

“Layer 2 Core” Issues List

Norman Finn
Cisco Systems

Assumptions

- **We will assume that all Access Bridges are P802.1ad Provider Bridges.**
- **We will assume that the Core Bridges are something similar, but perhaps not identical.**
- **We will assume the tag structure of new-seaman-service-multiplexing-tag.ppt**

What must we do to define a “Layer 2 Core” for Provider Bridges?

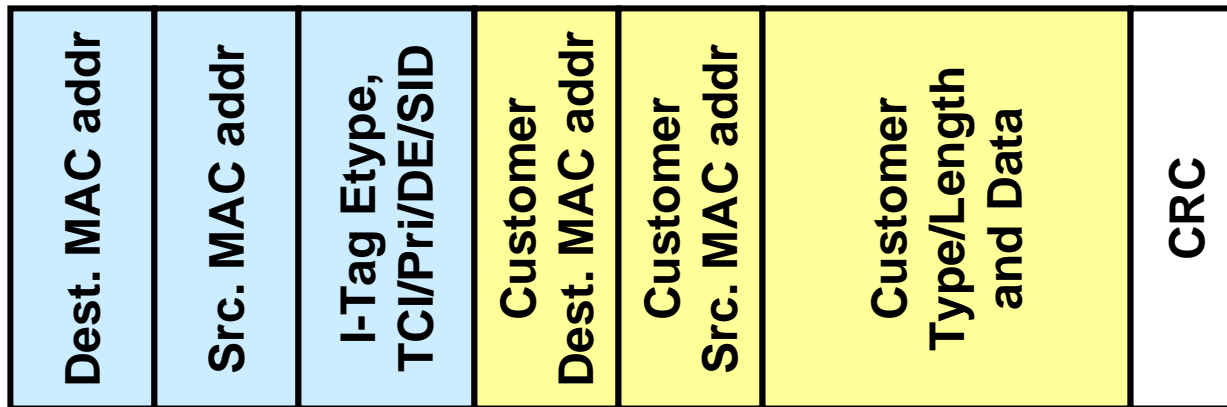
Partial list of decisions to be made

Cisco.com

- **Extending the I-Tag**
- **Which tags to use.**
- **VLAN filtering between core and access clouds**
- **Vectorized GMRP and GVRP in the core**
- **Isolate the core spanning tree from the access spanning trees, including mention of some of the methods available for achieving this**
- **Global looping issues and Protection switching**

Extending the I-Tag


Extending the I-Tag




- **Basic I-Tag Format**
- **There are extra bits in the TCI; what mischief might we get into?**

Extending the I-Tag

Optional
Customer MAC
makes frame
looks like this:

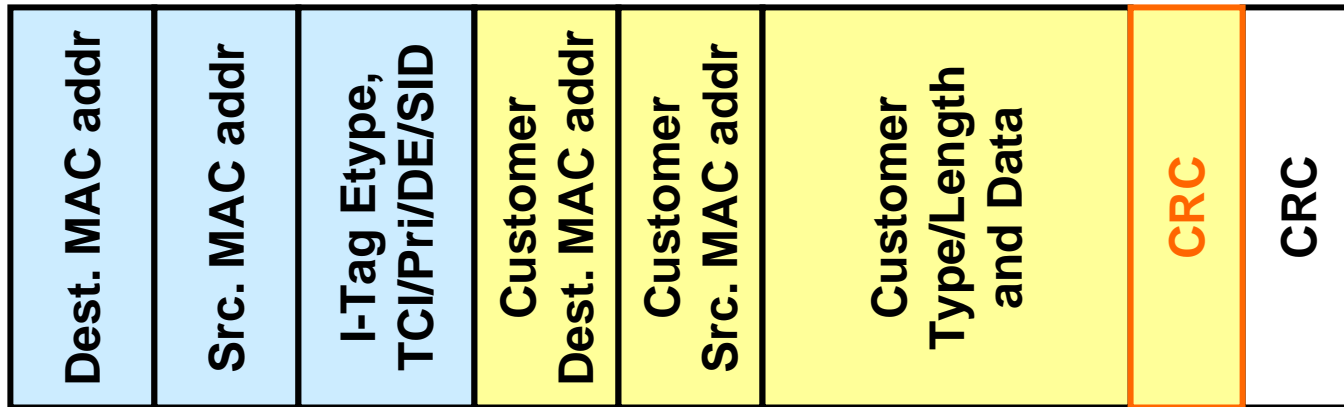


But frame is,
in effect, this:



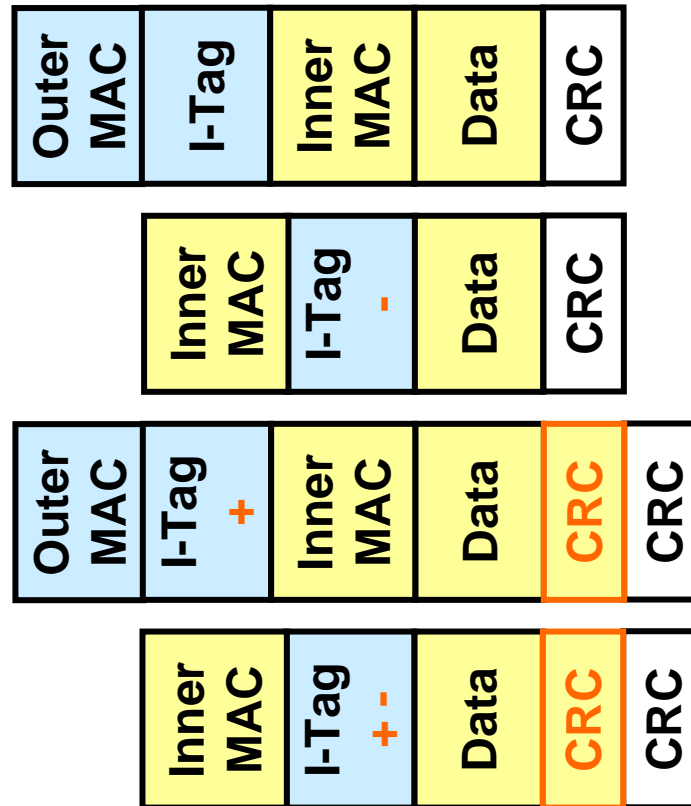
- “Customer MAC Address” could be optional, specified by a bit in the TCI.
(This is just a Q-tag with a big VLAN ID.)

Extending the I-Tag



- “Customer CRC” could be optional, specified by a bit in the TCI.

