Connectivity Fault Management

Norman Finn

DISCLAIMER

- IEEE P802.1ag is a work in progress.
- Some of the information in this slide deck reflects consensus among IEEE 802.1 and ITU-T Q.3/13 participants.
- Some of the information in this slide deck reflects only the opinions of the author.
- Until P802.1ag and/or Y.17ethoam are approved, there is no reliable means for distinguishing opinion from consensus.

What is Metro Ethernet Connectivity Fault Management?

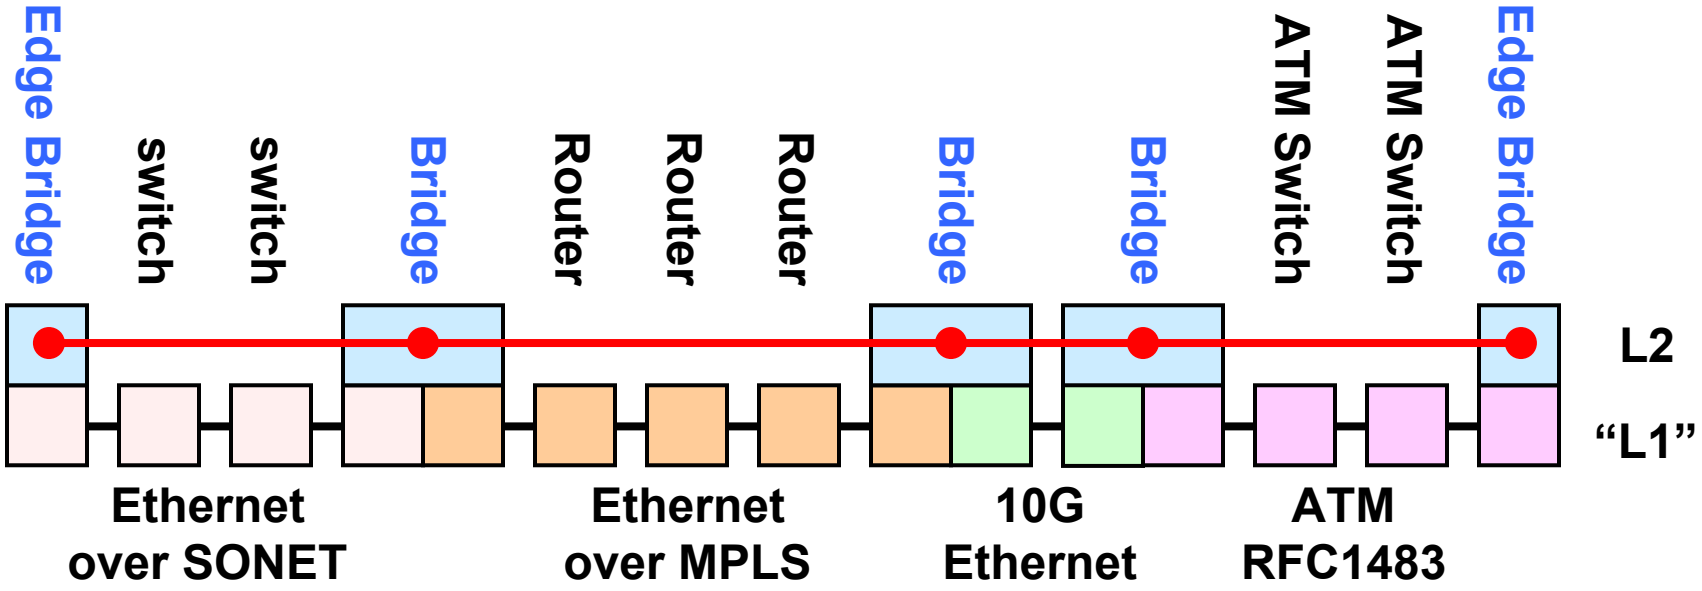

- **CFM: Standard Ethernet frames, distinguished from ordinary data frames only by destination MAC address and/or EtherType, and seen, relayed, and/or terminated by Provider Bridges.**

IEEE 802.3ah OAM versus 802.1ag CFM = ITU-T Q.3/13 EthOAM

IEEE 802.3ah OAM	Connectivity Fault Management
Operates on physical link only. Cannot pass through a bridge.	May be per-service or per-wire. Passes “end-to-end” through bridges.
Discovery, Variable request & response, Event notification, Information, Loopback mode.	Connectivity Verification, Traceroute, Ping, (Alarm Suppression?).
Single instantiation per physical link.	Multiple instances operating at multiple levels simultaneously.
Created by one committee.	Joint effort by IEEE 802.1, ITU-T.
Standard approved by IEEE 802.	Drafting in progress.
Both are TLV based, both will (likely) allow vendor-specific TLVs.	

The OAM Environment

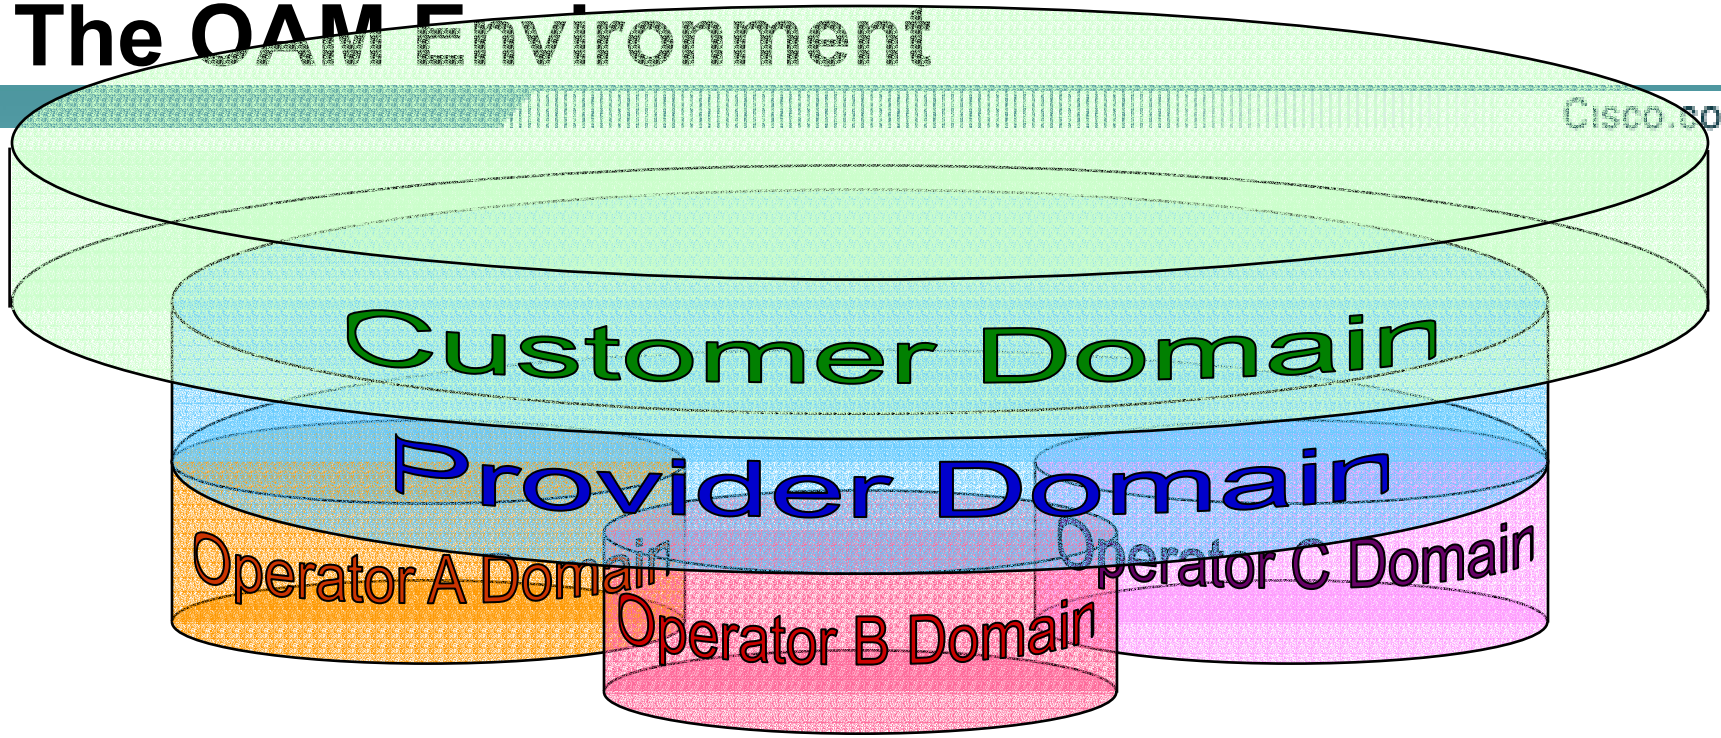

- **Customer contracts with Provider for end-to-end service.**
- **Provider contracts with Operator(s) to provide equipment and networks.**
- **Provider and Operator(s) may or may not be the same company or same division.**

