



IEEE Standard for

Local and metropolitan area networks—

Virtual Bridged Local Area Networks

**Amendment 13: Congestion
Notification**

IEEE Computer Society

Sponsored by the
LAN/MAN Standards Committee

IEEE
3 Park Avenue
New York, NY 10016-5997, USA

23 April 2010

IEEE Std 802.1Qau™-2010
(Amendment to
IEEE Std 802.1Q™-2005)

Currently in preview, click buy full version

IEEE Std 802.1Qau™-2010

(Amendment to
IEEE Std 802.1Q™-2005)

**IEEE Standard for
Local and metropolitan area networks—**

Virtual Bridged Local Area Networks

**Amendment 13: Congestion
Notification**

Sponsor

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

Approved 25 March 2010

IEEE-SA Standards Board

Abstract: This amendment specifies protocols, procedures and managed objects that support congestion management of long-lived data flows within network domains of limited bandwidth-delay product. This is achieved by enabling bridges to signal congestion to end stations capable of transmission rate limiting to avoid frame loss.

Keywords: congestion, congestion notification, LANs, local area networks, MAC Bridges, transparent bridging, VLANs

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

Copyright © 2010 by the Institute of Electrical and Electronics Engineers, Inc.
All rights reserved. Published 23 April 2010. Printed in the United States of America.

IEEE and 302 are registered trademarks in the U.S. Patent & Trademark Office, owned by The Institute of Electrical and Electronics Engineers, Incorporated.

PDF: ISBN 978-0-7381-6231-7 STD96051
Print: ISBN 978-0-7381-6232-4 STDPD96051

IEEE prohibits discrimination, harassment and bullying. For more information, visit <http://www.ieee.org/web/aboutus/whatis/policies/p9-26.html>.

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

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

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

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

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

Interpretations: Occasionally questions may arise regarding the meaning of portions of standards as they relate to specific applications. When the need for interpretations is brought to the attention of IEEE, the Institute will initiate action to prepare appropriate responses. Since IEEE Standards represent a consensus of concerned interests, it is important to ensure that any interpretation has also received the concurrence of a balance of interests. For this reason, IEEE and the members of its societies and Standards Coordinating Committees are unable to provide an instant response to interpretation requests except in those cases where the matter has previously received formal consideration. A statement, written or oral, that is not processed in accordance with the IEEE-SA Standards Board Operations Manual shall not be considered the official position of IEEE or any of its committees and shall not be considered to be, nor be relied upon as, a formal interpretation of the IEEE. At lectures, symposia, seminars, or educational courses, an individual presenting information on IEEE standards shall make it clear that his or her views should be considered the personal views of that individual rather than the formal position, explanation, or interpretation of the IEEE. 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. Recommendations to change the status of a stabilized standard should include a rationale as to why a revision or withdrawal is required.

Comments and recommendations on standards, and requests for interpretations should be addressed to:

Secretary, IEEE-SA Standards Board
445 Hoes Lane
Piscataway, NJ 08854
USA

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; +1 978 750 8400. Permission to photocopy portions of any individual standard for educational classroom use can also be obtained through the Copyright Clearance Center.

Introduction

This introduction is not part of IEEE Std 802.1Qau-2010, IEEE Standard for Local and Metropolitan Area Networks—Virtual Bridged Local Area Networks—Amendment 13: Congestion Notification.

This standard amends IEEE Std 802.1Q™-2005, providing congestion notification capabilities useful to Virtual Bridged Local Area Networks to support congestion management of long-lived data flows within network domains of limited bandwidth-delay product. Congestion notification mechanisms defined in this standard include the following:

- a) The ability of bridges and end stations to create Congestion Notification Domains for certain priority levels by signaling using the Link Layer Discovery Protocol defined in IEEE Std 802.1AB™-2009.^a
- b) The ability for bridges to use priority remapping to automatically defend a Congestion Notification Domain against sources that are not aware of congestion notification.
- c) Mechanisms by which bridges detect the congestion state of specified output queues, and send Congestion Notification Messages to the sources of a sampling of the frames in the queue.
- d) Mechanisms by which an end station responds to Congestion Notification Messages by stopping, increasing, decreasing, or disabling control of the rate of output for frames.
- e) A set of managed objects to provide controls for these capabilities on both bridges and end stations.

