



Verified Scalability Guide for Cisco APIC, Release 6.1(5) and Cisco Nexus 9000 Series ACI-Mode Switches, Release 16.1(5)

Overview 2

New and Changed Information 2

General Scalability Limits 2

Multiple Fabric Options Scalability Limits 7

Cisco Multi-Site Scalability Limits 8

Fabric Topology, SPAN, Tenants, Contexts (VRFs), Equal Cost Multipath (ECMP), External EPGs, Bridge Domains, Endpoints, and Contracts Scalability Limits 9

VMM Scalability Limits 35

Layer 4 to Layer 7 Services Scalability Limits 37

AD, TACACS, RBAC Scalability Limits 37

Cisco Mini ACI and Virtual APIC Small Profile Scalability Limits 38

QoS Scalability Limits 38

PTP Scalability Limits 39

NetFlow Scale 40

Revised: May 7, 2026

Overview

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

- Cisco Application Policy Infrastructure Controller (Cisco APIC), releases 6.1(5)
- Cisco Nexus 9000 Series ACI-Mode Switches, releases 16.1(5)

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 this URL: <https://www.cisco.com/c/en/us/support/cloud-systems-management/multi-site-orchestrator/products-device-support-tables-list.html>.

New and Changed Information

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

Date	Changes
December 15, 2025	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, because it interconnects separate availability zones (fabrics) and managed by an independent APIC 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. The maximum number of remote leaf switch pairs that can be added to a single autonomous remote leaf (ARL) switch group is 5.

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.



Note The Cisco Virtual APIC supports all fabric sizes - default, medium, and large - when used along with the minimum requirements listed in the "Virtual Machine Prerequisites" section of [Deploying Cisco Virtual APIC Using VMware vCenter](#).

Default and medium sized fabrics require virtual APIC medium or large profile. For large fabrics, the virtual APIC large profile is required.

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

Configurable Options	Default Fabric	Medium Fabric	Large Fabric	
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	
<p>Number of external EPGs across all BLs</p> <p>This is calculated as a product of (Number of external EPGs)*(Number of border leaf switches for the L3Out).</p> <p>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):</p> <ul style="list-style-type: none"> • 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	<p>General limits:</p> <ul style="list-style-type: none"> • Single-tenant fabrics: 4,000 • Multi-tenant fabrics: 500 <p>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):</p> <ul style="list-style-type: none"> • Use case 1: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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

ACI Border Gateways (BGW)

With the Cisco ACI border gateway (BGW) solution, you can now have a seamless extension of Virtual Routing and Forwarding (VRFs), Bridge Domains (BDs), and security policies between fabrics. The Cisco ACI BGW is a node that interacts with nodes within a site and with nodes that are external to the site. The Cisco ACI BGW feature can be conceptualized as multiple site-local EVPN control planes and IP forwarding domains interconnected by a single common EVPN control and forwarding domain. For more information, see the [Cisco APIC Layer 3 Networking Configuration Guide, Release 6.1\(x\)](#) content.

Configurable Options	Scale Limit
Number of remote fabrics	3
Number of remote Border Gateways (BGWs)	6
Number of ACI Pods	3
Number of Border Gateways (BGWs) per Pod	2
Number of stretched BDs	2,000
Number of stretched VRFs	500
Number of endpoints in the ACI fabric (Proxy database entries), including type-2 routes from remote fabrics	100,000 (36,000 MAC + 36,000 IPv4 + 28,000 IPv6)
Number of type-5 routes across all fabrics (local and remote)	29,000 (20,000 IPv4 + 9,000 IPv6 = 38,000 LPM)
Number of contract policy rules on the BGW	64,000
Number of ESGs across all stretched VRFs	5,000
Number of ESGs on a single VRF or Tenant	3,000
Number of VXLAN BD selectors	2,000
Number of external subnet selectors across all stretched VRFs	10,000
Number of external subnet selectors under an ESG	1,000
Number of external EPG selectors	1,000
Number of External EPG Subnets from user tenant L3Out	8,000

