MARVELL

WHITE PAPER

The OAM Jigsaw Puzzle

Bounding the Vast Horizon of Operations, Administration and Maintenance

Tal Mizrahi, Ilan Yerushalmi
Switching Product Definition Architecture
Marvell

April 2011

ABSTRACT

The acronym "OAM" has had different meanings and has been used in different contexts. The most widely used meaning is Operations, Administration, and Maintenance, referring to detection and diagnosis of link failures in a communication network. OAM has existed for a while, dating back to traditional telephony protocols and to TDM-based protocols such as SDH/ATM. As carriers and providers shifted towards Ethernet and packet-based networks, OAM protocols became a necessity for these networks. OAM in packet networks has continuously evolved in the last few years, with several different and competing standards developed by IEEE, ITU-T, and IETF.

This white paper presents a "bird's eye view" of the existing and evolving OAM standards in the context of Ethernet, IP and MPLS, and summarizes the main OAM functions defined in each of these standards. This paper also introduces the Marvell® OAM approach, a generic and flexible strategy for OAM-enabled networking ASICs.

OAM Overview

What is OAM? The purpose of OAM is to monitor availability and quality of customers' services. OAM is a set of tools that detect and report connection failures and performance degradation. These tools can also localize the defect, and in some cases take the necessary action to switch to a redundant connection, thus maintaining service performance.

OAM protocols have been defined for Ethernet networks, IP, MPLS and PWE3. They also are being defined for the MPLS-TP. A snapshot of contemporary OAM protocols is illustrated in Figure 1, new OAM protocols emerging from the major standard organizations.

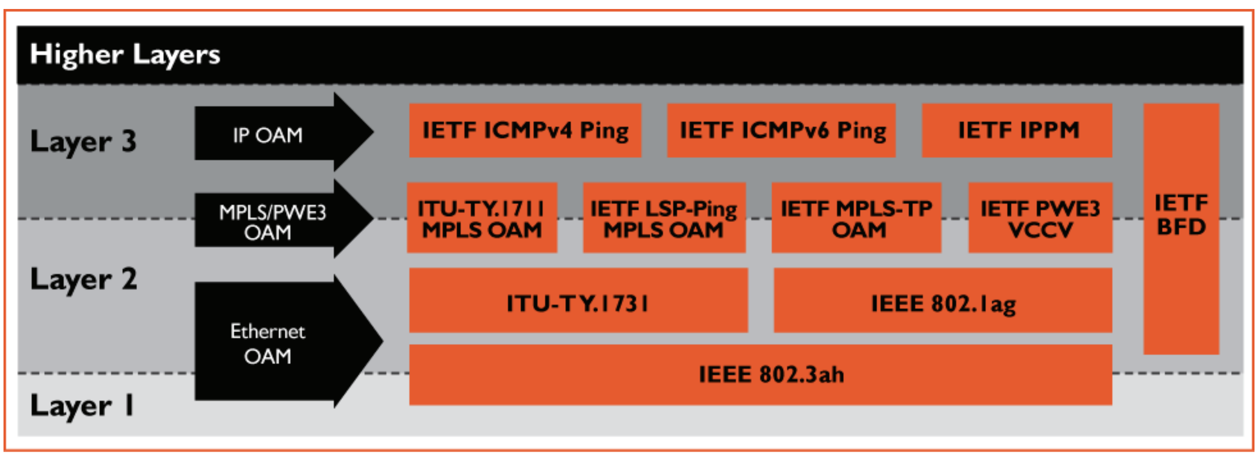

Figure 1. OAM Protocol Stack

What are the main building blocks of OAM? OAM standards generally define a system-level concept of operation and a set of OAM mechanisms. Typical mechanisms that are defined in various OAM protocols are:

- **Keepalive** verification—A keepalive message is sent from one node to another to verify the connectivity between the two nodes. Keepalive mechanisms are used for detecting connection failures. Once detected, the defect can be reported to a network management system. When **protection switching** mechanisms are deployed, once a fault is detected, the system can automatically switch to an alternate path.

