IP Flow Information Export
Internet-Draft
Intended status: Standards Track
Expires: September 13, 2012

S. Kashima
K. Nakata
NTT
A. Kobayashi
NTT East

March 12, 2012

Information Elements for Data Link Layer Traffic Measurement
draft-kashima-ipfix-data-link-layer-monitoring-07

Abstract

This document describes Information Elements related to 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 Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on September 13, 2012.

Table of Contents

1. Introduction

Ethernet [IEEE802.1D] and VLAN (Virtual LAN) [IEEE802.1Q-2005] technologies used to be used only in Local Area Networks. Recently, they have been used in Wide Area Networks, e.g., L2-VPN services. Accordingly, the IEEE802.1Q standard has been enhanced to IEEE802.1ad [IEEE802.1ad-2005] and IEEE802.1ah [IEEE802.1ah-2008]. And, Ethernet in data center also has been enhanced for server virtualization and I/O consolidation.

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

IPFIX/PSAMP helps to resolve these problems. However, the PSAMP Information Model [RFC5477] and the IPFIX Information Model [RFC5101] are not yet enough for Information Elements related to data link layer, e.g., Ethernet header forms. This document describes the Information Elements related to data link layers that enable a more sophisticated traffic measurement method.

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 Summary

Provider Bridge [IEEE802.1ad-2005] and Provider Backbone Bridge [IEEE802.1ah-2008], which are standards for the Wide-Area Ethernet, are described below.

- In Provider Bridge [IEEE802.1ad-2005], 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.
- In Provider Backbone Bridge [IEEE802.1ah-2008], 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 on IEEE802.1ad

[IEEE802.1ad-2005] 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 B-PCP, B-DEI, I-PCP, and I-DEI.

Provider Backbone Bridge enables a wide-area Ethernet service to be improved from a flat network to a hierarchical network co-existing Provider Bridge and Provider Backbone Bridge.

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

2.2. Data Center Bridging Summary

In data center networks, Ethernet needs to be enhanced to provide the flexibility, mobility for server virtualization, and I/O consolidation. In IEEE802.1 Data Center Bridging Task Group, several Ethernet header formats are proposed to enable a simplifying networks and server managements.

The one of the enhanced methods is Bridge Port Extension [IEEE802.1BR], which brings a traffic exchange point to upper bridges. Bridge Port Extension introduces a Ethernet format named Extension Tag (E-TAG) in addition to existing Service VLAN Tag (S-TAG) and Customer VLAN Tag (C-TAG) to move the policy enhancement to upper bridges in data center network. On the other hand, the complexity for traffic measurement would be increased, because multiple Ethernet header formats as shown in Appendix B co-exist in the same network.

2.3. Multiple Path Ethernet Summary

Inside data center networks, Ethernet need to be enhaced to support multiple path for effective utilization of bandwidth. Existing Ethernet does not support multiple path because it is based on STP (Spanning Tree Protocol) [IEEE802.1D].

Two solutions for multiple path Ethernet are discussed for standardization. One is Shortest Path Bridging [IEEE802.1aq] in IEEE 802, the other is TRILL (Transparent Interconnection of Lots of Links) [RFC6325] in IETF. Both solutions realize multiple path with IS-IS [ISO_IEC.10589_2002] and Layer 2 over Layer 2 encapsulation. Shortest Path Bridging uses IEEE802.1ah [IEEE802.1ah-2008] header as an encapsulation technology and TRILL uses an unique header as shown in Section 3 of [RFC6325].

In multiple path Ehternet network, a traffic measurement based on MAC address and VLAN Tag is more important than single path Ethernet network in which traffic measurement is based on only VLAN Tag.

2.4. VXLAN Summary

Inside and between data center networks, the existing Ethernet and VLAN technologies have two problems. One is that 12 bit of VLAN ID does not have enough length for identifying a lot of tenants. The other is that it is difficult to place same VLAN in multiple data centers.

The solution for the problems is VXLAN [I-D.mahalingam-dutt-dcops-vxlan]. VXLAN has 24 bit length of identifier named VXLAN Segment ID/VXLAN Network Identifier (VNI) and uses Layer 2 over Layer 3 encapsulation with the frame format as show in Section 5 of [I-D.mahalingam-dutt-dcops-vxlan].

3. New Information Elements

The following Information Elements whose ElementId are 31 from to 349 are necessary for enabling the IPFIX/PSAMP traffic measurement for data link layer, which is not limited to Ethernet because the method can be applied to other data link protocols as well.

The following Information Elements whose ElementId are 350 from to 354 are necessary for enabling the IPFIX/PSAMP traffic measurement for IEEE 802.1ah.

The following Information Elements whose ElementId are TBD01 from to TBD14 are octet counter or packet length for layer 2, and are necessary for enabling the IPFIX/PSAMP traffic measurement for data link layer.

Note that these are proposed IDs, subject to approval by IANA.

ID	Name
312	dataLinkFrameSize
315	dataLinkFrameSection
347	dataLinkFrameType
348	sectionOffset
349	sectionObservedOctets
350	dot1qServiceInstanceTag
351	dot1qServiceInstanceId
352	dot1qServiceInstancePriority
353	dot1qCustomerDestinationMacAddress
354	dot1qCustomerSourceMacAddress
TBD01	l2octetDeltaCount
TBD02	postMCastL2OctetDeltaCount
TBD03	postL2OctetDeltaCount
TBD04	minimumL2TotalLength
TBD05	maximumL2TotalLength
TBD06	l2octetTotalCount
TBD07	droppedL2OctetDeltaCount
TBD08	droppedL2OctetTotalCount
TBD09	ignoredL2OctetTotalCount
TBD10	notSentL2OctetTotalCount
TBD11	postL2OctetTotalCount
TBD12	postMCastL2OctetTotalCount
TBD13	l2octetDeltaSumOfSquares
TBD14	l2octetTotalSumOfSquares

3.1. 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

Status: current

3.2. dataLinkFrameSection

Description:

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

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

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

When the sectionObservedOctets 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 C.

The data link layer is defined in [ISO_IEC.7498-1_1994].

Abstract Data Type: octetArray

ElementId: 315

Status: current

3.3. dataLinkFrameType

Description:

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

The following data link types are defined here.

- 0x01 ETHERNET

Further values may be assigned by IANA.

