




Virtual LANs

Fulvio Riso


Politecnico di Torino





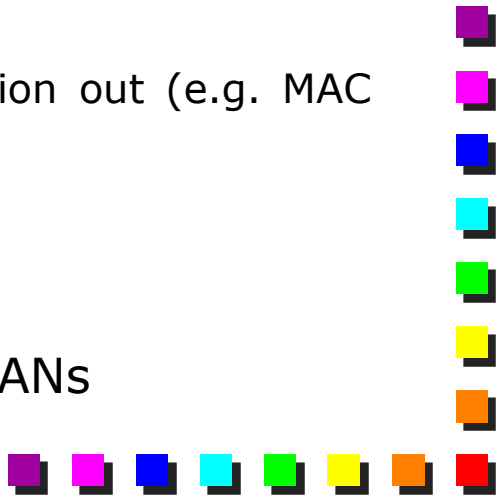
Introduction: L2 or L3?

- So far, we concentrated on L2
- Shall we stay with L2 or better moving to L3?



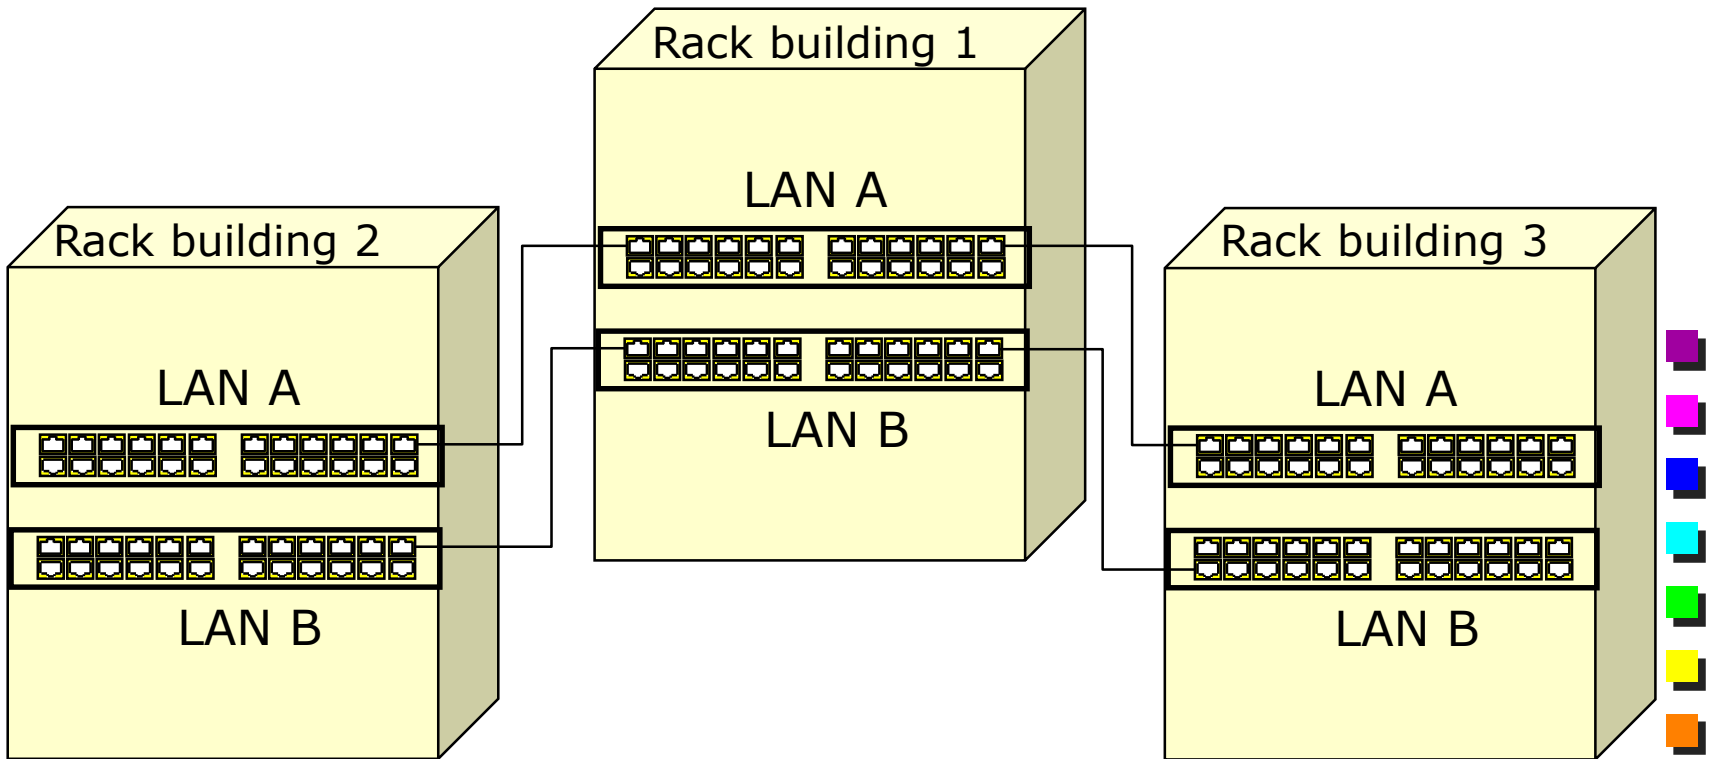




One or multiple LANs across a campus?

- Ok, so it's better to keep the L2 as long as we can
 - As far as the network is able to operate as a single L2 entity (remember scalability issues in L2 networks!)
 - But... a single, gigantic LAN, or multiple LANs?
 - Performance
 - A single LAN has too much broadcast traffic (not filtered by switches)
 - Flooded traffic (e.g. due to frequent STP reconfiguration)
 - Privacy, Security
 - Do not want a station to leak some information out (e.g. MAC Flooding attack)
 - Management
 - Smaller network, simple (and uniform) policies
 - Better to partition different users in different LANs
- 


Multiple LANs across a campus: how?

- Different physical networks (full separation)
 - N networks = N links + N devices
 - Waste of resources







Virtual LANs (1)



