

Address ACI Fault Codes: F606347, F606350 F606391

Contents

[Introduction](#)

[Background Information](#)

[Fault F606347: Port Group Addition or Deletion on the VM Controller Fails](#)

[Description](#)

[Recommended Action](#)

[Step 1 — Verify APIC-to-vCenter Connectivity](#)

[Step 2 — Verify vCenter Credentials and Privileges](#)

[Step 3 — Verify ACI and vCenter Version Compatibility](#)

[Step 4 — Check the VMM Controller Operational State and Event Log](#)

[Step 5 — Review the Affected EPG and VMM Domain Association](#)

[Step 6 — Collect Diagnostics and Contact TAC if the Fault Persists](#)

[Additional Details](#)

[Fault F606350: LACP Lag Policy Update at the DVS Fails](#)

[Description](#)

[Recommended Action](#)

[Additional Details](#)

[Fault F606391: LLDP/CDP Adjacency Not Found for Physical Adapters](#)

[Description](#)

[Recommended Action](#)

[Step 1 — Validate LLDP/CDP Configuration on the DVS](#)

[Step 2 — Validate LLDP/CDP on the Physical Leaf Switch](#)

[Step 3 — Validate LLDP/CDP on the Physical Switch Connected to the Host](#)

[Step 4 — Verify APIC Adjacency State After Changes](#)

[Additional Details](#)

[Future Prevention](#)

Introduction

This document describes next steps for the remediation of the following Cisco Application Centric Infrastructure (ACI) VMware Virtual Machine Manager (VMM) integration faults: fault F606347 (port group addition or deletion on the VM Controller fails), fault F606350 (LACP Lag Policy update failure at the Distributed Virtual Switch), and fault F606391 (Link Layer Discovery Protocol/Cisco Discovery Protocol adjacency information not found for physical adapters on the host).

Background Information

These faults arise in fabrics that use ACI VMM Domain integration with VMware vCenter and a Distributed Virtual Switch (DVS). ACI continuously synchronizes policies — including port group lifecycle, Link Aggregation Control Protocol (LACP) lag policies, and physical uplink topology — with the DVS through the vCenter API. When that synchronization fails or prerequisite discovery information is missing, ACI raises these faults to surface the condition for operator review.

Fault F606347: Port Group Addition or Deletion on the VM Controller Fails

Description

This fault is raised when ACI fails to add or delete a port group on a VM Controller (for example, VMware vCenter) as part of EPG-to-VMM Domain policy synchronization. When an EPG is associated with or dissociated from a VMM Domain, the APIC instructs the VM Controller to create or remove the corresponding port group on the Distributed Virtual Switch (DVS). If the Finite State Machine (FSM) that manages this operation does not complete successfully, ACI raises fault F606347 on the affected VMM domain controller object.

```
"Code" : "F606347",  
"Description" : "[FSM:FAILED]: Addition or Deletion of Port Group for: (uni/tn-<TENANT>/ap-<APP-PROFILE>  
"Dn" : "uni/vmmp-<VM-Provider>/dom-<VMM-NAME>/ctrlr-[<VMC>]/fault-F606347"
```

Recommended Action

This fault is most commonly caused by communication or compatibility issues between the ACI version and the VM Controller version. Work through the following steps before contacting Cisco Technical Assistance Center (TAC).

Step 1 — Verify APIC-to-vCenter Connectivity

The port group operation is executed over the vCenter API. If the APIC cannot reach the VM Controller, the FSM times out and the fault is raised.

1. In the APIC GUI, navigate to **VM Networking > VMware > [DVS Domain] > Controllers > [vCenter Controller]** and confirm the operational state is **online**.
2. Identify the APIC that is the VMM leader for the domain and verify basic network reachability. From that APIC, ping and attempt an HTTPS connection to the vCenter:

```
<#root>
```

```
apic1#
```

```
show vmware domain name <VMM-NAME> | grep " Leader"
```

```
<VMM-NAME>    apic2    Leader
```

```

apic2#

ping <VC-IP>

PING <VC-IP> (<VC-IP>) 56(84) bytes of data.
64 bytes from <VC-IP>: icmp_seq=1 ttl=63 time=0.312 ms
^C

apic2#

curl -k -X POST -H 'Accept: application/json' --basic \
  -u <USERNAME>@vsphere.local:<PASSWORD> \
  https://<VC-IP>/rest/com/vmware/cis/session

```

A successful HTTPS response confirms that the APIC can authenticate to vCenter. A connection failure or authentication error indicates a network or credential problem that must be resolved before the port group operation can succeed.

