Maintenance Point Architecture

Mick Seaman

This note supports my task group ballot comments on P802.1ag/D3.0. In particular it provides positive suggestions for an improved, OSI terminology conformant architecture, including diagrams. Much of this is in sympathy with what is in D3.0, but I there are significant differences. This is a work in progress.

1. Bridge Architecture

Before discussing maintenance points this note clarifies and extends some of the terms and conventions used to describe the architecture of bridges and bridge networks.

1.1 Ports and Bridge Ports

The term ‘Port’ and ‘Bridge Port’ is currently used in two quite different ways in 802.1 standards.

First a ‘Port’ and particularly a ‘Bridge Port’ is a protocol or ‘interface’ stack. It supplies an entity above it with a service, given the service provided by communication media. In particular a Bridge Port provides the MAC Relay Entity with service. In this sense the term Bridge Port has also been used to include higher layer entities also attached to that interface stack — that usage should be deprecated.

Second a Port is used to mean a ‘service access point’ or SAP. This is the sense used in .1X and .1AE when referring to the Controlled and Uncontrolled Ports of .1X.

Confusion between these two uses was a natural consequence of the simple beginnings of the bridge architecture, where the interface stack presented a single service access point directly supported by the MAC. Sometimes it is easier not to be specific about whether the service provided by a SAP or the entire interface stack, i.e. the means of providing the service, is meant.

I recommend that Port only be used in the first, “interface stack”, sense in .1ag, and ‘service access point’ be used when the other meaning is explicitly required. Port should not be used to describe the internal operation of a CFM shim — we could do without a third meaning.

1.2 Bridge Architectural Diagrams

Architectural diagrams of bridges and bridged networks are conventionally drawn with the MAC Relay Entity below the level of the service provided. This serves to emphasize the near invisibility of the relay function that supports the provided service, as in Figure 1, where protocol entities C1 and C2 communicate over the service provided at service access points (SAPs) 1 and 2.

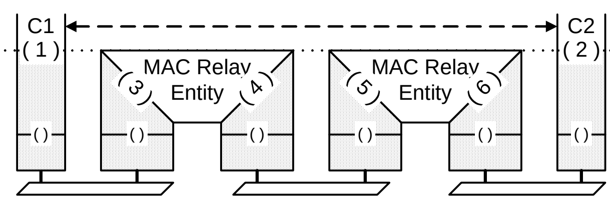

Figure 1—Bridged network architecture

Bridges, and hence the SAPs for the services — the ISS (Internal Sublayer Service) and EISS (Enhanced Internal Sublayer Service) — that they support, are not directly addressed by end stations communicating through the bridged network. The SAPs (1,2,3,4,5, and 6) shown in Figure 1 are all peers. An alternative depiction, Figure 2, aligns these, and proves convenient when the principle purpose of the diagram is to illustrate the layered interface stack provided by each Bridge Port†1. This is especially true if shims are to be added to the top of the stack.

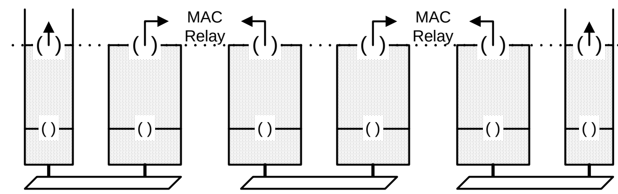

Figure 2—Service provided by Bridge Ports

2. Maintenance Point Entities

It is important to differentiate between a protocol entity within an interface stack and the functionality provided by that entity. Dynamic rearrangement of protocol stacks without service interruption is largely beyond our descriptive and implementation capabilities, so it is useful to describe a protocol entity, static and with fixed interfaces to the rest of the system, as a container for potentially dynamic functionality. The need for persistence of the protocol entity shows up most clearly — as do most things that might otherwise be considered internal properties of a system and independent of external behavior — when management is considered.

Any given bridge implementation can be deployed in a number of different places within a network, so it is useful to consider one or more Maintenance Point Entities (MPEs) as a fixed part of each Bridge Port, and to configure them with Maintenance End Point (MEP) or Maintenance Intermediate Point (MIP) functionality as required by its role within the network design.