The View From Above

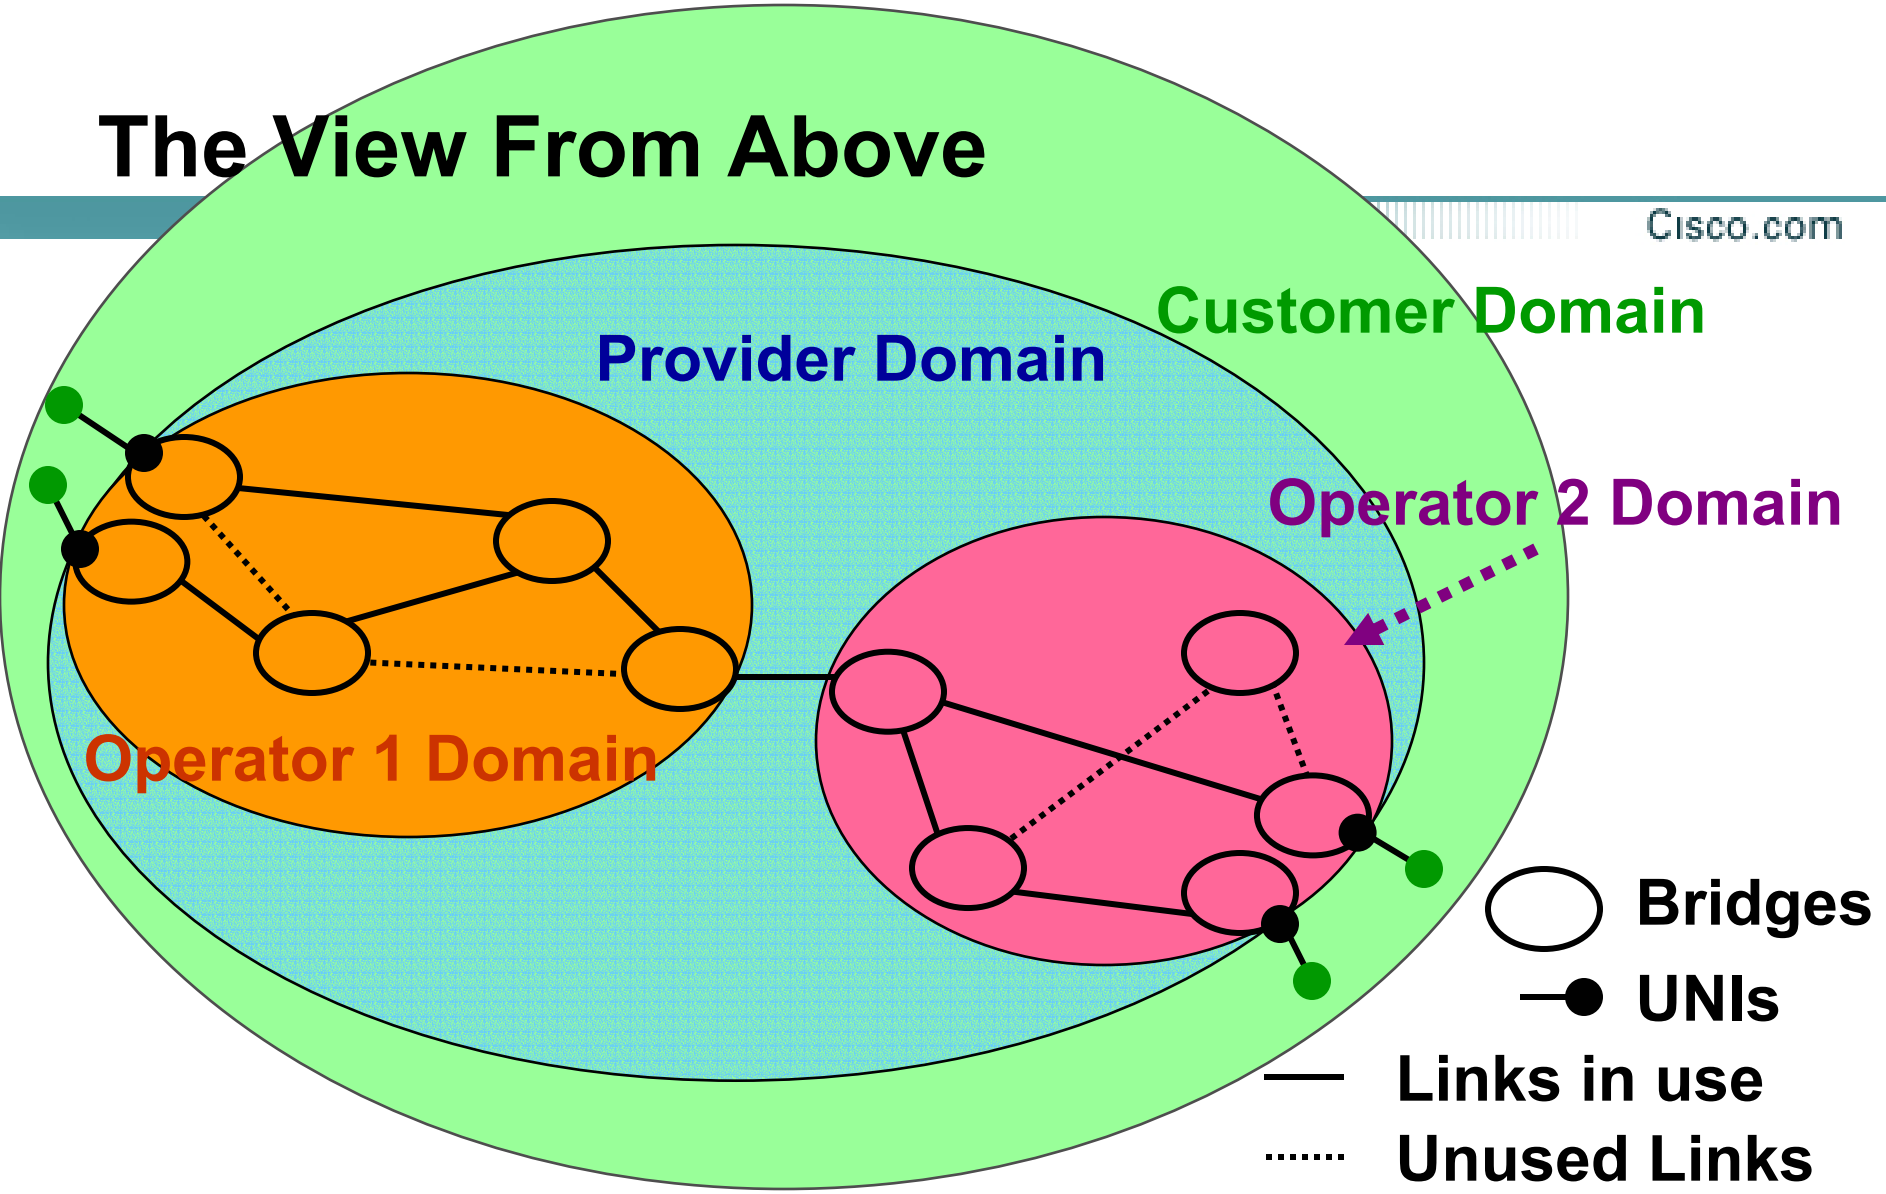

- Physical view of a network using CFM

The View From Above

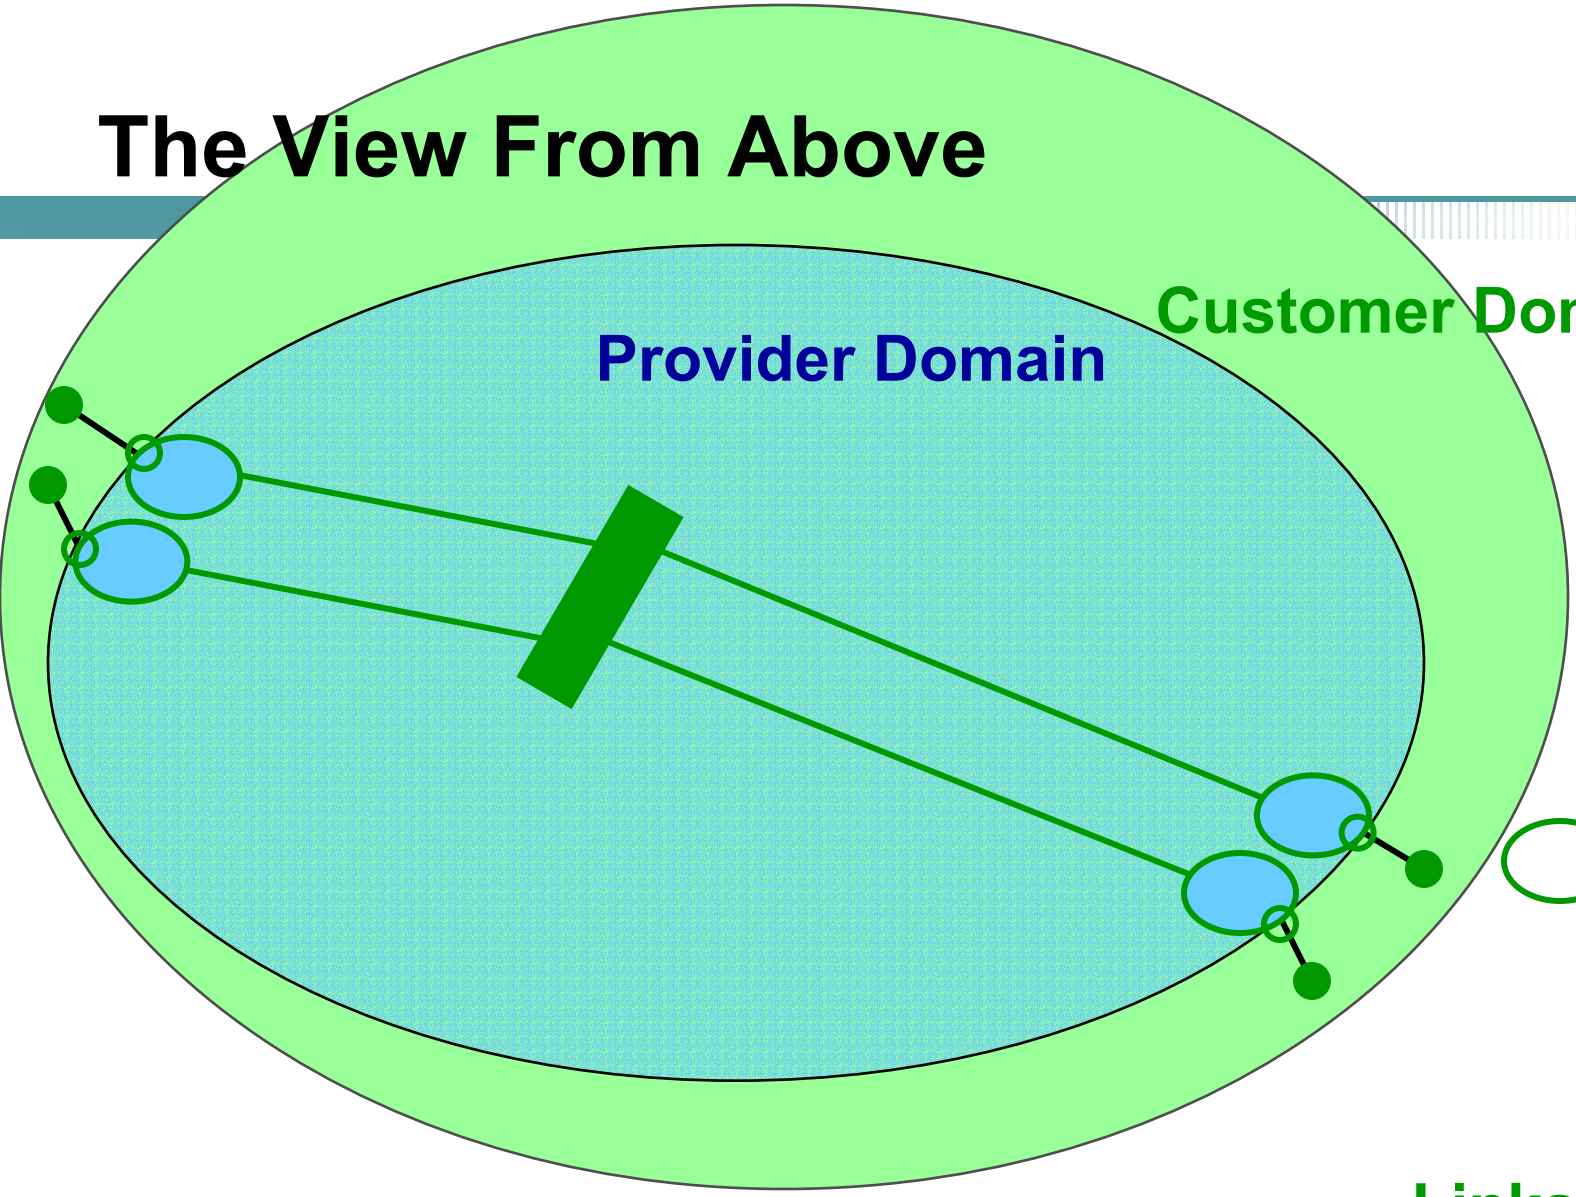

Bridges

MIPs

MEPs

Links in use

- Customer's view from above