Administration Department

Engineering Department


Without VLAN



Administration Department

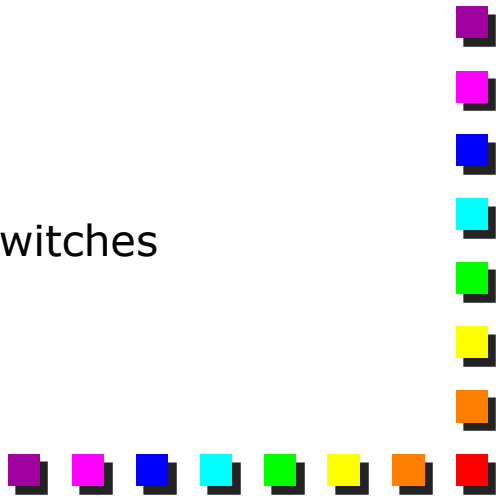
Engineering Department

With VLANs

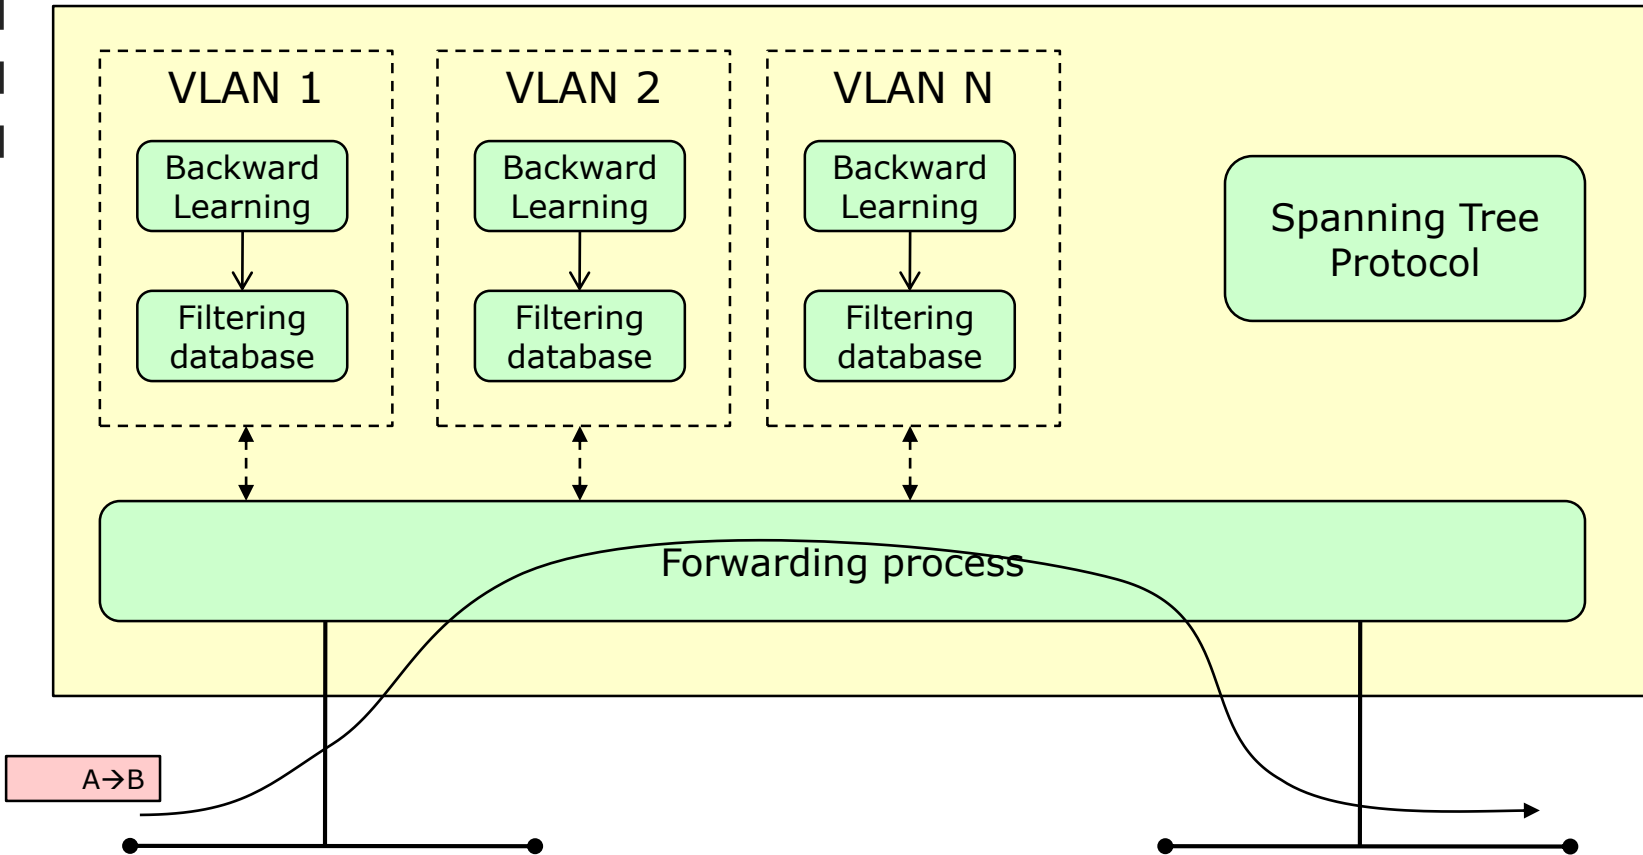




Virtual LANs (2)

- Single physical infrastructure
 - Same devices, same cabling
 - No switches in which only a few ports are used
 - No need to have multiple fibers (for different LANs) in the backbone
 - Different LANs
 - Different broadcast domains
 - E.g., Ethernet frames cannot be propagated on another VLAN
 - No broadcast between LANs
 - No MAC flooding attacks
 - No ARP spoofing
 - Created through a proper (logic) separation on switches
 - Intra-switch or inter-switch
- 

VLAN: switch architecture



VLAN: forwarding database

MAC Filtering DB (VLAN1) MAC Filtering DB (VLAN2) MAC Filtering DB (VLAN3)

MAC	Port
H1	1
H4	4

MAC	Port
H2	2
H5	4

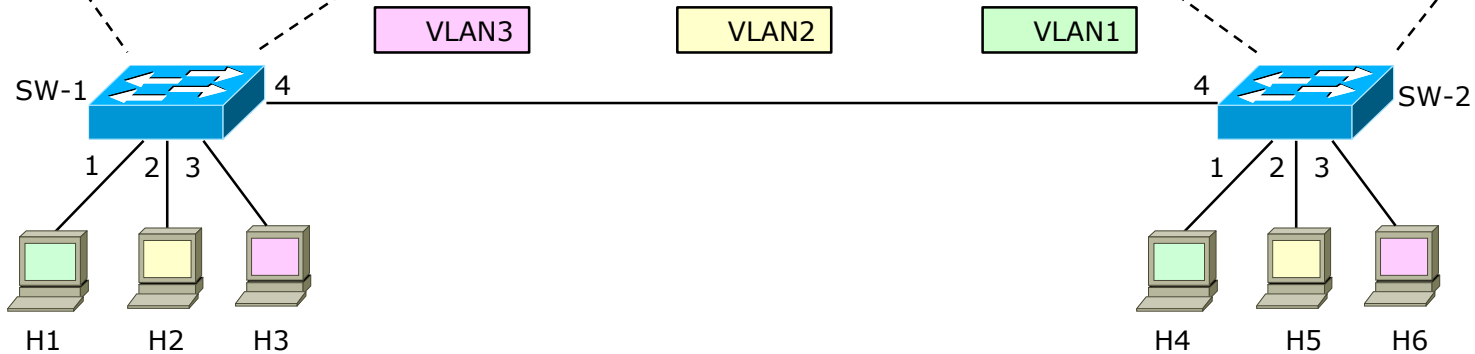
MAC	Port
H3	3
H6	4

MAC Filtering DB (VLAN1) MAC Filtering DB (VLAN2) MAC Filtering DB (VLAN3)

MAC	Port
H1	4
H4	1

MAC	Port
H2	4
H5	2

MAC	Port
H3	4
H6	3



Real implementations: unique filtering database (usually made with a TCAM, which is a single entity in the network device)



Interconnecting VLANs (1)

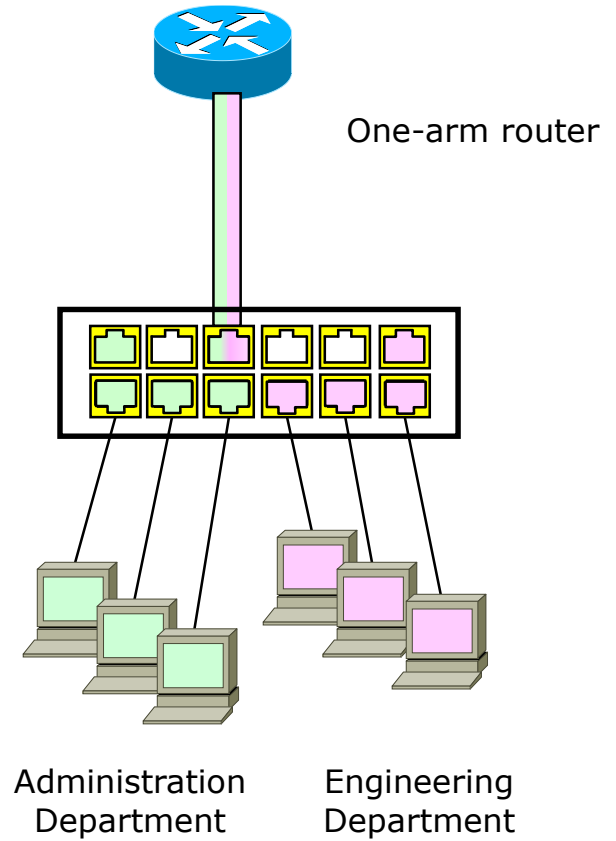


- L2 data cannot cross VLANs
 - An Ethernet station cannot send an L2 frame to another station in a different VLAN
 - VLANs are different broadcast domains

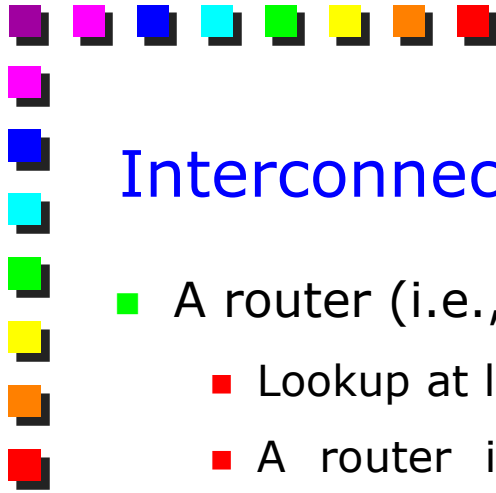
Beware:

L2 data cannot cross VLANs!



