Verified Scalability Guide for Cisco APIC, Release 6.0(2) and Cisco Nexus 9000 Series ACI-Mode Switches, Release 16.0(2)

Overview

This guide contains the maximum verified scalability limits for Cisco Application Centric Infrastructure (Cisco ACI) parameters in the following releases:

- Cisco Application Policy Infrastructure Controller (Cisco APIC), Release 6.0(2)

- Cisco Nexus 9000 Series ACI-Mode Switches, Release 16.0(2)

These values are based on a profile where each feature was scaled to the numbers specified in the tables. These numbers do not represent the theoretically possible Cisco ACI fabric scale.

Note The verified scalability limits for Cisco Multi-Site previously included as part of this guide are now listed in the Cisco Nexus Dashboard Orchestrator (NDO) release-specific documents available at the following URL: https://www.cisco.com/c/en/us/support/cloud-systems-management/multi-site-orchestrator/products-device-support-tables-list.html.

The verified scalability limits for Cisco Cloud APIC previously included as part of this guide are now listed in the Cloud APIC release-specific documents available at the following URL: https://www.cisco.com/c/en/us/support/cloud-systems-management/cloud-application-policy-infrastructure-controller/products-tech-notes-list.html.

New and Changed Information

These changes have been made to this document since the initial release:

Date	Changes
October 21, 2024	Updated General Scalability L4-L7 Devices topic to 1,200 physical or virtual devices and Layer 4 to Layer 7 topic by changing configurable options to add "concrete devices."
October 8, 2024	Updated High Dual Stack profile values in Fabric Topology section.
August 20, 2024	Updated VRFs Per Leaf Scale value and the External EPGs values in Fabric Topology section.
August 19, 2024	Updated "Per Fabric Scale" values of the IP SLA probes in the Fabric Topology section.
June 26, 2024	Updated General Scalability section with a Note about fabric having 32 GB of RAM. Updated Fabric Topology content with notes about EP scale, fabric-wide, fabric having 32 GB of RAM.
	Updated Fabric Topology content with updated scale guidelines for the number of Consumers and Providers for the same contract and Number of VLAN encapsulations per EPG.
March 13, 2023	First release of this document.

General Scalability Limits

- **L2 Fabric:** L2 Fabric in this document refers to an ACI fabric that contains only BDs with Scaled L2 Only mode (formerly known as Legacy mode). See **Bridging** > **Bridge Domain Options** > **Scaled L2 Only Mode - Legacy Mode** in APIC *Layer 2 Configuration Guide* for details about Scaled L2 Only mode.

- **L3 Fabric:** The ACI L3 fabric solution provides a feature-rich highly scalable solution for public cloud and large enterprise. With this design, almost all supported features are deployed at the same time and are tested as a solution. The scalability numbers listed in this section are multi-dimensional scalability numbers. The fabric scalability numbers represent the overall number of objects created on the fabric. The per-leaf scale numbers are the objects created and presented on an individual leaf switch. The fabric level scalability numbers represent APIC cluster scalability and the tested upper limits. Some of the per-leaf scalability numbers are subject to hardware restrictions. The per-leaf scalability numbers are the maximum limits tested and supported by leaf switch hardware. This does not necessarily mean that every leaf switch in the fabric was tested with maximum scale numbers.

- **Stretched Fabric:** Stretched fabric allows multiple fabrics (up to 3) distributed in multiple locations to be connected as a single fabric with a single management domain. The scale for the entire stretched fabric remains the same as for a single site fabric. For example a L3 stretched fabric will support up to 400 leaf switches total which is the maximum number of leaf switches supported on a single site fabric. Parameters only relevant to stretched fabric are mentioned in the tables below.

- **Multi-Pod:** Multi-Pod enables provisioning a more fault-tolerant fabric comprised of multiple Pods with isolated control plane protocols. Also, Multi-Pod provides more flexibility with regard to the full mesh cabling between leaf and spine switches. For example, if leaf switches are spread across different floors or different buildings, Multi-Pod enables provisioning multiple Pods per floor or building and providing connectivity between Pods through spine switches.

  Multi-Pod uses a single APIC cluster for all the Pods; all the Pods act as a single fabric. Individual APIC controllers are placed across the Pods but they are all part of a single APIC cluster.

- **Multi-Site**: Multi-Site is the architecture interconnecting and extending the policy domain across multiple APIC cluster domains. As such, Multi-Site could also be named as Multi-Fabric, since interconnects separate Availability Zones (Fabrics) and managed by an independent APIC controller cluster. A Cisco Nexus Dashboard Orchestrator (NDO) is part of the architecture and is used to communicate with the different APIC domains to simplify the management of the architecture and the definition of inter-site policies.

Leaf Switches and Ports

The maximum number of leaf switches is 400 per Pod and 500 total in Multi-Pod fabric. The maximum number of physical ports is 24,000 per fabric. The maximum number of remote leaf (RL) switches is 200 per fabric, with total number of BDs deployed on all remote leaf switches in the fabric not exceeding 60,000. The total number of BDs on all RLs is equal to the sum of BDs on each RL.

If Remote Leaf Pod Redundancy policy is enabled, we recommended that you disable the `Pre-emption` flag in the APIC for all scaled up RL deployments. In other words, you must wait for BGP CPU utilization to fall under 50% on all spine switches before you initiate pre-emption.

Breakout Ports

The N9K-C9336C-FX2 switch supports up to 34 breakout ports in both 10G or 25G mode.

