

Provider Bridge Layer 2 Protocols

Norman Finn, Cisco Systems

List of Relevant Layer 2 Protocols

Destination MAC Address	Name	Description
01-80-C2-00-00-2X	GARP ^a	IEEE 802.1D, Generic Attribute Registration Protocol. The “carrier” protocol over which GMRP and GVRP are implemented.
	GMRP ^a	IEEE 802.1D, GARP Multicast Registration Protocol. Prunes delivery of multicast MAC addresses back from ports that don't need to see them. L2 equivalent of IGMP.
	GVRP	IEEE 802.1Q, GARP VLAN Registration Protocol. Prunes each VLAN's broadcasts, multicasts, and unicast floods back from ports they don't need to go to.
01-80-C2-00-00-10	All Bridges ^a	IEEE 802.1D. Defined as an ordinary multicast address to be used to reach all bridges in a bridged LAN.
01-80-C2-00-00-04 - 01-80-C2-00-00-0F	Undefined 802.1 bridge addrs.	Reserved for use by 802.1. IEEE 802.1D states that a bridge will never forward a frame with one of these addresses.

Destination MAC Address	Name	Description
01-80-C2-00-00-00	STP	IEEE 802.1D, Standard Spanning Tree Protocol. Protocol packets called, “Bridge Protocol Data Units”, or BPDUs.
	RSTP	IEEE 802.1W, Rapid Spanning Tree protocol (RSTP). Same function as STP, but converges (typically) in tens of milliseconds, rather than tens of seconds.
	MSTP	IEEE 802.1S, Multiple Spanning Tree Protocol. Carries multiple STP instances on top of a single RSTP BPDUs.
01-80-C2-00-00-01	Pause	IEEE 802.3 Clause 31, Point-to-point Pause function. Used to implement L2 flow control on a whole physical link. Handled by hardware.

Destination MAC Address	Name	Description
01-80-C2-00-00-02	LACP	IEEE 802.3 Clause 43, Link Aggregation Control Protocol. Protocol to automatically establish groups of point-to-point links between two devices for load sharing.
	OAM	IEEE 802.3ah EFM Draft 1.3, Operations, Administration, and Maintenance.
	LLDP	IEEE 802.1ab Draft, Link Layer Discovery Protocol. Allows stations to exchange chassis and port information.
	Slow Proto- cols	Future IEEE 802 standard protocols which expect no more than about 1 packet per second are expected use this MAC address.
01-80-C2-00-00-03	802.1X	IEEE 802.1X, Port-Based Network Access Control. Port-level secure authentication, usually using a RADIUS server.
01-00-5E-XX-XX-XX	IGMP ^a	IETF RFCs 1112 and 2236, Internet Group Management Protocol. Layer 2.5 Multicast subscription protocol which runs between hosts and routers. Snooped by switches to control distribution of L2 multicast MAC addresses.

Destination MAC Address	Name	Description
00-00-5E-00-00-XX (Unicast address)	VRRP ^a	IETF RFC 2338, Virtual Router Redundancy Protocol. This unicast MAC address may move around. It may be used by two different MACs in two different locations on a bridged network, on different VLANs.

- a. This protocol's packets may be tagged with a VLAN ID.

What Does a Provider Bridge Do with Each Layer 2 Protocol?

- In the following table's "Tunnel" column:
 - "T" means that the protocol should always be Tunneled, and perhaps, inspected by the Provider;
 - "P" that the Provider Bridge and Customer Equipment should participate as Peers in the protocol; and
 - "D" that its frames should be Discarded and never generated.