Interconnecting VLANs (2)





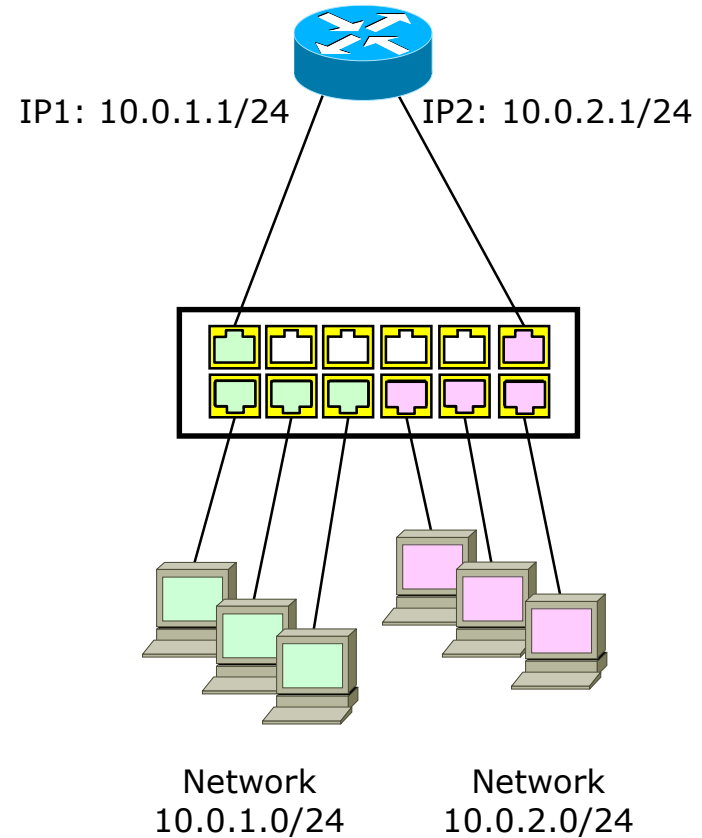
Interconnecting VLANs (3)

- A router (i.e., device operating at layer 3) is needed
 - Lookup at layer 3 (e.g., IP destination address)
 - A router is often used to enforce L3 (or even L4/7) layer protection (e.g. firewall)
 - The original L2 header is thrown away and a new one is created with other MAC addresses (src/dst)



VLANs and IP addresses

- Broadcast cannot cross the VLAN boundaries
 - Cannot use ARP to resolve the MAC address in another VLAN
- Hosts in different VLANs must belong to different IP networks






Associate frames to VLANs (1)

■ Problem

- How can we associate frames to VLANs?


■ VLANs on a single switch

- Simplest method: we can mark the ports on the switch
 - The received frame is associated to the VLAN the port belongs to
- Other methods exist
 - Presented later



Associate frames to VLANs (2)

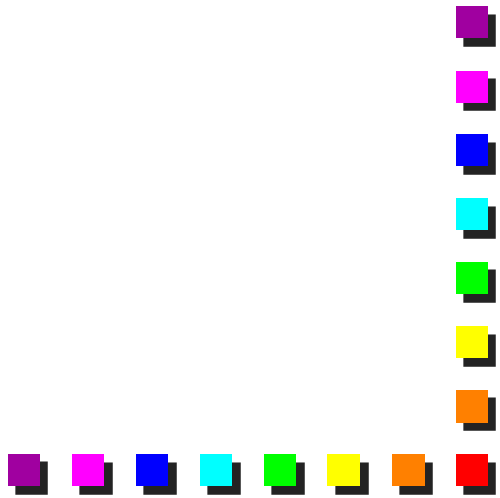
- VLAN on different switches
 - Problem: how to distinguish which VLAN a frame belongs to, as there is a single link between switches?
- Same problem for devices that belong to different VLANs
 - E.g., servers, routers



Note: the IDs in the frames are the MAC addresses of the involved stations




Associate frames to VLANs: tagging


- Required only on links that transport traffic of different VLANs
 - Old method: Tunneling
 - An Ethernet (Token Ring or FDDI) frame is encapsulated into another Ethernet frame
 - Proprietary solutions
 - E.g., ISL (Inter-Switch Link) by Cisco
 - Frame Tagging
 - An additional header is added to the MAC header
 - Standardized by IEEE 802.1Q
 - 4 additional bytes added to the frame
 - Basically, VLAN-ID plus a bunch of other info
- 

IEEE 802.1Q Tag Encoding (1)

VLAN in Ethernet encapsulation (default)






Ethertype for VLAN tagging




VLAN in Ethernet with LLC SNAP

Ethertype for VLAN tagging

IEEE 802.1Q Tag Encoding (2)

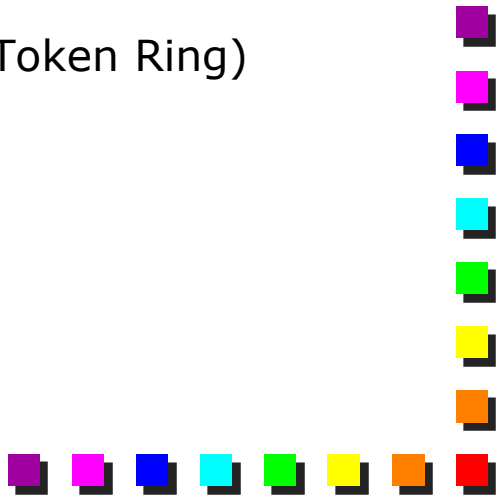


Frame is 4 bytes longer than the one generated by H1






IEEE 802.1Q Tag Encoding (3)

- It can be encapsulated in either Ethernet (DIX) or any link layer using LLC SNAP
 - In both cases, it uses the Ethertype 0x8100
 - The frame has IEEE 802.1Q tag
 - Called TPID (Tag Protocol Identifier)
 - PCP (Priority Code Point)
 - Refers to IEEE 802.1p priority
 - CFI (Canonical Format Indicator)
 - "1": MAC address in non-canonical format (e.g. Token Ring)
 - Usually set to "0" (e.g., Ethernet)
- 

IEEE 802.1Q Tag Encoding (4)

■ VID (VLAN Identifier)

- Values 1- 4094
- Usually, "1" refers to the default VLAN
- 0xFFF: reserved
- 0: the frame does not belong to any VLAN (or I don't know which VLAN this frame belongs to)
 - Used in case the user just wants to set the priority for her traffic





Modification to existing MACs


- Minor modifications
- New framing (for tagging) specified in 802.1Q
 - Independent from the technology of the Medium Access Control
- Maximum length of the frame has to be extended 4 bytes
 - E.g., Ethernet reaches 1522 bytes (from 1518)
 - Minimum length unchanged (still 64 bytes)
 - Hubs cannot handle frames > 1518 bytes




Link types: Access (1)

- *Access Links* receive and transmit *Untagged* frames
- Default configuration (on hosts, switches, servers, routers, etc)
- Usually used to connect end-stations to the network
 - Hosts do not need to change their frame format

Incoming traffic is associated to the VLAN configured on the port of the switch





Link types: Access (2)


- Given the following network
 - All ports are configured in "access mode"
 - SW-1 is configured with the RED VLAN on all its ports
 - SW-2 is configured with the GREEN VLAN on all its ports
- Can host H1 communicate with host H4?



Yes, because values configured on access ports are not propagated outside the switch!

Link types: Trunk (1)

- Trunk links receive and transmit Tagged frames
- Must be configured explicitly
 - Often used in switch-to-switch connections and to connect servers/routers
- Cannot have hubs on trunk ports
 - E.g. Ethernet hubs do not support frames > 1518 bytes





Link types: Trunk (2)

- Tagging on trunk ports

- Different possibilities

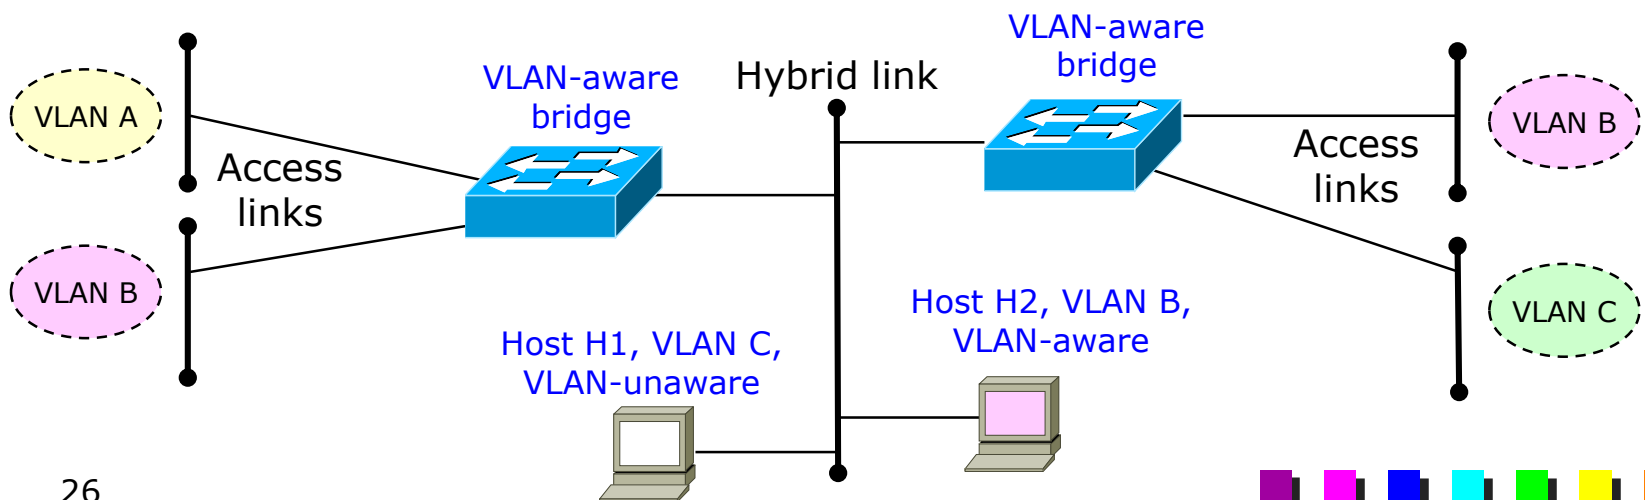
- Some switches tag the traffic belonging to all VLANs
- Other leave the traffic belonging to VLAN 1 untagged