Step 2 — Verify vCenter Credentials and Privileges

The vCenter account configured in the VMM Domain must be valid and must have sufficient permissions to create and delete port groups on the DVS.

1. In the APIC GUI, navigate to **VM Networking > VMware > [DVS Domain] > vCenter Credentials** and confirm the username and password are current.
2. Confirm that the vCenter user account has, at minimum, the following privileges on the DVS:
 - DVS: Create, delete, and modify port groups.
 - Network: Assign network policies to port groups.

Refer to the [ACI VMM Troubleshooting guide](#) for the complete list of required vCenter privileges.

Step 3 — Verify ACI and vCenter Version Compatibility

Incompatibilities between the ACI software version and the VM Controller version can cause the port group API call to fail silently or return an unexpected error that the APIC FSM cannot recover from.

1. Confirm that the vCenter version is listed as supported for the ACI release currently running in the fabric. Refer to the [ACI Compatibility Matrix](#) on Cisco.com.
2. If a recent upgrade of either ACI or vCenter preceded the onset of this fault, consult the ACI release notes for the upgraded version to identify known VMM integration issues or required minimum vCenter versions.
3. If the vCenter version is not compatible, upgrade vCenter (or ACI) to a supported combination. Refer to the [ACI VMM Troubleshooting guide](#) for version-specific known issues.

Step 4 — Check the VMM Controller Operational State and Event Log

1. In the APIC GUI, navigate to **VM Networking > VMware > [DVS Domain] > Controllers > [vCenter Controller]** and open the **Operational** tab. Review the **Events** and **Faults** sub-tabs for

concurrent VMM connectivity faults (for example, F606225 or F606327). If broader connectivity faults are present, resolve them first.

2. You can also query the fault directly via the APIC REST API to review the full fault description and the specific error text from the FSM:

```
<#root>
```

```
apic#
```

```
moquery -c faultInst -x 'query-target-filter=eq(faultInst.code,"F606347")'
```

The **description** field in the output contains the FSM error detail, including the VM Controller name, VM Domain, VM Provider, and the EPG that triggered the operation. Use this information to narrow the scope of the investigation to the specific EPG and VMM Domain involved.

Step 5 — Review the Affected EPG and VMM Domain Association

1. Identify the EPG named in the fault description (uni/tn-<TENANT>/ap-<APP-PROFILE>/epg-<EPG>).
2. In the APIC GUI, navigate to **Tenants > [Tenant] > Application Profiles > [App Profile] > Application EPGs > [EPG] > Domains** and confirm that the VMM Domain association exists and is in the correct state.
3. If the port group operation was triggered by an accidental configuration change, verify whether the EPG-to-VMM Domain association should exist. Removing and re-adding the association can reset the FSM and clear the fault if the underlying infrastructure issue has been resolved.

Step 6 — Collect Diagnostics and Contact TAC if the Fault Persists

If the fault does not clear after completing the steps above, collect the following information and open a case with Cisco TAC:


- APIC tech-support bundle: Navigate to **System > Troubleshooting > Tech Support** in the APIC GUI to generate and download the bundle.
- The full fault DN and description text from the `moquery` output in Step 4.
- The ACI software version (from **System > Controllers > [APIC] > Summary**) and the vCenter version.
- The timeframe of the first occurrence and whether the fault appeared after an upgrade or configuration change.

Additional Details

When an EPG is associated with a VMM Domain, ACI programs a corresponding port group on the DVS through the vCenter API. The Finite State Machine (FSM) task `CompEppDAddorDelExtPol` manages this lifecycle operation. The FSM attempts the port group add or delete and transitions through a set of states. If any state transition fails — for example, due to an API error returned by vCenter, a timeout, or an authentication failure — the FSM is marked as **FAILED** and fault F606347 is raised on the `vmmCtrlr` object for the affected VM Controller.

