VLANs: advanced topics

Fulvio Risso

Politecnico di Torino

Copyright notice

Part I

VLAN Stacking

VLAN Stacking – 802.1ad

- Enables a backbone provider (e.g., metro Ethernet operator) to transport L2 tagged traffic coming from customers
- Standardized with 802.1ad
  - Also know as Provider bridging, Stacked VLANs, QinQ (or Q-in-Q)
- Requires 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"

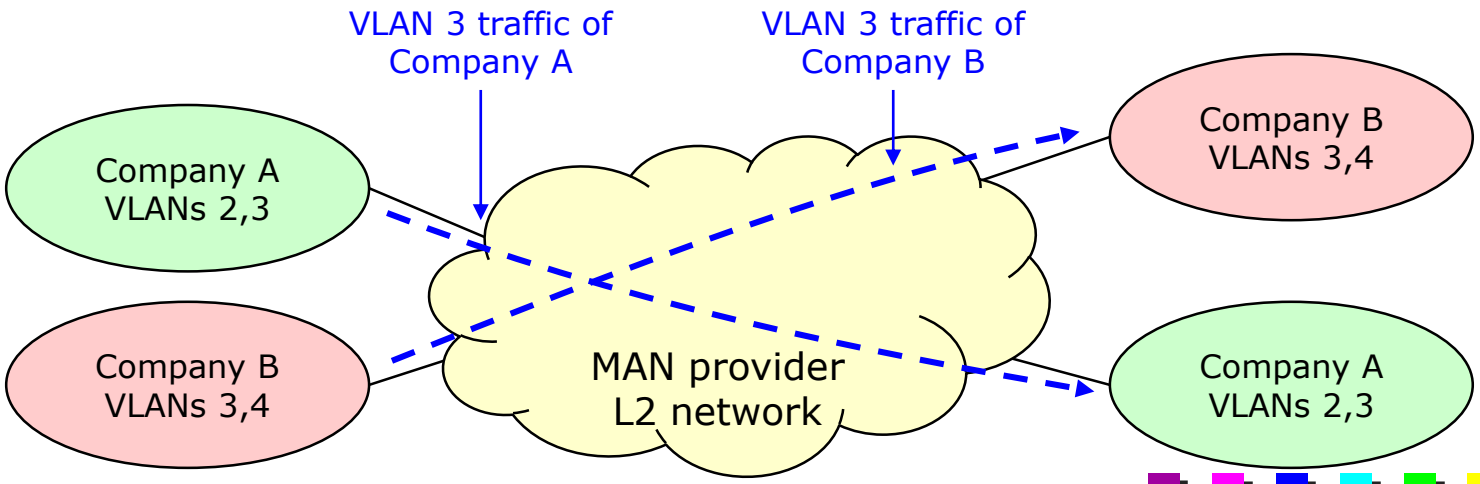

VLAN Stacking: frame format

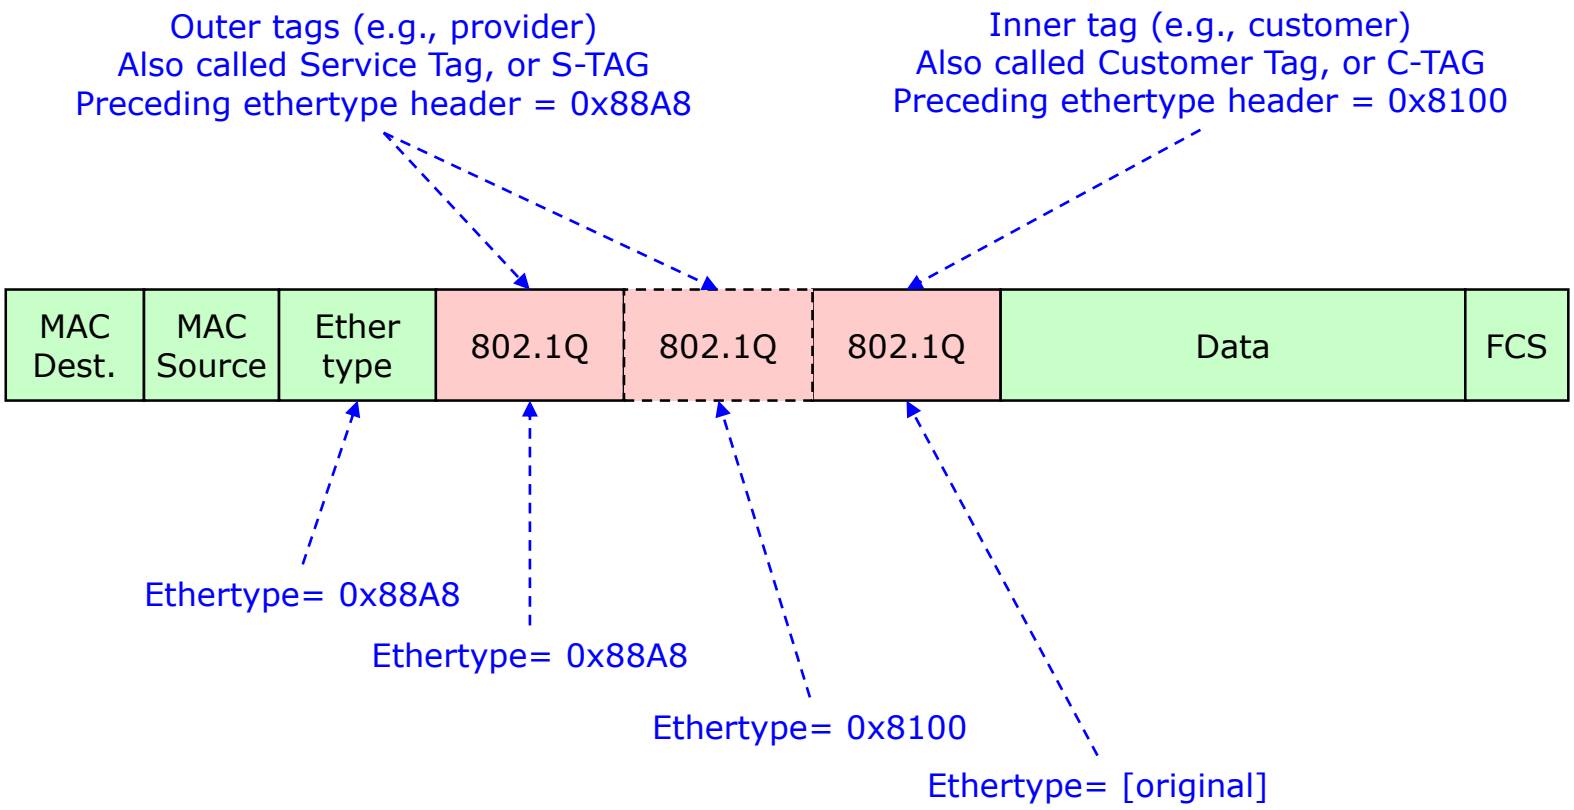

VLAN Stacking: example

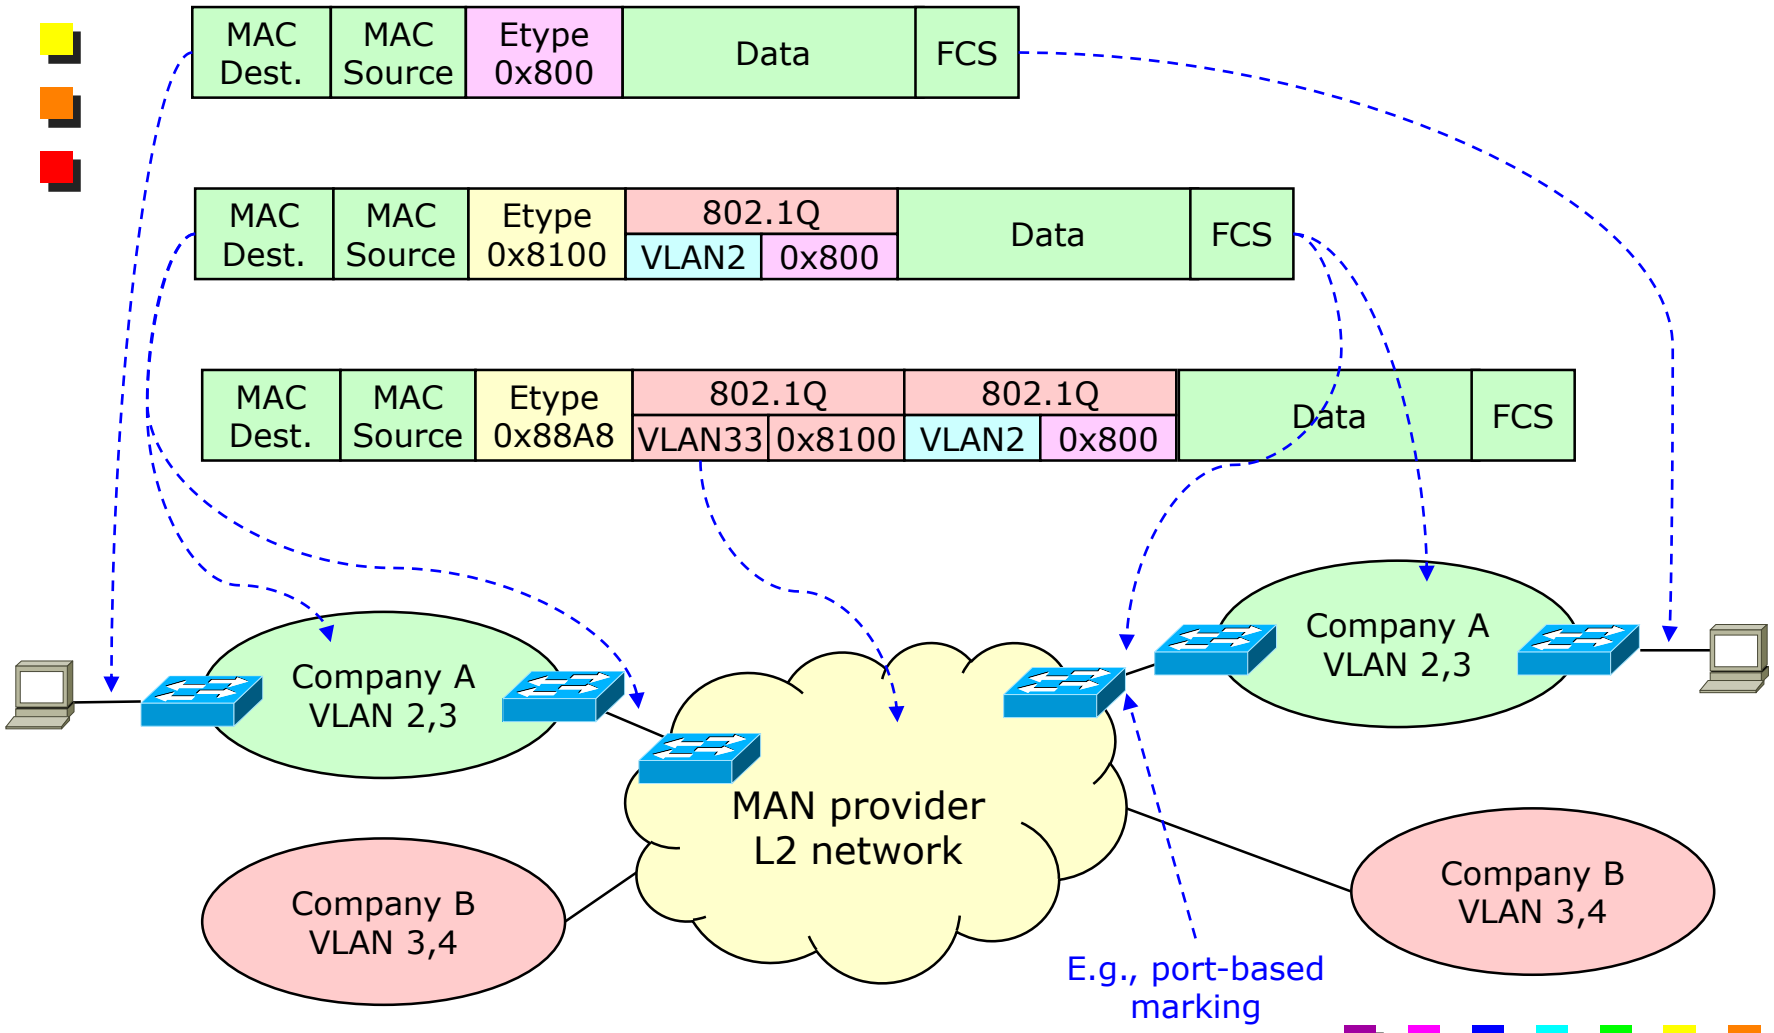

VLAN Stacking: example

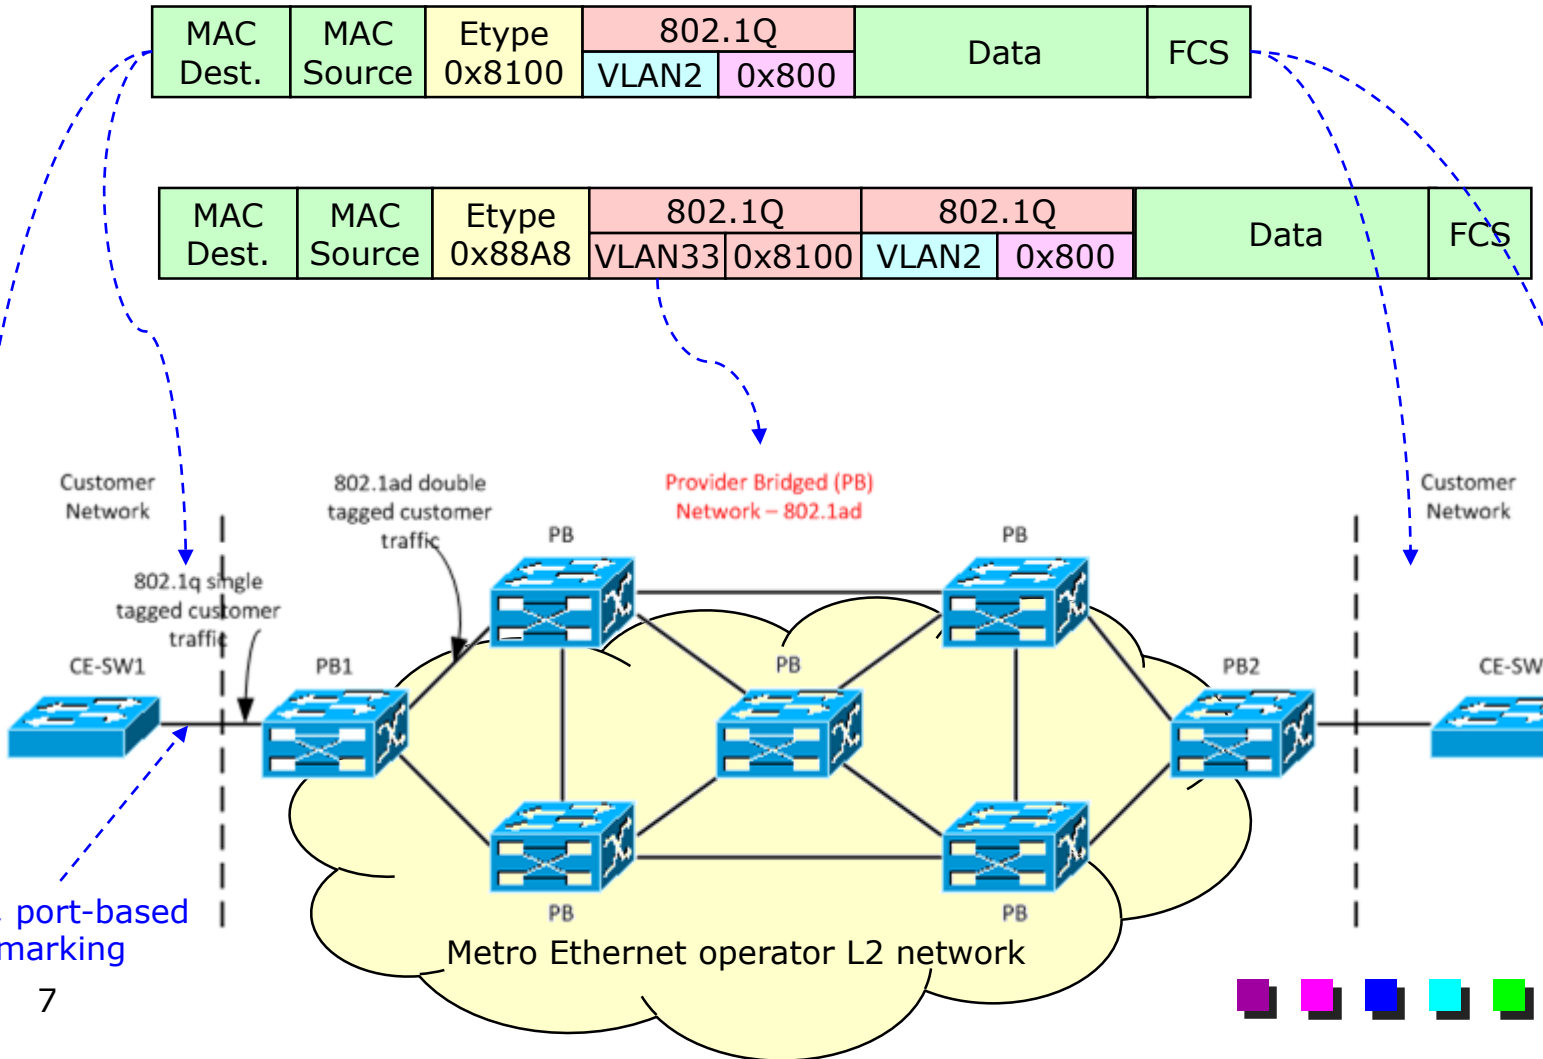

VLAN Stacking: advantages

- 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 Stacking: problems

