

Internet Engineering Task Force (IETF)
Request for Comments: 7133
Category: Standards Track
ISSN: 2070-1721

S. Kashima
NTT
A. Kobayashi, Ed.
NTT East
P. Aitken
Cisco Systems, Inc.
May 2014

Information Elements for Data Link Layer Traffic Measurement

Abstract

This document describes Information Elements related to the data link layer. They are used by the IP Flow Information Export (IPFIX) protocol for encoding measured data link layer traffic information.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc7133>.

Copyright Notice

Copyright (c) 2014 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction	4
1.1. Conventions Used in This Document	4
2. Extended Ethernet Technology	4
2.1. Wide-Area Ethernet Technology Summary	4
2.2. Virtual Ethernet Technology Summary	5
3. Modification and Addition of Information Elements	
Related to Data Link Layer	6
3.1. Existing Information Elements	7
3.1.1. dataLinkFrameSize	8
3.1.2. dataLinkFrameSection	9
3.1.3. layer2OctetDeltaCount	9
3.1.4. layer2OctetTotalCount	10
3.1.5. layer2FrameDeltaCount	10
3.1.6. layer2FrameTotalCount	11
3.2. New Information Elements	11
3.2.1. dataLinkFrameType	12
3.2.2. sectionOffset	12
3.2.3. sectionExportedOctets	13
3.2.4. dot1qServiceInstanceTag	13
3.2.5. dot1qServiceInstanceId	14
3.2.6. dot1qServiceInstancePriority	14
3.2.7. dot1qCustomerSourceMacAddress	15
3.2.8. dot1qCustomerDestinationMacAddress	15
3.2.9. postL2OctetDeltaCount	16
3.2.10. postMCastL2OctetDeltaCount	16
3.2.11. postL2OctetTotalCount	17
3.2.12. postMCastL2OctetTotalCount	17
3.2.13. minimumL2TotalLength	18
3.2.14. maximumL2TotalLength	18
3.2.15. droppedL2OctetDeltaCount	19
3.2.16. droppedL2OctetTotalCount	19
3.2.17. ignoredL2OctetTotalCount	20
3.2.18. notSentL2OctetTotalCount	20
3.2.19. layer2OctetDeltaSumOfSquares	21
3.2.20. layer2OctetTotalSumOfSquares	21
4. Modification of Existing Information Elements Related	
to Packet Section	22
4.1. ipHeaderPacketSection	22
4.2. ipPayloadPacketSection	23
4.3. mplsLabelStackSection	24
4.4. mplsPayloadPacketSection	25
5. Modification of Existing Information Elements Related	
to VLAN Tag	26
5.1. dot1qVlanId	26
5.2. dot1qPriority	27
5.3. dot1qCustomerVlanId	27

5.4. dot1qCustomerPriority	27
6. The Relationship between Ethernet Header Fields and Information Elements	28
7. Security Considerations	29
8. IANA Considerations	29
9. Acknowledgments	30
10. References	30
10.1. Normative References	30
10.2. Informative References	31
Appendix A. Frame Formats	32
Appendix B. Template Format Example	40

1. Introduction

Ethernet [IEEE802.1D] and VLAN (Virtual LAN) technologies had been used only in Local Area Networks. Recently, they have been used in Wide Area Networks, e.g., Layer 2 VPN (L2 VPN) services. Accordingly, carrier networks using VLAN technologies have been enhanced to Provider Bridged Networks and Provider Backbone Bridged Networks [IEEE802.1Q]. In addition, Ethernet in data centers has also been enhanced for server virtualization and input/output (I/O) consolidation.

While these innovations provide flexibility, scalability, and mobility to an existing network architecture, they increase the complexity of traffic measurement due to the existence of various Ethernet header formats. To cope with this, a more sophisticated method of traffic measurement is required.

IPFIX and Packet Sampling (PSAMP) help to resolve these problems. However, the PSAMP Information Model [RFC5477] and the IPFIX Information Model [RFC7011] don't yet contain enough Information Elements related to the data link layer, e.g., Ethernet header forms. This document describes existing and new Information Elements related to data link layers that enable a more sophisticated traffic measurement method.

Note that this document does not update [RFC5477] or [RFC7011] because IANA's IPFIX registry [IANA-IPFIX] is the ultimate Information Element reference, per Section 1 of [RFC7012].

1.1. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

2. Extended Ethernet Technology

2.1. Wide-Area Ethernet Technology Summary

Provider Bridge and Provider Backbone Bridge [IEEE802.1Q], which are standards for Wide-Area Ethernet, are described below.

- o In Provider Bridge [IEEE802.1Q], there are two VLAN IDs: Service VLAN Identifier (S-VID) and Customer VLAN Identifier (C-VID). S-VID is assigned to an Ethernet frame by a service provider, while C-VID is independently assigned to an Ethernet frame by a customer. Frame switching in a service provider network is based on only S-VID.

- o In Provider Backbone Bridge [IEEE802.1Q], new Ethernet fields, such as Backbone VLAN Identifier (B-VID) and Backbone Service Instance Identifier (I-SID), are introduced to overcome the limitations on the VLAN identifier space and to isolate the service provider and customer identifier spaces. Frame switching is based on a 12-bit B-VID, and customer identification is based on a 24-bit I-SID. A flexible network design has become possible because network management is separated from customer management. Other Ethernet fields that indicate quality of service (QoS) class are Backbone VLAN Priority Code Point (B-PCP), Backbone VLAN Drop Eligible Indicator (B-DEI), Backbone Service Instance Priority Code Point (I-PCP), and Backbone Service Instance Drop Eligible Indicator (I-DEI).

The Provider Backbone Bridge technologies have enhanced a Wide-Area Ethernet service from a flat network to a hierarchical network consisting of a Provider Bridged Network and Provider Backbone Bridged Network.

Frame formats used in Wide-Area Ethernet are shown in Appendix A.

2.2. Virtual Ethernet Technology Summary

There have been several challenges in the existing virtual switches environment in a data center. One is the lack of network management visibility: limited features on virtual switches make it difficult to monitor traffic among virtual machines (VMs). Another is the lack of management scalability and flexibility: increasing the number of VMs for multi-tenant architecture causes an increase in the number of virtual switches and in the number of the traffic control policies, which reach the limitations of network management scalability and flexibility.

In this situation, the IEEE 802.1 working group is standardizing virtual bridging technologies such as Edge Virtual Bridging (EVB), including two kinds of Edge Relays: Virtual Edge Bridge (VEB) and Virtual Edge Port Aggregator (VEPA) [IEEE802.1Qbg]. The VEB is a bridge that provides bridging among multiple VMs and the external bridging environment. The VEPA is a bridge-like device on a host that forwards all internal traffic to the adjacent EVB bridge and then distributes any traffic received from the adjacent EVB bridge to VMs. The VEPA makes all the VM-to-VM traffic visible to the EVB bridge so that the traffic can be monitored and so that the EVB bridge can apply filtering to the traffic.

To improve flexibility, a virtual link between a host system and EVB bridge is standardized as S-channel. S-channel allows a bridge to treat the traffic in the virtual link as if it comes in on a separate

port. For example, in the host, an S-channel may be attached to a VEB or a VEPA or directly to an internal port in order to apply each port-based filtering rule to the traffic. S-channel over the link between a host and its adjacent bridge uses Service VLAN Tag (S-TAG) [IEEE802.1Q]. When S-channel is in use, frames on the link carry an S-TAG to identify the S-channel.

On the other hand, Bridge Port Extension emulates single Extended Bridge from multiple physical switches and virtual switches, and it also simplifies network management. Also, it solves the lack of network management visibility by forwarding all traffic into a central Controlling Bridge using E-channel. E-channel over the link between a Bridge Port Extender and a Controlling Bridge uses E-TAG defined in [IEEE802.1BR].

