



Verified Scalability Guide for Cisco APIC, Release 5.2(1) and Cisco Nexus 9000 Series ACI-Mode Switches, Release 15.2(1)

[Overview](#) 2

[New and Changed Information](#) 2

[General Scalability Limits](#) 3

[Multiple Fabric Options Scalability Limits](#) 6

[Cisco ACI vPod Scalability Limits](#) 7

[Cisco Multi-Site Scalability Limits](#) 7

[Fabric Topology, SPAN, Tenants, Contexts \(VRFs\), External EPGs, Bridge Domains, Endpoints, and Contracts Scalability Limits](#) 7

[VMM Scalability Limits](#) 33

[Layer 4 - Layer 7 Scalability Limits](#) 35

[AD, TACACS, RBAC Scalability Limits](#) 36

[Cisco Mini ACI Fabric and Virtual APICs Scalability Limits](#) 36

[Cisco Cloud APIC Scalability Limits](#) 37

[Cisco ACI and UCSM Scalability](#) 39

[QoS Scalability Limits](#) 39

[PTP Scalability Limits](#) 39

[NetFlow Scale](#) 40

Revised: July 12, 2022,

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 5.2(1)
- Cisco Nexus 9000 Series ACI-Mode Switches, Release 15.2(1)

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 Multi-Site Orchestrator release-specific documents available at the following URL: <https://www.cisco.com/c/en/us/support/cloud-systems-management/multi-site-orchestrator/products-tech-notes-list.html>.

New and Changed Information

The following changes have been made to this document since initial release:

Date	Changes
May 11, 2022	Added dual-stack scale for "Number of L3 Outs".
April 29, 2022	Updated "Number of External Route Reflectors between Pods" scale. Updated "Number of External EPGs" and "Number of External EPGs per L3 Out" scale with examples for clarity.
April 21, 2022	Added "DHCP relay addresses per BD across all labels" scale.
March 25, 2022	Updated "PTP Scalability Limits" section.
March 20, 2022	Improved content organization of the "General Scalability" section.
March 15, 2022	Updated "Maximum number of Data Plane policers at the interface level" scale numbers.
February 25, 2022	Updated "Number of source EPGs in tenant SPAN sessions" if both Access and Tenant SPAN are configured.
December 17, 2021	Added NetFlow scale numbers.
June 7, 2021	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. An ACI Multi-Site Orchestrator 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 128 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. RL pod redundancy is supported for up to 32 remote leaf switches.

Breakout Ports

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

General Scalability Limits

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	80	200	300	500

Configurable Options	Default Fabric	Medium Fabric	Large Fabric	
Number of tier-2 leaf switches 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	100	100
Number of leaf switches per pod	80	200	200	400
Number of Pods	6	6	12	12
Number of tenants	1000	1000	3000	3000
Number of Layer 3 (L3) contexts (VRFs)	1000	1000	3000	3000

Table 2: General Scalability Limits Per Fabric

Configurable Options	Scale Limits
Number of spine switches per Pod	6
Number of spine switches in a Multi-Pod fabric	24
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: 4000 • Multi-tenant fabrics: 500 <p>Or one of the following 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 1400 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	<ul style="list-style-type: none"> • 200 VDS • 50 AVS • 50 Cisco ACI Virtual Edge
Number of Service Chains	1000
Number of L4 - L7 devices	30 managed or 50 unmanaged physical HA pairs 1200 virtual HA pairs (1200 maximum per fabric)
Number of ESXi hosts - VDS	3200
Number of ESXi hosts - AVS	3200 (only 1 AVS instance per host)
Number of ESXi hosts - AVE	3200 (only 1 AVE instance per host)

Configurable Options	Scale Limits
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)	1000 VRFs 60,000 routes in a fabric
Number of Routes in Overlay-1 VRF	1000

Multiple Fabric Options Scalability Limits

Stretched Fabric

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

Multi-Pod

Configurable Options	Per Fabric Scale
Maximum number of Pods	12
Maximum number of leaf switches per Pod	400
Maximum number of leaf switches overall	500
Maximum number of Route Reflectors for L3Out	24
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 <p>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</p>

Cisco ACI vPod Scalability Limits

Cisco ACI vPod Scalability Limits

Configurable Options	Scale
Number of vPods	6
Number of Cisco ACI Virtual Edge (AVE) instances per vPod	32
Number of Virtual Ethernet Ports (vETHs) per AVE in vPod	32
Number of EPGs per vPod	256
Number of EPGs across all vPods	864
Number of EPGs across all physical and virtual pods	15,000
Number of filters per ACI Virtual Edge	128
Number of contracts per ACI Virtual Edge *The total number of filters used by all contracts must not exceed the filter limit above	36

Cisco Multi-Site Scalability Limits

Cisco Multi-Site Orchestrator (MSO) does not require a specific version of APIC to be running in all sites. The APIC clusters in each site as well as the MSO 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 Multi-Site Orchestrator release are now available at the following URL: <https://www.cisco.com/c/en/us/support/cloud-systems-management/multi-site-orchestrator/products-tech-notes-list.html>.



Note Each site managed by the Multi-Site 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), External EPGs, Bridge Domains, Endpoints, and Contracts Scalability Limits

The following table shows the mapping of the "ALE/LSE Type" to the corresponding ToR switches. This information is helpful to determine which ToR switch is affected when we use the terms ALE v1, ALE v2, LSE, or LSE2 in remaining sections.



