CFM Issues

Norman Finn

The “3-view” of Connectivity Fault Management

The View From Above

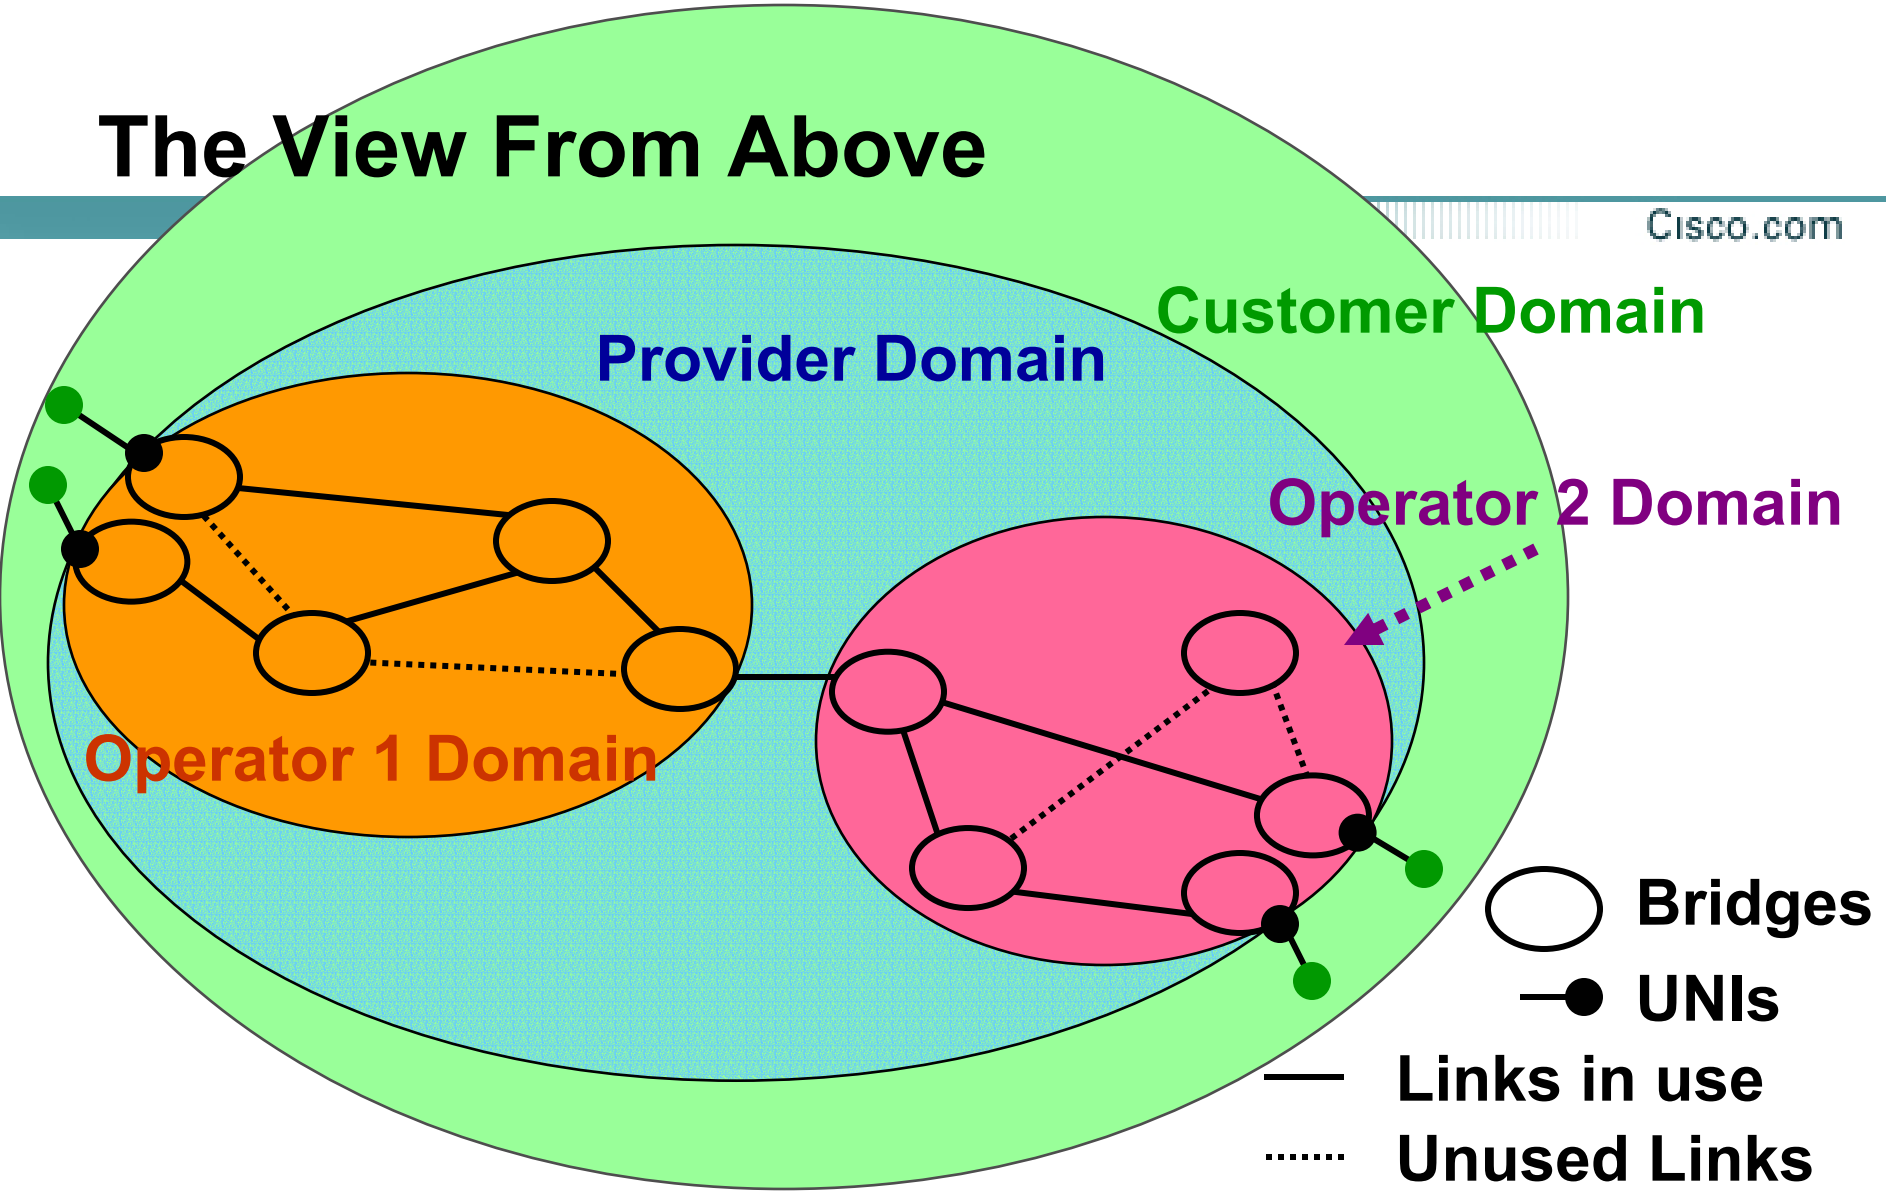

- Physical view of a network using CFM

The View From Above

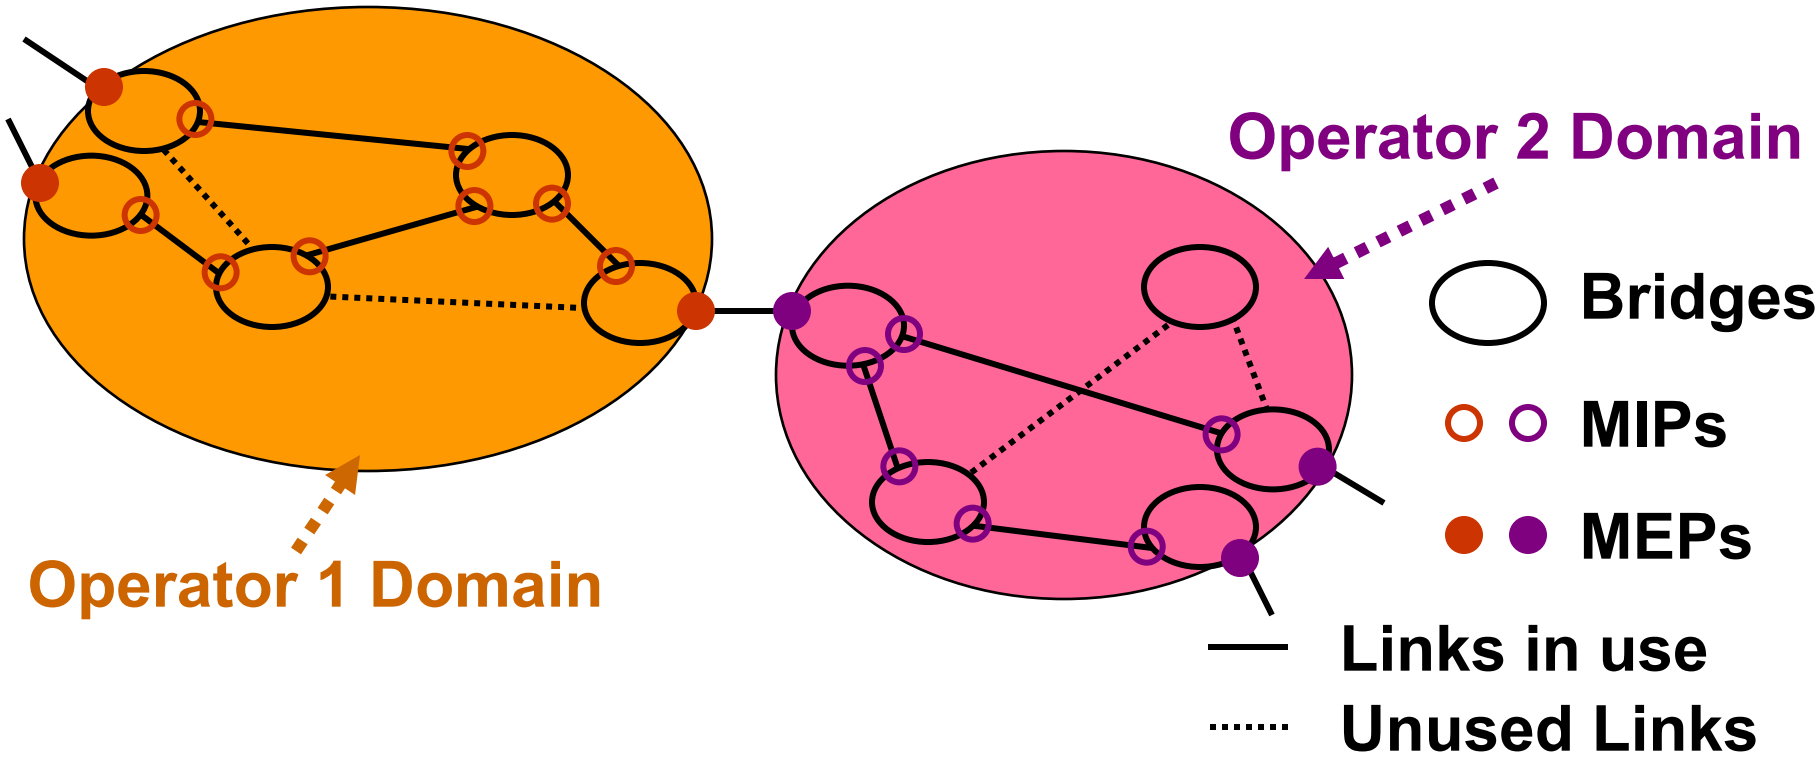

- Operators' view from above

The View From Above

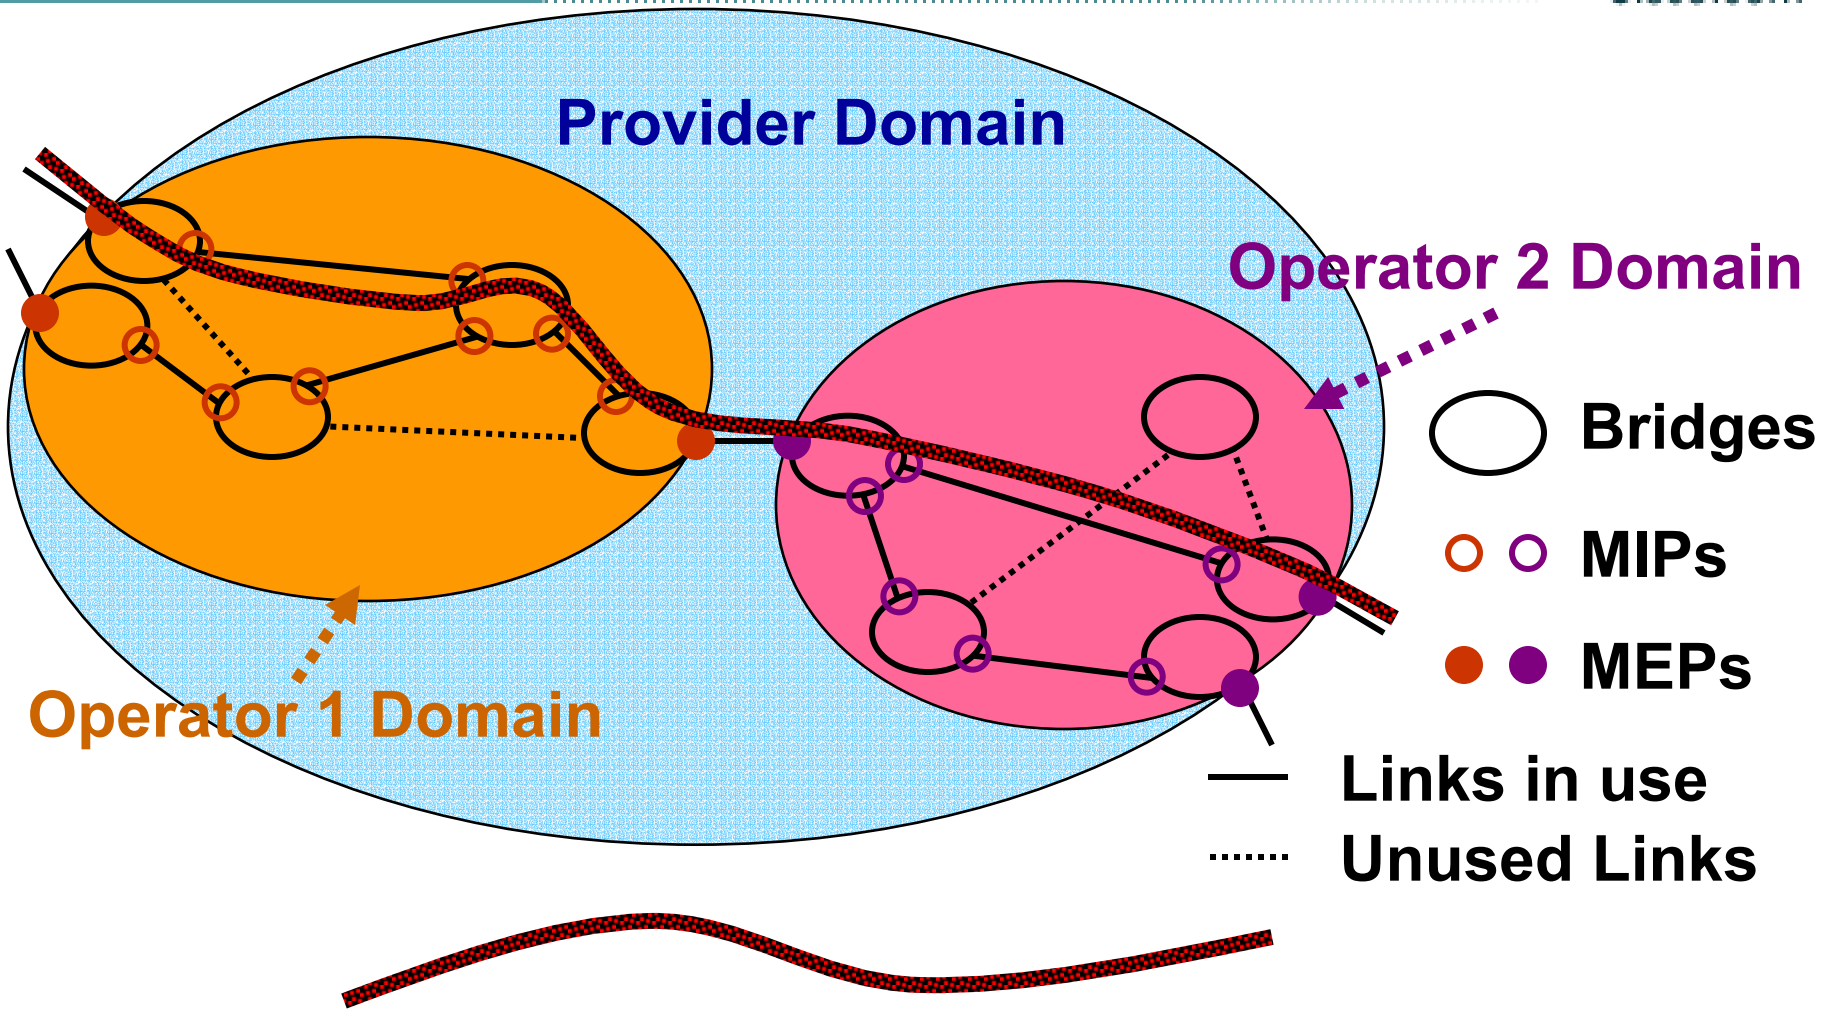

