Q-2022-Rev update

Status, activity, progress to date, brief report[1]

Mick Seaman
mickseaman@gmail.com

1. See https://www.ieee802.org/1/files/public/docs2025/x-seaman-mka-optimizations-0325-v04.pdf for detail. This presentation follows that note almost exactly.

Status, activity

A simple rollup of IEEE Std 802.1Q with all approved and published amendments is available:

https://www.ieee802.org/1/files/private/q-2022-rev-drafts/d1/802-1Q-2022-Rev-d1-6.pdf

Ongoing effort to align Introduction, Conformance, and PICS:

- Following up Q-2018-Rev comment on PICS inconsistency.
- Current standard only understandable by following the history, amendment by amendment, each over-writing, sometimes duplicating existing provisions.
- PICS basically for an 802.1D Bridge plus bolt-ons, independent of the structure of conformant systems.
- Many contradictions to make changes easy for amendments, each considered to be more important than any existing provisions.

Revision goals:

- Make standard understandable for new readers.
- Make standard maintainable.

Progress

Revised 1.3 Introduction:

- Added subheadings so list items can be added or removed without renumbering all subsequent unrelated items.
- Grouped topics,[1] rather than ordering by time of change.
- Added references to the amendments that introduced or changed items.[2]
- Added editor's comments re changes. These are conditional text so can be shown/not shown at the flick of a switch prior to book/PDF rebuild.
- First cut of 1.3 revision follows these slides.

No changes to 5. Conformance

- We may want to come back and reorder, but avoiding shaking up everything at the same time.

PICS restructuring in progress

- Organizing around 5.2 conformant systems and the components they contain, plus supporting interfaces and component options.

---

1. In some cases I may be the only 802.1 member, with the possible exception of Steve Haddock, who can remember the original logic.
2. It could be argued that the new reader does not need this information, but I believe it is required for maintenance and any informed discussion of potential change. I did not set out to incorporate the historical record in this way, but found it vital for checking myself. I have both correlated existing items with amendments, and checked through every amendment from 802.1Q–2005 on to make sure it is represented.

IEEE Standard for
Local and Metropolitan Area Networks—

Bridges and Bridged Networks

1. Overview

IEEE 802® Local Area Networks (LANs, 3.119)[6] of all types can be connected together with Media Access Control (MAC) Bridges (3.159) or Virtual Local Area Network (VLAN) Bridges (3.306), collectively known as Bridges (3.28). This standard defines the operation of Bridges and Bridged Networks. VLANs facilitate the administration of logical groups of stations. Stations in the same VLAN communicate as if they were on the same LAN, while traffic between VLANs is restricted. Management of VLAN Bridges and stations allows stations to be added to, removed from, or moved between VLANs.

This standard further extends the specification of VLAN Bridges to enable a service provider organization to use a common infrastructure of Bridges and LANs to offer the equivalent of separate LANs~~, Bridged,~~ or ~~Virtual~~ Bridged Networks to independent customer organizations.

The above update was missed when 1Q title changed from original 1998 and 2005 “Virtual Bridged Area networks”. Note however that this second paragraph of 1. and the one following duplicate information in 1.3, so there is an argument for removing both of them. We don’t want to end up with the whole of 1.3 transferred back into this introductory text. On the other hand I am not aiming for perfection here, only ease of future maintenance and clarity for the new reader.

This standard specifies protocols and protocol entities within the architecture of Bridges that provide capabilities for detecting, verifying, and isolating connectivity failures in Bridged Networks. These capabilities can be used in networks operated by multiple independent organizations, each with restricted management access to each other’s equipment.

1.1 Scope

This standard specifies Bridges that interconnect individual LANs, each supporting the IEEE 802 MAC Service using a different or identical media access control method, to provide Bridged Networks and VLANs.

1.2 Purpose

Bridges, as specified by this standard, allow the compatible interconnection of information technology equipment attached to separate individual LANs.

[6] IEEE and IEEE 802 are registered trademarks in the U.S. Patent & Trademark Office, owned by The Institute of Electrical and Electronics Engineers, Incorporated.

1.3 Introduction

For the purpose of compatible interconnection of information technology equipment using the IEEE 802 MAC Service supported by interconnected IEEE 802 standard LANs using different or identical media access control methods, this standard specifies the operation of VLAN-unaware MAC Bridges and VLAN Bridges and requirements to be satisfied by equipment claiming conformance[7] to this standard. It further specifies requirements for configuration and management of Bridges and Bridged Networks, and defines supporting SMIv2 (IETF STD 58)[8] Management Information Based (MIB) modules, YANG configuration and operational state models,[9] and TLVs for use with IEEE Std 802.1AB.[10]

Management, MIBs, and YANG were referenced (or not) in a haphazard way in 1.3 by the various 802.1Q amendments. The inserted text above should cover all those statements and remove the need for repetitious verbiage.As a consequence some feature specific references to MIB and YANG have been removed from the remainder of 1.3.

