Duplicate MAC Address Detection

Abstract

This document describes a duplicate hardware address detection protocol called DMAD. Hosts can use DMAD to determine if a group of link layer addresses, such as a MAC addresses, associated with a given Host are duplicate. DMAD allows a host to reserve a group of link layer addresses not used by any other host for future usage and it allows a host to release a group of reserved link layer addresses requested by any other host. DMAD achieves this goal for a host by maintaining a database of a group of link layer addresses used/reserved by all hosts in a broadcast domain.

Table Of Contents

1. Introduction

To communicate at the data link layer, each network interface controller is assigned with an Unique MAC (Media access control) address. These MAC addresses are of two types. One type is Burned and in MAC addresses and other type is configured MAC addresses. Burned in MAC addresses assigned during manufacture of network devices are typically unique. However there is a possibility of those burned in MAC addresses colliding among themshelves if not assigned properly by manufacturer. As configured MAC addresses are generated randomly, so there is a higher possibility of those configured MAC addresses colliding among themshelves.

This document describes a protocol to detect if a group of link layer addresses, such as a MAC address, associated with a given VM (Virtual Machine) Host and It's spawned VMs' are duplicate. The protocol is called DMAD (Duplicate MAC Address Detection). In DMAD, a Network device or Physical Host broadcasts a DMAD request message on a network. The DMAD request message contains the range of MAC addresses starting from a target MAC address. When other Physical Host or Network devices in same broadcast domain receive DMAD request message, each of the receiving Physical Host or Network device checks if at least a subset of the range of MAC addresses starting from the target hardware address is being used by the receiving Physical Host or the VMs spawned by the receiving Physical Host.

If the receiving Physical Host or Network device finds that at least a subset of the range of MAC addresses starting from the target hardware address is being used by it or VMs spawned by it, receiving Physical Host or Network device generates DMAD response message. Upon receiving DMAD response messages from other Physical Host or Network devices in same broadcast domain, the sending Network device or Physical Host determines whether at least a subset of the range of MAC addresses starting from the target hardware address is not available for use. Upon receiving DMAD response messages, the sending Physical Host or Network device also determines that at least a subset of the range of MAC addresses starting from the target hardware address are released by the host sending DAMD response message. The network device or sending Physical Host then reserves the released MAC addresses for usage by VMs executing on the Physical Host or Network device.

Since DMAD protocol helps determine duplicate MAC addresses from a group of MAC addresses for Physical Host or Hypervisor using single DMAD Request message, it avoids unnecessary packet transmissions for such activities. Determining whether MAC address is duplicate each time a new VM is spun-up can be time intensive and negatively impact the amount of time the VMs need to wait before being able to transmit and receive communication.

1.1. Terminology

This document uses the following terms:

o MAC Address: Link layer address assigned to the Medium access control (MAC) sublayer. MAC sublayer and the logical link control (LLC) sublayer together make up the data link layer.

o DMAD: Duplicate MAC Address Detection.

o NIC: A network interface card (NIC) is a circuit board or card that is installed in a computer so that it can be connected to a network.

o VNIC: A VNIC is a virtualized Network Interface Card,based on a physical NIC, used by a Virtual Machine as its network interface. A VNIC is assigned a MAC address.

o VM: A virtual machine (VM) is a software program or operating system that is an emulation of a computer system and is capable of performing tasks such as running applications and programs like a separate computer

o VM Guest: A guest virtual machine (guest VM) is the software component of a virtual machine (VM). The guest VM and the host VM are the two components that make up a virtual machine.

o VM Host: The host VM is the underlying hardware that provides computing resources -- such as processing power, memory, disk and network I/O (input/output).

o Virtualization Platform: Platform virtualization are emulators or hypervisors that emulate the whole physical computer machine, often providing multiple virtual machines on one physical platform.

o Centralised Entity: The entity guarantees that virtual machines are assigned unique MAC addresses within a given host system, responsible for managing and maintaining MAC address for VMs for a single host.

2. DMAD Request Message

The DMAD Request message is encapsulated in an Ethernet header. Unique ethertype value has been assigned for DMAD protocol.
An unique hardware type value has been assigned for DMAD protocol for ethernet. The DMAD Request packet use the fields.

