IEEE Std 802.1v-2001
(Amendment to IEEE Std 802.1Q, 1998 Edition)

IEEE Standard for
Local and metropolitan area networks—

Virtual Bridged Local Area Networks— Amendment 2: VLAN Classification by Protocol and Port

Sponsor

LAN/MAN Standards Committee
of the
IEEE Computer Society

Approved 17 March 2001

IEEE-SA Standards Board

Abstract: This amendment to IEEE Std 802.1Q, 1998 Edition describes enhancements to allow for classification of incoming packets by methods other than source port. Specifically, it defines rules for classification based on data-link layer protocol identification. The document identifies proposed changes to the text of IEEE Std 802.1Q, 1998 Edition that have arisen as a consequence of this activity. These are documented in the usual form for amendments to IEEE 802® standards; i.e., as an explicit set of editing instructions that, if correctly applied to the text of IEEE Std 802.1Q, 1998 Edition, will create a corrected document.
Keywords: local area networks, MAC Bridge management, media access control bridges, virtual LANs

The Institute of Electrical and Electronics Engineers, Inc.
3 Park Avenue, New York, NY 10016-5997, USA

Print: ISBN 0-7381-2818-X SH94917
PDF: ISBN 0-7381-2819-8 SS94917

IEEE Standards documents are developed within the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE-SA) Standards Board. The IEEE develops its standards through a consensus development process, approved by the American National Standards Institute, which brings together volunteers representing varied viewpoints and interests to achieve the final product. Volunteers are not necessarily members of the Institute and serve without compensation. While the IEEE administers the process and establishes rules to promote fairness in the consensus development process, the IEEE does not independently evaluate, test, or verify the accuracy of any of the information contained in its standards.

Use of an IEEE Standard is wholly voluntary. The IEEE disclaims liability for any personal injury, property or other damage, of any nature whatsoever, whether special, indirect, consequential, or compensatory, directly or indirectly resulting from the publication, use of, or reliance upon this, or any other IEEE Standard document.

The IEEE does not warrant or represent the accuracy or content of the material contained herein, and expressly disclaims any express or implied warranty, including any implied warranty of merchantability or fitness for a specific purpose, or that the use of the material contained herein is free from patent infringement. IEEE Standards documents are supplied "**AS IS**."

The existence of an IEEE Standard does not imply that there are no other ways to produce, test, measure, purchase, market, or provide other goods and services related to the scope of the IEEE Standard. Furthermore, the viewpoint expressed at the time a standard is approved and issued is subject to change brought about through developments in the state of the art and comments received from users of the standard. Every IEEE Standard is subjected to review at least every five years for revision or reaffirmation. When a document is more than five years old and has not been reaffirmed, it is reasonable to conclude that its contents, although still of some value, do not wholly reflect the present state of the art. Users are cautioned to check to determine that they have the latest edition of any IEEE Standard.

In publishing and making this document available, the IEEE is not suggesting or rendering professional or other services for, or on behalf of, any person or entity. Nor is the IEEE undertaking to perform any duty owed by any other person or entity to another. Any person utilizing this, and any other IEEE Standards document, should rely upon the advice of a competent professional in determining the exercise of reasonable care in any given circumstances.

Interpretations: Occasionally questions may arise regarding the meaning of portions of standards as they relate to specific applications. When the need for interpretations is brought to the attention of IEEE, the Institute will initiate action to prepare appropriate responses. Since IEEE Standards represent a consensus of concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason, IEEE and the members of its societies and Standards Coordinating Committees are not able to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration.

Comments for revision of IEEE Standards are welcome from any interested party, regardless of membership affiliation with IEEE. Suggestions for changes in documents should be in the form of a proposed change of text, together with appropriate supporting comments. Comments on standards and requests for interpretations should be addressed to:

Secretary, IEEE-SA Standards Board
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
USA

Note: Attention is called to the possibility that implementation of this standard may require use of subject matter covered by patent rights. By publication of this standard, no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE shall not be responsible for identifying patents for which a license may be required by an IEEE standard or for conducting inquiries into the legal validity or scope of those patents that are brought to its attention.

IEEE is the sole entity that may authorize the use of certification marks, trademarks, or other designations to indicate compliance with the materials set forth herein.

Introduction

[This introduction is not part of IEEE Std 802.1v-2001 (Amendment to IEEE Std 802.1Q, 1998 Edition), IEEE Standard for Local and metropolitan area networks—Virtual Bridged Local Area Networks—Amendment 2: VLAN Classification by Protocol and Port.]

The MAC Bridge standardization activities that resulted in the development of IEEE Std 802.1D, 1993 Edition and IEEE Std 802.1D, 1998 Edition introduced the concept of Filtering Services for both Unicast and Group-Addressed traffic in Bridged LANs, and mechanisms whereby filtering information in such LANs may be acquired and held in a Filtering Database. They also defined the use of multiple classes of expedited traffic, based on user priority, to support the transmission of time-critical information in a LAN environment.

Standardization activities also resulted in the development of IEEE Std 802.1Q, 1998 Edition, Virtual Bridged Local Area Networks, which allows for the implementation of Virtual Bridged LANs. The following capabilities are described in that document:

a) Virtual LAN Services in Bridged LANs.

b) The operation of the Forwarding Process that is required in order to support Virtual Bridged LANs.

c) The structure of the Filtering Database that is required in order to support Virtual Bridged LANs.

d) The nature of the protocols and procedures that are required in order to provide Virtual LAN services, including the definition of the frame formats used to represent VLAN identification information, and the procedures used in order to insert and remove VLAN identifiers and the headers in which they are carried.

e) The ability to support end to end signalling of user priority information regardless of the intrinsic ability of the underlying MAC protocols to signal user priority information.

f) The GARP VLAN Registration Protocol, GVRP, that allows distribution and registration of VLAN membership information. The protocol described makes use of the GARP protocol defined in IEEE Std 802.1D, 1998 Edition.

g) The management services and operations that are required in order to configure and administer Virtual Bridged LANs.

h) The classification rules to be used at the ingress to a bridge in order to classify incoming frames as belonging to a particular Virtual LAN. In this revision of the standard, these rules are limited to one of the following:

    1) Use of a pre-existing VLAN identifier in the frame

    2) Use of the incoming bridge port's PVID

This amendment describes the following additional capabilities:

i) A third type of rule for ingress classification:

    1) Use of a frame's data-link layer protocol identification to determine a VLAN to which the frame belongs.

This standard is part of a family of standards for local and metropolitan area networks. The relationship between the standard and other members of the family is shown below. (The numbers in the figure refer to IEEE standard numbers.)

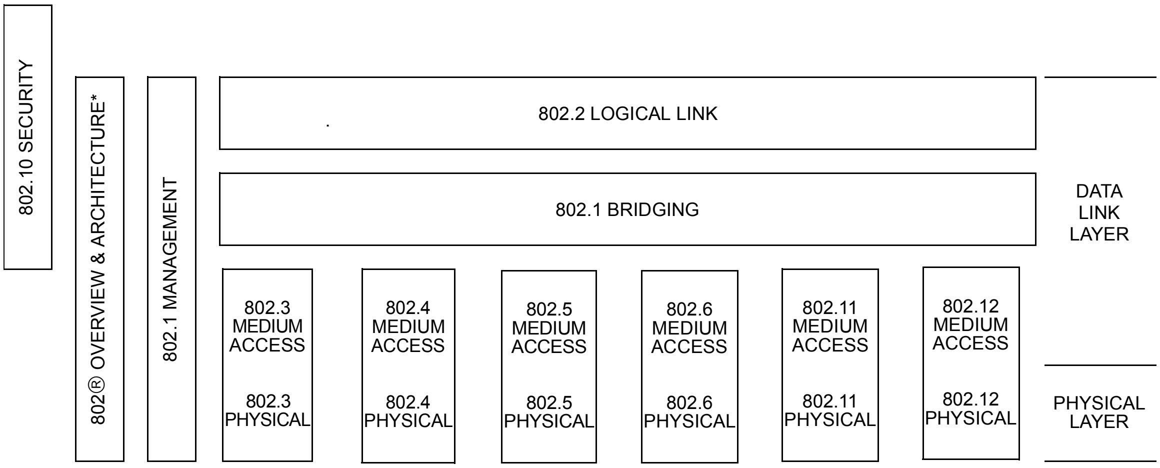

* Formerly IEEE Std 802.1A.

This family of standards deals with the Physical and Data Link Layers as defined by the International Organization for Standardization (ISO) Open Systems Interconnection Basic Reference Model (ISO/IEC 7498-1:1994). The access standards define several types of medium access technologies and associated physical media, each appropriate for particular applications or system objectives. Other types are under investigation.

The standards defining the technologies noted above are as follows:

- IEEE Std 802[1]: *Overview and Architecture*. This standard provides an overview to the family of IEEE 802 Standards. This document forms part of the IEEE 802.1 scope of work.

- ANSI/IEEE Std 802.1B and 802.1K [ISO/IEC 15802-2]: *LAN/MAN Management*. Defines an Open Systems Interconnection (OSI) management-compatible architecture, and services and protocol elements for use in a LAN/MAN environment for performing remote management.

- ANSI/IEEE Std 802.1D *Media Access Control* (*MAC) Bridges*. Specifies an architecture and protocol for the [ISO/IEC 15802-3]: interconnection of IEEE 802 LANs below the MAC service boundary.

- ANSI/IEEE Std 802.1E [ISO/IEC 15802-4]: *System Load Protocol*. Specifies a set of services and protocol for those aspects of management concerned with the loading of systems on IEEE 802 LANs.

[1] The 802 Architecture and Overview Specification, originally known as IEEE Std 802.1A, has been renumbered as IEEE Std 802. This has been done to accommodate recognition of the base standard in a family of standards. References to IEEE Std 802.1A should be considered as references to IEEE Std 802.

- ANSI/IEEE Std 802.1F *Common Definitions and Procedures for IEEE 802 Management Information*.

- ANSI/IEEE Std 802.1G [ISO/IEC 15802-5]: *Remote Media Access Control (MAC) Bridging*. Specifies extensions for the interconnection, using non-LAN systems communication technologies, of geographically separated IEEE 802 LANs below the level of the logical link control protocol.

- ANSI/IEEE Std 802.1H [ISO/IEC TR 11802-5] *Recommended Practice for Media Access Control (MAC) Bridging of Ethernet V2.0 in IEEE 802 Local Area Networks*.

- ANSI/IEEE Std 802.1Q *Virtual Bridged Local Area Networks*. Defines an architecture for Virtual Bridged LANs, the services provided in Virtual Bridged LANs, and the protocols and algorithms involved in the provision of those services.

- ANSI/IEEE Std 802.2 [ISO/IEC 8802-2]: *Logical Link Control*.

- ANSI/IEEE Std 802.3 [ISO/IEC 8802-3]: *CSMA/CD Access Method and Physical Layer Specifications*.

- ANSI/IEEE Std 802.4 [ISO/IEC 8802-4]: *Token Bus Access Method and Physical Layer Specifications*.

- ANSI/IEEE Std 802.5 [ISO/IEC 8802-5]: *Token Ring Access Method and Physical Layer Specifications*.

- ANSI/IEEE Std 802.6 [ISO/IEC 8802-6]: *Distributed Queue Dual Bus Access Method and Physical Layer Specifications*.

- ANSI/IEEE Std 802.10: *Interoperable LAN/MAN Security.* Currently approved: Secure Data Exchange (SDE).

- ANSI/IEEE Std 802.11: [ISO/IEC 8802-11] *Wireless LAN Medium Access Control (MAC) Sublayer and Physical Layer Specifications.*

- ANSI/IEEE Std 802.12: [ISO/IEC 8802-12] *Demand Priority Access Method, Physical Layer and Repeater Specification*.

- IEEE Std 802.15: *Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Wireless Personal Area Networks.*

- IEEE Std 802.16: *Standard Air Interface for Fixed Broadband Wireless Access Systems*.

- IEEE Std 802.17: *Resilient Packet Ring Access Method and Physical Layer Specifications.*

In addition to the family of standards, the following is a recommended practice for a common physical layer technology:

- IEEE Std 802.7: *IEEE Recommended Practice for Broadband Local Area Networks*.

The reader of this standard is urged to become familiar with the complete family of standards.

Conformance test methodology

An additional standards series, identified by the number IEEE 1802, has been established to identify the conformance test methodology documents for the IEEE 802 family of standards. Thus the conformance test documents for IEEE 802.3 are numbered IEEE 1802.3, the conformance test documents for IEEE 802.5 will be 1802.5, and so on. Similarly, ISO will use ISO/IEC 18802 to number conformance test standards for ISO/IEC 8802 standards.

