Issues brought up by P802.1ag Draft 6

Norman Finn

Rev. 1

Part 1: Rearranging the components of the Forwarding Process

P802.1ag Draft 6.0 Figure 22-1

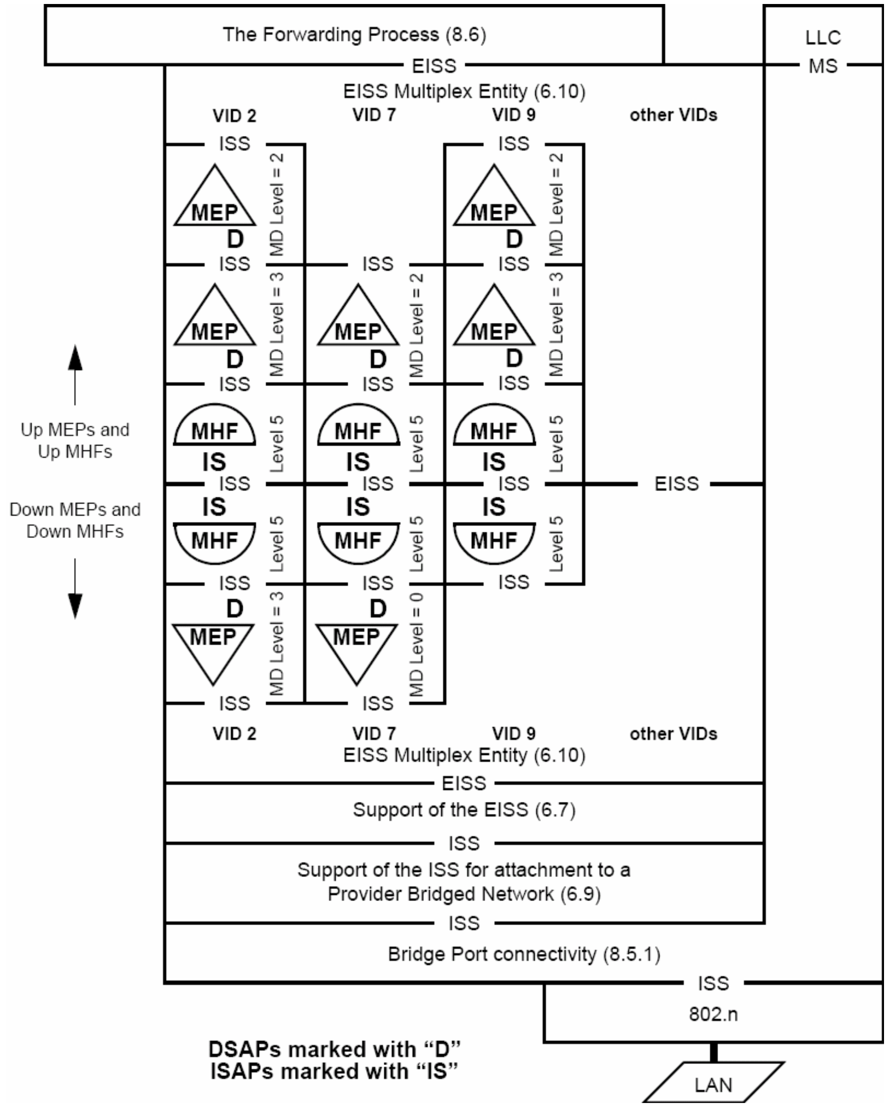

P802.1ag Draft 6.0 Figure 22-1, revised

- This is the **starting point**.
- CFM (19.2) entities **compressed** into single boxes (as shown by gray triangles and semicircles).
- Additional 802.1Q subclauses **added** (6.5, 6.9)

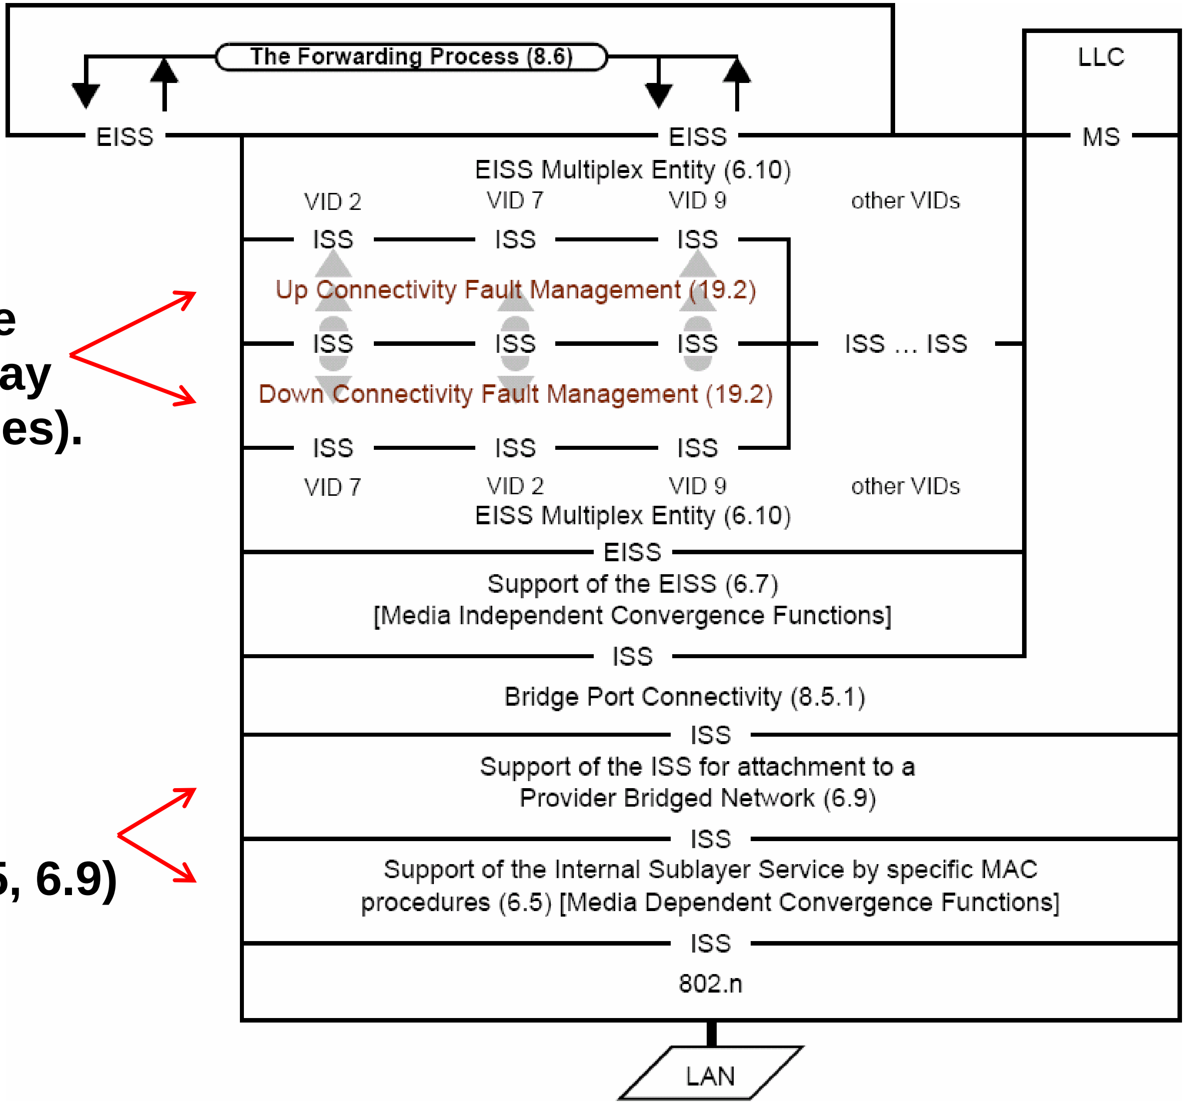

P802.1ag Draft 6.0 Figure 22-1, revised

- **Add an additional pair of ISS Multiplex Entities.**
- **Remove irrelevant (to this discussion) shims.**

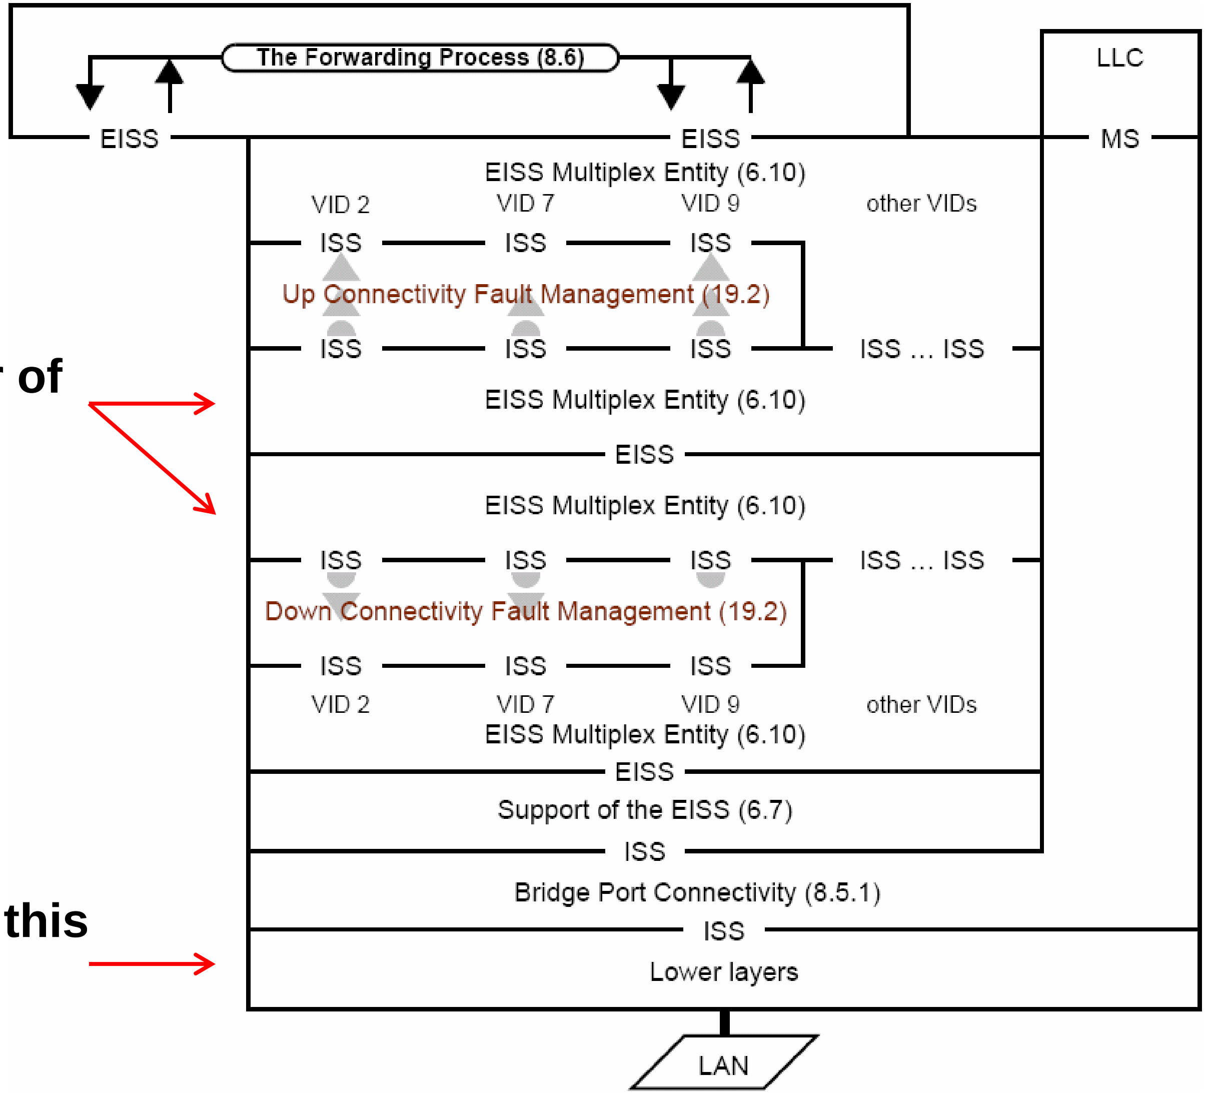

P802.1ag Draft 6.0 Figure 22-1, revised

- **Collapse the EISS-Multiplex shims and the CFM shims into a pair of shims, an Up shim and a Down shim.**