General Scalability Limits

| Note | For large fabrics, we recommend that all spines in the fabric have 32 GB of RAM. |

Table 1: Fabric Scale Limits Per Cluster Size

Configurable Options	Default Fabric	Medium Fabric	Large Fabric	
Number of APIC nodes	3	4	5 or 6	7
Number of leaf switches	85	200	300	500
Number of leaf switches per Pod	85	200	200	400
Number of tier-2 leaf switches per Pod in Multi-Tier topology **Note** The total number of leaf switches from all tiers must not exceed the "Number of leaf switches" listed above.	80	100	125	125
Number of Pods	6	6	25	25
Number of spine switches in a Multi-Pod fabric	24	24	50	50
Number of tenants	1,000	1,000	3,000	3,000
Number of Layer 3 (L3) contexts (VRFs)	1,000	1,000	10,000	10,000
Number of L3Outs	2,400	2,400	10,000	10,000

Configurable Options	Default Fabric	Medium Fabric	Large Fabric	
Number of external EPGs across all BLs This is calculated as a product of (Number of external EPGs)*(Number of border leaf switches for the L3Out). For example, this combination adds up to a total of 2000 external EPGs in the fabric (250 external EPGs * 2 border leaf switches * 4 L3Outs): • 250 External EPGs in `L3Out1` on `leaf1` and `leaf2` • 250 External EPGs in `L3Out2` on `leaf1` and `leaf2` • 250 External EPGs in `L3Out3` on `leaf3` and `leaf4` • 250 External EPGs in `L3Out4` on `leaf3` and `leaf4`	4,000	4,000	10,000	10,000

Table 2: General Scalability Limits Per Fabric

Configurable Options	Scale Limits
Number of spine switches per Pod	6
Number of FEXs	650 (maximum of 20 FEXs and 576 ports per leaf)
Number of contracts	10,000
Number of contract filters	10,000
Number of endpoint groups (EPGs)	15,000 (21,000 for L2 fabric)

Configurable Options	Scale Limits

Number of EPGs per tenant	General limits:
	
	  • Single-tenant fabrics: 4,000
	
	  • Multi-tenant fabrics: 500
	
	Or one of these two specific use cases within the same fabric (the EPGs must be deployed on local leaf switches only, not on remote leaf switches):
	
	  • Use case 1:
	
	    • Up to 10 tenants that have up to 700 EPGs per tenant, with the EPGs distributed across up to 100 leaf switches
	
	  • Use case 2:
	
	    • 1 tenant with up to 1,400 EPGs deployed on up to 100 leaf switches
	
	    For example, `tenant1` with `EPG1-1400` on `leaf1-100`
	
	    • 1 tenant with up to 800 EPGs deployed on a different set of up to 20 leaf switches
	
	    For example, `tenant2` with `EPG1401-2200` on `leaf101-120`
	
	    • 2 tenants with up to 800 EPGs per tenant deployed on a different set of up 20 leaf switches
	
	    For example, `tenant3` with `EPG2201-3000` and `tenant4` with `EPG 3001-3800` on `leaf121-140`

Number of bridge domains (BDs)	15,000 (21,000 for L2 fabric)
Number of vCenters	200 VDS
Number of Service Chains	1,000
Number of L4-L7 concrete devices	1,200 physical or virtual devices (1,200 maximum in total per fabric)
Number of ESXi hosts - VDS	3,200
Number of VMs	Depends on server scale
Number of configuration zones per fabric	30
L3 EVPN services over fabric WAN - GOLF (with and without OpFlex)	1,000 VRFs 60,000 routes in a fabric
Number of Routes in Overlay-1 VRF	1,000

Configurable Options	Scale Limits
Floating L3Out	6 anchor nodes
	32 non-anchor nodes

Multiple Fabric Options Scalability Limits

Stretched Fabric

Configurable Options	Per Fabric Scale
Number of fabrics that can be a stretched fabric	3
Number of Route Reflectors	6

Multi-Pod

Configurable Options	Per Fabric Scale
Number of Pods	25
Number of leaf switches per Pod	400
Number of leaf switches overall	500
Number of Route Reflectors for L3Out	50
Number of External Route Reflectors between Pods	• **For 1-3 Pods**: Up to 3 external route reflectors We recommend full mesh for external BGP peers instead of using external route reflectors when possible • **For 4 or more Pods**: Up to 4 external route reflectors We recommend using external route reflectors instead of full mesh We recommend that the external route reflectors are distributed across Pods so that in case of any failure there are always at least two Pods with external route reflectors still reachable

Cisco Multi-Site Scalability Limits

Cisco Nexus Dashboard Orchestrator (NDO) does not require a specific version of APIC to be running in all sites. The APIC clusters in each site as well as the NDO itself can be upgraded independently of each other and run in mixed operation mode as long as each fabric is running APIC, Release 3.2(6) or later.

As such, the verified scalability limits for your specific Cisco Nexus Dashboard Orchestrator release are now available at the following URL: https://www.cisco.com/c/en/us/support/cloud-systems-management/multi-site-orchestrator/products-device-support-tables-list.html.

> ✎
>
> **Note** Each site managed by the Cisco Nexus Dashboard Orchestrator must still adhere to the scalability limits specific to that site's APIC Release. For a complete list of all *Verified Scalability Guides*, see https://www.cisco.com/c/en/us/support/cloud-systems-management/application-policy-infrastructure-controller-apic/tsd-products-support-series-home.html#Verified_Scalability_Guides

Fabric Topology, SPAN, Tenants, Contexts (VRFs), Equal Cost Multipath (ECMP), External EPGs, Bridge Domains, Endpoints, and Contracts Scalability Limits

This content shows the mapping of the "Application Leaf Engine (ALE) and Leaf Spine Engine (LSE) type" to the corresponding leaf switches. The information is helpful to determine which leaf switch is affected when we use the terms LSE or LSE2 in the remaining sections.

> ✎
>
> **Note** The switches are listed as LSE or LSE2 for scalability purposes only. Check specific feature documentation for the full list of supported devices.