The data link layer is defined in [ISO_IEC.7498-1_1994].

Abstract Data Type: unsigned16

Data Type Semantics: identifier

ElementId: 347

Status: current

3.4. 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.

Abstract Data Type: unsigned16

Data Type Semantics: quantity

ElementId: 348

Status: current

3.5. sectionObservedOctets

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 sectionObservedOctets. In this case, octets in the packet section beyond the sectionObservedOctets MUST follow the [RFC5101] rules for padding (ie, be composed of zero (0) valued octets).

Abstract Data Type: unsigned16

Data Type Semantics: quantity

ElementId: 349

Status: current

3.6. dot1qServiceInstanceTag

Description:

This Information Element, which may have 16 octets length, represents the Backbone Service Instance Tag (I-TAG) Tag Control Information (TCI) field of an Ethernet frame as described in [IEEE802.1ah-2008]. It encodes the Backbone Service Instance Priority Code Point (I-PCP), Drop Eligible Indicator (I-DEI), Use Customer Addresses (UCA), Backbone Service Instance Identifier (I-SID), Encapsulated Customer Destination Address (C-DA), Encapsulated Customer Source Address (C-SA) and reserved fileds.

Abstract Data Type: octetArray

Data Type Semantics: identifier

ElementId: 350

Status: current

3.7. 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 section 9.8 of [IEEE802.1ah-2008]. The structure and semantics within the Tag Control Information field are defined in IEEE802.1ah.

Abstract Data Type: unsigned32

Data Type Semantics: identifier

ElementId: 351

Status: current

3.8. 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 section 9.8 of [IEEE802.1ah-2008]. The structure and semantics within the Tag Control Information field are defined in IEEE802.1ah.

Abstract Data Type: unsigned8

Data Type Semantics: identifier

ElementId: 352

Status: current

3.9. 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 section 9.8 of [IEEE802.1ah-2008]. The structure and semantics within the Tag Control Information field are defined in IEEE802.1ah.

Abstract Data Type: macAddress

Data Type Semantics: identifier

ElementId: 353

Status: current

3.10. 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 section 9.8 of [IEEE802.1ah-2008]. The structure and semantics within the Tag Control Information field are defined in IEEE802.1ah.

Abstract Data Type: macAddress

Data Type Semantics: identifier

ElementId: 354

Status: current

3.11. l2OctetDeltaCount

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.
memo: layer 2 version of octetDeltaCount (field #1)

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

ElementId: TBD01

Status: current

Units: octets

3.12. 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.
memo: layer 2 version of postMCastOctetDeltaCount (field #20)

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

ElementId: TBD02

Status: current

Units: octets

3.13. postL2OctetDeltaCount

Description:

The definition of this Information Element is identical to the definition of Information Element 'l2OctetDeltaCount', except that it reports a potentially modified value caused by a middlebox function after the packet passed the Observation Point.
memo: layer 2 version of postOctetDeltaCount (field #23)

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

ElementId: TBD03

Status: current

Units: octets

3.14. minimumL2TotalLength

Description:

Layer 2 length of the smallest packet observed for this Flow. The packet length includes the layer 2 header(s) length and the layer 2 payload length.
memo: layer 2 version of minimumIpTotalLength (field #25)

Abstract Data Type: unsigned64

ElementId: TBD04

Status: current

Units: octets

3.15. maximumL2TotalLength

Description:

Layer 2 length of the largest packet observed for this Flow. The packet length includes the layer 2 header(s) length and the layer 2 payload length.
memo: layer 2 version of maximumIpTotalLength (field #26)

Abstract Data Type: unsigned64

ElementId: TBD05

Status: current

Units: octets

3.16. l2octetTotalCount

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.
memo: layer 2 version of octetTotalCount (field #85)

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: TBD06

Status: current

Units: octets

3.17. 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 2payload.
memo: layer 2 version of droppedOctetDeltaCount (field #132)

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

ElementId: TBD07

Status: current

Units: octets

3.18. droppedL2OctetTotalCount

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.
memo: layer 2 version of droppedOctetTotalCount (field #134)

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: TBD08

Status: current

Units: octets

3.19. 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.
memo: layer 2 version of ignoredOctetTotalCount (field #165)

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: TBD09

Status: current

Units: octets

3.20. 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.
memo: layer 2 version of notSentOctetTotalCount (field #168)

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: TBD10

Status: current

Units: octets

3.21. postL2OctetTotalCount

Description:

The definition of this Information Element is identical to the definition of Information Element 'l2OctetTotalCount', except that it reports a potentially modified value caused by a middlebox function after the packet passed the Observation Point.
memo: layer 2 version of postOctetTotalCount (field #171)

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: TBD11

Status: current

Units: octets

3.22. 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.
memo: layer 2 version of postMCastOctetTotalCount (field #175)

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: TBD12

Status: current

Units: octets

3.23. l2octetDeltaSumOfSquares

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.
memo: layer 2 version of octetDeltaSumOfSquares (field #198)

Abstract Data Type: unsigned64

Data Type Semantics: deltaCounter

ElementId: TBD13

Status: current

Units: octets

3.24. l2octetTotalSumOfSquares

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.
memo: layer 2 version of octetTotalSumOfSquares (field #199)

Abstract Data Type: unsigned64

Data Type Semantics: totalCounter

ElementId: TBD14

Status: current

Units: octets

4. Modification of Existing Information Elements Related to Packet Section

--- This is open issue. ---

New Information Elements related to packet section (ie, sectionOffset and sectionObservedOctets) can be applied to not only dataLinkFrameSection but also all kinds of packet section. In this case, existing Information Elements Description should be modified as follows:

4.1. ipHeaderPacketSection

Description:

TBD

Abstract Data Type: octetArray

ElementId: 313

Status: current

4.2. ipPayloadPacketSection

Description:

TBD

Abstract Data Type: octetArray

ElementId: 314

Status: current

4.3. mplsLabelStackSection

Description:

TBD

Abstract Data Type: octetArray

ElementId: 316

Status: current

4.4. mplsPayloadPacketSection

Description:

TBD

Abstract Data Type: octetArray

ElementId: 317

Status: current

5. Modification of Existing Information Elements Related to VLAN Tag

Information Elements related to Backbone Service Instance Tag (I-TAG) and Backbone VLAN Tag (B-TAG) are required in order to monitor IEEE802.1ah traffic with IPFIX/PSAMP. New Information Elements related to I-TAG are added in section 3 because I-TAG has different structure and semantics from Service VLAN Tag (S-TAG) and Customer VLAN Tag (C-TAG). On the other hand, Information Elements related to B-TAG reuse existing Information Elements related to C-TAG and S-TAG because B-TAG has same structure and semantics with C-TAG and S-TAG, Though they reuse existing Information Elements, it required to modify existing Descriptions and Reference 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 as described in section 3.5.5 of [IEEE802.3-2005]. The structure and semantics within the Tag Control Information field are defined in IEEE P802.1Q. In case of a QinQ frame it represents the outer tag's VLAN identifier, in case of an IEEE 802.1ad frame it represents the Service VLAN identifier in the S-TAG Tag Control Information (TCI) field as described in [IEEE802.1ad-2005] and in case of an IEEE 802.1ah frame it represents the Backbone VLAN identifier in the B-TAG Tag Control Information (TCI) field as described in [IEEE802.1ah-2008].

Abstract Data Type: unsigned16

Data Type Semantics: identifier

ElementId: 243

Status: current

Reference:

(1) [IEEE802.3-2005]
(2) [IEEE802.1ad-2005]
(3) [IEEE802.1ah-2008]

5.2. dot1qPriority

Description:

The value of the 3-bit User Priority portion of the Tag Control Information field of an Ethernet frame as described in section 3.5.5 of [IEEE802.3-2005]. The structure and semantics within the Tag Control Information field are defined in IEEE P802.1Q. In case of a QinQ frame it represents the outer tag's 3-bit Class of Service (CoS) identifier, in case of an IEEE 802.1ad frame it represents the 3-bit Priority Code Point (PCP) portion of the S-TAG Tag Control Information (TCI) field as described in [IEEE802.1ad-2005] and in case of an IEEE 802.1ah frame it represents the 3-bit Priority Code Point (PCP) portion of the B-TAG Tag Control Information (TCI) field as described in [IEEE802.1ah-2008].

Abstract Data Type: unsigned8

Data Type Semantics: identifier

ElementId: 244

Status: current

Reference:

(1) [IEEE802.3-2005]
(2) [IEEE802.1ad-2005]
(3) [IEEE802.1ah-2008]

5.3. dot1qCustomerVlanId

Description:

In case of a QinQ frame it represents the inner tag's (*) VLAN identifier, in case of an IEEE 802.1ad frame it represents the Customer VLAN identifier in the C-TAG Tag Control Information (TCI) field as described in [IEEE802.1ad-2005] and in case of an IEEE 802.1ah frame it represents the Customer VLAN identifier in the C-TAG Tag Control Information (TCI) field in encapsulated IEEE 802.1ad or IEEE 802.1Q frame as described in [IEEE802.1ah-2008].

(*) Note: the 802.1Q tag directly following the outer one.

Abstract Data Type: unsigned16

Data Type Semantics: identifier

ElementId: 245

Status: current

Reference:

(1) [IEEE802.1ad-2005]
(2) [IEEE802.1Q-2005]
(3) [IEEE802.1ah-2008]

5.4. dot1qCustomerPriority

Description:

In case of a QinQ frame it represents the inner tag's (*) Class of Service (CoS) identifier, in case of an IEEE 802.1ad frame it represents the 3-bit Priority Code Point (PCP) portion of the C-TAG Tag Control Information (TCI) field as described in [IEEE802.1ad-2005] and in case of an IEEE 802.1ah frame it represents the 3-bit Priority Code Point (PCP) portion of the C-TAG Tag Control Information (TCI) field in encapsulated IEEE 802.1ad or IEEE 802.1Q frame as described in [IEEE802.1ad-2005]. (*) Note: the 802.1Q tag directly following the outer one.

Abstract Data Type: unsigned8

Data Type Semantics: identifier

ElementId: 246

Status: current

Reference:

(1) [IEEE802.1ad-2005]
(2) [IEEE802.1Q-2005]
(3) [IEEE802.1ah-2008]

6. The relationship between Ethernet header filelds and Information Elements

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

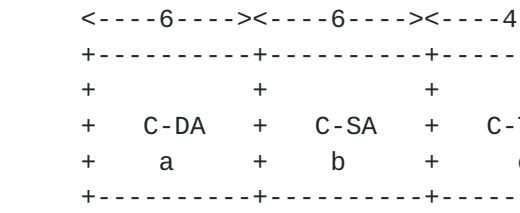

a.(Existing Information Element) destinationMacAddress 80
b.(Existing Information Element) sourceMacAddress 56
c.(Existing Information Elements) dot1qVlanId 243, dot1qPriority 244
d.(Existing Information Element) ethernetType 256

Figure 1: IEEE 802.1Q header fields

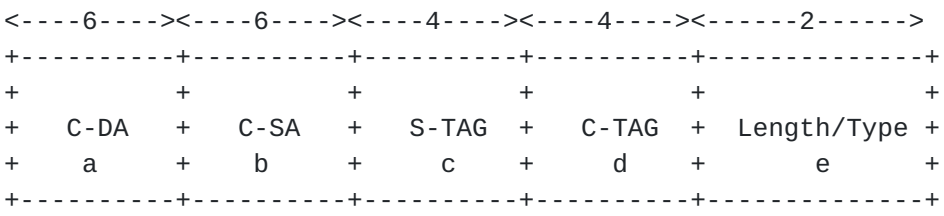

a.(Existing Information Element) destinationMacAddress 80
b.(Existing Information Element) sourceMacAddress 56
c.(Existing Information Elements) dot1qVlanId 243, dot1qPriority 244
d.(Existing Information Elements) dot1qCustomerVlanId 245, dot1qCustomerPriority 246
e.(Existing Information Element) ethernetType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