Common failure scenarios include:

- **ACI-to-vCenter version incompatibility** — an ACI or vCenter upgrade introduces a change in API behaviour that causes the port group operation to fail. This is one of the most common root causes and is addressed by aligning both products to a compatible version combination. Refer to the [ACI Virtualization Matrix](#) for details.
- **vCenter API timeout or transient error** — an overloaded or temporarily unavailable vCenter returns an error or does not respond within the FSM timeout. The operation is not automatically retried in all code paths; removing and re-adding the EPG-to-VMM Domain association manually triggers a fresh FSM run.
- **Insufficient vCenter privileges** — the vCenter service account does not have permission to create or delete port groups, causing the API call to return an authorization error.
- **Port group naming conflict** — a manually created port group with the same name as the one ACI is attempting to create already exists on the DVS, causing the operation to fail. Remove the conflicting port group or rename it before re-attempting the association.

 **Note:** Because the FSM state is preserved until the association is removed or a new trigger arrives, the fault may persist even after the underlying network or credential issue is resolved. If the fault remains after fixing the root cause, remove and re-add the EPG-to-VMM Domain association to force a new FSM execution.

Fault F606350: LACP Lag Policy Update at the DVS Fails

Description

This fault is raised when ACI attempts to update the LACP lag policy on the DVS through the vCenter API and the operation fails. ACI pushes LACP configuration to the DVS as part of VMM Domain policy synchronization, specifically when a LACP policy is associated with a VMM domain attached to the DVS. When the update cannot be applied, ACI raises fault F606350 on the affected leaf node.

```
"Code" : "F606350",  
"Description" : "Updating LACP Lag Policy at DVS failed.",  
"Dn" : "topology/pod-<podId>/node-<leafNodeId>/local/svc-policy/lem-id-0/uni/epp/fv-[uni/vmmp-VMware/do
```

Recommended Action

This task is automatically retried by ACI. A transient vCenter API delay or a momentary connectivity interruption between the APIC and vCenter can cause a single instance of this fault. In many cases the retry succeeds and the fault clears on its own.

If you observe repeated or persistent failures, take the following steps before contacting Cisco Technical Assistance Center (TAC):

1. Verify that the APIC can reach the vCenter server over the network. Navigate to **VM Networking > VMware > [DVS Domain] > Controllers > [DVS Controller]** in the Application Policy Infrastructure Controller (APIC) GUI and confirm the operational state is **online**.
2. Confirm that the vCenter credentials configured in the VMM Domain are valid and have not expired. Navigate to **VM Networking > VMware > [DVS Domain] > vCenter Credentials** and verify the username and password are correct.
3. Confirm that the vCenter user account associated with the VMM Domain has the required privileges. At minimum, the account must have DVS configuration and host network management permissions. Refer to the Cisco ACI VMware vSphere Integration Guide (available on Cisco.com) or the [ACI VMM Troubleshooting guide](#) for the complete list of required vCenter privileges.
4. Review the APIC system faults and event log for concurrent VMM connectivity faults (for example, F606225 or F606327) that indicate a broader vCenter API communication problem. If such faults are present, resolve the connectivity issue first.

1. You can use the following commands to confirm the apic Leader and from there test connectivity by nslookup if necessary, ping and HTTPS .

```

apic1# show vmware domain name shared-dvs | grep " Leader"
shared-vc      apic2      Leader
apic2# nslookup <VC-FQDN>
apic2# ping <VC-IP>
PING <VC-IP> (<VC-IP>) 56(84) bytes of data.
64 bytes from <VC-IP>: icmp_seq=1 ttl=63 time=0.237 ms
64 bytes from <VC-IP>: icmp_seq=2 ttl=63 time=0.406 ms
^C
--- <VC-IP> ping statistics ---
2 packets transmitted, 2 received, 0% packet loss, time 1000ms
rtt min/avg/max/mdev = 0.237/0.321/0.406/0.084 ms

apic2# curl -k -X POST -H 'Accept: application/json' --basic -u <USERNAME>@vsphere.local:<PA
  % Total    % Received % Xferd  Average Speed   Time    Time     Time  Current
                                 Dload  Upload  Total   Spent    Left  Speed
   0     0    0     0    0     0    0     0     0  0  0  0  0
100  408    0  408    0     0 1393     0  0  0  0  0 1397


```

5. Check the LACP policy attached to the VMM Domain Interface Policy Group. Navigate to **Fabric > Access Policies > Policies > Interface > Port Channel** and confirm that the LACP policy mode is compatible with the DVS uplink port group configuration in vCenter, refer to the "Teaming and ACI vSwitch Policy" section of the [ACI VMM Troubleshooting guide](#) to see the compatible combinations.
6. If the fault persists after verifying all of the above, collect the APIC tech-support file and contact Cisco TAC.
 - Navigate to **System > Troubleshooting > Tech Support** in the APIC GUI in order to generate and download the tech-support bundle.
 - Include the fault DN from the fault details and the timeframe of the recurring failures in the TAC case.