LSE Type	ACI-Supported leaf switches
LSE	• N9K-C93108TC-EX • N9K-C93108TC-EX-24 • N9K-C93180YC-EX • N9K-C93180YC-EX-24 • N9K-C93180LC-EX • N9K-C9336C-FX2 • N9K-C93216TC-FX2 • N9K-C93240YC-FX2 • N9K-C93360YC-FX2 • N9K-C9336C-FX2-E • N9K-C9364D-GX2A • N9K-C9348D-GX2A • N9K-C9400-SW-GX2A

LSE Type	ACI-Supported leaf switches
LSE2	• N9K-C93108TC-FX • N9K-C93108TC-FX-24 • N9K-C93180YC-FX • N9K-C93180YC-FX-24 • N9K-C9348GC-FXP • N9K-C93600CD-GX • N9K-C9364C-GX • N9K-C9316D-GX • N9K-C9332D-GX2B • N9K-C93180YC-FX3 • N9K-C93108TC-FX3P • N9K-C9358GY-FXP with 24GB of RAM • N9K-C93180YC-FX3H • N9K-C93108TC-FX3H

Note

- The High Policy, Multicast-Heavy, and High IPv4 EP Scale profiles are not supported on FXP switches.

- Full scale support for High Policy, Multicast-Heavy, and High IPv4 EP Scale profiles requires LSE2 with 32 GB of RAM.

- High IPv4 EP Scale—This profile is recommended to be used only for the ACI border leaf (BL) switches in Multi-Domain (ACI-SDA) Integration. It provides enhanced IPv4 EP and LPM scales specifically for these BLs and has specific hardware requirements.

- For maximum EP scale, fabric-wide, we recommend that all spines in the fabric have 32 GB of RAM.

- For full scale support of Maximum LPM scale profile, we recommend that all spines in the fabric have 32 GB of RAM.

For more details on Forwarding Scale Profiles and the list of supported devices, please refer to *Cisco APIC Forwarding Scale Profiles* at this url: https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/all/forwarding-scale-profiles/cisco-apic-forwarding-scale-profiles.html

Fabric Topology

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of PCs, vPCs	320 (with FEX HIF)	N/A
Number of encapsulations per access port, PC, vPC (non-FEX HIF)	3,000	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of encapsulations per FEX HIF, PC, vPC	100	N/A
Number of encapsulations per FEX	1,400	N/A
Number of member links per PC, vPC* *vPC total ports = 32, 16 per leaf	16	N/A
Number of ports x VLANs (global scope and no FEX HIF)	64,000 168,000 (when using legacy BD mode)	N/A
Number of ports x VLANs (FEX HIFs and/or local scope)	10,000	N/A
Number of static port bindings	60,000	700,000 (200,000 per tenant)
Number of VMACs	510	N/A
STP	All VLANs	N/A
Mis-Cabling Protocol (MCP)	2,000 VLANs per interface 12,000 logical ports (port x VLAN) per leaf	N/A
Mis-Cabling Protocol (MCP) (strict mode enabled on the port)	256 VLANs per interface 2,000 logical ports (port x VLAN) per leaf	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale

Number of endpoints (EPs)	Default profile or High LPM profile:
	
	• MAC: 24,000
	
	• IPv4: 24,000
	
	• IPv6: 12,000
	
	Maximum LPM profile:
	
	• MAC: 8,000
	
	• IPv4: 8,000
	
	• IPv6: 4,000
	
	IPv4 scale profile:
	
	• MAC: 48,000
	
	• IPv4: 48,000
	
	• IPv6: Not supported
	
	High Dual Stack scale profile:
	
	• LSE:
	
	    • MAC: 64,000
	
	    • IPv4: 64,000
	
	    • IPv6: 24,000
	
	• LSE2:
	
	    • MAC: 64,000
	
	    • IPv4: 64,000
	
	    • IPv6: 48,000
		16-slot and 8-slot modular spine switches:
		
		Max. 450,000 Proxy Database Entries in the fabric, which can be translated into any one of these:
		
		• 450,000 MAC-only EPs (each EP with one MAC only)
		
		• 225,000 IPv4 EPs (each EP with one MAC and one IPv4)
		
		• 150,000 dual-stack EPs (each EP with one MAC, one IPv4, and one IPv6)
		
		The formula to calculate in mixed mode is:
		
		#MAC + #IPv4 + #IPv6 <= 450,000
		
		NOTE: Four fabric modules are required on all spines in the fabric to support above scale.
		
		--- 	
		
		4-slot modular spine switches:
		
		Max. 360,000 Proxy Database Entries in the fabric, which can be translated into any one of these:
		
		• 360,000 MAC-only EPs (each EP with one MAC only)
		
		• 180,000 IPv4 EPs (each EP with one MAC and one IPv4)
		
		• 120,000 dual-stack EPs (each EP with one MAC, one IPv4, and one IPv6)
		
		The formula to calculate in mixed mode is:
		
		#MAC + #IPv4 + #IPv6 <= 360,000
		
		NOTE: Four fabric modules are required on all spines in the fabric to support above scale.