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 this 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	<ul style="list-style-type: none">• 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	<ul style="list-style-type: none"> • 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 • N9K-C9332D-H2R • N9K-C93400LD-H1 • N9K-C9364C-H1



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, 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
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 ¹	32	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)	<p>Default profile or High LPM profile:</p> <ul style="list-style-type: none"> • MAC: 24,000 • IPv4: 24,000 • IPv6: 12,000 <p>Maximum LPM profile:</p> <ul style="list-style-type: none"> • MAC: 8,000 • IPv4: 8,000 • IPv6: 4,000 <p>IPv4 scale profile:</p> <ul style="list-style-type: none"> • MAC: 48,000 • IPv4: 48,000 • IPv6: Not supported <p>High Dual Stack scale profile:</p> <ul style="list-style-type: none"> • LSE: <ul style="list-style-type: none"> • MAC: 64,000 • IPv4: 64,000 • IPv6: 24,000 • LSE2: <ul style="list-style-type: none"> • MAC: 64,000 • IPv4: 64,000 • IPv6: 48,000 	<p>16-slot and 8-slot modular spine switches:</p> <p>Max. 450,000 Proxy Database Entries in the fabric, which can be translated into any one of these:</p> <ul style="list-style-type: none"> • 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) <p>The formula to calculate in mixed mode is:</p> $\#MAC + \#IPv4 + \#IPv6 \leq 450,000$ <p>Note: Four fabric modules are required on all spines in the fabric to support above scale.</p>

Configurable Options	Per Leaf Scale	Per Fabric Scale
		<p>4-slot modular spine switches:</p> <p>Max. 360,000 Proxy Database Entries in the fabric, which can be translated into any one of these:</p> <ul style="list-style-type: none"> • 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) <p>The formula to calculate in mixed mode is:</p> $\#MAC + \#IPv4 + \#IPv6 \leq 360,000$ <p>Note: Four fabric modules are required on all spines in the fabric to support above scale.</p>

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of endpoints (EPs) (Continued)	<p>High Policy profile:</p> <ul style="list-style-type: none"> • LSE2 (except FXP switches): <ul style="list-style-type: none"> • MAC: 24,000 • IPv4: 24,000 • IPv6: 12,000 • LSE: <ul style="list-style-type: none"> • MAC: 16,000 • IPv4: 16,000 • IPv6: 8,000 <p>High IPv4 EP Scale profile:</p> <ul style="list-style-type: none"> • LSE: Not supported • LSE2 (with 32GB of RAM): <ul style="list-style-type: none"> • MAC: 24,000 • IPv4 local: 24,000 • IPv4 total: 280,000 • IPv6: 12,000 <p>Multicast Heavy profile:</p> <ul style="list-style-type: none"> • LSE: Not supported • LSE2 (except FXP switches): <ul style="list-style-type: none"> • MAC: 24,000 • IPv4 local: 24,000 • IPv4 total: 64,000 • IPv6: 4,000 	<p>Fixed spine switches:</p> <p>Max. 180,000 Proxy Database Entries in the fabric, which can be translated into any one of these:</p> <ul style="list-style-type: none"> • 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) <p>The formula to calculate in mixed mode is:</p> $\#MAC + \#IPv4 + \#IPv6 \leq 180,000$