The View From Above

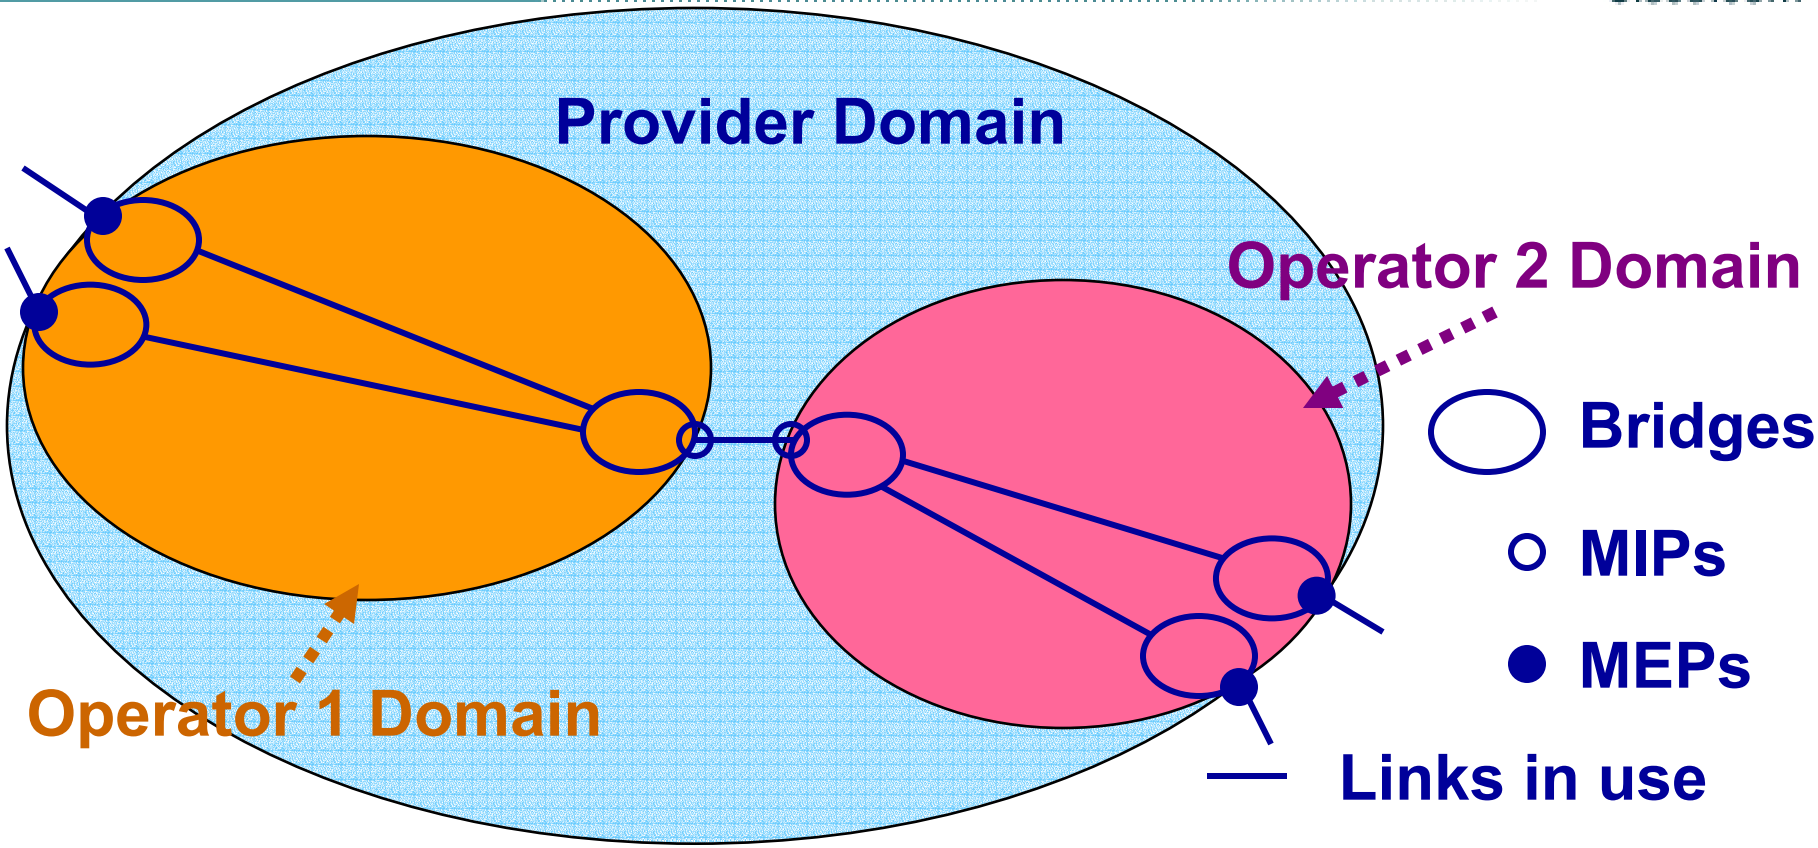

- Provider's view from above

The View From Above

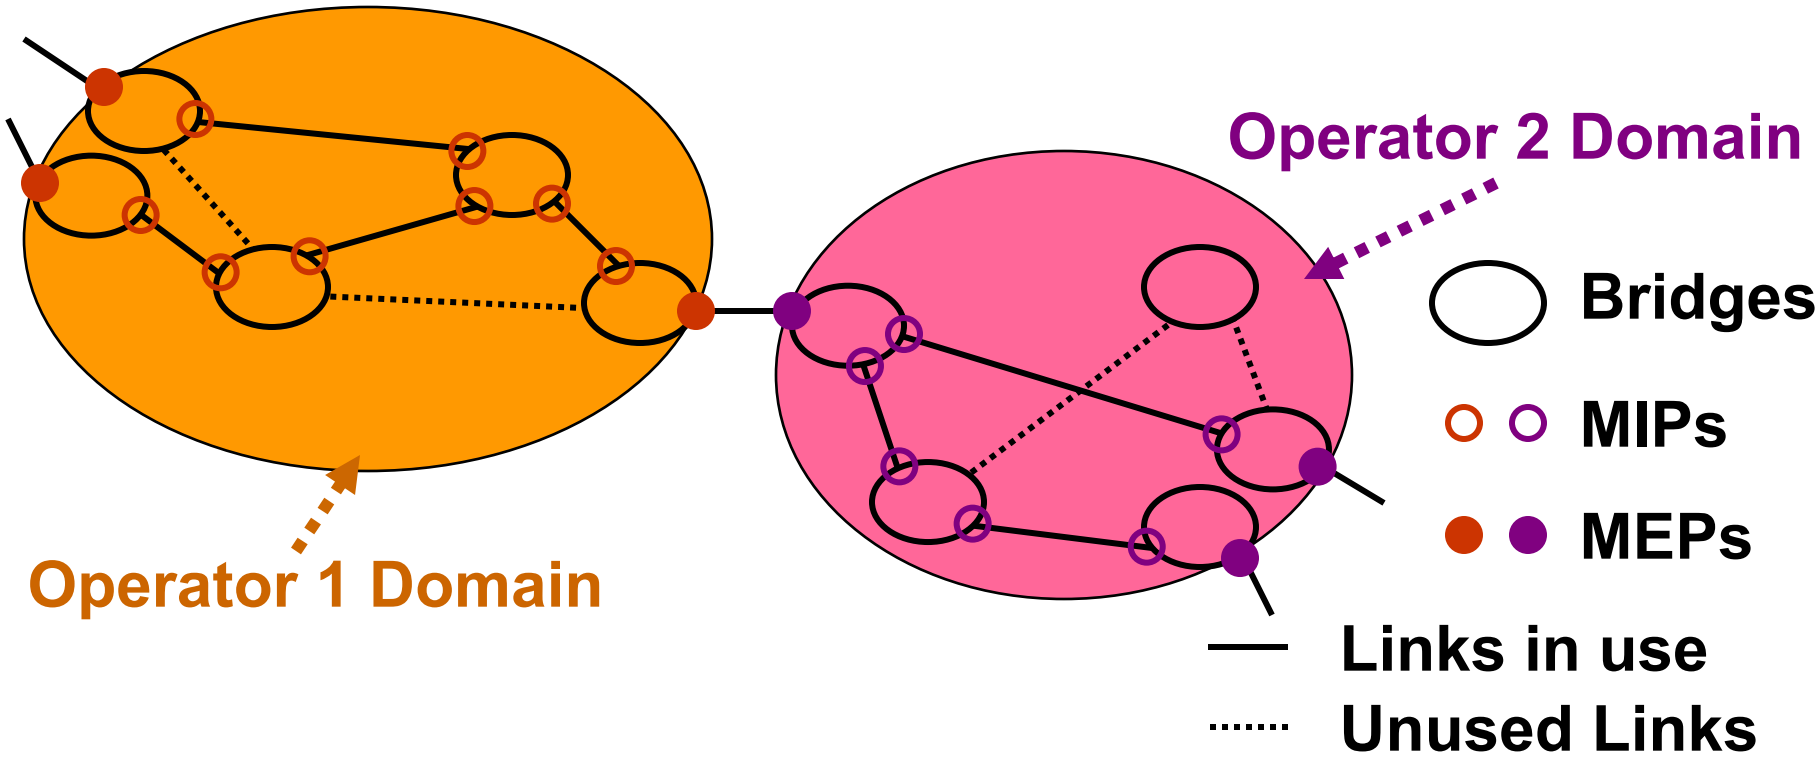

- Operators' view from above

The View From Above

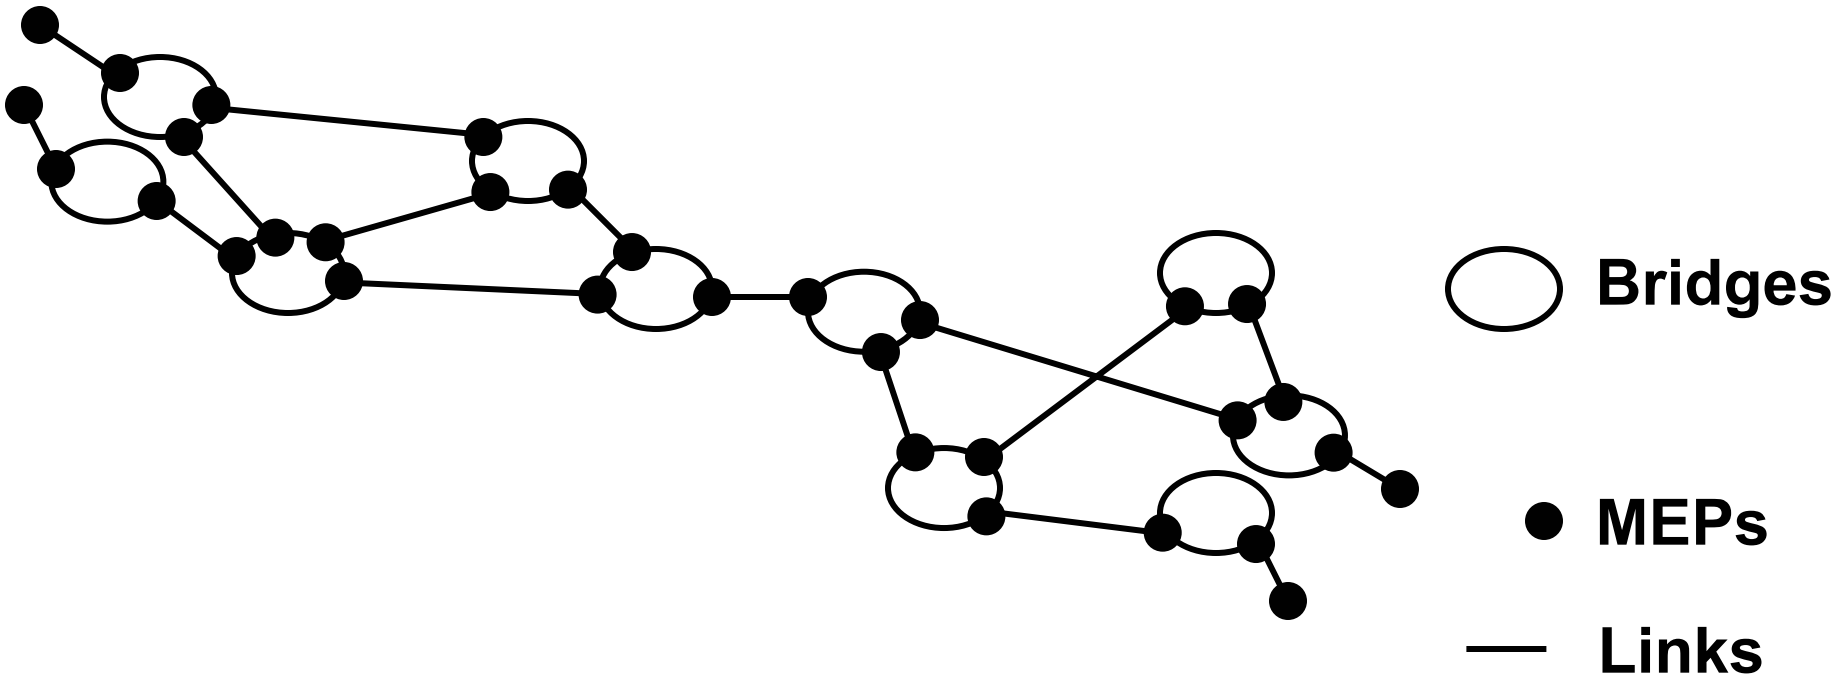

