

M A R V E L L

WHITE PAPER

Carrier Ethernet 2.0: A Chipmaker's Perspective

Tal Mizrahi
Uri Safrai
Marvell

June 2015

ABSTRACT

Over the past decade Ethernet has increasingly become a common and widely deployed technology in carrier networks. Carrier Ethernet 2.0 is a set of features and services that form the second generation of carrier networks, as defined by the Metro Ethernet Forum (MEF).

This white paper presents a brief overview of CE 2.0, and provides a chipmaker's perspective on CE 2.0, its main features, and its impact on network equipment silicon, with a focus on Marvell® Presteria®-DX devices.

Introduction

The Metro Ethernet Forum

The Metro Ethernet Forum (MEF) is an industry consortium, focused on the adoption of Carrier Ethernet networks and services. The forum is composed of service providers, carriers, network equipment vendors and other networking companies that share an interest in Metro Ethernet [9].

As opposed to other networking-related standard organizations, such as the Internet Engineering Task Force (IETF) and the IEEE 802.1 working group, that define networking **protocols**, the MEF is dedicated to defining how all the pieces of the puzzle fit together in Carrier Ethernet networks. The MEF defines network architectures, deployment scenarios and test suites. The MEF also defines the relationship and interaction between two main entities:

- Subscriber—the organization purchasing the Carrier Ethernet service.
- Service Provider—the organization providing the Carrier Ethernet service.

The MEF has a certification program that provides conformity testing to the MEF specifications.

What is Carrier Ethernet 2.0?

CE 2.0 is the second generation of services and networks defined by the MEF. In a nutshell, CE 2.0 was developed to address the challenges of carrier networks in the current decade, with a focus on 4G mobile backhaul networks. The core functionality of CE 2.0 is defined by **three key features** and **eight services**.

CE 2.0 Features

The three main features defined in CE 2.0 are:

- | | |
|----------------|---|
| Multi-CoS | Traffic is forwarded over Multiple Classes of Service (CoS). Furthermore, classes of service are associated with MEF-defined performance objectives and performance tiers, allowing consistent Quality of Service across multiple providers. |
| Managed | <p>CE 2.0 includes enhanced fault and traffic management capabilities.</p> <ul style="list-style-type: none"> • Operations, Administration and Maintenance (OAM): OAM ([7], [8]) is a set of mechanisms that allow providers and subscribers to detect failures and misconfigurations and to monitor the network performance. • Traffic management: Service providers use bandwidth enforcement (also known as policing) to guarantee that traffic, up to the contracted bandwidth, is allowed to be forwarded, whereas traffic that exceeds the Service Level Agreement (SLA) is discarded. CE 2.0 includes both ingress and egress bandwidth enforcement. |
| Interconnected | The MEF reference model focuses on two entities, the subscriber and the provider. CE 2.0 extends this model with a third entity, an operator. In cases where a |

service provider is not be able to reach all the subscriber sites, the provider can use a third party, an operator, to provide service to all the subscriber’s locations.

CE 2.0 Certification

Prior to CE 2.0, the MEF used to certify compliance to each specification document separately. Thus, certification was done separately for MEF 9 and MEF 14.

CE 2.0 is not defined in a single document or specification; it is a set of features and services that are defined in a number of MEF specifications. Hence CE 2.0 certification is provided on a **per-service** basis, rather than on a per-spec basis, as was done previously. The CE 2.0 Certification Blueprint [3] specifies the complete list of attributes to be verified in order to certify each CE 2.0 service. The MEF still provides certification to the pre-CE 2.0 feature-set, now referred to as CE 1.0.

The MEF provides two types of CE 2.0 certifications: for service providers and for equipment vendors.