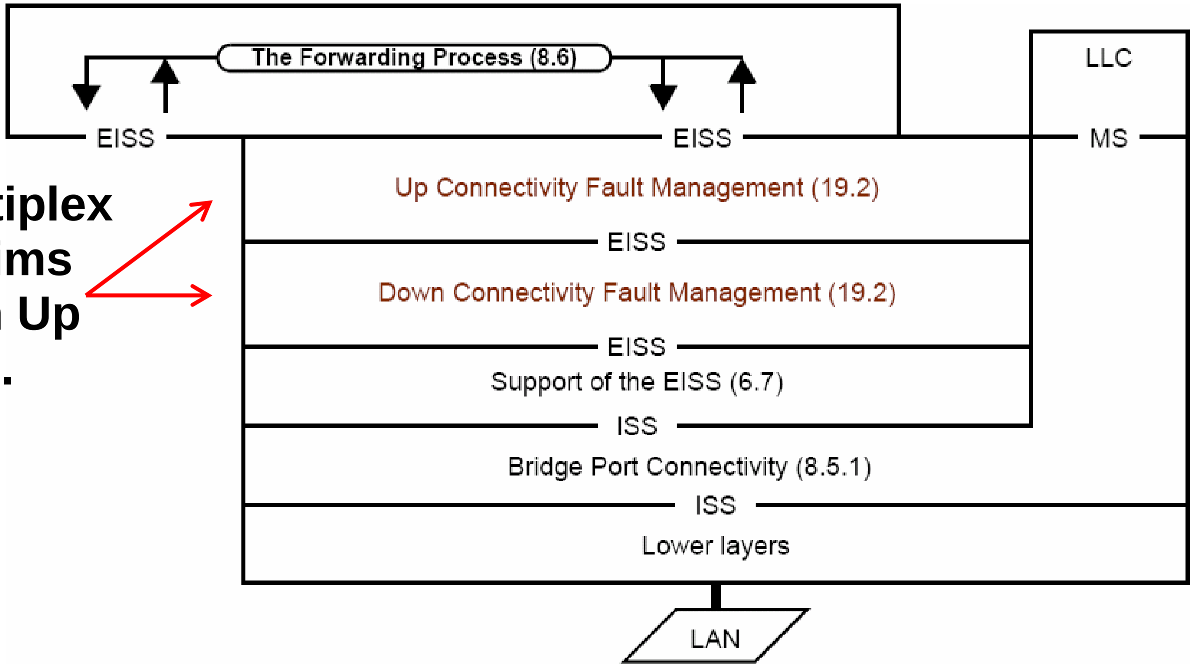

IEEE Std. 802.1Q-2005 Figure 8-9, revised

- **Delete the data bases (only for clarity).**

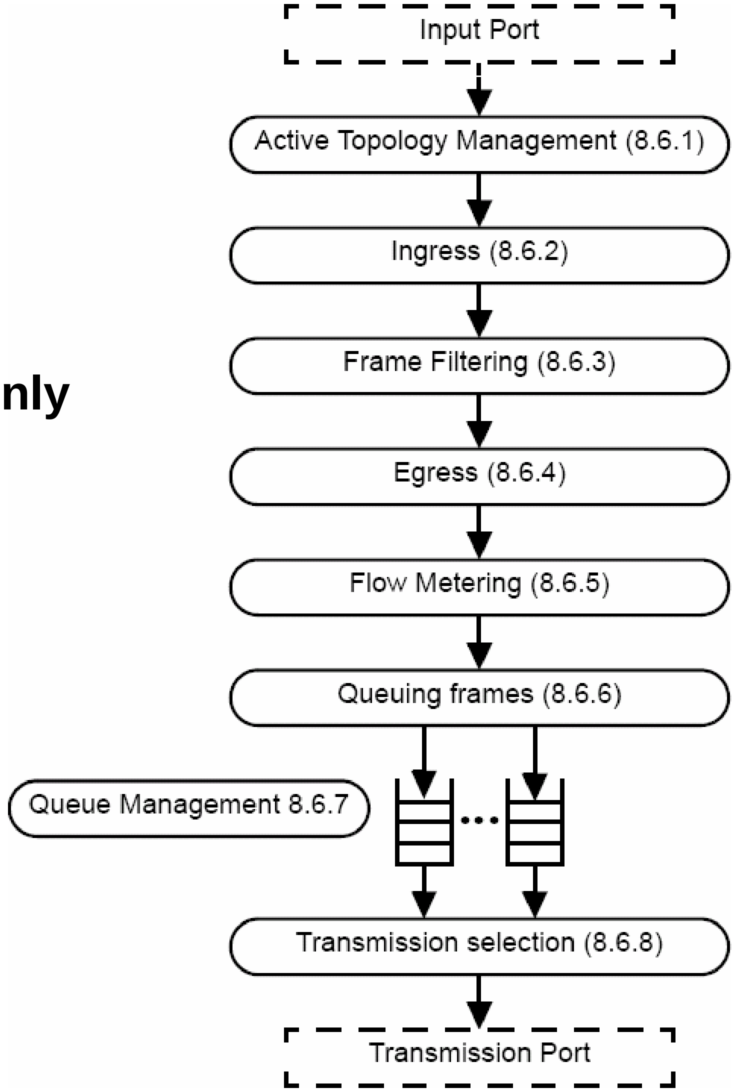

IEEE Std. 802.1Q-2005 Figure 8-9, revised

- **These functions can be per-port; the amount of information that each has about other ports is minimal.**

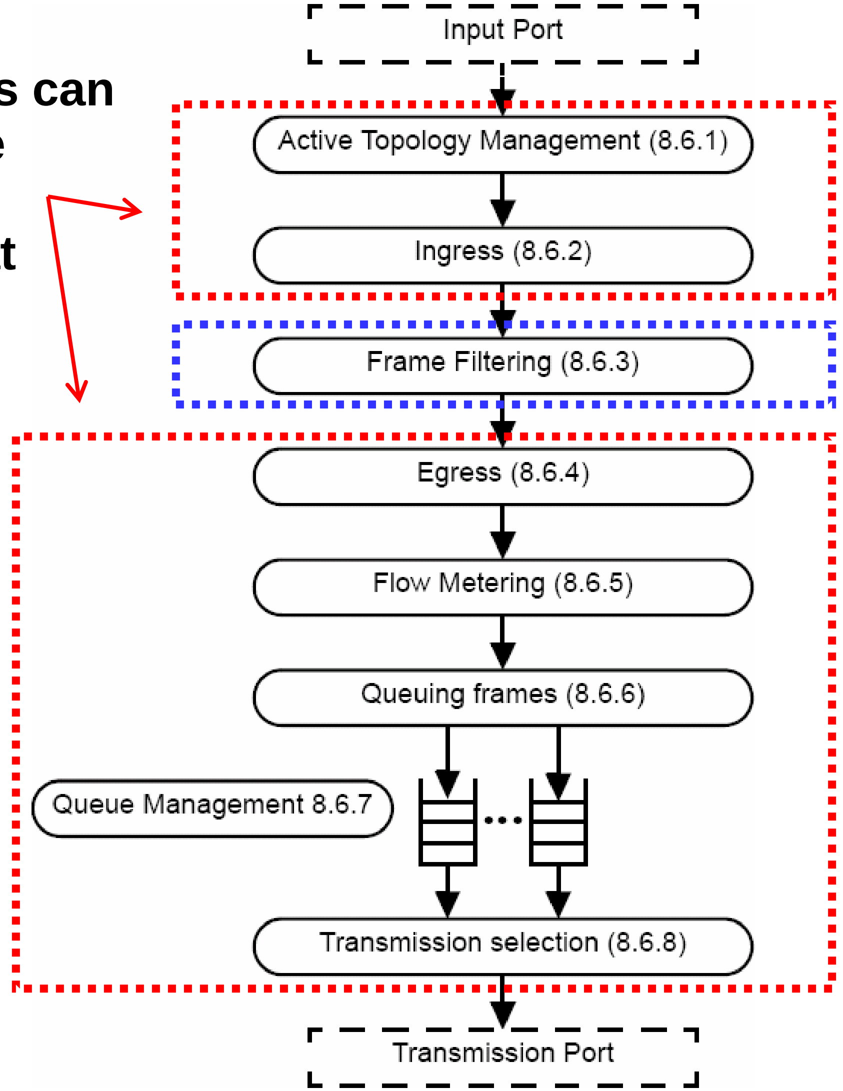

- **This function can be central; it clearly has information about all ports.**

IEEE Std. 802.1Q-2005 Figure 8-9, revised

- **Fold that picture in half.**

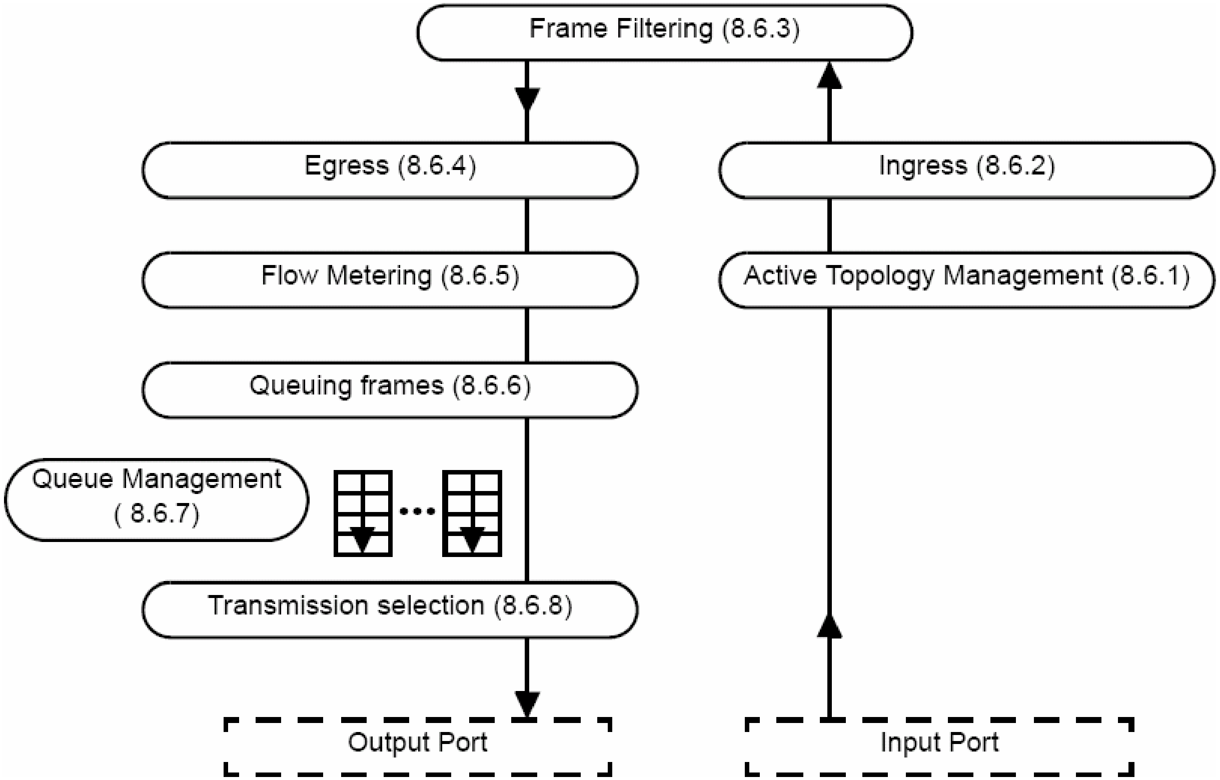

IEEE Std. 802.1Q-2005 Figure 8-9, revised

- **Show this folded diagram with multiple copies of per-port functions, and a single copy of the central function.**