Traffic monitoring over S-channel and E-channel is required in order to get visibility of VM-to-VM traffic and visibility of each channel's traffic on a virtual link.

Frame formats with E-TAG used in E-channel and S-TAG used in S-channel are shown in Appendix A. Though these frames carry special tags while on the link, those tags identify a virtual port (or for multicast in the downstream direction, a set of virtual ports) to which they are destined. These tag values only have local meaning, and the Flow would be reported as sent and arriving on the corresponding virtual ports. Therefore, IPFIX does not need to monitor data based on these tags.

3. Modification and Addition of Information Elements Related to Data Link Layer

The Information Elements listed in the upper section of Table 1 are necessary for enabling IPFIX and PSAMP traffic measurement for the data link layer, which is not limited to Ethernet because the method can be applied to other data link protocols as well.

Information Elements in the middle section of Table 1 are necessary for enabling the IPFIX and PSAMP traffic measurement for [IEEE802.1Q].

Information Elements in the lower section of Table 1 are octet counter or packet length for layer 2, and they are necessary for enabling IPFIX and PSAMP traffic measurement for the data link layer.

ID	Name
312	dataLinkFrameSize
315	dataLinkFrameSection
408	dataLinkFrameType
409	sectionOffset
410	sectionExportedOctets
411	dot1qServiceInstanceTag
412	dot1qServiceInstanceId
413	dot1qServiceInstancePriority
414	dot1qCustomerSourceMacAddress
415	dot1qCustomerDestinationMacAddress
352	layer2OctetDeltaCount
353	layer2OctetTotalCount
417	postL2OctetDeltaCount
418	postMCastL2OctetDeltaCount
420	postL2OctetTotalCount
421	postMCastL2OctetTotalCount
422	minimumL2TotalLength
423	maximumL2TotalLength
424	droppedL2OctetDeltaCount
425	droppedL2OctetTotalCount
426	ignoredL2OctetTotalCount
427	notSentL2OctetTotalCount
428	layer2OctetDeltaSumOfSquares
429	layer2OctetTotalSumOfSquares
430	layer2FrameDeltaCount
431	layer2FrameTotalCount

Table 1: Information Elements Related to Data Link Layer

3.1. Existing Information Elements

Some existing Information Elements are required for data link layer export. Their details are reproduced here from IANA's IPFIX registry [IANA-IPFIX]. Additions per this document appear between *.

Section 3.1.1 introduces the missing Data Type Semantics for the dataLinkFrameSize Information Element, which is held to be an interoperable change per #4 in Section 5.2 of [RFC7013].

Section 3.1.2 extends the definition of the dataLinkFrameSection Information Element with reference to the new sectionOffset Information Element, which is also an interoperable change per #4 in Section 5.2 of [RFC7013].

The layer2OctetDeltaCount Information Element reports the number of layer 2 octets since the previous report in incoming packets for this Flow, while the layer2OctetTotalCount Information Element reports the total number of layer 2 octets in incoming packets for this Flow. The layer2FrameDeltaCount Information Element reports the number of incoming layer 2 frames since the previous report for this Flow, while layer2FrameTotalCount Information Element reports the total number of incoming layer 2 frames for this Flow. All of these Information Elements are unchanged from the existing IANA [IANA-IPFIX] definitions, and are reproduced in Section 3.1.3 through Section 3.1.6 below for completeness.

Therefore, these changes do not introduce any backward-compatibility issues.

Per Section 5.2 of [RFC7013], for each of these changes, [RFC7133] has been appended to the requester in IANA's IPFIX registry [IANA-IPFIX], the Information Element's revision number has been incremented by one, and the Information Element's revision date column has been updated.

3.1.1.1. dataLinkFrameSize

Description:

This Information Element specifies the length of the selected data link frame.

The data link layer is defined in [ISO/IEC.7498-1:1994].

Abstract Data Type: unsigned16

Data Type Semantics: quantity

ElementId: 312

References: [ISO/IEC.7498-1:1994]

Status: current

3.1.2. dataLinkFrameSection

Description:

This Information Element carries *n* octets from the data link frame of a selected frame, starting *sectionOffset* octets into the frame.

*However, if no *sectionOffset* field corresponding to this Information Element is present, then a *sectionOffset* of zero applies, and the octets MUST be from the start of the data link frame.*

The *sectionExportedOctets* expresses how much data was observed, while the remainder is padding.

When the *sectionExportedOctets* field corresponding to this Information Element exists, this Information Element MAY have a fixed length and MAY be padded, or it MAY have a variable length.

When the *sectionExportedOctets* field corresponding to this Information Element does not exist, this Information Element SHOULD have a variable length and MUST NOT be padded. In this case, the size of the exported section may be constrained due to limitations in the IPFIX protocol.

Further Information Elements, i.e., *dataLinkFrameType* and *dataLinkFrameSize*, are needed to specify the data link type and the size of the data link frame of this Information Element. A set of these Information Elements MAY be contained in a structured data type, as expressed in [RFC6313]. Or a set of these Information Elements MAY be contained in one Flow Record as shown in Appendix B of [RFC7133].

The data link layer is defined in [ISO/IEC.7498-1:1994].

Abstract Data Type: `octetArray`

ElementId: 315

References: [RFC6313] [RFC7133] [ISO/IEC.7498-1:1994]

Status: current

3.1.3. layer2OctetDeltaCount

The *layer2OctetDeltaCount* Information Element is unchanged from the existing IANA [IANA-IPFIX] definition and is reproduced here for reference only.

Description

The number of layer 2 octets since the previous report (if any) in incoming packets for this Flow at the Observation Point. The number of octets includes layer 2 header(s) and layer 2 payload.

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

Units: octets

ElementId: 352

Status: current

3.1.4. layer2OctetTotalCount

The layer2OctetTotalCount Information Element is unchanged from the existing IANA [IANA-IPFIX] definition and is reproduced here for reference only.

Description:

The total number of layer 2 octets in incoming packets for this Flow at the Observation Point since the Metering Process (re-)initialization for this Observation Point. The number of octets includes layer 2 header(s) and layer 2 payload.

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

Units: octets

ElementId: 353

Status: current

3.1.5. layer2FrameDeltaCount

The layer2FrameDeltaCount Information Element is unchanged from the existing IANA [IANA-IPFIX] definition and is reproduced here for reference only.

Description:

The number of incoming layer 2 frames since the previous report (if any) for this Flow at the Observation Point.

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

Units: frames

ElementId: 430

Status: current

3.1.6. layer2FrameTotalCount

The layer2FrameTotalCount Information Element is unchanged from the existing IANA [IANA-IPFIX] definition and is reproduced here for reference only.

Description:

The total number of incoming layer 2 frames for this Flow at the Observation Point since the Metering Process (re-)initialization for this Observation Point.

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

Units: frames

ElementId: 431

Status: current

3.2. New Information Elements

The following new Information Elements have been added for data link layer monitoring.

In IANA's IPFIX registry [IANA-IPFIX], the Requester has been set to [RFC7133], the Information Element's Revision has been set to zero, and the Information Element's Date set to the date upon which the new Information Elements have been added to the registry. All other

columns that are not explicitly mentioned below (e.g., Units, Range, References) are not applicable and are to be left blank since the registry does not explicitly record "not applicable".

3.2.1. dataLinkFrameType

Description:

This Information Element specifies the type of the selected data link frame.

The following data link types are defined here:

- 0x01 IEEE802.3 ETHERNET [IEEE802.3]
- 0x02 IEEE802.11 MAC Frame format [IEEE802.11]

Further values may be assigned by IANA. Note that the assigned values are bits so that multiple observations can be OR'd together.

The data link layer is defined in [ISO/IEC.7498-1:1994].

Abstract Data Type: unsigned16

Data Type Semantics: flags

ElementId: 408

References: [IEEE802.3] [IEEE802.11] [ISO/IEC.7498-1:1994]

Status: current

3.2.2. sectionOffset

Description:

This Information Element specifies the offset of the packet section (e.g., dataLinkFrameSection, ipHeaderPacketSection, ipPayloadPacketSection, mplsLabelStackSection, and mplsPayloadPacketSection). If this Information Element is omitted, it defaults to zero (i.e., no offset).

If multiple sectionOffset Information Elements are specified within a single Template, then they apply to the packet section Information Elements in order: the first sectionOffset applies to the first packet section, the second to the second, and so on. Note that the "closest" sectionOffset and packet section

Information Elements within a given Template are not necessarily related. If there are fewer sectionOffset Information Elements than packet section Information Elements, then subsequent packet section Information Elements have no offset, i.e., a sectionOffset of zero applies to those packet section Information Elements. If there are more sectionOffset Information Elements than the number of packet section Information Elements, then the additional sectionOffset Information Elements are meaningless.

Abstract Data Type: unsigned16

Data Type Semantics: quantity

ElementId: 409

Status: current

3.2.3. sectionExportedOctets

Description:

This Information Element specifies the observed length of the packet section (e.g., dataLinkFrameSection, ipHeaderPacketSection, ipPayloadPacketSection, mplsLabelStackSection, and mplsPayloadPacketSection) when padding is used.

The packet section may be of a fixed size larger than the sectionExportedOctets. In this case, octets in the packet section beyond the sectionExportedOctets MUST follow the [RFC7011] rules for padding (i.e., be composed of zero (0) valued octets).

Abstract Data Type: unsigned16

Data Type Semantics: quantity

ElementId: 410

References: [RFC7011]

Status: current

3.2.4. dot1qServiceInstanceTag

Description:

This Information Element, which is 16 octets long, represents the Backbone Service Instance Tag (I-TAG) Tag Control Information (TCI) field of an Ethernet frame as described in [IEEE802.1Q]. It

encodes the Backbone Service Instance Priority Code Point (I-PCP), Backbone Service Instance Drop Eligible Indicator (I-DEI), Use Customer Addresses (UCAs), Backbone Service Instance Identifier (I-SID), Encapsulated Customer Destination Address (C-DA), Encapsulated Customer Source Address (C-SA), and reserved fields. The structure and semantics within the Tag Control Information field are defined in [IEEE802.1Q].

Abstract Data Type: octetArray

Data Type Semantics: default

ElementId: 411

References: [IEEE802.1Q]

Status: current

3.2.5. dot1qServiceInstanceId

Description:

The value of the 24-bit Backbone Service Instance Identifier (I-SID) portion of the Backbone Service Instance Tag (I-TAG) Tag Control Information (TCI) field of an Ethernet frame as described in [IEEE802.1Q].

Abstract Data Type: unsigned32

Data Type Semantics: identifier

ElementId: 412

References: [IEEE802.1Q]

Status: current

Range: The valid range is 0 - 16777215 (i.e., 24 bits).

3.2.6. dot1qServiceInstancePriority

Description:

The value of the 3-bit Backbone Service Instance Priority Code Point (I-PCP) portion of the Backbone Service Instance Tag (I-TAG) Tag Control Information (TCI) field of an Ethernet frame as described in [IEEE802.1Q].

Abstract Data Type: unsigned8

Data Type Semantics: identifier

ElementId: 413

References: [IEEE802.1Q]

Status: current

Range: The valid range is 0-7.

3.2.7. dot1qCustomerSourceMacAddress

Description:

The value of the Encapsulated Customer Source Address (C-SA) portion of the Backbone Service Instance Tag (I-TAG) Tag Control Information (TCI) field of an Ethernet frame as described in [IEEE802.1Q].

Abstract Data Type: macAddress

Data Type Semantics: default

ElementId: 414

References: [IEEE802.1Q]

Status: current

3.2.8. dot1qCustomerDestinationMacAddress

Description:

The value of the Encapsulated Customer Destination Address (C-DA) portion of the Backbone Service Instance Tag (I-TAG) Tag Control Information (TCI) field of an Ethernet frame as described in [IEEE802.1Q].

Abstract Data Type: macAddress

Data Type Semantics: default

ElementId: 415

References: [IEEE802.1Q]

Status: current

3.2.9. postL2OctetDeltaCount

Description:

The definition of this Information Element is identical to the definition of the layer2OctetDeltaCount Information Element, except that it reports a potentially modified value caused by a middlebox function after the packet passed the Observation Point.

This Information Element is the layer 2 version of postOctetDeltaCount (ElementId #23).

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

ElementId: 417

References: [RFC5477]

Status: current

Units: octets

3.2.10. postMCastL2OctetDeltaCount

Description:

The number of layer 2 octets since the previous report (if any) in outgoing multicast packets sent for packets of this Flow by a multicast daemon within the Observation Domain. This property cannot necessarily be observed at the Observation Point but may be retrieved by other means. The number of octets includes layer 2 header(s) and layer 2 payload.

This Information Element is the layer 2 version of postMCastOctetDeltaCount (ElementId #20).

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

ElementId: 418

References: [RFC5477]

Status: current

Units: octets

3.2.11. postL2OctetTotalCount

Description:

The definition of this Information Element is identical to the definition of the layer2OctetTotalCount Information Element, except that it reports a potentially modified value caused by a middlebox function after the packet passed the Observation Point.

This Information Element is the layer 2 version of postOctetTotalCount (ElementId #171).

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: 420

References: [RFC5477]

Status: current

Units: octets

3.2.12. postMCastL2OctetTotalCount

Description:

The total number of layer 2 octets in outgoing multicast packets sent for packets of this Flow by a multicast daemon in the Observation Domain since the Metering Process (re-)initialization. This property cannot necessarily be observed at the Observation Point but may be retrieved by other means. The number of octets includes layer 2 header(s) and layer 2 payload.

This Information Element is the layer 2 version of postMCastOctetTotalCount (ElementId #175).

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: 421

References: [RFC5477]

Status: current

Units: octets

3.2.13. minimumL2TotalLength

Description:

Layer 2 length of the smallest packet observed for this Flow. The packet length includes the length of the layer 2 header(s) and the length of the layer 2 payload.

This Information Element is the layer 2 version of minimumIpTotalLength (ElementId #25).

Abstract Data Type: unsigned64

ElementId: 422

References: [RFC5477]

Status: current

Units: octets

3.2.14. maximumL2TotalLength

Description:

Layer 2 length of the largest packet observed for this Flow. The packet length includes the length of the layer 2 header(s) and the length of the layer 2 payload.

This Information Element is the layer 2 version of maximumIpTotalLength (ElementId #26).

Abstract Data Type: unsigned64

ElementId: 423

References: [RFC5477]

Status: current

Units: octets

3.2.15. droppedL2OctetDeltaCount

Description:

The number of layer 2 octets since the previous report (if any) in packets of this Flow dropped by packet treatment. The number of octets includes layer 2 header(s) and layer 2 payload.

This Information Element is the layer 2 version of droppedOctetDeltaCount (ElementId #132).

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

ElementId: 424

References: [RFC5477]

Status: current

Units: octets

3.2.16. droppedL2OctetTotalCount

Description:

The total number of octets in observed layer 2 packets (including the layer 2 header) that were dropped by packet treatment since the (re-)initialization of the Metering Process.