CE 2.0 Services

CE 2.0 defines eight types of services, summarized in Table 1. While CE 1.0 includes three services, EPL, EVPL and EP-LAN, this set has been extended in CE 2.0 to include eight services, as shown in Table 1.

| | Port-based Service | VLAN-based Service |
|-----------------|--------------------|--------------------|
| E-Line | EPL | EVPL |
| E-LAN | EP-LAN | EVP-LAN |
| E-Tree | EP-Tree | EVP-Tree |
| E-Access | Access EPL | Access EVPL |

Table 1. CE 2.0 Services (services in dark shade were also present in CE 1.0).

This subsection surveys these eight service types [4].

We start by defining two commonly used terms in MEF specifications (depicted in Figure 1):

UNI The User-Network Interface is the interface connecting the Customer Equipment (CE) to the carrier Ethernet network.

EVC An Ethernet Virtual Connection is an association between two or more UNIs.




Figure 1. Ethernet Virtual Connection (EVC)

E-Line

An E-Line is a point-to-point EVC between two UNIs. An E-Line can have two possible flavors:

- **Port-based Ethernet Private Line (EPL)**. This mode (Figure 2) uses **all-to-one bundling**, where all traffic sent from a CE to a given port (UNI) is bundled as a single EVC.

- **VLAN-based Ethernet Virtual Private Line (EVPL).** This mode (Figure 3) performs **service-multiplexing** based on VLANs, i.e., each customer VLAN is bound to a separate EVC.




Figure 2. Ethernet Private Line (EPL): port-based all-to-one bundling




Figure 3. Ethernet Virtual Private Line (EVPL): VLAN-based service multiplexing

E-LAN

An E-LAN is a multipoint-to-multipoint EVC that can be either **port-based** or **VLAN-based**, as illustrated in Figure 4 and Figure 5, respectively.

Figure 4 illustrates a network in which the customer is spread over three sites, and an all-to-one bundling is used to bind customer traffic to the EVC. In Figure 5 two EVCs are used; a point-to-point EVC between site 1 and site 2, and a multipoint-to-multipoint EVC between sites 1, 3, and 4. VLAN-based multiplexing is used in site 1 to bind each VLAN to the respective EVC.




Figure 4. Ethernet Private LAN (EP-LAN): port-based all-to-one bundling




Figure 5. Ethernet Virtual Private LAN (EVP-LAN): VLAN-based service multiplexing

E-Tree

An E-Tree service is a rooted-multipoint EVC, where one UNI (the **root**) is permitted to send and receive traffic to other UNIs (**leafs**), but leaf-to-leaf communication is not permitted.

Figure 6 illustrates an example of a port-based E-Tree. The network in Figure 7 uses VLAN-based multiplexing of an E-Tree and an E-Line.




Figure 6. Ethernet Private Tree (EP-Tree): port-based all-to-one bundling




Figure 7. Ethernet Virtual Private Tree (EVP-Tree): VLAN-based service multiplexing

E-Access

An E-Access service defines interconnection between service providers; it allows an access provider (Figure 8) to reach out-of-franchise customer locations by using an **Operator Virtual Connection (OVC)**, operated by a retail service provider. E-Access appears to the customer as a conventional EVC, provided by the access service provider. E-Access can provide both port-based and VLAN-based services. The interface between the provider network and the operator network is called External Network to Network Interface (ENNI).




Figure 8. E-Access: Reaching out-of-franchise Customers via a Retail Service Provider

CE 2.0 from a Chipmaker's Perspective

As noted above, the MEF provides CE 2.0 certification for **services** and for **equipment**; no certification program has been defined for network equipment **silicon**. This fact is not surprising, as MEF specifications define functionality that affects both the data plane and the control plane, whereas silicons mostly implement the data plane functionality, and most of the control plane functionality is typically implemented in software.

Hence, CE 2.0 certification is a product of cooperation between network equipment vendors and chip vendors, allowing the combined product to be certified by the MEF.




Figure 9. CE 2.0 Certification Process: Cooperation between Silicon Vendor and Network Equipment Vendor

What does CE 2.0 require from packet processor chips?