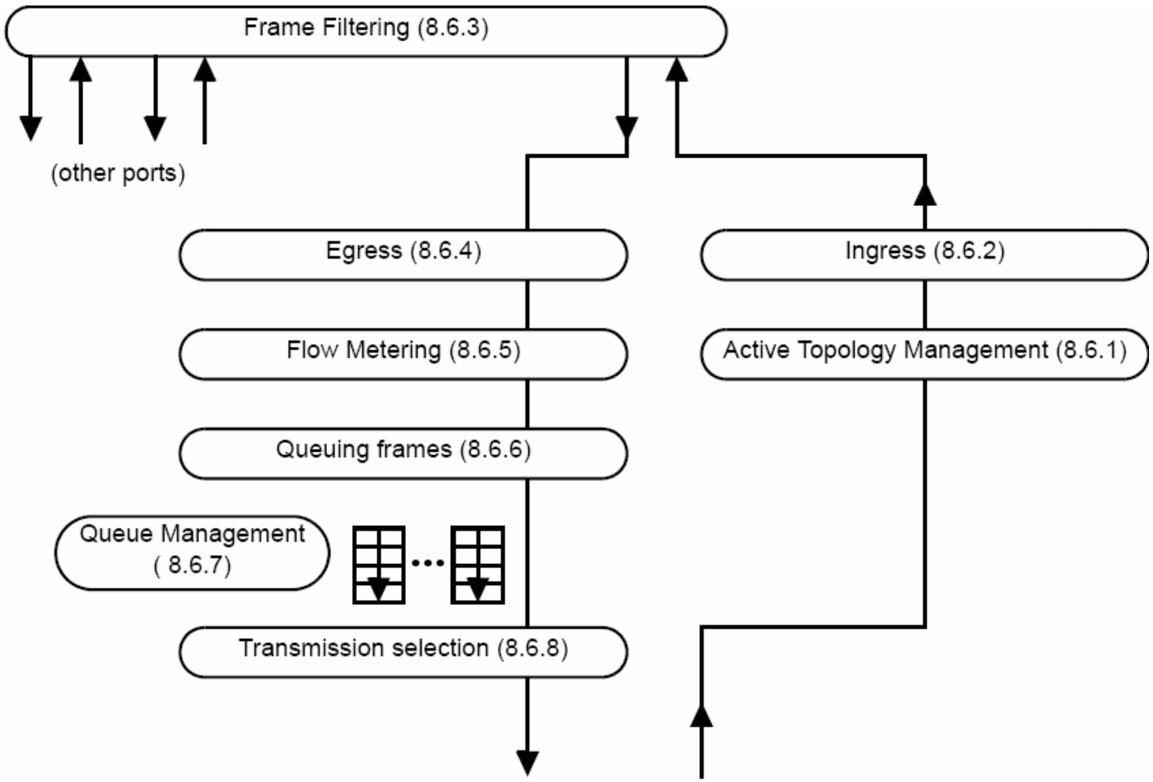

Combining the folded Forwarding Process into Revised P802.1ag diagram

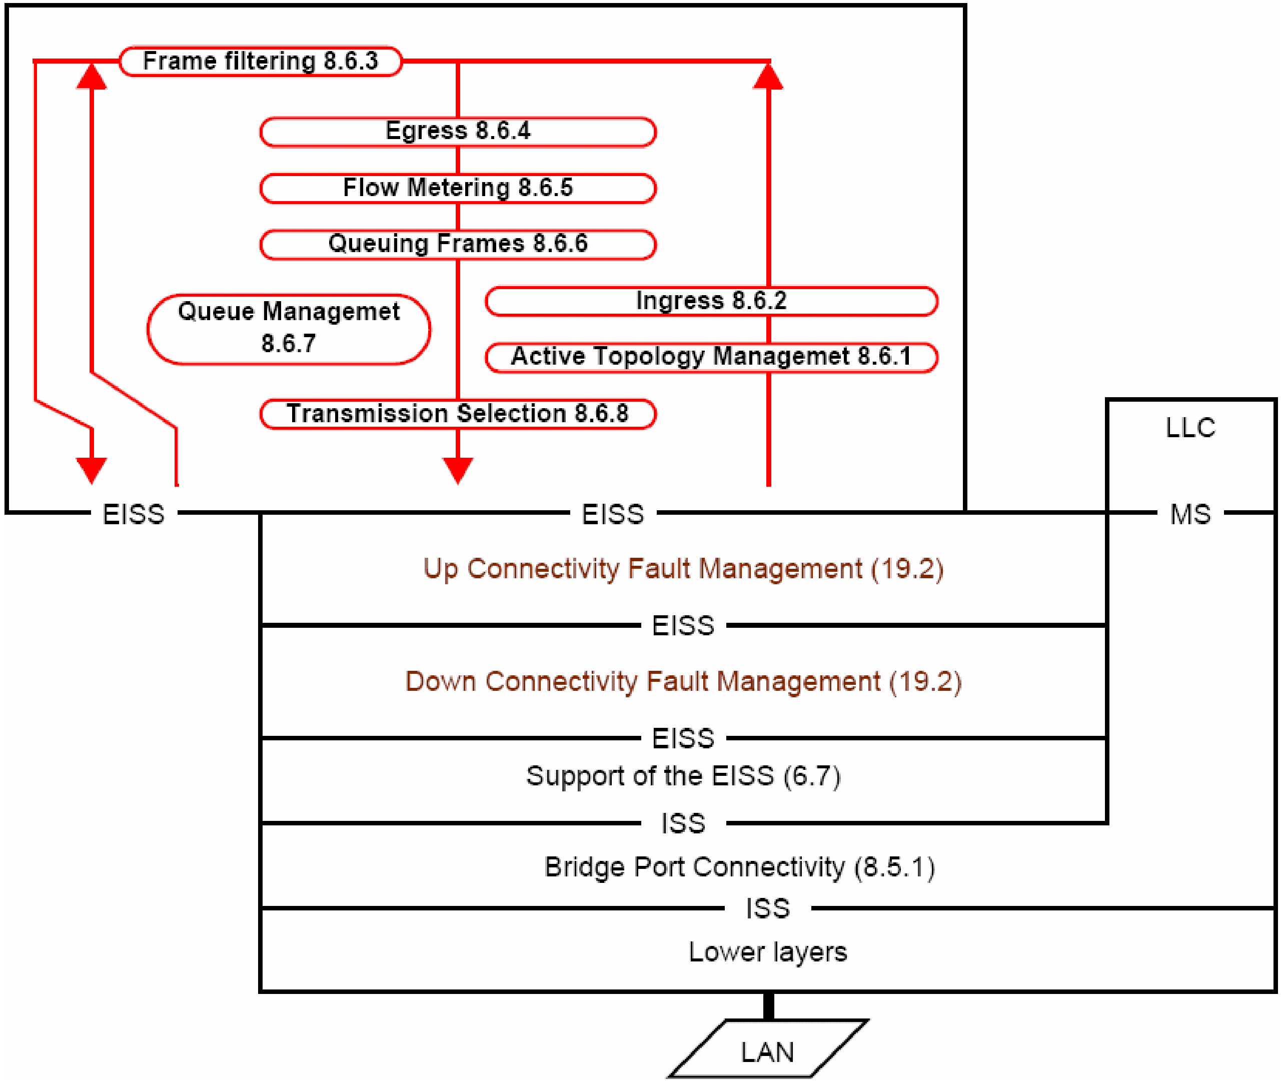

- **Combine the two diagrams.**

At last! The problem statement

1. In D6.0, the Down MEPs' reflected Loopback Replies, which can take a significant part of a Port's bandwidth, bypass the output queuing mechanism (while passed-through data is queued).
2. When a port is blocked in D6.0, presumably at the edge of a network by a means such as MSTP role restriction, Up MEPs are disconnected from their Maintenance Association. (In fact, when a Port is blocked, an Up MEP is not of much use at all.)

Problem 1

Queuing Mechanism

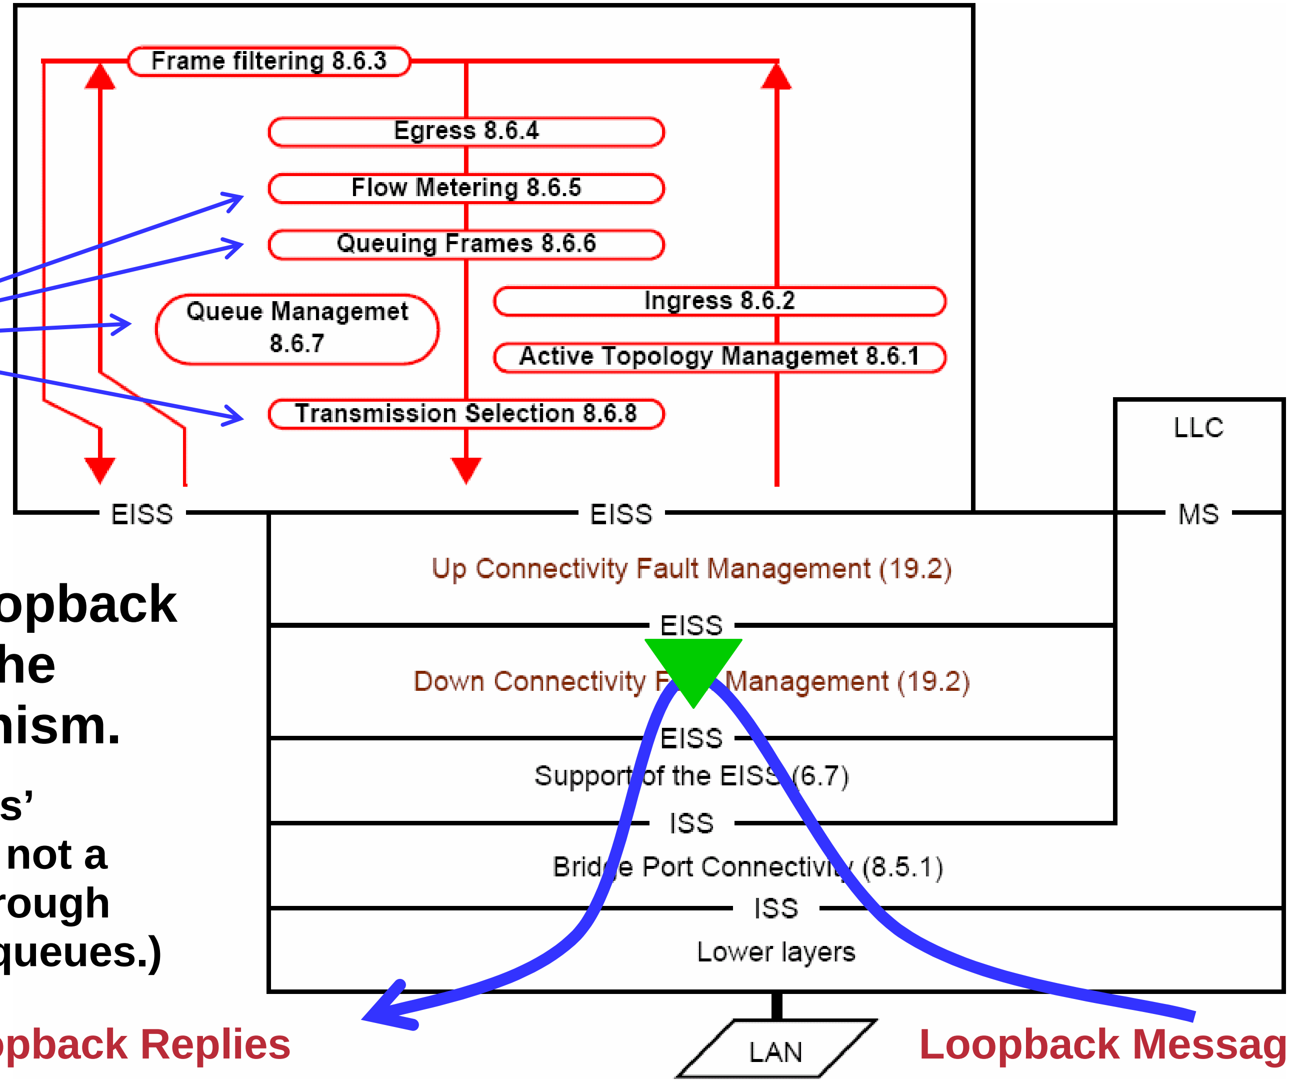

- **Down MEPs' Loopback path bypasses the queuing mechanism.**
- **(This Port's Up MEPs' Loopback paths are not a problem; they go through other Ports' output queues.)**