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 can be obtained from

Secretary, IEEE-SA Standards Board
445 Hoes Lane
Piscataway, NJ 08854
USA

Notice to users

Laws and regulations

Users of these documents should consult all applicable laws and regulations. Compliance with the provisions of this standard does not imply compliance to any applicable regulatory requirements. Implementers of the standard are responsible for observing or referring to the applicable regulatory requirements. IEEE does not, by the publication of its standards, intend to urge action that is not in compliance with applicable laws, and these documents may not be construed as doing so.

^aInformation of references can be found in Clause 2.

Copyrights

This document is copyrighted by the IEEE. It is made available for a wide variety of both public and private uses. These include both use, by reference, in laws and regulations, and use in private self-regulation, standardization, and the promotion of engineering practices and methods. By making this document available for use and adoption by public authorities and private users, the IEEE does not waive any rights in copyright to this document.

Updating of IEEE documents

Users of IEEE standards should be aware that these documents may be superseded at any time by the issuance of new editions or may be amended from time to time through the issuance of amendments, corrigenda, or errata. An official IEEE document at any point in time consists of the current edition of the document together with any amendments, corrigenda, or errata then in effect. In order to determine whether a given document is the current edition and whether it has been amended through the issuance of amendments, corrigenda, or errata, visit the IEEE Standards Association website at <http://ieeexplore.ieee.org/xpl/standards.jsp>, or contact the IEEE at the address listed previously.

For more information about the IEEE Standards Association or the IEEE standard development process, visit the IEEE-SA website at <http://standards.ieee.org>.

Errata

Errata, if any, for this and all other standards can be accessed at the following URL: <http://standards.ieee.org/reading/ieee/updates/errata/index.html>. Users are encouraged to check this URL for errata periodically.

Interpretations

Current interpretations can be accessed at the following URL: <http://standards.ieee.org/reading/ieee/interp/index.html>.

Patents

Attention is called to the possibility that implementation of this amendment may require use of subject matter covered by patent rights. By publication of this amendment, no position is taken with respect to the existence or validity of any patent rights in connection therewith. The IEEE is not responsible for identifying Essential Patent Claims for which a license may be required, for conducting inquiries into the legal validity or scope of Patent Claims or determining whether any licensing terms or conditions provided in connection with submission of a Letter of Assurance, if any, or in any licensing agreements are reasonable or non-discriminatory. Users of this amendment are expressly advised that determination of the validity of any patent rights, and the risk of infringement of such rights, is entirely their own responsibility. Further information may be obtained from the IEEE Standards Association.

Participants

At the time this amendment was submitted to the IEEE-SA for approval, the IEEE 802.1 Working Group had the following membership:

Tony Jeffree, Chair

Paul Congdon, Vice Chair

Norman Finn, Editor

Patricia Thaler, Chair, Data Center Bridging Task Group

Zehavit Alon
Ting Ao
Siamack Ayandeh
Jan Bialkowski
Jean-Michel Bonnamy
Paul Bottorff
Rudolf Brandner
Craig Carlson
Weiyang Cheng
Rao Cherukuri
Don Connor
Diego Crupnicoff
Claudio Desanti
Zheming Ding
Linda Dunbar
Donald Eastlake, III
Janos Farkas
Donald Fedyk
Robert Frazier
John Fuller
Geoffrey Garner
Anoop Ghanwani
Franz Goetz
Eric Gray
Craig Gunther
Mitch Gusat
Stephen Haddock
Asif Hazarika
Charles Hudson

Romain Insler
Pankaj Jha
Abhay Karandikar
Hal Keen
Yongbum Kim
Philippe Klein
Michael Krause
Vinod Kumar
Bruce Kwan
Kari Laihonen
Ashvin Lakshmikantha
John Lemon
Marina Lipshteyn
Yuanqiu Luo
Ben Mack-Crane
Menucher Menuchery
John Messenger
Gabriel Montenegro
John Morris
Eric Multanen
Paul Nikolich
Kevin Nolis
David Olsen
Vijay Pandey
Donald Parrell
Glenn Parsons
Joseph Pelissier
David Peterson
Mayim Porat
Karen Randall