Name	Tunnel	Possibilities
GARP	T/D	At this time, no specific proposals have been received by P802.1 for using these protocols to exchange information between Customer and Provider. Such proposals will most likely arise. In the meantime, we will assume that these protocols must be either tunneled through the Provider service or discarded by the Provider.
GMRP		
GVRP		
All Bridges		

Name	Tunnel	Possibilities
Undefined 802.1 bridge addr.	T/D	Given that the existing protocols in this class need to be treated differently by the UNI, it is difficult to predict what consequences new protocols may have.
MSTP		Clearly, Customers' BPDUs must not be confused with a Provider's BPDUs. Customer BPDUs must be either tunneled through the Provider's network or discarded. The Provider also needs to inspect them.
RSTP		
STP		
Pause	P	A bridged Provider network necessarily imposes significant latency penalties on all frames transported. It is unlikely, therefore, that transporting Customers' 802.3X Pause frames through a Provider's network can yield a useful service. The UNI must, therefore, either participate with the Customer Equipment to control the UNI link, or discard Pause frames.
LACP	T/P/D	In an opaque service, LACP packets should be discarded. One may imagine a Customer linking two Customer devices using two parallel point-to-point Provider Bridged services, perhaps obtained from two different Providers, and using LACP to obtain load sharing and fast failover across those services. One may also imagine, perhaps more easily, using LACP across the UNI to increase the bandwidth and reliability of the link between the Customer and the Provider. Thus, LACP may be tunneled, discarded, or participated in, across the UNI.

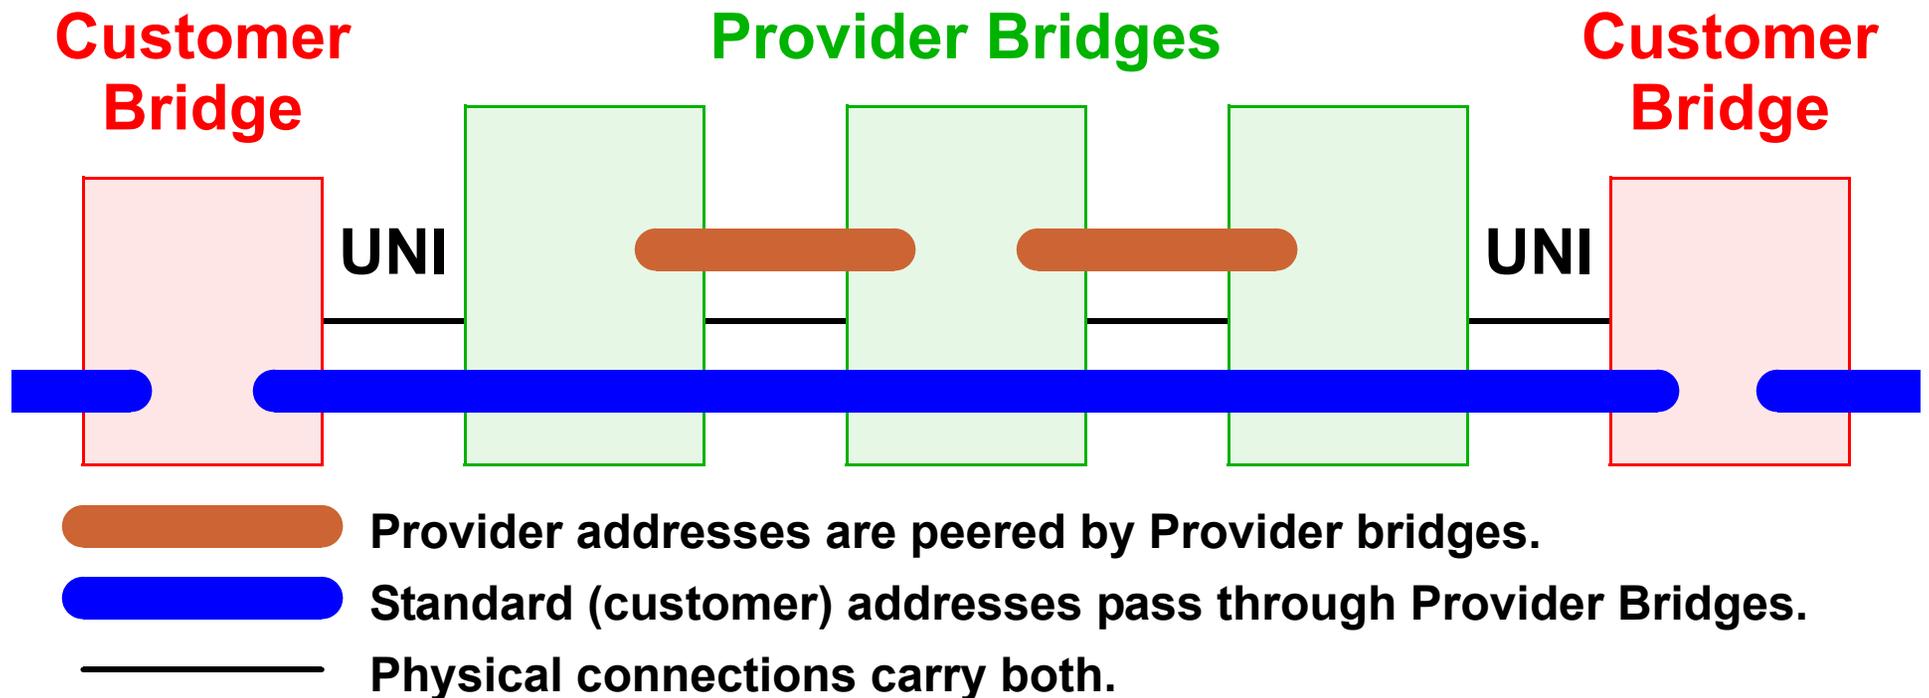
Name	Tunnel	Possibilities
OAM	P	The existing 802.3ah EFM Draft 1.3 OAM protocols are expected by P802.3ah to operate point-to-point across the UNI. Their assumption of which MAC addresses are to be used may need revision, based on the analysis of this present document.
802.1X	T/P/D	As for LACP, one may imagine scenarios for discarding, tunneling, or peering 802.1X packets. In an opaque service, 802.1X packets could be discarded. A Customer may well want to use 802.1X between ends of a Provider network to ensure that the Provider has not incorrectly connected one Customer's equipment to another's. A Customer and Provider may well want to use 802.1X across the UNI to authenticate each other.
IGMP	T	A Provider Bridge should, by default, transport IGMP packets through the Provider network. Enterprise switches commonly intercept, inspect, alter, generate, and/or discard IGMP packets. Such activities in Provider Bridges are for IETF, not IEEE, to define.
VRRP	T	VRRP packet should certainly be transparently transported through the Provider network. They are mentioned, here, as a warning to the implementor of a Provider Bridge that the behavior of the VRRP unicast MAC addresses does not conform to IEEE 802 standards, but VRRP is common enough that a Provider Bridge likely must accommodate it.

