the default VLAN (VLAN 1).
- For stable trunks: (1) turn off Dynamic Trunking Protocol (DTP) and use manual configuration; (2) use IEEE 802.1Q rather than ISL because it has better support for QoS and is an open standard; (3) use "switchport nonegotiate" to save time on bootup; (4) manually configure allowed VLANs.
- Manually configure ports in access mode except those that are specifically intended for a trunk link.
- Keep all data traffic off VLAN 1; only permit control protocols to run on VLAN 1 (DTP, VTP, STP BPDUs, PAgP, LACP, CDP, etc.).
- Avoid using Telnet because of security risks; enable SSH support on management VLANs.

Compiled from FLG p. 65-66

VTP Basics

- VTP is used in a End-to-end VLANs to auto-matically distribute VLANs to switches; this avoids the need for manually configuring all switches
- What happens if traffic arrives on a trunk for a VLAN that isn't configured on a (managed) switch?

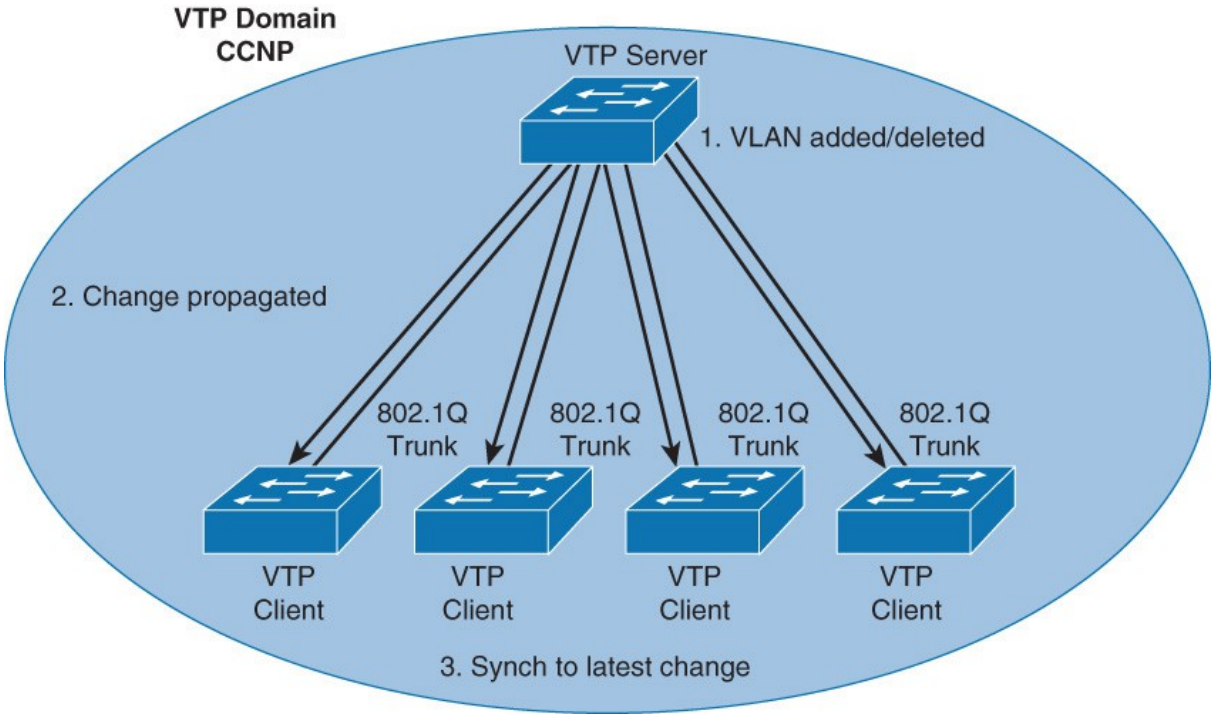

VTP Characteristics

- VTP is a Cisco proprietary protocol;
  GVRP and M(V)RP are open standard equivalents
- VTP will **only** communicate over trunk links.
- VTP always sends advertisements over *VLAN 1*,
  every 5 minutes or *whenever there's a change*,
  (FLG . 70 is wrong on both items; p. 75 is correct regarding frequency)
  (See: http://www.techexams.net/forums/ccie/41133-native-vlan-relation-stp-cdp-vtp-dtp.html)
  to a fixed multicast address (01:00:0C:CC:CC:CC),
  using an 802.3 frame (not Ethernet II)
- Key parameters for accepting a message:
  – version – revision number
  – domain name – password
  – MD5 hash (for above fields)