  Keepalive mechanisms can be divided into two categories:

  - Periodic messages that are sent proactively at a constant transmission rate; sometimes referred to as **continuity checks.**

- On-demand messages that are sent to verify a specific connection; sometimes referred to as **connectivity verification** mechanisms.
- **Performance measurement** mechanisms can include the following building blocks:
  - Packet loss measurement—Provides a mechanism that computes the packet loss rate between two nodes.
  - Delay measurement—Measures the packet delay and the delay variation between two nodes.
  - Throughput measurement—Measures the traffic throughput between two nodes.
- Path **discovery** and **fault localization**—These mechanisms allow a node to learn the topology of the network, and to localize link failures.

This white paper focuses on the building blocks listed above, sometimes known as forwarding plane OAM. Other aspects associated with the acronym OAM, such as management, fall outside the scope of this white paper.

	Description	Standard / RFC
ITU-T Ethernet OAM	OAM Functions and Mechanisms for Ethernet-based Networks Ethernet Protection Switching Ethernet Ring Protection Switching	ITU-T Y.1730, Y.1731 ITU-T G.8031 ITU-T G.8032
IEEE CFM	Connectivity Fault Management Management of Data-Driven and Data-Dependent Connectivity Faults	IEEE 802.1ag IEEE 802.1Qaw
MEF 17	Service OAM Requirements and Framework	MEF 17
IEEE 802.3 link level OAM	Ethernet on the First Mile (EFM)	IEEE 802.3ah
BFD	Bidirectional Forwarding Detection	RFC 5880, RFC 5881, RFC 5882, RFC 5583, RFC 5884, RFC 5885
ICMP Ping	Internet Control Message Protocol	RFC 792, RFC 4443
IPPM	IP Performance Metrics	RFC 2330, RFC 2678, RFC 2679, RFC 2680, RFC 2681, RFC 4656, RFC 5357.
IETF MPLS OAM (LSP Ping)	Operations and Management (OAM) for Multi-Protocol Label Switched (MPLS) Networks	RFC 4377, RFC 4378, RFC 4379, RFC 4687
MPLS-TP OAM	Requirements for OAM in MPLS MPLS Generic Associated Channel (G-ACh)	RFC 5860 RFC 5586
ITU-T MPLS OAM	Operation and Maintenance Mechanism for MPLS Networks Protection Switching for MPLS Networks Assignment of the 'OAM Alert Label' for MPLS OAM Functions	ITU-T Y.1710, Y.1711 ITU-T Y.1720 RFC 3429
PW VCCV	Pseudowire Virtual Circuit Connectivity Verification (VCCV)	RFC 5085

Table 1. Summary of OAM Standards in Packet Networks

What about Passive Optical Network (PON) OAM? While PON is outside the focus of this paper, it should not go unmentioned. EPON and 10G-EPON provide link-level OAM using the IEEE 802.3ah Ethernet on the First Mile (EFM), while GPON and XG-PON use Physical Layer OAM (PLOAM), defined in the ITU-T PON standards. Higher layers of the OAM protocol stack in PON-based networks take a very similar form to the one illustrated in Figure 1.

OAM Deployment At-a-Glance

How does this OAM "spaghetti" work in a network?

Figure 2 shows a typical multi-layer network OAM deployment, where different OAM protocols are used at different layers of the network.

Each layer in the network uses its own OAM protocol. The customer, the service provider and the operators deploy OAM protocols to monitor the network at their corresponding level of hierarchy. Different OAM layers must coexist, and at times must interoperate.