- Metro provider can support up to 4096 customers
  - Unless an additional tag is used
- Metro provider knows the MAC addresses of all customers
  - I.e., metro provider must know all the MAC addresses used by Company A and Company B
  - Not scalable (quicky reaching too many MAC addresses)
- Broadcast storms on one company could impair the traffic on other companies
  - Intermediate provider should enforce some QoS mechanism

Problems with duplicated MAC addresses

- A duplicate MAC address in the customer edge is not a problem as long as it belongs to different VLANs
- However, it appears undistinguishable when the frame enters in the metro network, as both frames are tagged with customer VLAN (inner VLAN tag is not used for forwarding)
  - The metro network will see the same MAC potentially coming from different sources

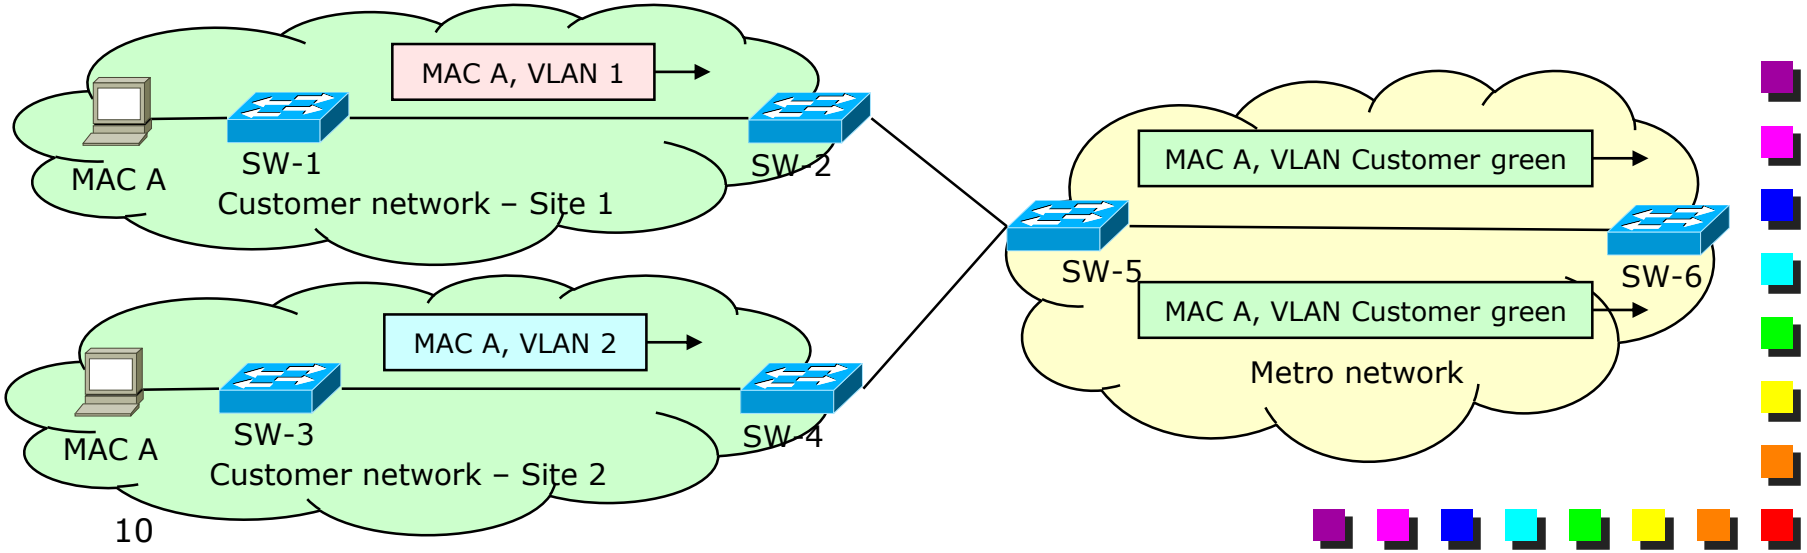

Provider Backbone Bridges (PBB)

- Also known as
  - MAC-in-MAC
  - IEEE 802.1ah
- Overcomes the limitations of the 802.1ad

Part II

Private VLANs

General concept (1)

- Sometimes, a huge number of (tiny) VLANs are required for **security**, **traffic isolation**, **performance**, etc.
  - Hotel rooms
  - Connectivity from a single basement switch to each apartment in private buildings
  - From different servers belonging to different customers within a single rack in the datacenter
  - Ethernet-based ADSL DSLAMs

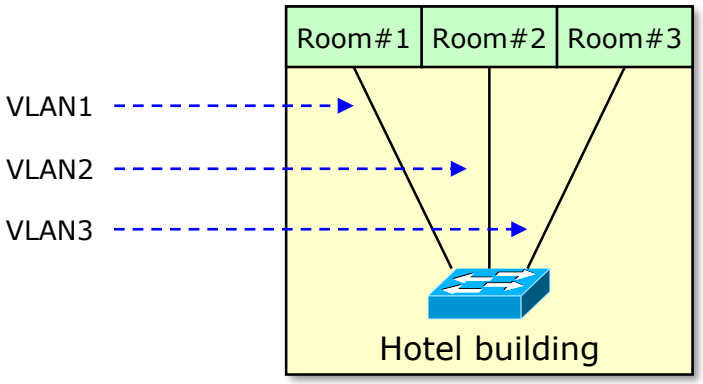

General concept (2)