This Information Element is the layer 2 version of droppedOctetTotalCount (ElementId #134).

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: 425

References: [RFC5477]

Status: current

Units: octets

3.2.17. ignoredL2OctetTotalCount

Description:

The total number of octets in observed layer 2 packets (including the layer 2 header) that the Metering Process did not process since the (re-)initialization of the Metering Process.

This Information Element is the layer 2 version of ignoredOctetTotalCount (ElementId #165).

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: 426

References: [RFC5477]

Status: current

Units: octets

3.2.18. notSentL2OctetTotalCount

Description:

The total number of octets in observed layer 2 packets (including the layer 2 header) that the Metering Process did not process since the (re-)initialization of the Metering Process.

This Information Element is the layer 2 version of notSentOctetTotalCount (ElementId #168).

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: 427

References: [RFC5477]

Status: current

Units: octets

3.2.19. layer2OctetDeltaSumOfSquares

Description:

The sum of the squared numbers of layer 2 octets per incoming packet since the previous report (if any) for this Flow at the Observation Point. The number of octets includes layer 2 header(s) and layer 2 payload.

This Information Element is the layer 2 version of octetDeltaSumOfSquares (ElementId #198).

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

ElementId: 428

References: [RFC5477]

Status: current

Units: octets

3.2.20. layer2OctetTotalSumOfSquares

Description:

The total sum of the squared numbers of layer 2 octets in incoming packets for this Flow at the Observation Point since the Metering Process (re-)initialization for this Observation Point. The number of octets includes layer 2 header(s) and layer 2 payload.

This Information Element is the layer 2 version of octetTotalSumOfSquares (ElementId #199).

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: 429

References: [RFC5477]

Status: current

Units: octets

4. Modification of Existing Information Elements Related to Packet Section

The new Information Elements related to packet section (i.e., sectionOffset and sectionExportedOctets) can be applied to not only dataLinkFrameSection but also to all kinds of packet section (i.e., ipHeaderPacketSection, ipPayloadPacketSection, mplsLabelStackSection, and mplsPayloadPacketSection defined in [RFC5477]). Therefore, existing Information Elements Descriptions should be modified as follows.

4.1. ipHeaderPacketSection

This Information Element is defined in [RFC5477]. The description has been updated from [RFC5477].

Description:

This Information Element carries a series of n octets from the IP header of a sampled packet, starting sectionOffset octets into the IP header.

However, if no sectionOffset field corresponding to this Information Element is present, then a sectionOffset of zero applies, and the octets MUST be from the start of the IP header.

With sufficient length, this element also reports octets from the IP payload. However, full packet capture of arbitrary packet streams is explicitly out of scope per the Security Considerations sections of [RFC5477] and [RFC2804].

The sectionExportedOctets expresses how much data was exported, while the remainder is padding.

When the sectionExportedOctets field corresponding to this Information Element exists, this Information Element MAY have a fixed length and MAY be padded, or it MAY have a variable length.

When the sectionExportedOctets field corresponding to this Information Element does not exist, this Information Element SHOULD have a variable length and MUST NOT be padded. In this case, the size of the exported section may be constrained due to limitations in the IPFIX protocol.

Abstract Data Type: octetArray

ElementId: 313

References: [RFC2804] [RFC5477]

Status: current

4.2. ipPayloadPacketSection

This Information Element is defined in [RFC5477]. The description is updated from [RFC5477].

Description:

This Information Element carries a series of n octets from the IP payload of a sampled packet, starting sectionOffset octets into the IP payload.

However, if no sectionOffset field corresponding to this Information Element is present, then a sectionOffset of zero applies, and the octets MUST be from the start of the IP payload.

The IPv4 payload is that part of the packet that follows the IPv4 header and any options, which [RFC0791] refers to as "data" or "data octets". For example, see the examples in [RFC0791], Appendix A.

The IPv6 payload is the rest of the packet following the 40-octet IPv6 header. Note that any extension headers present are considered part of the payload. See [RFC2460] for the IPv6 specification.

The sectionExportedOctets expresses how much data was observed, while the remainder is padding.

When the sectionExportedOctets field corresponding to this Information Element exists, this Information Element MAY have a fixed length and MAY be padded, or it MAY have a variable length.

When the sectionExportedOctets field corresponding to this Information Element does not exist, this Information Element SHOULD have a variable length and MUST NOT be padded. In this case, the size of the exported section may be constrained due to limitations in the IPFIX protocol.

Abstract Data Type: octetArray

ElementId: 314

References: [RFC0791] [RFC2460]

Status: current

4.3. mplsLabelStackSection

This Information Element is defined in [RFC5477]. The description is updated from [RFC5477].

Description:

This Information Element carries a series of n octets from the MPLS label stack of a sampled packet, starting sectionOffset octets into the MPLS label stack.

However, if no sectionOffset field corresponding to this Information Element is present, then a sectionOffset of zero applies, and the octets MUST be from the head of the MPLS label stack.

With sufficient length, this element also reports octets from the MPLS payload. However, full packet capture of arbitrary packet streams is explicitly out of scope per the Security Considerations sections of [RFC5477] and [RFC2804].

See [RFC3031] for the specification of MPLS packets.

See [RFC3032] for the specification of the MPLS label stack.

The sectionExportedOctets expresses how much data was observed, while the remainder is padding.

When the sectionExportedOctets field corresponding to this Information Element exists, this Information Element MAY have a fixed length and MAY be padded, or it MAY have a variable length.

When the sectionExportedOctets field corresponding to this Information Element does not exist, this Information Element SHOULD have a variable length and MUST NOT be padded. In this case, the size of the exported section may be constrained due to limitations in the IPFIX protocol.

Abstract Data Type: octetArray

ElementId: 316

References: [RFC2804] [RFC3031] [RFC3032] [RFC5477]

Status: current

4.4. mplsPayloadPacketSection

This Information Element is defined in [RFC5477]. The description is updated from [RFC5477].

Description:

The mplsPayloadPacketSection carries a series of n octets from the MPLS payload of a sampled packet, starting sectionOffset octets into the MPLS payload, as it is data that follows immediately after the MPLS label stack.

However, if no sectionOffset field corresponding to this Information Element is present, then a sectionOffset of zero applies, and the octets MUST be from the start of the MPLS payload.

See [RFC3031] for the specification of MPLS packets.

See [RFC3032] for the specification of the MPLS label stack.

The sectionExportedOctets expresses how much data was observed, while the remainder is padding.

When the sectionExportedOctets field corresponding to this Information Element exists, this Information Element MAY have a fixed length and MAY be padded, or it MAY have a variable length.

When the sectionExportedOctets field corresponding to this Information Element does not exist, this Information Element SHOULD have a variable length and MUST NOT be padded. In this case, the size of the exported section may be constrained due to limitations in the IPFIX protocol.

Abstract Data Type: octetArray

ElementId: 317

References: [RFC3031] [RFC3032]

Status: current

5. Modification of Existing Information Elements Related to VLAN Tag

The traffic measurement using IPFIX and PSAMP for a Provider Backbone Bridged Network requires the Information Elements related to Backbone Service Instance Tag (I-TAG) and Backbone VLAN Tag (B-TAG). The set of Information Elements related to I-TAG is added in Section 3, because I-TAG structure and semantics are different from that of Service VLAN Tag (S-TAG) and Customer VLAN Tag (C-TAG). The set of Information Elements related to B-TAG reuses the existing Information Elements, because B-TAG structure and semantics are identical to that of C-TAG and S-TAG. This section modifies existing descriptions and references related to C-TAG and S-TAG as follows.

5.1. dot1qVlanId

Description:

The value of the 12-bit VLAN Identifier portion of the Tag Control Information field of an Ethernet frame. The structure and semantics within the Tag Control Information field are defined in [IEEE802.1Q]. In Provider Bridged Networks, it represents the Service VLAN identifier in the Service VLAN Tag (S-TAG) Tag Control Information (TCI) field or the Customer VLAN identifier in the Customer VLAN Tag (C-TAG) Tag Control Information (TCI) field as described in [IEEE802.1Q]. In Provider Backbone Bridged Networks, it represents the Backbone VLAN identifier in the Backbone VLAN Tag (B-TAG) Tag Control Information (TCI) field as described in [IEEE802.1Q]. In a virtual link between a host system and EVB bridge, it represents the Service VLAN identifier indicating S-channel as described in [IEEE802.1Qbg].

In the case of a multi-tagged frame, it represents the outer tag's VLAN identifier, except for I-TAG.

Abstract Data Type: unsigned16

Data Type Semantics: identifier

ElementId: 243

Status: current

References: [IEEE802.1Q] [IEEE802.1Qbg]

5.2. dot1qPriority

Description:

The value of the 3-bit User Priority portion of the Tag Control Information field of an Ethernet frame. The structure and semantics within the Tag Control Information field are defined in [IEEE802.1Q]. In the case of a multi-tagged frame, it represents the 3-bit Priority Code Point (PCP) portion of the outer tag's Tag Control Information (TCI) field as described in [IEEE802.1Q], except for I-TAG.

Abstract Data Type: unsigned8

Data Type Semantics: identifier

ElementId: 244

Status: current

References: [IEEE802.1Q]

5.3. dot1qCustomerVlanId

Description:

The value represents the Customer VLAN identifier in the Customer VLAN Tag (C-TAG) Tag Control Information (TCI) field as described in [IEEE802.1Q].

Abstract Data Type: unsigned16

Data Type Semantics: identifier

ElementId: 245

Status: current

References: [IEEE802.1Q]

5.4. dot1qCustomerPriority

Description:

The value represents the 3-bit Priority Code Point (PCP) portion of the Customer VLAN Tag (C-TAG) Tag Control Information (TCI) field as described in [IEEE802.1Q].

Abstract Data Type: unsigned8

Data Type Semantics: identifier


ElementId: 246

Status: current

References: [IEEE802.1Q]


6. The Relationship between Ethernet Header Fields and Information Elements

The following figures show a summary of various Ethernet header fields and the Informational Elements that would be used to represent each of the fields.




- a.(Information Element) destinationMacAddress (80)
- b.(Information Element) sourceMacAddress (56)
- c.(Information Elements) dot1qVlanId (243), dot1qPriority (244)
- d.(Information Element) ethernetType (256)

Figure 1: Customer-Tagged Frame Header Fields



- a.(Information Element) destinationMacAddress (80)
- b.(Information Element) sourceMacAddress (56)
- c.(Information Elements) dot1qVlanId (243), dot1qPriority (244)
- d.(Information Elements) dot1qCustomerVlanId (245), dot1qCustomerPriority (246)
- e.(Information Element) ethernetType (256)

Figure 2: Service-Tagged Frame Header Fields



- a.(Information Element) destinationMacAddress (80)
 b.(Information Element) sourceMacAddress (56)
 c.(Information Elements) dot1qVlanId (243), dot1qPriority (244)
 d.(Information Elements) dot1qServiceInstanceTag (411), or
 a set of dot1qServiceInstanceId (412),
 dot1qServiceInstancePriority (413),
 dot1qCustomerSourceMacAddress (414)
 dot1qCustomerDestinationMacAddress (415),
 e.(Information Elements) dot1qCustomerVlanId (245),
 dot1qCustomerPriority (246)
 f.(Information Element) ethernetType (256)

Figure 3: Backbone-VLAN-Tagged Frame Header Fields

7. Security Considerations

Reporting more granular data may increase the risk of DoS attacks against a Collector. Protection against DoS attacks is discussed in Section 11.4 of [RFC7011].

The recommendations in this document do not otherwise introduce any additional security issues beyond those already mentioned in [RFC7011] and [RFC5477].

8. IANA Considerations

Existing IPFIX Information Elements [IANA-IPFIX] have been modified as indicated in Sections 3.1, 4, and 5.

Per Section 5.2 of [RFC7013], for each of these changes, [RFC7133] has been appended to the Requester in IANA's IPFIX registry [IANA-IPFIX], the Information Element's Revision number has been incremented by one, and the Information Element's revision Date column has been updated.

New IPFIX Information Elements [IANA-IPFIX] have been allocated as shown in Section 3.2.

9. Acknowledgments

Thanks to Brian Trammell and the IPFIX working group participants who contributed to mailing-list discussions throughout the development of this document. Special thanks to Pat Thaler for her help with the IEEE 802 aspects of this work.

10. References

10.1. Normative References

- [IEEE802.11] IEEE, "IEEE Standard for 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", IEEE Std 802.11-2012, March 2012.
- [IEEE802.1BR] IEEE, "IEEE Standard for Local and metropolitan area networks: Virtual Bridged Local Area Networks: Bridge Port Extension", IEEE Std 802.1BR-2012, July 2012.
- [IEEE802.1Q] IEEE, "IEEE Standard for Local and metropolitan area networks: Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2011, August 2011.
- [IEEE802.1Qbg] IEEE, "IEEE Standard for Local and metropolitan area networks: Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks: Amendment 21: Edge Virtual Bridging", IEEE Std 802.1Qbg-2012, July 2012.
- [IEEE802.3] IEEE, "IEEE Standard for Ethernet", IEEE Std 802.3-2012, December 2012.
- [RFC0791] Postel, J., "Internet Protocol", STD 5, RFC 791, September 1981.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2460] Deering, S. and R. Hinden, "Internet Protocol, Version 6 (IPv6) Specification", RFC 2460, December 1998.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, January 2001.
- [RFC5477] Dietz, T., Claise, B., Aitken, P., Dressler, F., and G. Carle, "Information Model for Packet Sampling Exports", RFC 5477, March 2009.
- [RFC6313] Claise, B., Dhandapani, G., Aitken, P., and S. Yates, "Export of Structured Data in IP Flow Information Export (IPFIX)", RFC 6313, July 2011.
- [RFC7011] Claise, B., Trammell, B., and P. Aitken, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information", STD 77, RFC 7011, September 2013.

10.2. Informative References

- [IANA-IPFIX] IANA, "IP Flow Information Export (IPFIX) Entities", <<http://www.iana.org/assignments/ipfix>>.
- [IEEE802.1D] IEEE, "IEEE Standard for Local and metropolitan area networks: Media Access Control (MAC) Bridges", IEEE Std 802.1D-2004, June 2004.
- [ISO/IEC.7498-1:1994] International Organization for Standardization, "Information technology -- Open Systems Interconnection -- Basic Reference Model: The Basic Mode", ISO Standard 7498-1:1994, June 1996.
- [RFC2804] IAB and IESG, "IETF Policy on Wiretapping", RFC 2804, May 2000.
- [RFC7012] Claise, B. and B. Trammell, "Information Model for IP Flow Information Export (IPFIX)", RFC 7012, September 2013.
- [RFC7013] Trammell, B. and B. Claise, "Guidelines for Authors and Reviewers of IP Flow Information Export (IPFIX) Information Elements", BCP 184, RFC 7013, September 2013.