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of Multicast Routes	<p>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</p> <p>Maximum LPM profile:</p> <ul style="list-style-type: none"> • 1,000 with (S,G) scale not exceeding 500 <p>High Dual Stack profile:</p> <ul style="list-style-type: none"> • LSE: 512 • LSE2: 32,000 with (S,G) scale not exceeding 16,000 <p>Multicast Heavy profile:</p> <ul style="list-style-type: none"> • LSE: not supported • LSE2 (with 32GB of RAM): 90,000 with (S,G) scale not exceeding 72,000 	128,000
Number of Multicast Routes per VRF	<p>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</p> <p>Maximum LPM profile:</p> <ul style="list-style-type: none"> • 1,000 with (S,G) scale not exceeding 500 <p>High Dual Stack profile:</p> <ul style="list-style-type: none"> • LSE: 512 • LSE2: 32,000 with (S,G) scale not exceeding 16,000 <p>Multicast Heavy profile:</p> <ul style="list-style-type: none"> • LSE: not supported • LSE2 (except FXP switches): 32,000 	32,000

Configurable Options	Per Leaf Scale	Per Fabric Scale
IGMP snooping L2 multicast routes <ul style="list-style-type: none"> • For IGMPv2, route scale is for (*, G) only • For IGMPv3, route scale is for both (S, G) and (*, G) <p>Note IGMP snooping entries are created per BD (2 receivers that join the same group from 2 different BDs consume 2 separate entries).</p>	Default (Dual Stack), IPv4, High LPM, High Policy, or High IPv4 EP scale profiles: 8,000 Maximum LPM profile: <ul style="list-style-type: none"> • 1,000 High Dual Stack profile: <ul style="list-style-type: none"> • LSE: 512 • LSE2: 32,000 Multicast Heavy profile: <ul style="list-style-type: none"> • 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
Number of SPAN sessions for all SPAN types (fabric, access, and tenant)	32 unidirectional or 16 bidirectional sessions in the filter mode 8 out of 32 unidirectional or 4 out of 16 bidirectional in the scale mode	N/A
Number of ports per SPAN session <p>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.</p> <p>The number of ports in SPAN session in the scale mode are not counted to this number.</p>	63 – total number of unique ports (fabric + access) across all types of span sessions	N/A
Number of ports in SPAN session in the scale mode	96	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of SPAN sources in each direction	$2 * (V + FP + AP1 + AP2 + (V * AP1) + AP3_v6) + AP3_v4 \leq 480$ Where: <ul style="list-style-type: none"> • 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 may 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 SPAN ACL filter TCAM entries SPAN filters are not supported in these: <ul style="list-style-type: none"> • Fabric ports • Fabric and tenant SPAN sessions • Spine switches 	<ul style="list-style-type: none"> • IPv4: 480 • IPv6: 240 Total number of TCAM entries is calculated using this formula: $(IPv4-filters) * (IPv4-filter-source-groups) + 2 * (IPv6-filters) * (IPv6-filter-source-groups) + 2 * (no-filter-source-groups)$	N/A
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

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of Data Plane policers at the interface level	<ul style="list-style-type: none"> • 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
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 <code>vsh_lc -c 'show system internal aclqos pppi copp tcam-usage'</code> command to check on LSE/LSE2 platforms	N/A
Number of SNMP trap receivers	10	10
IP SLA probes ²	200	<ul style="list-style-type: none"> • 1,500 (for PBR tracking) • 400 (for static route tracking)
First Hop Security (FHS) ³	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
Number of Type 7 Keys Note Keys were tested with OSPFv2 and OSPFv3 clients. For OSPFv2, the key-ID supported value is from 0 to 255 only. For OSPFv3 the maximum number of IPsec policies supported is 128.	20,000	20,000
Number of Keychains	2,000	2,000

¹vPC total ports = 64, 32 per leaf

²With 1 second probe time and 3 seconds of a timeout

³With any combination of BDs/EPGs/EPs within the supported limit

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 ¹	N/A	100 total, 2 per RL location
VRFs ¹	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)

¹Including both remote leaf and multi-pod

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.

All numbers are qualified with default ACI timer values unless explicitly called out.