The 33 Special Layer 2 MAC Addresses

- **There are 33 addresses: 16 in the BPDU block, 16 in the GARP block, and one “All Bridges” address.**
 - **An IEEE 802.1D bridge never forwards any frame sent to any address in the BPDU block.**
 - **An IEEE 802.1D bridge stops frames sent to the GARP block that it understands, and forwards as normal multicasts frames sent to the GARP block that it does not understand.**
 - **An IEEE 802.1D bridge both receives, and forwards as a normal multicast, frames addressed to the All Bridges MAC address.**
- **Existing Customer devices, of course, utilize the 33 special Layer 2 addresses defined in IEEE 802.1D.**
- **Some protocols are point-to-point and should never be carried over a multipoint-to-multipoint service.**

How do we Tunnel, Discard, or Peer?

- Let us assume that new Provider Bridges normally utilize a new set of 33 special Layer 2 MAC addresses.



- But, how do we address Peers across the UNI?

UNI Peering Plan 1: Configuration

- **Per-Port Configuration:**

- Provider bridge is configured to either tunnel, discard, or peer all 33 special addresses.

- **Per-Address Configuration:**

- Provider bridge is configured separately to either tunnel, discard, or peer each one of the 33 standard MAC addresses.

- Customer Equipment *may* require special configuration to deal with the fact that its peer is different for different protocols.

- **Per-Protocol Configuration:**

- Provider bridge is configured separately to either tunnel, discard, or peer each one of Layer 2 protocols.

