

Addressing ACI Fault Code F0321, F0323, F0325: unhealthy - cluster diverged or degraded leadership

Contents

[Introduction](#)

[Additional Details](#)

[Quick Start to Address Fault](#)

[1. Command "acidiag cluster"](#)

[2. APIC SSD Health](#)

[3. DME Processes Status](#)

[Next Steps:](#)

[1. APIC Connectivity Issues](#)

[2. DME Process Down](#)

[4. Check Core Files](#)

[3. Collect TechSupport and Upload to SR](#)

Introduction

This document describes next steps for remediation of the below fault:

```
"Code" : "F0321",  
"Description" : "Controller <id> is unhealthy because: Data Layer Partially Degraded  
Leadership",  
"Dn" : "topology/pod-<POD-ID>/node-<NODE-ID>/av/node-<NODE-ID>/fault-F0321",
```

```
"Code" : "F0321",  
"Description" : "Controller 3 is unhealthy because: Data Layer Partially Diverged"  
"Dn" : "topology/pod-<POD-ID>/node-<NODE-ID>/av/node-<NODE-ID>/fault-F0321",
```

```
"Code" : "F0325",  
"Description" : "Connectivity has been lost to the leader for some data subset(s) of a service  
on <node >, the service may have unexpectedly restarted or failed",  
"Dn" : "topology/pod-<POD-ID>/node-<NODE-ID>/av/node-<NODE-ID>/fault-F0325",
```

```
"Code" : "F0323",  
"Description" : "Lost connectivity to leader for some data subset(s) of Access <Service> on  
<controller >",  
"Dn" : "topology/pod-<POD-ID>/node-<NODE-ID>/av/node-<NODE-ID>/fault-F0323",
```

If you have an Intersight connected ACI fabric, a Service Request was generated on your behalf to indicate that instances of this fault was found within your Intersight-Connected ACI fabric.

This specific fault is raised when APIC Cluster is unhealthy. Data Layer Partially Diverged is seen when either one of shard/replica is down which is denoted by "\" in "acidiag rvread output". This fault can also be seen when replica or database is completely missing from the APIC denoted by "X". we need to fix any underlying issue and restore the health of the cluster.

This is being actively monitored as part of [Proactive ACI Engagements](#).

Additional Details

Please **DO NOT** try any **intrusive steps** such as power off or reload or decommission to troubleshoot the clustering issue, if the fabric is in production. Collect and Upload the TS files to the TAC case to find out the exact steps to restore the APIC Cluster.