Figure 2: IEEE 802.1ad header fields

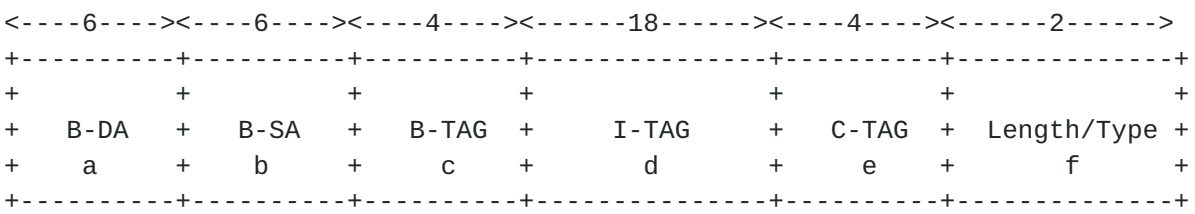

a.(Existing Information Element) destinationMacAddress 80
b.(Existing Information Element) sourceMacAddress 56
c.(Existing Information Elements) dot1qVlanId 243, dot1qPriority 244
d.(New Information Elements) defined in section 3.6, 3.7, 3.8, 3.9 and 3.10 of this draft
e.(Existing Information Elements) dot1qCustomerVlanId 245, dot1qCustomerPriority 246
f.(Existing Information Element) ethernetType 256

Figure 3: IEEE 802.1ah header fields

7. Security Considerations

The recommendations in this document do not introduce any additional security issues to those already mentioned in [RFC5101] and [RFC5477].

8. IANA Considerations

This document requests that the Information Element IDs are allocated as shown in section 3

In addition, the dataLinkFrameType Information Element requires the creation of new IANA registries.

And this document requests that the existing Information Element Description are modified as shown in section 4 and 5

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

9.2. Informative References

[I-D.mahalingam-dutt-dcops-vxlan]

C., Mahalingam, M., Duda, K., and P. Agarwal, "VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", draft-mahalingam-dutt-dcops-vxlan-01 (work in progress), February 2012.

[IEEE802.1BR]
IEEE Computer Society, "Draft Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks Amendment: Bridge Port Extension", IEEE Std P802.1BR/D1.1, February 2011.

[IEEE802.1D]
IEEE Computer Society, "IEEE Standards for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges", IEEE Std 802.1D-2004, June 2004.

[IEEE802.1Q-2005]
IEEE Computer Society, "IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2005, May 2006.

[IEEE802.1ad-2005]
IEEE Computer Society, "IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks Amendment 4: Provider Bridges", IEEE Std 802.1ad-2005, May 2006.

[IEEE802.1ah-2008]
IEEE Computer Society, "IEEE Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks Amendment 7: Provider Backbone Bridges", IEEE Std 802.1ah-2008, August 2008.

[IEEE802.1aq]
IEEE Computer Society, "Draft Standards for Local and Metropolitan Area Networks: Virtual Bridged Local Area Networks Amendment: Shortest Path Bridging", IEEE Std P802.1aq/D3.6, February 2011.

[IEEE802.3-2005]
IEEE Computer Society, "IEEE Standard for 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 802.3-2005, December 2005.

[ISO_IEC.10589_2002]
International Organization for Standardization, "Information technology -- Telecommunications and information exchange between systems -- Intermediate System to Intermediate System intra-domain routing information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode network service (ISO 8473)", ISO Standard 10589:2002, November 2002.

[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.

[RFC2863] McCloghrie, K. and F. Kastenholz, "The Interfaces Group MIB", RFC 2863, June 2000.

[RFC3954] Claise, B., "Cisco Systems NetFlow Services Export Version 9", RFC 3954, October 2004.

[RFC5101] Claise, B., "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information", RFC 5101, January 2008.

[RFC5475] Zseby, T., Molina, M., Duffield, N., Niccolini, S., and F. Raspall, "Sampling and Filtering Techniques for IP Packet Selection", RFC 5475, March 2009.

[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.

[RFC6325] Perlman, R., Eastlake, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011.