Configurable Options	Per Leaf Scale	Per Fabric Scale

Number of endpoints (EPs)
(Continued)
	High Policy profile:
	• LSE2 (except FXP switches):
	  • MAC: 24,000
	  • IPv4: 24,000
	  • IPv6: 12,000
	
	• LSE:
	  • MAC: 16,000
	  • IPv4: 16,000
	  • IPv6: 8,000
	
	High IPv4 EP Scale profile:
	• LSE: Not supported
	• LSE2 (with 32 GB of RAM):
	  • MAC: 24,000
	  • IPv4 local: 24,000
	  • IPv4 total: 280,000
	  • IPv6: 12,000
	
	Multicast Heavy profile:
	• LSE: Not supported
	• LSE2 (except FXP switches):
	  • MAC: 24,000
	  • IPv4 local: 24,000
	  • IPv4 total: 64,000
	  • IPv6: 4,000
		Fixed spine switches:
		Max. 180,000 Proxy Database Entries in the fabric, which can be translated into any one of these :
		• 180,000 MAC-only EPs (each EP with one MAC only)
		• 90,000 IPv4 EPs (each EP with one MAC and one IPv4)
		• 60,000 dual-stack EPs (each EP with one MAC, one IPv4, and one IPv6)
		
		The formula to calculate in mixed mode is:
		#MAC + #IPv4 + #IPv6 <= 180,000

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of Multicast Routes	Default (Dual Stack), IPv4 Scale, High LPM, High Policy or High IPv4 EP scale profiles: 8,000 with (S,G) scale not exceeding 4,000 Maximum LPM profile: • 1,000 with (S,G) scale not exceeding 500 High Dual Stack profile: • LSE: 512 • LSE2: 32,000 with (S,G) scale not exceeding 16,000 Multicast Heavy profile: • LSE: not supported • LSE2 (with 32 GB of RAM): 90,000 with (S,G) scale not exceeding 72,000	128,000
Number of Multicast Routes per VRF	Default (Dual Stack), IPv4 Scale, High LPM, High Policy or High IPv4 EP scale profiles: 8,000 with (S,G) scale not exceeding 4,000 Maximum LPM profile: • 1,000 with (S,G) scale not exceeding 500 High Dual Stack profile: • LSE: 512 • LSE2: 32,000 with (S,G) scale not exceeding 16,000 Multicast Heavy profile: • LSE: not supported • LSE2 (except FXP switches): 32,000	32,000

Configurable Options	Per Leaf Scale	Per Fabric Scale

IGMP snooping L2 multicast routes
 	
  • For IGMPv2, route scale is for (*, G) only
 	
  • For IGMPv3, route scale is for both (S, G) and (*, G)
 	
Note
IGMP snooping entries are created per BD (2 receivers that join the same group from 2 different BDs consume 2 separate entries).
	Default (Dual Stack), IPv4, High LPM, High Policy, or High IPv4 EP scale profiles: 8,000
	
	Maximum LPM profile:
	
	  • 1,000
	
	High Dual Stack profile:
	
	  • LSE: 512
	
	  • LSE2: 32,000
	
	Multicast Heavy profile:
	
	  • LSE: not supported
	
	  • LSE2: 32,000
		32,000

Number of IPs per MAC	4,096	4,096
Number of Host-Based Routing Advertisements	30,000 host routes per border leaf	N/A
SPAN	32 unidirectional or 16 bidirectional sessions (fabric, access, or tenant)	N/A
Number of ports per SPAN session **Note** This is also the total number of unique ports (fabric and access) that can be used as SPAN sources across all SPAN sessions combined	63 – total number of unique ports (fabric + access) across all types of span sessions	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of SPAN sources in each direction	$2 * (V + FP + AP1 + AP2 + (V * AP1) + AP3\_v6) + AP3\_v4 <= 480$ Where: • V: Number of source VLANs in Tenant SPAN. Each source EPG may contain multiple VLANs. • FP: Number of source ports in Fabric SPAN • AP1: Number of source ports in Access SPAN without any filters • AP2: Number of (VLAN, Port) pairs in Access SPAN with EPG/L3Out filters. Each EPG/L3Out ma contain multiple VLANs. • V * AP1: When both "V" and "AP1" are configured, additional entries are created for each (V, AP1) pair. • AP3_v6: Number of (IPv6 filter entry, Port) pairs in Access SPAN with Filter Group • AP3_v4: Number of (IPv4 filter entry, Port) pairs in Access SPAN with Filter Group	N/A
Number of VLAN encapsulations per EPG	If EPG has 3 VLAN encapsulations = 3 entries	If EPG has 3 VLAN encapsulations = 3 entries
Number of L4 Port Ranges	16 (8 source and 8 destination) First 16 port ranges consume a TCAM entry per range. Each additional port range beyond the first 16 consumes a TCAM entry per port in the port range. Filters with distinct source port range and destination port range count as 2 port ranges. You cannot add more than 16 port ranges at once.	N/A
Common pervasive gateway	256 virtual IPs per Bridge Domain	N/A
Number of Data Plane policers at the interface level	• 7 ingress policers • 3 egress policers	N/A
Number of Data Plane policers at EPG and interface level	128 ingress policers	N/A
Number of interfaces with Per-Protocol Per-Interface (PPPI) CoPP	63	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of TCAM entries for Per-Protocol Per-Interface (PPPI) CoPP	256 One PPPI CoPP configuration may use more than one TCAM entry. The number of TCAM entries used for each configuration varies in each protocol and leaf platform. Use `vsh_lc -c 'show system internal aclqos pppi copp tcam-usage'` command to check on LSE/LSE2 platforms	N/A
Number of SNMP trap receivers	10	10
IP SLA probes* *With 1 second probe time and 3 seconds of timeout	200	• 1500 (for PBR tracking) • 400 (for static route tracking)
First Hop Security (FHS)* With any combination of BDs/EPGs/EPs within the supported limit	2,000 endpoints 1,000 bridge domains	N/A
Number of Q-in-Q tunnels (both QinQ core and edge combined)	1,980	N/A
Number of TEP-to-TEP atomic counters (tracked by 'dbgAcPathA' object)	N/A	1,600

SR-MPLS

Configurable Options	Per Leaf Scale	Per Fabric Scale
EVPN sessions	4	100
BGP labeled unicast (LU) pairs	16	200
ECMP paths	16	N/A
Infra SR-MPLS L3Outs* * Including both, remote leaf and multi-pod	N/A	100 total, 2 per RL location
VRFs* * Including both, remote leaf and multi-pod	800	5,000
External EPGs	800	5,000 total, 100 per VRF
Interfaces	N/A	Same as fabric scale
Multi-pod remote leaf pairs	N/A	50 pairs (100 RLs total)

Tenants

Configurable Options	Per Leaf Scale	Per Fabric Scale
Contexts (VRFs) per tenant	128	128

VRFs (Contexts)

✎

Note   When deploying more than 1,000 VRFs, we recommend that all spines in the fabric have 32 GB of RAM.