Josef Roese
Derek Rohde
Dan Romascanu
Jessy Rouyer
Ali Sajassi
Joseph Salcedo
Panagiotis Salinas
John Sauer
James Scott
Michael Solomon
Koichiro Soto
Kamal Sharma
Masashi Shimizu
Nurit Sprecher
Kevin Stanton
Robert Sultan
Muneyoshi Suzuki
Michael Johas Teener
Paul Unbehagen
Manoj Wadekar
Yuehua Wei
Brian Weis
Martin White
Bert Wijnen
Robert Winter
Chien-Hsien Wu
Ken Young
Glen Zorn

The following members of the individual balloting committee voted on this standard. Balloters may have voted for approval, disapproval, or abstention.

| | | |
|----------------------|-------------------|-----------------------|
| Thomas Alexander | Atsushi Ito | Randall Safier |
| Richard Alfvén | Tony Jeffree | John Santhoff |
| Butch Anton | Shinkyō Kaku | John Sauer |
| Khin Mi Mi Aung | Stuart J. Kerry | Bartien Sayogo |
| Tomo Bogataj | Yongbum Kim | Michael Seaman |
| Edward Carley | Bruce Kraemer | Gil Shultz |
| Juan Carreon | Li Li | Amjad Soomro |
| Keith Chow | William Lumpkins | Matthew Squire |
| Charles Cook | Elvis Maculuba | Manikantan Srinivasan |
| Wael Diab | Mark Maloney | Kevin B. Stanton |
| Russell Dietz | Arthur Marris | Thomas Starai |
| Thomas Dineen | Peter Martini | Walter Struppler |
| Donald Eastlake, III | Gary Michel | Joseph Tardo |
| C. Fitzgerald | Jose Morales | William Taylor |
| Yukihiro Fujimoto | Joseph Moran | Michael Johas Teener |
| Ilango Ganga | Eric Multanen | Patricia Thaler |
| Reinhard Gloger | Michael S. Newman | Mark Rebe Uchida |
| Ron Greenthaler | Nick S. A. Nikjoo | Dmitri V. Ivanofiev |
| Randall Groves | Satoshi Obara | Prabh V. Ashney |
| C. Guy | Stephen Palm | Marjo Wentink |
| Rainer Hach | Glenn Parsons | Junpeng Wu |
| Stephen Haddock | Hayim Porat | Oren Yuen |
| David Hunter | Robert Robinson | Zhen Zhou |
| C. Huntley | Jessy V. Rouyer | |

When the IEEE-SA Standards Board approved this standard on 25 March 2010, it had the following membership:

Robert M. Grow, *Chair*
Richard H. Hulet, *Vice Chair*
Steve M. Mills, *Public Chair*
Judith German, *Secretary*

| | | |
|------------------|-----------------------|------------------------|
| John Barr | Mark Robinson | Ted Olsen |
| Karen Bartleson | Alexander Gelman | Glenn Parsons |
| Victor Berman | Jim Hughes | Ronald C. Petersen |
| Ted Burse | Young Kyun Kim | Narayanan Ramachandran |
| Richard DeBlasio | Joseph L. Koepfinger* | Jon Walter Rosdahl |
| Andy Drozd | John Kulick | Sam Sciacca |
| | David J. Law | |

*Member Emeritus

Also included are the following nonvoting IEEE-SA Standards Board liaisons:

Howard L. Wolfman, TAB Representative
Michael Janezic, NIST Representative
Satish K. Aggarwal, NRC Representative