The very long list in 1.3 that follows this brief introduction was impossible to manage. Any addition or removal would renumber many referenced items. Now broken up into logical groups, 1..3.1 ... to aid understanding, but mainly to make maintenance a practical proposition.

1.3.1 Bridge and Bridged Network architecture and operation[11]

This standard:~~To this end, it:~~

a) Positions the support of the MAC Service ~~VLANs~~ within an architectural description of the MAC Sublayer.

b) Specifies VLANs, supporting of multiple instances of the MAC Service over a common Bridged Network of LANs interconnected by Bridges.

c) Defines the principles of operation of ~~the~~ MAC Bridges and VLAN Bridges in terms of the support and preservation of the MAC Service, and the maintenance of quality of service (QoS).

d) Specifies an Enhanced Internal Sublayer Service (EISS) provided to the Media Access-Independent functions that provide frame relay in a VLAN Bridge.[12]

e) Establishes the principles and a model of Virtual Bridged Network operation.

f) Identifies the functions to be performed by Bridges, and provides an architectural model of the operation of a Bridge in terms of processes and entities that provide those functions.

g) Specifies reserved multicast address that are filtered by MAC Bridges and VLAN Bridges to facilitate discovery of, and communication with, their nearest neighbours.

h) Specifies a frame format that allows a VLAN Identifier (VID) and priority information to be carried by VLAN-tagged user data frames.

i) Specifies the rules that govern the addition or removal of VLAN tags to and from user data frames.

There is no 1.3 mention of frame classification into a VLAN, as introduced in 802.1v and now covered in 6.9 and 6.12.

j) ~~Establishes the requirements for VLAN Bridge Management in a Virtual Bridged Network, identifying managed objects and defining management operations.~~

k) ~~Defines SMIv2 (IETF STD 58)[7] Management Information Based (MIB) modules for the management of VLAN Bridge capabilities including spanning tree protocols and Provider Bridges.~~

[7] Clause 5.

[8] Information on references can be found in Clause 2.

[9] Clauses 17 and 48. Modules include their revision history. Augmented by each amendment.

[10] Annex D.

[11] Clauses 6, 7, 8, 9, and 11. Note that earlier revisions of this standard relied on provisions of IEEE Std 802.1D (now withdrawn) now incorporated in this standard, so clause order does not strictly reflect logical and descriptive dependency. Clauses 6 and 8 in particular also integrate provisions that support later amendments.The informative Annexes provide supporting information.

[12] The Internal Sublayer Service (ISS), which augments the specification of the MAC Service with elements necessary for MAC Bridge relay between LANs and supports the EISS in VLAN Bridges, is specified in IEEE Std 802.1AC.

l) ~~Define YANG configuration and operational state models (Clause 48) in support of Two-Port MAC Relays, Customer VLAN Bridges, and Provider Bridges, including Connectivity Fault Management (CFM) for those Bridges.~~

j), k), l) removed as part of moving generalized management statements to the first paragraph of 1.3. Note that at this point in the introduction there has been no hint as to what a Two-Port MAC Relay might be, nor any indication of the particular nature of any of the other items mentioned in l).

1.3.2 Automatic configuration of Bridge Network and VLAN topology[13]

This standard:

a) Establishes the requirements for automatic configuration of Bridged Network and VLAN topology.

b) Defines the operation of ~~the Multiple Spanning Tree Algorithm and Protocol (MSTP)~~ interoperable algorithms and protocols that provide simple and full connectivity throughout a Bridged Network comprising arbitrary interconnected Bridges and consistently assign frames for any given VLAN to a spanning tree active topology.

c) ~~Describes the protocols and procedures necessary to support interoperation between Multiple Spanning Tree (MST) and Single Spanning Tree (SST) Bridges in the same Virtual Bridged Networks.~~

d) ~~Specifies the requirements to be satisfied by equipment claiming conformance to this standard.~~

RSTP, MSTP, and SPB are defined. The deleted text assumes a detailed knowledge of the standard, not appropriate for an introduction.

This standard specifies protocols, procedures, and managed objects to support the Multiple Registration Protocol (MRP). MRP allows participants in an MRP Application to declare and register attributes with ~~other~~ neighbouring participants in a Bridged Network, the set of Four applications are defined—one to register VIDs [Multiple VLAN Registration Protocol (MVRP)], one to register MAC addresses [Multiple MAC Registration Protocol (MMRP)], one to register Streams and configure associated network resources [Multiple Stream Registration Protocol (MSRP)]~~, and one that provides the ability to flush learned MAC Address Entries held in the Filtering Database (FDB) of an I-component on a per-I-SID basis [Multiple I-SID Registration Protocol (MIRP)]~~. MVRP will furthermore provide for the rapid healing of network failures without interrupting services to unaffected VLANs. To this end, it specifies the following:

e) MRP and the operation of MRP entities.[14]

f) The generic frame formats used in MRP exchanges.

g) The MMRP application of MRP, and the frame formats that it uses.

h) The MVRP application of MRP, and the frame formats that it uses.