All numbers are applicable to dual stack unless explicitly called out.

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of Contexts (VRFs)	Maximum LPM scale profile: • 250 All other scale profiles: • 800	*See* Table 1: Fabric Scale Limits Per Cluster Size
Number of isolated EPGs	400	400
Border Leafs per L3Out	N/A	24 **Note** qualified with 100 VRFs + 16,000 IPv4 + 6400 IPv6 external prefixes
Number of vzAny Provided Contracts	Shared services: Not supported Non-shared services: 70 per Context (VRF)	N/A
Number of vzAny Consumed Contracts	Shared services: 16 per Context (VRF) Non-shared services: 70 per Context (VRF)	N/A
Number of Graphs Instances per device cluster	N/A	500
L3Out per context (VRF)	N/A	400

Configurable Options	Per Leaf Scale	Per Fabric Scale

Number of BFD neighbors	• Up to 256 sessions using these minimum BFD timers:
	
	  • `minTx:50`
	
	  • `minRx:50`
	
	  • `multiplier:3`
	
	• 257-2,000 sessions using these minimum BFD timers:
	
	  • `minTx:300`
	
	  • `minRx:300`
	
	  • `multiplier:3`
		N/A

Number of BGP neighbors	2,000 with up to 70,000 external prefixes with a single path	20,000

Number of OSPF neighbors	• Up to 700 with up to 10,000 external prefixes using these timers:
	
	  • `Hello` timer of `10` seconds
	
	  • `Dead` timer of `40` seconds
	
	  • No more than 300 OSPF neighbors per VRF
	
	• 701-2,000 with up to 35,000 external prefixe using these timers:
	
	  • `Hello` timer of `40` seconds
	
	  • `Dead` timer of `160` seconds
	
	  • No more than 300 OSPF neighbors per VRF
		12,000

Number of EIGRP neighbors	32	N/A
Number of subnets for Route Summarization	1,000	N/A
Number of static routes to a single SVI/VRF	5,000	N/A
Number of static routes on a single leaf switch	10,000	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale

Number of IP Longest Prefix Matches (LPM) entries
 	
Note
Except for the maximum LMP scale profile, the total of `(# of IPv4 prefixes) + 2*(# of IPv6 prefixes)` must not exceed the scale listed for IPv4 alone
	Default (Dual Stack) profile:
	
	• IPv4: 20,000 or
	
	• IPv6: 10,000
	
	• IPv6 wide prefixes (>= /84): 1,000
	
	   NOTE: This restriction only applies to EX models in LSE.
	
	IPv4 scale profile:
	
	• IPv4: 38,000
	
	• IPv6: Not supported
	
	High Dual Stack scale profile:
	
	• IPv4: 38,000 or
	
	• IPv6: 19,000
	
	• IPv6 wide prefixes (>= /84): 1,000
	
	   NOTE: This restriction only applies to EX models in LSE.
		N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale

Number of IP Longest Prefix Matches (LPM) entries
 	
(Continued)
 	
Note
Except for the maximum LMP scale profile, the total of `(# of IPv4 prefixes) + 2*(# of IPv6 prefixes)` must not exceed the scale listed for IPv4 alone
	High LPM Scale profile:
	
	• IPv4: 128,000 or
	
	• IPv6: 64,000
	
	• IPv6 wide prefixes (>= /84): 1,000
	
	  NOTE: This restriction only applies to EX models in LSE.
	
	Maximum LPM scale profile:
	
	• IPv4: 440,000
	
	• IPv6: 100,000
	
	  NOTE: This profile also supports the combination of 440,000 IPv4 + 100,000 IPv6 prefixes.
	
	High Policy profile:
	
	• LSE2 (except FXP switches):
	
	  • IPv4: 20,000 or
	
	  • IPv6: 10,000
	
	• LSE:
	
	  • IPv4: 8,000
	
	  • IPv6: 4,000
	
	    NOTE: This restriction only applies to EX models in LSE.
	
	High IPv4 EP Scale profile:
	
	• LSE2 (except FXP switches):
	
	  • IPv4: 40,000
	
	  • IPv6: 20,000
	
	• LSE: Not supported
	
	Multicast Heavy profile:
	
	• LSE2 (except FXP switches):
	
	  • IPv4: 20,000
	
	  • IPv6: 10,000
	
	• LSE: Not supported
		N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of Secondary addresses per logical interface	1	1
Number of L3 interfaces per Context	• 1,000 SVIs • 48 Routed interfaces • 100 sub-interfaces with or without port-channel	N/A
Number of L3 interfaces	• 1,000 SVIs • 48 Routed interfaces • 2,000 sub-interfaces with or without port-channel	N/A
Number of ARP entries for L3Outs	7,500	N/A
Shared L3Out	• IPv4 Prefixes: 2,000 or • IPv6 Prefixes: 1,000	• IPv4 Prefixes: 6,000 or • IPv6 Prefixes: 3,000
Number of L3Outs	800	*See* Table 1: Fabric Scale Limits Per Cluster Size

ECMP (Equal Cost MultiPath)