Since an MPE is a shim, i.e. it can be inserted into an existing protocol stack, it necessarily provides the same service interface as it uses, so we need two types of MPE—one that operates between ISS access points, and one that operates between EISS access points. Whenever it is necessary to distinguish between the two, the latter will be referred to as an Enhanced Maintenance Point Entity (EMPE). Most of the CFM functions are common to both MPEs and EMPEs, so there is little duplication in their specification, and avoiding the need for the reader to extrapolate from one case to another increases the clarity of the specification.

3. Service access and maintenance points

An OSI service is provided at an abstract interface, the service access point, within a system, not at a point on a wire (the wire being incapable of any protocol operation in support of the provided service). The domain service access

†1Current examples include .1Q-REV Figure 8-8 and a number of the figures in .1AE Clause 11.

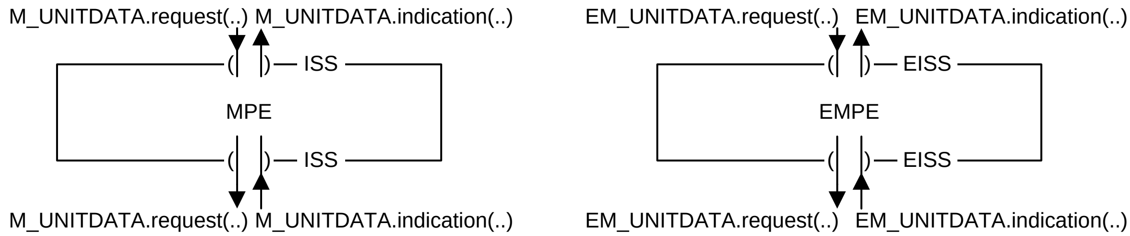

Figure 3—The two types of Maintenance Point Entity

point (DSAP) for a bridged service provided by a domain operator is illustrated, together with the first bridge supporting the service, in Figure 4.

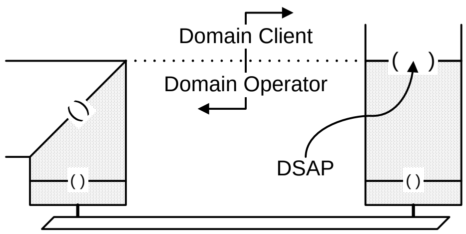

Figure 4—A Domain Service Access Point

The protocol entity using the service often resides in a system not under the control of the domain operator, so the operator cannot monitor service provision from the formal interface. Instead its first peer SAP within the domain that is selected as the domain service monitoring point (DSMP†1). A peer access point is chosen to allow a maintenance point entity supporting the DSMP access to the same service primitives and parameters as were present at the DSAP†2.

Connectivity between the DSMP and the DSAP can be monitored by using a MIP associated with the DSMP (see below) or separately as its own, inferior, domain. Although that domain spans administrations, the extent of the domain is constrained so that neither has to cede control or reveal private information to the other. See below.

Figure 5 shows DSAPs, DSMPs, and internal service access points (ISAPs) in the style of Figure 2, emphasizing the functionality of the Bridge Ports and the peer relationships amongst the SAPs they support.

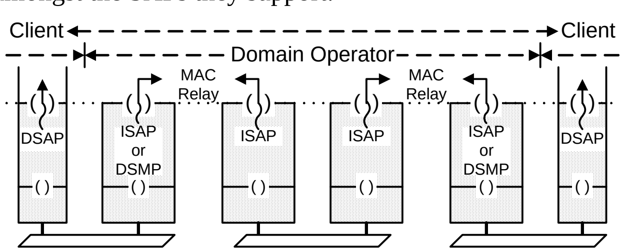

Figure 5—SAPs within a domain

†1I had thought to call this the 'proxy DSAP' but suspect that giving it a clearly distinct name will help when we are trying to be precise about the functionality within MPEs that support DSAPs and DSMPs.

†2Inevitably there will be those who argue that the DSMP should be 'lower down' in the operator's bridge port. The architectural point is that access to exactly the same parameters denotes a peer interface 'at the same height'. Those who fear internal breakage in the operator's bridge port should look to their implementation, not to the standard which has the job of clearly expressing externally observable consequences, not of constraining implementation. It follows that if unnatural descriptive acts in the standard appear to improve fault coverage, any implementation could also improve it without bending the standard around implementation constraints.