Loopback Replies

Loopback Messages

Problem 2

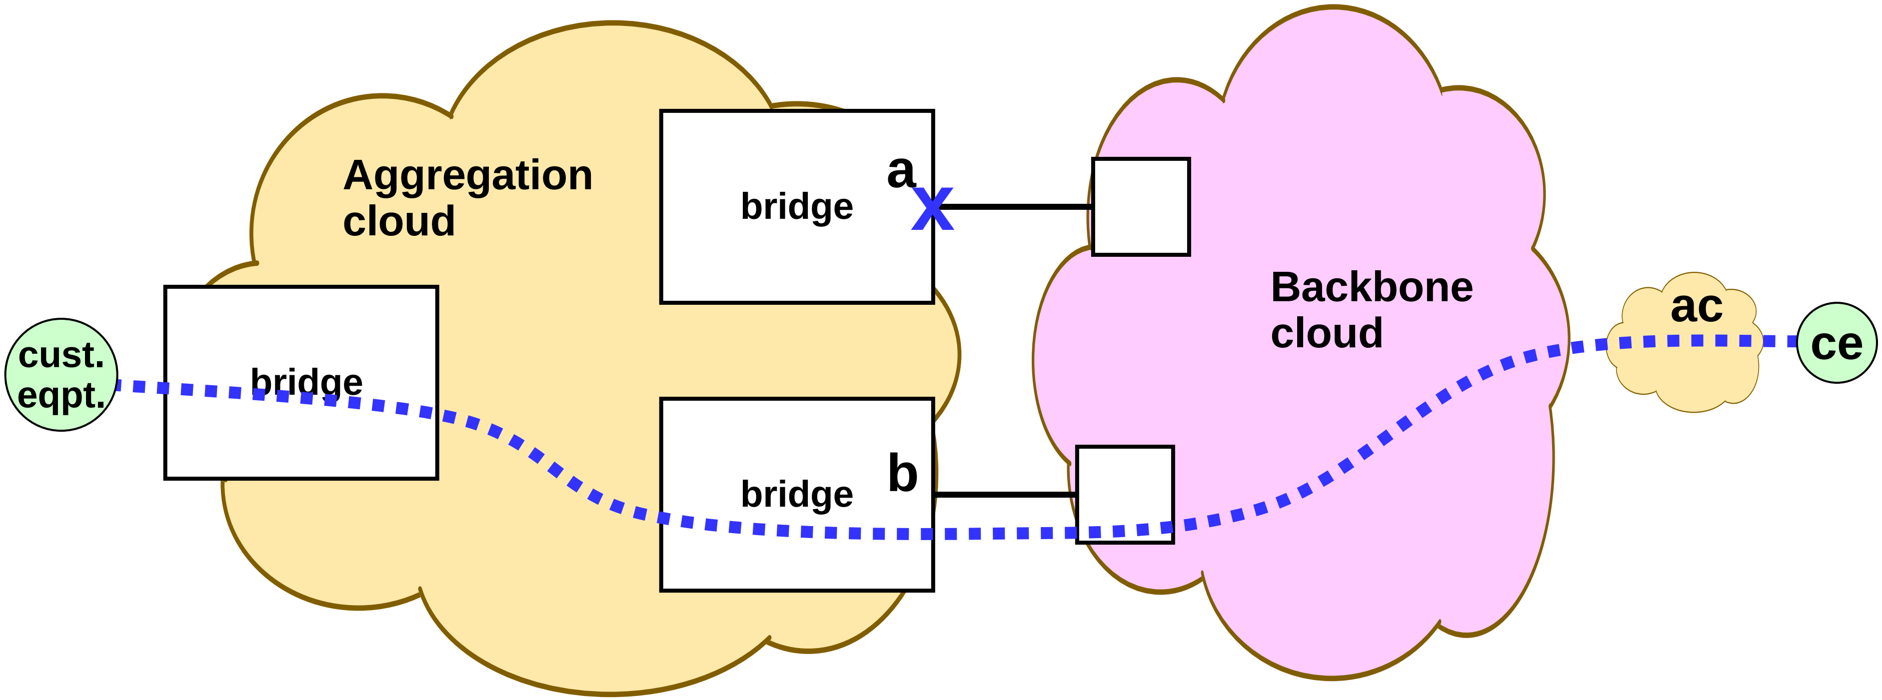

- Point to point service (MA in blue).
- Access cloud selects the bridge, a or b, to use for this service. Blue X marks the blocked port.

Problem 2

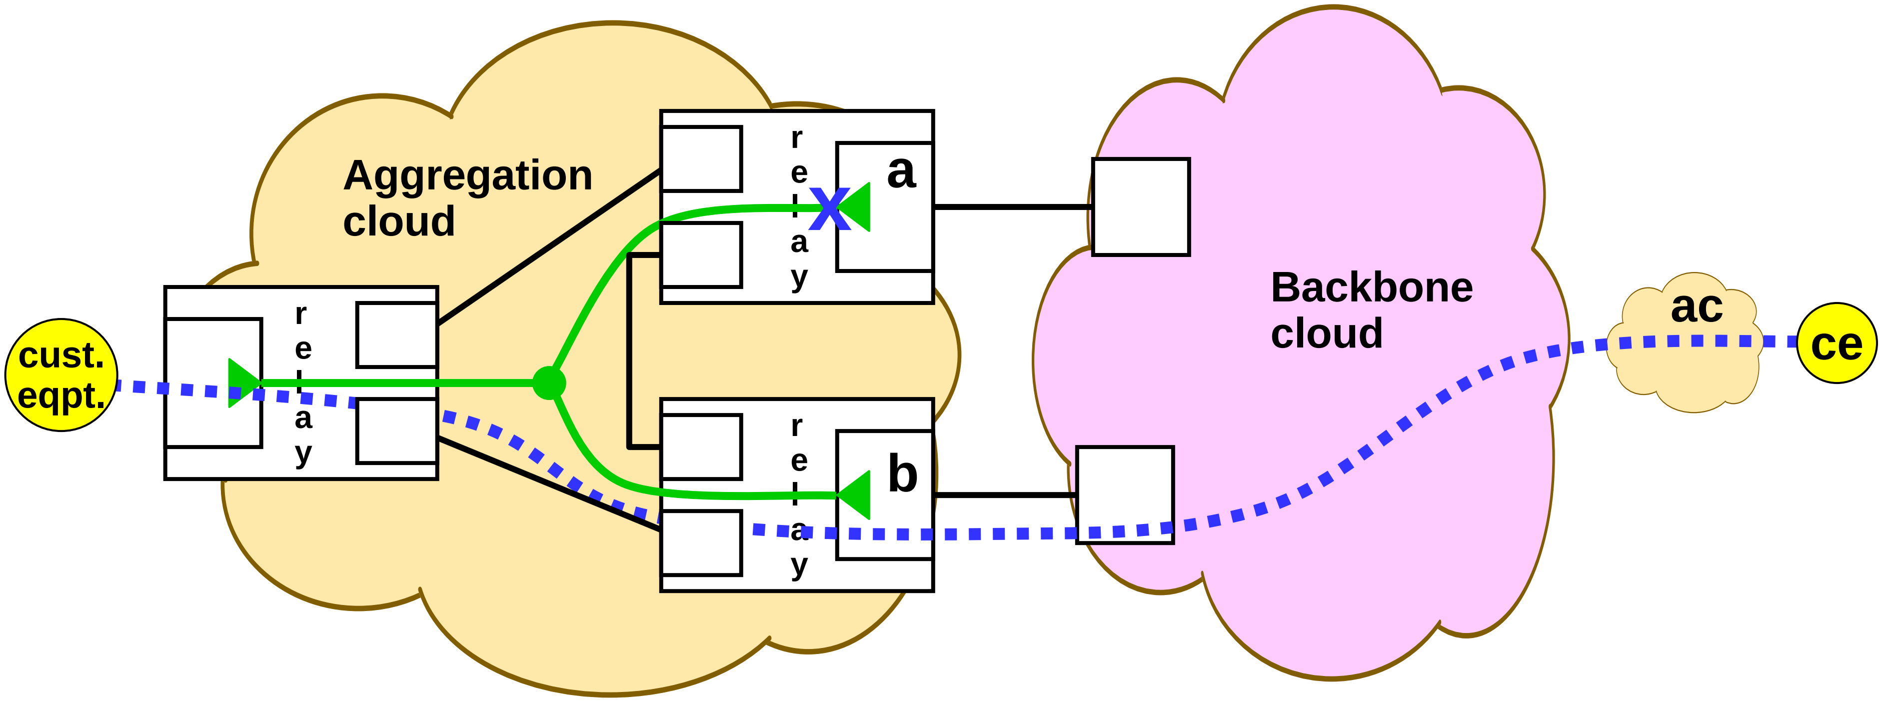

- Aggregation cloud operator creates a 3-point (green) service to support the 2-point (blue) service.
- Block (blue X) is in Relay, not outside the Port.
- How does Agg. cloud operator know whether Port a, or the Bridge it belongs to, are working?