Appendix A. Frame Formats in Wide-Area Ethernet Network

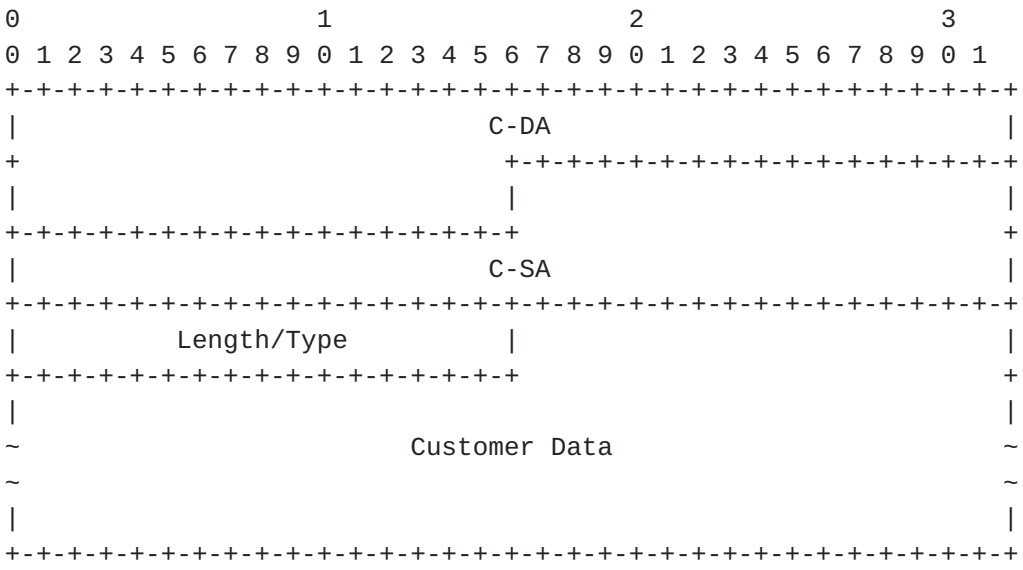

Figure A-1: IEEE802.1D Frame Format in Customer Bridged Network

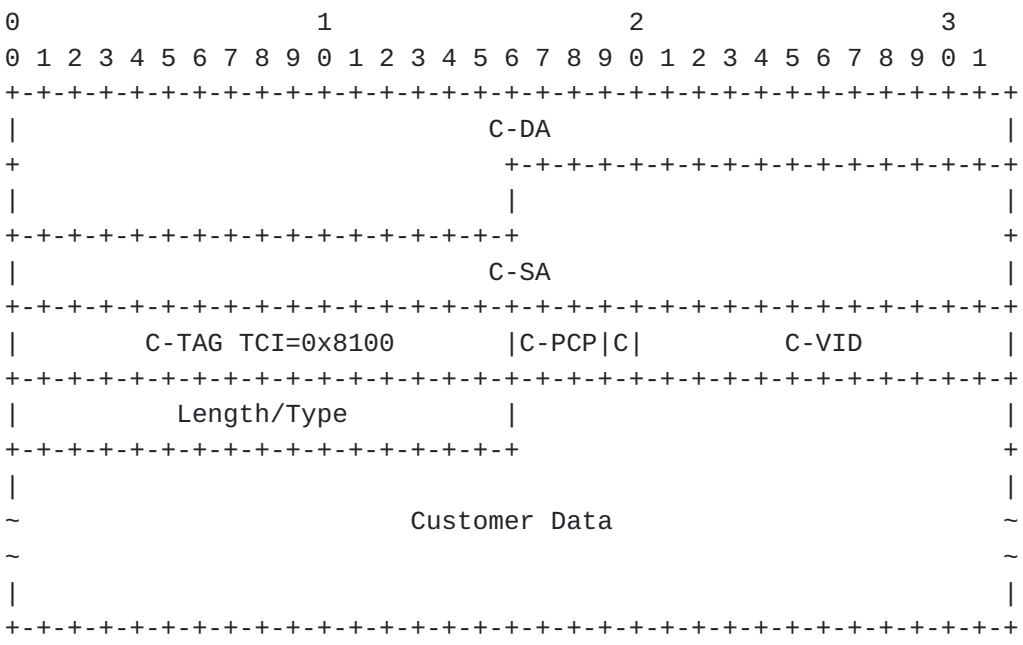

Figure A-2: IEEE802.1Q Frame Format in Customer Bridged Network

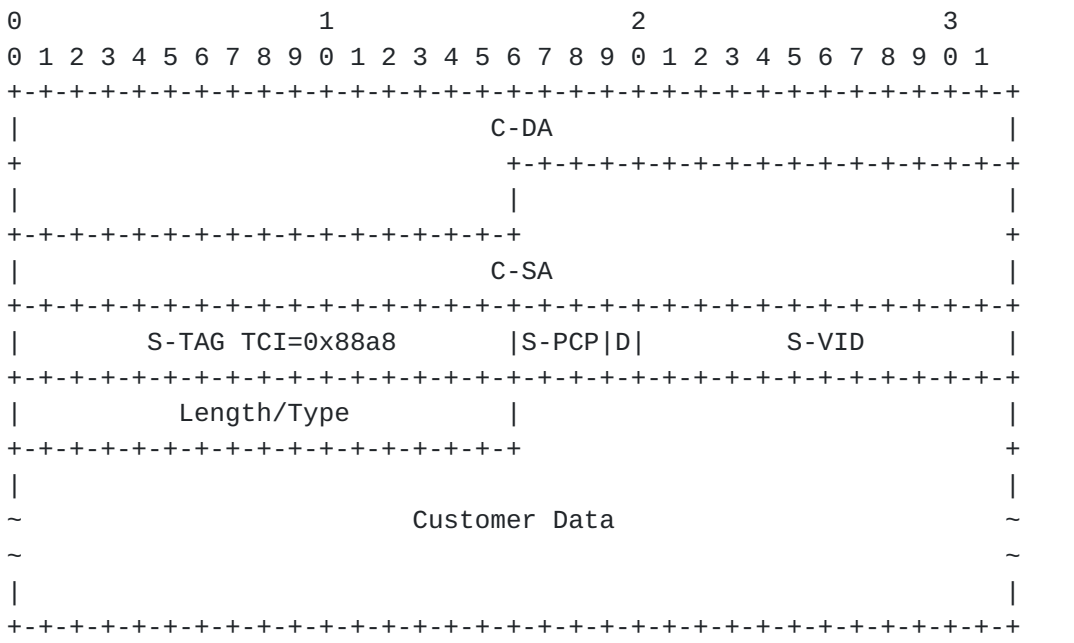

Figure A-3: IEEE802.1ad (no C-Tag) Frame Format in Provider Bridged Network

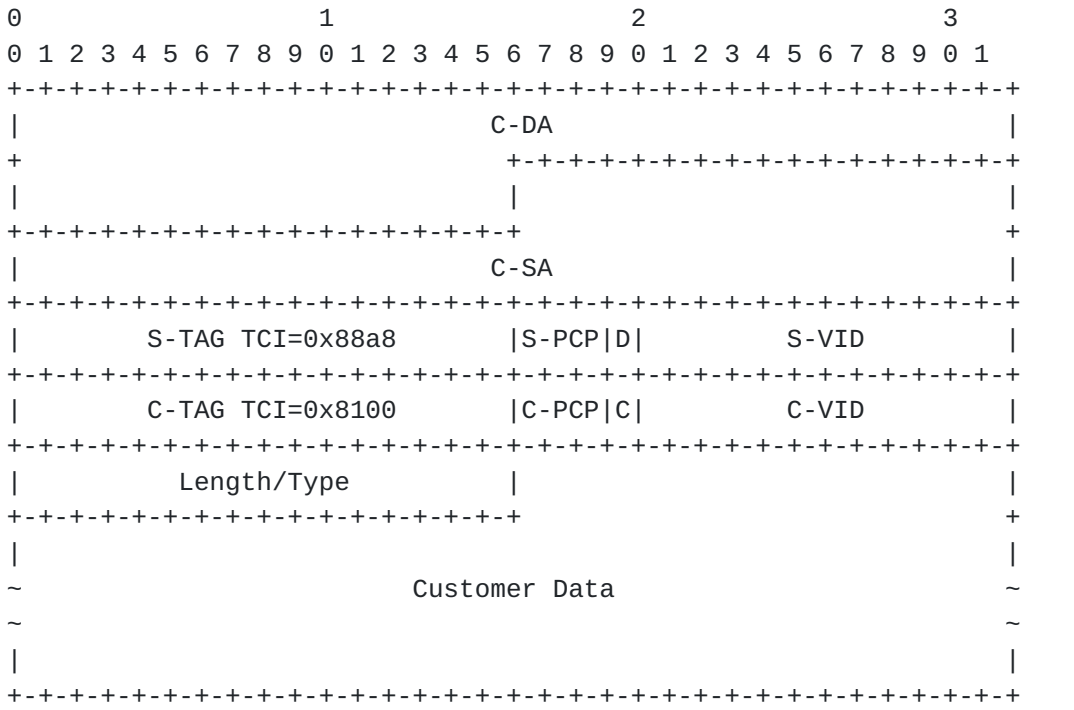

Figure A-4: IEEE802.1ad (C-Tagged) Frame Format in Provider Bridged Network

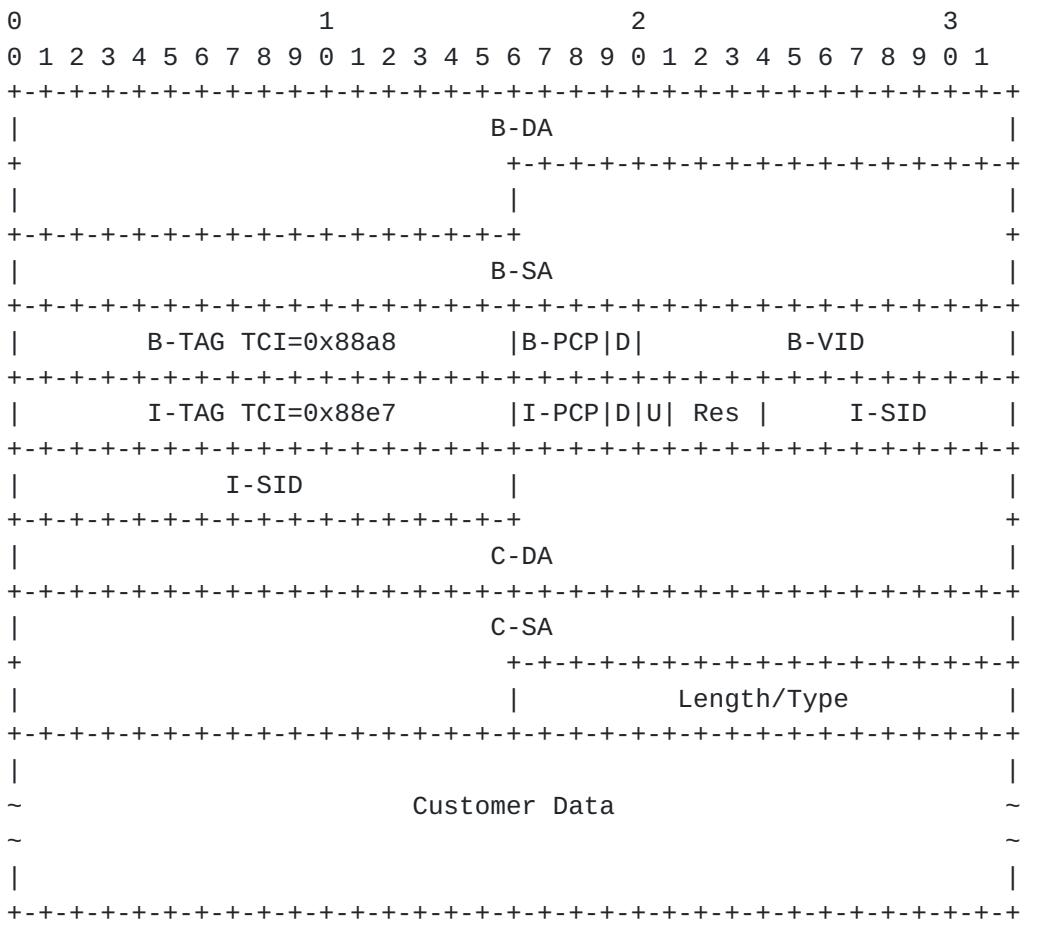

Figure A-5: IEEE802.1ah (no C-Tag) Frame Format in Provider Backbone Bridged Network