Michelle Turner
IEEE Standards Program Manager, Document Development

Kathryn Cush
IEEE Standards Program Manager, Technical Program Development

Currently in preview, click buy full version

Contents

| | | |
|---------|--|----|
| 1. | Overview | 2 |
| 1.1 | Scope | 2 |
| 2. | References | 3 |
| 3. | Definitions | 5 |
| 4. | Abbreviations | 7 |
| 5. | Conformance | 9 |
| 5.4.3 | VLAN-aware Bridge requirements for congestion notification | 9 |
| 5.10 | End station requirements for congestion notification | 9 |
| 6. | Support of the MAC Service in VLANs | 11 |
| 6.10.1 | Data indications | 11 |
| 8. | Principles of bridge operation | 13 |
| 8.6.6 | Queuing frames | 13 |
| 12. | Bridge management | 15 |
| 12.1.1 | Configuration management | 15 |
| 12.2 | Managed objects | 15 |
| 12.17 | Congestion notification managed objects | 15 |
| 12.17.1 | CN component managed object | 15 |
| 12.17.2 | CN component priority managed object | 16 |
| 12.17.3 | CN Port priority managed object | 17 |
| 12.17.4 | Congestion Point managed object | 18 |
| 12.17.5 | Reaction Point port priority managed object | 18 |
| 12.17.6 | Reaction Point group managed object | 19 |
| 17. | MIB Modules | 21 |
| 17.1 | The Internet Standard Management Framework | 21 |
| 17.2 | Structure of the MIB | 21 |
| 17.2.16 | Structure of the Congestion Notification MIB | 21 |
| 17.3 | Relationship to other MIB modules | 23 |
| 17.3.16 | Relationship of the Congestion Notification MIB to other MIB modules | 23 |
| 17.4 | Security considerations | 24 |
| 17.4.16 | Security considerations of the Congestion Notification MIB | 24 |
| 17.7 | MIB modules | 25 |
| 17.7.16 | Congestion Notification MIB module | 25 |
| 30. | Principles of congestion notification | 63 |
| 30.1 | Congestion notification design requirements | 63 |
| 30.2 | Quantized Congestion Notification protocol | 66 |
| 30.2.1 | The CP Algorithm | 66 |
| 30.2.2 | Basic Reaction Point algorithm | 67 |
| 30.2.3 | RP algorithm with timer | 68 |
| 30.3 | Congestion Controlled Flow | 69 |

| | | |
|---------|--|----|
| 30.4 | Congestion Notification Priority Value | 70 |
| 30.5 | Congestion Notification Tag | 70 |
| 30.6 | Congestion Notification Domain | 70 |
| 30.7 | Multicast data | 71 |
| 30.8 | Congestion notification and additional tags | 71 |
| 31. | Congestion notification entity operation | 73 |
| 31.1 | Congestion aware Bridge Forwarding Process | 73 |
| 31.1.1 | Congestion Point | 73 |
| 31.1.2 | Congestion Point ingress multiplexer | 74 |
| 31.2 | Congestion aware end station functions | 74 |
| 31.2.1 | Output flow segregation | 75 |
| 31.2.2 | Per-CNPV station function | 76 |
| 31.2.3 | Flow Select Database | 78 |
| 31.2.4 | Flow multiplexer | 78 |
| 31.2.5 | CNM demultiplexer | 78 |
| 31.2.6 | Input flow segregation | 78 |
| 31.2.7 | End station input queue | 79 |
| 31.2.8 | Reception selection | 79 |
| 32. | Congestion notification protocol | 81 |
| 32.1 | Congestion Notification Domain operations | 81 |
| 32.1.1 | Congestion Notification Domain defense | 81 |
| 32.1.2 | Automatic Congestion Notification Domain recognition | 83 |
| 32.1.3 | Variables controlling Congestion Notification Domain defense | 83 |
| 32.2 | CN component variables | 84 |
| 32.2.1 | cngMasterEnable | 85 |
| 32.2.2 | cngCnmTransmitPriority | 85 |
| 32.2.3 | cngDiscardedFrames | 85 |
| 32.2.4 | cngErroredPortList | 85 |
| 32.3 | Congestion notification per-CNPV variables | 85 |
| 32.3.1 | cncpDefModeChoice | 85 |
| 32.3.2 | cncpAlternatePriority | 86 |
| 32.3.3 | cncpAutoAltPri | 86 |
| 32.3.4 | cncpAdminDefenseMode | 86 |
| 32.3.5 | cncpCreation | 86 |
| 32.3.6 | cncpLldpInstanceChoice | 86 |
| 32.3.7 | cncpLldpInstanceSelector | 86 |
| 32.4 | CND defense per-Port per-CNPV variables | 87 |
| 32.4.1 | cnpdDefModeChoice | 87 |
| 32.4.2 | cnpdAdminDefenseMode | 87 |
| 32.4.3 | cnpdAutoDefenseMode | 88 |
| 32.4.4 | cnpdLldpInstanceChoice | 88 |
| 32.4.5 | cnpdLldpInstanceSelector | 88 |
| 32.4.6 | cnpdAlternatePriority | 88 |
| 32.4.7 | cnpdXmitCnpvCapable | 88 |
| 32.4.8 | cnpdXmitReady | 88 |
| 32.4.9 | cncpDoesEdge | 89 |
| 32.4.10 | cnpdAcceptsCnTag | 89 |
| 32.4.11 | cnpdRcvdCnpv | 89 |
| 32.4.12 | cnpdRcvdReady | 89 |
| 32.4.13 | cnpdIsAdminDefMode | 89 |
| 32.4.14 | cnpdDefenseMode | 90 |