Problem 2

Up MEPs' Maintenance Associations

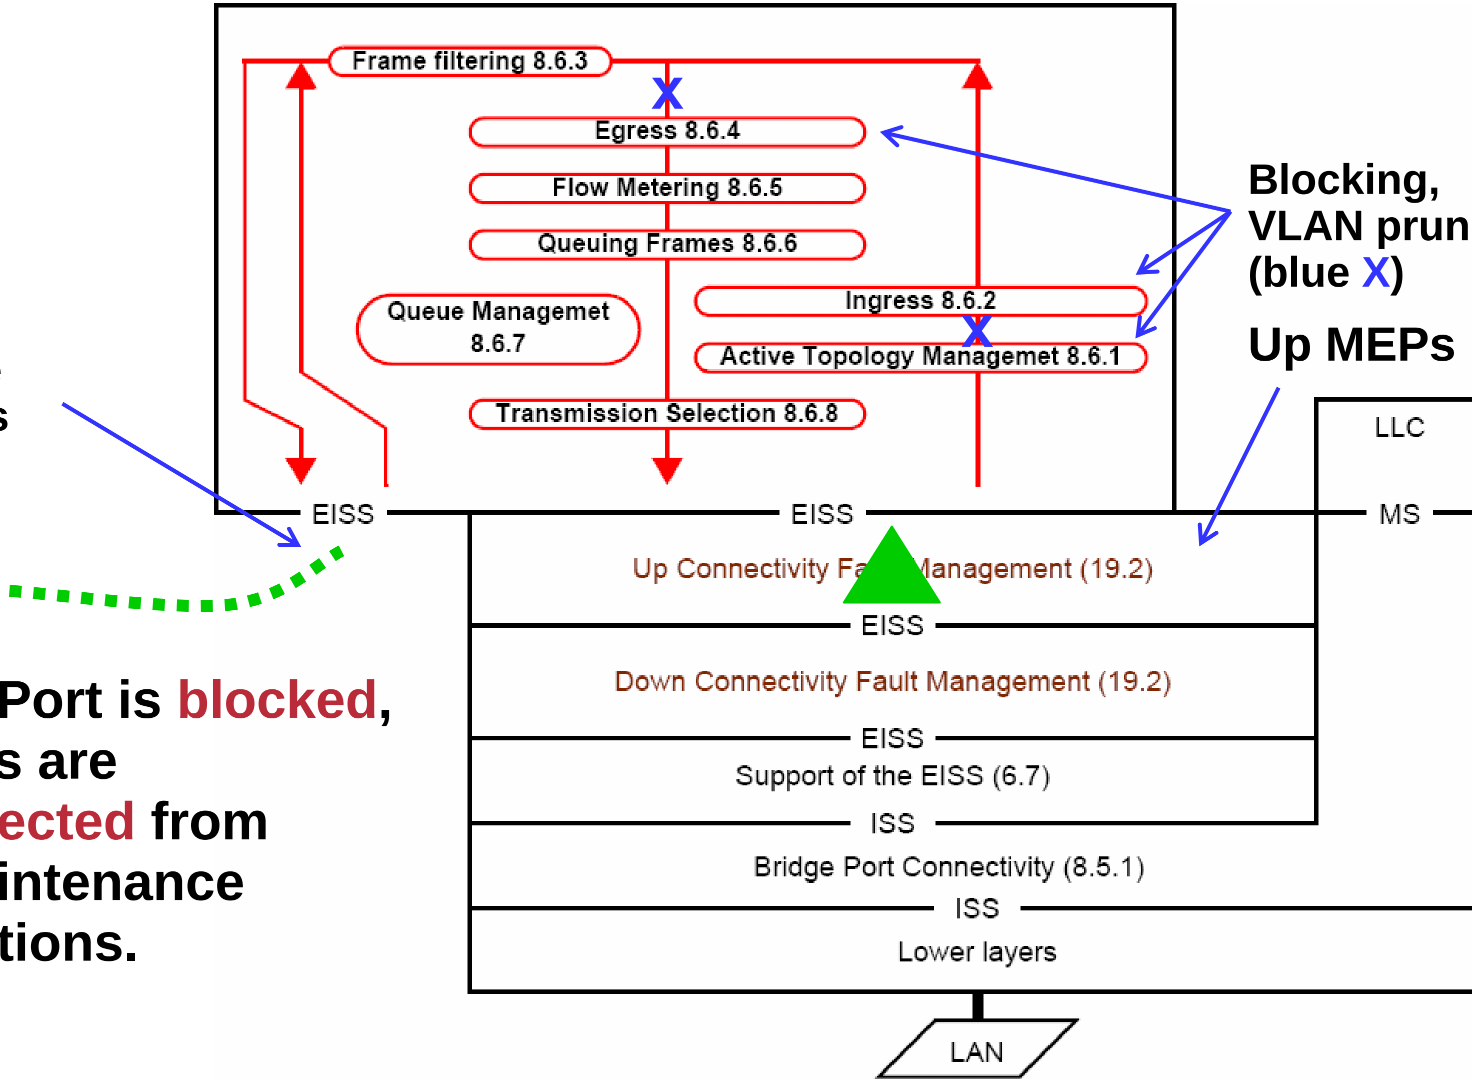

Blocking, VLAN pruning (blue X)

Up MEPs

1. When a Port is blocked, Up MEPs are disconnected from their Maintenance Associations.

The Proposed Solution: Step 1

- **Move per-port parts of the Forwarding Process down into the Bridge Port.**
- **This is a technical change from P802.1ag Draft 6.0.**

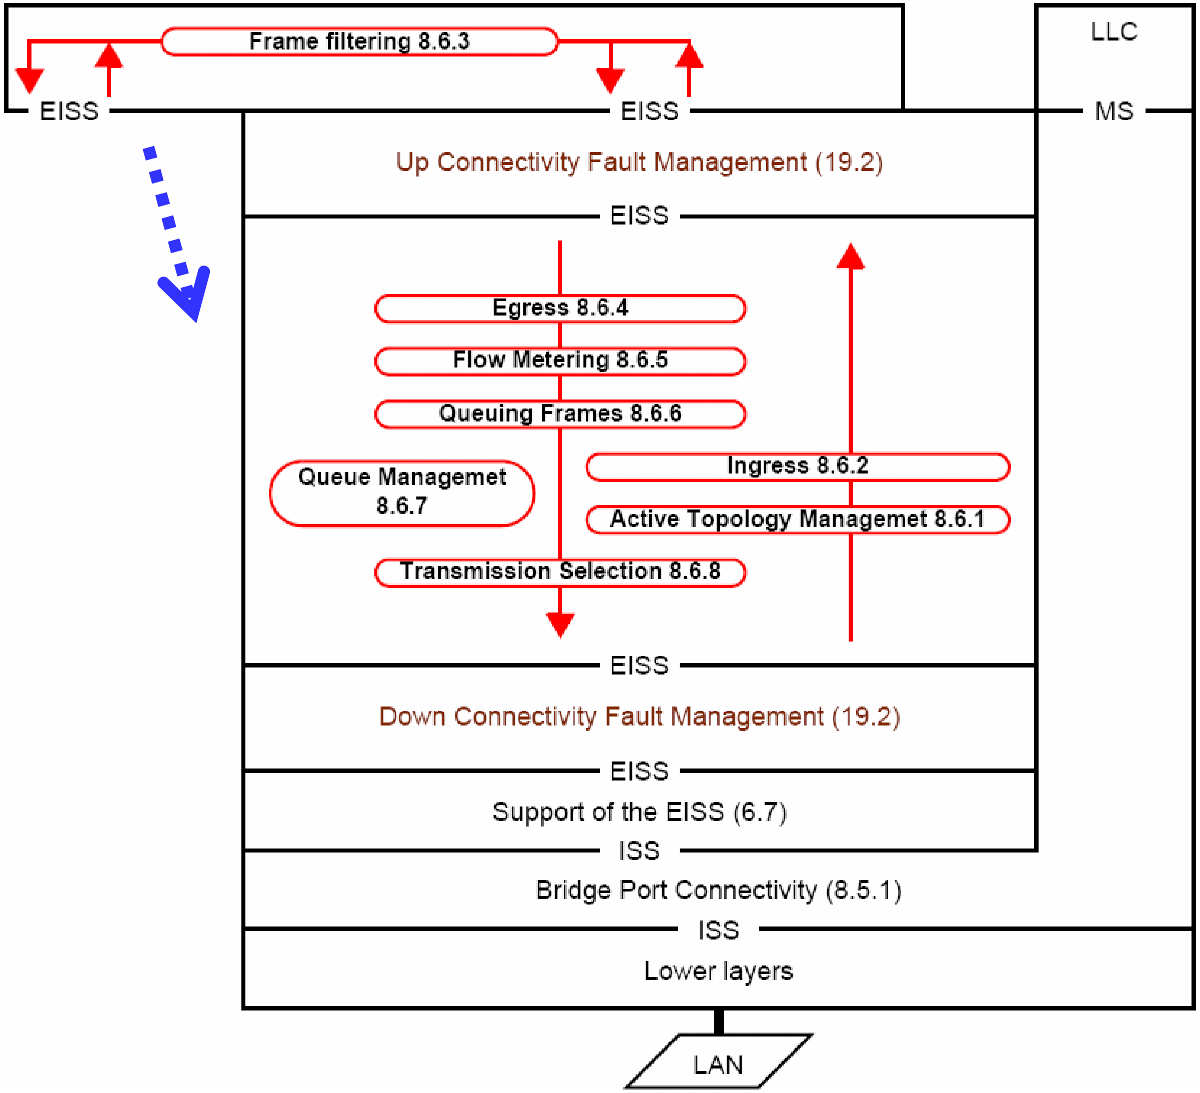

The Proposed Solution: Step 2

- **Separate** Queuing from ingress/egress filtering.
- **Move** Queuing down even further.

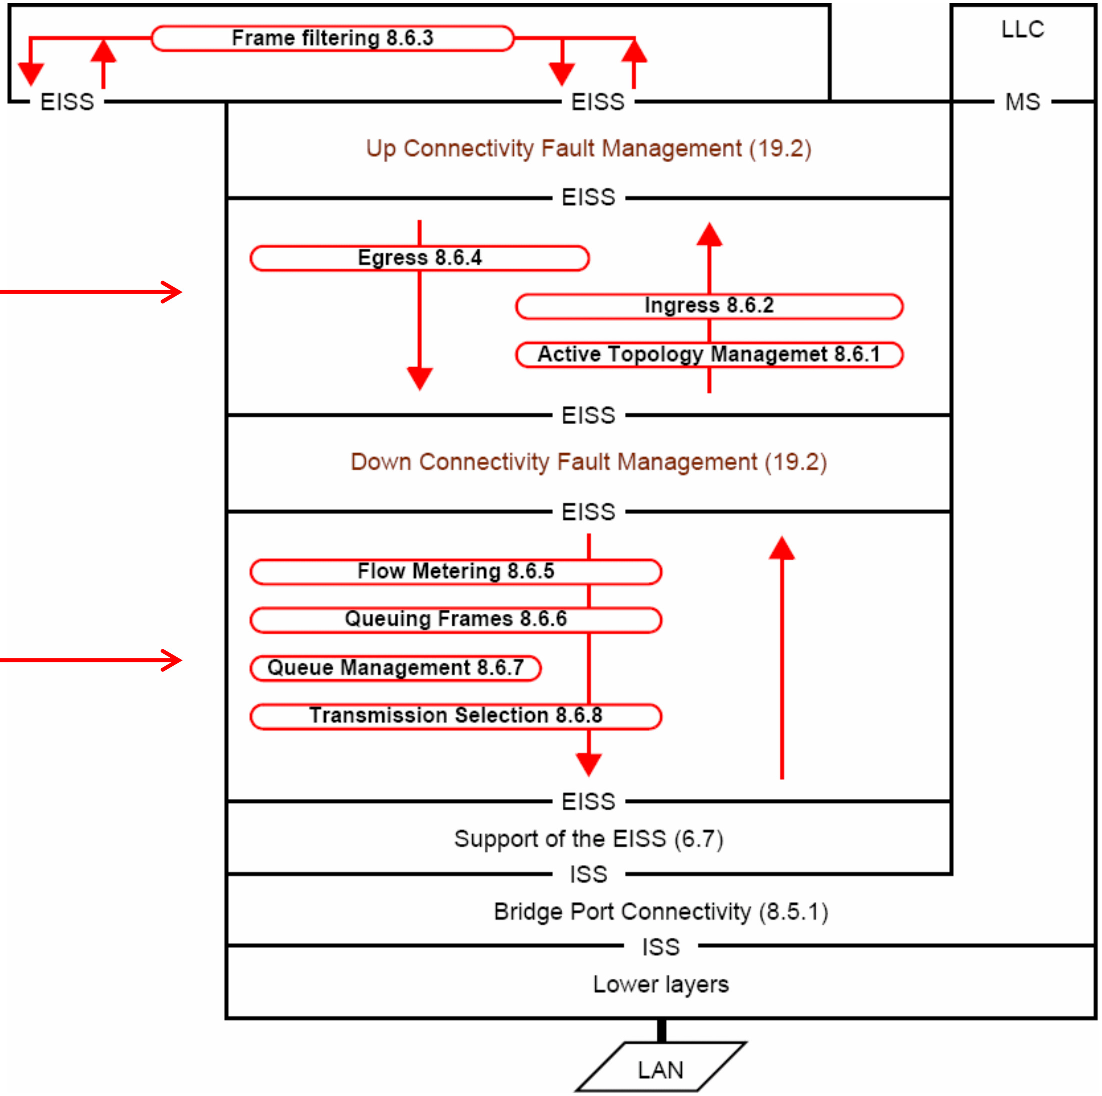

Problem 1 fixed

- **Down MEPs' Loopback path now goes through the queuing functions.**
- **Loopbacks' interference with other data is therefore controlled.**

- **(This Port's Up MEPs' Loopback paths are still not a problem; they still go through other Ports' output queues.)**

