Network working group
Internet Draft
Category: Standard Track

L. Xia
L. Yong
Weiguo Hao
Huawei
Anoop Ghanwani
Dell
Ram Krishnan
Brocade

Expires: April 2015
October 27, 2014

Layer 2 Gateway (L2GW)
draft-xia-nvo3-l2gw-02

Abstract

A Layer 2 Gateway (L2GW) is used for interconnecting a Layer 2 overlay network [NVO3FRWK] and a Layer 2 bridged network [IEEE802.1Q] to form a single Layer 2 virtual network. This draft describes data plane interconnection and control plane interworking at the L2GW.

Status of this Memo

Copyright Notice

Table of Contents

1. Introduction

Cloud computing and network virtualization are evolving in the direction of using network virtualization overlays over Layer 3 (NVO3). Some of the goals of NVO3 are -- fast and easy creation of tenant networks, support tenant system mobility, and improved manageability of all virtualized resources in the data center (DC).

Layer 2 (L2) overlay network in NVO3 means tenant systems are interconnected at L2, while the NVEs are interconnected using Layer 3 (L3). As a result, it forms a full mesh topology of overlay network, i.e. only one L2 hop between any pair of NVEs. On the other hand, L2 bridged network is used to refer to the L2 network as specified in IEEE 802.1Q [IEEE 802.1Q] in this draft.

In the first use case, involving DC network migration from physical tenant systems to virtual tenant systems, it is expected that the L2 overlay network may be used along with an existing L2 bridged network in a DC, and communication between them would be required. In the last use case, a L2 bridged network would be used to connect physical (non-virtualized) systems. These devices need to communicate to virtualized networks for information exchange. Some CPU-intensive applications such as big data analytics typically use physical servers rather than making of use of server virtualization.

To interconnect two networks that are implemented with different technologies (NVO3 and a bridged network), gateway functions are needed on the device(s)/system(s) that interconnect them. This device is referred to as a Layer 2 Gateway (L2GW) in this draft. The device can be thought of as implementing an NVE that connects the tenant systems in the L2 bridged network to tenant systems in the NVO3 network.

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].

1.2. Terminology

This document uses the terms defined in NVO3 framework [NVO3FRWK] and architecture [NVO3ARCH] documents.

2. L2GW Reference Model

The following figure shows a reference model where an L2GW provides an interconnection between an L2 overlay network and an L2 bridged network. It shows the case where two different technologies are used to implement a single L2 network.

```
                .........                   .........
           +---+         ...           ....          . +-------+
       TSs-+NVE|            +---------+              +-+Server|
           +---+ L2 Overlay |         |  L2 Bridge   . +-------+
             .    Network   |  L2GW   |   Network    .
             .              |         |              . +-------+
         ..+---+            +---------+              +-+Server|
       TSs-+NVE|         ...           ....        ... +-------+
           +---+.........                  ........
```

Figure 1: L2GW Reference Model

The L2GW can reside at the edge of the network providing direct connection to tenant systems, or reside at aggregation or core where the tenant systems attach to L2 switches. To connect with an L2 overlay network, an L2GW device physically connects to the underlay network on which the L2 overlay network is implemented and it functions as an NVE, providing termination for the L2 overlay network .

To provide node failure resilience, the reference model can further be shown as in Figure 2, where two L2GWs interconnect the two networks.

```
                .........                   .........
           +---+         ...           ....          . +-------+
       TSs-+NVE|            +---------+              +-+Server|
           +---+ L2 Overlay |  L2GW   |  L2 Bridge   . +-------+
             .    Network   +---------+   Network    .
             .                                       . +-------+
         ..+---+            +---------+              +-+Server|
       TSs-+NVE|         ...|  L2GW   |....        ... +-------+
           +---+.........   +---------+    ........
```

Figure 2: Redundant L2GW Model

Note that this draft assumes that L2GW device embeds an L2 NVE as well as IEEE802.1Q bridge functions.