VTP Characteristics

- Versions 1, 2, 3 exist; each ver. has new features
  – ver 1 is **not** compatible with ver 2, but will "jump"
  – ver 3 **is** backwards compatible with ver 2, but not ver 1
- Ver 1: Basic protocol to dynamically xfer VLANs
- Ver 2: Enhancements on FLG p. 73; *know* them!
- Ver 3: Covered briefly in FLG & v7 Lab!
  … so it *IS* covered in this course.
  Ref: FLG p. 74; v7 Lab 3-1: p. 12-13, 16

VTP modes

- Client: receive-only; cannot make any changes
  - but can *shutdown* a VLAN (!)
  - relays all messages from same domain
  - may adopt domain name from server if no password configured (but not in v3)
- Server: can make all changes (add/del/modify)
  - relays all messages from same domain
  - can *shutdown* and *suspend* a VLAN
  - may adopt domain name from server if no password configured (but not in v3)
  - v3 adds *Primary Server*, the only that can make changes
- Transparent: like server but with separate VLAN database
  - won't adopt domain name from any other server
- Off: like Transparent, except no forwarding; totally isolated!
  (on *Cisco* equipment, only available for IOS's that support VTP v3)

VTP Idiosyncracies

- Highest revision **trumps** all, regardless of mode!
- Ver 3: provides extended VLANs, but doesn't have 100% support (can't *shutdown* or *prune* them)
- *Ver 3: Primary Server*
  - configured in privilg'd exec and *not* global config
  - so not saved in either config or vlan.dat
  - so doesn't survive a reboot
- *Ver 3: only* the Primary Server can make changes
  - but must be manually configured *after* reboot
  - let's hope that no admins are in "panic" mode after a network crash, so that they don't make any mistakes!!

Beware: VLAN Nastiness with Cisco

- On Cisco switches, Server mode stored in vlan.dat trumps Transparent mode from startup-config !
- or: When trying to load a config with VTP Transparent mode, check for an existing vlan.dat because it will over-rule the Transparent mode!
- Check the log entries generated during bootup (ie. **`show logging`**)

Supplemental to textbook material!

```
...
Mar  1 00:00:43.889: %SPANTREE-5-EXTENDED_SYSID: Extended SysId enabled for type vlan
Mar  1 00:00:47.588: VTP mode mismatches between startup-config: transparent,
                     vlan database: server.
Mar  1 00:00:47.596: %SW_VLAN-4-BAD_STARTUP_VLAN_CONFIG_FILE:
                     Failed to configure VLAN from startup-config.  Fallback to use VLAN
                     configuration file from non-volatile memory
Mar  1 00:00:47.596: %SYS-5-CONFIG_I: Configured from memory by console
Mar  1 00:00:48.528: %SYS-5-RESTART: System restarted –
...
```

VTP Pruning

- Eliminates unnecessary propagation of broadcast & flooded traffic to links with no members in VLAN
- Eliminating this traffic does **not** eliminate the VLAN!
- VTP pruning is *Global*

