

Troubleshoot Ping Issue Between ND Mgmt and CIMC when Using Shared LOM

Contents

[Introduction](#)

[Background information](#)

[What is Shared LOM?](#)

[CIMC NIC Redundacy Mode](#)

[ND OS Bonding Mode](#)

[Internal structure and Packet sending behaviors](#)

[Ping issue patterns](#)

Introduction

This document describes ping issues that occurs when using Shared LAN On Motherboard (LOM).

Background information

There are known issues in early deployments of Nexus Dashboard (ND) where Cisco Integrated Management Controller (CIMC) and ND OS of the same node cannot ping each other.

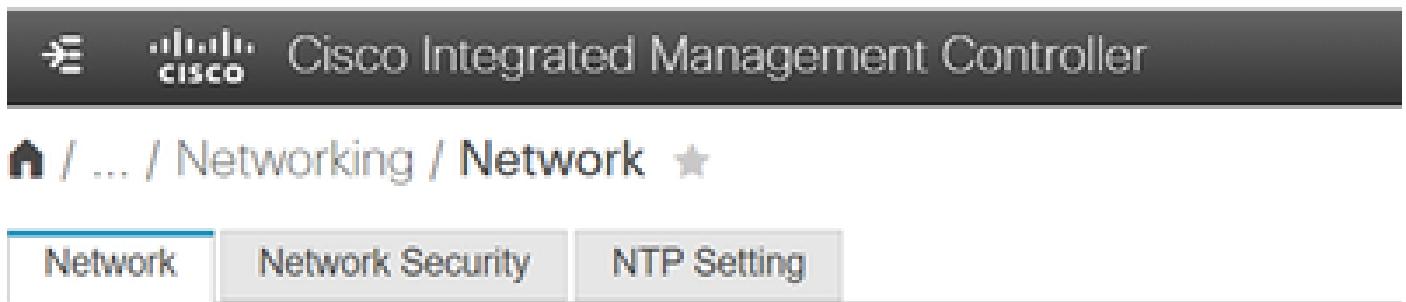
In consideration that there are no ND deployment-related requirements for CIMC - OS connectivity on the same node, the descriptions provided in this article are offered as a reference when solving issues related to connectivity behavior.

What is Shared LOM?

When using shared LOM, you can encounter a specific issue in the early deployment of ND where CIMC and ND OS of the same node cannot ping each other, and there were certain patterns of the ping issue.



The shared LOM has active/standby mode and active/active mode. It can be changed on CIMC GUI - **Admin** > **Networking** > **Network (tab)** > **NIC Properties (section)** section. Both active port and the standby port can receive packets, but only the active port can send packets.



▼ NIC Properties:

NIC Mode: Shared LOM ▼

NIC Redundancy: active-standby ▼

MAC Address: active-active
active-standby

CIMC NIC Redundancy Mode

There are two internal CIMC ports. Name them as CIMC#1 and CIMC#2 here. Recall that packets are only sent through active CIMC port when using shared LOM active/standby mode. By contrast, however, in the

active/active shared mode, both CIMC internal ports can send and receive packets.

It has been observed that the ICMP reply-sending port periodically alternates between ports. For ARP, the reply can only be sent through one port. This depends on the IP/MAC address of the sender. This behavior is similar to port channel load-balancing.

In this case, in CIMC dedicated mode, the CIMC MAC address ends in C0. With CIMC Shared-LOM in Active/Active mode, CIMC#1's MAC address ends with c1 and the CIMC#2 ends with c2.

Key points of CIMC behaviors:

- ARP reply can only be sent out through one active CIMC port
- ICMP reply can be periodically alternated between two active CIMC ports

ND OS Bonding Mode

We have two physical ports eth1-1, eth1-2 using for Nexus Dashboard OS as management ports. Although there are also two ports, mgmt0 and mgmt1, on ND OS level, they only work as active/standby mode. Unlike CIMC stand port, OS standby port cannot either send or receive packets. If packet arrived at standby OS port, it gets be discarded. Mac address of the mgmt ports end with C6. We can check the active port via command.

```
<#root>
```

```
root@pND1:~#
```

```
cat /proc/net/bonding/bond1
```

```
<<<
```

```
Ethernet Channel Bonding Driver: version 5.15.68.5
```

```
Bonding Mode: fault-tolerance (active-backup)
```

```
Primary Slave: None
```

```
Currently Active Slave: mgmt1
```

```
<<<
```

```
MII Status: up
```

```
MII Polling Interval (ms): 60
```

```
Up Delay (ms): 0
```

```
Down Delay (ms): 0
```

```
Peer Notification Delay (ms): 0
```

```
Slave Interface: mgmt1
```

```
MII Status: up
```

```
Speed: 1000 Mbps
```

```
Duplex: full
```

```
Link Failure Count: 1
```

```
Permanent HW addr: ec:01:d5:70:0d:c7
```

```
Slave queue ID: 0
```

```
Slave Interface: mgmt0
```

```
MII Status: up
```

```
Speed: 1000 Mbps
```