| | | |
|----------|--|-----|
| 32.5 | Congestion Notification Domain defense procedures | 90 |
| 32.5.1 | DisableCnpvRemapping() | 90 |
| 32.5.2 | TurnOnCnDefenses() | 90 |
| 32.5.3 | TurnOffCnDefenses() | 90 |
| 32.6 | Congestion Notification Domain defense state machine | 91 |
| 32.7 | Congestion notification protocol | 91 |
| 32.8 | Congestion Point variables | 92 |
| 32.8.1 | cpMacAddress | 93 |
| 32.8.2 | cpId | 93 |
| 32.8.3 | cpQSp | 93 |
| 32.8.4 | cpQLen | 93 |
| 32.8.5 | cpQLenOld | 93 |
| 32.8.6 | cpW | 93 |
| 32.8.7 | cpQOffset | 93 |
| 32.8.8 | cpQDelta | 93 |
| 32.8.9 | cpFb | 93 |
| 32.8.10 | cpEnqueued | 94 |
| 32.8.11 | cpSampleBase | 94 |
| 32.8.12 | cpDiscardedFrames | 94 |
| 32.8.13 | cpTransmittedFrames | 94 |
| 32.8.14 | cpTransmittedCnms | 94 |
| 32.8.15 | cpMinHeaderOctets | 94 |
| 32.9 | Congestion Point procedures | 94 |
| 32.9.1 | Random | 94 |
| 32.9.2 | NewCpSampleBase() | 95 |
| 32.9.3 | EM_UNITDATA.request (parameters) | 95 |
| 32.9.4 | GenerateCnmPdu() | 95 |
| 32.10 | Reaction Point per-Port per-CNPV variables | 97 |
| 32.10.1 | rpppMaxRps | 97 |
| 32.10.2 | rpppCreatedRps | 97 |
| 32.10.3 | rpppRpCentiseconds | 97 |
| 32.11 | Reaction Point group variables | 97 |
| 32.11.1 | rpgEnable | 98 |
| 32.11.2 | rpgTimeReset | 98 |
| 32.11.3 | rpgByteReset | 98 |
| 32.11.4 | rpgThreshold | 98 |
| 32.11.5 | rpgMaxRate | 98 |
| 32.11.6 | rpgAiRate | 98 |
| 32.11.7 | rpgHaiRate | 98 |
| 32.11.8 | rpgGd | 98 |
| 32.11.9 | rpgMinDecFac | 98 |
| 32.11.10 | rpgMinRate | 99 |
| 32.12 | Reaction Point timer | 99 |
| 32.12.1 | RpWhile | 99 |
| 32.13 | Reaction Point variables | 99 |
| 32.13.1 | rpEnabled | 99 |
| 32.13.2 | rpByteCount | 99 |
| 32.13.3 | rpByteStage | 99 |
| 32.13.4 | rpTimeStage | 100 |
| 32.13.5 | rpTargetRate | 100 |
| 32.13.6 | rpCurrentRate | 100 |
| 32.13.7 | rpFreeze | 100 |
| 32.13.8 | rpLimiterRate | 100 |
| 32.13.9 | rpFb | 100 |