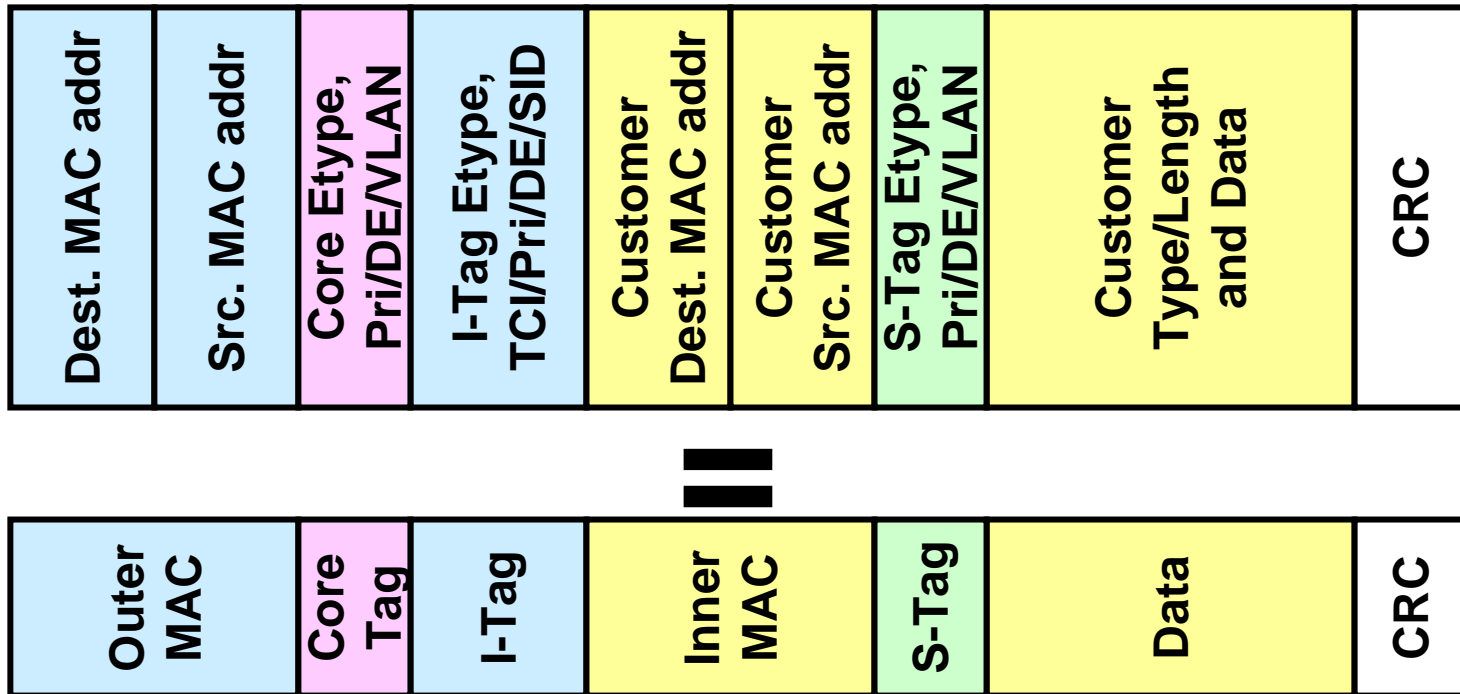
Extending the I-Tag



- I-Tag without MACs can be substituted for other VLAN tags.
- But, I-Tag without MACs cannot be stacked arbitrarily.