Participants

When the IEEE 802.1 Working Group approved this standard, it had the following membership:

Tony Jeffree, *Chair and Editor*
Neil Jarvis, *Vice Chair*
Mick Seaman, *Chair, Interworking Task Group*
David Delaney, *Editor*
Andrew Smith, *Editor*

Floyd Backes
Les Bell
Alan Chambers
Marc Cochran
Paul Congdon
Hesham El Bakoury
Norman W. Finn
Sharam Hakimi
Bob Hott

Toyayuki Kato
Hal Keen
Daniel Kelley
Keith Klamm
Bill Lidinsky
Yaron Nachman
Satoshi Obara
Luc Pariseau
Anil Rijsinghani

John J. Roese
Ted Schroeder
Benjamin Schultz
Rosemary V. Slager
Andrew Smith
Michel Soerensen
Robin Tasker
Manoj Wadekar
Robert Williams

The following members of the balloting committee voted on this standard:

Terry S. Arnold
James T. Carlo
Keith Chow
Robert S. Crowder
Guru Dutt Dhingra
Thomas J. Dineen
Christos Douligeris
Sourav K. Dutta
Philip H. Enslow
Changxin Fan
John W. Fendrich
Michael A. Fischer
Richard A. Froke
Robert J. Gagliano
Gautam Garai
Alireza Ghazizahedi
Tim Godfrey
Robert M. Grow
Chris G. Guy
Simon Harrison

Osamu Ishida
Raj Jain
Kamran Jamal
Neil A. Jarvis
Anthony A. Jeffree
Jack R. Johnson
Stuart J. Kerry
Daniel R. Krent
Stephen Barton Kruger
Joseph Kubler
David J. Law
William Lidinsky
Randolph S. Little
Ronald Mahany
Peter Martini
Bennett Meyer
David S. Millman
James F. Mollenauer
John E. Montague

Robert Mortonson
Robert O’Hara
Satoshi Obara
Roger Pandanda
Vikram Punj
Gary S. Robinson
Edouard Y. Rocher
James W. Romlein
Floyd E. Ross
Jaideep Roy
Rich Seifert
Leo Sintonen
Joseph S. Skorupa
Fred J. Strauss
Jonathan R. Thatcher
Mark-Rene Uchida
Scott A. Valcourt
John Viaplana
Paul A. Willis
Oren Yuen

Contents

Recognized as an
American National Standard (ANSI)

IEEE Std 802.1v-2001
(Amendment to IEEE Std 802.1Q, 1998 Edition)

IEEE Standard for Local and metropolitan area networks—

Virtual Bridged Local Area Networks— Amendment 2: VLAN Classification by Protocol and Port

Sponsor

LAN/MAN Standards Committee
of the
IEEE Computer Society

Approved 17 March 2001

IEEE-SA Standards Board

Abstract: This amendment to IEEE Std 802.1Q, 1998 Edition describes enhancements to allow for classification of incoming packets by methods other than source port. Specifically, it defines rules for classification based on data-link layer protocol identification. The document identifies proposed changes to the text of IEEE Std 802.1Q, 1998 Edition that have arisen as a consequence of this activity. These are documented in the usual form for amendments to IEEE 802® standards; i.e., as an explicit set of editing instructions that, if correctly applied to the text of IEEE Std 802.1Q, 1998 Edition, will create a corrected document.

Keywords: local area networks, MAC Bridge management, media access control bridges, virtual LANs

The Institute of Electrical and Electronics Engineers, Inc.
3 Park Avenue, New York, NY 10016-5997, USA

Print:	*ISBN 0-7381-2818-X*	*SH94917*
PDF:	*ISBN 0-7381-2819-8*	*SS94917*

5. Conformance

5.2 Options

Insert new list item l) and its sub-items to the end of the list, as follows:

A MAC Bridge for which conformance to this standard is claimed may

a) Support operation in Extended Filtering Mode (ISO/IEC 15802-3, 6.6.5) and the operation of GARP Multicast Registration Protocol (GMRP) (ISO/IEC 15802-3, Clause 10) as modified by Clause 10;

b) Support the ability for the Filtering Database to contain static and dynamic configuration information for more than one VLAN, by means of Static and Dynamic VLAN Registration Entries (The Filtering Database), up to a maximum of 4094 VLANs;

NOTE—The maximum number of VLANs that can be supported is 4094 rather than 4096, as the VID values 0 and FFF are reserved, as indicated in Table 9-2. As conformance to this standard is only with regard to externally visible protocol behavior, this limit on the number of VLANs that can be supported does not imply any such limitation with regard to the internal architecture of a Bridge.

c) On each Port, support both of the permissible values for the Acceptable Frame Types parameter, as defined in Acceptable Frame Types. If both values are supported, then the implementation shall support configuration of the parameter value via management;

d) Support the ability to enable and disable Ingress Filtering (8.4.5);

e) Support the ability to configure more than one VLAN whose untagged set includes that Port (8.8 and 8.11.9);

f) Support the management functionality defined in Clause 12;

g) Support more than one FID (6.4, 8.11.3, 8.11.7, and 8.11.8);

h) Support the ability to allocate more than one VID to each FID that is supported (6.4, 8.11.3, 8.11.7, and 8.11.8);

i) Support the ability to configure VLAN Learning Constraints via management (8.11.7 and 12.10.3);

j) Support the ability to configure fixed VID to FID allocations via management (8.11.7.1 and 12.10.3);

k) Support any other optional capabilities defined in ISO/IEC 15802-3, as modified by the provisions of this standard.

l) Support Port-and-Protocol-based VLAN classification, including multiple VID values per port, administrative control of the values of the multiple VIDs, and a Protocol Group Database. If this option is supported:

    1) One or more of the following Protocol Classifications and Protocol Template formats must be supported: Ethernet, RFC_1042, SNAP_8021H, SNAP_Other, or LLC_Other (8.6.1 and 8.6.2).

    2) The configuration of the contents of the Protocol Group Database may be supported.

6. Architectural overview

Change the wording in 6.3 as follows:

6.3 Relay

This is concerned with the mechanics of

a) Classifying each received frame as belonging to one and only one VLAN. This aspect of relay is determined by a set of MAC Bridge *ingress rules*;

b) Decisions related to where received frames should be forwarded. This aspect of relay is determined by a set of MAC Bridge *forwarding rules*;

c) Mapping frames for transmission through the appropriate outbound Ports, and in appropriate (VLAN-tagged or untagged) format. These aspects of relay are determined by a set of MAC Bridge *egress rules*;

d) The procedures used in order to add, modify, and remove tag headers, when relaying frames between LAN segments, in accordance with the details of the VLAN frame format (defined in Clause 9) used to carry VIDs (otherwise referred to as VLAN tags).

Clause 8 defines ingress, forwarding and egress rules, constituting a generic approach to the provision of VLAN functionality with respect to received VLAN-tagged frames, and ~~a Port-based approach~~ two approaches to the VLAN classification of received Priority-tagged and Untagged frames: Port-based VLAN classification and Port-and-Protocol-based VLAN classification. Clause 9 defines the format of the tag headers for different MAC methods, and the procedures for adding, modifying, and removing tag headers.

The ingress, forwarding, and egress rules allow Bridges to

e) Classify any received untagged frames or priority-tagged frames that are to be submitted to the Forwarding Process as belonging to a particular VLAN, as defined by the PVID or VID Set of ~~for~~ the receiving Port. The default PVID is specified in Table 9-2;

NOTE 1—This classification of untagged and priority-tagged frames is part of the functionality of the MAC relay entity (Figure 8-3, Figure 8-4), and is therefore only of significance for received frames that are potentially to be forwarded through other Ports of the Bridge (see Spanning Tree).

f) Classify any received VLAN-tagged frames that are to be submitted to the Forwarding Process as belonging to the VLAN identified by the VID carried in the tag header;

g) Make use of the VLAN classification thus associated with the received frame in order to take appropriate forwarding/filtering decisions;

h) Transmit frames in VLAN-tagged or untagged format, as defined for a given Port/VLAN pairing.

NOTE 2—This standard defines a default Port-based classification and an optional Port-and-Protocol-based classification for VLANs implemented using the procedures and VLAN frame format specified herein. End stations that transmit VLAN-tagged frames~~,~~ and ~~in the future,~~ Bridges capable of other classification methods, may actually do much of the VLAN classification of frames. More sophisticated tagging will be the rule for these devices, and bridges conformant to this standard will work with them. In this scenario, most or all LAN segments are likely to carry VLAN-tagged frames belonging to various VLANs, but each such LAN segment has its own “local default” VLAN (or set of such VLANs in the case of Port-and-Protocol-based classification). This “local default” defines the VLAN(s) to which untagged or priority-tagged frames are presumed to belong when received on the Ports of IEEE 802.1Q conformant bridges attached to that LAN segment.

NOTE 3—This standard uses the concept of a “VID Set” when Port-and-Protocol-based ingress classification rules are used. There is presumed to exist a mapping for any given untagged frame to one and only one of these VIDs. In the degenerate case, frames of all protocols are mapped onto the same, single PVID and this case is then identical to that defined by IEEE 802.1Q, 1998 Edition.

8. Principles of operation

8.1 Bridge operation

Change the wording of 8.1.2 as follows:

8.1.2 Filtering and relaying information

A Bridge filters frames, i.e., does not relay frames received by a Bridge Port to other Ports on that Bridge, in order to prevent the duplication of frames (ISO/IEC 15802-3, 6.3.4). The function that supports the use and maintenance of information for this purpose is

a) Calculation and configuration of Bridged LAN topology.

A Bridge also filters frames in order to reduce traffic in parts of the Bridged LAN that do not lie in the path between the source and destination of that traffic. The functions that support the use and maintenance of information for this purpose are

b) Permanent configuration of reserved addresses.

c) Explicit configuration of static filtering information.

d) Automatic learning of dynamic filtering information for unicast destination addresses through observation of source addresses of Bridged LAN traffic.

e) Ageing out of dynamic filtering information that has been learned.

f) Automatic addition and removal of dynamic filtering information as a result of GMRP protocol exchanges.

A Bridge classifies frames into traffic classes in order to expedite transmission of frames generated by critical or time-sensitive services. The function that supports the use and maintenance of information for this purpose is

g) Explicit configuration of traffic class information associated with the Ports of the Bridge.

A Bridge classifies untagged frames and priority-tagged frames as belonging to a particular VLAN in accordance with the *ingress rules* defined in the *ingress rules*. The function that supports the use and maintenance of information for this purpose is

h) Explicit configuration of the P~~ort~~VID and, for bridges supporting VLAN classification by Port and Protocol, VID Set (~~PVID,~~ 8.4.4) associated with each Port of the Bridge.

i) Explicit configuration of the Protocol Group Database (8.6.4) of the Bridge.

A Bridge may filter frames in order to prevent the injection of untagged and priority-tagged frames on a Port on which the reception of untagged and priority-tagged frames is disallowed. The function that supports the use and maintenance of information for this purpose is

j) Explicit configuration of the Acceptable Frame Types parameter (Acceptable Frame Types) associated with each Port of the Bridge.

A Bridge may filter frames in order to prevent the injection of traffic for a given VLAN on a Port on which that VLAN is disallowed. The function that supports the use and maintenance of information for this purpose is

k) Explicit configuration of the Enable Ingress Filtering parameter (8.4.5) associated with each Port of the Bridge.

A Bridge filters frames in order to confine traffic destined for a given VLAN to LAN segments that form a path from the source of the traffic to recipients that are members of that VLAN. The functions that support the use and maintenance of information for this purpose are

l) Automatic configuration of Dynamic VLAN Registration Entries by means of GVRP (8.11.5 and 11.2);

m) Explicit configuration of management controls associated with the operation of GVRP by means of Static VLAN Registration Entries (8.11.2 and 11.2);

n) Automatic learning of MAC Addresses in associated VLANs through the observation of network traffic (The Learning Process).

A Bridge adds and removes tag headers (9.3) from frames, and performs the associated frame translations that may be required, in accordance with the *egress rules* (8.8). The function that supports the use and maintenance of information for this purpose is

o) Explicit configuration of tagging requirements on egress for each Port (8.11.2 and 8.11.9).

8.4 Port States, Port parameters, Active Ports, and the active topology

Change the wording of 8.4.4 as follows:

8.4.4 Port VLAN identifier and VID Set

A VLAN Bridge supports Port-based VLAN classification, and may, in addition, support Port-and-Protocol-based VLAN classification.