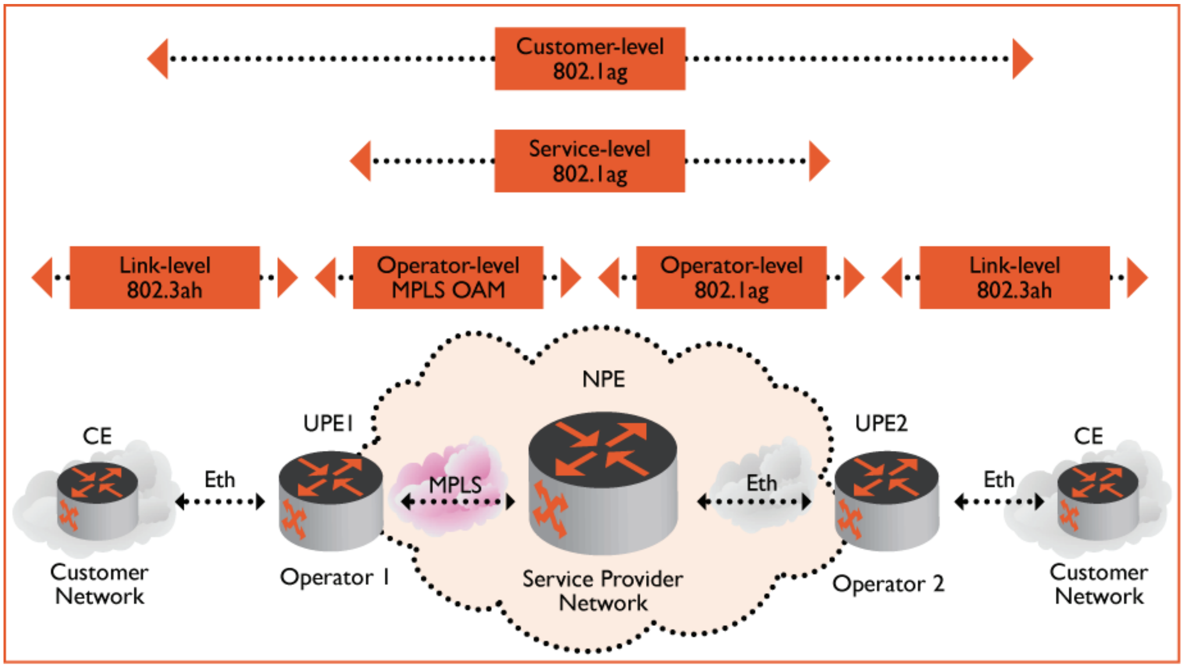

Figure 2. OAM Deployment at a Glance

As shown in Figure 2, MPLS OAM is used to monitor the MPLS network of Operator 1, while IEEE 802.1ag OAM is used to monitor the Ethernet based network of Operator 2. Ethernet OAM (802.1ag) is also deployed at the service provider level and at the customer level, and the link-level 802.3ah OAM is used to monitor the links between the Customer Equipment (CE) and the User-facing Provider Edges (UPE).

The increasing need for multiple coexisting OAM protocols has raised a few challenges for the networking industry:

- Network administrators must be presented with a user interface that transparently enables different OAM capabilities with minimal exposure to the difference between the protocols.
- Networking equipment vendors are seeking generic building blocks that can be reused by different OAM mechanisms, enabling efficient implementations and short time-to-market.

The major standard bodies—IEEE, ITU-T and IETF—have been addressing these challenges, and the wind of OAM has been blowing in the direction of consolidation and reuse. Bidirectional Forwarding Detection (BFD), for example, was conceived by IETF as a generic OAM architecture that can be used over different transport protocols. Another interesting example is the MPLS-TP OAM architecture, which reuses several existing building blocks, including IETF's MPLS OAM and BFD, and ITU-T's Y.1731.

Marvell's FlexOAM and OAM-VM Solutions

Given this highly complex and evolving OAM environment, networking equipment vendors are seeking generic solutions that provide high performance and scalability of OAM processing, while allowing flexibility for various OAM protocols and future changes.

Marvell provides a flexible and comprehensive OAM solution for networking system vendors:

- ASIC-based real-time OAM processing, including identifying, filtering and correct forwarding of OAM messages. Secondary processing is typically performed by the OAM application.
- Flexible parsing of OAM messages, allowing programmable packet formats and encapsulation types.
- Transmission and reception of keepalive messages, to offload the heavy burden of these mechanisms from the application CPU.

The Marvell Prestera®-DX family of packet processors offers a hybrid OAM solution comprised of two complementary building blocks (Figure 3):

- **FlexOAM:** FlexOAM provides real-time processing of OAM traffic. The FlexOAM mechanism provides real-time packet snooping and performs critical real-time processing of OAM traffic. It can also generate and transmit period OAM keepalive messages.
- **OAM-VM:** Marvell's OAM Virtual Machine (OAM-VM) is a software module that has the appearance to other applications as real hardware, while having all the flexibility of software. The OAM-VM can provide complementary OAM functionality, with full flexibility for proprietary extensions, and future compatibility.