Tag Combinations

Tag Combinations

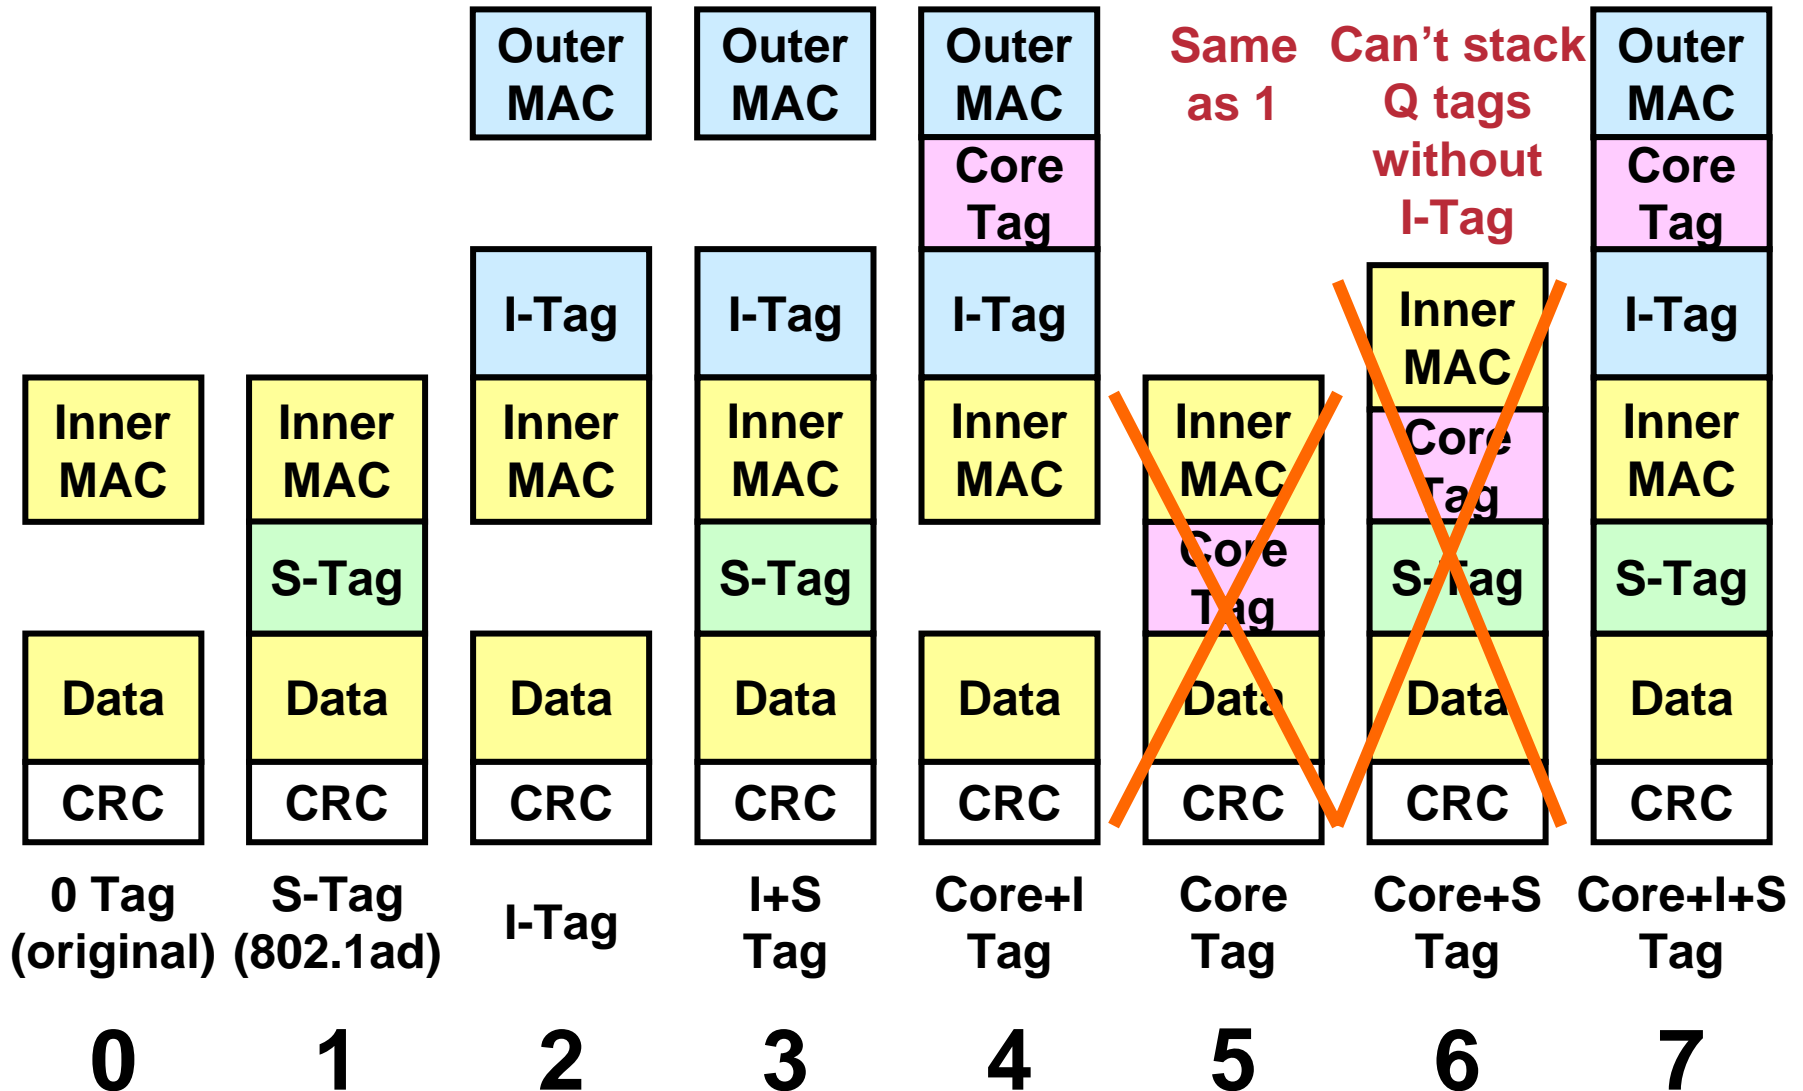


- **Format shorthand.**

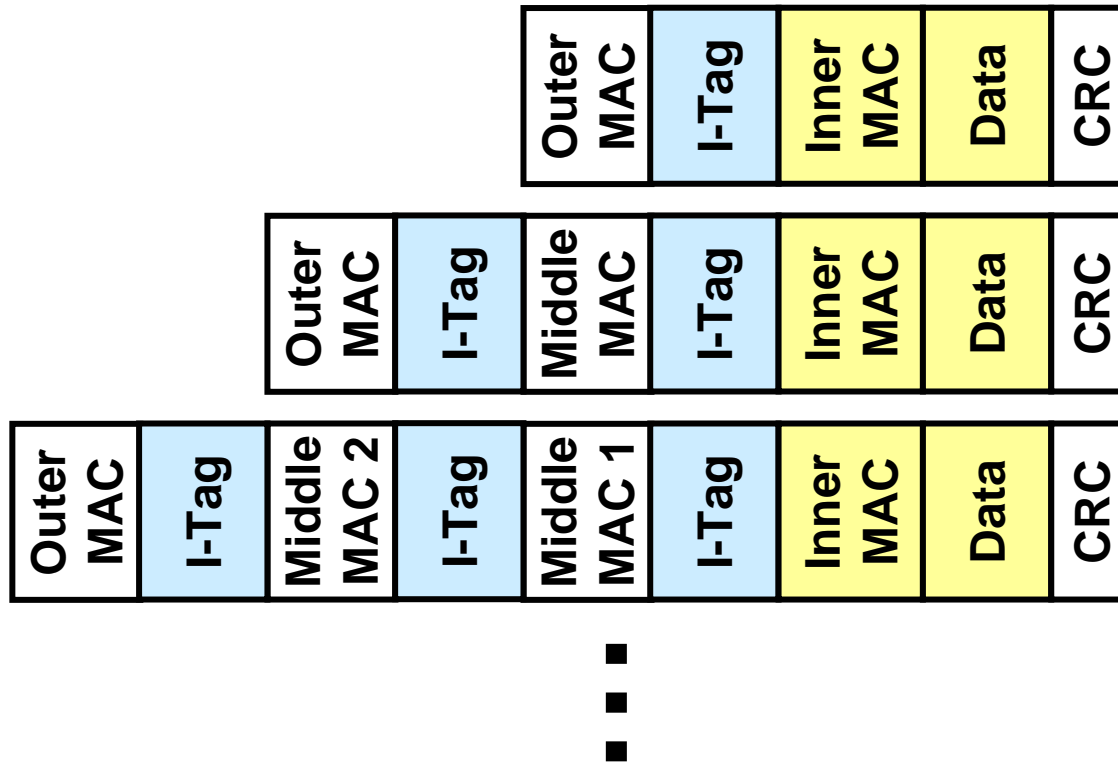
(Let's just forget the Customer VLAN tag for the moment; it is optional, anyway.)

Tag Combinations

Cisco.com



Tag Combinations

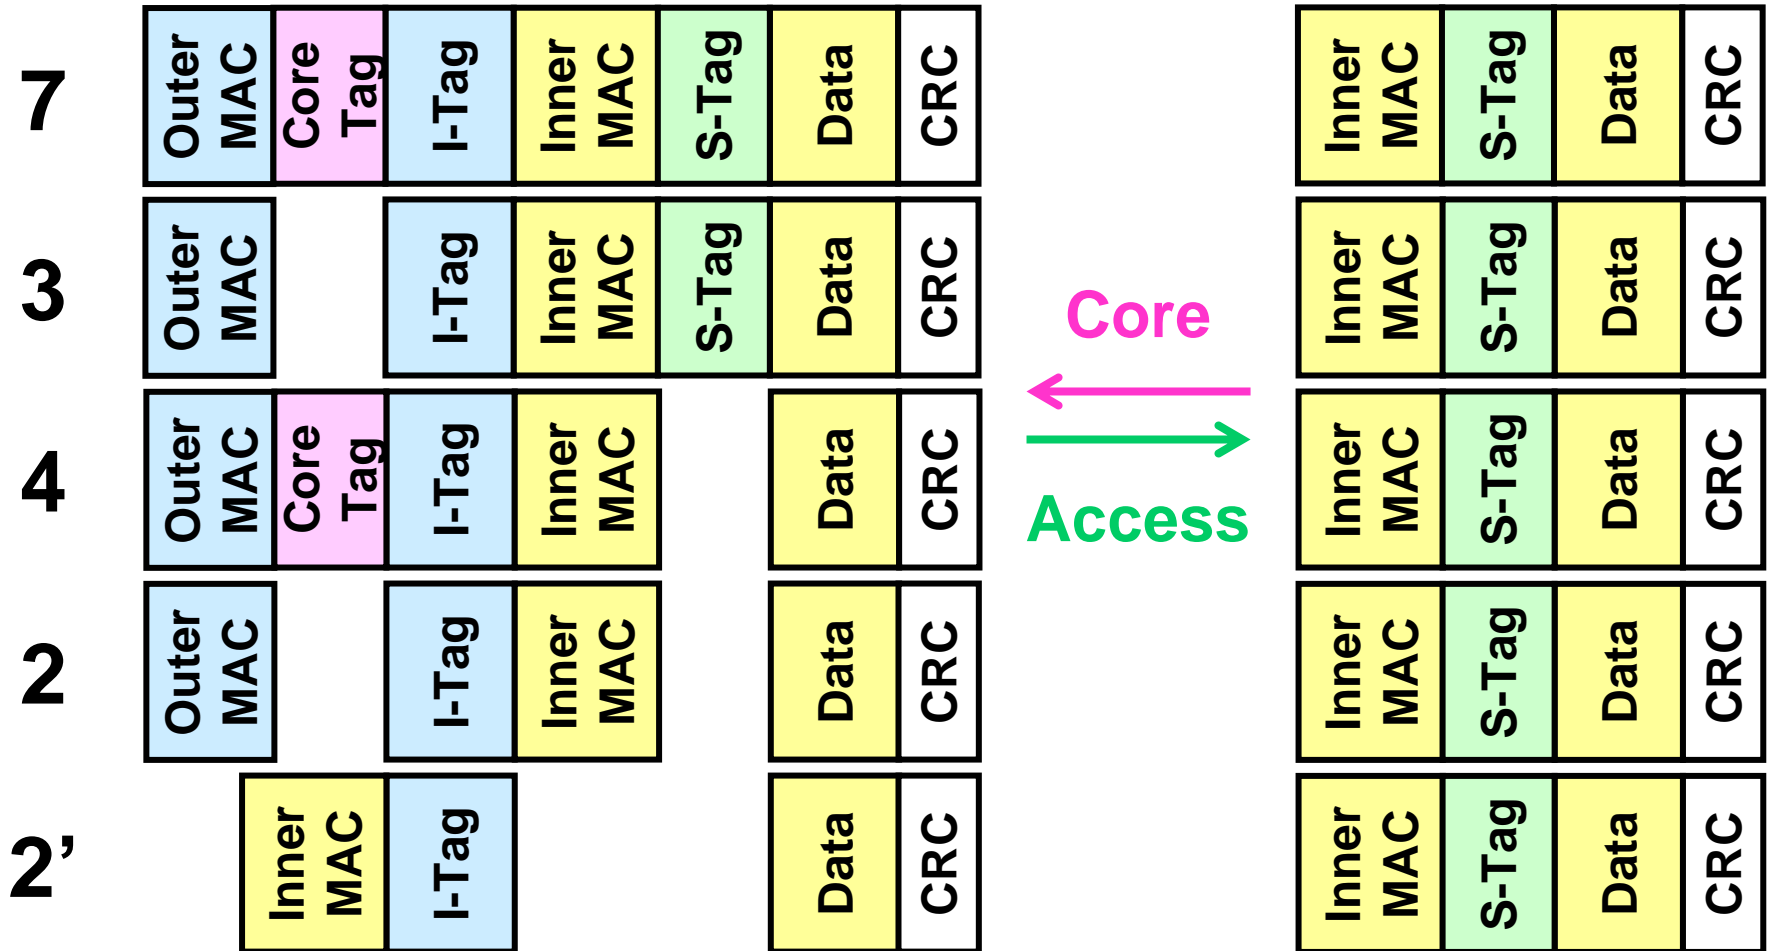


- **I-Tags may be stacked to any level.**
(Outer MACs insulate each level from the last.)