- Problem: too many VLANs to configure
  - Hard to configure (error prone), hard to maintain
  - In some cases not enough VLANs (hotels with > 4094 rooms)
  - Address exhaustion when using /30 IP networks
- Possible solution: Private VLANs
  - Proprietary, no standard exist
  - Also called “Port isolation”

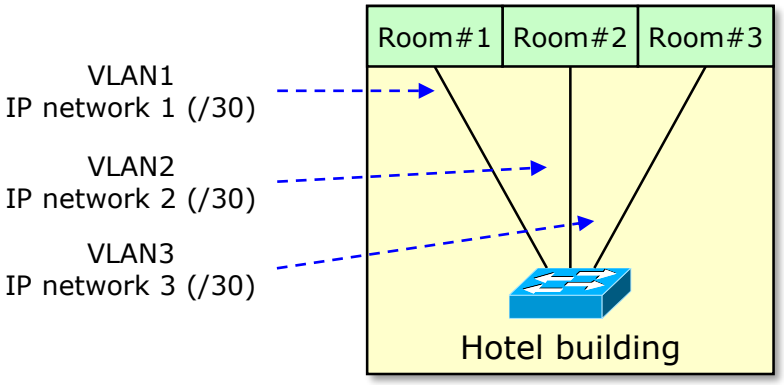

Private VLAN Edge (PVLAN Edge)

- Very simple segmentation mechanism implemented in Cisco switches, when full PVLAN is not needed
  - All the L2 traffic coming from **protected edge ports** cannot be forwarded to another protected edge port
  - All the L2 traffic coming from non-protected edge ports can be forwarded to any other port
  - Potentially interoperable, as configured on a per-switch fashion
- Consequence: direct peer-to-peer traffic between **protected** peers **through the switch** is blocked

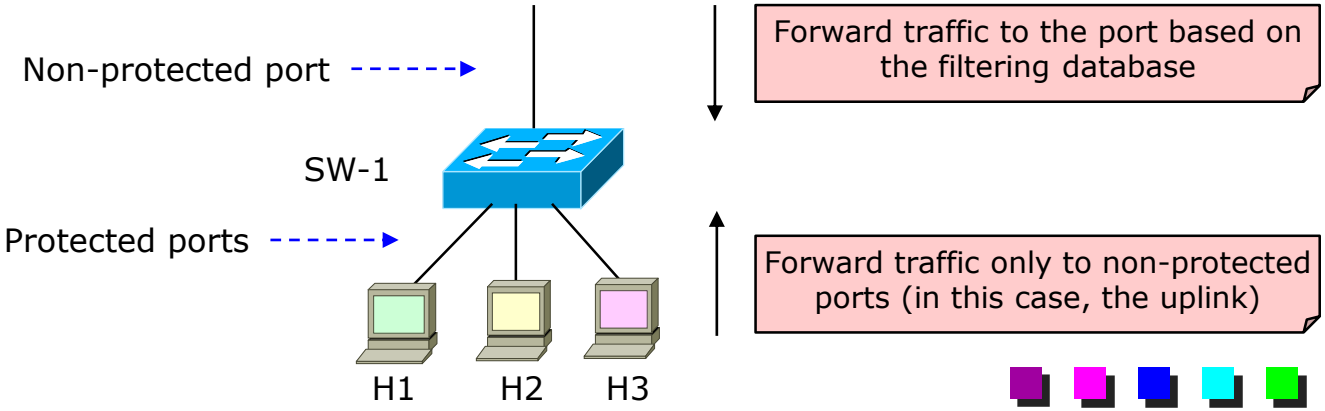

PVLAN Edge: possible usage

- Allows clients to get access to "shared" resources, while any direct (L2) communication between hosts is blocked
  - E.g., shared printer and default gateway

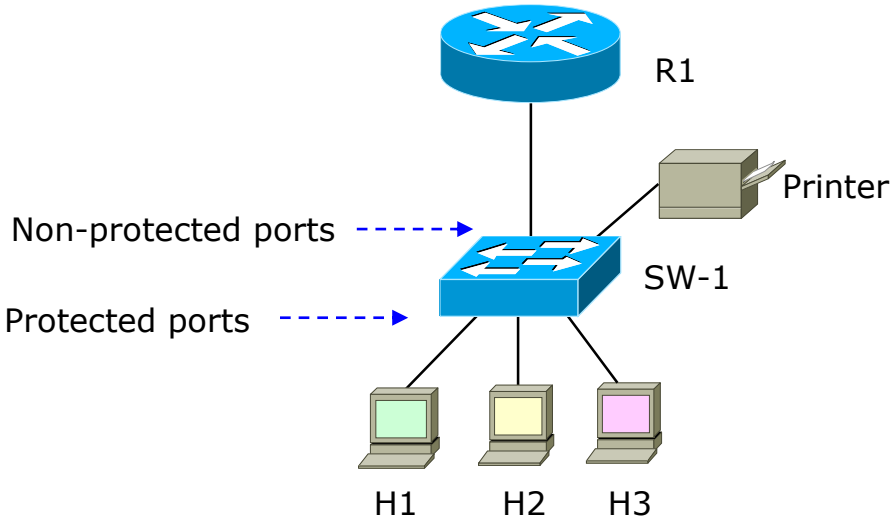

Communication between protected hosts

- PVLAN Edge provides isolation between peers at L2
- In theory, communication still possible at higher layers (router, firewall, etc.) depending on their configuration
  - In practice, not easy to obtain (device should answer to ARP requests on behalf of protected hosts)
  - Note that communication cannot be done through a traditional L2 switch: switches never forward a frame back to the same port it has been received from!

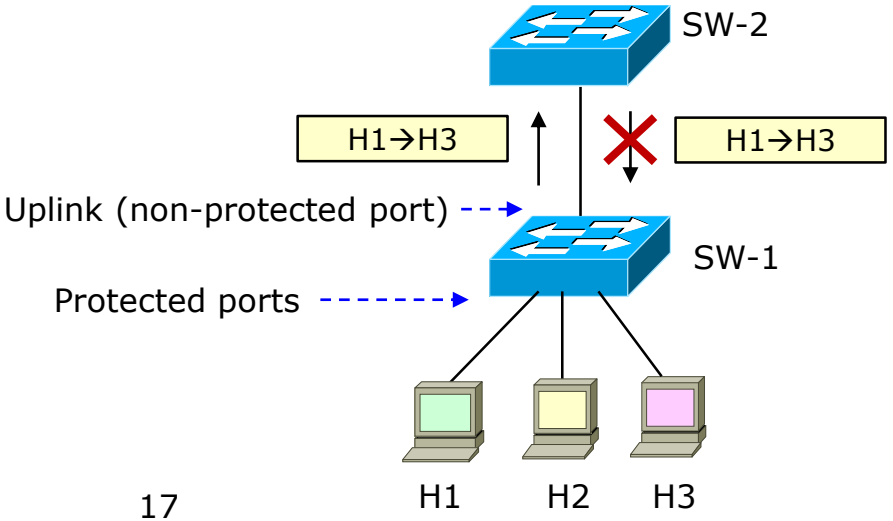

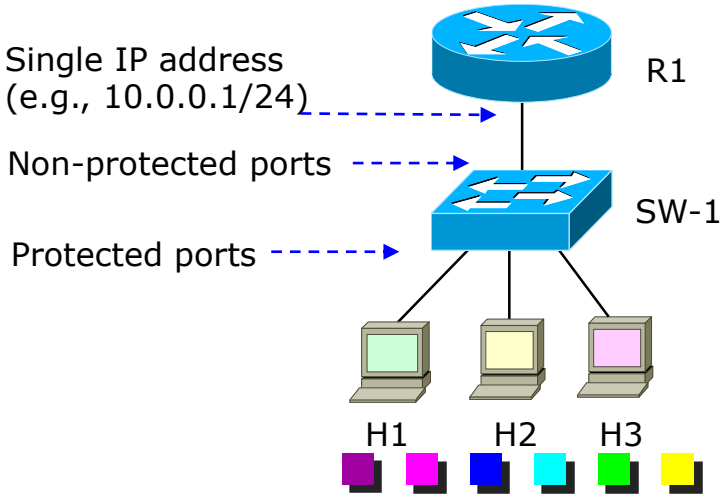

PVLAN Edges: example

```
Router
!
 interface FastEthernet 0
   ip address 10.0.0.1 255.255.255.0
!

Switch
!
 interface FastEthernet 0/0
   switchport mode access
!
 interface range FastEthernet 0/1 - 24
   switchport mode access
   switchport protected
   switchport block unicast
   switchport block multicast
!
```