3. General L2GW Operation Procedures

3.1. MAC Learning

The MAC addresses for an L2 virtual network created by interconnecting the two networks (the L2 overlay network and the L2 bridged network) needs to be distributed and/or learned at all NVEs that participate in that L2 virtual network. If NVE-NVA architecture is used, when an L2GW learns the MAC addresses from the bridged network, the L2GW should notify NVA of the MAC addresses. The NVA maintains the mapping of these MAC addresses from the L2GW, and informs the other NVEs of the mappings.

Similarly, if the NVA maintains the mappings between a tenant system's MAC address and NVE for an L2 virtual network, the NVA would be expected to inform those mappings of MAC addresses to NVEs to the L2GWs because the L2GWs also implement the functions of an NVE. The L2GW maintains the mapping of VNID from the L2 overlay network and VLAN ID in the bridged network. These mappings may be manually configured at the L2GW or may be configured via the NVA.

The L2GW maintains a forwarding table per virtual network which has all the MAC addresses learned from the bridged network as well as all of the MAC addresses it received from the NVA for that virtual network.

Upon receiving a packet from the overlay network, the L2GW decapsulates the packet, performs the table lookup, and may insert a VLAN ID (if the decapsulated frame doesn't already have one) or modify the VLAN ID (if one is already present) prior to forwarding it to the bridged network. If the destination MAC address of the decapsulated packet is unknown (i.e. not present in the forwarding table), the L2GW may choose to discard the packet or flood it on the VLAN depending on the configured policy.

Upon receiving a frame from the L2 bridge network, the L2GW encapsulates the frame prior to forwarding it to the remote NVE. If the frame's MAC DA is unknown to L2GW, it will be discarded or flooded to all the remote NVEs depending on the configured policy. Note that the outer VLAN ID on the packet may be removed before the encapsulation.

The two networks which are interconnected to form a single L2 virtual network MUST NOT have any overlapping MAC addresses; i.e. the same MAC address cannot appear in the both the L2 overlay network as well as the L2 bridged network.

3.2. ARP Handling

To avoid ARP flooding in the L2 overlay network, the L2GW may maintain an ARP cache locally and/or rely on NVA to maintain the ARP table. For the purpose of maintaining the ARP cache locally, the L2GW can snoop ARP requests from the bridged network and send ARP replies back.

If the L2 overlay network supports ARP flooding, the L2GW can simply flood ARP requests from one network to another.

3.3. Dual L2GWs

Two L2GWs may be used for network interconnection to support a network that is resilient to node failures. These two L2GWs may further operate in Active/Standby or Active/Active mode. In Active/Standby mode, only one of the L2GWs is actively passing traffic from one network to the other for a given L2 virtual network. In Active/Active mode, both L2GWs pass traffic from one network to the other for a given virtual network.

(TBD: Does this need to be restricted to only two L2GWs?)

In Active/Standby mode, to protect node failure, some protocol is necessary between the L2GWs to facilitate status exchange and determine which of them will operate in Active mode. The Active/Standby role may be configured or automatically selected based on an algorithm or policy. An L2GW should inform NVA about its role, i.e., Active or Standby, and the NVA should ensure that the active L2GW IP address is used in the mapping of (inner) MAC addresses to (outer) IP address.

In Active/Active mode, NVA/NVEs have two paths to the bridge network and vise versa. The NVEs in an overlay can choose one based on the policy.

The following presents the problems that need to be addressed and related solutions for Active/Active connection scenarios:

1. MAC flip-flop on remote NVEs

MAC learning on an L2GW can be performed either in data plane or control plane. When a local host h1 attaches to multiple L2GWs, address learning at the remote NVEs for a given host h1 may experience what we refer to as the MAC flip-flop problem where h1 appears behind the NVE of one L2GW and then subsequently appears

behind the NVE of the other L2GW, going back and forth in this manner.

In the data plane learning scenario, an anycast L2GW IP address that is shared among L2GWs may be used to avoid MAC flip-flop on remote devices (NVEs, L2GWs, etc). When a bridged network attaches to multiple L2GWs, any L2GW should use the shared anycast IP address, rather than its own IP address, as the ingress NVE IP address when it forwards NVO3 data frames into overlay network. Use of an anycast L2GW IP address makes the MAC addresses learnt by the remote devices appear to be behind a single source IP address rather than multiple different source IP addresses.