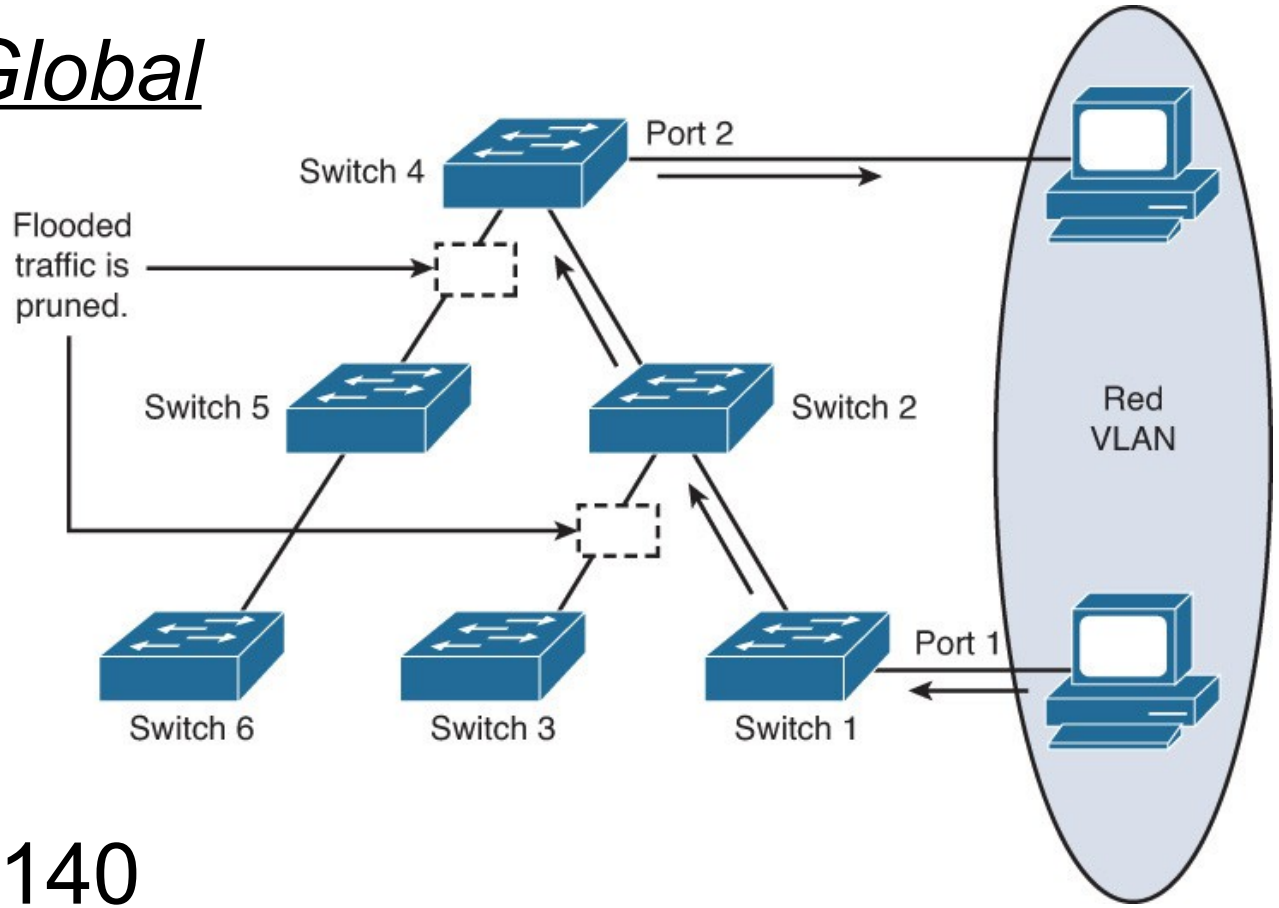

- Refs:
  - FLG, p. 74
  - Cert Guide, p. 140

VTP Messages

1.VTP Summary advertisement

Version (1 byte)	Type (Summary Adv) (1 byte)	Number of subset advertisements to follow (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision Number (4 bytes)			
Updater Identity (orginating IP address: 4 bytes)			
Update Time Stamp (12 bytes)			
MD5 Digest hash code (16 bytes)			