Quick Start to Address Fault

1. Command "acidiag cluster"

By running this command, it would do multiple checks including connectivity with the APICs. We should see all the test results return OK. If we notice anything other than OK, we will need to investigate the cause of it.

```
##### Sample output on a healthy cluster ##### apicl# acidiag cluster Admin password:
Running... Checking Wiring and UUID: OK Checking AD Processes: Running Checking All Apics in
Commission State: OK Checking All Apics in Active State: OK Checking Fabric Nodes: OK Checking
Apic Fully-Fit: OK Checking Shard Convergence: OK Checking Leadership Degration: Optimal leader
for all shards Ping OOB IPs: APIC-1: 10.197.204.149 - OK APIC-2: 10.197.204.150 - OK APIC-3:
10.197.204.151 - OK Ping Infra IPs: APIC-1: 10.0.0.1 - OK APIC-2: 10.0.0.2 - OK APIC-3: 10.0.0.3
- OK Checking APIC Versions: Same (5.2(4d)) Checking SSL: OK Full file system(s): None Done!
##### Sample output on a unhealthy cluster ##### apicl# acidiag cluster Admin password:
Running... Checking Wiring and UUID: switch(302) reports apic(3) has wireIssue: unapproved-ctrlr
Checking AD Processes: Running Checking All Apics in Commission State: OK Checking All Apics in
Active State: OK Checking Fabric Nodes: OK Checking Apic Fully-Fit: OK Checking Shard
Convergence: OK Checking Leadership Degration: Non optimal leader for shards :
3:1,3:2,3:4,3:5,3:7,3:8,3:10,3:11,3:13,3:14,3:16,3:17,3:19,3:20,3:22,3:23,3:25,3:26,3:28,3:29,3:
31,3:32,6:1,6:2,6:4,6:5,6:7,6:8,6:10,6:11,6:13,6:14,6:16,6:17,6:19,6:20,6:22,6:23,6:25,6:26,6:28
,6:29,6:31,6:32,9:1,9:2,9:4,9:5,9:7,9:8,9:10,9:11,9:13,9:14,9:16,9:17,9:19,9:20,9:22,9:23,9:25,9
:26,9:28,9:29,9:31,9:32,10:1,10:2,10:4,10:5,10:7,10:8,10:10,10:11,10:13,10:14,10:16,10:17,10:19,
10:20,10:22,10:23,10:25,10:26,10:28,10:29,10:31,10:32,11:1,11:2,11:4,11:5,11:7,11:8,11:10,11:11,
11:13,11:14,11:16,11:17,11:19,11:20,11:22,11:23,11:25,11:26,11:28,11:29,11:31,11:32,14:1,14:2,14
:4,14:5,14:7,14:8,14:10,14:11,14:13,14:14,14:16,14:17,14:19,14:20,14:22,14:23,14:25,14:26,14:28,
14:29,14:31,14:32,16:1,16:2,16:4,16:5,16:7,16:8,16:10,16:11,16:13,16:14,16:16,16:17,16:19,16:20,
16:22,16:23,16:25,16:26,16:28,16:29,16:31,16:32,22:1,22:2,22:4,22:5,22:7,22:8,22:10,22:11,22:13,
22:14,22:16,22:17,22:19,22:20,22:22,22:23,22:25,22:26,22:28,22:29,22:31,22:32,23:1,23:2,23:4,23:
5,23:7,23:8,23:10,23:11,23:13,23:14,23:16,23:17,23:19,23:20,23:22,23:23,23:25,23:26,23:28,23:29,
23:31,23:32,33:1,34:1,34:2,34:4,34:5,34:7,34:8,34:10,34:11,34:13,34:14,34:16,34:17,34:19,34:20,3
4:22,34:23,34:25,34:26,34:28,34:29,34:31,34:32,35:1,35:2,35:4,35:5,35:7,35:8,35:10,35:11,35:13,3
5:14,35:16,35:17,35:19,35:20,35:22,35:23,35:25,35:26,35:28,35:29,35:31,35:32,36:1,39:1,39:2,39:4
,39:5,39:7,39:8,39:10,39:11,39:13,39:14,39:16,39:17,39:19,39:20,39:22,39:23,39:25,39:26,39:28,39
:29,39:31,39:32 Ping OOB IPs: APIC-1: 10.197.204.184 - OK APIC-2: 10.197.204.185 - OK APIC-3:
10.197.204.186 - OK Ping Infra IPs: APIC-1: 10.0.0.1 - OK APIC-2: 10.0.0.2 - OK APIC-3: 10.0.0.3
- OK Checking APIC Versions: Same (5.2(3e)) Checking SSL: OK Full file system(s): None Done!
```

2. APIC SSD Health

Make sure APIC SSD are healthy and one of these faults are not raised on the ACI Fabirc - F2730, F2731 and F2732. Following are the commands to run on APIC CLI to find if any of these faults exist OR same could be verified on the GUI (System > Faults)

- show faults code F2730 controller
- show faults code F2731 controller

- show faults code F2732 controller

```
##### Example: # faultRecord ack : no cause : equipment-wearout changeSet :
available:unspecified, blocks:unspecified, capUtilized:0, device:Solid State Device,
fileSystem:/dev/sdb, firmwareVersion:Dxxxxxxx, mediaWearout:1, model:INTEL SSDSC2BB120G4,
mount:/dev/sdb, name:/dev/sdb, operSt:ok, serial:ABCDxxxxxxxxxxxxXYZ, used:unspecified
childAction : code : F2730 created : 2022-01-10T03:13:08.026+00:00 delegated : no descr :
Storage unit /dev/sdb on Node 3 with hostname apicl.cisco.com mounted at /dev/sdb has 1% life
remaining dn : topology/pod-2/node-3/sys/ch/p-[/dev/sdb]-f-[/dev/sdb]/fault-F2730 domain : infra
highestSeverity : warning lastTransition : 2022-01-10T03:13:08.026+00:00 lc : raised occur : 1
origSeverity : warning prevSeverity : warning rule : eqpt-storage-wearout-warning severity :
warning status : subject : equipment-wearout type : operational # faultRecord ack : no cause :
equipment-wearout changeSet : available:unspecified, blocks:unspecified, capUtilized:0,
device:Solid State Device, fileSystem:/dev/sdb, firmwareVersion:Dxxxxxxx, mediaWearout:1,
model:INTEL SSDSC2BB120G4, mount:/dev/sdb, name:/dev/sdb, operSt:ok, serial:ABCDxxxxxxxxxxxxXYZ,
used:unspecified childAction : code : F2731 created : 2022-01-10T03:13:08.026+00:00 delegated :
no descr : Storage unit /dev/sdb on Node 3 mounted at /dev/sdb has 1% life remaining dn :
topology/pod-2/node-3/sys/ch/p-[/dev/sdb]-f-[/dev/sdb]/fault-F2731 domain : infra
highestSeverity : major lastTransition : 2022-01-10T03:13:08.026+00:00 lc : raised occur : 1
origSeverity : major prevSeverity : major rule : eqpt-storage-wearout-major severity : major
status : subject : equipment-wearout type : operational
```