4. Domain levels

Figure 5 serves to introduce a point that is to easy to overlook, with the possibility of introducing endless confusion†3. A given domain, that is a connectivity association between peer SAPs, can be monitored both by its operator and its client. These monitoring activities are separated by associating a different domain level with each. The important point is that the 'domain level' is a property of the monitoring activity, not of the service provided or the equipment providing the service. A bridge, or a protocol layer in a bridge port, has no need to know its domain level — that is information particular to the Maintenance Point Entities added to the port.

5. Adding Maintenance Point Entities

Where should MPEs be added to the Bridge Ports shown in Figure 5—above or below the SAPs. The answer is simple, the user or Domain Client should add and/or be served by MPEs above the domain service boundary, while the service provider or Domain Operator should add and be served by MPEs below, i.e. on his side of, the boundary.

Figure 6 shows the MPEs for one of the client systems and part of the domain, each labelled with the functionality (MEP or MIP) that it provides†4. The null block, with the pass through function, that supports the DSAP to the left of the figure is included in the architectural diagram simply to align equivalent service points. The client system is not even aware of it.

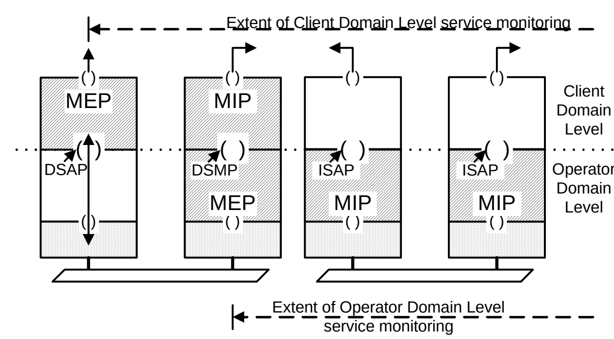

Figure 6—Adding Maintenance Point Entities

MIP functionality is bidirectional between its upper and lower SAPs, i.e. there is no difference in the specification required for the MIPs above or below DSMPs or ISAPs. Some MEP functions transmit messages into the domain,

†3Symptoms of the confusion include attempting to hold information or transmit and receive messages for the immediately superior domain in a MEP or MIP at any given domain level.

†4Every MPE that supports an ISAP provides MIP functionality, an MPE that uses a DSAP provides MEP functionality, an MPE that supports a DSMP provides MEP functionality. An MPE supporting a DSAP provides null functionality, an MPE using an ISAP provides null functionality.

receiving messages from other MIPs or MEPs within or attached to the same domain, while other functions signal out of the domain to the attached users. The way in which MEP functions are placed within an MPE therefore differs depending on whether that MPE uses the service provided at a DSAP or supports that provided at a DSMP. These internal details will be described after the actual functions, and after additional scenarios for MPE placement.

6. Concatenated networks

A service provider may contract for two operators to provide concatenated connectivity between service creation systems that the provider controls. Figure 7 shows one possible arrangement of physical equipment, for operators conveniently named L(eft) and R(ight).

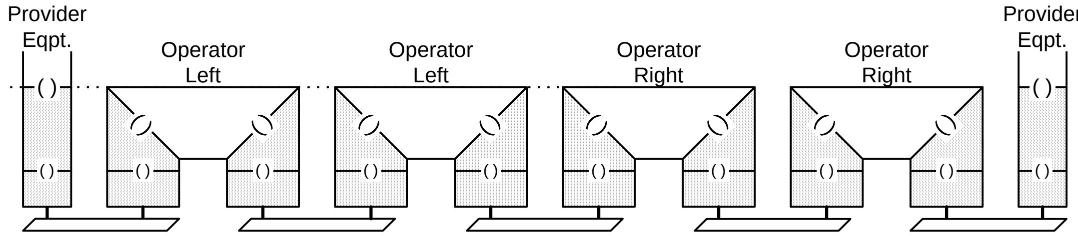

Figure 7—Concatenated networks

Figure 8 shows the Bridge and Provider Equipment Ports with MPEs configured at appropriate domain levels.