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of contexts (VRFs)	<p>Default (Dual Stack) scale profile:</p> <ul style="list-style-type: none"> Switches with 32GB of RAM: 2,000 Other switches: 800 <p>High Dual Stack, High LPM, High Policy scale profiles:</p> <ul style="list-style-type: none"> LSE2 switches with 32GB of RAM: 2,000 Other switches: 800 <p>Maximum LPM scale profile:</p> <ul style="list-style-type: none"> LSE2 switches with 32GB of RAM: 250 Other switches: not supported <p>Multicast heavy, IPv4 and High IPv4 EP scale:</p> <ul style="list-style-type: none"> All switch models: 800 	See Table 1: Fabric Scale Limits Per Cluster Size
Number of isolated EPGs	400	400
Border leaf switches per L3Out	N/A	24 Note qualified with 100 VRFs + 16,000 IPv4 + 6400 IPv6 external prefixes
Number of vzAny provided contracts	<p>Shared services: Not supported</p> <p>Non-shared services: 70 per Context (VRF)</p>	N/A
Number of vzAny consumed contracts	<p>Shared services: 16 per Context (VRF)</p> <p>Non-shared services: 70 per Context (VRF)</p>	N/A
Number of graph 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: <ul style="list-style-type: none"> • minTx:50 • minRx:50 • multiplier:3 257-2,000 sessions using these minimum BFD timers: <ul style="list-style-type: none"> • minTx:300 • minRx:300 • multiplier:3 	N/A
Number of BGP neighbors Note With BGP aggressive timers (1s or 3s), if the system is not stable such as if ARP or North-South floods keep the SUP busy, these timers may lead to BGP neighbor flaps. Cisco recommends that you use BGP default timers and leverage bidirectional forwarding detection (BFD) to get sub-second failure detection. BGP route-paths calculation example: 200 BGP routes over 5-way ECMP = 1,000 BGP route-paths	2,000 with up to 70,000 route-paths <ul style="list-style-type: none"> • Keepalive interval: 60s • Hold interval: 180s 1,000 with up to 30,000 route-paths <ul style="list-style-type: none"> • Keepalive interval: 20s • Hold interval: 60s 20 with up to 1,000 route-paths <ul style="list-style-type: none"> • Keepalive interval: 1s • Hold interval: 3s Note These numbers are mutually exclusive.	20,000

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of OSPF neighbors	<p>Up to 700 with up to 10,000 external prefixes using these timers:</p> <ul style="list-style-type: none"> • Hello timer of 10 seconds • Dead timer of 40 seconds • No more than 300 OSPF neighbors per VRF <p>701-2,000 with up to 35,000 external prefix using these timers:</p> <ul style="list-style-type: none"> • 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
<p>Number of IP Longest Prefix Matches (LPM) entries</p> <p>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</p>	<p>Default (Dual Stack) profile:</p> <ul style="list-style-type: none"> • IPv4: 20,000 or • IPv6: 10,000 • IPv6 wide prefixes ($\geq /84$): 1,000 <p>Note: This restriction only applies to EX models in LSE.</p> <p>IPv4 scale profile:</p> <ul style="list-style-type: none"> • IPv4: 38,000 • IPv6: Not supported <p>High Dual Stack scale profile:</p> <ul style="list-style-type: none"> • IPv4: 38,000 or • IPv6: 19,000 • IPv6 wide prefixes ($\geq /84$): 1,000 <p>Note: This restriction only applies to EX models in LSE.</p>	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
<p>Number of IP Longest Prefix Matches (LPM) entries (Continued)</p> <p>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</p>	<p>High LPM Scale profile:</p> <ul style="list-style-type: none"> • IPv4: 128,000 or • IPv6: 64,000 • IPv6 wide prefixes ($\geq /84$): 1,000 <p>Note: This restriction only applies to EX models in LSE.</p> <p>Maximum LPM scale profile:</p> <ul style="list-style-type: none"> • IPv4: 440,000 • IPv6: 100,000 <p>Note: This profile also supports the combination of 440,000 IPv4 + 100,000 IPv6 prefixes.</p> <p>High Policy profile:</p> <ul style="list-style-type: none"> • LSE2 (except FXP switches): <ul style="list-style-type: none"> • IPv4: 20,000 or • IPv6: 10,000 • LSE: <ul style="list-style-type: none"> • IPv4: 8,000 • IPv6: 4,000 <p>Note: This restriction only applies to EX models in LSE.</p> <p>High IPv4 EP Scale profile:</p> <ul style="list-style-type: none"> • LSE2 (except FXP switches): <ul style="list-style-type: none"> • IPv4: 40,000 • IPv6: 20,000 • LSE: Not supported <p>Multicast Heavy profile:</p> <ul style="list-style-type: none"> • LSE2 (except FXP switches): <ul style="list-style-type: none"> • IPv4: 20,000 • IPv6: 10,000 • LSE: Not supported 	<p>N/A</p>