In Port-based VLAN classification within a Bridge, the VID associated with an untagged or priority-tagged frame (i.e., a frame with no tag header, or a frame with a tag header that carries the null VLAN ID) is determined, based on the Port of arrival of the frame into the Bridge, as described in 8.6. This classification mechanism requires the association of a specific ~~VLAN ID, the~~ *Port VLAN Identifier*, or *PVID*, with each of the Bridge's Ports. In this case, t~~T~~he PVID for a given Port provides the VID for untagged and priority-tagged frames received through that Port.

In addition to the PVID, for bridges that implement Port-and-Protocol-based VLAN classification, the VID associated with an Untagged or Priority-tagged Frame is determined based on the Port of arrival of the frame into the bridge and on the protocol identifier of the frame, as described in 8.6. This classification mechanism requires the association of multiple VIDs with each of the Ports of the Bridge: this is known as the "VID Set" for that port. Each VID of a Port of a Bridge that supports Port-and-Protocol-based VLAN classification is also associated with a Protocol Group Identifier. A Protocol Group Identifier is not relevant in a Bridge that supports only Port-based VLAN classification. The contents of the VID Set for each port may be configured by management. The VID Set is in addition to the PVID value described above.

The PVID and VID Set for each Port shall contain ~~a~~ valid VID values~~, and shall not contain the value of the null VLAN ID~~ (Table 9-2).

NOTE 1—This rule ensures that the process of ingress classification of frames always associates a ~~non-null~~valid VID with each received frame. As a consequence, a VLAN-aware Bridge can never transmit priority-tagged frames; all frames transmitted are either untagged or carry a ~~non-null~~valid VID in their tag header.

The PVID and VID Set values may be configured by management, if management operations are supported by the implementation. If no PVID value has been explicitly configured for a Port, the PVID shall assume the value of the default PVID defined in Table 9-2 and the VID Set shall be empty.

NOTE 2—If a Bridge is configured so that, for any Port, all the members of the VID Set of that port assume the same VID, then it is impossible to tell from the frame relay behavior of the bridge whether the Bridge supports Port-based or Port-and-Protocol-based VLAN classification. In particular, the default frame relay behavior (the frame relay behavior before any administrative actions on the Bridge) of a Bridge that supports Port-and-Protocol-based VLAN classification is the same as the default frame relay behavior of a Bridge that supports only Port-based VLAN classification.

8.6 The ingress rules

Insert the following figure as Figure 8-8 in 8.6, and renumber subsequent figures accordingly:

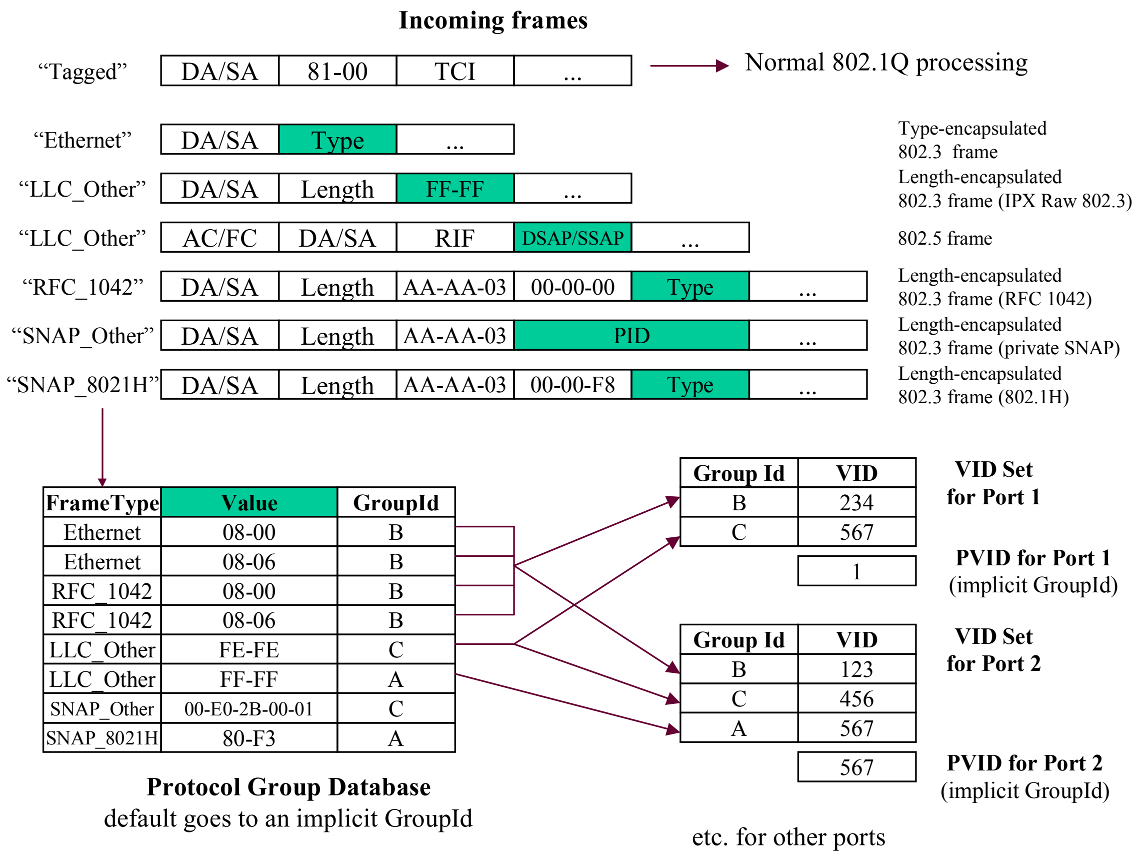
Incoming frames
"Tagged" DA/SA 81-00 TCI ...
Normal 802.1Q processing
"Ethernet" DA/SA Type ...
Type-encapsulated 802.3 frame
"LLC_Other" DA/SA Length FF-FF ...
Length-encapsulated 802.3 frame (IPX Raw 802.3)
"LLC_Other" AC/FC DA/SA RIF DSAP/SSAP ...
802.5 frame
"RFC_1042" DA/SA Length AA-AA-03 00-00-00 Type ...
Length-encapsulated 802.3 frame (RFC 1042)
"SNAP_Other" DA/SA Length AA-AA-03 PID ...
Length-encapsulated 802.3 frame (private SNAP)
"SNAP_8021H" DA/SA Length AA-AA-03 00-00-F8 Type ...
Length-encapsulated 802.3 frame (802.1H)
VID Set for Port 1
PVID for Port 1 (implicit GroupId)
VID Set for Port 2
PVID for Port 2 (implicit GroupId)
etc. for other ports

FrameType	Value	GroupId
Ethernet	08-00	B
Ethernet	08-06	B
RFC_1042	08-00	B
RFC_1042	08-06	B
LLC_Other	FE-FE	C
LLC_Other	FF-FF	A
SNAP_Other	00-E0-2B-00-01	C
SNAP_8021H	80-F3	A

Protocol Group Database
default goes to an implicit GroupId

VID Set for Port 1

Group Id	VID
B	234
C	567

PVID for Port 1 (implicit GroupId): 1

VID Set for Port 2

Group Id	VID
B	123
C	456
A	567

PVID for Port 2 (implicit GroupId): 567

NOTE—The PID shown in this figure is a Protocol Identifier, as defined in 5.3 of IEEE Std 802 . It is a 5-octet value, consisting of a 3-octet OUI value followed by a 2-octet locally administered identifier.

Figure 8-8—Example of operation of port-and-protocol based classification

Replace the existing text of 8.6, inserting new subclauses 8.6.1 through 8.6.4, and modifying the existing text of 8.6 to form subclause 8.6.5, as follows:

8.6.1 Protocol Classification

Protocol Classification is not defined for Bridges that support only Port-based VLAN classification.

For frames received from the E-ISS on bridge ports which implement Port-and-Protocol-based VLAN classification, the following procedures are followed in order to classify the frame's format and protocol. These procedures are described as if they assigned values to parameters that are used as input to the Ingress Rules.

The **detagged_frame_type** parameter indicates the frame format. The value is determined as follows:

a) If the frame is Untagged or Priority Tagged, this parameter is present and indicates the link-layer encapsulation format of the *Detagged Frame*. The Detagged Frame of an Untagged Frame is the Frame itself. The Detagged Frame of a Tagged Frame or Priority Tagged Frame is the Frame which results from untagging the Frame by the procedure described in 9.1. The value of detagged_frame_type is as follows:

   1) Ethernet, if the Detagged Frame uses Type-encapsulated 802.3 format

   2) RFC_1042, if the Detagged Frame is of the format specified by 10.5 in IEEE Std 802 for the encoding of an IEEE 802.3 Type Field in an 802.2/SNAP header (this supersedes the original definition, which appeared in RFC 1042)

   3) SNAP_8021H, if the Detagged Frame is of the format specified by IEEE Std. 802.1H for the encoding of an IEEE 802.3 Type Field in an 802.2/SNAP header

   4) SNAP_Other, if the Detagged Frame contains an LLC UI PDU with DSAP and SSAP fields equal to the LLC address reserved for SNAP and the 5-octet SNAP Protocol Identifier (PID) value is not in either of the ranges used for RFC_1042 or SNAP_8021H above

   5) LLC_Other, if the Detagged Frame contains both a DSAP and an SSAP address field in the positions specified by IEEE 802.2 Logical Link Control, but is not any of the formats described for LLC frames above

b) Else the parameter is not present.

The **ethertype** parameter is present if the detagged_frame_type parameter is present and has the value Ethernet, RFC_1042, or SNAP_8021H. Its value is the IEEE 802.3 Type Field present in the Detagged Frame.The value is determined as follows:

c) If the detagged_frame_type parameter is present and has the value Ethernet, RFC_1042, or SNAP_8021H, then this parameter is present and has the value of the IEEE 802.3 Type Field present in the Detagged Frame.

d) Else the parameter is not present.

The **llc_saps** parameter is present if the detagged_frame_type parameter is present and has the value LLC_Other. Its value is determined as follows:

e) If the detagged_frame_type parameter is present and has the value LLC_Other then this parameter is present and its value is the pair of LLC 802.2 DSAP and SSAP address field values.

f) Else the parameter is not present.

NOTE 1—A frame that is encapsulated using values of hex FF/FF in the position where an LLC header is to be expected (as defined by IEEE Std 802.2, 1998 Edition) is known as a "Novell IPX Raw" encapsulation. Such frames do not conform to IEEE Std 802.2 in that they do not include some of the other required LLC fields. For the purposes of this standard, they are treated as LLC_Other, regardless of whether they are legal LLC frames or not.

NOTE 2—Bridges are not required, for the purposes of this standard, to completely verify the format of frames as meeting IEEE Std 802.2 or not: they are only required to recognize the DSAP and SSAP fields of such frames.

The **snap_pid** parameter is present if the detagged_frame_type parameter is present and has the value SNAP_Other. Its value is determined as follows:

g) If the detagged_frame_type parameter is present and has the value SNAP_Other then the parameter is present and its value is the contents of the 5 octets following the LLC header, i.e., the PID field.

h) Else the parameter is not present.

8.6.2 Protocol Templates

Protocol Templates are not defined for Bridges that support only Port-based VLAN classification.

In a Bridge that supports Port-and-Protocol-based VLAN classification, a Protocol Template is a tuple that specifies a protocol to be identified in received frames. A Protocol Template has one of the following formats:

a) A value "Ethernet" and a 16-bit IEEE 802.3 Type Field value

b) A value "RFC_1042" and a 16-bit IEEE 802.3 Type Field value

c) A value "SNAP_8021H" and a 16-bit IEEE 802.3 Type Field value

d) A value "SNAP_Other" and a 40-bit PID value

e) A value "LLC_Other" and a pair of IEEE 802.2 LSAP values: DSAP and SSAP

A Protocol Template *matches* a Frame if

f) The Frame's detagged_frame_type is Ethernet, the Protocol Template is of type Ethernet, and the frame's IEEE 802.3 Type Field is equal to the value of the IEEE 802.3 Type Field of the Protocol Template, or

g) The Frame's detagged_frame_type is RFC_1042, the Protocol Template is of type RFC_1042 and the frame's IEEE 802.3 Type Field is equal to the IEEE 802.3 Type Field of the Protocol Template, or

h) The Frame's detagged_frame_type is SNAP_8021H, the Protocol Template is of type SNAP_8021H, and the frame's IEEE 802.3 Type Field is equal to the IEEE 802.3 Type Field of the Protocol Template, or

i) The Frame's detagged_frame_type is SNAP_Other, the Protocol Template is of type SNAP_Other, and the frame's snap_pid is equal to the PID of the Protocol Template, or

j) The Frame's detagged_frame_type is LLC_Other, the Protocol Template is of type LLC_Other, and the frame's llc_saps matches the value of the DSAP and SSAP of the Protocol Template.