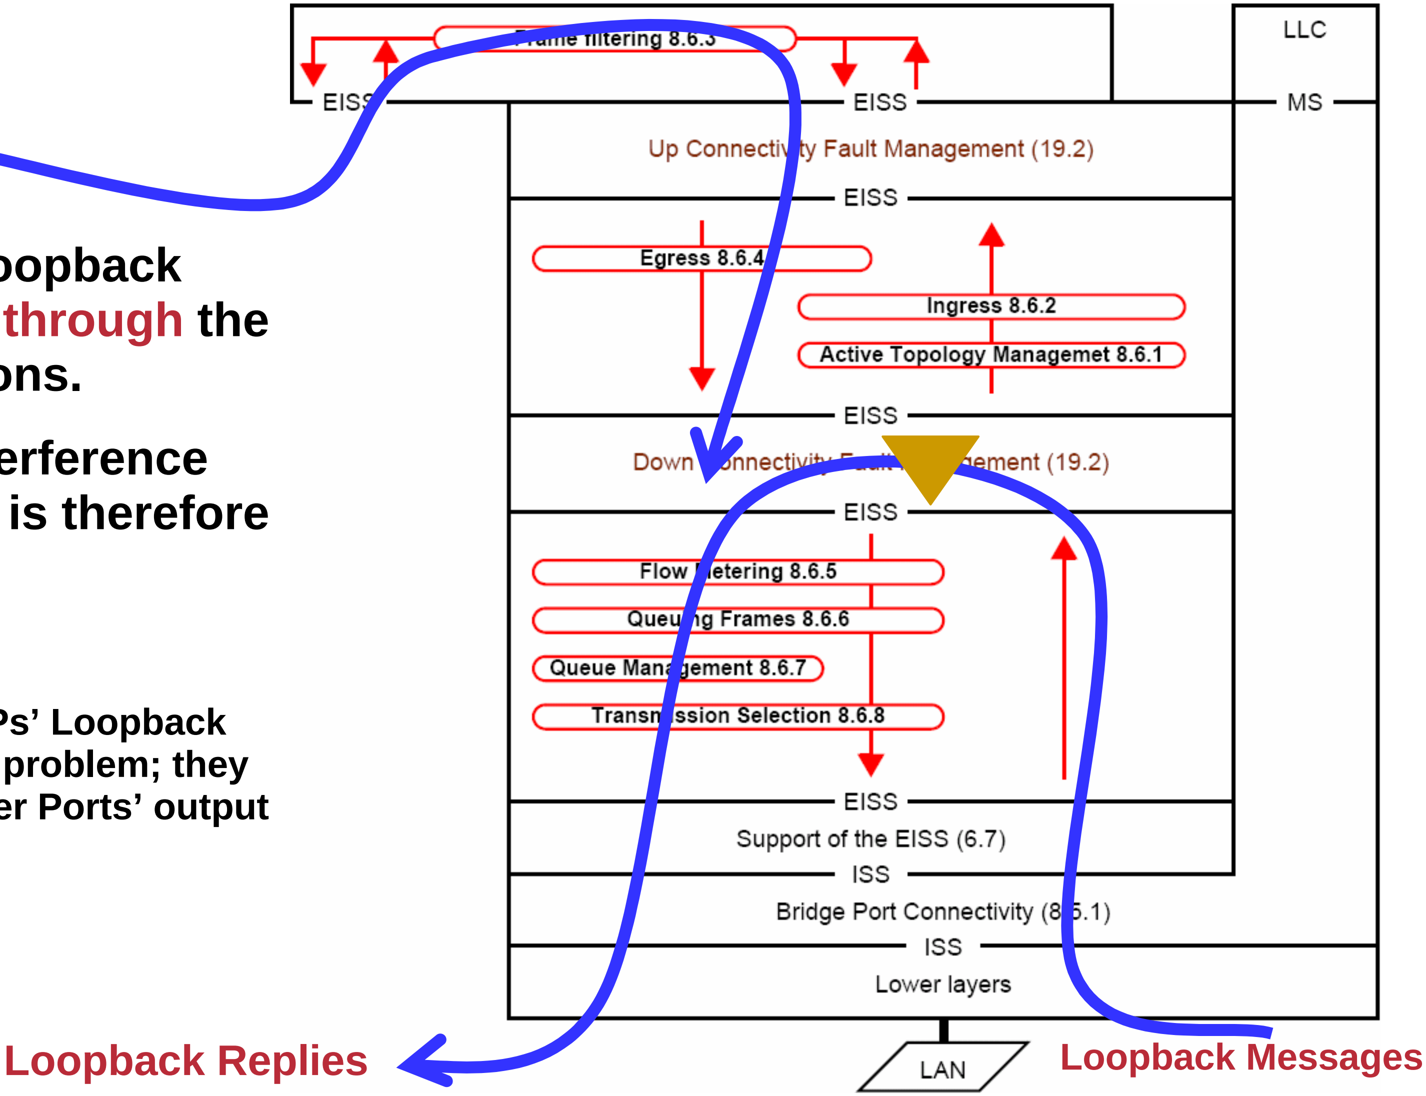

Problem 2 fixed

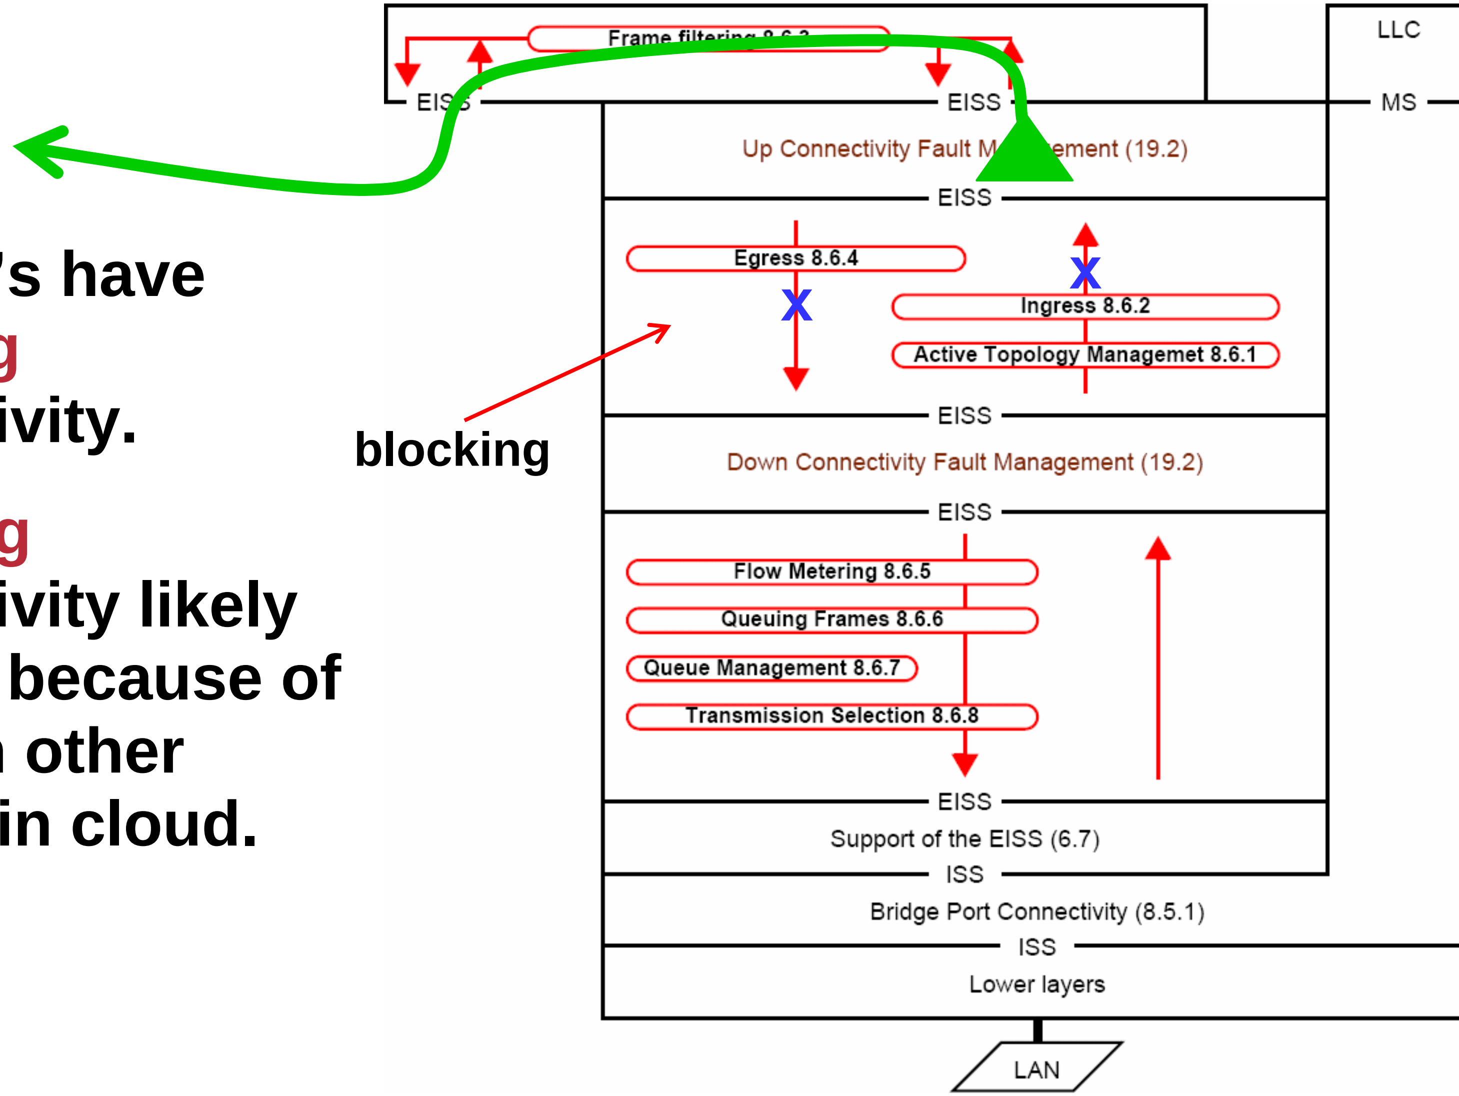

- Up MEP's have **outgoing** connectivity.
- **Incoming** connectivity likely **blocked** because of MVRP in other bridges in cloud.

Problem 2 fixed

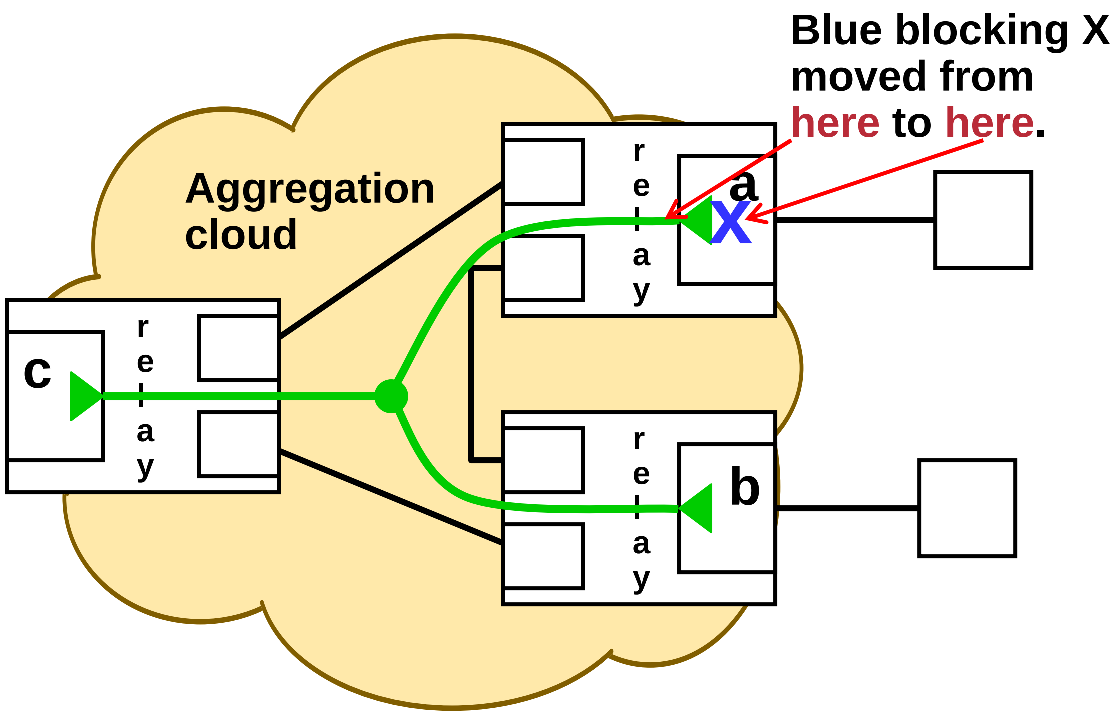

- CCMs from Up MEP a can be seen by b and c.
- CCMs from Up MEPs b and c (typically) cannot reach a, because of MVRP (or GVRP).

