Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.



# IEEE Standards for Local and Metropolitan Area Networks:

# Virtual Bridged Local Area Networks

Sponsor

**LAN MAN Standards Committee**
**of the**
**IEEE Computer Society**

Approved 8 December 1998

**IEEE-SA Standards Board**

**Abstract:** This standard defines an architecture for Virtual Bridged LANs, the services provided in Virtual Bridged LANs, and the protocols and algorithms involved in the provision of those services.
**Keywords:** local area networks, MAC Bridge management, media access control bridges, virtual LANs

---

The Institute of Electrical and Electronics Engineers, Inc.
345 East 47th Street, New York, NY 10017-2394, USA

Copyright © 1999 by the Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 8 March 1999. Printed in the United States of America.

*Print:* ISBN 0-7381-1537-1 SH94709
*PDF:* ISBN 0-7381-1538-X SS94709

*No part of this publication may be reproduced in any form, in an electronic retrieval system or otherwise, without the prior written permission of the publisher.*

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

**IEEE Standards** documents are developed within the IEEE Societies and the Standards Coordinating Committees of the IEEE Standards Association (IEEE-SA) Standards Board. Members of the committees serve voluntarily and without compensation. They are not necessarily members of the Institute. The standards developed within IEEE represent a consensus of the broad expertise on the subject within the Institute as well as those activities outside of IEEE that have expressed an interest in participating in the development of the standard.

Use of an IEEE Standard is wholly voluntary. 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.

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.

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 all 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 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.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# Introduction to IEEE Std 802.1Q-1998

(This introduction is not part of IEEE Std 802.1Q-1998, IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks.)

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

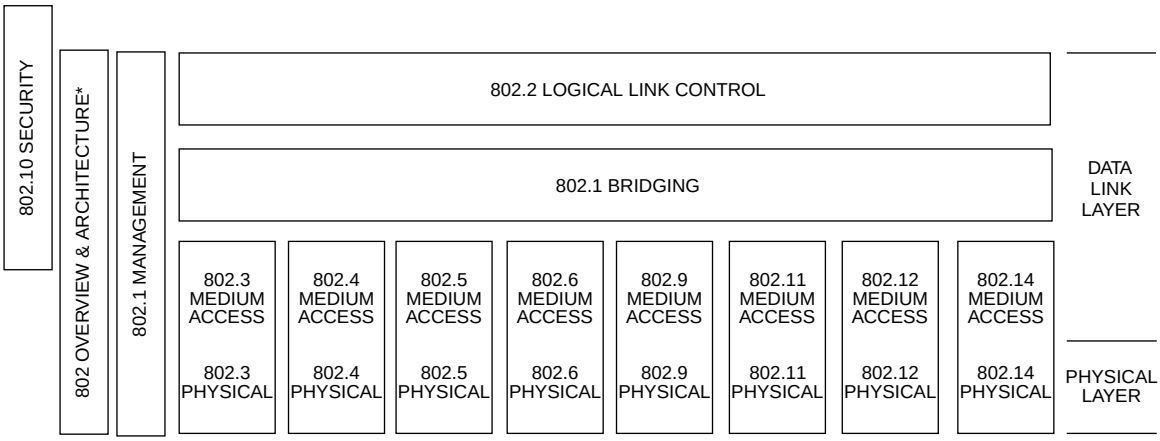


* Formerly IEEE Std 802.1A.

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

The standards defining the technologies noted above are as follows:

- IEEE Std 802 — *Overview and Architecture.* This standard provides an overview to the family of IEEE 802 Standards.
- ANSI/IEEE Std 802.1B and 802.1k [ISO/IEC 15802-2] — *LAN/MAN Management.* Defines an OSI management-compatible architecture, and services and protocol elements for use in a LAN/MAN environment for performing remote management.
- ANSI/IEEE Std 802.1D [ISO/IEC 15802-3] — *Media Access Control (MAC) Bridges.* Specifies an architecture and protocol for the interconnection of IEEE 802 LANs below the MAC service boundary.
- ANSI/IEEE Std 802.1E [ISO/IEC 15802-4] — *System Load Protocol.* Specifies a set of services and protocol for those aspects of management concerned with the loading of systems on IEEE 802 LANs.
- ANSI/IEEE Std 802.1F — *Common Definitions and Procedures for IEEE 802 Management Information*
- ANSI/IEEE Std 802.1G [ISO/IEC 15802-5] — *Remote Media Access Control (MAC) bridging.* Specifies extensions for the interconnection, using non-LAN communication technologies, of geographically separated IEEE 802 LANs below the level of the logical link control protocol.
- ANSI/IEEE Std 802.2 [ISO/IEC 8802-2] — *Logical link control*
- ANSI/IEEE Std 802.3 [ISO/IEC 8802-3] — *CSMA/CD access method and physical layer specifications*

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 3 of 211

- ANSI/IEEE Std 802.4 [ISO/IEC 8802-4] — *Token passing bus access method and physical layer specifications*
- ANSI/IEEE Std 802.5 [ISO/IEC 8802-5] — *Token ring access method and physical layer specifications*
- ANSI/IEEE Std 802.6 [ISO/IEC 8802-6] — *Distributed Queue Dual Bus (DQDB) access method and physical layer specifications*
- ANSI/IEEE Std 802.9 [ISO/IEC 8802-9] — *Integrated Services (IS) LAN Interface at the Medium Access Control (MAC) and Physical (PHY) Layers*
- ANSI/IEEE Std 802.10 — *Interoperable LAN/MAN Security*
- ANSI/IEEE Std 802.11 [ISO/IEC DIS 8802-11] — *Wireless LAN Medium Access Control (MAC) and physical layer specifications*
- ANSI/IEEE Std 802.12 [ISO/IEC 8802-12] — *Demand-priority access method, physical layer and repeater specifications*

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

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

The following additional working group has authorized standards projects under development:

- IEEE 802.14 — *Standard Protocol for Cable-TV Based Broadband Communication Network*

## Conformance test methodology

An additional standards series, identified by the number 1802, has been established to identify the conformance test methodology documents for the 802 family of standards. Thus, the conformance test documents for 802.3 are numbered 1802.3.

## IEEE Std 802.1Q-1998

The MAC Bridge standardization activities that resulted in the development of ISO/IEC 10038: 1993 introduced the concept of Filtering Services in Bridged LANs, and mechanisms whereby filtering information in such LANs may be acquired and held in a Filtering Database.

ISO/IEC 15802-3, a revision of ISO/IEC 10038, extends this concept of Filtering Services in order to define additional capabilities in Bridged LANs aimed at

a) The provision of expedited traffic capabilities, to support the transmission of time-critical information in a LAN environment;

b) The use of signalled user priority information as the basis for identifying expedited classes of traffic;

c) The provision of filtering services that support the dynamic definition and establishment of Groups in a LAN environment, and the filtering of frames by Bridges such that frames addressed to a particular Group are forwarded only on those LAN segments that are required in order to reach members of that Group;


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

d) The provision of a Generic Attribute Registration Protocol (GARP) that is used to support the mechanism for providing Group filtering capability, and is also made available for use in other attribute registration applications.

This standard makes use of the concepts and mechanisms of LAN Bridging that were introduced by ISO/IEC 15802-3, and defines additional mechanisms that allow the implementation of Virtual Bridged LANs. The following are described:

e) Virtual LAN Services in Bridged LANs;

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

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

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

i) The ability to support end-to-end signalling of user priority information regardless of the intrinsic ability of the underlying MAC protocols to signal user priority information;

j) The GARP VLAN Registration Protocol (GVRP) that allows distribution and registration of VLAN membership information (the protocol described makes use of the GARP protocol defined in ISO/IEC 15802-3);

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

This standard contains state-of-the-art material. The area covered by this standard is undergoing evolution. Revisions are anticipated within the next few years to clarify existing material, to correct possible errors, and to incorporate new related material. Information on the current revision state of this and other IEEE 802 standards may be obtained from

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

IEEE 802 committee working documents are available from

IEEE Document Distribution Service
AlphaGraphics #35 Attn: P. Thrush
10201 N. 35th Avenue
Phoenix, AZ 85051
USA

Copyright © 1999 IEEE. All rights reserved.



Dell et al., Ex. 1008
Page 5 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## Participants

The following is a list of participants in the Interworking activities of the IEEE 802.1 Working Group. Voting members at the time of publication are marked with an asterisk (*).

**William P. Lidinsky,** *Chair**
**Mick Seaman,** *Chair, Interworking Task Group**
**Editing Team:**
**Tony Jeffree*,** *Coordinating Editor*
**Anil Rijsinghani***, **Richard Hausman***, **Michele Wright***, **Paul Langille***, **P. J. Singh***

Steve Adams*
Stephen Ades
Ken Alonge
Floyd Backes*
John Bartlett*
Les Bell*
Avner Ben-Dor
Michael Berger*
James S. Binder*
David Brady
Martin Brewer
Bill Bunch*
Bob Cardinal
Paul Carroll*
Jeffrey Catlin*
Dennis Cave
Alan Chambers*
Steve Chan
David W. Chang*
Ken Chapman
Hon Wah Chin*
Chi Chong
Chris Christ*
Paul Congdon*
Glenn Connery*
David Cullerot*
Ted Davies*
Andy Davis
David Delaney*
Prakash Desai
Jeffrey Dietz*
Kurt Dobbins
Peter Ecclesine*
J. J. Ekstrom*
Norman W. Finn*
Yishai Fraenkel
Paul Frantz
Lars Henrik Frederiksen*
Anoop Ghanwani*
John Grinham
Steve Haddock
Sharam Hakimi*
John Hart*
Scott Harvell
Wayne Hathaway
Vic Hayes
David Head*
Gaby Hecht
Deepak Hegde*
Ariel Hendel
John Hickey
David Hollender
Steve Horowitz*
Michelle Hsiung
Rita Hunt
David Husak
Altaf Hussain*
Vipin K. Jain*
Neil Jarvis
Allen Kasey
Toyoyuki Kato*
Hal Keen*
Kevin Ketchum*
Keith Klamm*
Bruce Kling*
Walter Knitl
Dan Krent*
Paul Kummer
Paul Lachapelle*
Bill Lane
Johann Lindmeyr*
Gary Littleton
Robert D. Love
Andy Luque
Peter Martini
Keith McCloghrie
Martin McNealis
Milan Merhar*
John Messenger*
Colin Mick
Amol Mitra
Yaron Nachman*
Krishna Narayanaswamy*
Paul Nikolich
Lawrence Ng*
Henry Ngai*
Eugene O'Neil
Satoshi Obara*
Toshio Ooka*
Jörg Ottensmeyer*
Luc Pariseau*
Yonadav Perry
John Pickens*
Gideon Prat
Kirk Preiss
Steve Ramberg*
Shlomo Reches*
Dick Reohr
James Richmond*
Doug Ruby
Ray Samora
Ayman Sayed*
Rich Seifert
Lee Sendelbach*
Himanshu Shah*
Phil Simmons*
K. Karl Shimada
Fred Shu
Rosemary V. Slager*
Alexander Smith*
Andrew Smith*
Larry Stefani*
Stuart Soloway*
Sundar Subramaniam*
Richard Sweatt
Robin Tasker*
Fouad Tobagi
Naoki Tsukutari
Dhadesugoor Vaman
Steve Van Seters*
Dono van-Mierop*
John Wakerly*
Peter Wang*
Philip Wang
Y. C. Wang*
Trevor Warwick*
Bob Watson
Alan Weissberger
Glenn Wenig
Keith Willette*
Michael Witkowski*
Edward Wong*
Michael D. Wright*
Allen Yu*
Wayne Zakowski*


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

The following persons were on the balloting committee of IEEE Std 802.1Q:

Corey Anderson
Kit Athul
Thomas W. Bailey
Peter K. Campbell
James T. Carlo
David E. Carlson
Alan M. Chambers
R. Allan Cobb
Robert S. Crowder
Thomas J. Dineen
Peter Ecclesine
Philip H. Enslow
Changxin Fan
John W. Fendrich
Michael A. Fischer
Harvey A. Freeman
Gautam Garai
Harry Gold
Julio Gonzalez-Sanz
Stephen R. Haddock
Allen W. Hathaway
Donald N. Heirman
Raj Jain
Neil A. Jarvis
Anthony A. Jeffree
Robert W. Klessig
Stephen Barton Kruger
William G. Lane
David J. Law
Lanse M. Leach
Randolph S. Little
Peter Martini
Milan Merhar
John L. Messenger
Bennett Meyer
David S. Millman
John E. Montague
Wayne D. Moyers
Shimon Muller
Paul Nikolich
Charles Oestereicher
Roger Pandanda
John R. Pickens
Vikram Punj
Edouard Y. Rocher
James W. Romlein
Floyd E. Ross
Michael Salzman
Norman Schneidewind
Mick Seaman
Rich Seifert
Leo Sintonen
Michael A. Smith
Patricia Thaler
Geoffrey O. Thompson
Mark-Rene Uchida
Oren Yuen

When the IEEE-SA Standards Board approved this standard on 8 December 1998, it had the following membership:

**Richard J. Holleman,** *Chair* **Donald N. Heirman,** *Vice Chair*
**Judith Gorman,** *Secretary*

Satish K. Aggarwal
Clyde R. Camp
James T. Carlo
Gary R. Engmann
Harold E. Epstein
Jay Forster*
Thomas F. Garrity
Ruben D. Garzon
James H. Gurney
Jim D. Isaak
Lowell G. Johnson
Robert Kennelly
E. G. "Al" Kiener
Joseph L. Koepfinger*
Stephen R. Lambert
Jim Logothetis
Donald C. Loughry
L. Bruce McClung
Louis-François Pau
Ronald C. Petersen
Gerald H. Peterson
John B. Posey
Gary S. Robinson
Hans E. Weinrich
Donald W. Zipse

*Member Emeritus

Kristin M. Dittmann
*IEEE Standards Project Editor*

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 7 of 211

# Contents


1. Overview ..... 1

1.1 Scope ..... 1
1.2 VLAN aims and benefits ..... 1
1.3 Relationship with ISO/IEC 15802-3 ..... 2

2. References ..... 4

3. Definitions ..... 7

3.1 Ethernet Type-encoding ..... 7
3.2 Logical Link Control (LLC) encoding ..... 7
3.3 Frame ..... 7
3.4 Frame relay ..... 7
3.5 Independent Virtual Local Area Network (VLAN) Learning (IVL) ..... 7
3.6 Independent Virtual Local Area Network (VLAN) Learning (IVL) Bridge ..... 7
3.7 Legacy region ..... 8
3.8 Priority-tagged frame ..... 8
3.9 Shared Virtual Local Area Network (VLAN) Learning (SVL) ..... 8
3.10 Shared Virtual Local Area Network (VLAN) Learning (SVL) Bridge ..... 8
3.11 Shared Virtual Local Area Network (VLAN) Learning (SVL)/Independent Virtual Local Area Network (VLAN) Learning (IVL) Bridge ..... 8
3.12 Tagged frame ..... 8
3.13 Tag header ..... 8
3.14 Untagged frame ..... 9
3.15 Virtual Bridged Local Area Network (LAN) ..... 9
3.16 Virtual Local Area Network (VLAN) ..... 9
3.17 VLAN-aware ..... 9
3.18 VLAN-tagged frame ..... 9
3.19 VLAN-unaware ..... 9
3.20 Terms used in ISO/IEC 15802-3 ..... 9

4. Abbreviations ..... 10

5. Conformance ..... 11

5.1 Static conformance requirements ..... 11
5.2 Options ..... 12
5.3 Protocol Implementation Conformance Statement (PICS) ..... 12
5.4 MAC-specific bridging methods ..... 12

6. Architectural overview ..... 13

6.1 Configuration ..... 13
6.2 Distribution of configuration information ..... 13
6.3 Relay ..... 13
6.4 Filtering Database architecture ..... 14
6.5 VLAN classification ..... 15
6.6 Rules for tagging frames ..... 16
6.7 Spanning Tree ..... 16




Copyright © 1999 IEEE. All rights reserved.



Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

7. Support of the MAC Service in VLANs ..... 18

7.1 Enhanced Internal Sublayer Service provided within VLAN Bridges ..... 18
7.2 Support of the Internal Sublayer Service by IEEE Std 802.3 (CSMA/CD) ..... 22

8. Principles of operation ..... 23

8.1 Bridge operation ..... 23
8.2 Bridge architecture ..... 25
8.3 Model of operation ..... 26
8.4 Port States, Port parameters, Active Ports, and the active topology ..... 30
8.5 Frame reception ..... 32
8.6 The ingress rules ..... 33
8.7 The Forwarding Process ..... 34
8.8 The egress rules ..... 38
8.9 Frame transmission ..... 39
8.10 The Learning Process ..... 39
8.11 The Filtering Database ..... 40
8.12 Bridge Protocol Entity and GARP Protocol Entities ..... 53
8.13 Bridge Management ..... 53
8.14 Addressing ..... 53

9. Tagged frame format ..... 62

9.1 Overview ..... 62
9.2 Transmission and representation of octets ..... 65
9.3 Structure of the tag header ..... 65

10. Use of GMRP in VLANs ..... 73

10.1 Definition of a VLAN Context ..... 73
10.2 GMRP Participants and GIP Contexts ..... 73
10.3 Context identification in GMRP PDUs ..... 74
10.4 Default Group filtering behavior and GMRP propagation ..... 74

11. VLAN topology management ..... 76

11.1 Static and dynamic VLAN configuration ..... 76
11.2 GARP VLAN Registration Protocol ..... 77
11.3 Conformance to GVRP ..... 81
11.4 Procedural model ..... 83

12. VLAN Bridge Management ..... 93

12.1 Management functions ..... 93
12.2 Managed objects ..... 94
12.3 Data types ..... 94
12.4 Bridge Management Entity ..... 95
12.5 MAC entities ..... 98
12.6 Forwarding process ..... 98
12.7 Filtering Database ..... 102
12.8 Bridge Protocol Entity ..... 106
12.9 GARP Entities ..... 110
12.10 Bridge VLAN managed objects ..... 112

Copyright © 1999 IEEE. All rights reserved.



Dell et al., Ex. 1008
Page 9 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Annex A (normative) PICS proforma ..... 121

A.1 Introduction ..... 121
A.2 Abbreviations and special symbols ..... 121
A.3 Instructions for completing the PICS proforma ..... 122
A.4 PICS proforma for IEEE Std 802.1Q-1998 ..... 124
A.5 Major capabilities and options ..... 125
A.6 Relay and filtering of frames ..... 128
A.7 Maintenance of filtering entries in the Filtering Database ..... 130
A.8 Addressing ..... 131
A.9 Spanning Tree Algorithm ..... 132
A.10 Bridge Management ..... 136
A.11 Performance ..... 138
A.12 GARP and GMRP ..... 139
A.13 VLAN support ..... 140

Annex B (informative) Shared and Independent VLAN Learning ..... 143

B.1 Requirements for Shared and Independent Learning ..... 143
B.2 Configuring the Global VLAN Learning Constraints ..... 148
B.3 Interoperability ..... 150

Annex C (informative) MAC method dependent aspects of VLAN support ..... 151

C.1 The variables ..... 151
C.2 Bridging functions ..... 153
C.3 Frame formats ..... 156
C.4 Procedures for tagging, untagging, and relaying tagged frames ..... 162
C.5 Frame translations for different MAC methods ..... 166
C.6 Field definitions ..... 180

Annex D (informative) Background to VLANs ..... 182

D.1 Basic VLAN concepts ..... 182
D.2 Relationship with the Port-based VLAN model ..... 184

Annex E (informative) Interoperability considerations ..... 186

E.1 Requirements for interoperability ..... 186
E.2 Homogenous 802.1Q Bridged LANs ..... 187
E.3 Heterogeneous Bridged LANs: intermixing ISO/IEC 15802-3 (D) and 802.1Q (Q) Bridges ..... 189
E.4 Heterogeneous Bridged LANs: intermixing ISO/IEC 11802-5 and 802.1Q Bridges ..... 190
E.5 Heterogeneous Bridged LANs: intermixing 802.1Q Bridges with ISO/IEC 15802-3 Bridges ..... 195
E.6 Intermixing 802.1Q Version 1.0 Bridges with future 802.1Q Bridges ..... 196

Annex F (informative) Frame translation considerations ..... 198


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## Figures

Figure 6-1 VLAN architectural framework ..... 13
Figure 7-1 Relationships between MAC Entity, ISS, E-ISS, and MAC Relay Entity ..... 18
Figure 8-1 Example of a Bridged LAN ..... 26
Figure 8-2 Bridge ports ..... 27
Figure 8-3 VLAN Bridge architecture ..... 27
Figure 8-4 Relaying MAC frames ..... 28
Figure 8-5 Observation of network traffic ..... 29
Figure 8-6 Operation of inter-bridge protocol ..... 29
Figure 8-7 Operation of the GARP protocol ..... 30
Figure 8-8 Illustration of the detailed operation of the Forwarding Process ..... 35
Figure 8-9 Logical separation of points of attachment used by Higher Layer Entities and the MAC Relay Entity ..... 58
Figure 8-10 Effect of control information on the forwarding path ..... 59
Figure 8-11 Per-Port points of attachment ..... 59
Figure 8-12 Single point of attachment—relay permitted ..... 60
Figure 8-13 Single point of attachment—relay not permitted ..... 60
Figure 8-14 Ingress/egress control information in the forwarding path ..... 61
Figure 9-1 Tag header formats ..... 66
Figure 9-2 Ethernet-encoded TPID format ..... 67
Figure 9-3 SNAP-encoded TPID format ..... 67
Figure 9-4 Tag Control Information (TCI) format ..... 67
Figure 9-5 E-RIF Route Control (RC) field ..... 70
Figure 10-1 Example of GMRP propagation in a VLAN context ..... 75
Figure 11-1 Operation of GVRP ..... 78
Figure B-1 Connecting independent VLANs—1 ..... 144
Figure B-2 Connecting independent VLANs—2 ..... 145
Figure B-3 Duplicate MAC Addresses ..... 145
Figure B-4 Asymmetric VLAN use: “multi-netted server” ..... 146
Figure C-1 Services and environments ..... 152
Figure C-2 Heterogeneous Bridging functions ..... 153
Figure C-3 Tagged frames on 8802-5 Token Ring LANs ..... 157
Figure C-4 Tagged frames on FDDI LANs ..... 158
Figure C-5 Tagged frames on 802.3/Ethernet LANs ..... 159
Figure C-6 Translation between E-C-T/C,U and E-C-T/C,T ..... 167
Figure C-7 Translation between E-C-T/C,U and E-C-T/R,T ..... 168
Figure C-8 Translation between L-C-T/C,U and L-C-T/C,T ..... 169
Figure C-9 Translation between L-C-T/C,U and L-C-T/R,T ..... 170
Figure C-10 Translation between E-X-X/R,U and E-X-X/C,T ..... 171
Figure C-11 Translation between E-X-X/R,U and E-X-X/R,T (8802-5 & SR FDDI) ..... 172
Figure C-12 Translation between E-X-X/R,U and E-X-X/R,T (transparent FDDI) ..... 173
Figure C-13 Translation between L-X-X/R,U and L-X-X/C,T ..... 174
Figure C-14 Translation between L-X-X/R,U and L-X-X/R,T (8802-5 & SR FDDI) ..... 175
Figure C-15 Translation between L-X-X/R,U and L-X-X/R,T (transparent FDDI) ..... 176
Figure C-16 Relaying Ethernet Type-encoded tagged frames ..... 177
Figure C-17 Relaying LLC-encoded tagged frames ..... 178
Figure C-18 Relaying tagged frames between transparent and SR forms ..... 180
Figure C-19 SNAP-encoded Protocol Type format ..... 180
Figure D-1 Port-based VLANs ..... 182
Figure D-2 Hybrid Links ..... 183
Figure E-1 Static filtering inconsistency ..... 188
Figure E-2 Interoperability with ISO/IEC 15802-3 Bridges: example 1 ..... 189
Figure E-3 Interoperability with ISO/IEC 15802-3 Bridges: example 2 ..... 190
Figure E-4 Interoperability between Q versions 1 and 2 ..... 196

Copyright © 1999 IEEE. All rights reserved.



Dell et al., Ex. 1008
Page 11 of 211

## Tables


Table 1-1 Relationship between this standard and ISO/IEC 15802-3 ..... 3
Table 5-1 Support requirements for insertion, removal, and modification of tag headers ..... 11
Table 8-1 User priority regeneration ..... 33
Table 8-2 Recommended user priority to traffic class mappings ..... 36
Table 8-3 Outbound access priorities ..... 38
Table 8-4 Ageing time parameter value ..... 44
Table 8-5 Combining Static and Dynamic Filtering Entries for an individual MAC Address ..... 50
Table 8-6 Combining Static Filtering Entry and Group Registration Entry for “All Group Addresses” and “All Unregistered Group Addresses” ..... 50
Table 8-7 Forwarding or Filtering for specific group MAC Addresses ..... 51
Table 8-8 Determination of whether a Port is in a VLAN’s member set ..... 52
Table 8-9 Standard LLC address assignment ..... 54
Table 8-10 Reserved addresses ..... 55
Table 8-11 Addressing bridge management ..... 57
Table 9-1 802.1Q Ethernet Type allocations ..... 67
Table 9-2 Reserved VID values ..... 69
Table 11-1 GVRP Application address ..... 80


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# IEEE Standards for Local and Metropolitan Area Networks:

# Virtual Bridged Local Area Networks

## 1. Overview

IEEE 802 Local Area Networks (LANs) of all types can be connected together with Media Access Control (MAC) Bridges, as specified in ISO/IEC 15802-3[1]. This standard defines the operation of Virtual LAN (VLAN; see 3.16) Bridges that permit the definition, operation, and administration of VLAN topologies within a Bridged LAN (see 3.20) infrastructure.

### 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 methods, this standard specifies a general method for the operation of MAC Bridges that support the construction of VLANs (see 3.16). To this end it

a) Positions the function of VLANs within an architectural description of the MAC Sublayer;
b) 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 Independent functions that provide frame relay (3.4) in the VLAN Bridge;
d) Specifies the operation of the functions that provide frame relay in the VLAN Bridge;
e) Defines the structure, encoding, and interpretation of the VLAN control information carried in tagged frames (3.12) in a VLAN;
f) Specifies the rules that govern the insertion and removal of VLAN control information in frames;
g) Specifies the rules that govern the ability to carry data in Canonical format and Non-canonical format using different LAN MAC methods;

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

h) Establishes the requirements for, and specifies the means of, automatic configuration of VLAN topology information;
i) Defines the management functionality that may be provided in a VLAN Bridge in order to facilitate administrative control over VLAN operation;
j) Specifies the requirements to be satisfied by equipment claiming conformance to this standard.

### 1.2 VLAN aims and benefits

VLANs aim to offer the following benefits:

a) VLANs are supported over all IEEE 802 LAN MAC protocols, and over shared media LANs as well as point-to-point LANs.

[1]Information about references can be found in Clause 2.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 13 of 211

b) VLANs facilitate easy administration of logical groups of stations that can communicate as if they were on the same LAN. They also facilitate easier administration of moves, adds, and changes in members of these groups.
c) Traffic between VLANs is restricted. Bridges forward unicast, multicast, and broadcast traffic only on LAN segments that serve the VLAN to which the traffic belongs.
d) As far as possible, VLANs maintain compatibility with existing bridges and end stations.
e) If all Bridge Ports are configured to transmit and receive untagged frames (3.14), bridges will work in plug-and-play ISO/IEC 15802-3 mode. End stations will be able to communicate throughout the Bridged LAN.

NOTE—Whether a VLAN Bridge will operate in ISO/IEC 15802-3 mode depends upon the configuration of the various Port parameters (8.4) and the Filtering Database (8.11). A VLAN Bridge in its default configuration is transparent to untagged frames (3.14) but is not transparent to tagged frames (3.12), so the operation of such Bridges in the presence of tagged traffic differs from that of an ISO/IEC 15802-3 Bridge. If the configuration settings of VLAN Bridges are changed from the default values defined in this standard, then transparency with respect to untagged frames may also be affected.

## 1.3 Relationship with ISO/IEC 15802-3

This standard makes use of specific aspects of the MAC Bridge specification contained in ISO/IEC 15802-3; those aspects therefore become provisions of this standard. Table 1-1 shows how the relevant clauses of ISO/IEC 15802-3 are incorporated into this standard.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

**Table 1-1—Relationship between this standard and ISO/IEC 15802-3**

| ISO/IEC 15802-3 clause | Use in this standard |
|---|---|
| 5. Conformance | Provision of this standard, as extended by Clause 5 |
| 6. Support of the MAC Service | Provision of this standard, as extended by Clause 7 |
| 7. Principles of operation | Replaced by Clause 8 |
| 8. The spanning tree algorithm and protocol | Provision of this standard |
| 9. Encoding of bridge protocol data units | Provision of this standard |
| 10. GARP Multicast Registration Protocol (GMRP) | Provision of this standard, as modified by Clause 11 |
| 11. Example “C” code implementation of GMRP | Provision of this standard, as modified by Clause 11 |
| 12. Generic Attribute Registration Protocol (GARP) | Provision of this standard |
| 13. Example “C” code implementation of GARP | Provision of this standard |
| 14. Bridge management | Replaced by Clause 12 |
| 15. Management protocol | Not applicable |
| 16. Bridge performance | Provision of this standard |
| Annex A (normative) PICS proforma | Replaced by Annex A |
| Annex B (informative) Calculating Spanning Tree parameters | Provision of this standard |
| Annex C (normative) Source-routing transparent bridge operation | Provision of this standard |
| Annex D (normative) PICS proforma for source routing transparent bridge operation | Provision of this standard |
| Annex E (normative) Allocation of Object Identifier values | Not applicable |
| Annex F (informative) Target topology, migration, and interoperability | Provision of this standard |
| Annex G (informative) Preserving the integrity of FCS fields in MAC Bridges | Provision of this standard |
| Annex H (informative) Design considerations for Traffic Class Expediting and Dynamic Multicast Filtering | Provision of this standard |

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## 2. References

The following standards contain provisions which, through reference in this text, constitute provisions of this standard. At the time of publication, the editions indicated were valid. All standards are subject to revision, and parties to agreements based on this standard are encouraged to investigate the possibility of applying the most recent editions of the standards indicated below. Members of IEC and ISO maintain registers of currently valid International Standards.

ANSI X3.159-1989, American National Standards for Information Systems—Programming Language—C.[2]

IEEE Std 802-1990, IEEE Standards for Local and Metropolitan Area Networks: Overview and Architecture.[3]

IEEE Std 802.1F-1993, IEEE Standards for Local and Metropolitan Area Networks: Common Definitions and Procedures for IEEE 802 Management Information.

IEEE Std 802.3, 1998 Edition, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.

IEEE Std 802.3ac-1998, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Supplement to Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications: Frame Extensions for Virtual Bridged Local Area Network (VLAN) tagging on 802.3 Networks.

IEEE Std 802.9a-1995, IEEE Standards for Local and Metropolitan Area Networks: Supplement to Integrated Services (IS) LAN Interface at the Medium Access Control (MAC) and Physical (PHY) Layers: Specification of ISLAN 16-T.

IETF RFC 1042, Postel & Reynolds, A Standard for the Transmission of IP Datagrams over IEEE 802 Networks, February 1988.[4]

IETF RFC 1390, D. Katz, Transmission of IP and ARP over FDDI Networks, January 1993.

ISO 6937-2: 1994, Information technology—Coded graphic character set for text communication—Latin alphabet.[5]

ISO/IEC 7498-1: 1994, Information processing systems—Open Systems Interconnection—Basic Reference Model—Part 1: The Basic Model.

ISO/IEC 7498-4: 1989, Information processing systems—Open Systems Interconnection—Basic Reference Model—Part 4: Management framework.

---

[2]ANSI publications are available from the Sales Department, American National Standards Institute, 11 West 42nd Street, 13th Floor, New York, NY 10036, USA.

[3]IEEE publications are available from the Institute of Electrical and Electronics Engineers, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331, USA.

[4]Internet RFCs are retrievable by FTP at ds.internic.net/rfc/rfcnnnn.txt (where nnnn is a standard's publication number such as 1042), or call InterNIC at 1-800-444-4345 for information about receiving copies through the mail.

[5]ISO and ISO/IEC documents are available from the ISO Central Secretariat, 1 rue de Varembé, Case Postale 56, CH-1211, Genève 20, Switzerland/Suisse; and from the Sales Department, American National Standards Institute, 11 West 42nd Street, 13th Floor, New York, NY 10036, USA.


Copyright © 1999 IEEE. All rights reserved.

ISO/IEC 8802-2: 1998 [ANSI/IEEE Std 802.2, 1998 Edition], Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 2: Logical link control.[6]

ISO/IEC 8802-4: 1990 [ANSI/IEEE Std 802.4-1990], Information processing systems—Local area networks—Part 4: Token-passing bus access method and physical layer specifications.

ISO/IEC 8802-5: 1998 [ANSI/IEEE Std 802.5, 1998 Edition], Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 5: Token ring access method and physical layer specifications.

ISO/IEC 8802-6: 1994 [ANSI/IEEE Std 802.6, 1994 Edition], Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 6: Distributed Queue Dual Bus (DQDB) access method and physical layer specifications.

ISO/IEC 8802-9: 1996 [ANSI/IEEE Std 802.9, 1996 Edition], Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 9: Integrated Services (IS) LAN Interface at the Medium Access Control (MAC) and Physical (PHY) Layers.

ISO/IEC DIS 8802-11, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications.

ISO/IEC 8802-12: 1998 [ANSI/IEEE Std 802.12, 1998 Edition], Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Specific requirements—Part 12: Demand-Priority access method, physical layer and repeater specifications.

ISO/IEC 8824: 1990, Information technology—Open Systems Interconnection—Specification of Abstract Syntax Notation One (ASN.1) (Provisionally retained edition).

ISO/IEC 8825: 1990, Information technology—Open Systems Interconnection—Specification of Basic Encoding Rules for Abstract Syntax Notation One (ASN.1) (Provisionally retained edition).

ISO 9314-2: 1989, Information processing systems—Fibre Distributed Data Interface—Part 2: FDDI Token Ring Media Access Control (MAC).

ISO/IEC 9595: 1998, Information technology—Open Systems Interconnection—Common management information service.

ISO/IEC 9596: 1998, Information technology—Open Systems Interconnection—Common management information protocol—Part 1: Specification.

ISO/IEC 10038: 1993 [ANSI/IEEE Std 802.1D-1993], Information technology—Telecommunications and information exchange between systems—Local area networks—Media Access Control (MAC) bridges.

ISO/IEC 11802-5: 1997 [ANSI/IEEE Std 802.1H-1997], Information technology—Telecommunications and information exchange between systems—Local and Metropolitan Area Networks—Technical Reports and Guidelines—Part 5: Media Access Control Bridging of Ethernet V2.0 in IEEE 802 Local Area Networks.

[6]ISO [IEEE] and ISO/IEC [IEEE] documents are available from ISO Central Secretariat, 1 rue de Varembé, Case Postale 56, CH-1211, Genève 20, Switzerland/Suisse; and from the Institute of Electrical and Electronics Engineers, 445 Hoes Lane, P.O. Box 1331, Piscataway, NJ 08855-1331.

Copyright © 1999 IEEE. All rights reserved.



Dell et al., Ex. 1008
Page 17 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

ISO/IEC 15802-1: 1995, Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 1: Medium Access Control (MAC) service definition.

ISO/IEC 15802-2: 1995 [ANSI/IEEE Std 802.1B, 1995], Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 2: LAN/MAN Management.

ISO/IEC 15802-3: 1998 [ANSI/IEEE Std 802.1D, 1998 Edition], Information technology—Telecommunications and information exchange between systems—Local and metropolitan area networks—Common specifications—Part 3: Media Access Control (MAC) Bridges.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# 3. Definitions

The following terms are specific to this standard:

## 3.1 Ethernet Type-encoding

The use of the Type interpretation of an IEEE 802.3 Length/Type field value in a frame as a protocol identifier associated with the MAC Service user data carried in the frame.[7]

NOTES

1—The term *frame* is defined in 3.3.

2—Ethernet Type-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 ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390.

## 3.2 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.

## 3.3 Frame

A unit of data transmission on an IEEE 802 LAN MAC that conveys a protocol data unit (PDU) between MAC Service users. There are three types of frame; *untagged*, *VLAN-tagged*, and *priority-tagged*.

NOTE—The term *IEEE 802 LAN* is defined in ISO/IEC 15802-3. *Untagged frame* is defined in 3.14, *VLAN-tagged frame* is defined in 3.18, and *priority-tagged frame* is defined in 3.8.

## 3.4 Frame relay

The function of the Forwarding Process that forwards frames between the Ports of a Bridge.

NOTE—The operation of the Forwarding Process is defined in 8.7.

## 3.5 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.

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

## 3.6 Independent Virtual Local Area Network (VLAN) Learning (IVL) Bridge

A type of Bridge that supports only Independent VLAN Learning.

[7]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 © 1999 IEEE. All rights reserved.



Dell et al., Ex. 1008
Page 19 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## 3.7 Legacy region

A set of LAN segments interconnected such that there is physical connectivity between any pair of segments using only ISO/IEC 15802-3-conformant, VLAN-unaware MAC Bridges.

NOTE—In other words, if, in a Bridged LAN containing both ISO/IEC 15802-3 and IEEE 802.1Q Bridges, all the IEEE 802.1Q Bridges were to be removed, the result would be a set of one or more Bridged LANs, each with its own distinct Spanning Tree. Each of those Bridged LANs is a legacy region. The term VLAN-unaware is defined in 3.19.

## 3.8 Priority-tagged frame

A tagged frame whose tag header carries priority information, but carries no VLAN identification information.

NOTE—The term *tagged frame* is defined in 3.12; *tag header* is defined in 3.13. See also *VLAN-tagged frame* (3.18) and *untagged frame* (3.14).

## 3.9 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.

## 3.10 Shared Virtual Local Area Network (VLAN) Learning (SVL) Bridge

A type of Bridge that supports only Shared VLAN Learning.

## 3.11 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.

## 3.12 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. There are two types of tagged frames: VLAN-tagged frames and priority-tagged frames.

NOTE—The term *tag header* is defined in 3.13, *priority-tagged frame* is defined in 3.8, and *VLAN-tagged frame* is defined in 3.18. See also *untagged frame* (3.14).

## 3.13 Tag header

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

NOTE—The structure of the tag header is defined in 9.3.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### 3.14 Untagged frame

An *untagged frame* is a frame that does not contain 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.

NOTE—The term *tag header* is defined in 3.13. See also *tagged frame* (3.12).

### 3.15 Virtual Bridged Local Area Network (LAN)

A Bridged LAN in which the existence of one or more VLAN-aware Bridges allows the definition, creation, and maintenance of VLANs.

NOTE—The term *VLAN-aware* is defined in 3.17.

### 3.16 Virtual Local Area Network (VLAN)

A subset of the active topology of a Bridged Local Area Network. Associated with each VLAN is a VLAN Identifier (VID).

### 3.17 VLAN-aware

A property of Bridges or end stations that recognize and support VLAN-tagged frames.

### 3.18 VLAN-tagged frame

A tagged frame whose tag header carries both VLAN identification and priority information.

### 3.19 VLAN-unaware

A property of Bridges or end stations that do not recognize VLAN-tagged frames.

### 3.20 Terms used in ISO/IEC 15802-3

The following terms used in this standard are used in ISO/IEC 15802-3:

Active topology
Bridge Port
Bridged Local Area Network (also Bridged LAN)
GARP Participant
GARP Application
GIP Context
Group
IEEE 802 Local Area Network (also IEEE 802 LAN, or LAN)
Port

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 21 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## 4. Abbreviations

The following abbreviations are used in this standard:

| | |
|---|---|
| BPDU | Bridge Protocol Data Unit (ISO/IEC 15802-3) |
| CFI | Canonical Format Indicator (Annex F) |
| E-ISS | Enhanced Internal Sublayer Service (7.1) |
| FCS | Frame Check Sequence (7.1) |
| FID | Filtering Identifier (8.11.3, 8.11.7) |
| GARP | Generic Attribute Registration Protocol (ISO/IEC 15802-3) |
| GID | GARP Information Declaration (ISO/IEC 15802-3) |
| GIP | GARP Information Propagation (ISO/IEC 15802-3) |
| GMRP | GARP Multicast Registration Protocol (Clause 10, ISO/IEC 15802-3) |
| GVRP | GARP VLAN Registration Protocol (Clause 11) |
| ISS | Internal Sublayer Service (Clause 7, ISO/IEC 15802-3) |
| IVL | Independent VLAN Learning (3.5) |
| LAN | Local Area Network (IEEE Std. 802) |
| LS | Least-significant |
| LLC | Logical Link Control (ISO/IEC 8802-2) |
| MAC | Medium Access Control (IEEE Std. 802) |
| MIB | Management Information Base (ISO/IEC 7498-4) |
| MS | Most-significant |
| MSDU | MAC Service Data Unit (ISO/IEC 15802-1) |
| NCFI | Non-Canonical Format Indicator (Annex F) |
| PDU | Protocol Data Unit |
| PICS | Protocol Implementation Conformance Statement (Annex A) |
| PVID | Port VID (8.4.4) |
| RIF | Routing Information Field (ISO/IEC 8802-5) |
| STPID | SNAP-encoded Tag Protocol Identifier (9.3) |
| SVL | Shared VLAN Learning (3.9) |
| TCI | Tag Control Information (9.3) |
| TPID | Tag Protocol Identifier (9.3) |
| VID | VLAN Identifier (7.1, 8.4.4, 9.3) |
| VLAN | Virtual LAN (3.16) |


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# 5. Conformance

## 5.1 Static conformance requirements

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

a) Conform to the requirements of ISO/IEC 15802-3, as modified by the provisions of this standard;
b) Relay and filter frames as described in 8.1 and specified in 8.5, 8.6, 8.7, 8.8, and 8.9;
c) On each Port, support at least one of the permissible values for the Acceptable Frame Types parameter, as defined in 8.4.3;
d) Support the following on each Port that supports untagged and priority-tagged frames:
   1) A Port VLAN Identifier (PVID) value (8.4.4);
   2) The ability to configure at least one VLAN whose untagged set includes that Port (8.8 and 8.11.9);
   3) Configuration of the PVID value via management operations (12.10);
   4) Configuration of Static Filtering Entries via management operations (12.7).
e) Support the ability to insert tag headers into, modify tag headers in, and remove tag headers from relayed frames, as described and specified in 7.1 and Clause 9, as required by the value(s) of the Acceptable Frame Types parameter that are supported on each Port, and by the ability of each Port to be configured to transmit VLAN-tagged frames and/or untagged frames. These requirements are summarized in Table 5-1 for frames relayed between any pair of Ports;

**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. |

f) Support the ability to perform automatic configuration and management of VLAN topology information by means of Generic Attribute Registration Protocol (GARP) VLAN Registration Protocol (GVRP) (Clause 11) on all Ports;
g) Support the ability for the Filtering Database to contain static and dynamic configuration information for at least one VLAN, by means of Static and Dynamic VLAN Registration Entries (8.11);
h) Support at least one Filtering Identifier (FID) (6.4, 8.11.3, 8.11.7, and 8.11.8);
i) Support the ability to allocate at least one VID to each FID that is supported (6.4, 8.11.3, 8.11.7, and 8.11.8).

NOTE—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 © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 23 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## 5.2 Options

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

a) Support operation in Extended Filtering Mode (ISO/IEC 15802-3, 6.6.5) and the operation of GARP Multicast Registration Protocol (GMRP) (ISO/IEC 15802-3, Clause 10) as modified by Clause 10;
b) Support the ability for the Filtering Database to contain static and dynamic configuration information for more than one VLAN, by means of Static and Dynamic VLAN Registration Entries (8.11), up to a maximum of 4094 VLANs;

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

c) On each Port, support both of the permissible values for the Acceptable Frame Types parameter, as defined in 8.4.3. If both values are supported, then the implementation shall support configuration of the parameter value via management;
d) Support the ability to enable and disable Ingress Filtering (8.4.5);
e) Support the ability to configure more than one VLAN whose untagged set includes that Port (8.8 and 8.11.9);
f) Support the management functionality defined in Clause 12;
g) Support more than one FID (6.4, 8.11.3, 8.11.7, and 8.11.8);
h) Support the ability to allocate more than one VID to each FID that is supported (6.4, 8.11.3, 8.11.7, and 8.11.8);
i) Support the ability to configure VLAN Learning Constraints via management (8.11.7 and 12.10.3);
j) Support the ability to configure fixed VID to FID allocations via management (8.11.7.1 and 12.10.3);
k) Support any other optional capabilities defined in ISO/IEC 15802-3, as modified by the provisions of this standard.

## 5.3 Protocol Implementation Conformance Statement (PICS)

The supplier of an implementation that is claimed to conform to this standard shall complete a copy of the PICS proforma provided in Annex A and shall provide the information necessary to identify both the supplier and the implementation.

## 5.4 MAC-specific bridging methods

MAC-specific bridging methods may exist. Use of a MAC-specific bridging method and the method specified in this standard on the same LAN shall

a) Not prevent communication between stations in a Bridged LAN.
b) Preserve the MAC Service.
c) Preserve the characteristics of each bridging method within its own domain.
d) Provide for the ability of both bridging techniques to coexist simultaneously on a LAN without adverse interaction.

ISO/IEC 15802-3, Annex C, defines one such MAC-specific bridging method, source routing, and that method is also a provision of this standard. While this standard defines how source-routed frames can be transported in a VLAN environment, it does not attempt to specify VLAN aspects of the source routing bridging method itself.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# 6. Architectural overview

The architectural framework for VLANs is based on a three-layer model, consisting of the following layers:

a) Configuration;
b) Distribution of configuration information;
c) Relay.

The model is illustrated in Figure 6-1, and further described in the following subclauses.

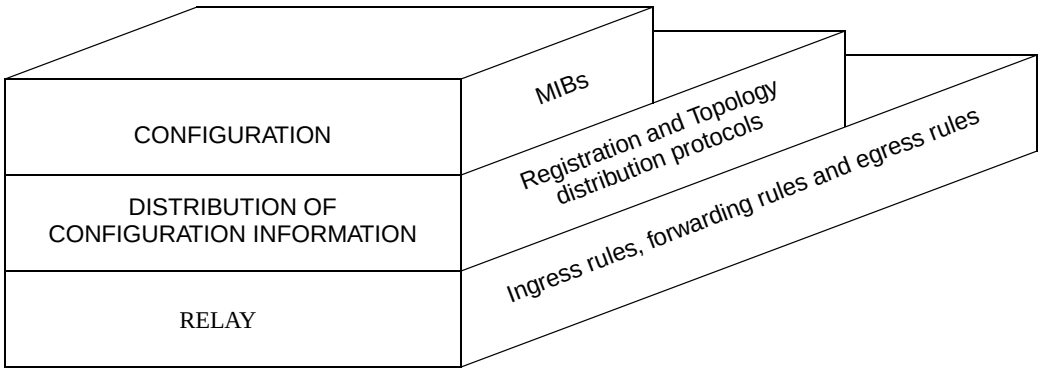


**Figure 6-1—VLAN architectural framework**

## 6.1 Configuration

Configuration is concerned with the following issues:

a) The means whereby the VLAN configuration is specified in the first place. This might be achieved via local and/or remote management mechanisms, via server mechanisms, via distribution protocols, or via other means. Determination of the configuration is outside the scope of this standard.
b) Assignment of VLAN configuration parameters.

Clause 12 defines the management operations that are standardized for use in the configuration of VLAN devices.

## 6.2 Distribution of configuration information

This is the process that allows information to be distributed in order for Bridges to be able to determine to which VLAN a given frame should be classified. Clause 11 defines a general mechanism for the distribution of VLAN membership information to all VLAN-aware devices in a Bridged LAN.

## 6.3 Relay

This is concerned with the mechanics of

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

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 25 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

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

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

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

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

The ingress, forwarding, and egress rules allow Bridges to

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

NOTE 1—This classification of untagged and priority-tagged frames is part of the functionality of the MAC relay entity (Figure 8-3, Figure 8-4), and is therefore only of significance for received frames that are potentially to be forwarded through other Ports of the Bridge (see 6.7).

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

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

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

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

## 6.4 Filtering Database architecture

The Filtering Database architecture defined in this standard recognizes that

a) For some configurations, it is necessary to allow address information learned in one VLAN to be shared among a number ofVLANs. This is known as *Shared VLAN Learning* (3.9);

b) For some configurations, it is desirable to ensure that address information learned in one VLAN is not shared with other VLANs. This is known as *Independent VLAN Learning* (3.5);

c) For some configurations, it is immaterial as to whether learned information is shared between VLANs or not.

NOTE 1—Annex B discusses the need for Shared and Independent VLAN Learning, and also some of the related interoperability issues.


Copyright © 1999 IEEE. All rights reserved.

Shared VLAN Learning is achieved by including learned information from a number of VLANs in the same Filtering Database; Independent VLAN Learning is achieved by including information from each VLAN in distinct Filtering Databases.

NOTE 2—The actual Filtering Database specification specifies a single Filtering Database that, through the inclusion of VLAN identification information in each database entry, can model the existence of one or more distinct Filtering Databases.

Within a given VLAN-Bridged LAN, there may be a combination of configuration requirements, so that individual VLAN Bridges may be called upon to share learned information, or not share it, according to the requirements of particular VLANs or groups of VLANs. The Filtering Database structure that is defined in this standard allows both Shared and Independent VLAN Learning to be implemented within the same VLAN Bridge; i.e., allows learned information to be shared between those VLANs for which Shared VLAN Learning is necessary, while also allowing learned information not to be shared between those VLANs for which Independent VLAN Learning is necessary. The precise requirements for each VLAN with respect to sharing or independence of learned information (if any) are made known to VLAN Bridges by means of a set of *VLAN Learning Constraints* (8.11.7.2), which may be configured into the Bridges by means of management operations. By analyzing the set of learning constraints for the VLANs that are currently active, the Bridge can determine

d) How many independent Filtering Databases are required in order to meet the constraints;
e) For each VLAN, which Filtering Database it will feed any learned information into (and use learned information from).

The manner in which this mapping of VLANs onto Filtering Databases is achieved is defined in 8.11.7; the result is that each VLAN is associated with exactly one Filtering Database.

The most general application of the Filtering Database specification in this standard is a Bridge that can support M independent Filtering Databases, and can map N VLANs onto each Filtering Database. Such a Bridge is known as an SVL/IVL Bridge (3.11).

The conformance requirements in this standard (5.1, 5.2) recognize that VLAN Bridges will be implemented with differing capabilities in order to meet a wide range of application needs, and that the full generality of the SVL/IVL approach is not always either necessary or desirable, as observed in the discussion in Annex B. In a given conformant implementation, there may be restrictions placed upon the number of Filtering Databases that can be supported, and/or the number of VLANs that can be mapped onto each Filtering Database. The full spectrum of conformant Filtering Database implementations is therefore as follows:

f) The SVL/IVL Bridge, as described above. Such Bridges provide support for M Filtering Databases, with the ability to map N VLANs onto each one;
g) Support for a single Filtering Database only. MAC Address information that is learned in one VLAN can be used in filtering decisions taken relative to all other VLANs supported by the Bridge. Bridges that support a single Filtering Database are referred to as SVL Bridges;
h) Support for multiple Filtering Databases, but only a single VLAN can be mapped onto each Filtering Database. MAC Address information that is learned in one VLAN cannot be used in filtering decisions taken relative to any other VLAN. Bridges that support this mode of operation are referred to as IVL Bridges.

## 6.5 VLAN classification

VLAN technology introduces the following three basic types of frame:

a) Untagged frames;
b) Priority-tagged frames; and
c) VLAN-tagged frames.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 27 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

An *untagged frame* or a *priority-tagged frame* does not carry any identification of the VLAN to which it belongs. Such frames are classified as belonging to a particular VLAN based on parameters associated with the receiving Port, or, through proprietary extensions to this standard, based on the data content of the frame (e.g., MAC Address, layer 3 protocol ID, etc.).

NOTE—For the purposes of VLAN identification, priority tagged frames, which, by definition, carry no VLAN identification information, are treated the same as untagged frames.

A *VLAN-tagged frame* carries an explicit identification of the VLAN to which it belongs; i.e., it carries a tag header that carries a non-null VID. Such a frame is classified as belonging to a particular VLAN based on the value of the VID that is included in the tag header. The presence of the tag header carrying a non-null VID means that some other device, either the originator of the frame or a VLAN-aware Bridge, has mapped this frame into a VLAN and has inserted the appropriate VID. Clause 7 describes how the insertion and removal of VLAN tags is achieved.

## 6.6 Rules for tagging frames

For a given VLAN, all frames transmitted on a given LAN segment by a VLAN-aware Bridge shall be tagged the same way on that segment. They shall be either

a) All untagged; or
b) All VLAN tagged with the same VID.

NOTE—In other words, a Bridge can transmit untagged frames for someVLANs and VLAN-tagged frames for other VLANs on a given link, but cannot transmit both formats for the same VLAN. The single format rule expressed here only applies to the frame transmission behavior of individual VLAN-aware devices; i.e., it does not express a requirement for VLAN-aware devices to police the behavior of other devices in order to enforce a single format on a segment for all attached devices.

## 6.7 Spanning Tree

This standard defines a VLAN environment that operates over a single Spanning Tree. All Bridges within a Bridged LAN infrastructure participate in a single Spanning Tree, as defined by ISO/IEC 15802-3, over which multiple VLANs can coexist. As a consequence, Bridges implemented in conformance with ISO/IEC 15802-3 can be integrated into a VLAN infrastructure based on the specification contained in this standard.

NOTE 1—There are some limitations on the intermixing of ISO/IEC 15802-3 Bridges and VLAN Bridges in the same Bridged LAN, as identified in Annex E.

The primary goals of Spanning Tree are as follows:

a) Elimination of loops in a bridged infrastructure;
b) Improved scalability in a large network;
c) Provision of redundant paths, which can be activated upon failure.

There are two important items to note with respect to Spanning Tree topologies.

First, the Spanning Tree formed in a VLAN environment need not be identical to the topology of the VLAN(s). All VLANs are aligned along the Spanning Tree from which they are formed; a given VLAN is defined by a subset of the topology of the Spanning Tree upon which it operates.

Second, the topology of the VLAN is dynamic. The structure of the VLAN may change due to new devices requesting or releasing the services available via the VLAN. The dynamic nature of VLANs has the advantages of flexibility and bandwidth conservation, at the cost of network management complexity.


Copyright © 1999 IEEE. All rights reserved.

NOTE 2—There is a choice to be made as to how many Spanning Trees operate in a VLAN environment, and how the VLANs in that environment map to those Spanning Trees. In all cases, a given VLAN maps to a single Spanning Tree; the mapping choice to be made with multiple Spanning Trees is whether there is one Spanning Tree per VLAN, or whether many VLANs map to each Spanning Tree. Although multiple Spanning Trees offer some advantages over a single Spanning Tree in VLAN environments, this standard avoids the added complexity of defining a mapping function of VLANs to Spanning Trees by defining a VLAN environment that operates over a single Spanning Tree. It is the intent of this standard not to preclude future extensions from using multiple Spanning Trees.

In order for coexistence to be correctly maintained with Bridges implemented in conformance with ISO/IEC 15802-3, the addition or removal of tag headers by a VLAN-aware Bridge is performed only upon frames submitted to the relay function of the Bridge that are potentially to be forwarded on other Ports. Frames that carry control information that is necessary for the establishment of Spanning Tree and other aspects of the connectivity and forwarding behavior of the Bridged LAN, such as BPDUs (ISO/IEC 15802-3, Clauses 8 and 9) and GVRP PDUs (11.2), are not forwarded by a VLAN-aware Bridge. Such frames are discarded by the Forwarding Process as a result of permanently configured static entries in the Filtering Database (see 8.2, 8.3, and 8.14).

NOTE 3—GARP PDUs destined for any GARP Applications are forwarded or filtered depending upon whether the application concerned is supported by the Bridge, as specified in 8.14.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 29 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# 7. Support of the MAC Service in VLANs

The provisions of ISO/IEC 15802-3, Clause 6, apply to this standard, with the additions and modifications defined in this clause.

## 7.1 Enhanced Internal Sublayer Service provided within VLAN Bridges

The Enhanced Internal Sublayer Service (E-ISS) is derived from the Internal Sublayer Service (ISS, defined in ISO/IEC 15802-3, 6.4) by augmenting that specification with elements necessary to the operation of the tagging and untagging functions of the MAC Bridge. 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.

Bridges that support these functions are known as *VLAN-aware Bridges* (3.17). The E-ISS defines the MAC Service provided to the relay function in VLAN-aware Bridges.

The relationships between the MAC Entity, the ISS, the E-ISS, and the MAC Relay Entity in a VLAN-aware Bridge are illustrated in Figure 7-1 and Figure 8-3.

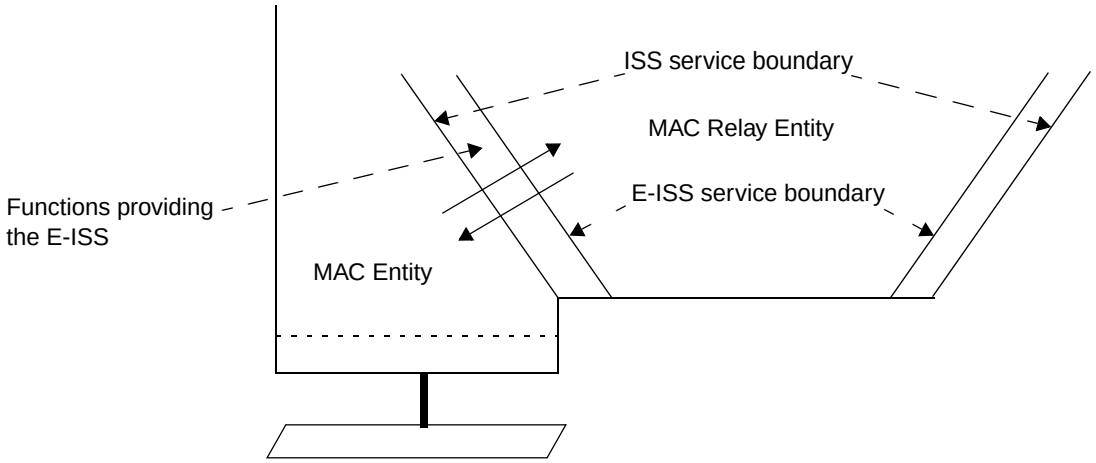


**Figure 7-1—Relationships between MAC Entity, ISS, E-ISS, and MAC Relay Entity**

### 7.1.1 E-ISS service definition

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,
frame_check_sequence,
canonical_format_indicator,
vlan_identifier,
rif_information (optional)
)


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Each data indication primitive corresponds to the receipt of a M_UNITDATA.indication primitive from the Internal Sublayer Service, as defined in ISO/IEC 15802-3, 6.4.

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 M_UNITDATA.indication primitive of the Internal Sublayer service.

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 **vlan_identifier** parameter carries the VLAN identifier associated with the indication.

The **rif_information** parameter is present if a tag header was present in the indication, and if that tag header contained a Routing Information Field (RIF). Its value is equal to the value of the RIF.

EM_UNITDATA.request (
frame_type,
mac_ action,
destination_address,
source_address,
mac_service_data_unit,
user_priority,
access_priority,
frame_check_sequence,
canonical_format_indicator,
vlan_classification,
rif_information (optional),
include_tag
)

A data request primitive is invoked in order to generate a M_UNITDATA.request primitive, as defined in the Internal Sublayer Service, ISO/IEC 15802-3, 6.4.

The **frame_type, mac_action, destination_address, source_address, mac_service_data_unit, user_priority, access_priority,** and **frame_check_sequence** parameters are as defined for the M_UNITDATA.request primitive of the Internal Sublayer service.