MIRP reference moved to the Provider Backbone Networks clause below. In the present clause there is no reason why the reader should know about ISIDs.

[13] Clause 7 provides an overview of the physical network topology , active topology, and VLAN topology. Clauses 13 and 14 specify the Rapid Spanning Tree Algorithm and Protocol (RSTP) and the Multiple Spanning Tree Algorithm and Protocol (MSTP) and their interoperability with Shortest Path Bridging (SPB). Clause 8.9 describes the consistent allocation of VIDs to spanning trees, shortest path trees, or Ethernet Switched Paths. Clauses 10 and 11 specify MRP, MMRP, and MVRP.

[14] MRP (Clause 10), added by IEEE Std 802.1ak-2007, replaced the Generic Attribute Registration Protocol (GARP), defined in IEEE Std 802.1D™-2004 [B14], that was used to support GVRP and GMRP in earlier revisions of IEEE Std 802.1Q. Similarly, GVRP and GMRP were replaced by MVRP (11.2) and MMRP (10.9–10.12), respectively.

1.3.3 Provider Bridges and Provider Bridged Networks[15]

To enable a service provider to use a Virtual Bridged Network to provide separate instances of the IEEE 802 MAC Service, MAC Internal Sublayer Service (ISS), and EISS to multiple independent customers, in a manner that does not require cooperation among the customers and that requires a minimum of cooperation between the customers and the provider of the MAC Service, this standard further specifies the operation of Provider Bridges. To this end, it:

a) Differentiates Customer VLANs (C-VLANs) that are under the administrative control of a single customer of a service provider, from the Service VLANs (S-VLANs) that are used by a service provider to support different customers.

b) Specifies VLAN tag formats for both C-VLANs and S-VLANs, allowing each to be distinguished and separately applied and administered by customers and by a service provider.

c) Specifies the functionality of a generic VLAN Bridge component within a system and the specific requirements of derived C-VLAN and S-VLAN components.

d) Specifies a C-VLAN Bridge as comprising a single C-VLAN component, and a Provider Bridge as encompassing Bridges that comprise a single S-VLAN component and no C-VLAN components (S-VLAN Bridge) or a single S-VLAN component and one or more C-VLAN components (Provider Edge Bridge).

e) Specifies parameters and mappings that allow the EISS to support traffic classes that comprise distinct aggregate flows supporting different QoS characteristics and provide independent guarantees to different customers, through support of priority and drop precedence marking.

f) Specifies the incorporation of flow metering, transmission queue management, and transmission selection algorithms within the forwarding process of a Bridge.

g) Positions the support of S-VLANs within the architectural description of the MAC Sublayer and specifies their relationship to media access method-dependent functions and to the media-independent functions used by customers to administer their networks, including the support of C-VLANs.

h) Allocates the reserved multicast addresses to media access method-dependent, provider network, and customer network functions, specifying the filtering to be applied in each type of VLAN Bridge component.

i) Defines the principles of network operation in terms of the support and preservation of the MAC Service, and the maintenance of QoS for each service instance, including the segregation of data belonging to different organizations.

j) Specifies customer interfaces to a Provider Bridged Network (PBN) in terms of the operation and configuration of the VLAN Bridge components of Provider Bridges, including interfaces that
   1) Provide access to a single service instance through a Bridge Port.
   2) Allow a customer to select among and identify service instances by Customer VLAN Identifier (C-VID).
   3) Allow a customer to select among and identify service instances by Service VLAN Identifier (S-VID).
   4) Support customer signaling of priority information on a frame by frame basis.
   5) Multiplex service instances over LANs that provide access to a provider network.
   6) Support fault tolerance through redundant provision of access LANs and equipment.

k) ~~Describes the functions to be performed within the PBN in order to support and maintain the connectivity provided to customer service instances.~~

l) ~~Establishes the requirements for Bridge Management in the PBN, identifying the managed objects and defining the management operations.~~

There is no management specification that addresses PBNs specifically, beyond management of the aspects of individual Bridges. The deleted list items above were probably 'aspirational'.

[15] Clauses 15 and 16. PBN capabilities were added by IEEE Std 802.1ad-2005.

m) Specifies performance requirements, and recommends default values and applicable ranges for the operational parameters of a Provider Bridge.

The following paragraph and list item were previously (i..e., in 802.1Q-2022) positioned after bf), between a description of MIRP for PBBNs and the description of PFC. I have not shown the move as a change remove and change add, but simply moved the text.