Problem 2 fixed

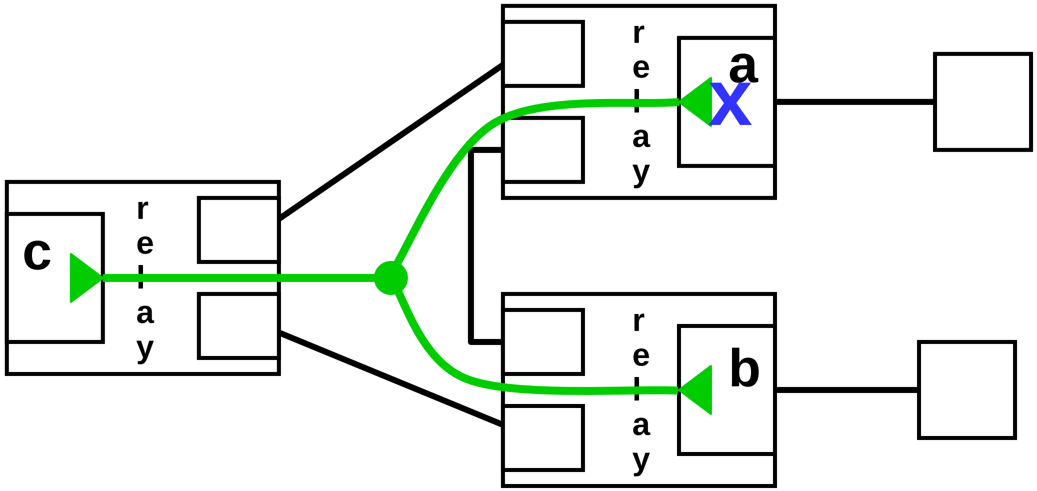

- So, MEP a’s CCMs advertise “I’m blocked”.
- MEP a’s CCM receiver is held in “happy” state.
- All MEPs treat “I’m blocked” just like “All OK”.
- If Port a goes Forwarding, its CCM receiver is released. Network has (3.5 * CCMInterval) seconds to deliver b’s and c’s CCMs to keep a happy.

Notes

- The network administrator is no longer in a “fool’s paradise” when a monitored port happens to be unused.
- Other ideas investigated, such as CFM passing through VLAN pruning, or through blocked ports, seem to lead to more problems.
- No further changes are required to the structure of Clauses 6 or 8, even in the final rolled-up 802.1Q.
- Whether the modified diagram of the Forwarding Process needs to be in Clause 8 is to be determined, but the text would need to change very little, if we did change the diagram.

Part 2: Miscellaneous technical issues

Other issues

- 17.2 Security: What is needed, here? 17.3 MIB: MIB needs to be revised. Please examine 12.14 (managed objects) and subclause 22.2 (creating MEPs and MHFs) for accuracy and completeness, so that future revisions of the MIB will be minimal.
- ifOperStatus: What is the relationship between the MAC_Operational status parameters of the two SAPs of a MEP? (Are inner/outer SAP the best names going forward for future shims?
- What status parameters are passed by EISS Multiplex Entity?
- Where is the DSAP? The ISAP?

ifOperStatus / MAC_Operational

- Draft 6.0 speaks of ifOperStatus, which is incorrect. MAC_Operational is the correct reference.
- Needs new control on Down MEP, passDefectPriority, that determines whether the MEP’s defect status controls MAC_Operational for the higher layers. Down MEP presents:

  upperSAP.MAC_Operational = lowerSAP.MAC_Operational
  & upperSAP.MAC_Enable &
  (highestDefectPri < passDefectPriority)

- Is this what Up MEPs should do?

  upperSAP.MAC_Operational = lowerSAP.MAC_Operational
  & upperSAP.MAC_Enable

EISS Multiplex Entity status parameters

- **Seems pretty clear that EISS Multiplex Entity going from one to many (as you pass up the stack) should present:**

  upperSAP.MAC_Operational = upperSAP.MAC_Enable & (OR of all lowerSAP.MAC_Operationals)

- **EISS Multiplex Entity going from many to one (as you pass up the stack) should present:**

  Each upperSAP.MAC_Operational =
  that upperSAP.MAC_Enable &
  the only lowerSAP.MAC_Operational

Where is the DSAP? The ISAP?

- Draft 6.0 attaches the DSAP to the MEP, and the ISAP to the MIP.
- Clause 18(.0) indicates that:

  The DSAP belongs to a Maintenance Domain.

  That DSAP can be assigned to a Service Instance.

  That Service Instance can be protected by a Maintenance Association (MA).

  A MEP belongs to the MA.

- If the DSAP exists before the MEP exists, it cannot be defined by the MEP.

Where is the DSAP? The ISAP?

- **First question: What is a Service Instance?**

  A Service Instance is an instance of the MAC Service.

  That implies two or more peer Service Access Points (SAPs).

- **If those SAPs are at the edge of a Maintenance Domain, they are Domain SAPs, or DSAPs.**

Where is the DSAP? The ISAP?

- Here is the whole shebang. Where are the Service Instances?

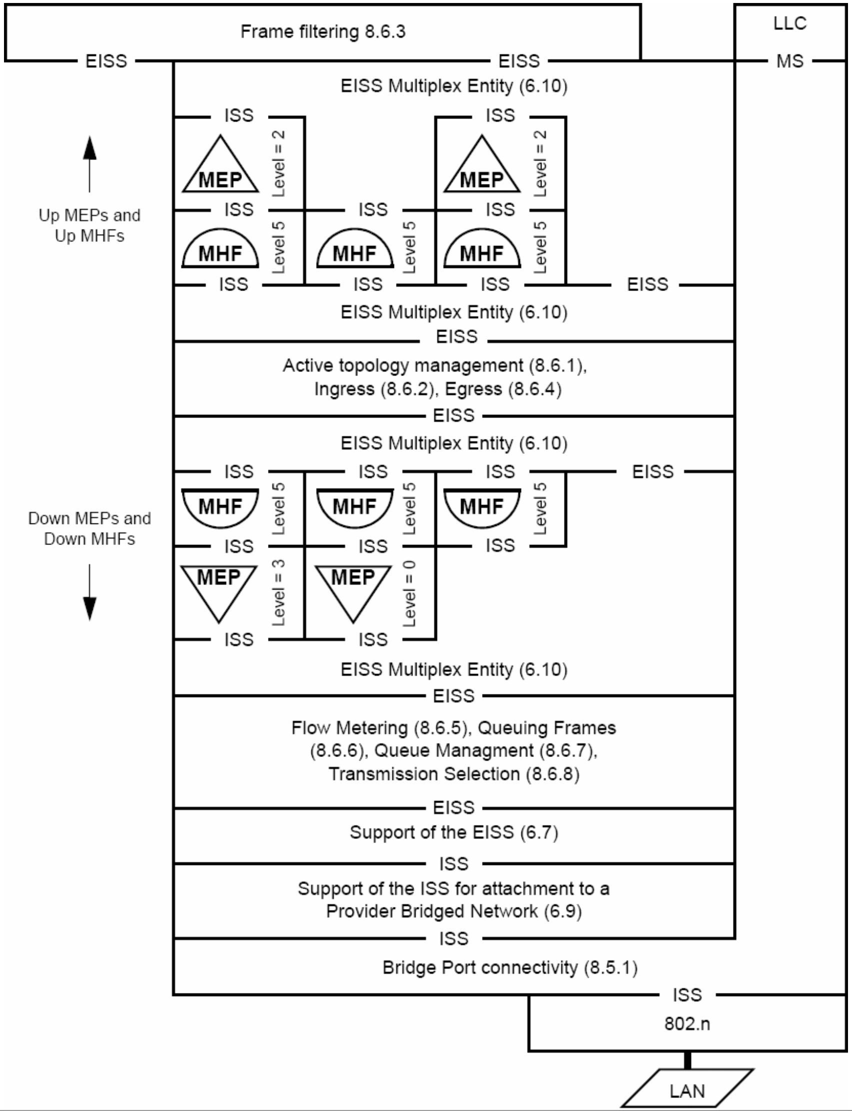

Where is the DSAP? The ISAP?

This service instance is offered to a customer on the LAN, and extends through the bridged network. This is the typical use of Service Instances.

- This is where single VLAN Service Instances can be said to start.

This service instance is offered to the Bridge, and extends into the network outside the Bridge. This might be one of a few Service Instances bought from a Provider.

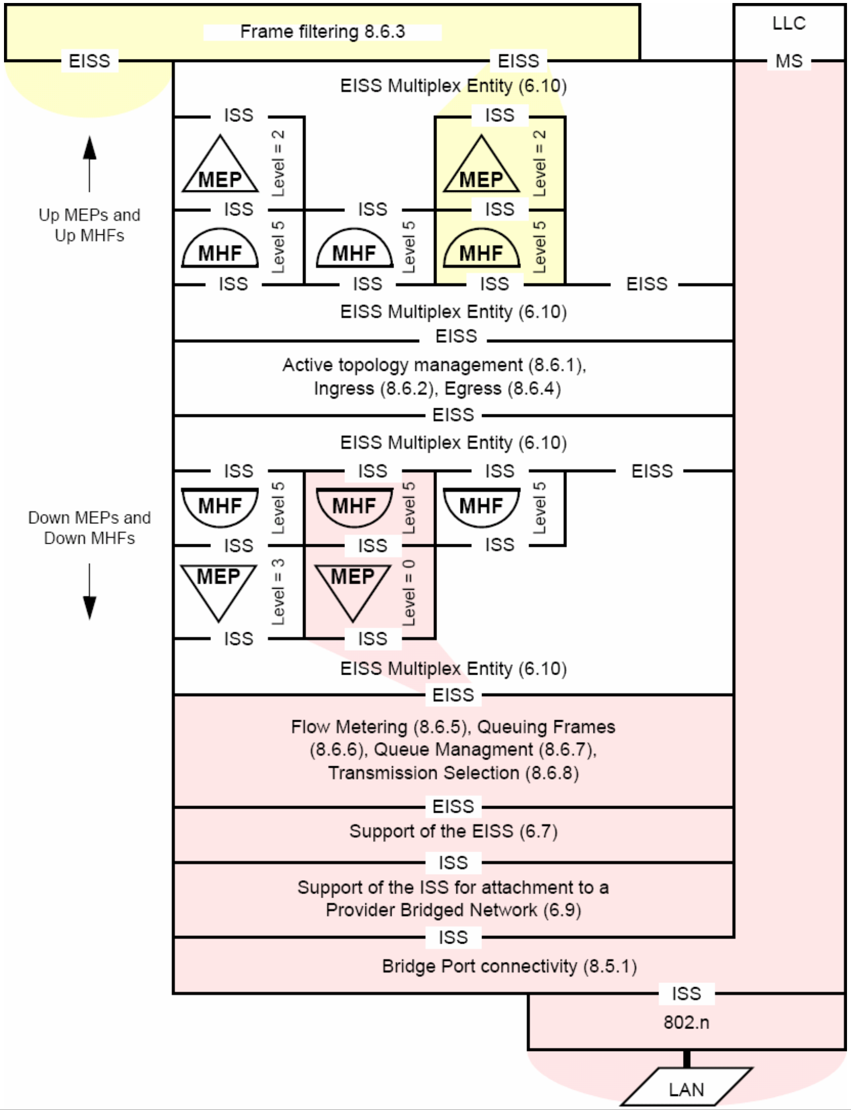

Where is the DSAP? The ISAP?