Which tag to use in the Core

What format do we use in the Core?


Cisco.com



• **Wrong question!**


What format do we use in the Core?

Cisco.com



- **Right question #1!**

What format do we use in the Core?



- **Both tag choices are reasonable:**

If outermost Core tag is an S-Tag, P802.1ad bridges can be Core bridges.

If outermost Core tag is an I-Tag, a new “.1Q++” bridge (with more functionality) is required.

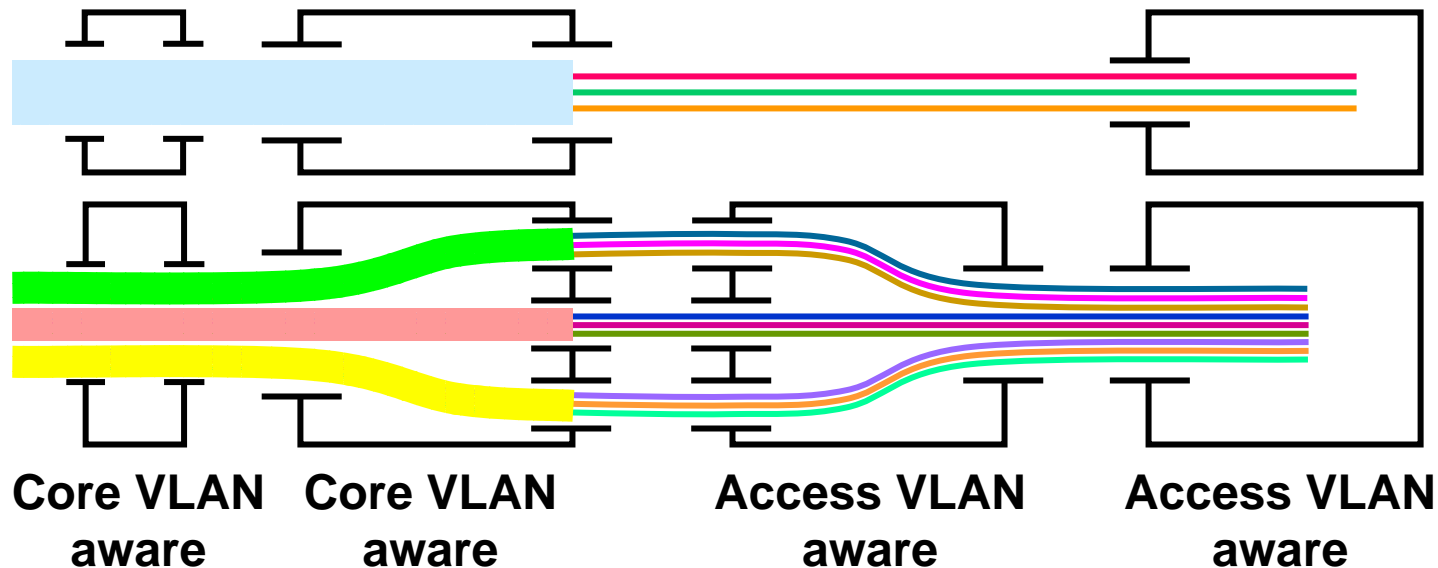
Positioning Core–Access VLAN and I-Tag Insertion and Removal

Core-Access VLAN Filtering

- The “Core–Access interface” allows bundling many Access VLANs (Service Instances) into one Core VLAN.
- Functions to be performed:
 - VLAN assignment and filtering.**
 - I-Tag addition and removal.**

Core-Access VLAN Filtering: Core VLAN assignment

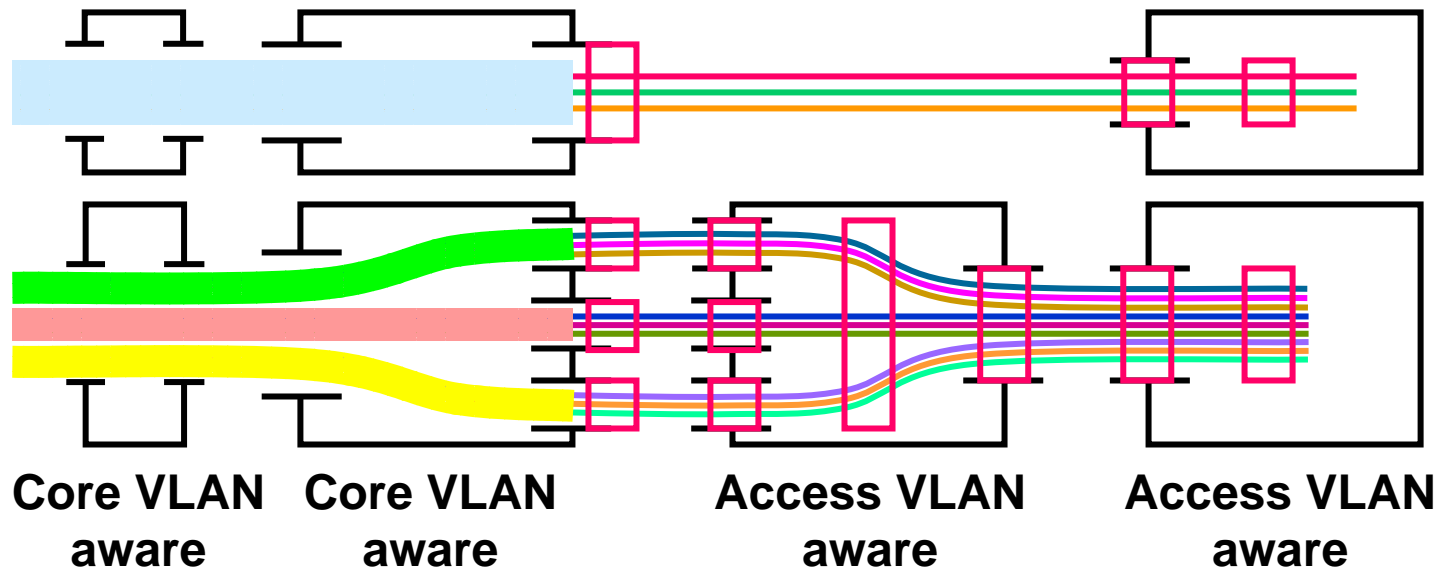
Cisco.com



- Reusing the P802.1ad model is attractive.
- Make the Core–Access model look like the Access–Customer model.