- Another reason of incompatibility between network devices of different vendors




Link types: Hybrid


- Hybrid links accepts both tagged and untagged frames
 - Differentiates frame according to the "type" field (0x8100 or not)
 - Some hosts may not be fully operational (e.g. Station A cannot understand tagged traffic directed to it)
- Trunk links are usually also Hybrid links
- May be used on ports on which both hosts and servers / routers / switches are connected
- In any case, very uncommon nowadays






Assigning hosts to VLANs

- Different methods to associate devices to the proper VLAN
 - Port-based VLANs
 - Transparent assignment
 - Per-user assignment (802.1x)
 - Cooperative assignment
 - Note: a station can be associated also to multiple VLANs
 - E.g., required in case of servers, routers
 - In this case, trunk links are required on the device
 - Frames are tagged directly by the device
 - Fourth assignment method: Configuration of Trunk Interfaces
 - Can be seen as an extension of the Cooperative Assignment
- 




Port-based VLANs (1)

- Most common choice in current networks
 - Each port can be configured as either access port or trunk port
 - Each access port is associated to a single VLAN
 - Each trunk port is associated to a group of allowed VLANs
- Default: all ports in Access mode, associated to VLAN 1



Port-based VLANs (2)

- Completely transparent to the user
 - Association is done on the switch
 - Maximum compatibility, since there is no need to configure hosts
- Different VLANs (e.g., privileges) depending on the actual physical network socket we connect to
- No seamless mobility at L3
 - Host will change the IP address when moved into another VLAN





Transparent assignment

- New criteria in transparent assignment
 - Per L3 protocol (802.1v; no longer useful)
 - Per MAC address
 - Configuration problems
 - Keep MAC database aligned (new host, host with new NIC card, ...)
 - Network administrator has full control on association user-VLAN
 - Allows seamless mobility
- Mainly historical



Per user-assignment (802.1x)

- 802.1x is a standard that enables the network port on the switch only if the user authenticates successfully
- Since the switch knows who is attached to the port, it can assign the proper VLAN to the user
 - E.g., if the switch detects that user U1 connects to the switch, it enables VLAN1
 - Assignment is *per-user*, not *per-host*
 - It looks similar to the per-port assignment, but the coloring is done based on the UserID





Cooperative assignment (1)

- Also known as “anarchic” VLAN assignment
- Users keep control of the VLAN assignment
 - User sets the VLAN on the network card
- Allows seamless mobility
 - User will attach always to the same VLAN anywhere in the campus
- What about a user joining the wrong VLAN?
 - Negligence or bad will
- Used mostly on devices that must be part of different VLANs
 - E.g. routers, servers

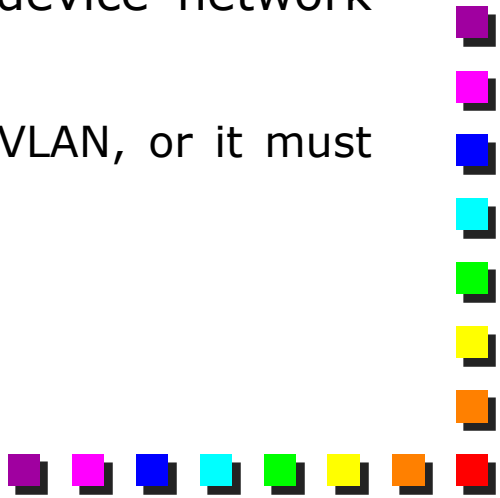


Cooperative assignment (2)

■ Requires


- The (manual?) configuration on all the PCs
- The usage of trunk interfaces
 - Frames are tagged by the user, which sets the right VLAN-ID in outgoing frames
 - In any case, the port on the switch has to be configured anyway with the list of allowed VLANs
 - Often we use "VLAN allow all"

■ Two way of configuring this feature on the device network card

- Depends if the device has to support *a single* VLAN, or it must belong to *multiple* VLANs
- 

Cooperative assignment: single VLAN per NIC


- Simple association of VLAN tagging to the incoming/outgoing traffic
 - Incoming/outgoing traffic is generated with 802.1Q tagging
 - Only one VLAN-ID per NIC interface is allowed (and specified by configuration)
 - Allowed on almost all network cards (e.g., the ones we have in our PCs)
 - We may have multiple cards in case multiple VLANs are required
 - Barely used



Coop. assignment: multiple VLANs per NIC (1)


■ Without VLANs in the host

- Two network interfaces
- Each one with its own IP configuration
- Each one belongs to a different LAN
 - E.g., receives only the broadcast associated with that VLAN




■ With VLANs in the host

- We need to create exactly the same environment that was available before VLANs
- Needed so that the software can operate exactly in the same way
- We had two NICs before, we need two NICs now as well




Coop. assignment: multiple VLANs per NIC (2)

- Requires the usage of virtual NICs
 - Multiple virtual network interfaces are created
 - Each one with its L3 configuration (e.g. IP address) and VLAN-ID
 - Only one VLAN-ID is allowed per virtual card
 - A maximum of N VLANs are allowed (N = number of V-NICs)
 - Widely used; mostly on servers and routers
 - Explicit support required from the NIC driver *and/or* the Operating System
- Important: IP addresses associated to the interfaces (either real or virtual) **must** belong to different IP networks




Trunk Interfaces and IP configuration



Virtual interface Vir1.2
MAC: 00:00:00:11:11:11
VLAN 2
IP: 10.0.1.1/24
DG: 10.1.1.254/24

Virtual interface Vir1.3
MAC: 00:00:00:11:11:11
VLAN 3
IP: 10.0.2.1/24
DG: 10.1.2.254/24



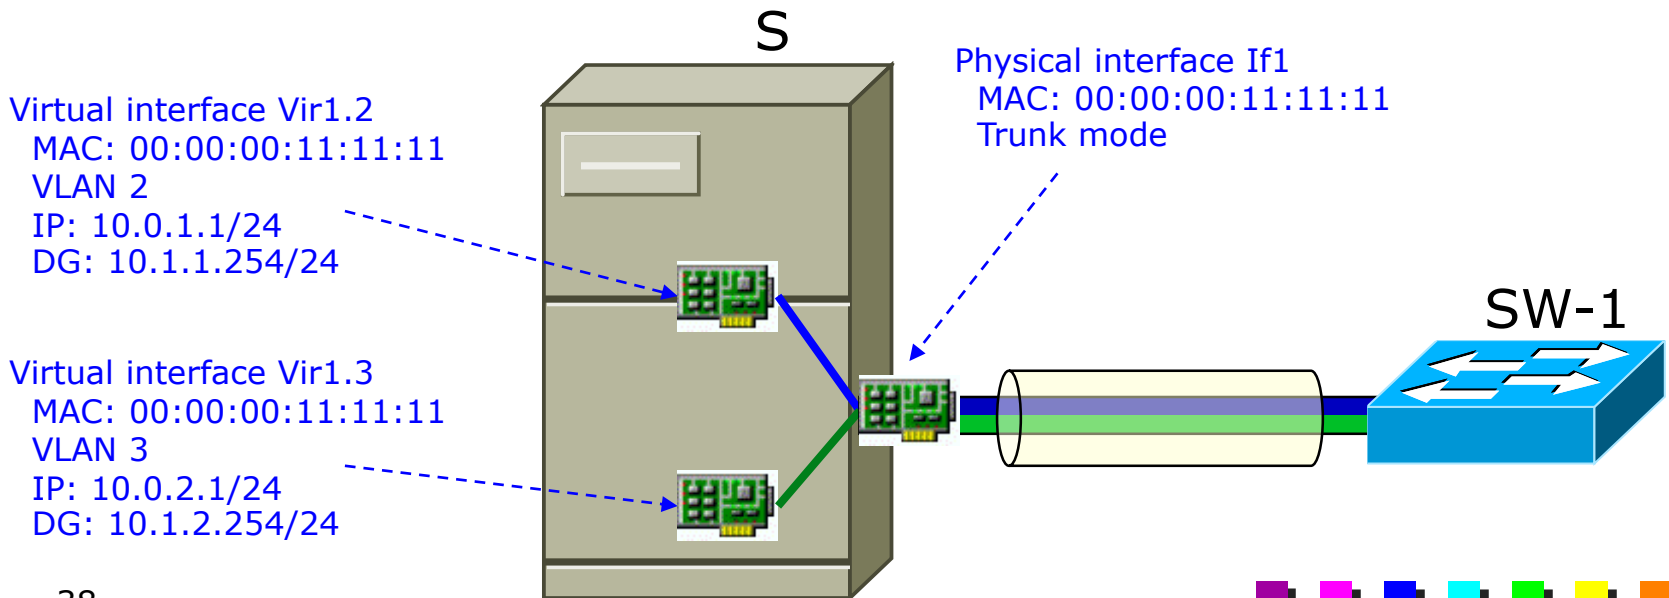
Each interface has its own configuration at IP level (two different IP networks)