CE 2.0 is defined by a number of MEF specifications, and includes a rich set of features and complex functionality. Hence, CE 2.0 has many implications on packet processor chips, including:

- Port-based and VLAN-based packet classification.
- VLAN tag manipulations.
- Service OAM.
- Provider tunneling technologies, such as VPLS and MPLS.
- Multi-CoS traffic differentiation.
- Traffic metering (policing).

Marvell CE 2.0 Features

The Marvell® Prestera®-DX family of packet processors includes a comprehensive feature set that addresses the requirements of Carrier Ethernet networks, and specifically of CE 2.0.

The Marvell Prestera-DX family includes two architectural building blocks that are key enablers of CE 2.0, **eBridge** and **FlexOAM**.




Figure 10. Marvell eBridge Architecture




Figure 11. Marvell FlexOAM Architecture

Marvell eBridge Architecture. The extended-bridging (eBridge) architecture is a unified architecture, implementing a hardware-based virtualization of interfaces and switching domains. It extends the traditional physical bridge port paradigm to a flexible paradigm supporting virtual interfaces called Extended Ports or **ePorts**. In addition, the architecture extends the 4K Virtual LAN (VLANs) defined by IEEE 802.1Q into larger switching/flooding domains called Extended VLANs or **eVLANs**.

The eBridge architecture is ideal for implementing CE 2.0:

- The eBridge architecture provides a straightforward and scalable abstraction for EVCs, using ePorts.
- eBridge allows flexible VLAN-based multiplexing and VLAN tag manipulations.
- eBridge enables an intuitive and scalable use of provider tunneling technologies such as VPLS, MPLS and PBB, as:
 - EVCs and tunneled EVCs are represented by the ePort abstraction, which is independently mapped into a physical interface (e.g., physical port, link aggregation group).
 - Switching/flooding domains are represented by the eVLAN abstraction, independently from the switching/flooding methods; an 802.1Q VLAN, a VPLS VSI, as well as switching/flooding domains of other provider technologies are supported using a single eVLAN abstraction that is not technology specific.

- eBridge flexible configuration of ePort/eVLAN abstraction attributes, and the independency of these attributes from the actual physical interface they are mapped to, allows for a single-pass line-rate any-to-any encapsulation conversion, including multi-target replications, where each copy may be encapsulated independently. This is an extremely powerful capability for ENNI and UNI nodes that interface multiple network encapsulation technologies.
- eBridge end-to-end Quality of Service (QoS) support - flexible ingress and egress QoS mapping tables enable network edge nodes that interface multiple QoS domains, a smooth and extremely flexible QoS translation upon network/technology boundary cross, thus allowing the provider to ensure end-to-end SLA preservation over multiple network technologies.

Marvell FlexOAM Architecture. FlexOAM provides generic hardware support for real-time processing of OAM traffic.

FlexOAM uses a programmable flow classification engine, allowing support for various existing OAM protocols, as well as proprietary or future ones.

Network failures, also known as loss of connectivity, are detected by the FlexOAM keepalive engine, which is in charge of reception and transmission of keepalive messages. FlexOAM performs accurate loss and delay measurement, using hardware-based packet counters and timestamping.

Since OAM is one of the key building blocks of CE 2.0, FlexOAM is an important CE 2.0 enabler.

One Architecture: Three Features, Eight Services

As shown in Table 2, the three major features of CE 2.0 - multi-CoS, manageability and interconnection, are addressed by the Marvell Prestera-DX family, with eBridge and FlexOAM as the key enablers.

