



Troubleshooting IOM Issues

This chapter contains the following sections:

- [IOM Terminology, on page 1](#)
- [Chassis Boot Sequence, on page 2](#)
- [Link Pinning and Failover Behavior, on page 3](#)
- [Recommended Solutions for IOM Issues, on page 4](#)

IOM Terminology

The following abbreviations and terms may be encountered while diagnosing IOM difficulties.

- HR - Host Receive Block
- NR - Network Receive Block
- SS- Switching Subsystem
- HI- Host Interface Block
- NI - Network Interface Block
- CI- CPU Interface Block
- BI- BMC Interface Block
- HIF- Host Interface
- NIF- Network Interface
- CIF- CPU Interface
- BIF- BMC Interface
- VIF- Virtual Interface
- VNTag- Virtual NIC Tag
- h2n- Host-to-Network direction. Used to describe traffic received on an HI, CI, or BI destined for an NI.
- n2h- Network-to-Host direction. Used to describe traffic received on an NI destined for an HI, CI, or BI

- Redwood - ASIC on the 2104 IOM. The basic functionality of the Redwood ASIC is to aggregate traffic to/from 8 host-facing 10G Ethernet ports connected to server adapter cards from/to 4 network-facing 10G Ethernet ports.
- Woodside - ASIC on the 2204 and 2208 IOM. Aggregates traffic to/from 32 host-facing 10G Ethernet ports from/to 4 or 8 network-facing 10G Ethernet ports.
- Chassis Management Switch (CMS) - a Marvell 88E6095 Ethernet switch integrated into the IOM.
- CMC- A CPU that controls the Redwood or Woodside ASIC and the CMS, runs the required IOM firmware, and perform other chassis management functionality.

Chassis Boot Sequence

The 2100 and 2200 series IOMs in the Cisco 5108 chassis are the only active components in the chassis itself. A single IOM is sufficient to bring the chassis up, though both would be needed in a cluster configuration.

Problems in the chassis and IOM can sometimes be traced by understanding the following boot sequence:

1. Power is applied.
2. The bootloader is invoked, the IOM memory is configured, scrubbed and ECC is enabled. The bootloader sets the IOM health LED to amber.
3. The kernel checksum and boot begins.

An alternate kernel is booted if the selected kernel's checksum fails, or if the selected kernel failed to boot the user process "OHMS" for the last two boots.

4. If the kernel can't be booted the IOM health LED blinks amber. The IOM will not be recognized by Cisco UCS Manager. If this is the only active IOM, the entire chassis will not be recognized

If the kernel boot is successful, the IOM Health LED is set to green. If the IOM is not recognized by Cisco UCS Manager, rule out a problem with the physical cabling between IOM and fabric interconnect. A single functioning physical connection should be enough for the chassis to be managed in Cisco UCS Manager. If the cabling is not the issue, the final possibility is that the firmware version on the IOM may be much older than the version of Cisco UCS Manager. This may or may not require the IOM be returned to Cisco.

5. Processes running on the Communications ASIC (either Redwood or Woodside) and the CMC Process Monitor (pmon) starts and restarts the following CMC platform processes:
 - platform_ohms - POST and run-time health monitoring
 - dmserver - device manager, caches seeprom data, scans I2C devices
 - ipmiserver - sends sensor and FRU data to UCS Manager
 - cmc_manager - set chassis info, respond to UCS Manager requests
 - cluster_manager - local cluster master and client data transfer
 - updated - listens for software update requests
 - thermal - chassis thermal management
 - pwrmgr- chassis power manager

- pppd - communication path to peer CMC over UART 2
- obfllogger - accepts client requests to log messages to OBFL flash
- rsyslogd - syslog, messages sent to UCS Manager controlled by level

If a failure at a stage 2 or 3 of the boot sequence is identified, the related components are the most likely causes and the IOM will almost certainly have to be returned. There is an HDMI console port on the IOMs that can directly monitor the IOM bootloader console, but its use is limited to Cisco's internal technicians, who would have access to the debugging software needed to make further changes such as loading in known functioning firmware images.

Table 1: Expected IOM and Chassis LED Behavior

LED	Status	LED State
IOM Health LED	Normal operation	Green
	Booting or minor error	Amber
	Major error	Blinking amber
Chassis OK LED	Booting	Off
	IOM Controlling	Green
Chassis FAIL LED	No error	Off
	Minor error	Amber
	Major error	Blinking amber

Link Pinning and Failover Behavior

Failures seen when a link between IOM and fabric interconnect (an IOM HIF port) goes down are more easily understood when the static route pinning applied to the servers in the chassis is understood. The quickest solution may be simply to reacknowledge the chassis, but understanding this topic will provide insight into when to apply that solution.

Table 2: Link Pinning on an IOM

Number of Active Fabric Links	Blade slot pinned to fabric link
1-Link	All the HIF ports are pinned to the active link
2-Link	1,3,5,7 to link-1 2,4,6,8 to link-2

Number of Active Fabric Links	Blade slot pinned to fabric link
4-Link	1,5 to link-1 2,6 to link-2 3,7 to link-3 4,8 to link-4
8-Link (Applies only to 2208XP)	1 to link-1 2 to link-2 3 to link-3 4 to link-4 5 to link-5 6 to link-6 7 to link-7 8 to link-8

Only 1,2,4 and 8 links are supported. 3,5,6, and 7 links are not valid configurations.

