

LACP Between Nexus 9000 and UCS Server



Document ID: 118851

Contributed by Viral Bhutta and Manuel Velasco, Cisco TAC Engineers.

Mar 18, 2015

Contents

Introduction

Prerequisites

Requirements

Components Used

Problem

Solution

Introduction

This document describes how to form a Link Aggregation Control Protocol (LACP) port-channel between the Nexus 9000 and the Unified Computing System (UCS) server.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

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

- Nexus 9300 that runs Version 6.1(2)I(3)3a
- UCS C240 M3 Series Server that runs Cisco Integrated Management Controller (CIMC) Version 2.0(3f)
- Cisco Virtual Interface Card (VIC) Version 1225
- Red Hat Enterprise Linux Version 7.0

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, make sure that you understand the potential impact of any command.

Problem

The Nexus 9000 is not able to form LACP port-channels with the UCS server. If *no lacp suspend-individual* is configured on the port-channel, then it goes to I state. If it is not configured, then it suspends the port due to no LACP Data Unit (LACPDU).

The LACPDU interface counters on the Nexus 9000 show that transmitted (Tx) counters get incremented but received (Rx) counters are not incremented.

Here is how to configure an access port on Nexus 9000 and disable the LACP suspend-individual:

```
interface port-channel1019
  switchport access vlan 4
  no lacp suspend-individual

interface Ethernet1/19
  description csm-b-ceph-001
  switchport access vlan 4
  spanning-tree port type edge
  spanning-tree bpduguard enable
  spanning-tree guard root
  channel-group 1019 mode active
```

This can also be configured as a trunk-edge port; however, the behavior is not different:

Nexus9396X-2# **show port-c sum**

```
Flags: D - Down P - Up in port-channel (members)
I - Individual H - Hot-standby (LACP only)
s - Suspended r - Module-removed
S - Switched R - Routed
U - Up (port-channel)
M - Not in use. Min-links not met
```

```
-----
Group Port- Type Protocol Member Ports
Channel
```

```
-----
1019 Po1019(SD) Eth LACP Eth1/19(I) ----> Indicates No PDUs are being received
```

This can be (s) for suspended if the LACP suspend-individual is enabled (by default, it is enabled). However, the reason should be the same as why it does not receive Protocol Data Units (PDUs) from the peer.

Nexus9396X-2# **show lacp interface e1/19**

```
Interface Ethernet1/19 is individual
Channel group is 1019 port channel is Po1019
PDUs sent: 63
PDUs rcvd: 0 ----> Indicates No PDUs are being received from UCS server.
Markers sent: 0
Markers rcvd: 0
Marker response sent: 0
Marker response rcvd: 0
Unknown packets rcvd: 0
Illegal packets rcvd: 0
Lag Id: [ [(0, 0-0-0-0-0-0, 0, 0, 0), (0, 0-0-0-0-0-0, 0, 0, 0)] ]
Operational as aggregated link since Thu Jan 1 00:00:00 1970
```

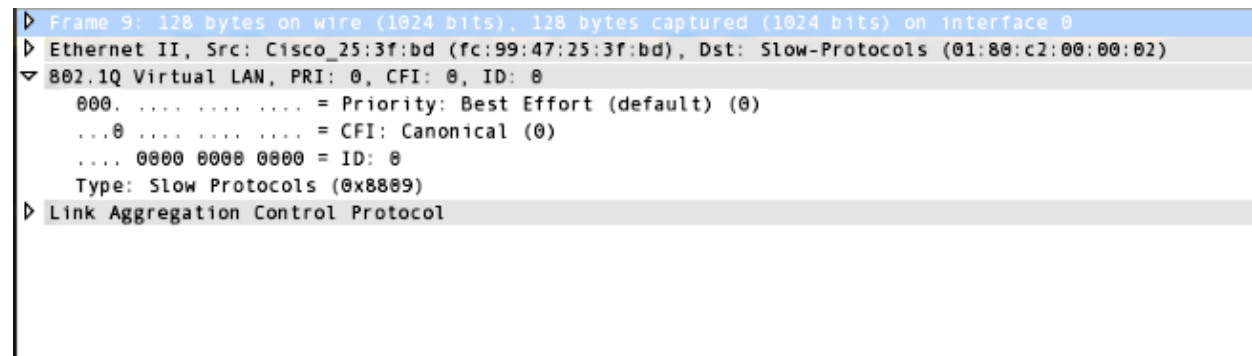
```
Local Port: Eth1/19 MAC Address= 7c-69-f6-10-59-d7
System Identifier=0x8000, Port Identifier=0x8000,0x113
Operational key=33787
LACP_Activity=active
LACP_Timeout=Long Timeout (30s)
Synchronization=IN_SYNC
Collecting=true
Distributing=true
Partner information refresh timeout=Long Timeout (90s)
Actor Admin State=125
Actor Oper State=125
Neighbor: 0x0
MAC Address= 0-0-0-0-0-0
System Identifier=0x0, Port Identifier=0x0,0x0
Operational key=0
LACP_Activity=unknown
LACP_Timeout=Long Timeout (30s)
Synchronization=NOT_IN_SYNC
```

```
Collecting=false
Distributing=false
Partner Admin State=0
Partner Oper State=0
Aggregate or Individual(True=1)= 1
```