| Feature | | Relevant Features in Marvell Prestera-DX Devices |
|----------------|---|--|
| Multi-CoS | ✓ | <ul style="list-style-type: none"> • The Marvell Prestera-DX advanced CoS architecture allows for multi-CoS traffic differentiation. • Shaping, scheduling, metering and tail-dropping mechanisms can be applied to differentiate between classes of service. |
| Managed | ✓ | <ul style="list-style-type: none"> • OAM: Marvell FlexOAM architecture provides flexible support for various contemporary and future OAM protocols. • Ingress/egress traffic management: the Marvell Prestera®-DX device includes both an Ingress Policer engine and an Egress Policer engine, allowing the definition of bandwidth profiles both on the ingress and egress of the EVC. • Traffic management per UNI, per EVC, per CoS: metering can be performed based on various criteria: per-flow, per-ePort, per-physical-port, per-eVLAN, or, can be flexibly configured to be based on any criteria or packet header field, allowing the bandwidth enforcement to be based on any of the three criteria defined in CE 2.0. |
| Interconnected | ✓ | The eBridge architecture and the flexible VLAN manipulation capabilities provide the necessary tools for topologies that span multiple providers, including E-Access. Every OVC in the E-Access network is represented by an ePort in the ENNI. |

Table 2. Supporting the CE 2.0 Features

Table 3 summarizes the Marvell Prestera-DX features that enable the CE 2.0 services.

| Service | | | Relevant Features in Marvell Prestera-DX Devices |
|----------|-------------|---|---|
| E-Line | EPL | ✓ | <ul style="list-style-type: none"> Interface virtualization using the eBridge architecture: ePorts and eVLANs. |
| | EVPL | ✓ | |
| E-LAN | EP-LAN | ✓ | <ul style="list-style-type: none"> Flexible VLAN translation and QinQ. |
| | EVP-LAN | ✓ | |
| E-Tree | EP-Tree | ✓ | <ul style="list-style-type: none"> Flexible tunneling and L2VPN support: MPLS, VPLS, PBB. |
| | EVP-Tree | ✓ | |
| E-Access | Access EPL | ✓ | <ul style="list-style-type: none"> E-Tree support using root/leaf indication [6]. |
| | Access EVPL | ✓ | |

Table 3. Enabling CE 2.0 Services

Conclusion

Carrier Ethernet 2.0 is the second generation of **features** and **services** in carrier networks. The CE 2.0 feature set encompasses both the **data plane** and the **control plane**. Hence, packet processor silicon cannot be CE 2.0 certified as-is; only a full solution that includes both silicon and software can receive the CE 2.0 certification.

Marvell Prestera-DX devices provide a wide set of features that enable a CE 2.0-compliant system, including **eBridge**, a flexible hardware-based interface virtualization architecture, and **FlexOAM**, an enhanced hardware-based OAM solution.

About the Authors

Tal Mizrahi

Feature Definition Architect



Tal Mizrahi is a feature definition architect at Marvell. With 15 years of experience in networking, network security, and ASIC design, Tal has served in various positions in the industry, including system engineer, team leader and, for the past 8 years, an architect for Marvell's networking product line. Tal received his BSc. and MSc. in Electrical Engineering from the Technion, Israel Institute of Technology. Tal is an active participant in the Internet Engineering Task Force (IETF), the Open Networking Foundation (ONF), and serves as the security editor of the IEEE 1588 working group. Tal is a MEF Carrier Ethernet 2.0 Certified Professional (MEF-CECP).

Uri Safrai

Software and Solution Architect

Uri has over 17 years of networking experience. Prior to his current position, Uri worked at Galileo Technology until its acquisition by Marvell in 2001, and since then has held a variety of technological positions at Marvell. At his former role as a switch architect of the Prestera® line of packet processors, Uri was involved with definition and micro-architecture of networking features, protocols, and various ASIC engines and mechanisms, as well as led the definition of the Prestera® eBridge virtualized data-plane architecture. Since 2010, Uri represents Marvell at the Metro Ethernet Forum (MEF), and recently joined the ONF Chipmakers Advisory Board (CAB).

