

Bridging Solution for the MAN: Service Separation

November 2002

Marc Holness, Nortel Networks

Problem Context

- An Ethernet Service Provider supplies Ethernet Layer 2 services among customer sites
- IETF's Provider Provisioned Virtual Private Networks (PPVPN) WG is defining L2-VPNs
 - PPVPN drafts do not assume that bridges, as defined by 802.1, are essential to providing the service
 - They assume that devices which learn MAC addresses and forward packets based on that learned information are essential, but these are not bridges
- How can bridges supply these services?

Objective

- Proposal of an Ethernet bridging solution for the MAN
 - Carrier service separation
- Solution alignment with 802 and 802.1 Architecture
- Solution interworking with MPLS/IP/EoS network



Carrier Grade Ethernet: Part II – Service Separation

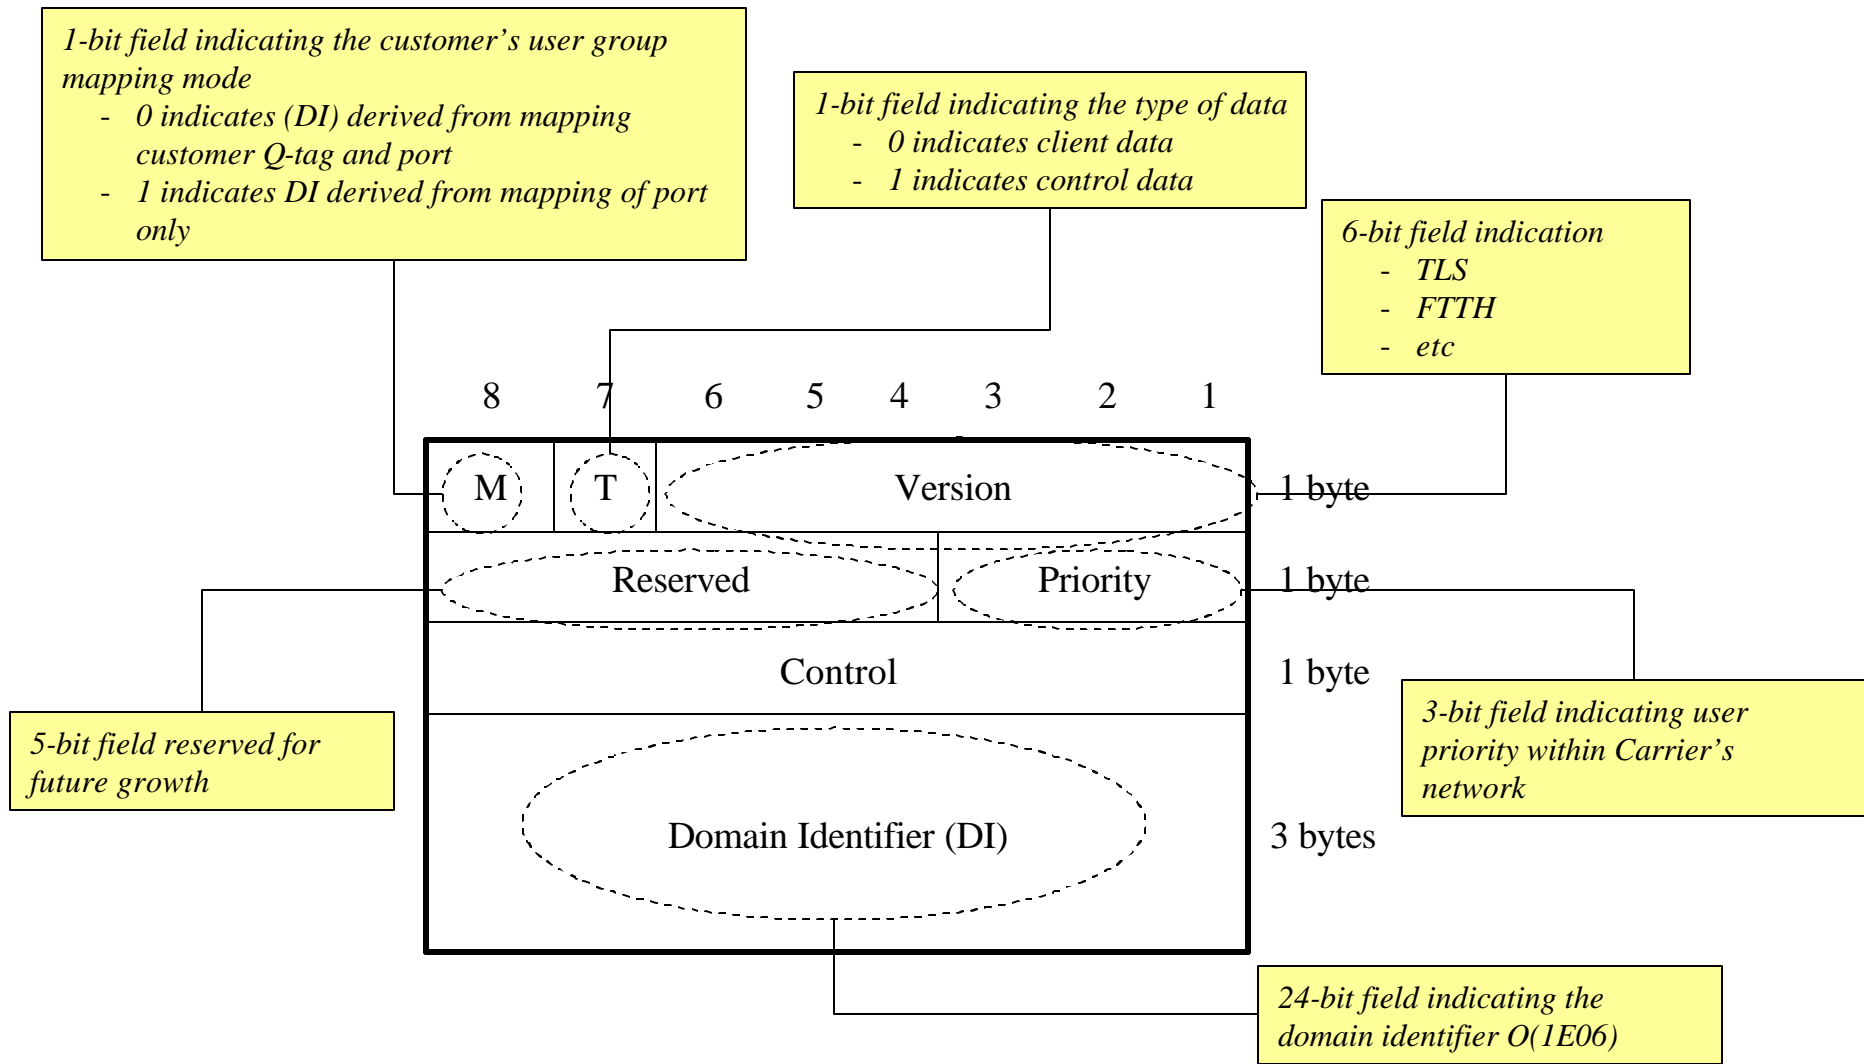
Carrier Service Identifier

- A service instance represents a set of customer sites participating in a service
 - Forms a closed user group within the carrier's network
- A carrier's network can support multiple service instances
 - $O(1E03)$ to $O(1E06)$
- Service identification is achieved by tagging the customer packet with a Carrier grade tag

Carrier Service Identification Tag

- Q-tag could provide 12 bits (4094) of domain identification. However
 - VLAN encapsulation is not currently a standard
 - 12 bits of domain identification does not support range partitioning in a way which allows a service provider to properly assign different layer 2 segments to more than 4K clients (especially in the context of a typical FTTH EFM deployment or TLS deployment)
- Define a service indicator tag with a space greater than 4K that can support realizable FTTH, TLS, etc deployments