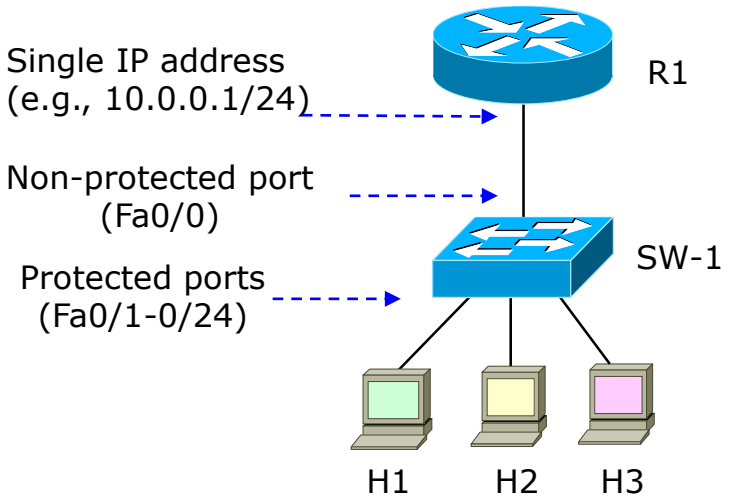

Usually, ports configured as protected are also configured not to receive unknown unicast (frame with destination MAC address not in switch's MAC table) and multicast frames flooding for added security.

PVLAN Edges across switches: the good

- Protection can be achieved by setting up another switch (SW-1) that has protected ports toward edge switches (SW-2 and SW-3)
  - H1-H6 are protected and cannot exchange any data
  - H1-H6 can connect to the corporate router and the printer

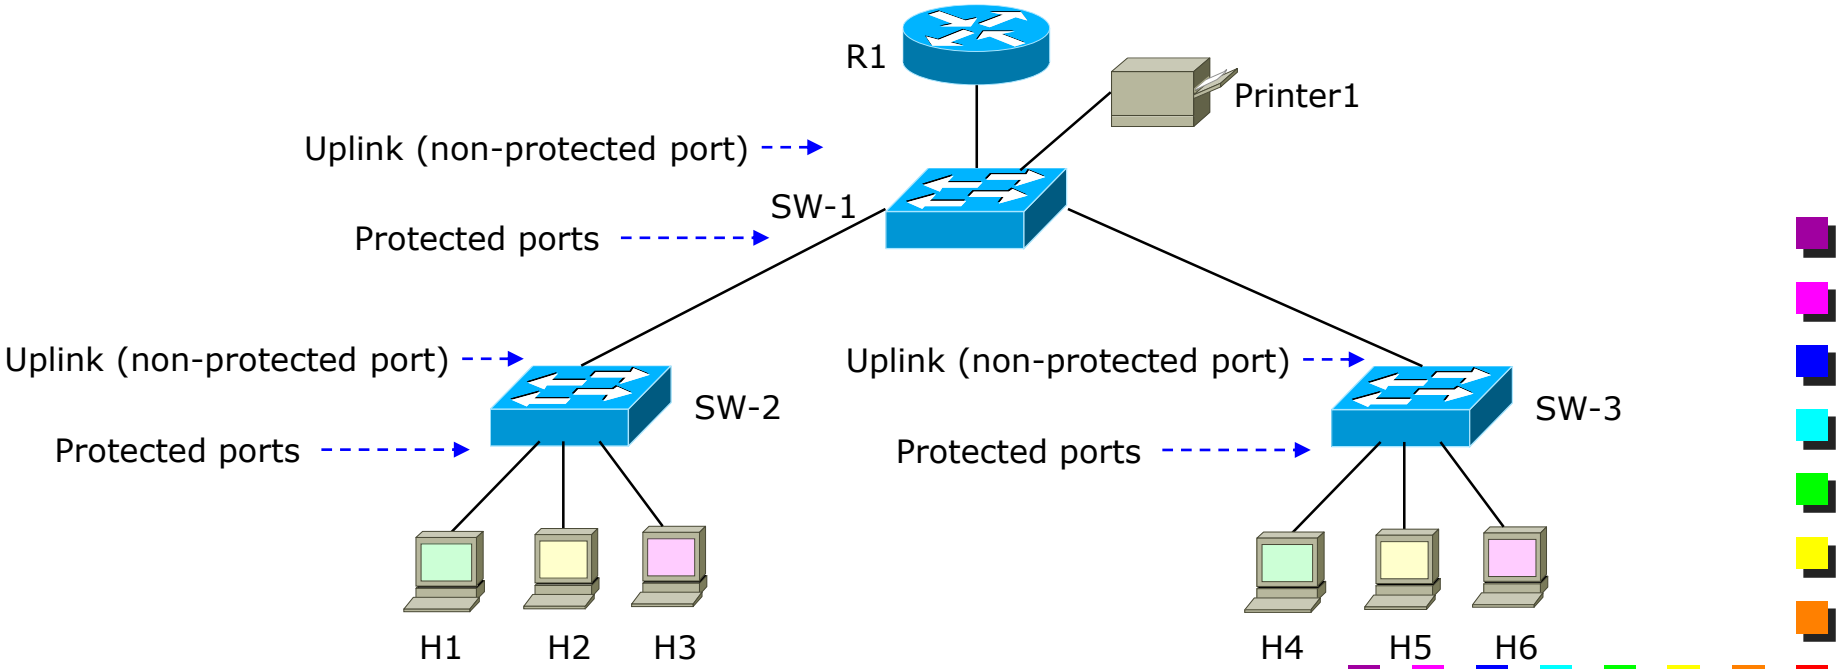

PVLAN Edges across switches: the bad

- Limited flexibility when deployng services across the network (e.g., a printer)
  - E.g., H1-H3 cannot access Printer 2 (traffic coming from the protected port on the right of SW-1 cannot be delivered to the protected port on the left)