NOTE—If a port does not support Protocol Templates of the Frame's detagged_frame_type then no match will occur.

8.6.3 Protocol Group Identifiers

Protocol Group Identifiers are undefined in a Bridge that supports only Port-based VLAN classification.

A Bridge that supports Port-and-Protocol-based VLAN classification shall support Protocol Group Identifiers.

A Protocol Group Identifier, shown as "Group Id" in Figure 8-8, designates a group of protocols that will be associated with one member of the VID Set of a Port. The association of protocols into groups is established by the contents of the Protocol Group Database, as described in 8.6.4. The identifier has scope only within a single bridge.

There is an implicit Protocol Group Identifier that is assigned to frames that match none of the entries in the

Protocol Group Database. Therefore, every incoming Frame can be assigned to a Protocol Group Identifier.

8.6.4 Protocol Group Database

A Protocol Group Database is not defined in Bridges that support only Port-based VLAN classification.

A Bridge that supports Port-and-Protocol-based VLAN classification, shall support a single Protocol Group Database. The Protocol Group Database groups together a set of one or more Protocols by assigning them the same Protocol Group Identifier (8.6.3). Each entry of the Protocol Group Database comprises the following:

a) A Protocol Template

b) A Protocol Group Identifier

The Protocol Group Database specifies a mapping from Protocol Templates to Protocol Group Identifiers: if two entries of the Protocol Group Database contain different Protocol Group Identifiers then their Protocol Templates must also be different.

The entries of the Protocol Group Database may be configured by management. A Bridge that supports Port-and-Protocol-based VLAN classification shall support at least one of the formats of Protocol Template.

An implicit Protocol Group Database entry exists that matches all frames: this entry is invoked for frames that do not match the template of any of the other entries. It references an implicit Protocol Group Identifier that selects the PVID on each port. In this way, it is ensured that all incoming Frames are matched by a Protocol Group Identifier and, hence, are assigned to a VID.

NOTE—If there are no entries in the Protocol Group Database, then the frame relay behavior of this Bridge is identical to the frame relay behavior of a Bridge having the same number of Ports that supports only Port-based VLAN classification.

8.6.5 VLAN Classification rules

If the vlan_identifier parameter carried in a received data indication is equal to the null VLAN ID (Table 9-2) and the Acceptable Frame Types parameter (8.4.3) for the Port through which the frame was received is set to the value *Admit Only VLAN-tagged frames*, then the frame shall be discarded.

Each frame received by a VLAN Bridge shall be classified as belonging to exactly one VLAN by associating a VID value with the received frame. The classification is achieved as follows:

a) If the vlan_identifier parameter carried in a received data indication is the null VLAN ID (Table 9-2), and the Bridge supports only Port-based VLAN classification, then the VID for the Frame is the unique PVID associated with the Port through which the frame was received (8.4.4). Otherwise:

   ~~1) If the implementation supports further VLAN classification rules in addition to Port-based classification (D.2.2), and if the application of these rules associates a non-null VID value with the frame, then that VID value is used.~~

   ~~2) If the implementation supports only Port-based classification, or if any additional classification rules supported are unable to associate a non-null VID with the frame, then the PVID value associated with the Port through which the frame was received is used (Port VLAN identifier and VID Set).~~

b) If the vlan_identifier parameter carried in a received data indication is the null VLAN ID and the Bridge supports Port-and-Protocol-based VLAN classification, then the VID for the Frame is selected from the VID Set of the port through which the Frame was received. The VID selected is the member of the VID Set (8.4.4) for which the associated Protocol Group Identifier (8.6.3) is equal to the Protocol Group Identifier of the Frame. If no matches are found then the VID for the frame is the PVID associated with the Port. Otherwise:

c) ~~If the vlan_identifier parameter carried in a received data indication is not the null VLAN ID (Table 9-2), then t~~The VID for the Frame is the vlan_identifier parameter value~~ is used~~.

NOTE 1—As defined in 7.1.2, the vlan_identifier parameter carries the null VLAN ID if the frame was not VLAN-tagged. There are two cases; either the frame was untagged, or the frame was tagged and the tag header carried a VID value equal to the null VLAN ID (i.e., a priority-tagged frame).

NOTE 2—VIDs of value FFF cannot be configured in any Filtering Database entry (see Table 9-2). Consequently, any incoming frame whose VLAN classification is FFF will be discarded by the Forwarding Process.

The VID value thus identified, known as the *VLAN classification* of the frame, is used as the value of the vlan_classification parameter of any corresponding data request primitives.

If the Enable Ingress Filtering parameter (8.4.5) for the Port through which the frame was received is set, and if the Port is not in the Member set (8.11.9) for the frame's VLAN classification, then the frame is discarded.

All frames that are not discarded as a result of the application of the ingress rules are submitted to the Forwarding Process and to the Learning Process. All frames that are discarded as a result of the application of the ingress rules are not submitted either to the Forwarding Process or to the Learning Process.

8.14 Addressing

Change the wording of 8.14.7 as follows:

8.14.7 Points of attachment and connectivity for Higher Layer Entities

Higher Layer Entities such as the Bridge Protocol Entity and GARP Protocol Entity (8.12), and Bridge Management (8.13) are modeled as being connected directly to the Bridged LAN via one or more points of attachment. From the point of view of their attachment to the Bridged LAN, Higher Layer Entities associated with a Bridge can be regarded as if they are distinct end stations, directly connected to one or more of the LAN segments served by the Bridge Ports, in the same way as any other end station is connected to the Bridged LAN. In practice, the Higher Layer Entities will, in many cases, share the same physical points of attachment used by the relay function of the Bridge, as stated in 8.14; however, from the point of view of the transmission and reception of frames by these functions, the behavior is the same as if they were contained in logically separate end stations with points of attachment "outside" the Port(s) with which they are associated. ~~Figure 8-9~~ Figure 8-10 is functionally equivalent to Figure 8-3, but illustrates this logical separation between the points of attachment used by the Higher Layer Entities and points of attachment used by the MAC Relay Entity.

Higher Layer Entities fall into two distinct categories:

a) Those entities, such as the Bridge Management Entity, that require only a single point of attachment to the Bridged LAN;

b) Those entities, such as Bridge Protocol Entities and GARP Participants, that require a point of attachment per Port of the Bridge.

The fundamental distinction between these two categories is that for the latter, it is essential for the operation of the entity concerned that it is able to associate received frames with the LAN segment on which those frames were originally seen by the Bridge, and that it is able to transmit frames to peer entities that are connected directly to that LAN segment. It is therefore essential that

c) It does not receive frames via a point of attachment associated with one Port that have been relayed by the Bridge from other Ports; and

d) Frames that it transmits via one point of attachment are not relayed by the Bridge to any other Ports.

For this reason, the MAC Addresses used to reach entities of this type are permanently configured in the Filtering Database in order to prevent the Bridge from relaying such frames received via any Port to any other Port of the Bridge, as defined in 8.14.3 and 8.14.6.

NOTE 1—The MAC Addresses used to address such entities are generally group MAC Addresses.

The MAC Relay Entity forwards a frame received on one Port through the other Port(s) of the Bridge, subject to the following control information permitting such forwarding to take place:

e) The Port state information (8.4) associated with the Port on which the frame was received;

f) The information held in the Filtering Database (8.11);

g) The Port state information (8.4) associated with the Port(s) on which the frame is potentially to be transmitted.

This is illustrated in ~~Figure 8-10~~ Figure 8-11, where the control information represented by the Port state and Filtering Database information is represented as a series of switches (shown in the open, disconnected state) inserted in the forwarding path provided by the MAC Relay Entity. For the Bridge to forward a given frame between two Ports, all three switches must be in the closed state. This figure also illustrates that the controls placed in the forwarding path have no effect upon the ability of a Higher Layer Entity to transmit and receive frames directly onto a given LAN segment via the point of attachment to that segment (e.g., from entity A to segment A); they only affect the path taken by any indirect transmission/reception (e.g., from entity A to segment B).

~~Figure 8-11~~ Figure 8-12 illustrates the state of the forwarding path with respect to frames destined for Higher Layer Entities that require per-Port points of attachment. The fact that the Filtering Databases in all Bridges are permanently configured to prevent relay of frames addressed to these entities means that they can receive frames only via their direct points of attachment (i.e., from segment A to entity A, and from segment B to entity B), regardless of the Port states.

~~Figure 8-12~~ Figure 8-13 illustrates the state of the forwarding path with respect to frames destined for a Higher Layer Entity that requires only a single point of attachment, for the case where the Port states and Filtering Database states permit relay of frames. Frames destined for the Higher Layer Entity that originate on LAN segment B are relayed by the Bridge, and are both received by the entity and transmitted on LAN segment A.

~~Figure 8-13~~ Figure 8-14 illustrates the state of the forwarding path with respect to frames destined for a Higher Layer Entity that requires only a single point of attachment, for the case where one of the Port states does not permit relay. Frames destined for the Higher Layer Entity that originate on LAN segment A are received by the entity; however, frames that originate on LAN segment B are not relayed by the Bridge, and can therefore only be received by the entity if there is some other forwarding path provided by other components of the Bridged LAN between segments A and B.

NOTE 2—If the Port state shown in ~~Figure 8-13~~ Figure 8-14 occurs as a result of the normal operation of the Spanning Tree (as opposed to being a result of equipment failure, or administrative control of Port state information), then such a path will exist, either via another Port of this Bridge (not shown in the diagram) connected to segment A, or via one or more Bridges providing a path between segments A and B. If there is no active Spanning Tree path from segment B to segment A, then the Bridged LAN has partitioned into two separate Bridged LANs, one on either side of this Port, and the Higher Layer Entity shown is only reachable via segment A.

In VLAN-aware Bridges, two more switches appear in the forwarding path, corresponding to the ingress and egress rules defined in 8.6 and 8.8, as illustrated in Figure 8-14.

As with Port state information, the configuration of the ingress and egress rules does not affect the reception of frames received on the same LAN segment as a Higher Layer Entity's point of attachment. For example, the reception of a frame by Higher Layer Entity A that was transmitted on LAN Segment A is unaffected by the ingress or egress configuration of either Port. However, for Higher Layer Entities that require only a single point of attachment, the ingress and egress configuration affects the forwarding path. For example, frames destined for Higher Layer Entity A that are transmitted on LAN Segment B would be subjected to the ingress rules that apply to Port B and the egress rules that apply to Port A.

The decision as to whether frames transmitted by Higher Layer Entities are VLAN-tagged or untagged depends upon the Higher Layer Entity concerned, and the connectivity that it requires.

h) Spanning Tree BPDUs transmitted by the Bridge Protocol Entity are not forwarded by Bridges, and must be visible to all other BPEs attached to the same LAN segment. Such frames shall be transmitted untagged;

NOTE 3—Any BPDUs or GVRP PDUs that carry a tag header are not recognized as well-formed BPDUs or GVRP PDUs and are not forwarded by the Bridge.

i) The definition of the GVRP application (11.2.3) calls for all GVRP frames to be transmitted untagged for similar reasons;

j) The definition of the GMRP application (Clause 10) calls for all GMRP frames originating from VLAN-aware devices to be transmitted VLAN-tagged, in order for the VID in the tag to be used to identify the VLAN context in which the registration applies;

k) It may be necessary for PDUs transmitted for Bridge Management (8.13) to be VLAN-tagged in order to achieve the necessary connectivity for management in a VLAN Bridged LAN. In order to access a Bridge Management entity located in a region of the network that is served only by a given set of VLANs, it may be necessary to communicate with that entity using frames VLAN-tagged with one of the VIDs concerned, unless one of those VIDs also happens to be the PVID (or a member of the VID Set that matches the appropriate Protocol) for the Port serving the management station.

9. Tagged frame format

Insert new wording into NOTE 1 and change Table 9-2 as follows:

9.3.2.3 VID format

The twelve-bit VLAN Identifier (VID) field uniquely identify the VLAN to which the frame belongs. The VID is encoded as an unsigned binary number. Table 9-1 identifies values of the VID field that have specific meanings or uses; the remaining values of VID are available for general use as VLAN identifiers.