- Physical view from above

The View From Above

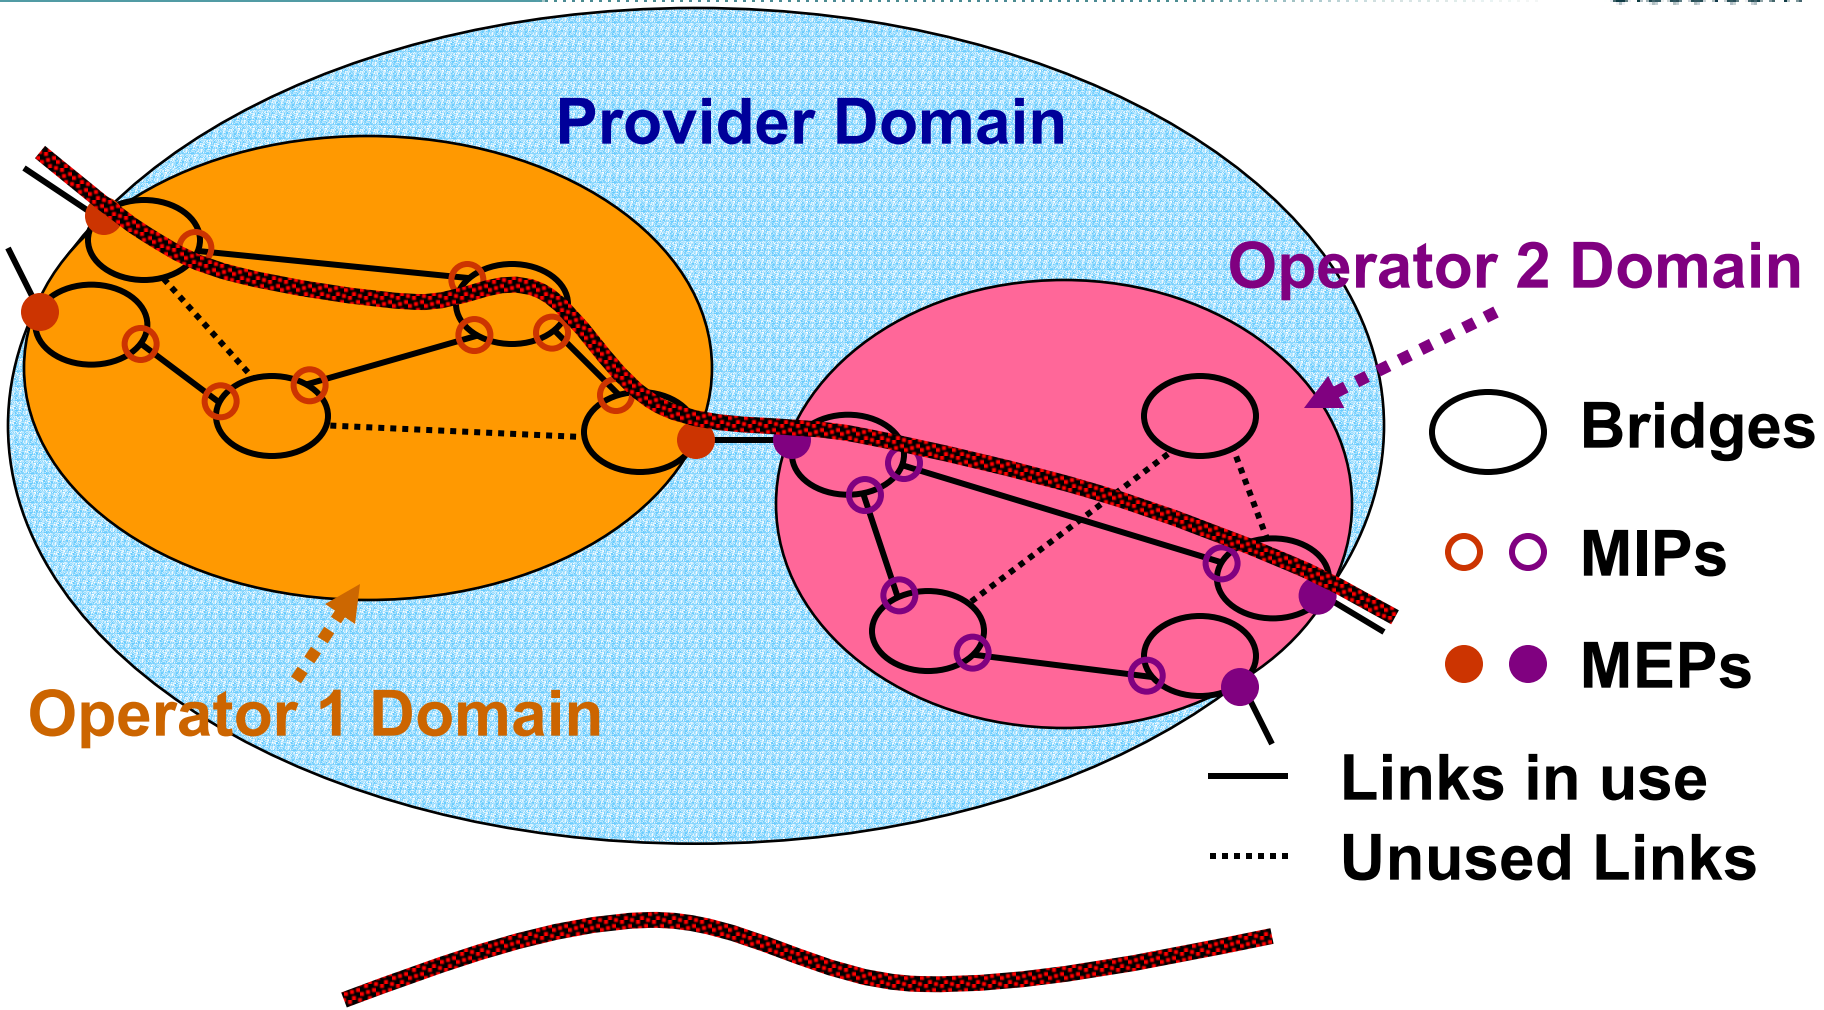

- Let's take a slice along one path …

Longitudinal View

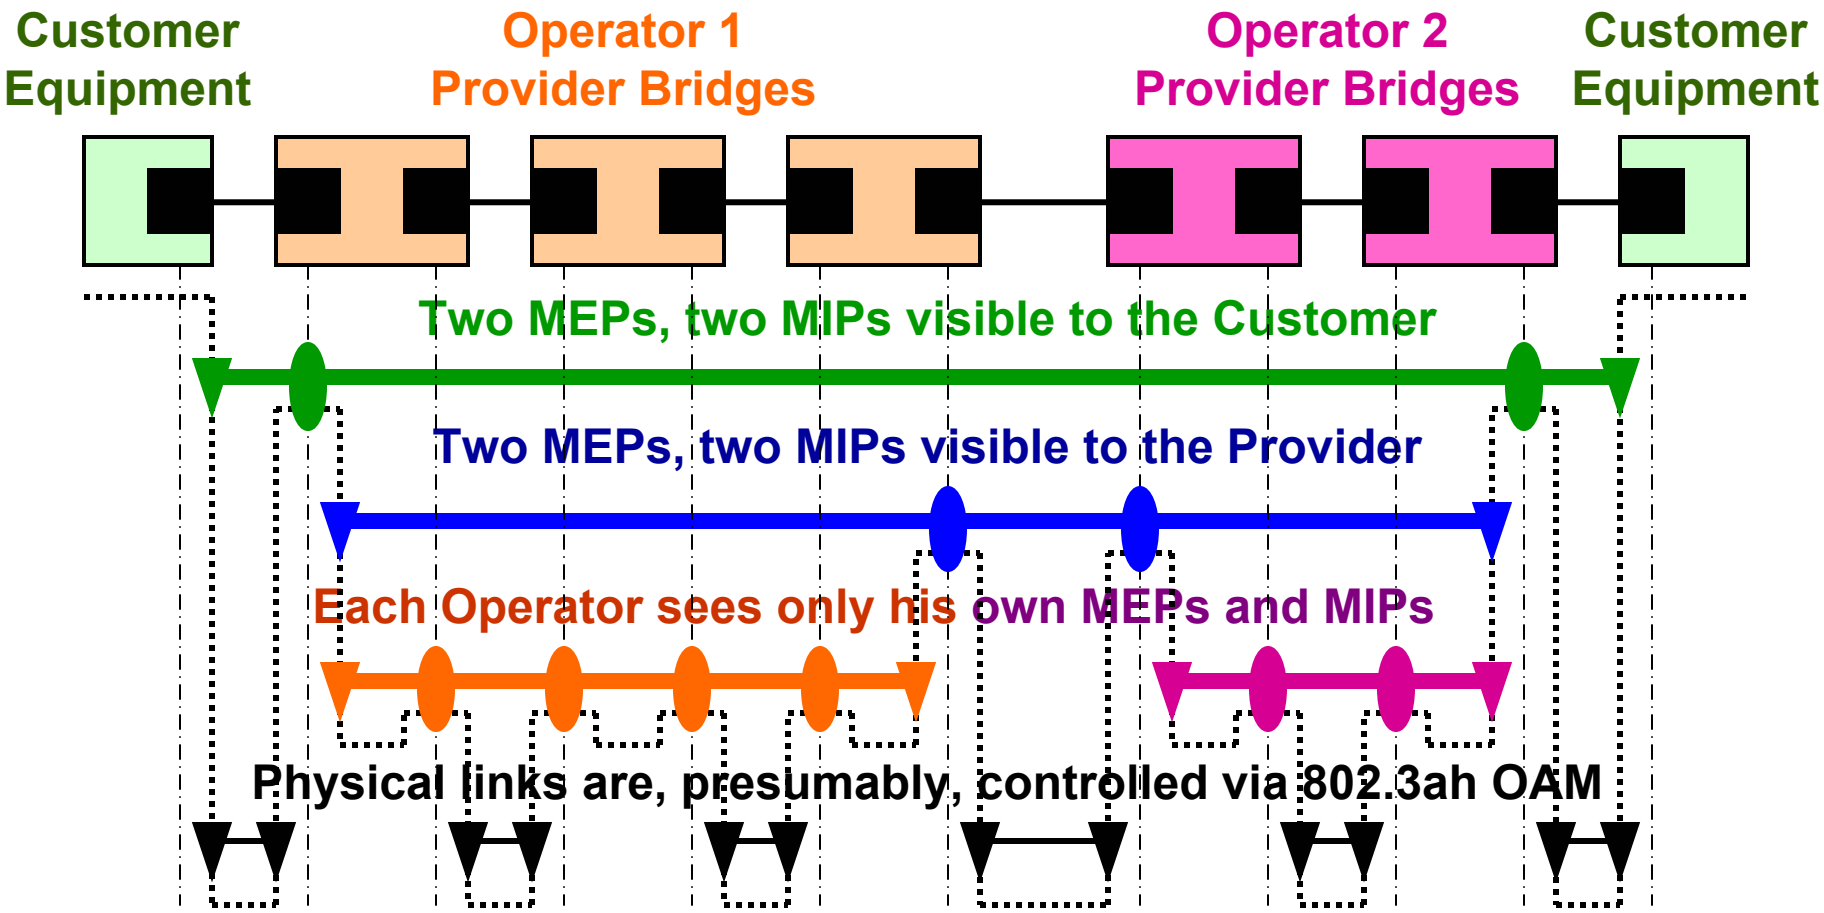

- We can see the MIP/MEP relationships.