Figure 1 depicts the DMAD Request Message.

```
 0                   1                   2
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          Hardware type                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|  Hardware address  |  Device ID           |
|  length            |   length           |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|        Operation Code                   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|   Sender hardware address (first 2 octets) |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          (next 2 octets)                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          (last 2 octets)                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|      Device Identifier  (first 2 octets)   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          (last 2 octets)                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|   Target hardware address (first 2 octets) |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          (next 2 octets)                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          (last 2 octets)                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|   Range for Target hardware addresses      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          Reserved                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
```

Figure 1: DMAD Request message.

Ethernet Header fields:

o Destination Ethernet Address : The Destination Ethernet Address identifies all hosts in the broadcast domain.It MUST be a valid MAC broadcast address.

o Source Ethernet Address: The Source Ethernet Address identifies the host sending DMAD Request message. It MUST be a MAC unicast address.

o Frame Type: Specifies the type of data that follows. An unique ethertype value is assigned.

DMAD Request Message fields:

o Hardware type : Type of Layer 2. The value for ethernet protocol is "DMAD-over-Ethernet". The value for ATM can be "DMAD-over-ATM". This field decides whether hardware address resolution is done for ethernet hardware address, ATM hardware or any other L2 hardware address.

o Hardware length: Hardware address's length in octets. The value is 6.

o Device ID length: Length (in octets) of device id used.

o Operation Code : For DMAD Request message. The value is "MAC-Verify-Request".

o Sender hardware address: Hardware address of the host sending DMAD request.

o Device identifier : Specifies a device identifier of the network device sending DMAD request message. The device identifier is an opaque string that users or devices can use to identify the network device sending the message. For example, the device identifier specifies a string such as ìRACK_01_BLADE_12î.

- Target hardware address : The target hardware address field specifies a hardware address that Host plans to reserve for future usage by it's VMs. Host sends DMAD request message to verify if the hardware address is duplicate or not. When Host wants to verify if it's NIC's hardware address is duplicate or not then Host sends a DMAD request message specifying it's NIC hardware address in target hardware address field.

- Range of Target hardware addresses: Range of hardware addresses is a set of addresses starting from Target hardware address. Host wants to verify if these addresses are duplicate or not using a single DMAD request message.

2.1 Target hardware address Field and Range for Target hardware addresses Field

Each network device in a broadcast domain maintains a mapping of device identifier of a network device to hardware address reserved by the network device. Each network device maintains database of above mentioned mapping of all Reserved hardware addresses reserved by all other network devices. So this database is called "non-available hardware addresses" database as these hardware addresses are not available for the network device that has stored this database.

When a network device is determining a Target hardware address to assign to a VM, the network device checks the above database to ensure that either the Target hardware address or any other hardware address out of a Range of hardware addresses starting from Target hardware address are not already assigned to another network device connected to the broadcast domain.

If the network device finds that all of the hardware addresses out of a Range of hardware addresses starting from Target hardware address are already assigned to other network devices connected to the broadcast domain then the networking device decides to reserve another Range of hardware addresses starting from another Target hardware address. So the network device prepares DMAD request message with the new Target hardware address value, a Range for the new Target hardware addresses value and sends DMAD request message to all hosts in the broadcast domain to figure out if the Range of hardware addresses are available for reservation or not.

3. DMAD Response message

The DMAD Response message is encapsulated in an Ethernet header.
Unique ethertype value has been assigned for DMAD protocol.
An unique hardware type value has been assigned for DMAD protocol
for ethernet. The DMAD Response packet use the fields.

Figure 2 depicts the DMAD Response message.

```
 0                   1                   2
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          Hardware type                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|  Hardware address  |  Device ID          |
|  length            |   length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|           Operation Code                 |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
| Reserved hardware address (first 2 octets) |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          (next 2 octets)                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          (last 2 octets)               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|      Device Identifier  (first 2 octets)   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          (last 2 octets)               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
| Released hardware address (first 2 octets) |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          (next 2 octets)                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|          (last 2 octets)               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|   Range for Reserved hardware addresses    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|   Range for Released hardware addresses    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
```

Figure 2: DMAD Response message.

Ethernet Header fields:

o Destination Ethernet Address : The Destination Ethernet Address identifies all hosts in the broadcast domain.It MUST be a valid MAC broadcast address.

o Source Ethernet Address: The Source Ethernet Address identifies the host sending DMAD Response message. It MUST be a MAC unicast address.

o Frame Type: Specifies the type of data that follows. An unique ethertype value is assigned.

DMAD Response Message fields:

o Hardware type : Type of Layer 2 hardware. The value for ethernet protocol is ìDMAD-over-Ethernet". The value for ATM can be "DMAD-over-ATM". This field decides whether hardware address resolution is done for ethernet hardware address, ATM hardware address or any other L2 hardware address.

o Hardware length: Hardware address's length in octets. The value is 6.

o Device ID length: Length (in octets) of Device id.

o Operation Code: For DMAD Response message. The value is "MAC-Verify-Response".

o Reserved hardware address: The Reserved hardware address field specifies a hardware address that Host has already reserved. Upon receiving DMAD request message, Host sends DMAD response message and informs other hosts that such addresses are not available for use.

o Device identifier : Specifies a device identifier of the network device sending DMAD response message. The device identifier is an opaque string that users or devices can use to identify the network device sending the message. For example, the device identifier specifies a string such as ìRACK_01_BLADE_12î.

o Released hardware address : The Released hardware address field specifies a hardware address that Host plans to release. Upon receiving DMAD response message, Host that has sent DMAD request message can reserve such released hardware address for future usage by it's own VM.

- o Range of Reserved hardware addresses: Range of Reserved hardware addresses is a set of addresses starting from above mentioned Reserved hardware address. Host uses this field to inform other hosts that a set of hardware addresses are not available for use and Host does this with help of single DMAD response message.

- o Range of Released hardware addresses: Range of released hardware addresses is a set of addresses starting from above mentioned Released hardware address. Host uses this field to informs other hosts that a set of hardware addresses are available for use and Host does this with help of single DMAD response message.

3.1 Reserved hardware address field and Range for Reserved hardware addresses field

Each network device in broadcast domain maintains a mapping of device identifiers of a network device to Reserved hardware addresses reserved by the network device. Since each network device maintains database of all such mapping of Reserved hardware addresses reserved by the network device, this database is called "available hardware addresses" database for the network that stores this database.

When a network device has already reserved a Range of hardware addresses and receives a DMAD request with a same range of hardware addresses then the network device informs the sending host and all other hosts in the broadcast domain about it's own set of Reserved hardware addresses. To inform other hosts, that these hardware addresses are not available, the network device prepares a DMAD Response message with Reserved hardware address value and Range for Reserved hardware addresses value and send the message to all hosts in the broadcast domain.

3.2 Released hardware address field and Range for Released hardware addresses field

When a network device has already reserved a Range of hardware addresses and receives a DMAD request with the same range of hardware addresses then the network device can choose to release some of itís reserved hardware addresses. The network device informs the sending host and other hosts in the broadcast domain about these released hardware addresses through a DMAD Response message. To inform sending host and other hosts, that these released hardware addresses are available, the network device prepares a DMAD Response message with Released hardware address value, Range for Released hardware addresses value and send it to all hosts in Broadcast domain.

When a network device decides to release a group of hardware addresses reserved by it, the network device removes those group of hardware addresses from it's "available hardware addresses" database and adds those group of hardware addresses to it's "non-available hardware addresses" database.

4. DMAD Request Message Processing

When a node receives a DMAD Request message and any of the following conditions apply, the node MUST silently discard the incoming message and is NOT REQUIRED to return any message to the sender of the unsupported message.

o The node does not recognize DMAD Request messages.

o The node has not explicitly enabled DMAD functionality.

o The node is a VM Guest. VM Guest can not process DMAD Request packet unless explicitly authorized to handle incoming DMAD Request packet for the VM Guest.

o The node is a VM Guest and incoming DMAD Request packet has Range for Target hardware addresses field value more than one. VM Guest can not process DMAD Request packet unless explicitly authorized to handle incoming DMAD Request packet on behalf of other VM Guests or VM Host.

o The node is not a VM Host or any other non-VM node and not authorised to handle incoming DMAD Request packet with Range for Target hardware addresses field value more than one.

o The Source Hardware Address of the incoming message is not a unicast address.

o The Destination Hardware Address of the incoming message is a multicast address.

o The value for Operation field is not "MAC-Verify-Request".

Otherwise, when a node receives a DMAD Request packet, it MUST format a DMAD Response packet as follows:

- o Set the Source Hardware Address to the Hardware address the node.

- o Set the Destination Hardware Address to Broadcast Hardware address.

- o Set EtherType as assigned by IEEE.

The responding node MUST do the following:

- o Set Hardware type to ìDMAD-over-Ethernet".

- o Set Hardware length to 6.

- o Set Device ID length to 4.

- o Set Operation to "MAC-Verify-Response".

- o Set Reserved hardware address field as described in section 5.1.

- o Set Device identifier field to an opaque string that users or devices can use to identify the network device sending the message. For example, ìRACK_01_BLADE_12î.

- o Set Released hardware address field as described in section 5.2.

- o Set Range for Reserved hardware addresses as described in section 5.1

- o Set Range for Released hardware addresses field as described in section 5.2

- o Forward the DMAD Response packet to its destination.

4.1 Target hardware address Field and Range for Target hardware addresses Field Processing

- o Upon receiving a DMAD Request message, a network device retrieves the Target hardware address value and the Range for Target hardware addresses value from DMAD Request message.

o The network device checks if any of these Range of hardware addresses starting with Target hardware address is present in network device's "available hardware addresses" database.

o If none of the Range of hardware addresses starting with Target hardware address matches with any of the addresses from the network device's "available hardware addresses" database then network device ignores the DMAD Request message.

o In response to determining that at least one hardware address assigned or reserved by the network device (As per "available hardware addresses" database) is within the range of hardware addresses starting the target hardware address from DMAD request message, the network device generates DMAD Response message.

o The DMAD Response message specifies a group of hardware addresses within the range of hardware addresses specified by the requesting network device through DMAD Request message.

o The DMAD Response message specifies a group of hardware addresses that has been reserved by the responding device through two fields Reserved hardware address field and Range for Reserved hardware addresses field.

o In response to determining that at least one hardware address assigned or reserved by the network device (As per "available hardware addresses" database) is within the range of hardware addresses starting the target hardware address as mentioned in the DMAD Request message, the network device sometime decides to release some of the reserved hardware addresses.

o When network device decides to release some of the hardware address, Network device add entries to it's "non-available hardware addresses" database specifying both Released hardware addresses and the device identifier of the network device that has sent DMAD Request message. At the sametime, these Released hardware addresses are removed from "available hardware addresses" database of the Host or the network device. The above mentioned operation is done for all released hardware addresses.

o The DMAD Response message specifies a group of hardware addresses that has been released by the responding device through two fields Released hardware address field and Range for Released hardware addresses field.

5. DMAD Response Message Processing

When a node receives a DMAD Response message and any of the following conditions apply, the node MUST silently discard the incoming message:

- The node does not recognize DMAD Response messages.

- The node has not explicitly enabled DMAD functionality.

- The node is a VM Guest. VM Guest can not process DMAD Response packet unless explicitly authorized to handle incoming DMAD Response packet for the VM Guest.

- The node is a VM Guest and the VM Guest is not authorised to handle incoming DMAD Response packet with Range for Reserved hardware addresses field value more than one.

- The node is a VM Guest and the VM Guest is not authorised to handle incoming DMAD Response packet with Range for Released hardware addresses field value more than one.

- The node is not a VM Host or any other non-VM node that is not authorised to handle incoming DMAD Response packet either with Range for Reserved hardware addresses field value more than one or with Range for Released hardware addresses field value more than one.

- The Source Hardware Address of the incoming message is not a unicast address.

- The Destination Hardware Address of the incoming message is a multicast address.

- The value for Operation field is not "MAC-Verify-Response".