The definition of the **canonical_format_indicator** parameter is as defined for the EM_UNITDATA.indication.

The **vlan_classification** parameter carries the VLAN classification assigned to the frame by the ingress rules (8.6).

The **rif_information parameter**, if present, carries the value of any RIF information to be associated with the request.

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.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 31 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### 7.1.2 Support of the E-ISS in VLAN-aware Bridges

#### 7.1.2.1 Data indication primitives

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

NOTE 1—The **mac_action** parameter only ever takes the value request_with_no_response for frames relayed by the Bridge. The **frame_check_sequence** parameter of the data indication carries the FCS value contained in the received frame. The original FCS associated with a frame is invalidated if there are changes to any fields of the frame, if fields are added or removed, or if bit ordering or other aspects of the frame encoding have changed. An invalid FCS is signalled in the E-ISS by an unspecified value in the frame_check_sequence parameter of the data request primitive. This signals the need for the FCS to be regenerated according to the normal procedures for the transmitting MAC. The options for regenerating the FCS under these circumstances are discussed in ISO/IEC 15802-3, Annex G.

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

c) If the received mac_service_data_unit parameter contained a tag header (9.3), then the value contained in the user_priority field of the tag header is used. Otherwise;
d) The value of the received user_ priority parameter, regenerated as defined in 8.5.1 and ISO/IEC 15802-3, 6.4, is used.

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

e) If the received mac_service_data_unit parameter contained a tag header (9.3), then the value(s) contained in the Canonical Format Indicator (CFI) (and Non-Canonical Format Indicator [NCFI], if present) field(s) of the tag header are used to determine this parameter value, in accordance with the definition of the CFI and NCFI field(s) in Clause 9. Otherwise;
f) 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;
g) The parameter carries the value True.

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

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

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

j) If the initial octets of the received mac_service_data_unit parameter contained a tag header (9.3), and that tag header contained a RIF field in which one or more route descriptors were present, then the value contained in the RIF field is used. Otherwise;
k) The parameter is not present.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

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

#### 7.1.2.2 Data request primitives

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 False, the value of the mac_service_data_unit parameter is determined as follows:

a) If the destination MAC method is the same as the MAC method on which the corresponding data indication was received, then the value used is equal to the value of the mac_service_data_unit parameter received in the data request. Otherwise;
b) The value used is equal to the value of the mac_service_data_unit parameter received in the data request, modified, if necessary, in accordance with the procedures described in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390.
c) 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 method, then the Bridge shall either
    1) Convert any embedded MAC Addresses in the mac_service_data_unit to the format appropriate to the destination MAC method; or
    2) Discard the EISS data request without issuing a corresponding ISS data request.

If the value of the **include_tag parameter** is True, then a tag header, formatted as necessary for the destination MAC method, is inserted as the first N octets of the mac_service_data_unit parameter. The values of the user_priority, canonical_format_indicator, vlan_classification, and rif_information (if present) parameters are used to determine the contents of the tag header, in accordance with the structure defined in 9.2 and 9.3. The value inserted after the tag header is determined as follows:

d) If the destination MAC method is the same as the MAC method on which the corresponding data indication was received, then the value used is equal to the value of the mac_service_data_unit parameter received in the data request. Otherwise;
e) The value used is equal to the value of the mac_service_data_unit parameter received in the data request, modified, if necessary, in accordance with the procedures described in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390.

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

f) 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;
g) 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.

NOTE—The original FCS associated with a frame is invalidated if there are changes to any fields of the frame, if fields are added or removed, or if bit ordering or other aspects of the frame encoding have changed. An invalid FCS is signalled in the E-ISS by an unspecified value in the frame_check_sequence parameter of the data request primitive. This signals the need for the FCS to be regenerated according to the normal procedures for the transmitting MAC. The options for regenerating the FCS under these circumstances are discussed in ISO/IEC 15802-3, Annex G.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 33 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## 7.2 Support of the Internal Sublayer Service by IEEE Std 802.3 (CSMA/CD)

In addition to the provisions of ISO/IEC 15802-3, 6.5.1, 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:

a) Use the procedure as described in ISO/IEC 15802-3, 6.5.1. This can result in tagged frames of less than 68 octets (but at least 64 octets) being transmitted; or
b) Include additional octets before the FCS field in order for the transmitted frame length for such frames to be 68 octets. This results in a minimum tagged frame length of 68 octets.

When a tagged frame of less than 68 octets in length is received on a CSMA/CD LAN segment, and is forwarded as an untagged frame, the provisions of ISO/IEC 15802-3, 6.5.1, result in additional octets being included before the FCS field on transmission in order that the transmitted frame length meets the minimum frame size requirements of IEEE Std 802.3, 1998 Edition, 3.2.7.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# 8. Principles of operation

This clause establishes the principles of operation of a VLAN-aware Bridge, by reference to a model of that operation, as follows:

a) Explains the principal elements of Bridge operation and lists the functions that support these.
b) Establishes an architectural model for a Bridge that governs the provision of these functions.
c) Provides a model of the operation of a Bridge in terms of the processes and entities that support the functions.
d) Details the addressing requirements in a Bridged LAN and specifies the addressing of entities in a Bridge.

The provisions of this clause replace the provisions of ISO/IEC 15802-3, Clause 7, in a VLAN-aware Bridge.

## 8.1 Bridge operation

The principal elements of Bridge operation are

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

### 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 order of frames shall be preserved as defined in 8.7.3.

The functions that support the relaying of frames and maintain the Quality of Service supported by the Bridge are

a) Frame reception.
b) Discard on received frame in error (ISO/IEC 15802-3, 6.3.2).
c) Frame discard if the frame_type is not user_data_frame, or if its mac_action parameter is not request_with_no_response (8.5, ISO/IEC 15802-3, 6.4).
d) Regeneration of user priority, if required (ISO/IEC 15802-3, 6.4).
e) Frame discard following the application of filtering information.
f) Frame discard on transmittable service data unit size exceeded (ISO/IEC 15802-3, 6.3.8).
g) Forwarding of received frames to other Bridge Ports.
h) Selection of traffic class, following the application of filtering information.
i) Queuing of frames by traffic class.
j) Frame discard to ensure that a maximum bridge transit delay is not exceeded (ISO/IEC 15802-3, 6.3.6).
k) Selection of queued frames for transmission.
l) Selection of outbound access priority (ISO/IEC 15802-3, 6.3.9).
m) Mapping of service data units and recalculation of Frame Check Sequence, if required (8.7.6, ISO/IEC 15802-3, 6.3.7).
n) Frame transmission.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 35 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### 8.1.2 Filtering and relaying information

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

a) Calculation and configuration of Bridged LAN topology.

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

b) Permanent configuration of reserved addresses.
c) Explicit configuration of static filtering information.
d) Automatic learning of dynamic filtering information for unicast destination addresses through observation of source addresses of Bridged LAN traffic.
e) Ageing out of dynamic filtering information that has been learned.
f) Automatic addition and removal of dynamic filtering information as a result of GMRP protocol exchanges.

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

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

A Bridge classifies untagged frames and priority-tagged frames as belonging to a particular VLAN in accordance with the *ingress rules* defined in 8.6. The function that supports the use and maintenance of information for this purpose is

h) Explicit configuration of the Port VID (PVID, 8.4.4) associated with each Port of the Bridge.

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

i) Explicit configuration of the Acceptable Frame Types parameter (8.4.3) associated with each Port of the Bridge.

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

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

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

k) Automatic configuration of Dynamic VLAN Registration Entries by means of GVRP (8.11.5 and 11.2);
l) Explicit configuration of management controls associated with the operation of GVRP by means of Static VLAN Registration Entries (8.11.2 and 11.2);
m) Automatic learning of MAC Addresses in associated VLANs through the observation of network traffic (8.10).


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

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

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

### 8.1.3 Bridge management

The functions that support Bridge Management control and monitor the provision of the above functions. They are specified in Clause 12.

## 8.2 Bridge architecture

### 8.2.1 Architectural model of a Bridge

Figure 8-1 gives an example of the physical topology of a Bridged LAN. The component LANs are interconnected by means of MAC Bridges; each Port of a MAC Bridge connects to a single LAN. Figure 8-2 illustrates a Bridge with two Ports, and Figure 8-3 illustrates the architecture of such a Bridge.

A Bridge is modeled as consisting of

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

### 8.2.2 MAC Relay Entity

The MAC Relay Entity handles the MAC method independent functions of relaying frames between Bridge Ports, filtering frames, and learning filtering information. It uses the Internal Sublayer Service provided by the separate MAC Entities for each Port. (The Internal Sublayer Service and its support are described in ISO/IEC 15802-3, 6.4 and 6.5.) Frames are relayed between Ports attached to different LANs.

### 8.2.3 Ports

Each Bridge Port transmits and receives frames to and from the LAN to which it is attached. An individual MAC Entity permanently associated with the Port provides the Internal Sublayer Service used for frame transmission and reception. The MAC Entity handles all the MAC method dependent functions (MAC protocol and procedures) as specified in the relevant standard for that IEEE 802 LAN MAC technology.

### 8.2.4 Higher Layer Entities

The Bridge Protocol Entity handles calculation and configuration of Bridged LAN topology.

The Bridge Protocol Entity and other higher layer protocol users, such as Bridge Management (8.1.3) and GARP application entities including GARP Participants (ISO/IEC 15802-3, Clause 12), make use of Logical Link Control procedures. These procedures are provided separately for each Port, and use the MAC Service provided by the individual MAC Entities.

Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

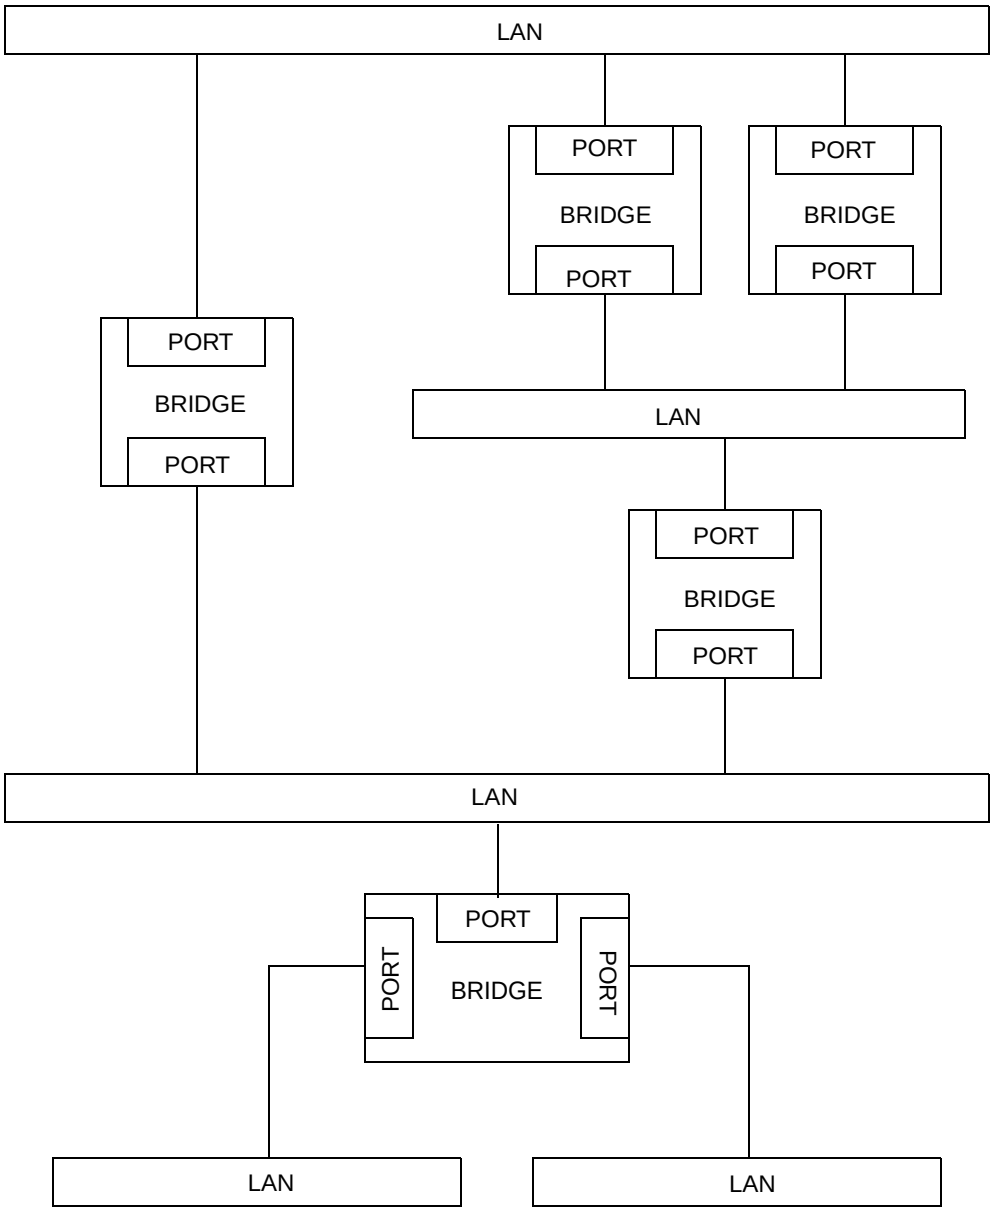


**Figure 8-1—Example of a Bridged LAN**

## 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.

Subclauses 8.5 and 8.9 specify the MAC Relay Entity's use of the Internal Sublayer Service. State information associated with each Port governs the Port's participation in the Bridged LAN. (Port States are specified in detail in ISO/IEC 15802-3, 8.4.)


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

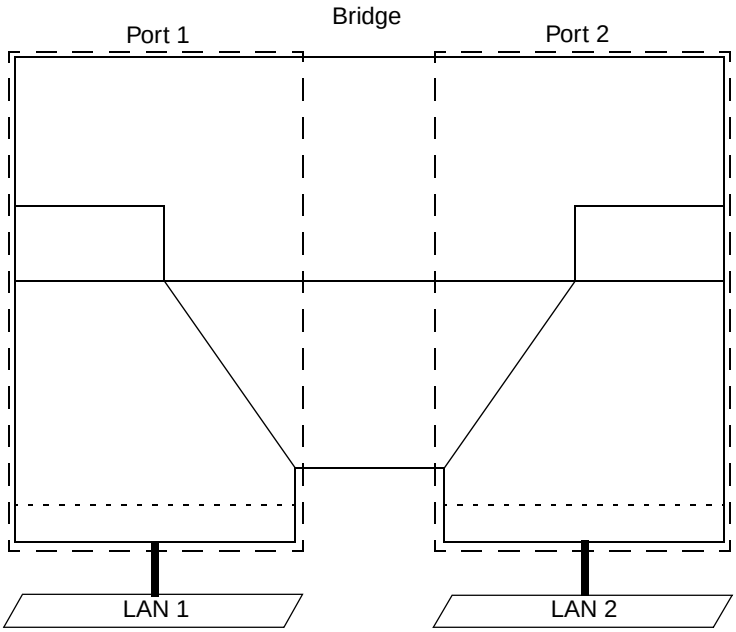


**Figure 8-2—Bridge ports**

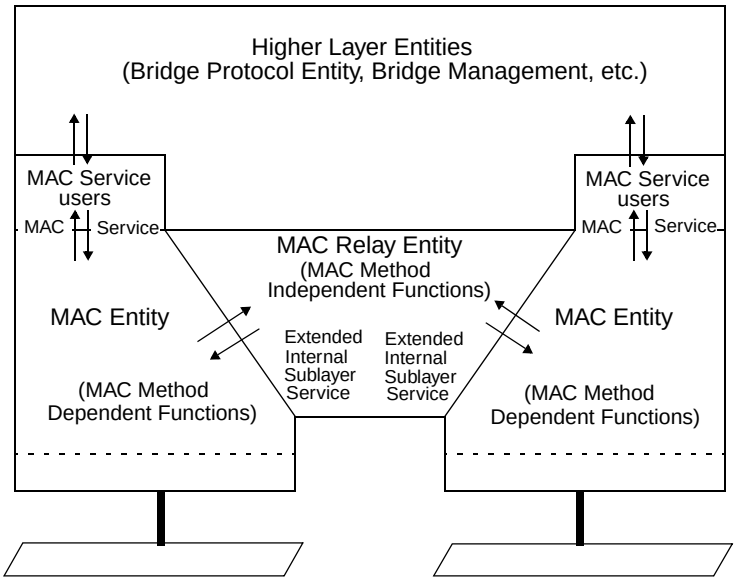


**Figure 8-3—VLAN Bridge architecture**

Copyright © 1999 IEEE. All rights reserved.



Dell et al., Ex. 1008
Page 39 of 211

Frames are accepted for transmission and delivered on reception to and from processes and entities that model the operation of the MAC Relay Entity in a Bridge. These are

a) The ingress rules (8.6), which classify received frames according to their VLAN membership, may filter frames based on the absence of a VID in the received frame (8.4.3), and may filter frames based on the frame's VLAN identifier (8.4.5);

b) The Forwarding Process (8.7), 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 (8.11) and on the state of the Bridge Ports (8.4);

c) The egress rules (8.8), which determine, for a given VLAN, through which Ports frames may be transmitted, and in what format;

d) The Learning Process (8.10), which, by observing the source addresses and VIDs of frames classified by the ingress rules, updates the Filtering Database (8.11), conditionally on the Port state (8.4);

e) The Filtering Database (8.11), which holds filtering information and supports queries by the Forwarding Process as to whether frames with given values of the destination MAC Address field and VID should be forwarded to a given Port.

Each Bridge Port also functions as an end station providing the MAC Service to LLC, which in turn supports operation of the Bridge Protocol Entity (8.12) and of other possible users of LLC, such as protocols providing Bridge Management (8.13).

Each Bridge Port shall support the operation of LLC Type 1 procedures in order to support the operation of the Bridge Protocol Entity. Bridge Ports may support other types of LLC procedures, which may be used by other protocols.

Figure 8-4 illustrates a single instance of frame relay between the Ports of a Bridge with two Ports.

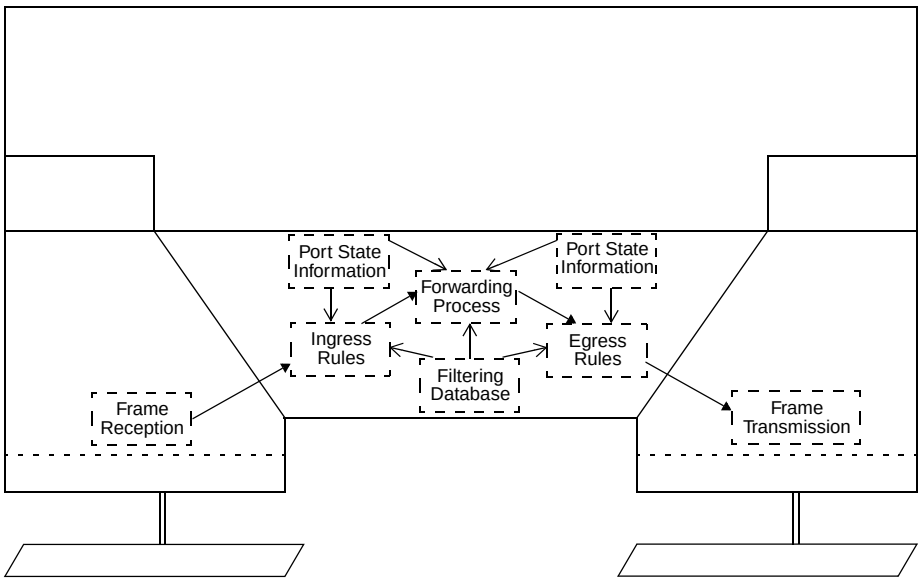


**Figure 8-4—Relaying MAC frames**


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Figure 8-5 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.

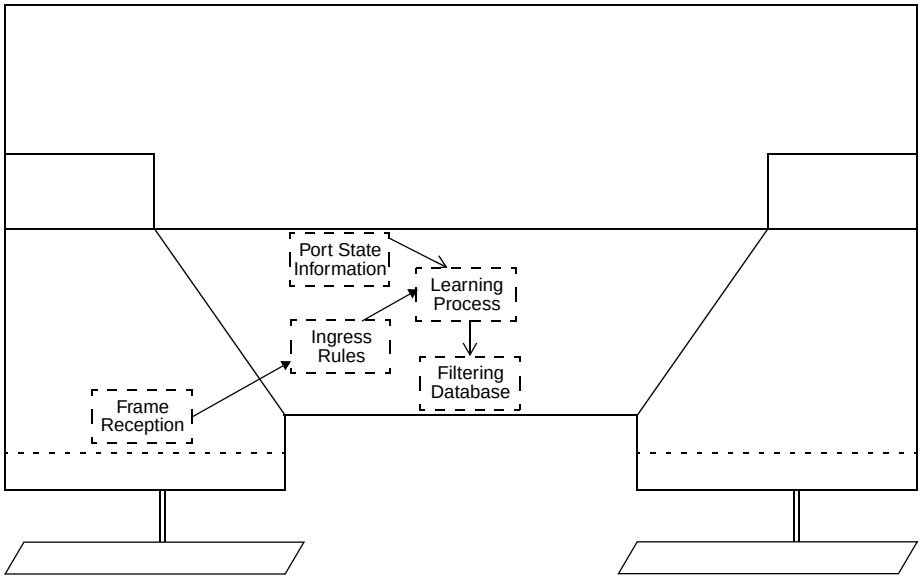


**Figure 8-5—Observation of network traffic**

Figure 8-6 illustrates the reception and transmission of Bridge Protocol Data Units by the Bridge Protocol Entity.

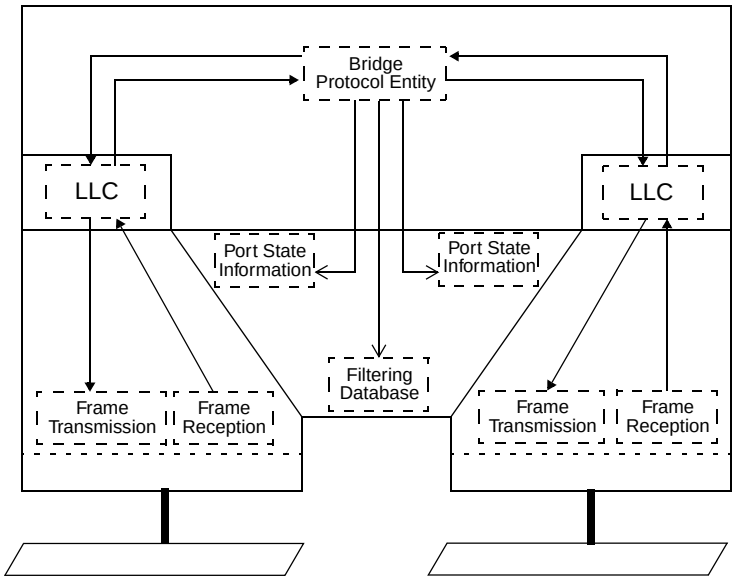


**Figure 8-6—Operation of inter-bridge protocol**

Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Figure 8-7 illustrates the reception and transmission of GARP Protocol Data Units by a GARP Protocol Entity (8.12).

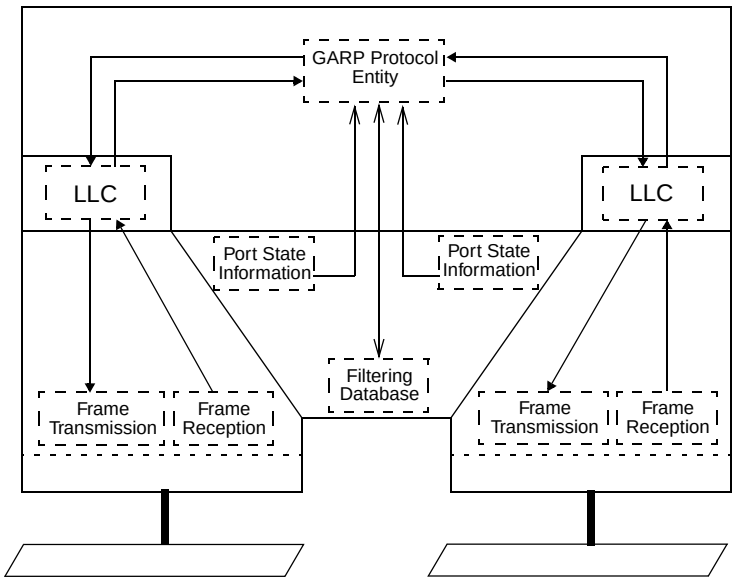


**Figure 8-7—Operation of the GARP protocol**

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

### 8.4.1 Forwarding states

State information associated with each Bridge Port governs whether or not it participates in relaying MAC frames. A Port can be disabled by management, in which case it plays no part in the operation of the Bridged LAN; a Port that is not disabled can be dynamically excluded from participation in frame relaying by operation of the Spanning Tree algorithm. If neither of these applies to a Port, it is described as *forwarding*.

The *active topology* of a Bridged LAN 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 (ISO/IEC 15802-3, Clause 8) is to construct an active topology that is simply connected relative to communication between any given pair of MAC Addresses used to address end stations on the LANs.

Figure 8-6 illustrates the operation of the Bridge 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 Bridged LAN. The Port states associated with the determination of the active topology are specified in detail in ISO/IEC 15802-3, 8.4.

Figure 8-4 illustrates the Forwarding Process's use of Port state information: first, for a Port receiving a frame, in order 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.

### 8.4.2 Learning states

The incorporation of end station location information in the Filtering Database by the Learning Process also depends on the active topology. If information associated with frames received on a Port is to be incorpo-


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

rated in the Filtering Database by the Learning Process, then the Port is described as being in a learning state; otherwise, it is in a non-learning state. Figure 8-5 illustrates the use of the Port state information for a Port receiving a frame, by the Learning Process, in order to determine whether the station location information is to be incorporated in the Filtering Database.

### 8.4.3 Acceptable Frame Types

Associated with each Port of a VLAN Bridge is an Acceptable Frame Types parameter that controls the reception of VLAN-tagged and non VLAN-tagged frames on that Port. Valid values for this parameter are

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

If this parameter is set to *Admit Only VLAN-tagged frames*, any frames received on that Port that carry no VID (i.e., untagged frames or priority-tagged frames) are discarded by the ingress rules (8.6).

Frames that are not discarded as a result of this parameter value are classified and processed according to the ingress rules that apply to that Port.

Each Port of the Bridge shall support at least one of these values, and may support both. Where both values are supported,

c) The implementation shall support the ability to configure the value of the parameter by means of the management operations defined in Clause 12; and
d) The default value of the parameter shall be *Admit All Frames*.

### 8.4.4 Port VLAN identifier

In Port-based VLAN classification within a Bridge, the VID associated with an untagged or priority-tagged frame (i.e., a frame with no tag header, or a frame with a tag header that carries the null VLAN ID) is determined, based on the Port of arrival of the frame into the Bridge, as described in 8.6. This classification mechanism requires the association of a specific VLAN ID, the *Port VLAN Identifier*, or *PVID*, with each of the Bridge's Ports.

The PVID for a given Port provides the VID for untagged and priority-tagged frames received through that Port. The PVID for each Port shall contain a valid VID value, and shall not contain the value of the null VLAN ID (Table 9-2).

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

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

### 8.4.5 Enable Ingress Filtering

An Enable Ingress Filtering parameter is associated with each Port. If the Enable Ingress Filtering parameter for a given Port is set, the ingress rules (8.6) shall discard any frame received on that Port whose VLAN classification does not include that Port in its Member set (8.11.9). If the parameter is reset for that Port, the ingress rules shall not discard frames received on that Port on the basis of their VLAN classification.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 43 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

The default value for this parameter is reset, i.e., Disable Ingress Filtering, for all Ports. The value of this parameter may be configured by means of the management operations defined in Clause 12, if management operations are supported by the implementation. If the implementation supports the ability to enable Ingress Filtering on any Port, then it shall also support the ability to disable Ingress Filtering on those Ports.

## 8.5 Frame reception

The individual MAC Entity associated with each Bridge Port examines all frames received on the LAN to which it is attached.

All error-free received frames give rise to EM_UNITDATA indication primitives, which shall be handled as follows.

NOTE—A frame that is in error, as defined by the relevant MAC specification, is discarded by the MAC Entity without giving rise to any EM_UNITDATA indication: see 7.2 and ISO/IEC 15802-3, 6.4.

Frames with EM_UNITDATA.indication primitive frame_type and mac_action parameter values of user_data_frame and request_with_no_response, respectively (7.2 and ISO/IEC 15802-3, 6.4), shall be submitted to the ingress rules (8.6).

Frames with other values of frame_type and mac_action parameters, (e.g., request_with_response and response frames), shall not be submitted to the ingress rules (8.6).

Frames with a frame_type of user_data_frame and addressed to the Bridge Port as an end station shall be submitted to the MAC Service user. Such frames carry either the individual MAC Address of the Port or a group address associated with the Port (8.14) in the destination address field. Frames submitted to the MAC Service user can also be submitted to the ingress rules (8.6), as specified above.

Frames addressed to a Bridge Port as an end station, and relayed to that Bridge Port from other Bridge Ports in the same Bridge by the Forwarding Process, shall also be submitted to the MAC Service user.

NOTE—The consequence of the above is that frames “relayed to that Bridge Port” are both submitted to that Port’s MAC Service user and transmitted on the LAN to which that Port is attached (see 8.14.7).

No other frames shall be submitted to the MAC Service user.

### 8.5.1 Regenerating user priority

The user_priority of received frames 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 user_priority (0 through 7). Each entry specifies, for the given value of received user_priority, the corresponding Regenerated user_priority value.

NOTE 1—IEEE 802 LAN technologies signal a maximum of eight user_priority values. Annex H.2 of ISO/IEC 15802-3 contains further explanation of the use of user_priority values and how they map to traffic classes.

Table 8-1 defines the default values of Regenerated user_priority for the eight possible values of the user_priority parameter received in a data indication; these values shall be used as the initial values of the corresponding entries of the User Priority Regeneration Table for each Port.

Optionally, the ability to modify the values in the User Priority Regeneration Table by management means may be supported, as described in Clause 12. If this capability is provided, the value of the table entries may


Copyright © 1999 IEEE. All rights reserved.

be independently settable for each reception Port and for each value of received user_priority, and the Bridge may have the capability to use the full range of values in the parameter ranges specified in Table 8-1.

NOTE 2—It is important to ensure that the regeneration and mapping of user priority within the Bridge is consistent with the end-to-end significance attached to that user priority in the Bridged LAN. 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 Bridged LAN, and should generate appropriate access priority values for each transmission MAC method. The user priority value regenerated via the User Priority Regeneration Table on reception is used:

— Via the traffic class table (8.7.3) to determine the traffic class for a given outbound Port, and
— Via fixed, MAC method-specific mappings (8.7.5) to determine the access priority that will be used for a given outbound MAC method.

Table 8-1 shows the default values for the regeneration of user priority. Table 8-2 shows the default values for the traffic class table, for all possible numbers of supported traffic classes. Table 8-3 shows the fixed mappings from user priority to access priority that are required for different outbound MAC methods.

**Table 8-1—User priority regeneration**

| User priority | Default regenerated user 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 |

## 8.6 The ingress rules

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

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

a) If the vlan_identifier parameter carried in a received data indication is the null VLAN ID (Table 9-2), then
   1) If the implementation supports further VLAN classification rules in addition to Port-based classification (D.2.2), and if the application of these rules associates a non-null VID value with the frame, then that VID value is used.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 45 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

2) If the implementation supports only Port-based classification, or if any additional classification rules supported are unable to associate a non-null VID with the frame, then the PVID value associated with the Port through which the frame was received is used (8.4.4).

b) If the vlan_identifier parameter carried in a received data indication is not the null VLAN ID (Table 9-2), then the vlan_identifier parameter value is used.

NOTE 1—As defined in 7.1.2, the vlan_identifier parameter carries the null VLAN ID if the frame was not VLAN-tagged. There are two cases; either the frame was untagged, or the frame was tagged and the tag header carried a VID value equal to the null VLAN ID (i.e., a priority-tagged frame).
NOTE 2—VIDs of value FFF cannot be configured in any Filtering Database entry (see Table 9-2). Consequently, any incoming frame whose VLAN classification is FFF will be discarded by the Forwarding Process.

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

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

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

## 8.7 The Forwarding Process

Frames submitted to the Forwarding Process after being received at any given Bridge Port (8.5) shall be forwarded through the other Bridge Ports subject to the constituent functions of the Forwarding Process. These functions enforce topology restrictions (8.7.1), use Filtering Database information to filter frames (8.7.2), queue frames (8.7.3), select queued frames for transmission (8.7.4), map priorities (8.7.5), and recalculate FCS if required (8.7.6).

The Forwarding Process functions are described in 8.7.1–8.7.6 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 is discarded without being transmitted at the other Ports.

NOTE—The model of operation of the Forwarding Process described in this standard is limited to the operation of the relay function of the MAC Bridge, and does not take into consideration what may occur in real implementations once frames are passed to the MAC for transmission. In some MAC implementations, and under some traffic conditions, a degree of indeterminacy may be introduced between the modeled description of the process of passing selected frames to the MAC for transmission and the actual sequence of frames as visible on the LAN medium itself. Examples can be found in the handling of access_priority in Token-Passing Bus MACs, or in the effect of different values for Token Holding Time in FDDI LANs. Such indeterminacy could result in apparent violation of the queuing/de-queueing and prioritizing rules described for the Forwarding Process, when observing traffic on the medium. As a consequence, in some implementations of this standard, it may prove to be impossible to test conformance to the standard simply by relating observed LAN traffic to the described model of the forwarding process; conformance tests would have to allow for the (permissible) behavior of the MAC implementations as well.

Figure 8-4 illustrates the operation of the Forwarding Process in a single instance of frame relay between the Ports of a Bridge with two Ports. Figure 8-8 illustrates the detailed operation of the Forwarding Process.

### 8.7.1 Enforcing topology restriction

Each Port is selected as a potential transmission Port if, and only if

a) The Port on which the frame was received was in a forwarding state (ISO/IEC 15802-3, 8.4), and
b) The Port considered for transmission is in a forwarding state, and


Copyright © 1999 IEEE. All rights reserved.

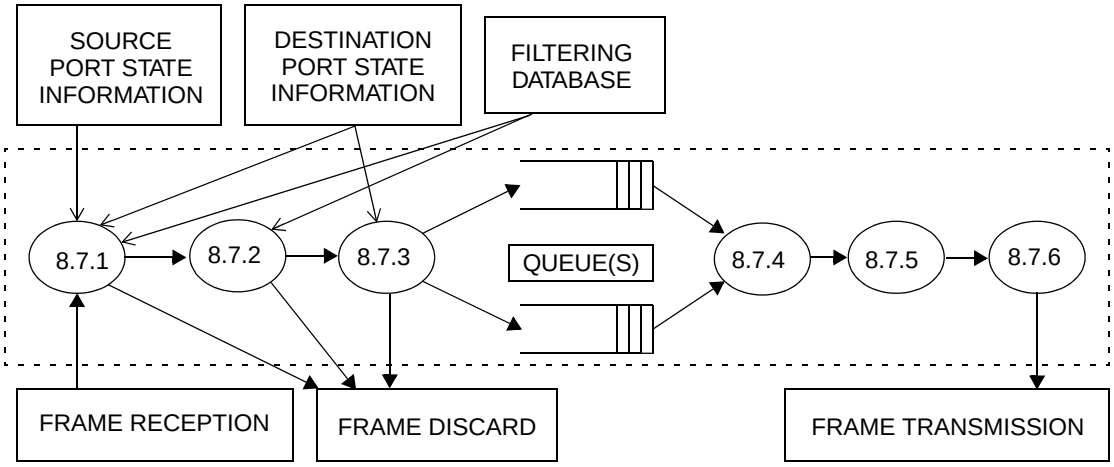


**Figure 8-8—Illustration of the detailed operation of the Forwarding Process**

c) The Port considered for transmission is not the same as the Port on which the frame was received, and
d) The size of the mac_service_data_unit conveyed by the frame does not exceed the maximum size of mac_service_data_unit supported by the LAN to which the Port considered for transmission is attached.

For each Port not selected as a potential transmission Port the frame shall be discarded.

### 8.7.2 Filtering frames

Filtering decisions are taken by the Forwarding Process on the basis of

a) The destination MAC Address carried in a received frame;
b) The VID associated with the received frame;
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 (8.11.6).

For each potential transmission Port selected as in 8.7.1, the frame shall be forwarded, or discarded (i.e., filtered), on the basis of this information, in accordance with the definition of the Filtering Database entry types (8.11.1, 8.11.3, and 8.11.4). The required forwarding and filtering behavior is summarized in 8.11.6, 8.11.8, Table 8-5, Table 8-6, and Table 8-7.

### 8.7.3 Queuing frames

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

a) Unicast frames with a given user_priority (regenerated as defined in 8.5.1) for a given combination of destination_address and source_address;
b) Group-addressed frames with a given user_priority (regenerated as defined in 8.5.1) for a given destination_address.

The Forwarding Process may provide more than one transmission queue for a given Bridge Port. Frames are assigned to storage queue(s) on the basis of their user_priority using a traffic class table that is part of the

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 47 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

state information associated with each Port. The table indicates, for each possible value of user_priority, the corresponding value of traffic class that shall be assigned. Values of user_priority range from 0 through 7. Queues correspond one-to-one with traffic classes.

NOTE 1—Annex H.2 of ISO/IEC 15802-3 contains further explanation of the use of user_priority values and how they map to traffic classes.

For management purposes, up to eight traffic classes are supported by the traffic class tables in order to allow for separate queues for each level of user_priority. Traffic classes are numbered 0 through N-1, where N is the number of traffic classes associated with a given outbound Port. Management of traffic class information is optional. Traffic class 0 corresponds to non-expedited traffic; non-zero traffic classes are expedited classes of traffic.

NOTE 2—In a given Bridge, it is permissible to implement different numbers of traffic classes for each Port. Ports associated with MAC methods that support a single transmission priority, such as CSMA/CD, can support more than one traffic class.

Where the Forwarding Process does not support expedited classes of traffic for a given Port, in other words, where there is a single traffic class associated with the Port, all values of user_priority map to traffic class 0. In bridges which support expedited traffic, the recommended mapping of user_priority to traffic class, for the number of traffic classes implemented, is as shown in Table 8-2. Each entry in the body of the table is the traffic class assigned to traffic with a given user_priority, for a given number of available traffic classes.

**Table 8-2—Recommended user priority to traffic class mappings**

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

NOTE—The rationale behind the choice of values shown in this table is discussed in Annex H.2 of ISO/IEC 15802-3. A consequence of the mapping shown is that frames carrying the default user priority are given preferential treatment relative to user priority 1 and 2 in Bridges that implement four or more Traffic Classes.

A frame queued by the Forwarding Process for transmission on a Port shall be removed from that queue on submission to the individual MAC Entity for that Port. No further attempt shall be made to transmit the frame on that Port even if the transmission is known to have failed.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

A frame queued by the Forwarding Process for transmission on a Port can be removed from that queue, and not subsequently transmitted, if the time for which buffering is guaranteed has been exceeded for that frame.

A frame queued for transmission on a Port shall be removed from that queue if that is necessary to ensure that the maximum bridge transit delay (ISO/IEC 15802-3, 6.3.6) will not be exceeded at the time at which the frame would subsequently be transmitted.

A frame queued for transmission on a Port shall be removed from that queue if the associated Port leaves the forwarding state.

Removal of a frame from a queue for any particular Port does not of itself imply that it shall be removed from a queue for transmission on any other Port.

### 8.7.4 Selecting frames for transmission

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 8.7.3.

Additional algorithms, selectable by management means, may be supported as an implementation option so long as the requirements of 8.7.3 are met.

### 8.7.5 Mapping priority

The user_priority parameter in an EM_UNITDATA.request primitive (7.1) shall be equal to the user_priority parameter in the corresponding data indication.

The mapping of user_priority to outbound access_priority is achieved via fixed, MAC method-specific mappings. The access_priority parameter in an EM_UNITDATA.request primitive (7.1) shall be determined from the user_priority in accordance with the values shown in Table 8-3 for the MAC methods that will carry the data request. The values shown in Table 8-3 are not modifiable by management or other means.

The table shows two columns for the 8802-5 MAC method. The mapping in the column marked “8802-5 (alternate)” is included in order to permit backwards compatibility with equipment manufactured in accordance with ISO/IEC 10038: 1993; however, the use of this mapping reduces the number of available access priority values to three. For this reason, it is recommended that the column marked “8802-5 (default)” is supported as the default mapping where backward compatibility is not an issue.

### 8.7.6 Recalculating FCS

Where a frame is being forwarded between two individual MAC Entities of the same IEEE 802 LAN type, and relaying the frame involves no changes to the data that is within the FCS coverage, the FCS received in the EM_UNITDATA.indication primitive may be supplied in the corresponding EM_UNITDATA.request primitive and not recalculated (7.1, 7.2, ISO/IEC 15802-3, 6.3.7).

Where a frame is being forwarded between two individual MAC Entities of different types, recalculation of the FCS is necessary if the differences between the LAN MAC methods is such that an FCS calculated according to the MAC procedures for the destination MAC method would differ from the FCS carried by the

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 49 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

**Table 8-3—Outbound access priorities**

| user_priority | Outbound Access Priority per MAC method | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| | 802.3 | 8802-4 | 8802-5 (default) | 8802-5 (alternate) | 8802-6 | 802.9a* | 8802.11 | 8802-12 | FDDI |
| **0** | 0 | 0 | 0 | 4 | 0 | 0 | 0 | 0 | 0 |
| **1** | 0 | 1 | 1 | 4 | 1 | 0 | 0 | 0 | 1 |
| **2** | 0 | 2 | 2 | 4 | 2 | 0 | 0 | 0 | 2 |
| **3** | 0 | 3 | 3 | 4 | 3 | 0 | 0 | 0 | 3 |
| **4** | 0 | 4 | 4 | 4 | 4 | 0 | 0 | 4 | 4 |
| **5** | 0 | 5 | 5 | 5 | 5 | 0 | 0 | 4 | 5 |
| **6** | 0 | 6 | 6 | 6 | 6 | 0 | 0 | 4 | 6 |
| **7** | 0 | 7 | 6 | 6 | 7 | 0 | 0 | 4 | 6 |

*In the absence of a definition, in ISO/IEC 15802-3, 6.5, of support by IEEE Std 802.9a-1995, it is assumed that for this MAC method, access priority 0 will map to “low.”

received frame, or if relaying the frame involves changes to the data that is within the FCS coverage. Where necessary, the FCS is recalculated according to the specific MAC procedures of the transmitting MAC entity.

NOTE—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 may be preferable in terms of its ability to protect against increased levels of undetected frame errors. ISO/IEC 15802-3, Annex G, 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.

FCS recalculation is necessary if any of the following conditions are true:

a) The algorithm used to determine the FCS differs between the MAC methods used by the two MAC entities;
b) The FCS coverage differs between the MAC methods used by the two MAC entities;
c) Relaying the frame between the two MAC entities involves changes to the data that is within the coverage of the FCS (e.g., the frame was tagged on one link, but not on the other).

## 8.8 The egress rules

Frames shall be filtered, i.e., discarded, if

a) For the frame’s VID, as determined by the ingress rules (8.6), the transmission Port is not present in the Member set (8.11.9); or
b) The value of the include_tag parameter, determined as shown below, is False, and the Bridge does not support the ability to translate embedded MAC Address information from the format indicated


Copyright © 1999 IEEE. All rights reserved.

by the canonical_format_indicator parameter to the format appropriate to the MAC method on which the data request will be carried.

NOTE 1—The meanings of the terms Canonical format and Non-canonical format are discussed in Annex F.

The value of the include_tag parameter in the data request primitive is determined as follows:

c) If, for the frame's VID, as determined by the ingress rules (8.6), the transmission Port is present in the untagged set (8.11.9), then the value False is used. Otherwise;
d) The value True is used.

NOTE 2—As all incoming frames, including priority-tagged frames, are classified as belonging to a VLAN by the ingress rules (8.6), the transmitting Port only transmits VLAN-tagged frames or untagged frames, and can never transmit priority-tagged frames. Hence, a station sending a priority-tagged frame via a VLAN Bridge will receive a response that is either VLAN-tagged or untagged, depending upon the state of the untagged set for the VLAN concerned.

The value of the canonical_format_indicator parameter of the data request primitive is equal to the value of that parameter as received in the corresponding data indication.

## 8.9 Frame transmission

The individual MAC Entity associated with each Bridge Port transmits frames submitted to it by the MAC Relay Entity.

Relayed frames are submitted for transmission by the Forwarding Process. The EM_UNITDATA.request primitive associated with such frames conveys the values of the source and destination address fields received in the corresponding EM_UNITDATA.indication primitive.

LLC Protocol Data Units are submitted by LLC as a user of the MAC Service provided by the Bridge Port. Frames transmitted to convey such Protocol Data Units carry the individual MAC Address of the Port in the source address field.

Each frame is transmitted subject to the MAC procedures to be observed for that specific IEEE 802 LAN technology. The values of the frame_type and mac_action parameters of the corresponding EM_UNITDATA.request primitive shall be user_data_frame and request_with_no_response respectively (7.2; ISO/IEC 15802-3, 6.5).

Frames transmitted following a request by the LLC user of the MAC Service provided by the Bridge Port shall also be submitted to the MAC Relay Entity.

## 8.10 The Learning Process

The Learning Process observes the source MAC Addresses of frames received on each Port and updates the Filtering Database conditionally on the state of the receiving Port. The VID associated with the frame is used to ensure that the address information is learned relative to the frame's VLAN.

Frames are submitted to the Learning Process by the ingress rules as specified in 8.6.

The Learning Process can deduce the Port through which particular end stations in the Bridged LAN can be reached by inspection of the source MAC Address field and VID of received frames. It records such information in the Filtering Database (8.11). It shall create or update a Dynamic Filtering Entry (8.11.3) associated with the frame's VID (8.11.9), associating the reception Port with the source MAC Address, if and only if

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 51 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

a) The Port on which the frame was received is in a state that allows learning (ISO/IEC 15802-3, 8.4), and
b) The source address field of the frame denotes a specific end station, i.e., is not a group MAC Address, and
c) The resulting number of entries would not exceed the capacity of the Filtering Database, and
d) The Member set (8.11.9) for the frame's VID includes at least one Port.

NOTE—If the Member set for a given VID is the empty set, then that VLAN is not currently active, and the Bridge will therefore filter all frames destined for that VLAN, regardless of their destination address. There is therefore no reason to include MAC Address filtering information in the Filtering Database for that VLAN until such a time as it becomes active.

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

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.

## 8.11 The Filtering Database

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 (8.7.1 and 8.7.2). 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.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Three entry types are used to represent dynamic filtering information. 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 (8.10), and are subject to ageing and removal by the Filtering Database. 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 (8.11.4; ISO/IEC 15802-3, 6.6.5; ISO/IEC 15802-3, Clause 10). 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).

Static Filtering Entries and Group Registration Entries comprise

h) A MAC Address specification;
i) A VLAN Identifier (VID);
j) 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

k) A MAC Address specification;
l) A locally significant Filtering Identifier (FID; see 8.11.7);
m) 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

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

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:

p) All Group Addresses, for which no more specific Static Filtering Entry exists;
q) 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 p) 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 (8.10). In Bridges that support Extended Filtering Services, the Filtering Database shall support the creation, updating, and removal of Group Registration Entries by GMRP (ISO/IEC 15802-3, Clause 10).

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 53 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Figure 8-4 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 8-5 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 ISO/IEC 15802-3-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 8-6 illustrates the operation of the Bridge Protocol Entity (8.12), 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:

r) 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
s) Of that total number, the total number of VLAN Registration Entries (static and dynamic) that the Filtering Database can support.

### 8.11.1 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 (8.11.6) 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).


Copyright © 1999 IEEE. All rights reserved.

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 (ISO/IEC 15802-3, Clauses 10, 12, and 12.9.1). 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.

### 8.11.2 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 8.11.9. The values that can be represented are
        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.

### 8.11.3 Dynamic Filtering Entries

A Dynamic Filtering Entry specifies

a) An individual MAC Address;
b) The FID, an identifier assigned by the MAC Bridge (8.11.7) 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* (3.9) takes place. Any pair of FIDs identifies two sets of VLANs between which *Independent VLAN Learning* (3.5) takes place. The allocation of FIDs by a Bridge is described in 8.11.7.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 55 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

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 (8.10). 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 8-4; 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 8-4—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 ISO/IEC 15802-3, Clause 8, 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 (ISO/IEC 15802-3, 8.3.5). 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.

### 8.11.4 Group Registration Entries

A Group Registration Entry specifies

a) A MAC Address specification, comprising
   1) A group MAC Address; or


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

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 (ISO/IEC 15802-3, Clause 10, 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.

### 8.11.5 Dynamic VLAN Registration Entries

A Dynamic VLAN Registration Entry specifies

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

### 8.11.6 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 8.11.1, 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.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 57 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

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 8.11.8, Table 8-6, and Table 8-7.

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.

### 8.11.7 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 (8.11.9) 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 (8.4.5) disabled;


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

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.

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

NOTE—In an SVL/IVL Bridge (3.11), a number of FIDs are supported, and one or more VID can be mapped to each FID. In an SVL Bridge (3.10), a single FID is supported, and all VIDs are mapped to that FID. In an IVL Bridge (3.6), a number of FIDs are supported, and only one VID can be mapped to each FID.

### 8.11.7.1 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 FID X; or
c) A dynamic allocation has been defined (as a result of applying the VLAN Learning Constraints), allocating this VID to FID X.

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.

### 8.11.7.2 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.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 59 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

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 (8.11.7),

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.

### 8.11.7.3 VLAN Learning Constraint inconsistencies and violations

The application of the rules specified in 8.11.7.2, 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 (8.11.7), 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 © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Learning Constraint violations are detected

d) When a VLAN that was hitherto not a member of the set of active VLANs (8.11.7) 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.

### 8.11.8 Querying the Filtering Database

If a frame is classified into a VLAN containing a given outbound Port in its member set (8.11.9), 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, 8.11.3) 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 8.11.7.

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 8-5 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 8-6 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 8-7 combines Static Filtering Entry and Group Registration Entry information for a specific group MAC Address with the Table 8-6 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.

Copyright © 1999 IEEE. All rights reserved.

**Table 8-5—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 8-6—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 |

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

d) The results shown in Table 8-7 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 8-7 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.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

**Table 8-7—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 8-6): | Not Registered | All Unregistered Group Addresses control elements for this VID and Port specify (Table 8-6): | 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) |

### 8.11.9 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 (8.11.2) and/or a Dynamic VLAN Registration Entry (8.11.5). This information defines, for that VLAN:

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 operation of the ingress rules (8.6) 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 8-8 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.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 63 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

**Table 8-8—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 |

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,

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 in determining the operation of the ingress rules (8.6) and the egress rules (8.8) 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*).

### 8.11.10 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—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.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## 8.12 Bridge Protocol Entity and GARP Protocol Entities

The Bridge Protocol Entity operates the Spanning Tree Algorithm and Protocol.

The Bridge Protocol Entities of Bridges attached to a given individual LAN in a Bridged LAN communicate by exchanging Bridge Protocol Data Units (BPDUs).

Figure 8-6 illustrates the operation of the Bridge 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.

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 (ISO/IEC 15802-3, Clauses 10 and 12.3).

The GARP Protocol Entities of Bridges attached to a given individual LAN in a Bridged LAN communicate by exchanging GARP Protocol Data Units (GARP PDUs).

Figure 8-7 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

Remote management facilities may be provided by the Bridge. Bridge Management is modeled as being performed by means of the Bridge Management Entity. The facilities provided by Bridge Management, and the operations that support these facilities, are specified in Clause 12.

Bridge Management protocols use the MAC Service provided by the Bridged LAN.

## 8.14 Addressing

All MAC Entities communicating across a Bridged LAN shall use 48-bit addresses. These may be Universally Administered Addresses, Locally Administered Addresses, or a combination of both.

### 8.14.1 End stations

Frames transmitted between end stations using the MAC Service provided by a Bridged LAN carry the MAC Address of the source and destination 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 end stations for the purpose of frame relay in the Bridged LAN.

The broadcast address and other group MAC Addresses apply to the use of the MAC Service provided by a Bridged LAN as a whole. In the absence of explicit filters configured via management as Static Filtering Entries, or via GMRP as Group Registration Entries (8.11, Clause 12, and ISO/IEC 15802-3, Clause 10), frames with such destination addresses are relayed throughout the Bridged LAN.

### 8.14.2 Bridge Ports

The individual MAC Entity associated with each Bridge Port shall have a separate individual MAC Address. This address is used for any MAC procedures required by the particular MAC method employed.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 65 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Frames that are received from the LAN to which a Port is attached and that carry a MAC Address for the Port in the destination address field are submitted to the MAC Service User (LLC) exactly as for an end station.

### 8.14.3 Bridge Protocol Entities and GARP Protocol Entities

Bridge Protocol Entities only receive and transmit BPDUs. These are only received and transmitted from other Bridge Protocol Entities (or where two Bridge Ports are connected to the same LAN, to and from themselves).

GARP Protocol Entities only receive and transmit GARP PDUs (ISO/IEC 15802-3, 12.11) that are formatted according to the requirements of the GARP Applications they support. These are only received and transmitted from other GARP Protocol Entities.

A Bridge Protocol Entity or a GARP Protocol Entity uses the DL_UNITDATA.request primitive (see ISO/IEC 8802-2) provided by the individual LLC Entities associated with each active Bridge Port to transmit BPDUs or GARP PDUs. Each PDU is transmitted on one selected Bridge Port. PDUs are received through corresponding DL_UNITDATA.indication primitives. The source_address and destination_address parameters of the DL_UNITDATA.request primitive shall both denote the standard LLC address assigned to the Bridge Spanning Tree Protocol. This identifies the Bridge Protocol Entity and the GARP Protocol Entity among other users of LLC.

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. The source and destination LLC address fields are set to the values supplied in the request primitive.

The value assigned to the Bridge Spanning Tree Protocol LLC address is given in Table 8-9.[8]

**Table 8-9—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 ISO/IEC 11802-1: 1997.

ISO/IEC 15802-3 defines a Protocol Identifier field, present in all BPDUs (ISO/IEC 15802-3, Clause 9) and GARP PDUs (ISO/IEC 15802-3, 12.11), which serves to identify different protocols supported by Bridge Protocol Entities and GARP Protocol Entities, within the scope of the LLC address assignment. This standard specifies a single value of the Protocol Identifier, defined in ISO/IEC 15802-3, Clause 9, for use in BPDUs. This value serves to identify BPDUs exchanged between Bridge Protocol Entities operating the Spanning Tree Algorithm and Protocol specified in ISO/IEC 15802-3, Clause 8. A second value of this protocol identifier for use in GARP PDUs is defined in ISO/IEC 15802-3, 12.11. This value serves to identify GARP PDUs exchanged between GARP Participants operating the GARP protocol specified in ISO/IEC 15802-3, Clause 12. Further values of this field are reserved for future standardization.

[8]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 © 1999 IEEE. All rights reserved.

A Bridge Protocol Entity or GARP Protocol Entity that receives a BPDU or a GARP PDU with an unknown Protocol Identifier shall discard that PDU.

A Bridge Protocol Entity that operates the Spanning Tree Algorithm and Protocol specified in ISO/IEC 15802-3, Clause 8, always transmits BPDUs addressed to all other Bridge Protocol Entities attached to the LAN on which the frame containing the BPDU is transmitted. A group address shall be used in the destination address field to address this group of Entities. This group address shall be configured in the Permanent Database (8.14.6) in order to confine BPDUs to the individual LAN on which they are transmitted.

A 48-bit Universal Address, known as the Bridge Group Address, has been assigned for this purpose. Its value is specified in Table 8-10. Bridges that use 48-bit Universally Administered Addresses shall use this address in the destination address field of all MAC frames conveying BPDUs.

**Table 8-10—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 |
| Reserved for future standardization | 01-80-C2-00-00-02 |
| Reserved for future standardization | 01-80-C2-00-00-03 |
| Reserved for future standardization | 01-80-C2-00-00-04 |
| Reserved for future standardization | 01-80-C2-00-00-05 |
| Reserved for future standardization | 01-80-C2-00-00-06 |
| Reserved for future standardization | 01-80-C2-00-00-07 |
| Reserved for future standardization | 01-80-C2-00-00-08 |
| Reserved for future standardization | 01-80-C2-00-00-09 |
| Reserved for future standardization | 01-80-C2-00-00-0A |
| Reserved for future standardization | 01-80-C2-00-00-0B |
| Reserved for future standardization | 01-80-C2-00-00-0C |
| Reserved for future standardization | 01-80-C2-00-00-0D |
| Reserved for future standardization | 01-80-C2-00-00-0E |
| Reserved for future standardization | 01-80-C2-00-00-0F |

A GARP Protocol Entity that

a) Operates the GARP protocol specified in ISO/IEC 15802-3, Clause 12; and
b) Supports a given GARP Application,

always transmits GARP PDUs addressed to all other GARP Protocol Entities that

c) Implement the same GARP Application; and
d) Are attached to the LAN on which the frame containing the GARP PDU is transmitted.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 67 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

A group MAC Address, specific to the GARP Application concerned, shall be used as the destination MAC Address field to address this group of GARP Protocol Entities. A set of 48-bit Universal Addresses, known as GARP Application addresses, have been assigned for that purpose. The values of the GARP Application addresses are defined in ISO/IEC 15802-3, 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 ISO/IEC 15802-3, Table 12-1.

In Bridges that do not support any GARP applications, the set of GARP Application addresses should not be configured in the Filtering Database (8.11) or the Permanent Database (8.11.10). In Bridges that support one or more GARP applications, the set of GARP Application addresses should be configured as Static Filtering Entries in the Filtering Database (8.11.1) and Permanent Database (8.11.10) as follows:

e) GARP Application addresses assigned to GARP Applications that are supported by the Bridge should be configured in order to confine GARP PDUs for that GARP Application to the individual LAN on which they are transmitted;

f) GARP Application addresses assigned to GARP Applications that are not supported by the Bridge should not be configured in the Filtering Database or Permanent Database.

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).

### 8.14.4 Bridge Management Entities

Bridge Management Entities transmit and receive protocol data units using the Service provided by the individual LLC Entities associated with each Bridge Port. Each of these in turn uses the MAC Service, which is provided by the individual MAC Entities associated with that Port and supported by the Bridged LAN as a whole.

As a user of the MAC Service provided by a Bridged LAN, the Bridge Management Entity may be attached to any point in the Bridged LAN. Frames addressed to the Bridge Management Entity will be relayed by Bridges if necessary to reach the LAN to which it is attached.

In order to ensure that received frames are not duplicated, the basic requirement in a single LAN or a Bridged LAN that a unique address be associated with each point of attachment shall be met.

A Bridge Management Entity for a specific Bridge is addressed by one or more individual MAC Addresses in conjunction with the higher layer protocol identifier and addressing information. It may share one or more points of attachment to the Bridged LAN with the Ports of the Bridge with which it is associated. It is recommended that it make use of the MAC Service provided by all the MAC Entities associated with each Bridge Port, i.e., that it be reachable through each Bridge Port using frames carrying the individual MAC Address of that Port in the destination address field.