Source: Official Cert Guide, p. 129

VTP Messages

2.VTP Subset advertisment

0	1	2	3
Version (1 byte)	Type (Subset Adv) (1 byte)	Subset sequence number (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Configuration Revision Number (4 bytes)			
VLAN Info Field 1 (see below)			
VLAN Info Field ...			
VLAN Info Field N			

0	1	2	3
Info Length	VLAN Status	VLAN Type	VLAN Name Length
VLAN ID		MTU Size	
802.10 SAID			
VLAN Name (padded with zeros to multiple of 4 bytes)			

Source: Official Cert Guide, p. 130

VTP Messages

3.Advertisement Request

0	1	2	3
Version (1 byte)	Type (Adv request) (1 byte)	Reserved (1 byte)	Domain name length (1 byte)
Management Domain Name (zero-padded to 32 bytes)			
Starting advertisement to request			

4.Mystery message: propagates VTP pruning option

VTP Configuration

- You are responsible for config method & cmds
- Don't forget about precautions for revision number

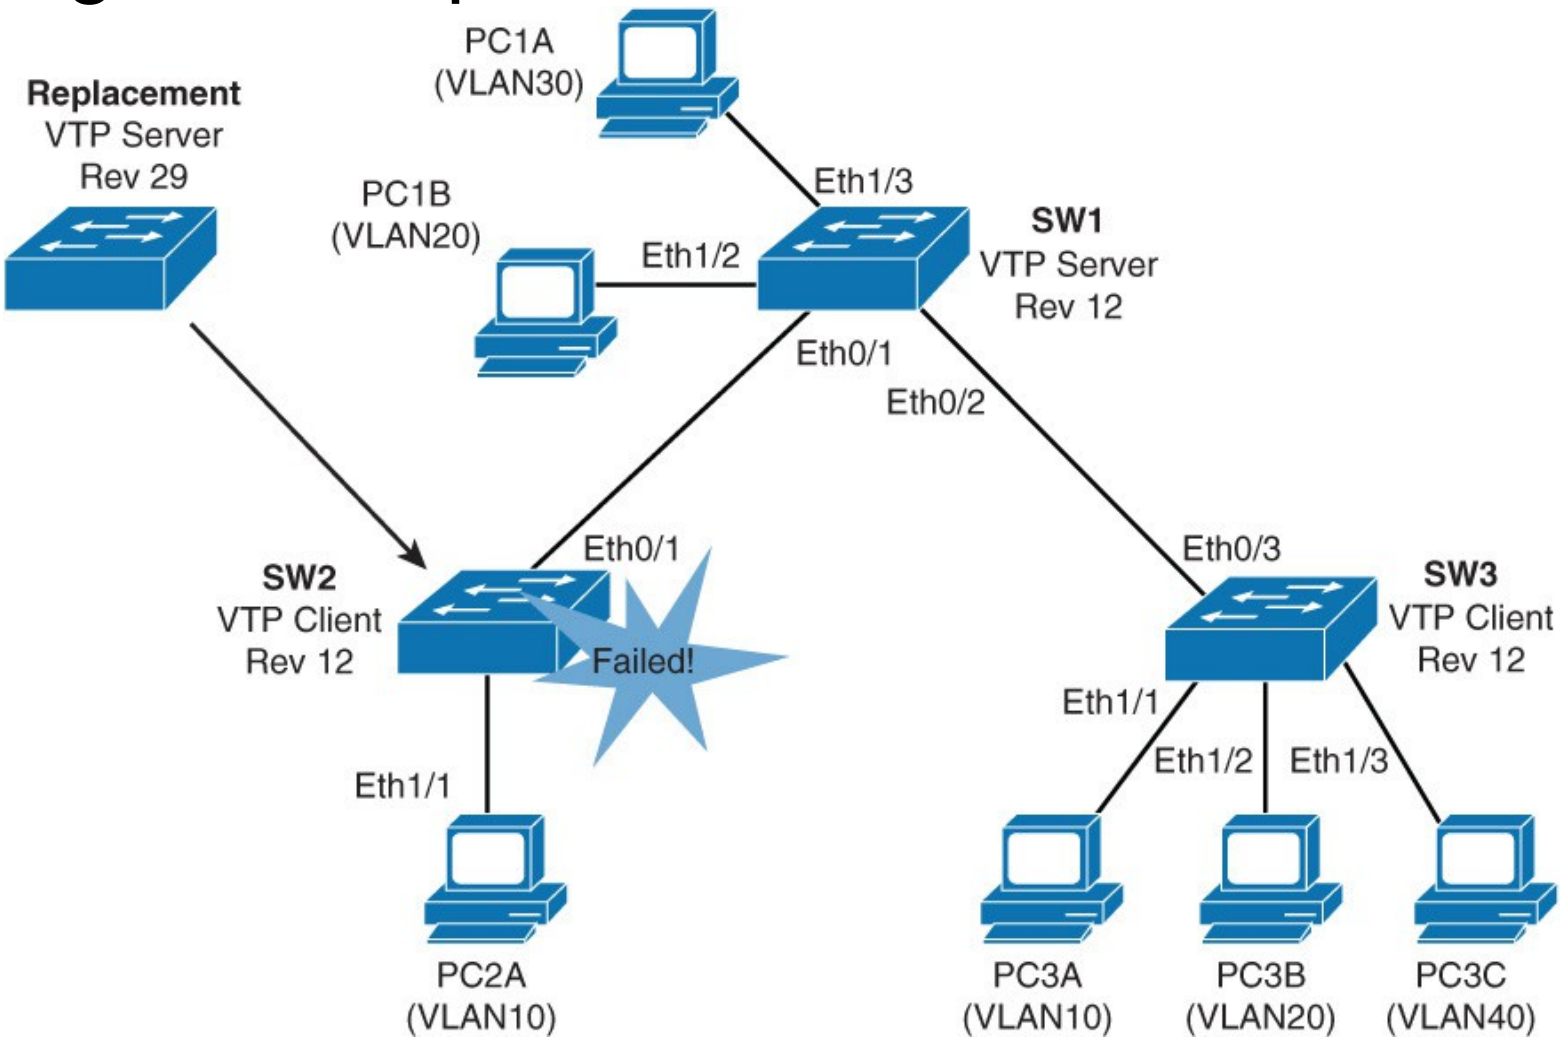

VTP Best Practices

- "Avoid, as much as possible, VLANs that span the entire network"
- When adding a new switch, ensure revision = 0:
  – change to fake domain name and back; or
  – change to Transparent and back to desired mode
- "Configure all switches to transparent VTP mode and manually add VLANs as needed, especially in a larger campus network. VTP configuration is usually good for small environments."
- … What are consequences of above statements to End-to-End vs Local VLAN debate?

Link Aggregation (LAG) Basics

- Aggregation is done to obtain higher bandwidth or higher utilization of existing links
- Cisco refers to all LAG as "EtherChannel"