Core-Access VLAN Filtering: I-Tag insertion / removal

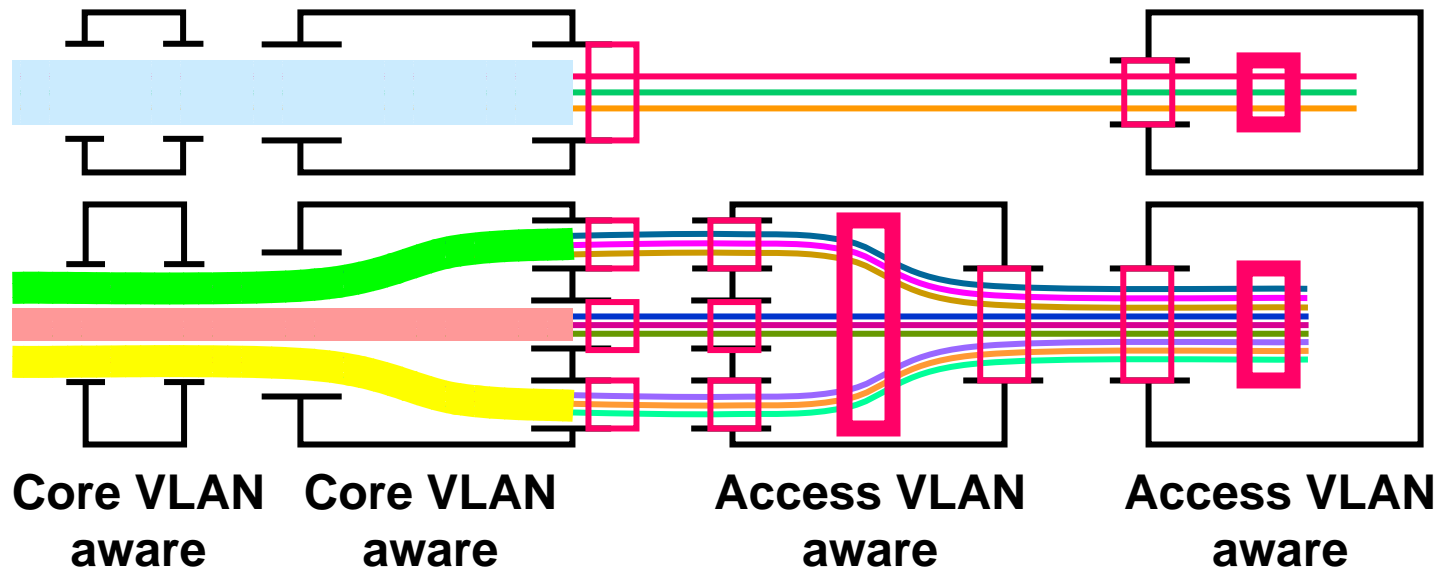
Cisco.com



- I-Tag insertion / removal requires knowledge of Access VLANs; Customer MAC addresses are tied to the context of an Access VLAN.
- Therefore, I-Tag insertion / removal could be modeled as being where the **red boxes** are shown, above.

Core-Access VLAN Filtering: I-Tag insertion / removal

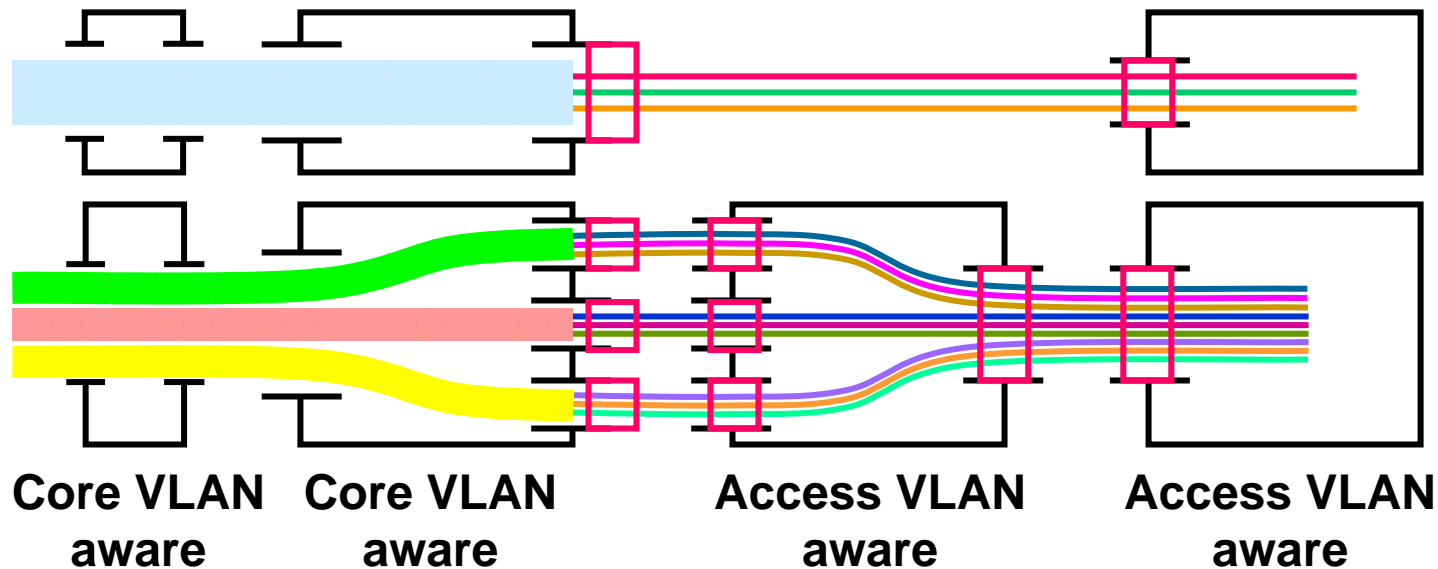
Cisco.com



- The outer MAC addresses in an I-Tag are the MAC addresses of insertion / removal points.
- Bridge Ports have MAC addresses, not Relay Functions.