Additional Details

ACI VMM integration uses the vCenter API to program DVS configuration on behalf of the fabric. When a LACP Policy is associated with a VMM Domain Interface Policy Group (**infraAccPortGrp**), ACI translates the policy into a DVS LACP group configuration and pushes it to vCenter. The push operation can fail for several reasons:

- **vCenter API timeout** — a slow or overloaded vCenter may not respond within the APIC's timeout window. The operation is retried automatically.
- **Insufficient privileges** — the vCenter service account configured in the VMM Domain does not have the permissions required to modify DVS uplink port group properties.
- **DVS version incompatibility** — the DVS version in vCenter does not support the LACP configuration being pushed. ACI requires DVS version 5.1 or later for LACP support.
- **LACP policy conflict** — an existing manual LACP configuration on the DVS uplink port group conflicts with the policy ACI is attempting to apply.

 **Note:** A single isolated instance of F606350 that clears after a retry does not indicate a persistent problem. Investigate only when the fault recurs repeatedly within a short window or does not clear within a few minutes.

Fault F606391: LLDP/CDP Adjacency Not Found for Physical Adapters

Description

This fault is raised when ACI cannot find Link Layer Discovery Protocol (LLDP) or Cisco Discovery Protocol (CDP) adjacency information for the physical network adapters (vmnics) on a host managed by a VMM Domain. ACI uses LLDP or CDP to discover which leaf switch port is physically connected to each vmnic on the host. Without this adjacency information, ACI cannot correctly map VM traffic from the DVS to the corresponding leaf port, which affects policy deployment and endpoint learning for virtual machines on that host.

```
"Code" : "F606391",  
"Description" : "LLDP/CDP Adjacency information not found for physical adapters on the host.",  
"Dn" : "topology/pod-<podId>/node-<leafNodeId>/local/svc-policy/lem-id-0/uni/epp/fv-[uni/vmmp-VMware/do
```

Recommended Action


This fault requires manual validation of LLDP or CDP configuration at three points in the path: the DVS in vCenter, the ESXi host, and the physical leaf switch. Work through the following steps in order.

Step 1 — Validate LLDP/CDP Configuration on the DVS

The DVS Discovery Protocol setting controls whether the DVS advertises and listens for LLDP or CDP frames, these protocols are mutually exclusive as mentioned in the [ACI VMM Troubleshooting guide](#). If this setting is disabled or set to **Advertise Only**, the APIC cannot read adjacency information from vCenter.

1. Log in to the vSphere Client and navigate to **Home > Networking > [DVS Name] > Configure > Settings > Properties**.

2. Locate the **Advanced** section and check the **Discovery Protocol** fields:
 - **Type** — set to **Link Layer Discovery Protocol** (recommended for ACI) or **Cisco Discovery Protocol**, depending on your environment.
 - **Operation** — must be set to **Both** or **Listen**. A setting of **Advertise** or **Disabled** prevents the DVS from receiving neighbor information, which means vCenter has no adjacency data to report to the APIC.
3. If the operation is set to **Advertise** or **Disabled**, change it to **Both** and save the settings. Allow several minutes for the APIC to re-query vCenter for updated adjacency data.

 **Note:** Changing the DVS Discovery Protocol setting is non-disruptive to VM traffic. It only affects the control-plane discovery information exchanged between the DVS and connected switches.

Step 2 — Validate LLDP/CDP on the Physical Leaf Switch

The leaf switch interface connected to the host (or to the upstream access switch that the host connects through) must have LLDP or CDP enabled. In ACI, LLDP and CDP are controlled by Interface Policies applied to the Interface Policy Group used on the relevant port.