3. DME Processes Status

Check if all DME process are running

Run `ps -aux | egrep "svc|nginx.bin|dhcp"`

Expected output below:

```
apicl# ps -ef | egrep "svc|nginx.bin|dhcp" root 3063 1 5 22:08 ? 00:04:40 /mgmt//bin/nginx.bin -
p /data//nginx/ root 8889 1 7 21:53 ? 00:06:43 /mgmt//bin/svc_ifc_appliancedirector.bin --x ifc
8891 1 1 21:53 ? 00:01:29 /mgmt//bin/svc_ifc_policydist.bin --x root 8893 1 2 21:53 ? 00:02:28
/mgmt//bin/svc_ifc_bootmgr.bin --x ifc 8894 1 1 21:53 ? 00:01:41 /mgmt//bin/svc_ifc_vmmmgr.bin -
-x ifc 8895 1 2 21:53 ? 00:02:14 /mgmt//bin/svc_ifc_topomgr.bin --x ifc 8901 1 2 21:53 ?
00:02:22 /mgmt//bin/svc_ifc_observer.bin --x root 8903 1 1 21:53 ? 00:01:40
/mgmt//bin/svc_ifc_plgnhandler.bin --x ifc 8914 1 1 21:53 ? 00:01:34
/mgmt//bin/svc_ifc_domainmgr.bin --x ifc 8915 1 2 21:53 ? 00:02:04 /mgmt//bin/svc_ifc_dbgr.bin -
-x ifc 8917 1 1 21:53 ? 00:01:34 /mgmt//bin/svc_ifc_edmgr.bin --x ifc 8918 1 1 21:53 ? 00:01:22
/mgmt//bin/svc_ifc_vtap.bin --x ifc 8922 1 2 21:53 ? 00:02:09 /mgmt//bin/svc_ifc_eventmgr.bin --
x ifc 8925 1 3 21:53 ? 00:03:15 /mgmt//bin/svc_ifc_reader.bin --x ifc 8929 1 1 21:53 ? 00:01:34
/mgmt//bin/svc_ifc_idmgr.bin --x ifc 8930 1 1 21:53 ? 00:01:26 /mgmt//bin/svc_ifc_licensemgr.bin
--x ifc 8937 1 3 21:53 ? 00:03:18 /mgmt//bin/svc_ifc_policymgr.bin --x ifc 8941 1 1 21:53 ?
00:01:34 /mgmt//bin/svc_ifc_scripthandler.bin --x root 11157 1 1 21:54 ? 00:01:29
/mgmt//bin/dhcpd.bin -f -4 -cf /data//dhcp/dhcpd.conf -lf /data//dhcp/dhcpd.lease -pf
/var/run//dhcpd.pid --no-pid bond0.3902 root 11170 1 4 21:54 ? 00:04:15
/mgmt//bin/svc_ifc_ae.bin --x admin 17094 16553 0 23:27 pts/0 00:00:00 grep -E
svc|nginx.bin|dhcp
```

You can check fault code F1419 for failed DMEs.

```
apicl# show faults code F1419 history ID : 4294971876 Description : Service policymgr failed on
apic bgl-aci02-apicl of fabric POD02 with a hostname bgl-aci02-apicl Severity : major DN : subj-
[topology/pod-1/node-1/sys/proc/proc- policymgr]/fr-4294971876 Created : 2022-03-
21T18:29:20.570+12:00 Code : F1419 Type : operational Cause : service-failed Change Set : id
(Old: 5152, New: 0), maxMemAlloc (Old: 1150246912, New: 0), operState (Old: up, New: down)
```

Action : creation Domain : infra Life Cycle : soaking Count Fault Occurred : 1 Acknowledgement
Status : no

Next Steps:

1. APIC Connectivity Issues

If there is loss of connectivity between apics one of possible reasons could be Wiring issues. Acidiag Cluster command will also show what kind of wiring issues is present on the link. Here are all possible wiring issues:

ctrlr-uuid-mismatch – APIC UUID mismatch (duplicate APIC ID)

fabric-domain-mismatch – Adjacent node belongs to a different fabric

wiring-mismatch – Invalid connection (Leaf to Leaf, Spine to non-leaf, Leaf fabric port to non-spine etc.)

adjacency-not-detected – No LLDP adjacency on fabric port

infra-vlan-mismatch – Infra VLAN mismatch between leaf and APIC.

pod-id-mismatch – Pod ID mismatch between APIC and Leaf

unapproved-ctrlr – The SSL handshake between APIC and connected leaf is not completed.

unapproved-serialnumber – Detected a node that is not present in Apic's DB.

2. DME Process Down

if output from DME process status section is not matching with the expected output. Try to start the DME by using 'acidiag start <DME>' for example if svc_ifc_eventmgr is missing try 'acidiag start eventmgr'

```
apicl# ps -aux | egrep "svc|nginx.bin|dhcp" root 5112 7.3 0.4 1033952 323180 ? Ssl Mar10 3073:27
/mgmt//bin/svc_ifc_appliancedirector.bin --x ifc 5117 1.7 0.6 1062664 439876 ? Ssl Mar10 720:52
/mgmt//bin/svc_ifc_topomgr.bin --x ifc 5118 2.1 2.2 2164512 1468200 ? Ssl Mar10 884:11
/mgmt//bin/svc_ifc_policymgr.bin --x ifc 5119 1.5 0.3 1115984 256904 ? Ssl Mar10 664:51
/mgmt//bin/svc_ifc_licensemgr.bin --x ifc 5120 1.5 0.5 1088252 356760 ? Ssl Mar10 666:26
/mgmt//bin/svc_ifc_edmgr.bin --x root 5121 1.6 0.6 1125948 423392 ? Ssl Mar10 698:11
/mgmt//bin/svc_ifc_bootmgr.bin --x ifc 5123 2.3 1.2 1474388 800564 ? Ssl Mar10 994:15
/mgmt//bin/svc_ifc_eventmgr.bin --x ifc 5126 1.5 8.2 6032524 5363184 ? Ssl Mar10 635:58
/mgmt//bin/svc_ifc_reader.bin --x root 5130 4.6 0.6 1092480 439580 ? Ssl Mar10 1927:08
/mgmt//bin/svc_ifc_ae.bin --x ifc 5132 1.6 0.8 1312136 567420 ? Ssl Mar10 689:43
/mgmt//bin/svc_ifc_vmmmgr.bin --x ifc 5133 1.5 0.5 1064176 346760 ? Ssl Mar10 659:03
/mgmt//bin/svc_ifc_domainmgr.bin --x ifc 5135 1.8 1.6 1736876 1099924 ? Ssl Mar10 770:39
/mgmt//bin/svc_ifc_observer.bin --x root 5141 1.5 0.7 1092948 458156 ? Ssl Mar10 663:41
/mgmt//bin/svc_ifc_plgnhandler.bin --x ifc 5146 2.0 0.6 1037676 397236 ? Ssl Mar10 857:43
/mgmt//bin/svc_ifc_idmgr.bin --x ifc 5148 1.3 0.3 650596 222336 ? Ssl Mar10 580:25
/mgmt//bin/svc_ifc_vtap.bin --x ifc 5160 1.6 0.6 1098280 453492 ? Ssl Mar10 669:17
/mgmt//bin/svc_ifc_scripthandler.bin --x root 7089 1.4 0.4 856360 315016 ? Ssl Mar10 592:04
/mgmt//bin/dhcpd.bin -f -4 -cf /data//dhcp/dhcpd.conf -lf /data//dhcp/dhcpd.lease -pf
/var/run//dhcpd.pid --no-pid bond0.3903 admin 29834 0.0 0.0 112800 1780 pts/1 S+ 17:22 0:00 grep
-E svc|nginx.bin|dhcp ifc 30432 1.4 0.6 894088 405968 ? Ssl Mar17 473:45
/mgmt//bin/svc_ifc_policydist.bin --x root 31215 2.8 5.2 4503880 3397276 ? Ssl Apr05 124:08
```

```
/mgmt//bin/nginx.bin -p /data//nginx/
```