- Let's take a slice along one path …

Longitudinal View

Cisco.com

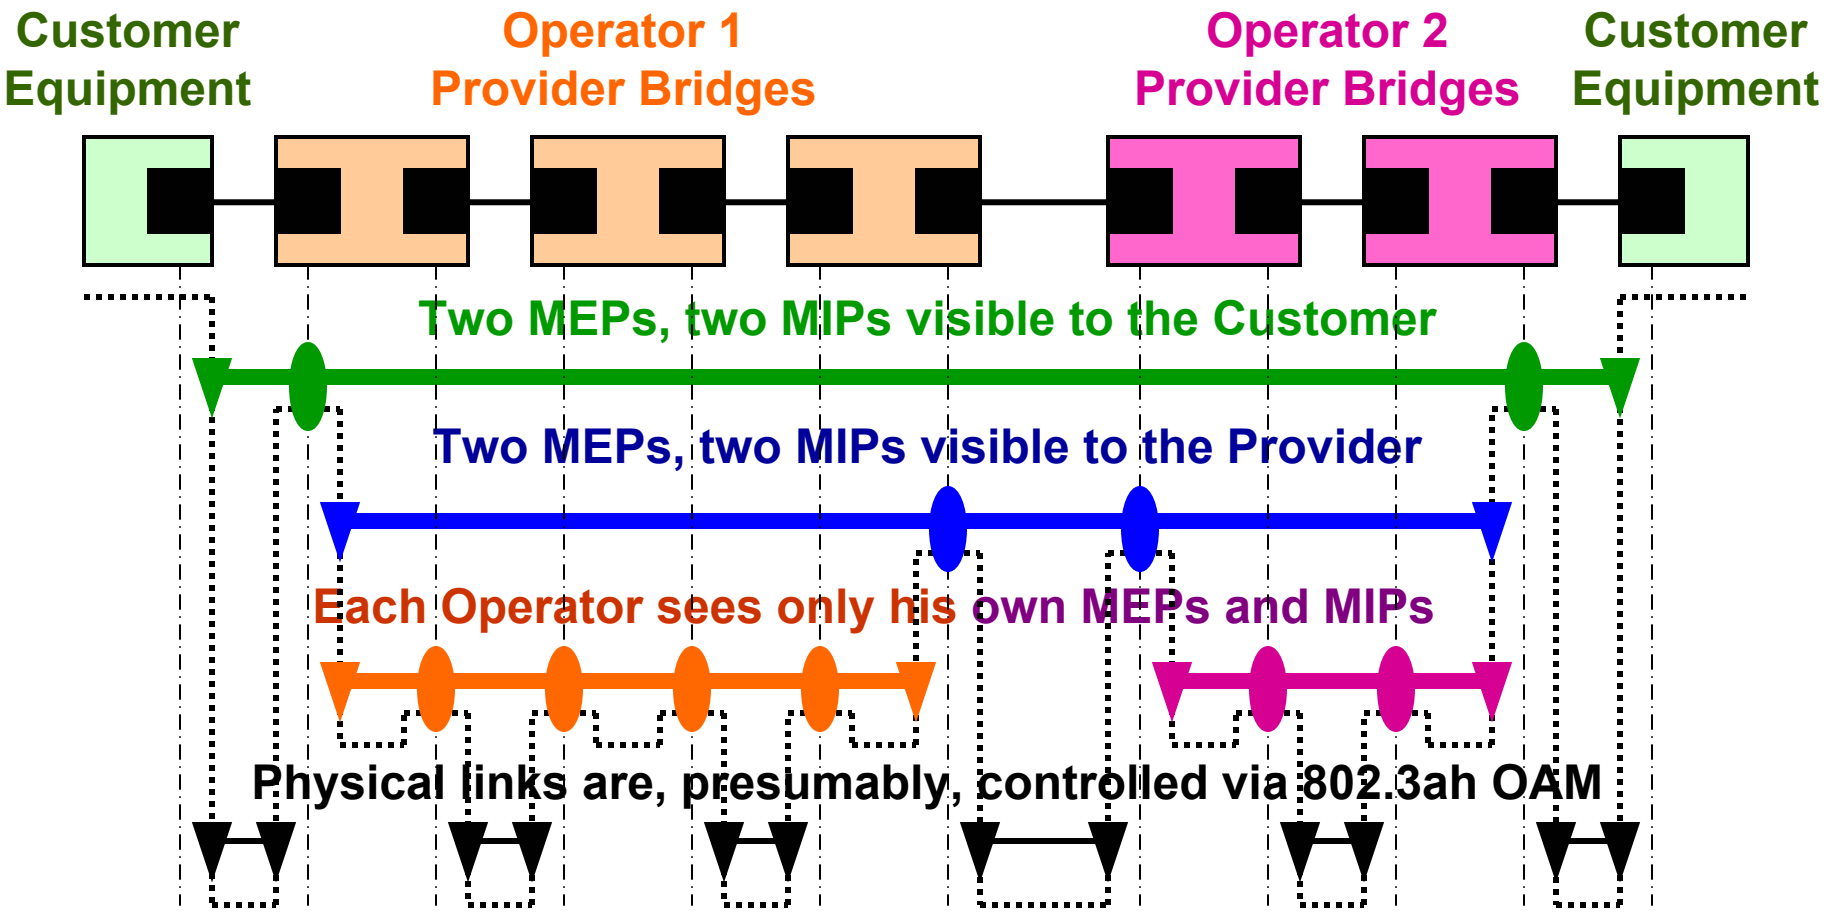

- We can see the MIP/MEP relationships.

Longitudinal View

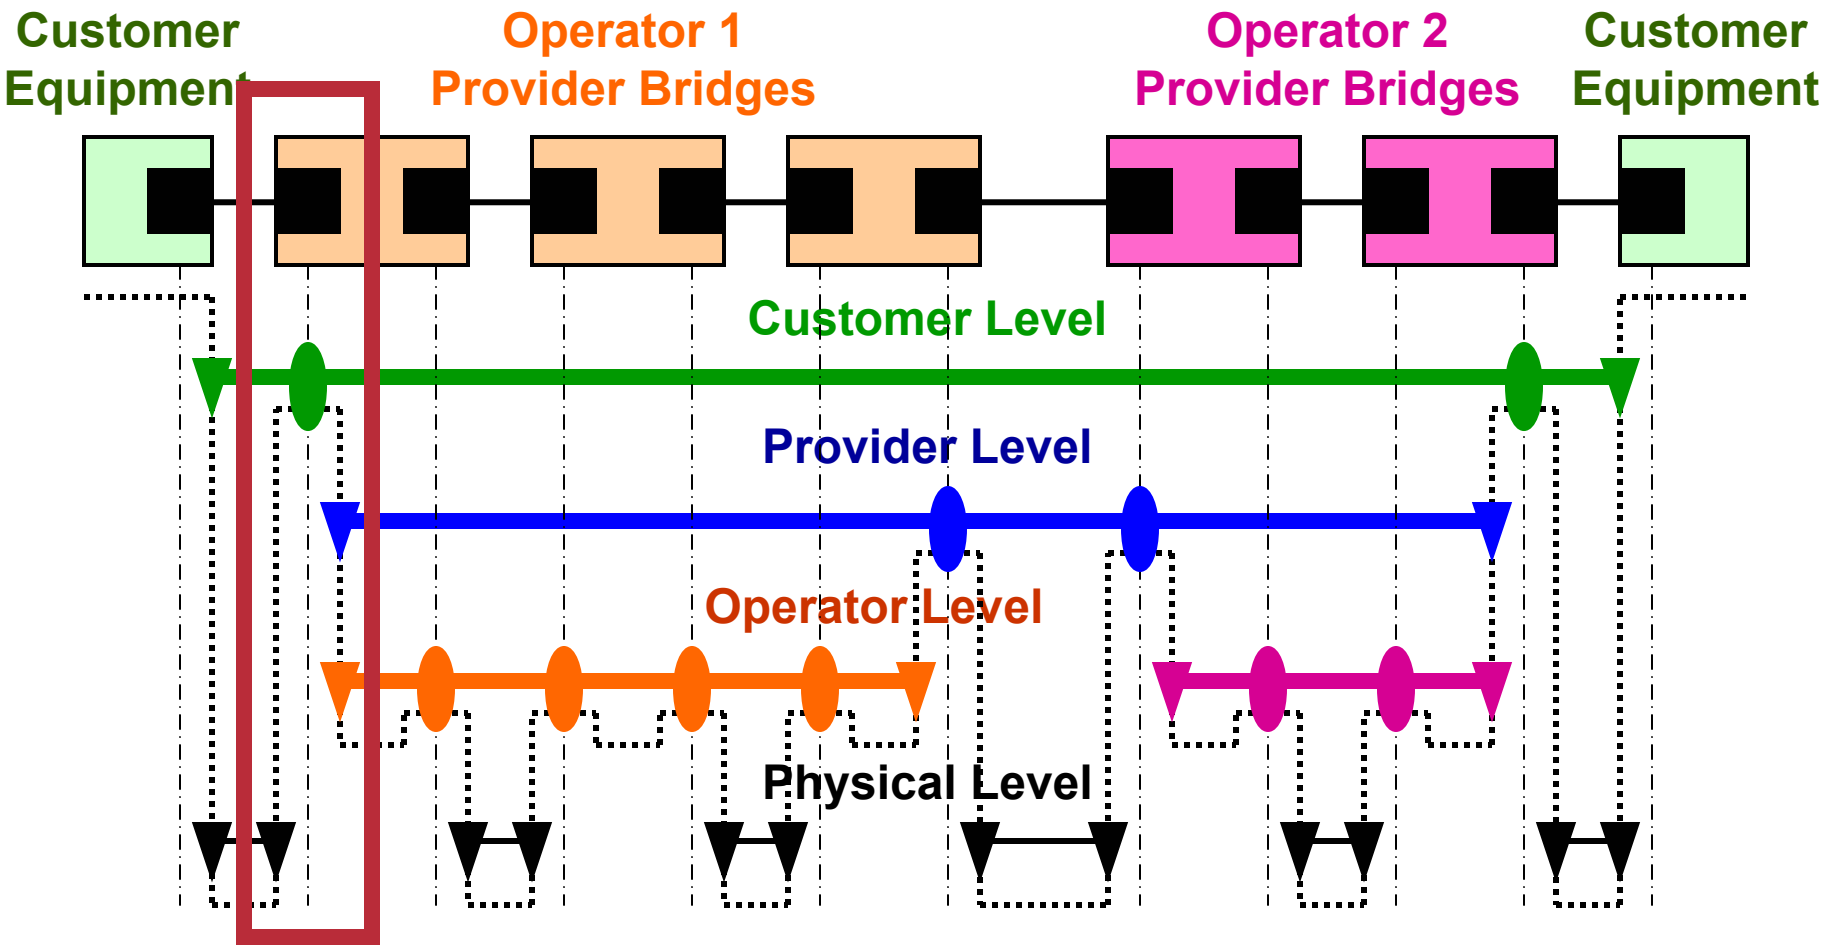

- Let's look at one port from an end view …

Longitudinal View

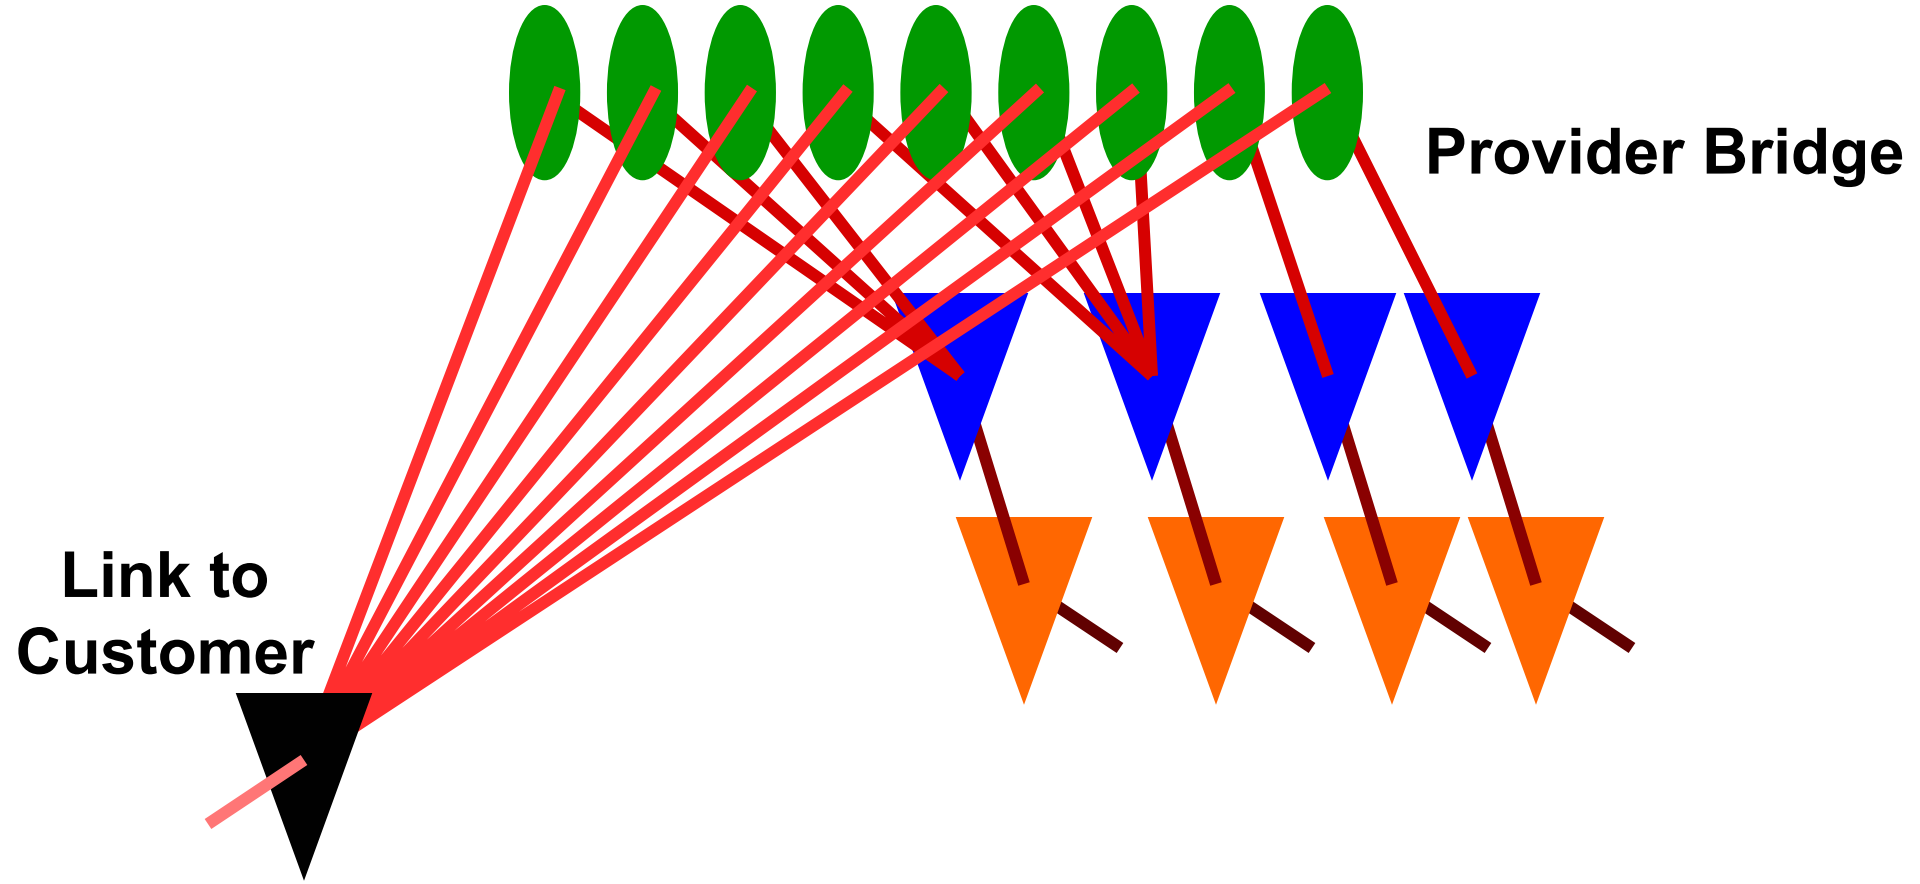

- From the end, we see the multiplexing.