Appendix A. Frame Formats

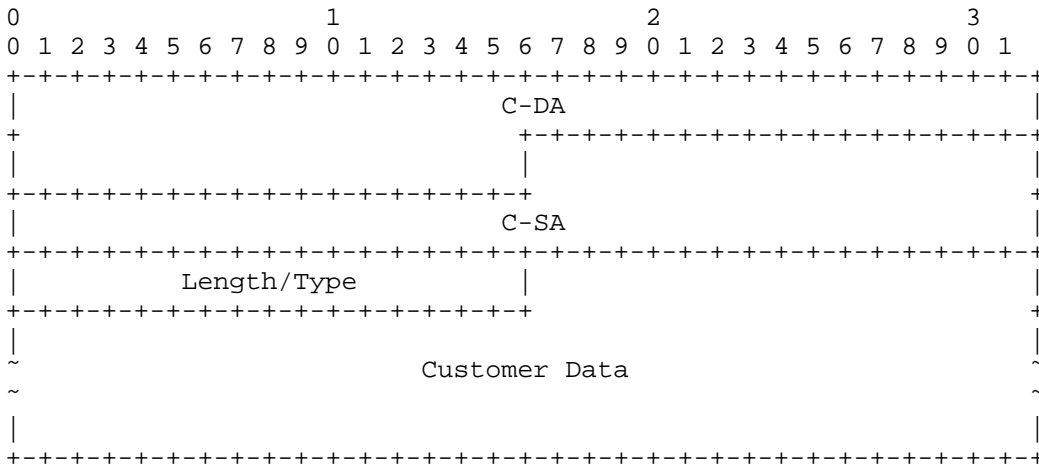


Figure A-1: Untagged Frame Format

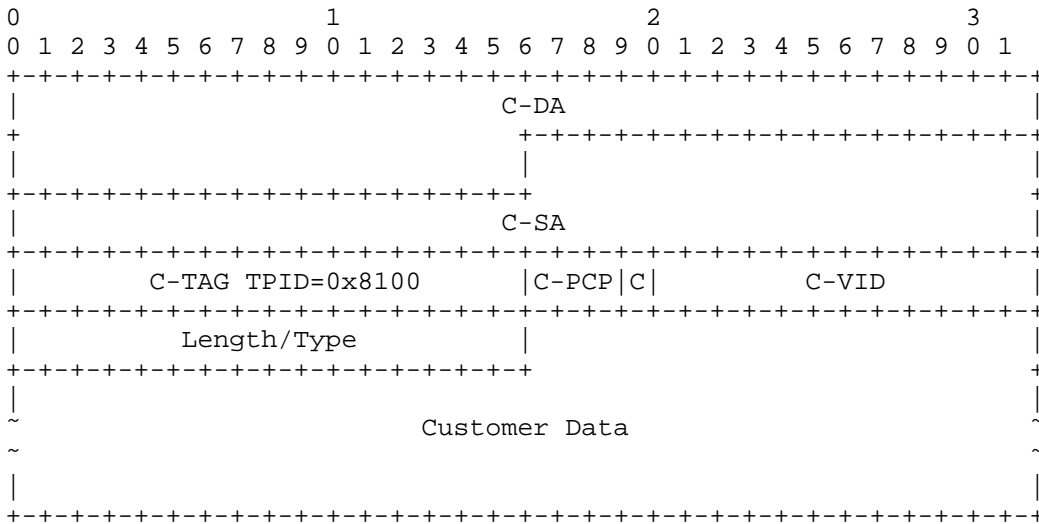


Figure A-2: C-TAG Tagging Frame Format

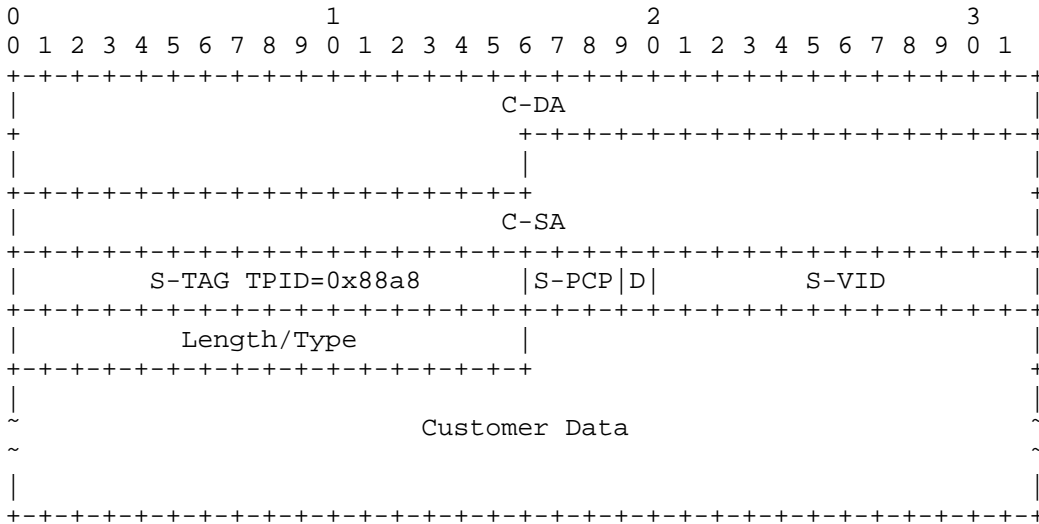


Figure A-3: S-TAG Tagging Frame Format in Provider Bridged Networks

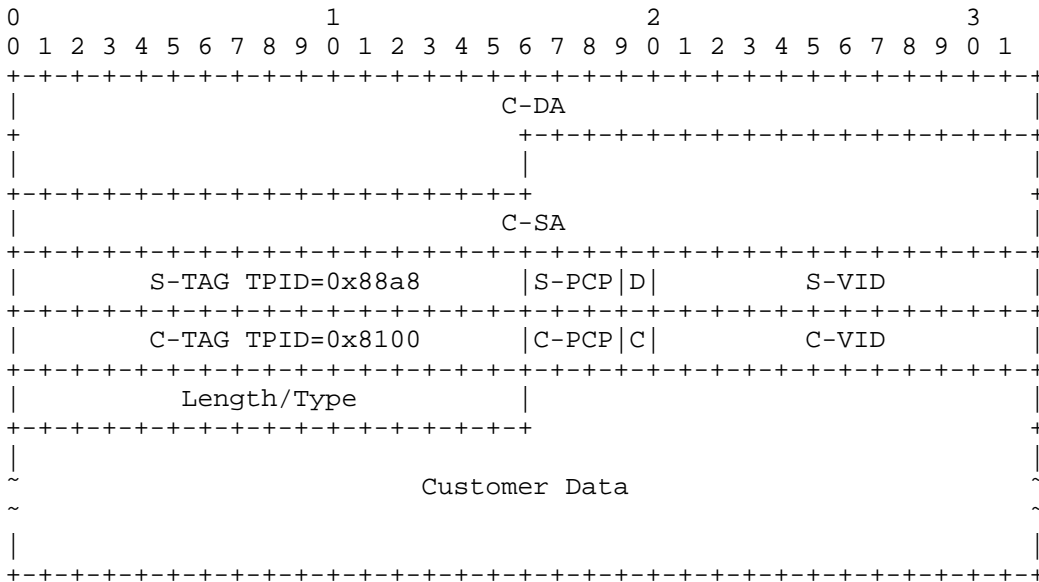


Figure A-4: S-TAG and C-TAG Tagging Frame Format in Provider Bridged Networks

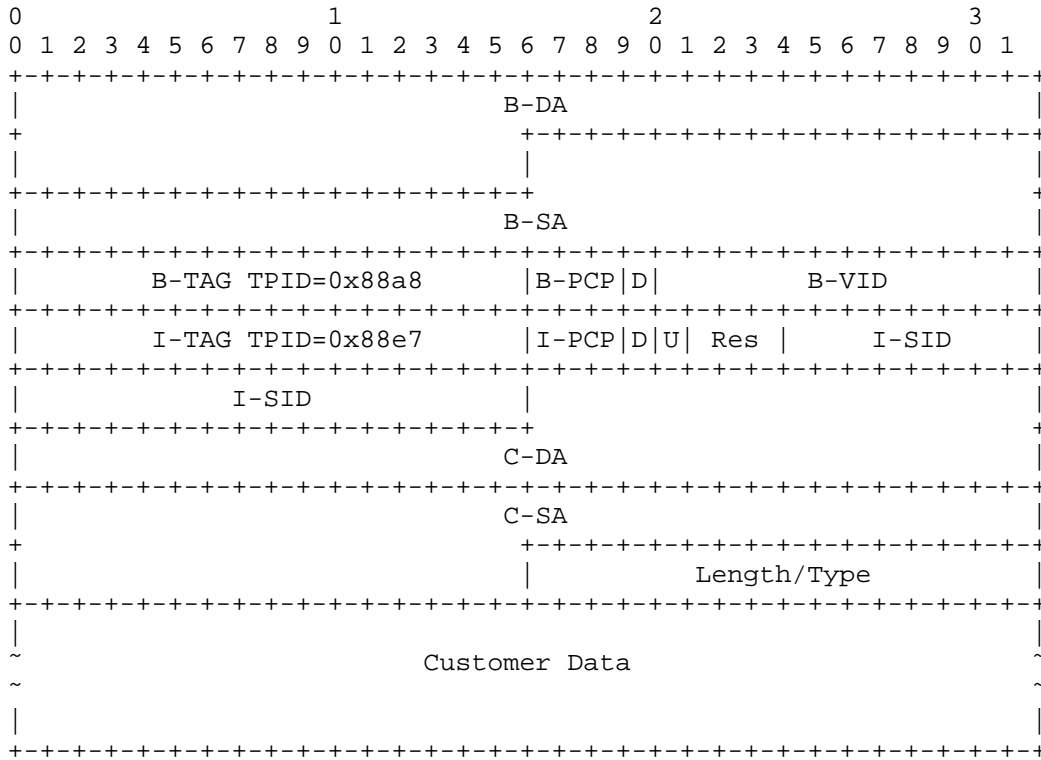