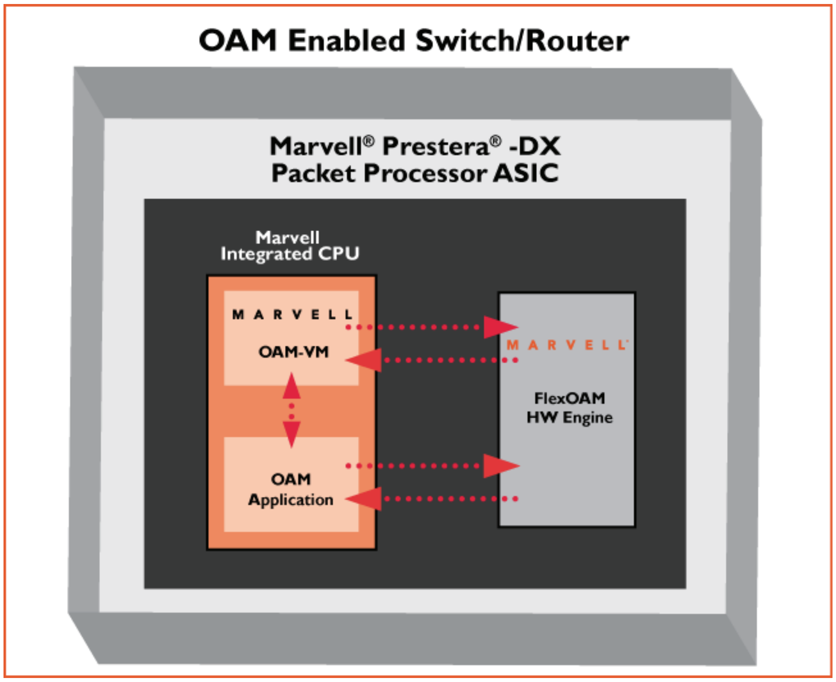

Figure 3. Marvell's OAM Solution

Marvell's integrated OAM solution can be deployed in one of two approaches:

- Fully integrated solution, comprising the FlexOAM hardware engine and an integrated CPU running the OAM-VM software stack, or
- Partially integrated solution, running the OAM-VM software stack in an external application CPU with or without the FlexOAM hardware.

Why is Marvell's integrated OAM solution better than existing ones? The FlexOAM and OAM-VM combine to create a synergic and powerful OAM engine, which is the cutting edge in ASIC-based OAM engines in terms of scalability, real-time performance, flexibility and future compatibility.

	Marvell's FlexOAM + OAM-VM	SW-Based OAM Solutions	HW-Based OAM Solutions
Scalable real-time performance	√		√
Highly accurate loss / delay measurement	√		√
Future compatiblity with upgradeable firmware	√	√	
Full flexibility for future / proprietary OAM protocols	√	√	

Table 2. Marvell's OAM Solution vs. Existing OAM Solutions

FlexOAM

Marvell's FlexOAM is both simple and intuitive. It provides generic hardware support for real-time processing of OAM traffic. The block diagram in Figure 4 includes the following building blocks:

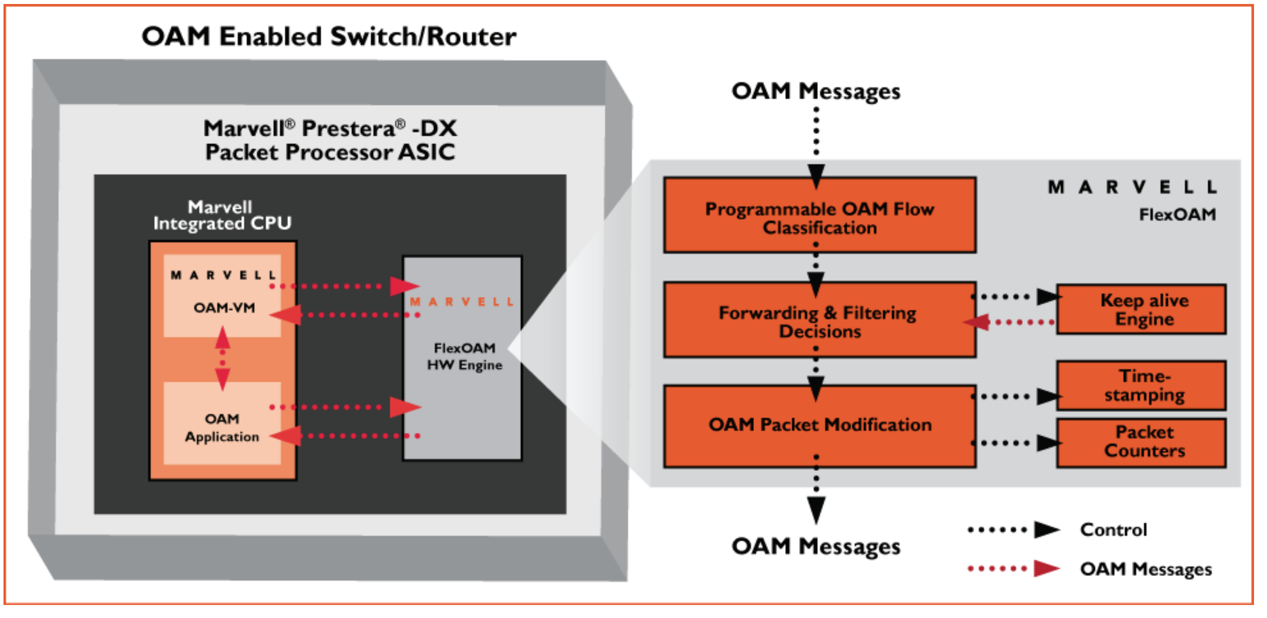

Figure 4. Marvell's FlexOAM Processing

- **Programmable OAM Flow Classification:** All incoming traffic is snooped by this block. OAM messages are subject to flow classification. The flow classification engine is programmable, providing the flexibility to classify flows of any OAM protocols, including tunnel-encapsulated packets and future packet formats.
- **Forwarding and Filtering Decisions:** At this point, FlexOAM takes a forwarding decision. An OAM packet can be forwarded transparently, trapped or mirrored to the application CPU, or it can be dropped. Packets are dropped based on various filtering criteria, such as source/destination addresses, OAM header sanity checks and MD level filtering[1].
- **Keepalive Engine:** This engine is in charge of reception and transmission of keepalive messages. The keepalive engine periodically transmits keepalive messages to each of the peer

nodes it is connected to. It can also detect a **loss of continuity** when no messages were received from a peer node for a predetermined period of time.

- **Timestamping:** OAM delay measurement messages use accurate timestamping of transmission/reception time.
- **Packet counters:** Transmission and reception packet counters are maintained per flow. These counters are used by OAM loss measurement mechanisms to compute the packet loss rate.
- **OAM Application:** The OAM application resides in the application CPU and is typically responsible for the OAM control plane. It transmits and receives OAM messages that do not require excessive CPU utilization and receives notifications from the FlexOAM mechanism in the packet processor.

[1]When multi-layer Ethernet OAM is deployed (see Figure 2), the MD (Maintenance Domain) level specifies the level of hierarchy in the OAM stack. The MD level is defined in IEEE 802.1ag.

OAM-VM

The OAM-VM enables networking equipment vendors to enjoy the best of both worlds: retaining software flexibility, while achieving guaranteed high performance. The OAM Virtual Machine (firmware-based) can be seen by the OAM application as hardware, while also having all the flexibility of software. The OAM-VM co-exists transparently with existing application software, with negligible impact on CPU performance.

The main advantages of the OAM-VM include:

- Allows networking system vendors to tune the OAM solution to the target network and topology.
- The upgradeable firmware-based solution enables better life time support.
- Future standards and enhancements are just a firmware upgrade away.

The OAM-VM was originally developed for systems that use legacy packet processors without FlexOAM support, but it is a very powerful supplement for FlexOAM-enabled packet processors, as well. Thus, the OAM-VM is capable of operating independently of the FlexOAM engine. In this case, the FlexOAM can provide full support of keepalive message transmission and reception, along with other performance sensitive OAM functions such as delay and loss measurement.