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of Secondary addresses per logical interface	1	1
Number of L3 interfaces per Context	<ul style="list-style-type: none"> • 1,000 SVIs • 48 Routed interfaces • 100 sub-interfaces with or without port-channel 	N/A
Number of L3 interfaces	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • IPv4 Prefixes: 2,000 or • IPv6 Prefixes: 1,000 	<ul style="list-style-type: none"> • IPv4 Prefixes: 6,000 or • IPv6 Prefixes: 3,000
Number of L3Outs	2,000	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	128	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: <ul style="list-style-type: none"> • 64,000 Note Should not exceed 32,000 in steady state, to allow room for make-before-break transitions (*) All other scale profiles: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • IPv4: 32 • IPv6: 12 IPv4 scale profile: <ul style="list-style-type: none"> • IPv4: 16 • IPv6: NA High Dual Stack scale profile: <ul style="list-style-type: none"> • IPv4: 16 • IPv6: 6 High LPM scale profile: <ul style="list-style-type: none"> • IPv4: 4 • IPv6: 1 Maximum LPM scale profile: <ul style="list-style-type: none"> • IPv4: 1.8 • IPv6: 1.8 	N/A



Note (*) For more information about managing the equal cost multipath scale, 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	<ul style="list-style-type: none"> Switches with 32GB of RAM: 2,000 Other switches: 800 	See Table 1: Fabric Scale Limits Per Cluster Size
Number of External EPGs per L3Out	250	<p>600</p> <p>The listed scale is calculated as a product of (Number of external EPGs per L3Out)*(Number of border leaf switches for the L3Out)</p> <p>For examples, 150 external EPGs on L3Out1 that is deployed on leaf1, leaf2, leaf3, and leaf4 adds up to a total of 600</p>
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: <ul style="list-style-type: none"> IPv4 (/32): 16,000 IPv6 (/128): 12,000 Combined number of host prefixes and endpoints can't exceed 12,000 IPv4 Scale profile: <ul style="list-style-type: none"> IPv4 (/32): 16,000 Combined number of host prefixes, multicast groups, and endpoints can't exceed 56,000 <ul style="list-style-type: none"> IPv6 (/128): 0 	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
<p>Number of host prefixes for External EPG Classification</p> <p>(Continued)</p> <p>Note Maximum combined number of IPv4/IPv6 host and LPM prefixes for External EPG Classification must not exceed 64,000</p>	<p>High Dual Stack Profile:</p> <ul style="list-style-type: none"> • LSE: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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 <p>High LPM Profile:</p> <ul style="list-style-type: none"> • 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 	<p>N/A</p>