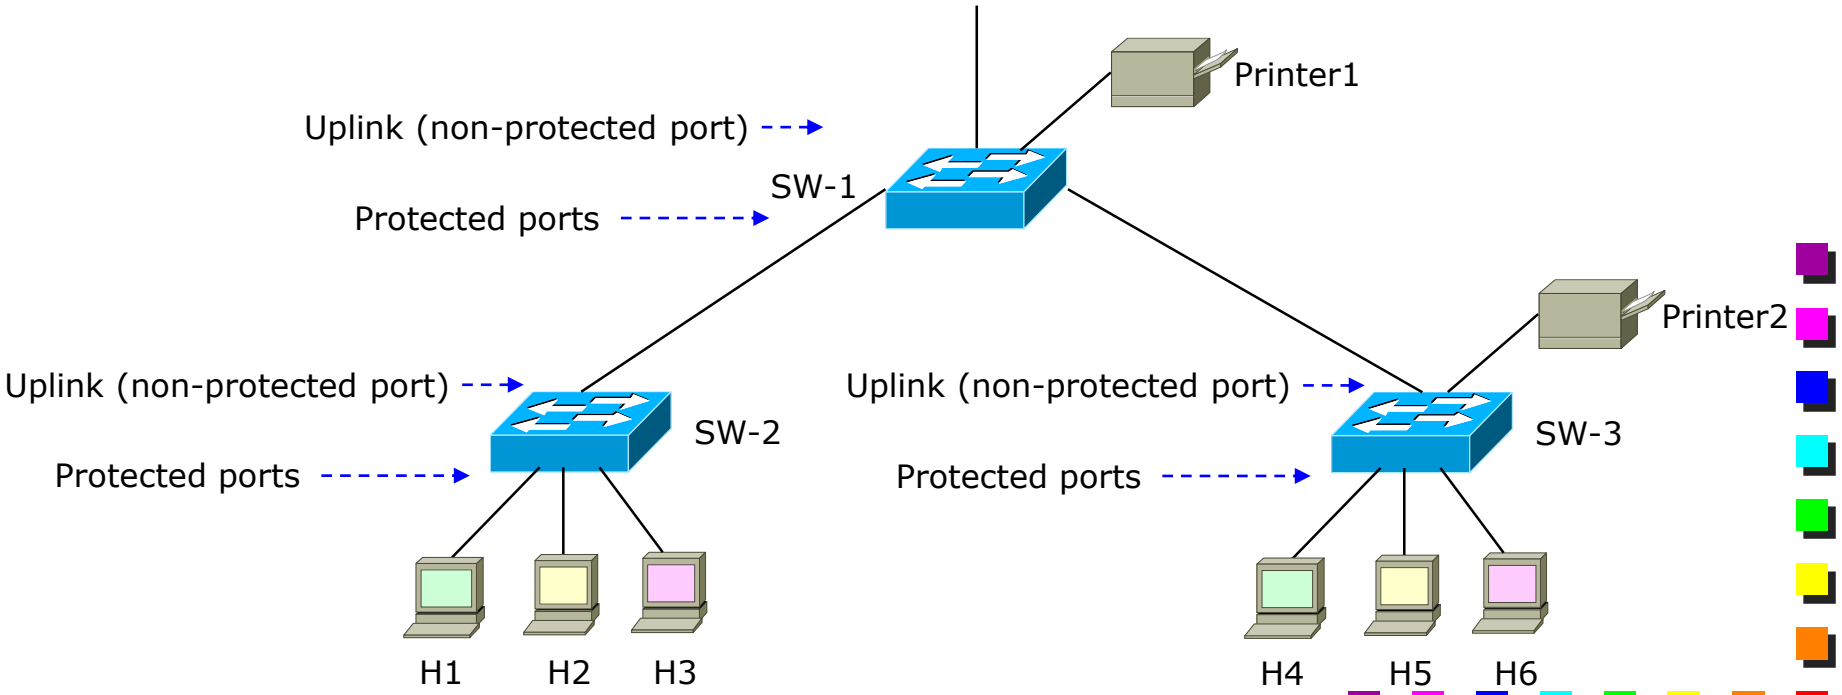

Private VLANs

- Cisco proprietary
- Extends PVLAN Edge with a more general concept
  - Appropriate when the simpler PVLAN Edge is not enough
- Each Private VLAN domain has:
  - **One** primary VLAN
  - At most **one** secondary **isolated** subdomain
  - An arbitrary number of secondary **community** subdomains

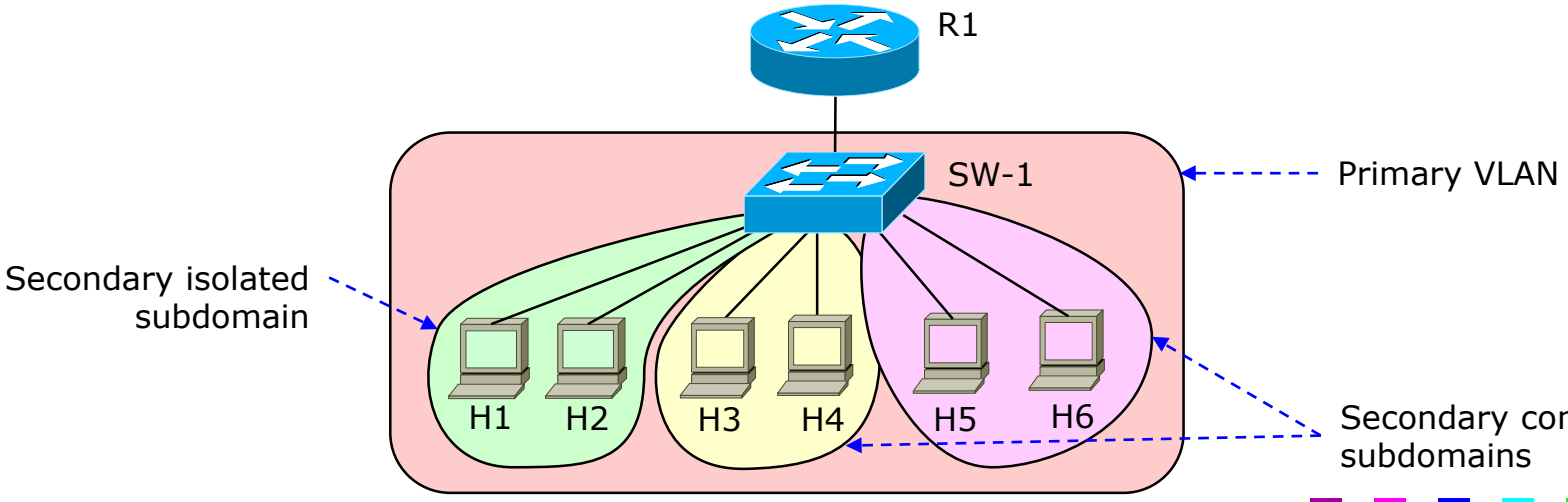

PVLAN: secondary subdomains

- Secondary VLANs provide L2 isolation between ports within the same private VLAN domain:
  - Isolated VLAN: ports within an isolated VLAN cannot communicate to each other
  - Community VLANs: ports within a community VLAN can communicate to each other, but not to other communities