In the control plane learning scenario (i.e. when NVA-NVE is used to learn address mappings), if an L2 bridged network is multi-homed to multiple L2GWs in Active/Active mode, each edge L2GW should announce the MAC addresses of its attached end systems to all other devices through NVE-NVA control plane protocol. For MAC addresses that originate from multiple L2GWs, remote devices will learn the MAC addresses as being associated with multiple ingress IP addresses and will generate multiple MAC forwarding entries in ECMP mode. All edge L2GWs should disable the data plane MAC learning function in their NVEs; they must still continue to learn MAC addresses from traffic received from the L2 bridged network. MAC address to NVE IP address association should be learned only through the control plane. The control plane must be aware of edge ports that are multi-homed to multiple L2GWs.

2. Duplicated traffic from remote device

Frame duplication may occur when BUM (broadcast, unknown unicast, multicast) traffics are forwarded bidirectionally between an L2 bridged network and a NVO3 network which have an Active/Active connection through multiple edge L2GWs. The Designated Forwarder (DF) election mechanism defined in [EVPN] can be used to resolve this issue. According to [EVPN], multi-homing functions cover two scenarios. For the MHN (Multi-Homed Network) scenario, DF election mechanism allows only one L2GW of an edge group to forward BUM traffics between NVO3 network and the L2 bridged network by two directions for each VN. The basic idea of DF is to elect one L2GW per VN from an edge group to be responsible for forwarding the BUM traffics. For the MHD (Multi-Homed Device) scenario, the only difference with MHN scenario is at the L2 bridged network side, MC-LAG mechanism guarantees BUM traffics coming from L2 bridged network only goes to one L2GW. DF mechanism is not needed in this direction.

3. Loops

Consider the case where a bridged network originates a frame that is sent as a BUM frame to the NVO3 network via an L2GW, say L2GW1, that is one of multiple gateways interconnecting the bridged network and the NVO3 network. This frame will be encapsulated and then forwarded through NVO3 network and reach the other L2GW, say L2GW2, that is also connected to the bridged network. In this case, if L2GW2 decapsulates the NVO3 frame and forwards it into the bridged network where the frame originated, the frame loops endlessly. This is why it is important to have only single designated forwarder for multicast traffic.

4. Unsynchronized information among member L2GWs

A local L2GW, say L2GW1 in an edge group, may have learned a VLAN and MAC to IP correspondence for a remote end system ES1 when ES1 sends a packet to local bridge. The returning traffic from local bridge may go to any other member L2GW of MC-LAG, for example L2GW2. To avoid flooding unicast traffic on L2GW2, MAC address should be synchronized among the edge L2GWs in an edge group.

Additionally, to ensure DF election consistency, dynamic joined VLAN through VLAN registration protocol (VRP, [IEEE 802.1ak] amendment to the [IEEE 802.1Q]) and dynamic joined multicast group through IGMP or MLD protocol should be synchronized among all L2GWs in an edge group.

4. L2CP Review and Applicability to L2 Overlay Network

This Section mainly discusses which L2CP (Layer 2 Control Protocol, specified in [IEEE 802.1Q]) should be supported by L2 overlay network and which should not, Section 5 specifies how L2GW should deal with L2CP frames.

L2CP protocols defined in [IEEE 802.1Q] are listed in Table 1:

MAC DA	Assignment	Protocol Type	L2CP Action	
			VLAN-based L2 services	PORT-based L2 services
01-80-C2-00-00-00	Nearest Customer Bridge	STP/RSTP/MSTP, LACP/LAMP	Filter	Pass