Configurable Options	Per Leaf Scale	Per Fabric Scale
<p>Number of host prefixes for External EPG Classification</p> <p>Note Maximum combined number of IPv4/IPv6 host and LPM prefixes for External EPG Classification must not exceed 64,000</p> <p>(Continued)</p>	<p>Maximum LPM profile:</p> <ul style="list-style-type: none"> • IPv4 (/32): 10,000 <p>Combined number of host prefixes, multicast groups, and endpoints can't exceed 10,000</p> <ul style="list-style-type: none"> • IPv6 (/128): 4,000 <p>Combined number of host prefixes and endpoints can't exceed 4,000</p> <p>High Policy profile:</p> <ul style="list-style-type: none"> • LSE: <ul style="list-style-type: none"> • IPv4 (/32): 16,000 <p>Combined number of host prefixes, multicast groups, and endpoints can't exceed 24,000</p> • IPv6 (/128): 8,000 <p>Combined number of host prefixes and endpoints can't exceed 8,000.</p> <ul style="list-style-type: none"> • LSE2 (except FXP switches): <ul style="list-style-type: none"> • IPv4 (/32): 16,000 • IPv6 (/128): 12,000 <p>Combined number of host prefixes and endpoints can't exceed 12,000</p> 	<p>N/A</p>

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: <ul style="list-style-type: none"> • LSE: Not supported • LSE2 (except FXP switches): <ul style="list-style-type: none"> • IPv4 (/32): 16,000 • IPv6 (/128): 12,000 Combined number of host prefixes and endpoints can't exceed 12,000 Multicast Heavy profile: <ul style="list-style-type: none"> • LSE: Not supported • LSE2 (except FXP switches): <ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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	<ul style="list-style-type: none"> • 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 uSeg EPG (bare metal)	4,000	N/A
MAC-based uSeg EPG (base metal)	4,000	N/A

Contracts

Configurable Options	Per Leaf Scale	Per Fabric Scale
Security TCAM size	<p>Default scale profile: 64,000</p> <p>IPv4 scale profile: 64,000</p> <p>High Dual Stack scale profile:</p> <ul style="list-style-type: none"> • LSE: 8,000 • LSE2: 128,000 <p>High LPM scale profile:</p> <ul style="list-style-type: none"> • LSE2 switches with 32GB of RAM: 32,000 • Other switches: 8,000 <p>Maximum LPM scale profile: 8,000</p> <p>High Policy profile:</p> <ul style="list-style-type: none"> • LSE: 100,000 • LSE2 (with 24GB of RAM): 140,000 • LSE2 (with 32GB of RAM): 256,000¹ <p>High IPv4 EP Scale profile:</p> <ul style="list-style-type: none"> • LSE: Not supported • LSE2 (except FXP switches): 64,000 <p>Multicast Heavy profile:</p> <ul style="list-style-type: none"> • LSE: Not supported • LSE2 (except FXP switches): 64,000 	N/A
<p>Software policy scale with Policy Table Compression enabled</p> <p>(Number of <code>actrlRule</code> Managed Objects)</p>	<p>Dual stack profile: 80,000 (except EX switches)</p> <p>High Dual Stack profile:</p> <ul style="list-style-type: none"> • LSE: Not supported • LSE2: 140,000 <p>High Policy profile:</p> <ul style="list-style-type: none"> • LSE (except EX switches): 100,000 • LSE2 (with 24GB of RAM) : 140,000 • LSE2 (with 32GB of RAM) : 256,000¹ 	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
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
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



Note For more information on scalability considerations, refer to scalability consideration section of [Cisco ACI Contract Guide White Paper](#).

If a vzAny is a shared services consumer, then any EPGs in the consumer VRF that are also shared services providers are treated as individual contract consumers. All ESGs in the consumer VRF are also treated as individual contract consumers since ESGs always get global pcTags.

¹High Policy profile scale number with LSE2 (with 32GB of RAM) is supported with the policy compression applied to 100% of the rules. If your contract configuration includes options where policy rule compression does not apply (for example contract subject labels, unidirectional rules, or contract subject exception), the achievable scale will be lower than the verified numbers shown above. See the [Cisco ACI Contract Guide White Paper](#) for details and guidance.