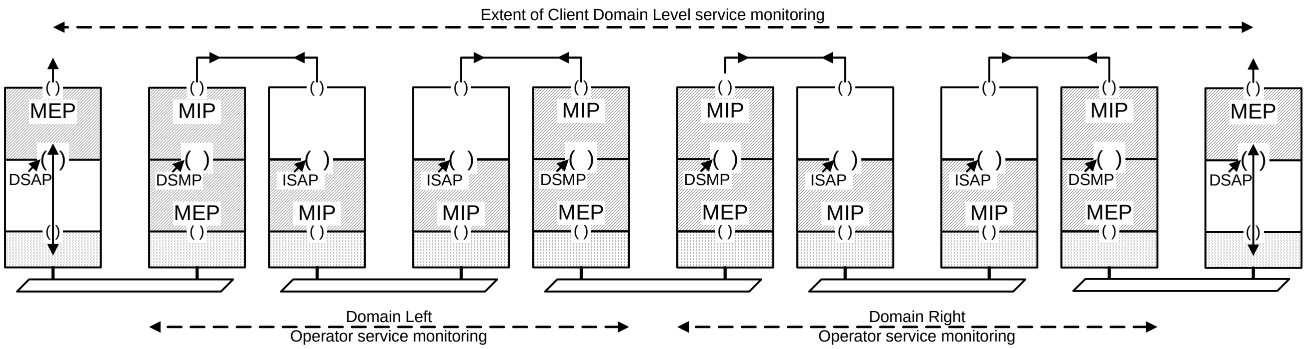

Figure 8—MPEs in concatenated networks

Figure 9 shows that we are not done yet. Left and Right have organized a small domain between them so that they can check that they have cross-connected the service correctly between them, and not simply crossed wires. From an architectural point of view the newly introduced inferior domain is supporting the service used by their two MEPs, but it can of course identify each of the VLANs supported (if the lowest SAP shown is an EISS-SAP).

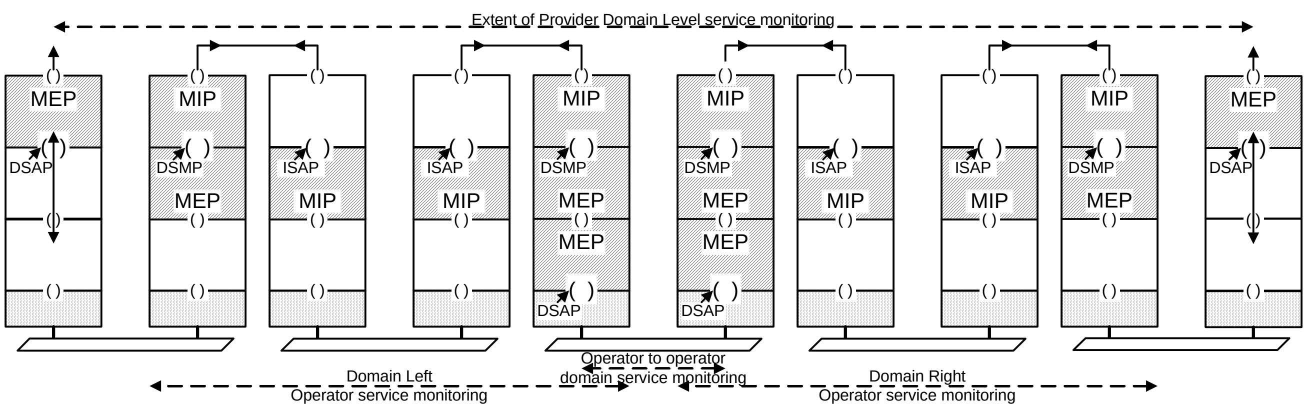

Figure 9—MPEs in concatenated networks

So far Bridge Ports have been shown without any content other than MPEs. To be clear about the operator to operator service monitoring domain in Figure 9, it is necessary to show where support of the EISS by the ISS (802.1Q clause 6.7) occurs. Figure 10 shows the full Bridge Port for each of the VLAN-aware Bridges supporting the provider's end to end service.