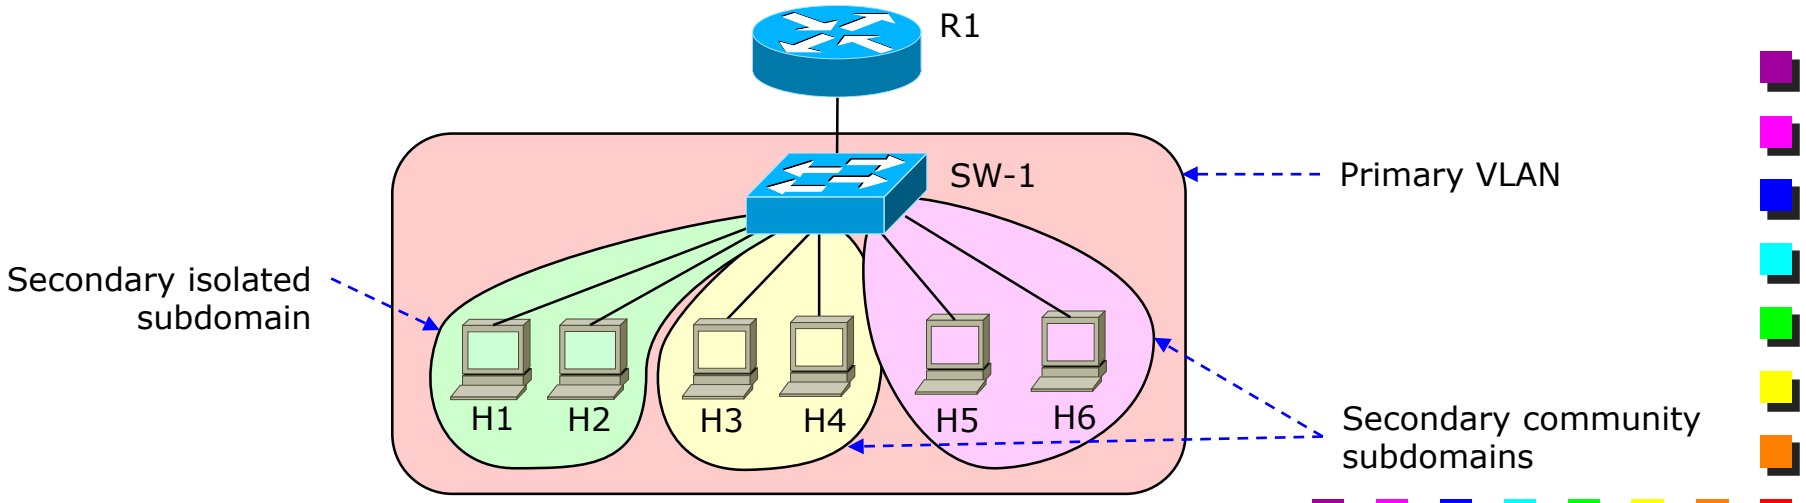

PVLANs port types

- Promiscuous port
  - Port belonging to the primary VLAN
- Isolated port
  - Port belonging to an isolated secondary VLAN
- Community port
  - Port belonging to a community secondary VLAN

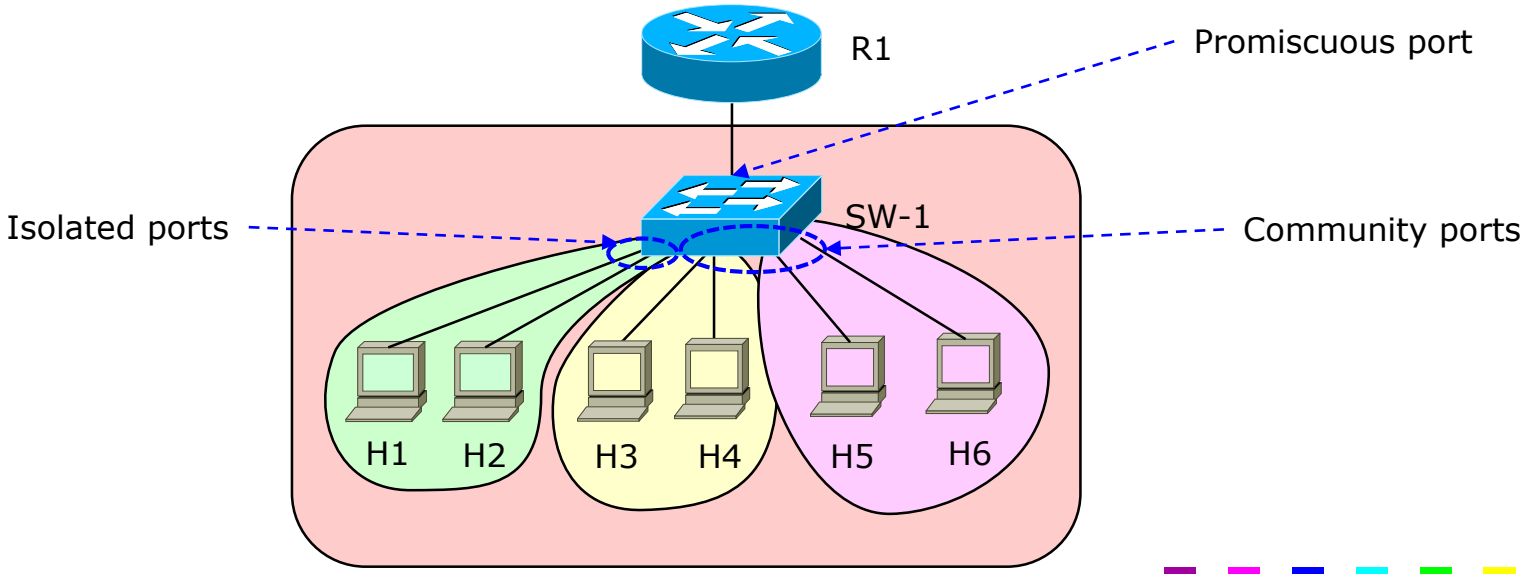

PVLANs: port types and forwarding rules

- Promiscuous port
  - Can communicate with all interfaces, including the community and isolated host ports that belong to the secondary VLANs that are associated with the same primary VLAN.
- Isolated port
  - Has complete L2 isolation from other ports within the same private VLAN domain, except for the promiscuous ports. Private VLANs block all traffic to isolated ports except traffic from promiscuous ports. Traffic received from an isolated port is forwarded only to promiscuous ports.
- Community port
  - Communicates with other ports in the same community VLAN and with promiscuous ports. These interfaces are isolated at Layer 2 from all other interfaces in other communities and from isolated ports within their private VLAN domain.

PVLAN forwarding rules

- Primary VLAN
  - Carries unidirectional traffic downstream from the promiscuous ports to the (isolated and community) host ports and to other promiscuous ports
- Isolated (secondary) VLAN
  - Carries unidirectional traffic upstream from the hosts toward the promiscuous ports
- Community (secondary) VLAN
  - Carries upstream traffic from the community ports to the promiscuous port gateways and to other host ports in the same community

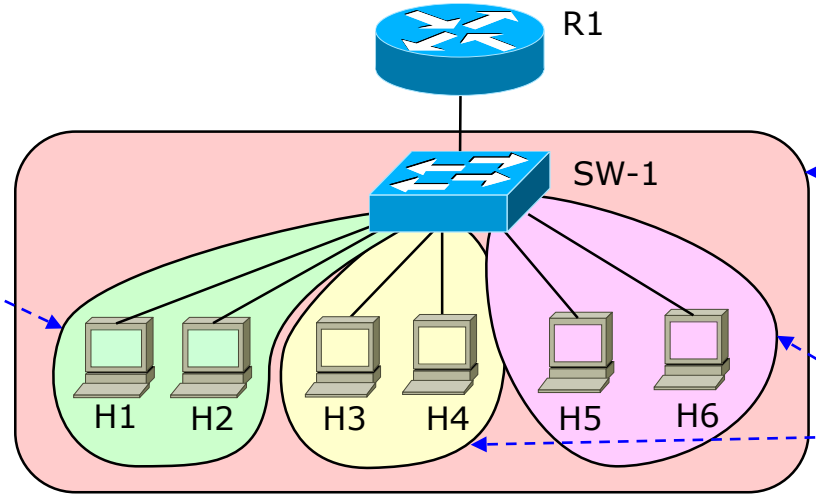

PVLANs and IP addresses

- The address space is assigned to the primary VLAN
  - All (DHCP) requests coming from secondary VLANs can be served by picking addresses from the same space(s)
  - In theory, all the hosts can belong to the same address space
- Reduces the IP addressing fragmentation that may come out when the address space has to be partitioned among the (many) different VLANs