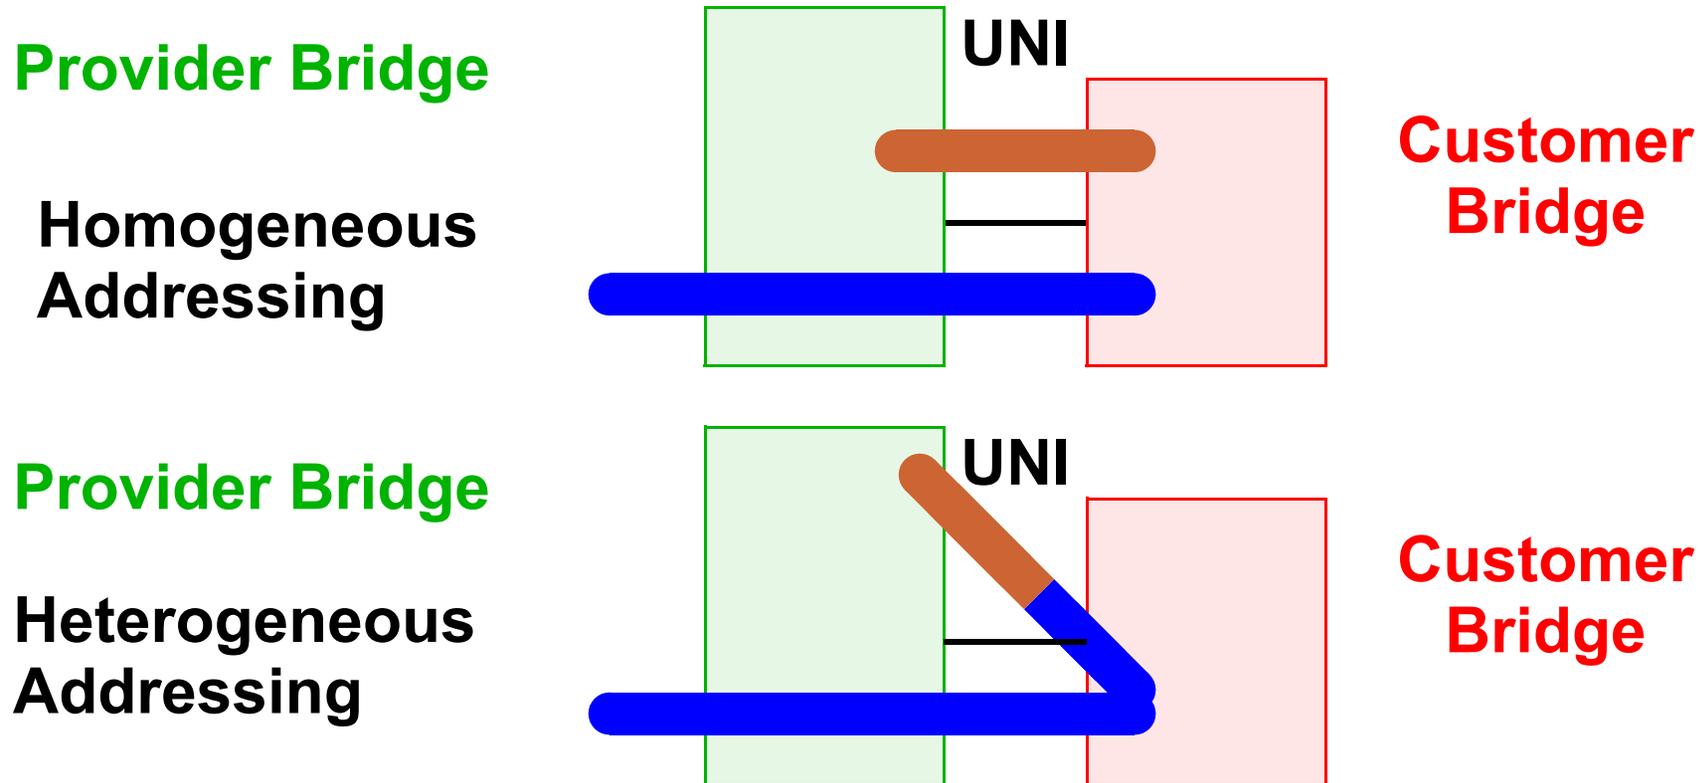
- Customer Equipment *may* require special configuration to deal with the fact that its peer is different for different protocols.