Longitudinal View

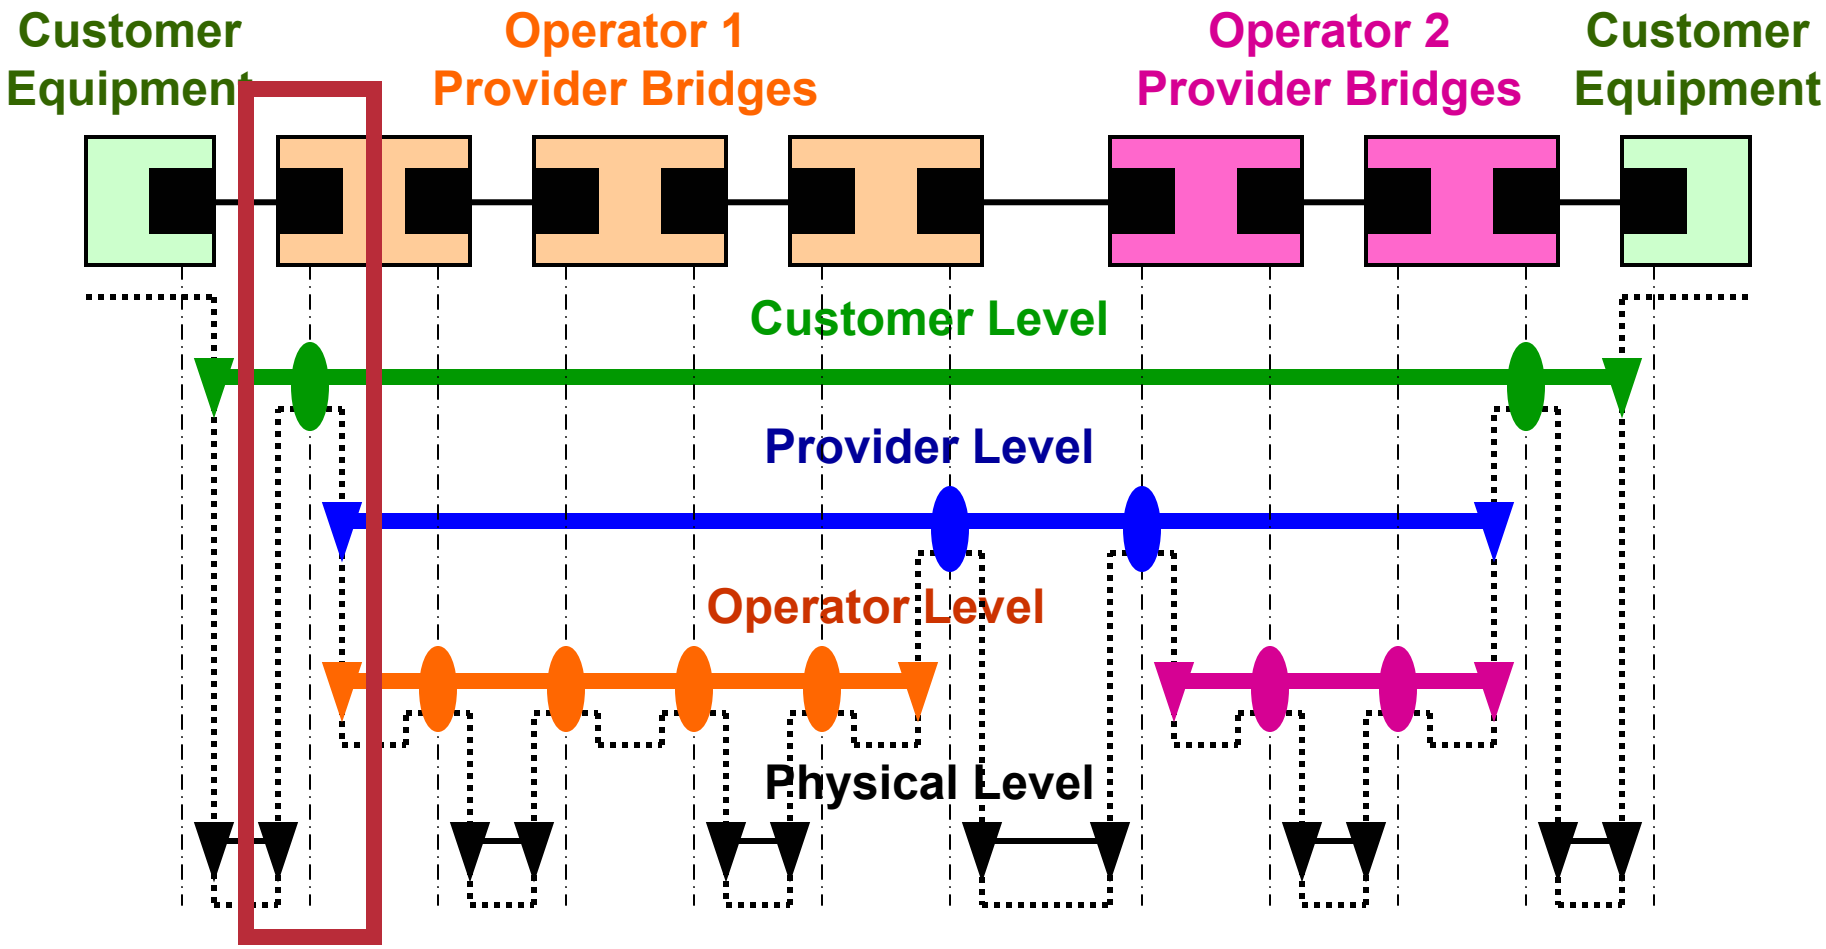

- Let's look at one port from an end view …

Longitudinal View

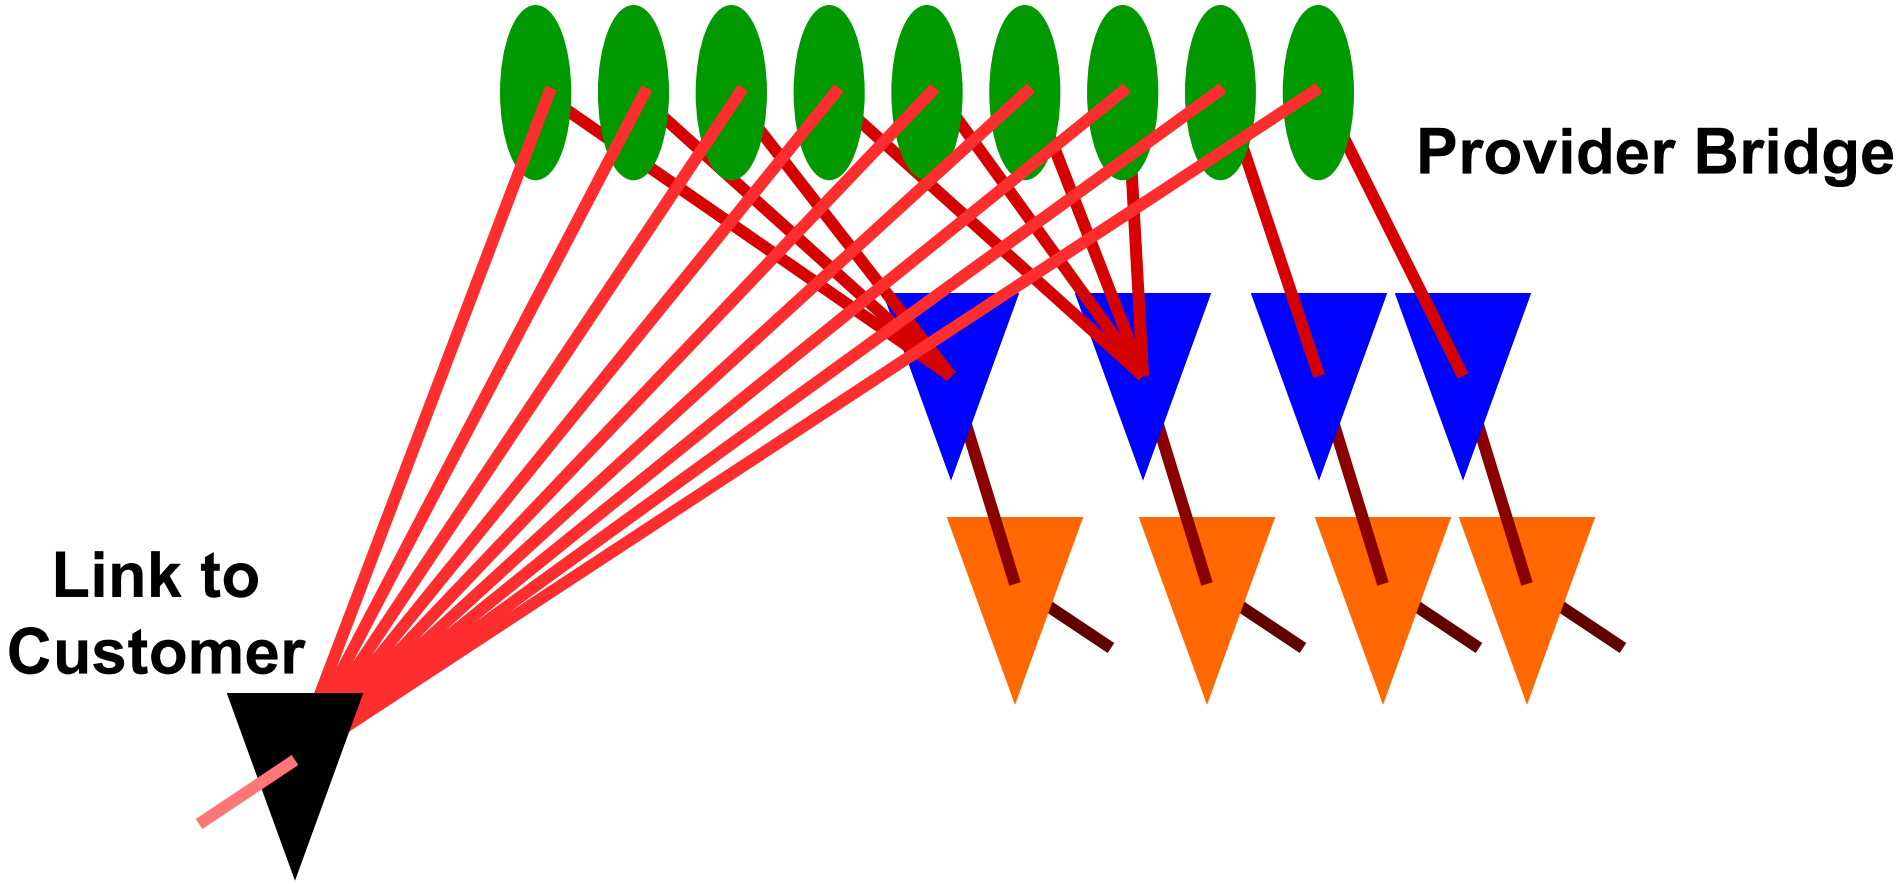

- From the end, we see the multiplexing.

CFM Packet Types

What's in a CFM Packet?

In every CFM packet, there is a:

- **Maintenance Level**
  - Customer level, Provider level, etc.
- **Globally Unique Service Identifier Label**
  - Perhaps based on a DNS name.
- **Transmitting MEP Identifier**
  - Unique only within the Service.
- **Transaction Identifier**
  - A sequence number.

CFM Packet Types

Cisco.com

- Continuity Check
- Traceroute
- Loopback (Ping)
- Alarm Indication?
- Other?

Continuity Check

Continuity Check

How does a MEP know that the other MEPs, and thus the whole service, are working?

- **It could "ping" each of the other MEPs.**

  But, a point-to-point ping requires knowing the MAC addresses of the other MEPs, which requires either configuration (which makes hardware changes difficult) or discovery (which requires yet another protocol).

  Also, point-to-point pings are inefficient.

Continuity Check

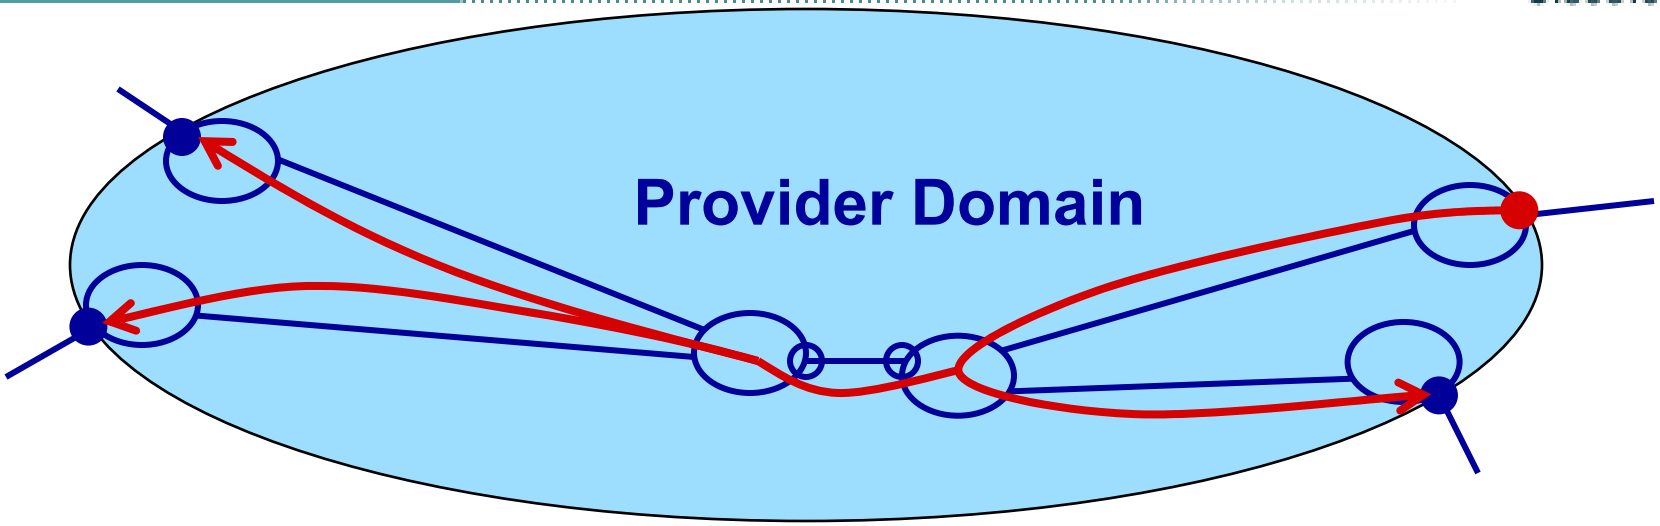

- Each MEP transmits a periodic multicast Continuity Check Message *inward* towards the other MEPs.
- For each MEP, there is 1 transmission per time period and *n* receptions.