Here is an example of an expected behavior:

1. There are four active links on each IOM to their respective fabric interconnects.
2. Link 4 between IOM-1 and its fabric interconnect (currently Active) is accidentally unplugged by a datacenter worker.
3. Connectivity through IOM-1 from blade slots 3,4,7, and 8 fails over to IOM-2 and the standby fabric interconnect. While you might think that only slots 4 and 8 would be affected, link 3 fails administratively because 3 link configurations are not supported. Data throughput is not lost, but the failure is noted in Cisco UCS Manager.
4. At this point you can either:
 - Resolve the connectivity issue by plugging link 4 back in. Normal configured operation will resume.
 - Re-acknowledge the chassis, and the configuration will re-establish pinning to work over 2 fabric links. If at a later time the links are replaced or repaired, a second re-acknowledgement will be needed.

Recommended Solutions for IOM Issues

The following table lists guidelines and recommended solutions for troubleshooting IOM issues.

Table 3: IOM Issues

Issue	Recommended Solution
<p>The IOM Health LED turns amber at initial bootup, and stays there.</p>	<p>Re-seat the affected IOM.</p> <p>Remove and replace the IOM.</p> <p>If both IOMs in a chassis are showing the same behavior, decommission the server or chassis and call Cisco TAC.</p>
<p>The IOM health LED blinks amber but never turns green.</p>	<p>Re-seat the affected IOM.</p> <p>Remove and replace the IOM.</p> <p>If both IOMs in a chassis are showing the same behavior, decommission the server or chassis and call Cisco TAC.</p>
<p>CMC receives chassis info from Cisco UCS Manager but one or more blades are either not responding or do not accept the chassis info.</p>	<p>Verify that the IOM firmware and Cisco UCS Manager are at the same software level.</p> <p>Re-seat the affected IOM.</p> <p>Check the POST results for the Redwood or Woodside ASIC in the IOM.</p> <p>Check for runtime link down status.</p> <p>Check for failed POST tests on the affected server. The most detail on chassis info is in the CMC Manager logs. The CMC Cluster state can be compared to the fabric interconnect with the following commands</p> <p>FI: show cluster state</p> <p>cmc connected directly to the IOM: show platform software cmctrl dmclient all</p>
<p>CMC never receives chassis info from Cisco UCS Manager.</p>	<p>Verify that the IOM firmware and Cisco UCS Manager are at the same software level.</p> <p>Verify that at least one physical cable between the IOM and fabric interconnect is functioning properly.</p> <p>Check for runtime link down status.</p> <p>Re-seat the affected IOM.</p>
<p>The link to one or more servers has been lost.</p>	<p>Verify that the affected servers are in the same pinning group. Isolate and replace the downed link if possible.</p> <p>Re-seat the affected server.</p> <p>Re-seat the affected IOM.</p> <p>Re-establish pinning to the affected servers by reacknowledging the chassis.</p>


```

00:00:00.000000 | 0|
| 3 |HI3 |E| Up |1|1|1|1|1|1|1|0|0|0|0| | 0.00| 02/03/2012 23:17:36.046137 | 02/03/2012
23:17:34.815303 | 24|
| 4 |HI4 |-| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 01/01/1970 00:00:00.000000 | 01/01/1970
00:00:00.000000 | 0|
| 5 |HI5 |-| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 01/01/1970 00:00:00.000000 | 01/01/1970
00:00:00.000000 | 0|
| 6 |HI6 |-| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 01/01/1970 00:00:00.000000 | 01/01/1970
00:00:00.000000 | 0|
| 7 |HI7 |E| Up |1|1|1|1|1|1|1|0|0|0|0| | 0.00| 02/03/2012 22:53:04.761879 | 02/03/2012
22:52:44.548148 | 17|
| 8 |HI8 |-| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 01/01/1970 00:00:00.000000 | 01/01/1970
00:00:00.000000 | 0|
| 9 |HI9 |-| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 01/01/1970 00:00:00.000000 | 01/01/1970
00:00:00.000000 | 0|
|10 |HI10 |-| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 01/01/1970 00:00:00.000000 | 01/01/1970
00:00:00.000000 | 0|
|11 |HI11 |E| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 02/03/2012 20:46:44.214237 | 02/03/2012
20:49:30.606932 | 3|
|12 |HI12 |-| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 01/01/1970 00:00:00.000000 | 01/01/1970
00:00:00.000000 | 0|
|13 |HI13 |-| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 01/01/1970 00:00:00.000000 | 01/01/1970
00:00:00.000000 | 0|
|14 |HI14 |-| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 01/01/1970 00:00:00.000000 | 01/01/1970
00:00:00.000000 | 0|
|15 |HI15 |E| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 02/03/2012 20:45:30.918631 | 02/03/2012
20:48:06.811009 | 3|
|16 |HI16 |-| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 01/01/1970 00:00:00.000000 | 01/01/1970
00:00:00.000000 | 0|
|17 |HI17 |-| Dn |0|0|1|1|0|0|0|0|0|0|0| | 0.00| 01/01/1970 00:00:00.000000 | 01/01/1970
00:00:00.000000 | 0|

```