Network traffic belonging to the two VLANs is separated and sent to two different virtual interfaces



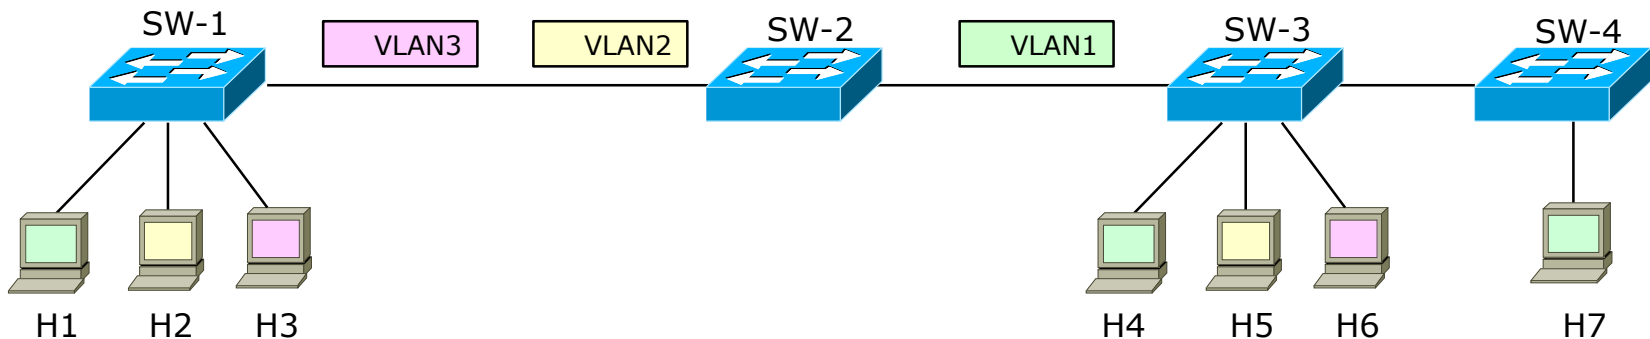
Note: duplicated MAC addresses

- Please note that duplicate MAC addresses are
 - Very common in modern LANs
 - Another common situation is host virtualization (e.g. virtual machines)
 - Do not cause troubles as soon as they belong to different VLANs
 - Switches MUST handle the filtering databases of different VLANs as distinct entities




Assigning VLANs to trunk links (1)

- Necessity to know which VLANs are handled on a given trunk link / switch
 - The switch needs to create the proper number of filtering DB
 - How can SW-2 know that it will have to forward VLANs 1-3?
- Possibility to optimize the number of filtering DB on the switch
 - E.g., FilteringDB for VLANs 2,3 are not needed on SW4
 - Useful to reduce the number of MAC entries on the switches




Assigning VLANs to trunk links (2)

- Possibility to optimize broadcast traffic
 - Avoiding to send *broadcast/flooded* traffic belonging to a VLAN on a switch where no such VLANs are present
 - Unicast (not flooded) traffic is always optimized by the filtering database




Assigning VLANs to trunk links (3)

- The idea: let each switch to know which VLANs are active on its ports
- Three solutions
 - Manual configuration
 - Proprietary mechanisms
 - GVRP



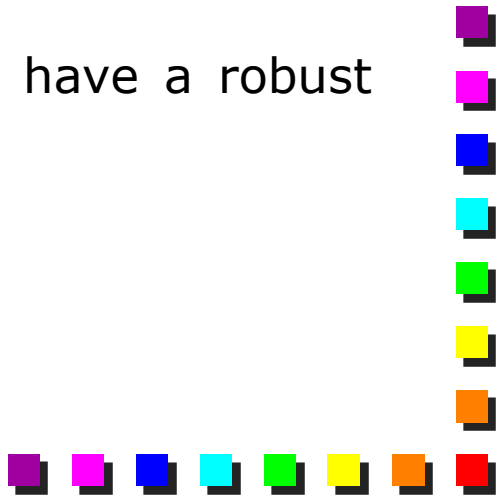
VLANs in the backbone: manual configuration

- Used in most networks
- Usually, VLANs are configured explicitly on each switch
 - Possible problems (related to STP) in case you want to optimize trunk ports and filter useless VLANs out
 - What about if the link between SW-1 and SW-2 is turned off?
 - Better to allow all VLANs on all links and avoid optimizations



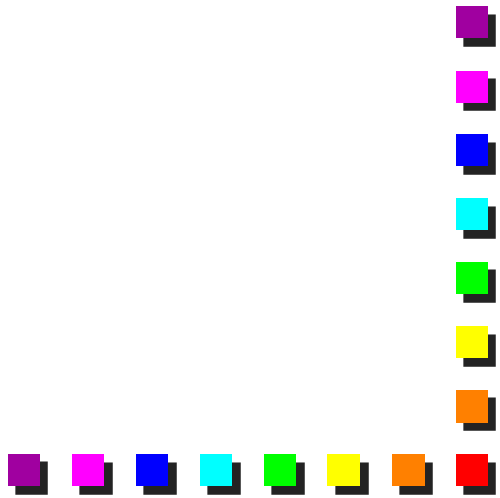


VLANs in the backbone: GVRP

- It propagates info about required VLANs on all the switches
 - Prunes switches that are not interested by some VLANs from the tree of that VLAN
 - Can filter the broadcast traffic of some VLANs on some switches
 - Handy (because automatic), but not widely used
 - It inserts a new level of intelligence in switches
 - Configuration required
 - New software (i.e. bugs)
 - Is it really needed (especially if you want to have a robust network)?
- 

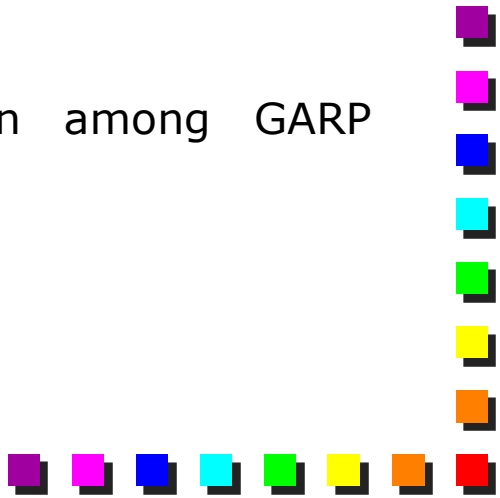



GVRP: GARP VLAN Registration Protocol

- A specialization of GARP: Generic Attribute Registration Protocol
 - Used to register or unregister VLAN related attributes
 - A switch registers the VLANs it “knows” with the switch on the other side of a trunk link
 - Remote switch learns the VLANs whose packets should be forwarded on the trunk link
 - Alternative to static definition of VLANs to be forwarded on a *Trunk Link*
 - Switch using GVRP are said GVRP-Aware
 - GVRP operates on the STP active topology
- 





GARP: Generic Attribute Registration Protocol

- Registers or unregisters various types of attributes into an entity within a switch called GID
 - GID (GARP Information Distribution)
 - Collection of state machines defining the current status of attribute registrations and declarations
 - Attribute registration relates to a port receiving a GARP PDU with the corresponding declaration
 - Also a port set in Blocking state by STP
 - GIP (GARP Information Propagation)
 - Entity in charge of propagating information among GARP Participants
 - Inside a single bridge
 - Among different bridges (based on LLC type 1)
- 