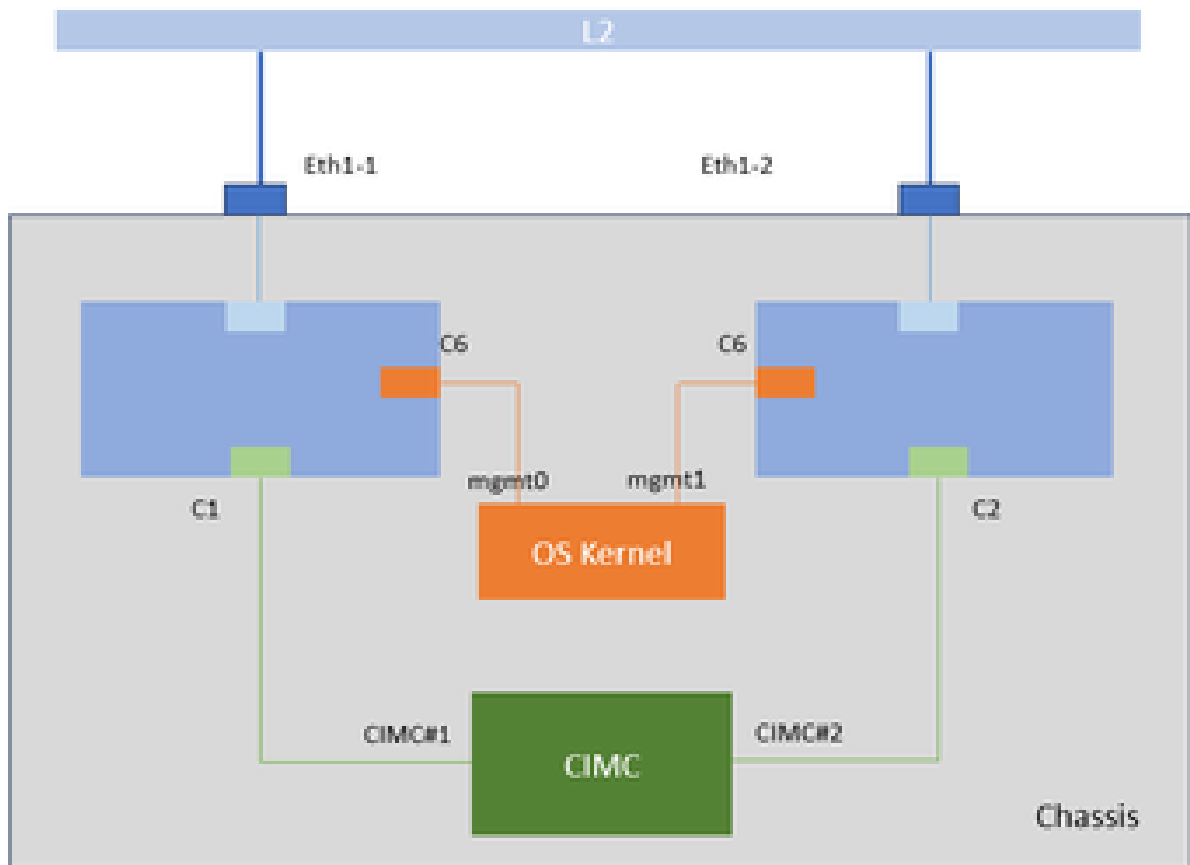
```
Duplex: full
```

```
Link Failure Count: 0
```

```
Permanent HW addr: ec:01:d5:70:0d:c6
```

