P802.1ad/D2.4
DRAFT Amendment to IEEE Std 802.1Q -1998
as amended by IEEE Std 802.1u -2001
and IEEE Std 802.1v -2001
and IEEE Std 802.1s -2002
September 27, 2004

IEEE P802.1ad/D2.4

Draft Standard for Local and Metropolitan Area Networks—

Virtual Bridged Local Area Networks — Amendment 4: Provider Bridges

Sponsor
LAN/MAN Standards Committee
of the
IEEE Computer Society

Prepared by the Interworking Task Group of IEEE 802.1

Abstract: This amendment enables a Service Provider to use the architecture and protocols of IEEE Std 802.1Q to offer the equivalent of separate LANs, Bridged Local Area Networks, or Virtual Bridged Local Area Networks to a number of users, while requiring no cooperation between the users, and minimal cooperation between each user and the provider.

Keywords: LANs, local area networks, metropolitan area networks, MAC Bridges, Bridged Local Area Networks, virtual LANs, Virtual Bridged Local Area Networks, Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP)

Copyright © 2004 by the Institute of Electrical and Electronics Engineers, Inc.
345 East 47th Street
New York, NY 10017, USA
All rights reserved.

All rights reserved.This document is an unapproved draft of a proposed IEEE Standard. As such, this document is subject to change. USE AT YOUR OWN RISK! Because this is an unapproved draft, this document must not be utilized for any conformance/compliance purposes. Permission is hereby granted for IEEE Standards Committee participants to reproduce this document for purposes of IEEE standardization activities only. Prior to submitting this document to another standards development organization for standardization activities, permission must first be obtained from the Manager, Standards Licensing and Contracts, IEEE Standards Activities Department. Other entities seeking permission to reproduce this document, in whole or in part, must obtain permission from the Manager, Standards Licensing and Contracts, IEEE Standards Activities Department.

IEEE Standards Department
Copyright and Permissions
445 Hoes Lane, P.O. Box 1331
Piscataway, NJ 08855-1331, USA

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 13 31
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.

Authorization to photocopy portions of any individual standard for internal or personal use is granted by the Institute of Electrical and Electronics Engineers, Inc., provided that the appropriate fee is paid to Copyright Clearance Center. To arrange for payment of licensing fee, please contact Copyright Clearance Center, Customer Service, 222 Rosewood Drive, Danvers, MA 01923 USA; (978) 750-8400. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Editors' Foreword

<<Notes>>

<<Throughout this document, all notes such as this one, presented between angle braces, are temporary notes inserted by the Editors for a variety of purposes; these notes and the Editors' Foreword will all be removed prior to publication and are not part of the normative text.>>

<<Comments and participation in 802.1 standards development

Comments on this draft are encouraged. **PLEASE NOTE: All issues related to IEEE standards presentation style, formatting, spelling, etc. are routinely handled between the 802.1 Editor and the IEEE Staff Editors prior to publication, after balloting and the process of achieving agreement on the technical content of the standard is complete.** Readers are urged to devote their valuable time and energy only to comments that materially affect either the technical content of the document or the clarity of that technical content. Comments should not simply state what is wrong, but also what might be done to fix the problem.>>

Full participation in the development of this draft requires individual attendance at IEEE 802 meetings. Information on 802.1 activities, working papers, and email distribution lists etc. can be found on the 802.1 Website:

http://ieee802.org/1/

Use of the email distribution list is not presently restricted to 802.1 members, and the working group has had a policy of considering ballot comments from all who are interested and willing to contribute to the development of the draft. Individuals not attending meetings have helped to identify sources of misunderstanding and ambiguity in past projects. Non-members are advised that the email lists exist primarily to allow the members of the working group to develop standards, and are not a general forum.

Comments on this document may be sent to the 802.1 email exploder, to the editors, or to the Chairs of the 802.1 Working Group and Interworking Task Group.

Mick Seaman
Chair, 802.1 Interworking Task Group
160 Bella Vista Ave
Belvedere
CA 94041
USA
Email:mick_seaman@ieee.org

Tony Jeffree
Chair, 802.1 Working Group
11A Poplar Grove
Sale
Cheshire
M33 3AX
UK
+44 161 973 4278 (Tel)
+44 161 973 6534 (Fax)
Email: tony@jeffree.co.uk

PLEASE NOTE: Comments whose distribution is restricted in any way cannot be considered, and may not be acknowledged.>>
<<A reference to the IEEE's patent policy will be added to this introductory text.>>

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

<<Overview: Draft text and accompanying information

This document currently comprises:

A cover page, identical to the title page.
The editors' introductory notes to each draft, briefly summarizing the progress and focus of each successive draft.
The title page for this amendment including an Abstract and Keywords. This title page will be retained for the period that the amendment is published as a separate document.
The amendment proper, documented in the usual form for amendments to 802 standards; i.e., as an explicit set of editing instructions that, if correctly applied to the text of 802.1Q, will create a corrected document.
An Annex Z comprising the editors' discussion of issues. This annex will be deleted from the document prior to sponsor ballot.
Editors' notes throughout the document, including requests for comment on specific issues and pointing deficiencies in the current draft.
IEEE boilerplate text.

The records of participants in the development of the standard, the introduction to 802 standards, and the introduction to this revision of the standard are not included, and will be added at an appropriate time.

During the early stages of draft development, 802.1 editors have a responsibility to attempt to craft technically coherent drafts from the resolutions of ballot comments and the other discussions that take place in the working group meetings. Preparation of drafts often exposes inconsistencies in editors instructions or exposes the need to make choices between approaches that were not fully apparent in the meeting. Choices and requests by the editors' for contributions on specific issues will be found in the editors' introductory notes to the current draft, at appropriate points in the draft, and in Annex Z. Significant discussion of more difficult topics will be found in the last of these.

The ballot comments received on each draft, and the editors' proposed and final disposition of comments, are part of the audit trail of the development of the standard and are available, along with all the revisions of the draft on the 802.1 website (for address see above).
>>

<<Editor's Introduction to the current draft.

P802.1ad/D2.4 is work in progress. Not all proposed dispositions from the D2.0 ballot have been included as yet. However the majority of the work required to properly include Drop Precedence has been completed, pending review by the group. Please review the revised Annex G in particular. Though strictly informative, it spells out a rationale for the technical changes that has to be (or be changed to be) more or less acceptable if we are to proceed with any degree of ceratinty that we know what we are doing. Other significant text is to be found near the end of Clause 6.

Clause 15 and the naming used to identify the functional roles of Ports at or near the edge of the network has also been revised significantly, in line with but hopefully an improvement on that tentatively proposed in the ballot disposition.

The intent is that comments received on the above will be taken into account, if possible, prior to issuing a ballot on a revised draft (to include other proposed dispositions on the D2.0 ballot) shortly after the October 2004 interim, closing in time to address comments in the November plenary meeting.

>>

<<Editor's Introduction to prior drafts (excerpts of continuing relevance).

P802.1ad/D2.0 has been/will be distributed for task group ballot as agreed at the November 2003 meeting. It remains work in progress, and known not to be completely consistent. The purpose of the task group ballot is to checkpoint the development so far and to provide a structured way of gathering further input from all participants. Those responding to the task group ballot will assist the process of further development by focusing on the major concepts, issues, and terminology. Comments on detailed wording, unless questions

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

about intent and opinions as to direction, are largely a waste of time. Comments to the effect that the draft is incomplete may be treated as “Abstain due to lack of expertise”.

The conformance clause has been largely reorganized as a result of the November 2003 and prior discussions, and clause 15 has also been updated to correspond to our latest thinking. To minimize uncertainty as to what is new, and what is being superseded (there never being enough time for the editor to finely polish each draft during task group balloting stage). The draft is probably best read for the first time in the following order

—Clause 5.9 (Provider Bridge Conformance)
—Clause 5 in its entirety (Conformance)
—Clause 15 (Support of the MAC Service by Provider Bridged Networks)
—Clause 8 (Principles of Bridge operation)
—The whole document in serial clause order

The editor has not had the time, or the inspiration, to resolve the difficulties around drop precedence (implicit or explicit) to his satisfaction, so nothing on this subject has yet been incorporated. Further comemnts are requested.

P802.1ad/D1.4, is work in progress. It does not complete the resolution of P802.1ad/D1 ballot comments. However, given the significant task group discussion and the many changes that have already been made to progress the resolution of those comments, it is time that we held another task group ballot. That ballot should test the extent to which prior comments have been successfully addressed, establish which comments are still outstanding, and allow the task group to take stock of the progress to date.

P802.1ad/D1.4 has been distributed prior to the November meeting to facilitate discussion at that meeting of which issues and anomalies should be fixed prior to task group ballot, largely to save everyone’s time on that ballot. The editor does not believe that the current draft is complete, though it should be far more internally consistent than prior drafts, nor that we should strive for completion prior to taking stock by way of a task group ballot. We are, after all, at task group ballot level, not working group ballot level on this draft. The suggested objective is to keep the remaining changes prior to ballot few enough (or on a best efforts basis) to commence the ballot by December 1st, and allow at least a week for collation of comments prior to the following interim (yet to be announced).

Clause 15 has been added, and builds further on our architectural decision (see note to P802.1ad/D3 below) as part of dealing with comments. It is clearly possible to expand significantly on the material in this clause, focusing on the transport characteristics of the customer interface. A note refers to the Bibliography for other aspects of the interface, and this is an appropriate place to reference MEF and ITU work in this area. How much we want to further describe the customer interface is a question for discussion, as is overlap with Clause 16.

P802.1ad/D1.3 is work in progress. It continues, but does not complete the task of providing text that may serve to assist resolution of P802.1ad/D1 ballot comments.

Clause 8 has been significantly revised to take into account our decision on the architecture (per Steve Haddock’s presentation and the working group resolution). The MAC dependent aspects of bridge functionality have been excised from clause 8 and placed in clause 6 as required to make the architecture work correctly. In doing the work it was apparent that changes needed to be applied after essential maintenance to the clause, so it has been presented (almost) in its entirety. The editor has also made significant efforts to replace out of date material with applicable text from the recently passed sponsor ballot draft of P802.1D.

Clause 9 has also been significantly revised as required to provide a secure base for this amendment.

A number of issues have been discovered with Clauses 8 and 9, both with respect to maintenance where defects in the standard 802.1Q text have been discovered, and with respect to technical choices on P802.1ad. The editor will raise these in the upcoming interim meeting.

P802.1ad/D1.1, is work in progress. It includes some, but by no means all of the proposed resolutions to the ballot on P802.1ad/D1. The most interesting new material is to be found in Clause 6.

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

P802.1ad/D1 was prepared by the P802.1ad editor, Mick Seaman, for task group ballot, comments received will be considered at the June 2003 interim meeting of IEEE 802.1.

Clause 16 of this draft is a detailed description of the ‘Principles of Provider Network Operation’. While this is based on discussion at the January and March 2003 meetings, this text has not been reviewed before and ballot comments are particularly requested. Other clauses have not changed significantly since the first draft, P802.1D0, and may well be found to be out of alignment with Clause 16. The intention is to resolve clause 16 issues and options, and then incorporate the necessary supporting changes in the other clauses.
>>

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

P802.1ad/D2.4
DRAFT Amendment to IEEE Std 802.1Q -1998
as amended by IEEE Std 802.1u -2001
and IEEE Std 802.1v -2001
and IEEE Std 802.1s -2002
September 27, 2004

IEEE P802.1ad/D2.4

Draft Standard for
Local and Metropolitan Area Networks—

Virtual Bridged Local Area Networks — Amendment 4: Provider Bridges

Sponsor
LAN/MAN Standards Committee
of the
IEEE Computer Society

Prepared by the Interworking Task Group of IEEE 802.1

Abstract: This amendment enables a Service Provider to use the architecture and protocols of IEEE Std 802.1Q to offer the equivalent of separate LANs, Bridged Local Area Networks, or Virtual Bridged Local Area Networks to a number of users, while requiring no cooperation between the users, and minimal cooperation between each user and the provider.

Keywords: LANs, local area networks, metropolitan area networks, MAC Bridges, Bridged Local Area Networks, virtual LANs, Virtual Bridged Local Area Networks, Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP)

Copyright © 2004 by the Institute of Electrical and Electronics Engineers, Inc.
345 East 47th Street
New York, NY 10017, USA
All rights reserved.

All rights reserved.This document is an unapproved draft of a proposed IEEE Standard. As such, this document is subject to change. USE AT YOUR OWN RISK! Because this is an unapproved draft, this document must not be utilized for any conformance/compliance purposes. Permission is hereby granted for IEEE Standards Committee participants to reproduce this document for purposes of IEEE standardization activities only. Prior to submitting this document to another standards development organization for standardization activities, permission must first be obtained from the Manager, Standards Licensing and Contracts, IEEE Standards Activities Department. Other entities seeking permission to reproduce this document, in whole or in part, must obtain permission from the Manager, Standards Licensing and Contracts, IEEE Standards Activities Department.

IEEE Standards Department
Copyright and Permissions
445 Hoes Lane, P.O. Box 1331
Piscataway, NJ 08855-1331, USA

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.

Authorization to photocopy portions of any individual standard for internal or personal use is granted by the Institute of Electrical and Electronics Engineers, Inc., provided that the appropriate fee is paid to Copyright Clearance Center. To arrange for payment of licensing fee, please contact Copyright Clearance Center, Customer Service, 222 Rosewood Drive, Danvers, MA 01923 USA; (978) 750-8400. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

Contents

Editors’ Foreword .... a

1. Overview 1

1.1 Scope 2

2. References 1

3. Definitions 2

4. Abbreviations 7

5. Conformance 8

5.1 Terminology 8
5.2 Conformant components and equipment 8
5.3 Protocol Implementation Conformance Statement (PICS) 9
5.4 VLAN-aware Bridge Requirements ~~Static conformance requirements .... 9~~
5.5 VLAN-aware Bridge Options 10
5.6 Customer-VLAN aware Bridge Conformance 11
5.7 Service-VLAN aware Bridge Conformance 12
5.8 VLAN Bridge Conformance 13
5.9 Provider Bridge Conformance 13
5.10 MAC-specific bridging methods 14

6. Support of the MAC Service in VLANs 16

6.1 Support of the MAC service 16
6.2 Preservation of the MAC service 16
6.3 Quality of service maintenance 17
6.4 Internal Sublayer Service 18
6.5 Support of the Internal Sublayer Service by specific MAC procedures 20
6.6 Enhanced Internal Sublayer Service 22
6.7 Support of the EISS ~~in VLAN-aware Bridges .... 23~~
6.8 Support of the ISS for attachment to a Provider Bridged Network 27
6.9 Support of the ISS within a system 28
6.10 Support of the ISS by additional technologies 28

7. Principles of network operation 30

8. Principles of bridge operation 32

8.1 Bridge operation 32
8.2 Bridge architecture 34
8.3 Model of operation 35
8.4 Port states and the active topology 39
8.5 Bridge Port Transmit and Receive 40
8.6 The Forwarding Process 41
8.7 The Learning Process 46
8.8 Protocol VLAN classification 47
8.9 The Filtering Database 50
8.10 MST configuration information 63
8.11 Spanning Tree Protocol Entity 64
8.12 GARP Entities 65

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

8.13 Bridge Management Entity 65
8.14 Addressing 65

9. Tagged frame format 74

9.1 Purpose of tagging 74
9.2 Representation and encoding of tag fields 74
9.3 Tag format 75
9.4 Tag Protocol Identifier (TPID) Formats 75
9.5 Tag Protocol Identification 75
9.6 Customer VLAN Tag Control Information 76
9.7 Service VLAN TAG Control Information 77
9.8 Embedded Routing Information Field (E-RIF) 77
9.9 Validation of received frames 79

10. Use of GMRP in VLANs 80

11. VLAN topology management 81

12. VLAN bridge management 82

13. The Multiple Spanning Tree Protocol (MSTP) 83

14. Use of BPDUs by MSTP 84

15. Support of the MAC Service by Provider Bridged Networks 86

15.1 Service transparency 86
15.2 Customer service interfaces 87
15.3 Port-based service interfaces 87
15.4 Customer-tagged service interface 88
15.5 S-tagged service interface 89
15.6 Service instance segregation 90
15.7 Service instance selection and identification 90
15.8 Service priority selection 90
15.9 Service access protection 90

16. Principles of Provider Bridged Network Operation 92

16.1 Provider Network Overview 92
16.2 Service Instance Segregation 93
16.3 Service Instance Selection and Identification 93
16.4 Service Instance Connectivity 94
16.5 Filtering and relaying of group addressed customer frames 95
16.6 Provider Learning of Customer End Station Addresses 96
16.7 Customer Network Topology Changes and Provider Address Learning 96
16.8 Detection of connectivity loops through attached networks 97
16.9 Provider Bridged Network Architecture 97

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Annex A (normative)PICS proforma 100

Annex B (informative)Shared and independent VLAN learning 101

Annex C (informative)Media access method dependent aspects of VLAN support 101

Annex D (informative)Background to VLANs 101

Annex E (informative)Interoperability considerations 101

E.1 Intermixing IEEE Std 802.1Q Version 1.0 bridges with future IEEE P802.1Q bridges 101

Annex G (informative)Priority and drop precedence 104

G.1 Traffic types 104
G.2 Managing latency and throughput 105
G.3 Traffic type to traffic class mapping 105
G.4 Traffic types and user priority values 107
G.5 Supporting drop precedence 108
G.6 Priority code point allocation 108
G.7 Interoperability 108

Annex Z (informative)COMMENTARY 111

Z.1 Frame Classification 111
Z.2 Editor’s notes of interest 111

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Figures

Figure 6-1 Internal organization of the MAC sublayer 16
Figure 8-1 A Bridged Local Area Network... 33
Figure 8-2 VLAN-aware Bridge Architecture... 35
Figure 8-3 Relaying MAC frames 37
Figure 8-4 Observation of network traffic... 37
Figure 8-5 Operation of Spanning Tree protocol... 38
Figure 8-6 Operation of GARP... 39
Figure 8-7 Management Port transmission and reception 39
Figure 8-8 Forwarding Process functions 41
Figure 8-9 Example of operation of port-and-protocol based classification 48
Figure 8-10 Logical points of attachment of the Higher Layer and Relay Entities............................ 68
Figure 8-11 Effect of control information on the forwarding path... 68
Figure 8-12 Per-Port points of attachment... 69
Figure 8-13 Single point of attachment—relay permitted 69
Figure 8-14 Single point of attachment—relay not permitted 70
Figure 8-15 Effect of Port State... 70
Figure 8-16 Effect of authorization 71
Figure 8-17 Ingress/egress control information in the forwarding path 71
Figure 9-1 C-TAG TCI format... 76
Figure 9-1 C-TAG TCI format... 77
Figure 9-1 E-RIF Route Control (RC) field.. 78
Figure 15-1 Internal organization of the MAC sublayer in a Provider Bridged Network.................. 86
Figure 15-2 Port-based interface to a Provider Bridged Network 87
Figure 15-3 Port-based interface to a Provider Bridged Network 88
Figure 15-4 Customer-tagged interface to a Provider Bridged Network.. 88
Figure 15-5 Customer-tagged interface to a Provider Bridged Network.. 88
Figure 15-6 Customer Edge Ports.. 89
Figure 15-7 S-tagged interface to a Provider Bridged Network.. 89
Figure 15-8 S-tagged interface to a Provider Bridged Network.. 89
Figure 16-1 Provider Bridged Network Architecture with CPE Examples.. 98

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Tables

Table 5-1 Support requirements for insertion, removal, and modification of tag headers 10
Table 6-2 FDDI and Token Ring access priorities.. 20
Table 6-3 FDDI and Token Ring user priority regeneration... 21
Table 6-4 Priority encoding.. 25
Table 6-5 Priority Code Point decoding... 25
Table 6-6 Priority regeneration 27
Table 8-1 Customer-VLAN aware Bridge Reserved addresses... 44
Table 8-2 Service-VLAN aware Bridge Reserved addresses 44
Table 8-3 Recommended priority to traffic class mappings 45
Table 8-1 Ageing time parameter value.. 55
Table 8-1 Combining Static and Dynamic Filtering Entries for an individual MAC Address............ 61
Table 8-2 Combining Static Filtering Entry and Group Registration Entry for “All Group Addresses” and “All Unregistered Group Addresses” 61
Table 8-3 Forwarding or Filtering for specific group MAC Addresses.. 62
Table 8-1 Determination of whether a Port is in a VLAN’s member set.. 63
Table 8-4 Standard LLC address assignment... 66
Table 9-2 Reserved VID values 76
Table 9-1 802.1Q Ethernet Type allocations 76
Table G-1 Traffic type to traffic class mapping... 106
Table G-2 Traffic type acronyms.. 107
Table G-3 Defining traffic types... 107
Table G-4 Priority encoding.. 110
Table G-5 Priority Code Point decoding... 110

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

IEEE P802.1ad/D2.4

Draft Standard for Local and Metropolitan Area Networks— Amendment 4 to 802.1Q Virtual Bridged Local Area Networks: Provider Bridges

Editorial Note

This amendment specifies changes to IEEE Std 802.1Q that enable a Service Provider to offer the equivalent of separate LANs, Bridged Local Area Networks, or Virtual Bridged Local Area Networks to a number of separate users. Changes are applied to the base text generated by applying the amendments IEEE Std 802.1s-2002, IEEE Std 802.1u-2001 and IEEE Std 802.1v-2001 to IEEE Std 802.1Q-1998. Text shown in bold italics in this amendment defines the editing instructions necessary to changes to this base text. Three editing instructions are used: ***change***, ***delete***, and ***insert***. ***Change*** is used to make a change to existing material. The editing instruction specifies the location of the change and describes what is being changed. Changes to existing text may be clarified using ~~strikeout~~ markings to indicate removal of old material, and underscore markings to indicate addition of new material). ***Delete*** removes existing material. ***Insert*** adds new material without changing the existing material. Insertions may require renumbering. If so, renumbering instructions are given in the editing instruction. Editorial notes will not be carried over into future editions of IEEE Std. 802.1Q.

<<References to IEEE Std 802.1D assume that the P802.1D/D3 has been approved prior to the successful sponsor ballot of this draft standard. If this is not the case text will be added to the above, to stating that such references are to IEEE Std 802.1D-1998 as amended by 802.1t-2000 and 802.1w-2001.>>

1. Overview

Delete the initial paragraph of clause 1 and insert the following.

IEEE 802 Local Area Networks (LANs) of all types can be connected together with Media Access Control (MAC) Bridges, as specified in IEEE Std. 802.1D. This standard specifies the operation of Bridges that permit the definition, operation, and administration of Virtual LANs (VLANs) within Virtual Bridged Local Area Networks.

This standard further extends the specification of VLAN-aware MAC Bridges to enable a service providing organization to use a common infrastructure of Bridges and LANs to offer the equivalent of separate LANs, Bridged, or Virtual Bridged Local Area Networks to independent customer organizations.

Change clause 1.1 as follows:

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

1.1 Scope

For the purpose of compatible interconnection of information technology equipment using the IEEE 802 MAC Service supported by interconnected IEEE 802 standard LANs using different or identical ~~MAC~~ media access control methods, this standard specifies ~~a general method for~~ the operation of MAC Bridges that support ~~the construction of~~ Virtual ~~V~~LANs (VLANs) ~~(see 3.16)~~. To this end it

a) Positions the support ~~function~~ of VLANs within an architectural description of the MAC Sublayer;

b) Defines the principles of operation of the VLAN-aware Bridge in terms of the support and preservation of the MAC Service, and the maintenance of Quality of Service;

~~Defines enhancements to the Support of the MAC Service, as described and defined in ISO/IEC 15802-3, for the purposes of VLAN Bridging;~~

c) Specifies an Enhanced Internal Sublayer Service provided to the Media Access Method Independent functions that provide frame relay ~~(<3.4>)~~ in ~~the~~ a VLAN-aware Bridge;

d) Establishes the principles and a model of Virtual Bridged Local Area Network operation;

e) Identifies the functions to be performed by VLAN-aware Bridges, and provides an architectural model of the internal operation of a Bridge in terms of Processes and Entities that provide those functions;

~~Specifies the operation of the functions that provide frame relay in the VLAN Bridge;~~

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

~~Defines the structure, encoding, and interpretation of the VLAN control information carried in tagged frames (3.12) in a VLAN;~~

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

~~Specifies the rules that govern the insertion and removal of VLAN control information in frames;~~

h) Specifies the rules that govern the ability to carry user data in either Canonical format and Non-canonical format in VLAN tagged frames ~~using different LAN MAC methods~~;

NOTE—The meanings of the terms *Canonical format* and *Non-canonical format* are discussed in Annex F.

i) Establishes the requirements for, and specifies the means of, automatic configuration of VLAN topology information;

j) Establishes the requirements for VLAN-aware Bridge Management in a Virtual Bridged Local Area Network, identifying managed objects and defining management operations;

~~Defines the management functionality that may be provided in a VLAN Bridge in order to facilitate administrative control over VLAN operation;~~

k) Defines the operation of the Multiple Spanning Tree algorithm and protocol (MSTP);

~~l) Defines the enhancements made to the Rapid Spanning Tree Algorithm and Protocol in order to create MSTP;~~

l) Describes the protocols and procedures necessary to support interoperation between MST and SST Bridges in the same Virtual Bridged ~~LAN~~ Local Area Network;

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

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

n) Defines terms and abbreviations used in the specification and description of the customer service instances provided, and of the MAC Bridges and LANs that support those service instances.

o) Specifies conformance requirements for the provision of customer service instances.

p) Specifies conformance requirements for provider MAC Bridging equipment:

   1) sited on the premises of a user organization to provide one or more instances of service through attachment to a LAN or LANs to that single organization

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

   2) physically secured within provider operated facilities, and connecting to equipment on the premises of one or more user organizations
   3) interconnecting other MAC Bridges within a single provider.

q) Positions the support of customer service instances within a layered architectural description.
r) Defines the provisioning of a service instance in terms of the administrative selection of the user provider interfaces connected by that service instance.
s) Specifies how Customer Bridges or stations can select from a number of multiple service instances accessed using a single LAN or multiple LANs.
t) Defines the principles of network operation in terms of the support and preservation of the MAC Service, and the maintenance of Quality of Service for each service instance, including the segregation of data belonging to different organizations.
u) Describes the components that compose a Provider Bridged Network.
v) Describes the physical and logical topology of a Provider Bridged Network and of the service instances supported by that network.
w) Describes the functions to be performed within the Provider Bridged Network in order to support customer interfaces and the location of those functions to Bridges and Bridge Ports with identified roles within the network.
x) Defines the principles of operation of the bridges that perform each role within the provider bridged network by reference to the principal elements of bridge operation, functions, processes and entities, specified in IEEE Std 802.1D and IEEE Std 802.1Q
y) Establishes the requirements for Bridge Management in the Provider Bridged Network, identifying the managed objects and defining the management operations.
z) Specifies performance requirements and recommends default values and applicable ranges for the operational parameters of a Provider Bridge.

2. References

Insert the following reference at the appropriate point:

<<A placeholder, it seems likely that we will have additional references.>>

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Replace Clause 3 with the following. This clause has been specified in its entirety to facilitate review of the amendment within the context of the final text. Prior references to a number of the definitions in IEEE Std 802.1D have been replaced by the full text of the definition from IEEE Std 802.1D-2003. Changes and additions to other references are specified by means of strike-through and underline notation.

3. Definitions

This Standard makes use of the following terms defined in IEEE Std 802.1D-2004

- Active topology
- Bridge Port
- GARP Participant
- GARP Application
- GIP Context
- Group
- Port

The following terms are specific to this standard or to this standard and IEEE Std 802.1D-2004:

3.1 Bridged Local Area Network: A concatenation of individual IEEE 802 LANs interconnected by MAC Bridges.

NOTE—Unless explicitly specified the use of the word ‘network’ in this Standard refers to a Bridged Local Area Network. The term Bridged Local Area Network is not otherwise abbreviated. The term Local Area Network and the abbreviation LAN are used exclusively to refer to an individual LAN specified by a MAC technology without the inclusion of Bridges. This precise use of terminology within this specification allows a Bridged Local Area Network to be distinguished from an individual LAN that has been bridged to other LANs in the network. In more general usage such precise terminology is not required, as it is a explicit goal of this standard that Bridges are transparent to the users of the MAC Service.

3.2 Boundary Port: A Bridge Port attaching an MST Bridge to a LAN that is not in the same region.

3.3 Common and Internal Spanning Tree (CIST): The single Spanning Tree calculated by STP and RSTP and the logical continuation of that connectivity through MST Bridges and Regions, calculated by MSTP to ensure that all LANs in the Bridged Local Area Network are simply and fully connected.

3.4 Common Spanning Tree (CST): The single Spanning Tree calculated by STP, RSTP, and by MSTP to connect MST Regions.

3.5 Customer Bridge: A MAC Bridge as specified by IEEE Std 802.1-2003 or a VLAN Bridge as specified by this Standard.

3.6 Customer Bridged Local Area Network: A network of Customer Bridges interconnected by IEEE 802 Local Area Networks.

3.7 Customer equipment: The physical embodiment of one or more customer systems.

3.8 Customer system: A system attached to Provider Bridged Network but not intended by the network provider to be under the control of the provider.

3.9 Customer VLANs: <<definition tbs.>>

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

3.10 Customer VIDs: <<definition tbs.>>

3.11 Customer tagged frames: <<definition tbs.>>

3.12 detagged frame: The detagged frame of an untagged frame is the frame itself. The detagged frame of a tagged frame or a priority-tagged frame is the frame that results from untagging the frame by the appropriate procedure.

NOTE—The procedure for untagging a frame is discussed in 9.1.

3.13 EtherType encoding: The use of the Type interpretation of an IEEE 802.3 Length/Type field value in a frame as a protocol identifier conveyed in the MAC Service Data Unit of the frame.[1]

NOTE 1—The term *frame* is defined in 3.15.

NOTE 2—EtherType-encoding can be used with MAC Service user data carried on non-IEEE 802.3 MACs by means of the SNAP-based encapsulation techniques specified in IEEE Std 802.1H, IETF RFC 1042, and IETF RFC 1390.

3.14 Expedited traffic : Traffic that requires preferential treatment as a consequence of jitter, latency, or throughput constraints, or as a consequence of management policy.

3.15 Frame: A unit of data transmission on an IEEE 802 LAN that conveys a MAC Protocol Data Unit (MPDU) and can cause a service indication with, at a minimum, destination and source MAC addresses and a MAC Service Data Unit (MSDU) or an MPDU that is the result of a service request with those parameters.

3.16 Frame relay: Forwarding of frames between the Ports of a Bridge.

3.17 Group: A Group associates

a) A group MAC address; and
b) A set of properties that define membership characteristics; and
c) A set of properties that define the forwarding/filtering behavior of a Bridge with respect to frames destined for members of that group MAC address;

with a set of end stations that all wish to receive information destined for that group MAC address. Members of such a set of end stations are said to be *Group members.*

A Group is said to *exist* if the properties associated with that Group are visible in an entry in the Filtering Database of a Bridge, or in the GARP state machines that characterize the state of the Group; a Group is said to *have members* if the properties of the Group indicate that members of the Group can be reached through specific Ports of the Bridge.

NOTE—An example of the information that Group members might wish to receive is a multicast video data stream.

3.18 IEEE 802 Local Area Network (LAN): IEEE 802 LANs (also referred to in the text simply as LANs) are LAN technologies that provide a MAC Service equivalent to the MAC Service defined in ISO/IEC 15802-1. IEEE 802 LANs include IEEE Std 802.3 (CSMA/CD), IEEE Std 802.5 (Token Ring), IEEE Std 802.11 (Wireless), and ISO 9314-2 (FDDI) LANs.

3.19 Independent Virtual Local Area Network (VLAN) Learning (IVL): Configuration and operation of the Learning Process and the Filtering Database such that, for a given set of VLANs, if a given individual MAC Address is learned in one VLAN, that learned information is not used in forwarding decisions taken for that address relative to any other VLAN in the given set.

[1] The use of Ethernet Type values as a means of protocol identification was defined in the specification of Ethernet V2.0 (The Ethernet, AA-K759B-TK, Digital Equipment, Intel, and Xerox Corps., Nov. 1982).

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

NOTE—In a Bridge that supports only IVL operation, the “given set of VLANs” is the set of all VLANs.

3.20 Independent Virtual Local Area Network (VLAN) Learning (IVL) Bridge: A Bridge that supports only Independent VLAN Learning.

3.21 Internal Spanning Tree (IST): The connectivity provided by the CIST within an MST Region.

3.22 Legacy region: A set of LANs connected such that there is physical connectivity between any pair of segments using only IEEE Std 802.1D conformant, VLAN-unaware MAC Bridges.

NOTE—If, in a Bridged Local Area Network containing both IEEE Std 802.1D and IEEE 802.1Q Bridges, all the IEEE 802.1Q Bridges were to be removed, the result would be one or more Bridged Local Area Networks, each with its own distinct Spanning Tree. Each of those Bridged LANs is a legacy region.

3.23 Logical Link Control (LLC) encoding: The use of LLC addressing information in a frame as a protocol identifier associated with the MAC Service user data carried in the frame. <<fix>>

3.24 MST Bridge: A Bridge capable of supporting the CST, and one or more MSTIs, and of selectively mapping frames classified in any given VLAN to the CST or a given MSTI.

3.25 MST Configuration Table: A configurable table that allocates each and every possible VLAN to the Common Spanning Tree or a specific Multiple Spanning Tree Instance.

3.26 MST Region: A set of LANs and MST Bridges physically connected via Ports on those MST Bridges, where each LAN’s CIST Designated Bridge is an MST Bridge, and each Port is either the Designated Port on one of the LANs, or else a non-Designated Port of an MST Bridge that is connected to one of the LANs, whose MCID matches exactly the MCID of the Designated Bridge of that LAN.

NOTE—It follows from this definition that the MCID is the same for all LANs and Ports in the Region, and that the set of MST Bridges in the region are interconnected by the LANs.

3.27 Multiple Spanning Tree Algorithm and Protocol (MSTP): The Multiple Spanning Tree Algorithm and Protocol described in Clause 13 of this standard.

3.28 Multiple Spanning Tree Bridge Protocol Data Unit (MST BPDU): The MST BPDU specified in Clause 14 of this standard.

3.29 Multiple Spanning Tree (MST) Configuration Identifier: A name for, revision level, and a summary of a given allocation of VLANs to Spanning Trees.

NOTE—Each MST Bridge uses a single MST Configuration Table and Configuration Identifier.

3.30 Multiple Spanning Tree Instance (MSTI): One of a number of Spanning Trees calculated by MSTP within an MST Region, to provide a simply and fully connected active topology for frames classified as belonging to a VLAN that is mapped to the MSTI by the MST Configuration Table used by the MST Bridges of that MST Region.

3.31 Network: <<definition tbs.>>

3.32 Network Provider: <<definition tbs.>>

3.33 Priority-tagged frame: A tagged frame whose tag header carries priority information, but carries no VLAN identification information.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

3.34 protocol group database: Specifies a group of protocols by assigning a unique protocol group identifier to all protocols of the same group.

3.35 protocol group identifier: Designates a group of protocols that are associated together when assigning a VID to a frame.

3.36 protocol template: A tuple of values that specify a data-link encapsulation format and an identification of the protocol layer above the data-link layer.

3.37 Provider Bridge: <<definition tbs.>>

3.38 Provider Bridged Network: <<definition tbs.>>

3.39 Provider Edge Bridge: <<definition tbs.>>

3.40 Provider Network Customer: <<definition tbs.>>

3.41 Provider Network Port: <<definition tbs.>>

3.42 Rapid Spanning Tree Algorithm and Protocol (RSTP): The Rapid Spanning Tree Algorithm and Protocol described in Clause 17 of IEEE Std 802.1w-2001.

3.43 Rapid Spanning Tree Bridge Protocol Data Unit (RST BPDU): The RST BPDU specified in Clause 9 of IEEE Std 802.1w-2001.

3.44 Service tagged frames: <<definition tbs.>>

3.45 Service VLANs: <<definition tbs.>>

3.46 Service VIDs: <<definition tbs.>>

3.47 Service Provider: <<definition tbs.>>

3.48 Shared Virtual Local Area Network (VLAN) Learning (SVL): Configuration and operation of the Learning Process and the Filtering Database such that, for a given set of VLANs, if an individual MAC Address is learned in one VLAN, that learned information is used in forwarding decisions taken for that address relative to all other VLANs in the given set.

NOTE—In a Bridge that supports only SVL operation, the “given set of VLANs” is the set of all VLANs.

3.49 Shared Virtual Local Area Network (VLAN) Learning (SVL) Bridge: A type of Bridge that supports only Shared VLAN Learning.

3.50 Shared Virtual Local Area Network (VLAN) Learning (SVL)/Independent Virtual Local Area Network (VLAN) Learning (IVL) Bridge: An SVL/IVL Bridge is a type of Bridge that simultaneously supports both Shared VLAN Learning and Independent VLAN Learning.

3.51 Single Spanning Tree (SST) Bridge: A Bridge capable of supporting only a single spanning tree, the CST. The single spanning tree may be supported by the Spanning Tree Algorithm and Protocol (STP) defined in IEEE Std 802.1D, 1998 Edition, or by the Rapid Spanning Tree Algorithm and Protocol (RSTP), defined in IEEE Std 802.1w-2001.

3.52 Spanning Tree: A simply and fully connected active topology formed from the arbitrary physical topology of connected Bridged Local Area Network components by relaying frames through selected bridge

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

ports and not through others. The protocol parameters and states used and exchanged to facilitate the calculation of that active topology and to control the bridge relay function.

3.53 Spanning Tree Algorithm and Protocol (STP): The Spanning Tree Algorithm and Protocol described in Clause 8 of IEEE Std 802.1D, 1998 Edition.

3.54 Spanning Tree Bridge Protocol Data Unit (ST BPDU): A Bridge Protocol Data Unit specified for use by the Spanning Tree Algorithm and Protocol, i.e. a Configuration or Topology Change Notification BPDU as described in Clause 9 of IEEE Std 802.1D, 1998 Edition.

3.55 Tagged frame: A *tagged frame* is a frame that contains a tag header immediately following the Source MAC Address field of the frame or, if the frame contained a Routing Information field, immediately following the Routing Information field.

3.56 Tag header: A tag header allows user priority information, and optionally, VLAN identification information, to be associated with a frame.

3.57 VLAN-aware Bridge: A component of a system, that recognises frames with a VLAN tag, and can insert or remove tag headers. A VLAN-aware Bridge is either a Customer-VLAN aware Bridge or a Service-VLAN aware Bridge.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

4. Abbreviations

Add the following abbreviations, in the appropriate collating sequence.

S-TAG	Service VLAN TAG
S-VID	Service VLAN ID
S-VLAN	Service VLAN
C-TAG	Customer VLAN TAG
C-VID	Customer VLAN ID
C-VLAN	Customer VLAN

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Delete the existing contents of Clause 5 with the exception of clause 5.4 MAC-specific bridging methods, insert replacement contents as shown below, adding the existing clause 5.4 at the end, appropriately numbered.

<<While this clause has been substantially reorganized, the intent is not to change the conformance requirements for an existing.1Q Bridge.>>

5. Conformance

This clause specifies the mandatory and optional capabilities provided by conformant implementations of this standard. An implementation can

a) compose all or part of the functionality of a system;
b) provide, as specified by this standard, one or more instances of the MAC Service to other functional entities whose specification is outside the scope of this standard;
c) provide, as specified by this standard, one or more instances of the MAC Internal Sublayer Service to other implementations or instances of the same implementation that conform to this standard.

Accordingly, and as detailed in 5.2, this clause specifies conformance requirements for common systems and for functional components within systems, possibly connected to other system components with interfaces that are not otherwise accessible.

5.1 Terminology

<<For consistency with existing 802.1 standards, requirements are expressed using approved ISO/IEC terminology, rather than IETF terminology. In short: “shall” is used for mandatory requirements, “may” to describe implementation or administrative choices (“may” and “may not” mean precisely the same thing), “should” for recommended choices (the behaviors described by “should” and “should not” are both permissible but not equally desirable choices). The generation of the PICS is largely driven by the mechanical editorial process of searching the draft standard for occurrences of the words shall, may, and should. The draft avoids needless repetition and apparent duplication of its formal requirements by using “is”/”is not”/“are”/”are not” for definitions and the logical consequences of conformant behavior. Behavior that is permitted but is neither always required nor directly controlled by an implementor or administrator, or whose conformance requirement is detailed elsewhere, is described by “can”. Behavior that never occurs in a conformant implementation or system of conformant implementations is described by “can not”. This terminology will be documented in this clause (5.1). Note the word “allow” has be introduced as a replacement for the cliche “Support the ability for”. And the word “capability” means “can be configured to”>>

5.2 Conformant components and equipment

This clause specifies requirements and options for the following core component

a) VLAN-aware Bridge component (5.4, 5.5);

for the following two components that use that core functionality

b) Customer-VLAN aware Bridge component (5.6);
c) Service-VLAN aware Bridge component (5.7);

and for the following systems that include instances of either or both of the above two components

d) VLAN Bridge (5.8);
e) Provider Bridge (5.9).

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

NOTE—A VLAN Bridge can also be referred to as a Customer Bridge.

5.3 Protocol Implementation Conformance Statement (PICS)

A claim of conformance specifies implementation of a Customer-VLAN aware Bridge component, or a Service-VLAN aware Bridge component, or a specific system. A component or system can support multiple claims in respect of its range of possible behavior.

The supplier of an implementation that is claimed to conform to this standard shall provide the information necessary to identify both the supplier and the implementation, and shall complete a copy of the PICS proforma provided in Annex A for that specific component or system, together with the further information and completed PICS(s) required to identify subcomponents.

NOTE 1—Customer-VLAN aware and Service-VLAN aware Bridge component PICS' both require completion of a PICS for a VLAN-aware Bridge component; the VLAN Bridge PICS requires a claim of conformance for a single Customer-VLAN aware component; the Provider Edge Bridge requires a claim of conformance for a single Service-VLAN aware component and, if Provider Edge Ports are supported, one or more claims for Customer-VLAN aware components.

NOTE 2—The claim of conformance that could be made to IEEE Std 802.1Q -2003 for an implementation of a VLAN-aware Bridge, is replaced by a claim of conformance to a VLAN Bridge. While the present and subsequent amendment(s) and or revision(s) of this standard has changed the presentation of the information, the technical requirements of conformance remain unchanged.

5.4 VLAN-aware Bridge Requirements ~~Static conformance requirements~~

An implementation of a VLAN-aware Bridge component shall

a) Conform to the relevant standard for the Media Access Control technology implemented at each Port in support of the MAC Internal Sublayer Service, as specified in 6.4, 6.5, and 6.9;
b) Support the MAC Enhanced Internal Sublayer Service at each Port, as specified in 6.6 and 6.7;
c) Implement an IEEE Std 802.2 conformant LLC class with Type 1 operation as required by 8.2;
d) Relay and filter frames as described in 8.1 and specified in 8.5, 8.6, 8.7, 8.8, and 8.9;
e) On each Port, support at least one of the permissible values for the Acceptable Frame Types parameter, as defined in 8.6.1;
f) Support the following on each Port that supports untagged and priority-tagged frames:
    1) A Port VLAN Identifier (PVID) value (8.6.1);
    2) Configuration of at least one VLAN whose untagged set includes that Port (8.6.1 and 8.10.9);
    3) Configuration of the PVID value via management operations (12.10);
    4) Configuration of Static Filtering Entries via management operations (12.7).
g) Allow tag headers to be inserted, modified, and removed from relayed frames, as specified in 8.1 and Clause 9, as required by the value(s) of the Acceptable Frame Types parameter supported on each Port, and by the ability of each Port to transmit VLAN-tagged and/or untagged frames. These requirements are summarized in Table 5-1 for frames relayed between any pair of Ports;
h) Allow automatic configuration and management of VLAN topology using the Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP) (Clause 11) on all Ports;
i) Allow static and dynamic configuration information for at least one VLAN, by means of Static and Dynamic VLAN Registration Entries in the Filtering Database (8.11);
j) Support at least one Filtering Identifier (FID) (6.4, 8.11.3, 8.11.7, and 8.11.8);
k) Allow allocation of at least one VID to each FID that is supported (6.4, 8.11.3, 8.11.7, and 8.11.8).

NOTE 1—Under some circumstances, the ability for VLAN Bridges to successfully interoperate depends upon the number of FIDs supported, and the number of VIDs that can be allocated to each FID. These circumstances are discussed in Annex B, along with the implications with respect to interoperability.;

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Table 5-1—Support requirements for insertion, removal, and modification of tag headers

		Reception Port receives as (and does not discard):		
		VLAN-tagged	**Priority-tagged**	**Untagged**
Transmission Port transmits as:	**Untagged**	Shall support removal of tag headers.	Shall support removal of tag headers.	N/A.
	VLAN-tagged	Shall support conversion of the tagged frame format if the format required for the destination MAC differs from the received format.	Shall support the insertion of a non-null Virtual LAN Identifier (VID) in tag headers, plus conversion of the tagged frame format if the format required for the destination MAC differs from the received format.	Shall support the insertion of tag headers of a format appropriate to the destination MAC, carrying a non-null VID.

5.5 VLAN-aware Bridge Options

An implementation of a VLAN-aware Bridge component may

a) Support MST operation (5.5.1);

b) Support Port-and-Protocol-based VLAN classification (5.5.2), including multiple VID values per port, administrative control of the values of the multiple VIDs, and a Protocol Group Database.

c) Support Extended Filtering Services (IEEE Std 802.1D-2003 Clause 6.6.5) and the operation of GARP Multicast Registration Protocol (GMRP) (IEEE Std 802.1D-2003 Clause 10) as modified by Clause 10 of this Standard;

d) Allow the Filtering Database to contain Static and Dynamic VLAN Registration Entries (8.9) for more than one VLAN, 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.

e) 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;

f) Support enabling and disabling of Ingress Filtering (8.6.1);

g) Allow configuration of more than one VLAN whose untagged set includes that Port (8.6.5, 8.9.9);

h) Support the management functionality defined in Clause 12;

i) Support more than one FID (8.9);

j) Allow allocation of more than one VID to each supported FID (8.9, 8.9.7);

k) Allow configuration of VLAN Learning Constraints (8.9.7, 12.10.3);

l) Allow configuration of fixed VID to FID allocations (8.9.7, 12.10.3);

m) Allow configuration of the Restricted_Group_Registration parameter (IEEE Std 802.1D-2003) for each Port of the Bridge;

n) Support the ability to configure the value of the Restricted_VLAN_Registration parameter (11.2.3.2.3) for each Port of the Bridge.

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

5.5.1 Multiple Spanning Tree (MST) Operation (Optional)

A VLAN-aware Bridge implementation in conformance to the provisions of this standard for an MST Bridge (5.5, 3.24, 8.3, 8.4, 8.6.2, 8.10, 8.11, 11.2, 11.3.1, 13, 14) shall:

1) Support the Multiple Spanning Tree Protocol (MSTP) as specified in Clause 13;
2) Support the CIST plus a stated maximum number of MSTIs, where that number is at least 2 (8.10) and at most 64 (13.14);

NOTE 2—In other words, a conformant MST Bridge supports a minimum of three spanning tree instances—the CIST and at least two additional MSTIs.

3) Support a stated maximum number of FIDs not less than the number of MSTIs (8.10);
4) Support the ability to associate each FID to a spanning tree (8.10.3);
5) Support the transmission and reception of MST Configuration Identifier information (8.10.2).
6) Support a Port State for each Port for each spanning tree instance supported (8.4, 13.34);
7) Support operation of spanning tree protocol for each spanning tree instance and Port (8.12, 13);
8) Use the Bridge Group Address as specified in 8.14.3;
9) Support the default values for Bridge Forward Delay and Bridge Priority parameters specified in 13.23;
10) Support the operation of GVRP in each supported spanning tree context (11.2.3.3, 11.2.3.4);
11) Support the VLAN bridge management functions for the bridge protocol entity for each supported spanning tree, independently (12.8)
12) Support, in particular, management of the bridge priority parameters, and of the port priority and path cost parameters for every port, independently for each supported spanning tree (12.8.1.1, 12.8.1.3, 13.24);
13) Support VLAN management functions for each supported spanning tree (12.10.1 and 12.11.1);
14) Support management of the MSTI configuration (12.12).

A VLAN-aware Bridge implementation in conformance to the provisions of this standard for an MST Bridge (5.5, 13) may

15) support a greater number of FIDs than spanning trees (8.10.7).

5.5.2 Port-and-Protocol-based VLAN Classification (Optional)

A VLAN-aware Bridge implementation in conformance to the provisions of this standard for Port-and-Protocol-based VLAN classification (5.5) shall

1) support one or more of the following Protocol Classifications and Protocol Template formats: Ethernet, RFC_1042, SNAP_8021H, SNAP_Other, or LLC_Other (8.6.1 and 8.6.2);

and may

2) support configuration of the contents of the Protocol Group Database.

5.6 Customer-VLAN aware Bridge Conformance

A Customer-VLAN aware Bridge component comprises a VLAN-aware Bridge component with the EISS on all Ports supported by the use of a Customer TAG (C-TAG) (6.7, 9.5).

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

5.6.1 Requirements

A conformant implementation of a Customer-VLAN aware Bridge component shall

a) comprise a single conformant VLAN-aware Bridge component;
b) recognize and use Customer VLAN TAGs;
c) filter the Reserved MAC Addresses specified in Table 8-1;
d) use the GARP Application Address specified in Table 11-1 for GVRP; and

shall not

e) use a VLAN-translation table on any Port;
f) use Service VLAN TAGs except in support of the functionality specified in clause 6.8;

5.6.2 Options

A conformant Customer-VLAN aware Bridge component may, in addition to options specified for a VLAN-aware Bridge component whose use is not specifically prohibited by 5.6.1

a) Comprise one or more Ports capable of operating as Customer Edge Ports (5.6.3).

5.6.3 Customer Edge Port (Optional)

Each Customer Edge Port shall

1) Allow support of the ISS as specified in clause 6.8 to facilitate priority selection when connected to a Provider Bridged Network.

5.7 Service-VLAN aware Bridge Conformance

A Service-VLAN aware Bridge component comprises a VLAN-aware Bridge component with the EISS on all Ports supported by the use of a Service VLAN TAG (C-TAG) (6.7, 9.5).

5.7.1 Requirements

A conformant implementation of a Service-VLAN aware Bridge component shall

a) comprise a single conformant Service-VLAN aware Bridge component; and
b) recognize and use Service VLAN TAGs;
c) filter the Reserved MAC Addresses specified in Table 8-2;
a) use the GARP Application Address specified in Table 15-tbs for GVRP;
b) not configure any of the GARP Application Addresses specified in IEEE Std 802.1D Table 12-1 in the Filtering Database (8.10) or Permanent Database (8.10.10);
c) allow the Acceptable Frame Types parameter (8.6.1) to be set to *Admit All Frames* for each Port;
d) allow the Enable Ingress Filtering parameter (8.6.1) to be set for each Port; and

shall not

e) recognize or use Customer VLAN TAGs;
f) allow support of the ISS as specified in clause 6.8 for any of its Ports;
g) use the GVRP Application Address specified in IEEE Std 802.1Q, Table 11-1.

5.7.2 Options

A conformant Service-VLAN aware Bridge component may implement any of the options specified for a VLAN-aware Bridge component whose use is not specifically prohibited by 5.7.1, and may

a) allow the translation of received VIDs through support of a VLAN Translation Table on each Port

5.8 VLAN Bridge Conformance

A VLAN Bridge is a system that comprises a single Customer-VLAN aware Bridge component with each Port capable of connecting to an 802 LAN.

5.8.1 Requirements

A conformant implementation of a VLAN Bridge shall

a) comprise a single conformant VLAN-aware Bridge component; and

shall not

b) allow support of the ISS as specified in clause 6.9 for any of its Ports.

5.8.2 Options

A conformant Service-VLAN aware Bridge component may implement any of the options specified for a VLAN-aware Bridge component whose use is not specifically prohibited by 5.8.1.

5.9 Provider Bridge Conformance

A Provider Bridge is a system that comprises a single Service-VLAN aware Bridge component. Each Port of the Service-VLAN aware Bridge connects to either

a) A Provider Network Port; or
b) A Customer Network Port.

Each Customer Network Port can connect either

1) Directly to a customer system; or to
2) A Customer-VLAN aware Bridge component that provides one or more Provider Edge Ports.

A Provider Bridge may comprise a number of Customer-VLAN aware Bridge components. No two customers are intentionally connected to the same Customer-VLAN aware Bridge.

NOTE 1—Each Service-VLAN aware Bridge Port has a single PVID and provides ingress and egress to or from a single service instance for frames that are not Service VLAN tagged. A Customer-VLAN aware Bridge component that provides service instance selection to attached customer systems on the basis of the C-VID is attached to the Service-VLAN aware component by a number of Customer Network Ports, one per service instance.

<<In time (absent ballot comments to the contrary) all this will be specified in Clauses 15 (for service interfaces) and 16 (provider network operation), using the terminology of this clause.>>

Management of a Provider Bridge is directly under the control of the Network Provider. No elements of the Provider Bridge are manageable directly by a Provider Network Customer.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Each Bridge Port shall be identified either as a Provider Edge Port, Customer Network Port, or Provider Network Port. Clause XX of this standard specifies how that administrative determination is made.

<<Clause XX is yet to be constructed. It is anticipated that the determination will be a combination of provisioning with verification or active determination based on secure identification of the attached customer system. A nice way to do support verification/determination would be to use 802.1X/MACsec Key Agreement to figure out if a LAN at the edge of the net leads to another of the provider's own Bridges, or to a particular user's equipment. Ideas on how to express the general requirement, with the 802.1X approach being one of a possible set of methods are solicited.>>

5.9.1 Requirements

A conformant implementation of a Provider Bridge shall

a) comprise a single conformant Service-VLAN Bridge component.

5.9.2 Options

A conformant implementation of a Provider Bridge may

a) support one or more Ports capable of being configured as Provider Edge Ports (5.9.3)

5.9.3 Provider Edge Port (Optional)

If a Provider Bridge supports Provider Edge Ports, the Provider Bridge shall

1) comprise an equal number of Customer-VLAN aware Bridge components
2) allow each Provider Edge Port to be connected to a separate or a shared Customer-VLAN aware Bridge component
3) connect each Customer-VLAN aware Bridge component to the single Service-VLAN aware component with as many point-to-point connections between their Ports as the maximum number of VLANs supported by the Customer-VLAN aware Bridge.

<<Allowing a single customer to use a number of Provider Edge Ports to connect to a single Customer-VLAN aware component in a Provider Bridge facilitates the use of Link Aggregation between the customer and the provider.>>

Renumber the existing clause 5.4 MAC-specific bridging methods as clause 5.10.

5.10 MAC-specific bridging methods

This amendment makes no changes to the text of this clause (5.10).

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

6. Support of the MAC Service in VLANs

Delete the introductory paragraph of Clause 6 and insert the following.

VLAN-aware MAC Bridges interconnect the separate IEEE 802 LANs that compose a Virtual Bridged Local Area Network by relaying and filtering frames between the separate MACs of the bridged LANs.

The position of a VLAN-aware Bridge's MAC Relay Entity (8.2) within the MAC Sublayer is shown in Figure 6-1.

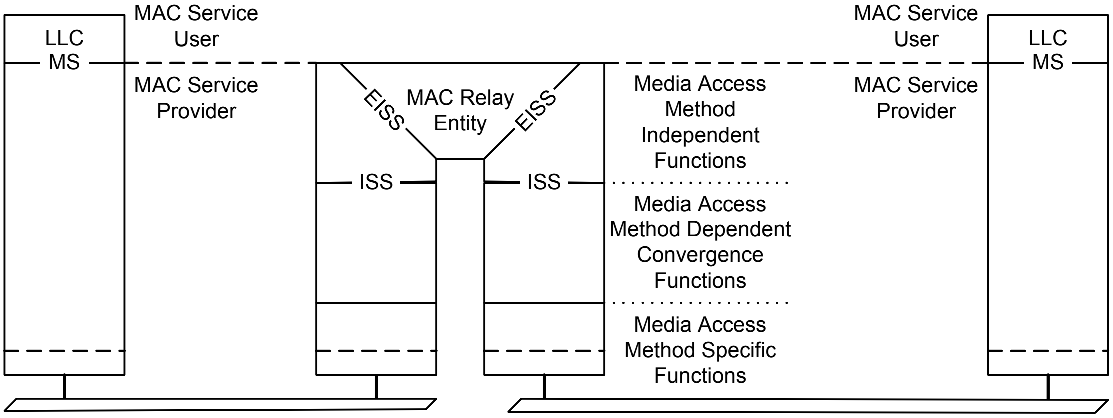
LLC
MS
MAC Service User
MAC Service Provider
MAC Relay Entity
EISS
EISS
ISS
ISS
Media Access Method Independent Functions
Media Access Method Dependent Convergence Functions
Media Access Method Specific Functions
MAC Service User
MAC Service Provider
LLC
MS

Figure 6-1—Internal organization of the MAC sublayer

The MAC Sublayer comprises:

a) Media access method specific functions[8] that realize transmission and reception of MAC Protocol Data Units (MPDUs);

b) Media access method dependent functions that use (a) to provide a media access method independent service;

c) Media access method independent functions that use a media independent service to provide the same or another media independent service.

A VLAN-aware Bridge's MAC Relay Entity forwards between the instances of the media independent Enhanced Internal Sublayer Service (EISS, 6.4). The EISS is provided by the functions specified in clause 6.7 using the media independent Internal Sublayer Service (IEEE Std 802.1D clause 6.4). The convergence functions that provide the ISS using the media specific functions for each 802 LAN MAC type are specified in 802.1D clause 6.5. The provisions of IEEE Std 802.1D, Clause 6, apply to this standard, with the additions and modification defined in this clause.

6.1 Support of the MAC service

This amendment makes no changes to clause 6.1.

6.2 Preservation of the MAC service

Replace the text of clause 6.2 with the following.

[8] The media access method specific functions together with media access method dependent convergence functions that realize a MAC Service for use in end stations are specified for each IEEE 802 LAN media access control method or 'MAC type' (e.g. 802.3, 802.11) by the relevant standard for that media access control method and are commonly referred to as "the MAC".

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

The MAC Service provided by each VLAN in a Virtual Bridged Local Area Network is similar to the service offered by a single LAN (6.3).

a) Frames transmitted between end stations carry the MAC Addresses of the peer-end stations in their destination and source address fields, not an address of a Bridge. Bridges are not directly addressed by communicating end stations, except as an end station for management purposes.
b) The MAC Addresses of end stations are not restricted by the network’s topology or configuration.
c) All MAC Addresses need to be unique within a VLAN, and within any set of VLANs for which filtering information is shared by a Bridge.

6.3 Quality of service maintenance

6.3.1 Service availability

<<The following text in 802.1Q (see 802.1s for Clause 6), may need work:>>

To maximize the service availability, no loss of service or delay in service provision should be caused by Bridges, except as a consequence of a failure, removal, or insertion of a network component, or as a consequence of the movement of an end station, or as a consequence of an attempt to perform unauthorized access. These are regarded as extraordinary events. The operation of any additional protocol necessary to maintain the quality of the MAC Service is thus limited to the configuration of the Bridged Local Area Network, and is independent of individual instances of service provision.

NOTE 1—This is true only in circumstances where admission control mechanisms are not present, i.e., where the Bridges provide a “best effort” service. The specification and applicability of admission control mechanisms in Bridges is outside the scope of this standard.

This amendment makes no changes to clauses 6.3.2, 6.3.3, 6.3.4, 6.3.5, 6.3.6, and 6.3.7.

<<Reviewers, please check that no changes are required.>>

This amendment makes no changes to clauses 6.3.8, 6.3.9, and 6.3.10.

<<Reviewers, please check that no changes are required. In particular to 6.3.10 Throughput. To what extent do we need to say anything about relative or absolute service guarantees and the presence or absence of mechanisms that insulate one Provider Network Customer from another? >>

Insert a new clause 6.4 as follows.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

6.4 Internal Sublayer Service

The Internal Sublayer Service (ISS) augments the specification of the MAC Service (ISO/IEC 15802-1) with elements necessary to the performance of the relay function. Within an end station, these additional elements are considered to be either below the MAC Service boundary, and pertinent only to the operation of the service provider; or local matters not forming part of the peer-to-peer nature of the MAC Service. The ISS excludes MAC-specific features and procedures whose operation is confined to an individual LAN.

NOTE 1—No new service primitives are defined. The frame_check_sequence is added to list of parameters associated with the MA_UNITDATA.request and MA_UNITDATA.indication primitives.

6.4.1 Service primitives and parameters

The ISS is specified by two unit-data primitives, an M_UNITDATA.indication and an M_UNITDATA.request, together with the parameters of those primitives. Each M_UNITDATA indication corresponds to the receipt of an error-free MAC frame from a LAN. A data request primitive is invoked to transmit a frame to an individual LAN.

NOTE 1—Detailed specifications of error conditions in received frames are contained in the relevant MAC standards; for example, FCS errors, length errors, non-integral number of octets.

```
M_UNITDATA.indication     (
                          destination_address,
                          source_address,
                          mac_service_data_unit,
                          priority,
                          frame_check_sequence
                          )

M_UNITDATA.request        (
                          destination_address,
                          source_address,
                          mac_service_data_unit,
                          priority,
                          frame_check_sequence
                          )
```

The **destination_address** parameter is the address of an individual MAC entity or a group of MAC entities. The **source_address** parameter is the individual address of the source MAC entity. The **mac_service_data_unit** parameter is the service user data. The default **priority** value is 0. Values 1 through 7 form an ordered sequence of user_priorities, with 1 being the lowest value and 7 the highest.

The **frame_check_sequence** parameter is explicitly provided with the M_UNITDATA.indication so that it can be used in a related M_UNITDATA.request. The parameter comprises the FCS value and sufficient information to determine whether the FCS value can be used. If the frame_check_sequence parameter is provided with an M_UNITDATA.request and the receiving and the transmitting service providers

a) Use the same algorithm to determine the FCS; and
b) Apply that algorithm to the same fields of the frame, i.e. the FCS coverage is the same; and
c) The data that is within the coverage of the FCS remains the same;

the transmitting service provider may use the supplied FCS value (7.1, 7.2, 6.3.7).

NOTE 2—There are two possibilities for recreating a valid FCS. The first is to generate a new FCS by algorithmically modifying the received FCS, based on knowledge of the FCS algorithm and the transformations that the frame has undergone between reception and transmission. The second is to rely on the normal MAC procedures to recalculate the FCS for the outgoing frame. The former approach can be preferable in terms of its ability to protect against increased levels of undetected frame errors. Annex G of IEEE Std 802.1D, 1998 Edition discusses these possibilities in more detail. The frame_check_sequence parameter of the Enhanced Internal Sublayer Service (7.1) is able to signal the validity, or otherwise, of the FCS; an unspecified value in this parameter in a data request indicates to the transmitting MAC that the received FCS is no longer valid, and the FCS must therefore be recalculated.

The identification of the LAN from which particular frames are received is a local matter and is not expressed as a parameter of the service primitive.

NOTE 3—The ISS specification in this standard differs from that in IEEE Std 802.1D-2004 as it omits the frame_type and access_priority parameters. The frame_type is not required as the receipt of a frames other than a user data frame does not cause a data indication, nor are such frames transmitted by the media independent bridge functions. The mapping of the ISS to particular access methods specified by this standard includes derivation of the access_priority parameter (for those media that require it) from the ISS priority parameter.

6.4.2 Status parameters

The Internal Sublayer Service also makes available status parameters that reflect the operational state and administrative controls over each instance of the service provided.

The **MAC_Enabled** parameter is TRUE if use of the service is permitted; and is otherwise FALSE. The value of this parameter is determined by administrative controls specific to the entity providing the service, as specified in 6.5.

The **MAC_Operational** parameter is TRUE if the entity providing the service is capable of transmitting and receiving frames and its use is permitted by management, i.e. MAC_Enabled is also TRUE. Its value is otherwise FALSE. The value of this parameter is determined by the specific MAC procedures, as specified in 6.5.

NOTE—These status parameters provide a common approach across MACs for handling the fact that:

a) A MAC can inherently be working or not;
b) If the MAC is working, its operational state can be administratively overridden.

6.4.3 Point-to-point parameters

The Internal Sublayer Service also makes available status parameters that reflect the point-to-point status of each instance of the service provided and provide administrative control over the use of that information.

If the **operPointToPointMAC** parameter is TRUE if the service is used as if it provides connectivity to at most one other system, if FALSE the service is used as if it can provide connectivity to a number of systems.

The **adminPointToPointMAC** parameter can take one of three values. If it is

a) **ForceTrue**, operPointToPointMAC shall be TRUE, regardless of any indications to the contrary generated by the service providing entity.
b) **ForceFalse**, operPointToPointMAC shall be FALSE.
c) **Auto**, operPointToPointMAC is determined by the service providing entity, as specified in 6.5.

The value of operPointToPointMAC is determined dynamically; i.e., it is re-evaluated whenever adminPointToPointMAC or the status of the service providing entity changes.

Renumber the existing clause 6.5 as clause 6.5.1, and insert a new heading and introductory paragraphs for clause 6.5 as follows:

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

6.5 Support of the Internal Sublayer Service by specific MAC procedures

This clause specifies support of the Internal Sublayer Service by MAC Entities that use specific IEEE 802 media access methods, including the mapping to the MAC protocol and procedures for each access method, and the encoding of the parameters of the service in MAC frames. The mapping is specified by reference to the IEEE 802 standards that specify each access method. The mapping draws attention to any special responsibilities of Bridges attached to LANs of that type. Control frames, i.e. frames that do not convey MAC user data, do not give rise to ISS data indications and are therefore not forwarded by a Bridge to any LAN other than that on which they originated.

Each MAC Entity examines all frames received on the LAN to which it is attached. All error-free received user data frames give rise to M_UNITDATA indication primitives. A frame that is in error, as defined by the relevant MAC specification, is discarded by the MAC Entity without giving rise to any M_UNITDATA indication.

Support of the ISS by the CSMA/CD access method is specified in clause 6.5 of IEEE Std 802.1D-2003 as amended by clause 6.5.1 of this Standard.

Support of the ISS by the Wireless LAN (802.11) access method is specified in clause 6.5 of IEEE Std 802.1D-2003.

Support of the ISS by the Token Ring (802.5) access method is specified in clause 4.1.13.2 of IEEE Std 802.5-1998. Following an M_UNITDATA.request, the user priority parameter referenced by that clause shall be set equal to the value of the priority parameter of the ISS as specified in this standard. The column marked “8802-5(default)” in Table 6-2 should be used to derive the access priority from the user priority, but the column marked “8802-5 (alternate)” may be used for backwards compatibility with equipment manufactured in accordance with ISO/IEC 10038: 1993. The use of this alternate mapping reduces the number of available access priority values to three. The frame type of each frame resulting from an M_UNITDATA.request shall be LLC, and the receipt of frames of other types shall not result in an M_UNITDATA.indication.

Table 6-2—FDDI and Token Ring access priorities

user priority	access priority		
	8802-5 (default)	**8802-5 (alternate)**	**FDDI**
0	0	4	0
1	1	4	1
2	2	4	2
3	3	4	3
4	4	4	4
5	5	5	5
6	6	6	6
7	6	6	6

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Following receipt of a frame, the priority parameter signaled in the corresponding M_UNITDATA.indication shall be regenerated from the user priority specified in clause 4.1.13.2 of IEEE Std 802.5-1998 using a User Priority Regeneration Table for the MAC instance. This specifies the regenerated priority for each of the eight possible received values (0 through 7). Table 6-3 defines default values. If these are modified by management, the value of the table entries may be independently set for each MAC instance and value of received user priority, and may use the full range of values in the parameter ranges specified.

NOTE 3—It is important that the regeneration and mapping of priority be consistent with the end-to-end significance of that priority in the network. Within a given Bridge, the values chosen for the User Priority Regeneration Table for a given Port should be consistent with the priority to be associated with traffic received through that Port across the rest of the network, and should generate appropriate access priority values for each media access method on transmission.

Table 6-3—FDDI and Token Ring user priority regeneration

Received user priority	Default regenerated priority	Range
0	0	0–7
1	1	0–7
2	2	0–7
3	3	0–7
4	4	0–7
5	5	0-7
6	6	0–7
7	7	0–7

Support of the ISS by the FDDI access methods is specified in clause 6.5 of IEEE Std 802.1D-2003. The user priority parameter referenced by that clause shall be set equal to the value of the priority parameter of the ISS as specified in this standard. The column marked “FDDI” in Table 6-2 should be used to derive the access priority from the user priority. A User Priority Regeneration Table for the MAC instance shall be used to determine the priority parameter of each M_UNITDATA.indication as specified above for Token Ring.

Change the initial paragraph of clause 6.5.1 as follows:

~~In addition to the provisions of 6.5.1 of IEEE Std 802.1D, 1998 Edition, o~~On receipt of an M_UNITDATA.request primitive that represents a tagged frame, the implementation is permitted to adopt either of the following approaches with regard to the operation of Transmit Data Encapsulation for frames whose length would, using the procedure as described, be less than 68 octets:

Replace the existing clause 6.4 with new clauses 6.6 and 6.7 as follows.

6.6 Enhanced Internal Sublayer Service

The Enhanced Internal Sublayer Service (EISS) is derived from the Internal Sublayer Service (ISS~~, defined in 6.4 of IEEE Std 802.1D, 1998 Edition~~) by augmenting that specification with elements necessary to the operation of the tagging and untagging functions of a VLAN-aware Bridge (3.57). Within the attached end station, these elements can be considered to be either below the MAC Service boundary, and pertinent only to the operation of the service provider; or local matters not forming part of the peer-to-peer nature of the MAC Service.

The EISS provides the same service status and point-to-point parameters as the ISS(6.4.2, 6.4.3).

6.6.1 ~~E-ISS service definition~~Service primitives

The unit-data primitives that define this service are

```
EM_UNITDATA.indication          (
                                ~~frame_type,~~
                                ~~mac_action,~~
                                destination_address,
                                source_address,
                                mac_service_data_unit,
                                ~~user_~~priority,
                                drop_eligible,
                                vlan_identifier,
                                frame_check_sequence,
                                canonical_format_indicator,
                                rif_information (optional)
                                )

EM_UNITDATA.request             (
                                ~~frame_type,~~
                                ~~mac_ action,~~
                                destination_address,
                                source_address,
                                mac_service_data_unit,
                                ~~user_~~priority,
                                drop_eligible,
                                vlan_identifier,
                                ~~access_priority,~~
                                frame_check_sequence,
                                canonical_format_indicator,
                                rif_information (optional),
                                include_tag
                                )
```

The ~~frame_type, mac_action,~~ **destination_address, source_address, mac_service_data_unit,** ~~**user_**~~**priority,** and **frame_check_sequence** parameters are as defined for the ISS.

The **drop_eligible** parameter provides guidance to the recipient of the service indication or of a indication corresponding to the service request, and takes the values True or False. If drop_eligible is True the parameters of the indication should be discarded in preference to discarding those corresponding to indications with comparable parameters and drop_eligible False.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

The **vlan_identifier** parameter carries the VLAN identifier (VID,).

The **include_tag** parameter carries a Boolean value. True indicates to the service provider that the mac_service_data_unit parameter of the data request shall include a tag header (9.3). False indicates that a tag header shall not be included.

The **canonical_format_indicator** parameter indicates whether embedded MAC Addresses carried in the mac_service_data_unit parameter are in Canonical format or Non-canonical format. The value False indicates Non-canonical format. The value True indicates Canonical format.

NOTE—The meanings of the terms Canonical format and Non-canonical format are discussed in Annex F.

The **rif_information** parameter is present if a tag header containing a Routing Information Field (RIF) is present (indication primitive) or requested (request primitive). Its value is equal to the value of the RIF.

6.7 Support of the EISS ~~in VLAN-aware Bridges~~

The EISS is supported by tagging and detagging functions that in turn use the ISS (6.4, 6.5). Any given instance of the EISS shall be supported by using one but not both of the following VLAN tag types:

a) Customer VLAN tag (C-TAG); or
b) Service VLAN tag (S-TAG);

selected as specified in Clause 9.5.

6.7.1 Data indication~~ primitive~~s

On receipt of a data indication from the Internal Sublayer Service, an EM_UNITDATA.indication primitive is invoked, with parameter values determined as follows:

The ~~**frame_type, mac_action,**~~ **destination_address, source_address,** and **frame_check_sequence** parameters carry values equal to the corresponding parameters in the received data indication.

The frame is determined to be tagged if the initial octets of the mac_service_data_unit parameter contain a VLAN tag of the type used to support the EISS (6.7, 9.3)

The value of the **mac_service_data_unit** parameter is as follows:

a) If the received mac_service_data_unit parameter contained a tag header (9.3), then the value used is equal to the value of the received mac_service_data_unit following removal of the tag header. Otherwise;
b) The value used is equal to the value of the received mac_service_data_unit.

The value of the **vlan_identifier** parameter is determined as follows:

c) If the initial octets of the received mac_service_data_unit parameter compose a VLAN tag (9.3), then the value contained in the VID field of the tag header is used. Otherwise;
d) A value equal to the null VLAN ID (as defined in Table 9-2) is used.

The values of the ~~**user_**~~**priority** and **drop_eligible** parameters are as follows:

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

e) If the received mac_service_data_unit parameter contained a tag header (9.3), then the values are determined using the values encoded in the priority code point and drop_eligible fields of the tag header as specified in 6.7.4 and Clause 9. Otherwise;

f) The value of the priority parameter is regenerated from the received ~~user_~~priority, as specified in 6.7.4~~ and 6.4 of IEEE Std 802.1D, 1998 Edition~~, and the drop_eligible parameter is False.

The value of the **canonical_format_indicator** parameter is determined as follows:

g) If the received mac_service_data_unit parameter is VLAN tagged (9.3), then the value is as specified in Clause 9. Otherwise;

h) If the MAC entity that received the data indication was an ISO/IEC 8802-5 Token Ring MAC, then the parameter carries the value False. Otherwise;

i) The parameter carries the value True.

The value of the **rif_information** parameter is determined as follows:

j) If the initial octets of the received mac_service_data_unit parameter compose a VLAN tag (9.3), then the value is as specified in Clause 9. Otherwise;

k) The parameter is not present.

NOTE 2—This field can be present only in tag headers received using 802.3/Ethernet or transparent FDDI MACs. The presence of one or more route descriptors indicates that there is source-routing information present in the received frame.

6.7.2 Data request~~ primitive~~s

On invocation of a data request primitive by a user of the E-ISS, an M-UNITDATA.request primitive is invoked, with parameter values as follows:

The~~ **frame_type, mac_action,**~~ **destination_address, source_address,** ~~**user_priority**~~, and ~~**access_**~~**priority** parameters carry values equal to the corresponding parameters in the received data request.

If the value of the **include_tag parameter** is True, then a tag header, formatted as necessary for the destination MAC type, is inserted as the initial octets of the mac_service_data_unit parameter. The values of the vlan_classification, priority, drop_eligible, canonical_format_indicator, and rif_information (if present) parameters are used to determine the contents of the tag header, in accordance with Clause 9.

If the value of the **include_tag parameter** is False, then no tag header is inserted.

The remaining octets of the **mac_service_data_unit** parameter are those accompanying the EISS-request. If the data request is a consequence of relaying a frame and the MAC type of the Port differs from that used to receive the frame, they are modified, if necessary, in accordance with the procedures described in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390. If 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 type, and the value of the **include_tag** parameter is False or the tag to be used is a Service VLAN tag (S-TAG), then the Bridge shall either

a) Convert any embedded MAC Addresses in the mac_service_data_unit to the format appropriate to the destination MAC type; or

b) Discard the EISS data request without issuing a corresponding ISS data request.

The value of the **frame_check_sequence** parameter is determined as follows:

c) If the frame_check_sequence parameter received in the data request is either unspecified or still carries a valid value, then that value is used. Otherwise;

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

d) The value used is either derived from the received FCS information by modification to take account of the conditions that have caused it to become invalid, or the unspecified value is used.

<<Is there any issue related to a Service Provider offering a point-to-point service and the values of the adminPointToPointMAC and operPointToPointMAC parameters? (see 802.1D 6.4.1)>>

6.7.3 Priority encoding

The priority and drop_eligible parameters are encoded in the priority code point (PCP) field of the VLAN tag using the Priority Encoding Table for the Port, and decoded from the PCP using the Priority Decoding Table. For each Port, the Priority Encoding Table has 16 entries, corresponding to each of the possible combinations of the eight possible values of priority (0 through 7) with the two possible values of drop_eligible (True or False). For each Port, the Priority Decoding Table has 8 entries, corresponding to each of the possible PCP values.

Alternate values for each table are specified as rows in Table 6-4 and Table 6-5, each alternative labelled by the number of distinct priorities that can be communicated, and the number of these for which drop precedence can be communicated. For example, the table entries 6P2D allow 6 distinct priorities with drop precedence for two. The combination of the priority and drop_eligible parameters is shown by the priority alone if drop_eligible is False, and by the priority followed by the letters “DE” if drop_eligible is True.

Table 6-4—Priority encoding

priority drop_eligible		7	7DE	6	6DE	5	5DE	4	4DE	3	3DE	2	2DE	0	0DE	1	1DE
PCP	**8P0D**	7	7	6	6	5	5	4	4	3	3	2	2	0	0	1	1
	7P1D	7	7	6	6	5	4	5	4	3	3	2	2	0	0	1	1
	6P2D	7	7	6	6	5	4	5	4	3	2	3	2	0	0	1	1
	5P3D	7	7	6	6	5	4	5	4	3	2	3	2	0	1	0	1

Table 6-5—Priority Code Point decoding

PCP		7	6	5	4	3	2	0	1
priority drop_eligible	**8P0D**	7	6	5	4	3	2	0	1
	7P1D	7	6	4	4DE	3	2	0	1
	6P2D	7	6	4	4DE	2	2DE	0	1
	5P3D	7	6	4	4DE	2	2DE	0	0DE

NOTE 1—Annex G discusses the selection of values for inclusion in Table 6-4 and Table 6-5.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

The values in the Priority Encoding Table and Priority Decoding Table may be modified by management, as described in Clause 12. If this capability is provided, the value of the table entries shall be independently settable for each Port. The values shall be constrained as follows:

a) The lower of any two priority values shall not map to a higher PCP than the higher priority.
b) The lower of any two PCP values shall not map to a higher priority than the higher PCP.
c) Any priority value combined with drop_eligible True shall not map to a higher PCP than the same priority value combined with drop_eligible False.
d) Any PCP value shall not map to a priority value combined with drop_eligible True if a lower PCP maps to the same priority value combined with drop_eligible False.

The default priority, 0, is taken as having a value intermediate between 2 and 1 for the purpose of these rules. The values may be further constrained, but if the tables can be modified the default values shown in Table 6-4 and Table 6-5, and the alternative sets of values in Table 6-4 and Table 6-5 shall be settable.

If the VLAN tag is a Service VLAN TAG (S-TAG) the drop_eligible parameter may also be encoded in and decoded from the DE bit in the S-TAG. If this capability is provided, it shall be independently manageable for each Port. If the Use DE Bit parameter is set for the port, the drop_eligible parameter is encoded in the S-TAGs of transmitted frames and shall be True for a received frame if the DE bit is set in the S-TAG or the Priority Decoding Table indicates drop_eligible True for the received PCP value.

6.7.4 Regenerating priority

The priority of received frames without a VLAN tag header is regenerated using priority information contained in the frame and the User Priority Regeneration Table for the reception Port. For each reception Port, the User Priority Regeneration Table has eight entries, corresponding to the eight possible values of priority (0 through 7). Each entry specifies, for the given value of received priority, the corresponding regenerated value.

NOTE 1—IEEE 802 LAN technologies signal a maximum of eight user_priority values. Annex G (informative) further explains the use of priority values and how they map to traffic classes.

Table 6-6 specifies default regenerated priority values for each of the eight possible values of the received priority. These default values shall be used as the initial values of the corresponding entries of the User Priority Regeneration Table for each Port.

The values in the User Priority Regeneration Table may be modified by management, as described in Clause 12. If this capability is provided, the value of the table entries shall be independently settable for each reception Port and for each value of received priority, and the Bridge shall have the capability to use the full range of values in the parameter ranges specified in Table 6-3.

NOTE 2—The regeneration and mapping of priority within the Bridge should be consistent with the end-to-end significance to that priority across the rest of the Bridged Local Area Network. The regenerated priority value is used:

— Via the traffic class table (8.6.6) to determine the traffic class for a given outbound Port, and
— Via fixed, MAC type-specific mappings (6.5) to determine the access priority that will be used for certain media access methods.

Insert a new clause, 6.8 as follows:

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

Table 6-6—Priority regeneration

Received priority	Default regenerated priority	Range
0	0	0–7
1	1	0–7
2	2	0–7
3	3	0–7
4	4	0–7
5	5	0-7
6	6	0–7
7	7	0–7

6.8 Support of the ISS for attachment to a Provider Bridged Network

This standard specifies both Customer Bridges (3.5) and Provider Bridges (3.37). The operation of Provider Bridges and Provider Bridged Networks is, by design, largely transparent to the operation of Customer Bridges and Customer Bridged Local Area Networks. Figure 15-1 illustrates the relationship of the service provided by the MAC Sublayer functionality of Provider Bridges to that used by a Customer Bridge.

The functions specified in this Clause (6.8) provide the ISS to the MAC Relay Entity of a Customer Bridge by making use of the ISS or the EISS. These functions shall be one of:

a) Null, i.e. each request by the user (Clause 6.4.2) of the provided ISS results in a request of the supporting ISS provider (Clause 6.5 of IEEE Std 802.1D) with identical parameters.
b) Service access priority selection, i.e. each request by the user (Clause 6.4.2) of the provided ISS results in a request of a supporting EISS provider (6.7) with the EISS priority set to the value in the Access Priority Table corresponding to the requested priority ISS request.

Specification of a null function allows a Customer Bridge without additional functionality to connect to a Provider Bridged Network.

Specification of the service access priority selection function allows the Customer Bridge Port to request priority handling of the frame from the provider network on the basis of the priority assigned within the Customer Bridged Local Area Network, while allowing for differences in cost and service level ascribed to individual values of priority within the provider and customer networks.

Following each request by the user of the provided ISS with the following parameters:

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

```
M_UNITDATA.request          (
                            destination_address,
                            source_address,
                            mac_service_data_unit,
                            priority,
                            frame_check_sequence,
                            )
```

a request is made of the underlying ESS with parameters as follows:

```
EM_UNITDATA.request         (
                            destination_address = M_UNITDATA.req.destination_address,
                            source_address = M_UNITDATA.req.source_address,
                            mac_service_data_unit = M_UNITDATA.req.mac_service_data_unit,
                            priority = Access Priority Table(M_UNITDATA.req.user_priority),
                            frame_check_sequence = unspecified,
                            canonical_format_indicator = False,
                            vlan_classification = Null,
                            rif_information (optional) = not present,
                            include_tag = True
                            )
```

NOTE—The user_priority of the original service request made of the customer network can be carried transparently and unchanged across the provider network in the Customer TAG.

Insert a new clause 6.9 as follows:

6.9 Support of the ISS within a system

A single system can comprise two or more instances of VLAN-aware Bridge components connected by one or more of their Ports. If those Ports are not otherwise accessible the ISS may be provided by means outside the scope of this specification. Each instance of such an implementation of the ISS shall support the MAC status and Point-to-point parameters. An M_UNITDATA.request at one of the Ports connected to such an instance of the ISS shall result in M_UNITDATA.indications with identical parameters at all other Ports connected to that instance.

For convenience of management such instances of ISS provision are assigned the LAN MAC Type 802.1. Management parameters common to MACs of this type are specified in Clause 12.

NOTE—The ISS can also be supported at Bridge Ports wholly contained within a system by any 802 LAN technology, subject to the system requirements of the particular media access method.

Insert a new clause, 6.10 as follows:

6.10 Support of the ISS by additional technologies

<<General guidance for support of the ISS by other technologies without restricting or overspecifying those technologies. A place to formally record the normative essentials that come out of any liaison activities. This editors' note solicits contributions.>>

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

7. Principles of network operation

This amendment makes no changes to Clause 7.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

8. Principles of bridge operation

This clause

a) Explains the principal elements of VLAN-aware Bridge operation and lists the supporting functions.
b) Establishes a Bridge architecture that governs the provision of these functions.
c) Provides a model of Bridge operation in terms of processes and entities that support the functions.
d) Details the addressing requirements in a Bridged Local Area Network.
e) Specifies the addressing of Entities in a Bridge.

8.1 Bridge operation

The principal elements of Bridge operation are

a) Relay and filtering of frames (8.1.1).
b) Maintenance of the information required to make frame filtering and relaying decisions (8.1.2).
c) Management of the above (Clause 12).

8.1.1 Relay

A MAC Bridge relays individual MAC user data frames between the separate MACs of the bridged LANs connected to its Ports. The functions that support relaying of frames and maintain the Quality of Service are

a) Frame reception.
b) Discard on received frame in error (6.3.2).
c) Discard of frames that do not carry user data (6.5).
d) Regeneration of priority, if required (6.7).
e) Priority and drop eligibility decoding from a VLAN TAG, if present (6.7).
f) Application of VLAN ingress rules to classify each received frame to a particular VLAN (8.6.1).
g) Frame discard to support management control over the active topology of each VLAN (8.6.1).
h) Frame discard to suppress loops in the physical topology of the network (8.6.2).
i) Frame discard following the application of filtering information (8.6.3).
j) Metering of frames, potentially discarding frames exceeding bandwidth limits and marking frames exceeding bandwidth guarantees as drop eligible (8.6.4).
k) Forwarding of received frames to other Bridge Ports (8.6.5).
l) Application of VLAN egress rules to determine if a VLAN tag is to be included in the frame (8.6.5).
m) Selection of traffic class and queuing of frames by traffic class (8.6.6).
n) Frame discard to ensure that a maximum bridge transit delay is not exceeded (6.3.6, 8.6.7).
o) Preferential discard of drop eligible frames to preserve QoS for other frames (8.6.7).
p) Selection of queued frames for transmission (8.6.8).
q) Mapping of service data units (6.3.7).
r) Frame discard on transmittable service data unit size exceeded (6.3.8).
s) Selection of outbound access priority (6.3.9).
t) Mapping of service data units and Frame Check Sequence recalculation, if required (6.3.7).
u) Frame discard if the service data unit cannot be mapped correctly.
v) Frame transmission.

Figure 8-1 gives an example of the physical topology of a Bridged Local Area Network.

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

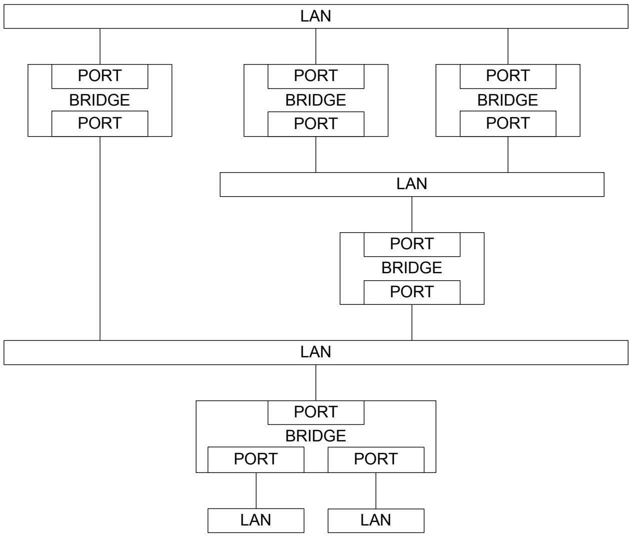

Figure 8-1—A Bridged Local Area Network

8.1.2 Filtering and relaying information

A Bridge maintains filtering and relaying information for the following purposes:

a) Duplicate frame prevention: to maintain a loop-free active topology for each VLAN;
b) Traffic segregation: to separate communication by different sets of network users;
c) Traffic reduction: to confine frames to the path(s) between their source and destination(s);
d) Traffic expediting: to classify frames in order to expedite time critical traffic;
e) Frame format conversion: to tag or untag as appropriate for the destination LAN and stations.

8.1.3 Duplicate frame prevention

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 (6.3.4). The functions that support the use and maintenance of information for this purpose are

a) Configuration and calculation of one or more spanning tree active topologies.
b) In MST Bridges, explicit configuration of the relationship between VIDs and spanning trees (8.10).

8.1.4 Traffic segregation

A Bridge can filter frames to confine them to LANs that belong to the VLAN to which they are assigned, and thus define the VLAN’s maximum extent (7.3). The functions that support the use and maintenance of information for this purpose are

a) Configuration of a PVID, to associate a VID with received untagged and priority-tagged frames;
b) In Bridges implementing Port-and-Protocol based classification, the configuration of a VID for untagged and priority-tagged frames for each protocol;
c) Enabling or disabling the application of Static VLAN Registration Entries to received frames through configuration of the Enable Ingress Filtering parameter;
d) Configuration of Static VLAN Registration Entries.

A Bridge can filter frames to partially partition a Virtual Bridged Local Area Network. Frames assigned to any given VLAN and addressed to specific end stations or groups of end stations can be excluded from relay

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

to certain Bridge Ports. The functions that support the use and maintenance of information for this purpose are:

e) Permanent configuration of Reserved Addresses (Table 8-1);
f) Configuration of Static Filtering Entries (8.9.1) and Group Registration Entries (8.9.4).

NOTE—The use of VLANs is generally less error prone and is preferred to filtering using destination addresses if a Bridged Local Area Network is to be partitioned for reasons of scale, efficiency, management, or security. Destination address filtering is the only mechanism available to Bridges that are not VLAN aware.

8.1.5 Traffic reduction

A Bridge can filter frames to confine them to LANs that either have end stations attached to their assigned VLAN or that connect those LANs, and thus define the current practical extent of the VLAN (7.4). LANs not attaching to or forming part of the path between the source and destination(s) of any given communication do not have to support the transmission of related frames, potentially improving the quality of the MAC service for other communications. The functions that support the use and maintenance of information for this purpose are:

a) Automatic learning of dynamic filtering information for unicast destination addresses through observation of source addresses of frames;
b) Ageing out or flushing of dynamic filtering information that has been learned to support the movement of end stations and changes in active topology;
c) Automatic inclusion and removal of Bridge Ports in the VLAN, through configuration of Dynamic VLAN Registration Entries by means of GVRP (8.9.5 and 11.2);
d) Explicit configuration of management controls associated with the operation of GVRP by means of Static VLAN Registration Entries (8.9.2 and 11.2);
e) Automatic configuration of Group Registration Entries by means of GMRP exchanges;
f) Explicit configuration of the management controls associated with the operation of GMRP by means of Group Registration Entries.

8.1.6 Traffic expediting

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

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

8.1.7 Conversion of frame formats

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 (8.6.5). The function that supports the use and maintenance of information for this purpose is

a) Explicit configuration of tagging requirements on egress for each Port (8.9.2 and 8.9.9).

8.2 Bridge architecture

A Bridge comprises

a) A MAC Relay Entity that interconnects the Bridge’s Ports;
b) At least two Ports;
c) Higher layer entities, including at least a Spanning Tree Protocol Entity.

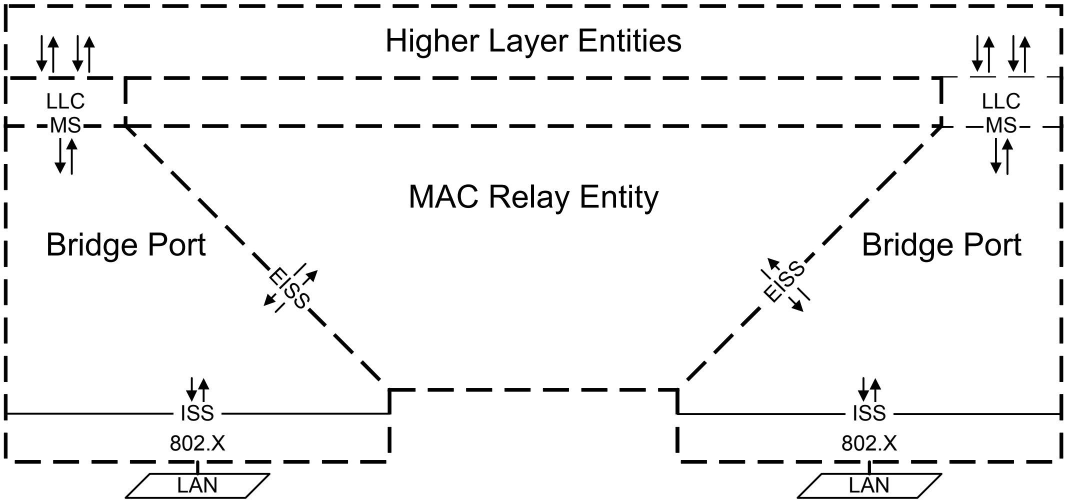

Figure 8-2—VLAN-aware Bridge Architecture

The MAC Relay Entity handles the media access method independent functions of relaying frames between Bridge Ports, filtering frames, and learning filtering information. It uses the Enhanced Internal Sublayer Service (EISS) (6.4, 6.5) provided by each Bridge Port.

Each Bridge Port also functions as an end station providing one or more instances of the MAC Service. Each instance of the MAC Service is provided to a distinct LLC Entity that supports protocol identification, multiplexing, and demultiplexing, for PDU transmission and reception by one or more higher layer entities.

NOTE 1—In most cases each Port provides a single instance of the MAC Service, to an LLC Entity that supports all the Higher Layer Entities that require a point of attachment to the Port. Further instances are only provided when the specifications of the Higher Layer Entities require the use of different instances of the MAC service or of different source addresses.

An LLC Entity for each Bridge Port shall use an instance of the MAC Service provided for that Port to support the operation of LLC Type 1 procedures in order to support the operation of the Spanning Tree Protocol Entity. Bridge Ports may support other types of LLC procedures for use by other protocols, such as protocols providing Bridge Management (8.13).

NOTE 2—For simplicity of specification this standard refers to a single LLC Entity that can provide both the procedures specified by IEEE Std 802.2 and Ethernet Type protocol discrimination in the cases where the media access method for the attached LAN supports the latter.

If the Bridge Port can be directly attached to an IEEE 802 LAN, an instance of the MAC for that LAN type is permanently associated with the Port, handles the media access method specific functions (MAC protocol and procedures), and provides an instance of the Internal Sublayer Service (ISS) as specified in Clause 6.4 to support frame transmission and reception by the other processes and entities that compose the Port.

Figure 8-2 illustrates a Bridge with two Ports, each directly connected to a LAN.

8.3 Model of operation

The model of operation is simply a basis for describing the functionality of the MAC Bridge. It is in no way intended to constrain real implementations of a MAC Bridge; these may adopt any internal model of operation compatible with the externally visible behavior that this standard specifies. Conformance of equipment to this standard is purely in respect of observable protocol.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

The processes and entities that model the operation of the MAC Relay Entity are

a) The Forwarding Process (8.6), which forwards received frames that are to be relayed to other Bridge Ports, filtering frames on the basis of information contained in the Filtering Database (8.9) and on the state of each Port (8.4);
b) The Learning Process (8.7), which, by observing the source addresses of frames received on each Port, updates the Filtering Database, conditionally on the Port State;
c) The Filtering Database (8.9), which holds filtering information and supports queries by the Forwarding Process as to whether frames with given values of the destination MAC Address field can be forwarded to a given Port.

In a VLAN-aware Bridge, these processes uses the Ingress Rules (8.6.1) for each Port, which comprise the

d) Acceptable Frames Types parameter;
e) VLAN Classification parameters;
f) Enable Ingress Filtering parameter;

and can include

g) The VID Translation Table.

NOTE—A Provider Bridge, comprising a Service-VLAN aware Bridge component, may use a VLAN Translation Table (5.7). A VLAN Bridge, comprising a Customer-VLAN aware Bridge component, does not (5.6).

To enforce a loop-free active topology (8.6.2) the Forwarding and Learning Processes make use of the following information for each Port

h) Port State (8.4)

in an MST Bridge there is a Port State for each spanning tree instance, together with the

i) MST Configuration Table.

To confine frames classified as belonging to a given VLAN to a subset of the active topology, and to determine the format of forwarded frames, the Forwarding Process uses the Egress Rules (8.6.5) for each Port, which comprise

j) VLAN Registration Entries in the Forwarding Database.

Figure 8-3 illustrates a single instance of frame relay between the Ports of a Bridge with two Ports.

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

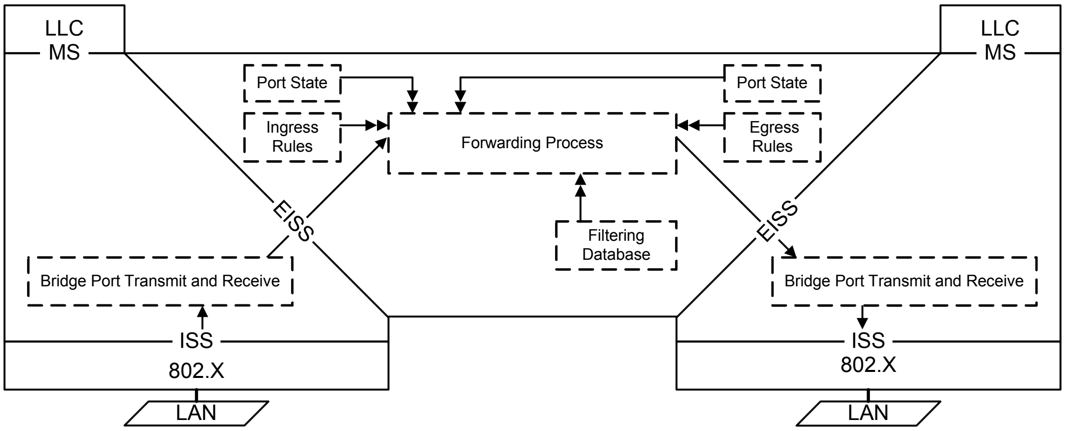

Figure 8-3—Relaying MAC frames

Figure 8-4 illustrates the inclusion of information carried by a single frame, received on one of the Ports of a Bridge with two Ports, in the Filtering Database.

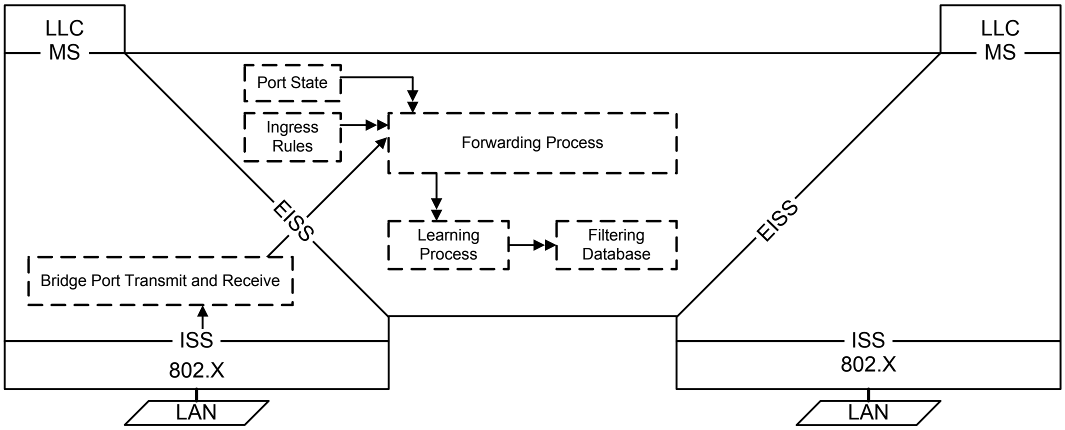

Figure 8-4—Observation of network traffic

The processes and entities that model the operation of a Bridge Port include

k) a Bridge Port Transmit and Receive Process, which delivers and accepts frames to and from the MAC Relay Entity and the LLC Entity or Entities and uses the service provided by an instance of the ISS to receive and transmit frames from and to the attached network;

l) the LLC Entity or Entities that support Higher Layer Entities.

If the Bridge Port is directly attached to an 802 LAN, the Bridge Port Transmit and Receive Process uses an instance of the ISS provided by an instance of the MAC for that LAN type as specified in Clause 6.4.

If the Bridge is a Customer-VLAN aware and the Port is providing connectivity to a Provider Bridge Network, the Bridge Port Transmit and Receive Process uses an instance of the ISS that supports a single point of attachment to that network as specified in Clause 6.6.

Higher Layer Entities that require only one point of attachment for the Bridge as a whole may attach to an LLC Entity that uses an instance of the MAC Service provided by a Management Port. A Management Port does not use an instance of the ISS to attach to a network, but uses the Bridge Port Transmit and Receive Process and the MAC Relay Entity to provide connectivity to other Bridge Ports and the attached LANs.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

NOTE—Management port functionality may also be provided by an end station connected to an 802 LAN that is wholly contained within the system that incorporates the Bridge. The absence of external connectivity to the LAN ensures that access to the management port through the bridge’s relay functionality can be assured at all times.

Figure 8-5 illustrates the operation of the Spanning Tree Protocol Entity.

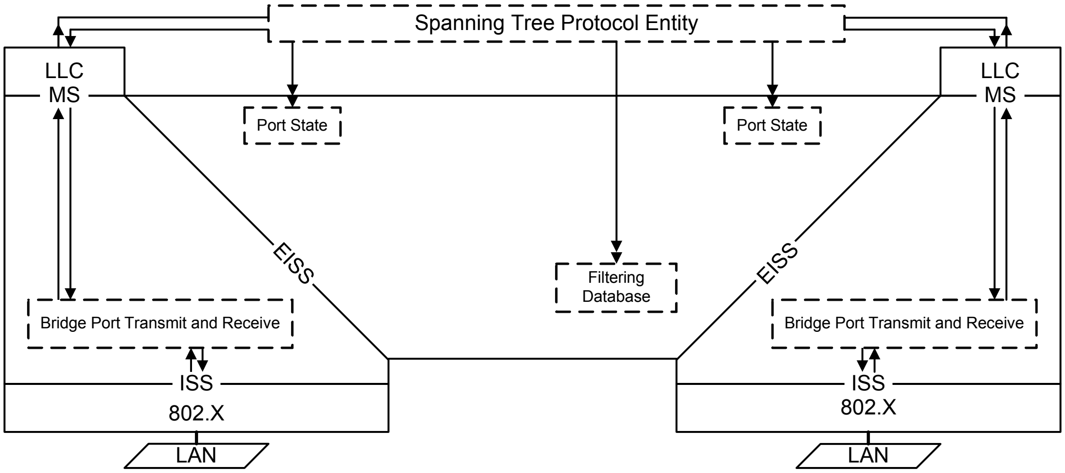
Spanning Tree Protocol Entity
LLC
MS
Port State
Port State
LLC
MS
EISS
Filtering
Database
EISS
Bridge Port Transmit and Receive
Bridge Port Transmit and Receive
ISS
802.X
ISS
802.X
LAN
LAN

Figure 8-5—Operation of Spanning Tree protocol

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Figure 8-6 illustrates the operation of the Generic Attribute Registration Protocol (GARP) Entity (8.11).

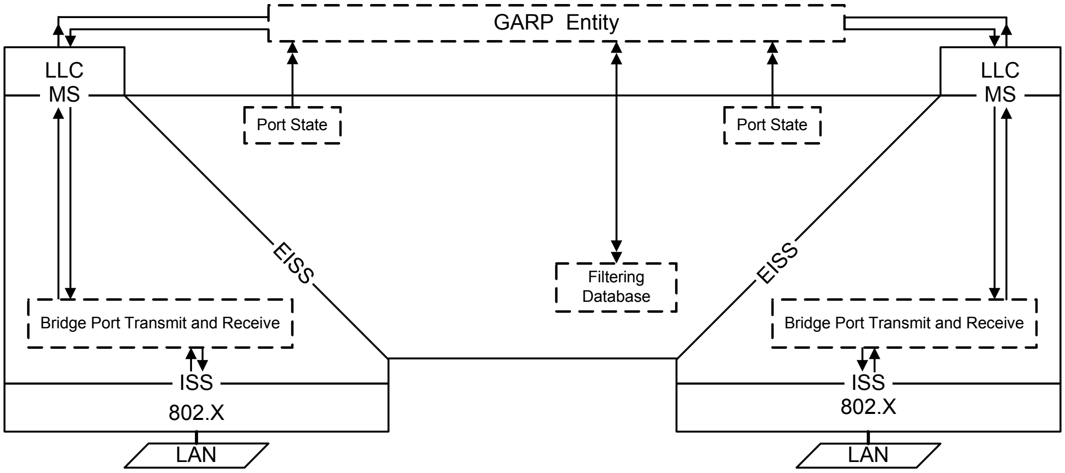

Figure 8-6—Operation of GARP

Figure 8-7 illustrates the reception and transmission by an Entity attached to a Management Port.

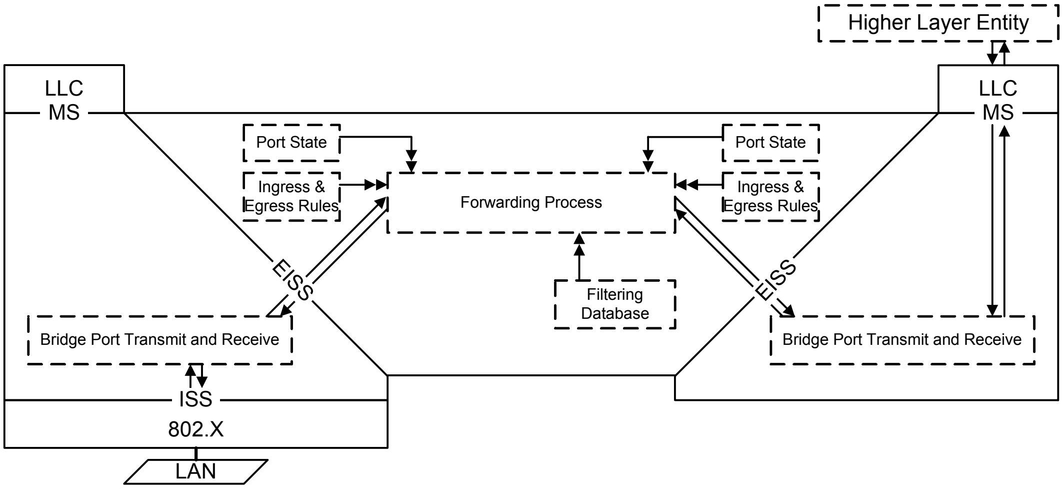

Figure 8-7—Management Port transmission and reception

8.4 Port states and the active topology

Each Bridge Port has an operational Port State that governs whether or not it forwards frames classified as belonging to a given VLAN, and whether or not it learns from their source addresses.

The *active topology* of a Bridged Local Area Network at any time is the set of communication paths formed by interconnecting the LANs and Bridges by the forwarding Ports. The function of the distributed Spanning Tree algorithm and the Rapid Spanning Tree Protocol (RSTP, IEEE Std 802.1D-2003 Clause 17) used by SST Bridges to execute that algorithm, is to construct an active topology that is simply connected relative to communication between any pair of end stations, irrespective of the VLAN classification of frames used. The Multiple Spanning Tree Protocol (MSTP, Clause 13) used by MST Bridges constructs multiple active topologies, each simply and fully connected for frames belonging to any given VLAN. The *forwarding* and *learning* performed by each Bridge Port for each spanning tree is dynamically managed by RSTP or MSTP to prevent temporary loops and reduce excessive traffic in the network while minimizing denial of service following any change in the *physical topology* of the network.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

RSTP constructs a single spanning tree, the Common Spanning Tree (CST), and maintains a single Port State for each Port. MSTP constructs multiple spanning trees, the Common and Internal Spanning Tree (CIST) and additional Multiple Spanning Tree Instances (MSTIs), and maintains a Port State for each spanning tree for each Port. An MST Bridge allocates all frames classified as belonging to a given VLAN to the CIST or to one of the MSTIs.

Any port that is not enabled, i.e. has MAC_Operational (6.4.2) False or has been excluded from the active topology by management setting of the Administrative Bridge Port State to Disabled (14.8.2.2), or has been dynamically excluded from forwarding and learning from MAC frames, is assigned the Port State *Discarding* for all spanning trees. Any Port that has learning enabled but forwarding disabled for frames allocated to a given spanning tree has the Port State *Learning* for that tree, and a Port that both learns and forwards frames the Port State *Forwarding*.

Figure 8-5 illustrates the operation of the Spanning Tree Protocol Entity, which operates the Spanning Tree algorithm and its related protocols, and its modification of Port state information as part of determining the active topology of the network.

Figure 8-3 illustrates the Forwarding Process’s use of the Port State: first, for a Port receiving a frame, to determine whether the received frame is to be relayed through any other Ports; and second, for another Port in order to determine whether the relayed frame is to be forwarded through that particular Port.

Figure 8-4 illustrates the use of the Port state information for a Port receiving a frame, in order to determine whether the station location information is to be incorporated in the Filtering Database.

8.5 Bridge Port Transmit and Receive

The Bridge Port Transmit and Receive process supports the attachment of the Bridge Port to a network. It provides an instance of the EISS for use by the MAC Relay Entity, and one or more instances of the MAC Service for use by Higher Layer Entities, by using an instance of the ISS, as specified in Clause 6.

NOTE 1—The ISS is typically provided directly by the MAC Entity for an IEEE 802 LAN.

Each ISS M_UNITDATA indication with a destination MAC address that is either the individual address of a MAC service access point (MSAP) provided by the Bridge Port or a group address used by the attached LLC Entity, shall cause an MA_UNITDATA indication at that MSAP.

Each EISS EM_UNITDATA request with a destination MAC address that is either the individual address of a MAC service access point (MSAP) provided by the Bridge Port or a group address used by the attached LLC Entity, shall cause an MA_UNITDATA indication at that MSAP.

NOTE 2—The consequence of the above is that frames relayed to a Bridge Port are both submitted to that Port’s MAC Service users and transmitted on the attached LAN (see 8.14.9).

No other frames shall be submitted to MAC Service users.

Each ISS M_UNITDATA indication and EISS EM_UNITDATA request is also processed as specified by Clause 6.7 (Support of the EISS) and shall cause or not cause an EISS EM_UNITDATA indication or an ISS M_UNITDATA request respectively, as required by that specification.

A single Port for a Bridge, known as the Management Port, may provide one or more instances of the MAC Service to higher layer entities without providing an point of attachment to a network through the ISS. If MAC_Enabled parameter for the Port’s EISS MILSAP is set True then the MAC_Operational status parameter shall be True if the Port is a Management Port, and shall take the same value as MAC_Operational for the ISS instance otherwise. If the adminPointToPointMAC parameter for the EISS is set to Auto, then

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

operPointToPointMAC shall be True if the Port is a Management Port and shall take the same value as operPointToPointMAC for the ISS instance otherwise.

The MAC status parameter MAC_Enabled for the EISS MILSAP may be subject to management control, independently of the value of the parameter for the ISS instance.

8.6 The Forwarding Process

Each frame submitted to the MAC Relay Entity shall be forwarded subject to the constituent functions of the Forwarding Process (Figure 8-8).

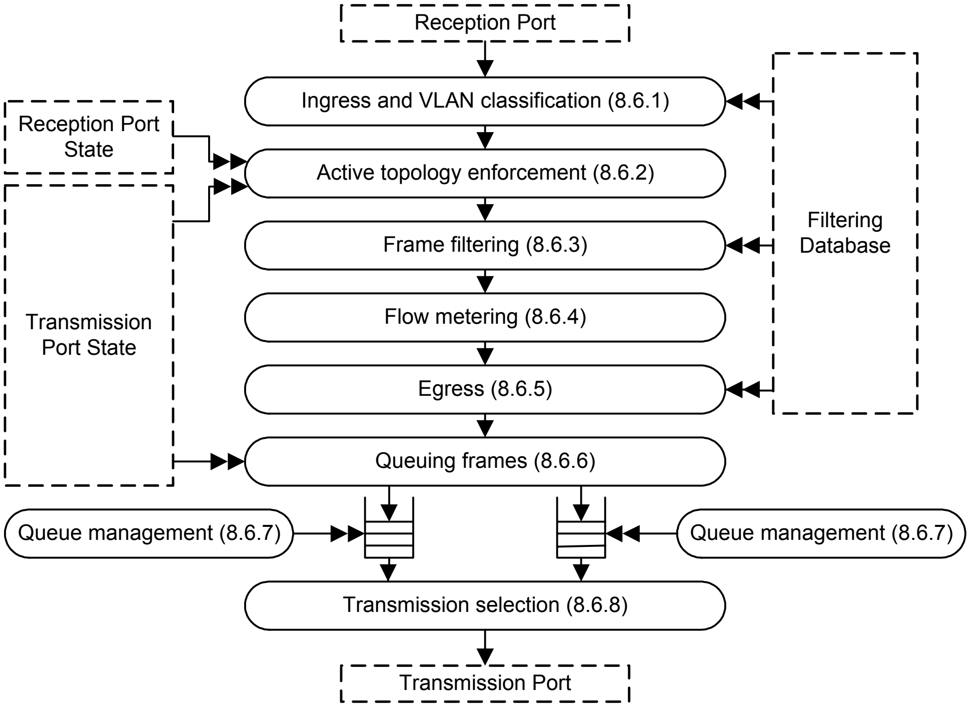

Figure 8-8—Forwarding Process functions

Each function is described in terms of the action taken for a given frame received on a given Port (termed “the reception Port”). The frame can be forwarded for transmission on some Ports (termed “transmission Ports”) and discarded without being transmitted at the other Ports.

Each EISS M_UNITDATA request made by the MAC Relay Entity at a transmission Port shall have the same destination_address, source_address, priority, mac_service_data_unit, and canonical_format_indicator parameters as those accompanying the EISS M_UNITDATA indication corresponding to the received frame. The ingress function performs VLAN classification to determine the requested vlan_identifier, and the egress function determines whether the transmitted frame is to be tagged or untagged.

NOTE 1—IEEE Standard 802.1Q-2003 Edition includes frame formatting and FCS recalculation functions within the Forwarding Process. This standard places those functions below the EISS interface, to allow the specification of additional methods for Bridge Port support of the EISS.

NOTE 2—The Forwarding Process models the Bridge relay function, and does not take into consideration what may occur once frames are passed to the Bridge Port for transmission. Conformant implementations of some media access methods can vary the transmission order in apparent violation of the transmission selection rules when observing frames on the medium. Historic examples include the handling of access_priority in Token-Passing Bus MACs, and the effect of different values for Token Holding Time in FDDI LANs. It may not be possible to test conformance to this standard for

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

some implementations simply by relating observed LAN traffic to the functionality of the forwarding process; tests also have to allow for the (conformant) behavior of the MAC.

Figure 8-3 illustrates the operation of the Forwarding Process in a single instance of frame relay between the Ports of a Bridge with two Ports.

8.6.1 Ingress

The Forwarding Process uses the vlan_identifier parameter conveyed in the EISS data indication to assign a VID to each and every received frame. The frame shall not be forwarded to any other Port or submitted to the Learning Process if it is discarded by any of the following ingress rules.

The frame shall be discarded if the vlan_identifier is FFF, reserved in Table 9-2 for implementation use.

Each Port shall support an Acceptable Frame Types parameter with at least one of the following values

a) *Admit Only VLAN-tagged frames*;
b) *Admit Only Untagged frames*;
c) *Admit All Frames*.

All three values may be supported, if so they shall be configurable using management operations defined in Clause 12, and the default shall be *Admit All Frames*.

A received frame with a vlan_identifer of zero, indicating an untagged or priority tagged frame, shall be discarded if *Admit Only VLAN-tagged frames* is set. Otherwise, if the Bridge and Port support port-and-protocol-based VLAN classification (8.8), the VID Set value for the Protocol Group Identifier selected by matching the received frame with a Protocol Template is assigned. If there is no such VID, or port-and-protocol based classification is not implemented, the value of the PVID parameter for the Port is assigned.

Each Port shall support a PVID parameter for port-based VLAN classification, and may support a VID Set for port-and-protocol-based classification (8.8). The PVID and VID Set shall contain valid VID values (Table 9-2) and may be configured by management. If they have not been explicitly configured, the PVID shall assume the value of the default PVID defined in Table 9-2 and the VID Set shall be empty.

A received frame with a vlan_identifer that is not zero is assigned that VID value.

NOTE 1—The default behavior of a Bridge that supports port-and-protocol-based classification is the same as that of a Bridge that supports only port-based classification, since all the Protocol Group Identifiers in the VID Set for each Port assign the same VID as the PVID.

NOTE 2—A Service-VLAN aware Bridge considers a received frame to be untagged if the initial octets of the MAC user data do not compose a Service VLAN tag header.

A Service-VLAN aware Bridge Port may implement a manageable VID Translation Table that specifies a value to be substituted for each of the possible 4094 VID values assigned or received as specified above. The translation shall occur prior to application of the Enable Ingress Filtering parameter and the other constituent functions of the Forwarding Process specified below. The default configuration of the table shall retain the original VID value. A Customer-VLAN aware Bridge shall not translate VIDs.

Each Port may support an Enable Ingress Filtering parameter. A frame received on a Port that is not in the Member Set (11.9) associated with the VID shall be discarded if this parameter is set. The default value for is reset, i.e. Disable Ingress Filtering, for all Ports. Any Port that supports setting this parameter shall also support resetting it. The parameter may be configured by the management operations defined in Clause 12.

8.6.2 Active topology enforcement

The Forwarding Process allocates each received frame to a spanning tree. If the reception Port State for that spanning tree is Forwarding or Learning, the source address and VID are submitted to the Learning Process. If the reception Port State is Forwarding, each Bridge Port, other than the reception Port, with a Port State of Forwarding for that tree is identified as a potential transmission Port.

An SST Bridge allocates all frames to a single spanning tree, the Common Spanning Tree (CST).

An MST Bridge allocates all frames with a given VID to the CIST or to a Multiple Spanning Tree Instance (MSTI). The allocation can be controlled by configuration of the MST Configuration Table (8.10.1) maintained by the Forwarding Process, subject to constraints (if any) imposed by the allocation of VIDs to FIDs (8.9.7). VIDs allocated to different spanning trees shall also be allocated to different FIDs. VIDs allocated to a given spanning tree may share the same FID.

8.6.3 Frame filtering

The Forwarding Process takes filtering decisions, i.e. reduces the set of potential transmission Ports (8.6.2), for each received frame on the basis of

a) Destination MAC Address;
b) VID;
c) The information contained in the Filtering Database for that MAC Address and VID;
d) The default Group filtering behavior for the potential transmission Port (8.9.6);

in accordance with the definition of the Filtering Database entry types (8.9.1, 8.9.3, and 8.9.4). The required behavior is summarized in 8.9.6, 8.9.8, Table 8-1, Table 8-2, and Table 8-3.

Each of the Reserved MAC Addresses specified in Table 8-1 shall be permanently configured in the Filtering Database in Customer-VLAN aware Bridges. Each of the Reserved MAC Addresses specified in Table 8-2 shall be permanently configured in the Filtering Database in Service-VLAN aware Bridges. The Filtering Database Entries for Reserved MAC Addresses shall specify filtering for all Bridge Ports and all VLANs. Management shall not provide the capability to modify or remove entries for Reserved MAC Addresses.

8.6.4 Flow metering

The Forwarding Process may apply a flow meter to frames received on a Bridge Port for which the set of potential transmission Ports (as determined by 8.6.1, 8.6.2, and 8.6.3 above) is not empty. The flow meter can set, but not clear, the drop_eligible parameter associated with each frame and may also discard frames on the basis of the following parameters for each received frame and previously received frames, and the timed elapsed since those frames were received:

a) Destination MAC Address
b) VID
c) Priority
d) The received value of the drop_eligible parameter
e) The mac_service_data_unit size

The flow meter shall not base its decision on the parameters of frames received on other Bridge Ports, or on any other parameters of those Ports.

NOTE—The flow meter described here can encompass a number of meters, each with a simpler specification. However given the breadth of implementation choice permitted, further structuring to specify, for example, that frames can bypass a meter or are subject only to one of a number of meters provides no additional information.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Table 8-1—Customer-VLAN aware Bridge Reserved addresses

Assignment	Value
Bridge Group Address	01-80-C2-00-00-00
IEEE Std 802.3, 1998 Edition, Full Duplex PAUSE operation	01-80-C2-00-00-01
IEEE Std. 802.3ad Slow_Protocols_Multicast address	01-80-C2-00-00-02
IEEE Std. 802.1X PAE address	01-80-C2-00-00-03
Reserved for future standardization - media access method specific	01-80-C2-00-00-04
Reserved for future standardization - media access method specific	01-80-C2-00-00-05
Reserved for future standardization - Service & Customer-VLAN aware Bridge specific	01-80-C2-00-00-06
Reserved for future standardization - Service & Customer-VLAN aware Bridge specific	01-80-C2-00-00-07
Reserved not allocated - example Service-VLAN aware Bridge Group Address	01-80-C2-00-00-08
Reserved for future standardization - Service & Customer-VLAN aware Bridge specific	01-80-C2-00-00-09
Reserved for future standardization - Service & Customer-VLAN aware Bridge specific	01-80-C2-00-00-0A
Reserved for future standardization - Service-VLAN aware Bridge specific	01-80-C2-00-00-0B
Reserved for future standardization - Service-VLAN aware Bridge specific	01-80-C2-00-00-0C
Reserved not allocated - example Service VLAN GVRP Address	01-80-C2-00-00-0D
Reserved for future standardization - Customer-VLAN aware Bridge specific	01-80-C2-00-00-0E
Reserved for future standardization - Customer-VLAN aware Bridge specific	01-80-C2-00-00-0F

Table 8-2—Service-VLAN aware Bridge Reserved addresses

Assignment	Value
IEEE Std 802.3, 1998 Edition, Full Duplex PAUSE operation	01-80-C2-00-00-01
IEEE Std. 802.3ad Slow_Protocols_Multicast address	01-80-C2-00-00-02
IEEE Std. 802.1X PAE address	01-80-C2-00-00-03
Reserved for future standardization - media access method specific	01-80-C2-00-00-04
Reserved for future standardization - media access method specific	01-80-C2-00-00-05
Reserved for future standardization - Service & Customer-VLAN aware Bridge specific	01-80-C2-00-00-06
Reserved for future standardization - Service & Customer-VLAN aware Bridge specific	01-80-C2-00-00-07
Reserved not allocated - example Service-VLAN aware Bridge Group Address	01-80-C2-00-00-08
Reserved for future standardization - Service & Customer-VLAN aware Bridge specific	01-80-C2-00-00-09
Reserved for future standardization - Service & Customer-VLAN aware Bridge specific	01-80-C2-00-00-0A

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

8.6.5 Egress

The Forwarding Process shall queue each received frame to each of the potential transmission Ports (8.6.2, 8.6.3) that is present in the Member Set (8.9.9) for the frame's VID.

The include_tag parameter to be used in the EM_UNITDATA request to transmit the frame (8.6.8) will be False if the transmission Port is present in the untagged set (8.9.9), and True otherwise.

NOTE 1—The Forwarding Process is modelled as receiving a frame as the parameters of a data indication and transmitting through supplying the parameters of a data request. Queueing a frame awaiting transmission amounts to placing the parameters of a data request on an outbound queue.

NOTE 2—As all incoming frames, including priority-tagged frames, are classified as belonging to a VLAN, the transmitting Port transmits VLAN-tagged frames or untagged frames. Hence, a station sending a priority-tagged frame via a VLAN Bridge will receive a response that is either VLAN-tagged or untagged.

8.6.6 Queuing frames

The Forwarding Process provides storage for queued frames, awaiting an opportunity to submit these for transmission. The order of frames received on the same Bridge Port shall be preserved for

a) unicast frames with a given VID, priority, and destination address and source address combination;
b) multicast frames with a given VID, priority, and destination address.

The Forwarding Process provides one or more queues for a given Bridge Port, each corresponding to a distinct traffic class. Each frame is mapped to a traffic class using the Traffic Class Table for the Port and the frame's priority. Traffic class tables may be managed. Table 8-3 shows the recommended mapping for the number of classes implemented. Up to eight traffic classes may be supported, allowing separate queues for each priority.

Table 8-3—Recommended priority to traffic class mappings

		Number of Available Traffic Classes							
		1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
Priority	**0 (Default)**	0	0	0	0	0	1	1	1
	1	0	0	0	0	0	0	0	0
	2	0	0	0	1	1	2	2	2
	3	0	0	0	1	1	2	3	3
	4	0	1	1	2	2	3	4	4
	5	0	1	1	2	2	3	4	5
	6	0	1	2	3	3	4	5	6
	7	0	1	2	3	4	5	6	7

NOTE—The rationale for these mappings is discussed in Annex G (informative).

NOTE—Different numbers of traffic classes may be implemented for different Ports. Ports with media access methods that support a single transmission priority, such as CSMA/CD, can support more than one traffic class.

Copyright © 2004 IEEE. All rights reserved.

8.6.7 Queue management

A frame queued for transmission on a Port shall be removed from that queue

a) following a transmit data request. No further attempt shall be made to transmit the frame on that Port even if the transmission is known to have failed.
b) if that is necessary to ensure that the maximum bridge transit delay (6.3.6) will not be exceeded at the time at which the frame would subsequently be transmitted.
c) if the associated Port leaves the Forwarding state.

The frame can be removed from the queue, and not subsequently transmitted

d) if the time for which buffering is guaranteed has been exceeded for that frame
e) by a queue management algorithm that attempts to improve the QoS provided by deterministically or probabilistically manging the queue depth based on the current and past queue depths.

Removal of a frame from a queue for any particular Port does not affect queuing of that frame for transmission on any other Port. The probability of removing a frame with drop_eligible set shall not be less than that of removing a frame with drop_eligible clear, all other conditions being equal.

NOTE—Applicable queue management algorithms include RED (random early discard), WRED (weighted random early discard) <<provide reference in Bibilography>>. If the Bridge supports drop precedence, i.e. is capable of decoding or encoding the drop_eligible parameter from or to the PCP field of VLAN TAGs (6.7.3), the algorithm should exhibit a higher probability of dropping frames with drop_eligible set.

8.6.8 Transmission selection

The following algorithm shall be supported by all Bridges as the default algorithm for selecting frames for transmission:

a) For each Port, frames are selected for transmission on the basis of the traffic classes that the Port supports. For a given supported value of traffic class, frames are selected from the corresponding queue for transmission only if all queues corresponding to numerically higher values of traffic class supported by the Port are empty at the time of selection;
b) For a given queue, the order in which frames are selected for transmission shall maintain the ordering requirement specified in 8.6.6.

Additional algorithms, selectable by management means, may be supported as an implementation option so long as the requirements of 8.6.6 are met.

8.7 The Learning Process

The Learning Process receives the source MAC Addresses and VIDs of received frames from the Forwarding Process, following active topology enforcement. It shall create or update a Dynamic Filtering Entry (8.9.3) that specifies the reception Port for the frame’s source address and VID, if and only if the source address is an Individual Address, i.e. is not a Group Address, the resulting number of entries would not exceed the capacity of the Filtering Database, and the filtering utility criteria for the receiving Bridge Port are met, as specified below.

If the Filtering Database is already filled to capacity, but a new entry would otherwise be made, then an existing entry may be removed to make room for the new entry.

The purpose of filtering utility criteria is to reduce the capacity requirements of the Filtering Database and to reduce the time for which service can be denied (6.3.1) by retaining filtering information learnt prior to a change in the physical topology of the network. Filtering utility criteria shall be applied to the learning and retention of information for each Filtering Identifier (FID) (8.9.7). Enhanced filtering utility criteria may be implemented for any Bridge Port as specified below (8.7.2), if implemented both the default (8.7.2) and the enhanced criteria shall be selectable by management.

8.7.1 Default filtering utility criteria

The default for a VLAN-aware Bridge configured as either a Customer-VLAN aware or a Service-VLAN aware Bridge is that the Member Set (8.9.9) for the frame's VID includes at least one Port.

NOTE—If the Member Set for a given VID is empty, that VLAN is not currently active, and the Bridge will filter all frames destined for that VLAN, regardless of their destination address.

8.7.2 Enhanced filtering utility criteria

The enhanced criteria are satisfied if at least one VLAN that uses the FID includes the reception Port and at least one other Port with a Port State of Learning or Forwarding in its Member Set, and:

a) operPointToPointMAC is false for the reception Port; or
b) ingress to the VLAN is permitted through a third Port.

NOTE —The third port can, but is not required to be in the Member Set.

Figure 8-5 illustrates the operation of the Learning Process in the inclusion of station location information carried by a single frame, received on one of the Ports of a Bridge, in the Filtering Database.

8.8 Protocol VLAN classification

The Forwarding Process uses the following procedures on ingress (8.6.1) to assign a VID to each frame received from a Port that implements port-and-protocol-based VLAN classification.

The **detagged_frame_type** parameter indicates the frame format. Its 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 802-2001 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, 1997 Edition, 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.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

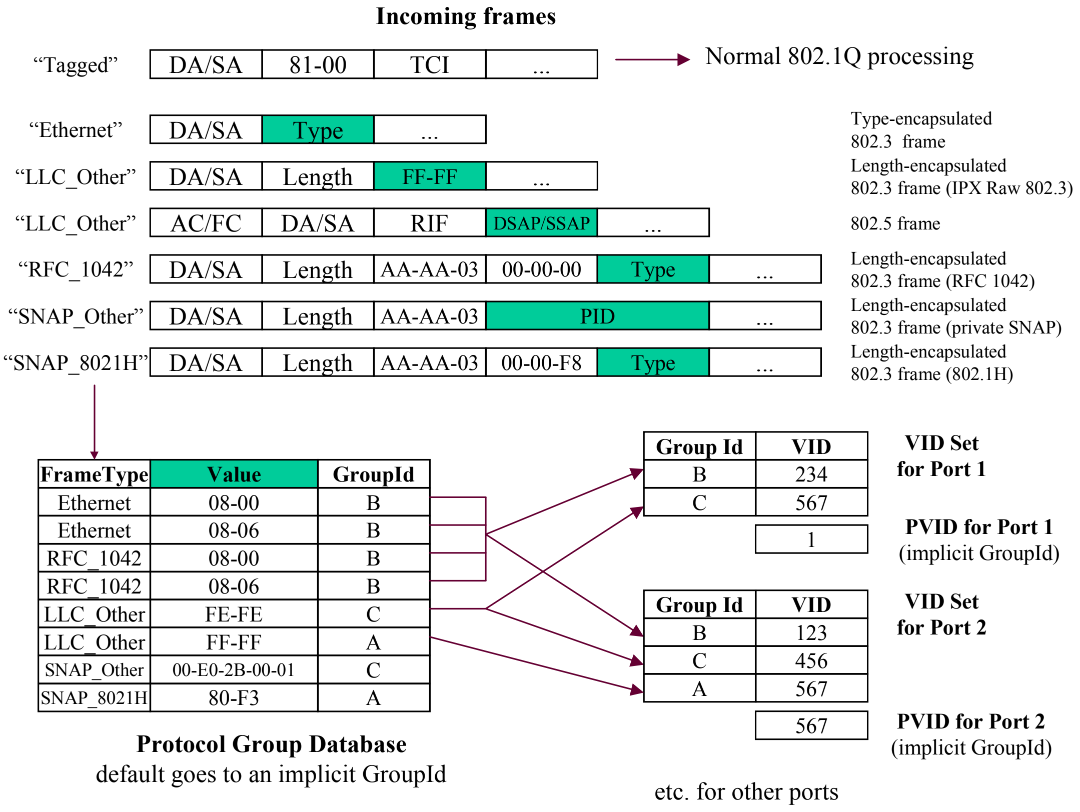
Incoming frames
"Tagged" DA/SA 81-00 TCI ...
Normal 802.1Q processing
"Ethernet" DA/SA Type ...
Type-encapsulated 802.3 frame
"LLC_Other" DA/SA Length FF-FF ...
Length-encapsulated 802.3 frame (IPX Raw 802.3)
"LLC_Other" AC/FC DA/SA RIF DSAP/SSAP ...
802.5 frame
"RFC_1042" DA/SA Length AA-AA-03 00-00-00 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)
Protocol Group Database
default goes to an implicit GroupId
VID Set for Port 1
PVID for Port 1 (implicit GroupId)
1
VID Set for Port 2
PVID for Port 2 (implicit GroupId)
567
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

Group Id	VID
B	234
C	567

Group Id	VID
B	123
C	456
A	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 8-9—Example of operation of port-and-protocol based classification

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, 1998 Edition, 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.

8.8.1 Protocol Templates

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.

8.8.2 Protocol Group Identifiers

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 8-9, 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 8.8.3. 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.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

8.8.3 Protocol Group Database

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 (8.8.2). 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.

8.9 The Filtering Database

This amendment does not change this clause.

<<Changes are required to allow wild card VIDs in Filtering Entries. As discussed at the September 2003 interim, there is currently no way that the instruction to put entries corresponding to the Reserved Addresses in the Filtering Database can be followed, without adding 4094 entries per address.>>

The Filtering Database supports queries by the Forwarding Process as to whether frames received by the Forwarding Process, with given values of destination MAC Address parameter and VID, are to be forwarded through a given potential transmission Port (8.6.2 and 8.6.5). It contains filtering information in the form of filtering entries that are either

a) Static, and explicitly configured by management action; or
b) Dynamic, and automatically entered into the Filtering Database by the normal operation of the bridge and the protocols it supports.

Two entry types are used to represent static filtering information. The Static Filtering Entry represents static information in the Filtering Database for individual and for group MAC Addresses. It allows administrative control of

c) Forwarding of frames with particular destination addresses; and
d) The inclusion in the Filtering Database of dynamic filtering information associated with Extended Filtering Services, and use of this information.

The Filtering Database shall contain entries of the Static Filtering Entry type.

The Static VLAN Registration Entry represents all static information in the Filtering Database for VLANs. It allows administrative control of

e) Forwarding of frames with particular VIDs;
f) The inclusion/removal of tag headers in forwarded frames; and
g) The inclusion in the Filtering Database of dynamic VLAN membership information, and use of this information.

The Filtering Database may contain entries of the Static VLAN Registration Entry type.

Static filtering information is added to, modified, and removed from the Filtering Database only under explicit management control. It shall not be automatically removed by any ageing mechanism. Management of static filtering information may be carried out by use of the remote management capability provided by Bridge Management (8.13) using the operations specified in Clause 12.

Three entry types are used to represent dynamic filtering information:

h) Dynamic Filtering Entries are used to specify the Ports on which individual MAC Addresses have been learned. They are created and updated by the Learning Process (8.7), and are subject to ageing and removal by the Filtering Database.
i) Group Registration Entries support the registration of group MAC Addresses. They are created, updated, and removed by the GMRP protocol in support of Extended Filtering Services (8.9.4; 6.6.5 of IEEE Std 802.1D, 1998 Edition; Clause 10 of IEEE Std 802.1D, 1998 Edition), subject to the state of the Restricted_Group_Registration management control (10.3.2.3 of IEEE Std 802.1D, 1998 Edition). If the value of this control is TRUE, then the creation of a Group Registration Entry is not permitted unless a Static Filtering Entry exists that permits dynamic registration for the Group concerned.
j) Dynamic VLAN Registration Entries are used to specify the Ports on which VLAN membership has been dynamically registered. They are created, updated, and removed by the GVRP protocol, in support of automatic VLAN membership configuration (Clause 11), subject to the state of the Restricted_VLAN_Registration management control (11.2.3.2.3). If the value of this control is TRUE, then the creation of a Dynamic VLAN Registration Entry is not permitted unless a Static VLAN Registration Entry exists that permits dynamic registration for the VLAN concerned.

Static Filtering Entries and Group Registration Entries comprise

k) A MAC Address specification;
l) A VLAN Identifier (VID);
m) A Port Map, with a control element for each outbound Port to specify filtering for that MAC Address specification and VID.

Dynamic Filtering Entries comprise

n) A MAC Address specification;
o) A locally significant Filtering Identifier (FID; see 8.9.7);
p) A Port Map, with a control element for each outbound Port to specify filtering for that MAC Address specification in the VLAN(s) allocated to that FID.

Static and Dynamic VLAN Registration Entries comprise

q) A VLAN Identifier;
r) A Port Map, with a control element for each outbound Port to specify filtering for the VLAN.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Dynamic filtering information may be read by use of the remote management capability provided by Bridge Management (8.13) using the operations specified in Clause 12.

The Filtering Services supported by a Bridge (Basic and Extended Filtering Services) determine the default behavior of the Bridge with respect to the forwarding of frames destined for group MAC Addresses. In Bridges that support Extended Filtering Services, the default forwarding behavior for group MAC Addresses, for each Port, and for each VID, can be configured both statically and dynamically by means of Static Filtering Entries and/or Group Registration Entries that can carry the following MAC Address specifications:

s) All Group Addresses, for which no more specific Static Filtering Entry exists;
t) All Unregistered Group Addresses (i.e., all group MAC Addresses for which no Group Registration Entry exists), for which no more specific Static Filtering Entry exists.

NOTE 1—The All Group Addresses specification s) above, when used in a Static Filtering Entry with an appropriate control specification, provides the ability to configure a Bridge that supports Extended Filtering Services to behave as a Bridge that supports only Basic Filtering Services on some or all of its Ports. This might be done for the following reasons:

— The Ports concerned serve “legacy” devices that wish to receive multicast traffic, but are unable to register Group membership;

— The Ports concerned serve devices that need to receive all multicast traffic, such as routers or diagnostic devices.

The Filtering Database shall support the creation, updating and removal of Dynamic Filtering Entries by the Learning Process (8.7). In Bridges that support Extended Filtering Services, the Filtering Database shall support the creation, updating, and removal of Group Registration Entries by GMRP (Clause 10 of IEEE Std 802.1D, 1998 Edition).

Figure 8-3 illustrates use of the Filtering Database by the Forwarding Process in a single instance of frame relay between the Ports of a Bridge with two Ports.

Figure 8-4 illustrates the creation or update of a dynamic entry in the Filtering Database by the Learning Process. The entries in the Filtering Database allow MAC Address information to be learned independently for each VLAN or set of VLANs, by relating a MAC Address to the VLAN or set of VLANs on which that address was learned. This has the effect of creating independent Filtering Databases for each VLAN or set of VLANs that is present in the Bridged LAN.

NOTE 2—This standard specifies a single Filtering Database that contains all Filtering Database entries; however, the inclusion of VIDs and FIDs in the filtering entries effectively provides distinct IEEE Std 802.1D-style Filtering Databases per VLAN or set of VLANs.

NOTE 3—The ability to create VLAN-dependent Filtering Database entries allows a VLAN Bridge to support

— Multiple end stations with the same individual MAC Address residing on different VLANs;

— End stations with multiple interfaces, each using the same individual MAC Address,
as long as not more than one end station or interface that uses a given MAC Address resides in a given VLAN.

Figure 8-5 illustrates the operation of the Spanning Tree Protocol Entity (8.11), which operates the Spanning Tree Algorithm and Protocol, and its notification of the Filtering Database of changes in active topology signaled by that protocol.

There are no standardized constraints on the size of the Filtering Database in an implementation for which conformance to this standard is claimed. The PICS Proforma in Annex A requires the following to be specified for a given implementation:

u) The total number of entries (Static Filtering Entries, Dynamic Filtering Entries, Group Registration Entries, Static VLAN Registration Entries, and Dynamic VLAN Registration Entries) that the implementation of the Filtering Database can support, and
v) Of that total number, the total number of VLAN Registration Entries (static and dynamic) that the Filtering Database can support.

8.9.1 Static Filtering Entries

A Static Filtering Entry specifies

a) A MAC Address specification, comprising
   1) An Individual MAC Address; or
   2) A group MAC Address; or
   3) All Group Addresses, for which no more specific Static Filtering Entry exists; or
   4) All Unregistered Group Addresses, for which no more specific Static Filtering Entry exists.
b) The VID of the VLAN to which the static filtering information applies;
c) A Port Map, containing a control element for each outbound Port, specifying that a frame with a destination MAC Address and VID that meets this specification is to be
   1) Forwarded, independently of any dynamic filtering information held by the Filtering Database; or
   2) Filtered, independently of any dynamic filtering information; or
   3) Forwarded or filtered on the basis of dynamic filtering information, or on the basis of the default Group filtering behavior for the outbound Port (8.9.6) if no dynamic filtering information is present specifically for the MAC Address.

All Bridges shall have the capability to support the first two values for the MAC Address specification, and all three values for each control element for all Static Filtering Entries (i.e., shall have the capability to support a1, a2, c1, c2, and c3 above).

A Bridge that supports Extended Filtering Services shall have the capability to support all four values for the MAC Address specification and all three control element values for all Static Filtering Entries.

For a given MAC Address specification, a separate Static Filtering Entry with a distinct Port Map may be created for each VLAN from which frames are received by the Forwarding Process.

In addition to controlling the forwarding of frames, Static Filtering Entries for group MAC Addresses provide the Registrar Administrative Control values for the GMRP protocol (Clauses 10, 12, and 12.9.1 of IEEE Std 802.1D, 1998 Edition). Static configuration of forwarding of specific group addressed frames to an outbound port indicates Registration Fixed on that port: a desire to receive frames addressed to that Group even in the absence of dynamic information. Static configuration of filtering of frames that might otherwise be sent to an outbound port indicates Registration Forbidden. The absence of a Static Filtering Entry for the group address, or the configuration of forwarding or filtering on the basis of dynamic filtering information, indicates Normal Registration.

8.9.2 Static VLAN Registration Entries

A Static VLAN Registration Entry specifies

a) The VID of the VLAN to which the static filtering information applies;
b) A Port Map, consisting of a control element for each outbound Port, specifying
   1) The Registrar Administrative Control values for the GVRP protocol (Clause 11) for the VLAN specified. In addition to providing control over the operation of GVRP, these values can also directly affect the forwarding behavior of the Bridge, as described in 8.9.9. The values that can be represented are

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

i) Registration Fixed; or
ii) Registration Forbidden; or
iii) Normal Registration.

2) Whether frames destined for the VLAN specified are to be VLAN-tagged or untagged when forwarded through this Port.

All Bridges shall be capable of supporting all values for each control element for all Static VLAN Registration Entries; however, the ability to support more than one untagged VLAN on egress on any given Port is optional (see 5.1 and 5.2).

NOTE—In other words, it shall be possible to configure any VLAN as untagged on egress, but it is an implementation option as to whether only a single untagged VLAN per Port on egress is supported, or whether multiple untagged VLANs per Port on egress are supported.

A separate Static VLAN Registration Entry with a distinct Port Map may be created for each VLAN from which frames are received by the Forwarding Process.

8.9.3 Dynamic Filtering Entries

A Dynamic Filtering Entry specifies

a) An individual MAC Address;
b) The FID, an identifier assigned by the MAC Bridge (8.9.7) to identify a set of VIDs for which no more than one Dynamic Filtering Entry can exist for any individual MAC Address;

NOTE 1—An FID identifies a set of VLANs among which *Shared VLAN Learning* (3.26) takes place. Any pair of FIDs identifies two sets of VLANs between which *Independent VLAN Learning* (3.8) takes place. The allocation of FIDs by a Bridge is described in 8.9.7.

c) A Port Map that specifies forwarding of frames destined for that MAC Address and FID to a single Port.

NOTE 2—This is equivalent to specifying a single port number; hence, this specification is directly equivalent to the specification of dynamic entries in ISO/IEC 10038: 1993.

Dynamic Filtering Entries are created and updated by the Learning Process (8.7). They shall be automatically removed after a specified time, the Ageing Time, has elapsed since the entry was created or last updated. No more than one Dynamic Filtering Entry shall be created in the Filtering Database for a given combination of MAC Address and FID.

Dynamic Filtering Entries cannot be created or updated by management.

NOTE 3—Dynamic Filtering Entries may be read by management (Clause 12). The FID is represented in the management Read operation by any one of the VIDs that it represents. For a given VID, the set of VIDs that share the same FID may also be determined by management.

The ageing out of Dynamic Filtering Entries ensures that end stations that have been moved to a different part of the Bridged LAN will not be permanently prevented from receiving frames. It also takes account of changes in the active topology of the Bridged LAN that can cause end stations to appear to move from the point of view of the bridge; i.e., the path to those end stations subsequently lies through a different Bridge Port.

The Ageing Time may be set by management (Clause 12). A range of applicable values and a recommended default is specified in Table 8-1; this is suggested to remove the need for explicit configuration in most cases. If the value of Ageing Time can be set by management, the Bridge shall have the capability to use values in the range specified, with a granularity of 1 s.

Table 8-1—Ageing time parameter value

Parameter	Recommended default value	Range
Ageing time	300.0 s	10.0–1 000 000.0 s

NOTE 4—The granularity is specified in order to establish a common basis for the granularity expressed in the management operations defined in Clause 12, not to constrain the granularity of the actual timer supported by a conformant implementation. If the implementation supports a granularity other than 1 s, then it is possible that the value read back by management following a Set operation will not match the actual value expressed in the Set.

The Spanning Tree Algorithm and Protocol specified in Clause 8 of IEEE Std 802.1D, 1998 Edition includes a procedure for notifying all Bridges in the Bridged LAN of topology change. It specifies a short value for the Ageing Timer, to be enforced for a period after any topology change (8.3.5 of IEEE Std 802.1D, 1998 Edition). While the topology is not changing, this procedure allows normal ageing to accommodate extended periods during which addressed end stations do not generate frames themselves, perhaps through being powered down, without sacrificing the ability of the Bridged LAN to continue to provide service after automatic configuration.

8.9.4 Group Registration Entries

A Group Registration Entry specifies

a) A MAC Address specification, comprising
   1) A group MAC Address; or
   2) All Group Addresses, for which no more specific Static Filtering Entry exists; or
   3) All Unregistered Group Addresses, for which no more specific Static Filtering Entry exists.
b) The VID of the VLAN in which the dynamic filtering information was registered;
c) A Port Map, consisting of a control element for each outbound Port, which specifies forwarding (Registered) or filtering (Not registered) of frames destined to the MAC Address and VID.

Group Registration Entries are created, modified and deleted by the operation of GMRP (Clause 10 of IEEE Std 802.1D, 1998 Edition, as modified by Clause 10 of this standard). No more than one Group Registration Entry shall be created in the Filtering Database for a given combination of MAC Address specification and VID.

NOTE—It is possible to have a Static Filtering Entry which has values of Forward or Filter on some or all Ports that mask the dynamic values held in a corresponding Group Registration Entry. The values in the Group Registration Entry will continue to be updated by GMRP; hence, subsequent modification of that entry to allow the use of dynamic filtering information on one or more Ports immediately activates the true GMRP registration state that was hitherto masked by the static information.

The creation of Group Registration Entries is subject to the Restricted_Group_Registration management control (10.3.2.3 of IEEE Std 802.1D, 1998 Edition). If the value of this control is TRUE, a dynamic entry for a given Group may only be created if a Static Filtering Entry already exists for that Group, in which the Registrar Administrative Control value is Normal Registration.

8.9.5 Dynamic VLAN Registration Entries

A Dynamic VLAN Registration Entry specifies

a) The VID of the VLAN to which the dynamic filtering information applies;

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

b) A Port Map with a control element for each outbound Port specifying whether the VLAN is registered on that Port.

A separate Dynamic VLAN Registration Entry with a distinct Port Map may be created for each VLAN from which frames are received by the Forwarding Process.

The creation of Dynamic VLAN Registration Entries is subject to the Restricted_VLAN_Registration management control (11.2.3.2.3). If the value of this control is TRUE, a dynamic entry for a given VLAN may only be created if a Static VLAN Registration Entry already exists for that VLAN, in which the Registrar Administrative Control value is Normal Registration.

8.9.6 Default Group filtering behavior

Forwarding and filtering of group-addressed frames may be managed by specifying defaults for each VLAN and outbound Port. The behavior of each of these defaults, as modified by the control elements of more explicit Filtering Database entries applicable to a given frame’s MAC Address, VLAN classification, and outbound Port is as follows:

NOTE 1—As stated in 8.9.1, for a given MAC Address there may be separate Static Filtering Entries with a distinct Port Map for each VLAN.

a) *Forward All Groups.* The frame is forwarded, unless an explicit Static Filtering Entry specifies filtering independent of any dynamic filtering information.
b) *Forward Unregistered Groups.* The frame is forwarded, unless
    1) An explicit Static Filtering Entry specifies filtering independent of any dynamic filtering information; or
    2) An explicit Static Filtering Entry specifies forwarding or filtering on the basis of dynamic filtering information, and an applicable explicit Group Registration Entry exists specifying filtering; or
    3) An applicable explicit Static Filtering Entry does not exist, but an applicable Group Registration entry specifies filtering.
c) *Filter Unregistered Groups.* The frame is filtered unless
    1) An explicit Static Filtering Entry specifies forwarding independent of any dynamic filtering information; or
    2) An explicit Static Filtering Entry specifies forwarding or filtering on the basis of dynamic filtering information, and an applicable explicit Group Registration Entry exists specifying forwarding; or
    3) An applicable explicit Static Filtering Entry does not exist, but an applicable Group Registration entry specifies forwarding.

In Bridges that support only Basic Filtering Services, the default Group filtering behavior is Forward All Groups for all Ports of the Bridge, for all VLANs.

NOTE 2—Forward All Groups corresponds directly to the behavior specified in ISO/IEC 10038: 1993 when forwarding group MAC Addressed frames for which no static filtering information exists in the Filtering Database. Forward All Groups makes use of information contained in Static Filtering Entries for specific group MAC Addresses, but overrides any information contained in Group Registration Entries. Forward Unregistered Groups is analogous to the forwarding behavior of a Bridge with respect to individual MAC Addresses. If there is no static or dynamic information for a specific group MAC Address, then the frame is forwarded; otherwise, the frame is forwarded in accordance with the statically configured or dynamically learned information.

In Bridges that support Extended Filtering Services, the default Group filtering behavior for each outbound Port for each VLAN is determined by the following information contained in the Filtering Database:

d) Any Static Filtering Entries applicable to that VLAN with a MAC Address specification of All Group Addresses or All Unregistered Group Addresses;

e) Any Group Registration Entries applicable to that VLAN with a MAC Address specification of All Group Addresses or All Unregistered Group Addresses.

The means whereby this information determines the default Group filtering behavior is specified in 8.9.8, Table 8-2, and Table 8-3.

NOTE 3—The result is that the default Group filtering behavior for each VLAN can be configured for each Port of the Bridge via Static Filtering Entries, determined dynamically via Group Registration Entries created/updated by GMRP (Clause 10), or both. For example, in the absence of any static or dynamic information in the Filtering Database for All Group Addresses or All Unregistered Group Addresses, the default Group filtering behavior will be Filter Unregistered Groups on all Ports, for all VLANs. Subsequently, the creation of a Dynamic Group Registration Entry for All Unregistered Group Addresses indicating “Registered” for a given VLAN on a given Port would cause that Port to exhibit Forward Unregistered Groups behavior for that VLAN. Similarly, creating a Static Filtering Entry for All Group Addresses indicating “Registration Fixed” on a given Port for that VLAN would cause that Port to exhibit Forward All Groups behavior.

Hence, by using appropriate combinations of “Registration Fixed,” “Registration Forbidden,” and “Normal Registration” in the Port Maps of Static Filtering Entries for the All Group Addresses and All Unregistered Group Addresses address specifications, it is possible, for a given Port and VLAN, to

— Fix the default Group filtering behavior to be just one of the three behaviors described above; or

— Restrict the choice of behaviors to a subset of the three, and allow GMRP registrations (or their absence) to determine the final choice; or

— Allow any one of the three behaviors to be adopted, in accordance with any registrations received via GMRP.

8.9.7 Allocation of VIDs to FIDs

The allocation of VIDs to FIDs within a Bridge determines how learned individual MAC Address information is used in forwarding/filtering decisions within a Bridge; whether such learned information is confined to individual VLANs, shared among all VLANs, or confined to specific sets of VLANs.

The allocation of VIDs to FIDs is determined on the basis of

a) The set of *VLAN Learning Constraints* that have been configured into the Bridge (by means of the management operations defined in Clause 12);

b) Any fixed mappings of VIDs to FIDs that may have been configured into the Bridge (by means of the management operations defined in Clause 12);

c) The *set of active VLANs* (i.e., those VLANs on whose behalf the Bridge may be called upon to forward frames). A VLAN is active if either of the following is true:
   1) The VLAN’s Member set (8.9.9) contains one Port that is in a forwarding state, and at least one other Port of the Bridge is both in a forwarding state and has Ingress Filtering (8.6.1) disabled;
   2) The VLAN’s Member set contains two or more Ports that are in a forwarding state.

d) The capabilities of the Bridge with respect to the number of FIDs that it can support, and the number of VIDs that can be allocated to each FID.

A VLAN Bridge shall support a minimum of one FID, and may support up to 4094 FIDs. For the purposes of the management operations, FIDs are numbered from 1 through N, where N is the maximum number of FIDs supported by the implementation.

A VLAN Bridge shall support the ability to allocate at least one VID to each FID, and may support the ability to allocate more than one VID to each FID.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

The number of VLAN Learning Constraints supported by a VLAN Bridge is an implementation option.

NOTE—In an SVL/IVL Bridge (3.28), a number of FIDs are supported, and one or more VID can be mapped to each FID. In an SVL Bridge (3.27), a single FID is supported, and all VIDs are mapped to that FID. In an IVL Bridge (3.9), a number of FIDs are supported, and only one VID can be mapped to each FID.

An MST Bridge shall support the ability to allocate at least one FID to each spanning tree, and may support the ability to allocate more than one FID to each spanning tree

NOTE—In other words, the number of FIDs supported by the Bridge must be greater than or equal to the number of spanning trees supported by the Bridge.

An MST Bridge shall ensure that the maximum supported numbers of FIDs and VLANs can be associated unambiguously. This requires either 1) a number of fixed VID to FID allocations at least equal to the maximum number of VLANs supported; or 2) one I Constraint entry per FID supported and one S Constraint entry per MSTI supported, or both. (8.9.7.1).

8.9.7.1 Fixed and dynamic VID to FID allocations

A Bridge may support the ability to define fixed allocations of specific VIDs to specific FIDs, via an allocation table that may be read and modified by means of the management operations defined in Clause 12. For each VID supported by the implementation, the allocation table indicates one of the following:

a) The VID is currently not allocated to any FID; or
b) A fixed allocation has been defined (via management), allocating this VID to a specific FID; or
c) A dynamic allocation has been defined (as a result of applying the VLAN Learning Constraints), allocating this VID to a specific FID.

For any VID that has no fixed allocation defined, the Bridge can dynamically allocate that VID to an appropriate FID, in accordance with the current set of VLAN Learning Constraints.

8.9.7.2 VLAN Learning Constraints

There are two types of VLAN Learning Constraint:

a) A Shared Learning Constraint (or S Constraint) asserts that Shared VLAN Learning shall occur between a pair of identified VLANs. S Constraints are of the form {A S B}, where A and B are VIDs. An S constraint is interpreted as meaning that Shared VLAN Learning shall occur between the VLANs identified by the pair of VIDs;
b) An Independent Learning Constraint (or I Constraint) asserts that a given VLAN is a member of a set of VLANs amongst which Independent VLAN Learning shall occur. I Constraints are of the form {A I N}, where A is a VID and N is an Independent Set Identifier. An I Constraint is interpreted as meaning that Independent VLAN Learning shall occur among the set of VLANs comprising VLAN A and all other VLANs identified in I Constraints that carry the same Independent Set Identifier, N.

A given VID may appear in any number (including zero) of S Constraints and/or I Constraints.

NOTE 1—S Constraints are

— *Symmetric*: e.g., {A S B} and {B S A} both express an identical constraint;

— *Transitive*: e.g., {A S B}, {B S C} implies the existence of a third constraint, {A S C};

— *Reflexive*: e.g., {A S A} is a valid S Constraint.

I Constraints are not

— *Symmetric*: e.g., {A I 1} and {1 I A} express different constraints;

— *Transitive*: e.g., ({A I 1}, {B I 1}, {B I 2}, {C I 2}) does not imply either {A I 2} or {C I 1}.

The allocation of VIDs to FIDs shall be such that, for all members of the set of active VLANs (8.9.7),

c) A given VID shall be allocated to exactly one FID;
d) If a given VID appears in an I Constraint, then it shall not be allocated to the same FID as any other VID that appears in an I Constraint with the same Independent Set Identifier;
e) If a given VID appears in an S Constraint (either explicit, or implied by the transitive nature of the specification), then it shall be allocated to the same FID as the other VID identified in the same S Constraint;
f) If a VID does not appear in any S or I Constraints, then the Bridge may allocate that VID to any FID of its choice.

NOTE 2—The intent is that the set of Learning Constraints is defined on a global basis; i.e., that all VLAN-aware Bridges are configured with the same set of constraints (although individual constraints may well be defined and distributed by different managers/administrators). Any Bridge therefore sees the complete picture in terms of the Learning Constraints that apply to all VLANs present in the Bridged LAN, regardless of whether they all apply to VLANs that are present in that particular Bridge. This standard provides the definition, in Clause 12, of managed objects and operations that model how individual constraints can be configured in a Bridge; however, the issue of how a distributed management system might ensure the consistent setting of constraints in all Bridges in a Bridged LAN is not addressed by this standard.

8.9.7.3 VLAN Learning Constraint inconsistencies and violations

The application of the rules specified in 8.9.7.2, coupled with any fixed allocations of VIDs to FIDs that may have been configured, can result in the Bridge detecting Learning Constraint inconsistencies and/or violations (i.e., can result in situations where there are inherent contradictions in the combined specification of the VLAN Learning Constraints and the fixed allocations, or the Bridge’s own limitations mean that it cannot meet the set of VLAN Learning Constraints that have been imposed upon it).

A Bridge detects a Learning Constraint inconsistency if

a) The VLAN Learning Constraints, coupled with any fixed VID to FID allocations, are such that, if any given pair of VLANs became members of the set of active VLANs (8.9.7), the result would be a simultaneous requirement for Independent VLAN Learning and for Shared VLAN Learning for those two VLANs. Such an inconsistency would require the Bridge to allocate that pair of VIDs both to the same FID and to different FIDs.

Learning Constraint inconsistencies are detected when a management operation (12.10.3) attempts to set a new Learning Constraint value, or to modify the fixed VID to FID allocations. If the new constraint or allocation that is the subject of the operation is inconsistent with those already configured in the Bridge, then the management operation shall not be performed and an error response shall be returned.

A Bridge detects a Learning Constraint violation if

b) The Bridge does not support the ability to map more than one VID to any given FID, and the VLAN Learning Constraints indicate that two or more members of the active set of VLANs require to be mapped to the same FID; or
c) The number of FIDs required in order to correctly configure the Bridge to meet the VLAN Learning Constraints and fixed VID to FID allocations for all members of the active set of VLANs exceeds the number of FIDs supported by the Bridge.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Learning Constraint violations are detected

d) When a VLAN that was hitherto not a member of the set of active VLANs (8.9.7) becomes active, either as a result of management action or as a result of the operation of GVRP, resulting in the Bridge no longer being able to support the defined set of constraints and/or fixed allocations for the set of active VLANs; or
e) When other management reconfiguration actions, such as defining a new Learning Constraint or fixed VID to FID allocation, results in the Bridge no longer being able to support the defined set of constraints and/or fixed allocations for the set of active VLANs.

On detection of a violation, the Bridge issues the Notify Learning Constraint Violation management notification (12.10.3.10), in order to alert any management stations to the existence of the violation. There is the potential for a single change in configuration to result in more than one VLAN whose constraints cannot be met; in such cases, multiple notifications are generated.

8.9.8 Querying the Filtering Database

If a frame is classified into a VLAN containing a given outbound Port in its member set (8.9.9), forwarding or filtering through that Port is determined by the control elements of filtering entries for the frame’s destination MAC Address and for VLANs with the same VID or Filtering Identifier (FID, 8.9.7) as the frame’s VLAN.

Each entry in the Filtering Database for a MAC Address comprises

a) A MAC Address specification;
b) A VID or, in the case of Dynamic Filtering Entries, an FID;
c) A Port Map, with a control element for each outbound Port.

For Dynamic Filtering Entries, the FID that corresponds to a given VID is determined as specified in 8.9.7.

For a given VID, a given individual MAC Address specification can be included in the Filtering Database in a Static Filtering Entry, a Dynamic Filtering Entry, both or neither. Table 8-1 combines Static Filtering Entry and Dynamic Filtering Entry information for an individual MAC Address to specify forwarding, or filtering, of a frame with that destination MAC Address and VID through an outbound Port.

NOTE 1—The use of FID in this table for Static Filtering Entries, and the text in parentheses in the headings, reflects the fact that, where more than one VID maps to a given FID, there may be more than one Static Filtering Entry that affects the forwarding decision for a given individual MAC Address. The effect of all Static Filtering Entries for that address, and for VIDs that correspond to that FID, is combined, such that, for a given outbound Port:
— IF <any static entry for any VIDs that map to that FID specifies Forwarding> THEN <result = Forwarding>
— ELSE IF <any static entry for any VIDs that map to that FID specifies Filtering> THEN <result = Filtering>
— ELSE <result = Use Dynamic Filtering Information>

Table 8-2 specifies the result, Registered or Not Registered, of combining a Static Filtering Entry and a Group Registration Entry for the “All Group Addresses” address specification, and for the “All Unregistered Group Addresses” address specification for an outbound Port.

Table 8-3 combines Static Filtering Entry and Group Registration Entry information for a specific group MAC Address with the Table 8-2 results for All Group Addresses and All Unregistered Group Addresses to specify forwarding, or filtering, of a frame with that destination group MAC Address through an outbound Port.

Where a given VID is allocated to the same FID as one or more other VIDs, it is an implementation option as to whether

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

Table 8-1—Combining Static and Dynamic Filtering Entries for an individual MAC Address

Filtering Information	Control Elements in any Static Filtering Entry or Entries for this individual MAC Address, FID, and outbound Port specify:				
	Forward (Any Static Filtering Entry for this Address/ FID/Port specifies Forward)	Filter (No Static Filtering Entry for this Address/ FID/Port specifies Forward)	Use Dynamic Filtering Information (No Static Filtering Entry for this Address/FID/Port specifies Forward or Filter), or no Static Filtering Entry present. Dynamic Filtering Entry Control Element for this individual MAC Address, FID and outbound Port specifies:		
			Forward	Filter	No Dynamic Filtering Entry present
Result	Forward	Filter	Forward	Filter	Forward

Table 8-2—Combining Static Filtering Entry and Group Registration Entry for “All Group Addresses” and “All Unregistered Group Addresses”

Filtering Information	Static Filtering Entry Control Element for this group MAC Address, VID, and outbound Port specifies:				
	Registration Fixed (Forward)	Registration Forbidden (Filter)	Use Group Registration Information, or no Static Filtering Entry present. Group Registration Entry Control Element for this group MAC Address, VID and outbound Port specifies:		
			Registered (Forward)	Not Registered (Filter)	No Group Registration Entry present
Result	Registered	Not Registered	Registered	Not Registered	Not Registered

d) The results shown in Table 8-3 directly determine the forwarding/filtering decision for a given VID and group MAC Address (i.e., the operation of the Bridge with respect to group MAC Addresses ignores the allocation of VIDs to FIDs); or

e) The results for a given MAC Address and VID are combined with the corresponding results for that MAC Address for each other VID that is allocated to the same FID, so that if the Table 8-3 result is Forward in any one VLAN that shares that FID, then frames for that group MAC Address will be forwarded for all VLANs that share that FID (i.e., the operation of the Bridge with respect to group MAC Addresses takes account of the allocation of VIDs to FIDs).

NOTE 2—In case d), the implementation effectively operates a single FDB per VLAN for group MAC Addresses. In case e), the implementation combines static and registered information for group MAC Addresses in accordance with the VID to FID allocations currently in force, in much the same manner as for individual MAC Addresses.

8.9.9 Determination of the member set and untagged set for a VLAN

The VLAN configuration information contained in the Filtering Database for a given VLAN may include a Static VLAN Registration Entry (8.9.2) and/or a Dynamic VLAN Registration Entry (8.9.5). This information defines, for that VLAN:

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Table 8-3—Forwarding or Filtering for specific group MAC Addresses

				Static Filtering Entry Control Element for this group MAC Address, VID and outbound Port specifies:				
				Registration Fixed (Forward)	Registration Forbidden (Filter)	Use Group Registration Information, or no Static Filtering Entry present. Group Registration Entry Control Element for this group MAC Address, VID and outbound Port specifies:		
						Registered (Forward)	Not Registered (Filter)	No Group Registration Entry present
All Group Addresses control elements for this VID and Port specify (Table 8-2):	Not Registered	All Unregistered Group Addresses control elements for this VID and Port specify (Table 8-2):	Not Registered	Forward	Filter	Forward	Filter	Filter (Filter Unregistered Groups)
			Registered	Forward	Filter	Forward	Filter	Forward (Forward Unregistered Groups)
	Registered			Forward	Filter	Forward (Forward All Groups)	Forward (Forward All Groups)	Forward (Forward All Groups)

a) The member set, consisting of the set of Ports through which members of the VLAN can currently be reached;

b) The untagged set, consisting of the set of Ports through which, if frames destined for the VLAN are to be transmitted, they shall be transmitted without tag headers. For all other Ports (i.e., all Ports that are not members of the untagged set), if frames destined for the VLAN are to be transmitted, they shall be transmitted with tag headers.

NOTE 1—As the ingress rule checking function of the Forwarding Process (8.6.1) always associates a non-null VLAN ID with an incoming frame, all frames (including received frames that were priority-tagged and carried the null VLAN ID in their tag header) will be transmitted with or without a tag header in accordance with the membership of the untagged set for their VID.

For a given VID, the Filtering Database can contain a Static VLAN Registration Entry, a Dynamic VLAN Registration Entry, both or neither. Table 8-1 combines Static VLAN Registration Entry and Dynamic VLAN Registration Entry information for a VLAN and Port to give a result *member*, or *not member*, for the Port. The member set for a given VLAN consists of all Ports for which the result is member.

Membership of the untagged set for a given VLAN is derived from Static VLAN Registration Entry information contained in the Filtering Database as follows:

c) If there is no Static VLAN Registration Entry for the VLAN, then the untagged set is the empty set; otherwise,

Table 8-1—Determination of whether a Port is in a VLAN’s member set

Filtering Information	**Static VLAN Registration Entry Control Element for this VID and Port specifies:**				
	Registration Fixed	**Registration Forbidden**	**Normal Registration, or no Static VLAN Registration Entry present. Dynamic VLAN Registration Entry Control Element for this VID and Port specifies:**		
			Registered	**Not Registered**	**No Dynamic VLAN Registration Entry present**
Result	Member	Not member	Member	Not member	Not member

d) The untagged set is equal to the set of Ports for which the Port Map in the Static VLAN Registration Entry indicates that frames are to be transmitted untagged.

The untagged set and the member set for a given VLAN are used by the Forwarding Process in applying the ingress rules (8.6.1) and the egress rules (8.6.5) for that VLAN.

The initial state of the Permanent Database contains a Static VLAN Registration Entry for the VLAN corresponding to the Default PVID (Table 9-2). The Port Map in this entry specifies Registration Fixed and forwarding untagged for all Ports of the Bridge. This entry may be modified or removed from the Permanent Database by means of the management operations defined in Clause 12 if the implementation supports these operations.

NOTE 2—This causes the default tagging state for the PVID to be untagged, and for all other VIDs to be tagged, unless otherwise configured; however, the management configuration mechanisms allow any VID (including the PVID) to be specified as VLAN-tagged or untagged on any Port. Under normal circumstances, the appropriate configuration for the PVID would be untagged on an access Port or a hybrid Port, and VLAN-tagged on a trunk Port (Annex D discusses the terms *access Port*, *hybrid Port*, and *trunk Port*).

8.9.10 Permanent Database

The Permanent Database provides fixed storage for a number of Static Filtering Entries and Static VLAN Registration Entries. The Filtering Database shall be initialized with the Filtering Database Entries contained in this fixed data store.

Entries may be added to and removed from the Permanent Database under explicit management control, using the management functionality defined in Clause 12. Changes to the contents of Static Filtering Entries or Static VLAN Registration Entries in the Permanent Database do not affect forwarding and filtering decisions taken by the Forwarding Process or the egress rules until such a time as the Filtering Database is re-initialized.

NOTE 1—This aspect of the Permanent Database can be viewed as providing a “boot image” for the Filtering Database, defining the contents of all initial entries, before any dynamic filtering information is added.

NOTE 2—Subclause 10.3.2.3 of IEEE Std 802.1D, 1998 Edition defines an initial state for the contents of the Permanent Database, required for the purposes of GMRP operation.

8.10 MST configuration information

In order to support multiple spanning trees, an MST Bridge has to be configured with an unambiguous assignment of VIDs to spanning trees. This is achieved by:

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

a) Ensuring that the allocation of VIDs to FIDs (8.9.7) is unambiguous; and
b) Ensuring that each FID supported by the Bridge is allocated to exactly one Spanning Tree.

The first of these requirements is met by configuring a set of VLAN learning constraints and/or fixed VID to FID mappings that are self-consistent, and which define an I Constraint, an S Constraint, or a fixed VID to FID allocation for all VIDs supported by the Bridge.

The second requirement is met by means of the FID to MSTI Allocation Table (8.10.3).

The combination of the VID to FID allocations and the FID to MSTI allocations defines a mapping of VIDs to spanning trees, represented by the MST Configuration Table (8.10.1).

8.10.1 MST Configuration Table

The MST Configuration Table defines, for each VID, the MSTID of the spanning tree instance to which the VID is allocated.

The MST Configuration Table cannot be configured directly; configuration of the table occurs as a consequence of configuring the relationships between VIDs and FIDs (8.9.7), and between FIDs and spanning trees (8.10.3).

8.10.2 MST Configuration Identification

For two or more MST Bridges to be members of the same MST Region (3.26), it is necessary for those Bridges to be directly connected together (i.e., interconnected only by means of LANs, without intervening Bridges that are not members of the region), and for them to support the same MST Region configuration. Two MST Region configurations are considered to be the same if the correspondence between VIDs and spanning trees is identical in both configurations.

NOTE 1—If two adjacent MST Bridges consider themselves to be in the same MST Region despite having different mappings of VIDs to spanning trees, then the possibility exists of undetectable loops arising within the MST Region.

In order to ensure that adjacent MST Bridges are able to determine whether they are part of the same MST Region, the MST BPDU supports the communication of an MST Configuration Identifier (13.7).

NOTE 2—As the MST Configuration Identifier is smaller than the mapping information that it summarizes, there is a small but finite possibility that two MST Bridges will assume that they have the same MST Region Configuration when this is not actually the case. However, given the size of the identifier, this standard assumes that this possibility is sufficiently small that it can safely be ignored. Appropriate use of the Configuration Name and Revision Level portions of the identifier can remove the possibility of an accidental match between MST Configuration Identifiers that are derived from different configurations within a single administrative domain (see 13.7).

8.10.3 FID to MSTI Allocation Table

The FID to MSTI Allocation Table defines, for all FIDs that the Bridge supports, the MSTID of the spanning tree instance to which the FID is allocated. An MSTID of zero is used to identify the CIST.

NOTE—The management operations defined in Clause 12.12 make use of the concept of an MSTI List to instantiate/de-instantiate MST instances, and will only permit the allocation of FIDs to MSTIDs that are present in the MSTI List.

8.11 Spanning Tree Protocol Entity

Figure 8-5 illustrates the operation of the Spanning Tree Protocol Entity including the reception and transmission of frames containing BPDUs, the modification of the state information associated with individual Bridge Ports, and notification of the Filtering Database of changes in active topology.

A given MST bridge is not required to support all of the spanning trees that exist within the MST bridged network. That is, the number of spanning trees operated by the Spanning Tree Protocol Entity in a given bridge may be different from the number operated by that in another bridge. However, as a direct consequence of the conditions stated in 8.10.2, the number of instances of the Spanning Tree Protocol operated by a given MST Bridge is the same for all Bridges that are members of a given MST Region.

8.12 GARP Entities

The GARP Protocol Entities operate the Algorithms and Protocols associated with the GARP Applications supported by the Bridge, and consist of the set of GARP Participants for those GARP Applications (Clause 10 and 12.3 of IEEE Std 802.1D, 1998 Edition).

Figure 8-6 illustrates the operation of a GARP Protocol Entity including the reception and transmission of frames containing GARP PDUs, the use of control information contained in the Filtering Database, and notification of the Filtering Database of changes in filtering information.

8.13 Bridge Management Entity

Bridges should be managed using SNMP.

<<Expand, but briefly, on the above. Input requested.>>

8.14 Addressing

All MAC Entities communicating across a Bridged Local Area Network use 48-bit addresses. These can be Universally Administered Addresses or a combination of Universally Administered and Locally Administered Addresses.

8.14.1 End stations

Frames transmitted between end stations using the MAC Service carry the MAC Address of the source and destination peer end stations in the source and destination address fields of the frames, respectively. The address, or other means of identification, of a Bridge is not carried in frames transmitted between peer users for the purpose of frame relay in the network.

In the absence of explicit filters configured via management as Static Filtering Entries, or via GMRP as Group Registration Entries (8.9, IEEE Std 802.1D-2003 Clause 10), frames with a destination address of the broadcast address or any other group address that is not a Reserved Address (8.6.3) are assigned to a VLAN and relayed throughout that VLAN.

8.14.2 Bridge Ports

A separate individual MAC Address is associated with each instance of the MAC Service provided to an LLC Entity. That MAC Address is used as the source address of frames transmitted by the LLC Entity.

Media access method specific procedures can require the transmission and reception of frames that use an individual MAC Address associated with the Bridge Port, but neither originate from nor are delivered to a MAC Service user. Where an individual MAC Address is associated with the provision of an instance of the MAC Service by the Port, that address can be used as the source and or destination address of such frames, unless the specification of the media access method specific procedures requires otherwise.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

8.14.3 Use of LLC by Spanning Tree Protocol and GARP Entities

Both Spanning Tree Protocol and GARP Entities uses the DL_UNITDATA.request and DL_UNITDATA.indication primitives (ISO/IEC 8802-2) provided by individual LLC Entities associated with each Bridge Port to transmit and receive. The source_address and destination_address parameters of the DL_UNITDATA.request shall both denote the standard LLC address assigned to the Bridge Spanning Tree Protocol (Table 8-4). Each DL_UNITDATA request primitive gives rise to the transmission of an LLC UI command PDU, which conveys the BPDU or GARP PDU in its information field.[9]

Table 8-4—Standard LLC address assignment

Assignment	Value
Bridge spanning tree protocol	01000010

Code Representation: The least significant bit of the value shown is the right-most. The bits increase in significance from right to left. It should be noted that the code representation used here has been chosen in order to maintain consistency with the representation used elsewhere in this standard; however, it differs from the representation used in IEEE Standard 802.1D-2003.

IEEE Std 802.1D defines a Protocol Identifier field, present in all BPDUs (IEEE Std 802.1D, Clause 9) and GARP PDUs (IEEE Std 802.1D, 12.11), which serves to identify different protocols supported within the scope of the LLC address assignment. Further values of this field are reserved for future standardization. A Spanning Tree Protocol Entity or GARP Protocol Entity that receives a BPDU or a GARP PDU with an unknown Protocol Identifier shall discard that PDU.

8.14.4 Reserved MAC Addresses

Any frame with a destination address that is a Reserved MAC Address shall not be forwarded by a Bridge. Reserved MAC Addresses for Customer-VLAN aware Bridges and for Service-VLAN aware Bridges are specified in Table 8-1 and Table 8-2 respectively. These Group MAC Addresses are reserved for assignment to standard protocols, according to the criteria for such assignments (Clause 5.5 of ISO/IEC TR 11802-2).

8.14.5 Group MAC Addresses for spanning tree protocols

A Spanning Tree Protocol Entity uses an instance of the MAC Service provided by each of the Bridge's Ports to transmits frames addressed to the Spanning Tree Protocol Entities of all the other Bridges attached to the same LAN as that Port. A 48-bit universally administered Group Address, known as the Bridge Group Address, has been assigned (Table 8-1) for use by Customer-VLAN aware Bridges for this purpose.

It is essential to the operation of the spanning tree protocols that frames with this destination address both reach all the peer protocol entities attached to the same LAN and do not reach protocol entities attached to other LANs. The Bridge Group Address is therefore included in the block of Customer-VLAN aware Bridge Reserved MAC Addresses and is always filtered by Customer Bridges (8.6.3).

Since a network of Service-VLAN aware Bridges needs to appear to attached Customer Bridges as if it were a single LAN, frames addressed to the Bridge Group Address are forwarded. Service-VLAN aware Bridges also use a spanning protocol to provide one or more loop-free active topologies, so a distinct 48-bit universally administered Group Address, the Service-VLAN aware Bridge Group Address, that can be

[9] ISO/IEC TR 11802-1: 1997, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Technical reports and guidelines—Part 1: The structure and coding of Logical Link Control addresses in Local Area Networks, contains the full list of standard LLC address assignments, and documents the criteria for assignment.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

confined to the LANs that their Bridge Ports attach, has been assigned (Table 8-2). The Service-VLAN aware Bridge Group Address is included in both the Customer-VLAN aware and Service-VLAN aware Bridge Reserved MAC Addresses and is always filtered by all Bridges (8.6.3).

The source MAC address field of frames transmitted by Spanning Tree Protocol Entities contains the individual MAC Address for the Bridge Port used to transmit the frame.

8.14.6 Group MAC Addresses for GARP Applications

A GARP Entity that supports a given GARP Application transmits frames addressed to all other GARP Entities that implement the same GARP Application. The peers of each such entity bound a region of the network that contains no peers, commonly a single LAN in the case where all Bridges attached to the LAN implement the application. A distinct universally administered 48-bit Group Address is assigned to each GARP application. Filtering Database Entries for each GARP Application address assigned to an application that is supported by a Customer-VLAN aware Bridge should be configured in the Filtering Database so as to confine frames for that application to the peer region, while addresses for applications that are not supported should not be included.

A set of 48-bit Universal Addresses, known as GARP Application addresses, have been assigned for use by Customer Bridges. The values of the GARP Application addresses are defined in IEEE Std 802.1D, Table 12-1. These group MAC Addresses are reserved for assignment to standard protocols, according to the criteria for such assignments (Clause 5.5 of ISO/IEC TR 11802-2).

NOTE—Table 11-1 allocates a group MAC Address for use by the GVRP application; however, the value allocated in that table is one of the GARP Application addresses reserved by IEEE Std 802.1D, Table 12-1.

Since a network of Provider Bridges needs to appear to attached Customer Bridges as if it were a single LAN, Customer Bridge GARP Application Addresses are always forwarded by Provider Bridges. Certain GARP Applications may also be used by Provider Bridges, so distinct 48-bit universally administered Group Addresses that are Customer Reserved MAC Addresses but not Provider Bridged Reserved MAC Addresses are assigned for such use. One such address, the Provider Bridge GVRP Address, is assigned by this standard (Table 8-1).

The source address field of MAC frames conveying BPDUs or GARP PDUs contains the individual MAC Address for the Bridge Port through which the PDU is transmitted (8.14.2).

8.14.7 Bridge Management Entities

Bridges should be managed using SNMP, which uses IP to support management communication. If implemented, the IP stack and IP Address used shall be supported by a single LLC Entity attached to a Bridge Port. The Port should be a Management Port for the Bridge, but may be a Port attached to a LAN.

NOTE—A 48-bit universally administered Group Address, known as the All LANs Bridge Management Group Address with a value of 01-80-C2-00-00-10 was assigned and recorded in the 1990 Edition of this Standard. That address should not be used for bridge management, or for any other purpose.

8.14.8 Unique identification of a Bridge

A unique 48-bit Universally Administered MAC Address, termed the Bridge Address, shall be assigned to each Bridge. The Bridge Address may be the individual MAC Address of a Bridge Port, in which case use of the address of the lowest numbered Bridge Port (Port 1) is recommended.

NOTE—The Rapid Spanning Tree Protocol (RSTP) (IEEE Std 802.1D-2003, Clause 17) and the Multiple Spanning Tree Protocol (MSTP) (Clause 13) require that a single unique identifier be associated with each Bridge. That identifier is derived from the Bridge Address as specified in 9.2.5 of IEEE Std 802.1D-2003.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

8.14.9 Points of attachment and connectivity for Higher Layer Entities

The Higher Layer Entities in a Bridge, such as the Spanning Tree Protocol Entity (8.11), GARP Entities (8.12), and Bridge Management (8.13), are modeled as attaching directly to one or more individual LANs connected by the Bridge’s Ports, in the same way that any distinct end station is attached to the network. While these entities and the relay function of the Bridge use the same individual MAC entities to transmit and receive frames, the addressing and connectivity to and from these entities is the same as if they were attached as separate end stations “outside” the Port or Ports where they are actually attached. Figure 8-10 is functionally equivalent to Figure 8-2, but illustrates this logical separation between the points of attachment used by the Higher Layer Entities and those used by the MAC Relay Entity.

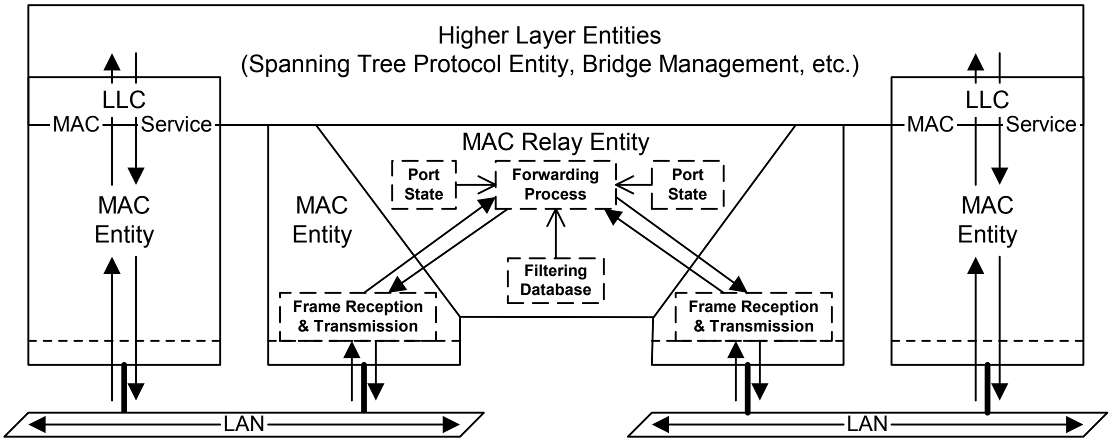

Figure 8-10—Logical points of attachment of the Higher Layer and Relay Entities

Figure 8-11 depicts the information used to control the forwarding of frames from one Bridge Port to another (the Port States and the content of the Filtering Database) as a series of switches (shown in the open, disconnected state) inserted in the 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. While showing Higher Layer Entities sharing the point of attachment to each LAN used by each Bridge Port to forward frames, this figure further illustrates a point made by Figure 8-10: controls placed in the forwarding path have no effect upon the ability of a Higher Layer Entity to transmit and receive frames to or from a given LAN using a direct attachment to that LAN (e.g., from entity A to LAN A), they only affect the path taken by any indirect transmission or reception (e.g., from entity A to or from LAN B).

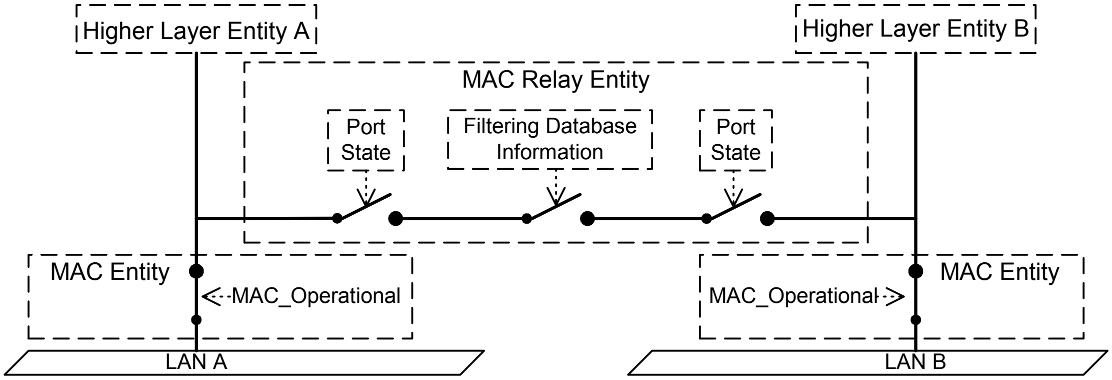

Figure 8-11—Effect of control information on the forwarding path

The functions provided by Higher Layer Entities can be categorized as requiring either

a) a single point of attachment to the Bridged Local Area Network, providing connectivity to stations attached to the network at any point (subject to administrative control), as does Bridge Management; or
b) a distinct point of attachment to each individual LAN attached by a Bridge Port, providing connectivity only to peer entities connected directly to that LAN, as do the Spanning Tree Protocol Entity and the GARP Entity.

In the latter case it is essential that the function associate each received and transmitted frame with a point of attachment. Frames transmitted or received via one point of attachment are not to be relayed to and from other Ports and attached LANs, so the MAC Addresses (8.14.4) used to reach these functions are permanently configured in the Filtering Database.

NOTE 1 —Addresses used to reach functions with distinct points of attachment are generally group MAC Addresses.

NOTE 2—A single higher layer entity can incorporate both a function requiring a single point of attachment and a function requiring distinct points of attachment. The two functions are reached using different MAC addresses.

Figure 8-12 illustrates forwarding path connectivity for frames destined for Higher Layer Entities requiring per-Port points of attachment. Configuration of the Permanent Database in all Bridges to prevent relay of frames addressed to these entities means that they receive frames only via their direct points of attachment (i.e., from LAN A to entity A, and from LAN B to entity B), regardless of Port states.

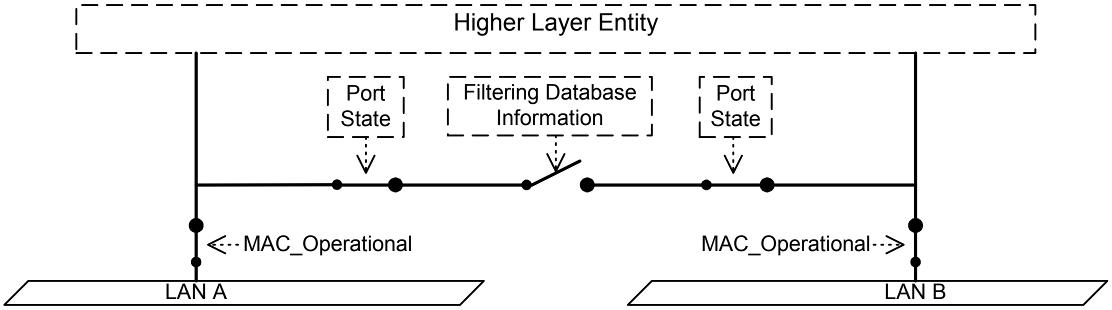

Figure 8-12—Per-Port points of attachment

Figures 8-13 and 8-14 illustrate forwarding path connectivity for frames destined for a Higher Layer Entity requiring a single point of attachment. In both figures the Filtering Database permits relay of frames, as do the Port states in Figure 8-13 where frames received from LAN B are relayed by the Bridge to the entity and to LAN A.

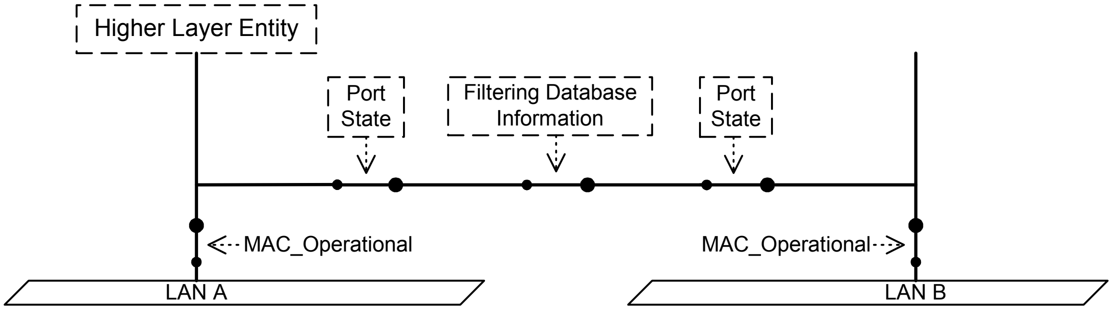

Figure 8-13—Single point of attachment—relay permitted

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

In Figure 8-14 frames received from LAN A are received by the entity directly, but frames received from LAN B are not relayed by the Bridge, and will only be received by the entity if another forwarding path is provided between LANs A and B. If the Discarding Port state shown resulted from spanning tree computation (and not from disabling the Administrative Bridge Port State) such a path will exist via one or more Bridges. If there is no active Spanning Tree path from B to A the network has partitioned into two separate Bridged Local Area Networks, and the Higher Layer Entity shown is reachable only via LAN A.

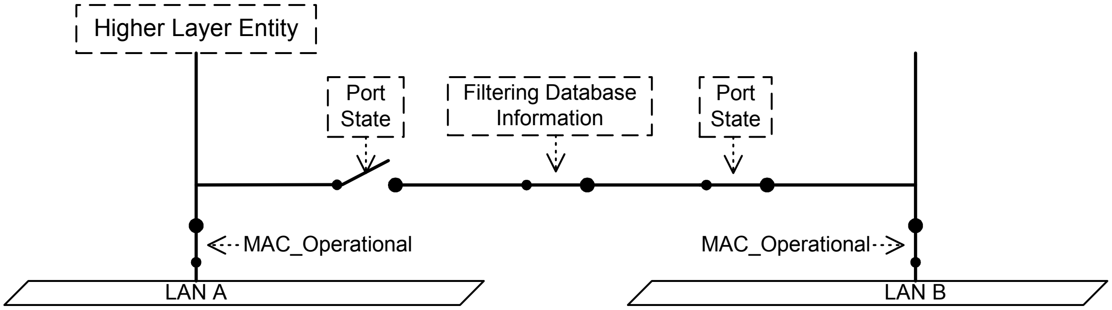

Figure 8-14—Single point of attachment—relay not permitted

Specific Higher Layer Entities can take notice of the Administrative Bridge Port State, as required by their specification. The Spanning Tree Protocol Entity is one such example — BPDUs are never transmitted or received on Ports with an Administrative Bridge Port State of Disabled.

If a Bridge Port's MAC Entity is not operational, a Higher Layer Entity directly attached at the Port will not be reachable, as Figure 8-15 illustrates. The Spanning Tree Protocol Entity ensures that the Port State is Discarding if the MAC_Operational (6.4.2) is FALSE even if the Administrative Bridge Port State is Enabled.

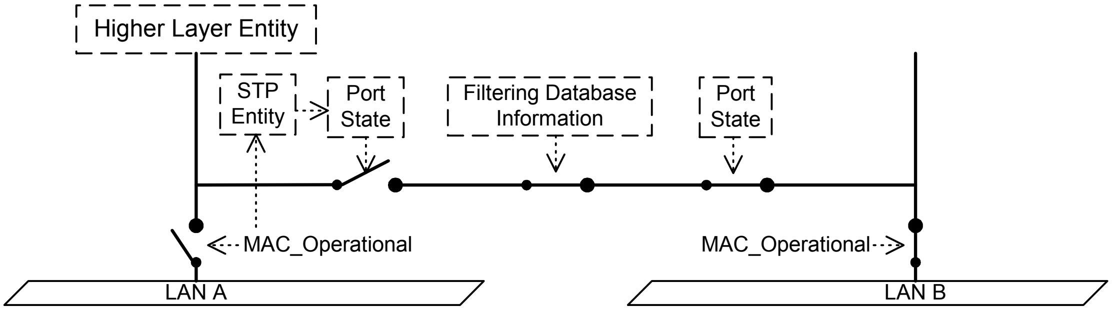

Figure 8-15—Effect of Port State

The connectivity provided to Higher Layer Entities and to the LANs that compose a Bridged Local Area Network can be further controlled by a Bridge Port operating as a network access port (IEEE Std 802.1X). The operation of Port-based access control has the effect of creating two distinct points of access to the LAN. One, the *uncontrolled Port,* allows transmission and reception of frames to and from the attached LAN regardless of the authorization state; the other, the *controlled Port*, only allows transmission following authorization. If the port is not authorized, the Spanning Tree Protocol Entity, which uses the controlled port (as does the MAC Relay Entity) will be unable to exchange BPDUs with other Bridges attached to LAN A, and will set the Bridge Port State to Discarding.

NOTE—If the Spanning Tree Protocol Entity was not aware of the Unauthorized state of the Port, and believed that it was transmitting and receiving BPDUs it might assign a Bridge Port State of Forwarding. Following authorization a temporary loop in network connectivity might then be created.

Figure 8-16 illustrates the connectivity provided to Higher Layer Entities if the MAC entity is physically capable of transmitting and receiving frames, i.e. MAC_Operational is TRUE, but AuthControlledPortStatus

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

is Unauthorized. Higher Layer Entity A and the PAE (the port access entity that operates the authorization protocol) are connected to the uncontrolled port and can transmit and receive frames using the MAC entity associated with the Port, which Higher Layer Entity B cannot. None of the three entities can transmit or receive to or from LAN B.

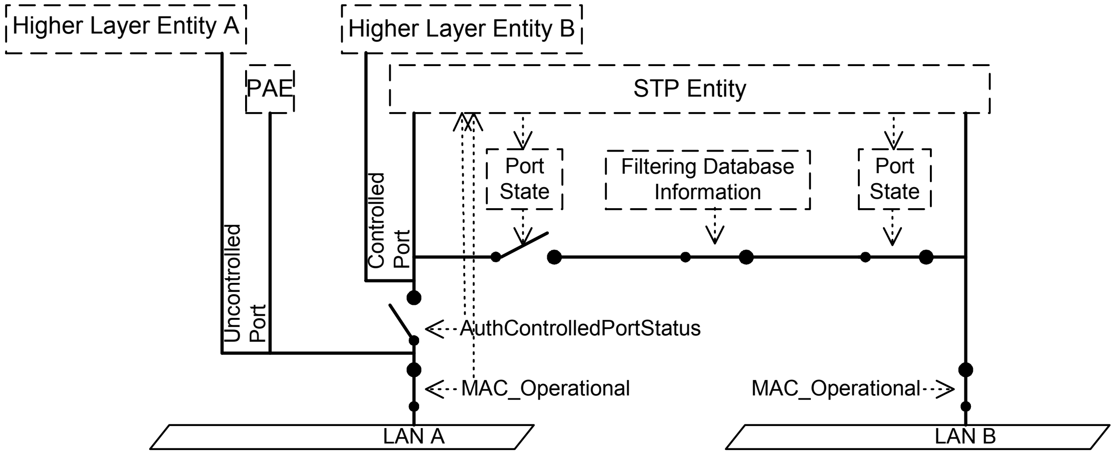

Figure 8-16—Effect of authorization

NOTE—The administrative and operational state values associated with the MAC, the Port’s authorization state, and the Bridge Port State, equate to the ifAdminStatus and ifOperStatus parameters associated with the corresponding interface definitions; see IETF RFC 2233.

8.14.10 VLAN attachment and connectivity for Higher Layer Entities

In VLAN-aware Bridges, two more switches appear in the forwarding path, corresponding to the actions taken by the Forwarding Process (8.6.1 and 8.6.5) in applying the ingress and egress rules (8.6.1 and 8.6.5), as illustrated in Figure 8-17.

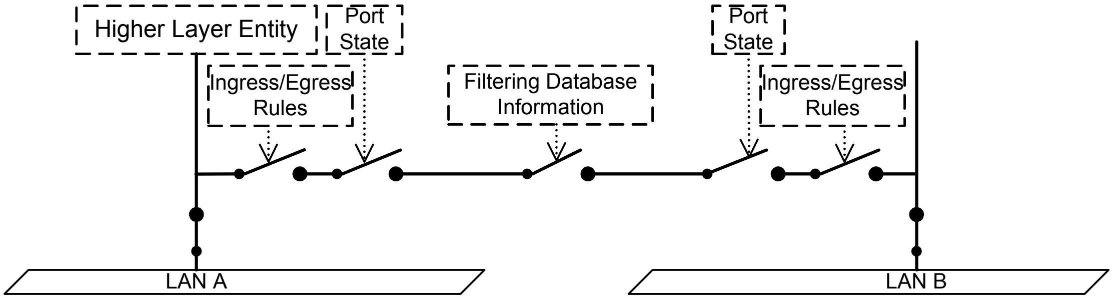

Figure 8-17—Ingress/egress control information in the forwarding path

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 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 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 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

a) Spanning Tree BPDUs transmitted by the Spanning Tree Protocol Entity are not forwarded by Bridges, and must be visible to all other Spanning Tree Protocol Entities attached to the same LAN. 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.

b) The definition of the GVRP application (11.2.3) calls for all GVRP frames to be transmitted untagged for similar reasons;
c) 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;
d) 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 Virtual Bridged Local Area Network. This is normally achieved by routing a packet containing the PDU to the routed subnet associated with the VLAN. Transmission of the packet through the router interface to that VLAN and subsequent forwarding of the resulting frame by VLAN-aware Bridges ensures that the frame is correctly VLAN-tagged, as required.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Delete the existing contents of Clause 9, and insert replacement contents as shown below.

9. Tagged frame format

This clause specifies the format of the VLAN tags added to and removed from user data frames by the tag encoding and decoding functions that support the Enhanced Internal Sublayer Service (EISS, 6.4). It:

a) reviews the purpose of VLAN tagging, and the functionality provided;
b) specifies generic rules for the representation of tag fields and their encoding in the octets of a MAC Service Data Unit (MSDU);
c) specifies a general tag format, comprising a Tag Protocol Identifier (TPID), Tag Control Information (TCI), with additional information as signaled in the TCI, and;
d) specifies the format of the TPID for each 802 media access control method;
e) describes the types of VLAN tag that can be used;
f) documents the allocation of Ethertype values to identify the types of tag specified in this standard;
g) specifies the format of the TCI and additional information for each tag type.

Further analysis of the frame formats and the format translations that can occur when frames are tagged or untagged when relayed between different media access control methods can be found in Annex C.

9.1 Purpose of tagging

Tagging a frame with a VLAN tag:

a) Allows a VLAN Identifier (VID) to be conveyed, facilitating consistent VLAN classification of the frame throughout the network and enabling segregation of frames assigned to different VLANs;
b) Allows priority (6.4, 6.6) to be conveyed with the frame when using 802 LAN media access control methods that provide no inherent capability to signal priority;
c) Can support the use of differing media access control methods within a single network, by:
    1) Signaling the bit order of MAC address information conveyed in the MSDU;
    2) Signaling the presence or absence of an embedded routing information field (E-RIF) (6.4) carried in the MSDU of MAC types that provide no inherent capability to convey routing information.

9.2 Representation and encoding of tag fields

In this clause, octets are numbered starting from 1 and increasing in the order in which they are encoded in the sequence of octets that comprise a MAC Service Data Unit (MSDU).

Where bits in consecutive octets are used to encode a binary number in a single field, the lower octet number encodes the more significant bits of the field, and the least significant bit of the lower octet number and the most significant bit of the next octet both form part of the field.

Where the value of a field comprising a sequence of octets is represented as a sequence of two-digit hexadecimal values separated by hyphens (e.g., A1-5B-03), the leftmost hexadecimal value (A1 in this example) appears in the lowest numbered octet of the field and the rightmost hexadecimal value (03 in this example) appears in the highest numbered octet of the field.

The bits in an octet are numbered from 1 to 8, where 1 is the least significant bit.

Copyright © 2004 IEEE. All rights reserved.

When the terms *set* and *reset* are used in the text to indicate the values of single-bit fields, *set* is encoded as a binary 1 and *reset* as a binary 0 (zero).

When the encoding of a field or a number of fields is represented using a diagram:

a) Octets are shown with the lowest numbered octet nearest the top of the page, the octet numbering increasing from the top to bottom; or
b) Octets are shown with the lowest numbered octet nearest the left of the page, the octet numbering increasing from left to right;
c) Within an octet, bits are shown with bit 8 to the left and bit 1 to the right.

9.3 Tag format

Each VLAN tag comprises the following sequential information elements:

a) A Tag Protocol Identifier (TPID) (9.4);
b) Tag Control Information (TCI) that is dependent on the tag type (9.5, 9.6, 9.7);
c) Additional information, if and as required by the tag type and TCI.

The tag encoding function supports each EISS (6.4) instance by using an instance of the Internal Sublayer Service (ISS) to transmit and receive frames, and encodes the above information in the first and subsequent octets of the MSDU that will accompany an ISS M_UNITDATA.request, immediately prior to encoding the sequence of octets that comprise the corresponding EISS M_UNITDATA.request's MSDU. On reception the tag decoding function is selected by the TPID, and decodes the TCI and additional information octets (if present) prior to issuing an EISS M_UNITDATA.indication with an MSDU that comprises the subsequent octets.

9.4 Tag Protocol Identifier (TPID) Formats

The TPID includes an Ethernet Type value that is used to identify the frame as a tagged frame and to select the correct tag decoding functions.

Where the ISS instance used to transmit and receive tagged frames is provided by a media access control method that can support Ethernet Type encoding directly (i.e. is an 802.3 or 802.11 MAC) or is media access method independent (e.g. 6.6), the TPID is Ethernet Type encoded, i.e. is two octets in length and comprises solely the assigned Ethernet Type value.

Where the ISS instance is provided by a media access method that cannot directly support Ethernet Type encoding (i.e. is an 802.5 or FDDI MAC), the TPID is encoded according to the rule for a Subnetwork Access Protocol (IEEE Std 802-2001 Clause 10) that encapsulates Ethernet frames over LLC, and comprises the SNAP header (AA-AA-03) followed by the SNAP PID (00-00-00) followed by the two octets of the assigned Ethernet Type value.

9.5 Tag Protocol Identification

Two types of tags are specified:

a) a Customer VLAN TAG (C-TAG), for general use by VLAN Bridges (5.1);
b) a Service VLAN TAG (S-TAG), reserved for use by Provider Bridges (5.3, 5.5).

A distinct Ethertype has been allocated (Table 9-1) for use in the TPID field (9.4) of each tag type so they can be distinguished from each other, and from other protocols.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Table 9-1—802.1Q Ethernet Type allocations

Tag Type	Name	Value
Customer VLAN TAG	802.1Q Tag Protocol Type (802.1QTagType)	81-00
Service VLAN TAG	802.1Q Service Tag Type (802.1QSTagType)	<<to be assigned>>

9.6 Customer VLAN Tag Control Information

The C-TAG TCI field (Figure 9-1) is two octets in length, and encodes the vlan_identifier, priority, and drop_eligible parameters of the corresponding EISS M_UNITDATA.request as unsigned binary numbers, and a Canonical Format Indicator (CFI) as a single bit.

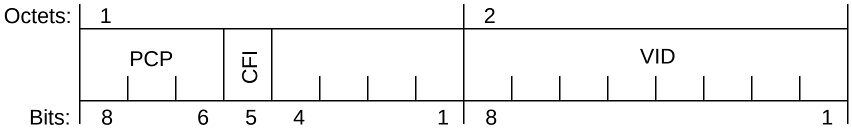

Figure 9-1— C-TAG TCI format

The VLAN Identifier is encoded in a 12 bit field. A VLAN-aware Bridge may not support the full range of VID values, but shall support the use of all VID values in the range 0 through a maximum N, less than or equal to 4094 and specified for that implementation. Table 9-2 identifies VID values that have specific meanings or uses.

Table 9-2—Reserved VID values

VID value (hexadecimal)	Meaning/Use
0	The null VLAN ID. Indicates that the tag header contains only 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.
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.

NOTE 1—There is a distinction between the range of VIDs that an implementation can support, and the maximum number of active VLANs supported at any one time. An implementation supports only 16 active VLANs, for example, may use VIDs chosen from anywhere in the identifier space, or from a limited range. The latter can result in difficulties where different Bridges in the same network support different maximums. It is recommended that new implementations of this standard support the full range of VIDs, even if the number of active VLANs is limited.

The priority and drop_eligible parameters are conveyed in the three bit Priority Code Point (PCP) field as specified in 6.7.3.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

If the CFI is reset, all MAC Address information that may be present in the MSDU is in Canonical format and the tag comprises solely the TPID and TCI fields, i.e. the tag does not contain an Embedded Routing Information Field (E-RIF).

The information conveyed by a CFI bit that is set depends on the media access control method that directly provides the ISS instance used by the tagging or detagging function, as follows:

a) If the TPID is Ethernet Type encoded (802.3, 802.11, or media access method independent provision), an E-RIF () follows the TCI. The NCFI bit in the E-RIF is reset if MAC address information present in the MSDU is in Canonical format and set otherwise (Non-canonical format).
b) If an 802.5 MAC is used, all MAC Address information present in the MSDU is in Non-canonical format.
c) If an FDDI MAC is used and the frame is source routed (i.e., the RII bit in the frame's source MAC Address field is set indicating that a RIF follows the source MAC Address), all MAC Address information present in the MSDU is in Non-canonical format.
d) In an FDDI MAC is used and the frame is not source routed (i.e., the RII bit in the frame's source MAC Address field is reset), an E-RIF () follows the TCI as for (a) above.

NOTE 2—A decision to use native source-routing on FDDI or to use an embedded routing information field in the VLAN tag depends on local knowledge in a Bridge or end station of the capabilities of the other stations attached to the FDDI LAN. The VLAN tag E-RIF allows source-routing information to be transparently "tunneled" across LANs that do not support source routing and through MAC Bridges and VLAN-aware Bridges that discard native source-routed frames.

NOTE 3—An E-RIF is never present when the ISS is directly provided by an 802.5 MAC.

9.7 Service VLAN TAG Control Information

The S-TAG TCI field (Figure 9-1) is two octets in length, and encodes the priority, drop_eligible, and vlan_identifier parameters of the corresponding EISS M_UNITDATA.request as unsigned binary numbers.

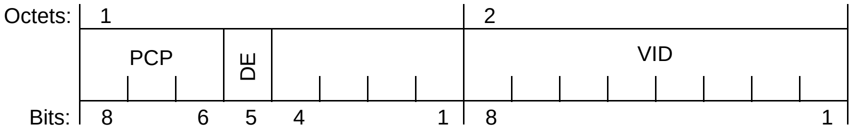

Figure 9-1— C-TAG TCI format

The semantics and structure of the S-TAG is identical to that of the C-TAG with the exception that bit 5, the DE Bit, does not convey a CFI (9.6). The priority and drop_eligible parameters are conveyed in the three bit Priority Code Point (PCP) field and the DE Bit as specified in 6.7.3.

NOTE—Although a Service VLAN tag never includes a routing information field or an indication of canonical or non-canonical address formats, the MSDU of the associated frame can contain a Customer VLAN tag with these elements.

9.8 Embedded Routing Information Field (E-RIF)

The E-RIF that can appear in a Customer VLAN tag may be used to encode the rif_information parameter of the corresponding EISS M_UNITDATA.request, and is a modification of the RIF as defined in IEEE Std 802.1D, 1998 Edition, C.3.3.2. It comprises:

a) A two-octet Route Control Field (RC) (Figure 9-1), that comprises the following fields:
   1) Routing Type (RT);
   2) Length (LTH);

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

3) Direction Bit (D);

4) Largest Frame (LF);

5) Non-canonical Format Indicator (NCFI), set to indicate that all MAC Address information that may be present in the frames MSDU is in Canonical format, and reset otherwise.

b) Zero or more octets of Route Descriptors (up to a maximum of 28 octets), as defined by the Route Control Field.

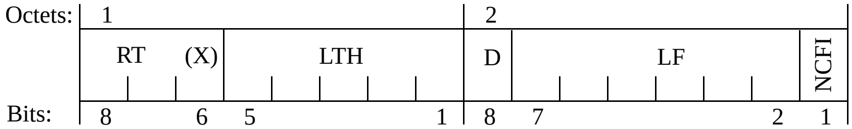

Figure 9-1—E-RIF Route Control (RC) field

The structure and semantics associated with the RT, LTH, D, LF, and Route Descriptor fields are as defined in IEEE Std 802.1D, 1998 Edition, C.3.3.2 with the following exceptions.

On receipt of an EISS M_UNITDATA.request without a rif_information parameter (or with a null rif_information parameter), with the canonical_format_indicator parameter False and the include_tag parameter True, a tagging function that uses an FDDI MAC shall encode an RT value of 010 to indicate a transparent frame, i.e. a frame that does not use the media access method dependent capabilities of any particular MAC to carry routing information and shall encode 00000 in the LTH to indicate that zero octets of Route Descriptor follow.

On receipt of an ISS M_UNITDATA.indication conveying an E-RIF RT value of 010 or 011 the tag decoding function shall discard any Routing Descriptor whose presence is indicated by the LTH field.

On receipt of an EISS M_UNITDATA.request with a rif_information parameter (or with a null rif_information parameter), with the canonical_format_indicator parameter False and the include_tag parameter True, the tagging function that uses an FDDI MAC shall encode an RT value of 000 to indicate a specifically routed frame.

NOTE 1—In effect, the RT bits in the E-RIF encode the state of the RII bit that would appear in an equivalent frame in a source-routed environment. RT values 01X encode RII reset; RT values 00X or 1XX encode RII set.

NOTE 2—the definition of the E-RIF and its use within tag headers does not affect the definition of the RIF used in untagged frames in a source routing environment. The use of an RT value of 01X to indicate a transparent frame applies only to RT values carried in the E-RIF; values of 01X appearing in the RIF of a normal source-routed frame (whether tagged or untagged) are never interpreted in this way. A LTH of 00000 never is never encoded in an untagged frame. The bit position corresponding to the NFCI is reserved in untagged frames and is transmitted as 0.

NOTE 3—The use of a zero length in conjunction with the transparent RT indicator ensures that there is no possibility of such frames being misinterpreted as valid Specifically Routed frames by devices that support source routing.

Use of the Customer VLAN tag formats specified in this clause together with the tag encoding and decoding rules specified in this clause and in clause 6.4 ensure that the interpretation and use of the RT field by VLAN-aware devices is unambiguous, and does not conflict with use by non-VLAN-aware devices. Specifically routed frames always carry an RT value of 000 though they can be received with an RT of 0XX in untagged frames. A VLAN-aware station that generates tagged frames with an E-RIF in order to communicate with stations connected to source-routed LANs shall use an RT of 000.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

9.9 Validation of received frames

The functions that support an instance of the EISS (6.6, 6.7, 8.5) examine the initial octets of the MSDU parameter that accompanying each ISS indication for the supported VLAN tag type. If the MSDU comprises fewer octets than are required to compose the TPID, or the initial octets differ from the TPID, the received frame is identified as a valid untagged frame. Otherwise the received frame is identified as VLAN tagged, but shall be discarded if required by the validation checks (9.9.1, 9.9.2) specified for the supported tag type.

9.9.1 Validating Customer VLAN tags

The received frame shall be discarded if:

a) there are fewer than two octets following the TPID; or
b) the frame is Ethertype encoded, the CFI bit is set, and the encoding of the E-RIF does not meet the requirements of Clause 9.8.

9.9.2 Validating Service VLAN tags

The received frame shall be discarded if there are fewer than two octets following the TPID.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

10. Use of GMRP in VLANs

This amendment makes no changes to Clause 10.

11. VLAN topology management

This amendment makes no changes to Clause 11.

<<The enhancements to GVRP to accommodate the VID translation table are trivial and will be done here.>>

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

12. VLAN bridge management

<<This amendment will clearly make changes to Clause 12. These are omitted from the current draft until we are sure what we are trying to manage. The intent is that provider bridge management will result in a few additions to the existing bridge management, not a new section or extensive replacement.>>

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

13. The Multiple Spanning Tree Protocol (MSTP)

This amendment makes no changes to Clause 13.

<<Use of the restrictedRole parameter requires a modest enhancement to MSTP. Enhancing RSTP in the same way is also desirable, as the utility of the functionality provided goes beyond provider networks and beyond the use of MSTP. When restrictedRole is set for a Port, updtRolesTree() will not select that port as a Root Port, but only as a Designated, Alternate, Backup Port. This is an improvement over just discarding Provider BPDUs received on ingress ports, as it detects some misconfigurations and resolves them without either causing loops or cutting off connectivity entirely. It is useful in the provider network architecture described in clause 16.8. Outside the provider network context it can be used to cut out some potential temporary information propagation paths that will never be included in the final active topology — such as all connectivity from one half of the campus being provided to the other half through some wiring closet — and thus speeds reconfiguration in worst case scenarios. Proprietary technology with similar goals and functionality has been deployed.>>

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

14. Use of BPDUs by MSTP

This amendment makes no changes to Clause 14.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Insert the following Clause after Clause 14.

15. Support of the MAC Service by Provider Bridged Networks

Provider Bridges interconnect the separate MACs of the IEEE 802 LANs that compose a Provider Bridged Network, relaying and filtering frames to provide connectivity between all the LANs that provide the customer interfaces for a given service instance. The position of the Provider Bridge supported S-VLAN aware Bridging function within the MAC Sublayer is shown in Figure 15-1.

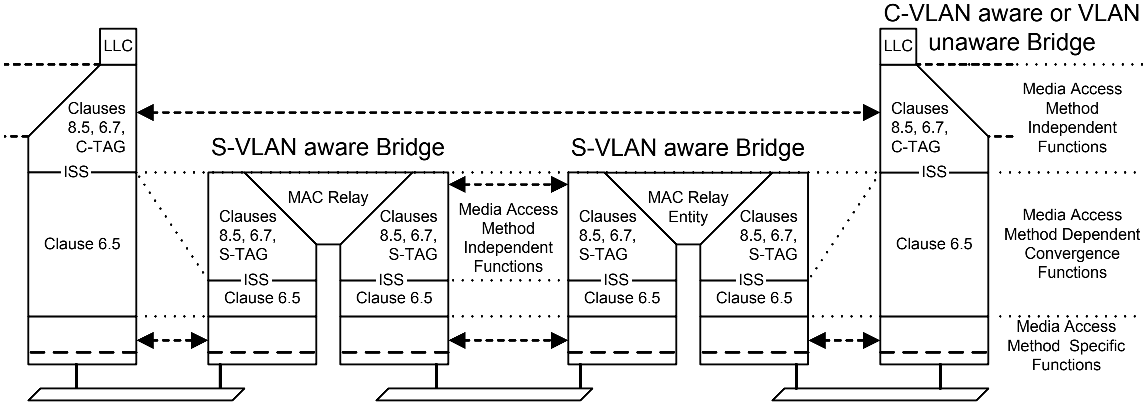
C-VLAN aware or VLAN unaware Bridge
LLC
LLC
Clauses 8.5, 6.7, C-TAG
Clauses 8.5, 6.7, C-TAG
S-VLAN aware Bridge
S-VLAN aware Bridge
ISS
MAC Relay
MAC Relay Entity
Clauses 8.5, 6.7, S-TAG
Clause 6.5
Media Access Method Independent Functions
Media Access Method Dependent Convergence Functions
Media Access Method Specific Functions

Figure 15-1—Internal organization of the MAC sublayer in a Provider Bridged Network

This clause discusses the following aspects of provider service provision

a) Service transparency
b) Customer service interfaces
c) Service instance segregation
d) Service instance selection and identification
e) Service priority selection
f) Service access protection

NOTE—This standard makes use of term ‘service’ as defined by the OSI Reference Model (IS0 7498). In this sense a service comprises a set of primitives and associated parameters, provided by one protocol layer in the architectural model to the protocol layer above, and the causal relationships between the primitives invoked by an upper layer protocol entity in one system with those resulting indications to a peer entity in another system. The term ‘service’ used by network providers, while including layering concepts, goes far beyond this formal definition, and commonly specifies some or all of the following: interfacing considerations across multiple protocol layers (including physical connectors, for example); selection of interface points; interfacing equipment; quality of service guarantees and measurement methods; charging methods and responsibilities; connectivity verification and other management tools; and regulatory issues. Many of these aspects lie outside the scope of this standard; the reader is referred to the Bibliography which includes references to completed and ongoing work in the MEF (Metro Ethernet Forum) and the ITU.

15.1 Service transparency

The operation of Provider Bridges and the networks they compose is, by design, largely transparent to Customer Bridges and Customer Bridged Local Area Networks as illustrated by Figure 15-1.

The service provided by Provider Bridges to attached Customer Bridges is transparent to the use of the MAC Service provided to end stations on LANs attached to the Customer Bridged LANs and transparent to the operation of media access method independent functions by the Customer Bridges.

The service is not transparent to the operation of media access method dependent convergence functions, specified in clause 6.5 of this standard, or to the operation of the media access method specific functions specified by standards for each media access method. Media access method dependent and specific functions operate between Bridges, whether Customer Bridges or Provider Bridges, attached to the same LAN. Where these functions make use of standard Group MAC Addresses, those addresses are included in the Reserved Addresses that are always filtered by Customer Bridges (Table 8-1) and by Provider Bridges (Table 8-2).

Frames transmitted and received by media access method independent functions particular to Provider Bridged Network operation are not forwarded by Customer Bridges between provider networks. Where these frames are addressed using standard Group MAC Addresses, those addresses are included in the Reserved Addresses that are always filtered by Customer Bridges (Table 8-1). In addition such frames may be filtered from Provider Bridged Ports that are connected to Customer Bridge Ports.

15.2 Customer service interfaces

A network provider can offer one or more types of customer service interfaces, each providing different capabilities for service selection, priority selection, and service access protection (<ref>). Some services interfaces are provided by the network provider operating systems that include C-VLAN aware Bridge components, or by a customer operating systems that include S-VLAN aware Bridge components. In all cases segregation of different customer instances is achieved at an interface wholly under the control of the provider by authentication and authorization of the attached customer systems, and by verification of customer provided parameters that provide service instance selection.

NOTE—The term “service access protection” has been introduced to describe the capability to provide access to a service over one or more access LANs with redundancy and rapid failover in case of failure of an access LAN or attached equipment.

Access to a given provider service instance can be provided through different types of customer interface.

15.3 Port-based service interfaces

The customer service interfaces that can be provided by a network of Provider Bridges are specified by reference to a Customer Network Port provided by the S-VLAN aware component of a Provider Bridge. The Customer Network Port provides a single service instance, as illustrated in Figure 15-2 and Figure 15-3. The attached customer system can be a bridge, a router, or an end-station.

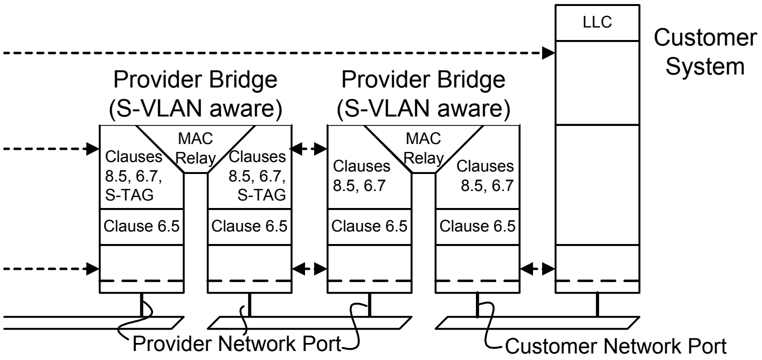

Figure 15-2—Port-based interface to a Provider Bridged Network

This interface is Port-based, i.e. customers select between and identify different service instances by associating each with a different Customer Network Port. Frames transmitted through a Customer Network Port by a C-VLAN aware customer system do not include an S-VID (3.46), but can be priority tagged.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

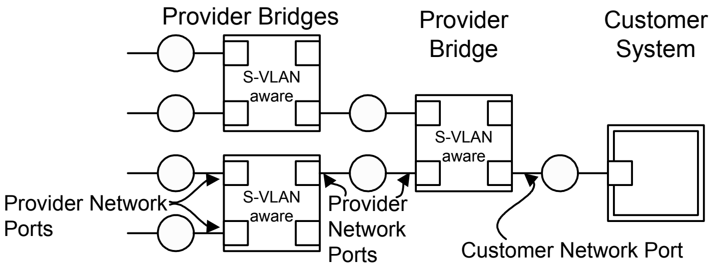

Figure 15-3—Port-based interface to a Provider Bridged Network

NOTE—The terms Customer Network Port, Customer Edge Port, Provider Network Port, and Provider Edge Port, do not refer to the ownership of equipment, or necessarily to differently implemented Ports, but to Ports that are configured to fulfil the requirements of precise roles within a structured provider network design. These requirements are developed in 15.6, 15.7, 15.8, and 15.9 below and in clause 16. All ‘Network’ Ports are part of S-VLAN aware Bridge components, and all ‘Edge’ Ports part of C-VLAN aware components, while all ‘Customer’ Ports receive data from a single customer inbound to the network and transmit data outbound to the network to a single customer. See Clause 3 for definitions.

15.4 Customer-tagged service interface

A C-TAGged service interface can be provided by a Provider Edge Bridge comprising one or more C-VLAN aware Bridge components attached to Port-based interfaces provided by a single S-VLAN aware component, as illustrated by Figure 15-4 and Figure 15-5.

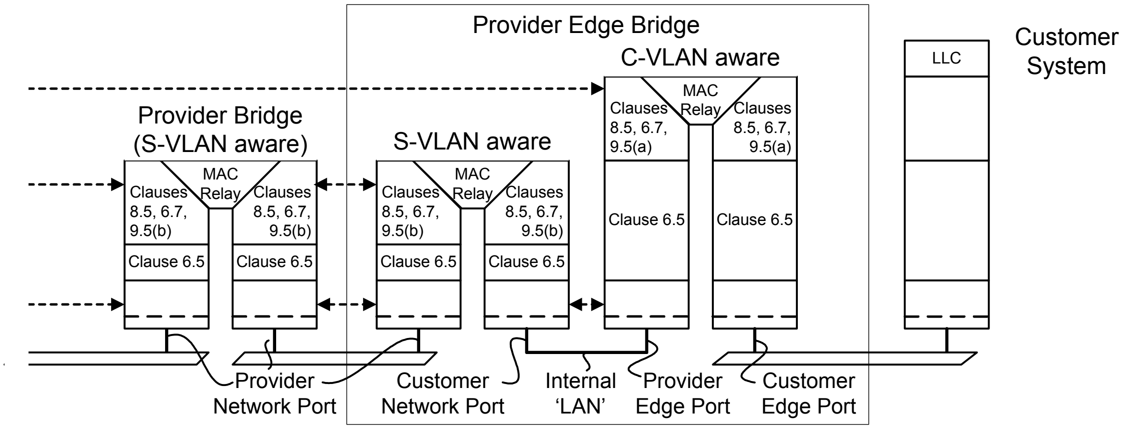

Figure 15-4—Customer-tagged interface to a Provider Bridged Network

The Customer-tagged service interface allows service instance selection and identification by C-VID. Each frame from the customer system is assigned to a C-VLAN and presented at one or more internal Port-based interfaces, each supporting a single service instance that the customer desires to carry that C-VLAN.

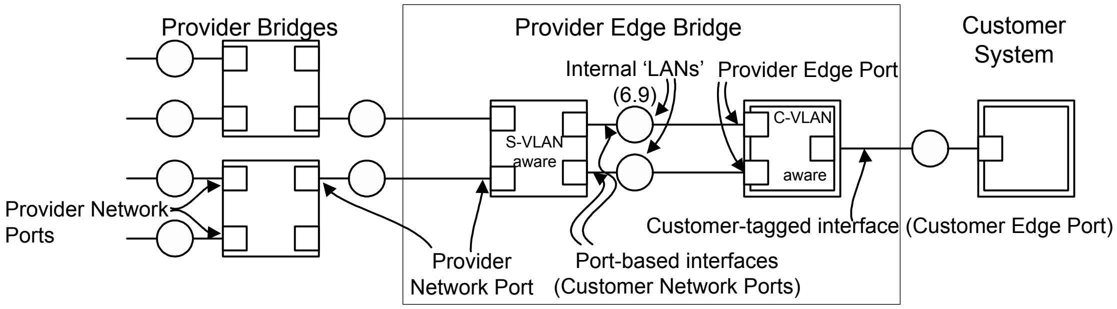

Figure 15-5—Customer-tagged interface to a Provider Bridged Network

Similarly frames from the provider network are assigned to an internal interface or ‘LAN’ on the basis of the S-VID. Since each internal interface supports a single service instance, the S-TAG can be, and is, removed at this interface. If multiple C-VLANs are supported by this service instance, the frames will have been C-TAGged with the possible exception of frames for a single C-VID. The C-VLAN aware component applies a PVID to untagged frames received on each internal ‘LAN’, allowing full control over the delivery of frames for each C-VLAN through the Customer Edge Port.

Each Provider Edge Bridge can support multiple Customer Edge Ports for multiple customers, each supported by a dedicated C-VLAN aware component as illustrated in Figure 15-6.

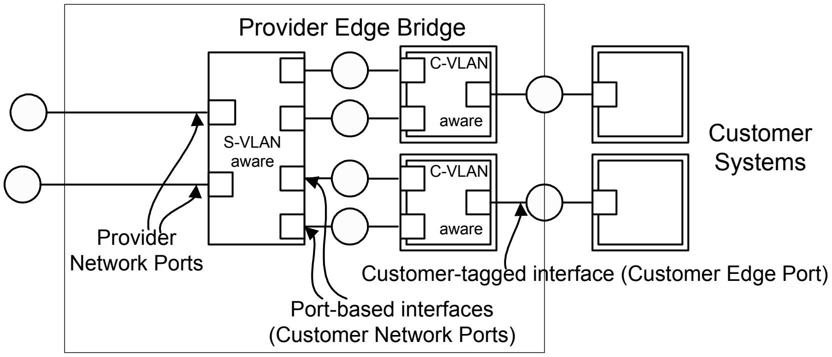

Figure 15-6—Customer Edge Ports

15.5 S-tagged service interface

An S-tagged service interface can be provided to a Provider Bridge or Provider Edge Bridge operated by a customer as illustrated by Figure 15-7 and Figure 15-8. The customer’s Bridges can in turn provide customer-tagged interfaces within the customer’s own network as described above (15.4).

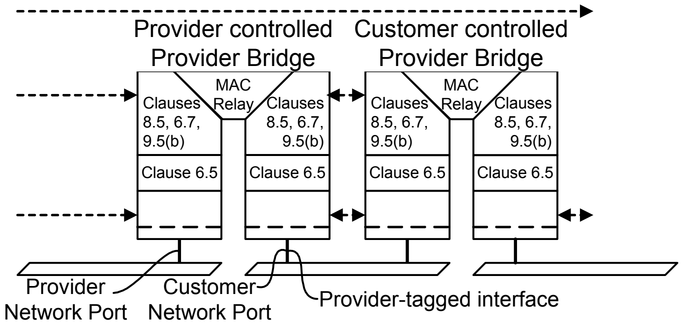

Figure 15-7—S-tagged interface to a Provider Bridged Network

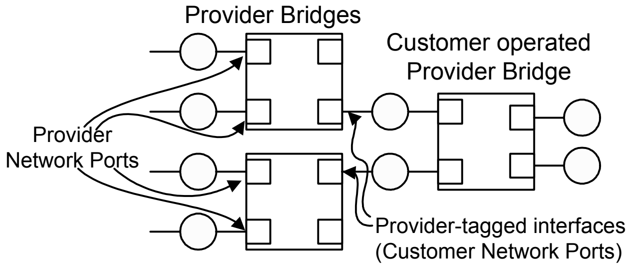

Figure 15-8—S-tagged interface to a Provider Bridged Network

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

15.6 Service instance segregation

Segregation of data frames associated with different MAC Service instances is achieved by supporting each service instance with a separate Service VLAN (S-VLAN) and ensuring that:

a) No customer data frames are transmitted through a Provider Network Port untagged, i.e. without a Service VLAN TAG (S-TAG).
b) No frames are accepted, i.e. received and relayed, from any customer system without first being subject to service instance selection.
c) No frames are delivered to any customer system without explicit service instance identification.
d) Prior to transmission through a Provider Network Port, customer data frames are received through a Customer Network Port within the provider network that is exclusively accessed by a single customer. The S-VIDs of all frames received through that Customer Network Port correspond to service instances that the customer is permitted to access.
e) Provider Bridges and the S-VLAN aware component of each of the Provider Edge Bridges within the provider network can only be directly controlled by the provider.
f) Only frames that have been transmitted through a Provider Network Port can be received through a Provider Network Port within the provider network.

15.7 Service instance selection and identification

Service instance selection is provided for Port-based service interfaces and Customer-tagged service interfaces by configuring a Customer Network Port with a PVID value corresponding to the S-VID used to identify the service instance and Acceptable Frame Types of *Admit Only untagged frames*.

If the interface is Port-based, the Customer Network Port is directly accessed by the customer.

If the interface is Customer-tagged, the Customer Network Port is internal to a Provider Edge Bridge, and frames are directed to the Port by including it in the Member Set for each of the C-VIDs assigned to frames to be conveyed by the service instance. Frames for at most one C-VLAN can be conveyed untagged over a single service instance, this is achieved by configuring a Static VLAN Registration Entry (8.9.2) for the Provider Edge Port specifying that frames with the VLAN’s C-VID are transmitted untagged.

NOTE—A Provider Edge Bridge can also provide a Port-based interface by configuring a Customer Edge Port with *Admit All Frames*, attaching the associated C-VLAN aware Bridge component to a single internal Customer Network Port with *Admit Only Untagged frames* set and the PVID value set as specified above.

Service instance selection is provided by the attached customer system for S-tagged interfaces. The Customer Network Port is configured with Enable Ingress Filtering and the Port is only included in the Member Set for S-VLANs corresponding to service instances that the customer is permitted to use.

15.8 Service priority selection

<<Some of the material currently in clause 16 is independent of how the service is provided, and should be relocated or duplicated here.>>

15.9 Service access protection

<<Some of the material currently in clause 16 is independent of how the service is provided, and should be relocated or duplicated here.>>

Copyright © 2004 IEEE. All rights reserved.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Insert the following Clause after the new proposed Clause 15.

16. Principles of Provider Bridged Network Operation

This clause establishes the principles and a model of Provider Bridged Network operation. It provides the context necessary to understand how the

a) operation of individual Provider Bridges (Clause 8),
b) configuration and management of individual Provider Bridges (Clause 12), and
c) management of Spanning Tree and VLAN Topologies within a Provider Network (Clauses 7, 11, 13)

support, preserve, and maintain the quality of each of the instances of the MAC Service offered to the Customers of the Provider Network (Clauses 6, 15) including:

d) independence of a particular instance of the MAC Service offered by a Provider from other instances of the service;
e) selection and identification of MAC Service instances by a customer system (Clause 15, 8.8);
f) maintenance of service availability in the event of the failure, restoration, removal, or insertion of particular LAN components connecting a customer network to a provider network.

A Provider Bridged Network is a Virtual Bridged Local Area Network that comprises Provider Bridges under the administrative control of a single provider, and includes the LANs attached to those Provider Bridges. The principal elements of provider network operation are those specified in Clause 7 for Virtual Bridged Local Area Networks in general, as amended by this Clause.

NOTE 1—Unless explicitly stated the use of the term ‘provider network’ in this Standard refers to a Provider Bridged Network. The term Provider Bridged Network is used exclusively to refer to networks configured and managed in accordance with the requirements of this Clause and comprising only (a) Provider Bridges (as specified in this Standard) and (b) communications media and equipment providing the Enhanced Internal Layer Service (as specified in Clause 15 and IEEE Standard 802.1Q Clause 6.4).While the requirements of Clause 15 are generally applicable to similar services, a generalized framework for all the network designs that could meet the requirements of Clause 15 is outside the scope of this Standard. While such a general framework might prove useful in the context of other equipment and services, the goal of this clause is to describe specific best practice for Provider Bridged Networks, to ensure that the requirements for Provider Bridge functionality are clear. Conformance of a Provider Bridge implementation to this standard does not require that the implementation be used as specified in this Clause, merely that it is capable of being so used.

NOTE 2—Within a provider network, an instance or instances of the MAC Service are reserved for the provider’s own use to configure and manage the network. All frames associated with such service instances, and that are not confined to an individual LAN, are subject to service instance selection, segregation, and identification as specified below.

16.1 Provider Network Overview

The principal elements of Provider Bridge Network operation comprise:

a) Service instance segregation within the provider network for customer frames (16.2).
b) Service instance selection on ingress to the provider network for each customer frame, and service instance identification for each customer frame on egress (16.3).
c) Resource allocation and configuration to provide service instance connectivity (16.4).
d) Filtering and relaying of group addressed customer frames on ingress to and within the provider network in support of protocols operated by customer systems (16.5).

and may also include:

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

e) Management of customer end station address learning (16.6).

f) Detection and signalling of changes in customer network active topology to facilitate relearning of customer end station location information (16.7).

g) Prevention of connectivity loops formed through attached networks (16.8).

16.2 Service Instance Segregation

Segregation of data frames associated with different MAC Service instances is achieved by supporting each service instance with a separate Service VLAN (S-VLAN) and ensuring that, within the provider network:

a) No customer data frames are transmitted untagged, i.e. without a Service VLAN TAG (S-TAG).

b) No frames are accepted, i.e. received and relayed, from any customer system without first being subject to service instance selection.

c) No frames are delivered to any customer system without explicit service instance identification.

16.3 Service Instance Selection and Identification

On ingress to a Provider Bridge the VLAN classification rules (8.9) are used to determine a VID for each customer data frame, by examining all frames for a Service Tag (5.8).

NOTE 1—Formally it is not a frame that is examined, but the parameters supplied with an ISS (Internal Sublayer Service) indication (see <Figure 6-1>, 6.4). An ISS M_UNITDATA.indication will only give rise to an EISS EM_UNITDATA indication with a non-NULL VID if the initial octets of the mac_service_data_unit (msdu) parameter of the M_UNITDATA.indication contain a Service Tag. A Customer TAG that occupies the initial octets, and any tag in subsequent octets, of the ISS msdu will remain as part of the EISS msdu. The distinction between examining the frame and the msdu is important if the service is being provided with underlying functionality that both uses and provides the ISS, e.g. MACsec.

The S-VID of the frame within the provider network is equal by default to the VID determined by the classification rules, but can be selected by managing the VID Translation Table for the Port (an optional capability (5.8, 5.9.2, 8.8)).

NOTE 3—If a VID Translation Table is used and the classification rules determine that the frame is untagged, the S-VID assigned will be the value in the table corresponding to the PVID assigned to the Port, i.e. the PVID is the assigned value prior to translation.

The Enable Ingress Filtering parameter (8.6.1) is set for all potential ingress Ports and the frame is discarded if the S-VID is not in the Member set.

NOTE 4—If a VID Translation Table is used, the ingress filtering check is applied to the S-VID resulting from translating the VID determined by the classification rules.

The network provider associates each potential ingress Provider Bridge Port with a customer point of attachment (see XX) and this information is used to determine the PVID, the contents of the VID Translation Table, and the Member set for the Port. The network provider configures the Member set so as to ensure that inappropriate access to any service instance cannot occur.

The optional VID Translation Table allows a network provider to assign S-VIDs independently from the VIDs used by a customers to identify MAC Service instances, and can permit a customer to identify the same service instance by different VIDs at different customer points of attachment.

NOTE 5—The means used by a network provider and a customer to determine the VIDs used by the customer to identify a given service instance are outside the scope of this Standard.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

The VID Translation Table is only used at Provider Bridge Ports that are designated as ingress Ports by the network provider. This ensures that S-VIDs assigned to service instances within the network do not change if the active topology of the network changes due to failure, removal, addition, or management of a network component.

NOTE 6—Use of the VID Translation Table allows service instances to be supported by concatenated provider networks, each with differing S-VID to service instance assignments, as translated on ingress to and egress from each individual provider network. The total number of service instances offered by a network provider or a number of cooperating network providers can be increased by this means, up to the point where each Bridge within the concatenated network carries the maximum possible number of S-VLANs. The means used to determine the S-VID to be associated with each service instance and LAN interconnecting individual provider networks is outside the scope of this standard. The S-VIDs used by configuration protocols, such as GVRP, within each provider network are those determined by the VID Translation Table after ingress and prior to service identification and egress.

NOTE 7—Use of the VID Translation Table also allows service provider networks conformant to this specification to be connected with alternate technologies outside the scope of this specification, but offering an equivalent Virtual Bridged Local Area Network service at Ports attached to LANs, each LAN providing segregated communications for up to the maximum number of service instances identifiable by a VID.

NOTE 8—The Enable Ingress Filtering parameter is not typically used within an individual provider network, as it limits the ability of the network to carry service instances following changes in the active topology. However it can be used to limit the reachability and accessibility of service instances used by the provider for network management of collocated equipment and to restrict service instances carried from one provider network to another.

A customer system attached to a provider network identifies the MAC Service instance for each received frame in the same way as frames transmitted using the same customer point of attachment, but not necessarily in the same way that the service instance is selected or identified at other customer points of attachment to the network.

NOTE 9—The VID used by the customer to select the service instance is not conveyed independently of the S-VID within the provider network. In consequence the VID Translation Table cannot be configured such that two distinct VIDs select the same service instance. This restriction supports the requirement (established in prior revisions of IEEE Standard 802.1Q) for unique and universal identification of the VLANs on any given LAN (i.e. on any given LAN all frames for a given VLAN are either tagged or untagged, and if tagged use the same VID). A single service instance can convey frames for multiple Customer VIDs — if no Service Tag is present all frames including their Customer VIDs, treated as user data by the ingress rules, will be assigned to a single service instance.

NOTE 10—The Provider Bridge capabilities mandated in this standard can be realized by equipment conformant to 802.1Q-2003 with the ability to select a different value of Ethertype to identify the VLAN tag. The optional capabilities do not require any port to simultaneously encode two distinct VIDs, except where one is incorporated in user data.

NOTE 11—Identification and separation of customer points of attachment and of internal communication capabilities are provided by Bridge Ports. If a given media access technology is capable of providing logical separation between users on a single media instance, attachment to the logically separated communications capabilities is represented as occurring through distinct Bridge Ports. The use of security associations that partition the traffic on a LAN by establishing exclusive groups is one way in which such logical separation might be realized on a common medium.

16.4 Service Instance Connectivity

The VLAN Topology of each S-VLAN is established by the mechanisms introduced in Clause 7.1 and Figure 7-1. The network provider can use and configure MSTP to provide a number of independent spanning tree active topologies and can assign each S-VLAN independently to one of these so as to best use the resources in the network. GVRP running in the context of each spanning tree active topology configures the extent of each S-VLAN to the subset of that active topology necessary to support connectivity between the customer points of attachment to the MAC Service instance provided, and can reconfigure that connectivity as required if the spanning tree active topology changes.

NOTE 1—Autoconfiguration of the extent of each S-VLAN is accomplished by the network provider configuring the GARP Administrative Control ‘Registration Fixed’ for the S-VLAN on each ingress Provider Bridge Port where the

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

corresponding MAC Service Instance can be selected. There are no direct customer manageable objects within a Provider Bridge, and the means by which a network provider agrees the service instances to be provided at any customer attachment point are outside the scope of this Standard.

The operation of MSTP within a provider network is independent of the operation of any spanning tree protocol within attached customer networks. This is achieved by using the Provider Bridge Group Address (Table 8-1) as the destination address of all MSTP BPDUs transmitted on Provider Bridge Ports and by setting the restrictedRole parameter (<ref>) for Provider Bridge Ports that provide ingress to the network. Frames received by Provider Bridge Ports and addressed to the Bridge Group Address are subject to service instance selection and relay in the same way as customer data frames.

NOTE 2—Where service instance selection is achieved by attached customer systems using either (a) distinct customer attachment points per service instance and thus implicitly service tagging all frames received by the provider network ingress port; or (b) explicit service tagging for each service instance; customer BPDUs addressed to the Bridge Group Address will be transparently conveyed, allowing the customer to use an instance of MSTP or RSTP that is completely independent of the provider network to establish and maintain full and loop-free connectivity of the customer connected networks and services. Use of GVRP by the customer is recommended so that the customer can limit the transmission of frames assigned to C-VLANs to the provider service instances required for C-VLAN connectivity.

The operation of GVRP within a provider network is independent of the operation of any configuration protocol within attached customer networks. The Provider GVRP Address (Table 8-1) is used as the destination address of all GARP PDUs transmitted in support of the GVRP Application. Frames received by Provider Bridge Ports and addressed to the GVRP Address (802.1D Table 12-1) are relayed by a Provider Bridge in the same way as are frames for unsupported GARP Applications. The GARP Administrative Control for each S-VLAN is either ‘Registration Fixed’ or ‘Registration Forbidden’ on all ingress Ports that are customer points of attachment to the provider network, so no information is received from any Provider GVRP PDU that has been erroneously transmitted by a customer system.

16.5 Filtering and relaying of group addressed customer frames

The set of end stations addressed by an invocation of the MAC Service with a Group destination MAC address can include any or all the stations attached to a Virtual Bridged Local Area Network, or can be limited by filtering in Bridges to those stations attached to the same individual LAN as the station issuing the service request (see clause 15, 8.9), to allow a protocol entity to communicate exclusively with those of its peers that are attached to the same LAN.

NOTE—Within this standard the term Local Area Network and the abbreviation LAN are used exclusively to refer to an individual LAN specified by a MAC technology without the inclusion of Bridges (3.18).

Media access method dependent and specific protocols are confined to the operation of a single LAN (15.1). Standard Group MAC addresses that support such protocols are specified as Provider Bridge Reserved Addresses (Table 8-2) and are always filtered (8.6.3).

Media access method independent customer protocols are always forwarded by Provider Bridges. Where a provider network offers a customer interface, such as a customer-tagged interface, at a port of a Bridge configured as a Customer Bridge, the service offered may appear to specific customer protocols in one of four ways

a) Transparent to the operation of the protocol between customer systems using the service provided, appearing as an individual LAN without Bridges; or,
b) Discarding frames, acting as a non-participating barrier to the operation of the protocol; or,
c) Peering, with a local protocol entity at the point of provider network ingress and egress, participating in and terminating the operation of the protocol; or,
d) Participation in individual instances of customer protocol.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

The applicability of the Transparent option (a) depends on the quality of service guarantees required by the protocol and their support by the particular provider network, and that of the Peering option (c) on the operation of the protocol without (or with acceptably small) requirements for coordinated management of its parameters between provider and customer. The operation of the bridge relay function is the same for Discarding (b) and Peering (c), the difference being whether a protocol entity for the specific protocol is attached to and supported by the provider bridge port.

16.6 Provider Learning of Customer End Station Addresses

Customer data frames for any given MAC Service instance are restricted to that part of the provider network that supports the VLAN topology of the associated S-VLAN as described above (16.7), and are further restricted by learning the source addresses of frames as described in Clauses 7 and 8.

In a Provider Bridged Network that commonly provides interfaces to each customer at a small fraction of the total number of customer interfaces provided, the requirement for learning customer end station addresses can be much reduced by applying enhanced filtering utility criteria (8.7.2). In particular learning can be restricted to the ingress and egress Provider Bridge Ports of each S-VLAN that connects only two customer points of attachment, or to the customer systems attached to those Ports.

16.7 Customer Network Topology Changes and Provider Address Learning

A single customer network with its own physical connectivity can be connected to a provider network at more than one customer point of attachment. Alternatively separate customer networks can be connected using distinct MAC service instances provided by one or more network providers. Such network topologies are designed to protect against the failure or removal of network components. The customer networks use their own instances of spanning tree protocol to configure and partition their active topology, so that the provider connectivity does not result in a data loop. Reconfiguration of a customer's active topology can result in the apparent movement of customer end stations from the point of view of the provider bridges as described in Clause 13.7, however the requirement for mutual independence of the distinct MAC Service instances that can be supported by a single provider spanning tree active topology does not permit either the direct receipt of provider topology change notifications from customer systems or the use of received customer spanning tree protocol topology change notifications to stimulate topology change signalling on a provider spanning tree. Instead an ingress Provider Bridge Port that detects a possible change in the reachability of end stations communicating on a given S-VLAN transmits a Customer Change Notification (CCN) BPDU (<ref>) tagged with that S-VID and a destination MAC Address of the CCN Address specified in Table 8-1.

<<The format of the CCN BPDU will have to be defined in Clause 9, and its processing in Clause 13 or a new Clause constructed mainly for the purpose. Input on details is requested, the editor being under the impression that we have agreed a need but not fully worked out mechanisms yet (the recommendations in http://www.ieee802.org/1/files/public/docs2003/unlearn-signaling.pdf are a fair summary of agreement to date). In particular it seems likely that: (1) most customer networks will not be using spanning tree to manage back up connectivity— (1a) many customers will use a private wire service to connect routers that will manage the connectivity (1b) customers selecting service by C-VID are unlikely to spare expensive connectivity 1 for 1; (2) most changes in active topology of a customer network will not actually involve end stations moves from the point of view of the provider— so flushing addresses in the provider network just because otherwise harmless customer BPDUs are seen could be a feature whose main attribute is that it is always turned off if we don't get it right. Leaving the generation of the CCN BPDU to customer vendor imagination could impose a significant load on providers, with the result that the BPDU is always ignored. The intention above is that the CCN Address be one of the addresses that is currently reserved and always filtered by Bridges, so customer bridges will discard it but Provider Bridges can pass the BPDU and snoop on it. No retry mechanism is used within the provider network to provide CCN signalling reliability (having the Provider Bridges pass and snoop the BPDU lessens the processing load and, appropriately implemented, should maximize the rate at which it progresses through the provider net, as well as allowing it to be forwarded by bridges that ignore it) whether the change detecting ingress port implements a retry mechanism depends on how detection takes place.

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

Does the functionality described here represent any advance on snooping customer BPDUs?

In addition to the above we have to decide whether this is an optional or mandatory capability.>>

16.8 Detection of connectivity loops through attached networks

The transmission and reception of MSTP BPDUs through provider network ingress ports will detect accidental direct connection of those ports, or their interconnection by a network that is transparent to frames with the Provider Bridge Group Address as the destination MAC address. However a network provider cannot rely on any customer network relaying such frames, and should develop a policy and mechanisms to deal with potential data loops that can arise if the attached customer systems do not correctly operate its own instance or instances of spanning tree protocol.

NOTE 1—Use of the restrictedRole parameter at ingress ports ensures that receipt of BPDUs addressed to the Provider Bridge Group Address cannot disrupt internal connectivity within the provider network.

NOTE 2—Customer BPDUs convey the Bridge Group Address in the destination MAC address field, and are subject to service instance selection and relay by the provider network in the same way as other customer data frames.

NOTE 3—Specification of provider policies, mechanisms, and heuristics used to detect or minimize the impact of data loops created by customer systems is not addressed this standard. They can include, but are not limited to, bandwidth limitation, charging policies, detection of the repetitive movement of the apparent location of customer stations, and customer agreement to allow the use of provider loop detection protocols by not filtering the associated frames.

NOTE 4—A data loop is not the only possible cause of excess bandwidth consumption by a given customer of a provider network, and the network provider is usually required to meet service guarantees to other customers irrespective of the cause of the excess bandwidth demand. Data loops are not a unique threat to satisfactory overall network performance. Their distinct characteristic is consumption of discretionary bandwidth without benefitting any customer. The customer that creates the loop can suffer particularly serious network degradation or excess cost as the provider limits the total bandwidth consumed by that customer. It is therefore in the interests of each individual customer and the provider to raise service satisfaction by preventing and detecting loops.

16.9 Provider Bridged Network Architecture

The architecture of a Provider Bridged Network is illustrated by Figure 16-1.

Customer equipment (CE) attaches, via one or more customer interface LANs (C-LANs), to provider operated Customer Premises Equipment (CPE), which in turn connects via one or more Access LANs (A-LANs), to Provider Premises Equipment (PPE) (i.e. to Provider Bridges and Service LANs (S-LANs) that are secured so that only the network provider can manage the reception, transmission, and relay of frames between Provider Bridges).

The arbitrary physical network topology of the PPE and the internal S-VLAN connectivity that it provides to support segregated instances (16.2) of the MAC Service, is designed and managed (16.4) by the network provider to meet bandwidth and service availability requirements at the Provider Interface Ports (PIPs). Application of the service VLAN ingress and egress rules at the PIPs in support of service instance selection and identification (16.3) ensures that frames cannot be transmitted or received on any service instance by any customers equipment without prior agreement with the provider.

While the application of the ingress and egress rules, together with the use of the MSTP restrictedRole and GVRP registration controls (16.4), permit providers to allow direct attachment of customer operated equipment to access LANs connected to PPE PIPs there are commonly other reasons, such as OAM&P support of Access LANs, why CPE is mandated. Locating provider bridging functions within CPE, as illustrated by but not limited to the examples CPE1 through CPE3 in Figure 16-1, can be used to partition and enhance provider network access functionality in support of:

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

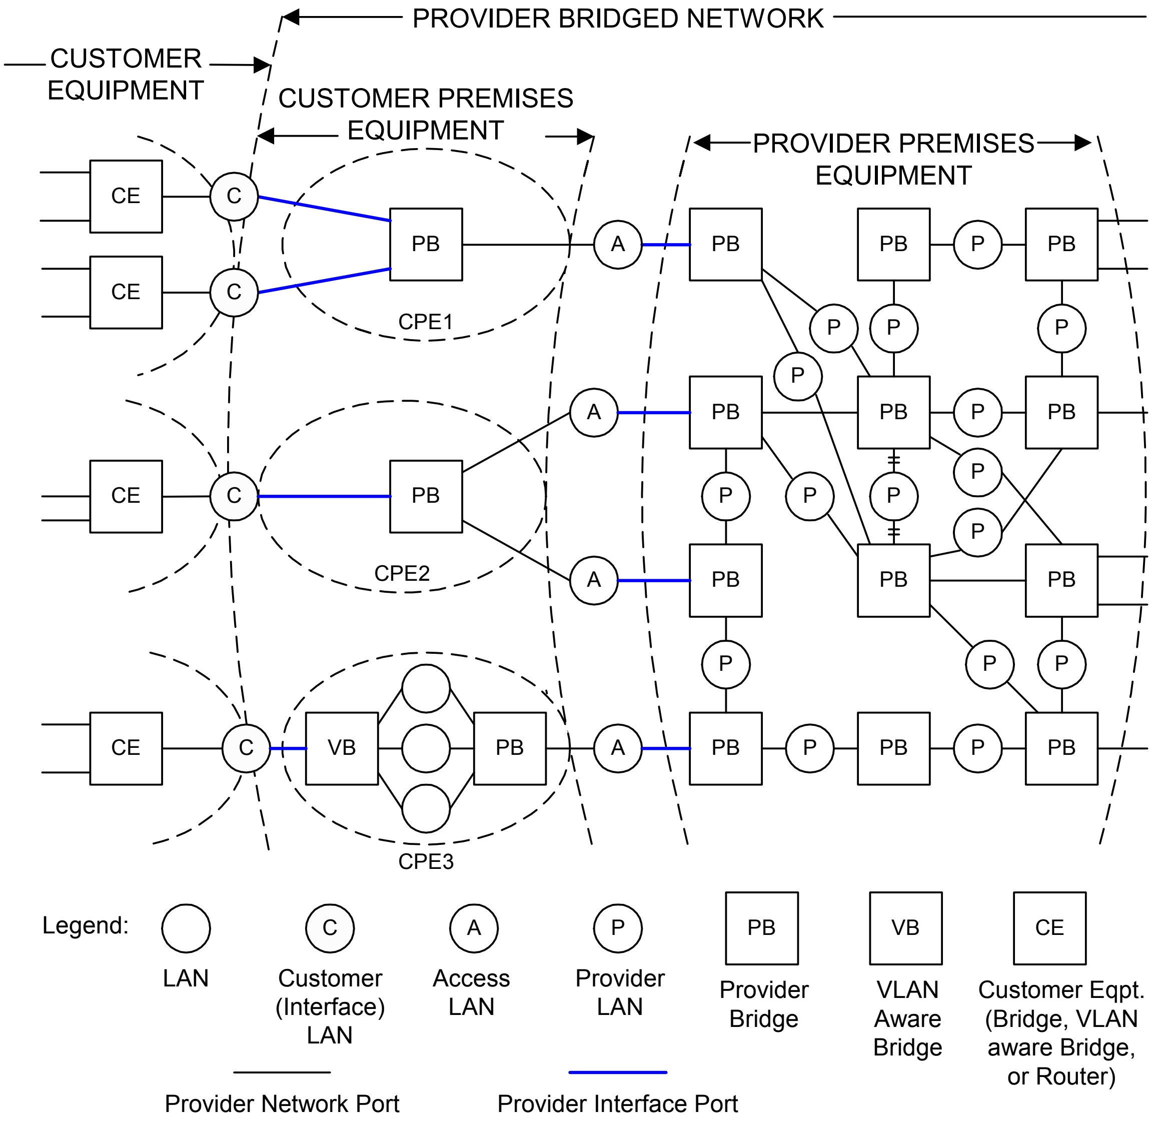

Figure 16-1—Provider Bridged Network Architecture with CPE Examples

a) service instance multiplexing on a single access LAN;
b) provision of resilient, and optionally physically route diverse, access;
c) reduced management of customer use of multiple service instances;
d) selective multiplexing of customer VLANs onto service instances by C-VID;
e) limiting the functionality of PPE Provider Bridges;
f) reliable identification of the customer point of attachment.

without requiring customer systems to handle provider tagged frames or understand internal details of the provider network.

CPE1 uses physically separate customer interface LANs to provide separate service instances to customer systems, provider tagging these service instances to multiplex them over a single access LAN. CPE1 can be managed by the network provider to use S-VIDs that do not require use of a VID Translation Table by the PPE PIP. Alternatively the latter can translate to remove the need for such management, with all CPE using the same S-VIDs in the same way.

CPE2 uses two access LANs to provide resilient connectivity to PIPs on distinct PPE Provider Bridges. Participation by the CPE in the provider spanning tree(s) protects against failure of an access LAN, of one of the PPE PIPs, or of some of the internal physical connectivity within the provider network. Use of the MSTP restrictedRole parameter by PPE PIPs ensures that the CPE cannot disrupt the S-VLANs-tag topologies within the PPE. Where multiple service instances are provided at a single customer point of attachment, both access links can be used.

CPE3 supports arbitrary selection by user tagging on the customer interface LAN of multiple provider tagged service instances on the access LAN. Individual user VLANs can be conveyed on one or many provider service instances, and this distribution of user VLANs can be changed in response to changes in the connectivity offered by the service instances through customer systems at other points of customer attachment. The traffic for a particular user VLAN can be the majority of that carried on a specific service provider instance; specifying that frames for that user VLAN are untagged on the real or logical connection between the VLAN aware Bridge and Provider Bridge functions of CPE3 that corresponds to that service instance allows those frames to be carried through the provider network without a C-VID following the S-VID if the overhead of conveying the additional octets is a concern.

NOTE—The VLAN aware Bridge functionality, VB in the figure, needs to participate in the customer's spanning trees to realize all the functionality described. While this may be done by agreeing a fairly simple set of rules concerning Port Path Cost assignment on the basis of whether assignment of a user VLAN to a particular provider service instance is to be preferred, retained as a backup option, or forbidden, it may well be more satisfactory for the customer to operate the functionality shown in CPE3. This functionality would then connect directly to a PPE PIP, without intervening CPE, or connect through simple or dual connected CPE using provider tagging on the customer interface LAN to select the service instance.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Annex A (normative)

PICS proforma

<<An update PICS (or an additional PICS to cover Provider Bridges) will be prepared once the material in the main body of the document is stable. The PICS does not create requirements or options but simply provides a checklist for those defined elsewhere.>>

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

Annex B (informative)

Shared and independent VLAN learning

This Amendment makes no changes to Annex B.

Annex C (informative)

Media access method dependent aspects of VLAN support

This Amendment makes no changes to Annex C.

Annex D (informative)

Background to VLANs

This Amendment makes no changes to Annex D.

Annex E (informative)

Interoperability considerations

This Amendment makes no changes to E.1, E.2, E.3, E.4 and E.5.

E.1 Intermixing IEEE Std 802.1Q Version 1.0 bridges with future IEEE P802.1Q bridges

This Amendment makes no changes to E.6.1 and E.6.2.

<<Check this.>>

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Annex F

(informative)

Frame translation considerations

This Amendment makes no changes to Annex F.

Annex G

(informative)

Differences between 802.1s and 802.1w state machines

This Amendment makes no changes to Annex G.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Annex G (informative)

Priority and drop precedence

This standard allows priorities, flow metering, queue assignment, and queue service disciplines to be managed to best support the goals of network administrators. This annex documents the rationale for the recommended and default priority to traffic class mappings in Table 8-3, and the encoding of priority and drop eligibility in Table 6-4 and Table 6-5.

Classification of user data frames into a small number of behavior aggregates, together with aggregate dependent forwarding behavior in each Bridge, allows signalling of application requirements to the network. Frame classification, aggregate bandwidth metering and policing, and drop precedence marking, also facilitate network scaling and provision of services to independent customers through allocation of those functions to appropriate Bridges in the network.

While there are many possible ways to classify frames and to specify forwarding behaviors, it is widely appreciated that a set of well known and easily understood defaults can facilitate interoperability and the deployment of services. The defaults described in this annex and supported by this standard were chosen to support integrated and differentiated services, to minimize the burden of management, to reduce the possibility of misconfiguration and out of order frame delivery, and to provide useful service without management in many networks.

This standard mandates support for strict priority frame transmission (8.6.6), but permits the use of additional traffic class based transmission selection algorithms. The default assignments of frames to traffic classes on the basis of frame priority, as described in this annex, also support the use of frame priority to select general traffic class based forwarding behavior.

NOTE — The bibliography (Annex <ref>) provides references to the IETF work on integrated and differentiated services.<<Do this.>>

G.1 Traffic types

A full description of the QoS needs of applications and network services is too complex to be represented by a simple number 0 through 7. The pragmatic aim of traffic classification is to simplify requirements to preserve the high-speed, low-cost characteristics of Bridges. At the margin, potential bandwidth efficiency is traded for simplicity and higher speed operation—historically a good decision in the LAN.

The following list of traffic types, each of which can benefit from simple segregation from the others, are of general interest:

a) Network Control — characterized by a guaranteed delivery requirement to support configuration and maintenance of the network infrastructure.
b) Internetwork Control — in large networks comprising separate administrative domains there is typically a requirement to distinguish traffic supporting the network as a concatenation of those domains from the Network Control of the immediate domain.
c) Voice — characterized by less than 10 ms delay, and hence maximum jitter (one way transmission through the LAN infrastructure of a single campus).
d) Video — characterized by less than 100 ms delay, or other applications with low latency as the primary QoS requirement.
e) Critical Applications — characterized by having a guaranteed minimum bandwidth as their primary QoS requirement and subject to some form of admission control to ensure that one system or application does not consume bandwidth at the expense of others. The admission control mechanism

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

can range from pre-planning of the network requirement at one extreme to bandwidth reservation per flow at the time the flow is started at the other.

f) Excellent Effort — or “CEO’s best effort,” the best-effort type services that an information services organization would deliver to its most important customers.
g) Best Effort — for default use by un-prioritized applications with fairness only regulated by the effects of TCP’s dynamic windowing and retransmission strategy.
h) Background — bulk transfers and other activities that are permitted on the network but that should not impact the use of the network by other users and applications.

G.2 Managing latency and throughput

Use of priorities and queuing by traffic classes, each class encompassing one or more priorities, facilitates improvement and management of latency and throughput, allowing QoS goals to be supported at higher levels of network loading than would otherwise be possible.

Congestion, resulting in QoS degradation, is not equally likely at all Bridges in a network. Transient traffic patterns are likely to result in congestion in only a few Bridges at a time, while over an extended period momentary congestion is more likely to occur in the network core than at Bridge Ports attached to one or a relatively small number of end stations. Use of fewer traffic classes for those Ports can lower the cost of implementation and management, and this standard facilitates the use of Bridges supporting differing numbers of classes within a single network that delivers a consistent set of QoS parameters for each frame priority level. Although the number of traffic classes supported by each Bridge Port along the path taken by a given flow of data can vary, the default mappings of priorities to classes ensures that frame ordering is preserved as required by Clause 8.6.6.

With few classes the focus is on meeting latency requirements—the bandwidth surplus required in a bursty data environment to guarantee sub-10 ms delays without distinct traffic classification is uneconomically large. As the number of traffic classes that can be used increases, the focus shifts to managing throughput.

The simple default queue servicing policy defined in this standard, strict priority, supports latency management. Active management of bandwidth sharing necessarily requires some management.

G.3 Traffic type to traffic class mapping

Table G-1 groups the traffic types introduced above to match the number of traffic class queues supported by a Bridge Port. Each grouping of types is shown as {*Distinguishing type*, Type, Type, . . .}. The “distinguishing type” is not treated in any way differently in a Bridge, but is italicized here to illustrate, for any given number of queues, which traffic types have driven the allocation of types to classes.

The step by step breaking out of traffic types as more classes are available proceeds as follows:

a) With a single queue, there are no choices. All traffic is Best Effort.
b) To support integrated services in the presence of bursty best effort data, it is necessary to segregate all the time-critical traffic. The amount of high priority traffic will be restricted by the need to support low latency for Voice, which becomes the defining type for the additional queue.
c) Two queues may be adequate for Bridge Ports attaching to end stations. The stability of the network as a whole may be unaffected by the performance of configuration protocols on those Ports, and in-band management of the Bridge itself is likely to occur through another Port. For Bridges within the network infrastructure a further queue is used to isolate Network Control from the user data traffic.
d) Traffic for business Critical Applications is separated from Best Effort to allow a bandwidth guarantee to be provided.

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

Table G-1—Traffic type to traffic class mapping

Number of queues	Traffic types
1	{*Best Effort,* Background, Excellent effort, Critical Applications, Voice, Video, Internetwork Control, Network Control}

2	{*Best Effort,* Background, Excellent effort, Critical Applications}
	{*Voice,* Video, Internetwork Control, Network Control}

3	{*Best Effort,* Background, Excellent effort, Critical Applications}
	{*Voice,* Video}
	{*Network Control,* Internetwork Control}

4	{*Best Effort,* Background}
	{*Critical Applications,* Excellent effort}
	{*Voice,* Video}
	{*Network Control,* Internetwork Control}

5	{*Best Effort,* Background}
	{*Critical Applications,* Excellent effort}
	{*Voice,* Video}
	{*Internetwork Control*}
	{*Network Control*}

6	{*Background*}
	{*Best Effort*}
	{*Critical Applications,* Excellent effort}
	{*Voice,* Video}
	{*Internetwork Control*}
	{*Network Control*}

7	{*Background*}
	{*Best Effort*}
	{*Excellent effort*}
	{*Critical Applications*}
	{*Voice,* Video}
	{*Internetwork Control*}
	{*Network Control*}

8	{*Background*} {*Best Effort*} {*Excellent effort*} {*Critical Applications*}
	{*Video*} {*Voice*} {*Internetwork Control*} {*Network Control*}

e) The queue separation so far provided can support a large network. The next queue is allocated to distinguishing Internetwork Control traffic from local Network Control.

f) Background is separated from Best Effort to minimize the effect of bulk transfers on ordinary network use.

g) Excellent Effort is separated from Critical Applications, either to provide a simple superior service based on policy controlled access or to provide an additional segregated bandwidth guarantee.

h) The final provides increased network utilization as the higher bandwidth traffic associated with Video is no longer given the same latency guarantee as Voice.

The above is an illustrative rather than definitive description of the logic of allocating traffic types to classes. The mappings in Table 8-3 support the assignment of other semantics to each of the traffic types identified by priority values, e.g. the identification of all three of the illustrative types “Video”, “Critical Applications”, and “Excellent Effort” with assured forwarding classes that provide segregated bandwidth guarantees. However alternate semantics should take into account the service provided by Bridges with limited traffic class queuing, e.g. of the foregoing only “Video” would receive priority treatment by default at a Bridge Port supporting queuing for only two classes.

The order of assignment of traffic types to queues has also taken into account requirements to encode a drop eligible bit within the priority field for certain traffic types, as described below (G.5).

G.4 Traffic types and user priority values

Table G-2 shows the correspondence between traffic types and priority values used to select the defaults in Table 8-3. The default priority used for transmission by end stations is 0. Changing this default would result in confusion, and likely interoperability problems. At the same time the default traffic type is definitely Best Effort. 0 is thus used both for default priority and for Best Effort, and Background is associated with a user priority value of 1. This means that the value 1 effectively communicates a lower priority than 0.

Table G-2—Traffic type acronyms

user_priority	Acronym	Traffic type
1	BK	Background
0 (Default)	BE	Best Effort
2	EE	Excellent Effort
3	CA	Critical Applications
4	VI	“Video,” < 100 ms latency and jitter
5	VO	“Voice,” < 10 ms latency and jitter
6	IC	Internetwork Control
7	NC	Network Control

Table G-3 summarizes Table G-1, showing just the defining traffic types. By maintaining the groupings of types established for a given number of queues for all less numbers, the table preserves the order of frames of any given type, independent of the number of queues.

Table G-3—Defining traffic types

<table>
<tr><th>Number of queues</th><th colspan="8">Defining traffic type</th></tr>
<tr><td>1</td><td colspan="8">BE</td></tr>
<tr><td>2</td><td colspan="4">VO</td><td colspan="4">BE</td></tr>
<tr><td>3</td><td colspan="2">NC</td><td colspan="2">VO</td><td colspan="4">BE</td></tr>
<tr><td>4</td><td colspan="2">NC</td><td colspan="2">VO</td><td colspan="2">CA</td><td colspan="2">BE</td></tr>
<tr><td>5</td><td>NC</td><td>IC</td><td colspan="2">VO</td><td colspan="2">CA</td><td colspan="2">BE</td></tr>
<tr><td>6</td><td>NC</td><td>IC</td><td colspan="2">VO</td><td colspan="2">CA</td><td>BE</td><td>BK</td></tr>
<tr><td>7</td><td>NC</td><td>IC</td><td colspan="2">VO</td><td>CA</td><td>EE</td><td>BE</td><td>BK</td></tr>
<tr><td>8</td><td>NC</td><td>IC</td><td>VO</td><td>VI</td><td>CA</td><td>EE</td><td>BE</td><td>BK</td></tr>
</table>

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

G.5 Supporting drop precedence

It is often desirable to meter traffic from different users to ensure fairness or to meet bandwidth guarantees, however dropping all traffic in excess of a committed rate is likely to result in severe under-utilization of the networks, since most traffic sources are bursty in nature. At the same time it is burdensome to meter traffic at all points in the network where bandwidth contention occurs. One solution is to mark those frames in excess of the committed rate as drop eligible on admission to the network.

This standard allows drop eligibility to be conveyed separately from priority in Service VLAN TAGs (S-TAGs) so that all of the previously introduced traffic types can be marked as drop eligible. To provide compatibility with previous revisions of this standard while allowing drop eligibility to be conveyed in Customer VLAN-TAGs (C-TAGs), this standard also allows a subset of the priorities to be conveyed along with drop eligibility marking for some of those priorities within the Priority Code Point (PCP) field of both S-TAGs and C-TAGs.

G.6 Priority code point allocation

Four or more priorities can be conveyed with a single level of drop eligibility. Table 6-4 and Table 6-5 specify encoding and decoding for five through eight priorities. The tables are consistent with the following step by step reduction in the number of distinct priorities to provide drop eligibility for certain traffic types:

a) If eight distinct priorities are required, drop eligibility cannot be encoded in the PCP.

b) Drop eligibility in support of QoS maintenance for traffic conforming to a committed rate is most effective when used to support time critical traffic. If seven priorities, one of which can be marked as drop eligible are required, then the traffic class queuing distinction between Voice and Video is sacrificed to providing drop eligibility for the combined traffic types. This does not preclude marking all Video traffic as drop eligible upon ingress to a network, so as to provide the same guarantee to Voice as a distinct priority.

c) The distinctions between Critical Applications and Excellent Effort, and between Best Effort and Background traffic types is removed to provide drop eligibility for Critical Applications and for Best Effort.

d) Although the use of four priorities, each with drop eligibility, is possible, it is not recommended. Combining Network Control with Internetwork Control could only serve to increase the guarantees provided to the latter at the expense of the former, which if not delivered threatens the stability of the overall network in any case. Moreover both traffic types should be supportable with guaranteed bandwidth if the network is to be operated successfully.

Choosing first to combine Video and Voice, and then Critical Applications with Excellent Effort (for six queues), provides consistency with the allocation of priorities to traffic classes in the absence of drop eligibility. Bridges that do not implement drop eligibility, but are configured to use the same number or fewer traffic classes, will not misorder frames. If such a bridge is configured to use only five traffic classes, and in accordance with Table 8-3, it will not misorder frames with a priority code point encoded using any of the alternatives provided by Table 6-4.

G.7 Interoperability

Encoding of drop eligibility within the PCP, as opposed to separately within the VLAN TAG (as is possible for the S-TAG), provides interoperability with Bridges that interpret all VLAN TAGs as including the CFI bit specified for the C-TAG and do not recognize the DE bit in the S-TAG. It also provides compatibility of PCP encoding with certain uses of the MPLS EXP bits.

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

However the requirement to provide different combinations of priorities with drop eligibility within the confines of the PCP, means that priority and drop eligibility information can be lost for frames traversing a network if the combinations used on individual LANs differ. Use of the DE bit does not suffer from this problem.

If Bridges attached to the same LAN encode and decode the PCP differently then incorrect priority values can be attributed and subsequent misordering of frames can occur. Misordering will not occur, with the recommended priority to traffic class mappings of Table 8-3 and the recommended PCP encoding and decagons in Table 6-4 and Table 6-5, if the Bridge performing the incorrect decoding assumes fewer priorities than are actually encoded or if all Bridges subsequently transited by the frame use the same number or fewer traffic classes than those used for the encoding. However incorrect decoding will in all probability affect other service guarantees that the network is intended to support. If a Bridge can be used in a network that encodes drop eligibility in the PCP, and there is any likelihood of the Bridge being brought into service prior to network dependent service level configuration, then five priorities three with drop eligibility (5P3D encoding and decoding) should be used. Bridges that do not support drop precedence should be configured to support five or fewer traffic classes in the same circumstances.

The use of separate Priority Encoding and Priority Decoding Tables for each Bridge Port allows adaptation between the Priority Drop Precedence (PD) scheme in one domain of the network and the PD scheme used in another to be accomplished in only one of any pair of Bridges, each serving as boundary of its domain, connected by a point-to-point LAN. However if more than two Bridges are attached to a LAN, all need to use the same encoding so that each of its recipients can assign the correct priority to the frame.

The default PCP encoding and decoding, as documented in Table 6-4 and Table 6-5, are reproduced in Table G-4 and Table G-5, with the addition of the default allocation of priorities to traffic classes to the latter.

The discussion of traffic types in G.1 above, and the suggested association of each with a priority value, differs from the similar discussion in IEEE Std 802.1D-2004 Annex G and prior revisions of that standard. The latter was developed contemporaneously with IETF Intserv and predates Diffserv. The discussion in this Annex better aligns with current practice, in particular Voice is associated with priority 5, matching the setting of the relevant bits for Expedited Forwarding (EF) in the DSCP (Differentiated Services Code Point) for IP and in the common use of the EXP bits for MPLS. Standards for DSCPs are believed to be the prime reference for use of priority by end stations, and there is no direct change to the behavior of Bridge implementations conforming to this standard as a result of this change.

The priority to traffic class mappings in Table 8-3 differ in one minor respect from those specified in prior revisions of this standard and in IEEE Std 802.1D-2004 and its prior revisions. Priority value 2 was previously described as ‘Spare’ and positioned lower than 0 (Best Effort) in priority order. This change may result in networks including Bridges conformant to prior revisions of this standard, and implementing four or more traffic classes, providing less than expected priority to traffic described in this annex as Excellent Effort, and misordering drop eligible traffic for Critical Applications. The change was necessitated by the requirement to provide best use of the PCP when encoding drop eligibility, and facilitated by the low demand for two distinct priorities of lesser importance than Best Effort.

Table G-4—Priority encoding

priority drop_eligible		7	7DE	6	6DE	5	5DE	4	4DE	3	3DE	2	2DE	0	0DE	1	1DE
PCP	8P0D	7	7	6	6	5	5	4	4	3	3	2	2	0	0	1	1
	7P1D	7	7	6	6	5	4	5	4	3	3	2	2	0	0	1	1
	6P2D	7	7	6	6	5	4	5	4	3	2	3	2	0	0	1	1
	5P3D	7	7	6	6	5	4	5	4	3	2	3	2	0	1	0	1

Table G-5—Priority Code Point decoding

<table>
<tr><th colspan="2">PCP</th><th>7</th><th>6</th><th>5</th><th>4</th><th>3</th><th>2</th><th>0</th><th>1</th></tr>
<tr><td rowspan="4">priority drop_eligible</td><td>8P0D</td><td>7</td><td>6</td><td>5</td><td>4</td><td>3</td><td>2</td><td>0</td><td>1</td></tr>
<tr><td>7P1D</td><td>7</td><td>6</td><td>4</td><td>4DE</td><td>3</td><td>2</td><td>0</td><td>1</td></tr>
<tr><td>6P2D</td><td>7</td><td>6</td><td>4</td><td>4DE</td><td>2</td><td>2DE</td><td>0</td><td>1</td></tr>
<tr><td>5P3D</td><td>7</td><td>6</td><td>4</td><td>4DE</td><td>2</td><td>2DE</td><td>0</td><td>0DE</td></tr>
<tr><td rowspan="8">number of traffic classes</td><td>1</td><td colspan="8">BE</td></tr>
<tr><td>2</td><td colspan="4">VO</td><td colspan="4">BE</td></tr>
<tr><td>3</td><td colspan="2">NC</td><td colspan="2">VO</td><td colspan="4">BE</td></tr>
<tr><td>4</td><td colspan="2">NC</td><td colspan="2">VO</td><td colspan="2">CA</td><td colspan="2">BE</td></tr>
<tr><td>5</td><td>NC</td><td>IC</td><td colspan="2">VO</td><td colspan="2">CA</td><td colspan="2">BE</td></tr>
<tr><td>6</td><td>NC</td><td>IC</td><td colspan="2">VO</td><td colspan="2">CA</td><td>BE</td><td>BK</td></tr>
<tr><td>7</td><td>NC</td><td>IC</td><td colspan="2">VO</td><td>CA</td><td>EE</td><td>BE</td><td>BK</td></tr>
<tr><td>8</td><td>NC</td><td>IC</td><td>VO</td><td>VI</td><td>CA</td><td>EE</td><td>BE</td><td>BK</td></tr>
</table>

Copyright © 2004 IEEE. All rights reserved.

This is an unapproved IEEE Standards Draft, subject to change.

Annex Z (informative)

COMMENTARY

<<Editor's Note: This is a temporary Annex.>>

Z.1 Frame Classification

See the editors' note to clause 8.9 for the present.

Z.1.1 Disposition

Z.2 Editor's notes of interest

Relocated here from where they originally appeared, these may have continuing value and get turned into commentary points, or even get fleshed out into a permanent annex.

<<The biggest customer 'security' risk is not that the ingress controls etc. intended for a particular customer point of attachment are set up incorrectly, as these can be automated, but that physical cabling errors are made in the field. Apart from training the provider's own craft, third parties often have to be used to make cross connects etc. The use of security like mechanisms to securely identify and check what is being connected to what, or better yet to install the appropriate parameters on the basis of what has been connected to what, is a large step forward for operational practice. Quick and accurate verification of where the physical connection goes to several miles away without organizing personnel at both ends of the link cuts the number of truck rolls. This from experience. P802.1ab should have an eventual role to play here in addition to 'security' developments. As always we need to refrain from speculative references to standards in progress, so the intent of clause XX —when it gets written—will be to be self contained and complete, if necessarily weaker than desirable in consequence.>>

<<In Clause 16: The terms 'customer' and 'customer point of attachment' have been used with some care (even if this is not apparent) in the preceding text. Network deployment experience has shown that it is necessary to be very specific about who can do what where, surfacing questions that can be missed or avoided by attempting to be sufficiently but superficially architectural. The easy case is where a single customer contracts for the provided service instance and provides the attaching equipment and secure arrangements for its identification and connection. A more complex but by no means unreal case is where one of the service provider's customers arranges and contracts for connections to be used between another customer and several others. Such cases play havoc with simple expressions of such notions as "a customer should only be able to connect to his or her own service instances". Fortunately, since this amendment/ standard is not the place to set out or develop the necessary business relationships that underpin successful service deployment, most of the real life discussion that occurs is an attempt to bridge terminological gaps that conceal furious agreement. Consequently no harm is done here by not spelling out the ways in which agreements to provide and gain access to services can be formulated and require secure support. Finding forms of words that convey our general intent and do not cause negative comment is sufficient.>>

<<MSTP is currently limited to supporting 64 spanning trees. From the point of view of the specification this could easily be increased to support the full VLAN range (circa 4094) by allowing any BPDU to contain a subset of the MSTIs. Ensuring that this capability would be implemented well and not degenerate into a BPDU per MSTI is less easy and it may be that 64 spanning trees is quite adequate for some time since each can support any of the VLANs, and 64 POPs (located in colos or COs) is a significant network. Input requested. Agreed in March 2003 802.1 meeting that 64 is sufficient for time being, any increase to be a separate project. This note to move to Annex Z.>>

Copyright © 2004 IEEE. All rights reserved.
This is an unapproved IEEE Standards Draft, subject to change.

<<In Clause 16: The provision of additional mechanisms to convey the spanning tree information pertinent to a single C-VLAN on the service instance selected for that VLAN has been deliberately left outside the scope of the current draft, and needs a compelling user demand if it is to be specified. One to one sparing of provider service instances seems an expensive choice, while one for many can be handled as for the editor's note immediately prior to this. Protecting the network against loops set up by customers (as opposed to picking paths from a number of choices) has to be handled independently of protocol (see below).>>

<<The text in 16.9 does not state how Provider Premises Equipment is secured, and hence avoids inconclusive references to work in progress. Obviously physical security plays or can play a role. In this context the term 'Provider Premises Equipment' simply means 'equipment and communications capabilities that have been secured by the provider'. Any other term that can be used to mean the same thing without long explanations or future looking statements would be an acceptable title for this part of the diagram. While the use of physical security has its limitations and there is always the possibility that the competition can dig up the street outside your office to access the fiber, the assurances provided by a physically secured Provider Bridged Network are equivalent to those provided by other technologies ("SONET-equivalent security"). Of course such layer 2 security will be augmented by higher layer methods for protecting sensitive traffic and by farewells. This is not to downplay the heightened risks posed by shared media.>>