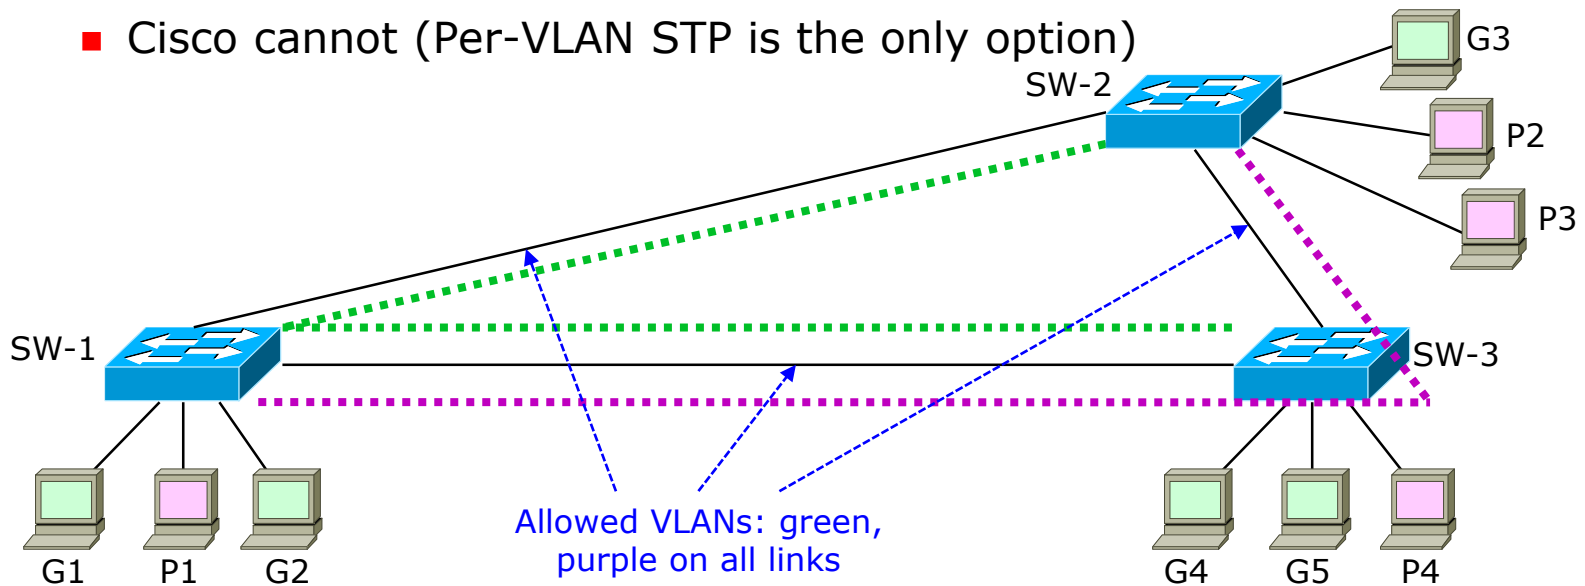


GVRP Packet Format



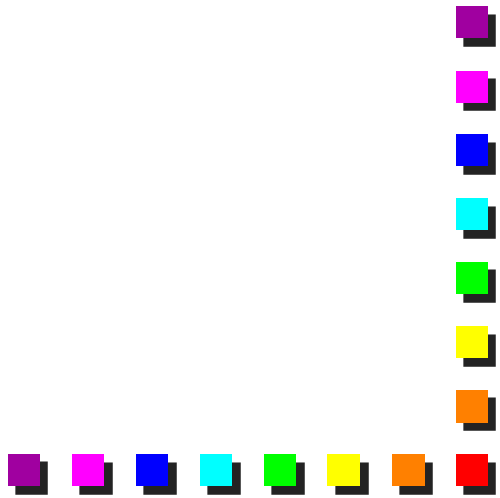
VLANs and Spanning Tree

- In theory, they are completely independent
 - First, Spanning Tree is computed in order to disable loops
 - Then, VLANs are used on the resulting topology
 - Unique forwarding tree for all the VLANs
- Almost all vendors offer Per-VLAN Spanning Tree
 - Most vendors can turn back to an unique STP via configuration
 - Cisco cannot (Per-VLAN STP is the only option)



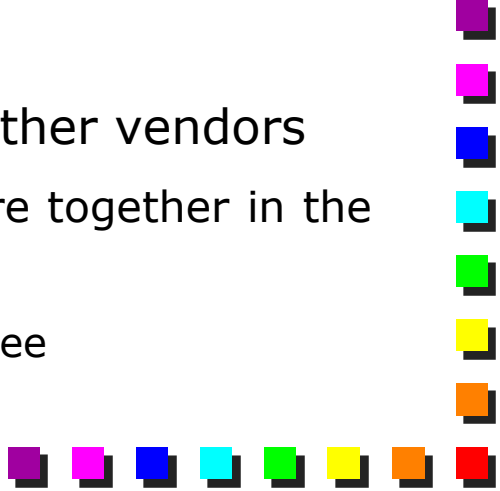


Per-VLAN Spanning Tree




- Allows multiple spanning trees in the network
 - Network optimization
 - Requires a per-VLAN configuration of ST parameters
 - Bridge priority, at least, for differentiating the root bridge among different VLANs (otherwise it will result the same tree for all VLANs)
 - Other parameters (Hello Time, Max Age, ...)
 - Increase the load on the CPU of the switches
 - N instances of ST at the same time
 - Mostly proprietary solutions
- 



Cisco and Per-VLAN STP


- Cisco offers different solutions for per-VLAN STP
 - PVST: old solution, which uses a proprietary technology to transport frames on trunk links
 - ISL, which operates by tunnelling frames instead of the newest tagging proposed by 802.1Q
 - PVST+: similar to the previous solution but uses 802.1Q tagging on trunk links
 - Rapid-PVST+: applies the per-VLAN STP idea to the RSTP, achieving fast convergence
 - Uses 802.1Q tagging on trunk links
 - Interoperability problems between Cisco and other vendors
 - Broadcast storm if devices with/without PVST are together in the same network
 - Some VLANs may not have a valid Spanning Tree
- 

PVST: example of the new BridgeID



Priority 32768 for VLAN 2

```
Cisco(config)# spanning-tree vlan 2 priority 32768
```

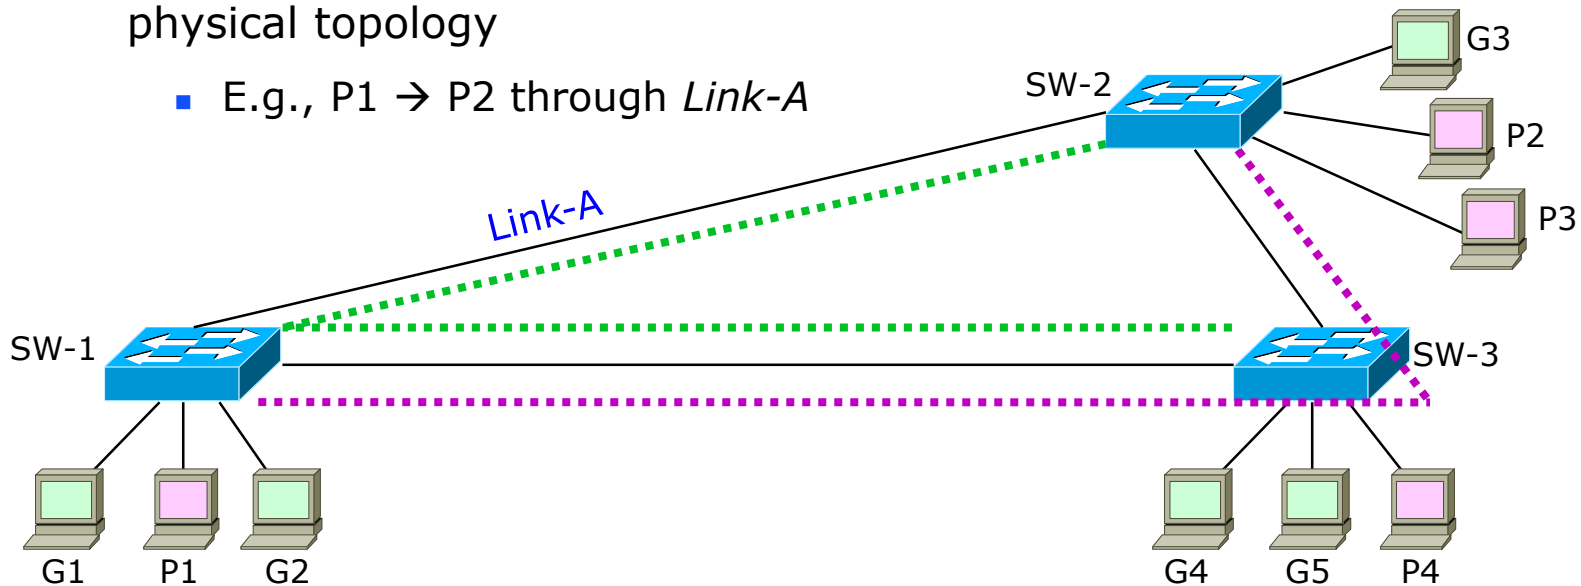


Priority 28672 for VLAN 3

```
Cisco(config)# spanning-tree vlan 3 priority 28672
```

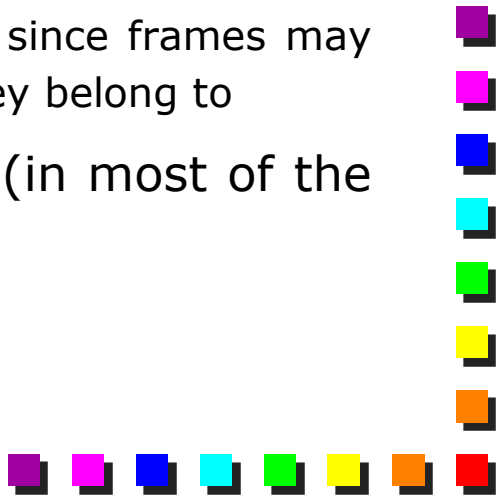
PVST: some notes on traffic optimization