Acronyms

| | |
|-----------------|--|
| BFD | Bidirectional Forwarding Detection |
| CE 2.0 | Carrier Ethernet 2.0 |
| CE | Customer Equipment |
| CoS | Class of Service |
| ENNI | External Network to Network Interface |
| EPL | Ethernet Private Line |
| EVPL | Ethernet Virtual Private Line |
| EP-LAN | Ethernet Private LAN |
| EVP-LAN | Ethernet Virtual Private LAN |
| EP-Tree | Ethernet Private Tree |
| EVP-Tree | Ethernet Virtual Private Tree |
| EVC | Ethernet Virtual Connection |
| MPLS | Multiprotocol Label Switching |
| MPLS-TP | MPLS Transport Profile |
| NNI | Network-to-Network Interface |
| OAM | Operations, Administration and Maintenance |
| OVC | Operator Virtual Connection |
| PBB | Provider Backbone Bridging |
| PE | Provider Edge |
| QoS | Quality of Service |
| SLA | Service Level Agreement |
| UNI | User Network Interface |
| VLAN | Virtual Local Area Network |
| VPLS | Virtual Private LAN Services |
| VSI | Virtual Switching Instance |

References

- [1] Metro Ethernet Forum, "A New Generation of Carrier Ethernet – Overview Presentation", 2012.
http://metroethernetforum.org/Assets/CE2.0/MEF_New_Generation_of_Carrier_Ethernet_-_Overview_2012-01-23.pdf
- [2] Metro Ethernet Forum, "Carrier Ethernet 2.0 Services - Technical Foundation Document", 2012.
http://metroethernetforum.org/Assets/CE2.0/CE_2-0_Technical_Foundation_Document_v2_2012-06-08.pdf
- [3] Metro Ethernet Forum, "Carrier Ethernet 2.0 Certification Blueprint", version 1.1, 2012.
http://metroethernetforum.org/Assets/CE2.0/Carrier_Ethernet_2.0_Certification_Blueprint_-_VERSION_1.1.pdf
- [4] Metro Ethernet Forum, "An Overview of MEF 6.1, 6.1.1, 10.2, 10.2.1 Carrier Ethernet Definitions and Attributes", 2012.
http://metroethernetforum.org/Assets/Technical_Specifications/PPT/Overview_of_MEF6-1_10-2_and_10-2-1_v8-Mar-27-2012.ppt
- [5] Metro Ethernet Forum, "An Overview of the MEF", 2011.
http://metroethernetforum.org/Assets/Presentation/Overview_of_the_Work_of_the_MEF_20130610.pptx
- [6] Jiang, Y., Yong, L., Paul, M., Jounay, F., Balus, F., Henderickx, W., Sajassi, A., "Ethernet-Tree (E-Tree) Support in Virtual Private LAN Service (VPLS)", draft-ietf-l2vpn-vpls-pe-etree-03 (work in progress), IETF, 2014.
<http://tools.ietf.org/html/draft-ietf-l2vpn-vpls-pe-etree>
- [7] T. Mizrahi, N. Sprecher, E. Bellagamba, Y. Weingarten, "An Overview of Operations, Administration, and Maintenance (OAM) Tools", RFC 7276, 2014.
<http://tools.ietf.org/html/rfc7276>
- [8] T. Mizrahi, I. Yerushalmi, "The OAM Jigsaw Puzzle", technical white paper, Marvell, 2011.
http://www.marvell.com/switching/assets/Marvell_OAM_Puzzle_001_white_paper.pdf
- [9] Wikipedia, "Metro Ethernet Forum", 2014.
http://en.wikipedia.org/wiki/Metro_Ethernet_Forum

Marvell Semiconductor, Inc.

5488 Marvell Lane
Santa Clara, CA 95054, USA

Tel: 1.408.222.2500
www.marvell.com

Copyright © 2015. Marvell International Ltd. All rights reserved.
Marvell, the Marvell logo and Prestera are registered trademarks of Marvell or its affiliates. Other names and brands may be claimed as the property of others.