UNI Peering Plan 1: Configuration

- **This style of configuration allows a given protocol to either tunnel, peer, or be discarded.**
- **It does not allow a protocol, e.g. 802.1X, to *both* peer and tunnel.**
- **These configuration options do not allow a Customer both to check the Provider's connections, by running 802.1X transparently through the Provider, and to check the UNI, by peering 802.1X across the UNI.**
- **However, this style of configuration, for the most part, allows existing equipment to run the existing protocols with some flexibility as to tunneling vs. peering.**

UNI Peering Plan 2: Addressing

- If both the devices and the standards are modified to operate using either the Standard 33 MAC addresses, the Provider's 33 MAC addresses, *or both*, then simultaneous use of a Layer 2 protocol is possible.



UNI Peering Plan 2: Addressing

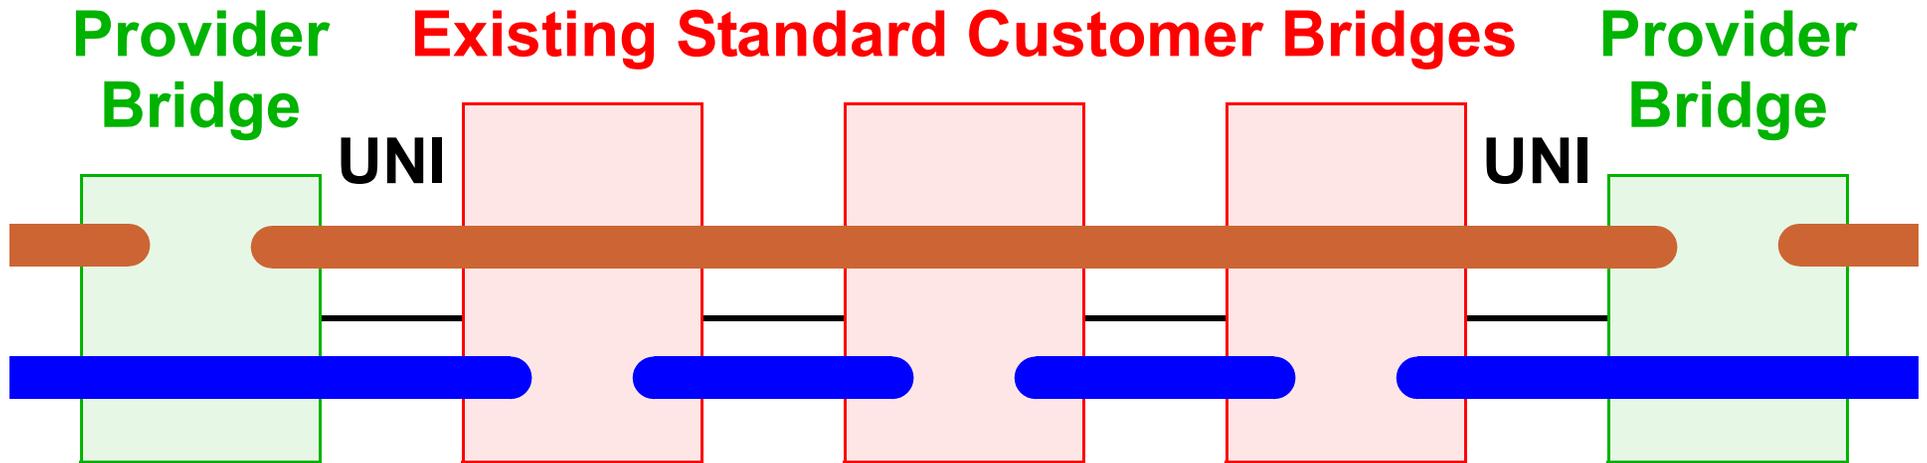
- **Homogeneous Addressing**

- Customer and Provider communicate as peers across the UNI using the Provider's set of 33 special multicast MAC addresses.
- Source MAC addresses and/or addressing information carried in the particular protocol PDUs must be used by the Customer Equipment to distinguish between end-to-end and Provider-Customer conversations.