Figure A-5: B-TAG and I-TAG Tagging Frame Format in Provider Backbone Bridged Networks

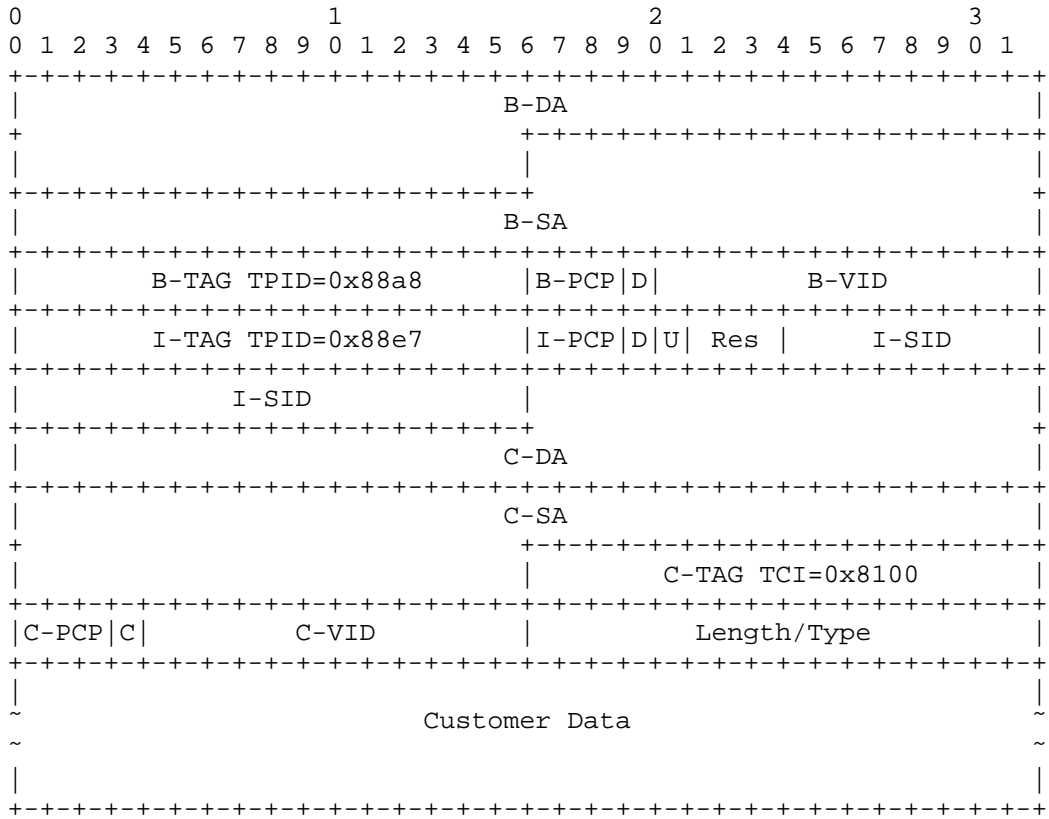


Figure A-6: B-TAG, I-TAG, and C-TAG Tagging Frame Format in Provider Backbone Bridged Networks

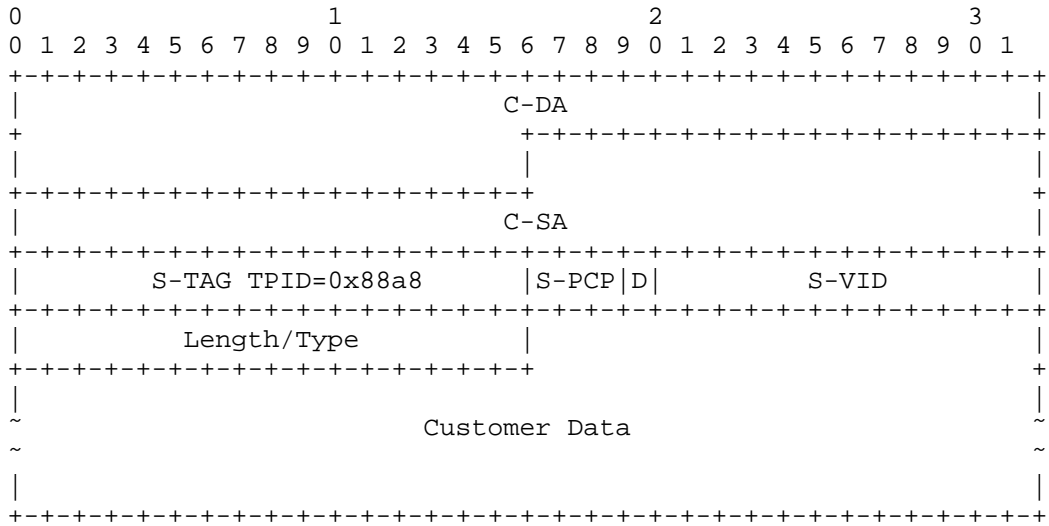


Figure A-7: S-TAG Tagging Frame Format for S-channel over the Link between an End Station and Its Adjacent Bridge

Note: The frame format in Figure A-7 is identical to the format in Figure A-3.