Alarm Indication Signal

What is AIS?

- When a MEP discovers an error, presumably a Continuity Check failure, it transmits periodic multicast AIS in the opposite direction from that in which it runs the CCs.
- The AIS includes a Lifetime TLV that causes it to expire after a certain time.
- AIS transmission ceases if the failure is repaired via Spanning Tree.

Alarm Indication Signal

Cisco.com

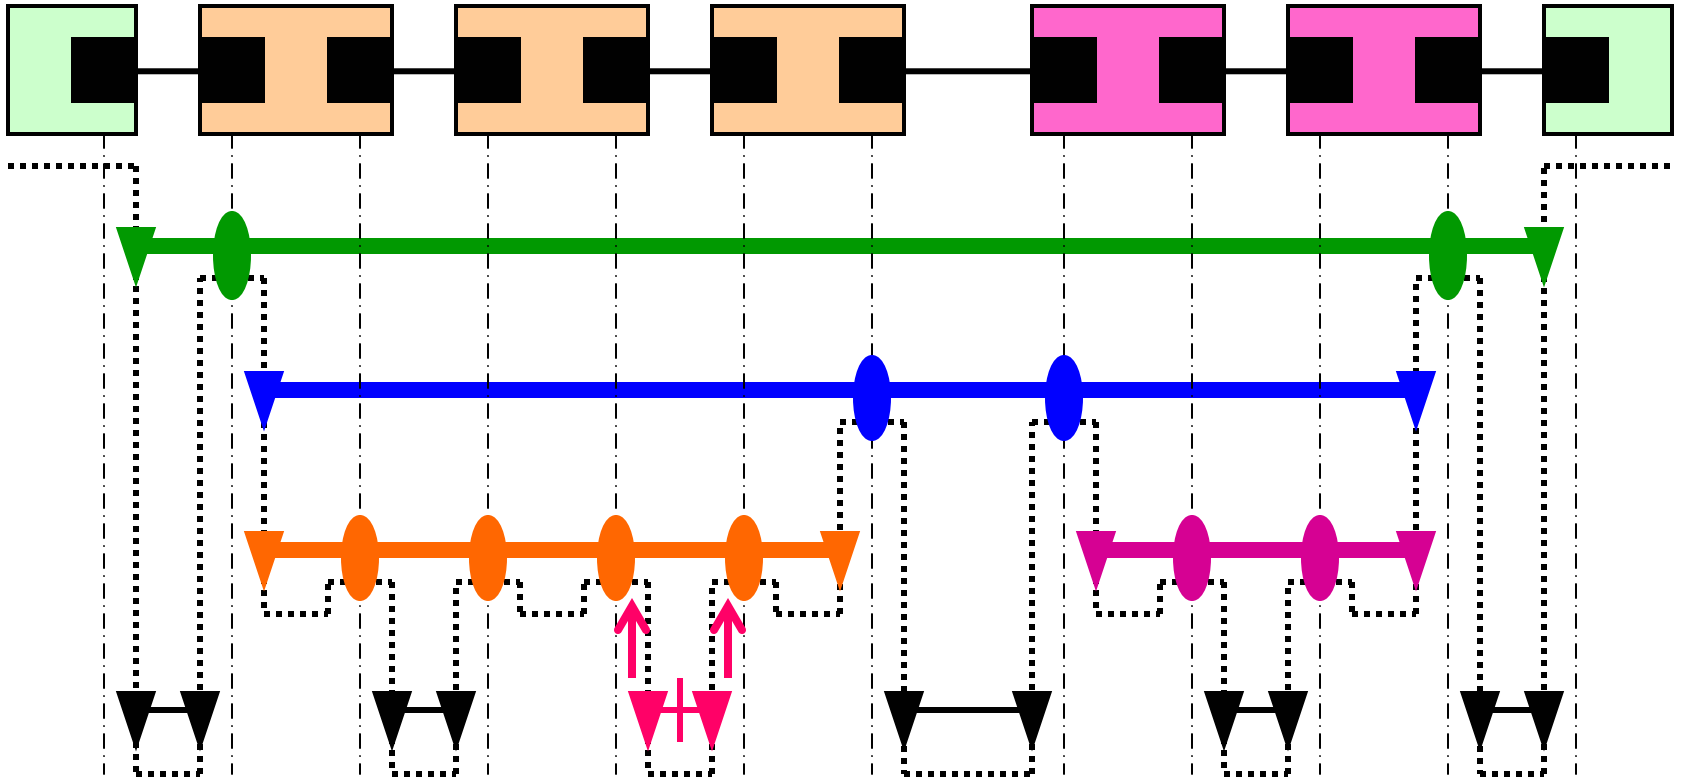

- **Failure causes MEP to generate AISs in *opposite* direction of Continuity Check.**

Alarm Indication Signal

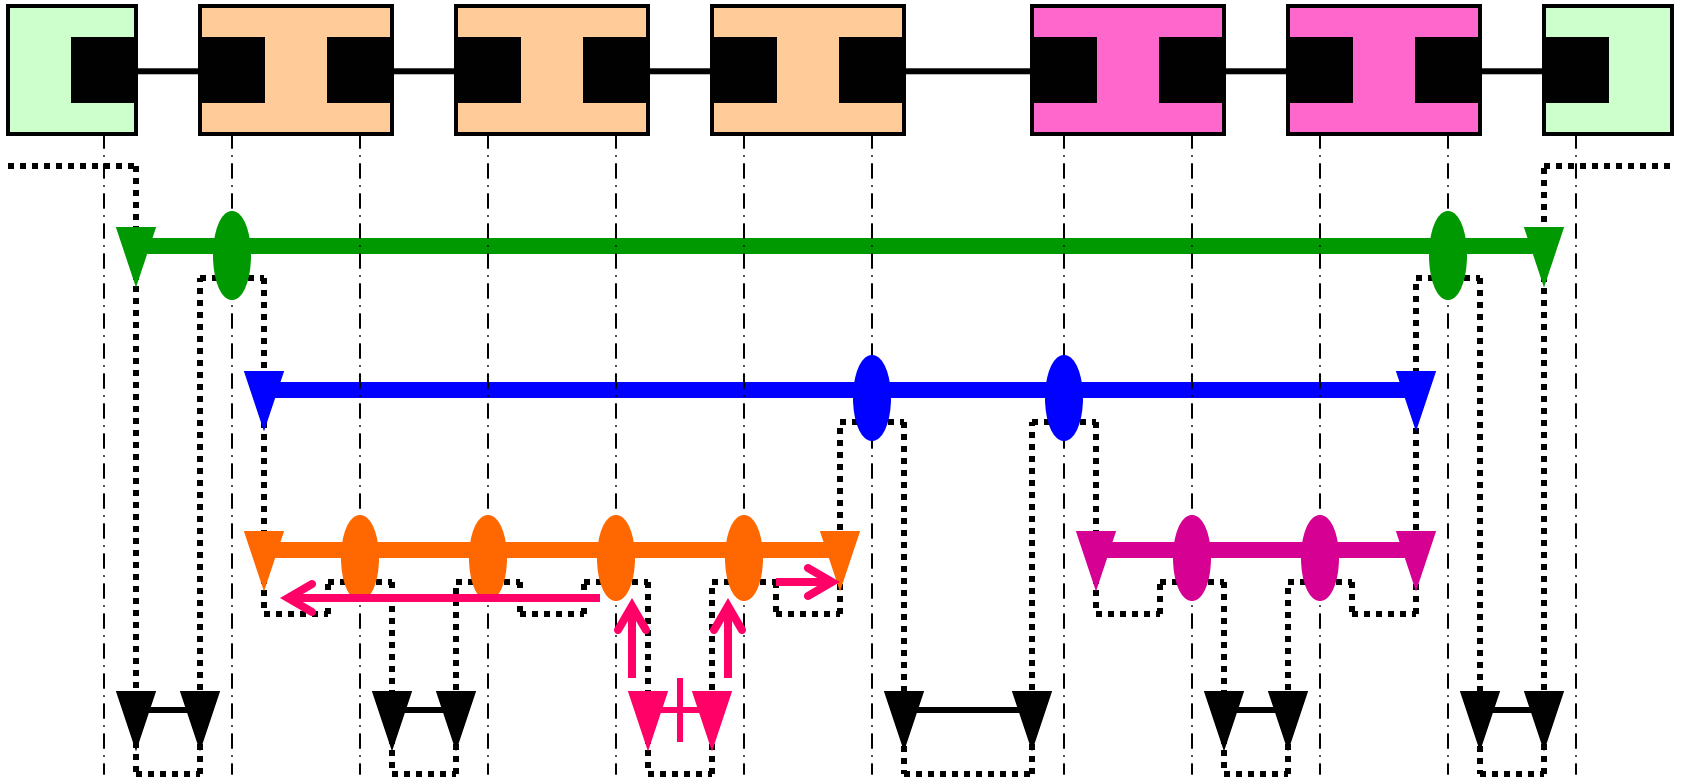

- MIPs receive AISs and promote them to their own ME Levels for propagation to MEPs.

Alarm Indication Signal

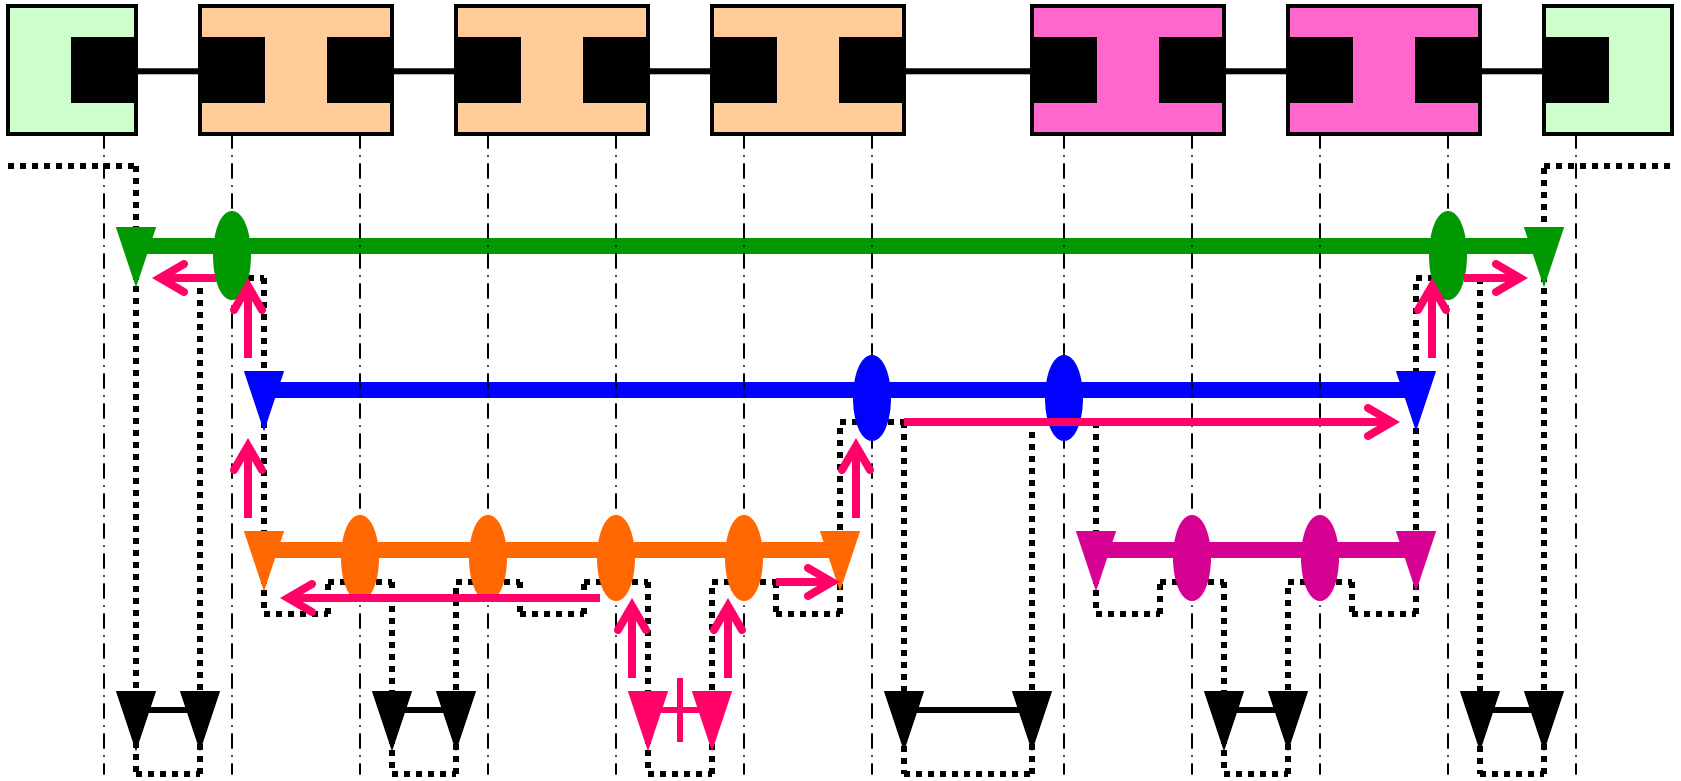

- MEP may propagate or terminate AIS.
- Propagation is shown.

Alarm Indication Signal

What is AIS Good for?

- **Receiving MEPs may catalogue AIS and wait to see whether their own CCs report a failure.**

  If Spanning Tree repairs error, none need be generated.

  If not, Network Administrator has already been notified of the problem; error trap from MEP is suppressed in order to reduce unnecessary alarms.

Alarm Indication Signal

What is AIS Good for?

- **Receiving MEPs may delay the propagation of AISs.**

  This gives Spanning Tree time to correct the problem before alarming higher levels.

- **Receiving MEPs may take action to correct the failure reported by AIS.**

  This assumes that there is no Spanning Tree to correct the problem.

Alarm Indication Signal

Multiplexing

- There may be a one-to-many relationship between lower and upper MEPs/MIPs on a given interface.

  For example, one physical link may carry many service instances.

  One Provider service instance may carry many Customer VLANs.

- This causes a multiplication of AISs.

Longitudinal View: The correct picture

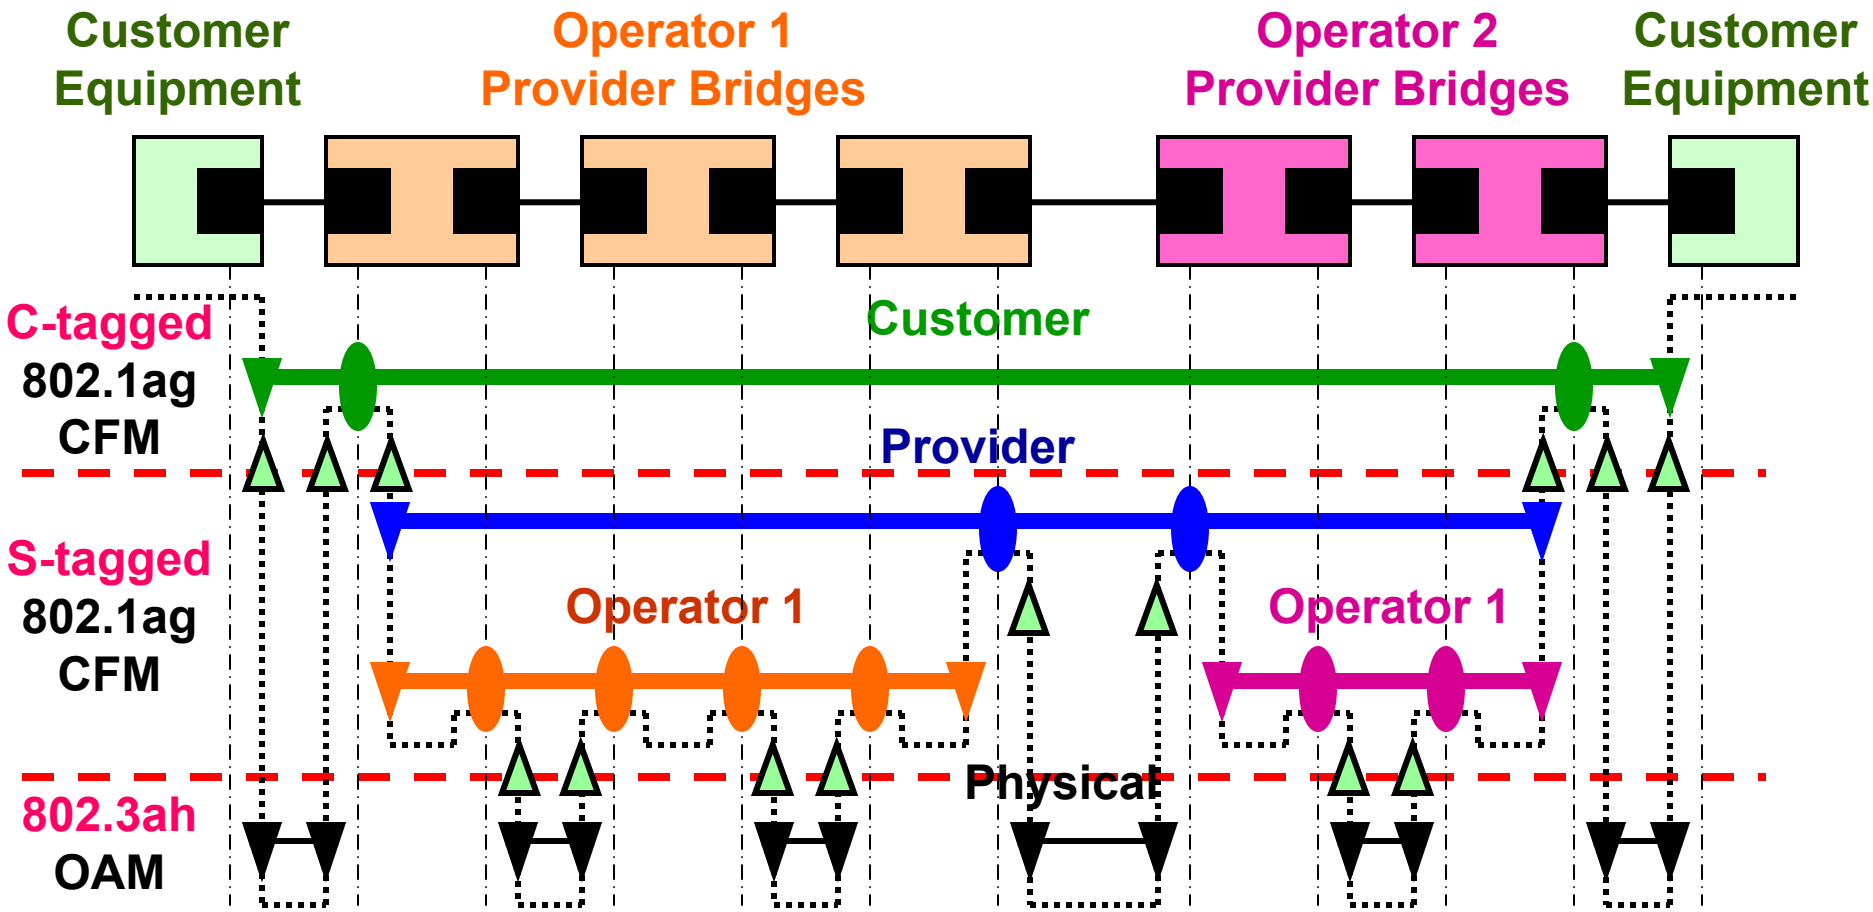

- **AIS translation** is required at OAM method and/or multiplexing boundaries.

Alarm Indication Signal (slightly wrong picture)

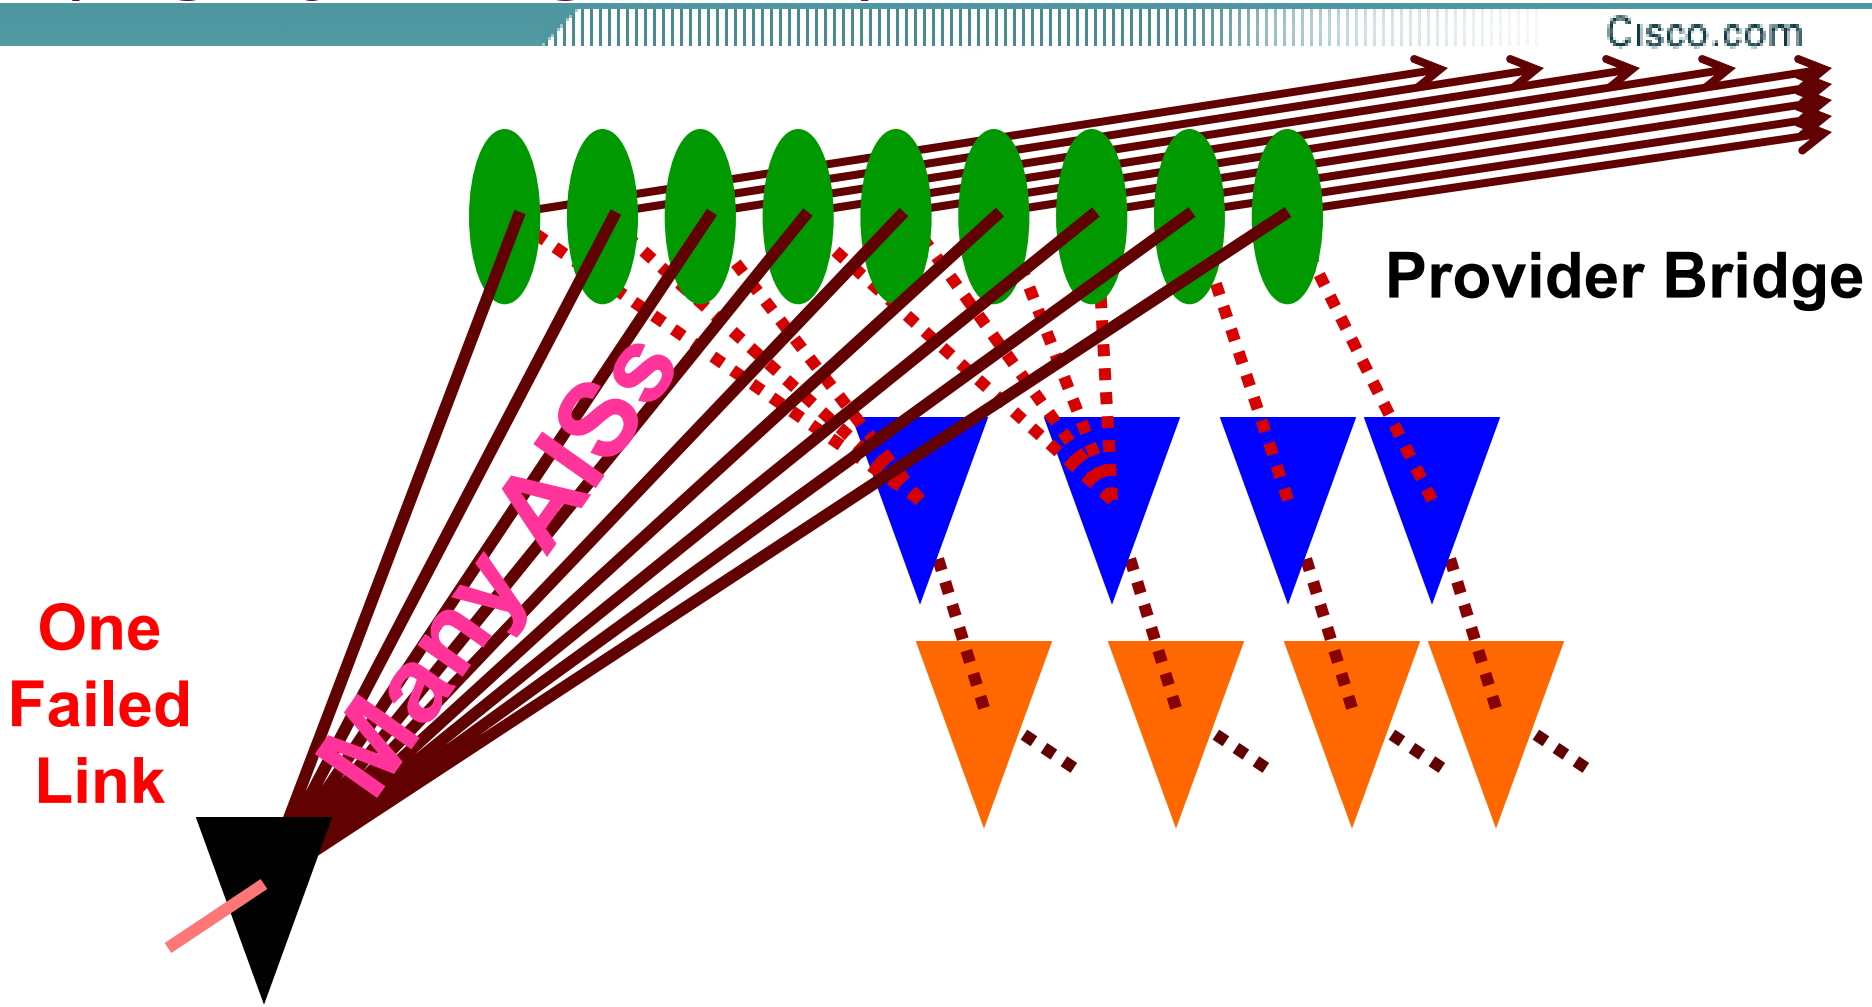

- One physical failure, many AISs.

Alarm Indication Signal (right picture)

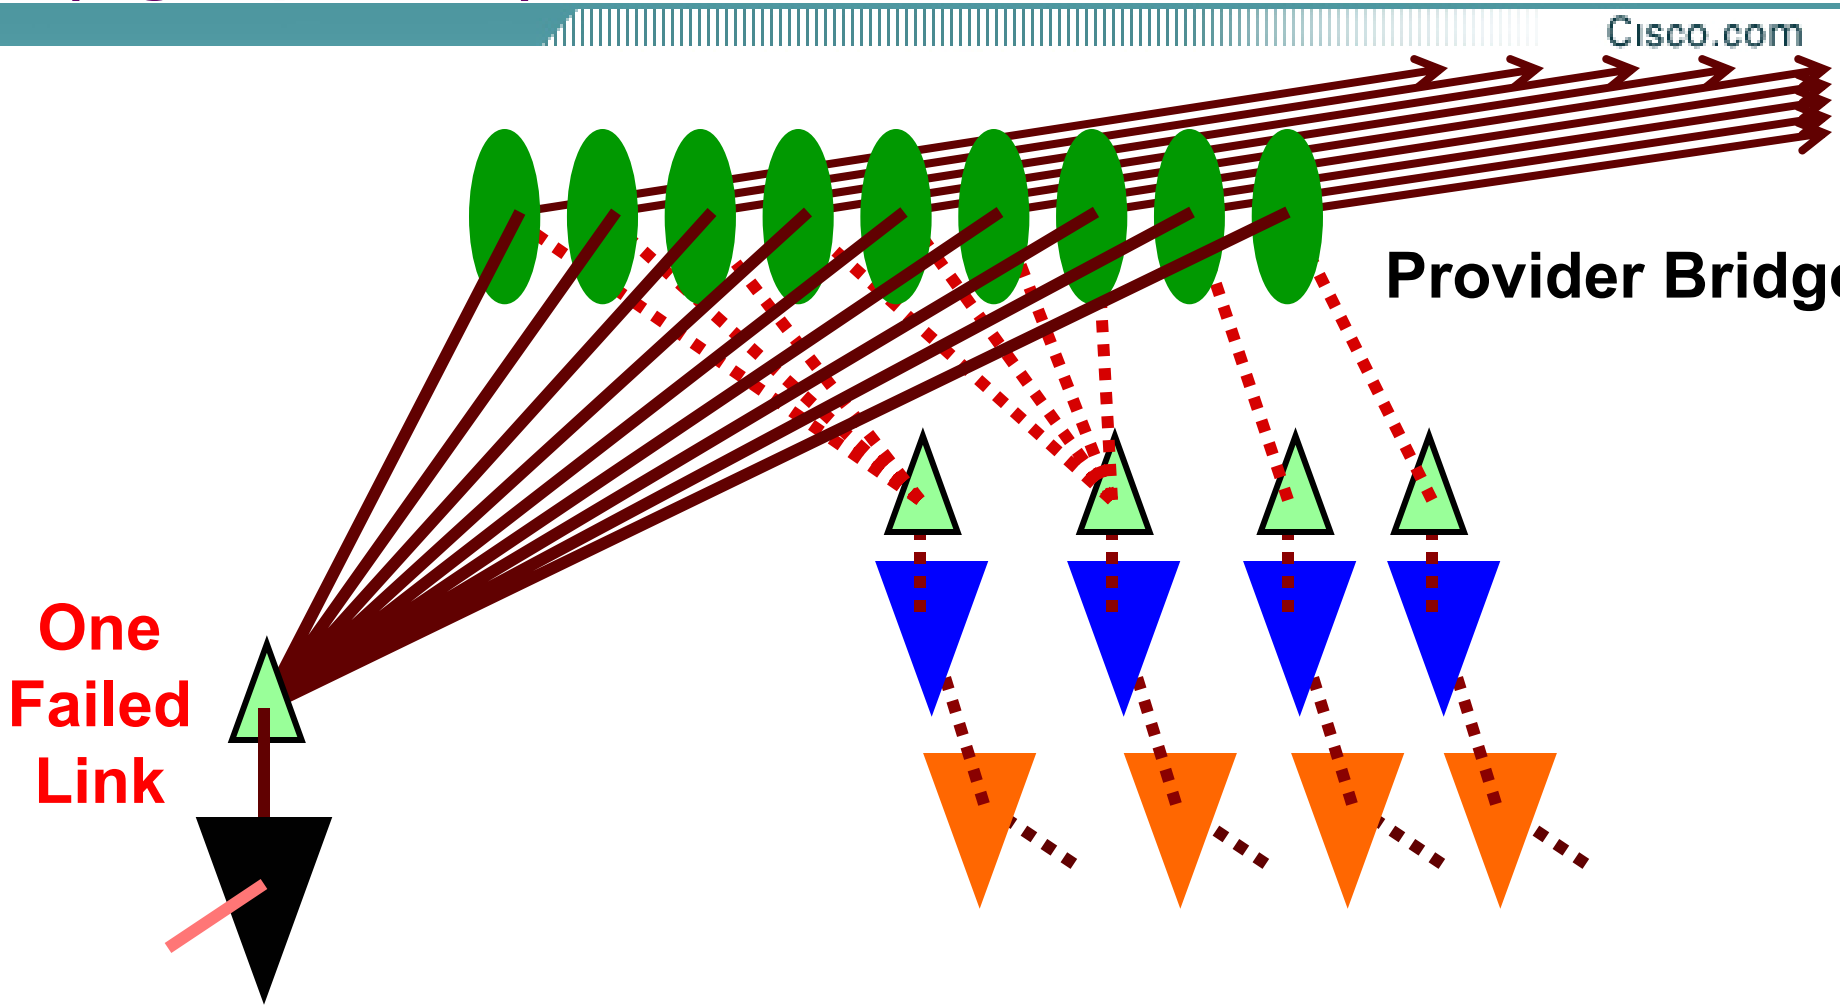

- An AIS Translation module is required to do the multiplexing.

AIS Translation

- **The AIS is theoretically generated by the MEP and sent up its “back door”.**
- **There is no point in doing this, and the feature should not be enabled, unless there is a MIP above the MEP to “promote” the AIS to the MIP’s ME Level.**

AIS Translation

- Therefore, the AIS generated by the MEP has no physical significance until it is promoted.
- So, the AIS Translation function is buried inside the port's CFM code, and is not necessarily a PDU translator.

AIS Translation

- Using the (perhaps fictional) AIS Translation function, a failure of an MPLS Pseudowire, reported by MPLS's or the Pseudowire's own integrity functions, may generate an Ethernet CFM AIS.
- Similarly, AIS from an underlying medium could actually be translated into CFM AIS.

Network Edge Problem #1

What happens when the Provider-Customer link fails?

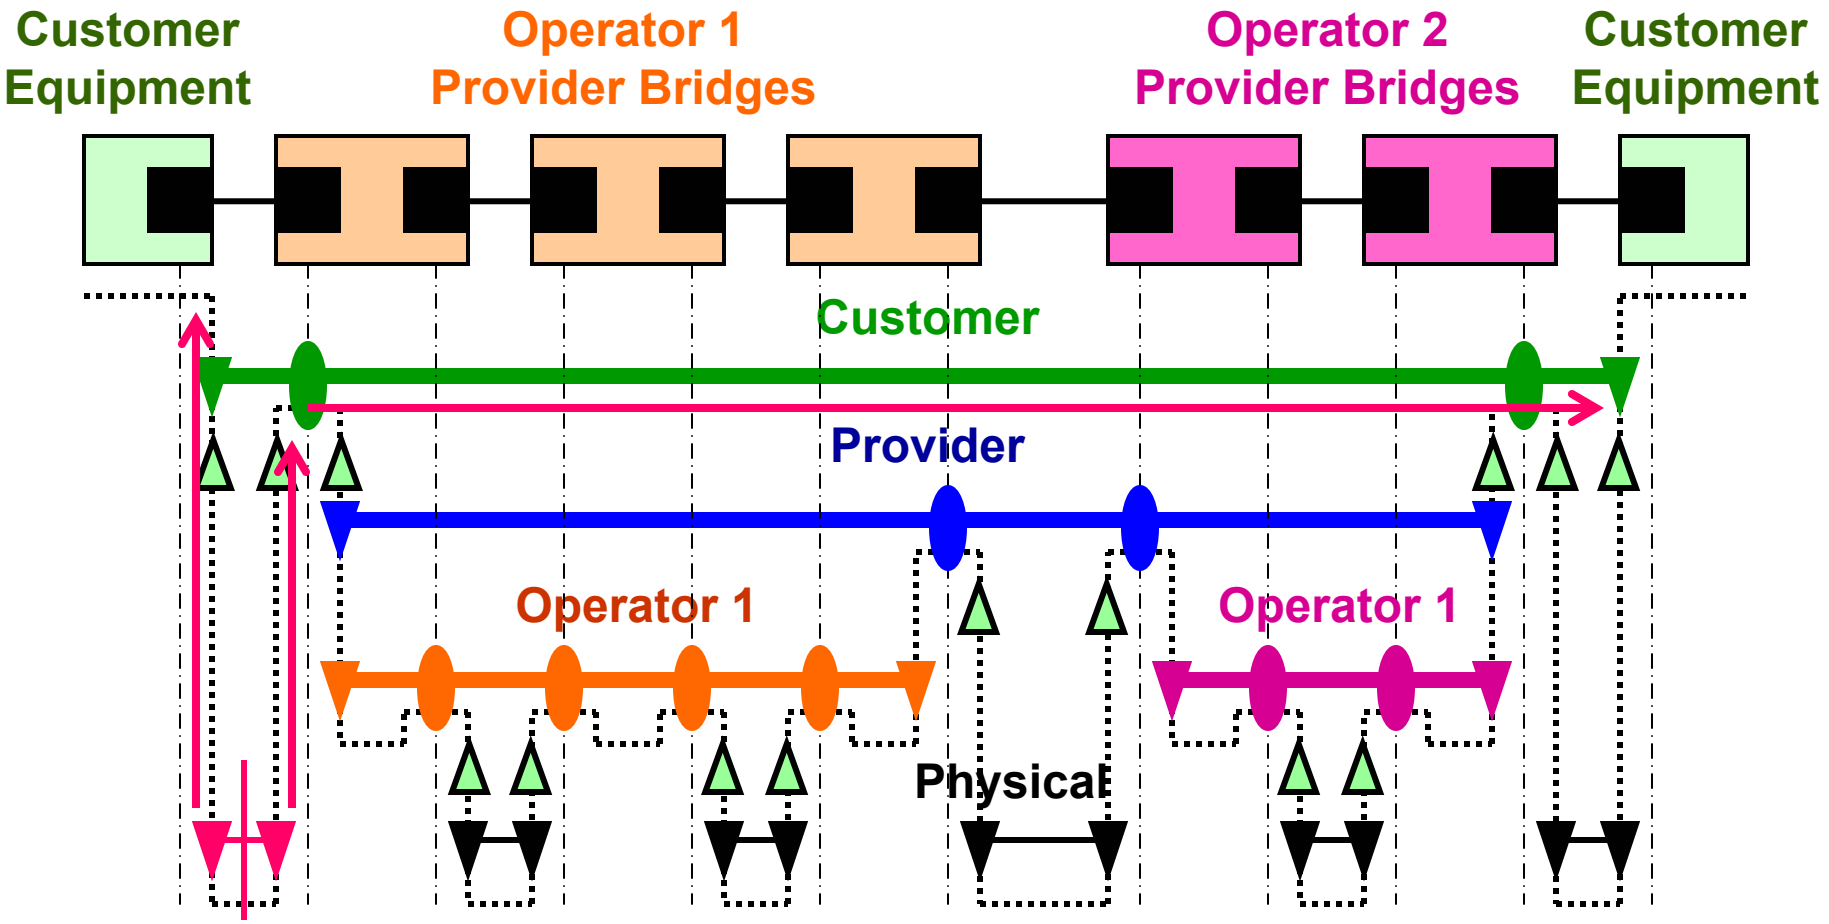

- AIS is sent at the Customer ME Level.

What happens when the Provider-Customer link fails?

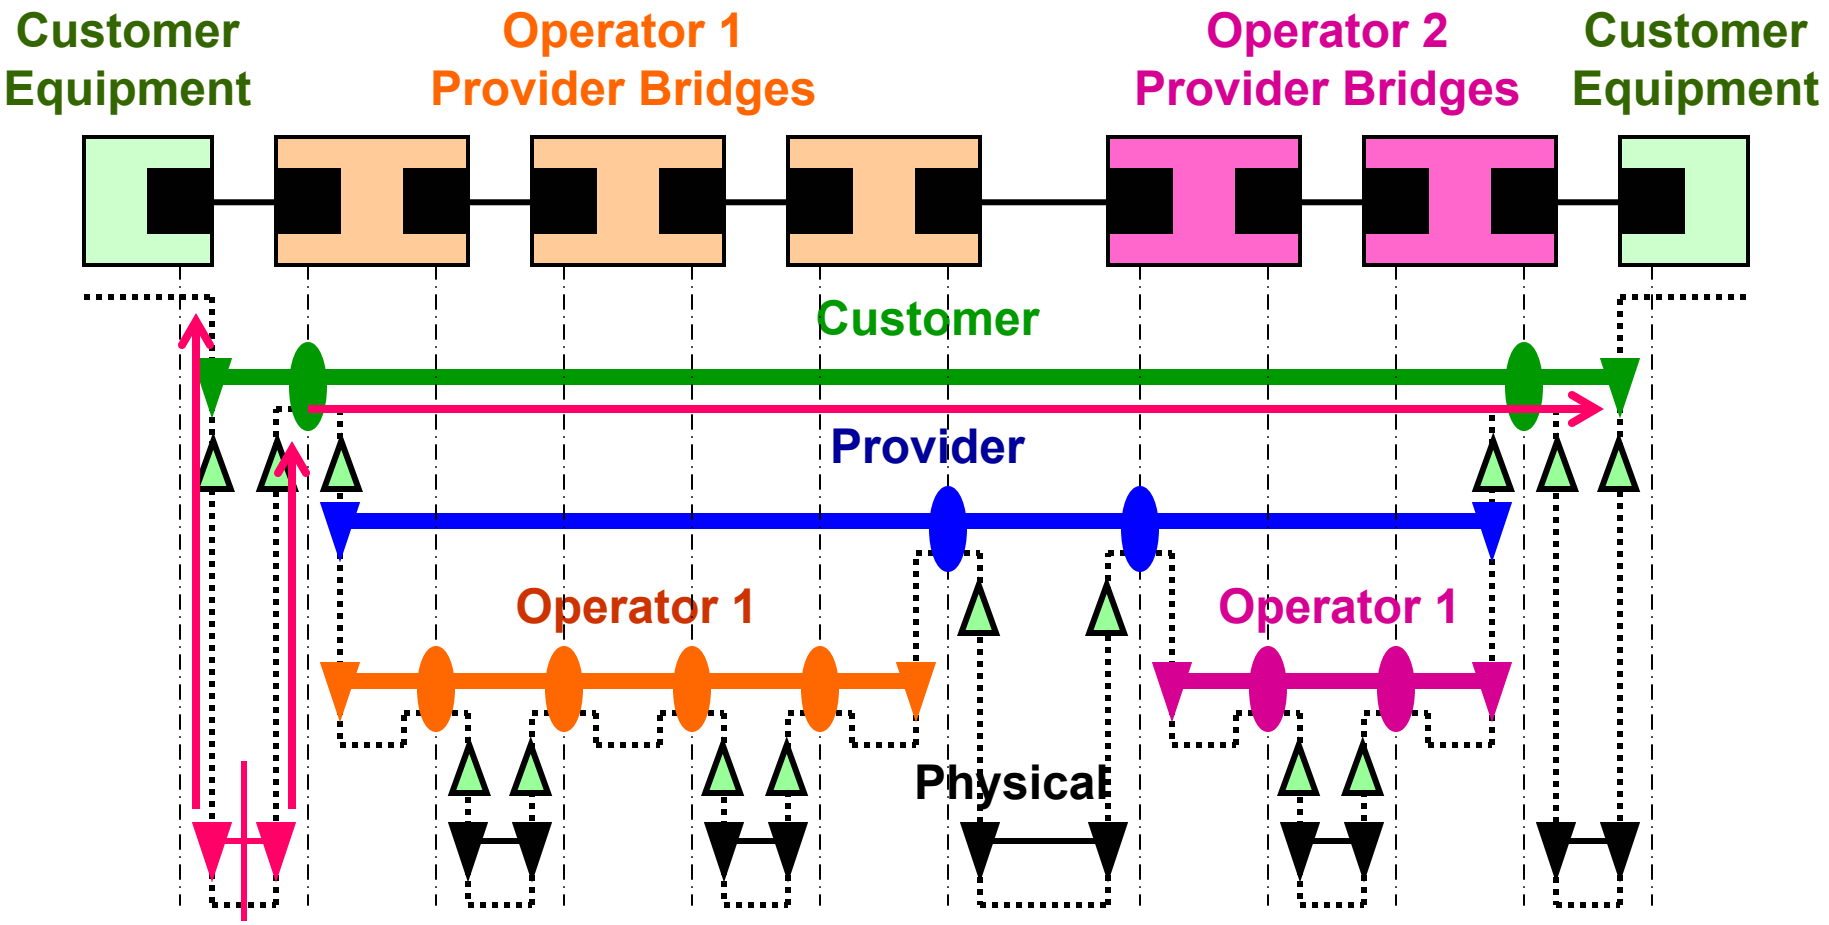

- But, the Provider doesn't know a thing!

What happens when the Provider-Customer link fails?

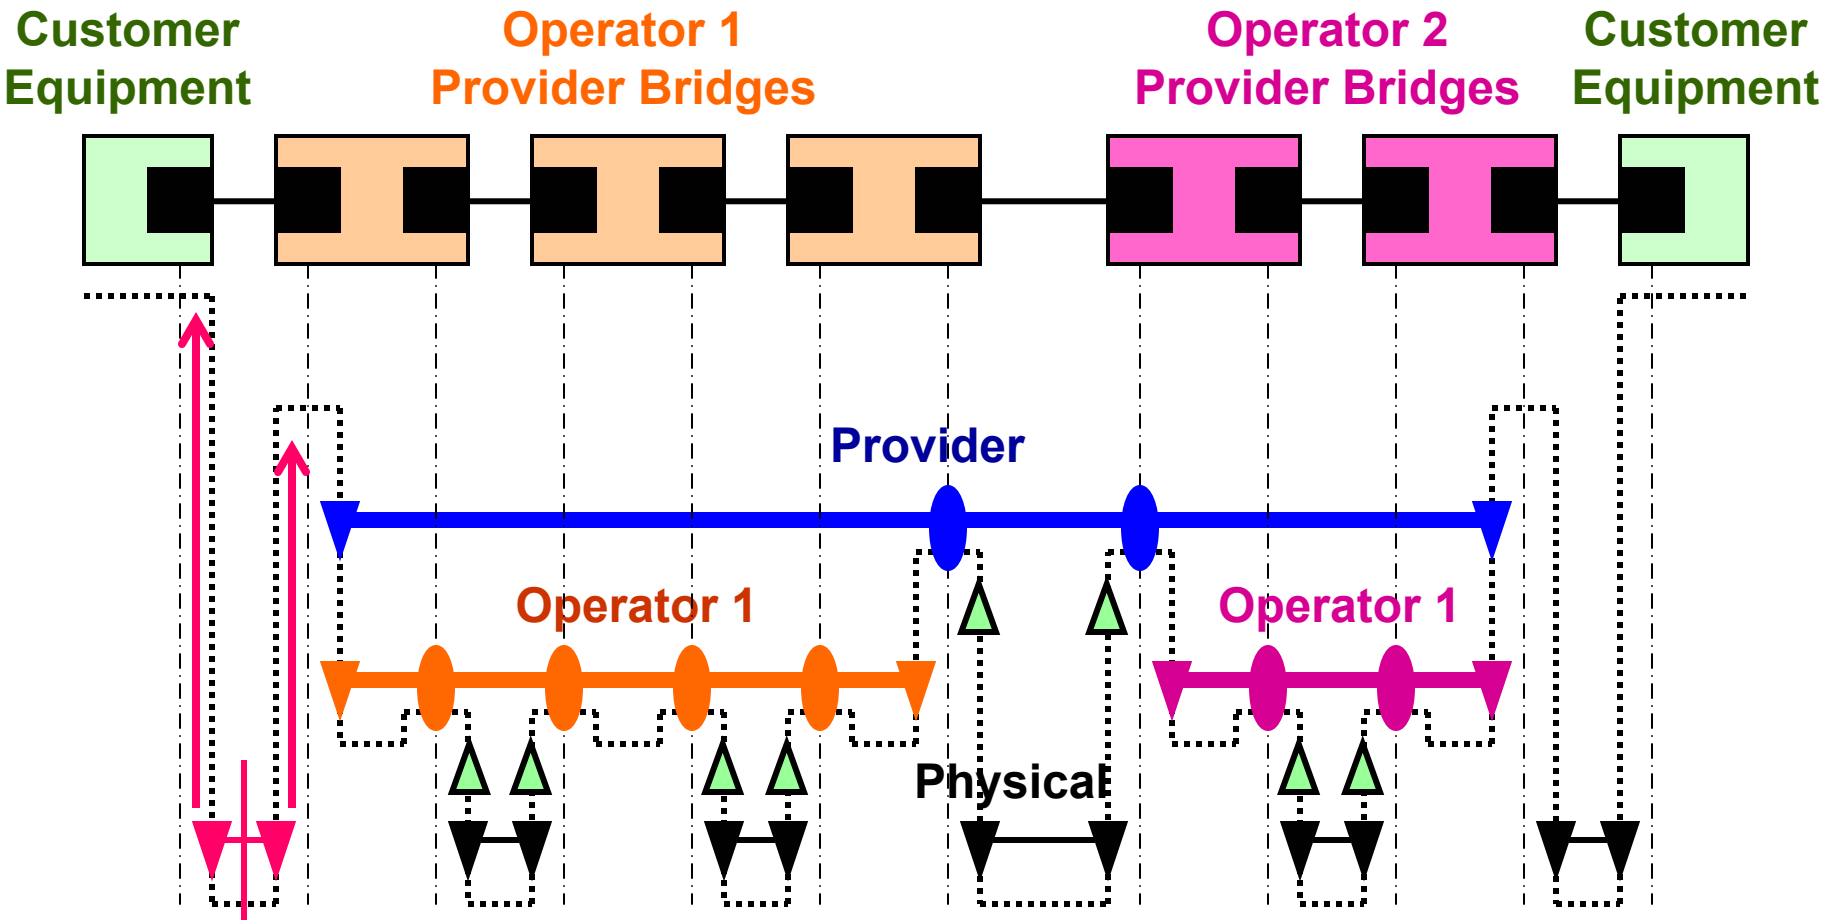

- What if there is no Customer ME Level?