- Optimizes the network load across the network
 - If spanning trees are well balanced, all the links are utilized
 - No longer have links that are in place but not used
- Does not optimize the network load inside a VLAN
 - The VLAN traffic is still bounded to a specific forwarding tree
 - We cannot use the “shortest path” that may be available on the physical topology
 - E.g., P1 → P2 through *Link-A*






Additional considerations on PVST

- Do we really need PVST?
 - PVST can optimize link utilization, but is this needed?
 - From practical experience
 - Bandwidth is no longer a problem in most of modern LANs
 - Having multiple topologies does not provide significant advantages in the network
 - Instead, multiple topologies
 - Complicate the troubleshooting
 - Difficult to understand the path of the traffic, since frames may cross different links depending on the VLAN they belong to
 - Consequence: better to stay with a single ST (in most of the cases)
 - “Keep it simple” is often the best strategy!
- 

VLANs and network isolation (1)


- Network isolation is not complete, even with VLANs
 - Although frames cannot cross the border of a VLAN, links are shared, hence a problem on a *link*, caused by the traffic of one VLAN, may affect other VLANs






VLANs and network isolation (2)

- For example, VLANs do not protect from broadcast storms
 - In fact, *broadcast traffic* is sent on the entire network
 - Except on the edge ports, since those are assigned to a specific VLAN
 - A trunk link may be saturated by a broadcast storm on a VLAN
 - Other VLANs do not receive that broadcast but...
 - ... the trunk link is congested and it may be unable to transport the traffic of other VLANs
- Per-VLAN QoS may be required
 - E.g., “Round-robin” service model based on VLAN ID, which guarantees a minimum amount of bandwidth to each VLAN






VLANs and network switches

- Two types of switches
 - VLAN-Aware: handle tagged and untagged frames
 - VLAN-Unaware: do not accept tagged frames
 - May discard frames (if too big)
 - Low-end devices
 - Availability on the market
 - Almost all professional products can handle VLAN tagging
 - Almost all domestic products do not have VLAN support
 - VLANs are no longer a “plug and play” technology
 - STP was (with some limitations)
 - This is one of the reasons VLANs are not supported on domestic switches
 - Typical users are not skilled enough to configure them
- 


Mixing VLAN-aware/unaware switches (1)




Problems:

- 1) H1-H3 cannot exchange data with H4-H6
- 2) H1-H3 can exchange data only with S1
- 3) H4-H4 exchange data only with S2

Mixing VLAN-aware/unaware switches (2)



Mixing VLAN-aware/unaware switches (3)



VLAN-unaware switches may be OK in the access side (e.g., in order to add new ports), provided that all clients belong to the same VLAN

Corollary

It is pretty common to have VLAN-unaware switches in corporate networks.

Network managers typically deploy only professional switches (with VLAN support) but often end users have some limitations (e.g., necessity to attach multiple hosts on a single network socket) and tend to sort those problems out by themselves, which usually means they buy the cheapest switch on the market, which does not have VLAN support.

Therefore, it is important that the network manager takes into account those situations (even if he does not know exactly where those switches may be installed) in order to prevent possible misbehaviour of the network.



Configuring VLANs on Cisco switches (1)

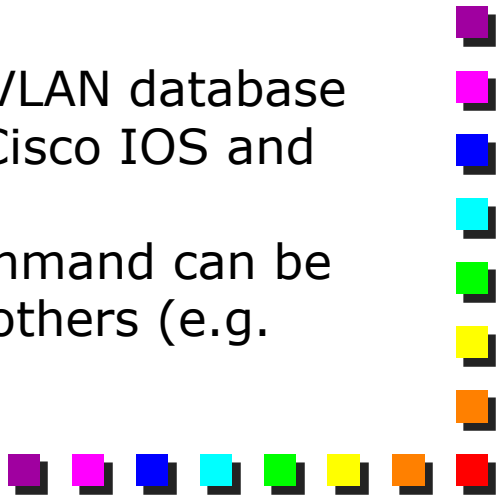
■ VLAN creation

```
Switch# vlan database
Switch(vlan)#vlan 2 name Administration
  VLAN 2 added:
    Name: Administration

Switch(vlan)#exit
  APPLY completed.
  Exiting....
switch#
```

Note: the command for adding an entry in the VLAN database changes according the different version of the Cisco IOS and given device in use.

In more modern devices, the *vlan database* command can be issued also in standard configuration mode. In others (e.g. Cisco 6500) the command is even different.






Configuring VLANs on Cisco switches (2)

■ VLAN port association

- Default behavior: a port is considered Access and associated to a default VLAN
- The switch has a VLAN-unaware behavior

```
Switch# configure terminal
Switch(config)#interface FastEthernet 0/1
Switch(config-if)#switchport access vlan 2
Switch(config-if)#exit
Switch# show vlan brief
```

VLAN	Name	Status	Ports
1	default	active	Fa0/2, Fa0/3, Fa0/4
2	Administration	active	Fa0/1





Configuring VLANs on Cisco switches (3)

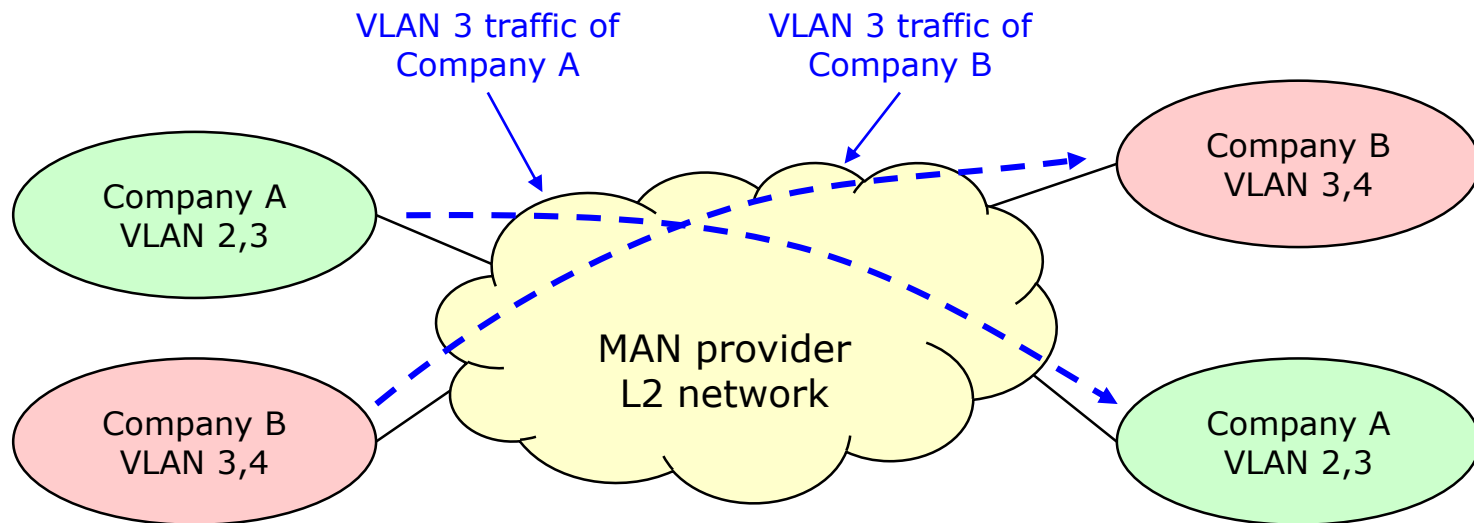
- Configuration of the trunk port

```
Switch# configure terminal
Switch(config)# interface FastEthernet 0/2
Switch(config-if)# switchport mode trunk
Switch(config-if)# switchport trunk allowed vlan
                    add 1,2 [or "all"]

Switch(config-if)# exit
Switch#
```


VLAN Tag Stacking (1)

- Possibility to add multiple VLAN tags to an Ethernet frame
 - Original 802.1Q specs allow a single VLAN tag
 - Possibility to define a “VLAN tag stack”
- Useful when an L2 network has to transport L2 traffic from different entities
 - E.g., Metro Ethernet