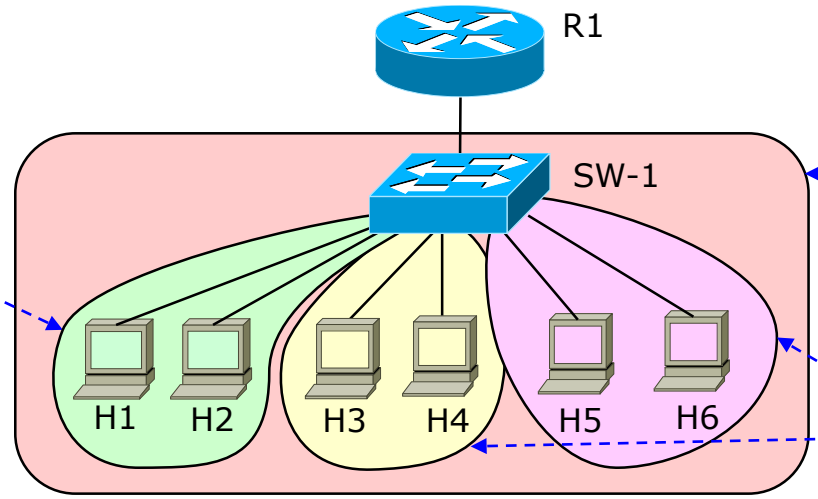

PVLANs across switches

- PVLANs can be extended across multiple devices by trunking the primary, isolated, and community VLANs to other devices that support private VLANs
- However, only the primary VLAN is visible outside the domain
  - VLANS 201-203 are propagated across the infrastructure, but are not visible from R1

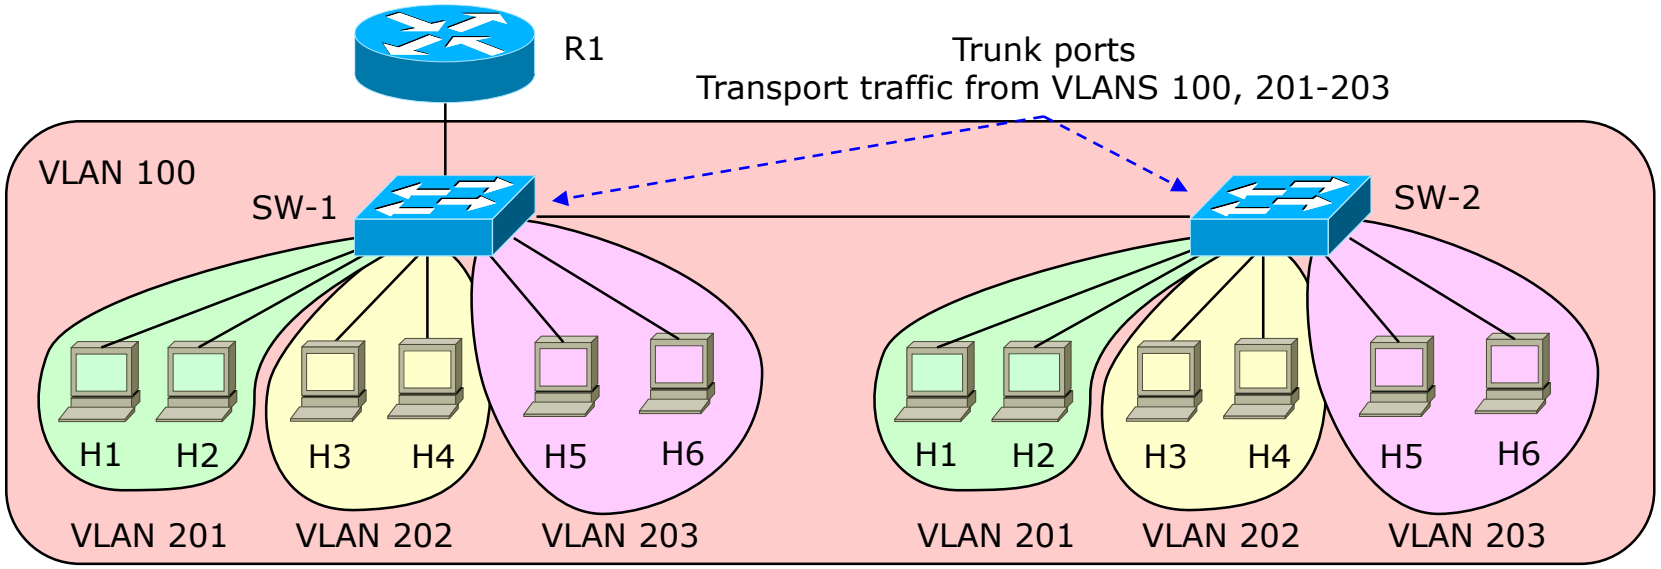

PVLAN: example

```
vlan 1000
  private-vlan primary
  private-vlan association 1001,1002,1003
!
vlan 1001
  private-vlan community
!
vlan 1002
  private-vlan community
!
vlan 1003
  private-vlan isolated
!
interface FastEthernet0/0
  switchport mode private-vlan promiscuous
  switchport private-vlan mapping 1000 1001,1002,1003
!
interface FastEthernet0/1
  switchport mode private-vlan host
  switchport private-vlan host-association 1000 1001
!
interface FastEthernet0/2
  switchport mode private-vlan host
  switchport private-vlan host-association 1000 1001
!
interface FastEthernet0/3
  switchport mode private-vlan host
  switchport private-vlan host-association 1000 1002
!
! ...
interface FastEthernet 0/7
  switchport trunk encapsulation dot1q
  switchport mode trunk
```

Only primary VLANs are propagated on promiscuous ports In this case this is irrelevant, as the port is "VLAN access"

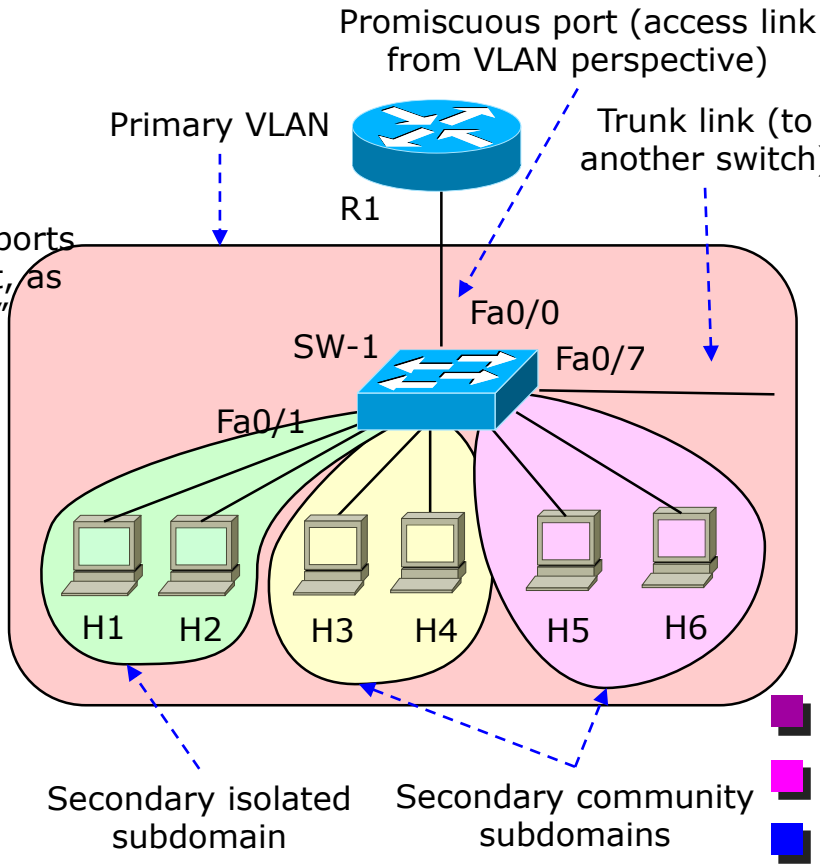

No difference in configuring ports in community or isolated subdomains

All VLANs are propagated on trunk ports (requires VTP)

Conclusions

- VLANs are extensively used in all big deployments nowadays, thanks to their flexibility and the capability to isolate traffic
- "Vanilla" VLANs are being extended in several directions in order to accommodate the above objectives, depending on specific use cases
  - E.g., hierarchy in network virtualization: VLAN stacking
  - E.g., Metro Ethernet: Private Backbone Bridges
  - E.g., full isolation: private VLANs
- Other extensions are possible, depending on specific use cases
- In general, not 100% interoperable