5.1 Reserved hardware address field and Range for Reserved hardware addresses field Processing

- Network devices broadcast DMAD response messages.Thus, each network device receives DMAD response messages regardless of whether the network device has sent a corresponding DMAD request message.

o Upon receiving DMAD Response message, a Host or network device retrieves Reserved hardware addresses, Range of Reserved hardware addresses and the device identifier value from DMAD Response message.

o Then the Host or the network device add entries to "non-available hardware addresses" database specifying both Reserved hardware addresses and the device identifier of the network device that has sent DMAD Response message. This is done for all reserved hardware addresses within Range of Reserved hardware addresses starting from Reserved hardware address.

o The above operation is done by all devices upon receiving DMAD Response message irrespective of whether the device has initiated corresponding DMAD Request or device has not initiated corresponding DMAD Request.

o When a device has initiated a DMAD Request message and receive corresponding DMAD response message specifying the group of Hardware addresses within the range of hardware addresses, the requesting network device determines that at least a subset of the hardware addresses are available and excluding the group of hardware addresses specified by the responding network device. Since the above mentioned subset of the hardware addresses are not reserved by any networking device, the Host or the network device successfully reserves the Range Of hardware addresses.

o Once the Host or the network device successfully reserves the Range Of hardware addresses, the Host or the network device add entries to "available hardware addresses" database specifying both Reserved hardware address and the device identifier of the network device. This is done for all reserved hardware addresses.

o When DMAD Response message has not been received for a DMAD Request message within a stipulated timeframe as decided by a network device, the Host or network device successfully reserves the Range of hardware addresses starting from Target hardware address as mentioned in the corresponding DMAD Request message.

o Then the Host or network device add entries to "available hardware addresses" database specifying both Target hardware address and the device identifier of the network device. This is done for all range of hardware addresses mentioned in the DMAD Request message.

5.2. Released hardware address field and Range for Released hardware addresses field Processing

- o Network devices broadcast DMAD Response messages. Thus, each network device receives DMAD Response messages regardless of whether the network device sent a corresponding DMAD Request message.

- o Upon receiving DMAD Response message, the Host that has sent DMAD Request message, verifies if the Response message contains information about Released hardware address then the host or the network device retrieves Released hardware addresses, Range for Released hardware addresses and the device identifier value from DMAD Response message.

- o Then the Host or the network device successfully reserves the Range of hardware addresses starting from Released hardware address.

- o Since the Host or the network device successfully reserves the above addresses, the Host or the network device add entries to "available hardware addresses" database specifying both Reserved hardware address and the device identifier of the network device. This is done for all released hardware addresses.

- o Upon receiving DMAD Response message, network device that has not sent a corresponding DMAD Request message, choose to ignore information about the Released hardware addresses.

6. Security Considerations

DMAD protocol has not proposed any additional security mechanisms (e.g authentication) to the ARP family of protocols. So there are known security issues relating to its use (e.g., host impersonation). This specification makes this existing ARP vulnerability no worse.

Fraudulent DMAD packets may appear on the broadcast network due to malicious hosts and these packets may interfer with the correct operation of other hosts. For example, it is easy for a host to answer all DMAD Requests with Replies giving a sequence of reserved hardware addresses, thereby claiming ownership of every addresses on the network. It is also possible for a host to answer all DMAD Requests with DMAD Replies giving a sequence of released hardware addresses, thereby claiming to have released certain hardware addresses on the network when those addresses are used by some other hosts.

The following are legitimate uses of DMAD:

- to determine if a group of link layer addresses, such as a MAC address, associated with a given Host and It's spawned VMs' are duplicate or not.

- to reserve a group of link layer addresses not used by any other hosts.

- to release a group of link layer addresses requested by any other hosts.

- to maintain a database of a group of link layer addresses used by other hosts upon receipt of DMAD Response message.

However, malicious users can use DMAD Request to obtain information of hardware addresses used by different hosts for illegitimate uses. For example, a malicious party can use DMAD Request message to discover hardware addresses used by different hosts. Having discovered reserved hardware addresses, the malicious party may be able to release those hardware addresses through DMAD Response message.

Understanding this risk, network operators establish policies that restrict access to DMAD functionality. In order to enforce these policies, nodes that support DMAD functionality MUST support the following configuration options:

- Enable/disable DMAD functionality. By default, DMAD functionality is disabled on an interface.

- For each interface that has enabled DMAD functionality, define the Device Identifier from which DMAD Request messages and DMAD Response messages are permitted.

- For each interface, determine whether DMAD Request messages and DMAD Response messages are accepted by looking at Device Identifier field.

When a node receives a DMAD Request message and the node is not configured to support DMAD, the node MUST silently discard the message. It is also same for DMAD Response message.

In order to protect local resources, implementations SHOULD rate-limit both incoming DMAD Request messages and incoming DMAD Response messages.

7. Use Cases

In the scenarios listed below, Hosts can use DMAD to determine if a group of link layer addresses, such as a MAC addresses, associated with a given Host and It's spawned VMs' are duplicate or not. In all scenarios, DMAD allows a host to reserve a group of link layer addresses not used by any other hosts and it allows a host to release a group of link layer addresses requested by any other hosts. DMAD achieves this goal for a host by maintaining a database of a group of link layer addresses used by other hosts. In the scenarios listed below, it is relevant to make sure that the hardware address is not a duplicate one prior to it's usage.

- The link layer address, such as a MAC address is a locally administered address and is assigned to a device by a network administrator, overriding the burned-in hardware address.

- The link layer address, such as a MAC address is a Randomly generated one by a virtualization platform and a centralised entity is not available for the virtualization platform to assign MAC addresses or keep track of assigned MAC addresses to different Hosts.

- Each virtualization platform has its own rules about how to generate random Hardware addresses and there is no centralised entity for such group of virtualization platforms.

- Multiple centralised entities are available to assign MAC addresses or keep track of assigned MAC addresses to different Hosts. There is no communication between such centralised entities about MAC address generation and subsequent assignment.

- When duplicate Hardware Address detection is required for a group of MAC addresses.

8. IEEE Considerations

This document requests the following actions from IEEE:

o Add an entry to the "Ethertype" registry, representing
the DMAD protocol.

9. Acknowledgements

TBD

10. References

10.1. Normative References

[RFC826] Plummer, D., "Ethernet Address Resolution Protocol: Or Converting Network Protocol Addresses to 48.bit Ethernet Address for Transmission on Ethernet Hardware", STD 37, RFC 826, DOI 10.17487/RFC0826, November 1982, <https://www.rfc-editor.org/info/rfc826>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC5342] Eastlake 3rd, D., "IANA Considerations and IETF Protocol Usage for IEEE 802 Parameters", RFC 5342, DOI 10.17487/RFC5342, September 2008, https://tools.ietf.org/html/rfc5342

[RFC5494] Arkko, J., Pignataro, C., "IANA Allocation Guidelines for the Address Resolution Protocol (ARP)", RFC 5494, DOI 10.17487/RFC5494, April 2009, https://www.rfc-editor.org/info/rfc5494

10.2. Informative References

[RFC2225] Laubach, M., Halpern, J., "Classical IP and ARP over ATM", RFC 2225, DOI 10.17487/RFC2225, April 1998, https://www.rfc-editor.org/info/rfc2225

[RFC2390] Bradley, T., Brown, C., Malis, A., "Inverse Address Resolution Protocol", RFC 2390, DOI 10.17487/RFC2390, September 1998, https://www.rfc-editor.org/info/rfc2390

[RFC2625] Rajagopal, M., Bhagwat, R., Rickard, W., "IP and ARP over Fibre Channel", RFC 2625, DOI 10.17487/RFC2625, June 1999, https://www.rfc-editor.org/info/rfc2625

Authors' Addresses

Manoj Nayak
Juniper Networks
Bangalore, KA 560103
India

Email: manojnayak@juniper.net

Ron Bonica
Juniper Networks
Herndon, Virginia 20171
USA

Email: rbonica@juniper.net

Rafik Puttur
Juniper Networks
Bangalore, KA 560103
India

Email: rafikp@juniper.net