Proposal:
Port State in Continuity Check Message

Cisco.com

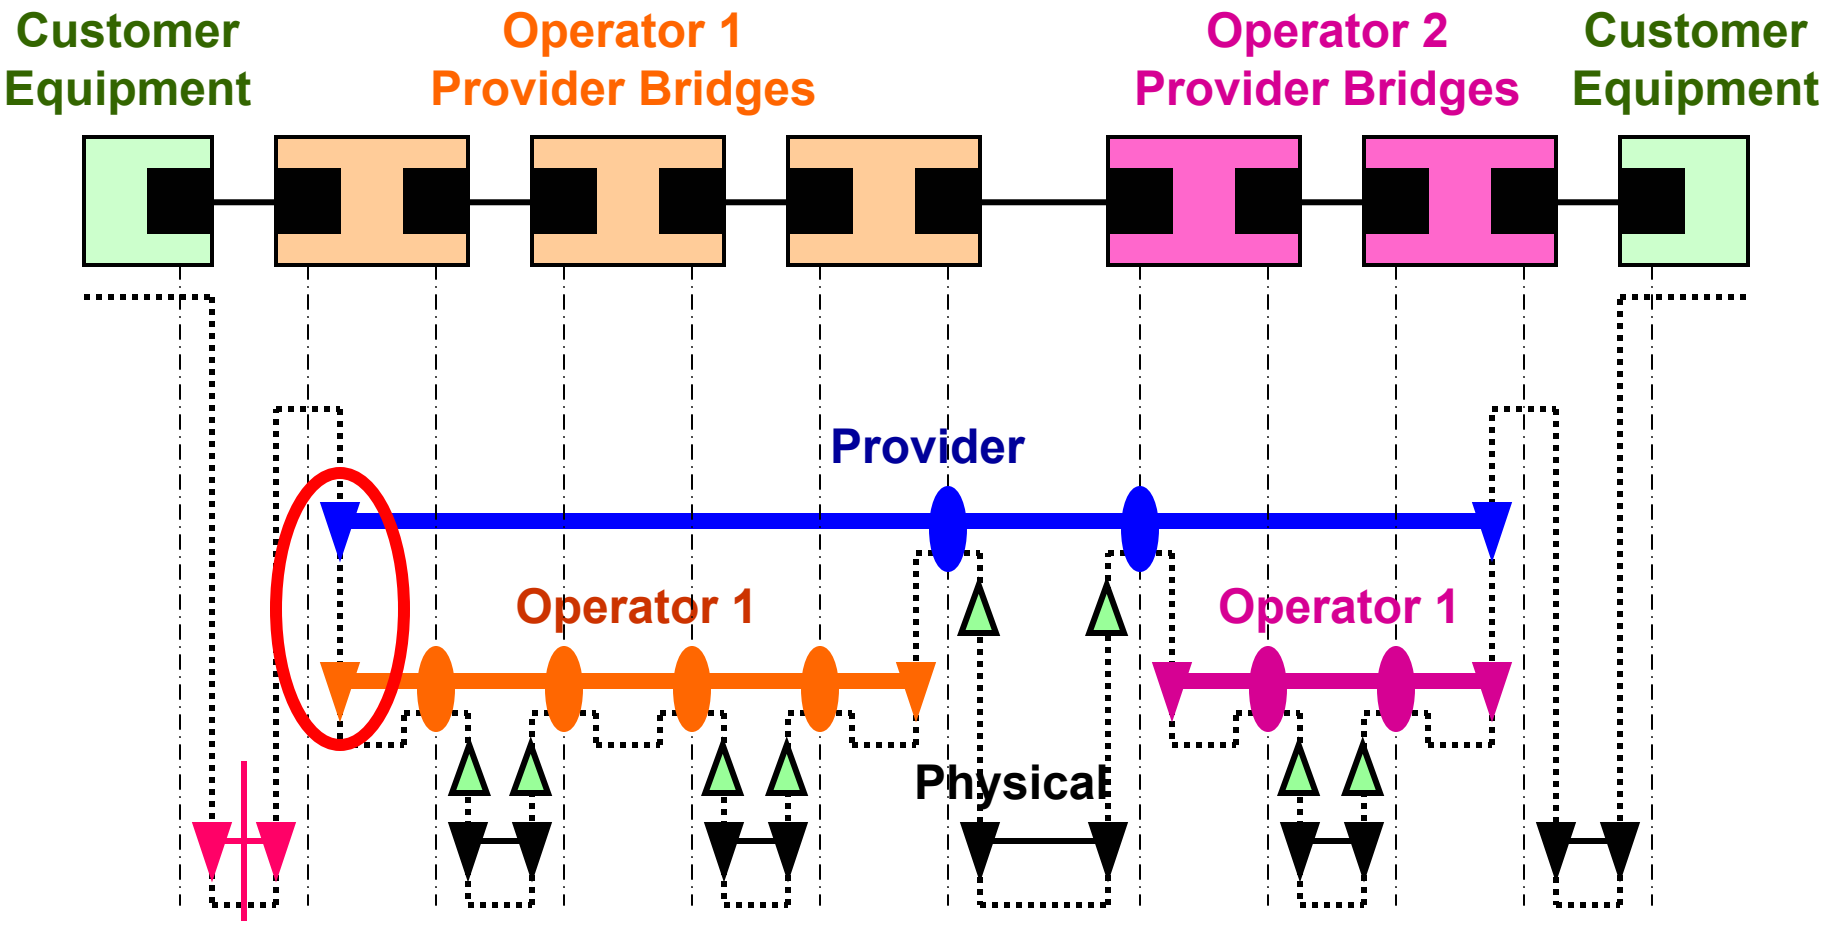

- The MEPs transmit port state in the CCMs.

Proposal:
Port State in Continuity Check Message

- Transmitting the state of the port to which the MEP is attached is not an unreasonable task for the MEP to perform.
- State would be (at least) ifOperStatus = UP or DOWN.
- States reported could include ExcessiveErrors, 3ahReportsFailure, DownForMaintenance, etc.
- AIS can still be used at Customer Level.

Network Edge Problem #2

The IEEE/ITU-T Model

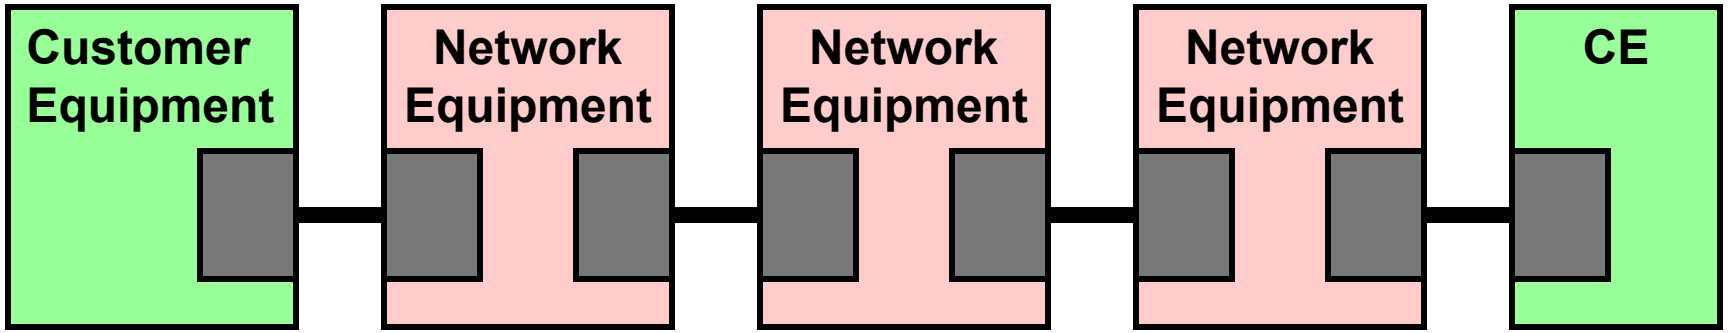

- Use 802.3ah OAM on each physical link.

The IEEE/ITU-T Model

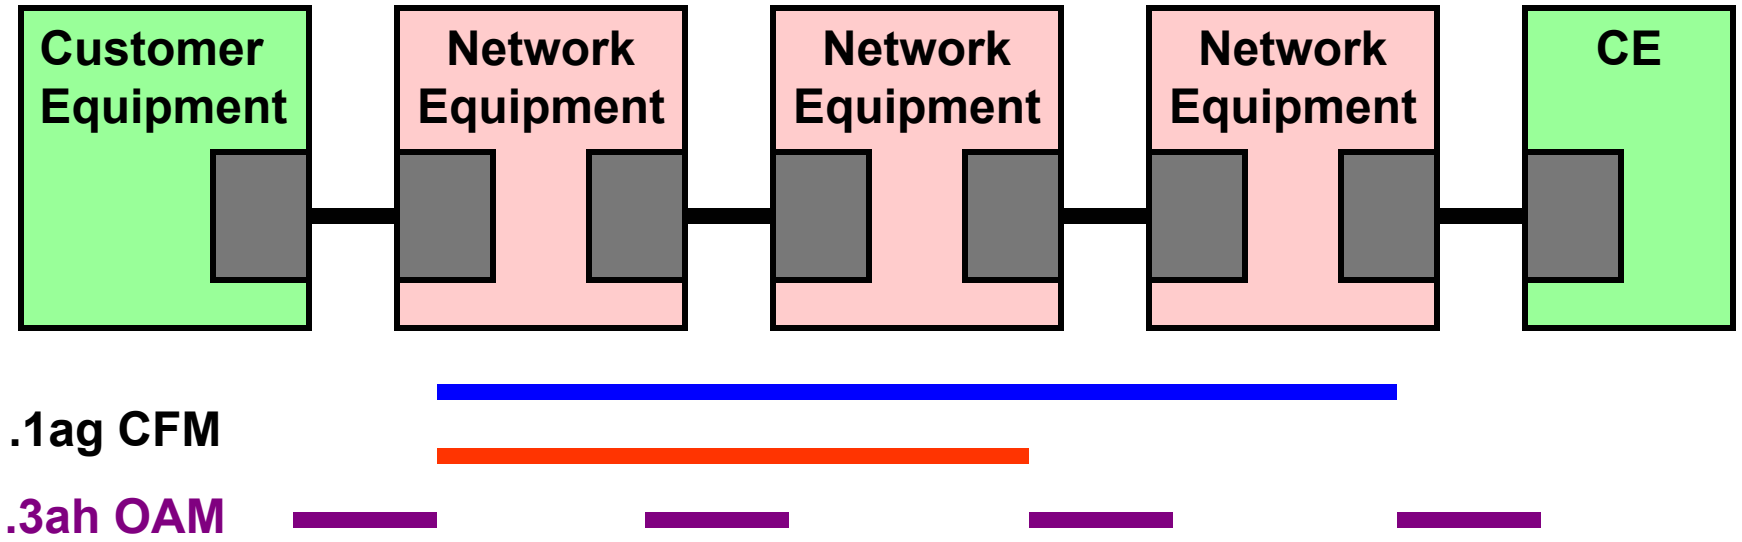

- Use 802.1ag CFM end-to-end.

Let's look at one end of the network

- The edge bridge must be the CFM Network Edge and the 802.3ah OAM Active port.
- An EB is perfectly capable of handling that role.

Let's look at one end of the network

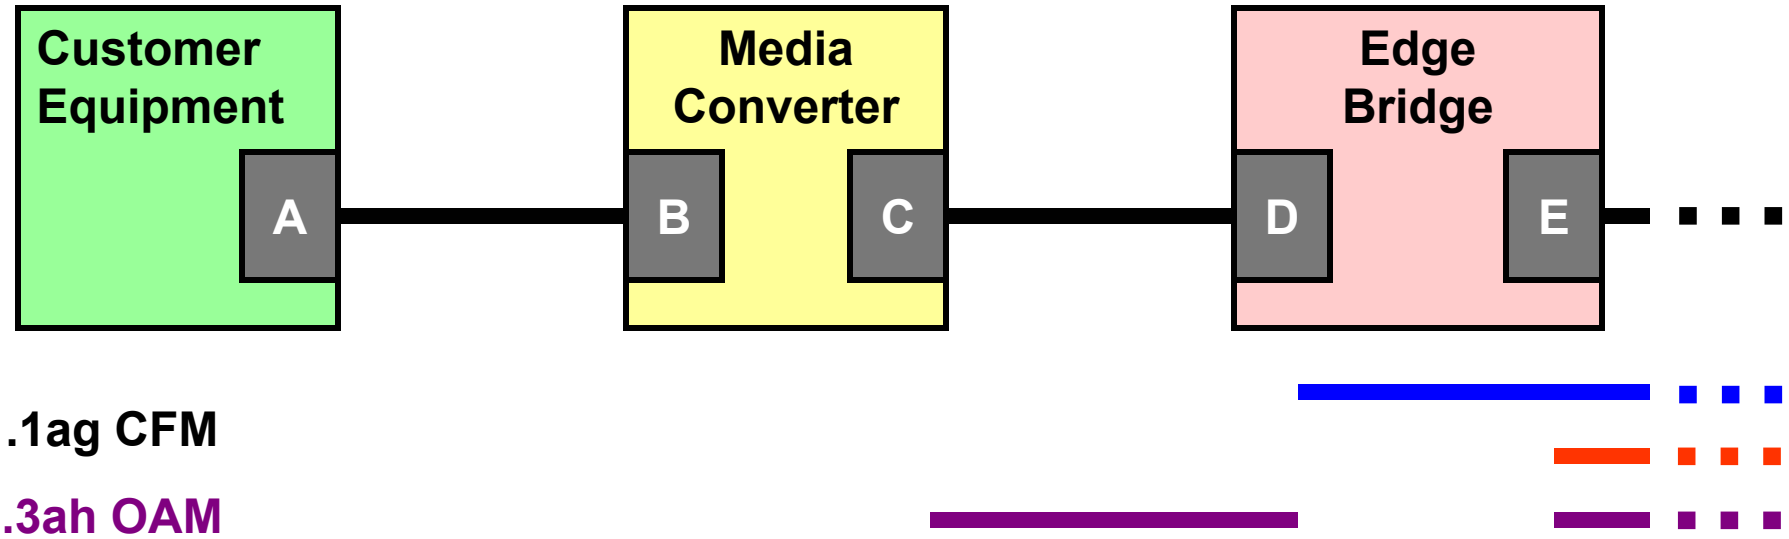

- But what happens if we drop in the Media Converter device?