This standard specifies a standard group MAC Address for public use which serves to convey management requests to the Bridge Management Entities associated with all Bridge Ports attached to a Bridged LAN. A management request that is conveyed in a MAC frame carrying this address value in the destination address field will generally elicit multiple responses from a single Bridge. This address is known as the All LANs Bridge Management Group Address and takes the value specified in Table 8-11.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

**Table 8-11—Addressing bridge management**

| Assignment | Value |
|---|---|
| All LANs Bridge Management Group Address | 01-80-C2-00-00-10 |

### 8.14.5 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 Spanning Tree Protocol (ISO/IEC 15802-3, Clause 8) requires that a single unique identifier be associated with each Bridge. That identifier is derived from the Bridge Address as specified in ISO/IEC 15802-3, 8.5.1.3, 8.5.3.7, and 9.2.5.

### 8.14.6 Reserved addresses

Frames containing any of the group MAC Addresses specified in Table 8-10 in their destination address field shall not be relayed by the Bridge. They shall be configured in the Permanent Database. Management shall not provide the capability to modify or remove these entries from the Permanent or the Filtering Databases. 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).

### 8.14.7 Points of attachment and connectivity for Higher Layer Entities

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

Higher Layer Entities fall into two distinct categories:

a) Those entities, such as the Bridge Management Entity, that require only a single point of attachment to the Bridged LAN;
b) Those entities, such as Bridge Protocol Entities and GARP Participants, that require a point of attachment per Port of the Bridge.

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

c) It does not receive frames via a point of attachment associated with one Port that have been relayed by the Bridge from other Ports; and
d) Frames that it transmits via one point of attachment are not relayed by the Bridge to any other Ports.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 69 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

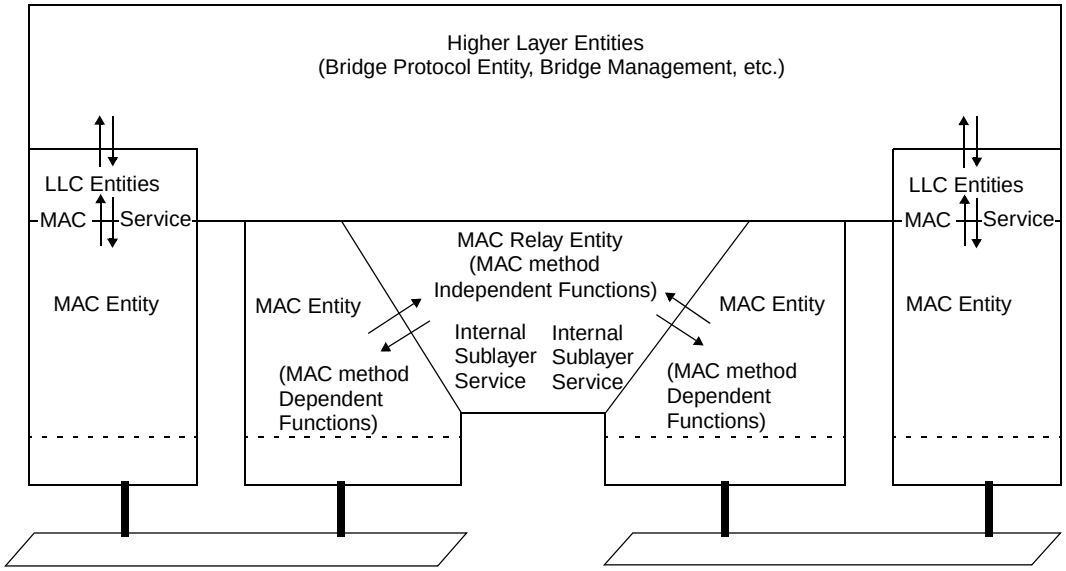


**Figure 8-9—Logical separation of points of attachment used by Higher Layer Entities and the MAC Relay Entity**

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

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

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

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

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

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

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

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


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

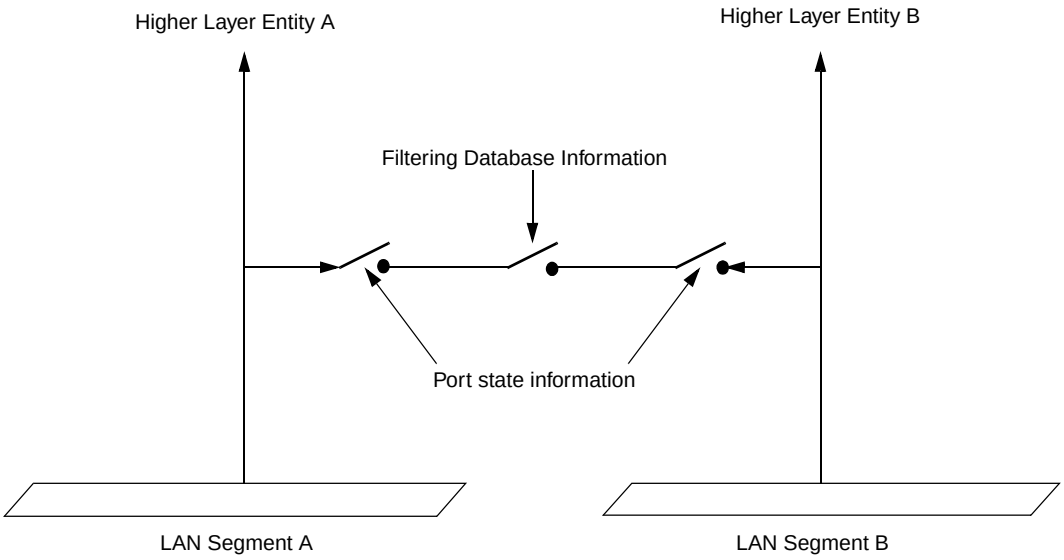


**Figure 8-10—Effect of control information on the forwarding path**

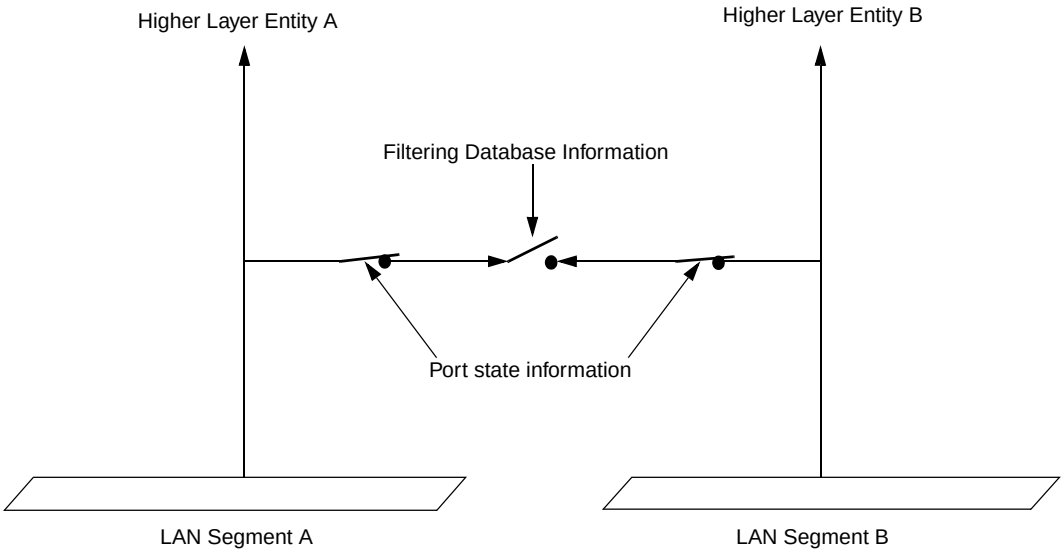


**Figure 8-11—Per-Port points of attachment**

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

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

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 71 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

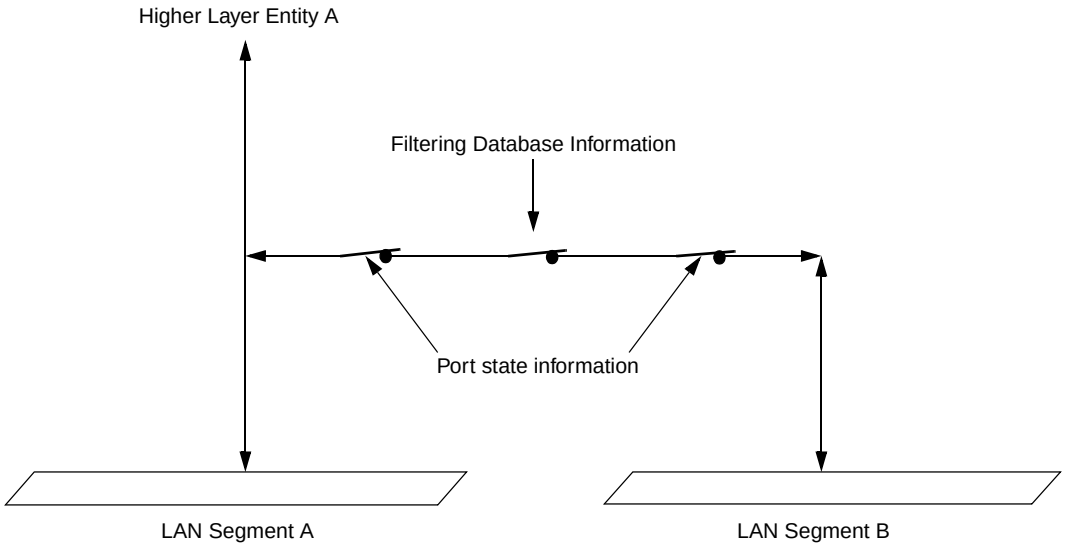


**Figure 8-12—Single point of attachment—relay permitted**

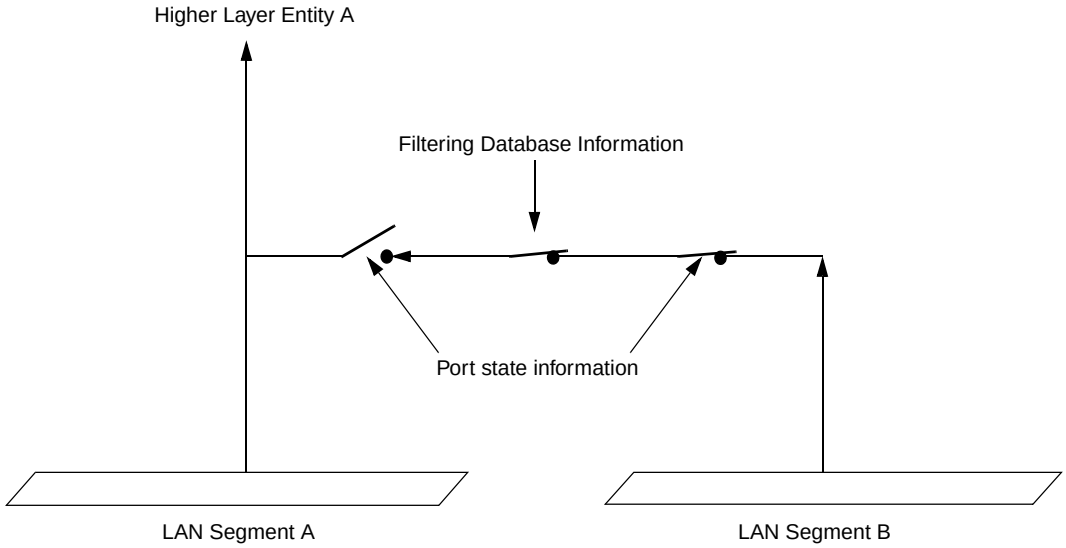


**Figure 8-13—Single point of attachment—relay not permitted**

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

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


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

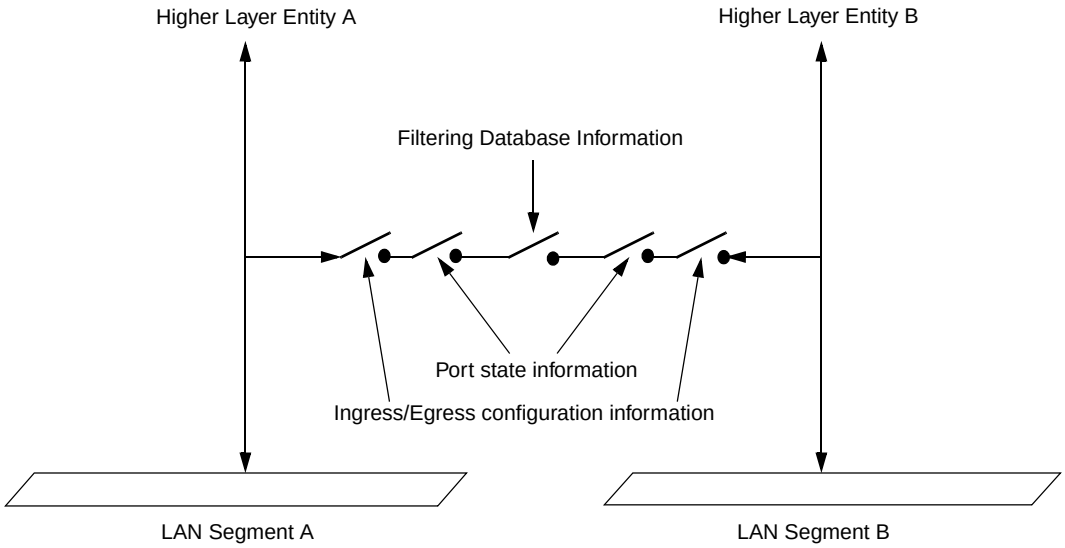


**Figure 8-14—Ingress/egress control information in the forwarding path**

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 segment as a Higher Layer Entity's point of attachment. For example, the reception of a frame by Higher Layer Entity A that was transmitted on LAN Segment A is unaffected by the ingress or egress configuration of either Port. However, for Higher Layer Entities that require only a single point of attachment, the ingress and egress configuration affects the forwarding path. For example, frames destined for Higher Layer Entity A that are transmitted on LAN Segment B would be subjected to the ingress rules that apply to Port B and the egress rules that apply to Port A.

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

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

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

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

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

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

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 73 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# 9. Tagged frame format

Tagging of frames is performed for the following purposes:

a) To allow user_priority information to be added to frames carried on IEEE 802 LAN MAC types that have no inherent ability to signal priority information at the MAC protocol level;
b) To allow a frame to carry a VID;
c) To allow the frame to indicate the format of MAC Address information carried in MAC user data;
d) To allow VLANs to be supported across different MAC types.

This clause describes the tag format used for tagging frames, as follows:

e) Subclause 9.1 gives an overview of tagging;
f) Subclause 9.2 defines the data representations that are used in the descriptions of the tag field formats;
g) Subclause 9.3 describes the structure of the tag header.

Further analysis of the frame formats, the format translations that can occur when frames are tagged or untagged when relayed between different MAC methods, and a description of the tagging/untagging procedure can be found in Annex C.

The description of the tagged frame structure, both here and in Annex C, is based on two generic frame formats:

h) The frame format used in IEEE Std 802.3 MACs, and which is used with minor variations in other MACs where the native Link Layer protocol identification mechanism is based on a choice between the Type interpretation and Length interpretation of the Length/Type field. The Type interpretation is used where a Type value provides the protocol identification; the Length interpretation is used where LLC addressing provides the protocol identification. LAN MAC methods that make use of this frame format are referred to in this standard as 802.3/Ethernet MAC methods;
i) The frame format used in ISO/IEC 8802-5 and FDDI MACs, and used with minor variations in other MACs where the native Link Layer protocol identification mechanism is based on LLC addressing, and where the frame may also be able to carry source-routing information. LAN MAC methods that make use of this frame format are referred to in this standard as Token Ring/FDDI MAC methods.

For MACs other than 802.3, 8802-5, and FDDI, the approach used is to apply these frame formats, with appropriate modification to the overall frame structure, as appropriate to the MAC concerned. For example:

j) MACs such as 8802-4 and 8802-6 that use LLC as the native Link Layer protocol identification would adopt format i). If they do not provide native support for source routing, the variant of this format that is used in transparent FDDI LANs would be used;
k) MACs such as 8802-12 that can support compatibility with 802.3 and 8802-5 MACs would adopt either format h) or format i), depending upon which compatibility mode was in operation.

## 9.1 Overview

Tagging a frame requires

a) The addition of a tag header to the frame. This header is inserted immediately following the destination MAC Address and source MAC Address (and routing, if present) fields of the frame to be transmitted;


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

b) If the source and destination MAC methods differ, tagging the frame may involve translation or encapsulation of the remainder of the frame, as specified in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390;
c) Recomputation of the Frame Check Sequence (FCS).

When relaying a tagged frame between 802.3/Ethernet MACs, a Bridge may adjust the PAD field such that the minimum size of a transmitted tagged frame is 68 octets (7.2).

The tag header carries the following information:

d) The Tag Protocol Identifier (TPID) appropriate to the MAC method concerned, as described in 9.3.1. This protocol identifier identifies the frame as a tagged frame, conforming to the tagging format described in this standard.
e) Tag Control Information (TCI) as described in 9.3.2. The TCI consists of the following elements:
    1) User_priority, as described in 9.3.2.1. This field allows the tagged frame to carry user_priority information across Bridged LANs in which individual LAN segments may be unable to signal priority information (e.g., 802.3/Ethernet segments).
    2) Canonical Format Indicator (CFI), as described in 9.3.2.2. This field is used
        i) In Token Ring/source-routed FDDI MAC methods, to signal the bit order of address information carried in the encapsulated frame; and
        ii) In 802.3/Ethernet and transparent FDDI MAC methods, to signal the presence or absence of a RIF field, and, in combination with the Non-canonical Format Indicator (NCFI) carried in the RIF, to signal the bit order of address information carried in the encapsulated frame.

NOTE 1—The meaning of Canonical format as applied to MAC Addresses, and the implications of the format of addresses on the requirements for frame translation, are discussed in Annex F.

    3) VLAN Identifier (VID), as described in 9.3.2.3. This field uniquely identifies the VLAN to which the frame belongs.
f) In 802.3/Ethernet and FDDI MAC methods, an Embedded Source-Routing Information Field (E-RIF) is included, if required by the state of the CFI flag in the TCI. If present, in addition to providing the ability to carry source-routing information, this field includes a further flag, the NCFI, that signals the bit order of address information carried in the encapsulated frame. The structure of this field, as used in this context, is described in 9.3.3.

NOTE 2—The ability of the tag header to carry embedded source-routing information using 802.3/Ethernet and FDDI MAC methods does not imply a requirement on the part of a pure 802.3/Ethernet Bridge or a transparent FDDI Bridge to support source routing. This capability is provided simply to allow traffic that originates in, and is destined for, a source-routed environment to transit as VLAN-tagged traffic across a non-source-routed environment. “Tunnelling” of source-routed frames across transparent media in this manner is still required to follow the rules for source routing as defined in ISO/IEC 15802-3 (see 1.3, 5.4); in particular, Bridges that support only transparent operation are not permitted to forward any frames received that have the RII bit set in the source MAC Address field. Any use of the capability of using the E-RIF to carry real source-routing information across transparent LANs can therefore only be made by Bridges and/or end stations that support source routing. Once RIF information has been encapsulated in this way, transparent Bridges can treat the frames as transparent frames, and forward/filter them accordingly. This standard does not specify any forwarding decisions based on the E-RIF.

The structure of the tagged frame allows the following types of information to be identified and carried in tagged frames across all MAC methods:

g) Ethernet Type-encoded (E) and LLC-encoded (L) information (see 3.1, 3.2);

NOTE 3—The distinction between E and L is not represented in the tag header itself, but is identifiable by examination of the data carried in the tagged frame; in 802.3/Ethernet MAC methods, by examining the value in the Length/Type

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 75 of 211

field, and in Token Ring/FDDI MAC methods, by the presence/absence of the SNAP-based protocol identifiers used in the encapsulation formats described in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390.

h) Frames in which any MAC Addresses embedded in the MAC data are carried in Canonical (C) or Non-canonical (N) format;
i) Source-routed (R) and transparent (T) frames.

NOTE 4—These abbreviations are used here and in Annex C to refer to different types of data frame. For example, a frame carrying LLC-encoded information, Non-canonical embedded addresses and source-routing information, is abbreviated to L-N-R; a frame carrying Ethernet Type-encoded information with Canonical addresses and no source-routing information would be E-C-T.

Relaying a tagged frame requires

j) If the frame formats used on the source and destination MAC methods differ, translation of the tag header to the format appropriate for the destination MAC method;
k) If the source and destination MAC methods differ, relaying the frame may require translation of the remainder of the frame, as defined in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390;
l) Inclusion/adjustment of the PAD field, if necessary, where the destination MAC method is 802.3/Ethernet (7.2);
m) Re-computation of the Frame Check Sequence (FCS) if necessary.

Untagging a tagged frame requires

n) The removal of the tag header, retaining the RIF in the appropriate position in the final untagged frame if necessary;
o) If the frame formats used on the source and destination MAC methods differ, untagging the frame may require translation of the remainder of the frame, as defined in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390;
p) Adjustment of the PAD field, if necessary, where the destination MAC method is 802.3/Ethernet (7.2);
q) Re-computation of the Frame Check Sequence (FCS).

The frame translations defined in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 are applied, as necessary, to all frames carrying Ethernet Type-encoded information relayed by VLAN-aware Bridges. Use of the CFI flag in the tagged frame allows

r) The format of embedded MAC Address information to be signalled end-to-end across a VLAN without the need for MAC Address format translation by VLAN-aware Bridges while the frame is in tagged format, regardless of the MAC methods involved in carrying the tagged frame from source to destination;

NOTE 5—In other words, from the point of view of ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 translation, the representation of Ethernet Type-encoded information is always as would be expected for the underlying MAC method, and in tagged frames, the format of embedded address information is always as indicated by the CFI/NCFI.

s) Source-routing information to be carried end-to-end across a VLAN, regardless of the MAC methods involved in carrying the tagged frame from source to destination.

The primary purpose in allowing the distinction between Canonical and Non-canonical information to be represented in the tagged frame is to permit such information to be carried across a VLAN without the need for VLAN-aware Bridges to translate the format of embedded MAC Addresses en route; however, Bridges that untag frames may still need to take the address format into consideration, and perform the appropriate translation if necessary, and if the Bridge supports such translation. For example, in order to successfully untag an L-N-T frame for transmission onto an 802.3/Ethernet segment, it would be necessary to convert any


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

embedded MAC Addresses to Canonical format in order for that frame to be meaningful to the end stations on that segment.

The ability to support translation of embedded MAC Addresses between Canonical and Non-canonical formats (and vice versa) when transmitting an untagged frame is not required by this standard. In Bridges that do not support such translation capability on a given outbound Port, frames that may require such translation before being forwarded as untagged frames on that Port shall be discarded.

NOTE 6—In particular, this means that a Bridge that supports only one MAC type on all Ports is not required to support MAC Address translation when untagging a frame that originated on a different MAC type.

Tagging of frames occurs when an untagged frame is relayed by a Bridge onto a LAN segment for which that frame is required by the egress rules (8.8) to be transmitted in tagged format.

Untagging of frames occurs when a tagged frame is relayed by a Bridge onto a LAN segment for which that frame is required by the egress rules (8.8) to be transmitted in untagged format.

## 9.2 Transmission and representation of octets

In this clause, octets in a PDU (or a field of a PDU) are numbered starting from 1 and increasing in the order in which they are put into a MAC Service Data Unit (MSDU).

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

Where consecutive octets are used to represent a binary number, the lower octet number carries the most significant value.

Where the value of a field is represented in hexadecimal notation, 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.

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 PDU (or a field within a PDU) is represented using a diagram, the following representations are used:

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 Structure of the tag header

The tag header consists of the following components:

a) The Tag Protocol Identifier (TPID) as described in 9.3.1;
b) The Tag Control Information (TCI) as described in 9.3.2;
c) In 802.3/Ethernet and non-Source-Rrouted FDDI frames (i.e., FDDI frames in which the RII bit is reset), the E-RIF, if required by the state of the CFI.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 77 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

There are three forms of the tag header, depending upon the type of encoding used for the TPID and the underlying MAC type. The overall structure of the three forms of header is illustrated in Figure 9-1.

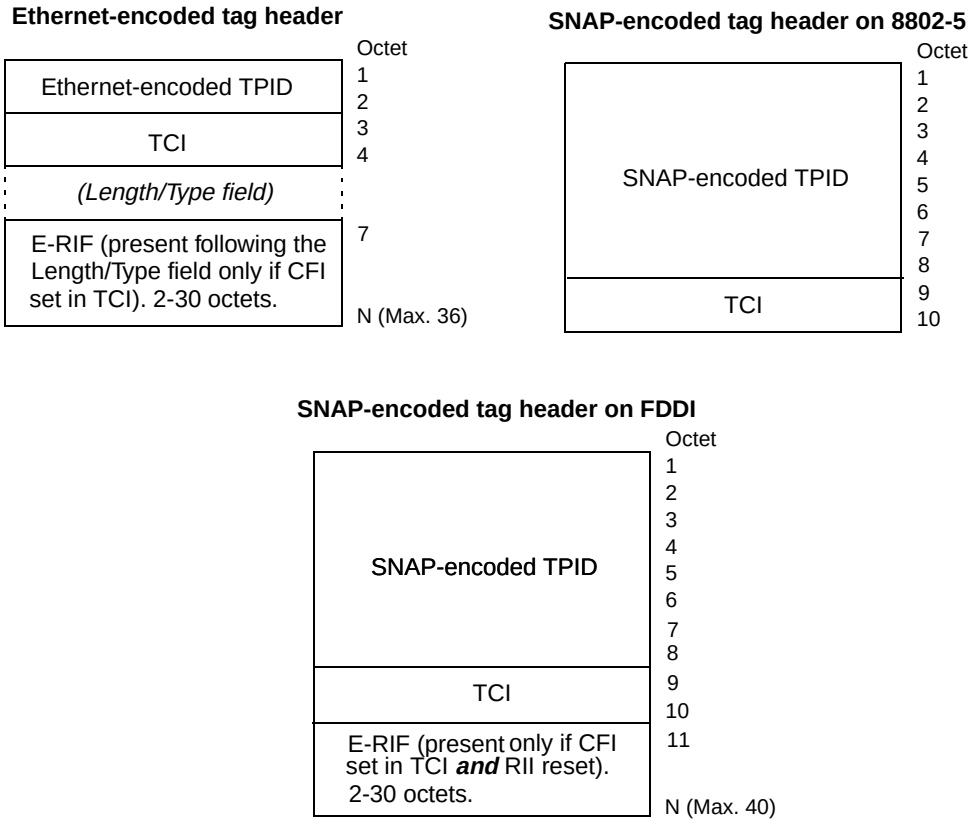


**Figure 9-1—Tag header formats**

The Ethernet-encoded form of the tag header is used where the tagged frame is to be transmitted using 802.3/Ethernet MAC methods.

The two SNAP-encoded forms of the tag header are used where the tagged frame is to be transmitted on Token Ring and FDDI MAC methods:

d) On 8802-5 MACs, any RIF information, if present, appears in the normal position in the frame, i.e., directly following the destination MAC Address field, and is not part on the tag header.

e) On FDDI MACs, RIF information, if present, may be carried either in the normal position in the frame (i.e., directly following the destination MAC Address field), or within the E-RIF field in the tag header.

NOTE—The SNAP-encoded frame format on FDDI allows FDDI LANs to be operated in transparent mode and/or in source routing mode.

### 9.3.1 Tag Protocol Identifier (TPID) format

The structure of the TPID field takes two forms, depending upon whether the field is Ethernet encoded (9.3.1.1) or SNAP encoded (9.3.1.2). The TPID carries an Ethernet Type value (802.1QTagType), which identifies the frame as a tagged frame. The value of 802.1QTagType is defined in Table 9-1.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

**Table 9-1—802.1Q Ethernet Type allocations**

| Name | Value |
|---|---|
| 802.1Q Tag Protocol Type (802.1QTagType) | 81-00 |

#### 9.3.1.1 Ethernet-encoded TPID

The Ethernet-encoded TPID (ETPID) field is two octets in length. The ETPID carries the value of the 802.1QTagType, as defined in Table 9-1.

Figure 9-2 illustrates the structure of the ETPID.

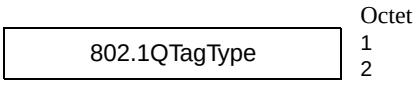


**Figure 9-2—Ethernet-encoded TPID format**

#### 9.3.1.2 SNAP-encoded TPID

The SNAP-encoded TPID (STPID) is eight octets in length, encoded in SNAP format, as follows:

a) Octets numbered 1 through 3 carry the standard SNAP header, consisting of the hexadecimal value AA-AA-03;
b) Octets numbered 4 through 6 carry the SNAP PID, consisting of the hexadecimal value 00-00-00;
c) Octets 7 and 8 carry the 802.1QTagType, as defined in Table 9-1.

Figure 9-3 illustrates the structure of the STPID.

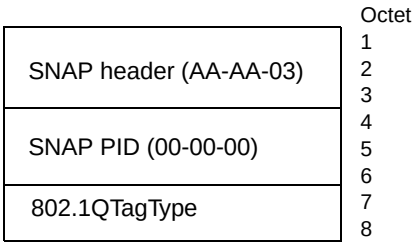


**Figure 9-3—SNAP-encoded TPID format**

### 9.3.2 Tag Control Information (TCI) format

The TCI field is two octets in length, and contains user_priority, CFI and VID (VLAN Identifier) fields. Figure 9-4 illustrates the structure of the TCI field.

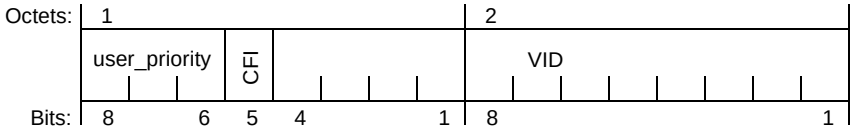


**Figure 9-4—Tag Control Information (TCI) format**

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 79 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 9.3.2.1 user_priority

The user_priority field is three bits in length, interpreted as a binary number. The user_priority is therefore capable of representing eight priority levels, 0 through 7. The use and interpretation of this field is defined in ISO/IEC 15802-3.

#### 9.3.2.2 CFI format

The Canonical Format Indicator (CFI) is a single bit flag value. CFI reset indicates that all MAC Address information that may be present in the MAC data carried by the frame is in Canonical format.

The meaning of the CFI when set depends upon the variant of the tag header in which it appears.

a) In a SNAP-encoded tag header transmitted using 8802-5 MAC methods, CFI has the following meanings:
   1) When set, indicates that all MAC Address information that may be present in the MAC data carried by the frame is in Non-canonical format (N);
   2) When reset, indicates that all MAC Address information that may be present in the MAC data carried by the frame is in Canonical format (C).

b) In an Ethernet-encoded tag header, transmitted using 802.3/Ethernet MAC methods, CFI has the following meanings:
   1) When set, indicates that the E-RIF field is present in the tag header, and that the NCFI bit in the RIF determines whether MAC Address information that may be present in the MAC data carried by the frame is in Canonical (C) or Non-canonical (N) format;
   2) When reset, indicates that the E-RIF field is not present in the tag header, and that all MAC Address information that may be present in the MAC data carried by the frame is in Canonical format (C).

c) In a SNAP-encoded tag header transmitted using FDDI MAC methods, CFI has the following meanings:
   1) When the frame takes the source-routed form (i.e., the RII bit is set in the frame's source MAC Address field and a RIF follows the source MAC Address), the interpretation of the CFI bit is as defined in a) for SNAP-encoded tag headers transmitted using 8802-5 MAC methods. The E-RIF field is not present in this form;
   2) When the frame takes the transparent form (i.e., the RII bit is reset in the frame's source MAC Address field and there is no RIF following the source MAC Address), the interpretation of the CFI bit and the presence or absence of the E-RIF is as defined in b) for Ethernet-encoded tag headers transmitted using 802.3/Ethernet MAC methods.

NOTE 1—The decision as to whether the source-routed form or the transparent form is used on FDDI is a local matter, and depends upon local knowledge in a Bridge or end station as to whether the FDDI LAN is capable of supporting source-routed traffic. The transparent form allows source-routing information to be transparently "tunneled" across LANs that do not support source routing; i.e., LANs where there may be intermediate transparent Bridges in the transmission path that would discard source-routed frames.
NOTE 2—In order to correctly relay frames between differing media types, the MAC Relay function of the Bridge needs to know the MAC type associated with each port. The means by which this information is provided to the MAC Relay function is a local matter.

#### 9.3.2.3 VID format

The twelve-bit VLAN Identifier (VID) field uniquely identify the VLAN to which the frame belongs. The VID is encoded as an unsigned binary number. Table 9-2 identifies values of the VID field that have specific meanings or uses; the remaining values of VID are available for general use as VLAN identifiers.

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


Copyright © 1999 IEEE. All rights reserved.

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

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

**Table 9-2—Reserved VID values**

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

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

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

### 9.3.3 Embedded RIF format

The E-RIF that can appear in Ethernet-encoded tag headers, and in the transparent form of SNAP-encoded tag headers on FDDI, is a modification of the RIF as defined in ISO/IEC 15802-3, C.3.3.2. When present, it immediately follows the Length/Type field in the 802.3/Ethernet tagged frame, or immediately follows the TCI field in an FDDI frame. It consists of two components:

a) A two-octet Route Control Field (RC);
b) Zero or more octets of Route Descriptors (up to a maximum of 28 octets), as defined by RC.

The structure and semantics associated with the Route Descriptors are as defined in ISO/IEC 15802-3, C.3.3.2.

Figure 9-5 illustrates the format of the RC component, as used in the E-RIF. The fields of the RC, and their usage, are defined in the following subclauses.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 81 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

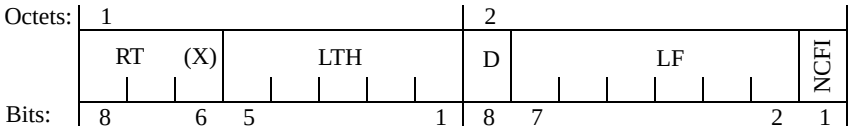


**Figure 9-5—E-RIF Route Control (RC) field**

Note that 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.

NOTE—The use of E-RIF fields in 802.3/Ethernet and FDDI frames is further discussed in 9.3.3.6.

### 9.3.3.1 Definition of the Routing Type (RT) field in the E-RIF

The definition of this field is as defined in ISO/IEC 15802-3, C.3.3.2, with the addition that an RT value of 01X indicates a transparent frame. The value of the rightmost bit, X, is ignored; i.e., it is the binary value 01 in bits 8 and 7 that is used to signal a transparent 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.

The transparent frame value indicates that, with the exception of the NCFI (9.3.3.5), the remainder of the E-RIF shall be discarded if the frame is forwarded using 8802-5 MAC methods. An E-RIF containing an RT value indicating a transparent frame contains no route descriptors, and is therefore exactly 2 octets in length (however, the value of the LTH field is set to zero in this case; see 9.3.3.2).

The least significant bit of the RT field, marked (X) is reserved, as defined in ISO/IEC 15802-3, C.3.3.2.

The following rules 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:

a) Where an untagged, source-routed frame is received from a Token Ring/FDDI LAN and is relayed as a tagged frame either on 802.3/Ethernet or on Token Ring/FDDI, if the received value of the RT field was 0XX, then the value of the RT field in the E-RIF or RIF in the tagged frame shall be transmitted as 000 (i.e., any 01X is converted to 000);
b) Where an untagged, transparent frame is received from a Token Ring LAN and is relayed as a tagged frame either on 802.3/Ethernet or on FDDI, then the tag header will carry an E-RIF in which the value of the RT field shall be 010;
c) Where a VLAN-aware end station on Token Ring/FDDI generates source-routed, tagged frames in the source-routed form (i.e., where the RIF appears in the normal position for a source-routed frame and there is no E-RIF in the tag header), then it shall not transmit RT values of 010 or 011 in the RIF;
d) Where a VLAN-aware end station on 802.3/Ethernet or FDDI generates source-routed, tagged frames in the transparent form (i.e., where there is source-routing information present that is carried in the E-RIF in the tag header, but there is no RIF in the normal position for a source-routed frame) then it shall not use RT values of 010 or 011 in the E-RIF;

NOTE 2—In other words, when source-routed frames are tunnelled across a transparent environment, the state of the RT bits signal that the E-RIF carries real source-routing information. An equivalent frame generated in a source-routed environment would have RII set, and the source-routing information would appear as a RIF in the normal position.

e) “X” in these rules is taken to mean “Ignored upon receipt, transmitted as zero.”


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

NOTE 3—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.
NOTE 4—In addition to its use as a means of tunnelling source-routed frames across transparent LANs, the above rules for the use of RT values in the E-RIF also provide the possibility of stations attached to transparent LANs using source routing to communicate with stations attached to source-routed LANs. In particular, bullet d) states the rule that makes such communication possible.

#### 9.3.3.2 Definition of the Length (LTH) field in the E-RIF

This field is as defined in ISO/IEC 15802-3, C.3.3.2, with the exception that if the value of the RT field in the E-RIF is 01X, indicating a transparent frame, then the LTH field shall carry a value of 0.

NOTE—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.

#### 9.3.3.3 Definition of the Direction Bit (D) field in the E-RIF

This field is as defined in ISO/IEC 15802-3, C.3.3.2.

#### 9.3.3.4 Definition of the Largest Frame (LF) field in the E-RIF

This field is as defined in ISO/IEC 15802-3, C.3.3.2, and is used accordingly for all tagged frames transmitted using 802.3/Ethernet or FDDI MAC methods that carry an E-RIF, whether source-routed or transparent. For frames transmitted using 802.3/Ethernet MAC methods, the value of this field shall indicate a largest frame size of 1470 octets or less.

#### 9.3.3.5 Definition of the NCFI field in the E-RIF

The Non-canonical Format Indicator field of the E-RIF has the following meanings:

a) When reset, indicates that all MAC Address information that may be present in the MAC data carried by the frame is in Non-canonical format (N);
b) When set, indicates that all MAC Address information that may be present in the MAC data carried by the frame is in Canonical format (C).

In source-routed frames on Token Ring/FDDI MAC methods, this bit in the RIF is reserved, and its value is preserved across Bridges; its value is normally reset.

Where a source-routed frame is received from a Token Ring/FDDI LAN and relayed as a tagged frame containing an E-RIF on 802.3/Ethernet or FDDI (i.e., where RIF information carried in the normal position for a source-routed frame is embedded in an E-RIF), the received value of this field is replaced by the appropriate N or C value.

Where a tagged frame containing an E-RIF is relayed from an 802.3/Ethernet or FDDI LAN onto a Token Ring or FDDI LAN as a source-routed frame (i.e., when the RIF information carried in the E-RIF is restored to its normal position in the frame), this bit in the RIF is reset in the frame transmitted onto the destination LAN.

#### 9.3.3.6 E-RIF usage in tagged frames on 802.3/Ethernet and FDDI

There are three cases where a tagged frame will carry an E-RIF:

a) It is a transparent frame that carries E-N or L-N information;
b) It is a source-routed frame that carries E-N or L-N information;

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 83 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

c) It is a source-routed frame that carries L-C information.

Case a) occurs if

d) An untagged frame containing E-N or L-N data, and with no RIF (RII reset), is received from an 8802-5 LAN and is transmitted as a VLAN-tagged frame on 802.3/Ethernet or FDDI; or

e) A tagged frame with CFI set, and with no RIF (RII reset), is received from an 8802-5 LAN and is transmitted as a VLAN-tagged frame on 802.3/Ethernet or FDDI.

Case b) occurs if

f) An untagged frame containing E-N or L-N data, and with a RIF (RII set), is received from a Token Ring/FDDI LAN and is transmitted as a VLAN-tagged frame on 802.3/Ethernet, or on an FDDI LAN that does not support source routing; or

g) A tagged frame with CFI set, and with a RIF (RII set), is received from a Token Ring/FDDI LAN and is transmitted as a VLAN-tagged frame on 802.3/Ethernet, or on an FDDI LAN that does not support source routing.

Case c) occurs if

h) An untagged frame containing L-C data, and with a RIF (RII set), is received from a Token Ring/ FDDI LAN and is transmitted as a VLAN-tagged frame on 802.3/Ethernet, or on an FDDI LAN that does not support source routing; or

i) A tagged frame with CFI reset, and with a RIF (RII set), is received from a Token Ring/FDDI LAN and is transmitted as a VLAN-tagged frame on 802.3/Ethernet, or on an FDDI LAN that does not support source routing.

The other possible frame representations on Token Ring/FDDI, namely, transparent frames carrying either L-C or E-C information, are represented on 802.3/Ethernet media as VLAN-tagged frames with CFI reset (i.e., the RIF is not present).

In case a), the RIF is created from scratch as part of the frame translation. The RIF in this case consists of only 2 octets, with field values as follows:

j) RT is set to the binary value 010, to indicate a transparent frame;

k) LTH is set to 0;

l) D and LF fields are set to 0;

m) NCFI is reset to indicate that the format is Non-canonical.

In cases b) and c), the RIF contained in the frame to be translated is used unmodified, with the two exceptions that

n) Received RT values of 0XX are translated to 000 in the RT field of the E-RIF;

o) The NCFI field is set appropriately to indicate the format (Canonical or Non-canonical) of the data carried in the frame.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# 10. Use of GMRP in VLANs

The GARP Multicast Registration Protocol, GMRP, defined in Clause 10 of ISO/IEC 15802-3, allows the declaration and dissemination of Group membership information, in order to permit GMRP-aware Bridges to filter frames destined for group MAC Addresses on Ports through which potential recipients of such frames cannot be reached. The specification in ISO/IEC 15802-3 calls for the propagation of GMRP registrations only in the GARP Information Propagation (GIP) Context known as the *Base Spanning Tree Context* (ISO/IEC 15802-3, Clauses 10 and 12.3.4); i.e., propagation of GMRP information occurs among the set of Ports of a Bridge that are part of the *active topology* (ISO/IEC 15802-3, 7.4) of the Spanning Tree resulting from operation of the Spanning Tree Algorithm and Protocol defined in ISO/IEC 15802-3, Clause 8. This GIP Context is identified by a GIP Context Identifier of 0.

In Bridged LAN environments that support the definition and management of VLANs in accordance with this standard, the operation of GMRP as specified in ISO/IEC 15802-3 is extended to permit GMRP to operate in multiple GIP contexts, defined by the set of VLANs that are active in the Bridged LAN; these are known as *VLAN Contexts.*

The use of GMRP in a VLAN Context allows GMRP registrations to be made that are specific to that VLAN; i.e., it allows the Group filtering behavior for one VLAN to be independent of the Group filtering behavior for other VLANs. The following subclauses define the extensions to the definition of GMRP that permit its use in VLAN contexts.

With the exception of the extensions defined in this standard, the operation of GMRP and the conformance requirements associated with GMRP are as defined in ISO/IEC 15802-3.

## 10.1 Definition of a VLAN Context

The GIP Context Identifier used to identify a VLAN Context shall be equal to the VID used to identify the corresponding VLAN.

The set of Ports of a Bridge defined to be part of the active topology for a given VLAN Context shall be equal to the set of Ports of a Bridge for which the following are true:

a) The Port is a member of the *Member set* (8.11.9) for that VLAN; and
b) The Port is one of the Ports of the Bridge that are part of the active topology for the spanning tree that supports that VLAN.

NOTE—For the purposes of this standard, a single spanning tree is used to support all VLANs; however, the above definition has been deliberately worded so as not to preclude its use in a context where a mapping exists between M instances of spanning tree and N VLANs.

## 10.2 GMRP Participants and GIP Contexts

For each Port of the Bridge, a distinct instance of the GMRP Participant can exist for each VLAN Context supported by the Bridge. Each GMRP Participant maintains its own set of GARP Applicant and Registrar state machines, and its own Leave All state machine. There is no GMRP Participant associated with the Base Spanning Tree Context.

A given GARP Participant, operating in a given GIP Context, manipulates only the Port Filtering Mode and Group Registration Entry information for the context concerned. In the case of Group Registration Entries, the GIP Context Identifier value corresponds to the value of the VID field of the entry.

Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## 10.3 Context identification in GMRP PDUs

Implementations of GMRP conformant to the specification of GMRP in ISO/IEC 15802-3 exchange PDUs in the Base Spanning Tree Context; such PDUs are transmitted and received by GMRP Participants as untagged frames.

Implementations of GMRP in VLAN Bridges apply the same ingress rules (8.6) to received GMRP PDUs that are defined for the reception Port. Therefore

a) GMRP frames with no VLAN classification (i.e., untagged or priority-tagged GMRP frames) are discarded if the Acceptable Frame Types parameter (8.4.3) for the Port is set to *Admit Only VLAN-tagged frames.* Otherwise, they are classified according to the PVID for the Port;
b) VLAN-tagged GMRP frames are classified according to the VID carried in the tag header;
c) If Ingress Filtering (8.4.5) is enabled, and if the Port is not in the Member set (8.11.9) for the GMRP frame's VLAN classification, then the frame is discarded.

The VLAN classification thus associated with received GMRP PDUs establishes the VLAN Context for the received PDU, and identifies the GARP Participant instance to which the PDU is directed.

GMRP PDUs transmitted by GMRP Participants are VLAN classified according to the VLAN Context associated with that Participant. GMRP Participants in VLAN Bridges apply the same egress rules that are defined for the transmission Port (8.8). Therefore

d) GMRP PDUs are transmitted through a given Port only if the value of the Member Set for the Port for the VLAN concerned indicates that the VLAN is registered on that Port;
e) GMRP PDUs are transmitted as VLAN-tagged frames or as untagged frames in accordance with the state of the Untagged Set (8.11.9) for that Port for the VLAN concerned. Where VLAN-tagged frames are transmitted, the VID field of the tag header carries the VLAN Context Identifier value.

## 10.4 Default Group filtering behavior and GMRP propagation

The propagation of GMRP registrations within a VLAN context has implications with respect to the choice of default Group filtering behavior within a Bridged LAN. As GMRP frames are transmitted only on outbound Ports that are in the Member set (8.11.9) for the VLAN concerned, propagation of Group registrations by a given Bridge occurs only towards regions of the Bridged LAN where that VLAN has been (statically or dynamically) registered. This is illustrated in Figure 10-1; dotted lines in the diagram show those regions of the LAN where propagation of registrations for Group M in VLAN V does not occur. Consequently, the Filtering Databases of the lower two Bridges will not contain any Dynamic Group Registration Entry for Group M in VLAN V.

The action of these two Bridges on receipt of frames, on either of their lower Ports, destined for Group M and VLAN V, will depend upon the Default Group Filtering Behavior adopted by their upper Ports, which are the Ports that are in the Member set for VLAN V. If the Default Group Filtering Behavior is either Forward All Groups or Forward Unregistered Groups, then these Bridges will forward the frames. If the Default Group Filtering Behavior is Filter Unregistered Groups, then these Bridges will filter the frames. In the scenario shown, the choice of Default Group Filtering Behavior is therefore crucial with respect to whether or not end station S, or any other station that is "outside" the VLAN, is able to send frames to members of the Group. The choice between Filter Unregistered Groups and the other default behaviors therefore has the effect of defining VLANs that are closed to external unregistered traffic (Filter Unregistered Groups) or open to external unregistered traffic (Either of the other default behaviors).


Copyright © 1999 IEEE. All rights reserved.

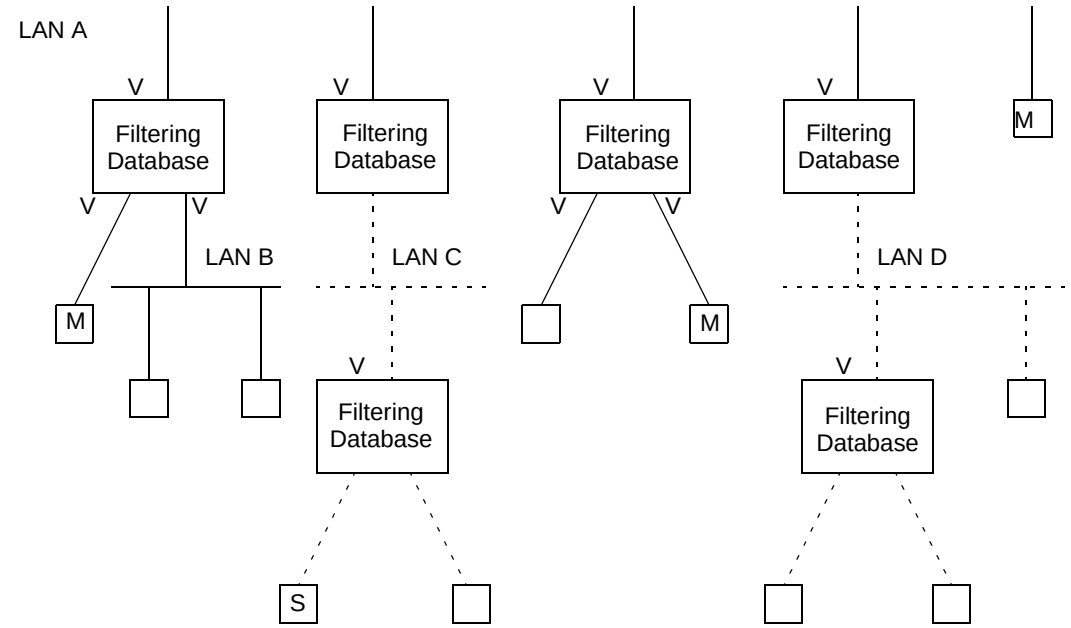


**Figure 10-1—Example of GMRP propagation in a VLAN context**

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 87 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# 11. VLAN topology management

The egress rules (8.8) defined for the Forwarding Process in VLAN Bridges rely on the existence of configuration information for each VLAN that defines the set of Ports of the Bridge through which one or more members are reachable. This set of Ports is known as the Member Set (8.11.9), and its membership is determined by the presence or absence of configuration information in the Filtering Database, in the form of Static and Dynamic VLAN Registration Entries (8.11.2, 8.11.5).

Reliable operation of the VLAN infrastructure requires VLAN membership information held in the Filtering Database to be maintained in a consistent manner across all VLAN-aware Bridges in the Bridged LAN, in order to ensure that frames destined for end station(s) on a given VLAN can be correctly delivered, regardless of where in the Bridged LAN the frame is generated. Maintenance of this information by end stations that are sources of VLAN-tagged frames can allow such stations to suppress transmission of such frames if no members exist for the VLAN concerned.

This standard defines the following mechanisms that allow VLAN membership information to be configured:

a) Dynamic configuration and distribution of VLAN membership information by means of the GARP VLAN Registration Protocol (GVRP), as described in 11.2;
b) Static configuration of VLAN membership information via Management mechanisms, as described in Clause 12, which allow configuration of Static VLAN Registration Entries.

These mechanisms provide for the configuration of VLAN membership information as a result of

c) Dynamic registration actions taken by end stations or Bridges in the bridged LAN;
d) Administrative actions.

## 11.1 Static and dynamic VLAN configuration

The combined functionality provided by the ability to configure Static VLAN Registration Entries in the Filtering Database, coupled with the ability of GVRP to dynamically create and update Dynamic VLAN Registration Entries, offers the following possibilities with respect to how VLANs are configured on a given Port:

a) *Static configuration only.* The management facilities described in Clause 12 are used to establish precisely which VLANs have this Port in their Member set, and the GVRP management controls are used to disable the operation of the GVRP protocol on that Port. Hence, any use of GVRP by devices reachable via that Port is ignored, and the Member set for all VLANs can therefore only be determined by means of static entries in the Filtering Database.
b) *Dynamic configuration only.* The operation of GVRP is relied upon to establish Dynamic VLAN Registration Entries that will dynamically reflect which VLANs are registered on the Port, their contents changing as the configuration of the network changes. The GVRP management controls are set to enable the operation of the GVRP protocol on that Port.
c) *Combined static and dynamic configuration*. The static configuration mechanisms are used in order to configure some VLAN membership information; for other VLANs, GVRP is relied upon to establish the configuration. The GVRP management controls are set to enable the operation of the GVRP protocol on that Port.

All of the above approaches are supported by the mechanisms defined in this standard, and each approach is applicable in different circumstances. For example:

d) Use of static configuration may be appropriate on Ports where the configuration of the attached devices is fixed, or where the network administrator wishes to establish an administrative boundary


Copyright © 1999 IEEE. All rights reserved.

outside of which any GVRP registration information is to be ignored. For example, it might be desirable for all Ports serving end user devices to be statically configured in order to ensure that particular end users have access only to particular VLANs.

e) Use of dynamic configuration may be appropriate on Ports where the VLAN configuration is inherently dynamic; where users of particular VLANs can connect to the network via different Ports on an ad hoc basis, or where it is desirable to allow dynamic reconfiguration in the face of Spanning Tree topology changes. In particular, if the “core” of the Virtual Bridged LAN contains redundant paths that are pruned by the operation of Spanning Tree, then it is desirable for Bridge Ports that form the core network to be dynamically configured.

f) Use of both static and dynamic configuration may be appropriate on Ports where it is desirable to place restrictions on the configuration of some VLANs, while maintaining the flexibility of dynamic registration for others. For example, on Ports serving mobile end user devices, this would maintain the benefits of dynamic VLAN registration from the point of view of traffic reduction, while still allowing administrative control over access to some VLANs via that Port.

## 11.2 GARP VLAN Registration Protocol

The GARP VLAN Registration Protocol (GVRP) defines a *GARP Application* that provides the VLAN registration service defined in 11.2.2. GVRP makes use of GARP Information Declaration (GID) and GARP Information Propagation (GIP), which provide the common state machine descriptions and the common information propagation mechanisms defined for use in GARP-based applications. The GARP architecture, GID, and GIP are defined in ISO/IEC 15802-3, Clause 12.

GVRP provides a mechanism for dynamic maintenance of the contents of Dynamic VLAN Registration Entries for each VLAN, and for propagating the information they contain to other Bridges. This information allows GVRP-aware devices to dynamically establish and update their knowledge of the set of VLANs that currently have active members, and through which Ports those members can be reached.

### 11.2.1 GVRP overview

The operation of GVRP is closely similar to the operation of GMRP (ISO/IEC 15802-3, Clause 10), which is used for registering Group membership information. The primary differences are as follows:

a) The attribute values carried by the protocol are 12-bit VID values, rather than 48-bit MAC Addresses and Group service requirement information;

b) The act of registering/deregistering a VID affects the contents of Dynamic VLAN Registration Entries (8.11.5), rather than the contents of Group Registration Entries (8.11.4).

GVRP allows both end stations and Bridges in a Bridged LAN to issue and revoke declarations relating to membership of VLANs. The effect of issuing such a declaration is that each GVRP Participant that receives the declaration will create or update a Dynamic VLAN Registration Entry in the Filtering Database to indicate that VLAN is registered on the reception Port. Subsequently, if all Participants on a segment that had an interest in a given VID revoke their declarations, the Port attached to that segment is set to Unregistered in the Dynamic VLAN Registration Entry for that VLAN by each GVRP Participant attached to that segment.

Figure 11-1 illustrates the architecture of GVRP in the case of a two-Port Bridge and an end station.

As shown in the diagram, the GVRP Participant consists of the following components:

c) The GVRP Application, described in 11.2.3;

d) GARP Information Propagation (GIP), described in ISO/IEC 15802-3, Clause 12;

e) GARP Information Declaration, described in ISO/IEC 15802-3, Clause 12.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 89 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

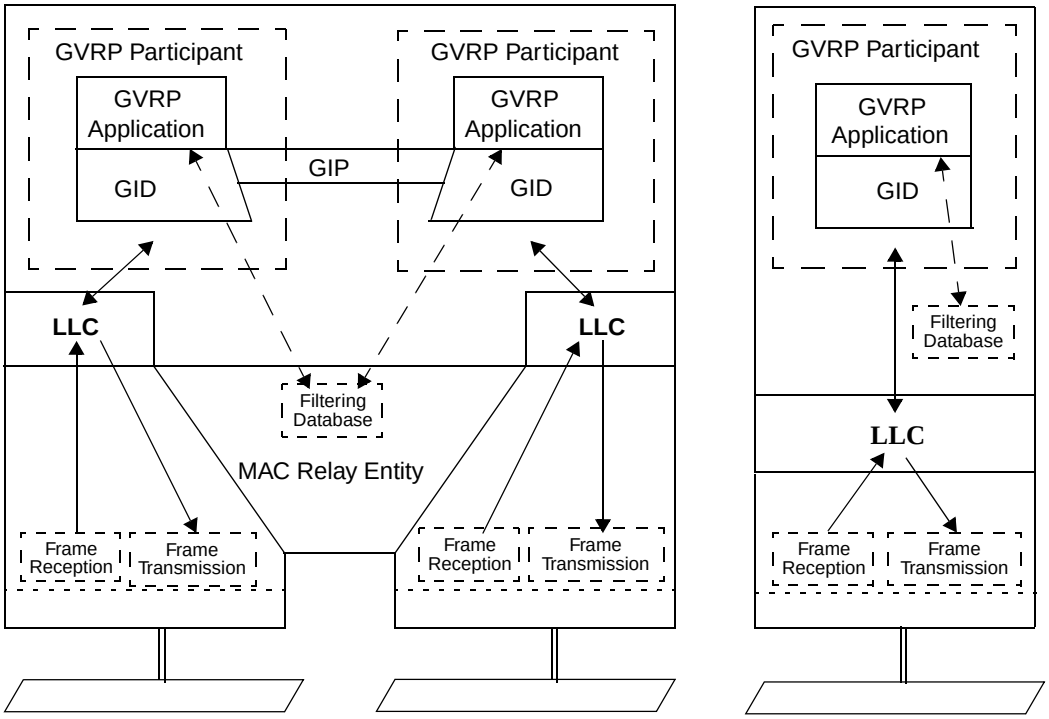


**Figure 11-1—Operation of GVRP**

### 11.2.1.1 Behavior of end stations

VLAN-aware end stations participate in GVRP protocol activity, as appropriate for the set of VLANs of which they are currently members. GVRP provides a way for such an end station to ensure that the VLAN(s) of which it is a member are registered for each Port on any LAN segment to which the end station is attached. GVRP also provides for that VID information to be propagated across the Spanning Tree to other VLAN-aware devices, as described in 11.2.1.2.

Incoming VLAN membership information (from all other devices on the same LAN segment) allows such end stations to “source prune” (i.e., discard at source; see ISO/IEC 15802-3, 10.2.2) any traffic destined for VLANs that currently have no other members in the Bridged LAN, thus avoiding the generation of unnecessary traffic on their local LAN segments. This is illustrated in Figure 11-1 by a Filtering Database shown as being present in the end station.

NOTE—Non-VLAN-aware end stations have no need to register VLAN membership via GVRP; indeed, this would be impossible for them to achieve if truly VLAN-unaware, as they would have no knowledge of the set of VLANs in which they participate. Their VLAN registration requirements are taken care of by means of the configuration of PVIDs (and possibly other VLAN classification mechanisms) and the propagation of registered VLAN IDs by the Bridges.

### 11.2.1.2 Behavior of Bridges

VLAN-aware Bridges register and propagate VLAN memberships on all Bridge Ports that are part of the active topology of the underlying Spanning Tree. Incoming VID registration and de-registration information is used to update the Dynamic VLAN Registration Entries associated with each VLAN. Any changes in the state of registration of a given VID on a given Port are propagated on Ports that are part of the active topol-


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

ogy of the Spanning Tree, in order to ensure that other GVRP-aware devices in the Bridged LAN update their Filtering Databases appropriately.

The Filtering Databases in all GVRP-aware devices are thus automatically configured such that the Port Map in the Dynamic VLAN Registration Entry for a given VID indicates that a given Port is registered if one or more members of the corresponding VLAN are reachable through the Port.

NOTE—The information that determines whether frames destined for each VLAN are transmitted VLAN-tagged or untagged is carried in Static VLAN Registration Entries (8.11.2); if no such entry exists for a VLAN, then it is assumed that frames for that VLAN are transmitted VLAN-tagged on all Ports. Therefore, if the configuration information held in the Filtering Database for a given VLAN consists only of information configured by the operation of GVRP (i.e., only a Dynamic VLAN Registration Entry), then all traffic for that VLAN will be VLAN-tagged on transmission.