- **Heterogeneous Addressing**

- Provider sends standard Multicast MAC addresses, but Customer sends Provider Multicast MAC addresses, across the UNI.
- **Problem:** Customer cannot distinguish between multicasts sourced from a device on the other side of the Provider network, and multicasts sourced by its local Provider Bridge, in the first conversation initiated by either one.

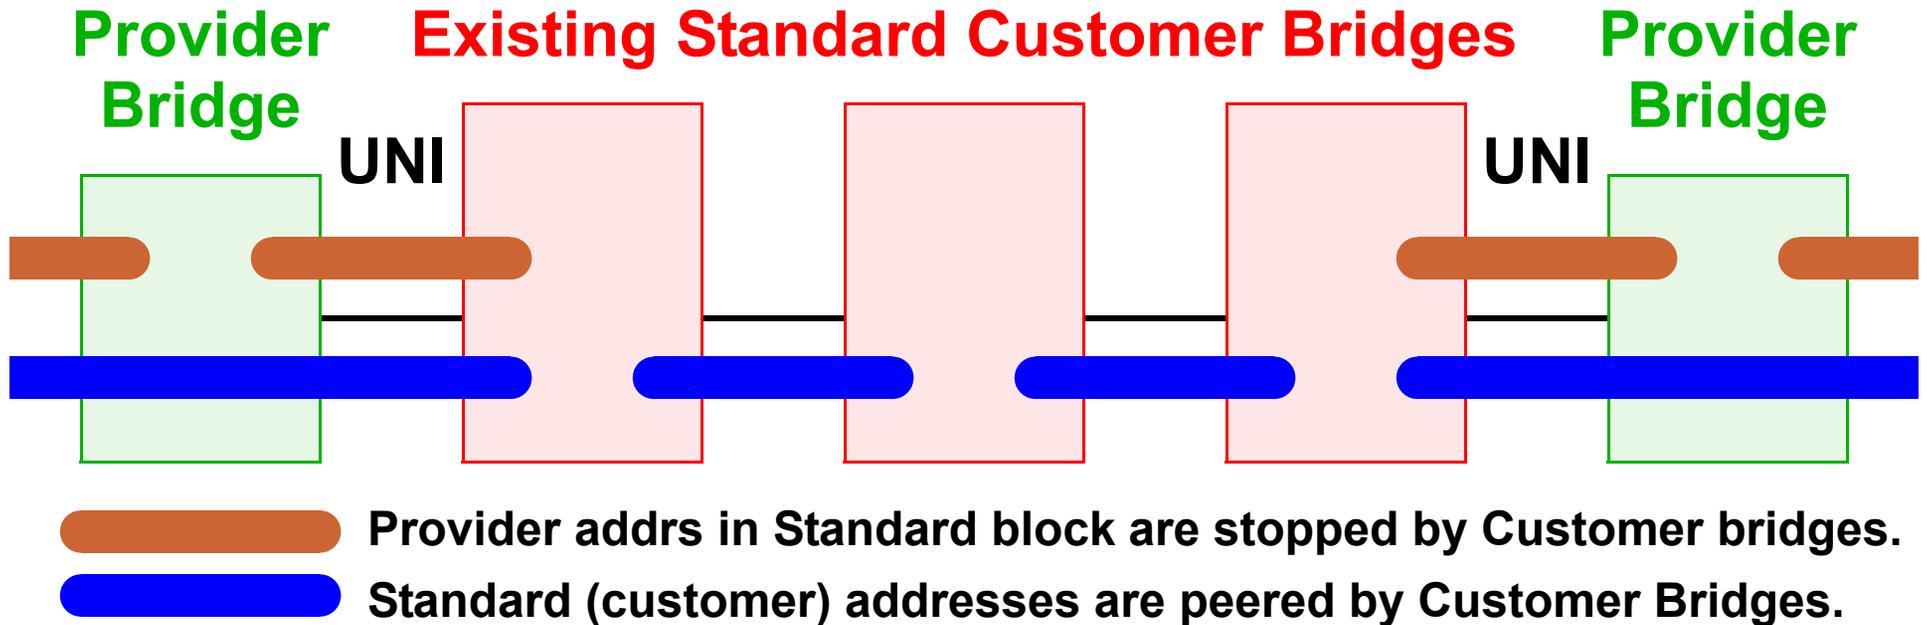
What Happens If ...



-  Provider addresses pass through Customer bridges.
-  Standard (customer) addresses are peered by Customer Bridges.

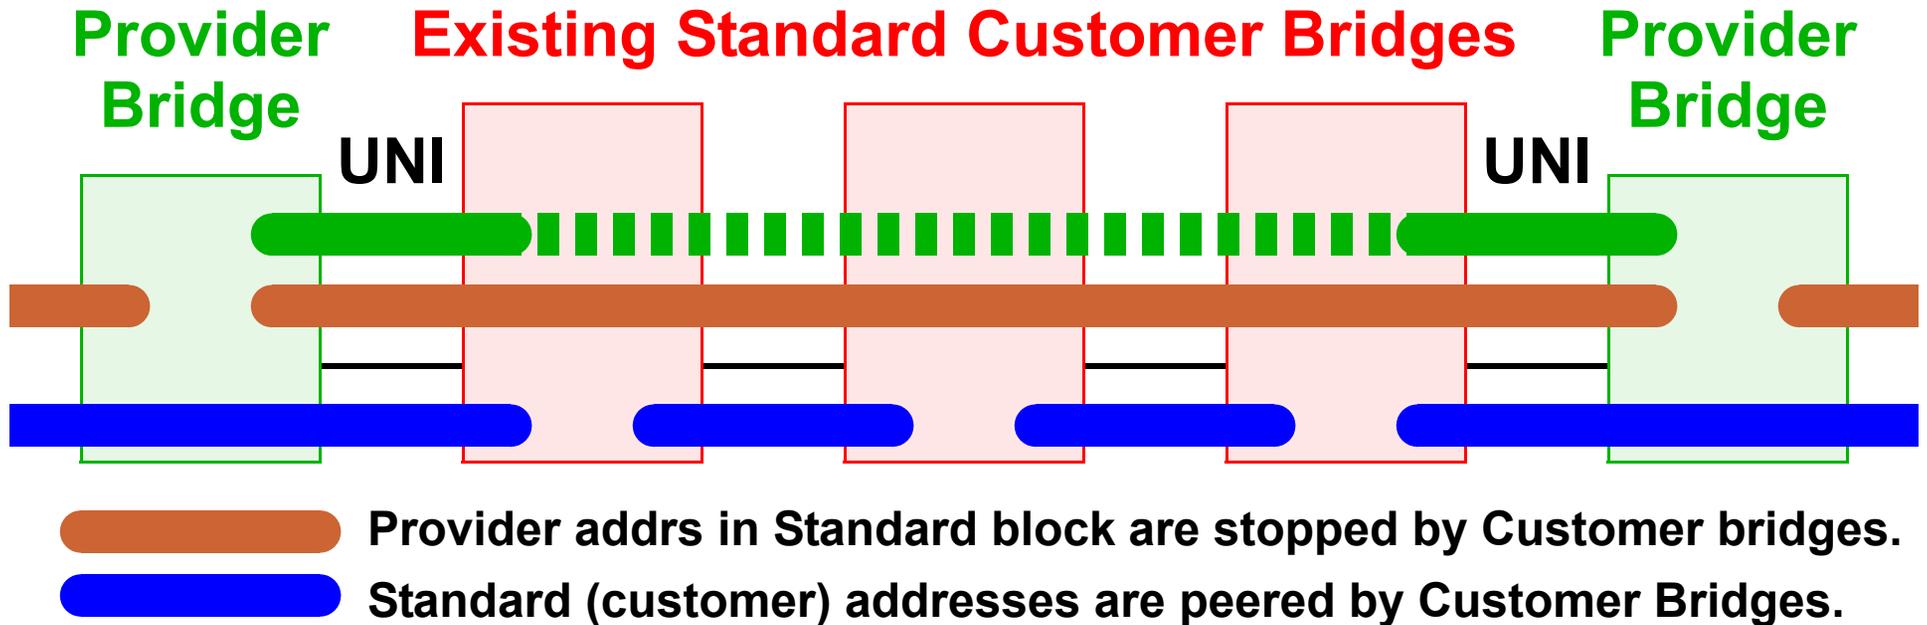
- If the Provider Bridge emits a frame destined to one of the 33 Provider addresses towards a dual-homed Customer that does not understand them, then that frame *may* traverse the Customer's network and hit one or more other Provider Bridges, or even hit another Provider's Provider Bridge.
- This frame may allow the Provider to do additional error checking, but may also annoy the Customer, and may be filtered.