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum ECMP for BGP	64	N/A
Maximum ECMP for OSPF	64	N/A
Maximum ECMP for Static Route	64	N/A
Number of ECMP groups	8,000 **Note** Should not exceed 4,000 in steady state, to allow room for make-before-break transitions (*)	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of ECMP members	Maximum LPM scale profile: • 64,000 **Note** Should not exceed 32,000 in steady state, to allow room for make-before-break transitions (*) All other scale profiles: • 32,000 **Note** Should not exceed 16,000 in steady state, to allow room for make-before-break transitions (*)	N/A
Average number of paths (ECMP) per prefix at maximum LPM scale **Note** Across all prefixes, the average number of equal cost next-hops (ECMP) must not exceed the specified number. Some prefixes may have a higher number of paths as long as it's compensated by other prefixes that have a lower number of paths.	Default (Dual Stack), High Policy and Multicast Heavy profiles: • IPv4: 32 • IPv6: 12 IPv4 scale profile: • IPv4: 16 • IPv6: NA High Dual Stack scale profile: • IPv4: 16 • IPv6: 6 High LPM scale profile: • IPv4: 4 • IPv6: 1 Maximum LPM scale profile: • IPv4: 1.8 • IPv6: 1.8	N/A

Note (*) For more information about managing the equal cost multipath scale, please see *Understand and Manage ECMP Scale in Cisco ACI* at this URL: https://www.cisco.com/c/en/us/solutions/collateral/data-center-virtualization/application-centric-infrastructure/manage-ecmp-scale-aci-wp.html.

External EPGs

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of External EPGs	800	*See* Table 1: Fabric Scale Limits Per Cluster Size
Number of External EPGs per L3Out	250	600 The listed scale is calculated as a product of (Number of external EPGs per L3Out)*(Number of border leaf switches for the L3Out) For examples, 150 external EPGs on `L3Out1` that is deployed on `leaf1`, `leaf2`, `leaf3`, and `leaf4` adds up to a total of 600
Number of LPM Prefixes for External EPG Classification **Note** Maximum combined number of IPv4/IPv6 host and LPM prefixes for External EPG Classification must not exceed 64,000	Refer to LPM scale section.	N/A

Number of host prefixes for External EPG Classification
 	
Note
Maximum combined number of IPv4/IPv6 host and LPM prefixes for External EPG Classification must not exceed 64,000
	Default Profile:
	
	• IPv4 (/32): 16,000
	
	• IPv6 (/128): 12,000
	
	  Combined number of host prefixes and endpoints can't exceed 12,000.
	
	IPv4 Scale profile:
	
	• IPv4 (/32): 16,000
	
	  Combined number of host prefixes, multicast groups, and endpoints can't exceed 56,000.
	
	• IPv6 (/128): 0
		N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale

Number of host prefixes for External EPG Classification
 	
(Continued)
 	
Note
Maximum combined number of IPv4/IPv6 host and LPM prefixes for External EPG Classification must not exceed 64,000
	High Dual Stack Profile:
	
	• LSE:
	
	    • IPv4 (/32): 64,000
	
	    Combined number of host prefixes, multicast groups, and endpoints can't exceed 64,000.
	
	    • IPv6 (/128): 24,000
	
	    Combined number of host prefixes and endpoints can't exceed 24,000.
	
	• LSE2:
	
	    • IPv4 (/32): 64,000
	
	    Combined number of host prefixes, multicast groups, and endpoints can't exceed 64,000.
	
	    • IPv6 (/128): 48,000
	
	    Combined number of host prefixes and endpoints can't exceed 48,000.
	
	High LPM Profile:
	
	• IPv4 (/32): 24,000
	
	Combined number of host prefixes, multicast groups, and endpoints can't exceed 32,000.
	
	• IPv6 (/128): 12,000
	
	Combined number of host prefixes and endpoints can't exceed 12,000.
		N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale

Number of host prefixes for External EPG Classification
 	
Note
Maximum combined number of IPv4/IPv6 host and LPM prefixes for External EPG Classification must not exceed 64,000
 	
(Continued)
	Maximum LPM profile:
	
	• IPv4 (/32): 10,000
	
	Combined number of host prefixes, multicast groups, and endpoints can't exceed 10,000.
	
	• IPv6 (/128): 4,000
	
	Combined number of host prefixes and endpoints can't exceed 4,000.
	
	High Policy profile:
	
	• LSE:
	
	    • IPv4 (/32): 16,000
	
	    Combined number of host prefixes, multicast groups, and endpoints can't exceed 24,000
	
	    • IPv6 (/128): 8,000
	
	    Combined number of host prefixes and endpoints can't exceed 8,000.
	
	• LSE2 (except FXP switches):
	
	    • IPv4 (/32): 16,000
	
	    • IPv6 (/128): 12,000
	
	    Combined number of host prefixes and endpoints can't exceed 12,000.
		N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale

Number of host prefixes for External EPG Classification
 	
Note
Maximum combined number of IPv4/IPv6 host and LPM prefixes for External EPG Classification must not exceed 64,000
 	
(Continued)
	High IPv4 EP Scale profile:
	
	• LSE: Not supported
	
	• LSE2 (except FXP switches):
	
	  • IPv4 (/32): 16,000
	
	  • IPv6 (/128): 12,000
	
	    Combined number of host prefixes and endpoints can't exceed 12,000.
	
	Multicast Heavy profile:
	
	• LSE: Not supported
	
	• LSE2 (except FXP switches):
	
	  • IPv4 (/32): 16,000
	
	    Combined number of host prefixes, multicast groups, and endpoints can't exceed 154,000.
	
