Cisco UCS Director Release Notes for Cisco ACI APIC Connector Pack, Release 6.7.x.x

First Published: 2019-07-04

Last Modified: 2020-11-20

Release Notes for Cisco ACI APIC Connector Pack

Cisco UCS Director

Cisco UCS Director delivers unified, highly secure management for supported compute, network, storage, and virtualization platforms and for the industry's leading converged infrastructure solutions, which are based on the Cisco Unified Computing System (Cisco UCS) and Cisco Nexus platforms. Cisco UCS Director extends the unification of computing and network layers through Cisco UCS to provide data center administrators with comprehensive visibility and management capabilities for compute, network, storage, and virtualization. For more information, see Cisco UCS Director on Cisco.com.

Revision History

Release	Date	Description
6.7.2.1	July 4, 2019	Created for Release 6.7.2.1.
6.7.3.1	November 07, 2019	• Support for PC/vPC Leaf Policy • Support for Route Tag Policy • Support for Creating VRF in APIC • Support for Access Port Selector • Support for VMM Domain • Enhancements to add Domain to an EPG
6.7.3.2	December 12, 2019	• Support for newer Versions of Cisco Application Centric Infrastructure Controllers and Cisco Application Centric Infrastructure Multi-Site Controllers

Release	Date	Description
6.7.4.1	June 18, 2020	• Introduction of Interface Policies • Enhancements to First-Hop Security (FHS) • Introduction of Port Policies • Enhancements to MACsec • Support for Virtual Switched Port Analyzer (VSPAN) • Support for Flow Record • Introduction of NetFlow Monitor Policy • Enhancements to Enhanced Interior Gateway Routing Protocol (EIGRP) • Support for APIC Monitoring Policy • Introduction of BGP Timers Policy • Support for Snoop Policy • Support for DHCP Policy • Enhancements to VRF • Support for Flow Control Policy • Support for Slow Control Policy • Support for Data Plane Policing • Enhancements to the Interface Policy
6.7.4.2	November 20, 2020	• Support for newer Versions of Cisco Application Centric Infrastructure Controllers and Cisco Application Centric Infrastructure Multi-Site Controllers

Connector Packs

Connector packs help you perform connector level upgrade in Cisco UCS Director without impacting other connectors and without having to upgrade the entire software version. After claiming Cisco UCS Director in Cisco Intersight, as a system administrator, you can view information on new versions of connector packs that are available for upgrade. The top header pane of the user interface displays a Download icon indicating that new connector pack versions are available. You can select and upgrade the connector packs in Cisco UCS Director.

Cisco Application Centric Infrastructure APIC Connector Pack

Cisco Application Centric Infrastructure (ACI) Cisco Application Policy Infrastructure Controller (APIC) connector pack provides support for new versions of Cisco ACI APIC. Using this connector pack, you can upgrade to the latest version of Cisco ACI APIC. During the upgrade process, Cisco UCS Director is restarted automatically.

Upgrading Connector Packs

Before you begin

- You must have system administrator privileges in Cisco UCS Director.
- Cisco UCS Director has been claimed in Cisco Intersight. For information on claiming a device, see the integrated guided walkthrough titled *Learn How to Claim a Device* available within the **Online Help** menu in the Cisco Intersight user interface.
- Cisco UCS Director is successfully connected to Cisco Intersight.
- Take a snapshot of Cisco UCS Director before you initiate the upgrade.

Procedure

Step 1 On the header, click **New Upgrades Available**.

The **Available System Upgrades** screen appears and will display all available connector packs for upgrade along with version information. Upon login, if you clicked **Yes** to the pop-up message, then the very same upgrade screen appears.

Note The **New Upgrades Available** icon is visible on the header only when new versions of the current running connector packs are available for upgrade.

Step 2 Check the check box of a connector pack from the list.

You can check the check boxes of multiple connector packs.

Step 3 Click **Upgrade**.

Step 4 In the **Confirm Upgrade** dialog box, click **Yes**.

After you confirm that the connector version must be upgraded, the validation process is initiated. If the validation process completes successfully, the upgrade process is initiated and the **System Upgrade Status** screen displays the upgrade status. After the upgrade process is successful, the **Logout** option is enabled.

Step 5 Click **Logout**.

While upgrading a base platform pack that includes changes to all infrastructure components, all Cisco UCS Director services are restarted. As a result, after clicking **Logout**, the screen could appear to be unresponsive for a few minutes. After all the services are restarted, and the upgrade process is complete, you can login to Cisco UCS Director .

What to do next

You can view the upgrade reports by choosing **Administration** > **System** > **System Updates**. From this screen, you can double-click on a report, and view additional details on the upgrade process. For more information, see Viewing Connector Pack Upgrade Information.

Viewing Connector Pack Upgrade Information

Procedure

Step 1 Choose **Administration** > **System**.

Step 2 On the **System** page, click **System Updates**.
Information such as upgrade request ID, user that initiated the upgrade, upgrade start time and end time, and the upgrade status are displayed.

Step 3 Select a connector pack and choose **View Details** to view details such as connector pack name, upgraded version, and prior version.

Step 4 Click **State History** to view the various states of the connector pack upgrade process. For example, upgrade request received, upgrade process initiated or upgrade process completed.

Step 5 Click **Stages** to view the entire lifecycle of the connector pack upgrade request.

New and Changed Features

This section provides an overview of the significant new and changed features in this release. This section does not provide an exhaustive list of all enhancements included in this release.

New and Changed Features in Release 6.7.2.1

Support for Netflow Monitor Policy and End Point Retention Policy

This connector pack release introduces:

- Logical NetFlow monitoring policy—You can create a logical NetFlow monitoring policy to associate a flow record with the monitoring policy. The monitor policy identifies packet flows for ingress IP packets and provides statistics based on these packet flows.
- Endpoint retention policy—You can create an endpoint retention policy to set the hold interval, bounce entry aging interval, local endpoint aging interval, remote endpoint aging interval, and move frequency

for endpoints. You can associate an endpoint retention policy to a bridge domain during bridge domain creation.

Enhancements to Bidirectional Forwarding Detection

You can use Bidirectional Forwarding Detection (BFD) to detect the number of times a sub-second failure occurs in the forwarding path between the ACI fabric border leaf switches configured to support peering router connections. This connector pack release introduces the following policy and profile to support BFD:

- BFD Interface Policy—In Cisco UCS Director, you can enable or disable the admin, control, and echo admin state for the BFD interface policy. You can set the detection multiplier, minimum transmit interval, minimum receive interval, and echo receive interval for the BFD interface policy.
- BFD Interface Profile—In the profile, you can set if authentication is required or not for accessing the profile and choose a BFD interface policy for the profile.

Enhancements to Hot Standby Router Protocol

This connector pack introduces the following policies and profiles to support Hot Standby Router Protocol (HSRP) in Cisco APIC account. The HSRP protocol acts as the gateway for the endpoints behind the Layer 2 switches.

- HSRP Interface Policy—While defining the HSRP interface policy, you can enable or disable BFD and control sourcing of hellos from the burned-in MAC address (BIA) to identify devices. You can set the minimum time to delay HSRP group initialization after an interface comes up, and set the time period to delay HSRP group initialization after the router has reloaded.
- HSRP Group Policy—In the HSRP group policy, you can enable or disable preemption for a group, set the priority for HSRP to define the active and standby routers, and choose MD5 or simple authentication method. Also, you can set the hello interval, hold interval, minimum preemption delay, the delay time for resuming the preemptive action after reloading the active HSRP leaf, the maximum amount of time allowed for the HSRP client to prevent preemption, and the timeout value for authentication.
- HSRP Interface Profile—In Cisco UCS Director, you can choose a specific version of HSRP interface policy for the logical interface profile.
- HSRP Interface Group to Interface Profile—You can define the ID, name, and type of the HSRP interface group along with the mode in which the IP address can be obtained for the group. You can provide a virtual Media Access Control (MAC) address and an IP address that is shared among a group of configured routers. You can also provide comma separated multiple IPv4 or IPv6 addresses that can be used as secondary virtual IP addresses.

Support for Dynamic Host Configuration Protocol

This connector pack release provides support for creating a DHCP relay policy with a unique name. A DHCP relay policy is used to dynamically assign IP address when the DHCP client and server are in different subnets. The DHCP relay profile contains one or more providers. Starting with this release, you can add providers to the DHCP relay policy at the tenant level and at the infrastructure level.

You can add a DHCP relay label to logical interface profile by assigning a DHCP policy to an infra or a tenant.

CoPP Leaf/Spine Policy Support for APIC Fabric Switch

Starting with this release, you can create and manage the APIC control plane policing (CoPP) leaf and spine policy to be applied on Cisco ACI leaf/spine switches. While creating a CoPP leaf/spine policy, choose **CoPP has custom values** as the type of profile if you wish to set policy for each protocol separately.

Enhancements to Routed Interface

From this release, you can configure the secondary IP address and BGP peer connectivity profile for the routed interface, routed sub-interface interface, and SVI interface.

Introduction of New REST APIs

This release introduces REST APIs for the following features:

- NetFlow Monitor Policy
  - CREATE_APIC_NETFLOW_MONITOR_POLICY
  - UPDATE_APIC_NETFLOW_MONITOR_POLICY
  - DELETE_APIC_NETFLOW_MONITOR_POLICY
- End Point Retention Policy
  - ADD_APIC_TENANT_BRIDGE_DOMAIN
  - MODIFY_APIC_TENANT_BRIDGE_DOMAIN
  - ADD_APIC_ENDPOINT_RETENTION_POLICY
  - UPDATE_APIC_ENDPOINT_RETENTION_POLICY
  - REMOVE_APIC_ENDPOINT_RETENTION_POLICY
- BFD Interface Policy
  - CREATE_APIC_BFD_INTERFACE_POLICY
  - DELETE_APIC_ BFD_INTERFACE _POLICY
  - UPDATE_APIC_ BFD_INTERFACE _POLICY
- BFD Interface Profile
  - CREATE_APIC_BFD_INTERFACE_PROFILE
  - DELETE_APIC_ BFD_INTERFACE _PROFILE
  - UPDATE_APIC_ BFD_INTERFACE _PROFILE
- HSRP Interface Policy
  - CREATE_APIC_HSRP_INTERFACE_POLICY
  - DELETE_APIC_HSRP_INTERFACE_POLICY
  - UPDATE_APIC_HSRP_INTERFACE_POLICY

- HSRP Group Policy
  - CREATE_HSRP_GROUP_POLICY
  - DELETE_HSRP_GROUP_POLICY
  - UPDATE_HSRP_GROUP_POLICY
- HSRP Interface Profile
  - CREATE_APIC_HSRP_INTERFACE_PROFILE
  - DELETE_APIC_HSRP_INTERFACE _PROFILE
  - UPDATE_APIC_HSRP_INTERFACE _PROFILE
- HSRP Interface Group to Interface Profile
  - CREATE_APIC_HSRP_INTERFACE_GROUP
  - UPDATE_APIC_HSRP_INTERFACE_GROUP
  - DELETE_APIC_HSRP_INTERFACE_GROUP
- DHCP Relay Policy
  - CREATE_DHCP_RELAY_POLICY
  - DELETE_DHCP_RELAY_POLICY
- DHCP Relay Policy Providers
  - ADD_PROVIDERS_TO_DHCP_RELAY_POLICY
  - DELETE_PROVIDERS_FROM_DHCP_RELAY_POLICY
- DHCP Relay Label
  - ADD_APIC_TENANT_DHCP_RELAY_LABEL_TO_INTERFACE_PROFILE
  - DELETE_APIC_TENANT_DHCP_RELAY_LABEL_FROM_INTERFACE_PROFILE
- CoPP Leaf Policy
  - APIC CoPP Leaf Policy
    - CREATE_APIC_FABRIC_COPP_LEAF_POLICY
    - DELETE_APIC_FABRIC_COPP_ LEAF_POLICY
    - UPDATE_APIC_FABRIC_COPP_ LEAF_POLICY
  - Associate Custom Values to Copp Leaf Policy
    - ASSOCIATE_CUSTOM_VALUES_TO_APIC_COPP_LEAF_POLICY
  - CoPP Spine Policy
    - CREATE_APIC_FABRIC_COPP_SPINE_POLICY

- DELETE_APIC_FABRIC_COPP_SPINE_POLICY
- UPDATE_APIC_FABRIC_COPP_SPINE_POLICY

- Routed Interface
  - APIC_TENANT_INTERFACE_TO_LOGICAL_INTERFACE_PROFILE_CREATE
  - APIC_TENANT_INTERFACE_TO_LOGICAL_INTERFACE_PROFILE_DELETE
- Routed Interface—BGP Peer Connectivity
  - ADD_BGP_PEER_CONNECTIVITY_PROFILE_TO_INTERFACE_OF_LOGICAL_INTERFACE_PROFILE
  - REMOVE_BGP_PEER_CONNECTIVITY_PROFILE_FROM_INTERFACE_OF_LOGICAL_INTERFACE_PROFILE
  - UPDATE_BGP_PEER_CONNECTIVITY_PROFILE_FROM_INTERFACE_OF_LOGICAL_INTERFACE_PROFILE
  - Routed Interface—Secondary IP Address
    - Add_APIC_TENANT_SECONDARY_ADDRESS_TO_INTERFACE
    - DELETE_APIC_TENANT_SECONDARY_ADDRESS_TO_INTERFACE

New and Changed Features in Release 6.7.3.1

Support for PC/vPC Leaf Policy

With this connector pack release, you can create a port channel interface leaf policy and virtual port channel leaf policy. The port channel (PC) and virtual port channel (vPC) leaf policy is a template to dictate port behavior and is associated to an Access Entity Profiles (AEP). After creating the policy, you can associate the following to it :

- Netflow monitor policy
- Virtual destination groups
- Virtual source groups
- Override policy group

Support for Route Tag Policy

This connector pack release provides support for creating a route tag policy with a tag value which is used to prevent routing loops.

When a transit route is redistributed into Open Shortest Path First (OSPF) or Enhanced Interior Gateway Routing Protocol (EIGRP), the route is tagged with the tag value specified in the route tag policy to prevent routing loops. If a route is received on an OSPF or EIGRP L3Out with this tag value, the route is dropped.

Support for Creating VRF in APIC

Starting with this release, you can define IPv4 unicast address family or IPv6 unicast address family as the EIGRP address family type, to configure an EIGRP routing instance. You can then add an APIC EIGRP to a Virtual routing and forwarding (VRF) object (called as private network in the APIC GUI).

Support for Access Port Selector

From this release, you can add an access port selector to a fabric interface profile. In Cisco UCS Director, choose **Physical** > **Network** > **Multi-Domain Managers** > **<APIC Account>** and click the **Fabric Interface Profiles** tab. Choose a profile and click **View Details** to view the access port selectors of the fabric interface profile. To add an access port selector to the fabric interface profile, click **Add** and complete the fields in the **Create Access Port Selector** screen.

Support for VMM Domain

This connector pack release provides support for creating a virtual machine manager (VMM) domain to integrate APIC with a third-party VMM (for example, VMware vCenter) to extend the benefits of ACI to the virtualized infrastructure. You can create VMM domains with one of the following virtual switches:

- VMware vSphere distributed switch (VDS)
- Cisco AVS
- Cisco AVE

Enhancements to add Domain to an EPG

Starting with this release, you can configure a default port binding type for all new vEthernet port profiles. For VMM type domain profiles, you can choose either **AVE** or **native** as the switching mode. You can choose **auto**, **VLAN**, or **VXLAN** as the encapsulation node for VMM type domain profiles.

Introduction of New REST APIs

This release introduces REST APIs for the following features:

- Route Tag Policy
  - CREATE_APIC_ROUTE_TAG_POLICY
  - DELETE_APIC ROUTE_TAG _POLICY
  - UPDATE_APIC ROUTE_TAG _POLICY
- Netflow Monitor Policy
  - CREATE_APIC_NETFLOW_MONITOR_POLICY_TO_PC_VPC_INTERFACE_POLICY_GROUP
  - DELETE_APIC_NETFLOW_MONITOR_POLICY_FROM_PC_VPC_INTERFACE_POLICY_GROUP
- Override Policy Group
  - CREATE_OVERRIDE_POLICY_GROUP_TO_INTERFACE_POLICY_GROUP
  - UPDATE_OVERRIDE_POLICY_GROUP_TO_INTERFACE_POLICY_GROUP
  - DELETE_OVERRIDE_POLICY_GROUP_FROM_INTERFACE_POLICY_GROUP

- VDestination Groups
  - ADD_FABRIC_VDESTINATION_GROUP_TO_APIC_PC_VPC_INTERFACE_POLICY
  - REMOVE_FABRIC_VDESTINATION_GROUP_FROM_APIC_PC_VPC_INTERFACE_POLICY
- VSource Groups
  - ADD_FABRIC_VSOURCE_GROUP_TO_APIC_PC_VPC_INTERFACE_POLICY
  - REMOVE_FABRIC_ VSOURCE _GROUP_FROM_APIC_PC_VPC_INTERFACE_POLICY

New and Changed Features in Release 6.7.3.2

Support for newer Versions of Cisco Application Centric Infrastructure Controllers and Cisco Application Centric Infrastructure Multi-Site Controllers

This release introduces support for the following:

- Cisco Application Centric Infrastructure Controllers 4.1(x) and 4.2(x)
- Cisco Application Centric Infrastructure Multi-Site Controllers 2.1(x) and 2.2(x)

For more details, refer the Cisco UCS Director Compatibility Matrix, Release 6.7.

New and Changed Features in Release 6.7.4.1

Introduction of Interface Policies

This connector pack release introduces the following interface policies to manage an APIC account in Cisco UCS Director:

- Spanning Tree Interface Policy—You can create a spanning tree interface policy which dictates the behavior of southbound leaf port Spanning Tree features. To create this policy, click **Add** in the **Spanning Tree Interface Policy** tab of an APIC account and enter a name and short description. While configuring this policy, you can also choose to enable a Bridge Protocol Data Unit (BPDU) filter or a BPDU guard, or both.
- Fibre Channel Interface Policy—You can create a fibre channel interface policy by defining the trunk mode, port mode, speed, auto max speed, fill pattern, and receive buffer credit for the fibre channel interface.
- Layer 2 Interface Policy—You can configure a Layer 2 interface policy to allow a per port-VLAN. A per-port VLAN allows the same VLAN to be used for different endpoint groups provided that the in-bound traffic is on a different port. While creating a Layer 2 interface policy, you can choose one of the following options to enable encapsulation: **corePort**, **disabled**, **doubleQtagPort**, and **edgePort**. You can also enable reflective relay on a port, port channel, or virtual port channel as a Layer 2 interface policy on the switch. Reflective relay allows inbound traffic to return on the same port.

Enhancements to First-Hop Security (FHS)

Starting with this connector pack release, you can associate an FHS policy to a tenant while adding a bridge domain to a Virtual routing and forwarding (VRF). Using the FHS policy, you can control address assignment

and derived operations, such as Duplicate Address Detection (DAD) and Address Resolution (AR) in a service provider environment.

Introduction of Port Policies

This connector pack release introduces the following port policies to manage an APIC account in Cisco UCS Director:

- Port Security Policy—You can configure a port security policy to protect ACI fabric resources from being flooded with unknown MAC addresses by limiting the number of MAC addresses learned per port. The violation action of this policy is always set to the protect mode.

  In the protect mode, MAC learning is disabled and MAC addresses are not added to the Content Addressable Memory (CAM) table. MAC learning is re-enabled after the configured timeout value.

- Port Channel Member Policy—You can create a port channel member policy to set a specific transmit rate and a specific priority for an ACI fabric resource. You can choose between Normal and Fast as the transmit rate, and set a priority with a value between 1 and 65535.

Enhancements to MACsec

Starting with this connector pack release, you can define and associate a MACsec KeyChain policy and a MACsec access parameters policy to a MACsec interface policy. A MACsec KeyChain policy consists of configuration details specific to keychain definition, while the MACsec access parameters policy consists of information specific to the MACsec configuration.

In a MACsec Key policy, you can define the key name, provide a pre-shared key for either 128-bit cipher suites or 256-bit cipher suites, and set a key validity period. In the MACsec access parameters policy, you can define the MACsec functionality such as key server priority for the MACsec Key Agreement (MKA) server selection, expiry time for Secure Association Key (SAK), and set a security policy as **Must secure mode** or **Should secure mode**. In the **Must secure mode**, only encrypted traffic is allowed on the link while in the **Should secure mode** both clear and encrypted traffic are allowed on the link.

Support for Virtual Switched Port Analyzer (VSPAN)

Starting with this connector pack release, you can start or stop VSPAN sessions on demand to copy relevant traffic from a virtual switch to a destination group. The destination group can be local ports or remote devices. While creating a VSPAN session in Cisco UCS Director, you have a provision to start or stop the session and to associate a specific destination group with the session.

In a VSPAN session, you can add an EPG or a client endpoint device as a source, set the traffic capture direction as incoming, outgoing or both, and select a specific tenant. Then, you can associate a source path to the VSPAN source by choosing a source type and a static path.

Support for Flow Record

Starting with this connector pack release, Cisco UCS Director has the provision to define NetFlow record at fabric and tenant levels. A NetFlow record lets you define a flow and the statistics to be collected for each flow. You can define match parameters to identify packets in the flow and define collect parameters that the NetFlow gathers from the flow.

Introduction of NetFlow Monitor Policy

Starting with this connector pack release, you can create a NetFlow monitor policy at a fabric resource level and associate it with a flow record. A NetFlow record lets you define a flow and the statistics to be collected for each flow.

You can create an external collector reachability (also known as NetFlow exporter) by clicking the **Add** action available under the **NetFlow Exporters** tab of an APIC account. Then, you can associate a NetFlow exporter with a fabric NetFlow monitor policy, to specify the destination for the data collected for the flow.

You can deploy the NetFlow monitor policy on an existing bridge domain by associating the NetFlow monitor policy with a bridge domain.

Enhancements to Enhanced Interior Gateway Routing Protocol (EIGRP)

EIGRP Address Family Context Policy (eigrpCtxAfPol) contains the configuration for a specific address family in a given VRF. Starting with this connector pack release, you can configure an eigrpCtxAfPol policy within multiple tenant protocol policies and apply the policy to one or more VRFs within the tenant.

You can enable this policy on a VRF with a relation in the VRF-per-address family. If there is no relation to a given address family or the specified eigrpCtxAfPol in the relation does not exist, the default VRF policy created under the common tenant is used for that address family.

Support for APIC Monitoring Policy

Starting with this connector pack release, you can create a monitoring policy as a default policy to be applied to all tenants in an APIC account to monitor EPGs, application profiles, services, and so on.

Introduction of BGP Timers Policy

Starting with this connector pack release, you can create a Border Gateway Protocol (BGP) Timers policy at a tenant level in Cisco UCS Director. After a BGP peer is established, keepalive messages are sent to the neighbor once in every keepalive interval. If no keepalive message was received within a hold interval, the BGP peer is considered down. The keepalive interval and hold interval must be in the range of 0 to 3600. You can set the Stale Interval timer to delete the stale routes in case the session is not re-established within the specified interval.

Support for Snoop Policy

Starting with this connector pack release, you can create the following policies at a tenant level:

- IGMP Snoop Policy—IGMP snooping is the process of listening to Internet Group Management Protocol (IGMP) network traffic. This allows a network switch to listen to the IGMP conversation between hosts and routers, and to filter multicasts links that do not need them. So with this policy controlling which ports receive specific multicast traffic.
- MLD Snoop Policy—Multicast Listener Discovery (MLD) snooping enables the efficient distribution of IPv6 multicast traffic between hosts and routers. MLD snooping provides the benefit of conserving bandwidth on those segments of the network where no node has expressed interest in receiving multicast traffic. This reduces the bandwidth usage instead of flooding the bridge domain, and also helps hosts and routers save unwanted packet processing.

Support for DHCP Option Policy

Starting with this connector pack release, you can add DHCP option policy to a tenant. To add a DHCP option to a DHCP option policy, choose a DHCP option policy, click **Add** and complete the fields available in the **Add DHCP Option** screen.

Support for BGP Route Target Profile

Starting with this connector pack release, you can create BGP route target profile for each VRF to identify routes that can be imported or exported from the VRF. After creating a BGP route target profile, you can add a BGP route target to the BGP route target profile. A route target is a type of BGP extended community.

Enhancements to VRF

Starting with this connector pack release, you can add BGP context per address family, OSPF context per address family, SNMP context, and community profile to VRFs.

Support for Flow Control Policy

Starting with this connector pack release, you can create a priority flow control (PFC) policy at a fabric resource level. This policy specifies under what circumstances QoS-level PFC is applied to FCoE traffic. To create a PFC policy on the interfaces, click **Add** in the **Flow Control Policy** tab of an APIC account, and enter a name and short description, and choose one of the following states to which the PFC policy will be applied on the local port: **Auto**, **On**, and **OFF**.

Support for Slow Drain Policy

Starting with this connector pack release, you can create a slow drain policy at a fabric resource level for handling FCoE packets that cause traffic congestion on an ACI fabric. To create a slow drain policy, click **Add** in the **Slow Drain Policy** tab of an APIC account, and enter a name and short description. You can specify the action to be taken during FCoE traffic congestion and the number of pause frames to trigger a congestion clear action in order to address the FCoE traffic congestion.

Support for Data Plane Policing

Starting with this connector pack release, you can create a Data Plane Policing (DPP) policy at a fabric resource level to manage bandwidth consumption on ACI fabric access interfaces. DPP policies can apply to egress traffic, ingress traffic, or both.

DPP policies can be single-rate, dual-rate, and color-aware. It can be applied to physical or virtual layer 2 connections for virtual or physical devices such as servers or hypervisors, and on layer 3 connections for routers. To create a DPP policy, click **Add** in the **Slow Data Plane Policing** tab of an APIC account and complete the fields in the **Create APIC Fabric Data Plane Policing Policy** screen.

Enhancements to the Interface Policy

This connector pack release introduces the following alias reports to view alias name of the respective policies. Alias is the alternative name that is applied to objects, which can be changed, unlike the name of an interface policy.

- CDP Interface Policy
- Link Level Policy
- LLDP Interface Policy

- MACsec Access Interface Policy
- Port Channel Member Policy
- Port Channel Policy
- Spanning Tree Interface Policy
- Storm Control Policy

Introduction of New REST APIs

This connector pack release introduces REST APIs for all the features described in this section. To view information on the list of APIs, see the Cisco UCS Director REST API Getting Started Guide, Release 6.7.

New and Changed Features in Release 6.7.4.2

Support for newer Versions of Cisco Application Centric Infrastructure Controller and Cisco Application Centric Infrastructure Multi-Site Controller

This release introduces support for the following:

- Cisco Application Centric Infrastructure Controller 5.0(x)
- Cisco Application Centric Infrastructure Multi-Site Controller 3.0(x)

For more details, refer the Cisco UCS Director Compatibility Matrix, Release 6.7.

Open and Resolved Bugs

The open and resolved bugs for this release are accessible through the Cisco Bug Search Tool. This web-based tool provides you with access to the Cisco bug tracking system, which maintains information about bugs and vulnerabilities in this product and other Cisco hardware and software products.

Note You must have a Cisco.com account to log in and access the Cisco Bug Search Tool. If you do not have one, you can register for an account.

For more information about the Cisco Bug Search Tool, see the Bug Search Tool Help & FAQ.

Open Bugs in Release 6.7.2.1

There are no open bugs in this release.

Open Bugs in Release 6.7.3.1

All open bugs for this release are available in the Cisco Bug Search Tool.

Bug ID	Headline
CSCvr75784	Fabric Network Creation is failing for default VLAN range (md0) for DCNM.

Open Bugs in Release 6.7.3.2

There are no open bugs in this release.

Open Bugs in Release 6.7.4.1

There are no open bugs in this release.

Open Bugs in Release 6.7.4.2

There are no open bugs in this release.

Resolved Bugs in Release 6.7.2.1

All resolved bugs for this release are available in the Cisco Bug Search Tool.

Bug ID	Headline
CSCvo61432	CMDB change record management for APIC & MSC

Resolved Bugs in Release 6.7.3.1

There are no resolved bugs in this release.

Resolved Bugs in Release 6.7.3.2

There are no resolved bugs in this release.

Resolved Bugs in Release 6.7.4.1

All resolved bugs for this release are available in the Cisco Bug Search Tool.

Bug ID	Headline
CSCvr94333	Default task "Create Subnet to APIC Bridge Domain" should have option to enable RA feature.
CSCvs08628	Add domain to EPG task not setting the Switching mode value with AVE type.

Resolved Bugs in Release 6.7.4.2

All resolved bugs for this release are available in the Cisco Bug Search Tool.

Bug ID	Headline
CSCvv79098	**createAPICContractSubject** task fails in validation if APIC contract name contains colon (:).
CSCvv70584	**Create Bridge Domain** task also configures FHS policies in addition to creating bridge domains.
CSCvu93245	**Create APIC Bridge Domain** task does not enable unicast routing even when the **Unicast Routing** checkbox is checked.
CSCvu82258	Taking long time to fetch details for **APIC Device Fabric Interface Policy group identity** input type.
CSCvv22459	**collect inventory** built-in task should have more granularity.
CSCvw01088	**Add Static Route To Logical Node in APIC** task is not validating IPv6 address.

Communications, Services, and Additional Information

- To receive timely, relevant information from Cisco, sign up at Cisco Profile Manager.
- To get the business impact you're looking for with the technologies that matter, visit Cisco Services.
- To submit a service request, visit Cisco Support.
- To discover and browse secure, validated enterprise-class apps, products, solutions and services, visit Cisco Marketplace.
- To obtain general networking, training, and certification titles, visit Cisco Press.
- To find warranty information for a specific product or product family, access Cisco Warranty Finder.

Cisco Bug Search Tool

Cisco Bug Search Tool (BST) is a web-based tool that acts as a gateway to the Cisco bug tracking system that maintains a comprehensive list of defects and vulnerabilities in Cisco products and software. BST provides you with detailed defect information about your products and software.