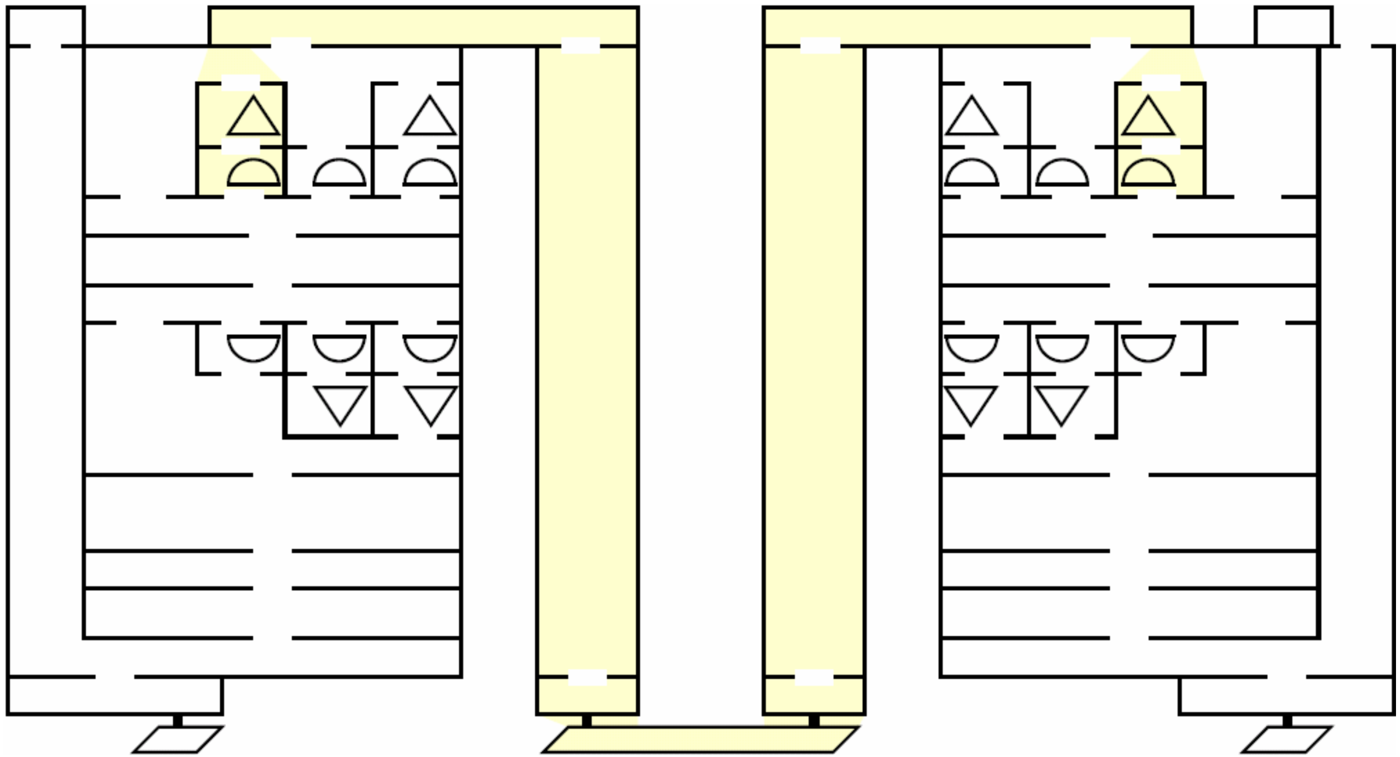

- **This is a typical VLAN-based Service Instance.**

This service instance is offered to a customer on the LAN, and extends through the bridged network. This is the typical use of Service Instances.

Where is the DSAP? The ISAP?

This service instance is passing through the Bridge, extending both into and out of this Port.

- This is a Service Instance (at MD Level 5) using the same VLAN as the yellow Service Instance in the previous slide.

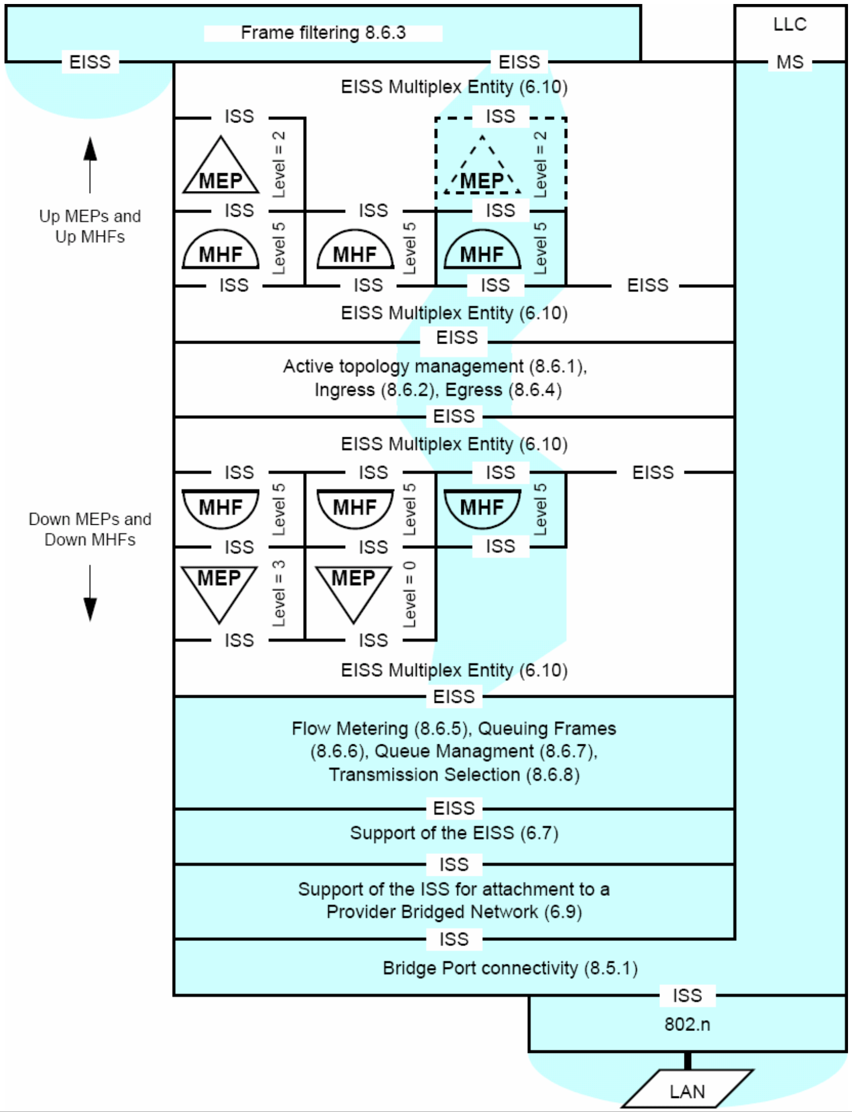

Where is the DSAP? The ISAP?

- There are different places to identify Service Instances extending outside the Bridge.

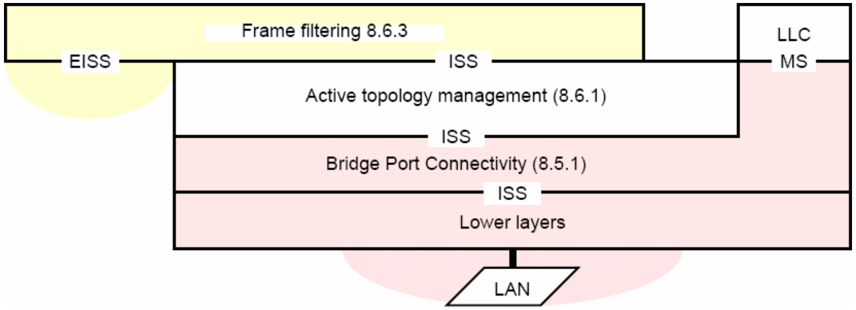

This Service Instance is offered is on a single VLAN.

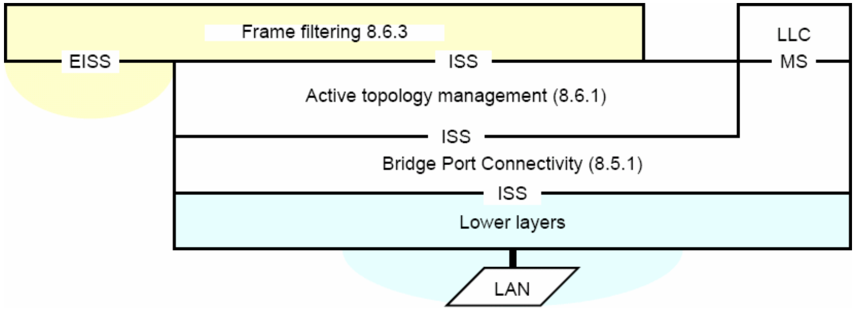

This Service Instance is a physical link.

Where is the DSAP? The ISAP?

- This is where the per-VLAN pass/don't pass decisions are made, so it seems DSAPs should be here, or adjacent to here.

- When the Service Instance is the whole LAN, this is the appropriate place.

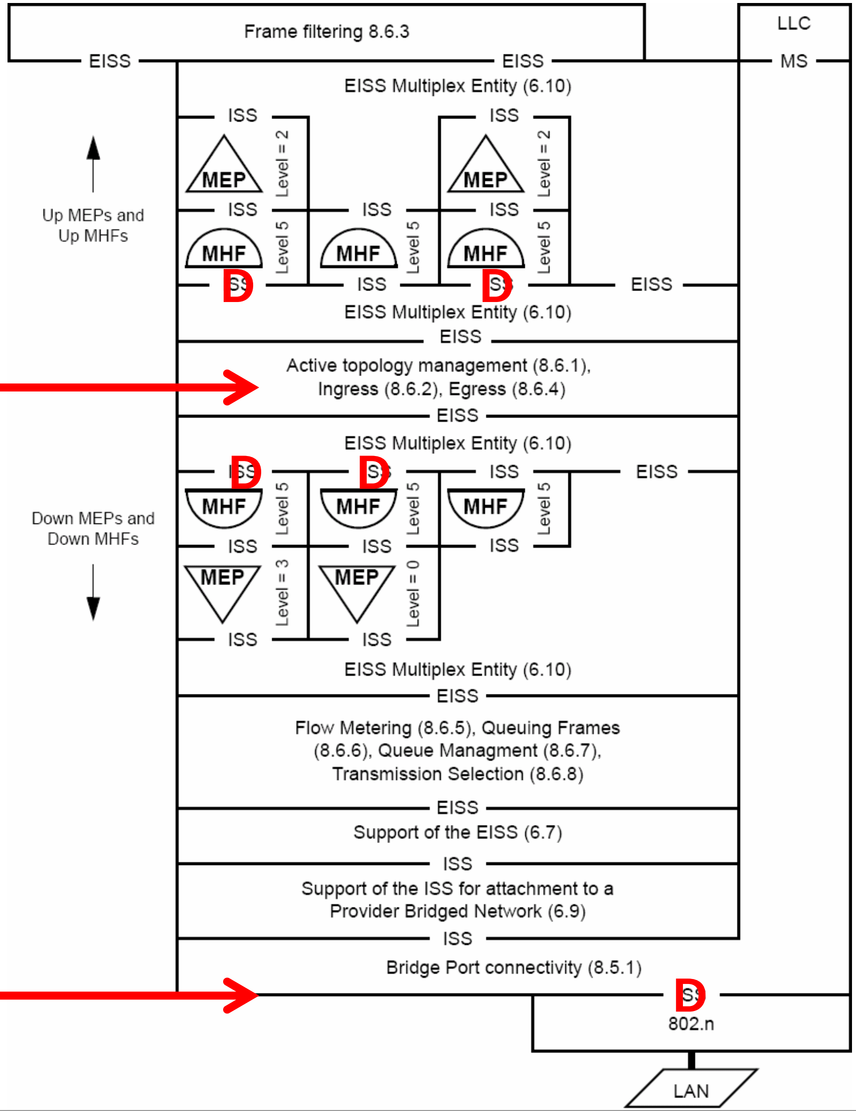

Where is the DSAP? The ISAP?

- **Apparently, there are two ISAPs on a Port, if both MHFs are on that Port; you just can't tell them apart, because they have the same MAC address.**