	  • IPv6 (/128): 4,000
	
	    Combined number of host prefixes and endpoints can't exceed 4,000.
		N/A

Bridge Domains

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of BDs	1,980 Legacy mode: 3,500	15,000
Number of BDs with Unicast Routing per Context (VRF)	1,000	1,750
Number of subnets per BD	1,000, cannot be for all BDs	1,000 per BD
Number of EPGs per BD	3,960	4,000
BD with Flood in Encapsulation: maximum number of replications (= EPG VLANs * ports)	The sum of all EPG VLANs * ports (i.e., VLAN "replications") for all EPG in a given BD with Flood in Encapsulation enabled must be less than 1,500	N/A
Number of L2 Outs per BD	1	1

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of BDs with Custom MAC Address	1,000	1,000
Number of EPGs + L3Outs per Multicast Group	128	128
Number of BDs with L3 Multicast enabled	1,750	1,750
Number of VRFs with L3 Multicast enabled	64	300
Number of L3Outs per BD	16	N/A
Number of static routes behind pervasive BD (EP reachability)	N/A	450
DHCP relay addresses per BD across all labels	16	N/A
DHCP Relay: maximum number of replications (= EPG VLANs * ports)	The maximum number of VLAN encapsulations * ports in a BD with DHCP relay enabled should be less than 1,500	N/A
ICMPv6 ND: maximum number of replications (= EPG VLANs * ports)	The maximum number of VLAN encapsulations * ports in a BD should be less than 1,500	N/A
Number of external EPGs per L2 out	1	1
Number of PIM Neighbors	1,000	1,000
Number of PIM Neighbors per VRF	64	64
Number of L3Out physical interfaces with PIM enabled	32	N/A

Endpoint Groups (Under App Profiles)

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of EPGs	Normally 3,960; if legacy mode 3,500	15,000
Maximum amount of encapsulations per EPG	1 Static leaf binding, plus 10 Dynamic VMM	N/A
Maximum Path encap binding per EPG	Equals to number of ports on the leaf	N/A
EPGs with Flood in Encapsulation: maximum number of replications (= EPG VLANs * ports)	The sum of all EPG VLANs * ports (i.e., VLAN "replications") for all EPG with Flood in Encapsulation enabled in a given BD must be less than 1,500	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum amount of encapsulations per EPG per port with static binding	One (path or leaf binding)	N/A
Number of domains (physical, L2, L3)	100	N/A
Number of VMM domains	N/A	200 VDS
Number of native encapsulations	• One per port, if a VLAN is used as a native VLAN. • Total number of ports, if there is a different native VLAN per port.	Applicable to each leaf independently
Number of 802.1p encapsulations	• 1, if path binding then equals the number of ports. • If there is a different native VLAN per port, then it equals the number of ports.	Applicable to each leaf independently
Can encapsulation be tagged and untagged?	No	N/A
Number of Static endpoints per EPG	Maximum endpoints	N/A
Number of Subnets for inter-context access per tenant	4,000	N/A
Number of Taboo Contracts per EPG	2	N/A
IP-based EPG (bare metal)	4,000	N/A
MAC-based EPG (bare metal)	4,000	N/A

Contracts

Configurable Options	Per Leaf Scale	Per Fabric Scale
Security TCAM size	Default scale profile: 64,000 IPv4 scale profile: 64,000 High Dual Stack scale profile: • LSE: 8,000 • LSE2: 128,000 High LPM scale profile: 8,000 Maximum LPM scale profile: 8,000 High Policy profile: • LSE: 100,000 • LSE2 (with 24GB of RAM): 140,000 • LSE2 (with 32 GB of RAM): 256,000 High IPv4 EP Scale profile: • LSE: Not supported • LSE2 (except FXP switches): 64,000 Multicast Heavy profile: • LSE: Not supported • LSE2 (except FXP switches): 64,000	N/A
Software policy scale with Policy Table Compression enabled (Number of `actrlRule` Managed Objects)	Dual stack profile: 80,000 (except EX switches) High Dual Stack profile: • LSE: Not supported • LSE2: 140,000 High Policy profile: • LSE (except EX switches): 100,000 • LSE2 (with 24GB of RAM) : 140,000 • LSE2 (with 32 GB of RAM) : 256,000	N/A
Approximate TCAM calculator given contracts and their use by EPGs	Number of entries in a contract X Number of Consumer EPGs X Number of Provider EPGs X 2	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of consumers (or providers) of a contract that has more than 1 provider (or consumer)	100	100
Number of consumers (or providers) of a contract that has a single provider (or consumer)	1,000	1,000
Scale guideline for the number of Consumers and Providers for the same contract	N/A	Number of consumer EPGs * number of provider EPGs * number of filters in the contract <= 50,000 This scale limit is per contract. If the limit is exceeded, the configuration is rejected. If 90% of the limit is reached, fault returns.
Number of rules for consumer/provider relationships with in-band EPG	400	N/A
Number of rules for consumer/provider relationships with out-of-band EPG	400	N/A

Endpoint Security Groups (ESG)

Configurable Options	Scale
Number of ESGs per Fabric	10,000
Number of ESGs per VRF	4,000
Number of ESGs per Tenant	4,000
Number of L2 MAC Selectors per Leaf	5,000
Number of L3 IP Selectors per Leaf	5,000

Fiber Channel over Ethernet N-Port Virtualization (FCoE NPV)

Configurable Options	Per Leaf Scale
Number of VSANs	32
Number of VFCs configured on physical ports and FEX ports	151
Number of VFCs on port-channel (PC), including SAN port-channel	7
Number of VFCs on virtual port-channel (vPC) interfaces, including FEX HIF vPC	151

Configurable Options	Per Leaf Scale
Number of FDISC per port	255
Number of FDISC per leaf	1,000

Fiber Channel N-Port Virtualization (FC NPV)

Configurable Options	Per Leaf Scale
Number of FC NP Uplink interfaces	48
Number of VSANs	32
Number of FDISC per port	255
Number of FDISC per leaf	1,000
Number of SAN port-channel, including VFC port-channel	7
Number of members in a SAN port-channel	16