Endpoint Security Groups (ESG)

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of ESGs	4,000	10,000
Number of ESGs per VRF	4,000	4,000
Number of ESGs per Tenant	4,000	4,000

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of MAC-based Selectors (e.g., Tag Selectors with Policy Tags on MAC Addresses or VM tags)	5,000	N/A
Number of IP-based Selectors (e.g., IP subnet selectors or Tag selectors with Policy Tags on IP addresses)	5,000	N/A
Number of external subnet selectors	64,000	N/A
Number of external subnet selectors under an ESG	8,000	N/A
Number of external EPG selectors	1,300	N/A

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
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	30,000 The qualified number combination of each Resolution Immediacy is listed below: Immediate: 30,000 out of 30,000 On-Demand: 10,000 out of 30,000 Pre-provision: 10,000 out of 30,000
Number of endpoints per vCenter	10,000	30,000 The qualified number combination of each Resolution Immediacy is listed below: Immediate: 30,000 out of 30,000 On-Demand: 10,000 out of 30,000 Pre-provision: 10,000 out of 30,000
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
Total number of endpoints	10,000	30,000 The qualified number combination of each Resolution Immediacy is listed below: Immediate: 30,000 out of 30,000 On-Demand: 10,000 out of 30,000 Pre-provision: 10,000 out of 30,000

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

Nutanix

Configurable Options	Per Fabric Scale
Total Number of Prism Central	10
Total Number of Nutanix domain instances	10
Number of EPGs per Prism Central	500
Number of EPGs per Nutanix domain	500
Number of endpoints per Prism Central (or Nutanix domain)	1,000
Number of VM Attribute Tags per Prism Central	500
Intra EPG isolation support per Prism Central	300 EPGs

Layer 4 to Layer 7 Services Scalability Limits

Configurable Options	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)

Configurable Options	Per Fabric Scale
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 and Virtual APIC Small Profile Scalability Limits

Mini ACI and Virtual APIC small profile supports the scale limits listed in the table. For details on virtual APIC small profile, see the requirements listed for small profile in the "Virtual Machine Prerequisites" section of [Deploying Cisco Virtual APIC Using VMWare vCenter](#) document.

For details on Mini ACI, see Cisco Mini ACI Fabric here: <https://www.cisco.com/c/en/us/td/docs/switches/datacenter/aci/apic/sw/kb/cisco-mini-aci-fabric.html>.

Configurable Options	Mini ACI Scale Limits	vAPIC Small Profile Scale Limits
Number of Pods	1	3
Number of spine switches	2	12
Number of leaf switches	4	40
Number of tenants	25	25
Number of VRFs	25	25
Number of bridge domains (BDs)	1,000	1,000
Number of endpoint groups (EPGs)	1,000	1,000
Number of endpoints	20,000	20,000
Number of contracts	2,000	2,000
Number of service graph instances	20	20
Number of L4-L7 logical device clusters	3 Physical or 10 Virtual	3 Physical or 10 Virtual
Number of multicast groups	200	200
Number of BGP + OSPF sessions	25	25
GOLF VRF, Route Scale	N/A	N/A

QoS Scalability Limits

This 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
Custom QoS Policy via a Contract	38

PTP Scalability Limits

This 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	Per Leaf Scale
Number of exporters	2
Number of monitor policies under bridge domains	EX switches: 100 All other models: 350 ¹
Number of monitor policies under L3Outs	EX switches: 100 All other models: 350 ¹
Number of records per collect interval	EX switches: 20,000 All other models: 1,000,000 ²

¹The total number of monitor policies under bridge domains and L3Outs must not exceed 350 (100 for EX switches).

²For more information, see *Cisco APIC and NetFlow*.



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA 95134-1706
USA

Asia Pacific Headquarters
CiscoSystems(USA)Pte.Ltd.
Singapore

Europe Headquarters
CiscoSystemsInternationalBV
Amsterdam,TheNetherlands

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