The OAM-VM provides complementary functionality for the FlexOAM engine and can provide additional flexibility and future compatibility beyond that of the FlexOAM.

The OAM-VM can be used for:

- Deep packet inspection that is difficult or inefficient to perform in the packet processor.
- Complex or stateful protocols.
- Non-standard or future OAM protocols.

The OAM application and the OAM-VM reciprocate to create a full OAM software package. The OAM-VM typically performs tasks that are performance sensitive, and require guaranteed and deterministic CPU resources, while the OAM application is responsible for the rest of the OAM-related activities.

The OAM-VM runs in fixed time slices called OAM-VM ticks. The value of a time slice varies according to the application needs and is typically equal to the shortest keepalive transmission period. For example, for IEEE 802.1ag Continuity Check Messages (CCMs), the VM tick is typically configured to 3.33 milliseconds. Invocation of the OAM-VM in fixed time slices guarantees deterministic performance and CPU utilization.

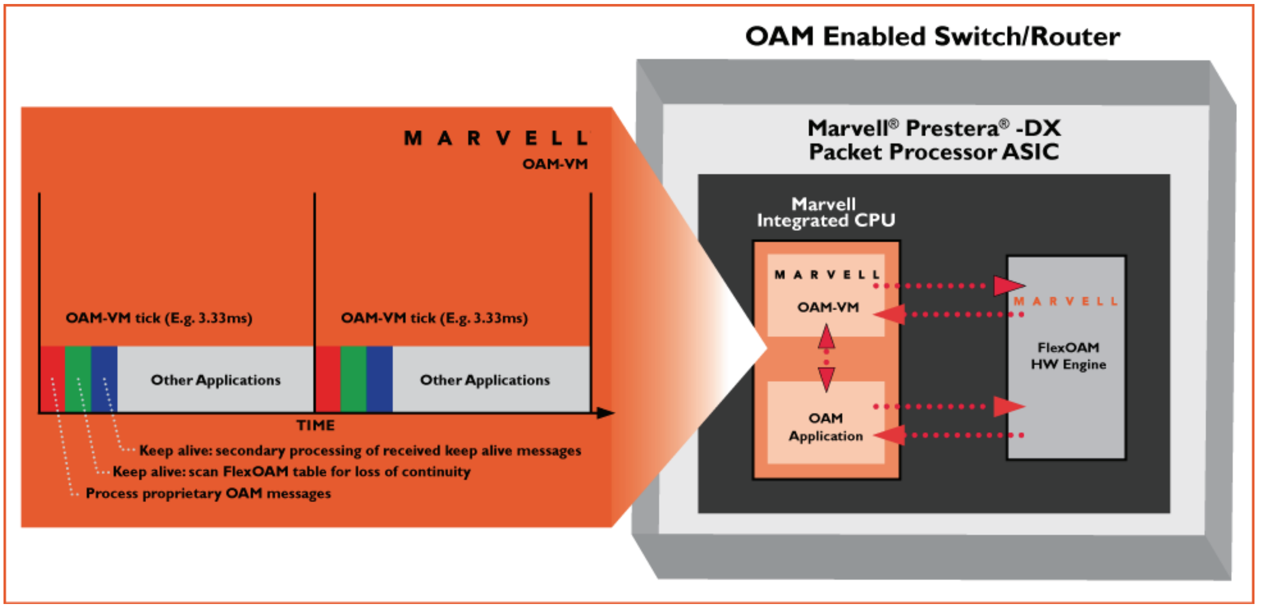

Figure 5. OAM-VM Time Slicing Technique

The OAM-VM can perform various OAM related tasks. Some key examples include:

- **Keepalive message secondary processing:** The FlexOAM engine receives and processes OAM keepalive messages. However, the FlexOAM engine can selectively forward a subset of the keepalive messages for further processing by the OAM-VM.