This standard allows an S-tagged service interface connecting two independently administered PBNs to be used to handle traffic (identified by a single S-VID) for a given customer attached to one PBN as if the customer were directly attached to the other PBN using a Port-based or C-tagged service interface. To this end, it:

n) Specifies the use of a Port-mapping S-VLAN component to associate selected S-VIDs registered on an external port with distinct internal ports, each of which supports a separate service interface.[16]

1.3.4 Connectivity Fault Management[17]

This standard specifies protocols, procedures, and managed objects to support Connectivity Fault Management (CFM). These allow discovery and verification of the path, through Bridges and LANs, taken for frames addressed to and from specified network users, and support detection and isolation of a connectivity fault to a specific Bridge or LAN. To this end, it:

a) Defines Maintenance Domains, Maintenance Associations (MAs), their constituent Maintenance Points (MPs), and the managed objects required to create and administer them.

b) Describes the protocols and procedures used by MPs to detect and diagnose connectivity faults within a Maintenance Domain.

~~This standard specifies CFM protocols, procedures, and managed objects that provide confirmation of successful transmission of frames conveying specified data. This capability supports diagnosis of faults sensitive to, or caused by, particular data patterns, and their isolation to part of the transmission path. Connectivity verification can be carried out from any single point with bridged connectivity to MPs on the path, can isolate failures to communicate in a specific direction, and can be carried out while service is being provided to other users of the data path. To this end, it:~~

~~c) Defines the extensions to CFM capabilities defined by Clause 18 through Clause 22 to facilitate diagnosis and isolation of faults sensitive to, or caused by, particular data patterns in frames transmitted by a service user.~~

~~d) Describes the protocols and procedures for data-driven and data-dependent connectivity fault management (DDCFM).~~

1.3.5 Provider Backbone Bridges and Provider Backbone Bridged Networks[18]

To allow scaling of Provider Networks to at least 2^{24} S-VLANs, this standard further specifies the operation of Provider Backbone Bridges (PBBs) by means of an architecture and Bridge protocols compatible and interoperable with PBN protocols and equipment, allowing interconnection of multiple PBNs. To this end, it:

a) Introduces BEBs that, by exchanging backbone frames that encapsulate the addresses, VLAN tags, and data of customer frames, support the virtual, media-independent equivalent of a number of independent instances of the service provided by media-dependent frame transmission procedures.

b) Extends the parameters of the ISS and EISS to include a connection identifier, capable of referencing the backbone addresses and other parameters, used to convey customer frames from one BEB to all, or one of, the other BEBs supporting a particular backbone service instance.

[16] Remote Customer Service Interface (RCSI) capability added to 16.2 by IEEE Std 802.1Qbc-2011.
[17] Clauses 18, 19, 20, 21, 22. CFM capabilities were added by IEEE Std 802.1ag-2007.
[18] Clauses 25and 26. Provider Backbone Bridge capabilities were added by IEEE Std 802.1ah-2008.

c) Specifies the format of the Backbone Service Instance tag (I-TAG) that encapsulates the customer addresses, and introduces a Backbone Service Instance Identifier (I-SID) that allows each BEB to support a number of backbone service instances and permits the unambiguous identification of up to 2^{24} backbone service instances within a single Provider Backbone Bridged Network (PBBN).

d) Provides a model of BEB operation in terms of VLAN Bridge components that allows the use of Provider Bridges as Backbone Core Bridges (BCBs), with PBBN traffic carried as frames containing I-TAGs on particular Backbone VLANs (B-VLANs) potentially coexisting with PBN traffic carried as frames without I-TAGs on other B-VLANs.

e) Specifies the interfaces that a PBBN can provide to transport service frames. These comprise a Port-based service interface that assigns all received untagged and priority-tagged frames to a single S-VLAN transported over a single backbone service instance, an S-tagged service interface capable of mapping individual S-VLANs to different backbone service instances, and an I-tagged service interface capable of mapping frames from one set of backbone service instances to another.

f) Describes the use of redundant Bridges and access LANs to protect backbone service access against failure of any of those systems or components.

g) Specifies the management of BEBs in terms of the model of operation [item c) above], making use of defined management objects for the individual VLAN Bridge components, and adding managed objects to facilitate service creation.

h) Describes the use of CFM to detect and isolate faults in the connectivity provided to individual S-VLANs across the PBBN, in the connectivity provided to the group of S-VLANs supported by a single backbone service instance (identified by an I-SID), and in the connectivity provided to individual B-VLANs within the backbone itself.

i) Specifies extensions to MSTP to allow network administrators to protect against loops through peered PBBNs without requiring coupling of spanning trees that operate independently for each PBBN.

1.3.6 Provider Backbone Bridged Network traffic engineered paths[19]

This standard specifies protocols, procedures, and managed objects to allow support of provisioning systems that explicitly select traffic engineered paths within Provider Backbone Bridged Networks (PBBNs) by allowing a network operator to disable unknown destination address forwarding, source address learning and spanning tree protocols for administratively selected VIDs, while allowing other network control protocols to dynamically determine active topologies for other services. These interoperable capabilities are supported by management of individual Bridges by Simple Network Management Protocol (SNMP) using an SMIv2 MIB, by extensions to the other control protocols specified in this standard, by the use of CFM with the addresses and VIDs that specify traffic engineered connections, and by 1:1 path protection switching capable of load sharing. To this end, it:

a) Enables construction of active topologies by an external agent that is responsible for setting up Ethernet Switched Paths (ESPs) by splitting the B-VID space between distributed spanning tree protocols and provisioned control.

b) Supports discard of frames with unknown destination addresses for B-VIDs under provisioned control.

c) Supports the operation of Continuity Check, Loopback, and Linktrace protocols on provisioned traffic engineered paths.

d) Supports 1:1 protection switching capable of load sharing for Traffic Engineering service instances (TESIs).

e) Supports protection of a group of TESIs that traverses a sequence of LANs and intervening Bridges using a method that does not require the modification of data or control frames.

[19] Provider Backbone Traffic Engineering (PBB-TE) capabilities were addressed by IEEE Std 802.1Qay-2009, with substantive additions to the PBBN and CFM Clauses (25, 26, and 18, 19, 20, 21, 22). IEEE Std 802.1Qbf-2011 added infrastructure segment protection.

f) ~~This standard specifies protocols, procedures, and managed objects to support~~ Supports topology change signaling to alter the binding (held in an I-Component) of Customer addresses to backbone addresses on a per-I-SID basis~~. This is accomplished by extending the use of MRP. To this end, it specifies the~~, specifying the MIRP application of MRP~~ and the frame formats that it uses~~.[20]

NOTE—MIRP can only trigger the flushing of learned MAC address information; it does not propagate the registration of I-SIDs.~~ The name Multiple I-SID Registration Protocol is chosen because MIRP is a Multiple Registration Protocol (MRP) application and can be extended to perform I-SID registrations.~~

g) ~~Provides required extension to SNMP management by SMIv2 MIB modules.~~

This standard does not specify operation of ESPs through multiple Provider Backbone Bridge Traffic Engineering (PBB-TE) Regions. All the Backbone Edge Bridges (BEBs) specified for use in a PBB-TE Region are combined I type and B type Backbone Edge Bridges (IB-BEBs).

1.3.7 Auto Attach for PBBNs[21]

This standard specifies the protocols, procedures, and management objects for auto attachment of network devices to backbone service instances (BSIs) by adding an Auto Attach Protocol (AAP) over the Link Layer Discovery Protocol (LLDP). The AAP simplifies the deployment and administration of PBBNs ~~networks~~ by automatically coupling end stations or Bridges to BSIs.

To this end, it:

a) Introduces the concepts of Auto Attach Protocol (AAP), Auto Attach System (AAS) on devices supporting Auto Attach[22]; non-PBB-based Auto Attach Devices (AADs) and PBB-based Auto Attach BEBs (AABs).

b) Describes the process of discovery and advertising of capabilities for AADs and AABs supporting the AAP acting as an AAS.

c) Describes the bindings between Backbone Service Instance Identifiers (I-SIDs) and VLAN Identifiers.

d) Specifies IEEE 802.1 Organizationally Specific LLDP TLVs for discovery and capabilities advertising, and attachment of C-VLANs to Backbone Service Instances.

e) ~~Specifies SNMP MIB support for PBBN Auto Attach.~~

f) ~~Specifies YANG extensions to the LLDP YANG modules for the PBBN Auto Attach LLDP TLVs.~~

g) ~~Specifies SNMP MIB extensions to the LLDP EVB extension MIB module for the PBBN Auto Attach LLDP TLVs.~~

Removed as per comment near the beginning of this clause (1.3).

[20] Clause 39, added by IEEE Std 802.1Qbe-2011.

[21] Clause 50 and supporting management and LLDP capabilities were added by IEEE Std 802.1Qcj-2023.

[22] The term “Auto Attach” is used interchangeably with “PBBN Auto Attach” where it is clear from the context that “PBBN Auto Attach” is discussed.

1.3.8 Shortest Path Bridging[23]

This standard specifies Shortest Path Bridging (SPB) of unicast and multicast frames, specifying protocols to calculate multiple active topologies that can share learned station information, and support of a VLAN by multiple, per-topology, Shortest Path VLAN Identifiers (SPVIDs). To this end, it:

a) Describes the use of shortest paths to increase throughput and minimize transit delay, while introducing a negligible rate of frame misordering.

b) Requires that active topologies calculated by spanning tree protocols and Shortest Path Tree (SPT) protocols be stable, predictable, and reproducible to maintain the characteristics of the MAC Service provided.

c) Requires, except in the case of SPB using Equal Cost Multiple Paths (ECMP), active topologies that are reverse path congruent and unicast-multicast congruent to permit learning of station location from the source addresses of all frames and simplify the detection and management of faults.

~~NOTE 4—ECMP operation does not provide (nor does this standard attempt to define for ECMP VLANs) reverse path congruence and unicast-multicast congruence as these concepts cease to have utility in an ECMP context.~~

NOTE above provides a snippet that is far too detailed for this Introduction. Clause 44 (ECMP) provides more easily understood detail.

d) Specifies the calculation of symmetric sets of SPTs, each rooted at a Bridge within an SPT Region comprising Bridges operating compatible protocols and configurations.

e) Specifies the use of Bridge Protocol Data Units (BPDUs) to identify and bound SPT Regions and to ensure loop-free interoperability with regions using the Rapid Spanning Tree Algorithm and Protocol (RSTP) and MSTP.

f) Specifies both Shortest Path Bridging VID (SPBV) and Shortest Path Bridging MAC (SPBM) modes:

   1) for SPBV, identifying each SPT by SPVID and locating end stations by source MAC address learning.
   2) for SPBM, identifying each SPT by VID and source MAC address and distributing end station location information explicitly.

g) Supports management selection of the Common Spanning Tree (CST), a Multiple Spanning Tree Instance (MSTI), or SPB for support of any given VLAN within an SPT Region.

h) Specifies a protocol that automatically assigns SPVIDs for each VLAN supported by SPBV.

i) Supports load sharing by Equal Cost Trees (ECTs) through the calculation of multiple SPT Sets, with each shortest path VLAN being assigned to one SPT Set.

j) Specifies Intermediate System to Intermediate System Protocol for Shortest Path Bridging (ISIS-SPB): the use of and extensions to the Intermediate System to Intermediate System (IS-IS) Protocol to calculate SPTs for both SPBV and SPBM.

k) Describes the addressing of ISIS-SPB entities and specifies the group MAC addresses they use.

l) Specifies the use of loop prevention (for SPBV and for multicast frames for SPBM) and loop mitigation (for unicast frames for SPBM).

m) Specifies an Agreement Protocol that prevents loops, specifying the necessary state information and computation as part of ISIS-SPB and communicating agreement information for the CIST and (as a compact Digest) for SPTs in each BPDU.

n) Specifies load spreading by distributing unicast traffic over the set of available equal cost paths and assigning multicast traffic flows to a variety of trees.

o) Specifies a flow filtering tag (F-TAG) containing a flow hash used in unicast ECMP traffic distribution and a TTL (time-to-live) field used to mitigate the effects of traffic loops resulting from transient conditions or control software errors or faults.

[23] Clauses 27 and 28, added by IEEE_Std_802.1aq-2012. Clause 44, ECMP support added by IEEE Std 802.1Qbp-2014.

1.3.9 IS-IS based explicit path control, bandwidth reservation, and redundancy[24]

This standard ~~also~~ specifies ~~further~~ protocol extensions, procedures, and managed objects to IS-IS for providing capabilities beyond Shortest Path Bridging (SPB) for Bridged Networks. These extensions involve explicit path control, bandwidth reservation, and redundancy (protection, restoration) for data flows. Thus, this standard specifies bridging on explicit paths for unicast and multicast frames, specifying protocols to determine multiple active topologies. To this end, it:

a) Describes the use of explicit trees, e.g., to improve resiliency and decrease the probability of congestion.
b) Requires that active topologies calculated by one or multiple entities external to the routing protocol are such that the characteristics of the MAC Service are provided.
c) Supports management selection of explicit trees for support of any given VLAN within an SPT Region.
d) Specifies Intermediate System to Intermediate System Path Control and Reservation (ISIS-PCR): the use of and extensions to the Intermediate System to Intermediate System (IS-IS) protocol to establish explicit trees.
e) Specifies the use of ISIS-PCR for recording bandwidth assignments.
f) Specifies redundancy for ISIS-SPB and ISIS-PCR.

1.3.10 Two-Port MAC Relays (TPMRs)[25]

This standard specifies the function of a Two-Port MAC Relay (TPMR), along with protocols and procedures that support its operation. A TPMR is a type of Bridge that has only two externally accessible Bridge Ports, and supports a subset of the functionality of a MAC Bridge. A TPMR is transparent to all frame-based media-independent protocols, except those explicitly addressed to it and those that are destined for reserved MAC addresses that the relay function of the TPMR is defined not to forward. It is remotely manageable through at least one of its external MACs, and signals a failure of either MAC's LAN through the other MAC. A TPMR should only be attached to point-to-point LANs. ~~The conformance requirements for a TPMR are stated in 5.13 and 5.15.~~

1.3.11 Time and loss-sensitive stream transmission[26]

To enable the end-to-end management of resource reservation for QoS guaranteed streams, this standard further specifies protocols, procedures, and managed objects, usable by existing higher layer mechanisms, that allow network resources to be reserved for specific traffic streams traversing a Bridged Network. To this end, it:

a) Specifies the use of Dynamic Reservation Entries (8.8.7) in the FDB to control the forwarding of frames associated with a particular Stream.
b) Specifies a Stream Reservation Protocol (SRP). SRP facilitates the registration, deregistration, and maintenance of stream reservation information in relevant Bridges to establish end-to-end stream paths.

This standard allows Bridges to provide performance guarantees for time-sensitive (i.e., bounded latency and latency variation) loss-sensitive real-time audio/video (AV) data stream transmission (AV traffic). It specifies priority regeneration and controlled bandwidth queue draining algorithms. VLAN tag encoded priority values are allocated, in aggregate, to segregate frames among queues that support AV traffic and queues that support non-AV traffic, allowing simultaneous support of both AV traffic and other bridged traffic over and between wired and wireless Local Area Networks (LANs). To this end, it:

[24] Clause 45, added by IEEE Std 802.1Qca-2015.
[25] Capabilities added by IEEE Std_802.1aj-2009, including Clause 23, MAC Status Propagation Protocol (MSPP).
[26] Clause 34, added by IEEE Std 802.1Qav-2009, and Clause 35, added by IEEE Std 802.1Qat-2010.

c) Defines status parameters that allow the boundaries of an ~~Stream Reservation Protocol (~~SRP~~— see Clause 35~~) domain~~ (35.1.4)~~ to be identified and maintained.

d) Specifies how the priority information in frames received at SRP domain boundary ports is regenerated.

~~NOTE 1—The priorities in frames transmitted from outside an SRP domain to a Bridge inside an SRP domain are remapped in order to ensure that traffic that is not associated with a reservation does not disrupt traffic that is associated with a reservation. Hence, traffic entering an SRP domain that uses Priority Code Point values associated with reserved traffic classes will be remapped to Priority Code Point values that are not associated with reserved traffic classes.~~[27]

NOTE above, and below, far too much detail for this Introduction.

e) Specifies how priority information is used to determine the traffic classes to be used for time-sensitive streams.

f) Defines a credit-based shaper algorithm to shape traffic in accordance with stream reservations.

~~NOTE 2—The credit-based shaper algorithm operates on the outbound queues; the mechanisms specified for the support of time-sensitive AV traffic do not involve any form of ingress metering or policing.~~

This standard also specifies:

g) ~~Defines e~~Enhancements for scheduled traffic to allow transmissions scheduled relative to a known timescale.

h) ~~Defines f~~Frame preemption to interrupt transmission of preemptable frames by express frames.

~~This standard specifies protocols, procedures, and managed objects that:~~

i) ~~Allow for the~~ Protocols, procedures, and managed objects for filtering and policing of individual traffic streams.[28]

j) ~~Allow for~~ Asynchronous Traffic Shaping (ATS) over full-duplex links with constant bit data rates.[29]

1.3.12 Network configuration for time-sensitive applications[30]

This standard specifies ~~enhancements to~~ protocols, procedures, and managed objects for the configuration of network resources for time-sensitive applications that require timely, high probability, delivery of frames without end station retransmission. To this end, it:

a) Describes three approaches to network configuration: fully distributed, centralized network/distributed user, and fully centralized.

b) ~~Specifies managed objects for forwarding and queuing enhancements for time-sensitive streams (FQTSS).~~

c) Describes Centralized User Configuration (CUC) and Centralized Network Configuration (CNC) entities.

d) ~~Specifies managed objects for configuration of Bridges by a CNC.~~

e) ~~Specifies YANG configuration and operational state models (Clause 48) in support of Scheduled Traffic, Frame Preemption, Per-Stream Filtering and Policing, and CUC configuration.~~

Deleted list items are all covered, for the purposes of this Introduction, by the first sentence under this heading "managed objects ..configuration of network resources ... time-sensitive .."and by the text in the paragraph following the 1.3 heading. In general there is obviously a tension between the desire of each amendment project to describe in detail what that project has set out to do (and in most cases has done), and the fact that the purpose of an amendment project is to modify the text of the complete standard, within which the subject of the amendment may have already been addressed (less technical refinement). We might make more use of the 'Introduction' that precedes the 'Contents' in amendments for effectively recording the Scope of the Amendment. Temporary text of that nature has also been preserved for amendment roi

[27] Notes in text, tables, and figures, are given for information only and do not contain requirements needed to implement the standard.

[28] Additions to 8.6 by IEEE Std 802.1Qci-2017, with subsequent amendment.

[29] Added by IEEE Std 802.1Qcr-2020, which included revisions to flow classification and metering(8.6.5), transmission selection (8.6.5), and ATS support in end stations (Clause 47).

[30] Clause 46, added by IEEE Std 802.1Qcc-2018, revised by IEEE Std 802.1Qdj-2024.

lled up into the current 802.1Q Revision in the ‘Introduction’ for this Revision.

1.3.13 Congestion notification[31]

This standard specifies protocols, procedures, and managed objects to support congestion notification. These allow a Virtual Bridged Network or a portion thereof, with a limited bandwidth-delay product, to transfer long-lived data flows with a significantly reduced chance of frame loss compared to a network without congestion notification. To this end, it:

a) Defines a means for VLAN Bridges that support congestion notification to form Congestion Managed Domains within a Virtual Bridged Network.
b) Defines a means for detecting congested queues in end stations and VLAN Bridges, for signaling such congestion to the end stations sourcing the frames causing the congestion, and for those end stations to control the rate of transmission of those frames.

1.3.14 Congestion Isolation[32]

This standard specifies protocols, procedures, and managed objects that support the isolation of congesting data flows within data center environments. This is achieved by enabling systems to ~~individually~~ identify individual flows creating congestion, isolate those flows to the congesting queue, and signal to neighbors. This mechanism reduces head-of-line blocking for non-congesting flows sharing the same traffic class. Congestion Isolation is used with higher layer protocols that utilize end-to-end congestion control in order to reduce packet loss and latency. To this end, it:

a) Defines a means for VLAN-aware Bridges that support congestion isolation to identify flows that are creating congestion.
b) Defines a means for adjusting transmission selection for frames of congesting flows.
c) Provides a means for discovering peer VLAN-aware Bridges and stations that support congestion isolation.
d) Defines a means for signaling congestion isolation to supporting peer Bridges and stations.
e) Defines a means for recognizing a system’s level and port orientation within the topology relative to the edge.

1.3.15 Priority-based Flow Control[33]

This standard specifies protocols, procedures, and managed objects to support Priority-based Flow Control (PFC). These allow a Virtual Bridged Network, or a portion thereof, to enable flow control per traffic class on IEEE 802 point-to-point full-duplex links. To this end, it:

a) Defines a means for a system to inhibit transmission of data frames on certain priorities from the remote system on the link.

1.3.16 Enhanced Transmission Selection[34]

This standard specifies protocols, procedures, and managed objects for enhancement of transmission selection to support allocation of bandwidth among traffic classes. When the offered load in a traffic class does not use its allocated bandwidth, Enhanced Transmission Selection (ETS) will allow other traffic classes

[31] Clauses 30, 31, 32, and 33, added by IEEE Std_802.1Qau-2010.
[32] Clause 49, added by IEEE Std 802.1Qcz-2023.
[33] Clause 36, added by IEEE Std 802.1Qbb-2011.
[34] Clause 37 (ETS) and Clause 38 (DCBX) added by IEEE Std 802.1Qaz-2011.

to use the available bandwidth. Bandwidth is used by traffic classes subject to ETS when there are no frames to be transmitted for traffic classes subject to strict priority or credit-based shaper algorithms. It defines the Data Center Bridging eXchange protocol (DCBX), which controls the application of ETS and PFC.

1.3.17 Edge Virtual Bridging (EVB)[35]

This standard specifies protocols, procedures, and managed objects that:

a) Provide for the discovery, configuration, and control of a pair of direct-attached Port-mapping S-VLAN components to extend the operation of a Customer Bridge to remote ports and enable coexistence of multiple services on station-resident ports (e.g., embedded bridging).

b) Provide for discovery, configuration, and operation of reflective relay (8.6.1) for a Bridge Port.

c) Provide for discovery of, and coordinated configuration of, edge relays (ERs) and other devices that utilize the reflective relay service.

d) ~~Provide for dynamic profile-driven port configuration.~~

e) Provide for Network Virtualization Overlays over Layer 3 (NVO3)-related port configuration.[36]

By all means be ‘dynamic’, and ‘profile-driven’ sounds good as well, but the word ‘dynamic’ does not appear anywhere in the EVB and related Clauses (40-43), and the above is the only instance of ‘profile-driven’ in the whole of the standard. So as a reader I am at a loss to understand for what I should be thankful.

Miscellaneous, without a home—for removal if none can be found

a) ~~Provides for the use of IEEE 802.11™ media as links internal to, as well as links providing access to, a Bridged Network or Virtual Bridged Network.~~

I can find no text anywhere in 802.1Q which addresses the above, previously item ch) in 802.1Q-2022 specifically. It was introduced by IEEE Std 802.1Qbz-2016 ‘Enhancements to Bridging of IEEE 802.11 Media’ which made useful clarifications in the general area of protocol identification and stacked tag handling, which are applicable to 802.11 and other media with a predominantly LLC encoding history, but did not differentiate as to the use of types of 802.11 link.

[35] Clauses 40, 41, 42, and 43, added by IEEE Std 802.1Qbg-2012.

[36] Clauses 40 and 41 revised by IEEE Std 802.1Qcy-2019.