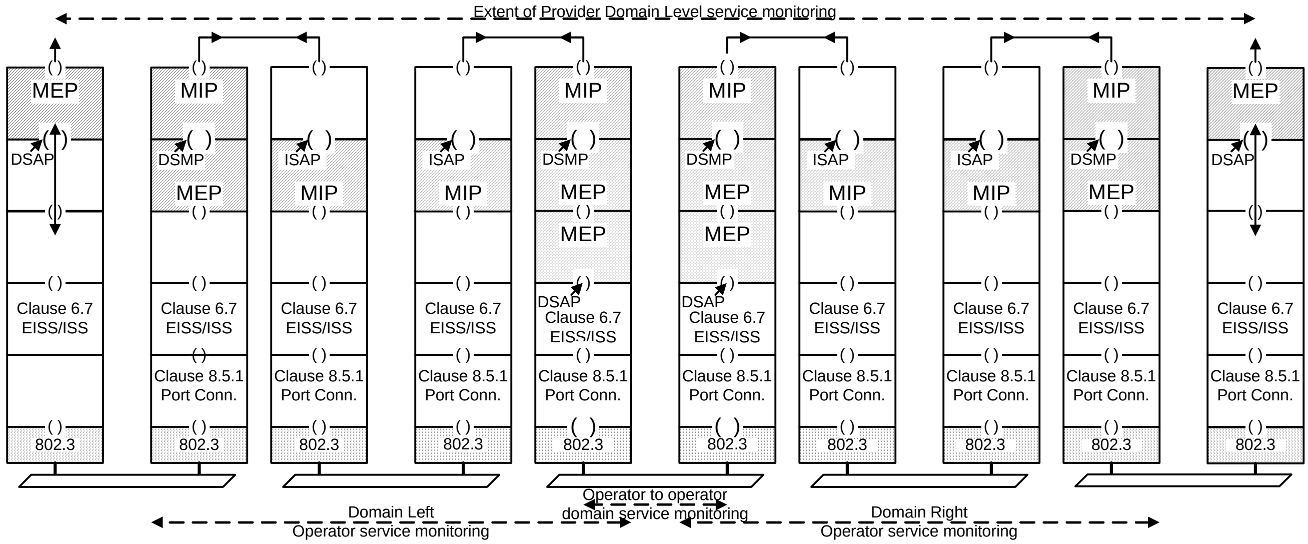

Figure 10—Concatenated networks showing complete Bridge Port functionality

Figure 11 shows an alternative, where the operator to operator monitoring is at the level of the ISS, effectively monitoring the 'trunk' connection between the two operators, but unable to provide information on the connectivity of a single service instance supported for the network provider.

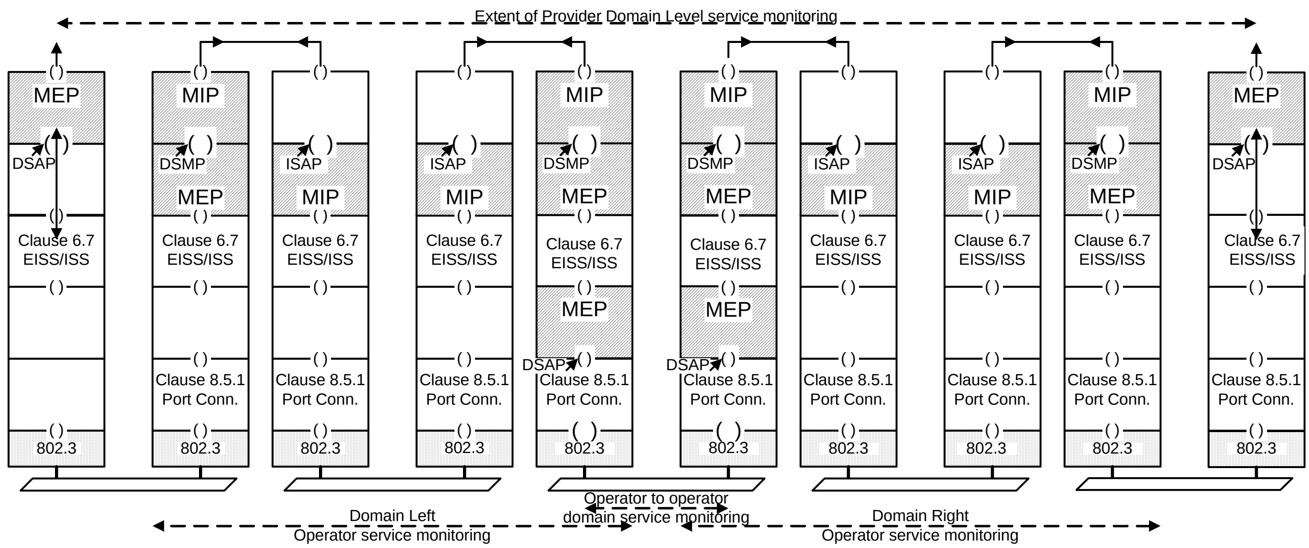

Figure 11—Concatenated networks showing complete Bridge Port functionality