  Furthermore, the FlexOAM performs sanity checks on received keepalive messages and traps invalid messages to the application CPU. The OAM-VM performs further processing of these packets and diagnoses the anomaly. For example, an IEEE 802.1ag CCM message that is received with an unexpected "transmission period" field is considered a defect condition. The OAM-VM reports these defect conditions to the OAM application.
- **Keepalive loss of continuity:** The OAM-VM periodically scans the FlexOAM engine and detects flows that have encountered a loss of continuity defect. The OAM-VM updates the OAM application when a connection fails, and when it is restored.
- **Processing of proprietary OAM messages:** Some applications use proprietary performance-sensitive OAM protocols, which require unique processing that may not be supported by the FlexOAM engine. These packets can be identified by the FlexOAM and can be further processed by the OAM-VM, enabling both the flexibility and the guaranteed performance.

Conclusion

- OAM has been continuously evolving over the last few years.
- Service providers, operators and networking equipment vendors are facing the following challenges:
  - Achieving fast time-to-market, albeit the dynamically evolving environment
  - Reducing effort and cost by reusing a set of building blocks for different OAM mechanisms
  - Making the network operation, administration and maintenance a friendly and intuitive user experience
- As packet-based networks are rapidly growing, the need for scalable solutions for maintaining service availability are constantly increasing.
- Marvell's FlexOAM meets these challenges using:
  - Hybrid real-time OAM packet generation and processing
  - Flexible OAM packet parsing
  - Scalable performance

About the Authors

Ilan Yerushalmi
Feature Definition Architect

Ilan Yerushalmi is a networking architect with over 15 years of experience in networking systems design. Ilan received his BSc in Electrical Engineering from the Technion, Israel Institute of Technology. He is currently a Feature Definition Architect in the Switching product line at Marvell, and takes part in defining the technologies, features, and architecture of Marvell's current and future networking products.

Tal Mizrahi
Feature Definition Architect

Tal Mizrahi is a Feature Definition Architect at Marvell's networking product line. With over 10 years of experience in networking, networking security, and ASIC design, Tal has served in various positions in the industry, including a system engineer, a team leader and an Architect. Tal received his BSc. and MSc. in Electrical Engineering from the Technion, Israel Institute of Technology.

Acronyms

ATM	Asynchronous Transfer Mode
BFD	Bidirectional Forwarding Detection
CC	Continuity Check
CCM	Continuity Check Message
CE	Customer Equipment
CFM	Connectivity Fault Management.
CV	Connectivity Verification
EPON	Ethernet Passive Optical Network
G-ACh	Generic Associated Channel
GPON	Gigabit Passive Optical Network
ICMP	Internet Control Message Protocol
MD	Maintenance Domain
MPLS	Multiprotocol Label Switching
MPLS-TP	MPLS Transport Profile
NPE	Network-facing Provider Edge
OAM	Operations, Administration, and Maintenance
OAM-VM	OAM Virtual Machine
PE	Provider Edge
PON	Passive Optical Network
PW	Pseudowire
PWE3	Pseudowire Emulation Edge-to-Edge
SDH	Synchronous Digital Hierarchy
TDM	Time Division Multiplexing
UPE	User-facing Provider Edge
VCCV	Virtual Circuit Connectivity Verification
XG-PON	10-Gigabit Passive Optical Network

References

[1] Mizrahi, T., Sprecher, N., Bellagamba, E., Weingarten, Y., "An Overview of Operations, Administration, and Maintenance (OAM) Mechanisms", draft-ietf-opsawg-oam-overview (work in progress), Jan 2011.

[2] Sprecher, N., Bellagamba, E., Weingarten, Y., "MPLS-TP OAM Analysis", draft-ietf-mpls-tp-oam-analysis (work in progress), Jan 2011.

[3] Andersson, L., Van Helvoort, H., Bonica, R., Romascanu, D., Mansfield, S., "Guidelines for the use of the OAM acronym in the IETF", draft-ietf-opsawg-mpls-tp-oam-def (work in progress), Sep 2010.

[4] Postel, J., "Internet Control Message Protocol", STD 5, RFC 792, September 1981.

[5] Conta, A., Deering, S., and M. Gupta, "Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification", RFC 4443, March 2006.

[6] Paxson, V., Almes, G., Mahdavi, J., and Mathis, M., "Framework for IP Performance Metrics", RFC 2330, May 1998.