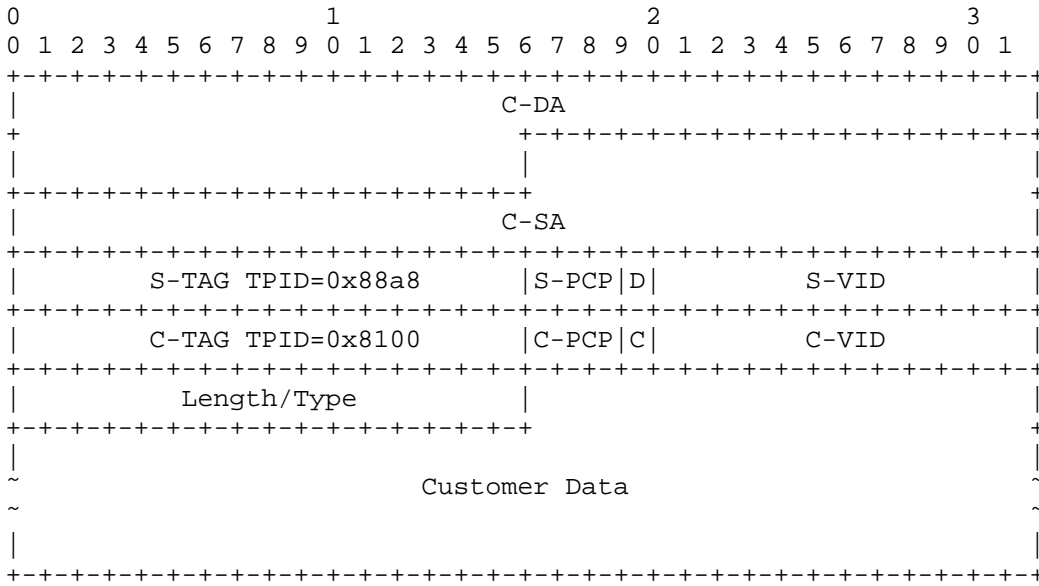


Figure A-8: S-TAG and C-TAG Tagging Frame Format over the Link between an End Station and Its Adjacent Bridge

Note: The frame format in Figure A-8 is identical to the format in Figure A-4.

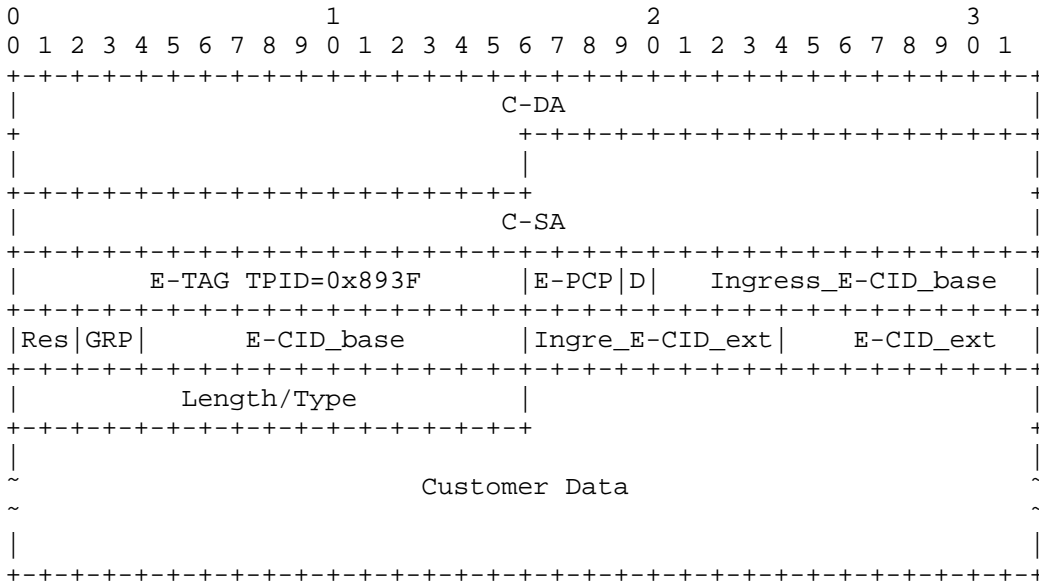


Figure A-9: E-TAG Tagging Frame Format over the Link between a Controlling Bridge and a Bridge Port Extender

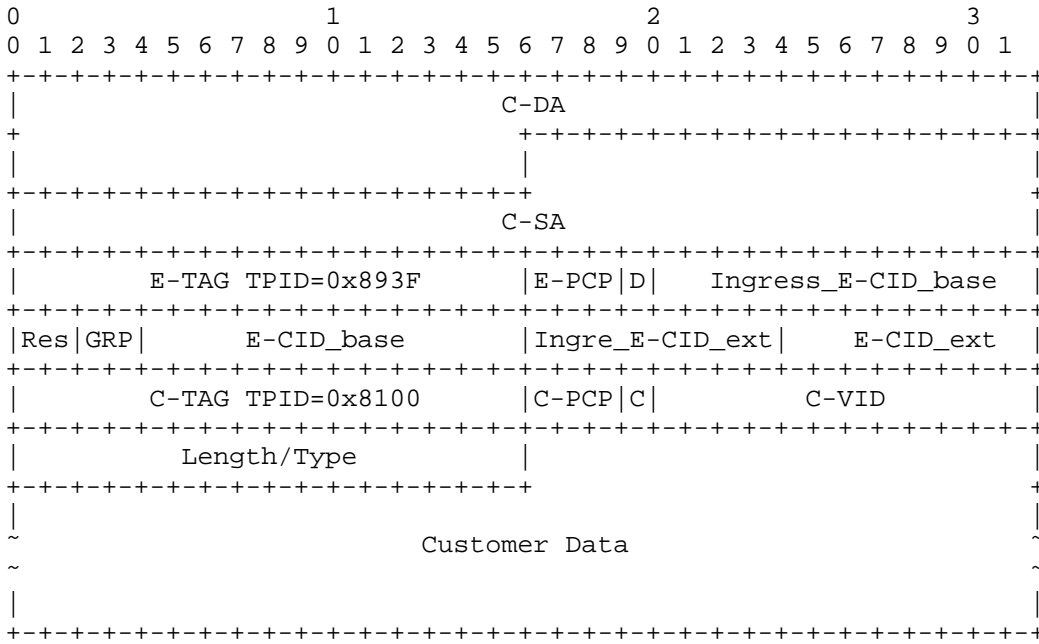


Figure A-10: E-TAG and C-TAG Tagging Frame Format over the Link between a Controlling Bridge and a Bridge Port Extender

Appendix B. Template Format Example

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9
Set ID (2)										Length																													
Template ID (256)										Field Count (8)																													
ingressInterface (10)										Field Length (4)																													
egressInterface (14)										Field Length (4)																													
observationTimeSeconds (322)										Field Length (8)																													
dataLinkFrameSize (312)										Field Length (2)																													
dataLinkFrameSection (315)										Field Length (65535)																													
dataLinkFrameType (408)										Field Length (2)																													
sectionOffset (409)										Field Length (2)																													
sectionExportedOctets (410)										Field Length (2)																													

Figure B-1: Template Format Example

Authors' Addresses

Shingo Kashima
Nippon Telegraph and Telephone Corporation
1-5-1 Otemachi
Chiyoda-ku, Tokyo 100-8116
Japan

Phone: +81 3 6838 5267
EMail: kashima@nttv6.net

Atsushi Kobayashi
Nippon Telegraph and Telephone East Corporation
3-19-2 Nishi-shinjuku
Shinjuku-ku, Tokyo 163-8019
Japan

Phone: +81 3 5359 4351
EMail: akoba@nttv6.net

Paul Aitken
Cisco Systems, Inc.
96 Commercial Quay
Commercial Street, Edinburgh EH6 6LX
United Kingdom

Phone: +44 131 561 3616
EMail: paitken@cisco.com