01-80-C2-00-00-01	IEEE MAC Specific Control Protocols	PAUSE	Filter	Filter
01-80-C2-00-00-02	IEEE 802 Slow Protocols	LACP/LAMP, Link OAM, ESMC	Filter	Filter
01-80-C2-00-00-03	Nearest non-TPRM Bridge	Port Authentication, LACP/LAMP	Filter	Filter
01-80-C2-00-00-04	IEEE MAC Specific Control Protocols		Filter	Filter
01-80-C2-00-00-05 01-80-C2-00-00-06 01-80-C2-00-00-09 01-80-C2-00-00-0A	Reserved for Future Standardization		Filter	Filter
01-80-C2-00-00-07	MEF ELMI	E-LMI	Filter	Filter
01-80-C2-00-00-08	Provide Bridge Group		Filter	Filter
01-80-C2-00-00-0B 01-80-C2-00-00-0C	Reserved for Future Standardization		Filter	Pass
01-80-C2-00-00-0D	Provider Bridge MVRP		Filter	Pass

01-80-C2-00-00-0E	Nearest Bridge, Individual LAN Scope	LLDP, PTP Peer Delay	Filter	Filter
01-80-C2-00-00-20 through 01-80-C2-00-00-2F		GARP/MRP Block	Pass	Pass

Table 1 L2CP protocols specification

Note:

Different L2CP protocols can use the same MAC DA in above block of 32 addresses, but be differentiated by protocol identifier. MAC DA determines the intended recipient device for the frame;

Filter represent the L2CP action of peer or discard;

Based on whether L2 interface is VLAN-aware, L2 services can divided into two categories: VLAN-based L2 services, PORT-based L2 services. L2CP action (peer, discard, pass) for these two L2 services is also different;

Whether the L2CP frames are peered or discarded is further determined by the configuration of L2 interface.

Further analysis about whether a L2CP protocol is necessary and how it is processed in NVO3 supported L2 VN, is provided in the following sub sections.

4.1. STP/RSTP/MSTP

The Spanning Tree Protocol (STP) is a L2 protocol that ensures a loop-free topology for any bridged Ethernet local area network. The basic function of STP is to prevent bridge loops and the broadcast storm that results from them. Rapid spanning Tree Protocol (RSTP) and Multiple Spanning Tree Protocol (MSTP) are all the enhanced xSTP protocols.

L2 overlay network does not need xSTP protocols to prevent bridge loops because it has its own mechanism for it, i.e., NVA, control plane mechanisms, full mesh + split horizon, etc. So, the process of xSTP frames in L2 VN is:

Be in line with L2CP protocols' specification of Table 1 from IEEE in the L2 sub-networks attached to L2 NVEs;

xSTP frames are filtered in L2 NVEs and should not go into L2 overlay network.

4.2. PAUSE

[IEEE 802.3-2005] has specified a L2 flow control mechanism through using the PAUSE frame. This frame uses L2CP MAC DA of 01-80-C2-00-00-01 to be sent to the node at the other end of the link for informing it to halt the frame transmission for a specified period of time.

When L2 NVE is co-located in Hypervisor, PAUSE frame is not necessary in one device. When they are separated, PAUSE frame is only used in layer 2 network between L2 NVE and Hypervisor, there is no need to overlay PAUSE frame between L2 NVEs. For the underlay network of NVO3 network, L2 PAUSE mechanism is still used between two adjacent switches for flow control.

4.3. LACP/LAMP

Link Aggregation [IEEE 802.1AXbk-2012] is a mechanism for making multiple point-to-point links between a pair of devices appear to be a single logical link between those devices. Link Aggregation Control Protocol (LACP) and Link Marker Control Protocol (LAMP) operate between exactly two peer devices for the purpose of creating, verifying, and monitoring the logical link created by aggregating individual links. Specific L2CP frames, known as Link Aggregation Control Protocol Data Units (LACPDUs), are exchanged between the peer devices on each individual link in the aggregation. The protocol identifier used by LACP is an Ethertype with a value of 0x8809 (the ''Slow Protocols'' Ethertype) and subtype values 01 (for LACP) and 02 (for LAMP). Note that LACP is used to represent LACP and LAMP in the following text.

LACP uses 3 different L2CP MAC DAs to determine the scope of propagation of LACPDUs within a bridged LAN, as Table 2 follows:

Assignment	L2CP MAC DA	Peered or discarded by

Nearest Customer Bridge	01-80-C2-00-00-00	End Station, Customer Bridge, Provider Edge Bridge
IEEE 802 Slow Protocols	01-80-C2-00-00-02	End Station, Customer Bridge, Provider Edge Bridge, Provider Bridge
Nearest non-TPRM Bridge	01-80-C2-00-00-03	Bridges except for Two Port MAC Relay

Table 2 LACP specification of L2CP MAC DAs

Base on the summary of Table 2, LACPDUs with the L2CP MAC DA of 01-80-C2-00-00-02 are peered or discarded by every node, so this kind of LACPDUs will not be overlaid across the L2 overlay network. For 01-80-C2-00-00-00, it is possible that LACPDUs need to be overlaid across Provider Bridge and L2 NVEs of L2 overlay network to reach the other end Custom Bridge, L2 overlay network maybe need to support to overlay this kind of LACP frame between L2 NVEs. How the L2 overlay network support LACP frame of 01-80-C2-00-00-03 is TBD.

4.4. Link OAM

Lin OAM defined is defined in [IEEE 802.3ah], as mechanisms for monitoring and troubleshooting Ethernet access links. Specifically it defines tools for discovery, remote failure indication, remote and local loopbacks and status and performance monitoring.

The Link OAM frames using L2CP MAC DA of 01-80-C2-00-00-02 are peered or discarded by every node, so this kind of frame will not be overlaid across the L2 overlay network.

4.5. Port Authentication

[IEEE 802.1X] is an IEEE Standard for Port-based Network Access Control (PNAC). It is part of the IEEE 802.1 group of networking protocols. It provides an authentication mechanism to devices wishing to attach to a LAN or WLAN.

Whether or not the L2 overlay network needs to overlay this L2CP frames is TBD.

4.6. E-LMI

Ethernet Local Management Interface (E-LMI) [MEF-16] is a protocol between the customer edge (CE) device and the provider edge (PE) device. It runs only on the PE-CE UNI link and notifies the CE of connectivity status and configuration parameters of Ethernet services available on the CE port. E-LMI interoperates with an OAM protocol, such as Connectivity Fault Management (CFM), that runs within the provider network to collect OAM status. CFM runs at the provider maintenance level (UPE to UPE with inward-facing MEPs at the UNI). E-LMI relies on the OAM Ethernet Infrastructure (EI) to interwork with CFM for end-to-end status of Ethernet virtual connections (EVCs) across CFM domains.

The LLDP frames using L2CP MAC DA of 01-80-C2-00-00-07 are peered or discarded by every node except for the Two Port MAC Relay (TPMR) bridge, so this kind of frame will not be overlaid across the L2 overlay network.

4.7. LLDP

The Link Layer Discovery Protocol (LLDP) is a vendor-neutral link layer protocol in the Internet Protocol Suite used by network devices for advertising their identity, capabilities, and neighbors on an IEEE 802 local area network, principally wired Ethernet. The protocol is formally referred to by the IEEE as Station and Media Access Control Connectivity Discovery specified in standards document [IEEE 802.1AB].

The LLDP frames using L2CP MAC DA of 01-80-C2-00-00-0E are peered or discarded by every node, so this kind of frame will not be overlaid across the L2 overlay network.

4.8. PTP Peer Delay

PTP Peer Delay frame is specified in [IEEE 1588-2008] to carry PTP peer time information. It uses L2CP MAC DA of 01-80-C2-00-00-0E and peered or discarded by every node, so this kind of frame will not be overlaid across the L2 overlay network.

4.9. ESMC

Ethernet Synchronization Messaging Channel (ESMC) is specified in [ITU-T Rec. G.8264] for conveying clock information between Synchronous Ethernet (SyncE) bridges.

The ESMC frames using L2CP MAC DA of 01-80-C2-00-00-02 are peered or discarded by every node, so this kind of frame will not be overlaid across the L2 overlay network.

4.10. GARP/MRP Block

Multiple Registration Protocol (MRP), which replaced Generic Attribute Registration Protocol (GARP), is a generic registration framework defined by the [IEEE 802.1ak] amendment to the [IEEE 802.1Q] standard. MRP allows bridges, switches or other similar devices to be able to register and de-register attribute values, such as VLAN identifiers and multicast group membership across a large LAN. MRP operates at the Data Link Layer.

The block of L2CP MAC DA from 01-80-C2-00-00-20 to 01-80-C2-00-00-2F is used for MRP protocol. Now, only 01-80-C2-00-00-20 is for Multiple MAC Registration Protocol (MMRP) and 01-80-C2-00-00-21 is for Multiple VLAN Registration Protocol (MVRP), other L2CP MAC DA of the block are all reserved for future use. Protocol using one address of this block is passed by all the intervening bridges that does not participate in the protocol using this address, and peered or discarded by the bridge that participate in the protocol at last. In order to send the MRP frames to all related nodes (i.e., NVEs, bridges, etc) in one L2 overlay network, the MRP frames may require to be overlaid across the L2 overlay network.

5. L2CP Processing in L2GWs

For all L2CP protocols, several differences exist between L2 overlay network and L2 bridge network on how to process them. As the demarcation point between L2 overlay network and L2 bridge network, L2GW keeps the same action to all L2CP frames as before at the L2 bridge network side on the one hand, but maybe processes some L2CP frames differently at the L2 overlay network side on the other hand. The following sub sections will describe the L2CP process in L2GW.

5.1. L2CP Frames Filtered (Peered or Discarded) in L2GW

Although xSTP protocols using Nearest Customer Bridge address of 01-80-C2-00-00-00 indicate that it can be overlaid across L2 overlay network, they still are not necessary for L2 overlay network because L2 overlay network has its own mechanism to prevent bridge loops. So xSTP frames will be filtered by the L2GW and not go into the L2 overlay network.

Based on the analysis of section 3.3, LACP/LAMP frames using IEEE 802 Slow Protocols of 01-80-C2-00-00-02 are not necessary for L2

overlay network. So, LACP/LAMP frames will be filtered by the L2GW and not go into the L2 overlay network. ESMC frames using the same MAC DA will also be filtered by L2GW.

For Link OAM frames, if OAM functions are necessary for the whole L2 network which interconnects L2 bridge network and L2 overlay network, L2GW needs to support the interworking of OAM as well. This means that L2GW should peer the Link OAM frames of L2 bridge network and perform some actions between NVEs in L2 overlay network. The detailed operation is TBD.

Other L2CP protocols that are filtered by L2GW and do not go into L2 overlay network include PAUSE, E-LMI, LLDP, PTP Peer Delay. The basic reason is that they all require to be processed hop by hop in L2 network strictly, but overlay network breaks this rule.

The action of ''filter'' can be ''peer'', or ''discard''. It depends on the specific service requirement, i.e., does L2GW need to participate in the L2CP protocol, etc. How to determine the specific action is TBD.

5.2. L2CP Frames Passed through L2GW

Excepting for the aforementioned L2CP protocols filtered by L2GW, the left L2CP protocols need to be passed through L2GW. They include:

LACP/LAMP frames using IEEE 802 Slow Protocols of 01-80-C2-00-00-00;

GARP/MRP series protocols (i.e., MMRP, MVRP) using the MAC DA block of 01-80-C2-00-00-20 through 01-80-C2-00-00-2F.

All these kinds of L2CP frames are passed through L2GW and traverse across the L2 overlay network and L2 bridge network to arrive the bridges that participate in the L2CP protocols. For MRP protocols, another necessary operation of L2GW is to use the pre-provisioned VLAN to virtual network instance (VNI) mappings in NVE locally or by getting from NVA to map these MRP frames into corresponding VNIs.

6. Other Interworking Cases

There are other L2 bridge network technologies that use L2 Control Plane protocols such as Provider Bridge [IEEE802.1AD] or Provider Backbone Bridge [PBB] [IEEE802.1AH]. The use case of L2 Overlay

Network interworking with these types of bridge networks is for the further study.

Note that VPLS [RFC4761] [RFC4762], EVPN [EVPN], Shortest Path Bridging [IEEE SPB] and TRILL [RFC6325] are also technologies for L2 private network implementation. These technologies rely on the control plane protocol and aim for service provider network. SDN controller interworking with such control plane protocol will be addressed in separate draft.