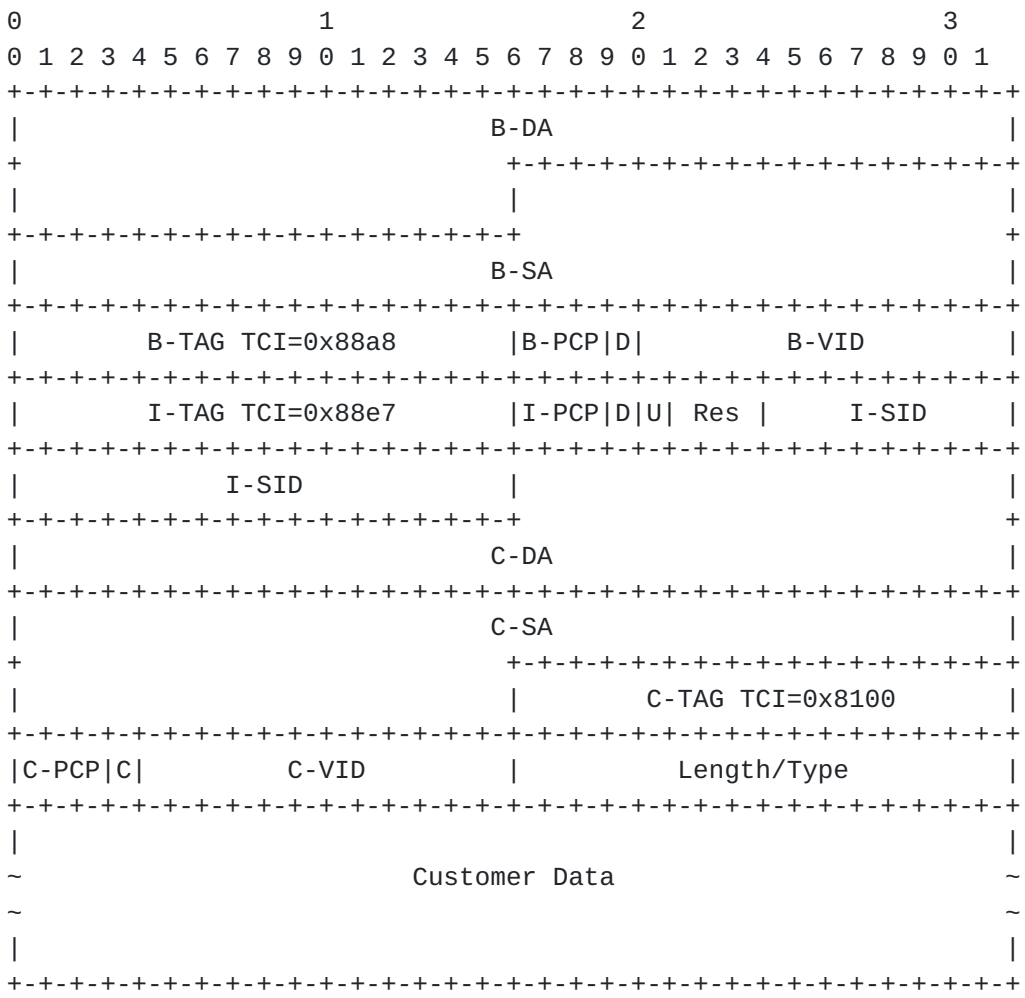

Figure A-6: IEEE802.1ah (C-Tagged) Frame Format in Provider Backbone Bridged Network