Let's look at one end of the network

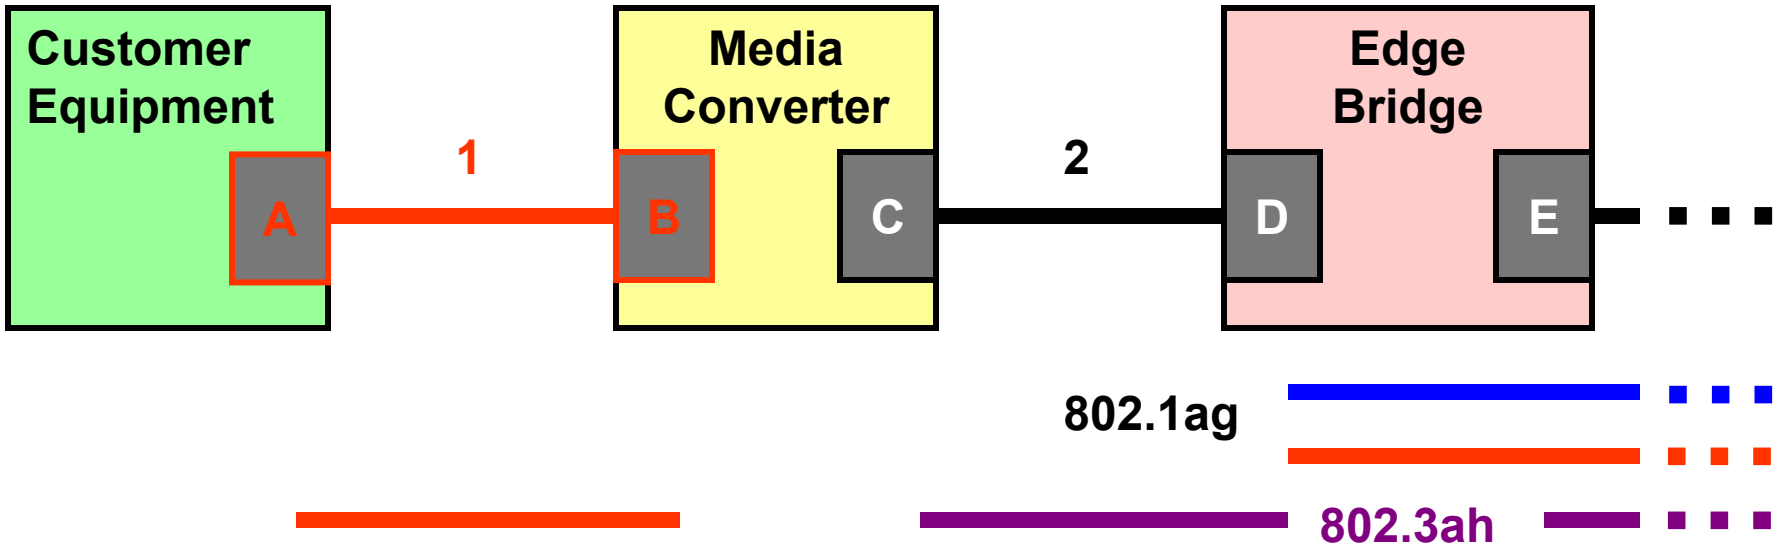

- This leaves the MC – to – Customer Link 1 unmanaged.

You cannot put the heavy CFM/OAM management points in the MC

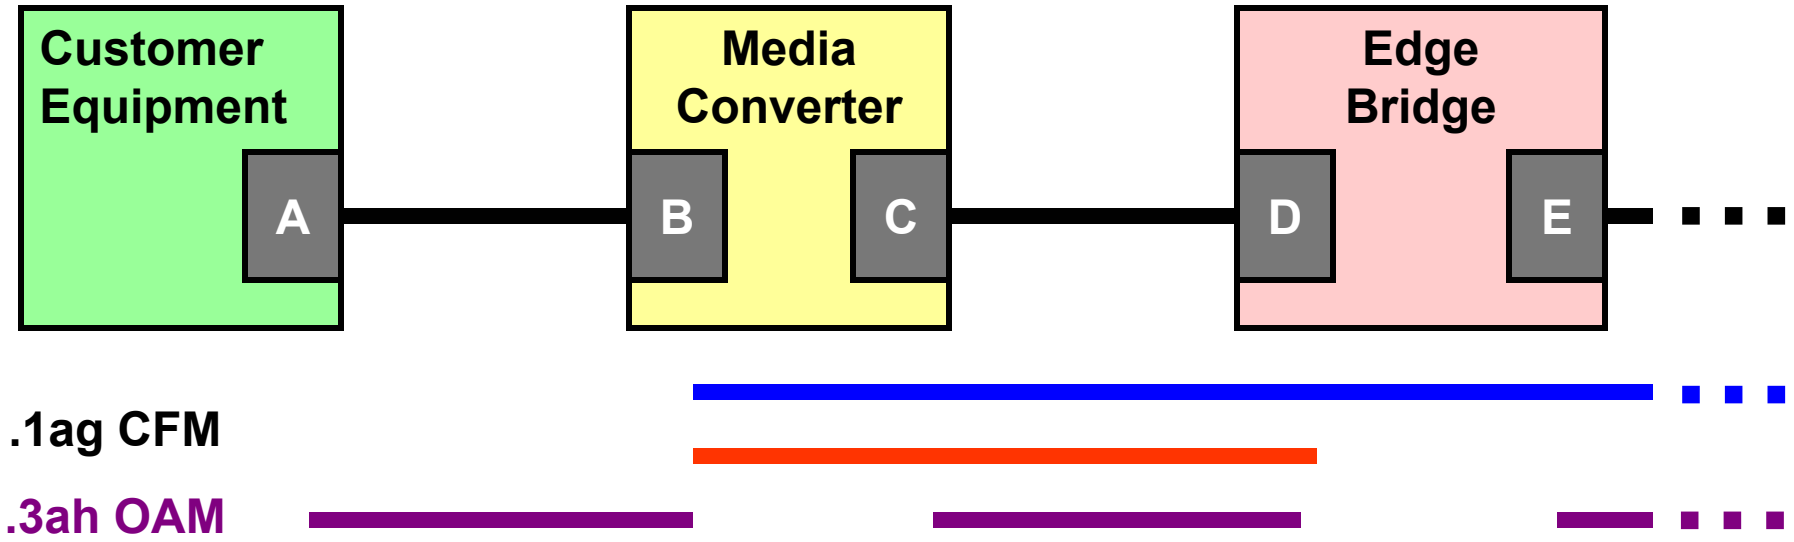

- Managing the endpoints and edges of the network requires significant configuration, intelligence, and accessibility, which are not Media Converter characteristics.

How 802.3ah OAM works

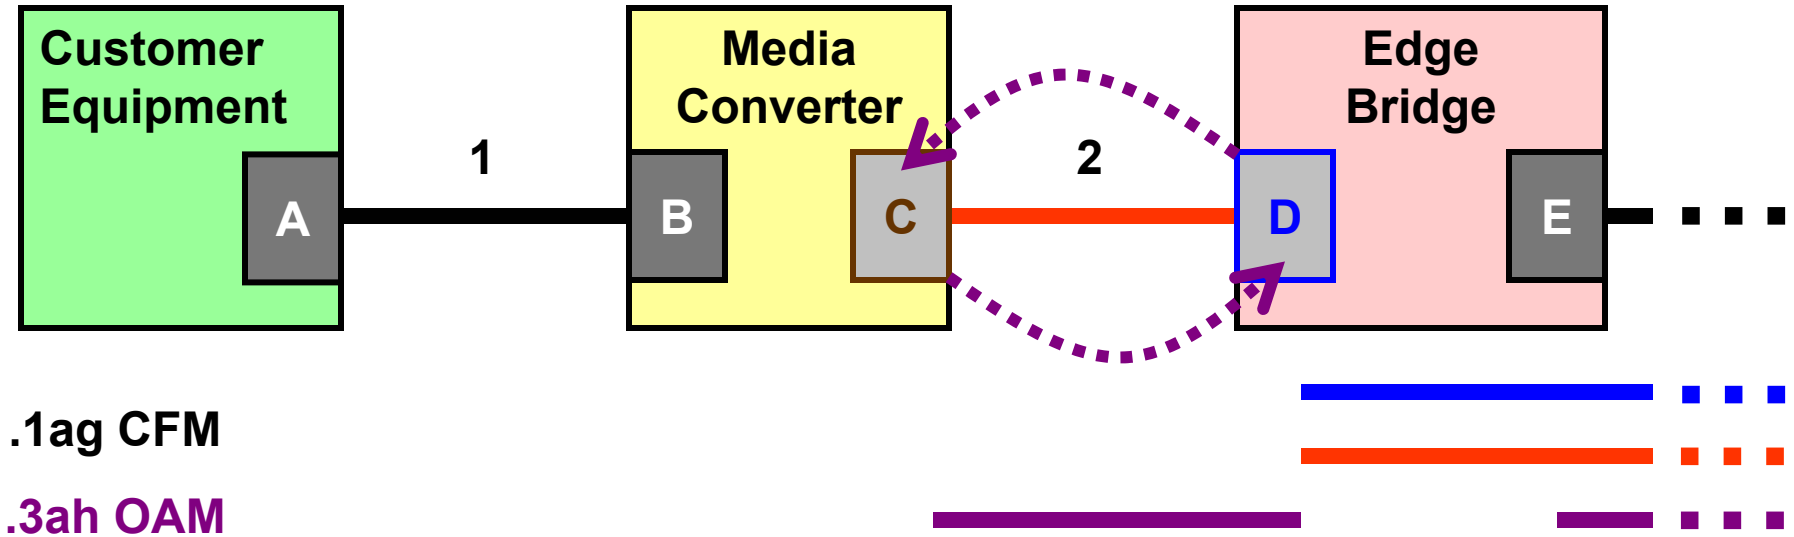

- Link 2 is managed by 802.3ah OAM.
- Port D is Active, Port C is Passive.

Part 1: The Simple Fix
Extend 802.3ah OAM

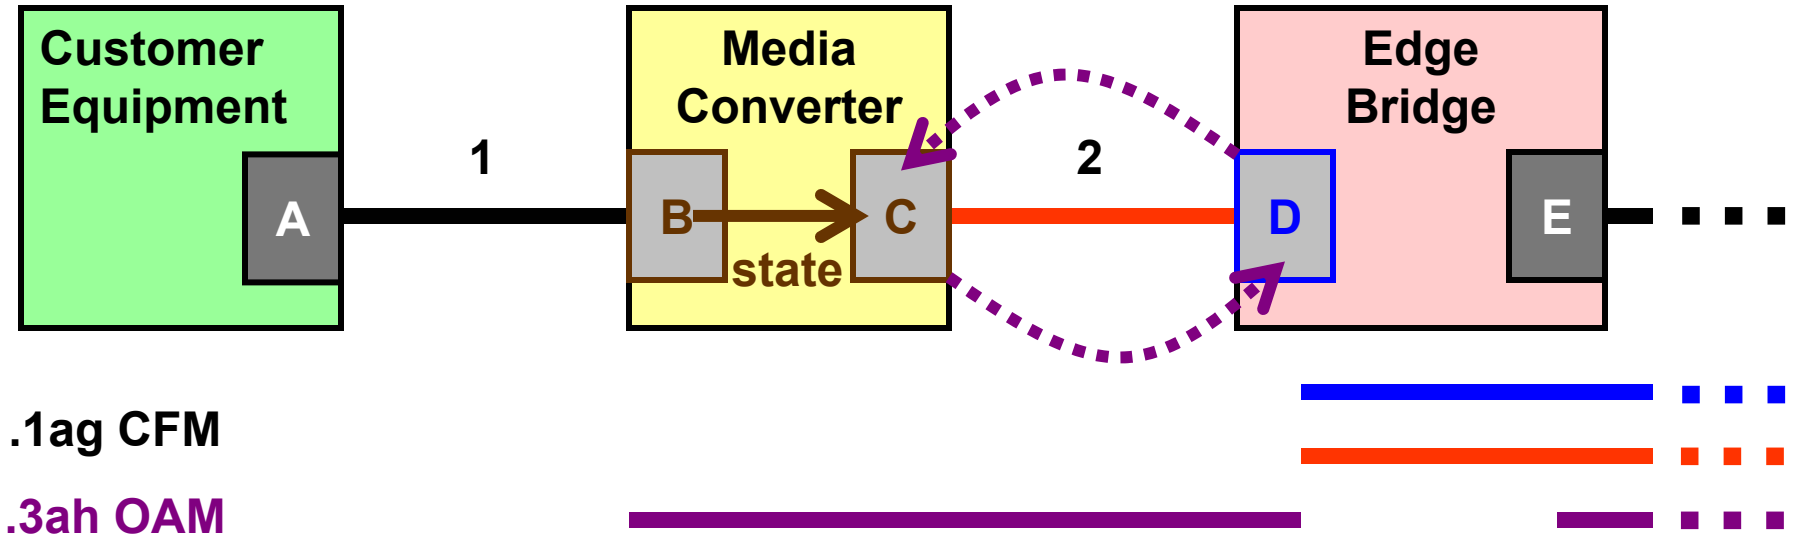

- We extend 802.3ah OAM.
- Port C reports Port B's state, as well as its own.

Part 2: What if …
Extend the reach of 802.3ah OAM

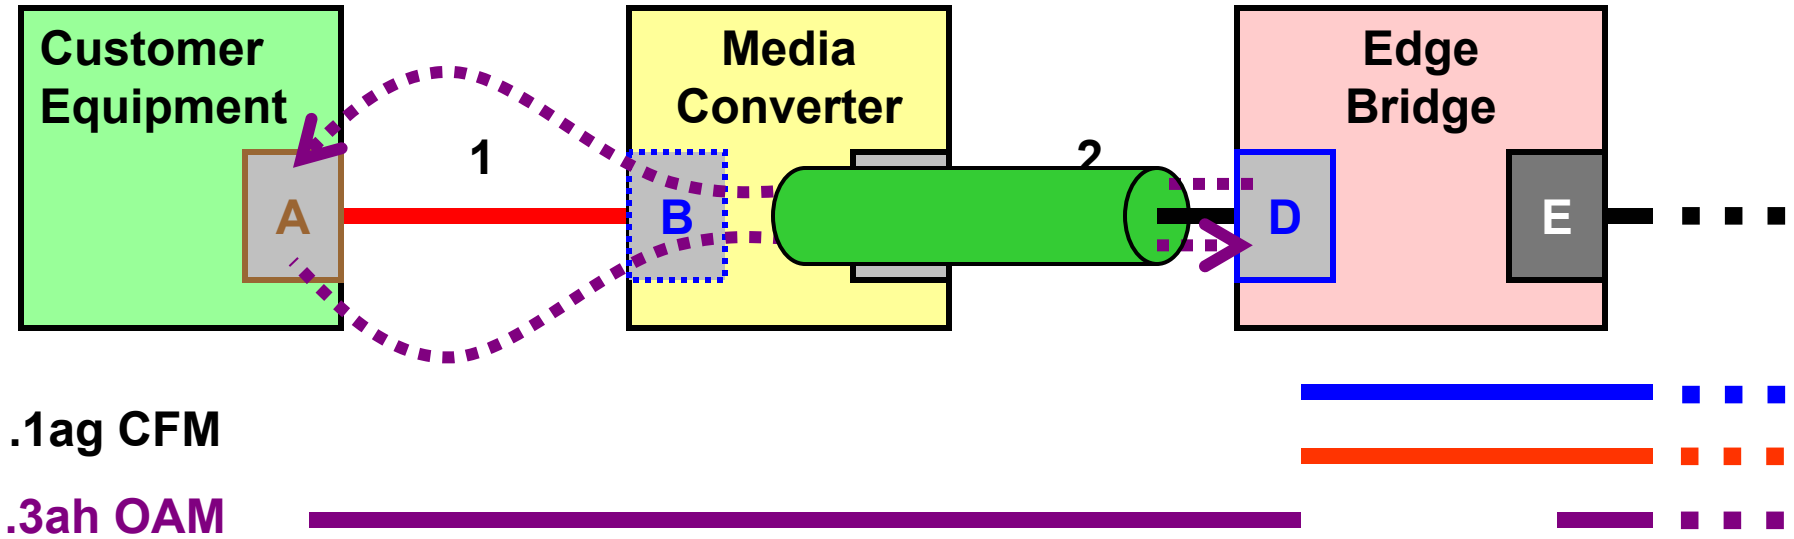

- 802.3ah OAM "from B to A" actually originate in Port D.
- The ED tunnels these packets through to Port B. Port B becomes Port D's Proxy.

What is the nature of the tunnel?

- **Plan A: Add a TLV to IEEE 802.3ah**

  Active to Passive, this TLV says, "Do not act on this OAM packet. Instead, tunnel it through to the "slave" interface (port B), remove this TLV, and emit the packet with your source MAC address.

  Passive to Active, this TLV says, "This packet came from the MAC address in the TLV, instead of the one in the Source MAC address field, which is the Active port's address."

What is the nature of the tunnel?

- **Plan B: Add a TLV to 802.1ag Loopback**

  Active to Passive, this TLV says, "take the 802.3ah OAM packet encapsulated in this TLV and transmit it from the 'slave' port."

  Passive to Active in the Loopback Reply Message, this TLV says, "The encapsulated OAM packet in this TLV came from the 'slave' port."

- **(This has problems with frame size, and may present a security hole.)**

And guess what?

- IEEE 802.1 could add this TLV to 802.3ah!
- Perhaps as part of the 802.1ag effort.
- Perhaps as part of the Media Converter effort.