A priority-tagged frame is a tagged frame whose tag header contains a VID value equal to the null VLAN ID.

NOTE 1—The specification of the ingress and egress rules for VLAN-Aware Bridges (8.6, 8.8) is such that a Bridge does not propagate priority-tagged frames; a received priority-tagged frame will acquire an appropriate VLAN classification on ingress as determined by the Ingress rules (8.6) and will, therefore, either be forwarded as an untagged frame~~,~~ or as a tagged frame tagged with that VLAN classification, depending upon the egress configuration for that VLAN. P~~p~~riority-tagged frames are therefore only ever generated by end stations.

A VLAN-tagged frame is a tagged frame whose tag header contains a VID value other than the null VLAN ID.

Table 9-1—Reserved VID values

VID value (hexadecimal)	Meaning/Use
0	The null VLAN ID. Indicates that the tag header contains only user_priority information; no VLAN identifier is present in the frame. This VID value shall not be configured as a PVID or a member of a VID Set, or configured in any Filtering Database entry, or used in any Management operation.
1	The default PVID value used for classifying frames on ingress through a Bridge Port. The PVID value of a Port can be changed by management ~~on a per-Port basis~~.
FFF	Reserved for implementation use. This VID value shall not be configured as a PVID or a member of a VID Set, configured in any Filtering Database entry, used in any Management operation, or transmitted in a tag header.

A Bridge may implement the ability to support less than the full range of VID values; i.e., for a given implementation, an upper limit, N, is defined for the VID values supported, where N is less than or equal to 4094. All implementations shall support the use of all VID values in the range 0 through their defined maximum VID, N.

NOTE 2—There is a distinction made here between the range of VID *values* (0 through N) that an implementation can support as identifiers for its active VLANs, and the maximum number of active VLANs (V) that it is able to support at any one time. An implementation that supports a maximum of, say, only 16 active VLANs (V=16) can support VIDs for those VLANs that are chosen from anywhere in the full VID number space (i.e., support N=4094), or from a subset of that number space (i.e., support N<4094). Therefore N is always greater than or equal to V.

11. VLAN topology management

Change wording in 11.2.1.3 as follows:

11.2.1.3 Use of the PVID and VID Set

The initial state of the Permanent Database contains a Static VLAN Registration Entry for the Default PVID, in which the Port Map indicates Registration Fixed on all Ports. This ensures that in the default state, where the value of every PVID ~~on all~~ of each Port~~s~~ is the Default PVID and where the VID Set of each Port is empty, membership of the Default PVID is propagated across the Bridged LAN to all other GVRP-aware devices. Subsequent management action may change both the Permanent Database and the Filtering Database in order to modify or remove this initial setting, and may change the PVID and/or VID Set value(s) on ~~each~~ any Port of the Bridge.

NOTE—In the absence of any modification of these initial settings, this ensures that connectivity is established across the Bridged LAN for the VLAN corresponding to the Default PVID.

12. VLAN Bridge Management

Change the wording in 12.1.1 as follows:

12.1.1 Configuration Management

Configuration Management provides for the identification of communications resources, initialization, reset and close-down, the supply of operational parameters, and the establishment and discovery of the relationship between resources. The facilities provided by Bridge Management in this functional area are

a) The identification of all Bridges that together make up the Bridged LAN and their respective locations and, as a consequence of that identification, the location of specific end stations to particular individual LANs.

b) The ability to remotely reset, i.e., reinitialize, specified Bridges.

c) The ability to control the priority with which a Bridge Port transmits frames.

d) The ability to force a specific configuration of the spanning tree.

e) The ability to control the propagation of frames with specific group MAC Addresses to certain parts of the configured Bridged LAN.

f) The ability to identify the VLANs in use, and through which Ports of the Bridge and for which Protocols frames destined for a given VLAN may be received and/or forwarded.

12.10 Bridge VLAN managed objects

Change the wording in 12.10.1 as follows:

12.10.1 Bridge VLAN Configuration managed object

The Bridge VLAN Configuration managed object models operations that modify, or enquire about, the overall configuration of the Bridge's VLAN resources. There is a single Bridge VLAN Configuration managed object per Bridge.

The management operations that can be performed on the Bridge VLAN Configuration managed object are

a) Read Bridge VLAN Configuration (12.10.1.1);

b) Configure PVID and VID Set values (12.10.1.2);

c) Configure Acceptable Frame Types parameters (12.10.1.3);

d) Configure Enable Ingress Filtering parameters (12.10.1.4);

e) Reset VLAN Bridge (12.10.1.5);

f) Notify VLAN registration failure (12.10.1.6)~~.~~;

g) Configure Protocol Group Database (12.10.2.1);

h) Configure VLAN Learning Constraints (12.10.3).

12.10.1.1 Read Bridge VLAN Configuration

12.10.1.1.1 Purpose

To obtain general VLAN information from a Bridge.

12.10.1.1.2 Inputs

None.

12.10.1.1.3 Outputs

a) The 802.1Q VLAN Version number. Reported as “1” by devices that implement VLAN functionality according to this edition of the standard;

b) The optional VLAN features supported by the implementation:

   1) The maximum number of VLANs supported;

   2) Whether the implementation supports the ability to override the default PVID setting, and its egress status (VLAN-tagged or untagged) on each Port.

   3) For a Bridge that supports Port-and-Protocol-based VLAN classification, which of the Protocol Template formats (8.6.2) are supported by the implementation.

c) For each Port:

   1) The Port number;

   2) The PVID value (8.4.4) currently assigned on that Port;

   3) For a Bridge that supports Port-and-Protocol-based VLAN classification, whether the implementation supports Port-and-Protocol-based VLAN classification on that Port;

   4) For a Bridge that supports Port-and-Protocol-based VLAN classification on that Port, the maximum number of entries supported in the VID Set on that Port; the VID value and Protocol Group Identifier currently assigned to each entry in the VID Set (8.4.4) on that Port;

   5) The state of the Acceptable Frame Types parameter (8.4.3). The permissible values for this parameter are as follows:

      i) Admit only VLAN-tagged frames;

      ii) Admit all frames.

   6) The state of the Enable Ingress Filtering parameter (8.4.5); Enabled or Disabled.

d) For a Bridge that supports Port-and-Protocol-based VLAN classification: the contents of the Protocol Group Database comprising a set of {Protocol Template, Protocol Group Identifier} bindings (8.6.1, 8.6.3, and 8.6.4); the maximum number of entries supported in the Protocol Group Database.

12.10.1.2 Configure PVID and VID Set values

12.10.1.2.1 Purpose

To configure the PVID and VID Set value(s) (8.4.4) associated with one or more Ports.

12.10.1.2.2 Inputs

a) For each Port to be configured, a Port number and the PVID value to be associated with that Port.

b) In addition, for a Bridge that supports Port-and-Protocol-based VLAN classification: for each Port to be configured, a Port number, a Protocol Group Identifier, and a VID value for the member of the Port's VID Set that is to be configured.

12.10.1.2.3 Outputs

a) Operation status for each Port to be configured. This takes one of the following values:
   1) Operation rejected due to there being no spare VID Set entries on this Port; or
   2) Operation rejected due to the PVID or VID being out of the supported range for this Port; or
   3) Operation accepted.

~~None.~~

12.10.1.5 Reset VLAN Bridge

12.10.1.5.1 Purpose

To reset all statically configured VLAN-related information in the Bridge to its default state. This operation

a) Deletes all VLAN Configuration managed objects;

b) Resets the PVID associated with each Bridge Port to the Default PVID value (Table 9-2);

c) Removes all entries in the Protocol Group Database and removes all members of the VID Set on each port, for a Bridge that supports Port-and-Protocol-based VLAN classification;

d) Resets the Acceptable Frame Types parameter value associated with each Port to the default value (8.4.3).

12.10.1.5.2 Inputs

None.

12.10.1.5.3 Outputs

None.

12.10.2 VLAN Configuration managed object

Change the wording in 12.10.2 as follows:

12.10.2.1 Configure Protocol Group Database

To configure a Protocol Group Database (8.6.4) entry. This operation is not applicable to a Bridge that does not support Port-and-Protocol-based VLAN classification.

NOTE—Implementation of the Configure Protocol Group Database operation is not mandatory; conformant implementations may implement a fixed set of Protocol Group Database entries.

12.10.2.1.1 Inputs

a) A value representing the frame format to be matched: Ethernet, RFC_1042, SNAP_8021H, SNAP_Other or LLC_Other (8.6.1);

b) One of

   1) An IEEE 802.3 Type value, for matching frame formats of Ethernet, RFC_1042, or SNAP_8021H;

   2) A 40-bit Protocol ID (PID), for matching frame formats of SNAP_Other;

   3) A pair of IEEE 802.2 DSAP and SSAP address field values, for matching frame formats of LLC_Other;

c) A Protocol Group Identifier (8.6.3).

12.10.2.1.2 Outputs

a) Operation status. This takes one of the following values:

   1) Operation rejected due to there being no spare Protocol Group Database entries; or

   2) Operation rejected due to an unsupported frame format; or

   3) Operation rejected due to an unsupported value for an IEEE 802.3 Type value, PID, DSAP, or SSAP; or

   4) Operation accepted.

Annex A

(normative)

PICS proforma[1]

A.5 Major capabilities and options

Add items (23l) and (23m) to the following table.

Item	Feature	Status	References	Support
(1a)*	Communications Support		{D}6.5	
	Which MAC types are supported on Bridge Ports, implemented in conformance with the relevant MAC standards?			
(1a.1)*	CSMA/CD, IEEE Std 802.3	O.1		Yes [ ] No [ ]
(1a.2)*	Token Bus, ISO/IEC 8802-4	O.1		Yes [ ] No [ ]
(1a.3)*	Token Ring, ISO/IEC 8802-5	O.1		Yes [ ] No [ ]
(1a.4)*	FDDI, ISO 9314-2	O.1		Yes [ ] No [ ]
(1a.5)*	DQDB, ISO/IEC 8802-6	O.1		Yes [ ] No [ ]
(1a.6)*	ISLAN, ISO/IEC 8802-9	O.1		Yes [ ] No [ ]
(1a.7)*	ISLAN 16-T, IEEE 802.9a	O.1		Yes [ ] No [ ]
(1a.8)*	Demand Priority, ISO/IEC 8802-12 (IEEE Std 802.3 format)	O.1		Yes [ ] No [ ]
(1a.9)*	Demand Priority, ISO/IEC 8802-12 (ISO/IEC 8802-5 format)	O.1		Yes [ ] No [ ]
(1b)	Is LLC Type 1 supported on all Bridge Ports in conformance with ISO/IEC 8802-2?	M	8.2, 8.3, 8.14, ISO/IEC 8802-2	Yes [ ]
(1c)*	Is Source-Routing Transparent Bridge operation supported on any of the Bridge Ports? (If support is claimed, the PICS proforma detailed in ISO/IEC 15802-3, Annex D, shall also be completed).	O	{D}Annex C	Yes [ ] No [ ]
(2)	Relay and filtering of frames (A.6)	M	8.5, 8.9, 8.6, 8.7, 8.8	Yes [ ]
(2a)	Does the Bridge support Basic Filtering Services?	M	{D}6.6.5, 8.7.2	Yes [ ]
(2b)*	Does the Bridge support Extended Filtering Services?	O	{D}6.6.5, 8.7.2	Yes [ ] No [ ]
	If item (2b) is not supported, mark "N/A" and continue at (2e).			N/A[ ]
(2c)*	Does the Bridge support dynamic Group forwarding and filtering behavior?	2b:M	{D}6.6.5	Yes [ ] No [ ]

[1] *Copyright release for PICS proformas:* Users of this standard may freely reproduce the PICS proforma in this annex so that it can be used for its intended purpose and may further publish the completed PICS.