Core-Access VLAN Filtering: I-Tag insertion / removal

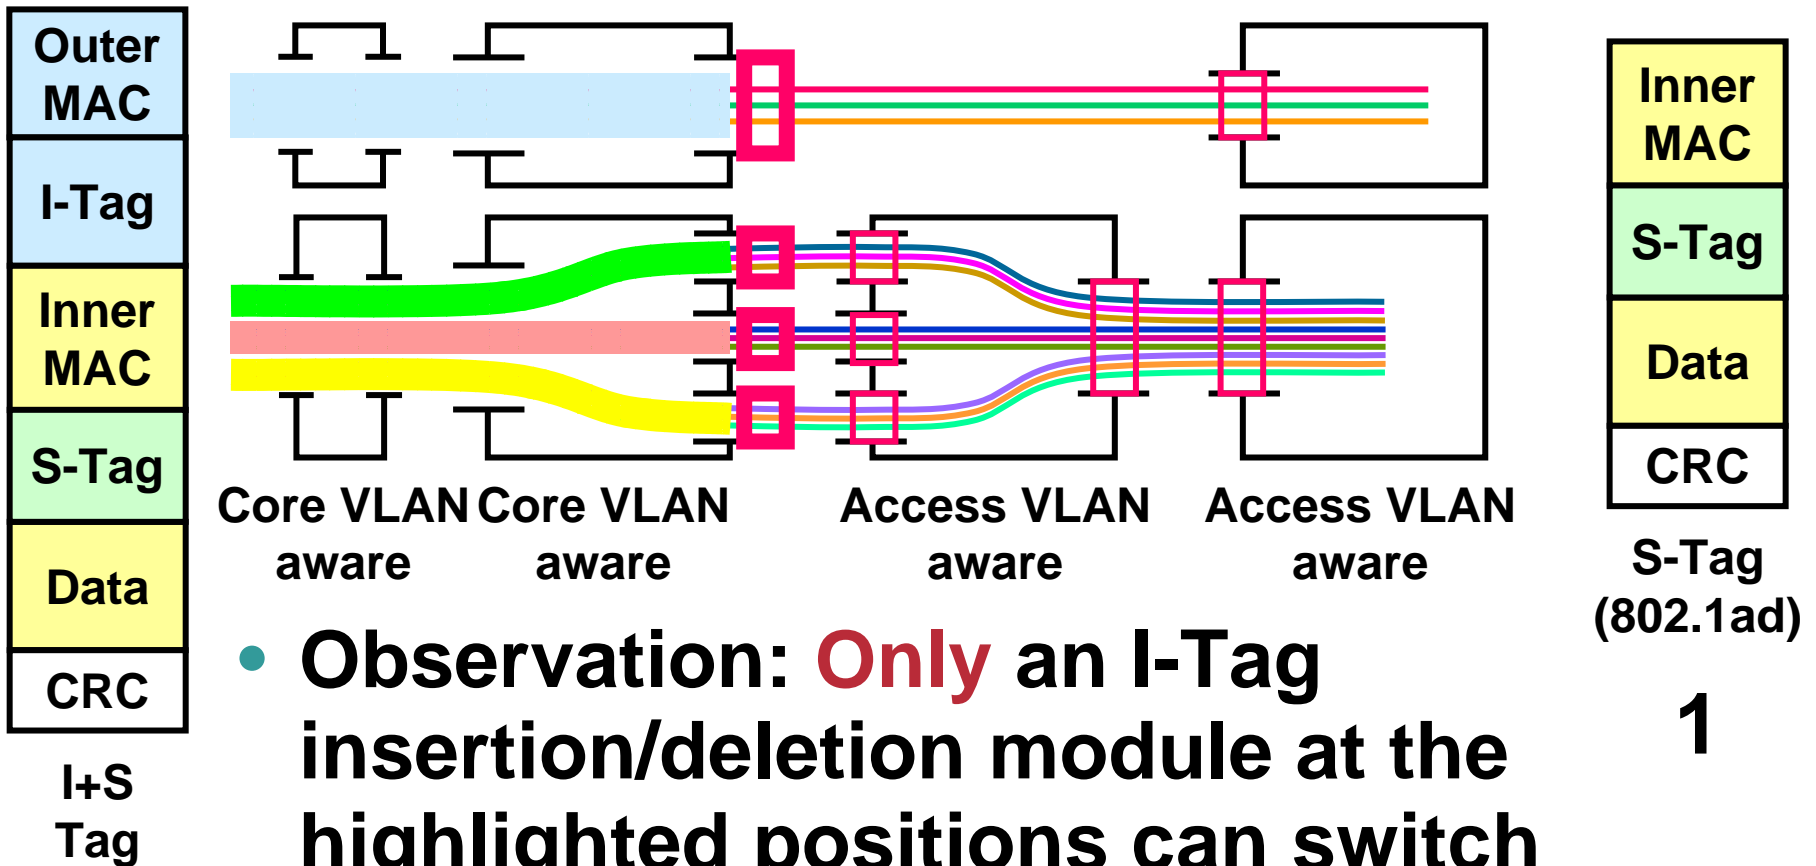
Cisco.com



- So, I-Tag insertion / removal is done only at Bridge Ports.

Core-Access VLAN Filtering: I-Tag insertion / removal

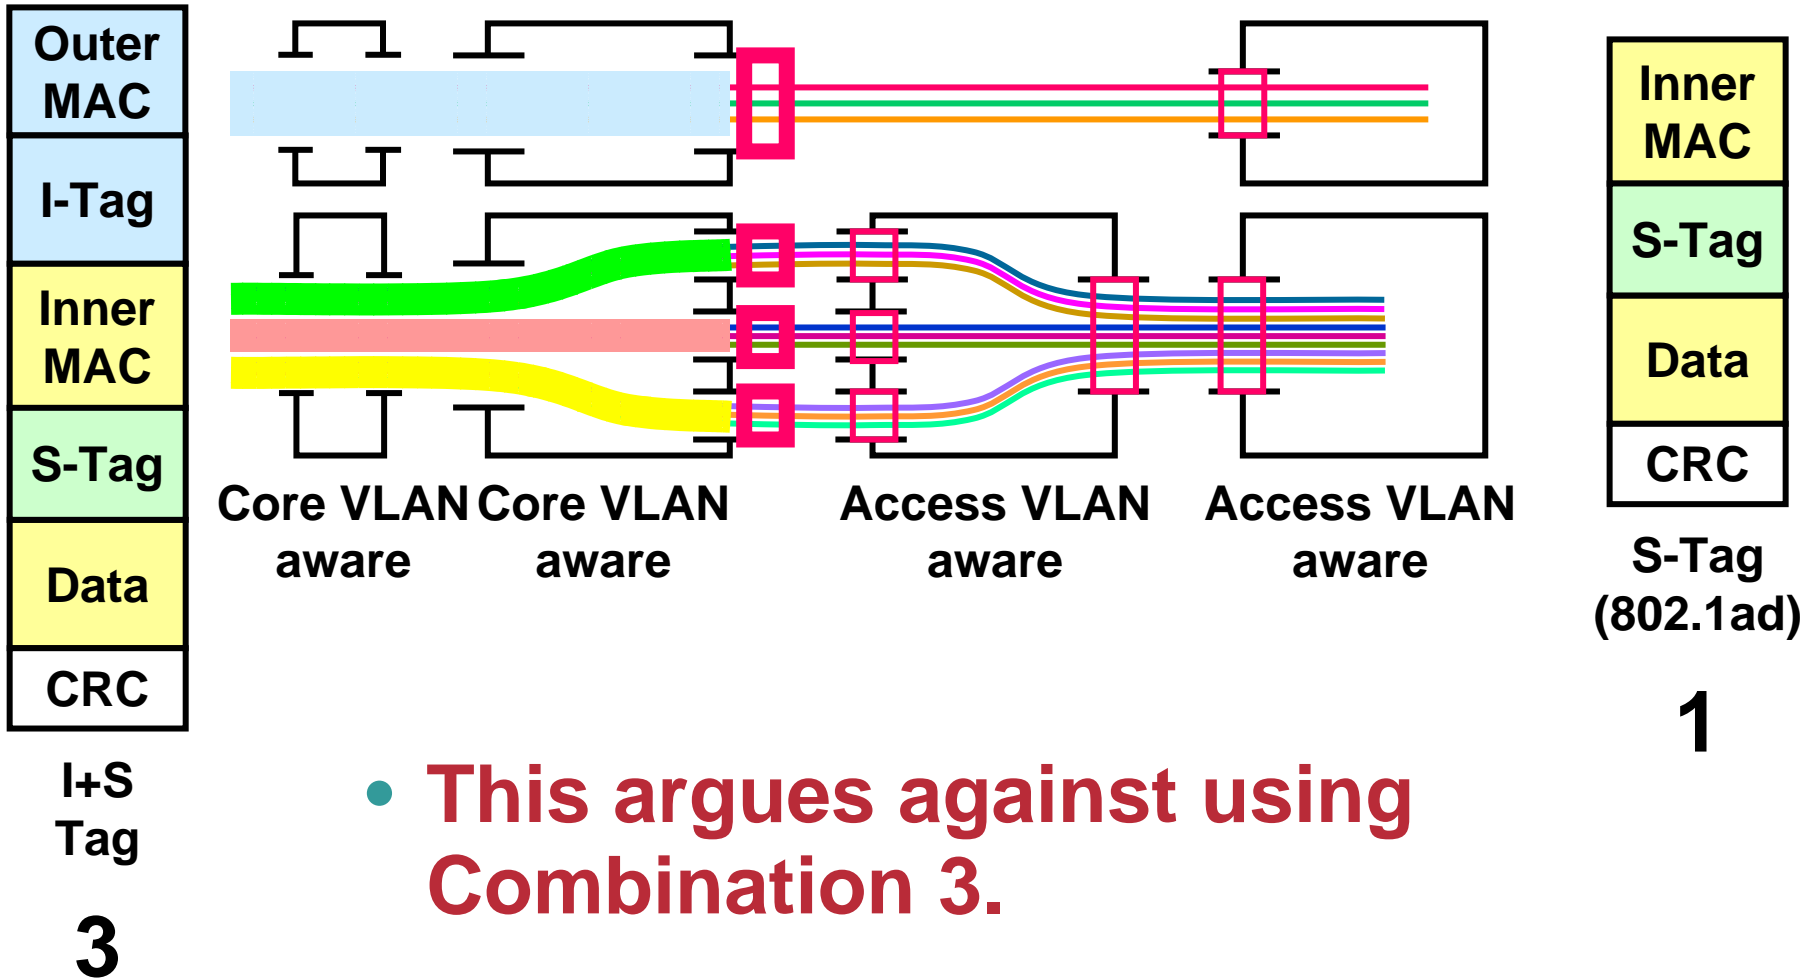
Cisco.com



- **Observation: Only an I-Tag insertion/deletion module at the highlighted positions can switch between Combination 1 (P802.1ad) and Combination 3.**