In the above output svc_ifc_dbgr.bin is missing when compared to expected output mentioned in DME process status section. We can start the process by using "acidiag restart dbgr"

```
apicl# acidiag start dbgr apicl# ps -aux | egrep "svc|nginx.bin|dhcp" root 5112 7.3 0.4 1033952
323240 ? Ssl Mar10 3073:43 /mgmt//bin/svc_ifc_appliancedirector.bin --x ifc 5117 1.7 0.6 1062664
439876 ? Ssl Mar10 720:56 /mgmt//bin/svc_ifc_topomgr.bin --x ifc 5118 2.1 2.2 2164512 1468200 ?
Ssl Mar10 884:16 /mgmt//bin/svc_ifc_policymgr.bin --x ifc 5119 1.5 0.3 1115984 256904 ? Ssl
Mar10 664:55 /mgmt//bin/svc_ifc_licensemgr.bin --x ifc 5120 1.5 0.5 1088252 356760 ? Ssl Mar10
666:30 /mgmt//bin/svc_ifc_edmgr.bin --x root 5121 1.6 0.6 1125948 423392 ? Ssl Mar10 698:15
/mgmt//bin/svc_ifc_bootmgr.bin --x ifc 5123 2.3 1.2 1474388 800784 ? Ssl Mar10 994:21
/mgmt//bin/svc_ifc_eventmgr.bin --x ifc 5126 1.5 8.2 6032524 5363184 ? Ssl Mar10 636:01
/mgmt//bin/svc_ifc_reader.bin --x root 5130 4.6 0.6 1092480 439580 ? Ssl Mar10 1927:18
/mgmt//bin/svc_ifc_ae.bin --x ifc 5132 1.6 0.8 1312136 567420 ? Ssl Mar10 689:46
/mgmt//bin/svc_ifc_vmmmgr.bin --x ifc 5133 1.5 0.5 1064176 346760 ? Ssl Mar10 659:07
/mgmt//bin/svc_ifc_domainmgr.bin --x ifc 5135 1.8 1.6 1736876 1099924 ? Ssl Mar10 770:43
/mgmt//bin/svc_ifc_observer.bin --x root 5141 1.5 0.7 1092948 458156 ? Ssl Mar10 663:45
/mgmt//bin/svc_ifc_plgnhandler.bin --x ifc 5146 2.0 0.6 1037676 397236 ? Ssl Mar10 857:48
/mgmt//bin/svc_ifc_idmgr.bin --x ifc 5148 1.3 0.3 650596 222336 ? Ssl Mar10 580:28
/mgmt//bin/svc_ifc_vtap.bin --x ifc 5160 1.6 0.6 1098280 453492 ? Ssl Mar10 669:21
/mgmt//bin/svc_ifc_scripthandler.bin --x root 7089 1.4 0.4 856360 315016 ? Ssl Mar10 592:07
/mgmt//bin/dhcpd.bin -f -4 -cf /data//dhcp/dhcpd.conf -lf /data//dhcp/dhcpd.lease -pf
/var/run//dhcpd.pid --no-pid bond0.3903 ifc 7609 126 0.5 987404 362824 ? Ssl 17:25 0:02
/mgmt//bin/svc_ifc_dbgr.bin --x <===== admin 7762 0.0 0.0 112800 1668 pts/1 S+ 17:26 0:00 grep -
E svc|nginx.bin|dhcp ifc 30432 1.4 0.6 894088 405968 ? Ssl Mar17 473:48
/mgmt//bin/svc_ifc_policydist.bin --x root 31215 2.8 5.2 4503880 3397252 ? Ssl Apr05 124:13
/mgmt//bin/nginx.bin -p /data//nginx/
```

After running "Acidiag start dbgr" the process started again. If you donot see process getting started please reachout to TAC for further troubleshooting.

4. Check Core Files

Run show core, if there are any core files upload them to the SR.

```
apicl# show core Node Module Creation-Time File-Size Service Process Original-Location Exit-Code
Death-Reason Last-Heartbeat -----
----- Ctrlr-Id Creation-Time File-Size Service Process
Original-Location Exit-Code -----
----- 1 2021-10-05T21:19:55.0 204534444 eventmgr 22453
/dmecoresh/svc_ifc_eventmgr.bin_log.22453 134 00-07:00 .tar.gz
```

Refer to the link for Core Colelction <https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/214520-guide-to-collect-tech-support-and-tac-re.html>

3. Collect TechSupport and Upload to SR

Capture APIC TS logs, and upload it to the SR, for further troubleshooting.

<https://www.cisco.com/c/en/us/support/docs/cloud-systems-management/application-policy-infrastructure-controller-apic/214520-guide-to-collect-tech-support-and-tac-re.html>