VMM Scalability Limits

VMware

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of vCenters (VDS)	N/A	200 (Verified with a load of 10 events/minute for each vCenter)
Datacenters in a vCenter	N/A	15
Total Number of VMM domain (vCenter, Datacenter) instances.	N/A	200 VDS
Number of EPGs per vCenter/vDS	N/A	5,000
Number of EPGs to VMware domains/vDS	N/A	5,000
Number of endpoints per VDS	10,000	10,000
Number of endpoints per vCenter	10,000	10,000
Support RBAC for VDS	N/A	Yes
Number of Microsegment EPGs with vDS	400	N/A
Number of VM Attribute Tags per vCenter	N/A	vCenter version 6.0: 500 vCenter version 6.5: 1,000

Microsoft SCVMM

Configurable Options	Per Leaf Scale (On-Demand Mode)	Per Leaf Scale (Pre-Provision Mode)	Per Fabric Scale
Number of controllers per SCVMM domain	N/A	N/A	5
Number of SCVMM domains	N/A	N/A	25
EPGs per Microsoft VMM domain	N/A	N/A	3,000
EPGs per all Microsoft VMM domains	N/A	N/A	9,000
EP/VNICs per HyperV host	N/A	N/A	100
EP/VNICs per SCVMM	3,000	10,000	10,000
Number of Hyper-V hosts	64	N/A	N/A
Number of logical switch per host	N/A	N/A	1
Number of uplinks per logical switch	N/A	N/A	4
Microsoft micro-segmentation	1,000	Not Supported	N/A

Microsoft Windows Azure Pack

Configurable Options	Per Fabric Scale
Number of Windows Azure Pack subscriptions	1,000
Number of plans per Windows Azure Pack instance	150
Number of users per plan	200
Number of subscriptions per user	3
VM networks per Windows Azure Pack user	100
VM networks per Windows Azure Pack instance	3,000
Number of tenant shared services/providers	40
Number of consumers of shared services	40
Number of VIPs (Citrix)	50
Number of VIPs (F5)	50

Layer 4 - Layer 7 Scalability Limits

Configurable Options (L4-L7 Configurations)	Per Fabric Scale
Number of L4-L7 concrete devices	1,200
Number of graph instances	1,000
Number of device clusters per tenant	30
Number of interfaces per device cluster	Any
Number of graph instances per device cluster	500
Deployment scenario for ASA (transparent or routed)	Yes
Deployment scenario for Citrix - One arm with SNAT/etc.	Yes
Deployment scenario for F5 - One arm with SNAT/etc.	Yes

AD, TACACS, RBAC Scalability Limits

Configurable Options	Per Fabric Scale
Number of ACS/AD/LDAP authorization domains	4 tested (16 maximum /server type)
Number of login domains	15
Number of security domains/APIC	15
Number of security domains in which the tenant resides	4
Number of priorities	4 (16 per domain)
Number of shell profiles that can be returned.	4 (32 domains total)
Number of users	8,000 local / 8,000 remote
Number of simultaneous logins	500 connections / NGNIX simultaneous REST logins

Cisco Mini ACI Fabric Scalability Limits

Property	Maximum Scale
Number of spine switches	2
Number of leaf switches	4

Property	Maximum Scale
Number of Pods	1
Number of tenants	25
Number of VRFs	25
Number of bridge domains (BDs)	1000
Number of endpoint groups (EPGs)	1000
Number of endpoints	20,000
Number of contracts	2000
Number of service graph instances	20
Number of L4-L7 logical device clusters	3 Physical or 10 Virtual
Number of multicast groups	200
Number of BGP+OSPF sessions	25
GOLF VRF, Route Scale	N/A

Cisco ACI and UCSM Scalability

The following table shows verified scalability numbers for Cisco Unified Computing System with Cisco ACI `ExternalSwitch` app.

Configurable Options	Scale
Number of UCSMs per APIC cluster	12
Number of VMM Domains per UCSM	4
Number of VLANs + PVLAN per UCSM	4,000
Number of vNIC Templates per UCSM	16

QoS Scalability Limits

The following table shows QoS scale limits. The same numbers apply for topologies with or without remote leafs as well as with COS preservation and MPOD policy enabled.

QoS Mode	QoS Scale
Custom QoS Policy with DSCP	7
Custom QoS Policy with DSCP and Dot1P	7
Custom QoS Policy with Dot1P	38

QoS Mode	QoS Scale
Custom QoS Policy via a Contract	38

PTP Scalability Limits

The following table shows Precision Time Protocol (PTP) scale limits.

Configurable Options	Scale (IEEE 1588 Default Profile)	Scale (AES67, SMPTE-2059-2)	Scale (Telecom Profile G.8275.1)
Number of leaf switches connected to a single spine with PTP globally enabled	288	40	N/A
Number of PTP peers per leaf switch	52	26	25
Number of ACI switches connected to the same tier-1 leaf switch (multi-tier topology) with PTP globally enabled	Within the range of the "Number of PTP peers per leaf switch" above	16	N/A
Number of access ports with PTP enabled on a leaf switch	Within the range of the "Number of PTP peers per leaf switch" above **Note** For improved performance on 1G interfaces with N9K-C93108TC-FX3P switches, the maximum number of 1G interfaces should not exceed 10	25 **Note** For improved performance on 1G interfaces with N9K-C93108TC-FX3P switches, the maximum number of 1G interfaces should not exceed 10 out of 25	24
Number of PTP peers per access port	PTP Mode Multicast (Dynamic/Master): 2 peers PTP Mode Unicast Master: 1 peer	PTP Mode Multicast (Dynamic/Master): 2 peers PTP Mode Unicast Master: 1 peer	1

NetFlow Scale

Configurable Options	Scale
Exporters per leaf switch	2
NetFlow monitor policies under BDs per leaf switch	100

Configurable Options	Scale
NetFlow monitor policies under L3Outs per leaf switch	120
Number of records per collect interval	20,000