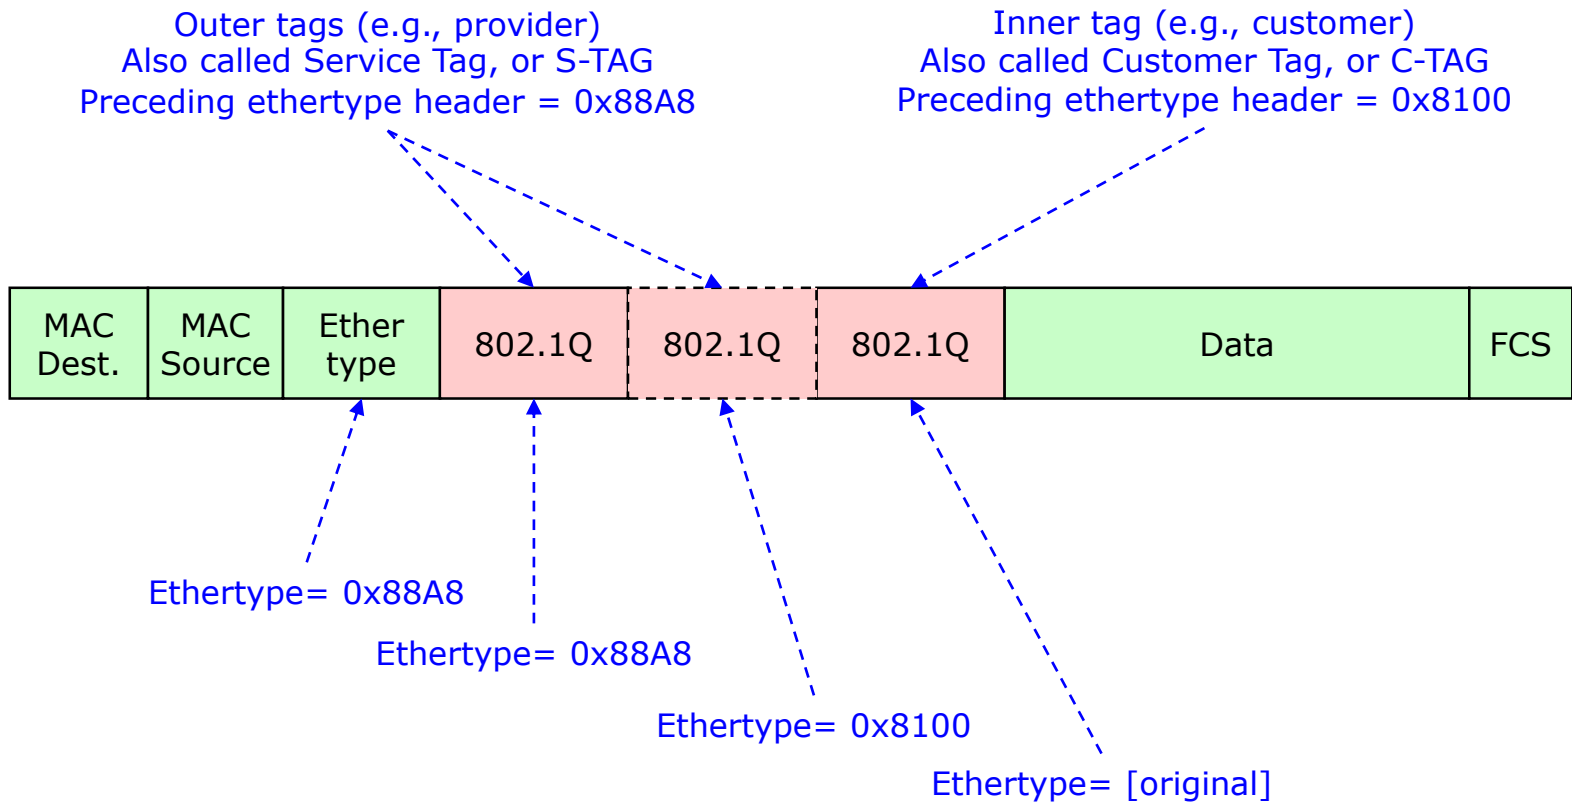


VLAN Tag Stacking (2)


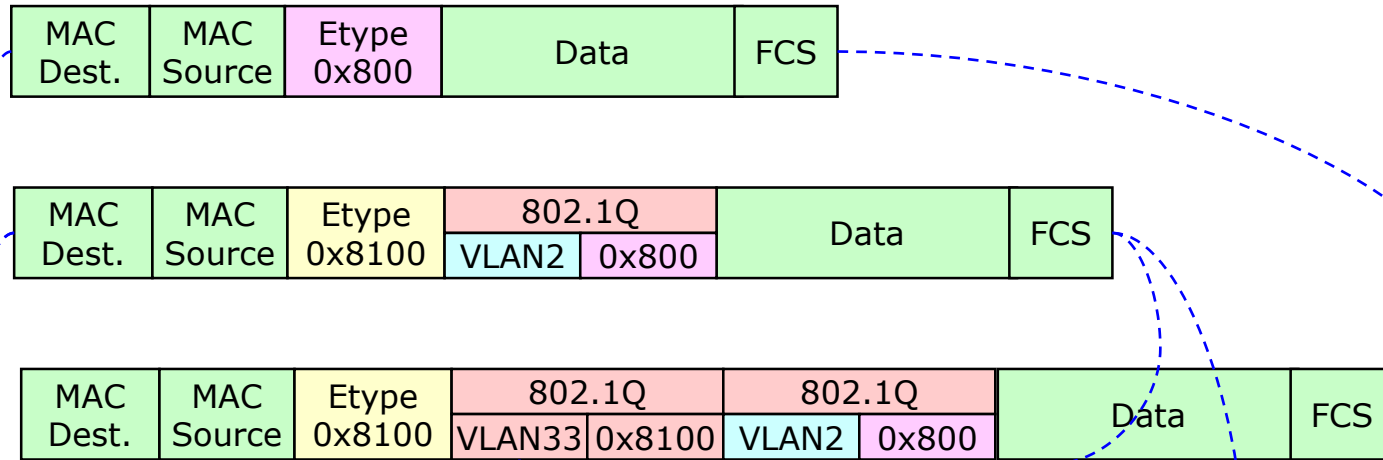
- Formally known as 802.1ad
- Also known as
 - Provider bridging
 - Stacked VLANs
 - QinQ (or Q-in-Q)



VLAN Tag Stacking: ethertype values




VLAN Tag Stacking: example






VLAN Tag Stacking: advantages

- Extends the range of VLAN-ID
 - Originally 12 bits, which may not be enough in large installations
 - Each stack has its own PRI field
 - Can define different priorities per stack
 - Much more flexible (and less disruptive) than defining another tagging format with a larger VLAN-ID
 - Allow different entities to define the same VLAN-ID
 - A common provider can transport all those frames with an additional level of stacking
 - E.g., metro Ethernet, private LANs in an airport, etc.
 - Easy to add/remove tags with *push* and *pop* operations
- 

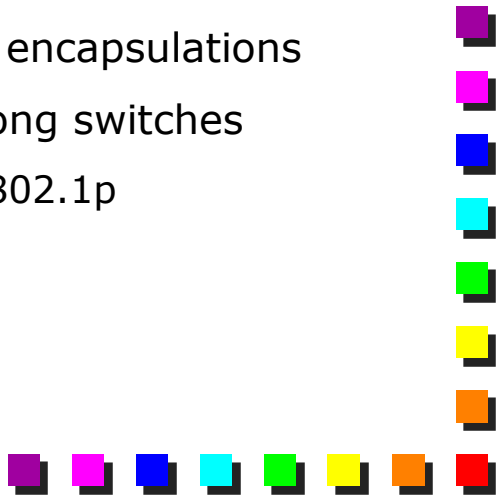


VLAN Tag Stacking: problems

- The intermediate provider should learn all the MAC addresses of all the entities
 - E.g., MAN provider must know all the MAC addresses used by Company A and Company B
 - Not scalable
 - Solutions based on forwarding traffic per-VLAN, instead of per-MAC
 - Broadcast storms on one company could impair the traffic on other companies
 - Intermediate provider should enforce some QoS mechanism
- 



Relevant standards (1)

- IEEE 802.1Q
 - Defines VLAN-aware bridges
 - Per port based VLAN assignment
 - For both “access” and “trunk” ports
 - Unique spanning tree
 - Multiple filtering database identified by FID (Filtering Identifier)
 - Only one entry per MAC address can be present in each filtering database
 - A MAC Address may be present in different filtering databases
 - Defines the VLAN tag (User Priority, CFI, VLAN-ID), encapsulations
 - GVRP for propagating VLAN related information among switches
 - Based on the more general GARP defined in IEEE 802.1p
 - GARP = Generic Attribute Registration Protocol
 - GVRP = GARP VLAN Registration Protocol
- 



Relevant standards (2)

- IEEE 802.3ac
 - Defines the new Ethernet frame format
 - Includes the VLAN tagging
 - Extends the maximum frame from 1518 to 1522 bytes
- IEEE 802.1p
 - Packet priority field, whose use is specified by IEEE 802.1p
- IEEE 802.1ad
 - VLAN Tag Stacking



Conclusions

- Pervasive technology
 - Network isolation useful in many cases (privacy, security, management, ...)
 - Not complete isolation, though
 - Broadcast storm in one VLAN
 - QoS enforcement not always easy
 - Requires a router for interconnection
 - Many form of incompatibilities between different vendors
 - Better to select a single vendor for the entire L2 network
- 