[7] Mahdavi, J., Paxson, V., "IPPM Metrics for Measuring Connectivity", RFC 2678, September 1999.

[8] Almes, G., Kalidindi, S., Zekauskas, M., "A One-way Delay Metric for IPPM", RFC 2679, September 1999.

[9] Almes, G., Kalidindi, S., Zekauskas, M., "A One-way Packet Loss Metric for IPPM", RFC 2680, September 1999.

[10] Almes, G., Kalidindi, S., Zekauskas, M., "A Round-trip Delay Metric for IPPM", RFC 2681, September 1999.

[11] Shalunov, S., Teitelbaum, B., Karp, A., Boote, J., and Zekauskas, M., "A One-way Active Measurement Protocol (OWAMP)", RFC 4656, September 2006.

[12] Hedayat, K., Krzanowski, R., Morton, A., Yum, K., and Babiarz, J., "A Two-Way Active Measurement Protocol (TWAMP)", RFC 5357, October 2008.

[13] Kompella, K., Swallow, G., "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.

[14] Nadeau, T., Morrow, M., Swallow, G., Allan, D., and Matsushima, S., "Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks", RFC 4377, February 2006.

[15] Allan, D., Nadeau, T., "A Framework for Multi-Protocol Label Switching (MPLS) Operations and Management (OAM)", RFC 4378, February 2006.

[16] Yasukawa, S., Farrel, A., King, D., and Nadeau, T., "Operations and Management (OAM) Requirements for Point-to-Multipoint MPLS Networks", RFC 4687, September 2006.

[17] Nadeau, T., Pignataro, C., "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires", RFC 5085, December 2007.

[18] Katz, D., Ward, D., "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.

[19] Katz, D., Ward, D., "Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)", RFC 5881, June 2010.

[20] Katz, D., Ward, D., "Generic Application of Bidirectional Forwarding Detection (BFD)", RFC 5882, June 2010.

[21] Katz, D., Ward, D., "Bidirectional Forwarding Detection (BFD) for Multihop Paths", RFC 5883, June 2010.

[22] Aggarwal, R., Kompella, K., Nadeau, T., and Swallow, G., "Bidirectional Forwarding Detection (BFD) for MPLS Label Switched Paths (LSPs)", RFC 5884, June 2010.

[23] Nadeau, T., Pignataro, C., "Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)", RFC 5885, June 2010.

[24] Ohta, H., "Assignment of the 'OAM Alert Label' for Multiprotocol Label Switching Architecture (MPLS) Operation and Maintenance (OAM) Functions", RFC 3429, November 2002.

[25] Vigoureux, M., Ward, D., Betts, M., "Requirements for OAM in MPLS Transport Networks", RFC 5860, May 2010.

[26] Bocci, M., Vigoureux, M., Bryant, S., "MPLS Generic Associated Channel", RFC 5586, June 2009.

[27] IEEE 802.1ag, "Connectivity Fault Management", December 2007.

[28] IEEE 802.1Qaw, "Management of Data Driven and Data Dependent Connectivity Faults", July 2009.

[29] IEEE 802.3ah, "Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks", clause 57, September 2004.

[30] ITU-T Y.1730, "Requirements for OAM functions in Ethernet-based networks and Ethernet services", January 2004.

[31] ITU-T Y.1731, "OAM Functions and Mechanisms for Ethernet-based Networks", February 2008.

[32] ITU-T Y.1710, "Requirements for Operation & Maintenance functionality for MPLS networks", November 2002.

[33] ITU-T Y.1711, "Operation & Maintenance mechanism for MPLS networks", February 2004.

[34] ITU-T Y.1720, "Protection switching for MPLS networks", December 2006.

[35] ITU-T G.8031, "Ethernet Protection Switching", June 2006.

[36] ITU-T G.8032, "Ethernet Ring Protection Switching", March 2010.

[37] MEF 17, "Service OAM Requirements & Framework

Marvell Semiconductor, Inc.
5488 Marvell Lane
Santa Clara, CA 95054, USA

Tel: 1.408.222.2500
www.marvell.com

OAM_Jigsaw_Puzzle_whitepaper-001 4/2011