Item	Feature	Status	References	Support
(2d)	Does the Bridge support the ability for static filtering information for individual MAC Addresses to specify a subset of Ports for which forwarding or filtering decisions are taken on the basis of dynamic filtering information?	2b:O	{D}6.6.5	Yes [ ] No [ ]
(2e)*	Does the Bridge support expedited traffic classes on any of its Ports?	O	8.1.2, 8.7.3	Yes [ ] No [ ]
(4)*	Does the Bridge support management of the priority of relayed frames?	O	{D}6.5, 8.5.1, 8.7.3, 8.7.5, Table 8-1, Table 8-2, Table 8-3	Yes [ ] No [ ]
(5)	Maintenance of filtering information (A.7)	M	8.10, 8.11	Yes [ ]
(7a)	Can the Filtering Database be read by management?	O	8.11	Yes [ ] No [ ]
(7c)*	Can Static Filtering Entries be created and deleted?	O	8.11.1	Yes [ ] No [ ]
(7g)	Can Static Filtering Entries be created and deleted in the Permanent Database?	O	8.11.10	Yes [ ] No [ ]
(7h)	Can Static Filtering Entries be created for a given MAC Address specification with a distinct Port Map for each inbound Port?	O	8.11.1	Yes [ ] No [ ]
(7i)	Can Group Registration Entries be dynamically created, updated and deleted by GMRP?	2c:M	8.11.4, {D}10	Yes [ ] N/A [ ]
(10)	Addressing (A.8)	M	8.14	Yes [ ]
(9a)*	Can the Bridge be configured to use 48-bit Universal Addresses?	O.3	8.14	Yes [ ] No [ ]
(9b)*	Can the Bridge be configured to use 48-bit Local Addresses?	O.3	8.14	Yes [ ] No [ ]
(13)*	Spanning Tree algorithm and protocol (A.9)	M	{D}8, {D}9	Yes [ ]
(16)*	Does the Bridge support management of the Spanning Tree topology?	O	{D}8.2	Yes [ ] No [ ]
(17)*	Does the Bridge support management of the protocol timers?	O	{D}8.10	Yes [ ] No [ ]
(19)* (20a)* (20b)*	VLAN Bridge Management Operations Are the Bridge Management Operations supported via a Remote Management Protocol? Are the Bridge Management Operations supported via a local management interface?	O 19:O.4 19:O.4	12 {D}5 {D}5	Yes [ ] No [ ] Yes [ ] No [ ] N/A [ ] Yes [ ] No [ ] N/A [ ]

Item	Feature	Status	References	Support
(23a)*	Does the implementation support, on each Port, one or more of the permissible combinations of values for the Acceptable Frame Types parameter?	M	5.1, 8.4.3	Yes [ ]
(23a.1)	State which Ports support: — Admit only VLAN-tagged frames; — Admit all frames.	M	5.1, 8.4.3	Ports: _____ Ports: _____
(23a.2)	On Ports that support both values, is the parameter configurable via management?	M	5.1, 8.4.3, 12.10	Yes [ ] N/A [ ]
(23b)	Does the implementation support the ability to insert tag headers into, modify tag headers in, and remove tag headers from relayed frames, as required by the capabilities of each Bridge Port?	M	5.1, 7.1, 9	Yes [ ]
(23c)	Does the implementation support the ability to perform automatic configuration and management of VLAN topology information by means of GVRP on all Ports?	M	5.1, 11	Yes [ ]
(23d)	Does the implementation support the ability for the Filtering Database to contain static and dynamic configuration information for at least one VLAN, by means of Static and Dynamic VLAN Registration Entries?	M	5.1, 8.11	Yes [ ]
(23d.1)	State the maximum number of VLANs supported by the implementation.	M	5.1, 8.11, 9.3.2.3	_____ VLANs
(23d.2)	State the range of VID values supported by the implementation.	M	8.11, 9.3.2.3	0 through _____
(23e)*	VLAN Learning support		5.1, 8.11.3, 8.11.7, 8.11.8	
(23e.1)	Does the implementation support at least one FID?	M		Yes [ ]
(23e.2)	Can the implementation allocate at least one VID to each FID supported?	M		Yes [ ]
(23e.4)	State the maximum number of FIDs that can be supported by the implementation.	M	8.11.7	____ FIDs
(23e.5)	State the maximum number of VIDs that can be allocated to each FID.	M	8.11.7	____ VIDs
(23e.6)	Does the implementation support configuration of VLAN Learning Constraints via management?	O	5.2, 8.11.7, 12.10.3	Yes [ ] No [ ]
(23e.7)	State the number of VLAN Learning Constraints that can be configured in the implementation.	23e.6:M	5.2, 8.11.7, 12.10.3	____ Constraints
(23e.8)	Does the implementation support configuration of VID to FID allocations via management?	O	5.2, 8.11.7.1, 12.10.3	Yes [ ] No [ ]

Item	Feature	Status	References	Support	
(23e.9)	Does the implementation take account of the allocation of VIDs to FIDs when making forwarding decisions relative to group MAC Addresses?	O	8.11.8	Yes [ ]	No [ ]
(23f)	On Ports that support untagged and priority-tagged frames, does the implementation support:		5.1, 8.4.4, 8.11.9, 12.10		
(23f.1)	— A PVID value?	M		Yes [ ]	N/A [ ]
(23f.2)	— The ability to configure one VLAN whose Untagged set includes that Port?	M		Yes [ ]	N/A [ ]
(23f.3)	— Configuration of the PVID value via management operations?	M		Yes [ ]	N/A [ ]
(23f.4)	— Configuration of Static Filtering Entries via management operations?	M		Yes [ ]	N/A [ ]
(23f.5)	— The ability to configure more than one VLAN whose Untagged set includes that Port?	O		Yes [ ] N/A [ ]	No [ ]
(23g)*	Does the implementation support the ability to enable and disable Ingress Filtering?	O	5.2, 8.4.5		
(23h)	Does the implementation support VLAN management operations?	19:O	5.2, 12.10.2, 12.10.3	Yes [ ]	No [ ]
(23i)	Is the minimum tagged frame length that can be transmitted on IEEE Std 802.3 Ports less than 68 (but 64 or more) octets?	1a.1:O	7.2	Yes [ ] N/A [ ]	No [ ]
(23j)*	When transmitting untagged frames and the canonical_format_indicator parameter indicates that the mac_service_data_unit may contain embedded MAC Addresses in a format inappropriate to the destination MAC method, which of the following procedures is adopted by the Bridge:		7.1, 7.1.2.2		
(23j.1)	Convert any embedded MAC Addresses in the mac_service_data_unit to the format appropriate to the destination MAC method.	O.7		Yes [ ]	No [ ]
(23j.2)	Discard the frame without transmission on that Port.	O.7		Yes [ ]	No [ ]
(23k)	Does the Bridge perform frame translations, where necessary, in accordance with the procedures described in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390?	TB:M	7.1, 7.1.2.2	Yes [ ] N/A [ ]	No [ ]
(23l)*	Does the implementation support Port-and-Protocol-based classification of frames on any or all Ports?	O	8.6	Yes [ ]	No [ ]
(23m)*	Does the implementation support a Protocol Group Database?	23l:M	8.6.4	Yes [ ]	N/A [ ]

Predicates:
TB = True if the Bridge supports translational Bridging; i.e., the Bridge supports 802.3/Ethernet MAC methods on one or more Ports and Token Ring/FDDI MAC methods on one or more Ports.

A.13 VLAN support

Change the wording in the table below as follows:

Item	Feature	Status	References	Support	
	Ingress rules				
(24a)	Can the PVID or the VID in any member of the VID Set for any Port be assigned the value of the null VLAN ID?	X	8.4.4, Table 9-2	No [ ]	
(24b)	Are frames discarded (or not discarded) in accordance with the settings of the Acceptable Frame Types parameters?	M	8.6	Yes [ ]	
(24c)	Are all frames received classified as belonging to exactly one VLAN, as defined in the ingress rules?	M	8.6	Yes [ ]	
(24d)	Is Ingress Filtering performed in accordance with the value of the Enable Ingress Filtering parameter?	M	8.6	Yes [ ]	
(24e)	Are all frames that are not discarded as a result of the application of the ingress rules submitted to the Forwarding Process and to the Learning Process?	M	8.6	Yes [ ]	
(24f)	State which Ports support Port-and-Protocol-based classification rules.	23l:M	8.6	Ports:	_____
(24f.1)	For each Port that supports Port-and-Protocol-based classification rules, is a VID Set supported?	23l:M	8.4.4	Port: Yes [ ]	_____ N/A [ ]
(24f.2)	For each Port that supports Port-and-Protocol-based classification rules, state how many entries are supported in the VID Set.	23l:M	8.4.4	Port: _____	_____ Entries
(24f.3)	For each Port that supports Port-and-Protocol-based classification rules, is the VID Set configurable via management?	23l:M	12.10.1.2	Port: Yes [ ]	_____ N/A [ ]

Item	Feature	Status	References	Support	
(24g.1)	State how many entries are supported in the Protocol Group Database.	23m:M	8.6.4	Entries	_____
(24g.2)	Is the Protocol Group Database configurable via management?	23m:O	12.10.2.1	Yes [ ]	No [ ]
(24g.3)	Does the Protocol Group Database support entries of format Ethernet?	23m:O	8.6.4	Yes [ ]	No [ ]
(24g.4)	Does the Protocol Group Database support entries of format RFC_1042?	23m:O	8.6.4	Yes [ ]	No [ ]
(24g.5)	Does the Protocol Group Database support entries of format SNAP_8021H?	23m:O	8.6.4	Yes [ ]	No [ ]
(24g.6)	Does the Protocol Group Database support entries of format SNAP_Other?	23m:O	8.6.4	Yes [ ]	No [ ]
(24g.7)	Does the Protocol Group Database support entries of format LLC_Other?	23m:O	8.6.4	Yes [ ]	No [ ]
(24g.8)	Does theProtocol Group Database support entries of at least one of the following formats: Ethernet, RFC_1042, SNAP_8021H, SNAP_Other, LLC_Other?	23m: M	8.6.4	Yes [ ]	
	Egress rules				
(25a)	Are frames discarded if the transmission Port is not present in the Member set for the frame's VID?	M	8.8, 8.11.9	Yes [ ]	
(25b)	Are frames discarded if the value of the include_tag parameter is False, and the Bridge does not support the ability to translate embedded MAC Address information from the format indicated by the canonical_format_indicator parameter to the format appropriate to the media type on which the data request will be carried?	23j.2:M	8.8	Yes [ ]	N/A [ ]
(25c)	Are frames transmitted as VLAN-tagged frames or as untagged frames in accordance with the value of the untagged set for the frame's VID?	M	8.8	Yes [ ]	
	Filtering Database				
(26a)	Does the implementation support Static VLAN Registration Entries as defined in 8.11.2?	M	8.11.2	Yes [ ]	
(26b)	Does the implementation support the creation of a separate Static VLAN Registration Entry with a distinct Port Map for each VLAN from which frames are received by the Forwarding Process?	O	8.11.2	Yes [ ]	No [ ]
(26c)	Does the implementation support Dynamic VLAN Registration Entries as defined in 8.11.5?	M	8.11.5	Yes [ ]	
(26d)	Does the implementation support the creation of a separate Dynamic VLAN Registration Entry with a distinct Port Map for each VLAN from which frames are received by the Forwarding Process?	O	8.11.5	Yes [ ]	No [ ]

Item	Feature	Status	References	Support
(26e)	Does the implementation allocate VIDs to FIDs in accordance with the specification in 8.11.7?	M	8.11.7, 8.11.7.2	Yes [ ]
(26f)	Does the implementation correctly detect Learning Constraint violations?	M	8.11.7.3	Yes [ ]
(26g)	Is determination of the Member set and the untagged set for a given VLAN achieved as defined in 8.11.9?	M	8.11.9	Yes [ ]
	Tagged frames			
(27a)	Do VLAN-tagged frames transmitted by the Bridge conform to the format defined in Clause 9 for the MAC type on which they are transmitted?	M	9	Yes [ ]
(27b)	Are all BPDUs transmitted untagged?	M	8.14.7	Yes [ ]
	VLAN use of GMRP. If item (2b) is not supported, mark N/A and continue at item (29a).			N/A [ ]
(28a)	Does the implementation of GMRP recognize the use of VLAN Contexts for the transmission and reception of GMRP PDUs?	2b:M	10, 10.1, 10.2 10.3	Yes [ ]
(28b)	Does the implementation of GMRP support the creation of distinct GMRP Participants for each VLAN context?	2b:M	10.2	Yes [ ]
(28c)	Does the implementation support the identification of VLAN contexts in transmitted GMRP PDUs by means of VLAN-tagged or untagged frames, in accordance with the member set and untagged set for the VLAN Context concerned?	2b:M	10.3	Yes [ ]
(28d)	Are GMRP PDUs transmitted only on Ports that are part of the active topology for the VLAN Context concerned?	2b:M	10.1	Yes [ ]
	VLAN Topology Management			
(29a)	Does the implementation support the creation, updating and removal of Dynamic VLAN Registration Entries in the Filtering Database under the control of GVRP?	M	11	Yes [ ]
(29b)	Does the Permanent Database contain an entry for the Default VID that defines Registration Fixed on all Ports?	O	11.2.1.3	Yes [ ] No [ ]
(29c)	Is the GVRP Application address used as the destination MAC Address in all GVRP protocol exchanges?	M	11, Table 11-1	Yes [ ]
(29d)	Are GVRP protocol exchanges achieved by means of LLC Type 1 procedures, using the LLC address for Spanning Tree protocol?	M	11, {D}12.4, {D}12.5, {D}Table 7-8	Yes [ ]
(29e)	Are GVRP protocol exchanges achieved using the GARP PDU formats, and the definition of the attribute type and value encodings defined for GVRP?	M	11, 11.2.3.1, {D}12.4, {D}12.5, {D}12.11	Yes [ ]

Item	Feature	Status	References	Support
(29f)	Does the implementation support the operation of the Applicant, Registrar, and Leave All state machines?	M	{D}12.8	Yes [ ]
(29g)	Does the Bridge propagate registration GVRP information only on Ports that are part of the active topology of the base Spanning Tree Context?	M	11, {D}12.3.3, {D}12.3.4	Yes [ ]
(29h)	Does the GVRP application operate as defined in Clause 11?	M	11	Yes [ ]

Annex D

(informative)

Background to VLANs

Change the wording in D.1 as follows:

D.1 Basic VLAN concepts

Figure D-1 shows a simple example of a Port-based VLAN. For untagged traffic, VLAN membership~~ for a Port-based VLAN is determined by the PVID assigned to the receiving Port.~~ is determined by one of several methods:

a) For a port on an access link that is using Port-based VLAN classification, membership is determined by the PVID assigned to the receiving Port.

b) For a port on an access link that is using Port-and-Protocol-based VLAN classification, membership is determined by the VID that is mapped from the link-layer protocol carried in the frame and is assigned to the receiving Port.