Continuity Check

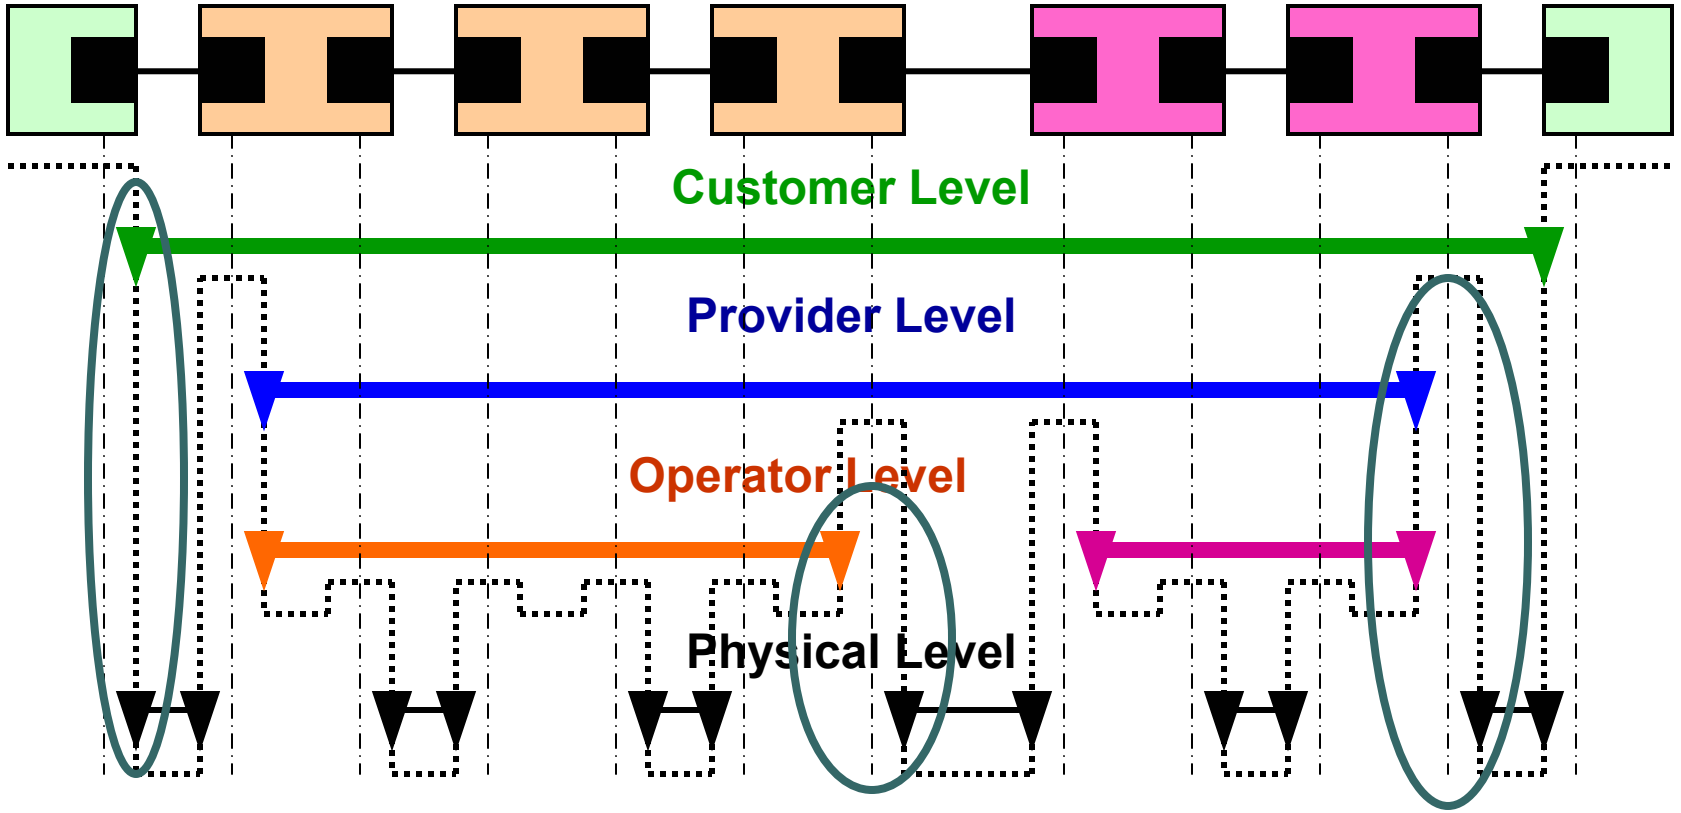

- Each MEP sends CCs to announce itself to the others.
- Some ports have multiple MEPs at different levels, or facing in different directions.

Continuity Check

What is in a Continuity Check Message?

- Destination MAC address is tied to the ME Level of the packet, to make it easier to confine CCs to a particular domain (Provider, Operator, etc.).
- ME Level, Service Identifier, MEP Identifier, Transaction Identifier.
- Lifetime.
- (Optional) Arbitrary data, software checksum.
- (Perhaps) Port state? Layer 3 management address? Transmission time?

Continuity Check

What does the receiver do with a CC? It …

- Verifies that the Service Identifier matches;
- Verifies that the MEP ID doesn't match;
- Catalogues the information by MEP ID;
- Refreshes the catalogued Lifetime;
- Checks the Transaction ID sequence; and
- (Optional) Processes optional info.

Continuity Check

What can CC detect?

- **Service cross-connect (Service ID mismatch).**
- **Duplicate MEP configurations (MEP ID match).**
- **Missing or unexpected MEPs (Optional check: unexpected MEPs may not be an error).**
- **Forwarding loops (duplicate Transaction IDs).**
- **Data loss (missing Transaction IDs or Lifetime expiration).**
- **Data corruption (bad optional data checksum).**
- **Bad frame size configuration (CC with optional data fails to reach other MEPs).**
- **Jitter (comparison of differences between transmission times and differences between reception times.)**

Continuity Check

What is the output from CC?

- Each MEP has a MIB that lists the configured list of other MEPs (if any) and the catalogue of active MEPs received from CCs.
- The catalogue also records the MAC addresses of the other MEPs in the service at this ME Level.
- Failures are normally reported to the network manager via SNMP traps.

Traceroute

Traceroute

- **This is not IP Traceroute!**

  Ethernet has no TTL (Time To Live) counter that is altered hop-by-hop in the data plane.

- **CFM Traceroute determines the path from a MEP to a given MAC address.**

  That given MAC address is normally another MEP's MAC address.

  That MAC address was learned from Continuity Check messages.

Traceroute

- MEP sends a Traceroute Message to a destination MAC address that will cause it to be stopped by any MEP or MIP at the ME Level of the Traceroute.

  This is easiest if Traceroute uses the same destination MAC address as the Continuity Checks of the next-lower level.

- Each receiving MEP or MIP sends a unicast Traceroute Reply to the original MEP.

Traceroute

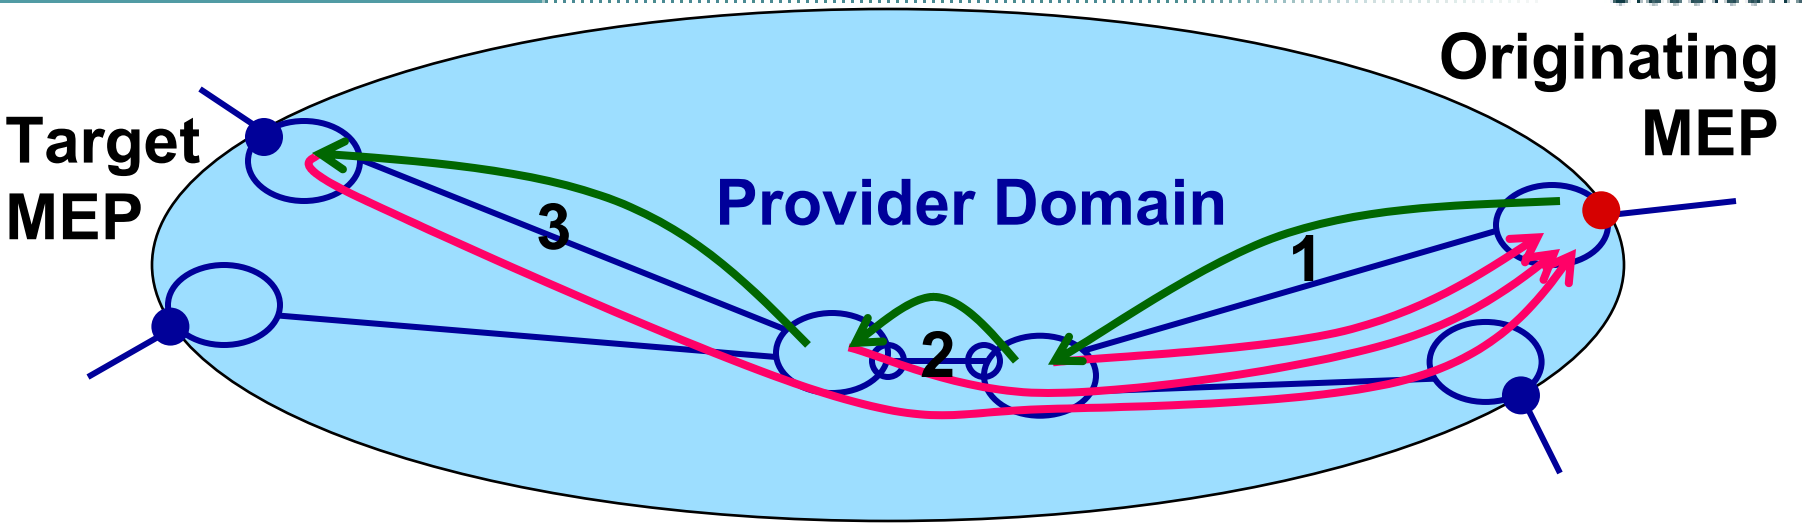

- Each receiving MEP sends a Traceroute Reply directly to the Originating MEP, and regenerates the Traceroute Message.
- If a MIP is not sure where to send the Traceroute Message, the trail ends.

Traceroute

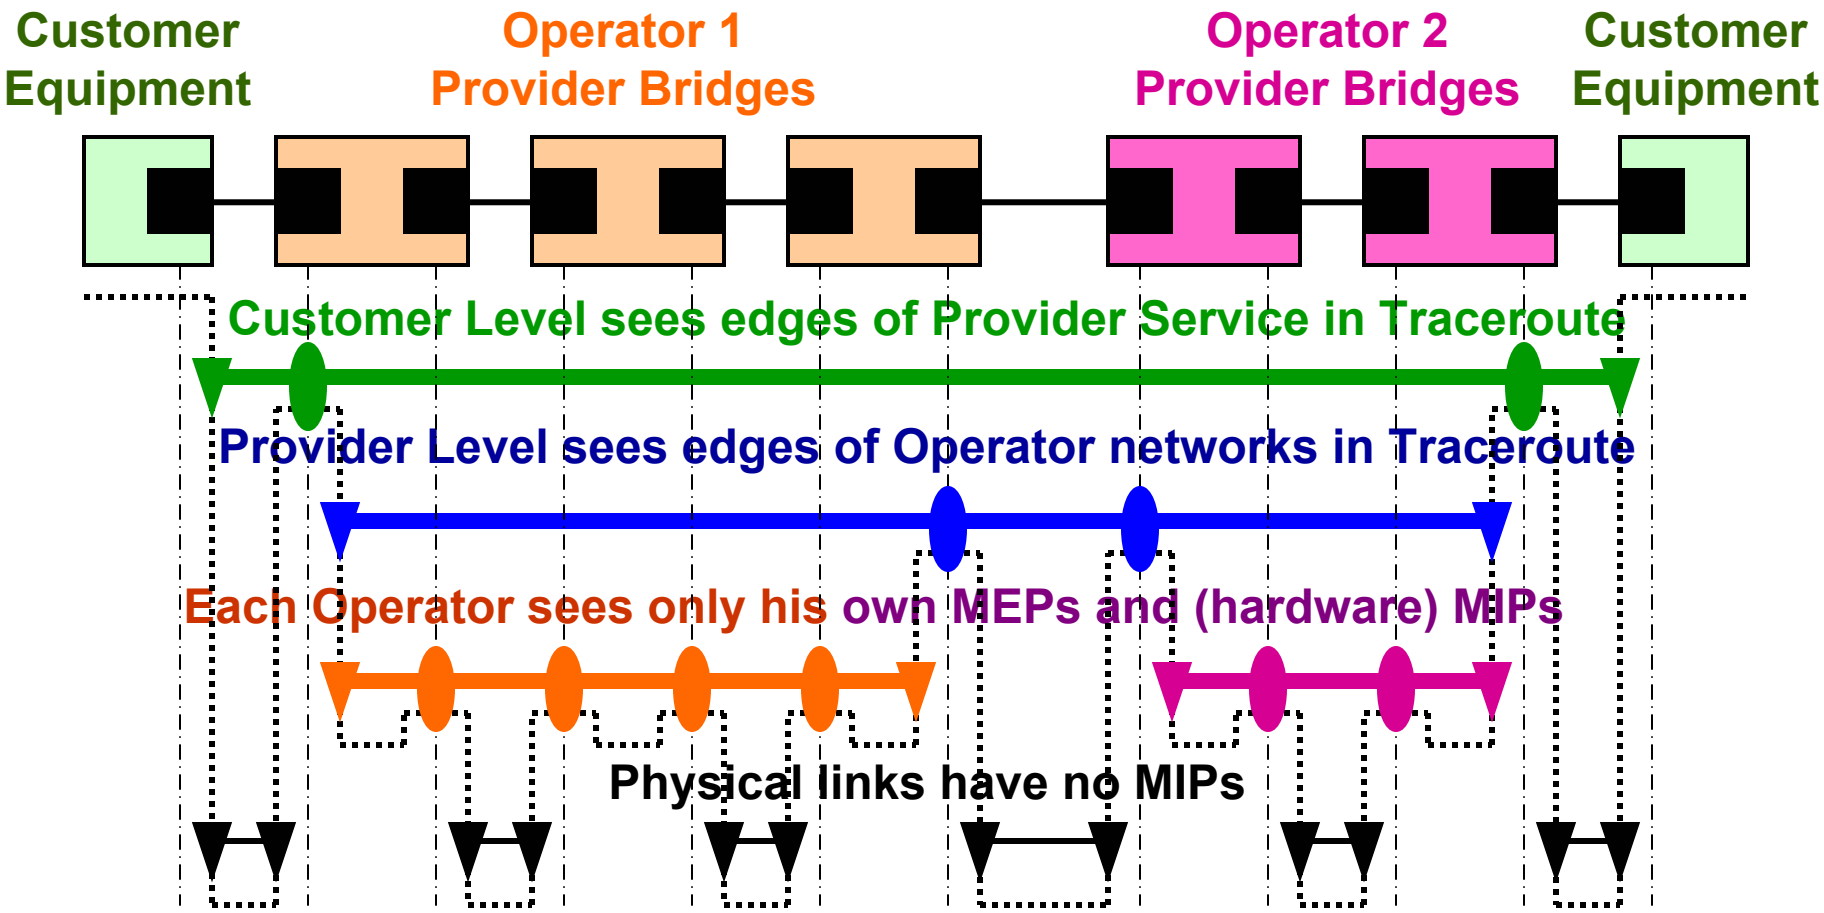

- MEPs at one ME Level are MIPs at some higher level.

Traceroute

But, what if the traced MAC address has been forgotten?

- MIPs can take copies of Continuity Checks and catalogue them, perhaps in software by source MAC address only in a “CC Database”.
- This database allows Traceroute to follow “forgotten” links, and thus return useful information.

Traceroute

Cisco.com

What is in a Traceroute Message?

- Destination MAC address is tied to the ME Level of the packet, to make it easier to confine CCs to a particular domain (Provider, Operator, etc.).
- ME Level, Service Identifier, MEP Identifier, Transaction Identifier.
- TTL. (To prevent Traceroute from looping forever, even if data plane has a loop.)
- Target MAC address.
- Originating MEP's MAC address.

Traceroute

What is in a Traceroute Reply?

- The Transaction ID and TTL from the Traceroute Message, so that the replies can be coordinated.
- Information about each MIP through which the Traceroute Message flowed.
  - There may be 1 or 2 in a single bridge.
  - Information includes the MIPs' MAC addresses.
- The identity of the next hop (in P802.1ab LLDP terms).
- Information about the reason for forwarding or not forwarding the Traceroute Message (e.g. MAC address forgotten, destination reached).

Traceroute

What is the output from Traceroute?

- **The results of the Traceroute fill a MIB for inspection by the Network Administrator.**
- **These results provider the MEP with a list of the MIPs and their MAC addresses for use in Loopback Messages.**

Loopback

Loopback

What is in a Loopback Message?

- Unicast MAC address of target.
- ME Level, Service Identifier, MEP Identifier, Transaction Identifier.
- (Optional) Arbitrary data, software checksum.
- (Perhaps) Transmission time? Reply time?

Loopback

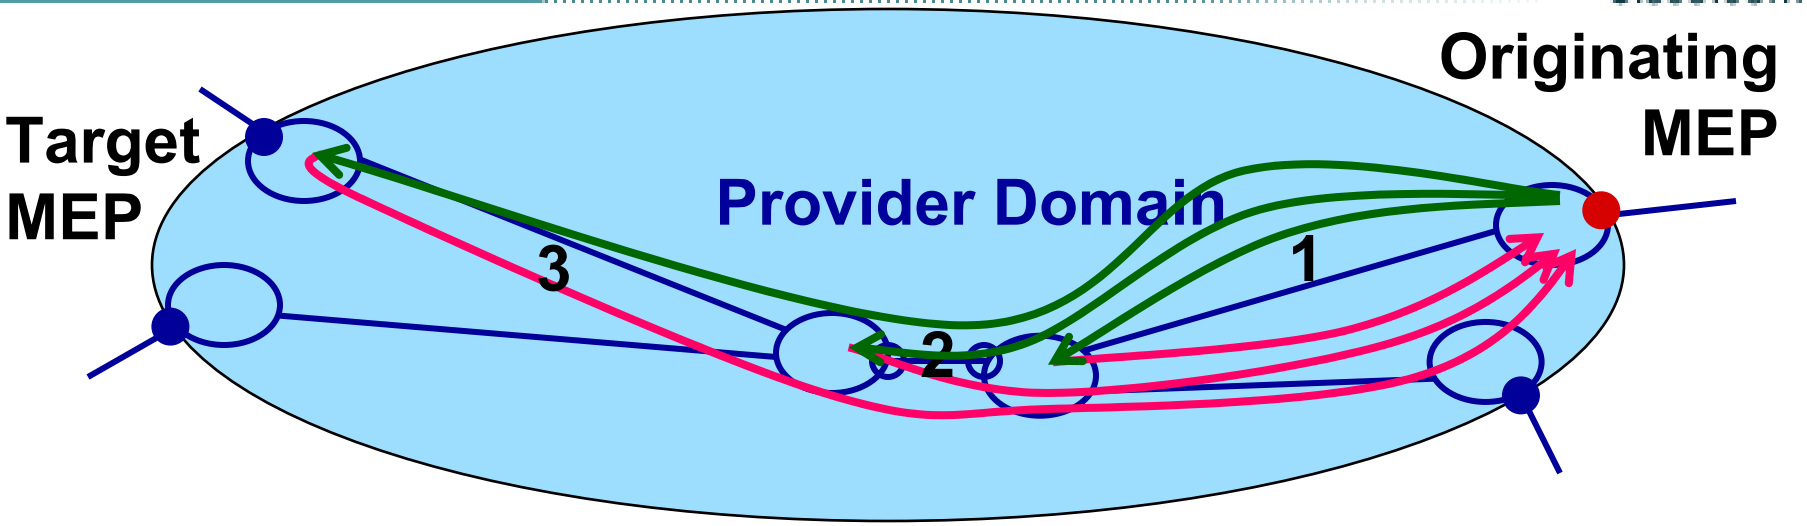

- MEP can send a Loopback to any MEP or MIP in the service.
- MEPs are discovered by Continuity Check, MIPs are discovered by Traceroute.

Loopback

What is Loopback Good for?

- Loopback follows the unicast data path, not the multicast path.
- Sending Loopback to successive MIPs can determine the location of a fault.
- Sending a high volume of Loopback Messages can test bandwidth, reliability, or jitter of a service.

Alarm Indication Signal

Alarm Indication Signal

What is AIS?

- When a MEP discovers an error, presumably a Continuity Check failure, it transmits periodic 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

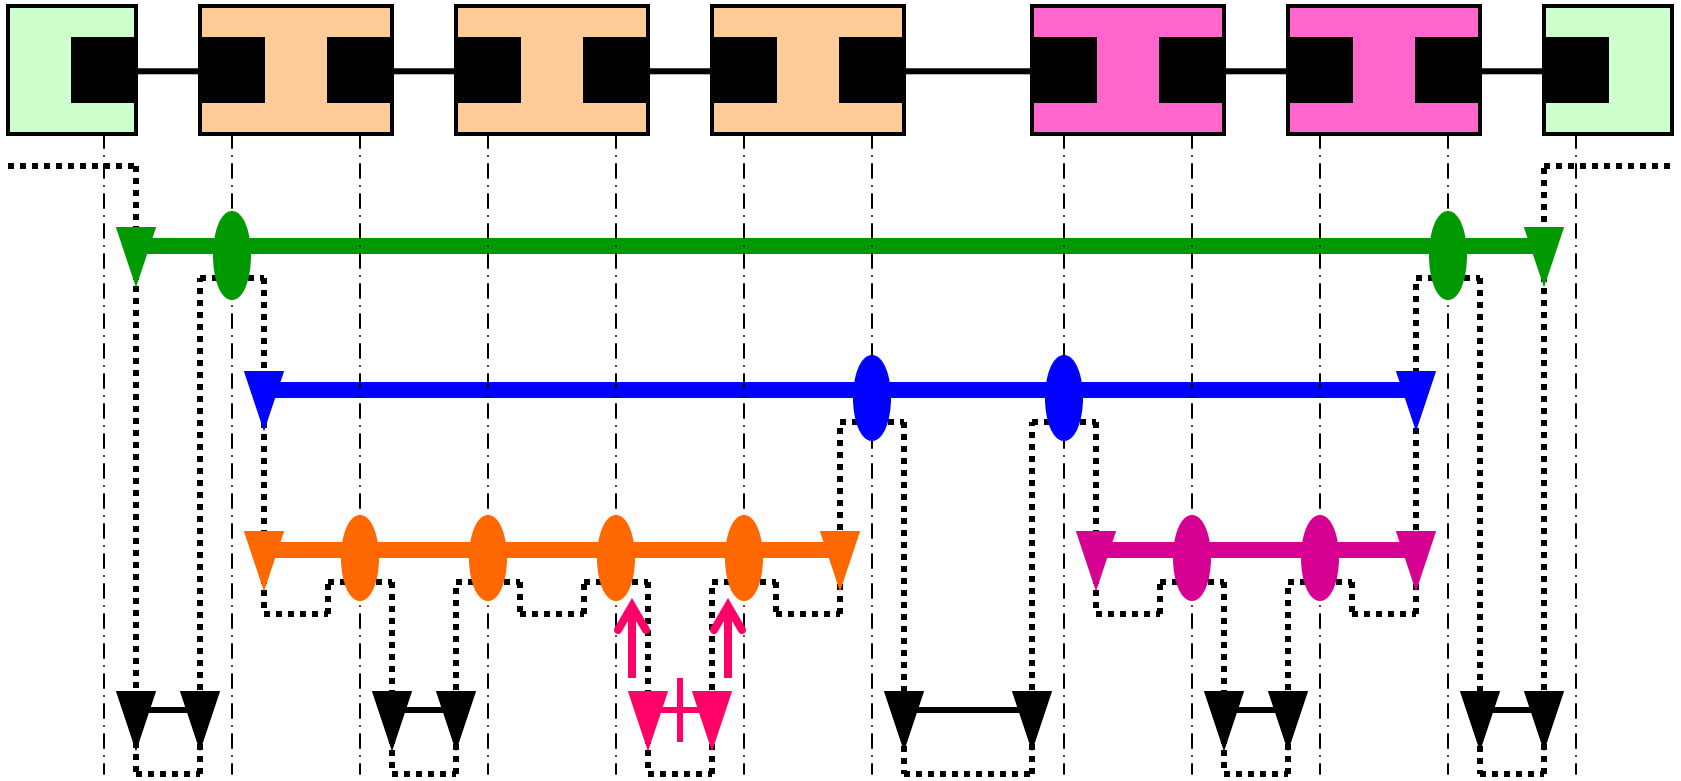

- **Failure causes MEP to generate AISs in *opposite* direction of Continuity Check.**