Example Carrier Service ID Tag

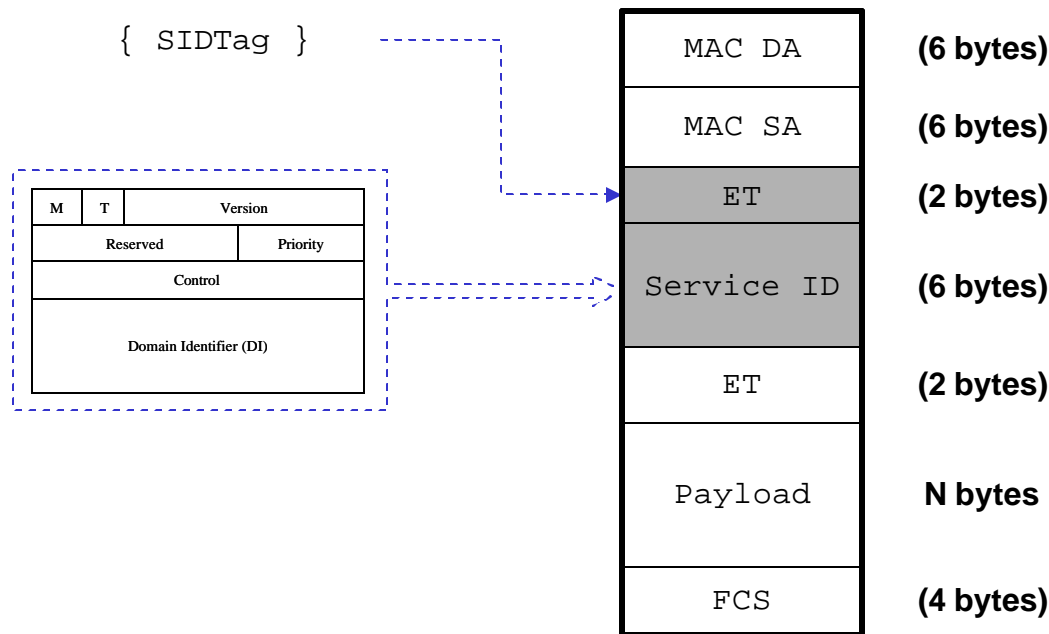


Carrier Service Tag Value Proposition

- Allows service providers to realize scalable Service identification within an IEEE 802 network
 - Does not require higher layer technologies to scale IEEE 802 solution
 - Removes dependency on higher layer protocols for VPN resolution
- Independence from control plane
 - No requirement for end-to-end control plane
 - Supports multiple networks running disparate control planes
- Enables cheap access devices that need not run complex protocols

802 Frame Format

- Use the EtherType (ET) field to indicate that a frame contains a Carrier service identification tag
 - Additional ET value is being proposed
 - New value that is being considered is SIDTag





Carrier Grade Ethernet: Service Provider Network Applicability

Approach

- Build on top of Norm's "Bridge Based Ethernet Provisioned" proposal (rev 2.0)
- Use same reference model; concepts of Islands, Inter-Island Trunks, and Service Instances

Enlarging Provider Networks

- Same direction as Norm

- ***Many techniques have been implemented, and even more discussed, to expand the size of a bridged network.***
 - Separate spanning trees at edges which run on top of central spanning tree.
 - Running two disconnected spanning trees in one bridge.
 - Substituting hop count for Max Age in Rapid Spanning Tree. (Standardized by 802.1y: 64k hops allowed across network!)
 - Enforce topology restrictions by some non-spanning tree means.
- **We will look at just one: Topology restrictions.**
 - **This also promises to interact well with the IETF solution(s).**