| | | |
|-----------------------|---|-----|
| 32.14 | Reaction Point procedures | 100 |
| 32.14.1 | ResetCnm | 101 |
| 32.14.2 | TestRpTerminate | 101 |
| 32.14.3 | TransmitDataFrame | 101 |
| 32.14.4 | ReceiveCnm | 101 |
| 32.14.5 | ProcessCnm | 102 |
| 32.14.6 | AdjustRates | 102 |
| 32.15 | RP rate control state machine | 102 |
| 32.16 | Congestion notification and encapsulation interworking function | 104 |
| 33. | Encoding of congestion notification Protocol Data Units | 107 |
| 33.1 | Structure, representation, and encoding | 107 |
| 33.2 | Congestion Notification Tag format | 107 |
| 33.2.1 | Flow Identifier | 108 |
| 33.3 | Congestion Notification Message | 108 |
| 33.4 | Congestion Notification Message PDU format | 109 |
| 33.4.1 | Version | 109 |
| 33.4.2 | ReservedV | 109 |
| 33.4.3 | Quantized Feedback | 110 |
| 33.4.4 | Congestion Point Identifier | 110 |
| 33.4.5 | cnmQOffset | 110 |
| 33.4.6 | cnmQDelta | 110 |
| 33.4.7 | Encapsulated priority | 110 |
| 33.4.8 | Encapsulated destination MAC address | 110 |
| 33.4.9 | Encapsulated MSDU length | 110 |
| 33.4.10 | Encapsulated MSDU | 111 |
| 33.4.11 | CNM Validation | 111 |
| 33.5 | Congestion Notification TLV | 111 |
| 33.5.1 | TLV type | 111 |
| 33.5.2 | TLV information string length | 112 |
| 33.5.3 | Per-priority CNPV indicators | 112 |
| 33.5.4 | Per-priority Ready indicators | 112 |
| Annex A (normative) | PICS Proforma | 113 |
| A.5 | Major capabilities | 113 |
| A.14 | Bridge Management | 113 |
| A.24 | Management Information Base (MIB) | 114 |
| A.25 | Congestion notification | 115 |
| Annex H (informative) | Bibliography | 119 |

List of figures

| | | |
|-------------|---|-----|
| Figure 30-1 | Congestion detection in QCN CP | 66 |
| Figure 30-2 | Sampling (reflection) probability in QCN CP as a function of $ Fb $ | 67 |
| Figure 30-3 | QCN RP operation | 67 |
| Figure 30-4 | Byte Counter and Timer interaction with Rate Limiter | 69 |
| Figure 30-5 | CP–RP peering in VLAN Bridged Network..... | 72 |
| Figure 30-6 | CP–RP peering in Provider Backbone Bridged Network | 72 |
| Figure 31-1 | Congestion Points and congestion aware queues in a bridge..... | 74 |
| Figure 31-2 | Congestion aware queue functions in an end station | 75 |
| Figure 31-3 | Per-CNPV station function | 77 |
| Figure 32-1 | Congestion Notification Domain defense state machine | 91 |
| Figure 32-2 | RP rate control state machine..... | 103 |
| Figure 32-3 | CP–RP peering in any hierarchical Bridged Network | 104 |
| Figure 33-1 | Congestion Notification TLV format..... | 111 |