7. Security Considerations

TBD.

8. IANA Considerations

The document does not require any IANA action.

9. References

9.1. Normative References

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

[RFC4761] Kompella, K. and Rekhter, Y. (Editors), "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling", RFC 4761, January 2007

[RFC4762] Lasserre, M. and Kompella, V. (Editors), "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling", RFC 4762, January 2007.

[RFC6325] Perlman, R., "RBridges: Base Protocol Specification", July 2011.

9.2. Informative References

[NVO3ARCH] Black, D, Narten, T., et al, "An Architecture for Overlay Networks (NVO3)", draft-narten-nvo3-arch-01, work in progress

[NVO3FRWK] LASSERRE, M., Motin, T., et al, "Framework for DC Network Virtualization", draft-ietf-nvo3-framework-07, work in progress.

[NVGRE] Sridharan, M., et al, "NVGRE: Network Virtualization using Generic Routing Encapsulation", draft-sridharan-virtualization-nvgre-03, work in progress

[VXLAN] Mahalingam, M., Dutt, D., etc, "VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks", draft-mahalingam-dutt-dcops-vxlan-05.txt, work in progress

[EVPN] Sajassi, A. and R. Aggarwal, "BGP MPLS Based Ethernet VPN", draft-ietf-l2vpn-evpn-07, May 2014

[EVPN-REQ] A. Sajassi, R. Aggarwal et. al., "Requirements for Ethernet VPN", RFC7209

[EVPN-MHN] Weiguo, Hao, Yizhou, Li, et al, "Multi-homed network in EVPN", draft-hao-l2vpn-evpn-mhn-00, work in progress

[802.1Q] IEEE, "Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q-2011, August, 2011.

[IEEE 802.3-2005] "Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications"

[IEEE 802.1AXbk-2012] "IEEE Standard for Local and metropolitan area networks--Link Aggregation Amendment 1: Protocol Addressing"

[IEEE 802.3ah] "IEEE Standard for Information technology--Local and metropolitan area networks--Part 3: CSMA/CD Access Method and Physical Layer Specifications Amendment: Media Access Control Parameters, Physical Layers, and Management Parameters for Subscriber Access Networks"

[IEEE 802.1X] "IEEE Standard for Local and metropolitan Area Networks. Port-based Network Access Control"

[IEEE 802.1AB] "IEEE Standard for Station and Media Access Control, Connectivity Discovery"

[MEF-16] Metro Ethernet Forum, MEF 16, Ethernet Local Management Interface (E-LMI), January 2006.

[IEEE 1588-2008] "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems"

[IEEE 802.1ak] "IEEE Standard for Local and metropolitan Area Networks - Virtual Bridged Local Area Networks, Amendment 7: Multiple Registration Protocol"

[IEEE 802.1AD], "Virtual Bridged Local Area Networks - Amendment 4: Provider Bridges", 2005

[PBB] Clauses 25 and 26 of "IEEE Standard for Local and metropolitan area networks - Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks", IEEE Std 802.1Q, 2013.

[IEEE802.1AH] IEEE Draft P802.1ah/D4.2 "Virtual Bridged Local Area Networks, Amendment 6: Provider Backbone Bridges", 2008

[IEEE SPB] "IEEE standard for local and metropolitan area networks: Media access control (MAC) bridges and virtual bridged local area networks -- Amendment 20: Shortest path bridging", IEEE 802.1aq, June 2012.

[ITU-T Rec. G.8264] "Distribution of Timing Through Packet Networks"

Authors' Addresses

Liang Xia (Frank)
Huawei Technologies

Email: frank.xialiang@huawei.com

Lucy Yong
Huawei Technologies, USA

Email: lucy.yong@huawei.com

Weiguo Hao
Huawei Technologies
101 Software Avenue,
Nanjing 210012
China

Phone: +86-25-56623144
EMail: haoweiguo@huawei.com

Anoop Ghanwani
Dell

Email: anoop@alumni.duke.edu

Ram (Ramki) Krishnan
Brocade

Email: ramk@brocade.com