Core-Access VLAN Filtering: I-Tag insertion / removal

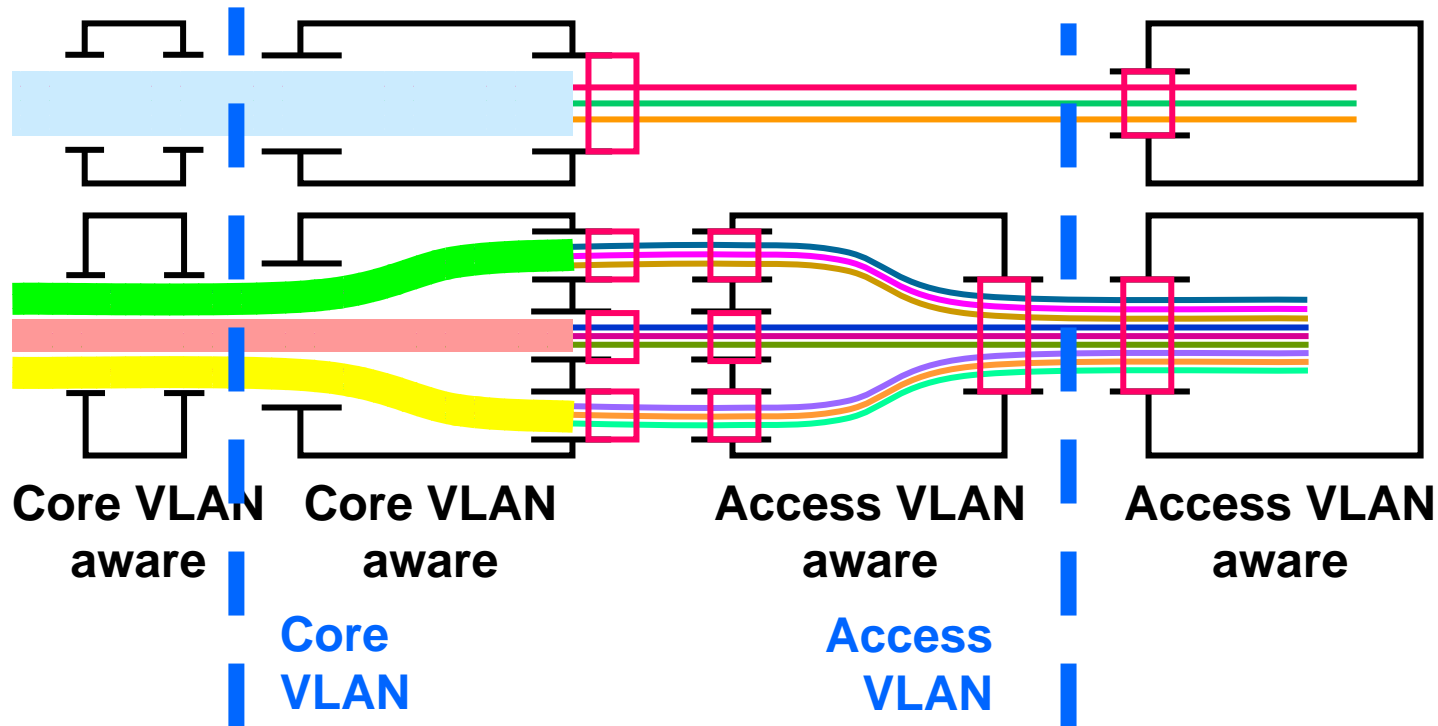
Cisco.com



Format of Frames on the Core–Access LAN

Core-Access VLAN Filtering: I-Tag insertion / removal

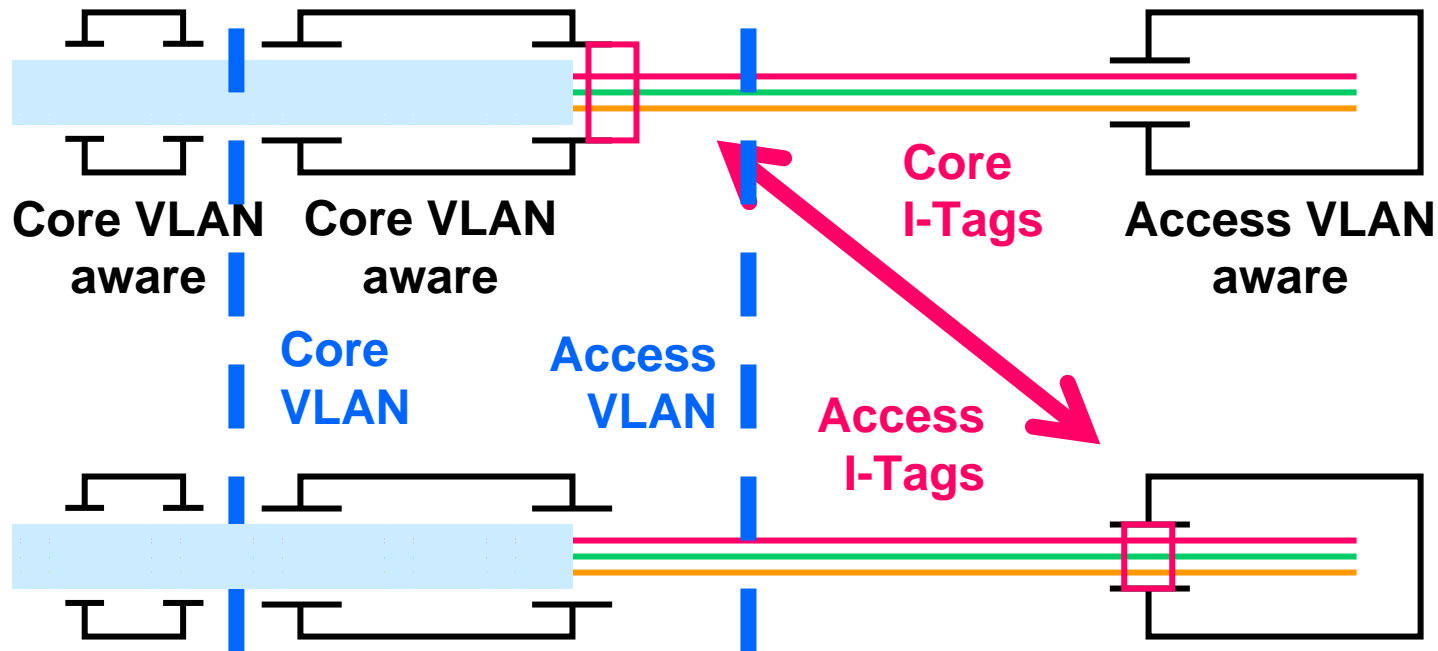
Cisco.com



- Because multiple virtual interfaces between the Access and Core VLAN-aware modules are equivalent to a simple service, there are only **two choices** for placing the **Access–Core interface**.

Core-Access VLAN Filtering: I-Tag insertion / removal


Cisco.com




- Because the virtual interfaces between the Access and Core VLAN-aware modules are, taken individually, the same as those in a simple service, we will omit the complex service.
- Now, there are only **two places** to put the **I-Tagging**.

Three choices for Core–Access frame Format


Cisco.com



- Same format as inside the Core.




- Same format as inside the Access network.



- I-Tagged frames, but I-Tag VLAN ID specifies an Access VLAN, not a Core VLAN.

Use Core Format


Cisco.com



- Same format as inside the Core.
- This is the MPLS/VPLS model. Most likely, the Core-facing Access port performs the complete translation.
 - The outermost Core VLAN is equivalent to an MPLS Label Switched Path.
 - The inner, Access VLAN is similar to the Pseudowire label.
- Multiple Core VLANs may cross the boundary.
- All Core bridges' ports are the same (no "UNIs").
- Access bridges' ports facing the Core are very complex.