List of tables

| | | |
|------------|--|-----|
| Table 12-2 | CN component priority managed object row elements | 16 |
| Table 12-1 | CN component managed object row elements | 16 |
| Table 12-3 | CN Port priority managed object row elements..... | 17 |
| Table 12-4 | Congestion Point managed object row elements..... | 18 |
| Table 12-6 | Reaction Point group managed object row elements..... | 19 |
| Table 12-5 | Reaction Point port priority managed object row elements..... | 19 |
| Table 17-1 | Structure of the MIB Modules | 21 |
| Table 17-2 | Variables, managed object tables, and MIB objects..... | 21 |
| Table 32-1 | LLDP instance selection managed object overrides..... | 84 |
| Table 32-2 | CND defense mode selection managed object overrides..... | 84 |
| Table 32-3 | Determining cnpdIsAdminDefMode and cnpdDefenseMode | 90 |
| Table 32-4 | Correspondence of QCN protocol and CCF message fields..... | 92 |
| Table 32-5 | NewCpSampleBase() return value as a function of cpFb..... | 95 |
| Table 33-3 | Congestion Notification Message Encapsulation: Type/Length Media | 108 |
| Table 33-1 | Congestion Notification Tag Encapsulation: Type/Length Media | 108 |
| Table 33-2 | Congestion Notification Tag Encapsulation: LLC Media | 108 |
| Table 33-5 | Congestion Notification Message PDU | 109 |
| Table 33-4 | Congestion Notification Message Encapsulation: LLC Media | 109 |

IEEE Standard for Local and metropolitan area networks—

Virtual Bridged Local Area Networks

Amendment 13: Congestion Notification

IMPORTANT NOTICE: This standard is not intended to ensure safety, security, health, or environmental protection in all circumstances. Implementers of the standard are responsible for determining appropriate safety, security, environmental, and health practices or regulatory requirements.

This IEEE document is made available for use subject to important notices and legal disclaimers. These notices and disclaimers appear in all publications containing this document and may be found under the heading “Important Notice” or “Important Notices and Disclaimers Concerning IEEE Documents.” They can also be obtained on request from IEEE or viewed at <http://standards.ieee.org/IPR/disclaimers.html>.

Editorial Note

This amendment specifies changes to the forwarding and queueing functions described in IEEE Std 802.1Q™. Changes are applied to the base text of IEEE Std 802.1Q-2005, as modified by those amendments that had been approved, but not incorporated into the base text of the standard, at the time that this amendment was approved, namely (in chronological order) IEEE Std 802.1ad™, IEEE Std 802.1ak™, IEEE Std 802.1ag™, IEEE Std 802.1ah™, IEEE Std 802.1Q-2005/Cor 1, IEEE Std 802.1ap™, IEEE Std 802.1Qaw™, IEEE Std 802.1Qay™, IEEE Std 802.1aj™, and IEEE Std 802.1Qav™. Text shown in bold italics in this amendment defines the editing instructions necessary to changes to this base text. Three editing instructions are used: **change**, **delete**, and **insert**. **Change** is used to make a change to existing material. The editing instruction specifies the location of the change and describes what is being changed. Changes to existing text may be clarified using ~~strikeout~~ markings to indicate removal of old material, and underline markings to indicate addition of new material. **Delete** removes existing material. **Insert** adds new material without changing the existing material. Insertions may require renumbering. If so, renumbering instructions are given in the editing instruction. Editorial notes will not be carried over into future editions of IEEE Std 802.1Q.

1. Overview

Insert the following after the initial paragraphs of Clause 1:

This standard specifies protocols, procedures, and managed objects that support congestion management of long-lived data flows within network domains with a bandwidth-delay product on the order of 5 Mbits or less. Such flows are typically encountered in data centers, backplane fabrics, computing clusters, and storage networks. This is achieved by enabling bridges to signal congestion to end stations capable of transmission rate limiting to avoid frame loss. This mechanism enables support for higher layer protocols that are highly loss or latency sensitive. VLAN tag encoded priority values are allocated to segregate frames subject to congestion control, allowing simultaneous support of both congestion controlled and other higher layer protocols. This standard does not specify communication or reception of congestion notification information to or from end stations outside the congestion controlled domain or encapsulation of frames from those end stations across the domain.

1.1 Scope

Insert the following at end of 1.1, relettering the bullet points so that they follow in order from those in the existing text.

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

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