Alarm Indication Signal

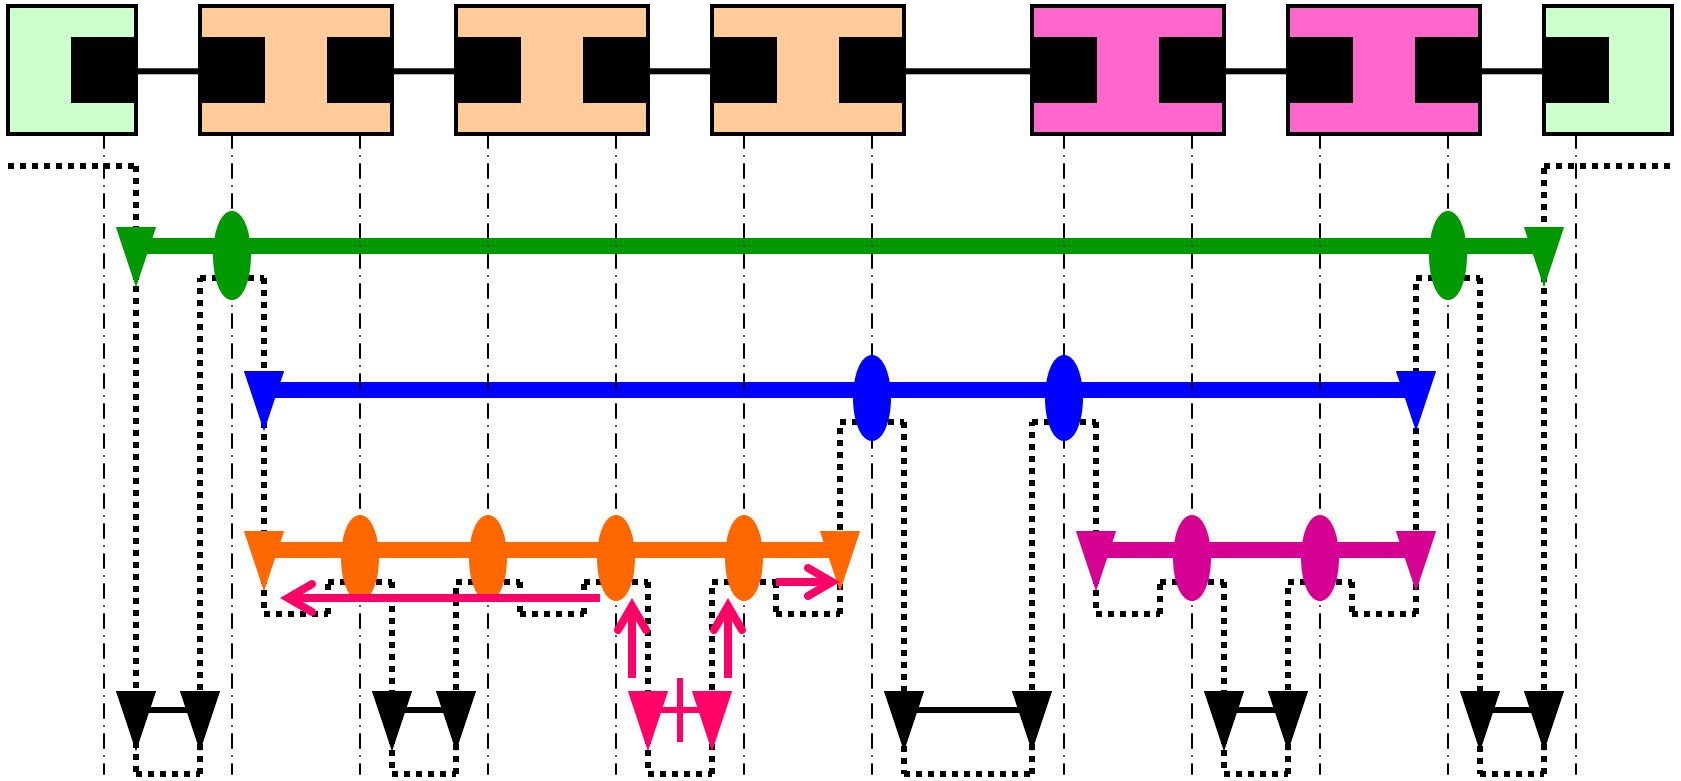

- **MIPs receive AISs and promote them to their own ME Levels for propagation to MEPs.**

Alarm Indication Signal

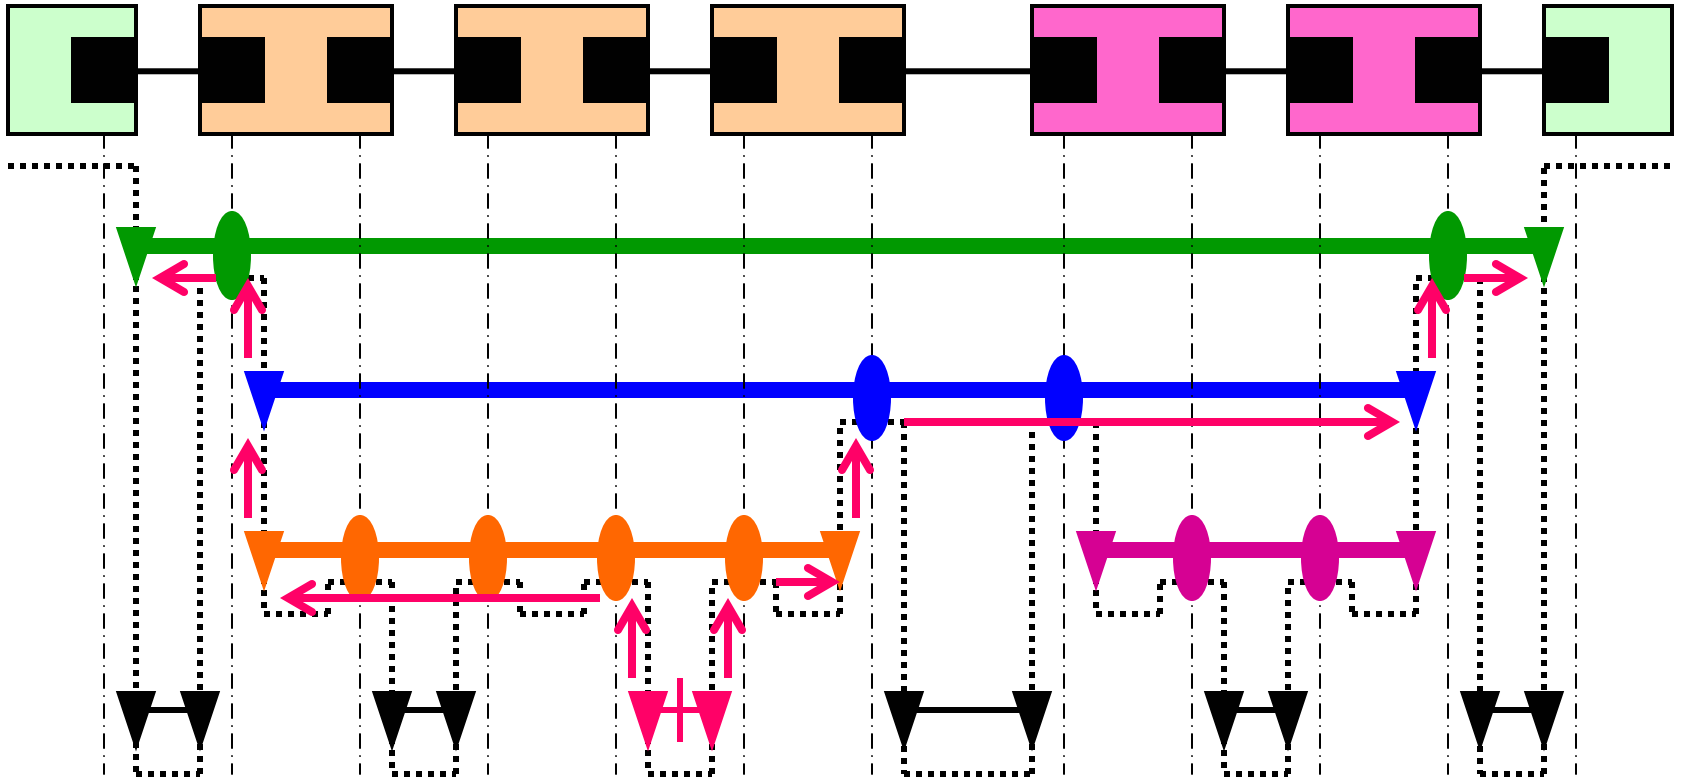

- MEP may propagate or terminate AIS.
- Propagation is shown.

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.

Alarm Indication Signal

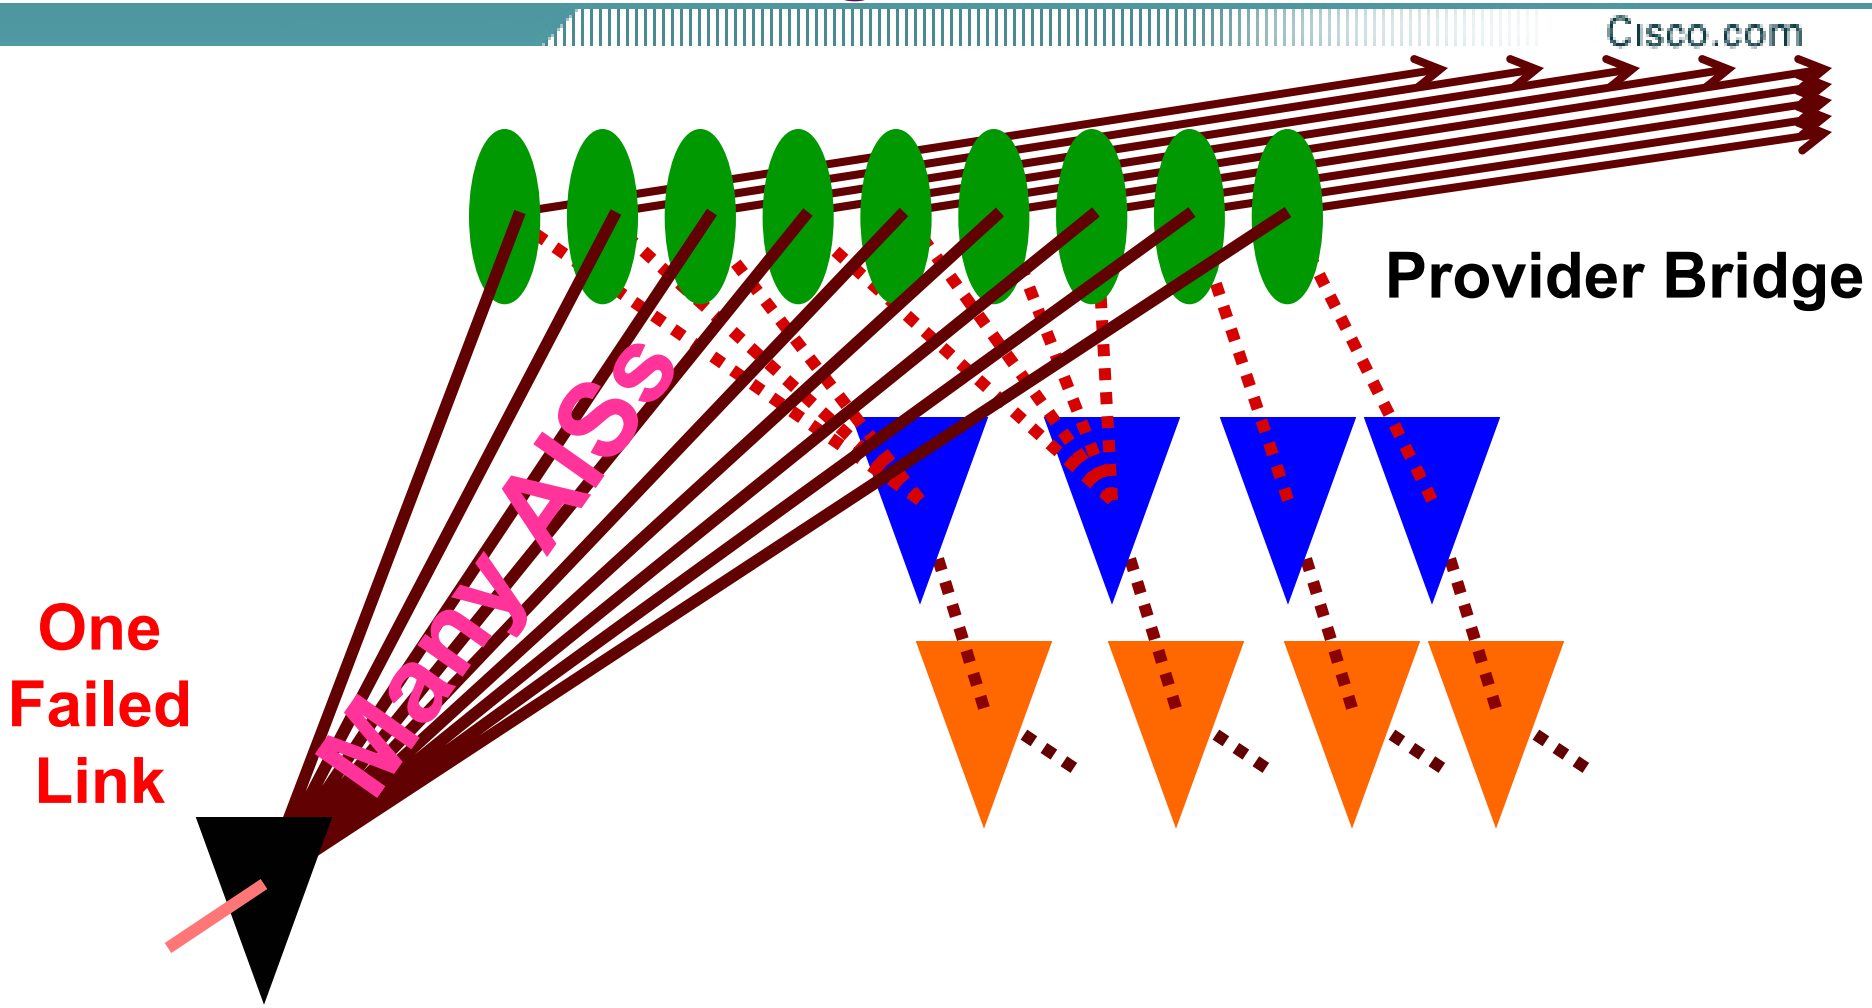

- One physical failure, many AISs.

Other CFM Packet Types

Other CFM Packet Types?

Cisco.com

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

- **Both IEEE 802.1 and ITU-T Study Group 13 Question 3, are cooperating.**
- **The extra meetings are enabling the rapid progression of the standard(s).**
- **It remains to be seen whether there will be problems caused by having two groups working on the same problem, but so far, there is great promise for the successful completion of the standards.**


CISCO SYSTEMS