Use Access Format


Cisco.com



- Same format as inside the Access network: P802.1ad.
- Access-facing Core ports are like a P802.1ad bridge's ports, plus the I-Tagging function.
- Core-facing Access ports are identical to internal P802.1ad (NNI) ports.
- 4k Customer Instances per Core–Access link is a limitation.

Use Intermediate Format

Cisco.com


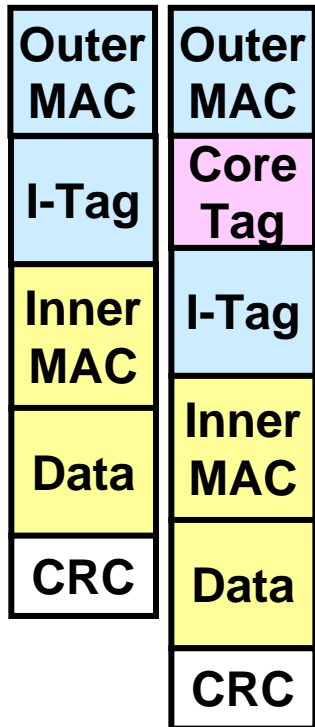


- **I-Tagged frames, but I-Tag VLAN ID specifies an Access VLAN, not a Core VLAN.**
- **1M Service Instances across boundary.**
- **Access-facing Core port is either:**
 - a P802.1ad simple port**
 - or a P802.1ad complex port that can extract the VLAN ID from an I-Tag.**

Three choices for Core–Access frame Format

Cisco.com

Core Format:
“.1Q++” or .1ad



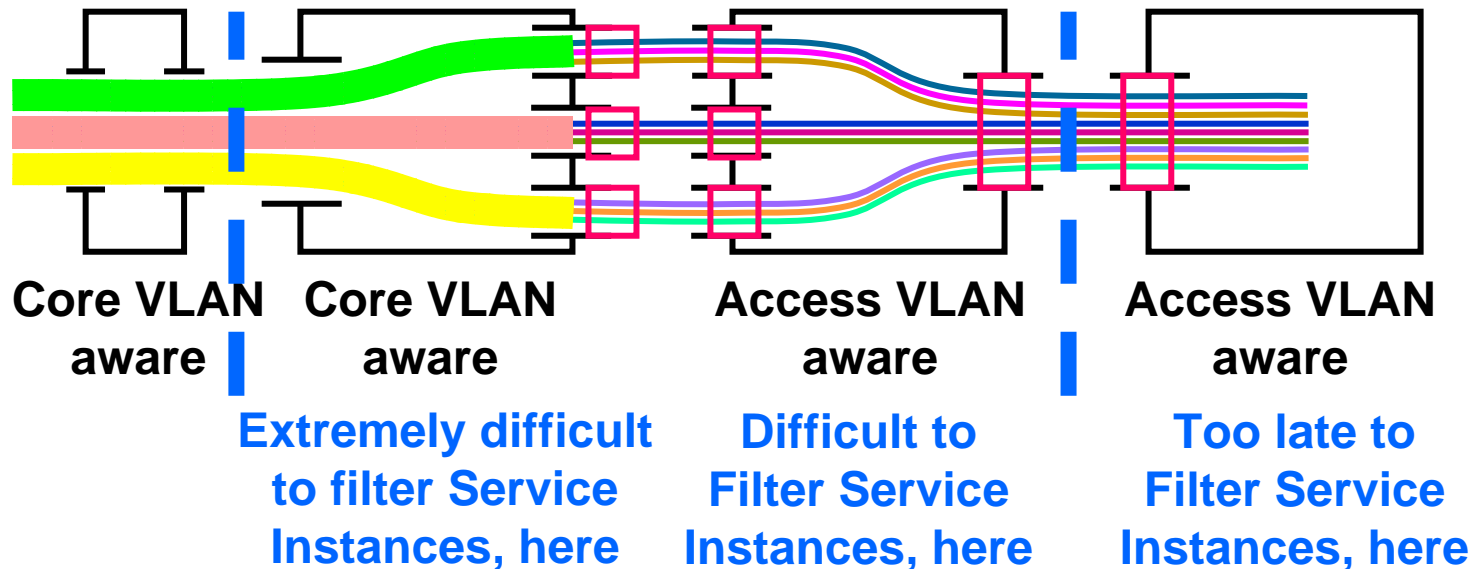
Access
Format:
802.1ad



VLAN Filtering across Core–Access boundary

Core-Access VLAN Filtering: I-Tag insertion / removal

Cisco.com




- If the Access network is interested in only **some** of the Service Instances in a given Core VLAN, it is not easy to avoid transmitting all of the SIs in a Core VLAN across the boundary LAN.

Separating Spanning Tree Domains

Separating Spanning Tree Domains


Cisco.com



- It is impossible to run one spanning tree over the whole world.
- The convergence time of the whole network must not be significantly slower than the convergence time of any one piece.

Separating Spanning Tree Domains


Cisco.com



- When there is a clear hierarchy, as in a Core-Access or Access-Customer interface, it is easy to assume that the Core-ward cloud is perfectly reliable, and let the Customer-ward cloud make all of the decisions about which VLANs are transmitted and which are blocked.

Separating Spanning Tree Domains

Cisco.com




- **In this asymmetrical case:**

The Access cloud decides which VLANs are to be passed and which blocked across the boundary.

Vectorized GVRP allows the Access cloud to communicate its choices to the Core cloud.

Separating Spanning Tree Domains



- If connecting two Access clouds, one cloud's VLAN labeling must be chosen.
- Both must agree on which VLANs cross which link.
- A change in one cloud's spanning tree state must not cause a change in the other cloud's spanning tree state.

Managing the Range of Service Instances in the Core

Limiting the Range of an SI in the Core

Cisco.com

- **Assuming a Core-Access relationship in which the Core is assumed to be reliable by the Access network:**

The Access network can decide which VLANs to pass and which to block across the boundary.

These decisions can be communicated to the Core via a Vectorized GVRP.

Limiting the Range of an SI in the Core

- There are two methods for limiting where in the Core a given Service Instance may go:

Vectorized GVRP to limit the outer (.1ad or .1Q++) VLAN.

Vectorized GMRP to limit the outer (multicast) MAC address.

- Using two tags is not an option: the advantages of layering are lost.

Limiting the Range of an SI in the Core

- **Given its configuration, or the decision of the Access network, the Access-facing Core port can use vectorized GVRP to attract the appropriate Core VLANs.**
- **If the Access network's block/pass decisions affect whole Core VLANs, then:**

SIs may be perfectly distributed across the core

Convergence time is the convergence time of the Access network plus the Core GVRP time.

Limiting the Range of an SI in the Core

- **If the Access network's block/pass decisions do not affect whole Core VLANs, or if the Core bridge does not relay the Access network's decisions this way, then:**

Bandwidth to the edge of the Core may be wasted delivering data that won't be delivered to the Access network.

But, the convergence time is just the convergence time of the Access network, with no addition Core GVRP time.

Limiting the Range of an SI in the Core


- **A Vectorized GMRP is easy to use, because each Core VLAN or each “Core Customer” can be assigned a range of multicast MAC addresses to allocate for this purpose.**
- **Rules must be established for the allocation and use of such MAC addresses in order to assure interoperability.**

Limiting the Range of an SI in the Core

- **Within one Core VLAN, it may be desirable to limit the distribution of a single Service Instance, or even a single multicast MAC address within a single Service Instance.**
- **This can be done by mapping specific Service Instances or multicast MAC addresses to specific outer multicast MAC addresses, and using a Vectorized GMRP, both across the Core–Access boundary and throughout the Core.**

Linking Networks Together


Linking clouds together



- **What prevents connecting clouds into a loop?**
- **Manual configuration is not perfect, but may be acceptable.**

Linking clouds together

Cisco.com



- **Can you use Alarm Indication Signals and Protection Switching to switch between two alternate paths through the clouds?**

CISCO SYSTEMS