Internal structure and Packet sending behaviors

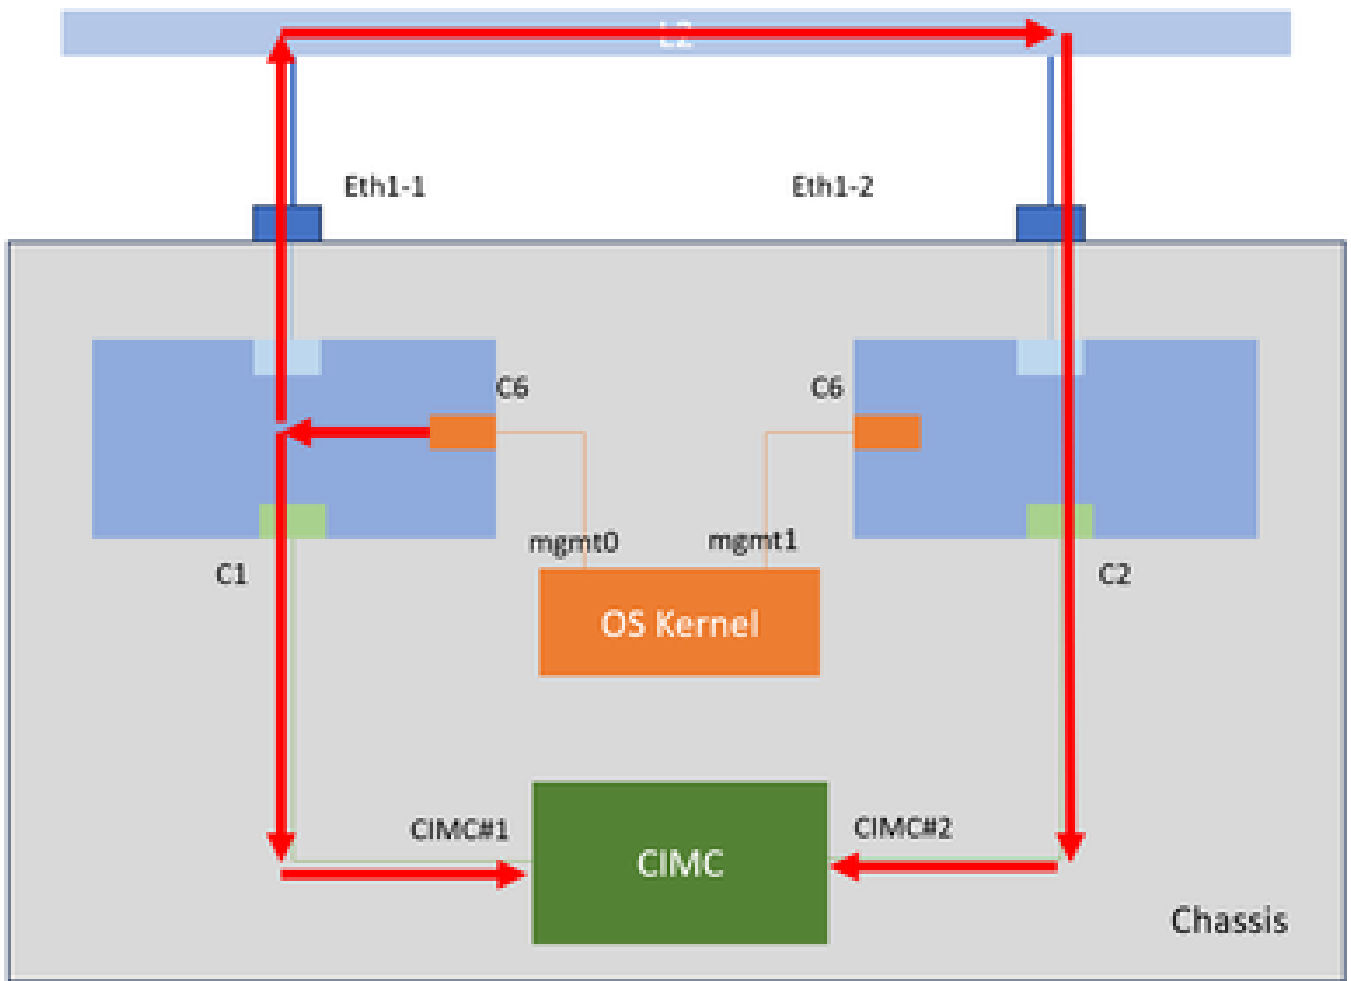
Here is the internal structure of ASIC including ND OS ports and CIMC ports, illustrated based on our lab test.



Packet sending behaviors:

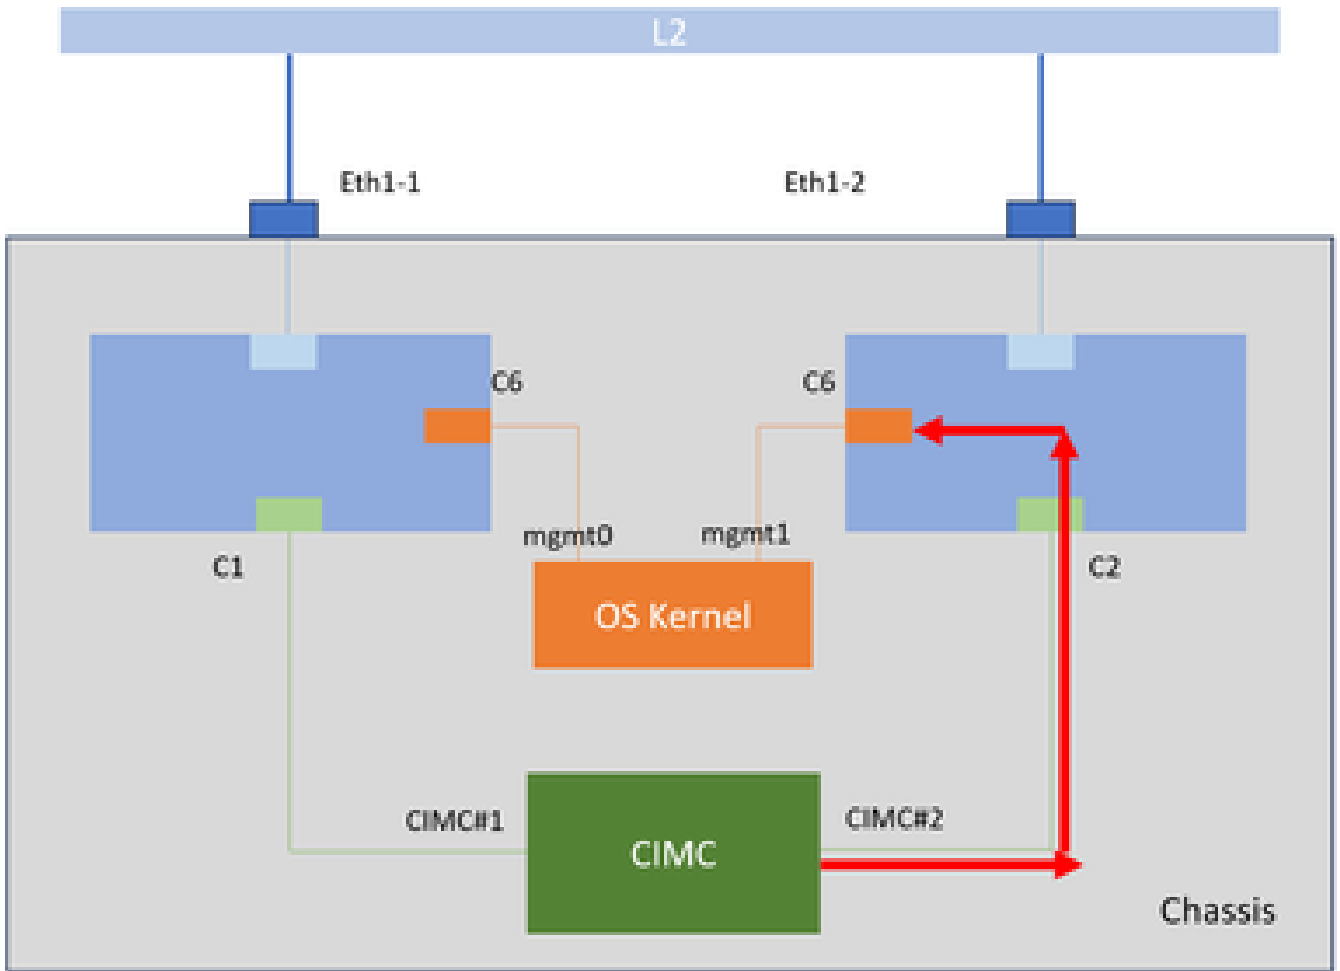
- Broadcast sent from mgmt reaches both CIMC internal ports

For example, ARP request (broadcast) sent from mgmt0 is sent through internal path to CIMC#1, also be sent through L2 to CIMC#2.



- Unicast sent from mgmt or CIMC can only be sent within internal path

For example, the ARP reply (unicast) sent from CIMC#2 is sent to mgmt1, even if mgmt1 is active port.



Ping issue patterns



Note: The behavior introduced in this section is based on the precondition that both physical ports eth1-1 and eth1-2 are linked up. No ping issue can be seen if only one link is linked up.

1. Shared LOM mode:Active/Standby

1-1. ND OS active and CIMC active are same port

Ping OK.

1-2. ND OS active and CIMC active are different port

Ping NG. ARP reply is received at mgmt standby port, which causes ping failure.

2. Shared LOM mode:Active/Active

2-1. ND OS active and CIMC arp reply port are same port

Ping success and Ping failure can be seen periodically.

2-2. ND OS active and CIMC arp reply port are different port

Ping NG. ARP reply is received at mgmt standby port, which causes ping failure.

Refer this video for details: