

UCS Fabric Interconnect Management Interfaces have Intermittent Connectivity Issues

Contents

[Introduction](#)

[Prerequisites](#)

[Requirements](#)

[Components Used](#)

[Problem: UCS Fabric Interconnect Management Interfaces have Intermittent Connectivity Issues.](#)

[Topology](#)

[Fabric Interconnect IPs on VLAN 10](#)

[Problematic Device Communication on VLAN 1](#)

[Troubleshoot](#)

[Root Cause](#)

[Solution](#)

[Additional Information](#)

[Related Information](#)

Introduction

This document describes how UCS Fabric Interconnect Management (Mgmt) Interfaces have experienced intermittent connectivity issues with communications to and fro a specific IP range.

Prerequisites

Requirements

Cisco recommends that you have knowledge of these topics:

- Unified Computing Systems Manager (UCSM) software
- Unified Computing System (UCS) Networking

Components Used

The information in this document is based on these software and hardware versions:

- 6200 FI
- UCS Manager 4.0

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Problem: UCS Fabric Interconnect Management Interfaces have Intermittent Connectivity Issues.

The UCS Fabric Interconnect Management Interfaces have intermittent connectivity loss, but only when communication is across a specific IP range. VLAN 10's IP Range 10.128.10.0/24 is used for Fabric Interconnect (FI) mgmt Interfaces and Virtual IP (VIP). When communication is to or from VLAN 1's IP range of 10.128.1.0/24 connectivity to and from the FIs breaks. So, any device on VLAN 1's IP range is not able to connect to UCSM, and can only ping one FI IP. At least one FI IP (of three, FI-A, FI-B, VIP) is always able to communicate.

Topology

Fabric Interconnect IPs on VLAN 10

FI-A: 10.128.10.84

FI-B: 10.128.10.85

VIP: 10.128.10.86

GW: 10.128.10.1

Problematic Device Communication on VLAN 1

Subnet 10.128.1.0/24

GW: 10.128.1.1

Troubleshoot

From the local-mgmt context of both Fabric Interconnects, it is able to reach its default (df) gateway (gw), 10.128.10.1. but no IP address on the VLAN 1 IP range of 10.128.1.0/24 is reachable to, or from, the Fabric Interconnects local-mgmt context.

At first, this appears to be an issue with routing at the gateway, and not a UCS issue, as this is simply mgmt interface on Fabric Interconnects and if it can reach the gateway and any other IP range. This presents as a Layer 3 route issue on the upstream network.

When traceroute runs from the Fabric Interconnect to a random IP range (and any other IP range not in VLAN 1's range) (for instance an IP from VLAN 20: 10.128.20.1), the first hop on the traceroute is the VLAN 10's Gateway of 10.128.10.1 and ping is successful.

When traceroute runs to the known, problematic IP range 10.128.1.x/24, the traceroute fails.

- This leads to the reason to investigate on the UCS side.
- This should not have failed, and the first-hop should be the mgmt interface's df gw, 10.128.10.1, just like it was for 10 out of 10 other IPs that were tested and were outside of VLAN 10 IP range, and were not in VLAN 1 IP range.

In order to investigate further, you ran an ethanalyzer to see what goes on and when VLAN 1's IP

range is pinged, ARP acts curious:

```
EWQLOVIUCS02-A(nxos)# ethalyzer local interface mgmt display-filter arp limit-captured-frames 0
```

Capturing on eth0

```
2019-12-17 11:45:50.807837 00:de:fb:a9:37:e1 -> ff:ff:ff:ff:ff:ff ARP Who has 10.128.1.77? Tell 10.128.0.142
```

```
2019-12-17 11:45:51.807835 00:de:fb:a9:37:e1 -> ff:ff:ff:ff:ff:ff ARP Who has 10.128.1.77? Tell 10.128.0.142
```

```
2019-12-17 11:45:52.807827 00:de:fb:a9:37:e1 -> ff:ff:ff:ff:ff:ff ARP Who has 10.128.1.77? Tell 10.128.0.142
```

```
2019-12-17 11:45:55.807829 00:de:fb:a9:37:e1 -> ff:ff:ff:ff:ff:ff ARP Who has 10.128.1.77? Tell 10.128.0.142
```

The expected behavior was to ask who has this VLAN 1 IP, but then tell the mgmt VLAN 10's Gateway.

However, when VLAN 1's IP range is pinged, ARP asks who has that IP and to tell 10.128.0.142, follow these:

- This should be "Tell 10.128.10.1" as it was with all the other IPs tested.
- You do not even recognize a 10.128.0.x Subnet or IP address.

This is an issue why the FI would tell 10.128.0.142, during the investigation of UCS domain it was found that this IP address was applied to server 1/5's CIMC:

```
EWQLOVIUCS02-B(local-mgmt)# show mgmt-ip-debug ip-tables
```

<SNIPPED>

Chain PREROUTING (policy ACCEPT 5303K packets, 360M bytes)

pkts	bytes	target	prot	opt	in	out	source	destination	
188	9776	cimcnat	tcp	--	*	*	0.0.0.0/0	0.0.0.0/0	tcp
dpt:443									
0	0	cimcnat	tcp	--	*	*	0.0.0.0/0	0.0.0.0/0	tcp
dpt:80									
0	0	DNAT	icmp	--	*	*	0.0.0.0/0	10.128.10.85	
to:127.6.1.1									
0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.85	tcp
dpt:2068 to:127.6.1.1:2068									
0	0	DNAT	udp	--	*	*	0.0.0.0/0	10.128.10.85	udp
dpt:623 to:127.6.1.1:623									
0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.85	tcp
dpt:22 to:127.6.1.1:22									
449	26940	DNAT	icmp	--	*	*	0.0.0.0/0	10.128.10.108	
to:127.6.1.2									

0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.108	tcp
dpt:2068 to:127.6.1.2:2068									
0	0	DNAT	udp	--	*	*	0.0.0.0/0	10.128.10.108	udp
dpt:623 to:127.6.1.2:623									
0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.108	tcp
dpt:22 to:127.6.1.2:22									
931	55860	DNAT	icmp	--	*	*	0.0.0.0/0	10.128.10.107	
to:127.6.1.3									
0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.107	tcp
dpt:2068 to:127.6.1.3:2068									
0	0	DNAT	udp	--	*	*	0.0.0.0/0	10.128.10.107	udp
dpt:623 to:127.6.1.3:623									
0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.107	tcp
dpt:22 to:127.6.1.3:22									
0	0	DNAT	icmp	--	*	*	0.0.0.0/0	10.128.10.104	
to:127.6.1.3									
0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.104	tcp
dpt:2068 to:127.6.1.3:2068									
0	0	DNAT	udp	--	*	*	0.0.0.0/0	10.128.10.104	udp
dpt:623 to:127.6.1.3:623									
0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.104	tcp
dpt:22 to:127.6.1.3:22									
920	55200	DNAT	icmp	--	*	*	0.0.0.0/0	10.128.10.106	
to:127.6.1.4									
0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.106	tcp
dpt:2068 to:127.6.1.4:2068									
0	0	DNAT	udp	--	*	*	0.0.0.0/0	10.128.10.106	udp
dpt:623 to:127.6.1.4:623									
0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.106	tcp
dpt:22 to:127.6.1.4:22									
912	54720	DNAT	icmp	--	*	*	0.0.0.0/0	10.128.10.105	
to:127.6.1.6									
0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.105	tcp
dpt:2068 to:127.6.1.6:2068									
0	0	DNAT	udp	--	*	*	0.0.0.0/0	10.128.10.105	udp
dpt:623 to:127.6.1.6:623									
0	0	DNAT	tcp	--	*	*	0.0.0.0/0	10.128.10.105	tcp
dpt:22 to:127.6.1.6:22									

0 0 DNAT icmp -- * * 0.0.0.0/0 10.128.0.142
to:127.6.1.5 <<---- Indicates that 10.128.0.142 is the OOB KVM IP address for server 1/5.

```

0 0 DNAT tcp -- * * 0.0.0.0/0 10.128.0.142 tcp
dpt:2068 to:127.6.1.5:2068

0 0 DNAT udp -- * * 0.0.0.0/0 10.128.0.142 udp
dpt:623 to:127.6.1.5:623

0 0 DNAT tcp -- * * 0.0.0.0/0 10.128.0.142 tcp
dpt:22 to:127.6.1.5:22

910 54600 DNAT icmp -- * * 0.0.0.0/0 10.128.10.102
to:127.6.1.7

0 0 DNAT tcp -- * * 0.0.0.0/0 10.128.10.102 tcp
dpt:2068 to:127.6.1.7:2068

0 0 DNAT udp -- * * 0.0.0.0/0 10.128.10.102 udp
dpt:623 to:127.6.1.7:623

0 0 DNAT tcp -- * * 0.0.0.0/0 10.128.10.102 tcp
dpt:22 to:127.6.1.7:22

908 54480 DNAT icmp -- * * 0.0.0.0/0 10.128.10.101
to:127.6.1.8

0 0 DNAT tcp -- * * 0.0.0.0/0 10.128.10.101 tcp
dpt:2068 to:127.6.1.8:2068

0 0 DNAT udp -- * * 0.0.0.0/0 10.128.10.101 udp
dpt:623 to:127.6.1.8:623

0 0 DNAT tcp -- * * 0.0.0.0/0 10.128.10.101 tcp
dpt:22 to:127.6.1.8:22

```

<SNIPPED>

Root Cause

The issue was a mistyped static CIMC IP address for server 1/5.

Additionally, it was put in a subnet of 255.255.248.0

This created an unwanted entry in the Fabric Interconnect's route table. One that would hit the condition before it hit the default route for all IPs in the range of 10.128.0.1 - 10.128.7.254

```
Linux(debug)# route -n
```

```
Kernel IP routing table
```

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
10.128.10.0	0.0.0.0	255.255.255.0	U	0	0	0	eth0
127.15.1.0	0.0.0.0	255.255.255.0	U	0	0	0	vlan4042
127.7.0.0	0.0.0.0	255.255.0.0	U	0	0	0	vlan4043
127.5.0.0	0.0.0.0	255.255.0.0	U	0	0	0	vlan4044

```

127.14.0.0      0.0.0.0      255.255.0.0   U    0    0    0  vlan4046
127.12.0.0      0.0.0.0      255.255.0.0   U    0    0    0  bond0
127.9.0.0       0.0.0.0      255.255.0.0   U    0    0    0  vlan4047
10.0.0.0        0.0.0.0      255.0.0.0     U    0    0    0  eth0    <<-----
Undesired route entry
10.0.0.0        0.0.0.0      255.0.0.0     U    0    0    0  eth0    <<-----
Undesired route entry
0.0.0.0         10.128.10.1  0.0.0.0        UG   0    0    0  eth0

```

Solution

The solution for this case is to browse through UCSM from an unaffected IP range and correct Server 1/5's CIMC Out of Band (OOB) static address. It is pulled from the OOB mgmt pool and is already set up. It should be used like every other server in the environment.

Additional Information

If the Fabric interconnect is rebooted, it works sometimes. The issue is with the managing Instance of that server. The undesired route table entry is only created on the Fabric Interconnect. When the managing instance was the same Fabric Interconnect as the Primary Fabric Interconnect, they are unable to reach the VIP or that Fabric Interconnect.

CIMC management IP assignment should always be within the same IP range as the Fabric Interconnect's OOB IP range.

Related Information

- [Cisco UCS Manager Administration Management Guide 3.1](#)
- [Technical Support & Documentation - Cisco Systems](#)