NOTE—Other criteria for VLAN membership, such as ~~protocol type or~~ MAC Address, could be used, but these are beyond the scope of this discussion.

For this configuration there needs to be a way to convey the VLAN information between the two bridges. This is done by adding a VLAN tag to every frame that is sent between the two bridges; such frames are known as VLAN-tagged frames. This connection between the two bridges is commonly known as a *Trunk Link*.

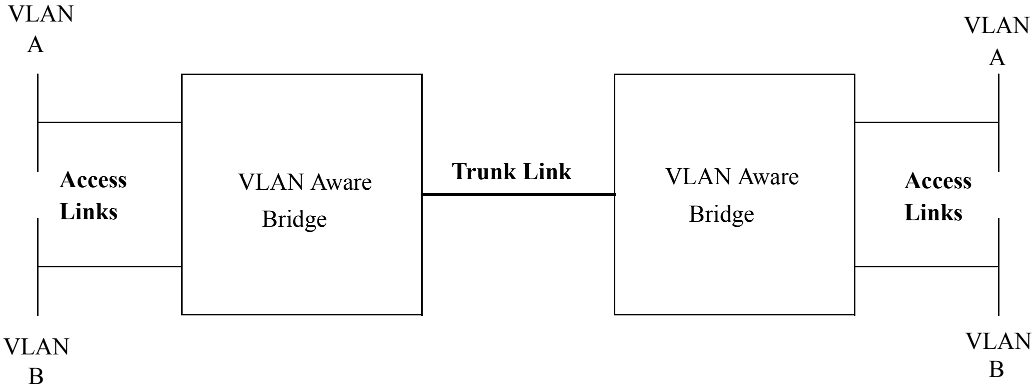

Figure D-1—Port-based VLANs

D.1.1 Trunk Links

A Trunk Link is a LAN segment used for multiplexing VLANs between VLAN Bridges. All the devices that connect to a Trunk Link must be *VLAN-aware*. VLAN-aware devices are devices that are able to understand VLAN membership and VLAN frame formats. Conversely, *VLAN-unaware* devices do not have an understanding of VLAN membership and VLAN frame formats. All frames, including end station frames, on a Trunk Link are VLAN-tagged, i.e., they carry a tag header that contains a non-null VLAN ID. Consequently, there are no VLAN-unaware end stations on an a Trunk Link. The Trunk Link in Figure D-1 is a point-to-point LAN segment; there are therefore exactly two VLAN-aware Bridges attached to this Trunk. A Trunk Link could also be a shared medium LAN segment that has many VLAN-aware Bridges attached to it.

D.2 Relationship of other VLAN styles to ~~with~~ the Port-based VLAN model

Change the items in D.2.2 and reletter appropriately.

D.2.2 Use of other VLAN styles

~~The current 802.1Q standard defines~~This standard defines the following classification rules:

a) A Port-based tagging rule, whereby all untagged and priority-tagged frames received by a Port are classified as belonging to the VLAN whose VID (the PVID) is associated with that Port.

b) A Port-and-Protocol-based rule, whereby all untagged frames received by a Port are classified, by inspecting the data-link layer framing and upper-layer protocol type, as belonging to one of the VLANs whose VID is associated with that Port.

Th~~is~~e Port-based style of operation should be viewed as the base level of a possible hierarchy of VLAN styles, each one able to classify untagged frames according to particular ingress rules. Examples of some other ~~such~~ ingress rules might include the following:

c) MAC Address-based classification; e.g., associating a set of MAC Addresses with a given VLAN ID in order to define the membership of the VLAN;

~~c)~~ ~~Protocol-based classification; e.g., allocating VLAN membership on the basis of the higher-layer protocol information carried in the frame;~~

d) Subnet-based classification; e.g., allocating VLAN membership on the basis of IP subnet addressing characteristics of frames.

For a given implementation, such rules might form a natural hierarchy; e.g., ~~using the above set,~~ IP Subnet-based tagging might take the highest priority. If the packet ~~was~~ is not an IP packet, then tagging is based on the protocol being used: IPX or LAT. If some other protocol is in use (not IP, IPX, or LAT), then the classification is based on MAC Addresses. If the addresses in the frame do not match the address-based classifications that are configured, then the Port-based rule is applied.

The result of such a hierarchy is that a given ingress rule defines the default that is applied if the higher priority rule fails to classify the frame, with the Port-based rule forming the lowest level, “catch-all” default.

NOTE—Clearly, if a given rule in the hierarchy is able to classify all possible frames, then all rules below that point in the hierarchy are effectively disabled.

The addition of further ingress rules in 802.1Q Bridges could be achieved

e) As proprietary extensions to the existing specification;

f) As future standardized extensions.

Given that the starting point for this standard is that all links are Hybrid Links, there is no need for such additional classification and tagging functionality to exist within the Bridges themselves; it would, for example, be possible to develop “tagging engines” that are capable of implementing more complex classifications than Port-based classification, and which are placed between the 802.1Q Bridge Port and an

Access Link. Such a device would provide a richer functionality in terms of VLAN classification style, while remaining compatible with Port-based VLAN operation.

Add new subclause D.3 as follows:

D.3 Example configurations for the Port-and-protocol-based VLAN model

One example scenario for the configuration of Port-and-Protocol based VLAN classification was shown in Figure 8-8. That example showed the use of Protocol Groups to tie together Protocol Templates for common treatment on several Ports of a bridge. Some further examples are presented here.

D.3.1 Example of per-protocol control

In the example in Figure D-3, Protocol Templates are used individually without any groupings. This method allows maximum granularity of control by mapping individual IEEE 802.3 Type values to their own VIDs. Here, each Protocol Template is assigned its own Protocol Group Identifier; each port's VID Set is set up to group protocols together into the same VLAN, as required. This allows a different grouping of protocols on different Ports of the bridge: here, Port 1 classifies Type-encapsulated 802.3-format IP and RFC 1042-format IP packets to the same VLAN (VID 234) and it classifies Type-encapsulated 802.3-format ARP and RFC 1042-format ARP packets to a different VLAN (VID 567). At the same time, Port 2 treats Type-encapsulated 802.3-format ARP and IP packets the same and classifies them to the same VLAN (VID 123) whilst treating RFC 1042 IP and ARP packets the same and using the default treatment to classify them to a different VLAN (VID 789, the PVID for this Port). The tradeoff here is that there is more per-Port configuration required to achieve this level of control and it can use up a larger number of VID Set entries, which may be a scarce resource in a Bridge, to achieve the desired result.

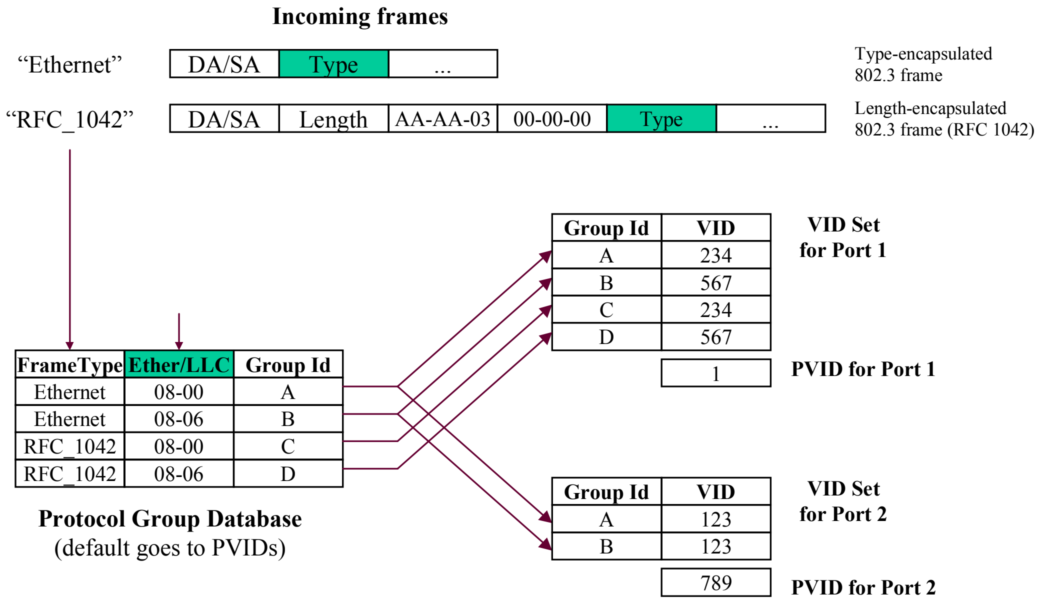

FrameType	Ether/LLC	Group Id
Ethernet	08-00	A
Ethernet	08-06	B
RFC_1042	08-00	C
RFC_1042	08-06	D

Group Id	VID
A	234
B	567
C	234
D	567

Group Id	VID
A	123
B	123

Figure D-3—Example of control down to the individual protocol level

D.3.2 Protocol-based frame filtering

Port and Protocol-based VLAN classification mechanism can also be used to set up protocol-based filters for selectively discarding frames based upon their upper-layer protocol: either of the approaches described above, using Protocol Groups or else using individual Protocol Templates, can be used to classify frames to a particular VLAN, created specially for the purpose. If the VLAN is configured to have no members in its Member Set, then any frames directed there will effectively be discarded. Note that this approach does use up a VID Set entry on each Port for which the filter is desired; this may favor using the Protocol Group approach.

Additionally, this mechanism can be extended for security analysis. Instead of totally discarding the frames that match protocol-based filters, a "mirroring" Port can be added to the Member Set for the filtering VLAN: this has the effect of trapping all frames which fail some sort of protocol-based access control test for future analysis by an analysis probe located on the mirroring Port.

Annex E

(informative)

Interoperability considerations

E.1 Requirements for interoperability

Change wording in E.1.2 as follows:

E.1.2 Configuration requirements for VLAN-tagging

802.1Q Bridges classify incoming untagged frames by applying either a Port-based tagging rule on ingress that uses the PVID for the receiving Port as the VLAN classification for such frames or a Port-and-Protocol-based rule that uses the frame's upper-layer protocol to select one of a port's VIDs. Maintaining consistent connectivity between any pair of end stations that are on the same VLAN, and where one or both of those end stations is VLAN-unaware, requires that

a) All VLAN-aware Bridge Ports that are connected to the same LAN segment apply a consistent set of ingress rules (8.6);

b) All VLAN-aware Bridge Ports that are connected to the same *legacy region* of a Bridged LAN apply a consistent set of ingress rules;

c) All VLAN-aware Bridge Ports that serve LAN segments to which members of the same VLAN are (or can be) attached apply a consistent set of ingress rules.

A legacy region of a Bridged LAN consists of any set of LAN segments that are physically interconnected via VLAN-unaware, ISO/IEC 15802-3 Bridges. A legacy region has the property that, by appropriate configuration of the Spanning Tree, a Spanning Tree path could be created between any pair of LAN segments in the region such that the path would pass only through VLAN-unaware Bridges.