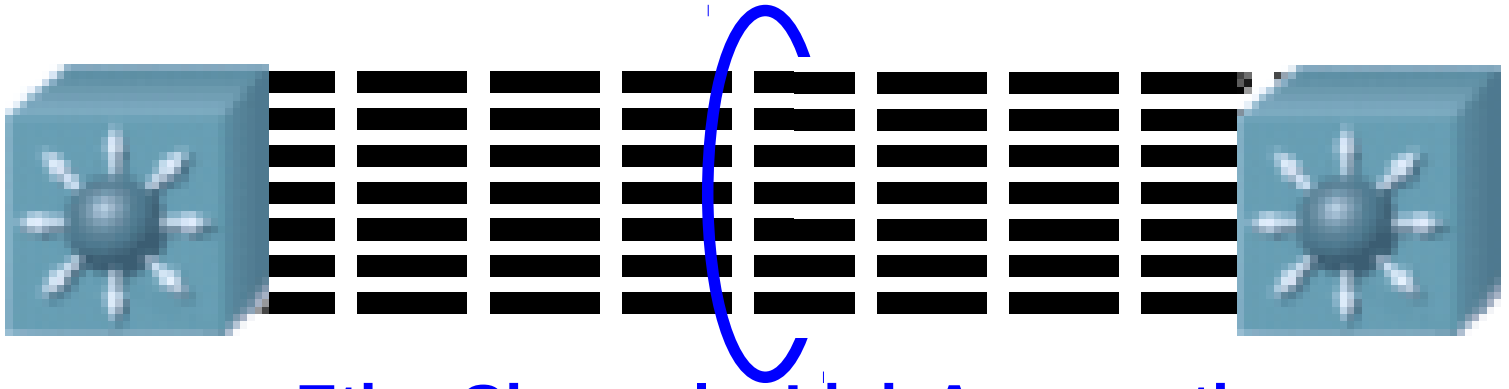

EtherChannel = Link Aggregation

- The choice of load-balancing algorithm is critical
- If not done right, LAG may be *no faster* than a single link!

LAG Characteristics

- There are three options for aggregating links:
  - LACP = IEEE 802.1AX or IEEE 802.3ad (older)
  - PAgP = Cisco's Port Aggregation Protocol
  - statically configured (**`channel-group xx mode on`**)
- *None* of the modes is compatible with the others
- *All* ports in the LAG must be configured identically
- LACP & PAgP modes operate similarly to DTP
- Frames are **never** split, they transit on only 1 link

LACP

	Active	Passive
Active	Yes	Yes
Passive	Yes	No

PAgP

	Desirable	Auto
Desirable	Yes	Yes
Auto	Yes	No

Static Persistance

	On	
On	Problem!	
	Problem!	Yes

Protocol Similarities & Differences

- Both protocols require a port to be configured identically to others before adding it to the bundle
- PAgP: Max 8 links per bundle (depending on HW)
- LACP: Max 16 links per bundle, but max 8 active at any one time
- Both need *power-of-two* links for optimal balancing
  ("an even number of links" as stated on FLG p. 102 is **not** correct; eg **6 links** won't be optimal)
- Both have the same load balancing options:
  - dst-ip, dst-mac, src-ip, src-mac, src-dst-ip, src-dst-mac
  - dst-port, src-port, src-dst-port (larger switches)
- Both use XOR for dual load-balance parameter, simple hash for single load-balance parameter

LAG Requirements

- All the following *must* be identical:
  – speed and duplex
  – if in access mode, same VLAN
  – if in trunk mode, same encapsulation, Native VLAN and allowed VLANs
  – LAG protocol (or statically configured)
  … for all links in the bundle **AND** on both sides of the bundle
- Non-identical links will be suspended and removed from the EtherChannel bundle until consistency is restored; log message is: %EC-5-CANNOT_BUNDLE2 …
- It is *not* required to have the same load-balancing option on both ends of a LAG!! Different may be *necessary*!

Vendor Specific Behaviour

- Which takes precedence: configuring links within a bundle or configuring the bundle? Contradictions? ... May vary depending on model & OS version.
- LACP: " Any port modification after the creation of the channel will also change all the other channel ports." [FLG p. 97]
- PAgP: "PAgP will automatically modify parameters of the EtherChannel if one of the ports in the bundle is modified." [FLG p. 98]
- "After you configure an EtherChannel, any configuration that you apply to the port channel interface affects the EtherChannel. Any configuration that you apply to the physical interfaces affects only the specific interface that you configured." [FLG p. 102]

Optimizing Load Balancing

- Consider the following scenarios and how to optimize the load balancing for each direction.
  - Hosts accessing a gateway (wired or wireless)
  - Rest of network accessing a file server
  - Many devices accessing Netflix over the internet
- Tools exist for capturing and analyzing network traffic (eg NetFlow, NBAR)
  – these tools identify categories of traffic
  – categories help determine the optimal load balancing option(s)

Reminder

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