Network Edge Problem #3

CFM in the Media Converter?

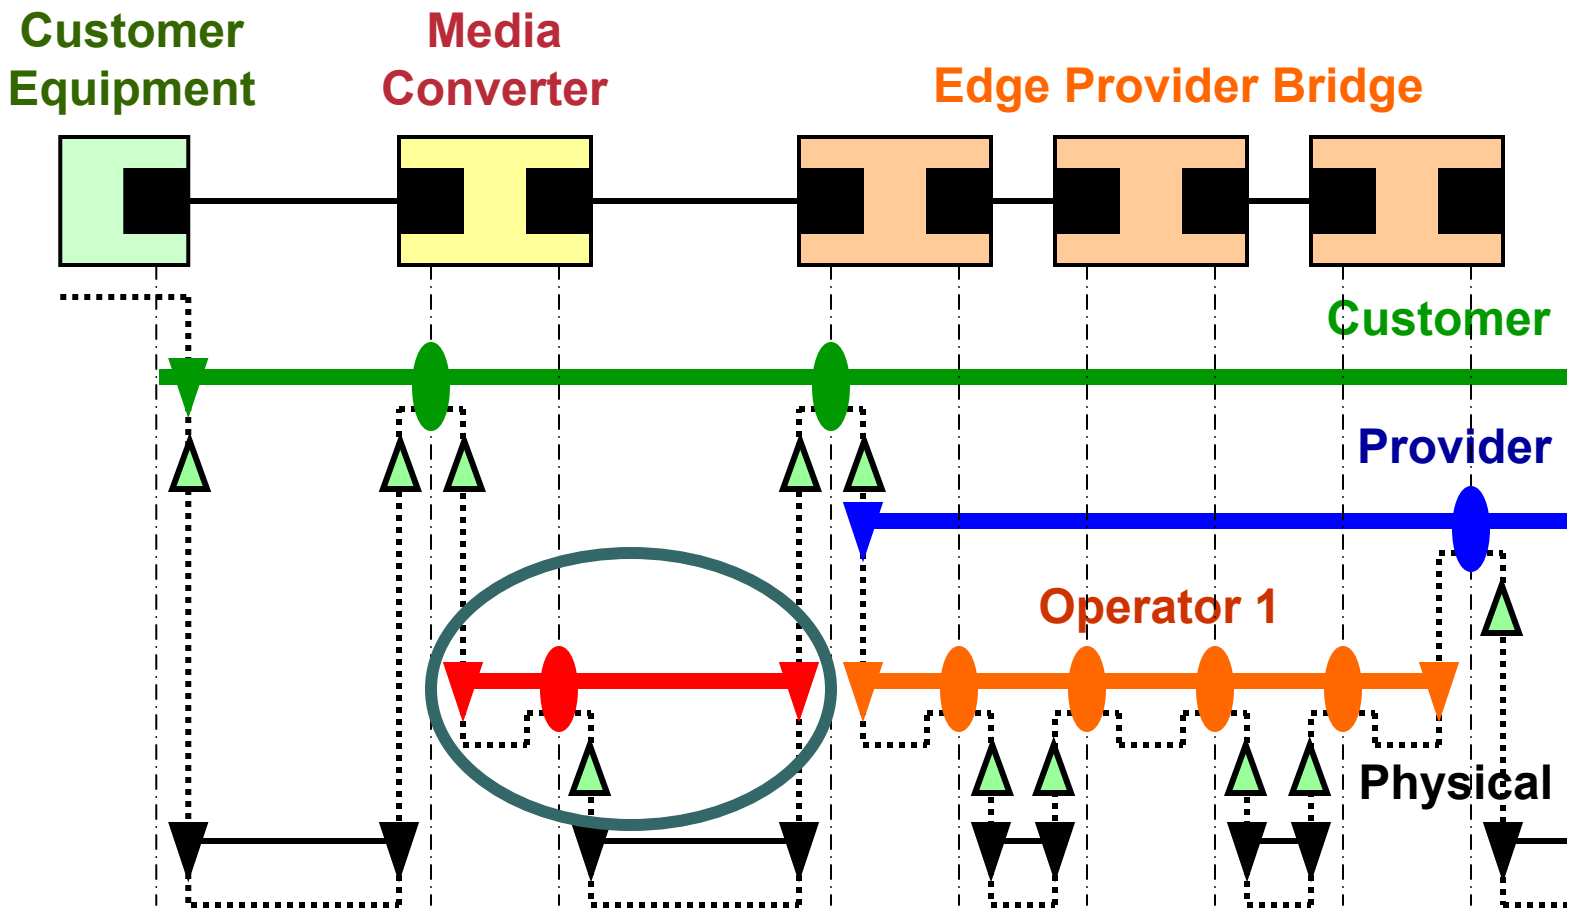

- You could add a new Operator-level Maintenance Entity.

CFM in the Media Converter?

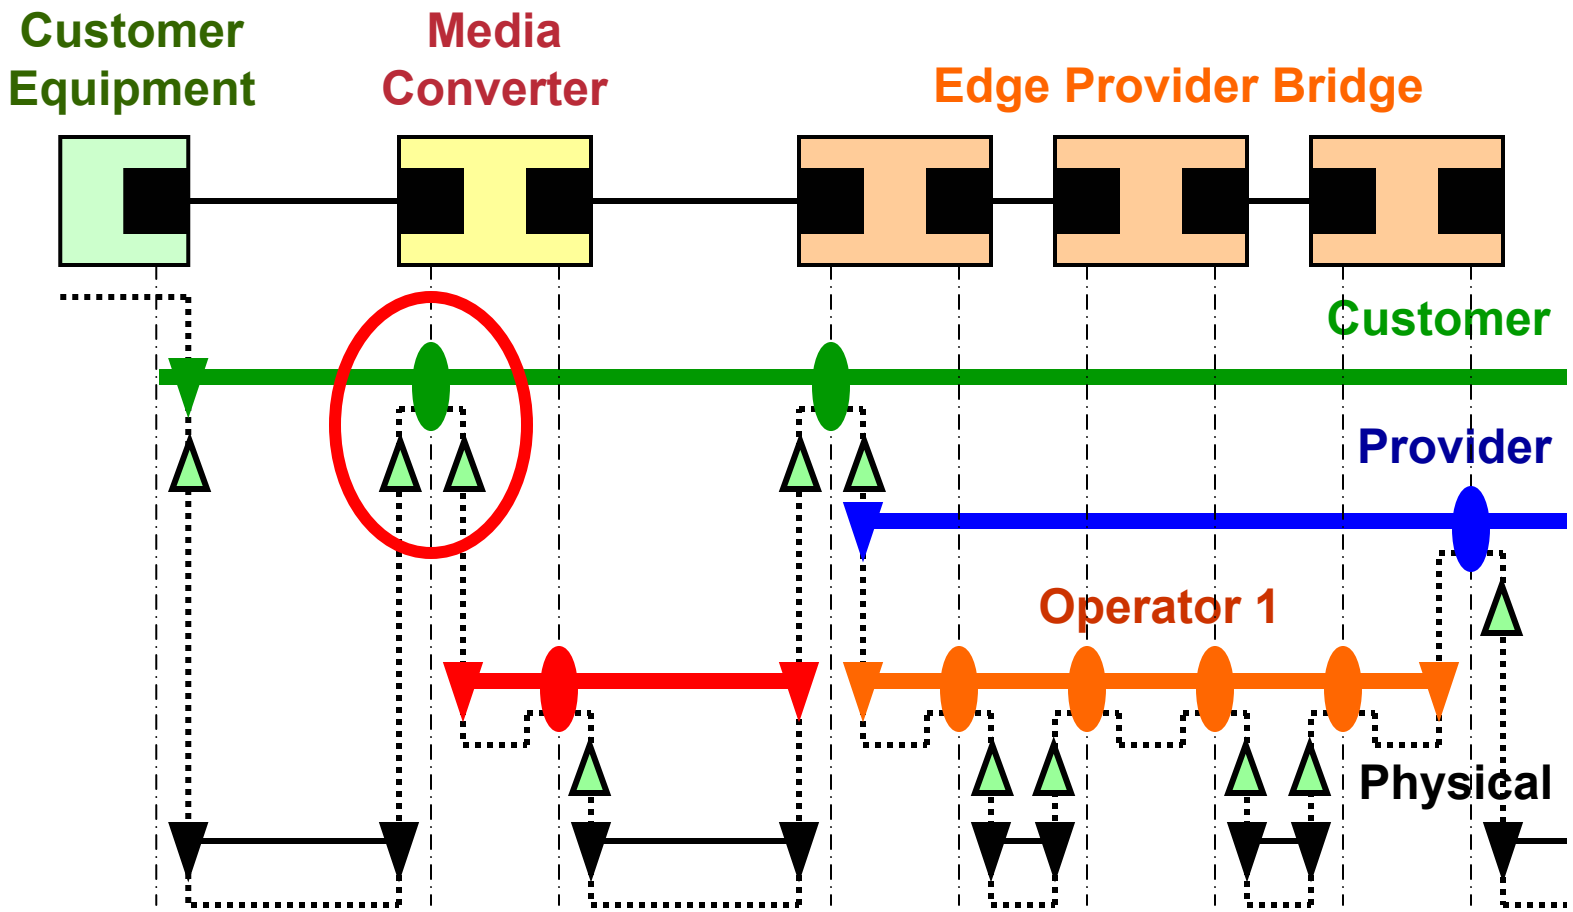

- But that requires C-VLAN awareness.

CFM between CE and Edge Bridge?

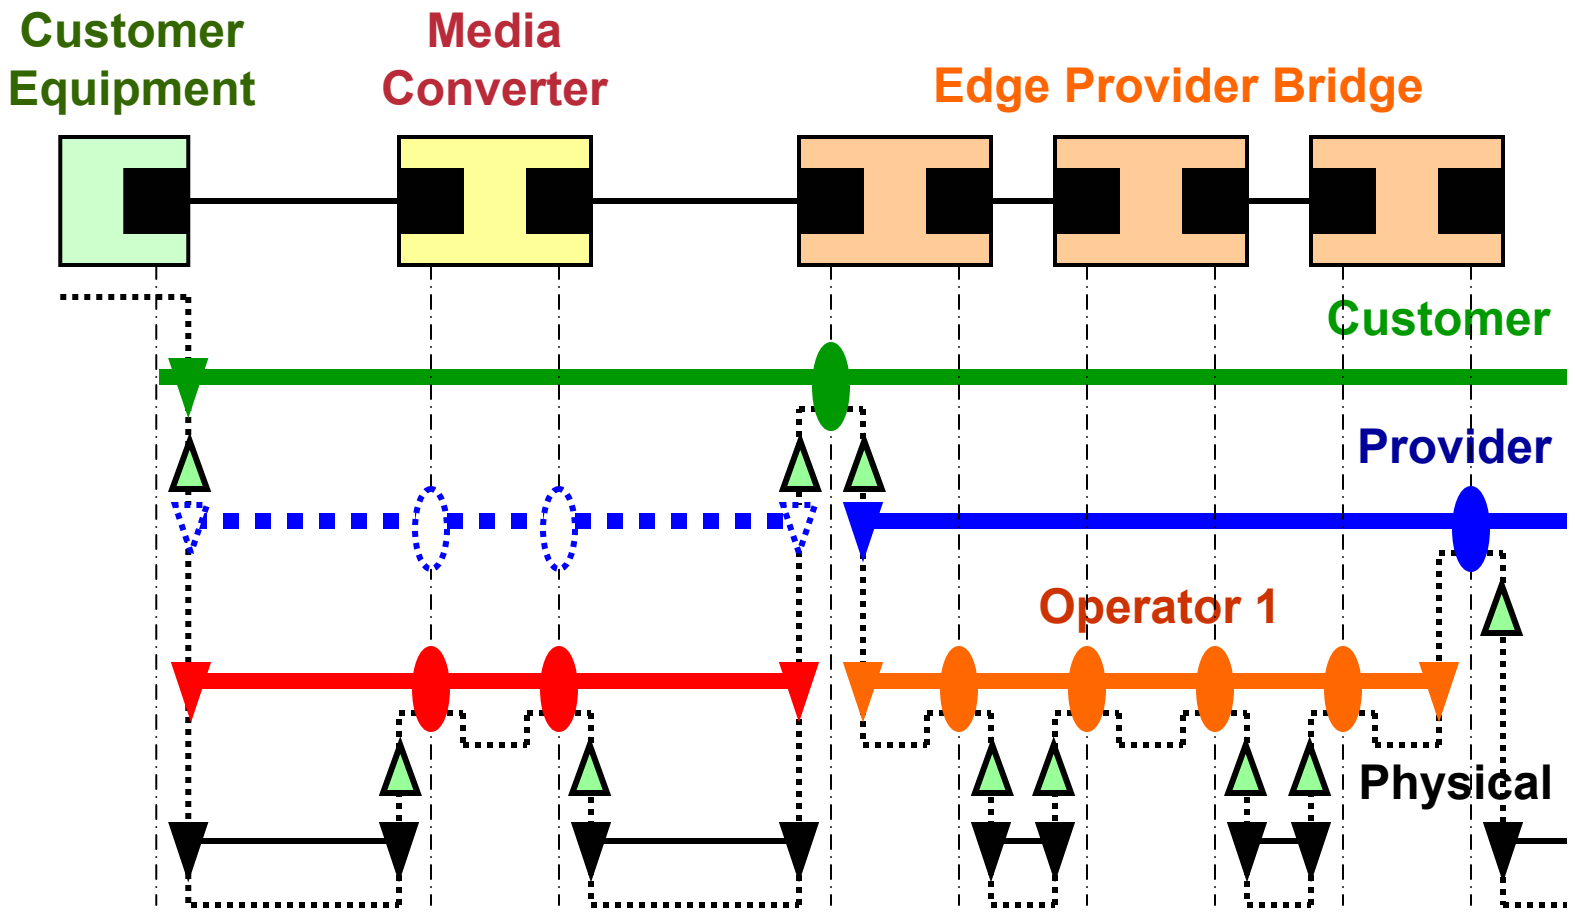

- **Operator** (or **Provider**) level ME.

CE-Provider Maintenance Entity

Cisco.com

- This allows the Media Converter to have only Maintenance Intermediate Points, which are much easier to implement and/or configure than Maintenance End Points.
- This requires (permits) the Customer to participate in the maintenance of the Media Converter.
- This may be appropriate for some Providers and/or Customers.

Other CFM Packet Types

Other CFM Packet Types?

Discovery

- **Continuity Check provides MEP discovery.**

  Loopback to a multicast destination requires jitter timers to avoid implosions, and hence is no faster than Continuity Check.

  IEEE 802.2 defined multicast Test and XID functions, which are the primary reason that these functions are disabled by default on most interfaces.

Other CFM Packet Types?

Discovery

- **MIP discovery may be done with Traceroute, one MEP at a time.**
- **A Traceroute for a multicast MAC target might trace an entire network in one large operation.**

  That is a heavy procedure.

  But, *N* Traceroutes for *N* MEPs would be even heavier.

Other CFM Packet Types?

Remote Defect Indication

- Familiar (to some) from use in other technologies.
- Ethernet does not support unidirectional links.
- RDI is meaningful only when unidirectional links are possible.

Other CFM Packet Types?

Performance Management

- It should be possible to piggyback additional TLVs on Continuity Check and Loopback in order to measure performance.
- An additional PM OpCode may be desired to mark packets that require hardware assistance (e.g. timestamps) to produce meaningful data.

Other CFM Packet Types?

Intrusive Loopback

- **Purpose is wire-speed bandwidth and forwarding verification.**
- **Similar to 802.3ah OAM Loopback, but loopback function would swap source and destination MAC addresses, and would operate on a per-Service, instead of a per-MAC, basis.**
- **Requires “hardware” implementation to be useful.**
- **Is extremely dangerous in a shared medium service.**

Summary

Summary

- 3-View would be useful in P802.1ag.
- AIS should be incorporated as a CFM PDU type.
- A Port State TLV should be incorporated into the Continuity Check Message.
- 802.1 should define the 802.3ah OAM “tunnel” TLV.
- The CE-Provider Maintenance Entity should be shown in P802.1ag.
- “Discovery”, “Remote Defect Indication” and “Intrusive Loopback” should not be a part of P802.1ag (or Y.17ethoam). “Performance Management” is questionable.


CISCO SYSTEMS