NOTE—In case b), Spanning Tree reconfiguration within the legacy region can change the logical connectivity between the VLAN Ports and the LAN segments that they (directly or indirectly) serve. Hence, a Spanning Tree reconfiguration could result in any end stations connected to the legacy region being serviced via any of the VLAN-aware Ports. In effect, such a reconfiguration reduces case b) to case a). Figure E-2 and Figure E-3 give examples of this type of configuration. In Figure E-2, the legacy region consists of all three LAN segments and both ISO/IEC 15802-3 Bridges. In Figure E-3, the legacy region consists of the ISO/IEC 15802-3 Bridge and both LAN segments to which it is attached. An example of case c) is where an end station attached to a leaf LAN segment is in the same VLAN as a server that is attached to a distinct LAN segment, i.e., all possible Spanning Tree paths between the two stations pass through a VLAN-aware region of the Bridged LAN.

The essence of what these rules express is that if a given untagged frame belongs on a given VLAN, then the classification and tagging behavior of any VLAN-aware Bridges that are required to tag that frame needs to be the same, regardless of the logical connectivity that is created by the Spanning Tree configuration of the Bridged LAN. Examples of the consequences of failure to apply these rules appear in E.3 and E.6.

E.2 Homogenous 802.1Q Bridged LANs

Change wording in E.2.2 as follows:

E.2.2 Consistent view of the “untagged VLAN(s)” on a given LAN segment

In the Port-based VLAN model defined in this standard, the PVID for a Port provides the VLAN classification for all frames on the attached LAN segment that are received in untagged form. Any LAN segment with more than one 802.1Q Bridge attached has such an “untagged VLAN” for each Bridge. No explicit requirement that these be consistent for all Bridges on the same LAN segment, nor mechanism to assure such, has been included in this standard.

Similarly, in the Port-and-Protocol-based VLAN model defined in this standard, one of the members of the VID Set for a Port provides the VLAN classification for all frames on the attached LAN segment that are received in untagged form with a particular upper-layer protocol. Any LAN segment with more than one such 802.1Q Bridge attached has such a set of “untagged VLANs” for each Bridge. No explicit requirement that these be consistent for all Bridges on the same LAN segment, nor mechanism to assure such, has been included in this standard.

Consider the case of a LAN segment to which are attached three VLAN-aware Bridges, each of which is using Port-based classification and is capable of transmission of untagged frames onto the LAN segment. An untagged frame placed on that segment by any one of the Bridges will be associated by each of the other two Bridges with their own configured PVID for their receiving port on that LAN. The 802.1Q VLAN model requires that each frame have a unique VLAN association, and that association is represented by a single, global VID value. Therefore, it follows that all 802.1Q Bridges on that LAN segment must make use of the same classification rules (in this case, the same PVID) for their ports connected to that LAN segment.

It has been suggested that in the special case of a direct point-to-point connection between two 802.1Q Bridges or other VLAN-aware devices, other rules might apply. No mechanism for identifying such links has yet been suggested.

This creates a configuration challenge for installers of Bridges that conform to this standard. Initial management configuration of the Bridges (the setting of PVIDs, VID Sets and Protocol Group Databases) must be made consistent among the Bridges, in a manner that takes into account the actual physical topology. Changes to the physical topology may require specific changes to the configuration of all affected switches. These requirements effectively disallow a plug-and-play installation as supported by ISO/IEC 15802-3 Bridged LANs, unless all Bridges are left with their default Port-based classification rules and with each ~~PVID configuration of~~ PVID = 1.

E.5 Heterogeneous Bridged LANs: intermixing 802.1Q Bridges with ISO/IEC 15802-3 Bridges

Change wording in E.5 as follows:

The specification in this standard for the use of GMRP in VLANs (11.2) makes use of VLAN-tagged frames to signal the GIP Context that applies to the registration information carried in GMRP PDUs. Devices that implement GMRP as specified in ISO/IEC 15802-3 will regard such frames as badly formed GMRP frames, and will therefore discard them on receipt. Using an ISO/IEC 15802-3 Bridge to interconnect two or more LAN regions containing 802.1Q devices that implement GMRP will therefore prevent GMRP information propagation between the 802.1Q regions, with attendant effects upon the forwarding behavior of both the

ISO/IEC 15802-3 and 802.1Q Bridges in the LAN. This configuration can be made to work if the ISO/IEC 15802-3 Bridge is statically configured with the following:

a) An All Groups entry in the Filtering Database, specifying Registration Fixed on all Ports, and

b) The GMRP Protocol Administrative Control parameters set to disable GMRP on all Ports.

As the Bridge no longer supports the GMRP application, it will forward GMRP PDUs on all Ports that are in Forwarding. The effect of this is to configure the ISO/IEC 15802-3 Bridge to behave in the same manner as an ISO/IEC 10038 Bridge.

Placing ISO/IEC 15802-3 Bridges around the periphery of an 802.1Q-based Bridged LAN works correctly, as long as, for a given ISO/IEC 15802-3 Bridge, the 802.1Q Bridges connected to the same segment(s) are configured to untag any VLANs that are relevant to the GMRP operation of the ISO/IEC 15802-3 Bridge. The ISO/IEC 15802-3 Bridge generates untagged GMRP frames, which the 802.1Q Bridges classify according to the value of the PVID for the reception Port; in a simple configuration of the 802.1Q Bridges, the Ports that connect to the ISO/IEC 15802-3 Bridge are configured for the PVID VLAN to be untagged on egress.

NOTE 1—There may be situations where more complex configurations are required, in which VLANs other than the PVID are configured untagged in order to maintain the correct ISO/IEC 15802-3 Bridge filtering behavior.

NOTE 2—For bridges that make VLAN assignments on untagged frames according to Port-and-Protocol-based classification rules, special care is necessary in configuration: since GMRP does not carry information about the particular protocol for which the Group membership is intended, it must be taken to apply to all protocols for which untagged traffic is carried on a particular link. Therefore, the VLAN indicated by a port's PVID is used to carry the GMRP frames associated with all of the VLANs which are members of the VID Set of that Port; in addition, the PVID must be configured to egress untagged on any other bridge port where any of the set of VIDs also egresses untagged and requires GMRP operation beyond that egress port.

The effect of this type of configuration is that all registrations propagated by a given ISO/IEC 15802-3 Bridge on a given (Port-based or Port-and-Protocol-based) VLAN are seen by all other ISO/IEC 15802-3 Bridges served by 802.1Q Bridges for which that VLAN is configured for untagged egress. The filtering behavior of the ISO/IEC 15802-3 Bridges is therefore governed only by the behavior of other devices (both ISO/IEC 15802-3 and 802.1Q) that are attached to the same VLAN.

Change wording in E.6 as follows:

E.6 Intermixing Port-based classification and Port-and-Protocol-based classification or future enhancements in 802.1Q ~~Version 1.0 Bridges with future 802.1Q Bridges~~

The discussion above on intermixing Q Bridges with D Bridges has a direct analogue in the mixing of bridges implementing only Port-based and Port-and-Protocol-based classification of frames in 802.1Q networks. ~~plan to provide a simple VLAN standard (Q version 1.0) initially, and later to provide extensions to the standard (Q version 2.0 on) which~~This and potential subsequent editions of 802.1Q extend~~s~~ the VLAN classification capabilities to support more sophisticated ingress rules for frame classification. ~~Some of the topology restrictions will probably be similar to the "Q intermixed with D" cases.~~

In VLAN configurations that use both Port-based and Port-and-Protocol-based VLAN classification, a Bridge that supports only Port-based VLAN classification will merge VLANs that would otherwise be classified separately by a Bridge that supports Port-and-Protocol-based VLAN classification. To get around this problem, it may be possible to dedicate specific Ports to specific protocols in Bridges that support only

Port-based VLAN classification, as in the example shown in Figure E-4. However, such solutions may not be possible where there are multiprotocol end stations in the network.

E.6.1 Example: Intermixing ~~Layer 3~~ Protocol-based Ingress Rules

Consider the case where Bridges implement a configuration mechanism to select between Port-based classification rules (a “Q-Port Bridge”) and Port-and-Protocol-based rules (a “Q-Port/Protocol~~v2”~~ Bridge”) ~~allowing for classification of frames by Protocol~~. This would allow, for example, for support for IP and IPX as distinct VLANs. The topology shown in Figure E-4~~following diagram~~ might apply when a Q-Port/Protocol~~v2~~ Bridge is added to a ~~version 1.0~~ topology, otherwise using only Q-Port bridges, to allow users of two protocols to participate in two separate VLANs.

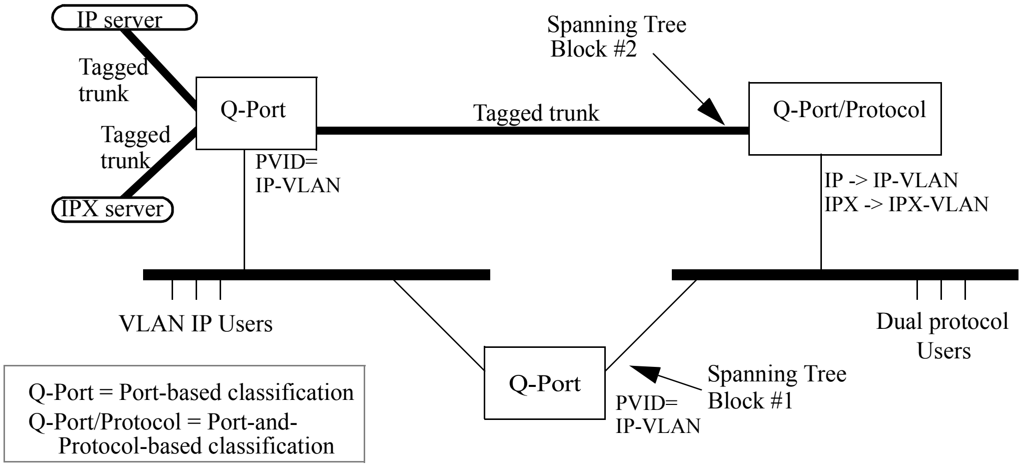

Figure E-4—Interoperability between ~~Q versions 1 and 2~~Port-based and Port-and-Protocol-based classification

Consider this network, when STP has blocked at point #1, and not at #2. The upper Q~~v1~~-Port Bridge~~switch~~ operates as expected, and the Q~~v2~~-Port/Protocol Bridge ~~switch~~ provides Port-and-P~~p~~rotocol-based classification for the frames received from the dual protocol users on the right-hand segment. IP-VLAN and IPX- VLAN frames are VLAN-tagged on the trunks to and from the uppermost Bridges and servers. But if a STP reconfiguration should result in a block at point #2, but not at #1, activating traffic through the lower Q~~v1~~-Port Bridge, dual protocol users will have all their traffic treated as part of the IP-VLAN. An immediate consequence is that the uppermost Bridge will no longer provide them access to the “IPX server~~.~~”. It is the fact that the lower Q-Port Bridge must have just *one* of the two VLANs configured for all untagged traffic, regardless of its protocol, that leads to this lack of connectivity.

E.6.2 Differing views of untagged traffic on a given LAN segment

Further challenges arise when one considers the case where several Q Bridges, some implementing Port-based and some implementing Port-and-Protocol-based classification~~of version 1 and some of version 2,~~ all attach to the same LAN segment. Again, the rule that any given frame exists in exactly one VLAN requires that all of these Bridges be configured with ~~the same~~consistent ingress rules. In this case, the Q~~1~~-Port Bridges will provide a least common capability, and this further requires common configuration of the PVID in the Q-Port and Q-Port/Protocol Bridges.

E.6.3 Interoperability with future 802.1D versions~~802.1Q Version 2.0~~ offering multiple spanning trees

~~Several very different~~ An architecture~~s~~ for multiple spanning tree support ~~have been discussed, but none define an architecture~~ is under development as IEEE P802.1s, to be integrated later into this standard. It has not yet been analyzed sufficiently for ~~analysis of~~ interoperability with the VLAN capabilities~~version 1.0~~ defined in this standard. The benefits of ~~such an~~this architecture have been discussed, and among these benefits are a potential relaxation of many, but not all, of the restrictions discussed in this Annex~~subclause~~. In particular, the multiple spanning tree model~~s~~ appears to offer easier integration within both homogenous environments and in networks intermixing D Bridges with multiple spanning tree VLAN-aware devices.