What Happens If ...



- **If Provider Bridge BPDU block is allocated from the existing standard BPDU block, instead of a whole new block, then Provider's control frames are stopped by Customer's bridges.**
 - **This is safer, but lessens Provider's chances for detecting errors.**

What Happens If...



- If Provider and Customer agree to a *third* set of MAC addresses just for use across a UNI?
- These addresses would be transmitted through a Customer's network, and would cause the problems described in Slide 15.

UNI Peering Plan 3: VLAN Tags

- **One VLAN could be used for all Provider/Customer Peering.**
- **This might be easier than using MAC addresses for some protocols.**
- **There are problems, however:**
 - **A VLAN must be selected and/or negotiated.**
 - **Some protocols already may have VLAN tags; they would have to be doubly tagged.**
 - **Existing standards and Provider port ASICs would have to treat the same special address differently for different VLANs.**

Is IEEE 802.3ah EFM Draft 1.3 OAM Compatible with these Models?

- **Perhaps...**

- **If OAM uses the existing Slow Protocols MAC address, it will be difficult for a Provider Bridge to distinguish it from a Customer BPDU, or worse, another Customer Slow Protocols packet, except in software. But, this would be inconsistent if other protocols use MAC addresses to distinguish UNI and tunneled protocols.**
- **If OAM uses the new Provider Slow Protocols MAC address, then it will be difficult for the Provider to determine whether the Customer knows about them. The Provider may see other Provider Bridges (or even another Provider's Provider Bridges!) through a dual-homed Customer. Some protocols may be filtered while others are not.**

How Can OAM be Fixed?

- **Retain current definition of OAM addresses.**
 - It is then difficult for the Provider Bridge to decide what to do with a Slow Protocols frame, and inconsistent with other protocols.
- **Heterogeneous Addressing: Customer sends to the Provider Slow Protocols address, Provider sends to the Customer Slow Protocols address.**
 - This has the problem with confusion in the Customer bridge between local and remote transmissions discussed in Slide 14.
- **Homogeneous Addressing: Customer and Provider use the same addresses, either Standard or Provider.**
 - As discussed on Slide 19, this makes it difficult, either for the Provider or the Customer, to do the right thing.

But Perhaps there *is* a Solution

- **If we select one of the existing 16 BPDU addresses address for a new UNI Test protocol which can determine whether this link does or does not cross a Provider UNI, then we can use any of the above solutions without fear.**
 - **Separate Provider, Customer, and UNI address sets should be used, if it is established that both devices understand the New World.**
 - **A configured solution using the standard 33 addresses should be used, if the Customer device does not understand Provider Bridges.**

How Can OAM be Fixed?

- **Or, if OAM uses the new UNI Test address, everything is solved.**
- **OAM could, perhaps, *be* the UNI Test protocol.**
- **Better yet, use LLDP!**