1. Identify the leaf port connected to the host. Navigate to **Fabric > Inventory > [Pod] > [Leaf Node] > Interfaces > Physical Interfaces** and locate the interface carrying the host's vmmic traffic.
2. Navigate to **Fabric > Access Policies > Interfaces > Leaf Interfaces > Policy Groups** and open the Interface Policy Group applied to that port.
3. Confirm that an **LLDP Interface Policy** is attached to the Policy Group with **Receive State: Enabled** and **Transmit State: Enabled**. If no LLDP policy is attached, the default policy is used, which has both states enabled.
4. If you are using CDP, confirm that a **CDP Interface Policy** is attached with **Admin State: Enabled**.
5. To confirm that the leaf is receiving LLDP neighbors on the expected interface, SSH to the leaf and run the following command:

```
<#root>
```

```
leaf101#
```

```
show lldp neighbors
```

The output lists each interface and its discovered neighbor. The host's vmmic or the upstream access switch must appear in the neighbor table for the expected interface. If the interface is missing from the output, the leaf is not receiving LLDP frames on that port, which indicates that LLDP is blocked upstream or disabled on the connected device.

6. If CDP is in use, run the following command in order to verify CDP neighbor discovery:

```
<#root>
```

```
leaf101#
```


```
show cdp neighbors
```

The host or upstream switch must appear in the output for the expected interface.

Step 3 — Validate LLDP/CDP on the Physical Switch Connected to the Host

If the host vmnics connect to an intermediate physical access switch (not directly to the ACI leaf), LLDP or CDP frames must be forwarded through that switch to reach the leaf. Verify the following on the intermediate switch:

- LLDP or CDP is enabled globally on the switch.
- LLDP or CDP is enabled on the interfaces facing both the host and the ACI leaf.
- The switch is not configured to filter or block LLDP/CDP Protocol Data Units (PDUs) on the relevant interfaces (for example, through a service policy or access control list).

 **Note:** LLDP is a link-local protocol. Standard Layer 2 switches forward LLDP PDUs transparently between ports in the same VLAN only when LLDP is not terminated on the switch itself. If the intermediate switch terminates LLDP, it becomes the LLDP neighbor for the leaf — not the host. In that case, ACI sees the intermediate switch as the neighbor, which means it cannot identify the host's vmnics. Either enable LLDP pass-through on the intermediate switch or connect the host directly to the ACI leaf.

Step 4 — Verify APIC Adjacency State After Changes

After making configuration changes, verify that the APIC can now resolve the host's physical uplink topology. In the APIC GUI, navigate to **VM Networking > VMware > [DVS Domain] > [DVS Name] > Hosts > [Host Name] > Physical Interfaces** and confirm that the **Discovered** field shows a leaf port for each vmnics. If the adjacency is correctly resolved, the fault clears automatically.

You can also query the APIC REST API in order to check the adjacency objects for a specific VMM domain:

```
<#root>
```

```
apic#
```

```
moquery -c compHv -x 'query-target-filter=eq(compHv.name,"hostname")'
```

The **compHv** object represents a hypervisor host within the VMM domain. Related **compNic** objects represent the physical adapters. When adjacency is resolved, the **peerDn** attribute of the **compNic** objects is populated with the DN of the corresponding leaf interface.

If the fault does not clear after validating all three configuration points above, collect the APIC tech-support file and contact Cisco TAC.

Additional Details

ACI VMM integration uses the vCenter API to retrieve LLDP and CDP neighbor data that vCenter collects from the DVS. The APIC reads this data in order to build a map of which host vmnics connects to which leaf

port. This mapping is used to:

- Program the correct leaf interface policies for VM traffic leaving a given host. When the Resolution Immediacy is configured as immediate or on demand, if host and leaf lose LLDP/CDP neighborship the policies are removed.
- Enforce microsegmentation and EPG membership for virtual endpoints based on their physical attachment point.
- Support ACI Virtual Edge (AVE) policy enforcement, which requires accurate knowledge of the host's physical uplink topology.

When adjacency information is missing, ACI raises fault F606391 to signal that it cannot validate the physical topology for the affected host. Virtual machine connectivity may still function in the interim — the fault does not immediately interrupt data forwarding — but policy deployment accuracy and endpoint learning reliability are degraded.

Future Prevention

To prevent fault F606391 from recurring after it has been resolved:

- Set the DVS Discovery Protocol Operation to **Both** as a standard build requirement for all DVS instances associated with ACI VMM Domains.
- Include LLDP and CDP enablement as part of the standard Interface Policy Group template applied to all leaf ports that connect to hosts running VMware ESXi.
- If using an intermediate access switch between the host and the ACI leaf, confirm that the switch vendor's LLDP forwarding behaviour is compatible with the ACI VMM discovery mechanism before deployment.