

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.

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.

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*

- 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;

- 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

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

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

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

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

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

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

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

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 15802-3¹. 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.

¹Information about references can be found in Clause 2.

- 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.

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

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.

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

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

⁴Internet 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.

⁵ISO 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.

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.

⁶ISO [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.

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.

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.

⁷The 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).

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.

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

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)

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.

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.

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*;

- 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 of VLANs. 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.

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.

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 some VLANs 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.

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.

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)  
)
```

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.

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.

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.

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.

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.

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).

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.




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.)




Figure 8-2—Bridge ports




Figure 8-3—VLAN Bridge architecture

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

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

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-

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.

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

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.

- 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




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

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.

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

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:

- The algorithm used to determine the FCS differs between the MAC methods used by the two MAC entities;
- The FCS coverage differs between the MAC methods used by the two MAC entities;
- 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

- 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
- 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