#### 11.2.1.3 Use of the PVID

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

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

### 11.2.2 VLAN registration service definition

The VLAN registration service allows MAC Service users to indicate to the MAC Service provider the set of VLANs in which they wish to participate; i.e., that the MAC Service user wishes to receive traffic destined for members of that set of VLANs. The service primitives allow the service user to:

a) Register membership of a VLAN;
b) De-register membership of a VLAN.

Provision of these services is achieved by means of GVRP and its associated procedures, as described in 11.2.3.

ES_REGISTER_VLAN_MEMBER (VID)

Indicates to the MAC Service provider that the MAC Service user wishes to receive frames destined for the VLAN identified by the VID parameter.

ES_DEREGISTER_VLAN_MEMBER (VID)

Indicates to the MAC Service provider that the MAC Service user no longer wishes to receive frames destined for the VLAN identified by the VID parameter.

The use of these services can result in the propagation of VID information across the Spanning Tree, affecting the contents of Dynamic VLAN Registration Entries (8.11.5) in Bridges and end stations in the Bridged LAN, and thereby affecting the frame forwarding behavior of those Bridges and end stations.

Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### 11.2.3 Definition of the GVRP Application

#### 11.2.3.1 Definition of GARP protocol elements

##### 11.2.3.1.1 GVRP Application address

The group MAC Address used as the destination address for GARP PDUs destined for GVRP Participants shall be the GVRP address identified in Table 11-1. Received PDUs that are constructed in accordance with the PDU format defined in ISO/IEC 15802-3, 12.11, and which carry a destination MAC Address equal to the GVRP address are processed as follows:

a) In Bridges and end stations that support the operation of GVRP, all such PDUs shall be submitted to the GVRP Participant associated with the receiving Port for further processing;
b) In Bridges that do not support the operation of GVRP, all such PDUs shall be submitted to the Forwarding Process.

**Table 11-1—GVRP Application address**

| Assignment | Value |
|---|---|
| GVRP address | 01-80-C2-00-00-21 |

NOTE—The GVRP Application Address has been allocated from the set of GARP Application addresses defined in ISO/IEC 15802-3, Table 12-1, using the MAC Address contained in the second entry of that table.

##### 11.2.3.1.2 Encoding of GVRP Attribute Types

The operation of GVRP defines a single Attribute Type (ISO/IEC 15802-3, 12.11.2.2) that are carried in GARP protocol exchanges; the VID Attribute Type. The VID Attribute Type is used to identify values of VLAN Identifiers (VIDs). The value of the Group Attribute Type carried in GVRP PDUs shall be 1.

##### 11.2.3.1.3 Encoding of GVRP Attribute Values

Values of instances of the VID Attribute Type shall be encoded as Attribute Values in GARP PDUs (ISO/IEC 15802-3, 12.11.2.6) as two octets, taken to represent an unsigned binary number, and equal to the hexadecimal value of the VLAN identifier that is to be encoded.

#### 11.2.3.2 Provision and support of the VLAN registration service

##### 11.2.3.2.1 End system VLAN membership declaration

The GVRP Application element of a GVRP Participant provides the dynamic registration and de-registration services defined in 11.2.2, as follows:

On receipt of an ES_REGISTER_VLAN_MEMBER service primitive, the GVRP Participant issues a GID_Join.request service primitive (ISO/IEC 15802-3, 12.3.2.1). The attribute_type parameter of the request carries the value of the VID Attribute Type (11.2.3.1.2) and the attribute_value parameter carries the value of the VID parameter carried in the ES_REGISTER_VLAN_MEMBER primitive.

On receipt of an ES_DEREGISTER_VLAN_MEMBER service primitive, the GVRP Participant issues a GID_Leave.request service primitive (ISO/IEC 15802-3, 12.3.2.1). The attribute_type parameter of the


Copyright © 1999 IEEE. All rights reserved.

request carries the value of the VID Attribute Type (11.2.3.1.2) and the attribute_value parameter carries the value of the VID parameter carried in the ES_REGISTER_VLAN_MEMBER primitive.

#### 11.2.3.2.2 VLAN membership registration

The GVRP Application element of a GVRP Participant responds to registration and de-registration events signalled by GID as follows:

On receipt of a GID_Join.indication (ISO/IEC 15802-3, 12.3.2.2) whose attribute_type is equal to the value of the VID Attribute Type (11.2.3.1.2), the GVRP Application element indicates the reception Port as Registered in the Port Map of the Dynamic VLAN Registration Entry for the VID indicated by the attribute_value parameter. If no such entry exists, and there is sufficient room in the Filtering Database, an entry is created.

On receipt of a GID_Leave.indication (ISO/IEC 15802-3, 12.3.2.2) whose attribute_type is equal to the value of the VID Attribute Type (11.2.3.1.2), the GVRP Application element indicates the reception Port as Unregistered in the Port Map of the Dynamic VLAN Registration Entry for the VID indicated by the attribute_value parameter. If marking this Port as Unregistered results in a Port Map that does not indicate any Port as Registered, the entry is deleted.

#### 11.2.3.2.3 Administrative controls

The provision of static control over the declaration or registration state of the state machines associated with the GVRP Application is achieved by means of the Registrar administrative control parameters provided by GARP (ISO/IEC 15802-3, 12.9.1). These administrative control parameters are represented as Static VLAN Registration Entries in the Filtering Database (8.11.2). Where management capability is implemented, these controls can be manipulated by means of the management functionality defined in 12.7.

The provision of static control over the ability of Applicant state machines to participate in protocol exchanges is achieved by means of the Applicant Administrative Control parameters associated with the operation of GARP (ISO/IEC 15802-3, 12.9.2). Where management capability is implemented, the Applicant Administrative Control parameters can be applied and modified by means of the management functionality defined in 12.9.

### 11.2.3.3 GIP context for GVRP

GVRP as defined by this standard operates in the Base Spanning Tree Context (ISO/IEC 15802-3, 12.3.1); i.e., GVRP operates only on the base Spanning Tree defined by ISO/IEC 15802-3. Consequently, all GVRP PDUs sent and received by GVRP Participants are transmitted as untagged frames.

## 11.3 Conformance to GVRP

This subclause defines the conformance requirements for implementations claiming conformance to GVRP. Two cases are covered; implementation of GVRP in MAC Bridges and implementation of GVRP in end stations. Although this standard is principally concerned with defining the requirements for MAC Bridges, the conformance requirements for end station implementations of GVRP are included in order to give useful guidance to such implementations. The PICS proforma defined in Annex A is concerned only with conformance claims with respect to MAC Bridges.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 93 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### 11.3.1 Conformance to GVRP in MAC Bridges

A MAC Bridge for which conformance to GVRP is claimed shall

a) Conform to the operation of the GARP Applicant and Registrar state machines, and the LeaveAll generation mechanism, as defined in ISO/IEC 15802-3, 12.8.1, 12.8.2, and 12.8.3;
b) Exchange GARP PDUs as required by those state machines, formatted in accordance with the generic PDU format described in ISO/IEC 15802-3, 12.11, and able to carry application-specific information as defined in 11.2.3, using the GVRP Application address as defined in Table 11-1;
c) Propagate registration information in accordance with the operation of GIP for the Base Spanning Tree Context, as defined in ISO/IEC 15802-3, 12.3.3 and 12.3.4;
d) Implement the GVRP Application component as defined in 11.2;
e) Forward, filter or discard MAC frames carrying any GARP Application address as the destination MAC Address in accordance with the requirements of 8.14.3.

### 11.3.2 Conformance to GVRP in end stations

An end station for which conformance to GVRP is claimed shall

a) Conform to the operation of one of
   1) The Applicant state machine, as defined in ISO/IEC 15802-3, 12.8.1; or
   2) The Applicant Only state machine, as defined in ISO/IEC 15802-3, 12.8.5; or
   3) The Simple Applicant state machine, as defined in ISO/IEC 15802-3, 12.8.6;
b) Exchange GARP PDUs as required by the GARP state machine(s) implemented, formatted in accordance with the generic PDU format described in ISO/IEC 15802-3, 12.11, and able to carry application-specific information as defined in 11.2.3, using the GVRP Application address as defined in Table 11-1;
c) Support the provision of end system registration and de-registration as defined in 11.2;
d) Discard MAC frames carrying any GARP Application address as the destination MAC Address in accordance with the requirements of 8.14.3.

An end station for which conformance to GVRP is claimed may optionally

e) Conform to the operation of the GARP Registrar state machine and the LeaveAll generation mechanism, as defined in ISO/IEC 15802-3, 12.8.2 and 12.8.3; and
f) Support the provision of VLAN registration and de-registration as defined in 11.2; and
g) Filter outgoing frames destined for group MAC Addresses in accordance with registered VLAN membership information, in a manner consistent with the operation of the filtering function of the forwarding process described in 8.7.2 and the operation of the egress rules defined in 8.8.

It is recommended that only those end stations that require the ability to perform Source Pruning (11.2.1.1) conform to the operation of the Applicant state machine (ISO/IEC 15802-3, 12.8.1).

For the reasons stated in ISO/IEC 15802-3, 12.7.9, it is recommended that end stations that do not require the ability to perform Source Pruning implement the Applicant Only state machine (ISO/IEC 15802-3, 12.8.5), in preference to the Simple Applicant state machine (ISO/IEC 15802-3, 12.8.6).

NOTE—End stations that implement only a) 2) and b) through d) are equivalent to the description of the Applicant Only Participant (ISO/IEC 15802-3, 12.7.7); those that implement a) 3) and b) through d) are equivalent to the description of the Simple Applicant Participant (ISO/IEC 15802-3, 12.7.8). Such end stations require only the ability to register membership of one or more VLANs, and revoke that membership at some later point in time; for this reason, there is no requirement to support the operation of the Registrar or Leave All state machines.
End stations that implement a) 1) and b) through g) are able to perform “source pruning” as described in 11.2.1.1; i.e., to suppress the transmission of frames destined for VLANs that currently have no membership. Consequently, such end stations need to support the full Applicant state machine, in combination with the Registrar and Leave All state machines.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## 11.4 Procedural model

### 11.4.1 Purpose

This section contains an example implementation of the GVRP application defined in this Clause. This “C” code description is included in order to demonstrate the structure of the GVRP application, and to show that a reasonably low overhead implementation can be constructed. The implementation has been designed with the intent of maximizing clarity and generality, not for compactness.

The example implementation is shown in two sections:

a) Header files for the GVRP application (11.4.2);
b) The GVRP application code (11.4.3).

The example implementation also references the GARP application independent “C” code contained in ISO/IEC 15802-3:

c) Header files for the GARP application independent code (ISO/IEC 15802-3, 13.2);
d) The GARP application independent code (ISO/IEC 15802-3, 13.3).

The separation shown in the documentation of the application dependent (GVRP) and application independent (GARP) aspects of the implementation gives a clear illustration of what is involved in implementing additional applications using the same basic GARP state machines. The code is intended to be largely self-documenting, by means of in-line comments.

### 11.4.2 GVRP application-specific header files

#### 11.4.2.1 gvr.h

```
/* gvr.h */
#ifndef gvr_h__
#define gvr_h__

#include “garp.h”
#include “gid.h”
#include “gip.h”

/*****************************************************************************
 * GVR : GARP VLAN REGISTRATION APPLICATION : GARP ATTRIBUTES
 *****************************************************************************
 */

typedef unsigned Vlan_id;

typedef enum {All_attributes, Vlan_attribute}
             Attribute_type;

/*****************************************************************************
 * GVR : GARP VLAN REGISTRATION APPLICATION : CREATION, DESTRUCTION
 *****************************************************************************
 */

extern Boolean gvr_create_gvr(int process_id, void **gvr);
   /*
    * Creates a new instance of GVR, allocating and initializing a control
    * block, returning True and a pointer to this instance if creation succeeds.
    * Also creates instances of GVD (the GARP VLAN database) and of GIP
    * (which controls information propagation).
    *
```

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 95 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

```
    * Ports are created by the system and added to GVR separately (see
    * gvr_added_port() and gvr_removed_port() below).
    *
    * The operating system supplied process_id is for use in subsequent calls
    * to operating system services. The system itself ensures the temporal
    * scope of process_id, guarding against timers yet to expire for destroyed
    * processes, etc. Although process_id will be implicitly supplied by many
    * if not most systems, it is made explicit in this implementation for
    * clarity.
    *
    */

extern void gvr_destroy_gvr(void *gvr);
   /*
    * Destroys an instance of GVR, destroying and deallocating the associated
    * instances of GVD and GIP, and any instances of GID remaining.
    */

extern void gvr_added_port(void *my_gvr, int port_no);
   /*
    * The system has created a new port for this application and added it to
    * the ring of GID ports. This function ensures that Static VLAN Entries
    * from the Permanent Database are represented in the GVD database (which
    * provides VLAN ID to GID index mapping) and have GID machines in the newly
    * added port (with the correct management control state). This can result
    * in the creation of new GID machines or modification of the state of
    * existing machines.
    *
    * Newly created ports are “connected” for the purpose of GARP information
    * propagation using the separate function gip_connect_port(). This should
    * be called after this function, gvr_added_port. It may cause GVRP/GIP
    * to propagate information from the static management controls through
    * other ports.
    *
    * It is assumed that new ports will be “connected” correctly before the
    * application continues as determined by the active topology of the network,
    * i.e., if stp_forwarding(port_no) gvr_connect_port(port_no);.
    *
    * As the system continues to run it should invoke gip_disconnect_port()
    * and gip_connect_port() as required to maintain the required connectivity.
    */

extern void gvr_removed_port(void *my_gvr, int port_no);
   /*
    * The system has removed and destroyed the GID port. This function should
    * provide any application-specific cleanup required.
    */

/*****************************************************************************
 * GVR : GARP VLAN REGISTRATION APPLICATION : JOIN, LEAVE INDICATIONS
 *****************************************************************************
 */

extern void gvr_join_indication(void *my_gvr, void *my_port,
                                unsigned joining_gid_index);
   /*
    *
    */

extern void gvr_join_leave_propagated(void *my_gvr, void *my_port,
                                      unsigned gid_index);
   /*
    *
    */
```


Copyright © 1999 IEEE. All rights reserved.

```
extern void gvr_leave_indication(void *my_gvr, void *my_port,
                                 unsigned leaving_gid_index);
   /*
    *
    */


/*****************************************************************************
 * GVR : GARP VLAN REGISTRATION APPLICATION : PROTOCOL AND MANAGEMENT EVENTS
 *****************************************************************************
 */

extern void gvr_rcv(void *my_gvr, void *my_port, void *pdu);
   /*
    * Process an entire received pdu for this instance of GVR.
    */

extern void gvr_tx(void *my_gvr, void *my_port);
   /*
    * Transmit a pdu for this instance of GVR.
    */

#endif /* gvr_h__ */
```

### 11.4.2.2 gvd.h

```
/* gvd.h */
#ifndef gvd_h__
#define gvd_h__

#include “sys.h”

/*****************************************************************************
 * GVD : GARP VLAN DATABASE
 *****************************************************************************
 *
 * The GARP VLAN Database maps VLAN IDs into compact GID indexes and
 * vice versa. It contains VLAN ID to index mappings for all the VLAN IDs
 * dynamically registered (except when there is a database overflow, which
 * should be an event that rarely occurs through appropriate
 * sizing) and for all those for which static controls exist.
 *
 * Taken together with the GID machines for each port (which are identified
 * by the GID indexes provided by GVD), GVD logically provides the Static VLAN
 * Registration Entries and the Dynamic VLAN Registration Entries of the abstract
 * Filtering Database (see 8.11).
 *
 * Static VLAN Registration Entries are included in this database (and have GID
 * machines defined) on an as-needed basis as ports are added to the GVR
 * Application. This example implementation assumes that the necessary information
 * is taken from Static Filtering Entries kept in a Permanent Database outside
 * the example implementation. Static VLAN Entries can also be added, changed, or
 * removed as the running system is managed.
 *
 * VLAN Registration Entries are added and removed by GVRP. Note that a
 * single VLAN ID will only give rise to one entry in this database, and one
 * GID machine per port. That machine provides the functionality for both
 * the Static VLAN Entry and the VLAN Registration Entry.
 */

extern Boolean gvd_create_gvd(int max_vlans, void **gvd);
   /*
    * Creates a new instance of gvd, allocating space for up to max_vlans
```

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 97 of 211
Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

```
    * VLAN IDs.
    *
    * Returns True if the creation succeeded together with a pointer to the
    * gvd information.
    */

extern void gvd_destroy_gvd(void *gvd);
   /*
    * Destroys the instance of gvd, releasing previously allocated database and
    * control space.
    */

extern Boolean gvd_find_entry(  void *my_gvd, Vlan_id key,
                              unsigned *found_at_index);

extern Boolean gvd_create_entry(void *my_gvd, Vlan_id key,
                              unsigned *created_at_index);

extern Boolean gvd_delete_entry(void *my_gvd,
                              unsigned  delete_at_index);

extern Boolean gvd_get_key(     void *my_gvd, unsigned index, Vlan_id *key);

#endif /* gvd_h__ */
```

### 11.4.2.3 gvf.h

```
/* gvf.h */
#ifndef gvf_h__
#define gvf_h__

#include “sys.h”
#include “prw.h”
#include “gvr.h”

/*****************************************************************************
 * GVF : GARP VLAN REGISTRATION APPLICATION PDU FORMATTING
 *****************************************************************************
 */

typedef struct
{  /*
    * This data structure saves the temporary state required to parse GVR
    * PDUs in particular. Gpdu provides a common basis for GARP Application
    * formatters, additional state can be added here as required by GVF.
    */

    Gpdu gpdu;

} Gvf;

typedef struct /* Gvf_msg_data */
{
    Attribute_type attribute;

    Gid_event   event;

    Vlan_id        key1;

} Gvf_msg;

extern void    gvf_rdmsg_init(Pdu *pdu, Gvf **gvf);
```


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

```
extern void    gvf_wrmsg_init(Gvf *gvf, Pdu  *pdu, int vlan_id);

extern Boolean gvf_rdmsg(     Gvf *gvf, Gvf_msg *msg);

extern Boolean gvf_wrmsg(     Gvf *gvf, Gvf_msg *msg);

#endif /* gvf_h__ */
```

#### 11.4.2.4 vfdb.h

```
/* vfdb.h */
#ifndef vfdb_h__
#define vfdb_h__

#include “sys.h”

/*****************************************************************************
 * VFDB : VLAN ACTIVE FILTERING DATABASE INTERFACE
 *****************************************************************************
 */

extern void vfdb_filter( int port_no, unsigned vlan_id);

extern void vfdb_forward(int port_no, unsigned vlan_id);

#endif /* vfdb_h__ */
```

### 11.4.3 GVRP application code

#### 11.4.3.1 gvr.c

```
/* gvr.c */

#include “gvr.h”
#include “gid.h”
#include “gip.h”
#include “garp.h”
#include “gvd.h”
#include “gvf.h”
#include “vfdb.h”


/*****************************************************************************
 * GVR : GARP VLAN REGISTRATION APPLICATION : IMPLEMENTATION SIZING
 *****************************************************************************
 */

enum {Max_vlans = 100};

enum {Number_of_gid_machines = Max_vlans};

enum {Unused_index = Number_of_gid_machines};

/*****************************************************************************
 * GVR : GARP VLAN REGISTRATION APPLICATION : CREATION, DESTRUCTION
 *****************************************************************************
 */
```

Copyright © 1999 IEEE. All rights reserved.



Dell et al., Ex. 1008
Page 99 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

```
typedef struct /* gvr */
{
    Garp      g;

    unsigned  vlan_id;

    void     *gvd; /* VLAN Registration Entry Database */

    unsigned  number_of_gvd_entries;

    unsigned  last_gvd_used_plus1;

} Gvr;


Boolean gvr_create_gvr(int process_id, void **gvr)
{  /*
    */
    Gvr *my_gvr;

    if (!sysmalloc(sizeof(Gvr), &my_gvr))
        goto gvr_creation_failure;

    my_gvr->g.process_id   = process_id;
    my_gvr->g.gid          = NULL;

    if (!gip_create_gip(Number_of_gid_machines, &my_gvr->g.gip))
        goto gip_creation_failure;

    my_gvr->g.max_gid_index = Number_of_gid_machines - 1;
    my_gvr->g.last_gid_used = Zero;

    my_gvr->g.join_indication_fn  = gvr_join_indication;
    my_gvr->g.leave_indication_fn = gvr_leave_indication;
    my_gvr->g.join_propagated_fn  = gvr_join_leave_propagated;
    my_gvr->g.leave_propagated_fn = gvr_join_leave_propagated;
    my_gvr->g.transmit_fn         = gvr_tx;
    my_gvr->g.added_port_fn       = gvr_added_port;
    my_gvr->g.removed_port_fn     = gvr_removed_port;

    if (!gvd_create_gvd(Max_vlans, &my_gvr->gvd))
        goto gvd_creation_failure;

    my_gvr->number_of_gvd_entries = Max_vlans;
    my_gvr->last_gvd_used_plus1   = 0;

    *gvr = my_gvr;    return(True);
gvd_creation_failure: gip_destroy_gip(my_gvr->g.gip);
gip_creation_failure: sysfree(my_gvr);
gvr_creation_failure: return(False);
}


void gvr_destroy_gvr(Gvr *my_gvr)
{
    Gid *my_port;

    gvd_destroy_gvd(my_gvr->gvd);
    gip_destroy_gip(my_gvr->g.gip);

    while ((my_port = my_gvr->g.gid) != NULL)
        gid_destroy_port(&my_gvr->g, my_port->port_no);
}
```



Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

```
void gvr_added_port(Gvr *my_gvr, int port_no)
{  /*
    * Query the Permanent Database for Static VLAN Entries with “Registration
    * Forbidden” or “Registration Fixed” for this Port. Repeat the following
    * steps until there are no more entries to be found.
    *
    * Check that the VLAN ID is represented in VLD. If not, create it, then
    * create GID machines for all the other Ports with control state “Normal
    * Registration” and create the GID machine for this Port. Change the
    * control state for this Port’s GID machine to forbidden or fixed as
    * required.
    *
    */
}


void gvr_removed_port(Gvr *my_gvr, int port_no)
{  /*
    * Provide any GVR specific cleanup or management alert functions for the
    * removed Port.
    */
}


/*****************************************************************************
 * GVR : GARP VLAN REGISTRATION APPLICATION : JOIN, LEAVE INDICATIONS
 *****************************************************************************
 */

void gvr_join_indication(Gvr *my_gvr, Gid *my_port, unsigned gid_index)
{  /*
    *
    */
    Vlan_id      key;

    gvd_get_key(my_gvr->gvd, gid_index, &key);
    vfdb_forward(my_port->port_no, key);
}


void gvr_join_leave_propagated(Gvr *my_gvr, Gid *my_port, unsigned gid_index)
{  /*
    * Nothing to be done since, unlike GMR with its Forward All Unregistered
    * Port mode, a join indication on one Port does not cause filtering to be
    * instantiated on another.
    */
}


void gvr_leave_indication(Gvr *my_gvr, Gid *my_port, unsigned leaving_gid_index)
{  /*
    *
    */
    Vlan_id  key;

    gvd_get_key(my_gvr->gvd, leaving_gid_index, &key);
    vfdb_filter(my_port->port_no, key);
}

/*****************************************************************************
 * GVR : GARP VLAN REGISTRATION APPLICATION : RECEIVE MESSAGE PROCESSING
 *****************************************************************************
 */
```

Copyright © 1999 IEEE. All rights reserved.



Dell et al., Ex. 1008
Page 101 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

```
static void gvr_db_full(Gvr *my_gvr, Gid *my_port)
{  /*
    * Placeholder for management alert functions indicating registrations
    * for more VLANs have been received than can be accepted.
    */
}

static void gvr_rcv_msg(Gvr *my_gvr, Gid *my_port, Gvf_msg *msg)
{  /*
    * Process one received message.
    *
    * A LeaveAll message never causes an indication (join or leave directly),
    * even for the point-to-point link protocol enhancements (where an
    * ordinary Leave does). No further work is needed here.
    *
    * A LeaveAllRange message is currently treated exactly as a LeaveAll
    * (i.e., the range is ignored).
    *
    * All the remaining messages refer to a single attribute (i.e., a single
    * registered VLAN). Try to find a matching entry in the gvd database.
    * If one is found, dispatch the message to a routine that will
    * handle both the local GID effects and the GIP propagation to other Ports.
    *
    * If no entry is found, Leave and Empty messages can be discarded, but
    * JoinIn and JoinEmpty messages demand further treatment. First, an attempt
    * is made to create a new entry using free space (in the database, which
    * corresponds to a free GID machine set). If this fails, an attempt may be
    * made to recover space from a machine set that is in an unused or less
    * significant state. Finally, the database is considered full and the received
    * message is discarded.
    *
    * Once (if) an entry is found, Leave, Empty, JoinIn, and JoinEmpty are
    * all submitted to GID (gid_rcv_msg()), which will generate and propagate
    * Join or Leave indications as necessary.
    *
    * JoinIn and JoinEmpty may cause Join indications, which are then propagated
    * by GIP.
    *
    * On a shared medium, Leave and Empty will not give rise to indications
    * immediately. However, this routine does test for and propagate
    * Leave indications so that it can be used unchanged with a point-to-point
    * protocol enhancement.
    *
    */

    unsigned gid_index  = Unused_index;

    if (   (msg->event == Gid_rcv_leaveall)
        || (msg->event == Gid_rcv_leaveall_range))
    {
        gid_rcv_leaveall(my_port);
    }
    else
    {
        if (!gvd_find_entry(my_gvr->gvd, msg->key1, &gid_index))
        {
            if (   (msg->event == Gid_rcv_joinin)
                || (msg->event == Gid_rcv_joinempty))
            {
                if (!gvd_create_entry(my_gvr->gvd, msg->key1, &gid_index))
                {
                    if (gid_find_unused(&my_gvr->g, Zero, &gid_index))
                    {
                        gvd_delete_entry(my_gvr->gvd, gid_index);
                        (void) gvd_create_entry(my_gvr->gvd, msg->key1,
```


Copyright © 1999 IEEE. All rights reserved.

```
                                            &gid_index);
                }
                else
                    gvr_db_full(my_gvr, my_port);
        }   }   }

        if (gid_index != Unused_index)
            gid_rcv_msg(my_port, gid_index, msg->event);
}   }


void gvr_rcv(Gvr *my_gvr, Gid *my_port, Pdu *pdu)
{  /*
    * Process an entire received pdu for this instance of GVR: initialize
    * the GVF pdu parsing routine, and, while messages last, read and process
    * them one at a time.
    */
    Gvf      gvf;
    Gvf_msg  msg;

    gvf_rdmsg_init(&gvf, pdu);

    while (gvf_rdmsg(&gvf, &msg))
        gvr_rcv_msg(my_gvr, my_port, &msg);
}


/*****************************************************************************
 * GVR : GARP VLAN REGISTRATION APPLICATION : TRANSMIT PROCESSING
 *****************************************************************************
 */

static void gvr_tx_msg(Gvr *my_gvr, unsigned gid_index, Gvf_msg *msg)
{  /*
    * Fill in msg fields for transmission.
    */

    if (msg->event == Gid_tx_leaveall)
    {
        msg->attribute = All_attributes;
    }
    else
    {
        msg->attribute = Vlan_attribute;

        gvd_get_key(my_gvr->gvd, gid_index, &msg->key1);
}   }


void gvr_tx(Gvr *my_gvr, Gid *my_port)
{  /*
    * Get and prepare a pdu for the transmission, if one is not available,
    * simply return; if there is more to transmit, GID will reschedule a call
    * to this function.
    *
    * Get messages to transmit from GID and pack them into the pdu using GVF
    * (GARP VLAN pdu Formatter).
    */
    Pdu         *pdu;
    Gvf          gvf;
    Gvf_msg      msg;
    Gid_event    tx_event;
    unsigned     gid_index;

    if ((tx_event = gid_next_tx(my_port, &gid_index)) != Gid_null)
```

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 103 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

```
    {
        if (syspdu_alloc(&pdu))
        {
            gvf_wrmsg_init(&gvf, pdu, my_gvr->vlan_id);

            do
            {
                msg.event = tx_event;

                gvr_tx_msg(my_gvr, gid_index, &msg);

                if (!gvf_wrmsg(&gvf, &msg))
                {
                    gid_untx(my_port);
                    break;
                }
            } while ((tx_event  = gid_next_tx(my_port, &gid_index))
                          != Gid_null);

            syspdu_tx(pdu, my_port->port_no);
}   }   }
```

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# 12. VLAN Bridge Management

This clause defines the set of managed objects, and their functionality, that allow administrative configuration of VLANs.

This clause

a) Introduces the functions of management to assist in the identification of the requirements placed on Bridges for the support of management facilities.
b) Establishes the correspondence between the Processes used to model the operation of the Bridge (8.3) and the managed objects of the Bridge.
c) Specifies the management operations supported by each managed object.

## 12.1 Management functions

Management functions relate to the users' needs for facilities that support the planning, organization, supervision, control, protection, and security of communications resources, and account for their use. These facilities may be categorized as supporting the functional areas of Configuration, Fault, Performance, Security, and Accounting Management. Each of these is summarized in 12.1.1 through 12.1.5, together with the facilities commonly required for the management of communication resources, and the particular facilities provided in that functional area by Bridge Management.

### 12.1.1 Configuration Management

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

a) The identification of all Bridges that together make up the Bridged LAN and their respective locations and, as a consequence of that identification, the location of specific end stations to particular individual LANs.
b) The ability to remotely reset, i.e., reinitialize, specified Bridges.
c) The ability to control the priority with which a Bridge Port transmits frames.
d) The ability to force a specific configuration of the spanning tree.
e) The ability to control the propagation of frames with specific group MAC Addresses to certain parts of the configured Bridged LAN.
f) The ability to identify the VLANs in use, and through which Ports of the Bridge frames destined for a given VLAN may be received and/or forwarded.

### 12.1.2 Fault Management

Fault Management provides for fault prevention, detection, diagnosis, and correction. The facilities provided by Bridge Management in this functional area are

a) The ability to identify and correct Bridge malfunctions, including error logging and reporting.

### 12.1.3 Performance Management

Performance Management provides for evaluation of the behavior of communications resources and of the effectiveness of communication activities. The facilities provided by Bridge Management in this functional area are

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 105 of 211

a) The ability to gather statistics relating to performance and traffic analysis. Specific metrics include network utilization, frame forward, and frame discard counts for individual Ports within a Bridge.

### 12.1.4 Security Management

Security Management provides for the protection of resources. Bridge Management does not provide any specific facilities in this functional area.

### 12.1.5 Accounting Management

Accounting Management provides for the identification and distribution of costs and the setting of charges. Bridge Management does not provide any specific facilities in this functional area.

## 12.2 Managed objects

Managed objects model the semantics of management operations. Operations upon an object supply information concerning, or facilitate control over, the Process or Entity associated with that object.

The managed resources of a MAC Bridge are those of the Processes and Entities established in 8.3 and ISO/IEC 15802-3, 12.2. Specifically,

a) The Bridge Management Entity (12.4 and 8.13).
b) The individual MAC Entities associated with each Bridge Port (12.5, 8.2, 8.5, and 8.9).
c) The Forwarding Process of the MAC Relay Entity (12.6, 8.2, and 8.7).
d) The Filtering Database of the MAC Relay Entity (12.7 and 8.11).
e) The Bridge Protocol Entity (12.8 and 8.12; ISO/IEC 15802-3, Clause 8).
f) GARP Participants (ISO/IEC 15802-3, Clause 12).

The management of each of these resources is described in terms of managed objects and operations below.

NOTE—The values specified in this clause, as inputs and outputs of management operations, are abstract information elements. Questions of formats or encodings are a matter for particular protocols that convey or otherwise represent this information.

## 12.3 Data types

This subclause specifies the semantics of operations independent of their encoding in management protocol. The data types of the parameters of operations are defined only as required for that specification.

The following data types are used:

a) Boolean.
b) Enumerated, for a collection of named values.
c) Unsigned, for all parameters specified as “the number of” some quantity, and for Spanning Tree priority values that are numerically compared. When comparing Spanning Tree priority values, the lower number represents the higher priority value.
d) MAC Address.
e) Latin1 String, as defined by ANSI X3.159, for all text strings.
f) Time Interval, an Unsigned value representing a positive integral number of seconds, for all Spanning Tree protocol timeout parameters;
g) Counter, for all parameters specified as a “count” of some quantity. A counter increments and wraps with a modulus of 2 to the power of 64.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

h) GARP Time Interval, an Unsigned value representing a positive integral number of centiseconds, for all GARP protocol time-out parameters.

## 12.4 Bridge Management Entity

The Bridge Management Entity is described in 8.13.

The objects which comprise this managed resource are

a) The Bridge Configuration (12.4.1).
b) The Port Configuration for each Port (12.4.2).

### 12.4.1 Bridge Configuration

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

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

a) Discover Bridge (12.4.1.1);
b) Read Bridge (12.4.1.2);
c) Set Bridge Name (12.4.1.3);
d) Reset Bridge (12.4.1.4).

#### 12.4.1.1 Discover Bridge

##### 12.4.1.1.1 Purpose

To solicit configuration information regarding the Bridge(s) in the Bridged LAN.

##### 12.4.1.1.2 Inputs

a) Inclusion Range, a set of ordered pairs of specific MAC Addresses. Each pair specifies a range of MAC Addresses. A Bridge shall respond if and only if
   1) For one of the pairs, the numerical comparison of its Bridge Address with each MAC Address of the pair shows it to be greater than or equal to the first, and
   2) Less than or equal to the second, and
   3) Its Bridge Address does not appear in the Exclusion List parameter below.

The numerical comparison of one MAC Address with another, for the purpose of this operation, is achieved by deriving a number from the MAC Address according to the following procedure. The consecutive octets of the MAC Address are taken to represent a binary number; the first octet that would be transmitted on a LAN medium when the MAC Address is used in the source or destination fields of a MAC frame has the most significant value, the next octet the next most significant value. Within each octet the first bit of each octet is the least significant bit.

b) Exclusion List, a list of specific MAC Addresses.

##### 12.4.1.1.3 Outputs

a) Bridge Address—the MAC Address for the Bridge from which the Bridge Identifier used by the Spanning Tree Algorithm and Protocol is derived (8.14.5; ISO/IEC 15802-3, 8.5.1.3).
b) Bridge Name—a text string of up to 32 characters, of locally determined significance.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 107 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

c) Number of Ports—the number of Bridge Ports (MAC Entities).
d) Port Addresses—a list specifying the following for each Port:
   1) Port Number—the number of the Bridge Port (ISO/IEC 15802-3, 8.5.5.1).
   2) Port Address—the specific MAC Address of the individual MAC Entity associated with the Port (8.14.2).
e) Uptime—count in seconds of the time elapsed since the Bridge was last reset or initialized (ISO/IEC 15802-3, 8.8.1).

#### 12.4.1.2 Read Bridge

##### 12.4.1.2.1 Purpose

To obtain general information regarding the Bridge.

##### 12.4.1.2.2 Inputs

None.

##### 12.4.1.2.3 Outputs

a) Bridge Address—the MAC Address for the Bridge from which the Bridge Identifier used by the Spanning Tree Algorithm and Protocol is derived (8.14.5; ISO/IEC 15802-3, 8.5.1.3).
b) Bridge Name—a text string of up to 32 characters, of locally determined significance.
c) Number of Ports—the number of Bridge Ports (MAC Entities).
d) Port Addresses—a list specifying the following for each Port:
   1) Port Number (ISO/IEC 15802-3, 8.5.5.1).
   2) Port Address—the specific MAC Address of the individual MAC Entity associated with the Port (8.14.2).
e) Uptime—count in seconds of the time elapsed since the Bridge was last reset or initialized (ISO/IEC 15802-3, 8.8.1).

#### 12.4.1.3 Set Bridge Name

##### 12.4.1.3.1 Purpose

To associate a text string, readable by the Read Bridge operation, with a Bridge.

##### 12.4.1.3.2 Inputs

a) Bridge Name—a text string of up to 32 characters.

##### 12.4.1.3.3 Outputs

None.

#### 12.4.1.4 Reset Bridge

##### 12.4.1.4.1 Purpose

To reset the specified Bridge. The Filtering Database is cleared and initialized with the entries specified in the Permanent Database, and the Bridge Protocol Entity is initialized (ISO/IEC 15802-3, 8.8.1).


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 12.4.1.4.2 Inputs

None.

#### 12.4.1.4.3 Outputs

None.

### 12.4.2 Port configuration

The Port Configuration object models the operations that modify, or inquire about, the configuration of the Ports of a Bridge. There are a fixed set of Bridge Ports per Bridge (one for each MAC interface), and each is identified by a permanently allocated Port Number.

The allocated Port Numbers are not required to be consecutive. Also, some Port Numbers may be dummy entries, with no actual LAN Port (for example, to allow for expansion of the Bridge by addition of further MAC interfaces in the future). Such dummy Ports shall support the Port Configuration management operations, and other Port-related management operations in a manner consistent with the Port being permanently disabled.

The information provided by the Port Configuration consists of summary data indicating its name and type. Specific counter information pertaining to the number of packets forwarded, filtered, and in error is maintained by the Forwarding Process resource. The management operations supported by the Bridge Protocol Entity allow for controlling the states of each Port.

The management operations that can be performed on the Port Configuration are

a) Read Port (12.4.2.1);
b) Set Port Name (12.4.2.2).

#### 12.4.2.1 Read Port

#### 12.4.2.1.1 Purpose

To obtain general information regarding a specific Bridge Port.

#### 12.4.2.1.2 Inputs

a) Port Number—the number of the Bridge Port (ISO/IEC 15802-3, 8.5.5.1).

#### 12.4.2.1.3 Outputs

a) Port Name—a text string of up to 32 characters, of locally determined significance.
b) Port Type—the MAC Entity type of the Port (IEEE Std 802.3; ISO/IEC 8802-4; ISO/IEC 8802-5; ISO/IEC 8802-6; ISO/IEC 8802-9; IEEE Std 802.9a-1995; ISO/IEC 8802-12 (IEEE Std 802.3 format); ISO/IEC 8802-12 (ISO/IEC 8802-5 format); ISO 9314; other).

#### 12.4.2.2 Set Port Name

#### 12.4.2.2.1 Purpose

To associate a text string, readable by the Read Port operation, with a Bridge Port.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 109 of 211

#### 12.4.2.2.2 Inputs

a) Port Number (ISO/IEC 15802-3, 8.5.5.1).
b) Port Name—a text string of up to 32 characters.

#### 12.4.2.2.3 Outputs

None.

## 12.5 MAC entities

The Management Operations and Facilities provided by the MAC Entities are those specified in the Layer Management standards of the individual MACs. A MAC Entity is associated with each Bridge Port.

## 12.6 Forwarding process

The Forwarding Process contains information relating to the forwarding of frames. Counters are maintained that provide information on the number of frames forwarded, filtered, and dropped due to error. Configuration data, defining how frame priority is handled, is maintained by the Forwarding Process.

The objects that comprise this managed resource are

a) The Port Counters (12.6.1).
b) The Priority Handling objects for each Port (12.6.2);
c) The Traffic Class Table for each Port (12.6.3).

### 12.6.1 The Port Counters

The Port Counters object models the operations that can be performed on the Port counters of the Forwarding Process resource. There are multiple instances (one for each VLAN for each MAC Entity) of the Port Counters object per Bridge.

The management operation that can be performed on the Port Counters is Read Forwarding Port Counters (12.6.1.1).

#### 12.6.1.1 Read forwarding port counters

#### 12.6.1.1.1 Purpose

To read the forwarding counters associated with a specific Bridge Port.

#### 12.6.1.1.2 Inputs

a) Port Number (ISO/IEC 15802-3, 8.5.5.1);
b) Optionally, VLAN Identifier (9.3.2.3).

If the VLAN Identifier parameter is supported, then the forwarding Port counters are maintained per VLAN per Port. If the parameter is not supported, then the forwarding Port counters are maintained per Port only.

#### 12.6.1.1.3 Outputs

a) Frames Received—count of all valid frames received (including BPDUs, frames addressed to the Bridge as an end station and frames that were submitted to the Forwarding Process, 8.5).


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

b) Optionally, Octets Received—count of the total number of octets in all valid frames received (including BPDUs, frames addressed to the Bridge as an end station, and frames that were submitted to the Forwarding Process).
c) Discard Inbound—count of valid frames received that were discarded by the Forwarding Process (8.7).
d) Forward Outbound—count of frames forwarded to the associated MAC Entity (8.9).
e) Discard Lack of Buffers—count of frames that were to be transmitted through the associated Port but were discarded due to lack of buffers (8.7.3).
f) Discard Transit Delay Exceeded—count of frames that were to be transmitted but were discarded due to the maximum bridge transit delay being exceeded (buffering may have been available, 8.7.3).
g) Discard on Error—count of frames that were to be forwarded on the associated MAC but could not be transmitted (e.g., frame would be too large, ISO/IEC 15802-3, 6.3.8).
h) If Ingress Filtering is supported (8.4.5), Discard on Ingress Filtering—count of frames that were discarded as a result of Ingress Filtering being enabled.
i) Optionally, Discard on Error Details—a list of 16 elements, each containing the source address of a frame and the reason why the frame was discarded (frame too large). The list is maintained as a circular buffer. The reasons for discard on error, at present, are
    1) Transmissible service data unit size exceeded; or
    2) Discard due to Ingress Filtering. The VID associated with the last discarded frame is recorded.

### 12.6.2 Priority handling

The Priority Handling object models the operations that can be performed upon, or inquire about, the Default User Priority parameter, the User Priority Regeneration Table parameter, and the Outbound Access Priority Table parameter for each Port. The operations that can be performed on this object are

a) Read Port Default User Priority (12.6.2.1);
b) Set Port Default User Priority (12.6.2.2);
c) Read Port User Priority Regeneration Table (12.6.2.3);
d) Set Port User Priority Regeneration Table (12.6.2.4);
e) Read Outbound Access Priority Table (12.6.2.5).

#### 12.6.2.1 Read Port Default User Priority

##### 12.6.2.1.1 Purpose

To read the current state of the Default User Priority parameter (ISO/IEC 15802-3, 6.4) for a specific Bridge Port.

##### 12.6.2.1.2 Inputs

a) Port number.

##### 12.6.2.1.3 Outputs

a) Default User Priority value—Integer in range 0–7.

#### 12.6.2.2 Set Port Default User Priority

##### 12.6.2.2.1 Purpose

To set the current state of the Default User Priority parameter (ISO/IEC 15802-3, 6.4) for a specific Bridge Port.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 111 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 12.6.2.2.2 Inputs

a) Port number;
b) Default User Priority value—Integer in range 0–7.

#### 12.6.2.2.3 Outputs

None.

### 12.6.2.3 Read Port User Priority Regeneration Table

#### 12.6.2.3.1 Purpose

To read the current state of the User Priority Regeneration Table parameter (8.5.1) for a specific Bridge Port.

#### 12.6.2.3.2 Inputs

a) Port number.

#### 12.6.2.3.3 Outputs

a) Regenerated User Priority value for Received User Priority 0—Integer in range 0–7.
b) Regenerated User Priority value for Received User Priority 1—Integer in range 0–7.
c) Regenerated User Priority value for Received User Priority 2—Integer in range 0–7.
d) Regenerated User Priority value for Received User Priority 3—Integer in range 0–7.
e) Regenerated User Priority value for Received User Priority 4—Integer in range 0–7.
f) Regenerated User Priority value for Received User Priority 5—Integer in range 0–7.
g) Regenerated User Priority value for Received User Priority 6—Integer in range 0–7.
h) Regenerated User Priority value for Received User Priority 7—Integer in range 0–7.

### 12.6.2.4 Set Port User Priority Regeneration Table

#### 12.6.2.4.1 Purpose

To set the current state of the User Priority Regeneration Table parameter (8.5.1) for a specific Bridge Port.

#### 12.6.2.4.2 Inputs

a) Port number;
b) Regenerated User Priority value for Received User Priority 0—Integer in range 0–7.
c) Regenerated User Priority value for Received User Priority 1—Integer in range 0–7.
d) Regenerated User Priority value for Received User Priority 2—Integer in range 0–7.
e) Regenerated User Priority value for Received User Priority 3—Integer in range 0–7.
f) Regenerated User Priority value for Received User Priority 4—Integer in range 0–7.
g) Regenerated User Priority value for Received User Priority 5—Integer in range 0–7.
h) Regenerated User Priority value for Received User Priority 6—Integer in range 0–7.
i) Regenerated User Priority value for Received User Priority 7—Integer in range 0–7.

#### 12.6.2.4.3 Outputs

None.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 12.6.2.5 Read Outbound Access Priority Table

##### 12.6.2.5.1 Purpose

To read the state of the Outbound Access Priority Table parameter (Table 8-3) for a specific Bridge Port.

##### 12.6.2.5.2 Inputs

a) Port number.

##### 12.6.2.5.3 Outputs

a) Access Priority value for User Priority 0—Integer in range 0–7.
b) Access Priority value for User Priority 1—Integer in range 0–7.
c) Access Priority value for User Priority 2—Integer in range 0–7.
d) Access Priority value for User Priority 3—Integer in range 0–7.
e) Access Priority value for User Priority 4—Integer in range 0–7.
f) Access Priority value for User Priority 5—Integer in range 0–7.
g) Access Priority value for User Priority 6—Integer in range 0–7.
h) Access Priority value for User Priority 7—Integer in range 0–7.

### 12.6.3 Traffic Class Table

The Traffic Class Table object models the operations that can be performed upon, or inquire about, the current contents of the Traffic Class Table (8.7.3) for a given Port. The operations that can be performed on this object are Read Port Traffic Class Table and Set Port Traffic Class Table.

#### 12.6.3.1 Read Port Traffic Class Table

##### 12.6.3.1.1 Purpose

To read the contents of the Traffic Class Table (8.7.3) for a given Port.

##### 12.6.3.1.2 Inputs

a) Port Number.

##### 12.6.3.1.3 Outputs

a) The number of Traffic Classes, in the range 1 through 8, supported on the Port;
b) For each value of Traffic Class supported on the Port, the value of the Traffic Class in the range 0 through 7, and the set of user_priority values assigned to that Traffic Class.

#### 12.6.3.2 Set Port Traffic Class Table

##### 12.6.3.2.1 Purpose

To set the contents of the Traffic Class Table (8.7.3) for a given Port.

##### 12.6.3.2.2 Inputs

a) Port number;
b) For each value of Traffic Class supported on the Port, the value of the Traffic Class in the range 0 through 7, and the set of user_priority values assigned to that Traffic Class.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 113 of 211

NOTE—If a Traffic Class value greater than the largest Traffic Class available on the Port is specified, then the value applied to the Traffic Class Table is the largest available Traffic Class.

#### 12.6.3.2.3 Outputs

None.

## 12.7 Filtering Database

The Filtering Database is described in 8.11. It contains filtering information used by the Forwarding Process (8.7) in deciding through which Ports of the Bridge frames should be forwarded.

The objects that comprise this managed resource are

a) The Filtering Database (12.7.1);
b) The Static Filtering Entries (12.7.2);
c) The Dynamic Filtering Entries (12.7.3);
d) The Group Registration Entries (12.7.4);
e) The Static VLAN Registration Entries (12.7.5);
f) The Dynamic VLAN Registration Entries (12.7.5);
g) The Permanent Database (12.7.6).

### 12.7.1 The Filtering Database

The Filtering Database object models the operations that can be performed on, or affect, the Filtering Database as a whole. There is a single Filtering Database object per Bridge.

The management operations that can be performed on the Database are:

a) Read Filtering Database (12.7.1.1);
b) Set Filtering Database Ageing Time (12.7.1.2);
c) Read Permanent Database (12.7.6.1);
d) Create Filtering Entry (12.7.7.1);
e) Delete Filtering Entry (12.7.7.2);
f) Read Filtering Entry (12.7.7.3);
g) Read Filtering Entry Range (12.7.7.4).

#### 12.7.1.1 Read Filtering Database

##### 12.7.1.1.1 Purpose

To obtain general information regarding the Bridge’s Filtering Database.

##### 12.7.1.1.2 Inputs

None.

##### 12.7.1.1.3 Outputs

a) Filtering Database Size—the maximum number of entries that can be held in the Filtering Database.
b) Number of Static Filtering Entries—the number of Static Filtering Entries (8.11.1) currently in the Filtering Database;


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

c) Number of Dynamic Filtering Entries—the number of Dynamic Filtering Entries (8.11.3) currently in the Filtering Database;
d) Number of Static VLAN Registration Entries—the number of Static VLAN Registration Entries (8.11.2) currently in the Filtering Database;
e) Number of Dynamic VLAN Registration Entries—the number of Dynamic VLAN Registration Entries (8.11.5) currently in the Filtering Database.
f) Ageing Time—for ageing out Dynamic Filtering Entries when the Port associated with the entry is in the Forwarding state (8.11.3).
g) If Extended Filtering Services are supported, Number of Group Registration Entries—the number of Group Registration Entries (8.11.4) currently in the Filtering Database;

#### 12.7.1.2 Set Filtering Database Ageing Time

##### 12.7.1.2.1 Purpose

To set the ageing time for Dynamic Filtering Entries (8.11.3).

##### 12.7.1.2.2 Inputs

a) Ageing Time.

##### 12.7.1.2.3 Outputs

None.

### 12.7.2 A Static Filtering Entry

A Static Filtering Entry object models the operations that can be performed on a single Static Filtering Entry in the Filtering Database. The set of Static Filtering Entry objects within the Filtering Database changes only under management control.

A Static Filtering Entry object supports the following operations:

a) Create Filtering Entry (12.7.7.1);
b) Delete Filtering Entry (12.7.7.2);
c) Read Filtering Entry (12.7.7.3);
d) Read Filtering Entry Range (12.7.7.4).

### 12.7.3 A Dynamic Filtering Entry

A Dynamic Filtering Entry object models the operations that can be performed on a single Dynamic Filtering Entry (i.e., one that is created by the Learning Process as a result of the observation of network traffic) in the Filtering Database.

A Dynamic Filtering Entry object supports the following operations:

a) Delete Filtering Entry (12.7.7.2);
b) Read Filtering Entry (12.7.7.3);
c) Read Filtering Entry Range (12.7.7.4).

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 115 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### 12.7.4 A Group Registration Entry

A Group Registration Entry object models the operations that can be performed on a single Group Registration Entry in the Filtering Database. The set of Group Registration Entry objects within the Filtering Database changes only as a result of GARP protocol exchanges.

A Group Registration Entry object supports the following operations:

a) Read Filtering Entry (12.7.7.3);
b) Read Filtering Entry Range (12.7.7.4).

### 12.7.5 A VLAN Registration Entry

A VLAN Registration Entry object models the operations that can be performed on a single VLAN Registration Entry in the Filtering Database. The set of VLAN Registration Entry objects within the Filtering Database changes under management control and also as a result of GARP protocol exchanges.

#### 12.7.5.1 Static VLAN Registration Entry object

A Static VLAN Registration Entry object supports the following operations:

a) Create Filtering Entry (12.7.7.1);
b) Delete Filtering Entry (12.7.7.2);
c) Read Filtering Entry (12.7.7.3);
d) Read Filtering Entry Range (12.7.7.4).

#### 12.7.5.2 Dynamic VLAN Registration Entry object

A Dynamic VLAN Registration Entry object supports the following operations:

a) Read Filtering Entry (12.7.7.3);
b) Read Filtering Entry Range (12.7.7.4).

### 12.7.6 Permanent Database

The Permanent Database object models the operations that can be performed on, or affect, the Permanent Database. There is a single Permanent Database per Filtering Database.

The management operations that can be performed on the Permanent Database are

a) Read Permanent Database (12.7.6.1);
b) Create Filtering Entry (12.7.7.1);
c) Delete Filtering Entry (12.7.7.2);
d) Read Filtering Entry (12.7.7.3);
e) Read Filtering Entry Range (12.7.7.4).

#### 12.7.6.1 Read Permanent Database

##### 12.7.6.1.1 Purpose

To obtain general information regarding the Permanent Database (8.11.10).


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 12.7.6.1.2 Inputs

None.

#### 12.7.6.1.3 Outputs

a) Permanent Database Size—maximum number of entries that can be held in the Permanent Database.
b) Number of Static Filtering Entries—number of Static Filtering Entries (8.11.1) currently in the Permanent Database;
c) Number of Static VLAN Registration Entries—number of Static VLAN Registration Entries (8.11.2) currently in the Permanent Database.

### 12.7.7 General Filtering Database operations

In these operations on the Filtering Database, the operation parameters make use of VID values, even when operating upon a Dynamic Filtering Entry (8.11.3) whose structure carries an FID rather than a VID. In this case, the value used in the VID parameter can be any VID that has been allocated to the FID concerned (8.11.7).

#### 12.7.7.1 Create Filtering Entry

#### 12.7.7.1.1 Purpose

To create or update a Static Filtering Entry (8.11.1) or Static VLAN Registration Entry (8.11.2) in the Filtering Database or Permanent Database. Only static entries may be created in the Filtering Database or Permanent Database.

#### 12.7.7.1.2 Inputs

a) Identifier—Filtering Database or Permanent Database.
b) Address—MAC Address of the entry (not present in VLAN Registration Entries).
c) VID—VLAN Identifier of the entry.
d) Port Map—a set of control indicators, one for each Port, as specified in 8.11.1 and 8.11.2.

#### 12.7.7.1.3 Outputs

None.

#### 12.7.7.2 Delete Filtering Entry

#### 12.7.7.2.1 Purpose

To delete a Filtering Entry or VLAN Registration Entry from the Filtering Database or Permanent Database.

#### 12.7.7.2.2 Inputs

a) Identifier—Filtering Database or Permanent Database.
b) Address—MAC Address of the desired entry (not present in VLAN Registration Entries).
c) VID—VLAN Identifier of the entry.

#### 12.7.7.2.3 Outputs

None.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 117 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 12.7.7.3 Read Filtering Entry

##### 12.7.7.3.1 Purpose

To read a Filtering Entry, Group Registration Entry, or VLAN Registration Entry from the Filtering or Permanent Databases.

##### 12.7.7.3.2 Inputs

a) Identifier—Filtering Database or Permanent Database.
b) Address—MAC Address of the desired entry (not present in VLAN Registration Entries).
c) VID—VLAN Identifier of the entry.
d) Type—Static or Dynamic entry.

##### 12.7.7.3.3 Outputs

a) Address—MAC Address of the desired entry (not present in VLAN Registration Entries).
b) VID—VLAN Identifier of the entry.
c) Type—Static or Dynamic entry.
d) Port Map—a set of control indicators as appropriate for the entry, as specified in 8.11.1 through 8.11.5.

#### 12.7.7.4 Read Filtering Entry range

##### 12.7.7.4.1 Purpose

To read a range of Filtering Database entries (of any type) from the Filtering or Permanent Databases.

Since the number of values to be returned in the requested range may have exceeded the capacity of the service data unit conveying the management response, the returned entry range is identified. The indices that define the range take on values from zero up to Filtering Database Size minus one.

##### 12.7.7.4.2 Inputs

a) Identifier—Filtering Database or Permanent Database.
b) Start Index—inclusive starting index of the desired entry range.
c) Stop Index—inclusive ending index of the desired range.

##### 12.7.7.4.3 Outputs

a) Start Index—inclusive starting index of the returned entry range.
b) Stop Index—inclusive ending index of the returned entry range.
c) For each index returned:
    1) Address—MAC Address of the desired entry (not present in VLAN Registration Entries).
    2) VID—VLAN Identifier of the entry.
    3) Type—Static or Dynamic entry.
    4) Port Map—a set of control indicators as appropriate for the entry, as specified in 8.11.1 through 8.11.5.

## 12.8 Bridge Protocol Entity

The Bridge Protocol Entity is described in 8.12 and ISO/IEC 15802-3, Clause 8.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

The objects that comprise this managed resource are

a) The Protocol Entity itself.
b) The Ports under its control.

### 12.8.1 The Protocol Entity

The Protocol Entity object models the operations that can be performed upon, or inquire about, the operation of the Spanning Tree Algorithm and Protocol. There is a single Protocol Entity per Bridge; it can, therefore, be identified as a single fixed component of the Protocol Entity resource.

The management operations that can be performed on the Protocol Entity are

a) Read Bridge Protocol Parameters (12.8.1.1);
b) Set Bridge Protocol Parameters (12.8.1.2).

#### 12.8.1.1 Read Bridge Protocol Parameters

##### 12.8.1.1.1 Purpose

To obtain information regarding the Bridge's Bridge Protocol Entity.

##### 12.8.1.1.2 Inputs

None.

##### 12.8.1.1.3 Outputs

a) Bridge Identifier—as defined in ISO/IEC 15802-3, 8.5.3.7.
b) Time Since Topology Change—count in seconds of the time elapsed since the Topology Change flag parameter for the Bridge (ISO/IEC 15802-3, 8.5.3.12) was last True.
c) Topology Change Count—count of the times the Topology Change flag parameter for the Bridge has been set (i.e., transitioned from False to True) since the Bridge was powered on or initialized.
d) Topology Change (ISO/IEC 15802-3, 8.5.3.12).
e) Designated Root (ISO/IEC 15802-3, 8.5.3.1).
f) Root Path Cost (ISO/IEC 15802-3, 8.5.3.2).
g) Root Port (ISO/IEC 15802-3, 8.5.3.3).
h) Max Age (ISO/IEC 15802-3, 8.5.3.4).
i) Hello Time (ISO/IEC 15802-3, 8.5.3.5).
j) Forward Delay (ISO/IEC 15802-3, 8.5.3.6).
k) Bridge Max Age (ISO/IEC 15802-3, 8.5.3.7).
l) Bridge Hello Time (ISO/IEC 15802-3, 8.5.3.9).
m) Bridge Forward Delay (ISO/IEC 15802-3, 8.5.3.10).
n) Hold Time (ISO/IEC 15802-3, 8.5.3.14).

#### 12.8.1.2 Set Bridge Protocol Parameters

##### 12.8.1.2.1 Purpose

To modify parameters in the Bridge's Bridge Protocol Entity in order to force a configuration of the spanning tree and/or tune the reconfiguration time to suit a specific topology.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 119 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 12.8.1.2.2 Inputs

a) Bridge Max Age—the new value (ISO/IEC 15802-3, 8.5.3.8).
b) Bridge Hello Time—the new value (ISO/IEC 15802-3, 8.5.3.9).
c) Bridge Forward Delay—the new value (ISO/IEC 15802-3, 8.5.3.10).
d) Bridge Priority—the new value of the priority part of the Bridge Identifier (ISO/IEC 15802-3, 8.5.3.7).

#### 12.8.1.2.3 Outputs

None.

#### 12.8.1.2.4 Procedure

The input parameter values are checked for compliance with ISO/IEC 15802-3, 8.10.2. If they do not comply, or the value of Bridge Max Age or Bridge Forward Delay is less than the lower limit of the range specified in ISO/IEC 15802-3, Table 8-3, then no action shall be taken for any of the supplied parameters. If the value of any of Bridge Max Age, Bridge Forward Delay, or Bridge Hello Time is outside the range specified in ISO/IEC 15802-3, Table 8-3, then the Bridge need not take action.

Otherwise, the Bridge's Bridge Max Age, Bridge Hello Time, and Bridge Forward Delay parameters are set to the supplied values. The Set Bridge Priority procedure (ISO/IEC 15802-3, 8.8.4) is used to set the priority part of the Bridge Identifier to the supplied value.