Service Instance, Islands and Inter-Island Trunks

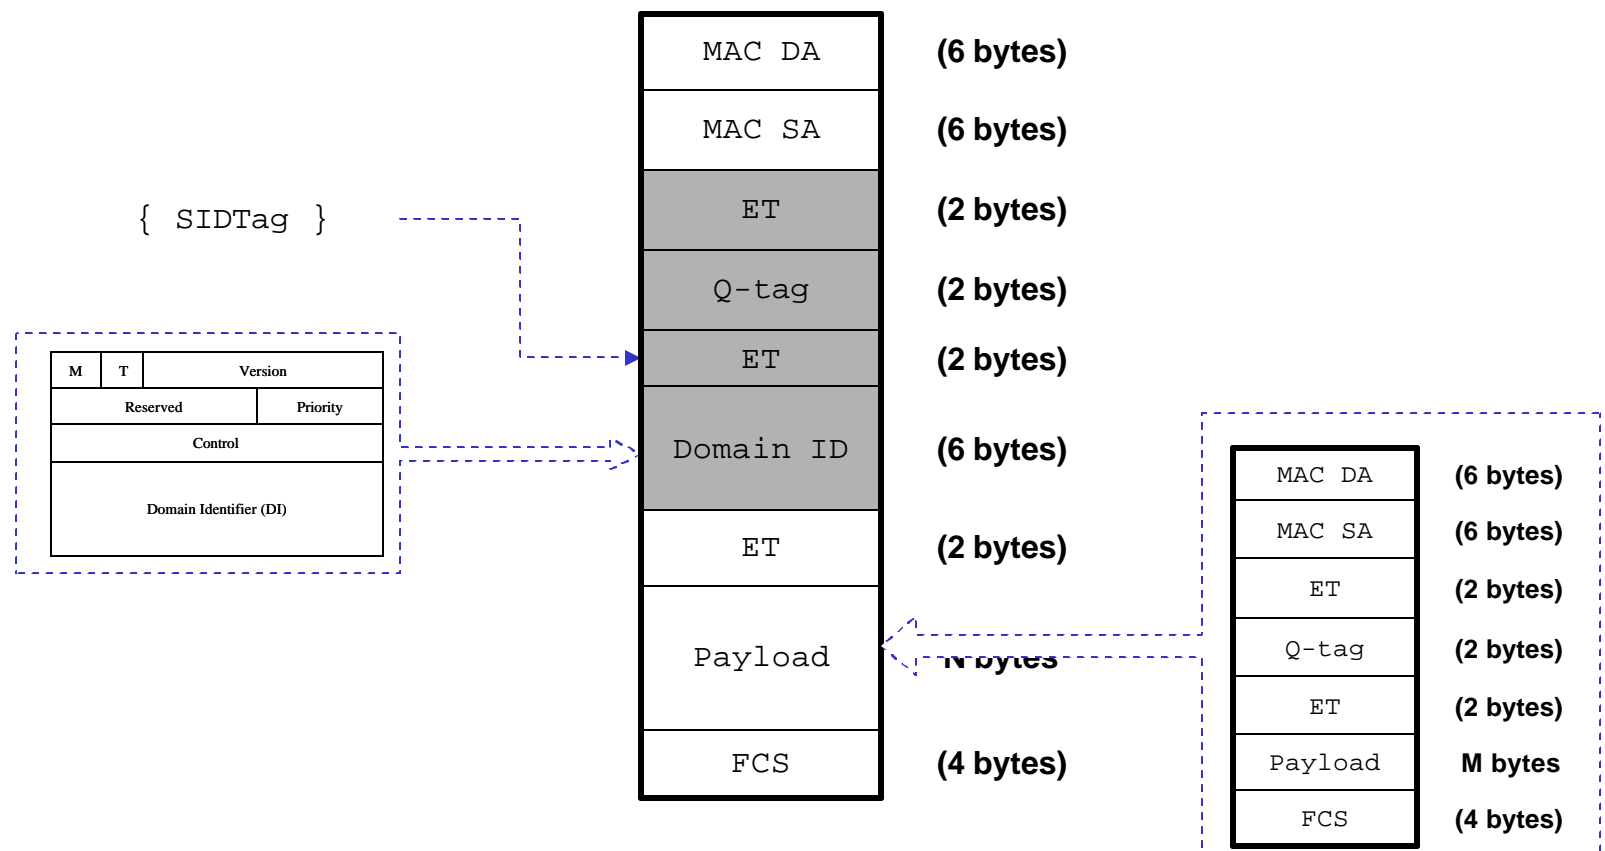
- Same terminology used by Norm; subtle difference in definition however
 - A “Service Instance” is the carrier analog of a VLAN in an 802.1Q enterprise network
 - It is uniquely identified by the domain identifier in the domain tag
 - It can uniquely identify the Service provided by the Service Provider’s network
 - The provider network carries “Service Instances” via “Islands” connected by “Inter-Island Trunks”
 - Islands
 - An Island consists of one or more bridges connected by normal LAN segments and/or Inter-island Trunks
 - Different Islands must be connect only via Inter-Island Trunks