Appendix B. Frame Formats in Data Center Network

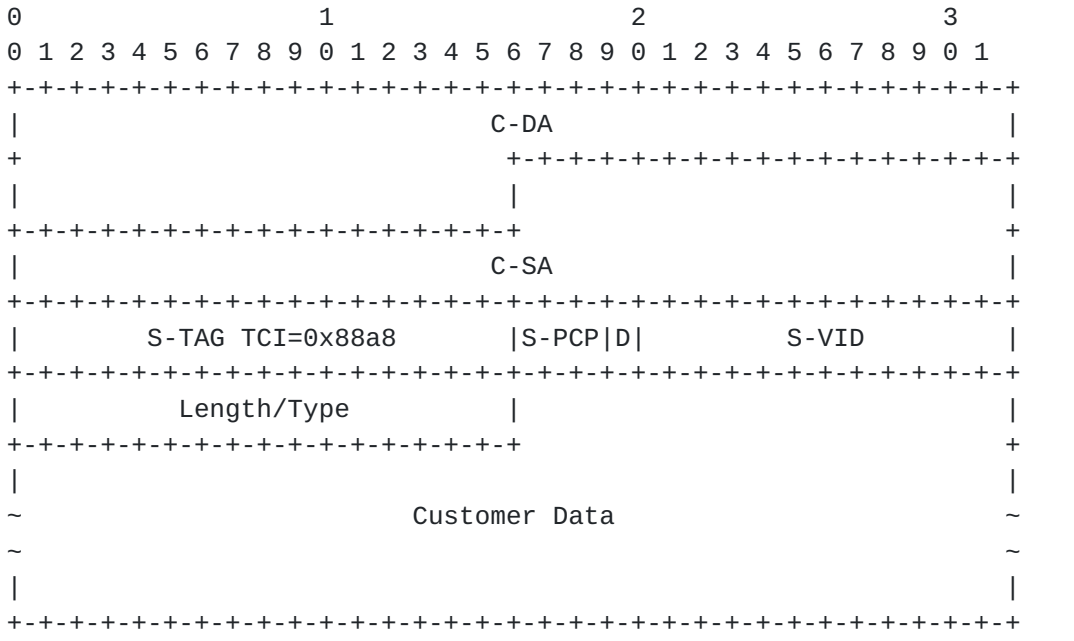

Figure B-1: IEEE802.1BR (S-TAG) Frame Format in Data Center Network

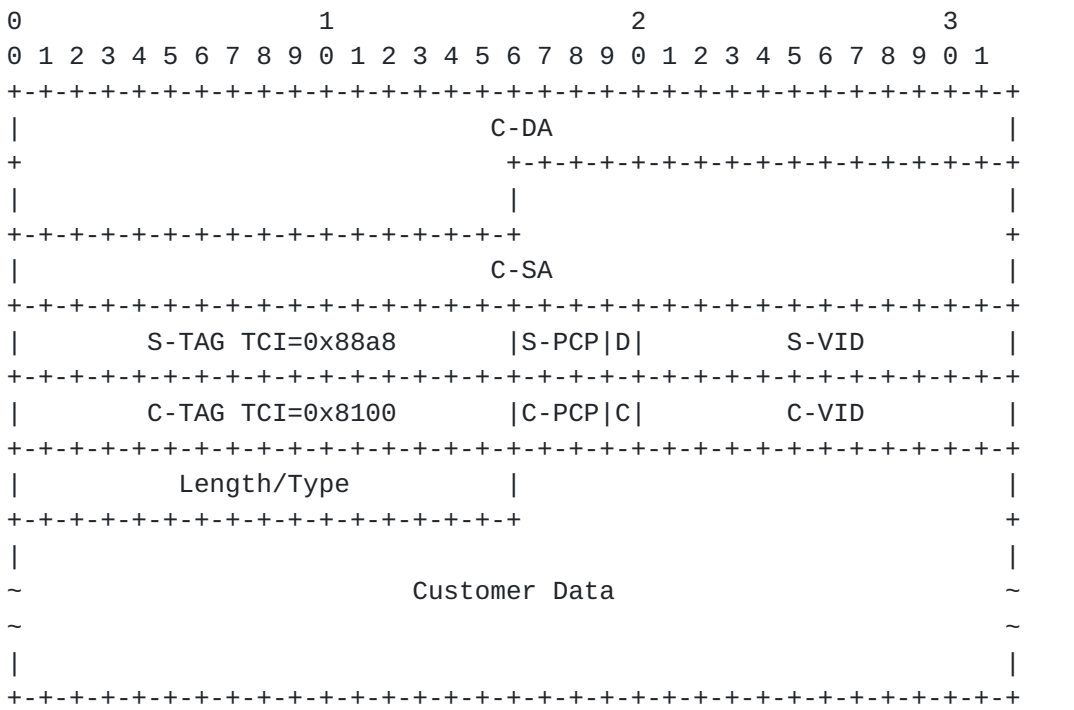

Figure B-2: IEEE802.1BR (S-TAG+C-TAG) Frame Format in Data Center Network

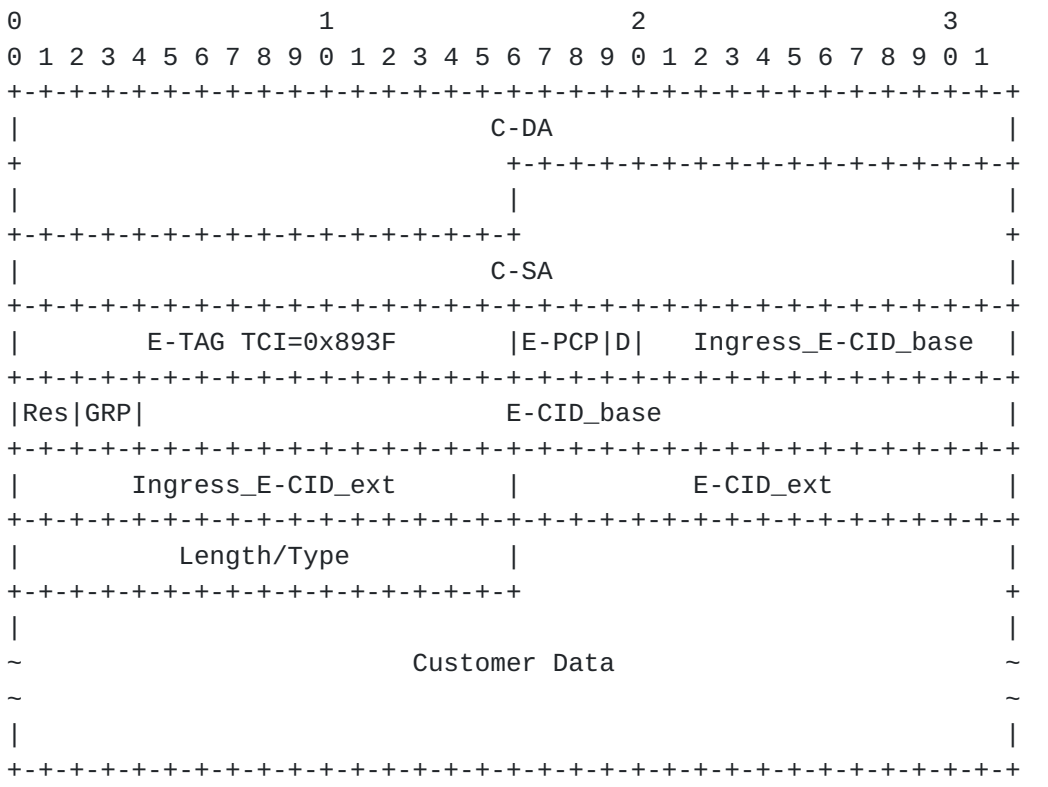

Figure B-3: IEEE802.1BR (E-TAG) Frame Format in Data Center Network

Figure B-4: IEEE802.1BR (E-TAG+C-TAG) Frame Format in Data Center Network

Appendix C. Template Formats Example

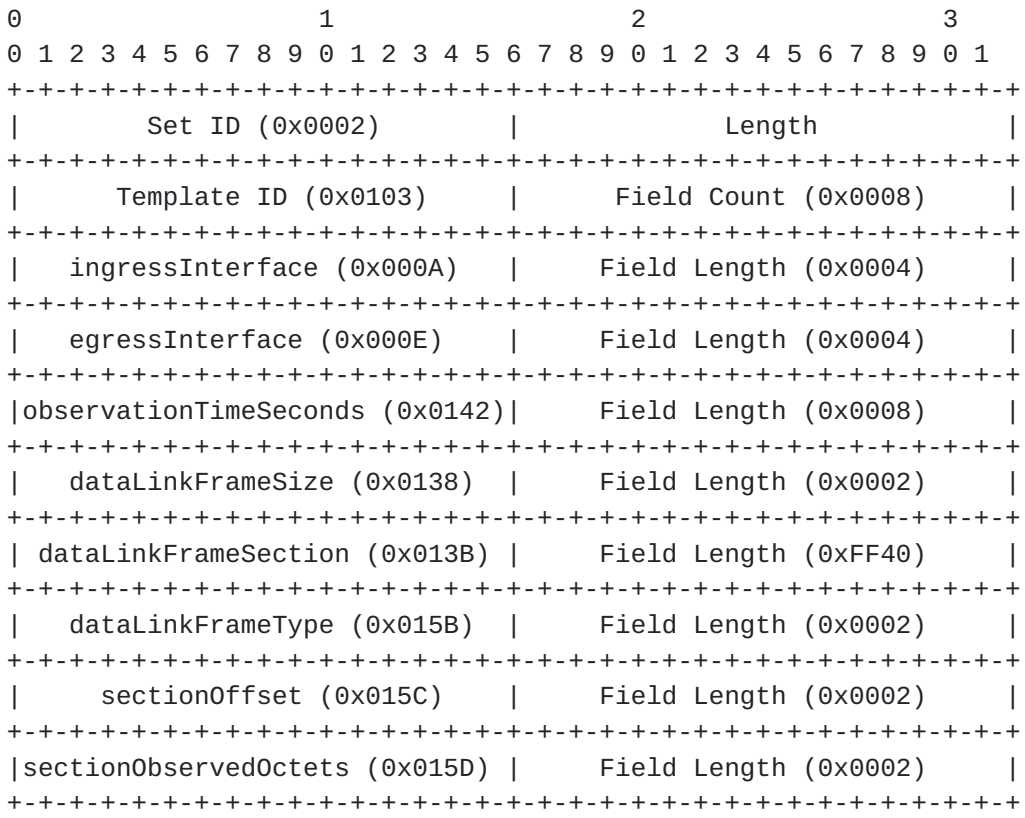

Figure C-1: Template Format Example

Authors' Addresses

Shingo Kashima
NTT Information Sharing Platform Lab.
Midori-Cho 3-9-11
Musashino-shi, Tokyo 180-8585
Japan

Phone: +81 422 59 3894
Email: kashima@nttv6.net

Kensuke Nakata
NTT Information Sharing Platform Lab.
Midori-Cho 3-9-11
Musashino-shi, Tokyo 180-8585
Japan

Phone: +81 422 59 6958
Email: nakata@nttv6.net

Atsushi Kobayashi
NTT EAST IT Innovation Department
25F 3-20-2 Nishi-shinjuku
Shinjuku-ku, Tokyo 163-1425
Japan

Phone: +81-3-5353-3636
Email: akoba@nttv6.net