Ethalyzer also does not capture PDUs that come from the peer device. However, a *tcpdump* that is taken from the server indicates that it sends as well as receives PDUs.

Solution

The lab setup showed that when the LACPDU was captured on the Nexus 9000, it was discovered that it sends LACPDU with the dot1q header as seen here:

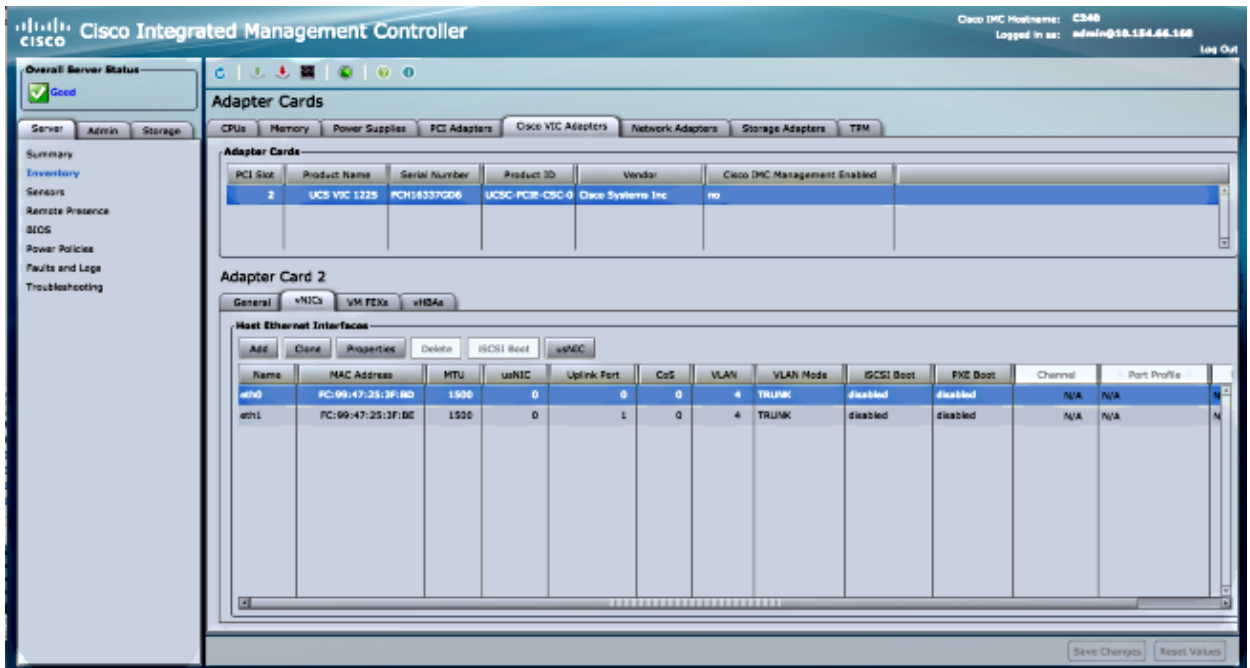
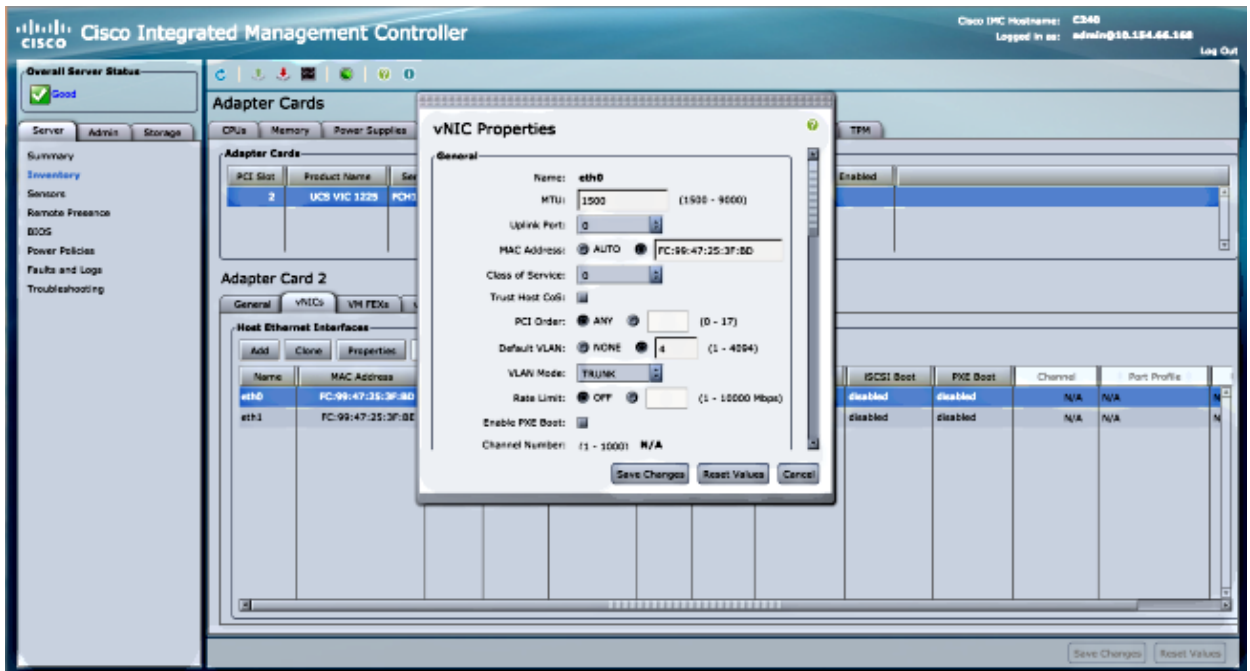


```
▶ Frame 9: 128 bytes on wire (1024 bits), 128 bytes captured (1024 bits) on interface 0
▶ Ethernet II, Src: Cisco_25:3f:bd (fc:99:47:25:3f:bd), Dst: Slow-Protocols (01:80:c2:00:00:02)
▼ 802.1Q Virtual LAN, PRI: 0, CFI: 0, ID: 0
    000. .... = Priority: Best Effort (default) (0)
    ...0 .... = CFI: Canonical (0)
    .... 0000 0000 0000 = ID: 0
    Type: Slow Protocols (0x8809)
▶ Link Aggregation Control Protocol
```

However, the VLAN ID is set to 0. Now, when you look at the configuration for Virtual Network Interface Card (vNIC), you see that by default the VLAN is set to *None* and the VLAN Mode is set to *access*. Now even in this mode, it sends LACPDU with the dot1q header. However, the Nexus 9000 does not identify VLAN 0 and hence this packet is dropped. This behavior is documented in UCS C-Series Rack Servers VIC Connectivity Options.

In order to make this work, you must configure the VLAN Mode as *Trunk* and also configure the default VLAN as a VLAN that is allowed on that port. In order to change this, you log into your server's CIMC IP address.

1. Click the *Server tab* in the left pane and then click *Inventory*.
2. Click *Cisco VIC adapters* in the right pane and then click *vNICs*.
3. Choose the interface and then click *Properties*. You should be able to make the changes here:



- Once the changes are made, **Save** changes. You must reset the server for the changes to take effect. After the reset, the port-channel should form correctly.

In newer releases of Nexus 9000, Version 6.1(2)I(3)4 and 7.0(3)I(1) and later, it ignores the dot1q header and packets reach Supervisor (SUP). This behavior is documented in Cisco bug ID CSCur69668.

Hence you should not encounter this issue in the newer releases.