### 12.8.2 Bridge Port

A Bridge Port object models the operations related to an individual Bridge Port in relation to the operation of the Spanning Tree Algorithm and Protocol. There are a fixed set of Bridge Ports per Bridge; each can, therefore, be identified by a permanently allocated Port Number, as a fixed component of the Protocol Entity resource.

The management operations that can be performed on a Bridge Port are

a) Read Port Parameters (12.8.2.1);
b) Force Port State (12.8.2.2);
c) Set Port Parameters (12.8.2.3).

#### 12.8.2.1 Read Port Parameters

#### 12.8.2.1.1 Purpose

To obtain information regarding a specific Port within the Bridge's Bridge Protocol Entity.

#### 12.8.2.1.2 Inputs

a) Port Number—the number of the Bridge Port.

#### 12.8.2.1.3 Outputs

a) Uptime—count in seconds of the time elapsed since the Port was last reset or initialized (ISO/IEC 15802-3, 8.8.1).
b) State—the current state of the Port (i.e., Disabled, Listening, Learning, Forwarding, or Blocking) (ISO/IEC 15802-3, 8.4, and 8.5.5.2).


Copyright © 1999 IEEE. All rights reserved.

c) Port Identifier—the unique Port identifier comprising two parts, the Port Number and the Port Priority field (ISO/IEC 15802-3, 8.5.5.1).
d) Path Cost (ISO/IEC 15802-3, 8.5.5.3).
e) Designated Root (ISO/IEC 15802-3, 8.5.5.4).
f) Designated Cost (ISO/IEC 15802-3, 8.5.5.5).
g) Designated Bridge (ISO/IEC 15802-3, 8.5.5.6).
h) Designated Port (ISO/IEC 15802-3, 8.5.5.7).
i) Topology Change Acknowledge (ISO/IEC 15802-3, 8.5.5.8).

#### 12.8.2.2 Force port state

##### 12.8.2.2.1 Purpose

To force the specified Port into Disabled (ISO/IEC 15802-3, 8.4.5) or Blocking (ISO/IEC 15802-3, 8.4.1).

##### 12.8.2.2.2 Inputs

a) Port Number—the number of the Bridge Port.
b) State—either Disabled or Blocking (ISO/IEC 15802-3, 8.4, 8.4.1, and 8.4.5).

##### 12.8.2.2.3 Outputs

None.

##### 12.8.2.2.4 Procedure

If the selected state is Disabled, the Disable Port procedure (ISO/IEC 15802-3, 8.8.3) is used for the specified Port. If the selected state is Blocking, the Enable Port procedure (ISO/IEC 15802-3, 8.8.2) is used.

#### 12.8.2.3 Set port parameters

##### 12.8.2.3.1 Purpose

To modify parameters for a Port in the Bridge's Bridge Protocol Entity in order to force a configuration of the spanning tree.

##### 12.8.2.3.2 Inputs

a) Port Number—the number of the Bridge Port.
b) Path Cost—the new value (ISO/IEC 15802-3, 8.5.5.3).
c) Port Priority—the new value of the priority field for the Port Identifier (ISO/IEC 15802-3, 8.5.5.1).

##### 12.8.2.3.3 Outputs

None.

##### 12.8.2.3.4 Procedure

The Set Path Cost procedure (ISO/IEC 15802-3, 8.8.6) is used to set the Path Cost parameter for the specified Port. The Set Port Priority procedure (ISO/IEC 15802-3, 8.8.5) is used to set the priority part of the Port Identifier (ISO/IEC 15802-3, 8.5.5.1) to the supplied value.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 121 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## 12.9 GARP Entities

The operation of GARP is described in ISO/IEC 15802-3, Clause 12.

The objects that comprise this managed resource are

a) The GARP Timer objects (12.9.1);
b) The GARP Attribute Type objects (12.9.2);
c) The GARP State Machine objects (12.9.3).

### 12.9.1 The GARP Timer object

The GARP Timer object models the operations that can be performed upon, or inquire about, the current settings of the timers used by the GARP protocol on a given Port. The management operations that can be performed on the GARP Participant are

a) Read GARP Timers (12.9.1.1);
b) Set GARP Timers (12.9.1.2).

#### 12.9.1.1 Read GARP Timers

##### 12.9.1.1.1 Purpose

To read the current GARP Timers for a given Port.

##### 12.9.1.1.2 Inputs

a) The Port identifier.

##### 12.9.1.1.3 Outputs

a) Current value of JoinTime—Centiseconds (ISO/IEC 15802-3, 12.10.2.1 and 12.12.1);
b) Current value of LeaveTime—Centiseconds (ISO/IEC 15802-3, 12.10.2.2 and 12.12.1);
c) Current value of LeaveAllTime—Centiseconds (ISO/IEC 15802-3, 12.10.2.3 and 12.12.1).

#### 12.9.1.2 Set GARP Timers

##### 12.9.1.2.1 Purpose

To set new values for the GARP Timers for a given Port.

##### 12.9.1.2.2 Inputs

a) The Port identifier;
b) New value of JoinTime—Centiseconds (ISO/IEC 15802-3, 12.10.2.1 and 12.12.1);
c) New value of LeaveTime—Centiseconds (ISO/IEC 15802-3, 12.10.2.2 and 12.12.1);
d) New value of LeaveAllTime—Centiseconds (ISO/IEC 15802-3, 12.10.2.3 and 12.12.1).

##### 12.9.1.2.3 Outputs

None.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### 12.9.2 The GARP Attribute Type object

The GARP Attribute Type object models the operations that can be performed upon, or inquire about, the operation of GARP for a given Attribute Type (ISO/IEC 15802-3, 12.11.2.2). The management operations that can be performed on a GARP Attribute Type are

a) Read GARP Applicant Controls (12.9.2.1);
b) Set GARP Applicant Controls (12.9.2.2).

#### 12.9.2.1 Read GARP Applicant Controls

##### 12.9.2.1.1 Purpose

To read the current values of the GARP Applicant Administrative control parameters (ISO/IEC 15802-3, 12.9.2) associated with all GARP Participants for a given Port, GARP Application and Attribute Type.

##### 12.9.2.1.2 Inputs

a) The Port identifier;
b) The GARP Application address (ISO/IEC 15802-3, Table 12-1);
c) The Attribute Type (ISO/IEC 15802-3, 12.11.2.5).

##### 12.9.2.1.3 Outputs

a) The current Applicant Administrative Control Value (ISO/IEC 15802-3, 12.9.2);
b) Failed Registrations—Count of the number of times that this GARP Application has failed to register an attribute of this type due to lack of space in the Filtering Database (12.10.1.6).

#### 12.9.2.2 Set GARP Applicant Controls

##### 12.9.2.2.1 Purpose

To set new values for the GARP Applicant Administrative control parameters (ISO/IEC 15802-3, 12.9.2) associated with all GARP Participants for a given Port, GARP Application and Attribute Type.

##### 12.9.2.2.2 Inputs

a) The Port identifier;
b) The GARP Application address (ISO/IEC 15802-3, Table 12-1);
c) The Attribute Type (ISO/IEC 15802-3, 12.11.2.5) associated with the state machine;
d) The desired Applicant Administrative Control Value (ISO/IEC 15802-3, 12.9.2).

##### 12.9.2.2.3 Outputs

None.

### 12.9.3 The GARP State Machine object

The GARP State Machine object models the operations that can be performed upon, or inquire about, the operation of GARP for a given State Machine.

The management operation that can be performed on a GARP State Machine is Read GARP State.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 123 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 12.9.3.1 Read GARP State

##### 12.9.3.1.1 Purpose

To read the current value of an instance of a GARP state machine.

##### 12.9.3.1.2 Inputs

a) The Port identifier;
b) The GARP Application address (ISO/IEC 15802-3, Table 12-1);
c) The GIP Context (ISO/IEC 15802-3, 12.3.4);
d) The Attribute Type (ISO/IEC 15802-3, 12.11.2.2) associated with the state machine;
e) The Attribute Value (ISO/IEC 15802-3, 12.11.2.6) associated with the state machine.

##### 12.9.3.1.3 Outputs

a) The current value of the combined Applicant and Registrar state machine for the attribute (ISO/IEC 15802-3, Table 12-6);
b) Optionally, Originator address—the MAC Address of the originator of the most recent GARP PDU that was responsible for causing a state change in this state machine (ISO/IEC 15802-3, 12.9.1).

## 12.10 Bridge VLAN managed objects

The following managed objects define the semantics of the management operations that can be performed upon the VLAN aspects of a Bridge:

a) The Bridge VLAN Configuration managed object (12.10.1);
b) The VLAN Configuration managed object (12.10.2);
c) The VLAN Learning Constraints managed object (12.10.3).

### 12.10.1 Bridge VLAN Configuration managed object

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

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

a) Read Bridge VLAN Configuration (12.10.1.1);
b) Configure PVID values (12.10.1.2);
c) Configure Acceptable Frame Types parameters (12.10.1.3);
d) Configure Enable Ingress Filtering parameters (12.10.1.4);
e) Reset VLAN Bridge (12.10.1.5);
f) Notify VLAN registration failure (12.10.1.6).

#### 12.10.1.1 Read Bridge VLAN Configuration

##### 12.10.1.1.1 Purpose

To obtain general VLAN information from a Bridge.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 12.10.1.1.2 Inputs

None.

#### 12.10.1.1.3 Outputs

a) The 802.1Q VLAN Version number. Reported as “1” by devices that implement VLAN functionality according to this edition of the standard;
b) The optional VLAN features supported by the implementation:
   1) The maximum number of VLANs supported;
   2) Whether the implementation supports the ability to override the default PVID setting, and its egress status (VLAN-tagged or untagged) on each Port.
c) For each Port:
   1) the Port number;
   2) the PVID value (8.4.4) currently assigned to that Port;
   3) the state of the Acceptable Frame Types parameter (8.4.3). The permissible values for this parameter are:
      i) Admit only VLAN-tagged frames;
      ii) Admit all frames.
   4) the state of the Enable Ingress Filtering parameter (8.4.5); Enabled or Disabled.

### 12.10.1.2 Configure PVID values

#### 12.10.1.2.1 Purpose

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

#### 12.10.1.2.2 Inputs

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

#### 12.10.1.2.3 Outputs

None.

### 12.10.1.3 Configure Acceptable Frame Types parameters

#### 12.10.1.3.1 Purpose

To configure the Acceptable Frame Types parameter (8.4.3) associated with one or more Ports.

#### 12.10.1.3.2 Inputs

a) For each Port to be configured, a Port number and the value of the Acceptable Frame Types parameter to be associated with that Port. The permissible values of this parameter are (as defined in 8.4.3):
   1) Admit only VLAN-tagged frames;
   2) Admit all frames.

#### 12.10.1.3.3 Outputs

None.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 125 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 12.10.1.4 Configure Enable Ingress Filtering parameters

##### 12.10.1.4.1 Purpose

To configure the Enable Ingress Filtering parameter(s) (8.4.5) associated with one or more Ports.

##### 12.10.1.4.2 Inputs

a) For each Port to be configured, a Port number and the value of the Enable Ingress Filtering parameter to be associated with that Port. The permissible values for the parameter are
   1) Enabled;
   2) Disabled.

##### 12.10.1.4.3 Outputs

None.

#### 12.10.1.5 Reset VLAN Bridge

##### 12.10.1.5.1 Purpose

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

a) Deletes all VLAN Configuration managed objects;
b) Resets the PVID associated with each Bridge Port to the Default PVID value (Table 9-2);
c) Resets the Acceptable Frame Types parameter value associated with each Port to the default value (8.4.3).

##### 12.10.1.5.2 Inputs

None.

##### 12.10.1.5.3 Outputs

None.

#### 12.10.1.6 Notify VLAN registration failure

##### 12.10.1.6.1 Purpose

To notify a manager that GVRP (11.2.3) has failed to register a given VLAN owing to lack of resources in the Filtering Database for the creation of a Dynamic VLAN Registration Entry (8.11.5).

##### 12.10.1.6.2 Inputs

None.

##### 12.10.1.6.3 Outputs

a) The VID of the VLAN that GVRP failed to register;
b) The Port number of the Port on which the registration request was received.


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### 12.10.2 VLAN Configuration managed object

The VLAN Configuration object models operations that modify, or enquire about, the configuration of a particular VLAN within a Bridge. There are multiple VLAN Configuration objects per Bridge; only one such object can exist for a given VLAN ID.

The management operations that can be performed on the VLAN Configuration are:

a) Read VLAN Configuration (12.10.2.1);
b) Create VLAN Configuration (12.10.2.2);
c) Delete VLAN Configuration (12.10.2.3);

#### 12.10.2.1 Read VLAN Configuration

##### 12.10.2.1.1 Purpose

To obtain general information regarding a specific VLAN Configuration.

##### 12.10.2.1.2 Inputs

a) VLAN Identifier: a 12-bit VID.

##### 12.10.2.1.3 Outputs

a) VLAN Name: A text string of up to 32 characters of locally determined significance;
b) List of Untagged Ports: The set of Port numbers for which this VLAN ID is a member of the Untagged set (8.11.9) for that Port;
c) List of Egress Ports: The set of Port numbers for which this VLAN ID is a member of the Member set (8.11.9) for that Port.

NOTE—The values of the Member set and the Untagged set are determined by the values held in VLAN Registration Entries in the Filtering Database (8.11.2, 8.11.5, and 8.11.9).

#### 12.10.2.2 Create VLAN Configuration

##### 12.10.2.2.1 Purpose

To create or update a VLAN Configuration managed object.

##### 12.10.2.2.2 Inputs

a) VLAN Identifier: a 12-bit VID;
b) VLAN Name: A text string of up to 32 characters of locally determined significance.

NOTE—Static configuration of the Member set and the Untagged set is achieved by means of the management operations for manipulation of VLAN Registration Entries (12.7.5).

##### 12.10.2.2.3 Outputs

None.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 127 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 12.10.2.3 Delete VLAN Configuration

##### 12.10.2.3.1 Purpose

To delete a VLAN Configuration managed object.

##### 12.10.2.3.2 Inputs

a) VLAN Identifier: a 12-bit VID;

##### 12.10.2.3.3 Outputs

None.

### 12.10.3 The VLAN Learning Constraints managed object

The VLAN Learning Constraints managed object models operations that modify, or enquire about, the set of VLAN Learning Constraints (8.11.7.2) and VID to FID allocations (8.11.7.1) that apply to the operation of the Learning Process and the Filtering Database. There is a single VLAN Learning Constraints managed object per Bridge. The object is modeled as a pair of fixed-length tables, as follows:

a) A Learning Constraint table in which each table entry either defines a single Learning Constraint or is undefined. For some of the operations that can be performed upon the table, an *entry index* is used; this identifies the number of the entry in the table, where index number 1 is the first, and N is the last (where the table contains N entries).

NOTE—The number of Learning Constraint table entries supported is an implementation option. This standard does not provide any distribution mechanism to ensure that the same set of constraints is configured in all Bridges; individual Bridges can be configured by use of the management operations defined in this subclause (for example, via the use of SNMP operating upon a VLAN Bridge MIB), but there is no in-built consistency checking to ensure that all Bridges have been provided with the same constraint information. Hence, any such consistency checking is the responsibility of the network administrator and the management applications employed in the LAN.

b) A VID to FID allocation table (8.11.7.1) with an entry per VID supported by the implementation. Each table entry indicates, for that VID, that there is currently
   1) No allocation defined; or
   2) A fixed allocation to FID X; or
   3) A dynamic allocation to FID X.

The management operations that can be performed on the VLAN Learning Constraints managed object are

c) Read VLAN Learning Constraints (12.10.3.1);
d) Read VLAN Learning Constraints for VID (12.10.3.2);
e) Set VLAN Learning Constraint (12.10.3.3);
f) Delete VLAN Learning Constraint (12.10.3.4);
g) Read VID to FID allocations (12.10.3.5);
h) Read FID allocation for VID (12.10.3.6);
i) Read VIDs allocated to FID (12.10.3.7);
j) Set VID to FID allocation (12.10.3.8);
k) Delete VID to FID allocation (12.10.3.9);
l) Notify Learning Constraint Violation (12.10.3.10);


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### 12.10.3.1 Read VLAN Learning Constraints

##### 12.10.3.1.1 Purpose

To read the contents of a range of one or more entries in the VLAN Learning Constraints table.

##### 12.10.3.1.2 Inputs

a) First Entry—Entry Index of first entry to be read;
b) Last Entry—Entry Index of last entry to be read.

##### 12.10.3.1.3 Outputs

a) List of Entries—for each entry that was read:
    1) The Entry Index;
    2) The type of the Learning Constraint: Undefined, S or I;
    3) The value of the Learning Constraint, which is one of:
        i) Undefined, indicating an empty element in the table;
        ii) An S Constraint value, consisting of a pair of VIDs;
        iii) An I Constraint value, consisting of a VID and an Independent Set Identifier.

NOTE—Where this operation is implemented using a remote management protocol, PDU size constraints may restrict the number of entries that are actually read to fewer than was requested in the input parameters. In such cases, retrieving the remainder of the desired entry range can be achieved by repeating the operation with a modified entry range specification.

#### 12.10.3.2 Read VLAN Learning Constraints for VID

##### 12.10.3.2.1 Purpose

To read all the VLAN Learning Constraints for a given VID.

##### 12.10.3.2.2 Inputs

a) VID—The VLAN Identifier to which the read operation applies.

##### 12.10.3.2.3 Outputs

a) All learning constraint values that identify the VID requested. Each value returned is either
    1) An S Constraint value, consisting of a pair of VIDs; or
    2) An I Constraint value, consisting of a VID and an Independent Set Identifier.

#### 12.10.3.3 Set VLAN Learning Constraint

##### 12.10.3.3.1 Purpose

To modify the contents of one of the entries in the VLAN Learning Constraints table.

##### 12.10.3.3.2 Inputs

a) Entry Index—Entry index of the entry to be set;
b) The type of the Learning Constraint: S or I;
c) The value of the Learning Constraint, which is either:
        i) An S Constraint value, consisting of a pair of VIDs; or

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 129 of 211

ii) An I Constraint value, consisting of a VID and an Independent Set Identifier.

#### 12.10.3.3.3 Outputs

a) Operation status. This takes one of the following values:
   1) Operation rejected due to inconsistent learning constraint specification (8.11.7.3)—The Set operation requested setting a constraint that is inconsistent with another constraint already defined in the constraint table. The operation returns the value of the constraint concerned; or
   2) Operation rejected due to inconsistent fixed VID to FID allocation (8.11.7.3)—The Set operation requested setting a constraint that is inconsistent with a fixed VID to FID allocation already defined in the allocation table. The operation returns the value of the fixed allocation concerned; or
   3) Operation rejected due to entry index exceeding the maximum index supported by the constraint table; or
   4) Operation accepted.

### 12.10.3.4 Delete VLAN Learning Constraint

#### 12.10.3.4.1 Purpose

To remove one of the entries in the VLAN Learning Constraints table. This operation has the effect of setting the value of the specified table entry to “Undefined.”

#### 12.10.3.4.2 Inputs

a) Entry Index—Entry index of the entry to be deleted.

#### 12.10.3.4.3 Outputs

a) Operation status. This takes one of the following values:
   1) Operation rejected due to entry index exceeding the maximum index supported by the constraint table; or
   2) Operation accepted.

### 12.10.3.5 Read VID to FID allocations

#### 12.10.3.5.1 Purpose

To read the contents of a range of one or more entries in the VID to FID allocation table.

#### 12.10.3.5.2 Inputs

a) First Entry—VID of first entry to be read;
b) Last Entry—VID of last entry to be read.

#### 12.10.3.5.3 Outputs

a) List of Entries—For each entry that was read:
   1) VID—The VLAN Identifier for this entry;
   2) Allocation Type—The type of the allocation: Undefined, Fixed or Dynamic;
   3) FID—The FID to which the VID is allocated (if not of type Undefined).

NOTE—Where this operation is implemented using a remote management protocol, PDU size constraints may restrict the number of entries that are actually read to fewer than was requested in the input parameters. In such cases, retrieving


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

the remainder of the desired entry range can be achieved by repeating the operation with a modified entry range specification.

#### 12.10.3.6 Read FID allocation for VID

##### 12.10.3.6.1 Purpose

To read the FID to which a specified VID is currently allocated.

##### 12.10.3.6.2 Inputs

a) VID—The VLAN Identifier to which the read operation applies.

##### 12.10.3.6.3 Outputs

a) VID—the VLAN Identifier to which the read operation applies;
b) Allocation Type—the type of the allocation: Undefined, Fixed or Dynamic;
c) FID—the FID to which the VID is allocated (if not of type Undefined).

#### 12.10.3.7 Read VIDs allocated to FID

##### 12.10.3.7.1 Purpose

To read all the VIDs currently allocated to a given FID.

##### 12.10.3.7.2 Inputs

a) FID—the Filtering Identifier to which the read operation applies.

##### 12.10.3.7.3 Outputs

a) FID—the Filtering Identifier to which the read operation applies
b) Allocation List—a list of allocations for this FID. For each element in the list:
    1) Allocation Type—the type of the allocation: Fixed or Dynamic;
    2) VID—the VID that is allocated.

#### 12.10.3.8 Set VID to FID allocation

##### 12.10.3.8.1 Purpose

To establish a fixed allocation of a VID to an FID.

##### 12.10.3.8.2 Inputs

a) VID—the VID of the entry to be set;
b) FID—the FID to which the VID is to be allocated.

##### 12.10.3.8.3 Outputs

a) Operation status. This takes one of the following values:
    1) Operation rejected due to inconsistent learning constraint specification (8.11.7.3)—The Set operation requested setting a fixed allocation that is inconsistent with a VLAN Learning Constraint. The operation returns the value of the VLAN Learning Constraint concerned; or

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 131 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

2) Operation rejected due to VID exceeding the maximum VID supported by the allocation table; or
3) Operation rejected due to FID exceeding the maximum ID supported by the implementation; or
4) Operation accepted.

#### 12.10.3.9 Delete VID to FID allocation

##### 12.10.3.9.1 Purpose

To remove a fixed VID to FID allocation from the VID to FID allocation table. This operation has the effect of setting the value of the specified table entry to "Undefined."

NOTE—If the VID concerned represents a currently active VLAN, then removal of a fixed allocation may result in the "Undefined" value in the table immediately being replaced by a dynamic allocation to an FID.

##### 12.10.3.9.2 Inputs

a) VID—VID of the allocation to be deleted.

##### 12.10.3.9.3 Outputs

a) Operation status. This takes one of the following values:
   1) Operation rejected due to VID exceeding the maximum value supported by the allocation table; or
   2) Operation accepted.

#### 12.10.3.10 Notify Learning Constraint Violation

##### 12.10.3.10.1 Purpose

To alert the Manager to the existence of a Learning Constraint violation (8.11.7.3). This is an unsolicited notification from the management entity of the Bridge, issued upon detection of the constraint violation.

NOTE—As indicated in 8.11.7.3, a single change in configuration, such as the registration of a new VID by GVRP or the addition of a new learning constraint, can give rise to more than one violation being notified, depending upon the set of learning constraints currently configured in the Bridge.

##### 12.10.3.10.2 Inputs

a) None.

##### 12.10.3.10.3 Outputs

a) Violation Type/Argument—one of
   1) Shared VLAN Learning not supported. The argument returned indicates the VIDs of a pair of active VLANs for which an S constraint exists.
   2) Independent VLAN Learning not supported. The argument returned indicates the VIDs of a pair of active VLANs for which I constraints exist that contain the same independent set identifier.
   3) Required FID range not supported. The argument returned indicates
      i) The VID that the Bridge is unable to allocate to an FID;
      ii) The maximum number of FIDs supported by the Bridge.

The violation type *Required FID range not supported* is detected only by IVL or IVL/SVL Bridges that support fewer than 4094 FIDs.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# Annex A

(normative)

# PICS proforma[1]

## A.1 Introduction

The supplier of a protocol implementation which is claimed to conform to this standard shall complete the following Protocol Implementation Conformance Statement (PICS) proforma.

A completed PICS proforma is the PICS for the implementation in question. The PICS is a statement of which capabilities and options of the protocol have been implemented. The PICS can have a number of uses, including use

a) By the protocol implementor, as a checklist to reduce the risk of failure to conform to the standard through oversight;
b) By the supplier and acquirer—or potential acquirer—of the implementation, as a detailed indication of the capabilities of the implementation, stated relative to the common basis for understanding provided by the standard PICS proforma;
c) By the user—or potential user—of the implementation, as a basis for initially checking the possibility of interworking with another implementation (note that, while interworking can never be guaranteed, failure to interwork can often be predicted from incompatible PICSs);
d) By a protocol tester, as the basis for selecting appropriate tests against which to assess the claim for conformance of the implementation.

## A.2 Abbreviations and special symbols

### A.2.1 Status symbols

M mandatory
O optional
*O.n* optional, but support of at least one of the group of options labelled by the same numeral n is required
X prohibited
pred: conditional-item symbol, including predicate identification: see A.3.4
¬ logical negation, applied to a conditional item's predicate

### A.2.2 General abbreviations

N/A not applicable
PICS Protocol Implementation Conformance Statement

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

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 133 of 211

## A.3 Instructions for completing the PICS proforma

### A.3.1 General structure of the PICS proforma

The first part of the PICS proforma, implementation identification and protocol summary, is to be completed as indicated with the information necessary to identify fully both the supplier and the implementation.

The main part of the PICS proforma is a fixed-format questionnaire, divided into several subclauses, each containing a number of individual items. Answers to the questionnaire items are to be provided in the rightmost column, either by simply marking an answer to indicate a restricted choice (usually Yes or No), or by entering a value or a set or range of values. (Note that there are some items where two or more choices from a set of possible answers can apply; all relevant choices are to be marked.)

Each item is identified by an item reference in the first column. The second column contains the question to be answered; the third column records the status of the item—whether support is mandatory, optional, or conditional: see also A.3.4 below. The fourth column contains the reference or references to the material that specifies the item in the main body of this standard, and the fifth column provides the space for the answers.

A supplier may also provide (or be required to provide) further information, categorized as either Additional Information or Exception Information. When present, each kind of further information is to be provided in a further subclause of items labelled *Ai* or *Xi*, respectively, for cross-referencing purposes, where *i* is any unambiguous identification for the item (e.g., simply a numeral). There are no other restrictions on its format and presentation.

A completed PICS proforma, including any Additional Information and Exception Information, is the Protocol Implementation Conformation Statement for the implementation in question.

NOTE—Where an implementation is capable of being configured in more than one way, a single PICS may be able to describe all such configurations. However, the supplier has the choice of providing more than one PICS, each covering some subset of the implementation's configuration capabilities, in case that makes for easier and clearer presentation of the information.

### A.3.2 Additional information

Items of Additional Information allow a supplier to provide further information intended to assist the interpretation of the PICS. It is not intended or expected that a large quantity will be supplied, and a PICS can be considered complete without any such information. Examples might be an outline of the ways in which a (single) implementation can be set up to operate in a variety of environments and configurations, or information about aspects of the implementation that are outside the scope of this standard but that have a bearing upon the answers to some items.

References to items of Additional Information may be entered next to any answer in the questionnaire, and may be included in items of Exception Information.

### A.3.3 Exception information

It may occasionally happen that a supplier will wish to answer an item with mandatory status (after any conditions have been applied) in a way that conflicts with the indicated requirement. No pre-printed answer will be found in the Support column for this: instead, the supplier shall write the missing answer into the Support column, together with an *Xi* reference to an item of Exception Information, and shall provide the appropriate rationale in the Exception item itself.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

An implementation for which an Exception item is required in this way does not conform to this standard.

NOTE—A possible reason for the situation described above is that a defect in this standard has been reported, a correction for which is expected to change the requirement not met by the implementation.

## A.3.4 Conditional status

### A.3.4.1 Conditional items

The PICS proforma contains a number of conditional items. These are items for which both the applicability of the item itself, and its status if it does apply—mandatory or optional—are dependent upon whether or not certain other items are supported.

Where a group of items is subject to the same condition for applicability, a separate preliminary question about the condition appears at the head of the group, with an instruction to skip to a later point in the questionnaire if the “Not Applicable” answer is selected. Otherwise, individual conditional items are indicated by a conditional symbol in the Status column.

A conditional symbol is of the form “**pred:** S” where **pred** is a predicate as described in A.3.4.2 below, and S is a status symbol, M or 0.

If the value of the predicate is true (see A.3.4.2), the conditional item is applicable, and its status is indicated by the status symbol following the predicate: the answer column is to be marked in the usual way. If the value of the predicate is false, the “Not Applicable” (N/A) answer is to be marked.

### A.3.4.2 Predicates

A predicate is one of the following:

a) An item-reference for an item in the PICS proforma: the value of the predicate is true if the item is marked as supported, and is false otherwise;

b) A predicate-name, for a predicate defined as a boolean expression constructed by combining item-references using the boolean operator OR: the value of the predicate is true if one or more of the items is marked as supported;

c) The logical negation symbol “¬” prefixed to an item-reference or predicate-name: the value of the predicate is true if the value of the predicate formed by omitting the “¬” symbol is false, and vice versa.

Each item whose reference is used in a predicate or predicate definition, or in a preliminary question for grouped conditional items, is indicated by an asterisk in the Item column.

### A.3.4.3 References to the text of ISO/IEC 15802-3

Many of the tables in the PICS Proforma refer to the text of ISO/IEC 15802-3 (ANSI/IEEE Std 802.1D). A short form reference, of the form {D}X, is used in the “References” columns of these tables to denote references to clauses, subclauses or tables in ISO/IEC 15802-3, where X is the clause, subclause or table identifier.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 135 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.4 PICS proforma for IEEE Std 802.1Q-1998

### A.4.1 Implementation identification

| | |
|---|---|
| Supplier | |
| Contact point for queries about the PICS | |
| Implementation Name(s) and Version(s) | |
| Other information necessary for full identification - e.g., name(s) and version(s) of machines and/or operating system names | |

NOTE 1—Only the first three items are required for all implementations; other information may be completed as appropriate in meeting the requirement for full identification.
NOTE 2—The terms Name and Version should be interpreted appropriately to correspond with a supplier's terminology (e.g., Type, Series, Model).

### A.4.2 Protocol summary, IEEE Std 802.1Q-1998

| | |
|---|---|
| Identification of protocol specification | IEEE Std 802.1Q-1998, IEEE Standards for Local and Metropolitan Area Networks: Standard for Virtual Bridged Local Area Networks |
| Identification of amendments and corrigenda to the PICS proforma which have been completed as part of the PICS | Amd. : Corr. :<br>Amd. : Corr. : |
| Have any Exception items been required? (See A.3.3: the answer Yes means that the implementation does not conform to IEEE Std 802.1Q-1998) | No [ ] Yes [ ] |

| | |
|---|---|
| Date of Statement | |


Copyright © 1999 IEEE. All rights reserved.

## A.5 Major capabilities and options

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (1a)* | Communications Support<br><br>Which MAC types are supported on Bridge Ports, implemented in conformance with the relevant MAC standards? | | {D}6.5 | |
| (1a.1)* | CSMA/CD, IEEE Std 802.3 | O.1 | | Yes [ ] No [ ] |
| (1a.2)* | Token Bus, ISO/IEC 8802-4 | O.1 | | Yes [ ] No [ ] |
| (1a.3)* | Token Ring, ISO/IEC 8802-5 | O.1 | | Yes [ ] No [ ] |
| (1a.4)* | FDDI, ISO 9314-2 | O.1 | | Yes [ ] No [ ] |
| (1a.5)* | DQDB, ISO/IEC 8802-6 | O.1 | | Yes [ ] No [ ] |
| (1a.6)* | ISLAN, ISO/IEC 8802-9 | O.1 | | Yes [ ] No [ ] |
| (1a.7)* | ISLAN 16-T, IEEE 802.9a | O.1 | | Yes [ ] No [ ] |
| (1a.8)* | Demand Priority, ISO/IEC 8802-12 (IEEE Std 802.3 format) | O.1 | | Yes [ ] No [ ] |
| (1a.9)* | Demand Priority, ISO/IEC 8802-12 (ISO/IEC 8802-5 format) | O.1 | | Yes [ ] No [ ] |
| (1b) | Is LLC Type 1 supported on all Bridge Ports in conformance with ISO/IEC 8802-2? | M | 8.2, 8.3, 8.14, ISO/IEC 8802-2 | Yes [ ] |
| (1c)* | Is Source-Routing Transparent Bridge operation supported on any of the Bridge Ports? (If support is claimed, the PICS proforma detailed in ISO/IEC 15802-3, Annex D, shall also be completed). | O | {D}Annex C | Yes [ ] No [ ] |
| (2) | Relay and filtering of frames (A.6) | M | 8.5, 8.9, 8.6, 8.7, 8.8 | Yes [ ] |
| (2a) | Does the Bridge support Basic Filtering Services? | M | {D}6.6.5, 8.7.2 | Yes [ ] |
| (2b)* | Does the Bridge support Extended Filtering Services?<br><br>If item (2b) is not supported, mark “N/A” and continue at (2e). | O | {D}6.6.5, 8.7.2 | Yes [ ] No [ ]<br><br>N/A[ ] |
| (2c)* | Does the Bridge support dynamic Group forwarding and filtering behavior? | 2b:M | {D}6.6.5 | Yes [ ] No [ ] |
| (2d) | Does the Bridge support the ability for static filtering information for individual MAC Addresses to specify a subset of Ports for which forwarding or filtering decisions are taken on the basis of dynamic filtering information? | 2b:O | {D}6.6.5 | Yes [ ] No [ ] |
| (2e)* | Does the Bridge support expedited traffic classes on any of its Ports? | O | 8.1.2, 8.7.3 | Yes [ ] No [ ] |
| (4)* | Does the Bridge support management of the priority of relayed frames? | O | {D}6.5, 8.5.1, 8.7.3, 8.7.5, Table 8-1, Table 8-2, Table 8-3 | Yes [ ] No [ ] |
| (5) | Maintenance of filtering information (A.7) | M | 8.10, 8.11 | Yes [ ] |
| (7a) | Can the Filtering Database be read by management? | O | 8.11 | Yes [ ] No [ ] |

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 137 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.5 Major capabilities and options *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (7c)* | Can Static Filtering Entries be created and deleted? | O | 8.11.1 | Yes [ ] No [ ] |
| (7g) | Can Static Filtering Entries be created and deleted in the Permanent Database? | O | 8.11.10 | Yes [ ] No [ ] |
| (7h) | Can Static Filtering Entries be created for a given MAC Address specification with a distinct Port Map for each inbound Port? | O | 8.11.1 | Yes [ ] No [ ] |
| (7i) | Can Group Registration Entries be dynamically created, updated and deleted by GMRP? | 2c:M | 8.11.4, {D}10 | Yes [ ]<br>N/A [ ] |
| (10) | Addressing (A.8) | M | 8.14 | Yes [ ] |
| (9a)* | Can the Bridge be configured to use 48-bit Universal Addresses? | O.3 | 8.14 | Yes [ ] No [ ] |
| (9b)* | Can the Bridge be configured to use 48-bit Local Addresses? | O.3 | 8.14 | Yes [ ] No [ ] |
| (13)* | Spanning Tree algorithm and protocol (A.9) | M | {D}8, {D}9 | Yes [ ] |
| (16)* | Does the Bridge support management of the Spanning Tree topology? | O | {D}8.2 | Yes [ ] No [ ] |
| (17)* | Does the Bridge support management of the protocol timers? | O | {D}8.10 | Yes [ ] No [ ] |
| (19)* | VLAN Bridge Management Operations | O | 12 | Yes [ ] No [ ] |
| (20a)* | Are the Bridge Management Operations supported via a Remote Management Protocol? | 19:O.4 | {D}5 | Yes [ ] No [ ]<br>N/A [ ] |
| (20b)* | Are the Bridge Management Operations supported via a local management interface? | 19:O.4 | {D}5 | Yes [ ] No [ ]<br>N/A [ ] |
| (23a)* | Does the implementation support, on each Port, one or more of the permissible combinations of values for the Acceptable Frame Types parameter? | M | 5.1, 8.4.3 | Yes [ ] |
| (23a.1) | State which Ports support:<br>— Admit only VLAN-tagged frames;<br>— Admit all frames. | M | 5.1, 8.4.3 | Ports: _____<br>Ports: _____ |
| (23a.2) | On Ports that support both values, is the parameter configurable via management? | M | 5.1, 8.4.3, 12.10 | Yes [ ] N/A [ ] |
| (23b) | Does the implementation support the ability to insert tag headers into, modify tag headers in, and remove tag headers from relayed frames, as required by the capabilities of each Bridge Port? | M | 5.1, 7.1, 9 | Yes [ ] |
| (23c) | Does the implementation support the ability to perform automatic configuration and management of VLAN topology information by means of GVRP on all Ports? | M | 5.1, 11 | Yes [ ] |


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.5 Major capabilities and options *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (23d) | Does the implementation support the ability for the Filtering Database to contain static and dynamic configuration information for at least one VLAN, by means of Static and Dynamic VLAN Registration Entries? | M | 5.1, 8.11 | Yes [ ] |
| (23d.1) | State the maximum number of VLANs supported by the implementation. | M | 5.1, 8.11, 9.3.2.3 | _____ VLANs |
| (23d.2) | State the range of VID values supported by the implementation. | M | 8.11, 9.3.2.3 | 0 through _____ |
| (23e)* | VLAN Learning support | | 5.1, 8.11.3, 8.11.7, 8.11.8 | |
| (23e.1) | Does the implementation support at least one FID? | M | | Yes [ ] |
| (23e.2) | Can the implementation allocate at least one VID to each FID supported? | M | | Yes [ ] |
| (23e.4) | State the maximum number of FIDs that can be supported by the implementation. | M | 8.11.7 | ____ FIDs |
| (23e.5) | State the maximum number of VIDs that can be allocated to each FID. | M | 8.11.7 | ____ VIDs |
| (23e.6) | Does the implementation support configuration of VLAN Learning Constraints via management? | O | 5.2, 8.11.7, 12.10.3 | Yes [ ] No [ ] |
| (23e.7) | State the number of VLAN Learning Constraints that can be configured in the implementation. | 23e.6:M | 5.2, 8.11.7, 12.10.3 | ____ Constraints |
| (23e.8) | Does the implementation support configuration of VID to FID allocations via management? | O | 5.2, 8.11.7.1, 12.10.3 | Yes [ ] No [ ] |
| (23e.9) | Does the implementation take account of the allocation of VIDs to FIDs when making forwarding decisions relative to group MAC Addresses? | O | 8.11.8 | Yes [ ] No [ ] |
| (23f) | On Ports that support untagged and priority-tagged frames, does the implementation support: | | 5.1, 8.4.4, 8.11.9, 12.10 | |
| (23f.1) | — A PVID value? | M | | Yes [ ] N/A [ ] |
| (23f.2) | — The ability to configure one VLAN whose Untagged set includes that Port? | M | | Yes [ ] N/A [ ] |
| (23f.3) | — Configuration of the PVID value via management operations? | M | | Yes [ ] N/A [ ] |
| (23f.4) | — Configuration of Static Filtering Entries via management operations? | M | | Yes [ ] N/A [ ] |
| (23f.5) | — The ability to configure more than one VLAN whose Untagged set includes that Port? | O | | Yes [ ] No [ ] N/A [ ] |
| (23g)* | Does the implementation support the ability to enable and disable Ingress Filtering? | O | 5.2, 8.4.5 | |

Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.5 Major capabilities and options *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (23h) | Does the implementation support VLAN management operations? | 19:O | 5.2, 12.10.2, 12.10.3 | Yes [ ] No [ ] |
| (23i) | Is the minimum tagged frame length that can be transmitted on IEEE Std 802.3 Ports less than 68 (but 64 or more) octets? | 1a.1:O | 7.2 | Yes [ ] No [ ] N/A [ ] |
| (23j)* | When transmitting untagged frames and the canonical_format_indicator parameter indicates that the mac_service_data_unit may contain embedded MAC Addresses in a format inappropriate to the destination MAC method, which of the following procedures is adopted by the Bridge: | | 7.1, 7.1.2.2 | |
| (23j.1) | Convert any embedded MAC Addresses in the mac_service_data_unit to the format appropriate to the destination MAC method. | O.7 | | Yes [ ] No [ ] |
| (23j.2) | Discard the frame without transmission on that Port. | O.7 | | Yes [ ] No [ ] |
| (23k) | Does the Bridge perform frame translations, where necessary, in accordance with the procedures described in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390? | TB:M | 7.1, 7.1.2.2 | Yes [ ] No [ ] N/A [ ] |

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

## A.6 Relay and filtering of frames

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (2f) | Are received frames with MAC method errors discarded? | M | {D}6.4, 8.5 | Yes [ ] |
| (2g) | Are correctly received frames submitted to the Learning Process? | M | 8.5 | Yes [ ] |
| (2h) | Are user data frames the only type of frame relayed? | M | 8.5 | Yes [ ] |
| (2i) | Are request with no response frames the only frames relayed? | M | 8.5 | Yes [ ] |
| (2j) | Are all frames addressed to the Bridge Protocol Entity submitted to it? | M | 8.5 | Yes [ ] |
| (2k) | Are user data frames the only type of frame transmitted? | M | 8.9 | Yes [ ] |
| (2l) | Are request with no response frames the only frames transmitted? | M | 8.9 | Yes [ ] |
| (2m) | Are relayed frames queued for transmission only under the conditions in 8.7.3? | M | 8.7.3, {D}8.4 | Yes [ ] |


Copyright © 1999 IEEE. All rights reserved.

## A.6 Relay and filtering of frames *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (2n) | Is the order of relayed frames preserved in accordance with the requirements of the forwarding process? | M | 8.7.3, 8.1.1 | Yes [ ] |
| (2o) | Is a relayed frame submitted to a MAC Entity for transmission only once? | M | 8.7.4, {D}6.3.4 | Yes [ ] |
| (2p) | Is a maximum bridge transit delay enforced for relayed frames? | M | 8.7.3 | Yes [ ] |
| (2q) | Are queued frames discarded if a Port leaves the Forwarding State? | M | 8.7.3 | Yes [ ] |
| (2r) | Is the user priority of relayed frames preserved where possible? | M | {D}6.4 | Yes [ ] |
| (2s) | Is the user priority set to the Default User Priority for the reception Port otherwise? | M | {D}6.4 | Yes [ ] |
| (2t) | Is the user priority regenerated by means of the User Priority Regeneration Table? | M | 8.5.1, Table 8-1 | Yes [ ] |
| (2u) | Is mapping of Regenerated User Priority to Traffic Class performed by means of the Traffic Class Table? | M | 8.7.3, Table 8-2 | Yes [ ] |
| (2v) | Is the access priority derived from the Regenerated User Priority as defined by the values in Table 8-3 for each outbound MAC method supported by the Bridge? | M | 8.7.5, Table 8-3 | Yes [ ] |
| (2w) | Does the Bridge generate an M_UNITDATA.indication primitive on receipt of a valid frame transmitted by the Bridge Port's local MAC entity? | MS1:X | {D}6.5.4, ISO 9314-2 | No [ ]<br>N/A [ ] |
| (2x) | Is only Asynchronous service used? | MS1:M | ISO 9314-2, 8.1.4 | Yes [ ]<br>N/A [ ] |
| (2y) | On receiving a frame from an FDDI ring for forwarding, does the bridge set the C indicator? | MS1:O | {D}6.5.4, ISO 9314-2, 7.3.8 | Yes [ ] No [ ]<br>N/A [ ] |
| (2z) | On receiving a frame from an FDDI ring for forwarding, does the bridge leave the C indicator unaltered? | MS1:O | {D}6.5.4, ISO 9314-2, 7.3.8 | Yes [ ] No [ ]<br>N/A [ ] |
| | If item 4 is not supported, mark "N/A" and continue at item (4d). | | | N/A [ ] |
| (4a)* | Can the Default User Priority parameter for each Port be set to any value in the range 0 through 7? | 4:O.6 | {D}6.4 | Yes [ ] No [ ] |
| (4b)* | Can the entries in the User Priority Regeneration Table for each Port be set to the full range of values shown in Table 8-1? | 4:O.6 | 8.5.1, Table 8-1 | Yes [ ] No [ ] |
| (4c)* | Can the entries in the Traffic Class Table for each Port be set to the full range of values shown in Table 8-2? | MS2:O | 8.7.3, Table 8-2 | Yes [ ] No [ ]<br>N/A [ ] |

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 141 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.6 Relay and filtering of frames *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| | If item 4 is supported, mark "N/A" and continue at item (4g) | | | N/A [ ] |
| (4d) | Does the Bridge support the recommended default value of the Default User Priority parameter for each Port? | ¬ 4:M | {D}6.4 | Yes [ ] |
| (4e) | Does the Bridge support the recommended default mappings between received user priority and Regenerated User Priority for each Port as defined in Table 8-1? | ¬ 4:M | 8.5.1, Table 8-1 | Yes [ ] |
| (4f) | Does the Bridge support the recommended default user_priority to traffic class mappings shown in Table 8-2 for each Port? | MS3:M | 8.7.3, Table 8-2 | Yes [ ]<br>N/A [ ] |
| (4g) | Is the Bridge able to use any values other than those shown in Table 8-3 when determining the access priority for the MAC methods shown? | X | 8.7.5, Table 8-3 | No [ ] |

Predicates:
MS1 = 1a.4 AND NOT (1a.1 OR 1a.2 OR 1a.3 OR 1a.5 OR 1a.6 OR 1a.7 OR 1a.8)
MS2 = 2e AND 4
MS3 = 2e AND NOT 4

## A.7 Maintenance of filtering entries in the Filtering Database

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (5a) | Are Dynamic Filtering Entries created and updated if and only if the Port State permits? | M | 8.10, 8.11.3, {D}8.4 | Yes [ ] |
| (5b) | Are Dynamic Filtering Entries created on receipt of frames with a group source address? | X | 8.10, 8.11.3 | No [ ] |
| (5c) | Does the Filtering Database support Static Filtering Entries? | M | 8.11.1 | Yes [ ] |
| (5d) | Can a Dynamic Filtering Entry be created that conflicts with an existing Static Filtering Entry? | X | 8.10, 8.11, 8.11.1, 8.11.3 | No [ ] |
| (5e) | Does the Filtering Database support Dynamic Filtering Entries? | M | 8.11.3 | Yes [ ] |
| (5f) | Does the creation of a Static Filtering Entry remove any conflicting information in a Dynamic Filtering Entry for the same address? | M | 8.11.1, 8.11.3 | Yes [ ] |
| (5g) | Does each Static Filtering Entry specify a MAC Address specification and a Port Map? | M | 8.11.1 | Yes [ ] |
| (5h) | Are Dynamic Filtering Entries removed from the Filtering Database if not updated for the Ageing Time period? | M | 8.11.3 | Yes [ ] |
| (5i) | Does each Dynamic Filtering Entry specify a MAC Address specification and a Port Map? | M | 8.11.3 | Yes [ ] |


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.7 Maintenance of filtering entries in the Filtering Database *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (5j) | Is the Filtering Database initialized with the entries contained in the Permanent Database? | M | 8.11.10 | Yes [ ] |
| | If item (2c) is not supported, mark N/A and continue at item (6a). | | | N/A [ ] |
| (5k) | Does each Group Registration Entry specify a MAC Address specification and a Port Map? | 2c:M | 8.11.4 | Yes [ ] |
| (5l) | Can the MAC Address specification in Group Registration Entries represent All Groups, All Unregistered Groups, or a specific group MAC Address? | 2c:M | 8.11.4 | Yes [ ] |
| (5m) | Are Group Registration Entries created, updated and removed from the Filtering Database in accordance with the specification of GMRP? | 2c:M | 8.11.4, {D}10 | Yes [ ] |
| (5n) | Are Group Registration Entries created, updated and removed from the Filtering Database by any means other than via the operation of GMRP? | 2c:X | 8.11.4, {D}10 | No [ ] |
| (6a) | State the Filtering Database Size. | M | 8.11 | ____ entries |
| (6b) | State the Permanent Database Size. | M | 8.11 | ____ entries |
| | If item (7c) is not supported, mark N/A and continue at item (8a). | | | N/A [ ] |
| (7d) | Can Static Filtering Entries be made for individual MAC Addresses? | **7c:**M | 8.11.1 | Yes [ ] |
| (7e) | Can Static Filtering Entries be made for group MAC Addresses? | **7c:**M | 8.11.1 | Yes [ ] |
| (7f) | Can a Static Filtering Entry be made for the broadcast MAC Address? | **7c:**M | 8.11.1 | Yes [ ] |
| (8a) | Can the Bridge be configured to use the default value of Ageing Time recommended in Table 8-4? | O | 8.11.3, Table 8-4 | Yes [ ] No [ ] |
| (8b) | Can the Bridge be configured to use any of the range of values of Ageing Time specified in Table 7-4? | O | 8.11.3, Table 8-4 | Yes [ ] No [ ] |

## A.8 Addressing

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (10a) | Does each Port have a separate MAC Address? | M | 8.14.2 | Yes [ ] |
| (10b) | Are all BPDUs transmitted to the same group address? | M | 8.14.3, {D}8.2 | Yes [ ] |
| | If item (9a) is not supported, mark N/A and continue at item (10d1). | | | N/A [ ] |

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 143 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.8 Addressing *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (10c) | Are all BPDUs transmitted to the Bridge Protocol Group Address when Universal Addresses are used? | **9a:**M | 8.14.3, {D}8.2 | Yes [ ] |
| (10d) | Is the source address of BPDUs the address of the transmitting Port? | **9a:**M | 8.14.3 | Yes [ ] |
| (10d1) | Is the LLC address of BPDUs the standard LLC address identified for the Spanning Tree Protocol? | M | 8.14.3, Table 8-9 | Yes [ ] |
| (10e) | Is the Bridge Address a Universal Address? | **M** | 8.14.5, {D}8.2 | Yes [ ]<br>N/A [ ] |
| (10f) | Are frames addressed to any of the Reserved Addresses relayed by the Bridge? | X | 8.14.6 | No [ ] |
| | If item (13) is not supported, mark N/A and continue at item (11c). | | | N/A [ ] |
| (11a) | Is Bridge Management accessible through each Port using the MAC Address of the Port and the LSAP assigned? | **13:**O | 8.14.4 | Yes [ ] No [ ] |
| (11b) | Is Bridge Management accessible through all Ports using the All LANs Bridge Management Group Address? | 13:O | 8.14.4 | Yes [ ] No [ ] |
| (11c) | Is the Bridge Address the Address of Port 1? | **9a:**O | 8.14.5 | Yes [ ] No [ ]<br>N/A [ ] |
| (11d)* | Are Group Addresses additional to the Reserved Addresses pre-configured in the Permanent Database? | O | 8.14.6 | Yes [ ] No [ ] |
| | If item (11d) is not supported, mark N/A and continue at item (12a). | | | N/A [ ] |
| (11e) | Can the additional pre-configured entries in the Filtering Database be deleted? | **11d:**O | 8.14.6 | Yes [ ] No [ ] |
| (12a) | Can a group MAC Address be assigned to identify the Bridge Protocol Entity? | **9b:**M | {D}8.2 | Yes [ ]<br>N/A [ ] |
| (12c) | Does each Port of the Bridge have a distinct identifier? | M | {D}8.2, {D}8.5.5.1 | Yes [ ] |

## A.9 Spanning Tree Algorithm

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (13a) | Are all the following Bridge Parameters maintained? | M | {D}8.5.3 | Yes [ ] |
| | Designated Root | | {D}8.5.3.1 | |
| | Root Cost | | {D}8.5.3.2 | |
| | Root Port | | {D}8.5.3.3 | |
| | Max Age | | {D}8.5.3.4 | |


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.9 Spanning Tree Algorithm *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| | Hello Time | | {D}8.5.3.5 | |
| | Forward Delay | | {D}8.5.3.6 | |
| | Bridge Identifier | | {D}8.5.3.7 | |
| | Bridge Max Age | | {D}8.5.3.8 | |
| | Bridge Hello Time | | {D}8.5.3.9 | |
| | Bridge Forward Delay | | {D}8.5.3.10 | |
| | Topology Change Detected | | {D}8.5.3.11 | |
| | Topology Change | | {D}8.5.3.12 | |
| | Topology Change Time | | {D}8.5.3.13 | |
| | Hold Time | | {D}8.5.3.14 | |
| (13b) | Are all the following Bridge Timers maintained? | M | {D}8.5.4 | Yes [ ] |
| | Hello Timer | | {D}8.5.4.1 | |
| | Topology Change Notification Timer | | {D}8.5.4.2 | |
| | Topology Change Timer | | {D}8.5.4.3 | |
| (13c) | Are all the following Port Parameters maintained for each Port? | M | {D}8.5.5 | Yes [ ] |
| | Port Identifier | | {D}8.5.5.1 | |
| | State | | {D}8.5.5.2, {D}8.4 | |
| | Path Cost | | {D}8.5.5.3 | |
| | Designated Root | | {D}8.5.5.4 | |
| | Designated Cost | | {D}8.5.5.5 | |
| | Designated Bridge | | {D}8.5.5.6 | |
| | Designated Port | | {D}8.5.5.7 | |
| | Topology Change Acknowledge | | {D}8.5.5.8 | |
| | Configuration Pending | | {D}8.5.5.9 | |
| | Change Detection Enabled | | {D}8.5.5.10 | |
| (13d) | Are all the following Timers maintained for each Port? | M | {D}8.5.6 | Yes [ ] |
| | Message Age Timer | | {D}8.5.6.1 | |
| | Forward Delay Timer | | {D}8.5.6.2 | |
| | Hold Timer | | {D}8.5.6.3 | |

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 145 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.9 Spanning Tree Algorithm *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (13e) | Are Protocol Parameters and Timers maintained, and BPDUs transmitted, as required on each of the following events? | M | {D}8.7, {D}8.9, {D}8.5.3, {D}8.5.4, {D}8.5.5, {D}8.5.6 | Yes [ ] |
| | Received Configuration BPDU | | {D}8.7.1 | |
| | Received Topology Change Notification BPDU | | {D}8.7.2 | |
| | Hello Timer Expiry | | {D}8.7.3 | |
| | Message Age Timer Expiry | | {D}8.7.4 | |
| | Forward Delay Timer Expiry | | {D}8.7.5 | |
| | Topology Change Notification Timer Expiry | | {D}8.7.6 | |
| | Topology Change Timer Expiry | | {D}8.7.7 | |
| | Hold Timer Expiry | | {D}8.7.8 | |
| (13f) | Do the following operations modify Protocol Parameters and Timers, and transmit BPDUs as required? | M | {D}8.8, {D}8.9, {D}8.5.3, {D}8.5.4, {D}8.5.5, {D}8.5.6 | Yes [ ] |
| | Initialization | | {D}8.8.1 | |
| | Enable Port | | {D}8.8.2 | |
| | Disable Port | | {D}8.8.3 | |
| | Set Bridge Priority | | {D}8.8.4 | |
| | Set Port Priority | | {D}8.8.5 | |
| | Set Path Cost | | {D}8.8.6 | |
| (13g) | Does the implementation support the ability to set the value of the Change Detection Enabled parameter to Disabled? | O | {D}8.5.5.10 | Yes [ ] No [ ] |
| (14a) | Does the Bridge underestimate the increment to the Message Age parameter in transmitted BPDUs? | X | {D}8.10.1 | No [ ] |
| (14b) | Does the Bridge underestimate Forward Delay? | X | {D}8.10.1 | No [ ] |
| (14c) | Does the Bridge overestimate the Hello Time interval? | X | {D}8.10.1 | No [ ] |
| (15a) | Does the Bridge use the specified value for Hold Time? | M | {D}8.10.2, {D}Table 8-3 | Yes [ ] |
| | If item (16) is not supported, mark N/A and continue at (17a). | | | N/A [ ] |
| (16a) | Can the relative priority of the Bridge be set? | **16:**M | {D}8.2, {D}8.5.3.7, {D}8.8.4 | Yes [ ] |


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.9 Spanning Tree Algorithm *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (16b) | Can the relative priority of the Ports be set? | **16:**M | {D}8.2, {D}8.5.5.1, {D}8.8.5 | Yes [ ] |
| (16c) | Can the path cost for each Port be set? | **16:**M | {D}8.2, {D}8.5.5.3, {D}8.8.6 | Yes [ ] |
| | If item (17) is not supported, mark N/A and continue at (18a). | | | N/A [ ] |
| (17a) | Can Bridge Max Age be set to any of the range of values specified? | **17:**M | {D}8.10.2, {D}8.5.3.8, {D}Table 8-3 | Yes [ ] |
| (17b) | Can Bridge Hello Time be set to any of the range of values specified? | **17:**M | {D}8.10.2, {D}8.5.3.9, {D}Table 8-3 | Yes [ ] |
| (17c) | Can Bridge Forward Delay be set to any of the range of values specified? | **17:**M | {D}8.10.2, {D}8.5.3.10, {D}Table 8-3 | Yes [ ] |
| (18a) | Do all BPDUs contain an integral number of octets? | M | {D}9.1.1 | Yes [ ] |
| (18b) | Are all the following BPDU parameter types encoded as specified? | M | {D}9.1.1, {D}9.2 | Yes [ ] |
| | Protocol Identifiers | | {D}9.2.1 | |
| | Protocol Version Identifiers | | {D}9.2.2 | |
| | BPDU Types | | {D}9.2.3 | |
| | Flags | | {D}9.2.4 | |
| | Bridge Identifiers | | {D}9.2.5 | |
| | Root Path Cost | | {D}9.2.6 | |
| | Port Identifiers | | {D}9.2.7 | |
| | Timer Values | | {D}9.2.8 | |
| (18c) | Do Configuration BPDUs have the format and parameters specified? | M | {D}9.3.1 | Yes [ ] |
| (18d) | Do Topology Change Notification BPDUs have the format and parameters specified? | M | {D}9.3.2 | Yes [ ] |
| (18e) | Are received BPDUs validated as specified? | M | {D}9.3.3 | Yes [ ] |

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyright © 1999 IEEE. All rights reserved.

## A.10 Bridge Management

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| | If item (19) is not supported, mark N/A and continue at (20c). | | | N/A [ ] |
| (19a) | Discover Bridge | **19:**M | 12.4.1.1 | Yes [ ] |
| (19b) | Read Bridge | **19:**M | 12.4.1.2 | Yes [ ] |
| (19c) | Set Bridge Name | **19:**M | 12.4.1.3 | Yes [ ] |
| (19d) | Reset Bridge | **19:**M | 12.4.1.4 | Yes [ ] |
| (19e) | Read Port | **19:**M | 12.4.2.1 | Yes [ ] |
| (19f) | Set Port Name | **19:**M | 12.4.2.2 | Yes [ ] |
| (19g) | Read Forwarding Port Counters | **19:**M | 12.6.1.1 | Yes [ ] |
| (19g.1) | Are the Forwarding Port Counters maintained per VLAN? | 19:O | | Yes [ ] No [ ] |
| (19g.2) | Does the implementation support the Discard on Error Details parameter? | 19:O | | Yes [ ] No[ ] |
| (19h) | Read Filtering Database | **19:**M | 12.7.1.1 | Yes [ ] |
| (19i) | Set Filtering Database Ageing Time | **19:**M | 12.7.1.2 | Yes [ ] |
| (19j) | Read Permanent Database | **19:**M | 12.7.6.1 | Yes [ ] |
| (19k) | Create Filtering Entry | **19:**M | 12.7.7.1 | Yes [ ] |
| (19l) | Delete Filtering Entry | **19:**M | 12.7.7.2 | Yes [ ] |
| (19m) | Read Filtering Entry | **19:**M | 12.7.7.3 | Yes [ ] |
| (19n) | Read Filtering Entry Range | **19:**M | 12.7.7.4 | Yes [ ] |
| (19o) | Read Bridge Protocol Parameters | **19:**M | 12.8.1.1 | Yes [ ] |
| (19p) | Set Bridge Protocol Parameters | **19:**M | 12.8.1.2 | Yes [ ] |
| (19q) | Read Port Parameters | **19:**M | 12.8.2.1 | Yes [ ] |
| (19r) | Force Port State | **19:**M | 12.8.2.2 | Yes [ ] |
| (19s) | Set Port Parameters | **19:**M | 12.8.2.3 | Yes [ ] |
| (19t) | Read Port Default User Priority | MS4:M | 12.6.2.1 | Yes [ ] N/A [ ] |
| (19u) | Set Port Default User Priority | MS4:M | 12.6.2.2 | Yes [ ] N/A [ ] |
| (19v) | Read Port User Priority Regeneration Table | MS5:M | 12.6.2.3 | Yes [ ] N/A [ ] |
| (19w) | Set Port User Priority Regeneration Table | MS5:M | 12.6.2.3 | Yes [ ] N/A [ ] |
| (19x) | Read Port Traffic Class Table | MS7:M | 12.6.3.1 | Yes [ ] N/A [ ] |
| (19y) | Set Port Traffic Class Table | MS7:M | 12.6.3.1 | Yes [ ] N/A [ ] |
| (19z) | Read Outbound Access Priority Table | MS6:M | 12.6.2.5 | Yes [ ] N/A [ ] |
| (19aa) | Read GARP Timers | MS8:M | 12.9.1.1 | Yes [ ] N/A [ ] |
| (19ab) | Set GARP Timers | MS8:M | 12.9.1.2 | Yes [ ] N/A [ ] |
| (19ac) | Read GARP Protocol Controls | MS8:M | 12.9.2.1 | Yes [ ] N/A [ ] |


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.10 Bridge Management *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (19ad) | Set GARP Protocol Controls | MS8:M | 12.9.2.2 | Yes [ ] N/A [ ] |
| (19ae) | Read GARP State | MS8:M | 12.9.3.1 | Yes [ ] N/A [ ] |
| (19af) | Read Bridge VLAN Configuration | 19:M | 12.10.1.1 | Yes [ ] N/A [ ] |
| (19ah) | Configure PVID values | 19:M | 12.10.1.2 | Yes [ ] N/A [ ] |
| (19ai) | Configure Acceptable Frame Types parameter | 23a.2:M | 12.10.1.3 | Yes [ ] N/A [ ] |
| (19aj) | Configure Enable Ingress Filtering parameters | 23g:M | 12.10.1.4 | Yes [ ] N/A [ ] |
| (19ak) | Reset Bridge VLAN Bridge. | 19:M | 12.10.1.5 | Yes [ ] N/A [ ] |
| (19al) | Notify VLAN Registration Failure | 19:M | 12.10.1.6 | Yes [ ] N/A [ ] |
| (19am) | Read VLAN Configuration | 19:M | 12.10.2.1 | Yes [ ] N/A [ ] |
| (19an) | Create VLAN Configuration | 19:M | 12.10.2.2 | Yes [ ] N/A [ ] |
| (19ao) | Delete VLAN Configuration | 19:M | 12.10.2.3 | Yes [ ] N/A [ ] |
| | If Item (23e.6) is not supported, mark N/A and continue at Item (19at). | | | N/A |
| (19ap) | Read VLAN Learning Constraints | 23e.6:M | 12.10.3.1 | Yes [ ] |
| (19aq) | Read VLAN Learning Constraints for VID | 23e.6:M | 12.10.3.2 | Yes [ ] |
| (19aq) | Set VLAN Learning Constraint | 23e.6:M | 12.10.3.3 | Yes [ ] |
| (19ar) | Delete VLAN Learning Constraint | 23e.6:M | 12.10.3.4 | Yes [ ] |
| (19as) | Notify Learning Constraint Violation | 23e.6:M | 12.10.3.10 | Yes [ ] |
| | If Item (23e.8) is not supported, mark N/A and continue at Item (20c). | | | N/A |
| (19at) | Read VID to FID allocations | 23e.8:M | 12.10.3.5 | Yes [ ] |
| | Read FID allocation for VID | 23e.8:M | 12.10.3.6 | Yes [ ] |
| | Read VIDs allocated to FID | 23e.8:M | 12.10.3.7 | Yes [ ] |
| | Set VID to FID allocation | 23e.8:M | 12.10.3.8 | Yes [ ] |
| | Delete VID to FID allocation | 23e.8:M | 12.10.3.9 | Yes [ ] |
| | If item (20a) is not supported, mark N/A and continue at (20e). | 23e.8:M | | N/A [ ] |
| (20c) | What Management Protocol standard(s) or specification(s) are supported? | **20a:**M | {D}5 | |
| (20d) | What standard(s) or specifications for Managed Objects and Encodings are supported? | **20a:**M | {D}5 | |

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyright © 1999 IEEE. All rights reserved.



Dell et al., Ex. 1008
Page 149 of 211

## A.10 Bridge Management *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| | If item (20b) is not supported, mark N/A and continue at A.11. | | | N/A [ ] |
| (20e) | What specification of the local management interface is supported? | **20b:**M | {D}5 | |

Predicates:
MS4=19 AND 4a
MS5=19 AND 4b
MS6=19 AND 4
MS7=19 AND 4c
MS8=19 AND 2b

## A.11 Performance

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (21a) | Specify a Guaranteed Port Filtering Rate, and the associated measurement interval *TF*, for each Bridge Port in the format specified below. | M | {D}16.1 | |
| (21b) | Specify a Guaranteed Bridge Relaying Rate, and the associated measurement interval *TR*, in the format specified below.<br><br>Supplementary information shall clearly identify the Ports. | M | {D}16.2 | |

| Guaranteed Bridge Relaying Rate | TR |
|---|---|
| _ _ _ _ _ _ _ _ _ _ frames per second | _ _ _ _ _ _ second(s) |

| Port number(s) or other identification | Guaranteed port filtering rate (specify for all ports) | $T_F$ (specify for all ports) |
|---|---|---|
| | _ _ _ _ _ _ _ _ _ _ frames per second | _ _ _ _ _ _ second(s) |
| | _ _ _ _ _ _ _ _ _ _ frames per second | _ _ _ _ _ _ second(s) |
| | _ _ _ _ _ _ _ _ _ _ frames per second | _ _ _ _ _ _ second(s) |
| | _ _ _ _ _ _ _ _ _ _ frames per second | _ _ _ _ _ _ second(s) |
| | _ _ _ _ _ _ _ _ _ _ frames per second | _ _ _ _ _ _ second(s) |


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

| Port number(s) or other identification | Guaranteed port filtering rate (specify for all ports) | $T_F$ (specify for all ports) |
|---|---|---|
| | _ _ _ _ _ _ _ _ _ _ frames per second | _ _ _ _ _ _ second(s) |
| | _ _ _ _ _ _ _ _ _ _ frames per second | _ _ _ _ _ _ second(s) |
| | _ _ _ _ _ _ _ _ _ _ frames per second | _ _ _ _ _ _ second(s) |

## A.12 GARP and GMRP

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| | If Item 2b is not supported, mark N/A and continue at item (22i). | | | N/A [ ] |
| (22a) | Is the GMRP Application address used as the destination MAC Address in all GMRP protocol exchanges? | 2b:M | {D}10.4.1, {D}Table 12-1 | Yes [ ] |
| (22b) | Are GMRP protocol exchanges achieved by means of LLC Type 1 procedures, using the LLC address for Spanning Tree protocol? | 2b:M | {D}12.4, {D}12.5, {D}Table 7-8 | Yes [ ] |
| (22c) | Are GMRP protocol exchanges achieved using the GARP PDU formats, and the definition of the attribute type and value encodings defined for GMRP? | 2b:M | 10, {D}10.3.1, {D}12.4, {D}12.5, {D}12.11 | Yes [ ] |
| (22d) | Does the implementation support the operation of the Applicant, Registrar, and Leave All state machines? | 2b:M | {D}12.8 | Yes [ ] |
| (22e) | Does the Bridge propagate registration GMRP information only on Ports that are part of the active topology of the GIP Context for the VLAN on which the registration was received? | 2b:M | 10, {D}12.3.3, {D}12.3.4 | Yes [ ] |
| (22f) | Are GARP PDUs received on Ports that are in the Forwarding State forwarded, filtered or discarded in accordance with the requirements for handling GARP Application addresses? | 2b:M | {D}7.12.3, {D}12.5 | Yes [ ] |
| (22g) | Does the GMRP application operate as defined in Clause 10 of ISO/IEC 15802-3, as modified by Clause 10 of this standard? | 2b:M | 10, {D}10, {D}10.3 | Yes [ ] |
| (22h) | Are received GARP PDUs that are not well formed for any GARP Applications supported, discarded? | 2b:M | 10, {D}10.3.1, {D}12.4, {D}12.5, {D}12.10, {D}12.11 | Yes [ ] |
| (22i) | Are all GARP PDUs that are<br>(a) Received on Ports that are in the Forwarding State, and are<br>(b) Destined for GARP applications that the Bridge does not support,<br>forwarded on all other Ports that are in Forwarding? | M | 8.14.3, {D}12.5 | Yes [ ] |

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 151 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.12 GARP and GMRP *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (22j) | Are any GARP PDUs that are<br>(a) Received on any Port, and<br>(b) Destined for GARP applications that the Bridge does not support,<br>submitted to any GARP Participants? | X | 8.14.3, {D}12.5 | No [ ] |
| (22k) | Are any GARP PDUs that are<br>(a) Received on any Ports that are not in the Forwarding State, and are<br>(b) Destined for GARP applications that the Bridge does not support,<br>forwarded on any other Ports of the Bridge? | X | 8.14.3, {D}12.5 | No [ ] |
| (22l) | Are any GARP PDUs that are<br>(a) Received on any Ports that are in the Forwarding State, and are<br>(b) Destined for GARP applications that the Bridge supports,<br>forwarded on any other Ports of the Bridge? | X | 8.14.3, {D}12.5 | No [ ] |
| (22m) | Are all GARP PDUs that are:<br>(a) Received on any Port, and<br>(b) Destined for GARP applications that the Bridge supports,<br>submitted to the appropriate GARP Participants? | M | 8.14.3, {D}12.5 | Yes [ ] |

## A.13 VLAN support

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| | Ingress rules | | | |
| (24a) | Can the PVID for any Port be assigned the value of the null VLAN ID? | X | 8.4.4, Table 9-2 | No [ ] |
| (24b) | Are frames discarded (or not discarded) in accordance with the settings of the Acceptable Frame Types parameters? | M | 8.6 | Yes [ ] |
| (24c) | Are all frames received classified as belonging to exactly one VLAN, as defined in the ingress rules? | M | 8.6 | Yes [ ] |
| (24d) | Is Ingress Filtering performed in accordance with the value of the Enable Ingress Filtering parameter? | M | 8.6 | Yes [ ] |
| (24e) | Are all frames that are not discarded as a result of the application of the ingress rules submitted to the Forwarding Process and to the Learning Process? | M | 8.6 | Yes [ ] |
| | Egress rules | | | |
| (25a) | Are frames discarded if the transmission Port is not present in the Member set for the frame's VID? | M | 8.8, 8.11.9 | Yes [ ] |


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.13 VLAN support *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (25b) | Are frames discarded if the value of the include_tag parameter is False, and the Bridge does not support the ability to translate embedded MAC Address information from the format indicated by the canonical_format_indicator parameter to the format appropriate to the media type on which the data request will be carried? | 23j.2:M | 8.8 | Yes [ ] N/A [ ] |
| (25c) | Are frames transmitted as VLAN-tagged frames or as untagged frames in accordance with the value of the untagged set for the frame's VID? | M | 8.8 | Yes [ ] |
| | Filtering Database | | | |
| (26a) | Does the implementation support Static VLAN Registration Entries as defined in 8.11.2? | M | 8.11.2 | Yes [ ] |
| (26b) | Does the implementation support the creation of a separate Static VLAN Registration Entry with a distinct Port Map for each VLAN from which frames are received by the Forwarding Process? | O | 8.11.2 | Yes [ ] No [ ] |
| (26c) | Does the implementation support Dynamic VLAN Registration Entries as defined in 8.11.5? | M | 8.11.5 | Yes [ ] |
| (26d) | Does the implementation support the creation of a separate Dynamic VLAN Registration Entry with a distinct Port Map for each VLAN from which frames are received by the Forwarding Process? | O | 8.11.5 | Yes [ ] No [ ] |
| (26e) | Does the implementation allocate VIDs to FIDs in accordance with the specification in 8.11.7? | M | 8.11.7, 8.11.7.2 | Yes [ ] |
| (26f) | Does the implementation correctly detect Learning Constraint violations? | M | 8.11.7.3 | Yes [ ] |
| (26g) | Is determination of the Member set and the untagged set for a given VLAN achieved as defined in 8.11.9? | M | 8.11.9 | Yes [ ] |
| | Tagged frames | | | |
| (27a) | Do VLAN-tagged frames transmitted by the Bridge conform to the format defined in Clause 9 for the MAC type on which they are transmitted? | M | 9 | Yes [ ] |
| (27b) | Are all BPDUs transmitted untagged? | M | 8.14.7 | Yes [ ] |
| | VLAN use of GMRP. If item (2b) is not supported, mark N/A and continue at item (29a). | | | N/A [ ] |
| (28a) | Does the implementation of GMRP recognize the use of VLAN Contexts for the transmission and reception of GMRP PDUs? | 2b:M | 10, 10.1, 10.2 10.3 | Yes [ ] |

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 153 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## A.13 VLAN support *(Continued)*

| Item | Feature | Status | References | Support |
|---|---|---|---|---|
| (28b) | Does the implementation of GMRP support the creation of distinct GMRP Participants for each VLAN context? | 2b:M | 10.2 | Yes [ ] |
| (28c) | Does the implementation support the identification of VLAN contexts in transmitted GMRP PDUs by means of VLAN-tagged or untagged frames, in accordance with the member set and untagged set for the VLAN Context concerned? | 2b:M | 10.3 | Yes [ ] |
| (28d) | Are GMRP PDUs transmitted only on Ports that are part of the active topology for the VLAN Context concerned? | 2b:M | 10.1 | Yes [ ] |
| | VLAN Topology Management | | | |
| (29a) | Does the implementation support the creation, updating and removal of Dynamic VLAN Registration Entries in the Filtering Database under the control of GVRP? | M | 11 | Yes [ ] |
| (29b) | Does the Permanent Database contain an entry for the Default VID that defines Registration Fixed on all Ports? | O | 11.2.1.3 | Yes [ ] No [ ] |
| (29c) | Is the GVRP Application address used as the destination MAC Address in all GVRP protocol exchanges? | M | 11, Table 11-1 | Yes [ ] |
| (29d) | Are GVRP protocol exchanges achieved by means of LLC Type 1 procedures, using the LLC address for Spanning Tree protocol? | M | 11, {D}12.4, {D}12.5, {D}Table 7-8 | Yes [ ] |
| (29e) | Are GVRP protocol exchanges achieved using the GARP PDU formats, and the definition of the attribute type and value encodings defined for GVRP? | M | 11, 11.2.3.1, {D}12.4, {D}12.5, {D}12.11 | Yes [ ] |
| (29f) | Does the implementation support the operation of the Applicant, Registrar, and Leave All state machines? | M | {D}12.8 | Yes [ ] |
| (29g) | Does the Bridge propagate registration GVRP information only on Ports that are part of the active topology of the base Spanning Tree Context? | M | 11, {D}12.3.3, {D}12.3.4 | Yes [ ] |
| (29h) | Does the GVRP application operate as defined in Clause 11? | M | 11 | Yes [ ] |


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# Annex B

(informative)

# Shared and Independent VLAN Learning

This standard provides for a variety of approaches to the implementation of VLAN Bridges from the point of view of the way that individual MAC Addresses are learned, and how that learned information is used in subsequent forwarding/filtering decisions. There are two mechanisms that are used as a basis for these variants:

a) Making use of address information learned across a number of VLANs in order to make learning decisions relative to any one of those VLANs. This is referred to as *Shared VLAN Learning* (SVL, 3.9);
b) Making use of address information learned in one VLAN only in order to make learning decisions relative to that VLAN, and ensuring that it is not used in learning decisions relative to any other VLAN. This is referred to as *Independent VLAN Learning* (IVL, 3.5).

These mechanisms lead to the SVL/IVL model for how a VLAN Bridge implements learning and filtering for MAC Addresses. Using the terminology of 8.11.7, an SVL/IVL Bridge supports multiple FIDs (which effectively equates to supporting multiple Filtering Databases), and multiple VLANs can use each FID. By varying the number of FIDs supported, and the number of VLANs that can share each FID, the following simplifications of the SVL/IVL model can be created:

c) *Shared VLAN Learning (SVL) only.* The implementation supports only one FID, so all VLANs share the same learned MAC Address information, regardless of which VLAN the information was learned in;
d) *Independent VLAN Learning (IVL) only.* Multiple FIDs are supported, but each FID can support only one VID, so each VLAN makes use only of MAC Address information learned within that VLAN.

All three approaches are permitted by this standard, and each has advantages in particular circumstances. The remainder of this annex discusses

e) The requirements for Independent VLAN Learning, Shared VLAN Learning, or both;
f) How Bridges are made aware of the requirement for particular VLANs to be “shared” or “independent”;
g) How Bridges based on one of these models can interoperate with Bridges based on a different model, in the same Bridged LAN.

## B.1 Requirements for Shared and Independent Learning

Under most circumstances, the SVL and IVL approaches work equally well, and Bridges adopting either approach can be freely intermixed within a Bridged LAN. There are, however, a small number of configuration cases where, in order to prevent undue flooding of unicast frames, and in some cases, to make communication between the affected end systems possible, it is necessary to make specific choices as to how Bridges that adopt these different learning models are deployed in a Bridged LAN. The following subclauses give examples of some of these configurations, and also provide a generic statement of the requirements that must be met in order for each learning model to be successfully deployed.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 155 of 211

### B.1.1 Connecting independent VLANs

Figure B-1 illustrates how a device that connects two VLANs together, and which therefore itself shares learning between those VLANs, creates a need for those VLANs to be independent in other Bridges.

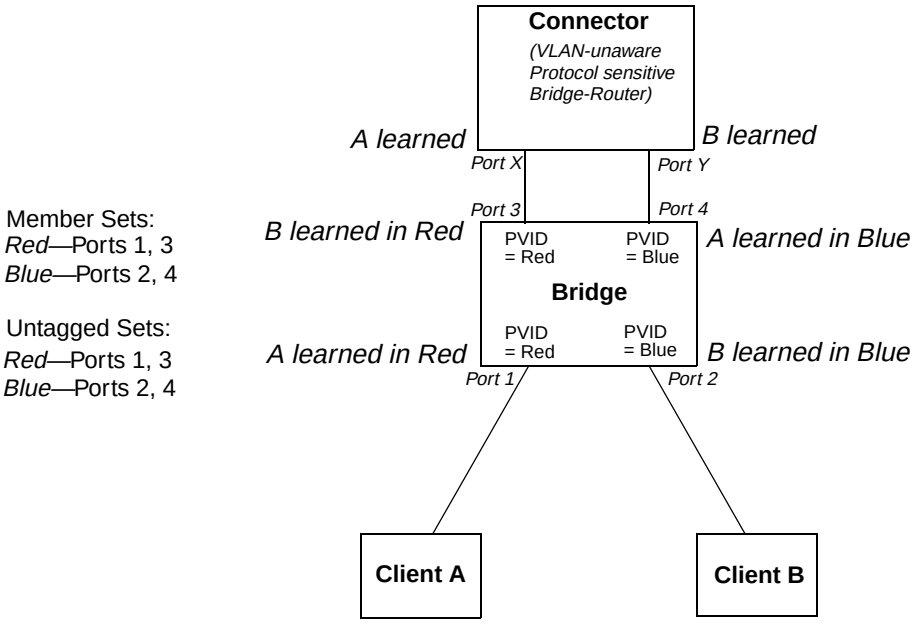


**Figure B-1—Connecting independent VLANs—1**

Clients A and B are connected via the protocol sensitive Bridge-Router (Connector), with an intervening VLAN-aware Bridge. The fact that all the Ports of the Bridge carry untagged traffic neatly conceals the fact that the Connector has the effect of connecting VLANs Red and Blue together with regard to bridged traffic. The Connector itself learns A and B in the same database, as it has no knowledge of VLANs Red and Blue. This prevents any traffic transmitted on the Red VLAN (Port X of the Connector) that is destined for A, from being bridged to Port Y and transmitted on the Blue VLAN.

The VLAN-aware Bridge must keep its learning separate for Red and Blue; otherwise, the addresses of the two clients would be alternately learned on diagonally opposite Ports as, for example, traffic sourced by A reenters the Bridge on Port 4 having previously been seen on Port 1.

NOTE—This example assumes that Spanning Tree is disabled in the Connector, so that the VLAN-aware Bridge does not attempt to suppress the loop that apparently exists if VLANs are not taken into account.

A simpler example can be constructed, with a single Port connecting the Connector and the VLAN-aware Bridge, if the Connector is itself VLAN-aware and transmits and receives only VLAN-tagged traffic. In this case, the Connector would allocate a single FID for use by Red and Blue. This is shown in Figure B-2.

### B.1.2 Duplicate MAC Addresses

The simplest example of a need for Independent VLAN Learning occurs where two (or more) distinct devices in different parts of the network reuse the same individual MAC Address, or where a single device is connected to multiple LAN segments, and all of its LAN interfaces use the same individual MAC Address. This is shown in Figure B-3.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

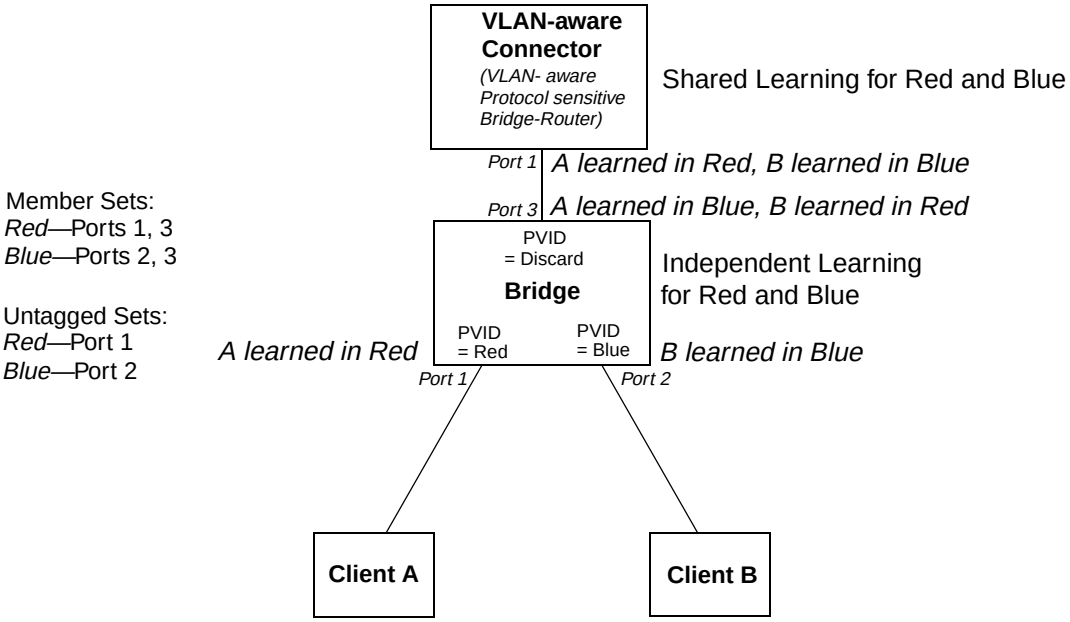


**Figure B-2—Connecting independent VLANs—2**

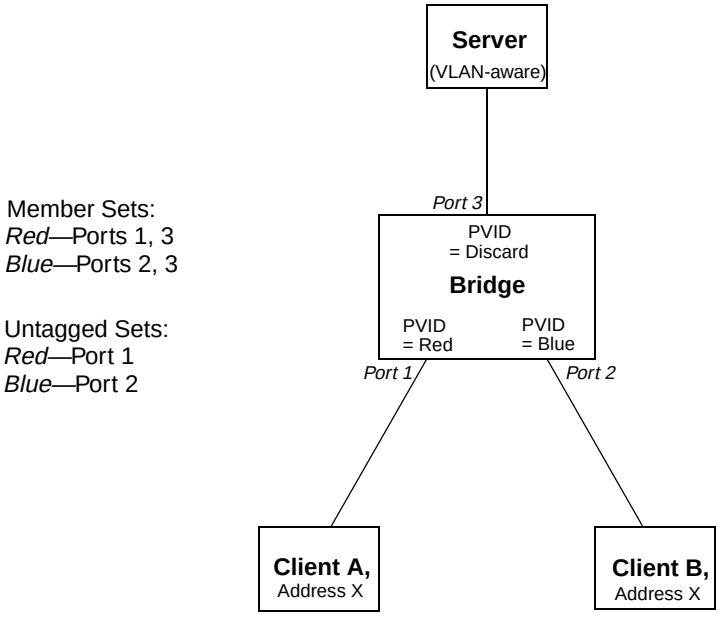


**Figure B-3—Duplicate MAC Addresses**

The example shows two clients with access to the same server; both clients are using the same individual MAC Address, X. If the Bridge shares learning between VLAN Red (which serves Client A) and VLAN Blue (which serves Client B), i.e., the Bridge uses the same FID for both VLANs, then Address X will appear to move between Ports 1 and 2 of the Bridge, depending upon which client has most recently transmitted a frame. Communication between these Clients and the server will therefore be seriously disrupted. Assignment of distinct FIDs for Red and Blue ensures that communication can take place correctly.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 157 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Hence, in order to construct this particular VLAN configuration, either an IVL Bridge or an SVL/IVL Bridge would be required.

### B.1.3 Asymmetric VLANs

A primary example of the requirement for Shared VLAN Learning is found in “asymmetric” uses of VLANs. Under normal circumstances, a pair of devices communicating in a VLAN environment will both send and receive using the same VLAN; however, there are some circumstances in which it is convenient to make use of two distinct VLANs, one used for A to transmit to B: the other used for B to transmit to A. An example of such an application of VLANs is shown in Figure B-4. Note that:

a) In the example, the server and both clients are assumed to be VLAN-unaware devices, i.e., they transmit and receive untagged frames only;
b) The ingress classification rules assumed by the example are as defined in this standard, i.e., Port-based classification only;
c) The configuration shown can only be achieved by management configuration of appropriate values in Static VLAN Registration Entries (8.11.9) in order to configure the indicated member sets and untagged sets.

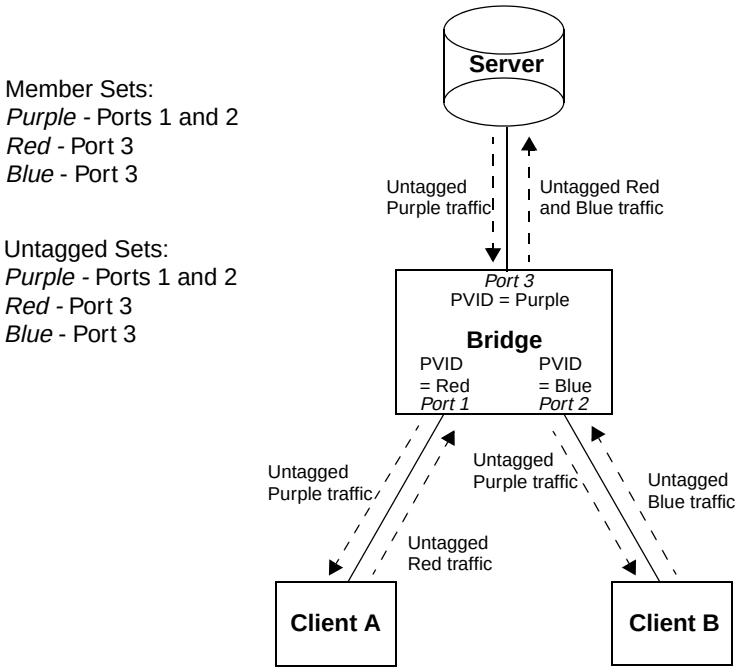


**Figure B-4—Asymmetric VLAN use: “multi-netted server”**

In the example, Port-based tagging and an asymmetric VLAN configuration is used in order to permit Clients A and B access to a common server, but to prohibit Clients A and B from talking to each other. Examples of where this type of configuration might be required are if the clients are on distinct IP subnets, or if there is some confidentiality-related need to segregate traffic between the clients.

Client A transmits to the server via Port 1, which will classify this traffic as belonging to VLAN Red; the Bridge therefore learns Client A’s MAC Address on Port 1 in VLAN Red. The Server transmits its responses to Client A via Port 3, which classifies the return traffic as belonging to VLAN Purple. If individual MAC Address learning is configured in the Bridge such that learning is independent between Red and Purple (Red


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

and Purple are allocated to distinct FIDs), then the Bridge will have no knowledge of A in VLAN Purple, and will therefore flood the server's responses to Client A on both Port 1 and Port 2. Conversely, if Red and Purple are defined to share the same FID, then the address information learned in Red will be available for use in forwarding the Purple traffic, and the responses to Client A are forwarded only through Port 1.

Similarly, there is a need in this configuration for Blue and Purple to share learning information; hence, in order for this configuration to achieve its objectives, the Red, Blue, and Purple VIDs must be allocated to the same FID in the Bridge.

Hence, in order to construct this particular VLAN configuration, either an SVL Bridge or an SVL/IVL Bridge would be required.

NOTE—The example has been deliberately simplified; in practical applications, the central Bridge would likely be replaced by a number of VLAN-aware Bridges, interconnected with links that would carry the traffic between clients and server as VLAN-tagged frames, with VLAN-tagging and untagging occurring only at the "edge" Ports of the Bridged LAN. An alternative approach to the one described here could also be achieved either by using a VLAN-aware server, or by use of more sophisticated Ingress classification rules.

### B.1.4 Generic constraints on SVL and IVL use

This subclause describes the general constraints on the mapping of VLANs to FIDs, from the point of view of a given Bridge that learns from or forwards frames on a set of VLANs (the Bridge's "active set" of VLANs). If

a) The individual MAC Addresses associated with each point of attachment to the active set of VLANs are unique (i.e., the "Duplicate MAC Address problem" is not present), and
b) There is no Bridge or Bridge-like device that takes frames from one VLAN in the active set and subsequently transmits them on another VLAN in the active set, then

every VLAN in the active set may share the same FID; in other words, individual MAC Address information learned in any one VLAN may be used in forwarding decisions taken relative to any of the others, so the SVL approach can be used in that Bridge.

Further, if

c) Each bidirectional, individual MAC-Addressed, conversation between pairs of end stations makes use of the same VLAN (ID) in both directions, then

every VLAN in the active set may be allocated a distinct FID (with the possibility of a little extra flooding as learning of addresses in one VLAN does not contribute to forwarding decisions for that address in any other VLAN). Under these circumstances, rule b) may also be relaxed and restated as follows:

d) Frames on one VLAN in the active set may be received by (up to) one Bridge and transmitted on another VLAN in the active set, provided that there is no loop in such VLAN to VLAN forwarding, e.g., for a set of VLANs Red, Blue, Green,...etc., there is no logical loop in copying frames between VLANs, such as copying from Red to Green by one Bridge, Green to Blue by another, and Blue back to Red by a third.

So

e) If rules a), b), and c) are true, and d) is false, for all VLANs in the active set, then either an SVL or an IVL Bridge can be deployed;
f) If rules a) and b) are true, and c) and d) are false for all VLANs in the active set, then only an SVL Bridge can be deployed;

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 159 of 211

g) If rules a) or b) or d) are false, and c) is true for all VLANs in the active set, then only an IVL Bridge can be deployed.

The above conditions are all on the basis that they apply “for all VLANs in the active set.” Clearly, in more complex scenarios, some VLANs in the active set will have requirements that dictate SVL behavior on the part of a given Bridge, while others will have requirements that dictate IVL behavior. Under such circumstances, an SVL/IVL Bridge is required, allowing those VLANs that need to be shared to be mapped to a single FID, while those that need to be independent are mapped to distinct FIDs. Needless to say, wherever an SVL or IVL Bridge can be deployed, it can successfully be replaced by an appropriately configured SVL/IVL Bridge.

## B.2 Configuring the Global VLAN Learning Constraints

Clause B.1 described the requirements that exist for the two approaches to learning in VLAN Bridges, closing with some generic rules for how to determine whether, for a given Bridge, SVL, IVL, or SVL/IVL can be successfully deployed. In VLAN Bridges, the set of requirements for Independent and/or Shared VLAN Learning is configured as a set of global VLAN Learning Constraint specifications, using the management tools defined in 12.10.3. Two types of VLAN Learning Constraint are defined in 8.11.7.2, which also defines how the set of constraints is used in order to derive a valid mapping of VIDs to FIDs.

The constraint specifications can be constructed on a modular basis. For example, the configuration shown in Figure B-4 has a requirement for Shared VLAN Learning between VLANs Red, Blue, and Purple. This could be expressed as follows:

{Red S Purple};
{Blue S Purple}

with {Red S Blue} being implied by the transitive nature of the S Constraint.

If we add a similar server access configuration in the same network that requires Red to share with Yellow and Orange, then this could be expressed as

{Red S Yellow};
{Orange S Yellow}

with {Red S Orange}, {Yellow S Blue}, {Yellow S Purple}, {Orange S Blue}, and {Orange S Purple} being implied by the transitive nature of the S Constraint.

Hence, Red, Blue, Purple, Yellow and Orange are all required to map to the same FID in order for the set of S Constraints (both explicit and implied) to be met. The constraints that express that requirement are built up from their constituent requirements; namely, for Red and Blue to share with Purple to meet one configuration need, and for Red to share with Yellow and Orange to meet another.

NOTE 1—The five VLANs in this example can be viewed as forming a Shared Set; i.e., a set of VLANs that have a mutual requirement to share learned information—all members of a Shared Set must map to the same FID. Any sequence of S Constraints defines one or more such Shared Sets. Any two Shared Sets can also map to the same FID as long as, for any pair of VLANs, one selected from each Shared Set, there are no I Constraints that require that pair of VLANs to learn independently. Hence, if there are only S Constraints defined, then all VLANs can be mapped to a single FID.

Similarly, the I Constraints can be added on a modular basis. Continuing from the above example, a Bridge-Router (Figure B-1) might be present in the Bridged LAN, which has the effect of connecting VLANs Indigo and Green together, thus creating a requirement for Indigo and Green to be independent. This could be expressed as


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

{Indigo I 1};
{Green I 1}

A separate independence requirement might be imposed by the fact that three stations, attached to Indigo, Vermilion, and Red VLANs, all make use of the same individual MAC Address. This could be expressed as:

{Indigo I 2};
{Vermilion I 2};
{Red I 2}

Hence, {Indigo, Vermilion, Red} have to be mutually independent (assigned to distinct FIDs), {Indigo, Green} have to be mutually independent, and {Red, Blue, Purple, Yellow, Orange} have to be shared.

The minimum number of FIDs required to satisfy this total constraint specification is three, e.g.:

FID A: Red, Blue, Purple, Yellow, Orange
FID B: Indigo
FID C: Green, Vermilion

although an equally valid allocation for 3 FIDs is

FID A: Green, Red, Blue, Purple, Yellow, Orange
FID B: Indigo
FID C: Vermilion

and an equally valid allocation, using the maximum number of FIDs that could be used for this set of constraints and VLANS is:

FID A: Red, Blue, Purple, Yellow, Orange
FID B: Indigo
FID C: Green
FID D: Vermilion

NOTE 2—It can clearly be seen from this example that it is possible to add further constraints that result in impossible VID to FID mappings; for example, if we were to add {Indigo S Red} or ({Yellow I 3}, {Blue I 3}), then the result is at least one pair of VLANs that have a requirement to both share the same FID and to use distinct FIDs at the same time. Such configurations are examples of Learning Constraint inconsistencies (8.11.7.3).

The assumption behind these constraint specifications is that they are applied globally, in the sense that all VLAN Bridges in a given Bridged LAN are configured with the same set of constraints. This is important, in order to ensure that each Bridge is in a position to determine whether or not, given its current active set of VLANs, it is capable of adopting a VID to FID mapping that will satisfy the specified constraints. If it cannot achieve such a mapping (for any of the reasons identified in 8.11.7.3), then it has detected a network misconfiguration that can only be resolved by management intervention. The managed object specification 12.10.3 provides a Notification for use in such circumstances, to alert a management station to the existence of the problem.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 161 of 211

## B.3 Interoperability

If the configuration of the Bridged LAN is such that it is not necessary to configure any VLAN Learning Constraints into the Bridges, i.e.:

a) There are no instances where two (or more) points of attachment to different LAN segments (and different VLANs) make use of the same individual MAC Address;
b) There are no instances where a Bridge receives frames on one VLAN and transmits them on another VLAN;
c) There is no asymmetric VLAN use, i.e., there is no pair of end stations for which bidirectional, unicast conversations make use of different VLANs for each direction of transmission,

then it is possible to freely intermix SVL, IVL, and SVL/IVL Bridges in that Bridged LAN, and they can all successfully interoperate.

If the configuration of the Bridged LAN requires one or more S Constraints (and no I Constraints) to be configured into the Bridges, then SVL and SVL/IVL Bridges can be used freely; however, IVL Bridges may only be used in locations where their active set of VLANs does not include any pair of VLANs for which an S Constraint (either explicit or implied) has been defined.

If the configuration of the Bridged LAN requires two or more I Constraints (and no S Constraints) to be configured into the Bridges, then IVL and SVL/IVL Bridges can be used freely; however, SVL Bridges may only be used in locations where their active set of VLANs does not include any pair of VLANs for which I Constraints with the same Independent Set Identifier have been defined.

If the configuration of the Bridged LAN requires both I Constraints and S Constraints to be configured into the Bridges, then SVL/IVL Bridges can be used freely; however,

d) SVL Bridges may only be used in locations where their active set of VLANs does not include any pair of VLANs for which I Constraints with the same Independent Set Identifier have been defined, and
e) IVL Bridges may only be used in locations where their active set of VLANs does not include any pair of VLANs for which an S Constraint (either explicit, or implied) has been defined.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# Annex C

(informative)

# MAC method dependent aspects of VLAN support

This annex examines the set of services, frame formats, and MAC methods involved in the provision of VLAN services across IEEE 802 LANs using different MAC methods, and the mapping/bridging functions necessary for that provision.

## C.1 The variables

End station MAC Service users make use of MAC data transmission services that convey the following types of information:

a) Ethernet Type-encoded (E) and LLC-encoded (L) information (see 3.1 and 3.2);
b) Frames (either E or L) in which any MAC Addresses embedded in the MAC data are carried in Canonical (C) or Non-canonical (N) format;

NOTE 1—The terms *Canonical format* and *Non-canonical format* are described in Annex F.

c) Frames that carry source-routing information (R), or frames that are Bridged transparently (T).

Hence, there are potentially eight combinations of these variables, corresponding to eight distinct services, as follows:

d) E-C-T (Ethernet Type-encoded, Canonical, transparent),
e) E-C-R (Ethernet Type-encoded, Canonical, source-routed),
f) E-N-T (Ethernet Type-encoded, Non-canonical, transparent),
g) E-N-R (Ethernet Type-encoded, Non-canonical, source-routed),
h) L-C-T (LLC-encoded, Canonical, transparent),
i) L-C-R (LLC-encoded, Canonical, source-routed),
j) L-N-T (LLC-encoded, Non-canonical, transparent),
k) L-N-R (LLC-encoded, Non-canonical, source-routed),

These services are supported over two basic LAN types:

l) 802.3/Ethernet (C);
m) Token Ring/FDDI (R).

There are two VLAN environments involved:

n) Untagged frames (U);
o) Tagged frames (T).

This leads to a total of 32 potential frame/encapsulation formats to consider (8 X 2 X 2). There are 96 possible one-way heterogeneous bridging functions between these various LAN/VLAN environments; 48 symmetrical (2-way) functions.

The combination of services and environments is illustrated in Figure C-1; italics indicate services that have no untagged representation on the MAC method concerned.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 163 of 211

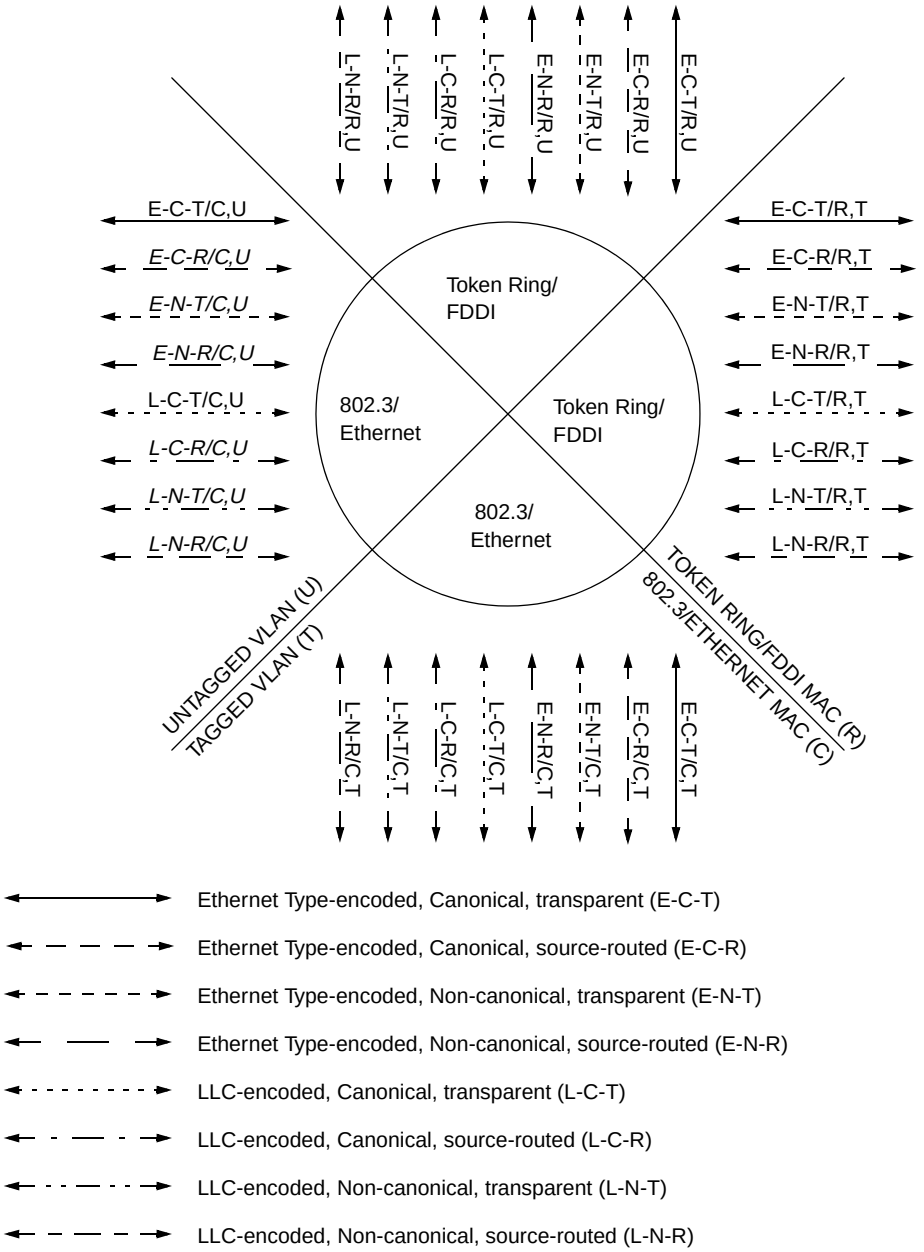


**Figure C-1—Services and environments**

Figure C-2 illustrates the heterogeneous Bridging functions involved. In this diagram, the bridging functions are labelled H, Q or Q+H depending upon whether the function involves ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 frame translation, 802.1Q VLAN-tagging/untagging/tag translation, or a combination of 802.1Q VLAN-tagging/untagging/tag translation and ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 frame translation.

NOTE 2—Figure C-2 is not intended to represent a real LAN, simply to illustrate the various Bridging functions.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

In both diagrams, the frame formats involved are identified by three initial letters that identify the service provided, from the list of 8 services above. The fourth letter indicates the MAC method that carries the frame (C or R), and the fifth letter indicates the type of VLAN (U or T).

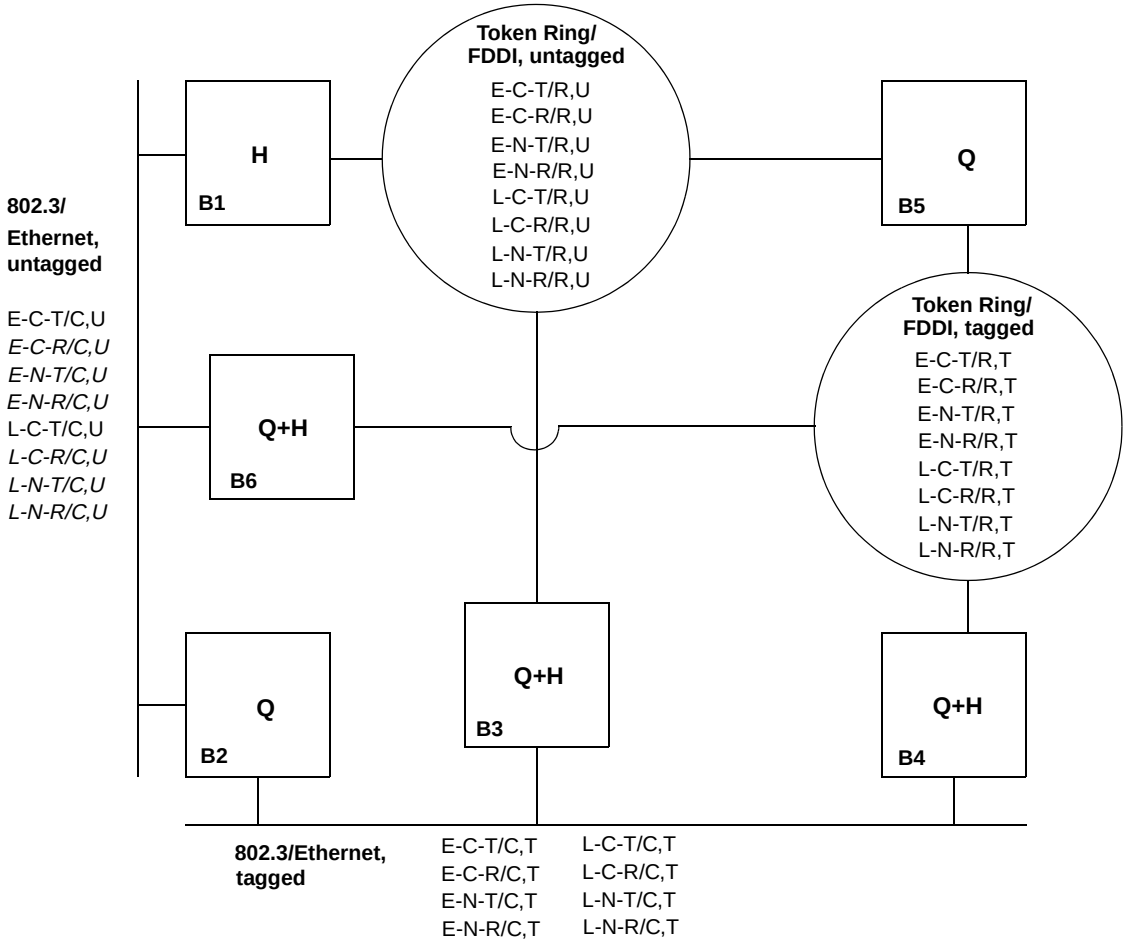


**Figure C-2—Heterogeneous Bridging functions**

## C.2 Bridging functions

### C.2.1 Bridging function B1

This function bridges between heterogeneous untagged VLAN environments. The following frame translations are involved:

a) ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 encapsulation/decapsulation for frames carrying Ethernet Type-encoded information;
b) ISO/IEC 15802-3 conversion for frames carrying LLC-encoded information.
c) Any requirements for translation from Canonical to Non-canonical address format, or vice versa, must be met if communication is to be maintained between end stations separated by this Bridging function.
d) Any source-routed traffic cannot be relayed between these environments, as there is no representation for source-routing information in untagged frames on 802.3/Ethernet LANs.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 165 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### C.2.2 Bridging function B2

This function involves VLAN entry (I to T) and exit (T to I); it bridges between untagged and tagged 802.3/Ethernet environments. The frame translations involved are as follows:

a) Insertion of Ethernet-encoded tag headers on VLAN entry;
b) Removal of Ethernet-encoded tag headers on VLAN exit;
c) Translation of Non-canonical MAC Addresses to Canonical on VLAN exit;
d) Any source-routed traffic cannot be relayed between these environments, as there is no representation for source-routing information in untagged frames on 802.3/Ethernet LANs.

NOTE—VLAN entry in Non-canonical format does not occur, as the native representation on 802.3/Ethernet is Canonical format. VLAN exit of Non-canonical information can occur only if the Bridge is capable of translating the representation of embedded MAC Addresses to their Canonical format.

### C.2.3 Bridging function B3

This function involves VLAN entry and exit; it bridges between tagged 802.3/Ethernet and untagged Ring environments. The following frame translations are involved:

a) For untagged Ethernet Type-encoded information on Token Ring/FDDI (VLAN entry): Removal of ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 encapsulation, and insertion of Ethernet-encoded tag header;
b) For tagged Ethernet Type-encoded information on 802.3/Ethernet (VLAN exit): Removal of tag header and insertion of ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 encapsulation;
c) For VLAN entry/exit with frames carrying LLC-encoded information: Insertion/removal of Ethernet-encoded tag header;
d) Translation of MAC Addresses to the format appropriate for the destination Ring on VLAN exit;
e) Any source-routing information present in the frame is preserved; the Token Ring/FDDI RIF is copied into the E-RIF in the Ethernet-encoded tag header on VLAN entry (with the CFI/NCFI flags set appropriately), and copied back on exit.

NOTE—VLAN entry (tagged 802.3/Ethernet from untagged Token Ring/FDDI) in Canonical format normally occurs only from ISO/IEC 9314-2, as the native representation on ISO/IEC 8802-5 is Non-canonical format, and for ISO/IEC 9314-2 is Canonical format. Hence, VLAN exit in Canonical format onto ISO/IEC 8802-5 can occur only if the Bridge is capable of translating the representation of embedded MAC Addresses; i.e., of converting the frame from Canonical on 802.3/Ethernet to Non-canonical on ISO/IEC 8802-5. Similarly, VLAN exit in Non-canonical format onto ISO/IEC 9314-2 can occur only if the Bridge is capable of converting the frame from Non-canonical on 802.3/Ethernet to Canonical on ISO/IEC 9314-2.

### C.2.4 Bridging function B4

This function bridges between tagged 802.3/Ethernet and Ring environments. The following frame translations are involved:

a) For tagged Ethernet Type-encoded information on Token Ring/FDDI to 802.3/Ethernet: Removal of ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 encapsulation, and conversion of the tag header to the Ethernet-encoded form;
b) For tagged Ethernet Type-encoded information on 802.3/Ethernet to Token Ring/FDDI: ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 encapsulation, and conversion of the tag header to the SNAP-encoded form;
c) For tagged frames carrying LLC-encoded information: conversion of the tag header between the SNAP-encoded and Ethernet-encoded forms;


Copyright © 1999 IEEE. All rights reserved.

d) Any source-routing information is preserved between these environments by copying between the Token Ring/FDDI RIF and the E-RIF in the Ethernet-encoded tag header.

NOTE 1—This Bridging function is not required to modify the format of embedded MAC Addresses.
NOTE 2—In FDDI LANs, source-routing information may be present either in an E-RIF within the tag header or in the normal position for a source-routed frame.

### C.2.5 Bridging function B5

This function involves VLAN entry and exit; it bridges between tagged and untagged Ring environments. The frame translations involved are as follows:

a) Insertion of SNAP-encoded tag header on VLAN entry;

b) Removal of SNAP-encoded tag header on VLAN exit;

c) Translation of MAC Addresses to the format appropriate for the destination Ring on VLAN exit;

d) Any source-routing information present in the frame is preserved.

NOTE 1—VLAN entry in Canonical format normally occurs only from ISO/IEC 9314-2, as the native representation on ISO/IEC 8802-5 is Non-canonical format, and for ISO/IEC 9314-2 is Canonical format. Hence, VLAN exit in Canonical format onto ISO/IEC 8802-5 can occur only if the Bridge is capable of translating the representation of embedded MAC Addresses; i.e., of converting the frame from Canonical on 802.3/Ethernet to Non-canonical on ISO/IEC 8802-5. Similarly, VLAN exit in Non-canonical format onto ISO/IEC 9314-2 can occur only if the Bridge is capable of converting the frame from Non-canonical on 802.3/Ethernet to Canonical on ISO/IEC 9314-2.
NOTE 2—In FDDI LANs, source-routing information may be present either in an E-RIF within the tag header or in the normal position for a source-routed frame.

### C.2.6 Bridging function B6

This function involves VLAN entry and exit; it bridges between untagged 802.3/Ethernet and tagged Ring environments. The following frame translations are involved:

a) For untagged Ethernet Type-encoded information on 802.3/Ethernet (VLAN entry): ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 encapsulation, and insertion of SNAP-encoded tag header;

b) For tagged Ethernet Type-encoded information on Token Ring/FDDI (VLAN exit): Removal of SNAP-encoded tag header and removal of ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 encapsulation;

c) For VLAN entry/exit with frames carrying LLC-encoded information: Insertion/removal of SNAP-encoded tag header;

d) Any source-routed traffic cannot be relayed between these environments, as there is no representation for source-routing information in untagged frames on 802.3/Ethernet LANs.

NOTE 1—VLAN entry in Non-canonical format does not occur, as the native representation on 802.3/Ethernet is Canonical format. VLAN exit of Non-canonical format can occur only if the Bridge is capable of translating the representation of embedded MAC Addresses; i.e., of converting the frame from Non-canonical format to Canonical format on 802.3/Ethernet.
NOTE 2—In FDDI LANs, source-routing information may be present either in an E-RIF within the tag header or in the normal position for a source-routed frame.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 167 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## C.3 Frame formats

The following abbreviations are used in the descriptions of the frame formats in this annex, with the following meanings:

| | |
|---|---|
| AC | Access Control field—in Token Ring frames only (see ISO/IEC 15802-5 and Note below) |
| RCI | Ring Control Information—AC (if present) plus FC fields |
| DA | Destination MAC Address |
| SA | Source MAC Address |
| PT | Ethernet Protocol Type |
| SPT | SNAP-encoded Ethernet Protocol Type (C.6.1) |
| TPID | Tag Protocol ID (9.3.1) |
| ETPID | Ethernet-encoded TPID (9.3.1.1) |
| STPID | SNAP-encoded TPID (9.3.1.2) |
| TCI | Tag Control Information (9.3.2) |
| CFI | Canonical Format Indicator (9.3.2.2) |
| NCFI | Non-canonical Format Indicator (9.3.3.5) |
| C | Canonical |
| N | Non-canonical |
| R | E-RIF present |
| VID | VLAN Identifier (9.3.2.3) |
| Len | IEEE Std 802.3-style Length/Type field (C.6.2) |
| LLC | LLC addressing and control information, as defined in ISO/IEC 15802-2 |
| RIF | Source-Routing Information Field (C.6.4) |
| E-RIF | Embedded RIF (9.3.3, C.6.4) |
| C-Data | MAC user data in which any embedded MAC Addresses are in Canonical format (C.6.3) |
| N-Data | MAC user data in which any embedded MAC Addresses are in Non-canonical format |
| PAD | Padding (C.6.5) |
| FCS | Frame Check Sequence |

In C.3.2, the possible frame formats are categorized by service type; in C.3.3 they are categorized by bearer MAC method and tagging method.

NOTE—The text in this annex makes the generalization of treating the FC fields in 8802-5 and FDDI as if they are the same, in order to simplify the descriptions as much as possible. In reality, there are detailed differences between FC fields in the two MAC methods. When translating between 8802-5 and FDDI, the most likely behavior is to propagate the “LLC frame” indication and the User Priority field between the FC octets on input and output.

### C.3.1 Structure of the tagged frame

Figure C-3 illustrates the frame formats used for carrying tagged Ethernet Type-encoded information and LLC-encoded information using 8802-5 Token Ring MAC methods.

Figure C-4 illustrates the frame formats used for carrying tagged Ethernet Type-encoded information and LLC-encoded information using FDDI MAC methods. Two forms of tagged frame are shown:

a) The *source-routed form*, in which the frame carries a RIF in the normal position, following the source MAC Address. This form can only be used on FDDI LANs that support source routing; and
b) The *transparent form*, in which an E-RIF is present in the tag header if the frame carries Non-canonical or source-routed information.

Figure C-5 illustrates the frame formats used for carrying tagged Ethernet Type-encoded information and LLC-encoded information on 802.3/Ethernet MAC methods.


Copyright © 1999 IEEE. All rights reserved.

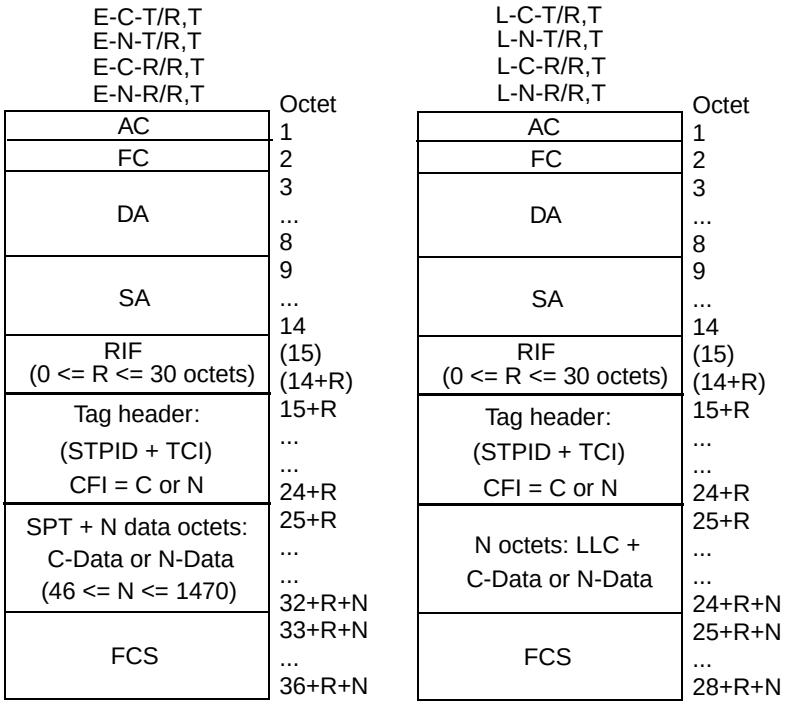


**Figure C-3—Tagged frames on 8802-5 Token Ring LANs**

As can be seen from these diagrams, the major differences between the tagged frame formats in 802.3/Ethernet and Token Ring/FDDI MAC methods are

c) The presence/absence of RCI (Ring Control Information);
d) The position of the RIF and E-RIF fields;
e) The encoding used to carry the Tag Protocol Identifier (2 octets for ETPID vs. 8 octets for STPID);
f) The encoding used to carry Ethernet Protocol Types (2 octets for PT vs. 8 octets for SPT);
g) The presence/absence of the Length/Type field;
h) The presence/absence of the PAD field.

The diagrams also illustrate the similarities between:

i) The format of tagged frames on 8802-5 and the source-routed form of tagged frames on FDDI;
j) The format of tagged frames on 802.3/Ethernet and the transparent form of tagged frames on FDDI.

## C.3.2 Frame formats by service type

### C.3.2.1 Frame formats for Ethernet Type-encoded service

#### C.3.2.1.1 Ethernet Type-encoded, Canonical, transparent

E-C-T/C,U: DA, SA, PT, C-Data, FCS
E-C-T/C,T: DA, SA, ETPID, TCI (CFI=C), PT, C-Data, FCS
E-C-T/R,U: RCI, DA, SA (RII reset), SPT, C-Data, FCS
E-C-T/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=C), SPT, C-Data, FCS

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 169 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

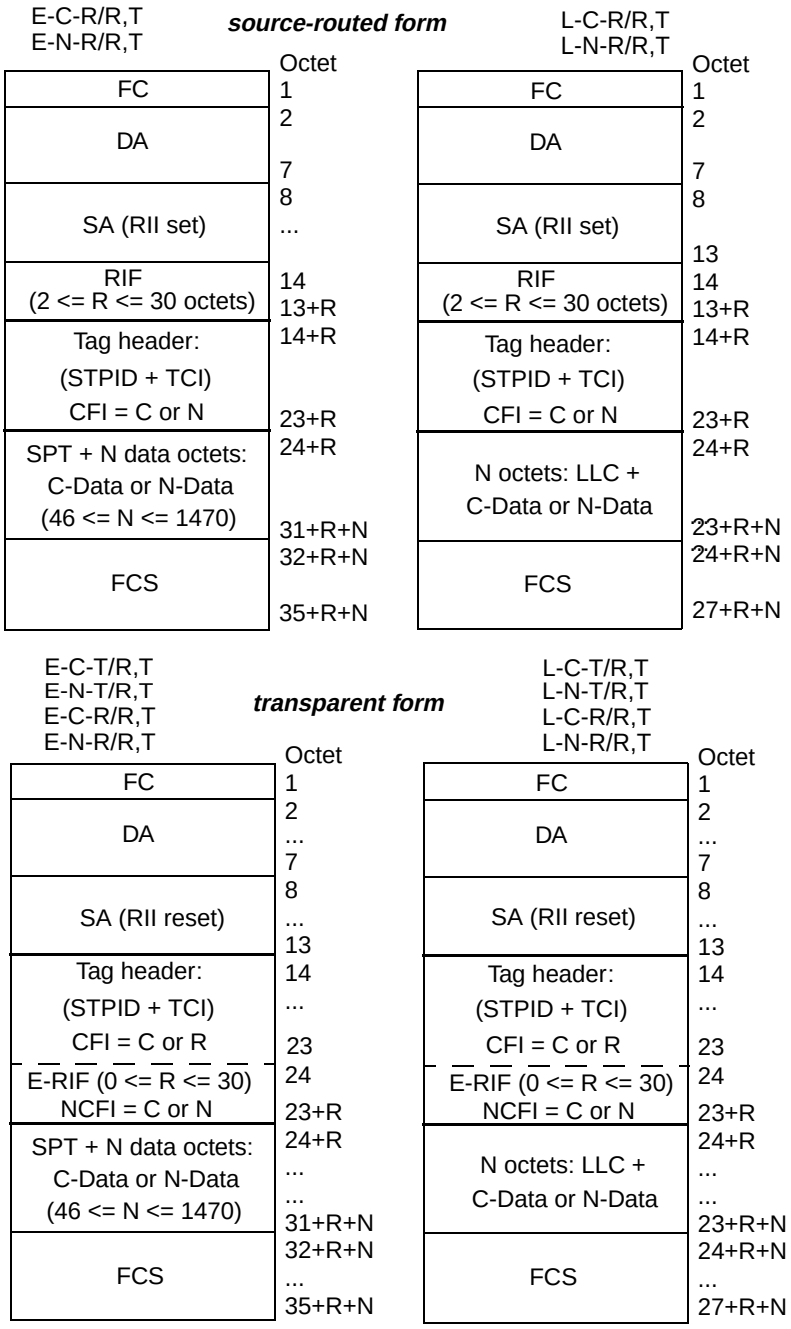


**Figure C-4—Tagged frames on FDDI LANs**

### C.3.2.1.2 Ethernet Type-encoded, Canonical, source-routed

E-C-R/C,U: No representation possible
E-C-R/C,T: DA, SA, ETPID, TCI (CFI=R), PT, E-RIF (NCFI=C), C-Data, FCS
E-C-R/R,U: RCI, DA, SA (RII set), RIF, SPT, C-Data, FCS
E-C-R/R,T: RCI, DA, SA (RII set), RIF, STPID, TCI (CFI=C), SPT, C-Data, FCS *(source-routed form)*
E-C-R/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=C), SPT, C-Data, FCS *(transparent form)*


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### C.3.2.1.3 Ethernet Type-encoded, Non-canonical, transparent

E-N-T/C,U: No representation possible
E-N-T/C,T: DA, SA, ETPID, TCI (CFI=R), PT, E-RIF (NCFI=N), N-Data, FCS
E-N-T/R,U: RCI, DA, SA (RII reset), SPT, N-Data, FCS
E-N-T/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI =N), SPT, N-Data, FCS
*(8802-5 Token Ring form)*
E-N-T/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=N), SPT, N-Data, FCS
*(FDDI form)*

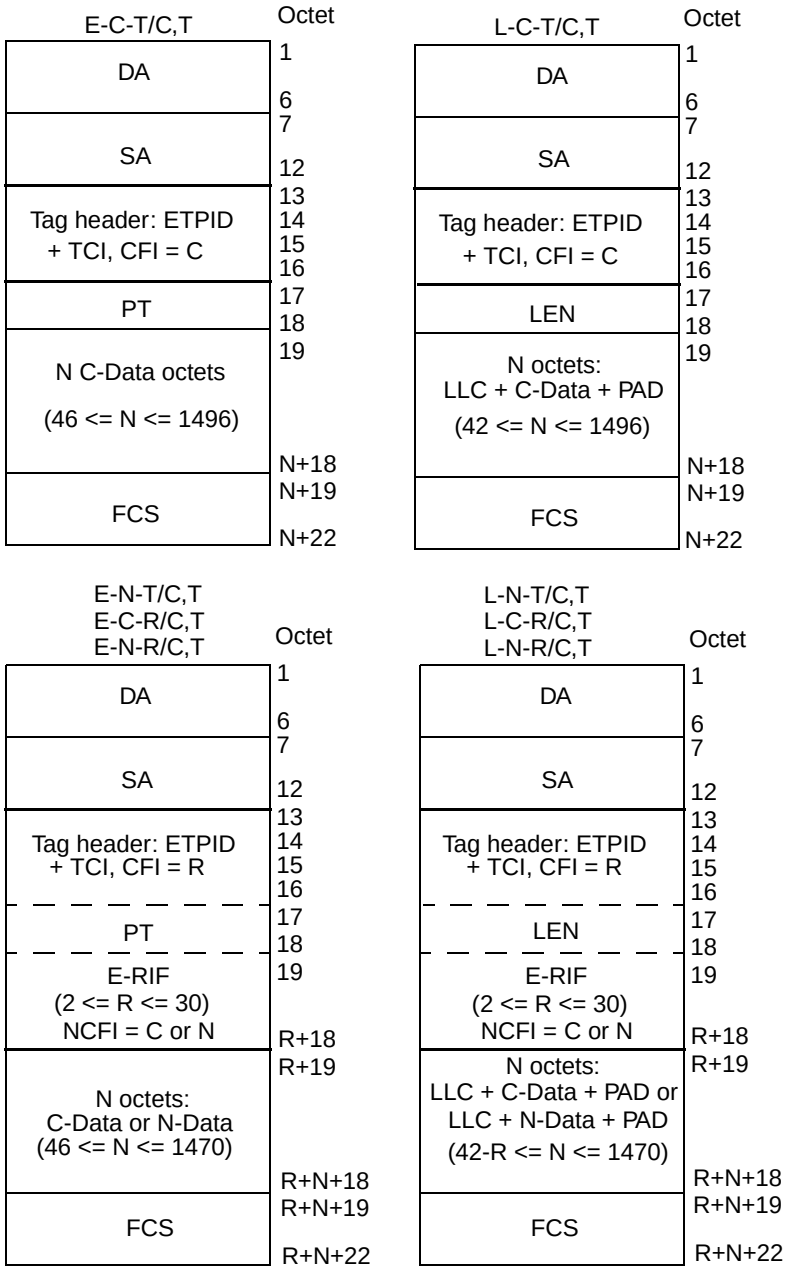


**Figure C-5—Tagged frames on 802.3/Ethernet LANs**

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 171 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### C.3.2.1.4 Ethernet Type-encoded, Non-canonical, source-routed

E-N-R/C,U: No representation possible
E-N-R/C,T: DA, SA, ETPID, TCI (CFI=R), PT, E-RIF (NCFI=N), N-Data, FCS
E-N-R/R,U: RCI, DA, SA (RII set), RIF, SPT, N-Data, FCS
E-N-R/R,T: RCI, DA, SA (RII set), RIF, STPID, TCI (CFI=N), SPT, N-Data, FCS
*(source-routed form)*
E-N-R/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=N), N-Data, FCS
*(transparent form)*

### C.3.2.2 Frame formats for LLC-encoded service

#### C.3.2.2.1 LLC-encoded, Canonical, transparent

L-C-T/C,U: DA, SA, LEN, LLC, C-Data, PAD, FCS
L-C-T/C,T: DA, SA, ETPID, TCI (CFI=C), LEN, LLC, C-Data, PAD, FCS
L-C-T/R,U: RCI, DA, SA (RII reset), LLC, C-Data, FCS
L-C-T/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=C), LLC, C-Data, FCS

#### C.3.2.2.2 LLC-encoded, Canonical, source-routed

L-C-R/C,U: No representation possible
L-C-R/C,T: DA, SA, ETPID, TCI (CFI=R), LEN, E-RIF (NCFI=C), LLC, C-Data, PAD, FCS
L-C-R/R,U: RCI, DA, SA (RII set), RIF, LLC, C-Data, FCS
L-C-R/R,T: RCI, DA, SA (RII set), RIF, STPID, TCI (CFI=C), LLC, C-Data, FCS
*(source-routed form)*
L-C-R/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=C), LLC, C-Data, FCS
*(transparent form)*

#### C.3.2.2.3 LLC-encoded, Non-canonical, transparent

L-N-T/C,U: No representation possible
L-N-T/C,T: DA, SA, ETPID, TCI (CFI=R), LEN, E-RIF (NCFI=N), LLC, N-Data, PAD, FCS
L-N-T/R,U: RCI, DA, SA (RII reset), LLC, N-Data, FCS
L-N-T/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=N), LLC, N-Data, FCS
*(8802-5 Token Ring form)*
L-N-T/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=N), LLC, N-Data, FCS
*(FDDI form)*

#### C.3.2.2.4 LLC-encoded, Non-canonical, source-routed

L-N-R/C,U: No representation possible
L-N-R/C,T: DA, SA, ETPID, TCI (CFI=R), LEN, E-RIF (NCFI=N), LLC, N-Data, PAD, FCS
L-N-R/R,U: RCI, DA, SA (RII set), RIF, LLC, N-Data, FCS
L-N-R/R,T: RCI, DA, SA (RII set), RIF, STPID, TCI (CFI=N), LLC, N-Data, FCS
*(source-routed form)*
L-N-R/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=N), LLC, N-Data, FCS
*(transparent form)*


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## C.3.3 Frame formats by MAC method type and tagging method

### C.3.3.1 Frame formats for 802.3/Ethernet MAC methods

#### C.3.3.1.1 802.3/Ethernet, untagged

E-C-T/C,U: DA, SA, PT, C-Data, FCS
E-C-R/C,U: No representation possible
E-N-T/C,U: No representation possible
E-N-R/C,U: No representation possible
L-C-T/C,U: DA, SA, LEN, LLC, C-Data, PAD, FCS
L-C-R/C,U: No representation possible
L-N-T/C,U: No representation possible
L-N-R/C,U: No representation possible

#### C.3.3.1.2 802.3/Ethernet, tagged

E-C-T/C,T: DA, SA, ETPID, TCI (CFI=C), PT, C-Data, FCS
E-C-R/C,T: DA, SA, ETPID, TCI (CFI=R), PT, E-RIF (NCFI=C), C-Data, FCS
E-N-T/C,T: DA, SA, ETPID, TCI (CFI=R), PT, E-RIF (NCFI=N), N-Data, FCS
E-N-R/C,T: DA, SA, ETPID, TCI (CFI=R), PT, E-RIF (NCFI=N), N-Data, FCS
L-C-T/C,T: DA, SA, ETPID, TCI (CFI=C), LEN, LLC, C-Data, PAD, FCS
L-C-R/C,T: DA, SA, ETPID, TCI (CFI=R), LEN, E-RIF (NCFI=C), LLC, C-Data, PAD, FCS
L-N-T/C,T: DA, SA, ETPID, TCI (CFI=R), LEN, E-RIF (NCFI=N), LLC, N-Data, PAD, FCS
L-N-R/C,T: DA, SA, ETPID, TCI (CFI=R), LEN, E-RIF (NCFI=N), LLC, N-Data, PAD, FCS

### C.3.3.2 Frame formats for Token Ring/FDDI MAC methods

#### C.3.3.2.1 Token Ring/FDDI, untagged

E-C-T/R,U: RCI, DA, SA (RII reset), SPT, C-Data, FCS
E-C-R/R,U: RCI, DA, SA (RII set), RIF, SPT, C-Data, FCS
E-N-T/R,U: RCI, DA, SA (RII reset), SPT, N-Data, FCS
E-N-R/R,U: RCI, DA, SA (RII set), RIF, SPT, N-Data, FCS
L-C-T/R,U: RCI, DA, SA (RII reset), LLC, C-Data, FCS
L-C-R/R,U: RCI, DA, SA (RII set), RIF, LLC, C-Data, FCS
L-N-T/R,U: RCI, DA, SA (RII reset), LLC, N-Data, FCS
L-N-R/R,U: RCI, DA, SA (RII set), RIF, LLC, N-Data, FCS

#### C.3.3.2.2 Token Ring/FDDI, tagged

E-C-T/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=C), SPT, C-Data, FCS
E-C-R/R,T: RCI, DA, SA (RII set), RIF, STPID, TCI (CFI=C), SPT, C-Data, FCS
*(source-routed form)*
E-C-R/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=C), SPT, C-Data, FCS
*(transparent form)*
E-N-T/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=N), SPT, N-Data, FCS
*(8802-5 Token Ring form)*
E-N-T/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=N), SPT, N-Data, FCS
*(FDDI form)*
E-N-R/R,T: RCI, DA, SA (RII set), RIF, STPID, TCI (CFI=N), SPT, N-Data, FCS
*(source-routed form)*
E-N-R/R,T: RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=N), N-Data, FCS
*(transparent form)*

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 173 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

| | |
|---|---|
| L-C-T/R,T: | RCI, DA, SA (RII reset), STPID, TCI (CFI=C), LLC, C-Data, FCS |
| L-C-R/R,T: | RCI, DA, SA (RII set), RIF, STPID, TCI (CFI=C), LLC, C-Data, FCS *(source-routed form)* |
| L-C-R/R,T: | RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=C), LLC, C-Data, FCS *(transparent form)* |
| L-N-T/R,T: | RCI, DA, SA (RII reset), STPID, TCI (CFI=N), LLC, N-Data, FCS *(8802-5 Token Ring form)* |
| L-N-T/R,T: | RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=N), LLC, N-Data, FCS *(FDDI form)* |
| L-N-R/R,T: | RCI, DA, SA (RII set), RIF, STPID, TCI (CFI=N), LLC, N-Data, FCS *(source-routed form)* |
| L-N-R/R,T: | RCI, DA, SA (RII reset), STPID, TCI (CFI=R), E-RIF (NCFI=N), LLC, N-Data, FCS *(transparent form)* |

## C.4 Procedures for tagging, untagging, and relaying tagged frames

The formal definition of the procedures whereby tag headers are added and removed, and tagged frames are relayed are embodied in Clauses 7 and 8. This informal description is included in order to add clarity to the formal definition of the process.

### C.4.1 Tagging

The following subclauses describe the translations that are performed when an untagged frame is relayed in tagged form.

#### C.4.1.1 MAC header information

The RCI, DA, SA and RIF fields (if supported in the source frame and/or destination MAC methods) are translated from their representation in the source frame into the equivalent representation in the destination frame in accordance with the procedures described in ISO/IEC 15802-3. This will result in

a) Preservation of the AC (Token Ring to Token Ring only) and FC fields (Token Ring/FDDI to Token Ring/FDDI only);
b) Translation of the DA and SA into their equivalent representation in the destination MAC methods;
c) Preservation of the RIF field, if present; either in its conventional position (Token Ring/FDDI to Token Ring/FDDI, source-routed form) or within the tag header (Token Ring/FDDI to 802.3/Ethernet or FDDI, transparent form).

NOTE—The ability of the tag header to carry source-routing information across 802.3/Ethernet LANs does not imply a requirement on the part of a pure 802.3/Ethernet Bridge to support source routing. This capability is provided simply to allow traffic that originates in, and is destined for, a source-routed environment to transit as tagged traffic across a non-source-routed environment. Similarly, this capability allows source-routed traffic to transit an FDDI network that is otherwise unable to support source routing.

#### C.4.1.2 Tag header insertion

The tag header is inserted immediately following the SA field (if no RIF is present in the destination frame) or immediately following the RIF field (if RIF is present in the destination frame). The header contains

a) An Ethernet-encoded TPID (destination MAC method is 802.3/Ethernet) or a Snap-encoded TPID (destination MAC method is Token Ring/FDDI);
b) A TCI field, as follows:
   1) The user_priority field is set in accordance with the procedure described in ISO/IEC 15802-3;


Copyright © 1999 IEEE. All rights reserved.

2) The CFI flag, indicating C/N (8802-5 Token Ring, and source-routed FDDI MAC methods), or C/[RIF present] (802.3/Ethernet and transparent FDDI MAC methods), in accordance with the format of the MAC user data;
3) The VID field is set to the VID of the VLAN to which the source frame belongs.

c) An E-RIF field, immediately following the Length/Type field (802.3/Ethernet and transparent FDDI MAC methods), if the frame is carrying Non-canonical data and/or source-routing information. The NCFI in the RIF indicates C or N, in accordance with the format of the MAC user data.

#### C.4.1.3 Ethernet Type-encoded data

If the MAC user data carries Ethernet Type-encoded data, i.e., the protocol identifier is an Ethernet Type value or an Ethernet Type value that has been SNAP encoded as described in ISO/IEC 15802-3 and ISO/IEC 11802-5, and if the frame is being relayed between differing MAC methods (802.3/Ethernet to or from Token Ring/FDDI), then the data is translated from its source format to the format appropriate to the destination MAC method in accordance with the procedures described in ISO/IEC 15802-3 and ISO/IEC 11802-5.

#### C.4.1.4 FCS

When tagging a frame and performing the attendant field translations, it is necessary to recompute the Frame Check Sequence (FCS) field of the tagged frame. As stated in ISO/IEC 15802-3, 6.3.7, the Bridge shall not introduce additional undetected frame errors as a result of such FCS recomputation.

NOTE—Where necessary in order to preserve the protection afforded by the original FCS, it is possible to incrementally compute the new FCS value, based on the original FCS, adjusted for the new fields added and the frame length. This technique, and other techniques for preserving FCS integrity, are discussed in ISO/IEC 15802-3, Annex G.

### C.4.2 Untagging

The following subclauses describe the frame translations that are performed when a received tagged frame is relayed in untagged format.

#### C.4.2.1 MAC header information

The RCI, DA, SA and RIF or E-RIF fields (if supported in the source frame and/or destination MAC methods) are translated from their representation in the source frame into the equivalent representation in the destination frame in accordance with the procedures described in ISO/IEC 15802-3. This will result in

a) Preservation of the AC (Token Ring to Token Ring only) and FC fields (Token Ring/FDDI to Token Ring/FDDI only);
b) Translation of the DA and SA into their equivalent representation in the destination MAC method;
c) Preservation of any source-routing information carried in the source frame, if present, and if the destination MAC method is a Token Ring/FDDI LAN that supports source routing. (If the source MAC method is 802.3/Ethernet or transparent FDDI and the tag header carries an E-RIF in which the RT field indicates a transparent frame, then the E-RIF is not considered to be carrying any source-routing information.)

#### C.4.2.2 Tag header

The tag header is removed.

#### C.4.2.3 Ethernet Type-encoded data

If the frame carries Ethernet Type-encoded data, i.e., the protocol identifier is an Ethernet Type value or an Ethernet Type value that has been SNAP encoded as described in ISO/IEC 15802-3 and ISO/IEC 11802-5,

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 175 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

and if the frame is being relayed between differing MAC methods (802.3/Ethernet to or from Token Ring/FDDI), then the data is translated from its source format to the format appropriate to the destination MAC method in accordance with the with the procedures described in ISO/IEC 15802-3 and ISO/IEC 11802-5.

#### C.4.2.4 Address translation

If the CFI/NCFI information in the tagged frame indicates that embedded addresses are being carried in a format inappropriate to the destination MAC method, then it is necessary either to translate the addresses from C to N or vice versa, or to discard the frame if such translation is not supported by the Bridge.

#### C.4.2.5 FCS

When removing a frame's VLAN tag and performing the attendant field translations, it is necessary to recompute the Frame Check Sequence (FCS) field of the tagged frame. As stated in ISO/IEC 15802-3, 6.3.7, the Bridge shall not introduce additional undetected frame errors as a result of such FCS recomputation.

NOTE—Where necessary in order to preserve the protection afforded by the original FCS, it is possible to incrementally compute the new FCS value, based on the original FCS, adjusted for the new fields added and the frame length. This technique, and other techniques for preserving FCS integrity, are discussed in ISO/IEC 15802-3, Annex G.

### C.4.3 Relaying tagged frames

The following subclauses describes the frame translations that are performed when a received tagged frame is relayed in tagged format.

#### C.4.3.1 MAC header information

The RCI, DA, and SA (if supported in the source frame and/or destination frame formats) are translated from their representation in the source frame into the equivalent representation in the destination frame in accordance with the procedures described in ISO/IEC 15802-3.

For source-routed Token Ring/FDDI to 802.3/Ethernet or transparent FDDI, the RIF field (if present) is translated into the E-RIF field of the destination frame.

For relay between source-routed Token Ring/FDDI environments, the RIF (if present) is copied into the RIF field of the destination frame.

For 802.3/Ethernet or transparent FDDI to source-routed Token Ring/FDDI, the E-RIF field, if present, is translated into the RIF of the destination frame, with the NCFI bit reset, unless the E-RIF indicates that the frame is a transparent frame, in which case, the E-RIF is discarded.

This will result in

a) Preservation of the AC (Token Ring to Token Ring only) and FC fields (Token Ring/FDDI to Token Ring/FDDI only);
b) Translation of the DA and SA into their equivalent representation in the destination MAC method;
c) Preservation of any information carried in the E-RIF or RIF field, if present and if it carries source-routing information.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

#### C.4.3.2 Tag header

If the source and destination MAC methods differ, the tag header is modified as follows:

a) The TPID field is set in accordance with the destination MAC method. An Ethernet-encoded TPID is used where the destination MAC method is 802.3/Ethernet; a Snap-encoded TPID is used where the destination MAC method is Token Ring/FDDI;
b) The information carried in the User Priority and VID fields in the TCI are copied unchanged into the destination frame's TCI.
c) If the source and destination MAC methods are of the same type, then the CFI (and RIF, if present in 802.3/Ethernet) are copied unchanged into the destination tag header.
d) If the source and destination MAC methods differ, then the CFI information in the source tag header is translated into the format appropriate for the destination tag header.

#### C.4.3.3 Ethernet Type-encoded data

If the frame carries Ethernet Type-encoded data, i.e., the protocol identifier is an Ethernet Type value or an Ethernet Type value that has been SNAP encoded as described in ISO/IEC 15802-3 and ISO/IEC 11802-5, and if the frame is being relayed between differing MAC methods (802.3/Ethernet to or from Token Ring/FDDI), then the data is translated from its source format to the format appropriate to the destination MAC method in accordance with the with the procedures described in ISO/IEC 15802-3 and ISO/IEC 11802-5.

#### C.4.3.4 FCS

When relaying tagged frames, if it is necessary to perform any attendant field translations, then it is necessary to recompute the Frame Check Sequence (FCS) field of the tagged frame. As stated in ISO/IEC 15802-3, 6.3.7, the Bridge shall not introduce additional undetected frame errors as a result of such FCS recomputation.

NOTE—Where necessary in order to preserve the protection afforded by the original FCS, it is possible to incrementally compute the new FCS value, based on the original FCS, adjusted for the new fields added and the frame length. This technique and other techniques for preserving FCS integrity are discussed in ISO/IEC 15802-3, Annex G.

### C.4.4 Padding and frame size considerations

#### C.4.4.1 Treatment of PAD fields in IEEE Std 802.3 frames

The minimum frame size constraint placed on 802.3/Ethernet frames requires frames to carry zero or more pad octets following the MAC client data, in order to ensure that no frame of total length less than 64 octets is transmitted on the medium. This means that frames whose overall length would otherwise be less than 64 octets in length have (64-len) octets of padding added after the MAC client data, where len is the size of the frame before padding.

When tagged frames are transmitted by a Bridge on an IEEE Std 802.3 MAC, there are two permissible approaches (7.2):

a) Keep the minimum frame size generated by the Bridge equal to 64 octets. This implies that the number of pad octets in a received untagged IEEE Std 802.3 frame would be reduced by up to 4 octets when that frame was tagged;
b) Adopt a minimum tagged frame length of 68 octets. This implies that the number of pad octets in a received untagged IEEE Std 802.3 frame would not be adjusted when tagging such frames; equally, if subsequently untagged, no pad adjustment would be necessary before transmission on 802.3/Ethernet.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 177 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

There is a similar choice to be made in end stations that generate tagged frames:

c) In some existing implementations, the decision as to whether pad octets are needed will be made at a point where it is impractical to distinguish between tagged and untagged frames. In these cases, the end station will use a minimum frame size of 64 octets for all frames;
d) In other cases, the padding decision will be taken at a point before it is known whether the frame will be transmitted tagged or untagged. In these cases, the end station will use a minimum tagged frame size of 68 octets, and a minimum of 64 octets for untagged frames.

The above approaches are all consistent with the IEEE Std 802.3 frame specification, as amended by IEEE Std 802.3ac-1998.

The implication of this is that, for correct operation on 802.3/Ethernet, all devices have to be capable of correctly handling tagged frames of less than 68 octets in length (C.4.4.3).

### C.4.4.2 Maximum PDU size

VLAN tagging of an untagged frame, or relaying frames in tagged frame format, can result in an increase in the length of the original frame. If transmission of a given frame in tagged frame format through a given destination Port would result in violation of the maximum PDU size for the destination MAC method, the Bridge discards the frame for that destination Port.

NOTE—Violation of the maximum PDU sizes for destination MAC methods can produce undefined results in Bridged LANs that contain devices that adhere strictly to these maxima, or in MAC methods where these maxima are inherently constrained by the operation of the MAC method itself (e.g., constrained by timing considerations in the MAC state machines).

IEEE Std 802.3ac-1998 defines an extension to the normal 802.3 maximum frame size for the specific purpose of accommodating the additional octets of the VLAN tag header. The example frame translations in this annex make use of this extension to the 802.3 frame size.

### C.4.4.3 Minimum PDU size

VLAN untagging of a tagged frame results in the original frame decreasing in length.

Where the destination MAC is CSMA/CD:

a) If untagging a given frame would result in violation of the minimum frame length requirements of CSMA/CD, the Bridge is required to adjust the PAD field to ensure that the frame length equals the minimum length of 64 octets (7.2 and C.4.4.1);
b) If a frame is transmitted in tagged frame format, the Bridge may adopt a minimum tagged frame length of either 64 or 68 octets, as an implementation option. If the latter is chosen, the Bridge adjusts the size of the PAD field on transmission for any tagged frame that is less than 68 octets in length (7.2, C.4.4.1).

## C.5 Frame translations for different MAC methods

Examples of the frame translations that can occur when an untagged frame is translated into a tagged frame, and when tagged frames are relayed, are illustrated in the following clauses.

Subclauses C.5.1 and C.5.2 describe the translations that can occur when untagged frames on 802.3/Ethernet, and Token Ring/FDDI are translated into the tagged frame format. C.5.3 describes the translations that


Copyright © 1999 IEEE. All rights reserved.

can occur when a tagged frame is relayed between differing MAC methods in tagged format. In each subclause, the following cases are shown:

a) The untagged frame carried Ethernet Type-encoded information;
b) The untagged frame carried LLC-encoded information.

NOTE—In developing the example translations, the field sizes on 802.3/Ethernet have been calculated using the IEEE Std 802.3ac-1998 extension to the standard maximum frame size (normally 1518 octets). IEEE Std 802.3ac-1998 allows the maximum frame size to be extended by 4 octets for the specific purpose of accommodating the tag header.

## C.5.1 Tagging of untagged 802.3/Ethernet frames

### C.5.1.1 Ethernet Type-encoded information on 802.3/Ethernet LAN to tagged frame format

Figure C-6 illustrates the translation between an untagged Ethernet Type-encoded frame on 802.3/Ethernet (E-C-T/C,U) and a tagged frame on 802.3/Ethernet (E-C-T/C,T).

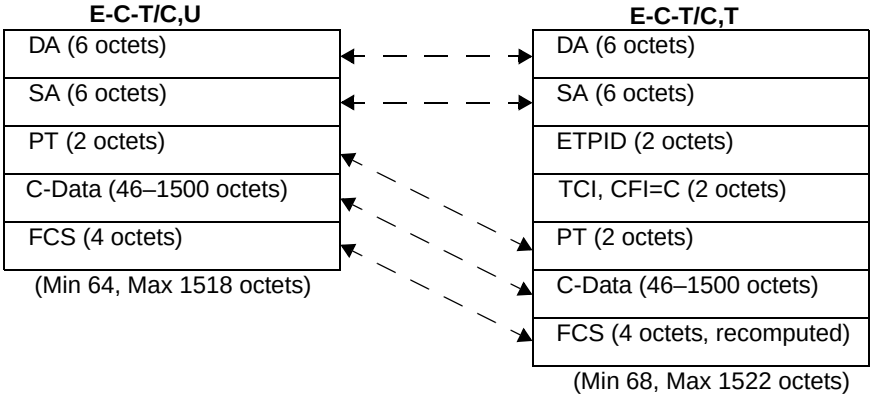


**Figure C-6—Translation between E-C-T/C,U and E-C-T/C,T**

The following translations are required in order to tag an E-C-T/C,U frame on 802.3/Ethernet:

a) The SA and DA fields are copied unchanged;
b) The ETPID and TCI are inserted, with CFI=C;
c) The PT and C-Data fields are copied unchanged;
d) The FCS is recomputed.

Removal of the tag involves the reverse of this process.

This form of tagging causes the original frame size to be increased by 4 octets.

Figure C-7 illustrates the translation between an untagged Ethernet Type-encoded frame on 802.3/Ethernet (E-C-T/C,U) and a tagged frame on Token Ring/FDDI (E-C-T/R,T).

The following translations are required in order to tag an E-C-T/C,U frame on a Token Ring/FDDI LAN:

e) The appropriate variant of the RCI field is added;
f) The DA and SA fields carry the same MAC Addresses as in the original frame;

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 179 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

NOTE—The meaning of the wording used in f) (and in other instances in this annex where this form of words is used) is that the MAC Addresses in the original and translated frames, when represented using the hexadecimal notation defined in Clause 5 of IEEE Standard 802, are the same.

g) The STPID and TCI are inserted, with CFI=C;
h) The PT is translated into the ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390-encoded form (SPT);
i) The C-Data field is copied unchanged;
j) The FCS is recomputed.

Removal of the tag involves the reverse of this process.

This form of tagging causes the original frame size to be increased by 17 octets for FDDI or 18 octets for Token Ring.

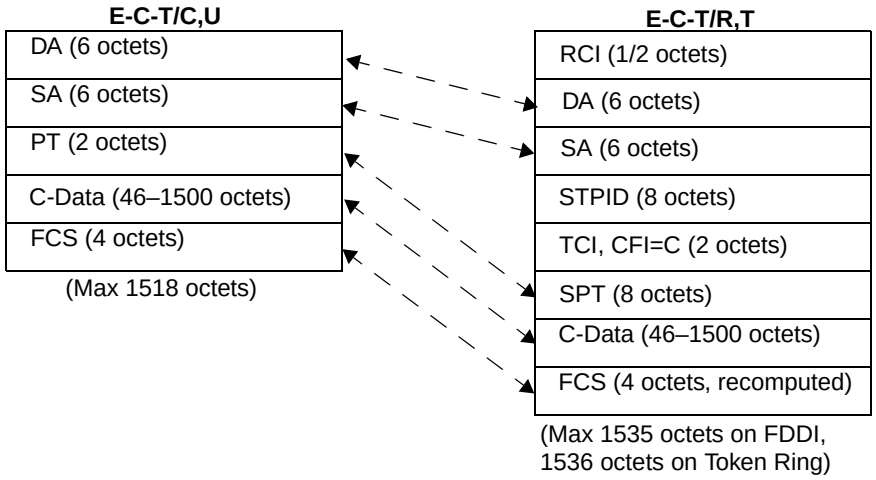


**Figure C-7—Translation between E-C-T/C,U and E-C-T/R,T**

NOTE—In translational (VLAN-unaware) bridging between 802.3/Ethernet and Ring LANs, an Ethernet Type-encoded frame increases in size by 7 octets on FDDI, and 8 octets on Token Ring.

Translations for E-C-R/C,U, E-N-T/C,U, and E-N-R/C,U to their equivalent tagged frame formats (E-C-R/C,T, E-N-T/C,T, and E-N-R/C,T on 802.3/Ethernet, and E-C-R/R,T, E-N-T/R,T and E-N-R/R,T on Token Ring/FDDI) cannot be shown, as there is no representation for such untagged frames on 802.3/Ethernet LANs. Similarly, translation of the tagged frames E-C-R/C,T, E-N-R/C,T, E-N-R/R,T, and E-C-R/R,T to untagged frames on 802.3/Ethernet is not possible, as it involves loss of the source-routing information. Translation of the remaining Non-canonical, transparent tagged frame formats into E-C-T/C,U is possible, but only if the Bridge is capable of translating Non-canonical data to its Canonical form.

### C.5.1.2 LLC-encoded information on 802.3/Ethernet to tagged frame format

Figure C-8 illustrates the translation between an untagged frame on 802.3/Ethernet carrying LLC-encoded information (L-C-T/C,U) and a tagged frame on 802.3/Ethernet (L-C-T/C,T).

Tagging an L-C-T/C,U frame on 802.3/Ethernet LANs requires the following frame translations:

a) The DA and SA fields are copied unchanged;
b) Insert ETPID and TCI fields, with CFI=C;


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

c) Len, LLC and C-Data fields are copied unchanged;
d) The PAD may either be copied unchanged (giving a minimum tagged frame size of 68 octets), or reduced by up to 4 octets (giving a minimum tagged frame size of 64 octets), as an implementation option;

NOTE—If the actual length of the data portion of the frame is inconsistent with the value held in LEN, then this inconsistency is not corrected by the tagging process. This is done in order that protocols which generate such inconsistency, and which require that inconsistency to be maintained for their correct operation, are not broken by this aspect of tagging.

e) Recompute the FCS.

Removal of the tag involves the reverse of this process.

This form of tagging causes the original frame size to be increased by 4 octets.

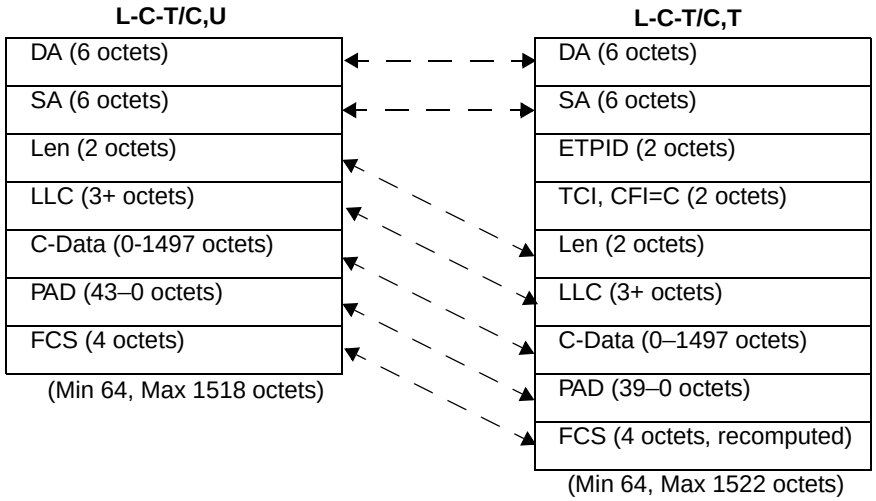


**Figure C-8—Translation between L-C-T/C,U and L-C-T/C,T**

Figure C-9 illustrates the translation between an untagged LLC-encoded frame on 802.3/Ethernet (L-C-T/C,U) and a tagged frame on Token Ring/FDDI (L-C-T/R,T).

Tagging in LLC-encoded format consists of the following frame translations:

f) The appropriate RCI field for the Ring MAC method concerned is added;
g) The DA and SA fields carry the same MAC Addresses as in the original frame;
h) Insert STPID and TCI fields, with CFI=C;
i) The Len field is removed;
j) Copy the LLC field unchanged;
k) The C-Data field is copied unchanged;
l) The PAD field is removed;
m) Recompute the FCS.

Removal of the tag (tagged frame to native 802.3/Ethernet frame) involves the reverse of this process.

This form of tagging causes the original frame size to be increased by 9 octets for FDDI or 10 octets for Token Ring.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 181 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

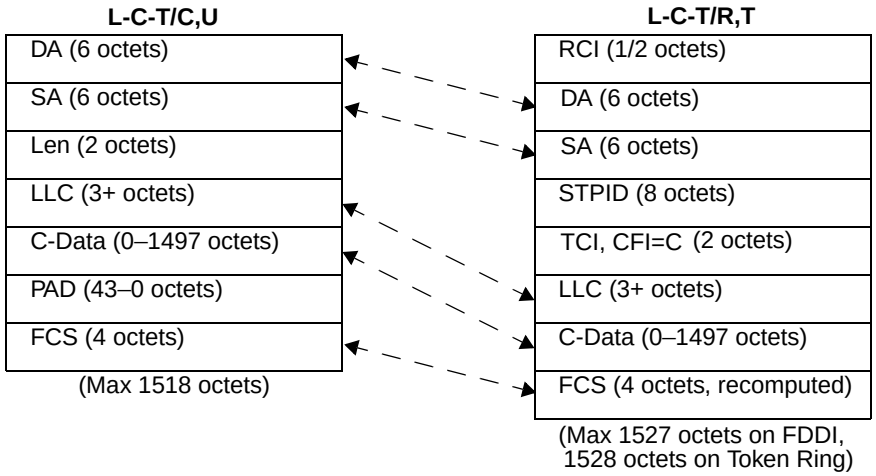


**Figure C-9—Translation between L-C-T/C,U and L-C-T/R,T**

NOTE—In translational (VLAN-unaware) bridging between 802.3/Ethernet and Ring LANs, an LLC-encoded frame reduces in size by 1 octet on FDDI, and does not change in length on Token Ring.

Translations for L-C-R/C,U, L-N-T/C,U, and L-N-R/C,U to their equivalent tagged frame formats (L-C-R/C,T, L-N-T/C,T, and L-N-R/C,T on 802.3/Ethernet, and L-C-R/R,T, L-N-T/R,T, and L-N-R/R,T on Token Ring/FDDI) cannot be shown, as there is no representation for such untagged frames on 802.3/Ethernet LANs. Similarly, translation of L-C-R/C,T, L-N-R/C,T, L-N-R/R,T, and L-C-R/R,T to untagged frames on 802.3/Ethernet is not possible, as it involves loss of the source-routing information. Translation of the remaining Non-canonical, transparent tagged frame formats into L-C-T/C,U is possible, but only if the Bridge is capable of translating Non-canonical data to its Canonical form.

## C.5.2 Translation of untagged Token Ring/FDDI frames

### C.5.2.1 Ethernet Type-encoded information on Token Ring/FDDI to tagged frame format

Figure C-10 illustrates the translation between an untagged Ethernet Type-encoded frame on Token Ring/FDDI (E-C-T/R,U, E-N-T/R,U, E-C-R/R,U or E-N-R/R,U) and a tagged frame on 802.3/Ethernet (E-C-T/C,T, E-N-T/C,T, E-C-R/C,T, or E-N-R/C,T).

Tagging requires the following frame translations:

a) Remove the RCI field;
b) The DA and SA fields carry the same MAC Addresses as in the original frame, with the RII bit reset;
c) Insert ETPID and TCI, with CFI=C (E-C-T/R,U) or R (all other frame types);
d) If the RII bit was set in the original frame, translate the RIF into the tag header E-RIF. For E-N-T/R,U, create a RIF with frame type = transparent. Set the E-RIF NCFI to C or N appropriately;
e) Translate the SPT into its corresponding PT;
f) Copy the Data field;
g) Recompute the FCS.

Removal of the tag involves the reverse of this process.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

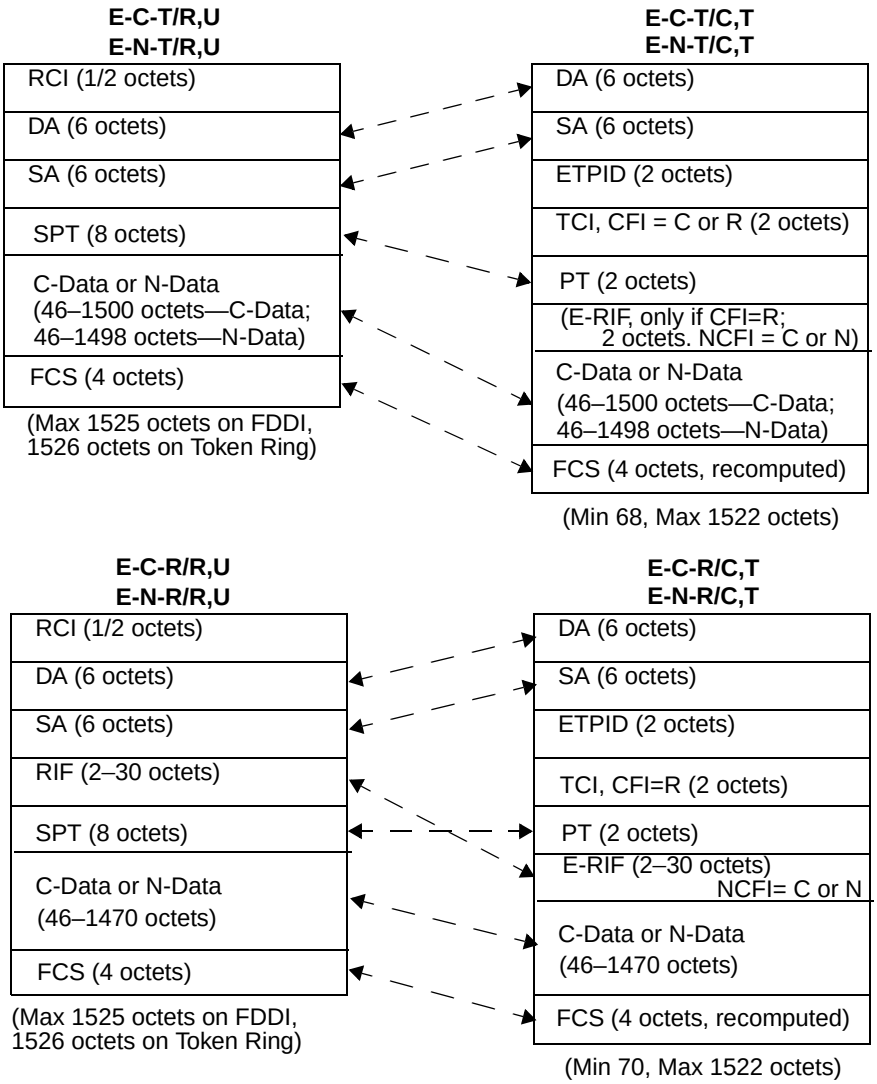


**Figure C-10—Translation between E-X-X/R,U and E-X-X/C,T**

NOTE 1—When removing the tag, if the CFI/NCFI indicates that embedded address information is in a form inappropriate for the destination MAC method, then it is necessary either to translate the address information or to discard the frame.

This form of tagging causes the original frame size to be reduced by 3 or 4 octets.

Figure C-11 illustrates the translation between an untagged Ethernet Type-encoded frame on Token Ring/FDDI (E-C-T/R,U, E-N-T/R,U, E-C-R/R,U, or E-N-R/R,U) and a tagged frame on 8802-5 Token Ring (E-C-T/R,T, E-N-T/R,T, E-C-R/R,T, or E-N-R/R,T). Figure C-11 also illustrates the translation of E-C-T/R,U, E-C-R/R,U, and E-N-R/R,U to tagged frames on FDDI media, the latter two translations illustrating the source-routed form of the tagged frame on FDDI.

Tagging requires the following frame translations:

h) Copy the RCI field;
i) The DA and SA fields carry the same MAC Addresses as in the original frame, with RII in the same state as in the original frame;

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 183 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

j) Copy the RIF field if present (RII set);
k) Insert STPID and TCI, setting the CFI to N or C appropriately;
l) Copy the SPT field;
m) Copy the Data field;
n) Recompute the FCS.

Removal of the tag (tagged frame to native Token Ring/FDDI frame) involves the reverse of this process.

NOTE 2—When removing the tag, if the CFI indicates that embedded address information is in a form inappropriate for the destination MAC method, then it is necessary either to translate the address information or to discard the frame.

This form of tagging causes the original frame size to be increased by 10 octets.

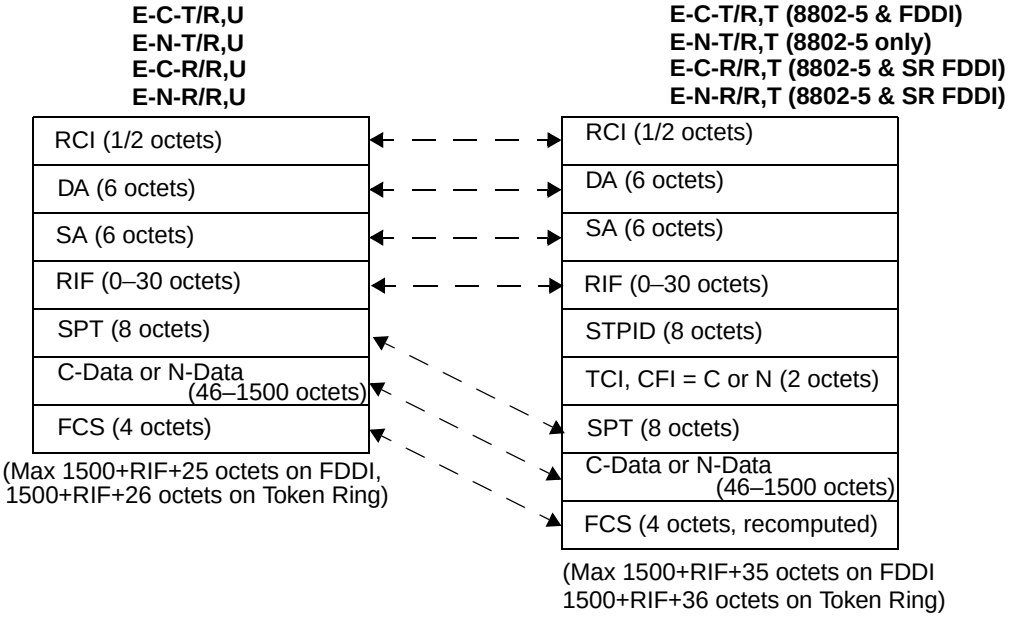


**Figure C-11—Translation between E-X-X/R,U and E-X-X/R,T (8802-5 & SR FDDI)**

Figure C-12 illustrates the translation between an untagged Ethernet Type-encoded frame on Token Ring/FDDI (E-N-T/R,U, E-C-R/R,U or E-N-R/R,U) and a tagged frame on FDDI (E-N-T/R,T, E-C-R/R,T or E-N-R/R,T). Note that the translation of E-C-T/R,U to E-C-T/R,T was dealt with in Figure C-11.

Tagging requires the following frame translations:

o) Copy the RCI field;
p) The DA and SA fields carry the same MAC Addresses as in the original frame, but with RII reset regardless of its state in the original frame;
q) Insert STPID and TCI, setting the CFI to R;
r) Translate the RIF field if present (RII set in source frame) to the E-RIF, otherwise create an E-RIF with RT indicating a transparent frame. Set the NCFI to C or N appropriately;
s) Copy the SPT field;
t) Copy the Data field;
u) Recompute the FCS.

Removal of the tag (FDDI tagged frame to native Token Ring/FDDI frame) involves the reverse of this process.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

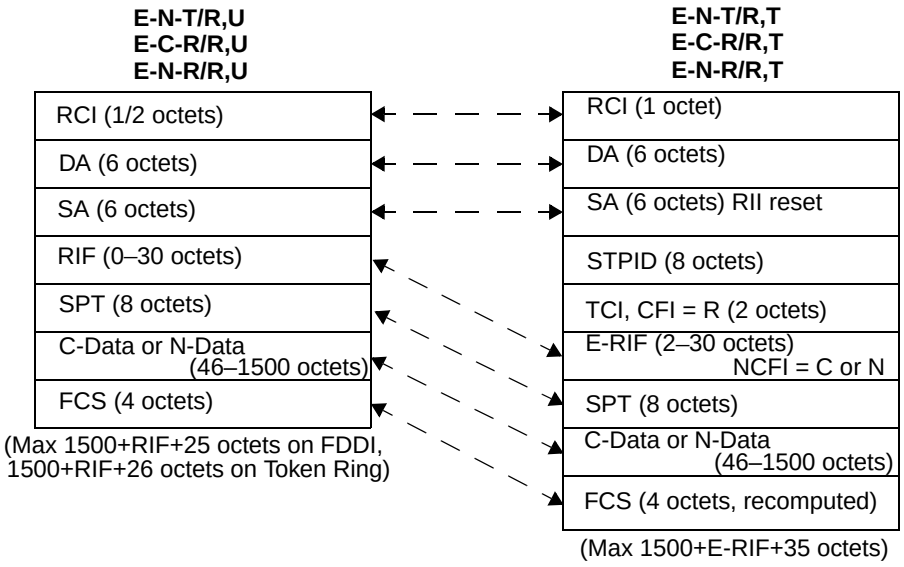


**Figure C-12—Translation between E-X-X/R,U and E-X-X/R,T (transparent FDDI)**

NOTE 3—When removing the tag, if the NCFI indicates that embedded address information is in a form inappropriate for the destination MAC method, then it is necessary either to translate the address information or to discard the frame.

This form of tagging causes the original frame size to be increased by 10 or 12 octets.

### C.5.2.2 LLC-encoded information on Token Ring/FDDI to tagged frame format

Figure C-13 illustrates the translation between an untagged LLC-encoded frame on Token Ring/FDDI (L-C-T/R,U, L-N-T/R,U, L-C-R/R,U, or L-N-R/R,U) and a tagged frame on 802.3/Ethernet (L-C-T/C,T, L-N-T/C,T, L-C-R/C,T, or L-N-R/C,T).

Tagging requires the following frame translations:

a) Remove the RCI field;
b) The DA and SA fields carry the same MAC Addresses as in the original frame, with the RII bit reset;
c) Insert ETPID and TCI, with CFI=C (L-C-T/R,U) or R (all other frame types);
d) If the RII bit was set in the original frame, translate the RIF into the tag header E-RIF. For L-N-T/R,U, create an E-RIF with frame type = transparent. Set the E-RIF NCFI to C or N appropriately;
e) Insert the Len field, with value equal to the number of LLC+Data octets;
f) Copy the LLC field;
g) Copy the Data field;
h) PAD field is inserted if Len is less than 46 (if a minimum tagged frame size of 68 is implemented) or if less than 42 (if a minimum tagged frame size of 64 is implemented);
i) Recompute the FCS.

Removal of the tag involves the reverse of this process.

NOTE 1—When removing the tag, if the CFI/NCFI information indicates that embedded address information is in a form inappropriate for the destination MAC method, then it is necessary either to translate the address information or to discard the frame.

Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

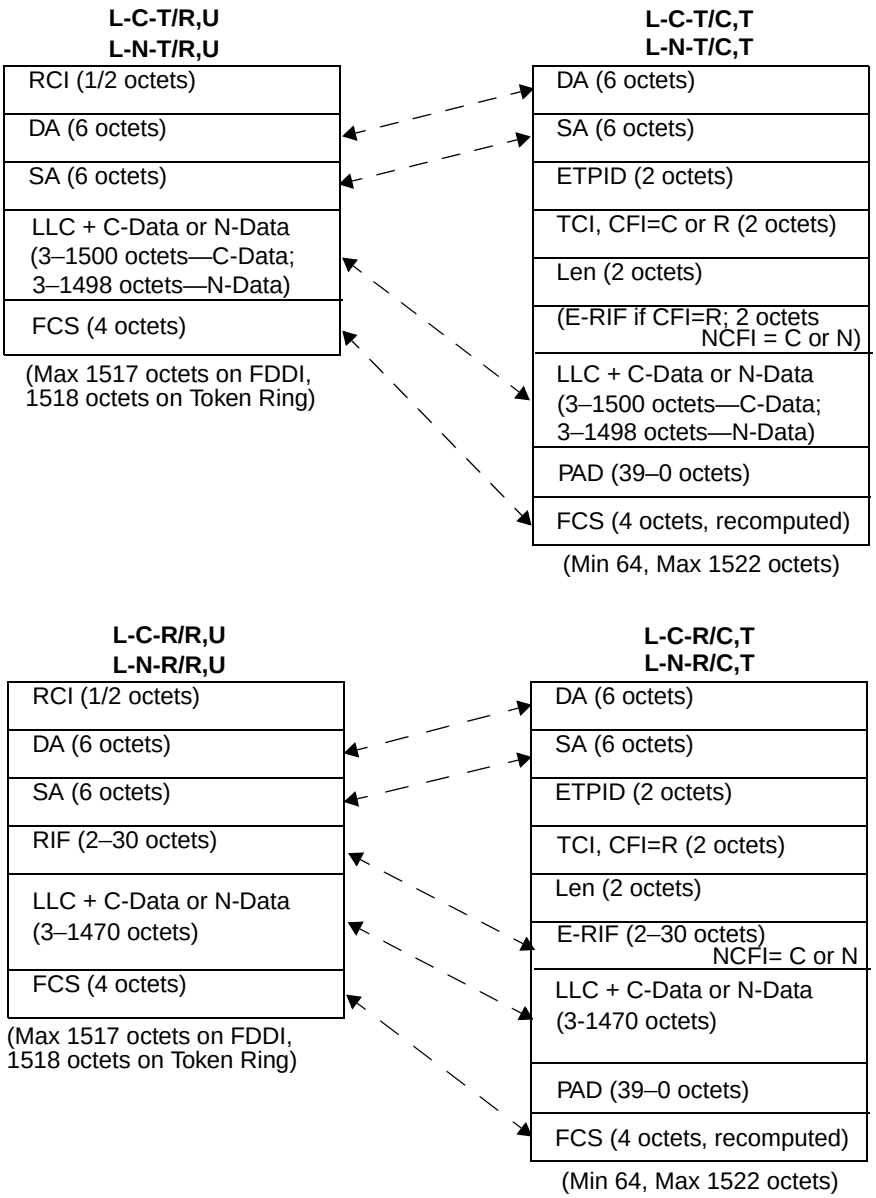


**Figure C-13—Translation between L-X-X/R,U and L-X-X/C,T**

This form of tagging causes the original frame size to be increased by 5 or 6 octets.

Figure C-14 illustrates the translation between an untagged LLC-encoded frame on Token Ring/FDDI (L-C-T/R,U, L-N-T/R,U, L-C-R/R,U, or L-N-R/R,U) and a tagged frame on 8802-5 Token Ring (L-C-T/R,T, L-N-T/R,T, L-C-R/R,T, or L-N-R/R,T). Figure C-14 also illustrates the translation of L-C-T/R,U, L-C-R/R,U, and L-N-R/R,U to tagged frames on FDDI media, the latter two translations illustrating the source-routed form of the tagged frame on FDDI.

Tagging requires the following frame translations:

j) Copy the RCI field;
k) The DA and SA fields carry the same MAC Addresses as in the original frame;


Copyright © 1999 IEEE. All rights reserved.

l) Copy the RIF field if present;
m) Insert STPID and TCI, setting the CFI to N or C appropriately;
n) Copy the LLC field;
o) Copy the Data field;
p) Recompute the FCS.

Removal of the tag (tagged frame to native Token Ring/FDDI frame) involves the reverse of this process.

NOTE 2—When removing the tag, if the CFI indicates that embedded address information is in a form inappropriate for the destination MAC method, then it is necessary either to translate the address information or to discard the frame.

This form of tagging causes the original frame size to be increased by 10 octets for FDDI and Token Ring.

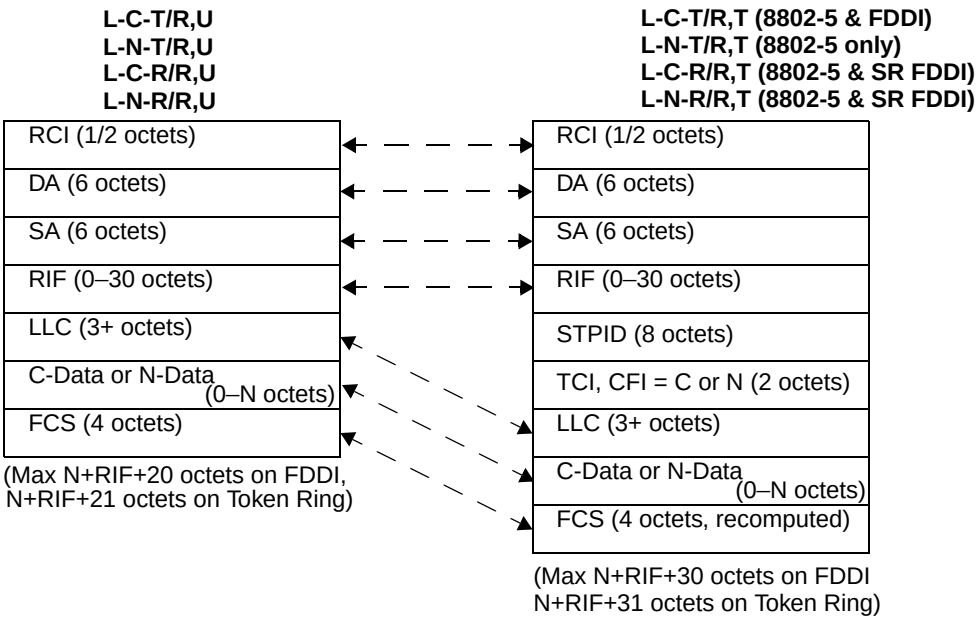


**Figure C-14—Translation between L-X-X/R,U and L-X-X/R,T (8802-5 & SR FDDI)**

Figure C-15 illustrates the translation between an untagged LLC-encoded frame on Token Ring/FDDI (L-N-T/R,U, L-C-R/R,U, or L-N-R/R,U) and a tagged frame on 8802-5 Token Ring (L-C-T/R,T, L-N-T/R,T, L-C-R/R,T, or L-N-R/R,T). Note that the translation of L-C-T/R,U to L-C-T/R,T was dealt with in Figure C-14.

Tagging requires the following frame translations:

q) Copy the RCI field;
r) The DA and SA fields carry the same MAC Addresses as in the original frame, but with RII reset regardless of its state in the original frame;
s) Insert STPID and TCI, setting the CFI to R;
t) Translate the RIF field if present (RII set in source frame) to the E-RIF, otherwise create an E-RIF with RT indicating a transparent frame. Set the NCFI to C or N appropriately;
u) Copy the LLC field;
v) Copy the Data field;
w) Recompute the FCS.

Removal of the tag (tagged frame to native Token Ring/FDDI frame) involves the reverse of this process.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 187 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

NOTE 3—When removing the tag, if the NCFI indicates that embedded address information is in a form inappropriate for the destination MAC method, then it is necessary either to translate the address information or to discard the frame.

This form of tagging causes the original frame size to be increased by 10 or 12 octets.

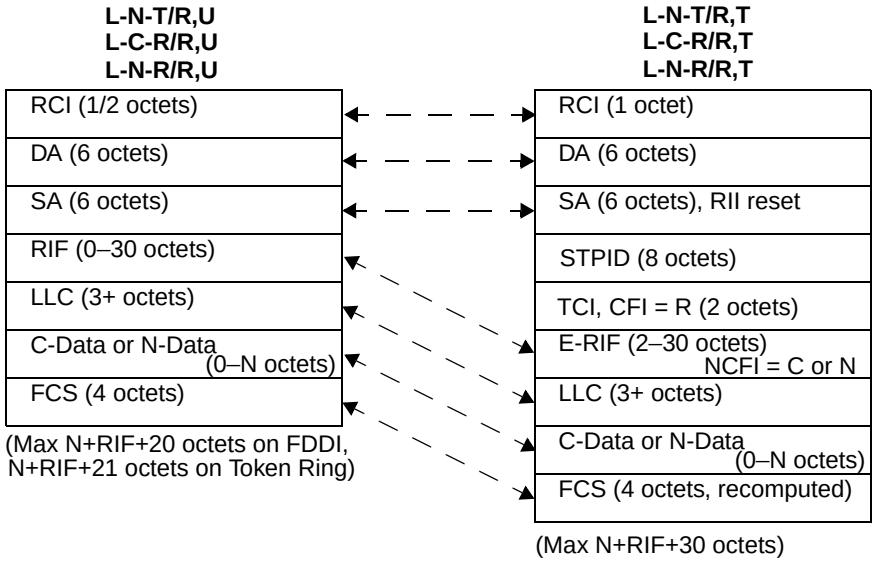


**Figure C-15—Translation between L-X-X/R,U and L-X-X/R,T (transparent FDDI)**

## C.5.3 Translation of tagged frames during relaying

The following subclauses show the frame translations that can occur when a tagged frame is relayed from 802.3/Ethernet to Token Ring/FDDI and vice versa. The translations that occur between the transparent FDDI tagged frame format and the SR form on Token Ring/FDDI are also shown.

### C.5.3.1 Tagged frames carrying Ethernet Type-encoded information

Figure C-16 illustrates the translation of tagged frames carrying Ethernet Type-encoded information between Token Ring/FDDI LANs and 802.3/Ethernet LANs.

Relaying Ethernet Type-encoded tagged frames from Token Ring/FDDI (SR form) to 802.3/Ethernet requires the following frame translations:

a) Remove the RCI field;
b) The DA and SA fields carry the same MAC Addresses as in the original frame, with the RII bit reset;
c) Replace the STPID with an ETPID;
d) The TCI field carries the same VID and Priority values as in the original frame (unless the relay function causes changes to user_priority or VID values). For E-C-T/C,T, CFI = C, otherwise CFI = R;
e) Convert the SPT field to a PT (ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 translation);
f) Copy the RIF, if present, into the tag header E-RIF. Create an E-RIF if the data type being carried is E-N-T/C,T. Set the NCFI in the E-RIF to C or N appropriately;
g) Copy the Data field;
h) Recompute the FCS.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

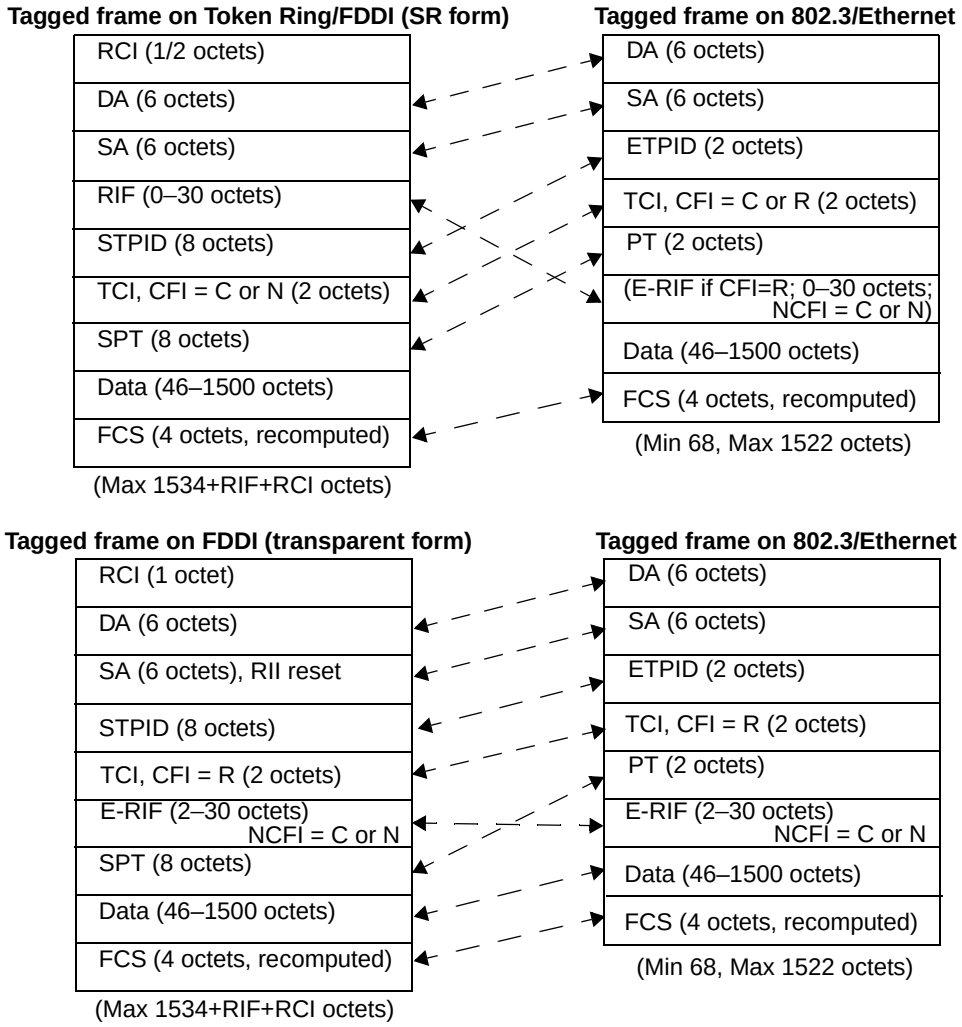


**Figure C-16—Relaying Ethernet Type-encoded tagged frames**

Relaying Ethernet Type-encoded tagged frames from FDDI (transparent form) to 802.3/Ethernet requires the following frame translations:

i) Remove the RCI field;
j) The DA and SA fields carry the same MAC Addresses as in the original frame, with the RII bit reset;
k) Replace the STPID with an ETPID;
l) The TCI field carries the same VID, Priority and CFI values as in the original frame (unless the relay function causes changes to user_priority or VID values);
m) Convert the SPT field to a PT (ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 translation);
n) Copy the E-RIF;
o) Copy the Data field;
p) Recompute the FCS.

Relaying from 802.3/Ethernet to Token Ring/FDDI involves the reverse of these processes.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 189 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### C.5.3.2 Tagged frames carrying LLC-encoded information

Figure C-17 illustrates the translation of tagged frames carrying LLC-encoded information between Token Ring/FDDI LANs and 802.3/Ethernet LANs.

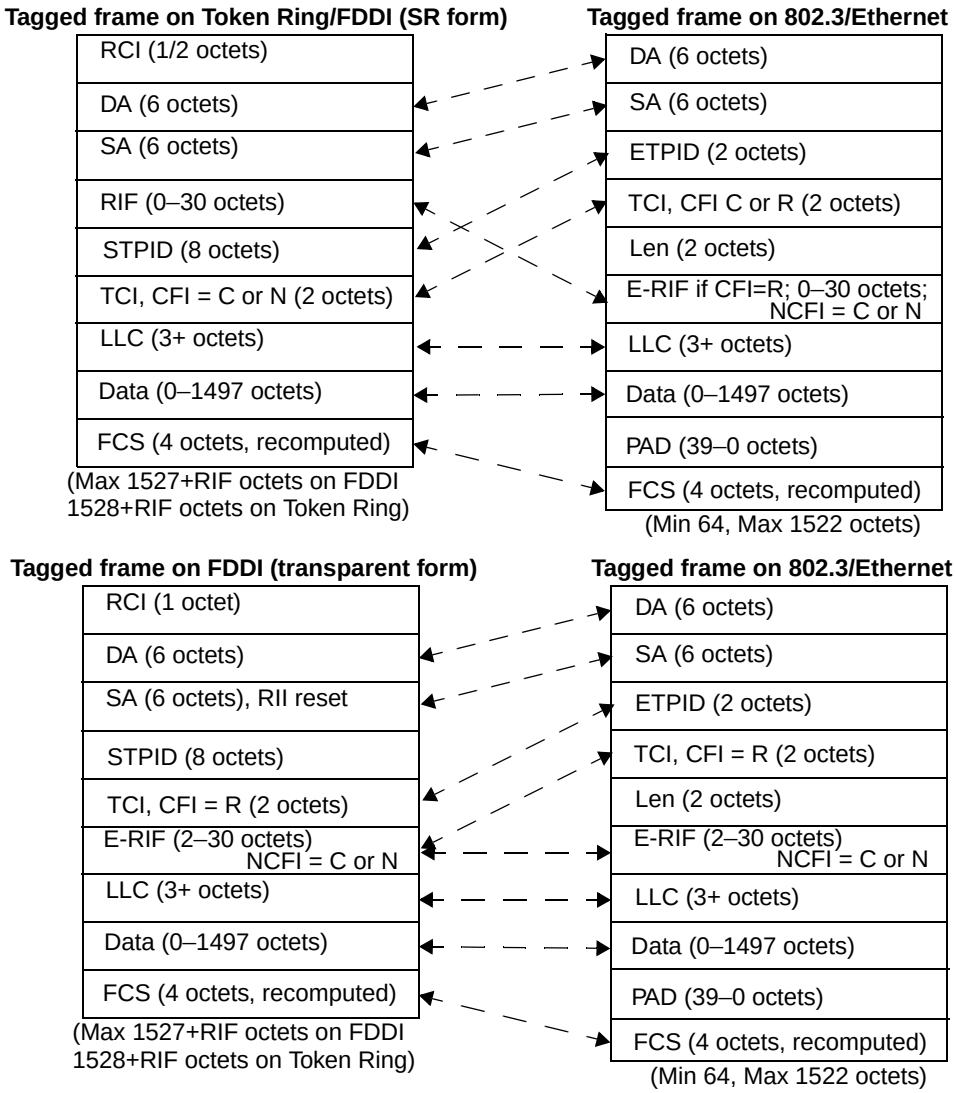


**Figure C-17—Relaying LLC-encoded tagged frames**

Relaying LLC-encoded frames in tagged format from Token Ring/FDDI (SR form) to 802.3/Ethernet requires the following frame translations:

a) Remove the RCI field;
b) The DA and SA fields carry the same MAC Addresses as in the original frame, with the RII bit reset;
c) Replace the STPID with an ETPID;
d) The TCI field carries the same VID and Priority values as in the original frame (unless the relay function causes changes to user_priority or VID values). For L-C-T/C,T, the CFI = C, otherwise CFI = R;
e) Insert a LEN field, equal to LLC+RIF (if present) +Data;


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

f) Copy the RIF, if present, into the tag header E-RIF. Create an E-RIF if the data type being carried is E-N-T/C,T. Set the NCFI bit to C or N appropriately;
g) Copy the LLC field;
h) Copy the Data field;
i) Recompute the FCS.

Relaying LLC-encoded frames in tagged format from FDDI (transparent form) to 802.3/Ethernet requires the following frame translations:

j) Remove the RCI field;
k) The DA and SA fields carry the same MAC Addresses as in the original frame, with the RII bit reset;
l) Replace the STPID with an ETPID;
m) The TCI field carries the same VID, Priority and CFI values as in the original frame (unless the relay function causes changes to user_priority or VID values);
n) Insert a LEN field, equal to E-RIF + LLC +Data;
o) Copy the E-RIF;
p) Copy the LLC field;
q) Copy the Data field;
r) Recompute the FCS.

Relaying from 802.3/Ethernet to Token Ring/FDDI involves the reverse of these process.

### C.5.3.3 Translation between transparent FDDI format and SR format

Figure C-18 illustrates the translation of tagged frames between transparent FDDI format and the corresponding SR format on Token Ring/FDDI LANs. The translation shown applies to X-N-T/R,T, X-C-R/R,T and X-N-R/R,T frames only; other than translation of the RCI field, there is no translation required for X-C-T/R,T frames between Token Ring/FDDI LANs.

Relaying tagged frames from FDDI (transparent form) to Token Ring/FDDI (SR form) requires the following frame translations:

a) Translate the RCI field;
b) The DA and SA fields carry the same MAC Addresses as in the original frame, with the RII bit set or reset to reflect the presence or absence of source-routing information in the E-RIF;
c) Translate source-routing information, if any, from the E-RIF form to the RIF, with the NCFI bit reset;
d) Copy the STPID;
e) The TCI field carries the same VID and Priority values as in the original frame (unless the relay function causes changes to user_priority or VID values). The CFI is set to C or N to match the NCFI value in the E-RIF;
f) Copy the SPT or LLC field;
g) Copy the Data field;
h) Recompute the FCS.

Relaying from Token Ring/FDDI (SR form) to FDDI (transparent form) involves the reverse of these processes.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 191 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

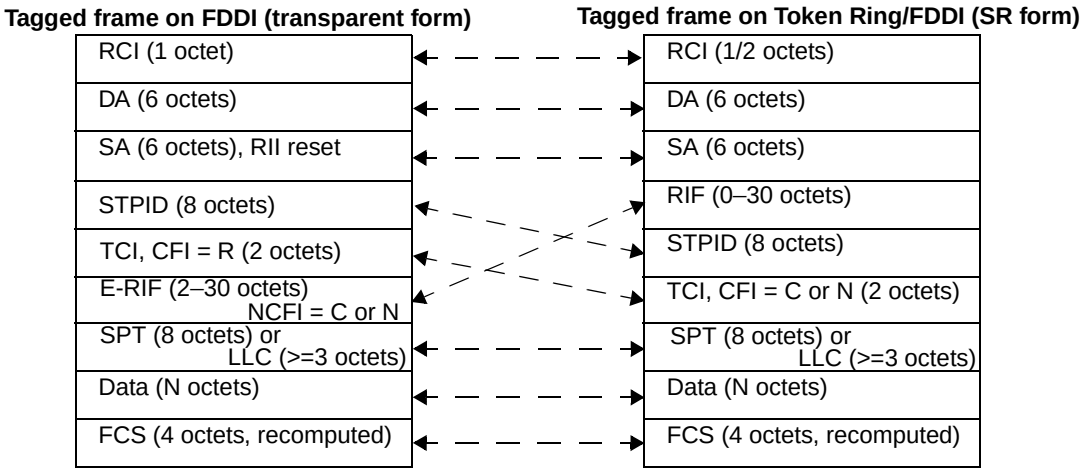


NOTE—Applies to X-N-T/R,T, X-C-R/R,T, and X-N-R/R,T frames only.

**Figure C-18—Relaying tagged frames between transparent and SR forms**

## C.6 Field definitions

Subclauses C.6.1 through C.6.5 describe the field structures that correspond to some of the field names that appear in abbreviated form in the frame format diagrams in this standard.

NOTE—These fields are defined in other standards, and are not part of the additional specification required for the tagged frame format. They are included here in order to simplify the frame descriptions that appear in this standard, not in order to redefine their structure.

### C.6.1 SNAP-encoded Protocol Type

The SNAP-encoded Protocol Type is eight octets in length, encoded in SNAP format. It consists of the standard SNAP header in the first three octets, followed by a SNAP PID consisting of the 00-00-00 OUI, followed by the Ethernet Type value to be encoded, as shown in Figure C-19.

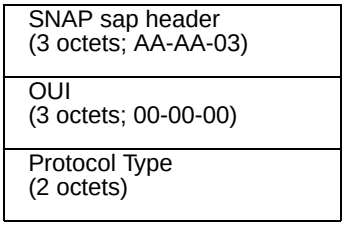


**Figure C-19—SNAP-encoded Protocol Type format**

### C.6.2 Len

This is the IEEE Std 802.3 Length/Type field; for the Length interpretation, it may take any value that is less than or equal to 1500. Values that exceed 1535 are interpreted as Ethernet Types. Values that exceed 1500 but are less than 1535 are undefined.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### C.6.3 C-Data and N-Data

This is the data field of the encapsulated frame:

a) N-Data refers to a data field that is carried in Canonical format regardless of the MAC method carrying the frame;
b) C-Data refers to a data field that is carried in Non-canonical format regardless of the MAC method carrying the frame.

### C.6.4 RIF and E-RIF

The RIF is the Source-Routing Information Field, as defined in ISO/IEC 15802-3, C.3.3.2. If the original (untagged) frame had a RIF, then the RIF field of the tagged frame takes its value.

The E-RIF is a modified form of the RIF that appears within the tag header in tagged frames on transparent LANs (802.3/Ethernet, and FDDI when used as a transparent LAN). The structure of the E-RIF is defined in 9.3.3.

### C.6.5 PAD

Zero or more padding octets, as required in order for the minimum frame size to be at least 64 octets.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 193 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# Annex D

(informative)

# Background to VLANs

The term VLAN has many different definitions throughout the communications industry. The model of VLANs defined in this standard supports most of these views, although it may not be immediately obvious as to how they are supported. The goal of this annex it to take some of the common terms that have been used in the description of VLANs and relate them to the model presented in this standard.

## D.1 Basic VLAN concepts

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

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

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

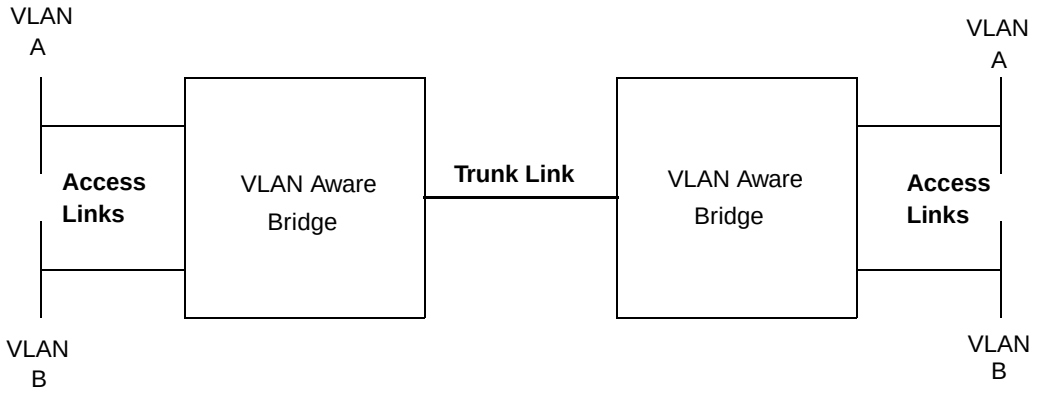


**Figure D-1—Port-based VLANs**

### D.1.1 Trunk Links

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


Copyright © 1999 IEEE. All rights reserved.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### D.1.2 Access Links

The other links in Figure D-1 are commonly known as *Access Links*. An Access Link is a LAN segment used to multiplex one or more VLAN-unaware devices into a Port of a VLAN Bridge. In simple terms this is an 802 LAN segment (IEEE Std 802.3, ISO/IEC 8802-5, etc.) with end stations attached, that is connected into a VLAN-aware Bridge. All frames on an Access Link carry no VLAN identification; i.e., there are no VLAN-tagged frames on an Access Link. Typically the Access Link is viewed as being on the edge of the VLAN network. The Access Link itself could consist of a number of LAN segments interconnected by ISO/IEC 15802-3-conformant Bridges (this is termed a *legacy region* in E.1.2). Like the Access Link, there are no VLAN-tagged frames transmitted in this legacy region.

### D.1.3 Hybrid Links

When VLAN-unaware end stations are added to a Trunk Link, the resultant link is commonly known as a Hybrid Link. A Hybrid Link is a LAN segment that has both VLAN-aware and VLAN-unaware devices attached to it. Consequently, a Hybrid Link can carry both VLAN-tagged frames and other (untagged or priority-tagged) frames. It must be borne in mind that, for a given VLAN, all frames transmitted by a given Bridge on a given hybrid link must be tagged the same way on that link. They must be either

a) All untagged; or
b) All tagged, carrying the same VLAN ID.

Note that a Bridge can transmit a mix of VLAN-tagged frames and untagged frames but they must be for different VLANs. In Figure D-2 all the frames for VLANs A and B are tagged on the hybrid link. All frames for VLAN C on the hybrid link are untagged.

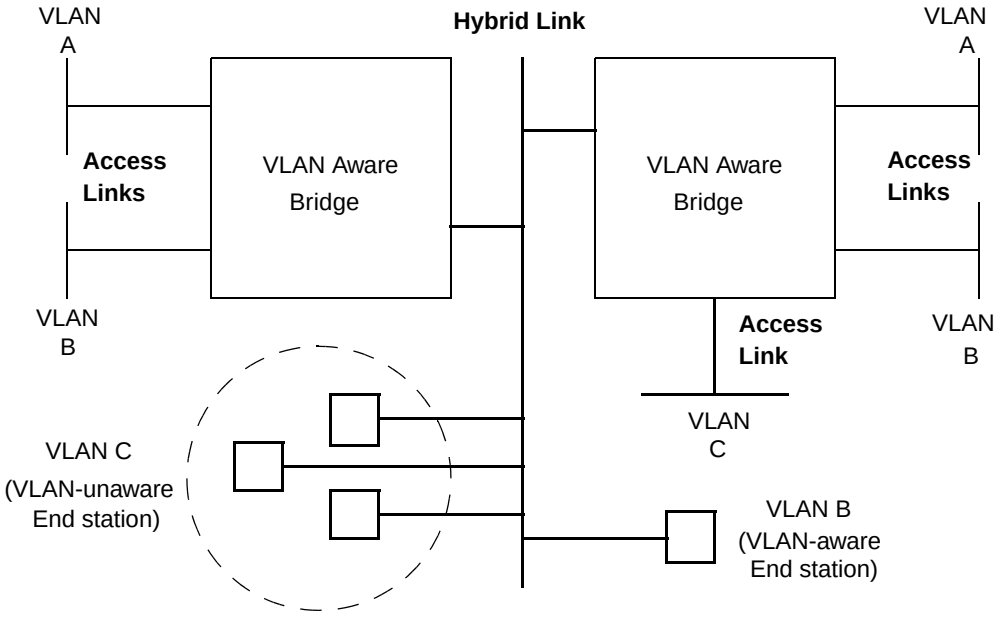


**Figure D-2—Hybrid Links**

On a Hybrid link the decision to tag or not to tag a frame is a function of the VLAN and not a function of the link itself, since both formats are allowed. The Hybrid link can be thought of as the general case of both Access and Trunk links.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 195 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

## D.2 Relationship with the Port-based VLAN model

### D.2.1 Link types

A Hybrid link in which all frames are VLAN-tagged is a Trunk Link. Conversely, a Hybrid Link that has no VLAN-tagged frames is an Access Link. The distinction between Access and Trunk becomes less important in an actual implementation where all types of frames have to be handled. The Acceptable frame Types parameter (8.4.3) allows control to be exerted over the reception of frames that do not carry VLAN identification information (i.e., untagged and priority-tagged frames). Depending upon the value of the parameter, the decision is taken as to whether to discard the frame or to add in VLAN information. The more general implementation will allow all types of frames (VLAN-tagged, priority-tagged, and untagged) to be present, hence all links are conceptually Hybrid Links.

The VLAN model defined in this standard essentially takes the view that all links are Hybrid Links. It is then up to the system administrator to determine, through appropriate application of the management functions available in the Bridge, whether all links remain operating as Hybrid Links, or whether particular links need to be configured as Access Links or Trunk Links, in other words, to explicitly configure some Ports to discard untagged and priority-tagged frames. This allows the description of the functionality of the Bridge to be kept relatively simple, while retaining the practical implementation and configuration mechanisms necessary in order for the other link types to be derived. The Acceptable frame Types parameter does not allow configuration of a Port such that VLAN-tagged frames are discarded; hence, from the practical point of view, the distinction between an Access Link and a Hybrid Link is not one of Port configuration; it is simply determined by the presence or absence of other devices on that link that generate VLAN-tagged frames.

### D.2.2 Use of other VLAN styles

This standard defines a Port-based tagging rule, whereby all untagged and priority-tagged frames received by a Port are classified as belonging to the VLAN whose VID (the PVID) is associated with that Port. This Port-based style of operation should be viewed as the base level of a possible hierarchy of VLAN styles, each one able to classify untagged frames according to particular ingress rules. Examples of such ingress rules might include

a) MAC Address-based classification; e.g., associating a set of MAC Addresses with a given VLAN ID in order to define the membership of the VLAN;
b) Protocol-based classification; e.g., allocating VLAN membership on the basis of the higher-layer protocol information carried in the frame;
c) Subnet-based classification; e.g., allocating VLAN membership on the basis of IP subnet addressing characteristics of frames.

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

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

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

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


Copyright © 1999 IEEE. All rights reserved.

d) As proprietary extensions to the existing specification;
e) As future standardized extensions.

Given that the starting point for this standard is that all links are Hybrid Links, there is no need for such additional classification and tagging functionality to exist within the Bridges themselves; it would, for example, be possible to develop “tagging engines” that are capable of implementing more complex classifications than Port-based classification, and which are placed between the 802.1Q Bridge Port and an Access Link. Such a device would provide a richer functionality in terms of VLAN classification style, while remaining compatible with Port-based VLAN operation.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 197 of 211

# Annex E

(informative)

# Interoperability considerations

VLAN-aware Bridges that conform to this standard are able to interoperate in Bridged LANs with other VLAN-aware Bridges. However, the VLAN-based filtering service defined in this standard, as provided in the context of a single spanning tree for the Bridged LAN, involves some constraints on the network topology and individual device configurations that differ from the set of constraints that apply to the building and configuration of Bridged LANs based only on ISO/IEC 15802-3.

In addition, VLAN-aware Bridges are able to interoperate with Bridges conformant with the ISO/IEC 15802-3 specification (or with the earlier ISO/IEC 10038 specification), as well as with both priority-aware and VLAN-aware end systems. Both the VLAN based filtering service and the tag insertion and removal service of 802.1Q cause constraints on intermixed network topologies and device configurations that again differ from the building and configuration of ISO/IEC 15802-3 standard networks.

The implications of certain device configurations may not be immediately apparent from the technical detail of this standard. In order to clarify the nature of the additional constraints, the following subclauses

a) Describe the basic requirements for interoperability;
b) Discuss those requirements in the context of homogeneous and heterogeneous configurations, with examples of some of the problems that can occur if these requirements are not adhered to.

## E.1 Requirements for interoperability

There are two primary aspects of the configuration of a Bridged LAN that are of concern from the point of view of interoperability:

a) Establishing a consistent view of the static filtering configuration of Bridges in the Bridged LAN;
b) Ensuring that untagged frames are VLAN-tagged (and that the tag is subsequently removed) consistently regardless of Spanning Tree reconfigurations.

### E.1.1 Static filtering requirements

Static filtering controls allow the network administrator to impose a level of control over the permitted connectivity in the Bridged LAN, by setting static MAC Address filters in the Filtering Databases of Bridges, and by controlling the extent of particular VLANs by manipulation of Static VLAN Registration Entries (8.11.2).

In order to ensure that end station to end station connectivity (or the lack of it) is consistent in all possible Spanning Tree configurations, any static filters need to be established taking account of the full mesh topology of the physical interconnections between Bridges in the Bridged LAN, not just the “normal” Spanning Tree topology to which the network is configured when all Bridges and LAN segments are operating correctly. An example of the consequences of failure to establish consistent controls for static VLAN filtering is given in E.2.1.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### E.1.2 Configuration requirements for VLAN-tagging

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

a) All VLAN-aware Bridge Ports that are connected to the same LAN segment apply a consistent set of ingress rules (8.6);
b) All VLAN-aware Bridge Ports that are connected to the same *legacy region* of a Bridged LAN apply a consistent set of ingress rules;
c) All VLAN-aware Bridge Ports that serve LAN segments to which members of the same VLAN are (or can be) attached apply a consistent set of ingress rules.

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

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

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

## E.2 Homogenous 802.1Q Bridged LANs

This standard requires new considerations in building a Bridged LAN in which all Bridges are VLAN-aware. The arbitrary plug and play capability of ISO/IEC 15802-3 in creating a network topology is restricted when making use of the VLAN extensions defined in this standard.

### E.2.1 Consistency of static VLAN filtering

In order for stations that are members of a given VLAN to be able to reach other members of the same VLAN elsewhere in the Bridged LAN, all Ports that are part of the Spanning Tree active topology (i.e., all Ports that are in a forwarding state) connecting the stations must be included in the Member Set (8.11.9) for the given VLAN. In order for this connectivity to be independent of any reconfiguration of the Spanning Tree topology, all paths among those stations, both forwarding and blocked, must have this characteristic. Use of management controls to manipulate the Member Set (e.g., filters for security) must be applied in a manner consistent with requirements of the full mesh topology of the Bridged LAN.

An inconsistency occurs, for example, if a VLAN is restricted from an active path, but not from a redundant path currently blocked by the operation of Spanning Tree. Should a Spanning Tree reconfiguration enable the previously blocked path, the restriction will no longer be in place. In the reverse, a Spanning Tree recon-

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 199 of 211

figuration may suddenly impose a restriction that had not previously existed. A common use of such management restriction will likely arise from managers who make use of an “access” port construct. An access port may be a port which is absent from the Member Set (8.11.9) in all VLANs but the untagged, default VLAN. Should such an access port become the active connection between two portions of the Bridged LAN as a result of a Spanning Tree reconfiguration, all VLANs but that one will be partitioned at that point in the topology.

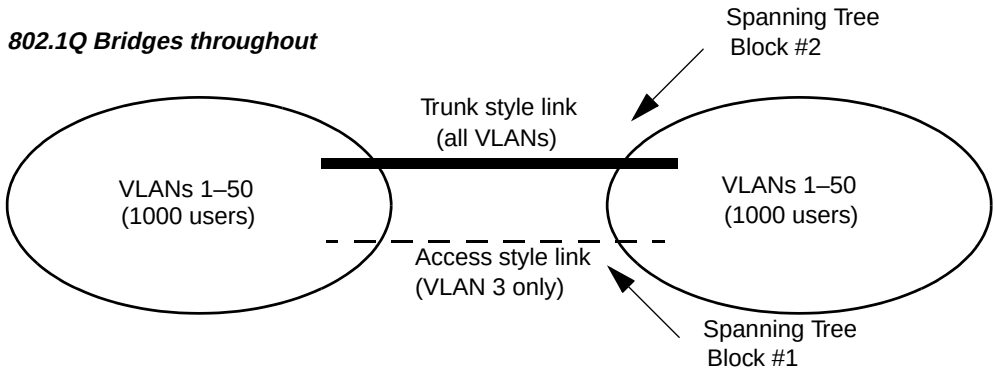


**Figure E-1—Static filtering inconsistency**

In Figure E-1, the trunk style link and access style link cause a loop through the left and right portions of the network. STP will block one or the other. Should the Spanning Tree block at point #1, all 2000 users may communicate on any of the 50 VLANs. However, should the Spanning Tree block at point #2, the left and right portions of the network will be partitioned on all VLANs excepting VLAN 3 (the VLAN carried via the access style link.)

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

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

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

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

This creates a configuration challenge for installers of Bridges that conform to this standard. Initial management configuration of the Bridges (the setting of PVIDs) must be made consistent among the Bridges, in a manner that takes into account the actual physical topology. Changes to the physical topology may require


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

specific changes to the configuration of all affected switches. These requirements effectively disallow a plug-and-play installation as supported by ISO/IEC 15802-3 Bridged LANs, unless all Bridges are left with their default PVID configuration of PVID = 1.

## E.3 Heterogeneous Bridged LANs: intermixing ISO/IEC 15802-3 (D) and 802.1Q (Q) Bridges

This clause discusses networks in which VLAN-aware Bridges that conform to this standard are intermixed with VLAN-unaware Bridges conformant to the ISO/IEC 15802-3 standard.

A principal limitation in intermixing Q Bridges with D Bridges is that the VLAN filtering services are not universally available throughout the Bridged LAN. Also, services for the insertion and removal of tags are not universally available. Further, spanning tree reconfigurations may cause filtering services, as well as tag insertion and removal services, to become available or become unavailable independent of actions of affected users.

### E.3.1 Example: Adding an 802.1Q Bridge to provide filtering to an ISO/IEC 15802-3 network

Example problems can be shown with the following topology diagrams. Figure E-2 includes one Q Bridge and two D Bridges:

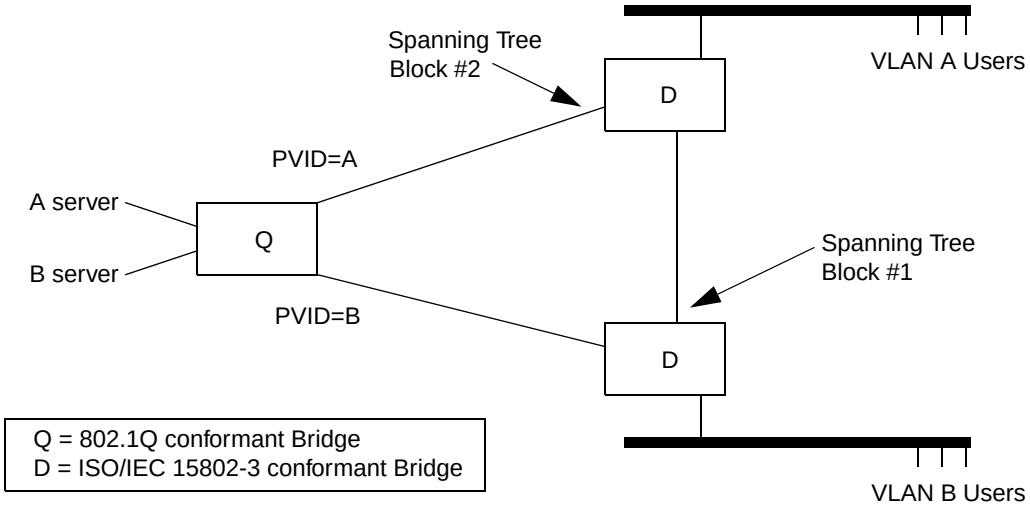


**Figure E-2—Interoperability with ISO/IEC 15802-3 Bridges: example 1**

If the Spanning Tree protocol determines to break the loop among the three Bridges by blocking at point #1, connectivity within each VLAN is as desired. However, should the block occur at point #2, traffic from VLAN A users will pass through both D Bridges, and be treated as VLAN B traffic upon arrival in the Q Bridge. Connectivity to the A server will be lost for the A users.

Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

### E.3.2 Example: Adding an ISO/IEC 15802-3 Bridge to a (previously) Homogenous 802.1Q Network

A similar problem, demonstrating the impact of placing a D Bridge within an otherwise homogenous Q topology, can be shown by the configuration in Figure E-3. Here we include two Q Bridges and add a single redundant D Bridge:

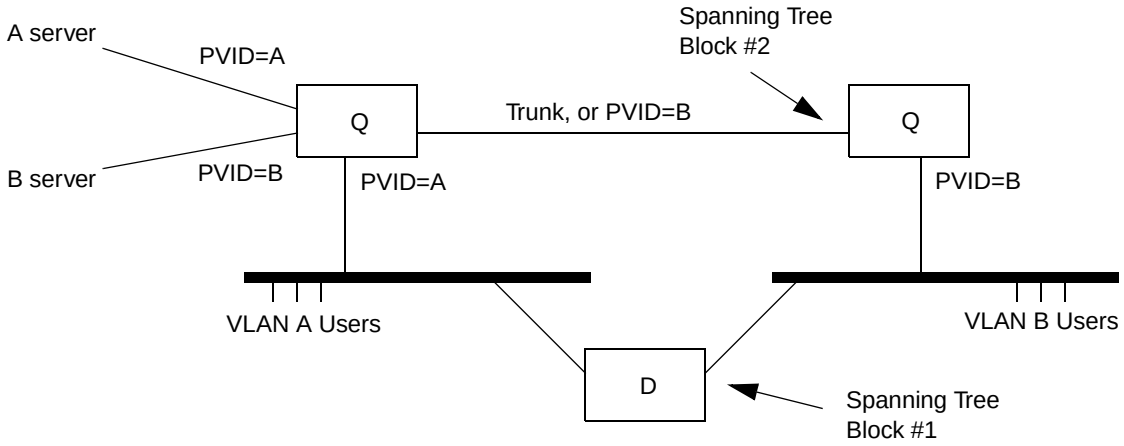


**Figure E-3—Interoperability with ISO/IEC 15802-3 Bridges: example 2**

If STP determines to break the loop among the three Bridges by blocking at point #1, connectivity within each VLAN is as desired. The two Q switches operate as expected. A and B VLAN frames are VLAN-tagged on arrival in either Q Bridge, and forwarded only to the appropriate servers. Now suppose an STP reconfiguration results in a block at point #2, but not at #1. This redirects VLAN B user traffic through the ISO/IEC 15802-3 Bridge. VLAN B users no longer have their traffic identifiably distinct from VLAN A. An immediate consequence is that the VLAN B users will no longer have access to the “B server.”

## E.4 Heterogeneous Bridged LANs: intermixing ISO/IEC 11802-5 and 802.1Q Bridges

Translating Bridges (i.e., Bridges that relay between Token Ring/FDDI and 802.3/Ethernet LANs) that implement the encapsulation techniques described in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390 can be intermixed with 802.1Q Bridges under certain limited conditions. In order to understand the limitations involved, it is necessary to describe what happens to the various tagged frame formats when passed through a Translating Bridge. The frame formats shown in the following subclauses use the notation that is defined in Annex C.

NOTE—The examples shown are not exhaustive; in particular, the examples do not make use of the transparent tagged frame format on FDDI. However, the examples illustrate the nature of the problems that can occur with such translations.

### E.4.1 LLC-encoded tagged frames relayed from 802.3/Ethernet to Token Ring or source-routed FDDI

A Transparent LLC-encoded frame on 802.3/Ethernet has the following form:

**L-C-T/C,T:** DA, SA, ETPID, TCI (CFI reset), LEN, LLC, C-Data, PAD, FCS


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

The Translating Bridge will recognize the ETPID as an Ethernet Type that requires translation into an SPT. This effectively translates the ETPID to an STPID. The translated frame therefore appears as follows:

RCI, DA, SA (RII reset), STPID, TCI (CFI reset), LEN, LLC, C-Data, PAD, FCS

Whereas a true translation of the tagged frame via an 802.1Q Bridge would result in

**L-C-T/R,T:** RCI, DA, SA (RII reset), STPID, TCI (CFI reset), LLC, C-Data, FCS

As can be seen, the resultant frame superficially appears to be a valid tagged frame; however, it includes spurious LEN and PAD fields, and is therefore not a valid tagged frame.

A source-routed LLC-encoded frame on 802.3/Ethernet has the following form:

**L-C-R/C,T:** DA, SA, ETPID, TCI (CFI set), LEN, RIF (NCFI=C), LLC, C-Data, PAD, FCS

The Translating Bridge generates

RCI, DA, SA (RII reset), STPID, TCI (CFI set), LEN, RIF (NCFI=C), LLC, C-Data, PAD, FCS

Whereas a true translation of the tagged frame via an 802.1Q Bridge would result in

**L-C-R/R,T:** RCI, DA, SA (RII set), RIF, STPID, TCI (CFI reset), LLC, C-Data, FCS

Again, the resultant frame superficially appears to be a valid tagged frame; however, it includes spurious LEN, RIF and PAD fields, the RIF information is not visible to any source routing Bridges attached to the Ring medium, and the CFI indicates Non-canonical data where the actual data is Canonical.

NOTE—Similar problems exist for the other two LLC-encoded formats, L-N-R and L-N-T, but the CFI correctly indicates Non-canonical data in the translated frame.

In both cases, the effect of the Translating Bridge is symmetrical; passing this invalid frame back through the Translating Bridge restores it to a valid tagged frame format on 802.3/Ethernet.

Consequently,

a) The translated frames cannot be correctly interpreted by VLAN-aware end stations attached to the Ring medium;
b) In both cases, any 802.1Q Bridge that attempts to untag the translated frames will generate invalid untagged frames;
c) Any 802.1Q Bridge that attempts to relay the translated frames back onto Token Ring/FDDI will generate invalid tagged frames;
d) In the case of the source-routed frame, any source routing Bridges attached to the Ring will treat the frame as a transparent frame;
e) As the effect of the Translating Bridge is symmetric, passing the frame through an even number of such translations before any 802.1Q device attempts to interpret the frame format results in correct operation of the 802.1Q devices.

### E.4.2 Ethernet Type-encoded tagged frames relayed from 802.3/Ethernet to Token Ring or source-routed FDDI

A Transparent Ethernet Type-encoded frame on 802.3/Ethernet has the following form:

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 203 of 211

**E-C-T/C,T:** DA, SA, ETPID, TCI (CFI reset), PT, C-Data, FCS

The Translating Bridge generates:

RCI, DA, SA (RII reset), STPID, TCI (CFI reset), PT, C-Data, FCS

Whereas a true translation of the tagged frame via an 802.1Q Bridge would result in

**E-C-T/R,T:** RCI, DA, SA (RII reset), STPID, TCI (CFI reset), SPT, C-Data, FCS

The resultant frame is superficially a valid tagged frame, but carries an untranslated Ethernet Type where there should be a SNAP-encoded Ethernet Type.

A source-routed Ethernet Type-encoded frame on 802.3/Ethernet has the following form:

**E-C-R/C,T:** DA, SA, ETPID, TCI (CFI set), PT, RIF (NCFI=C), C-Data, FCS

The Translating Bridge generates

**RCI,** DA, SA, STPID, TCI (CFI set), PT, RIF (NCFI=C), C-Data, FCS

Whereas a true translation of the tagged frame via an 802.1Q Bridge would result in

**E-C-R/R,T:** RCI, DA, SA (RII set), RIF, STPID, TCI (CFI reset), SPT, C-Data, FCS

Again, the resultant frame is superficially a valid tagged frame, but carries an untranslated Ethernet Type where there should be a SNAP-encoded Ethernet Type, and carries a spurious RIF field. The CFI is also incorrect.

NOTE—Similar problems exist for the other two Ethernet Type-encoded formats, E-N-R and E-N-T, but the CFI correctly indicates Non-canonical data in the translated frame.

In both cases, the effect of the Translating Bridge is symmetrical; passing this invalid frame back through the Translating Bridge restores it to a valid tagged frame format on 802.3/Ethernet.

Consequently,

a) The translated frames cannot be correctly interpreted by VLAN-aware end stations attached to the Ring medium;
b) Any 802.1Q Bridge that attempts to untag the translated frames will generate invalid untagged frames;
c) Any 802.1Q Bridge that attempts to relay the translated frames back onto Token Ring/FDDI will generate invalid tagged frames;
d) In the case of the source-routed frame, any source routing Bridges attached to the Ring will treat the frame as a transparent frame;
e) As the effect of the Translating Bridge is symmetric, passing the frame through an even number of such translations before any 802.1Q device attempts to interpret the frame format results in correct operation of the 802.1Q devices.

### E.4.3 LLC-encoded tagged frames relayed from Token Ring or source-routed FDDI to 802.3/Ethernet

A Transparent LLC-encoded frame on Token Ring/FDDI has the following form:


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

**L-C-T/R,T:** RCI, DA, SA (RII reset), STPID, TCI (CFI reset), LLC, C-Data, FCS

The Translating Bridge generates

DA, SA, ETPID, TCI (CFI reset), LLC, C-Data, PAD, FCS

Whereas a true translation of the tagged frame via an 802.1Q Bridge would result in

**L-C-T/C,T:** DA, SA, ETPID, TCI (CFI reset), LEN, LLC, C-Data, PAD, FCS

As can be seen, the resultant frame superficially appears to be a valid tagged frame; however, it is missing the LEN field, and is therefore not a valid tagged frame.

A source-routed LLC-encoded frame on Token Ring/FDDI has the following form:

**L-C-R/R,T:** RCI, DA, SA (RII set), RIF, STPID, TCI (CFI reset), LLC, C-Data, FCS

The Translating Bridge generates

DA, SA, ETPID, TCI (CFI reset), LLC, C-Data, PAD, FCS

Whereas a true translation of the tagged frame via an 802.1Q Bridge would result in

**L-C-R/C,T:** DA, SA, ETPID, TCI (CFI set), LEN, RIF (NCFI=C), LLC, C-Data, PAD, FCS

Again, the resultant frame superficially appears to be a valid tagged frame, but is missing the LEN field. The RIF information has been lost.

Similar problems exist for the other two LLC-encoded formats, L-N-R and L-N-T, but in addition, the CFI will be set, indicating the presence of a RIF in the tag header when no such RIF field exists.

In both cases, the effect of the Translating Bridge is almost symmetrical; passing the invalid frame back through the Translating Bridge restores it to a valid tagged frame format on 802.3/Ethernet, but with the loss of any source-routing information that may have been present, and the inclusion of a PAD field if the original frame on the Ring medium had been small.

Consequently,

a) The translated frames cannot be correctly interpreted by VLAN-aware end stations attached to the Ring medium;
b) Any 802.1Q Bridge that attempts to untag the translated frames will generate invalid untagged frames;
c) Any 802.1Q Bridge that attempts to relay the translated frames back onto Token Ring/FDDI will generate invalid tagged frames;
d) Any source-routing information is lost;
e) As the effect of the Translating Bridge is almost symmetric (the RIF is lost, and there may be a PAD included), passing the frame through an even number of such translations before any 802.1Q device attempts to interpret the frame format results in correct operation of the 802.1Q devices, as long as those devices are not sensitive to the presence of spurious PAD information.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 205 of 211

### E.4.4 Ethernet Type-encoded tagged frames relayed from Token Ring or source-routed FDDI to 802.3/Ethernet

A Transparent Ethernet Type-encoded frame carrying Canonical data on Token Ring/FDDI has the following form:

**E-C-T/R,T:** RCI, DA, SA (RII reset), STPID, TCI (CFI reset), SPT, C-Data, FCS

The Translating Bridge generates

DA, SA, ETPID, TCI (CFI reset), SPT, C-Data, FCS

Whereas a true translation of the tagged frame via an 802.1Q Bridge would result in

**E-C-T/C,T:** DA, SA, ETPID, TCI (CFI reset), PT, C-Data, FCS

As can be seen, the resultant frame superficially appears to be a valid tagged frame; however, the SPT has not undergone translation to a PT. End stations encountering this frame on 802.3/Ethernet would only be capable of interpreting it if they were able to recognize the SPT as an embedded Ethernet Type.

Translating Canonical source-routed Ethernet Type-encoded information produces similar results, but with the additional loss of the source-routing information.

A Transparent Ethernet Type-encoded frame carrying Non-canonical information on Token Ring/FDDI has the following form:

**E-N-T/R,T:** RCI, DA, SA (RII reset), STPID, TCI (CFI set), SPT, N-Data, FCS

The Translating Bridge generates:

DA, SA, ETPID, TCI (CFI set), SPT, N-Data, FCS

Whereas a true translation of the tagged frame via an 802.1Q Bridge would result in:

**E-N-T/C,T:** DA, SA, ETPID, TCI (CFI set), PT, RIF (NCFI=N), N-Data, FCS

Again, the resultant frame superficially appears to be a valid tagged frame, but the SPT has not been translated to a PT, and the RIF is not present in the tag header. Any 802.1Q device will therefore interpret the first N octets of the N-Data field as if it was the RIF.

Translating Non-canonical source-routed Ethernet Type-encoded information produces similar results, but with the additional loss of the source-routing information.

In both cases, the effect of the Translating Bridge is almost symmetrical; passing this invalid frame back through the Translating Bridge restores it to a valid tagged frame format on 802.3/Ethernet, but with the loss of any source-routing information that may have been present.

Consequently,

a) The translated frames cannot be correctly interpreted by VLAN-aware end stations attached to the Ring medium;
b) Any 802.1Q Bridge that attempts to untag translated frames carrying Non-canonical information will generate invalid untagged frames;


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

c) Any 802.1Q Bridge that attempts to relay untag translated frames carrying Non-canonical information back onto Token Ring/FDDI will generate invalid tagged frames;
d) Frames carrying Canonical information can be successfully untagged or relayed in tagged form using other MAC methods, as long as the 802.1Q Bridge is capable of correctly handling the embedded SPT;
e) Any source-routing information is lost;
f) As the effect of the Translating Bridge is symmetric (apart from the loss of RIF), passing the frame through an even number of such translations before any 802.1Q device attempts to interpret the frame format results in correct operation of the 802.1Q devices.

### E.4.5 Conclusions

Except for the limited case of relaying Canonical Ethernet Type-encoded information from 802.3/Ethernet to Token Ring/FDDI, the translation that a tagged frame undergoes when passing through a Translating Bridge renders the frame uninterpretable by any 802.1Q device (either end station or Bridge), unless the frame has passed through an even number of Translating Bridges before the 802.1Q device attempts to interpret the frame. In regions where an odd number of translations have occurred, source-routing information is rendered invisible to source routing Bridges in some cases. In the other cases, source-routing information is lost after the first translation.

Consequently, the use of Translating Bridges intermixed with 802.1Q Bridges is feasible only if

a) An even number of translations (or zero translations) is experienced by any tagged frame that is transmitted between any pair of 802.1Q-aware devices in the Bridged LAN;
b) The loss of source routing capability across some regions of the Bridged LAN is acceptable; specifically across regions where the first Translating Bridge encountered by a correctly formatted tagged frame will relay the frame from Token Ring/FDDI to 802.3/Ethernet;
c) End stations are not sensitive to receiving LLC-encoded frames that have PAD octets added to the LLC user data.

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

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

a) An All Groups entry in the Filtering Database, specifying Registration Fixed on all Ports, and
b) The GMRP Protocol Administrative Control parameters set to disable GMRP on all Ports.

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

Placing ISO/IEC 15802-3 Bridges around the periphery of an 802.1Q-based Bridged LAN works correctly, as long as, for a given ISO/IEC 15802-3 Bridge, the 802.1Q Bridges connected to the same segment(s) are configured to untag any VLANs that are relevant to the GMRP operation of the ISO/IEC 15802-3 Bridge.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 207 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

The ISO/IEC 15802-3 Bridge generates untagged GMRP frames, which the 802.1Q Bridges classify according to the value of the PVID for the reception Port; in a simple configuration of the 802.1Q Bridges, the Ports that connect to the ISO/IEC 15802-3 Bridge are configured for the PVID VLAN to be untagged on egress.

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

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

## E.6 Intermixing 802.1Q Version 1.0 Bridges with future 802.1Q Bridges

The discussion above on intermixing Q Bridges with D Bridges has a direct analogue in the plan to provide a simple VLAN standard (Q version 1.0) initially, and later to provide extensions to the standard (Q version 2.0 on) which extends the VLAN capabilities to support more sophisticated ingress rules for frame classification. Some of the topology restrictions will probably be similar to the “Q intermixed with D” cases.

### E.6.1 Example: Intermixing Layer 3 Ingress Rules

Consider the case where a “Qv2” switch capability is specified allowing for classification of frames by protocol. This would allow support for IP and IPX as distinct VLANs. The following diagram might apply when a Qv2 Bridge is added to a version 1.0 topology to allow users of two protocols to participate in two separate VLANs:

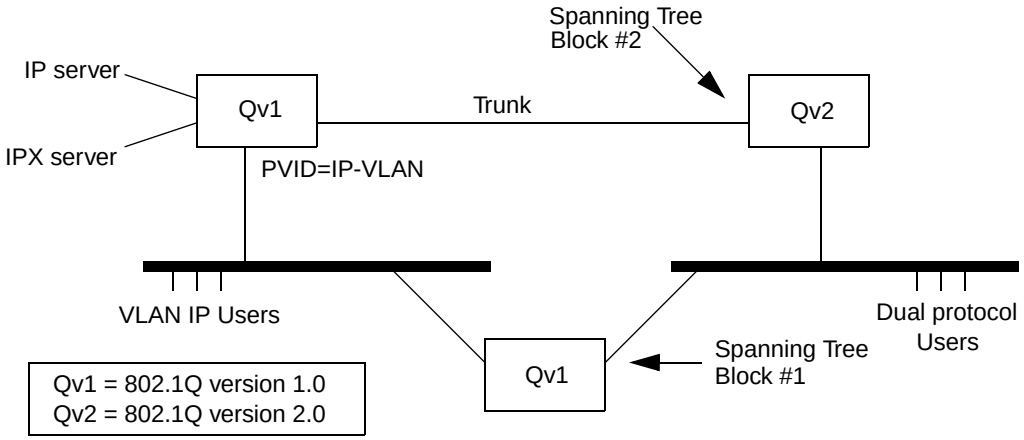


**Figure E-4—Interoperability between Q versions 1 and 2**

Consider this network, when STP has blocked at point #1, and not at #2. The upper Qv1 switch operates as expected, and the Qv2 switch provides protocol based classification for the frames received from the dual protocol users. IP and IPX VLAN frames are VLAN-tagged on the trunks to and from the uppermost Bridge and servers. But if a STP reconfiguration should result in a block at point #2, but not at #1, activating traffic through the lower Qv1 Bridge, dual protocol users will have all their traffic treated as part of the IP-VLAN. An immediate consequence is that the uppermost Bridge will no longer provide them access to the “IPX server.”


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

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

Further challenges arise when one considers the case where several Q Bridges, some of version 1 and some of version 2, all attach to the same LAN segment. Again, the rule that any given frame exists in exactly one VLAN requires that all of these Bridges be configured with the same ingress rules. In this case, the Q1 Bridges will provide a least common capability, and further require common configuration of the PVID.

### E.6.3 Interoperability with 802.1Q Version 2.0 offering multiple spanning trees

Several very different architectures for multiple spanning tree support have been discussed, but none define an architecture sufficiently for analysis of interoperability with the version 1.0 defined in this standard. The benefits of such an architecture have been discussed, and among these are relaxation of many, but not all, of the restrictions discussed in this subclause. In particular, the multiple spanning tree models appear to offer easier integration within both homogenous environments and in networks intermixing D Bridges with multiple spanning tree VLAN-aware devices.

Copyright © 1999 IEEE. All rights reserved.



Dell et al., Ex. 1008
Page 209 of 211

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

# Annex F

(informative)

## Frame translation considerations

When relaying frames between different MAC methods, there are a number of frame translations that may be required, depending upon the source and destination MAC methods concerned, the tagging state of the received and transmitted frames, and the data carried within the frames:

a) If the source and destination MAC methods differ, then the overall format of the received frame must be translated into the frame format required for the MAC onto which the frame will be transmitted. The details of the frame translation at this level is defined by the definition of the Internal Sublayer Service (ISO/IEC 15802-3, 6.4), the support of the internal sublayer service by specific MAC procedures (ISO/IEC 15802-3, 6.5), and the standards that define the specific MAC procedures concerned.
b) If the received frame is being relayed in tagged format between differing MAC methods, or if the tagging state is to change, then the tag header may require translation, insertion, or removal as defined in this standard.
c) If Ethernet Type-encoded data is being carried in the frame, then the format of that data may require translation as defined in ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390. These translations are essentially concerned with providing a standardized means of representing Ethernet Type-encoded frames on LANs that have no inherent ability within their MAC procedures to represent Ethernet Type values (i.e., LAN MACs where the “native” link layer protocol is LLC), and ensuring that the frame translation in the reverse direction results in the correct format on the 802.3/Ethernet LAN. The mechanisms specified in these standards are based around the use of SNAP encoding to carry Ethernet Type values in native LLC-encoded environments, and the use of translation tables to control the reverse translation if this should be necessary.

One aspect that is not fully covered in the standards mentioned is the issue of how differences in bit ordering between MAC methods can affect the data translation requirements. For the majority of data relayed by a Bridge, these differences in bit ordering are not an issue; although the bit ordering of MAC data “on the wire” may differ between MAC methods, the representation of that data within the relay function of the Bridge is independent of the order of transmission or reception adopted by the various MACs.

An exception can occur when the data carried in a frame includes MAC Address values; for example, in IP ARP packets. There are two formats used in LANs when representing MAC Address values in MAC user data:

d) Canonical format;
e) Non-canonical format.

The octet ordering is identical in both cases (the first octet put into the mac_service_data_unit is the left-most octet of the address when written down using hexadecimal notation as defined in Section 5 of IEEE Std 802); however, the bit ordering within each octet differs:

f) Canonical format: The least significant bit of each octet of the standard hexadecimal representation of the address (see IEEE Std 802-1990) represents the least-significant (LS) bit of the corresponding octet of the Canonical format of the address;
g) Non-canonical format: The most significant bit of each octet of the standard hexadecimal representation of the address (see IEEE Std 802-1990) represents the least-significant bit of the corresponding octet of the Canonical format of the address.


Copyright © 1999 IEEE. All rights reserved.


Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Ideally, a single format would be used in all places; however, both can be used, depending on the MAC method concerned. The native format used in MACs where the bit transmission order is LS bit first (e.g., IEEE Std 802.3 and ISO/IEC 8802-4), and also in FDDI, is the Canonical format; in MACs where the bit transmission order is most-significant (MS) bit first (e.g., ISO/IEC 8802-5), Non-canonical format is used. A further complication here is that the format used is not consistent across all higher layer protocols (some use a single format regardless of MAC method, others use the native format for the underlying MAC method), and there is no single standard that specifies which protocols carry embedded MAC Address information or the format in which they carry it.

Clearly, in order for the recipients of frames carrying embedded address information to be able to interpret MAC Address information correctly, it is necessary either to include information in the frame that specifies which format is in use, or for Bridges to modify the format of the information appropriately when relaying frames between regions of the network where the expected format differs for the protocol being carried.

In “Translating Bridges” (Bridges that relay between differing MAC methods, based on ISO/IEC 15802-3, ISO/IEC 11802-5, IETF RFC 1042, and IETF RFC 1390—i.e., in Bridging Function B1 in Figure C-2), the latter approach is the only option, as there is no way that additional information can be carried in the frame in order to identify the embedded format. Consequently, in order for such a Bridge to successfully translate embedded address information, it needs to be able to recognize the higher layer protocol carried in the frame, and act accordingly.

In VLAN-aware Bridges, the information carried in the CFI (and in the NCFI bit of the RIF, in 802.3/Ethernet frames) provides a direct indication of the format of embedded addresses. Hence, any considerations related to translating the format of embedded addresses can be confined to Bridges where a frame is received untagged and the tag header is inserted on transmission, or received tagged and the tag header is removed on transmission.

Copyrighted material licensed to Audrey Yang on 2022-11-02 for licensee's use only.
Copyrighted and Authorized by IEEE. Restrictions Apply.

Copyright © 1999 IEEE. All rights reserved.

Dell et al., Ex. 1008
Page 211 of 211