Cisco ACI Long-Lived Release 4.2(x)

Contents

This document provides information about Cisco ACI long-lived release 4.2(x).

Date	Description
June 10, 2021	This document was published.

About Cisco ACI

The Cisco® Application Centric Infrastructure (Cisco ACI®) is part of our intent-based networking framework to enable agility in the datacenter. It captures higher-level business and user intent in the form of a policy and translates this into the network constructs necessary to dynamically provision network, security, and infrastructure services.

Built on top of the industry-leading Cisco Nexus® 9000 platform, Cisco ACI uses a holistic systems-based approach, with tight integration between hardware and software, between physical and virtual elements, an open ecosystem model, and innovative Cisco Application-Specific Integrated Circuits (ASICs) to enable unique business value for modern data centers.

Cisco ACI is the industry's most secure, open, and comprehensive Software-Defined Networking (SDN) solution.

Cisco ACI enables automation that accelerates infrastructure deployment and governance, simplifies management to easily move workloads across a multifabric, multicloud framework, and proactively secures against risk arising from anywhere. It radically simplifies, optimizes, and expedites the application deployment lifecycle.

Modern data centers are dynamic. IT operations must meet the expectation of quality-of-service business needs in a rapidly changing environment. Cisco ACI transforms IT operations from reactive to proactive with a highly intelligent set of software capabilities that analyzes every component of the data center to ensure business intent, guarantee reliability, and identify performance issues in the network before they happen.

As application usage gets more pervasive across an enterprise's network, IT professionals are looking to build solutions for consistent policy and encryption from the campus to the datacenter. With Cisco ACI integrations with SDA/DNA Center and SD-WAN, customers can now automate and extend policy, security, assurance, and insights across their entire networking ecosystem.

Note: In this document, Cisco ACI followed by a release number represents Cisco APIC and Cisco Nexus 9000 series ACI-mode switches and their release numbers. For example, Cisco ACI 4.2(x) represents Cisco APIC, release 4.2(x) and Cisco Nexus 9000 series ACI-mode switches, release 14.2(x).

Cisco ACI Long-Lived Releases

Cisco ACI long-lived releases are software releases intended to help you stay on a given release on a long-term basis (up to approximately 18 months), while benefiting from frequent maintenance drops to ensure quality and stability. Cisco may support two long-lived releases at any given point of time. However, active maintenance will be focused primarily on the latest long-lived release. These releases will

be maintained for a longer time span than other releases. Long-lived releases are recommended for the deployment of widely adopted functions or for networks that will not be upgraded frequently.

All long-lived releases support upgrade or downgrade to the next or previous long-lived release. See the Cisco APIC Upgrade/Downgrade Support Matrix for confirmed support. Some release branches might be supported as long-lived releases while others might not be supported. For example, there might be three 4.x release branches: 4.1, 4.2, and 4.3. However, one of the three 4.x release branches might be supported as a long-lived release (4.2), while the other two release branches (4.1 and 4.3) might not be supported as long-lived releases.

Figure 1 displays a visual example of a long-lived release branch timeline.

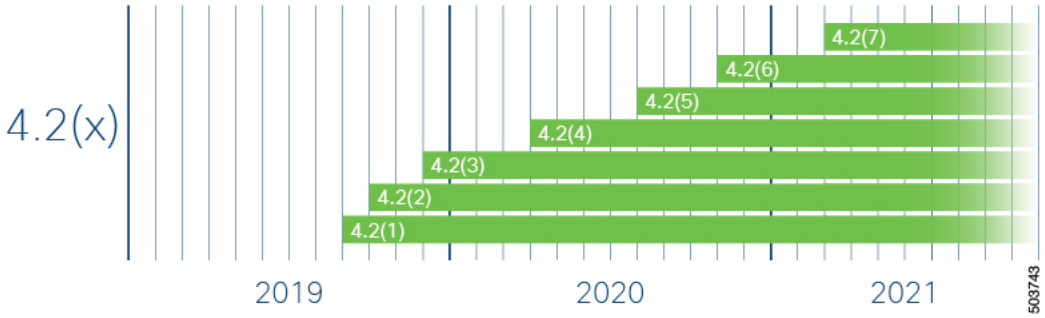

Figure 1. An Example of a Long-Lived Release Branch Timeline

Long-Lived Release Life Cycles

- The life cycle of a major long-lived release starts with the first customer shipment (FCS) of the first minor release.
- The major release then enters the maintenance release introduction phase, in which several releases are made available to address product defects.
- Afterward, the major release transitions to the mature maintenance phase. In this phase, the release receives defect resolutions only for severity 1 and severity 2 defects found by the customer. Defects found internally are addressed on a case-by-case basis.
- All long-lived releases support upgrade or downgrade to the next or previous long-lived release last maintenance version, respectively.

For additional information about long-lived releases, see the following:

- *Cisco APIC Installation, Upgrade, and Downgrade Guide*
- *Cisco APIC Upgrade Downgrade Support Matrix*
- *Recommended Cisco APIC and Cisco Nexus 9000 Series ACI-Mode Switches Releases*

Key Features in Release 4.2(x)

Feature name	Description	Release version
--	No new software features in this release.	4.2(7)
BGP neighbor shutdown	The BGP neighbor shutdown feature is similar to the neighbor shutdown command in NX-OS, which shuts down the corresponding BGP neighbor. Use this policy to disable and enable the BGP neighbor's admin state. Using this feature shuts down the BGP sessions without the need to delete the BGP peer configuration.	4.2(6)
IGMP and MLD packet forwarding through 802.1Q tunnels	IGMP and MLD packets can now be forwarded through 802.1Q tunnels.	4.2(6)
Link-level flow control	Link-level flow control is a congestion management technique that pauses data transmission until the congestion in the system is resolved. When a receiving device becomes congested, the device communicates with the transmitter by sending a pause frame. When the transmitting device receives a pause frame, the device stops the transmission of any further data frames based on the pause quanta value received in the link-level flow control feature pause frame.	4.2(6)
Applying a route map to interleak redistribution from direct subnets	Beginning in the Cisco APIC 4.2(6h) release, you can apply a route map to interleak redistribution from direct subnets (L3Out interfaces).	4.2(6h)
Deny action in the route-map for interleak redistribution for static routes and direct subnets	Beginning in the Cisco APIC 4.2(6h) release, you can configure the deny action in the route-map for interleak redistribution for static routes and direct subnets.	4.2(6h)
SSD write optimization	The SSD write strategy is optimized for improved performance and longer SSD life.	4.2(6)
Improved Precision Time Protocol support	You can now enable the Precision Time Protocol (PTP) on a leaf switch's front panel ports to connect the PTP nodes, clients, or grandmaster. The PTP implementation on fabric ports are still the same as the previous releases, except that the PTP parameters for fabric ports can now be adjusted. With this change, you can use the Cisco ACI fabric to propagate time synchronization using PTP with Cisco ACI switches as PTP boundary clock nodes. Prior to this release, the only approach Cisco ACI had was to use PTP only within the fabric for the latency measurement feature or to forward PTP multicast or unicast messages transparently as a PTP unaware switch from one leaf switch to another as a tunnel.	4.2(5)
Link flap policies	You can create a link flap policy in interface policies, which sets the state of an access port or fabric port to " error-disable" after the port flaps for specified number of times during a specified interval of time.	4.2(5)
UCSC-PCIE-IQ10GC Intel X710 Quad Port 10GBase-T network interface card support	You can now use the UCSC-PCIE-IQ10GC Intel X710 Quad Port 10GBase-T network interface card in the Cisco APIC M3/L3 servers for 10GBase-T connectivity to Cisco ACI leaf nodes.	4.2(5)
Upgrade enhancements	Various enhancements have been made to the upgrade process, including: • The restriction on the number of pods that you can upgrade in parallel has been relaxed so that you can upgrade multiple pods at the same time for pod nodes in Multi-Pod configurations. Switches in a Multi-Pod configuration that are part of the	4.2(5)

Feature name	Description	Release version
	same maintenance group can now be upgraded in parallel. • Upgrades or downgrades might be blocked if certain issues are present. • Additional information is provided in the GUI for each stage of the APIC upgrade or downgrade process. • The default concurrency in a group has changed from 20 to unlimited (the default number of leaf or spine switches that can be upgraded at one time is unlimited). • When upgrading nodes in an upgrade group using the GUI, Download Progress field is available in the Work pane, which provides a status on the progress of the download of the firmware for the node upgrade.	
Enhancements for remote leaf switches	Starting with release 4.2(4), the following enhancements have been introduced for remote leaf switches: • Support for 10 Mbps as a minimum bandwidth in the IPN • Support to create an 802.1Q tunnel between the remote leaf switch and the ACI main datacenter	4.2(4)
IGMP snooping version 2 group scale increase	IGMP snooping now supports 32,000 groups.	4.2(4)
Layer 3 multicast VRF scale increase	Layer 3 multicast VRF scale increase.	4.2(4)
Multipod leaf switch scale increase	A Cisco ACI Multi-Pod environment now supports up to 400 leaf switches per pod.	4.2(4)
Network Insights Base app is now prepackaged with Cisco APIC	The Network Insights Base app is now prepackaged with the Cisco APIC software.	4.2(4)
Support for 25 SCVMM domains with 10,000 endpoints	A Cisco ACI fabric can now have up to 25 SCVMM domains with 10k endpoints in pre-provisioned mode.	4.2(4)
Support for custom EPG names for VMM domains	You can give EPGs a custom name that carries over to a VMware vCenter port group or a Microsoft VM network. The feature is available for VMware vSphere Distributed Switch, Microsoft System Center Virtual Machine Manager (SCVMM), and Cisco ACI Virtual Edge. If you do not provide a custom name, the domain association assigns one.	4.2(4)
Support for VMware vSphere 7.0 with VMware vSphere Distributed Switch and Cisco ACI Virtual Edge	The 4.2(4o) release adds support for VMware vSphere 7.0 with the VMware vSphere Distributed Switch (VDS) and Cisco ACI Virtual Edge.	4.2(4o)
User lockout after continuous failed attempts to login	You can block a user from being able to log in after the user fails a configured number of login attempts. You can specify how many failed login attempts the user can have within a specific time period. If the user fails to log in too many times, then that user becomes unable to log in for a specified period of time.	4.2(4)
COOP Endpoint Dampening	When malicious or erroneous behavior causes unnecessary endpoint updates, the COOP process can become overwhelmed, preventing the processing of valid endpoint updates. The rogue endpoint detection feature of the leaf switch can prevent many erroneous updates from reaching the spine. In cases where the rogue endpoint detection is inadequate, the COOP process invokes endpoint dampening. To relieve pressure on COOP, the	4.2(3)

Feature name	Description	Release version
	spine asks all leaf switches to ignore updates from the misbehaving endpoint for a specified period.	
Enhancements for Match Prefix	Two new fields (From Prefix and To Prefix fields) are now available in the Match Prefix field to specify the mask range when you create a prefix match rule and enable aggregation.	4.2(3)
Filters-from-contract option in the service graph templates	The filters-from-contract option is available in the service graph templates using the Cisco APIC GUI. This option uses the specific filter of the contract subject where the service graph is attached, instead of the default filter for zoning-rules that do not include the consumer EPG class ID as the source or destination.	4.2(3)
Increased range for equal-cost multi-path (ECMP) routing paths	The range for the maximum number of equal-cost paths for eBGP and iBGP load sharing is now from 1 to 64, with a default value of 16.	4.2(3)
Incremental enhancements to the read-only admin user capability on spine and leaf switches	Switches running the 14.2(3) release now support L1 access (read-only privilege for an admin user) for the following things: • acidiag fnvread command • vsh_lc with the show commands • Tech support collection • show events command • PCAP under the visibility and troubleshooting section • BGP advertised and received routes (show bgp ipv4 unicast neighbor <neighbor ip> advertised-routes vrf <vrf name>) • CRC command to identify stomped CRC and genuine CRC • Read-only access to the log files, such as BGP, BFD, and IPv6 • tcpdump command	4.2(3)
Python SDK (Cobra) support for Python 3.x and Wheel	The Cisco APIC Python SDK adds support for Python 3.6 and later. A Wheel installation package is now included in addition to the egg files.	4.2(3)
Rogue EP Control in the First Time Setup wizard	The Rogue EP Control option is now part of the First Time Setup wizard.	4.2(3)
Stomped CRC errors and traditional CRC errors	CRC align errors in interface counters are now broken out into stomped CRC errors and traditional CRC errors. Stomped CRC errors refer to frames that were received and cut-through switched before the FCS trailer was received. Rather than rewriting the CRC field based on the corrupted frame, the switch will insert a special value into the CRC that indicates the frame should be stomped by the end device or the first device in the path that does store-and-forward switching. " CRC error" frames refer to corrupted frames that are dropped on the ingress interface and are not forwarded. You can view the split in error statistics in the Cisco APIC GUI or by directly querying the eqptIngrCrcErrPkts object. Additionally you can view the statistics directly on the switch by running the " show interface" command.	4.2(3)
Support for custom EPG names for VMM domains	You can now give EPGs a custom name that carries over to a VMware vCenter port group or a Microsoft VM network. The feature is available for VMware vSphere Distributed Switch, Microsoft System Center Virtual Machine Manager (SCVMM), and Cisco ACI Virtual Edge. If you do not provide a custom name, the domain association assigns a name in the format of tenant|app_profgile|epg_name for a port group or	4.2(3)

Feature name	Description	Release version
	tenant|application|epg|domain for a VM network. However, if you enter a custom name for the EPG, the same name is applied to the port group or VM network.	
Support for QoS MIBs	Selected OIDs from CISCO-CLASS-BASED-QOS-MIB and CISCO-SWITCH-QOS-MIB are added for leaf and spine switches.	4.2(3)
Remote leaf switch failover	In a multipod setup, if a remote leaf switch in a pod loses connectivity to the spine switch, the remote leaf switch now moves to another pod. This ensures that traffic continues to flow between endpoints of remote leaf switches that are connected to the original pod.	4.2(2)
Ability to pin EPGs to an uplink on a VMware DVS	You can configure up to 32 uplinks for each instance of Cisco ACI Virtual Edge (in native switching mode) or VMware VDS. You also can rename the uplinks and configure failover for them within endpoint groups (EPGs) associated with the VMware VDS or Cisco ACI Virtual Edge.	4.2(1)
avread CLI command	Cisco APIC Release 4.2.(1) introduces the new avread command, which provides the same information as the acidiag avread command, but in a tabular format.	4.2(1)
BGP neighbor shutdown	The BGP neighbor shutdown feature is similar to the neighbor shutdown command in NX-OS, which shuts down the corresponding BGP neighbor. Use this policy to disable and enable the BGP neighbor's admin state. Using this feature shuts down the BGP sessions without the need to delete the BGP peer configuration.	4.2(1)
BGP neighbor soft reset	The BGP neighbor soft reset feature provides automatic support for a dynamic soft reset of inbound and outbound BGP routing table updates that are not dependent upon stored routing table update information. Use this policy to enable the soft dynamic inbound reset and soft outbound reset.	4.2(1)
Blocking ACI upgrades or downgrades if faults are present	Beginning with release 4.2(1), when you attempt to trigger an upgrade or downgrade operation, the operation might be blocked if any faults on the fabric are detected, depending on the severity of the fault detected.	4.2(1)
cluster_health CLI command	Cisco APIC Release 4.2.(1) introduces the new cluster_health command, which enables you to verify the Cisco APIC cluster status.	4.2(1)
fd_vlan mismatch enhancement	If the same VLAN pool is being used on both a vPC and an orphan port, a fd_vlan mismatch will occur and a fault will be raised.	4.2(1)
Floating Layer 3 Outside network connection	You can configure a floating L3Out that allows a virtual router to move from under one leaf switch to another. The feature saves you from having to configure multiple L3Out interfaces to maintain routing when virtual machines move from one host to another. This feature is supported for VMware VDS.	
IPv6 multicast support	IPv6 multicast is now enabled with PIM6 protocol settings.	4.2(1)
Policy-based redirect backup policy	This feature enables you to configure a backup node for a policy-based redirect (PBR) policy. If an active node goes down, traffic gets routed through the backup node instead of getting routed through one of the other active nodes. The backup node avoids a situation in which the connection could be reset if, for example, the data paths through another active node are traversing stateful firewalls.	4.2(1)
Redistributing static routes to BGP with prefix list	For Cisco APIC releases before release 4.2(1), you can configure a route map policy for the redistribution of static routes into BGP using the Create Route Map/Profile feature, which defines the route map for BGP dampening	4.2(1)

Feature name	Description	Release version
	and route redistribution. This feature is used to set attributes, such as community, on certain static routes on one border leaf switch, and then, based on these attributes, configure these routes on other border leaf switches Beginning with Cisco APIC Release 4.2(1), this feature is extended for static routes. This allows you to configure a route map policy that will be applied while redistributing static routes into BGP.	
Route control on an aggregator route during import/export	When creating a subnet, the export route control subnet and import route control subnet allow Aggregate Export and Aggregate Import.	4.2(1)
Route control per BGP peer	Route control policies determine what routes are advertised out to the external network (export) or allowed into the fabric (import). Prior to Cisco APIC release 4.2(1), you configure these policies at the L3Out level, under the L3Out profile (l3extInstP) or through the L3Out subnet under the L3Out (l3extSubnet), so those policies apply to protocols configured for all nodes or paths included in the L3Out. With this configuration, there could be multiple node profiles configured in the L3Out, and each could have multiple nodes or paths with the BGP neighbor specified. Because of this, there is no way to apply individual policies to each protocol entity. Beginning with Cisco APIC release 4.2(1), the route control per BGP peer feature is introduced to begin to address this situation, where more granularity in route export and import control is needed.	4.2(1)
SDWAN integration enhancement	This release adds support for enabling returning traffic from a remote site that is destined for the ACI data center to receive differentiated services over the WAN. After the tenant admin registers the Cisco APIC to vManage, the Cisco APIC pulls the WAN-SLA policies and the WAN-VPN from vManage. Then, the Cisco APIC assigns a DSCP to each WAN-SLA policy and pushes a prefix list. The prefix list, which is taken from the EPG if the contract between this EPG and L3Out has WAN-SLA configured, enables quality of service on the returning traffic. The WAN-SLA policy and WAN-VPN are both available in the tenant common. Tenant admins map the WAN-VPNs to VRF instances on remote sites.	4.2(1)
Simplified ELAM output	This release adds an option to the Embedded Logic Analyzer Module (ELAM) tool that changes the output to a human-readable format, which enables you to find key information quickly and more efficiently. In addition, hexadecimal values have been converted to decimal values in some instances for improved readability. For backward compatibility, the existing usage of ELAM is kept intact.	4.2(1)
Storm control SNMP traps	This release supports triggering SNMP traps from Cisco ACI when storm control thresholds are met.	4.2(1)

Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at **https://www.cisco.com/go/offices**.

Printed in USA 01/21