802 Specification Impacts – Step (1)

- Use service provider Q-tag along with Service ID tag within each island
 - Provider Q-tag dictates the connectivity associated with customer sites within an Island
 - Provider Q-tag is localized to an Island
 - Carrier service ID tag used to denote the service instance
 - Carrier service ID tag can be localized to an Island or can be used ubiquitously across Service Provider network (includes Islands and Inter-Island Trunking network)

802 Specification Impacts – Step (1)

- Outer Q-tag used by interior bridging devices to forward traffic and support Spanning Tree specifications within Islands



Additional Value Proposition

- Full transparency is possible since carrier network elements do not have to deal with any flow identification process involving customer supplied information of any type
- End-to-end service identification is possible since the service tag can span the “Islands” and island inter-connecting network (e.g., IP, MPLS, etc.);
 - Minimizing service provider end-to-end service management costs
 - Facilitates end-to-end service management/control functions

Inter-Island Trunks

- Same as Norm with 1 notable exception:

- **Inter-Island Trunks**

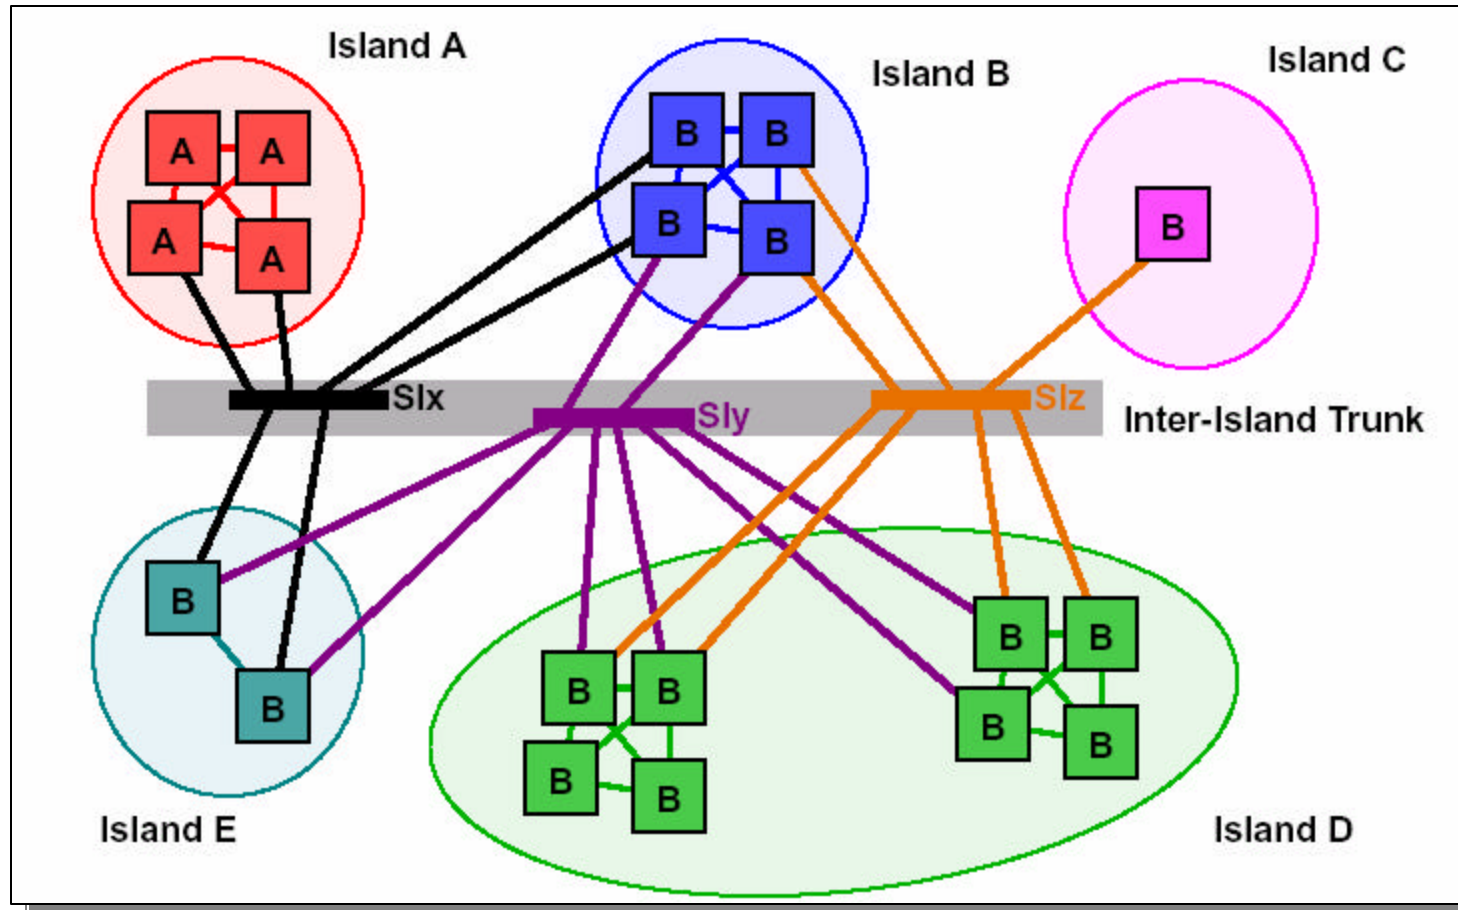
- An Inter-Island Trunk behaves, logically, like a Shared Medium LAN segment.
- Among different Islands, a given Service Instance is carried over, at most, one Inter-Island Trunk.
- The identification of a SI on an Inter-Island Trunk is dependent on the medium, e.g. Ethernet, or emulated Ethernet over layer 3 tunnels.
- Any bridge port connected to an Inter-Island Trunk must make 1:1 translations between the P-VLAN IDs used within the bridge's Island and the SI identifiers used on the IIT.

- **These rules prevent SI loops among Islands.**

- **Spanning trees prevent SI loops within an Island.**

- Any bridge port connected to an Inter-Island Trunk must make 1:1 translations between the service identifier used on the Inter-Island Trunk

SIs, Islands, and IITs



Interconnecting Islands

- Same direction as Norm

- **A separate MSTP Service Instance per Island carries MSTP BPDUs.**

- **Bridges in the same Island interchange BPDUs and guarantee that, for any given frame on a P-VLAN, only one copy will be transmitted to the Inter-Island Trunk.**

- **Similarly, by interchanging BPDUs, the bridges can guarantee that a frame on an Inter-Island Trunk will be delivered, at most, once to any LAN segment within the Island.**

- **Since we are guaranteed (See Slide 51) that no SI exists on more than one Inter-Island Trunk, and that there are no back doors, failing to receive the BPDUs from bridges in other Islands cannot cause a loop, but does manage to limit the size of any one Spanning Tree Instance.**

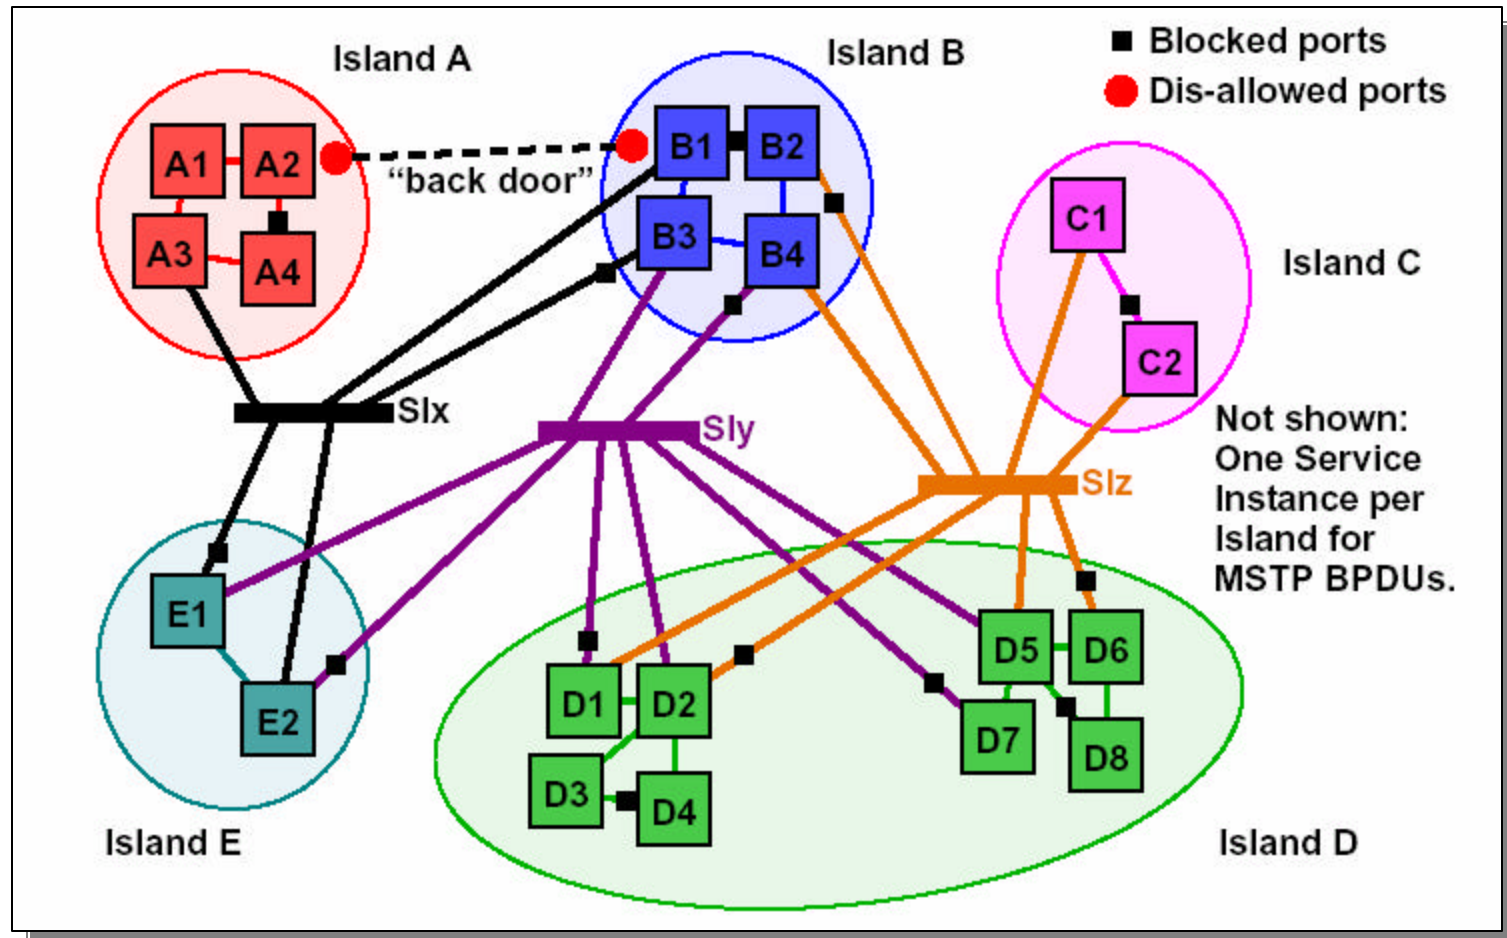
- The outer tag (e.g., P-VLAN) is used to identify the connectivity between the PEBs and IIT per Island

Interconnection Islands and Topology Change Notification

- Same direction as Norm

- **A new Topology Change Notice is required.**
 - Islands may have completely different MSTP configurations.
 - Forgetting all addresses is too much, but we must forget some.
 - The other “Island” may, in fact, be an IETF implementation.
 - This new Topology Change Notice must signal topology changes based on Service Instance (Inter-Island Trunk).

Interconnection Islands Example



Observations

- Service identifier can uniquely identify the Service across the Service Provider's network
- The Service Provider's network can support $O(1E06)$ services
- Islands can be composed of hierarchy bridges at the edge and transparent bridges* at the interior. Consequently, the Service Provider's interior devices per Island scale with the Service Providers network elements. In general, Islands have greater scaling properties
- Interworking with inter-island networks (e.g., IP, MPLS, Ethernet, etc.) can be supported using the service identifier tag
- Specification impact can be limited to the introduction of a service identifier tag