Note In the following table, the N9K-C9336C-FX2 and N9K-C93360YC-FX2 switches are listed as LSE for scalability limits purposes only; the switches support LSE2 platform features. Consult specific feature documentation for the full list of supported devices.

ALE/LSE Type	ACI-Supported ToR switches
ALE v2	<ul style="list-style-type: none">• N9K-C9396TX + N9K-M6PQ• N9K-C93128TX + N9K-M6PQ• N9K-C9396PX + N9K-M6PQ• N9K-C9372TX 64K• N9K-C9332PQ• N9K-C9372PX
LSE	<ul style="list-style-type: none">• N9K-C93108TC-EX• N9K-C93180YC-EX• N9K-C93180LC-EX• N9K-C9336C-FX2• N9K-C93216TC-FX2• N9K-C93240YC-FX2• N9K-C93360YC-FX2
LSE2	<ul style="list-style-type: none">• N9K-C93108TC-FX• N9K-C93180YC-FX• N9K-C9348GC-FXP• N9K-C93600CD-GX• N9K-C9364C-GX• N9K-C93180YC-FX3• N9K-C93108TC-FX3P

**Note**

- Unless explicitly called out, LSE represents both LSE and LSE2 and ALE represents both ALE v1 and ALE v2 in the rest of this document.
- The High Policy profile in LSE2 switches listed in the following sections is supported only on Cisco Nexus N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C9364C-GX, C93180YC-FX3, and C93108TC-FX3P switches with 32GB 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.

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)	3000	N/A
Number of encapsulations per FEX HIF, PC, vPC	20	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)	ALE v2: 9000 LSE and LSE2: 10,000	N/A
Number of static port bindings	ALE v2: 30,000 For LSE and LSE2: 60,000	600,000 (200,000 per tenant)
Number of VMACs	For ALE v2: 255 For LSE and LSE2: 510	N/A
STP	All VLANs	N/A
Mis-Cabling Protocol (MCP)	256 VLANs per interface 2000 logical ports (port x VLAN) per leaf	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of endpoints (EPs)		<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 the following:</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 as follows:</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> <hr/> <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 the following:</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 as follows:</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
		<p>Fixed spine switches:</p> <p>Max. 180,000 Proxy Database Entries in the fabric, which can be translated into any one of the following:</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 as follows:</p> <p>#MAC + #IPv4 + #IPv6 <= 180,000</p>

Configurable Options	Per Leaf Scale	Per Fabric Scale
	<p>Default (Dual Stack) profile:</p> <ul style="list-style-type: none"> • ALE v2: <ul style="list-style-type: none"> • MAC: 12,000 • IPv4: 12,000 or • IPv6: 6000 or • IPv4: 4000 IPv6: 4000 <p>Default profile or High LPM profile:</p> <ul style="list-style-type: none"> • LSE or LSE2: <ul style="list-style-type: none"> • MAC: 24,000 • IPv4: 24,000 • IPv6: 12,000 <p>IPv4 scale profile:</p> <ul style="list-style-type: none"> • LSE and LSE2: <ul style="list-style-type: none"> • MAC: 48,000 • IPv4: 48,000 • IPv6: Not supported • ALE v2: 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 • ALE v2: Not supported <p>High Policy profile:</p>	

Configurable Options	Per Leaf Scale	Per Fabric Scale
	<ul style="list-style-type: none"> • LSE2 (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C9364C-GX, C93180YC-FX3, and C93108TC-FX3P switches with 32GB of RAM only): <ul style="list-style-type: none"> • MAC: 24,000 • IPv4: 24,000 • IPv6: 12,000 • LSE (N9K-C9336C-FX2 and N9K-C93180YC-EX): <ul style="list-style-type: none"> • MAC: 16,000 • IPv4: 16,000 • IPv6: 8000 High IPv4 EP Scale profile: <ul style="list-style-type: none"> • LSE2 (N9K-C93180YC-FX, N9K-C93600CD-GX, and N9K-C93180YC-FX3 switches with 32GB of RAM only): <ul style="list-style-type: none"> • MAC: 24,000 • IPv4 local: 24,000 • IPv4 total: 280,000 • IPv6: 12,000 • Not supported on LSE Multicast Heavy profile: <ul style="list-style-type: none"> • LSE2 (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C93180YC-FX3, and N9K-C93108TC-FX3P switches with 32GB of RAM only): <ul style="list-style-type: none"> • MAC: 24,000 • IPv4 local: 24,000 • IPv4 total: 64,000 • IPv6: 4000 • Not supported on LSE 	

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: 8000 with (S,G) scale not exceeding 4000</p> <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 (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C93180YC-FX3, and N9K-C93108TC-FX3P switches with 32GB of RAM only): 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: 8000 with (S,G) scale not exceeding 4000</p> <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 (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C93180YC-FX3, and N9K-C93108TC-FX3P switches with 32GB of RAM only): 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: 8000 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	4096	4096
Number of Host-Based Routing Advertisements	30,000 host routes per border leaf	N/A
SPAN	ALE-based ToR switches: <ul style="list-style-type: none"> • 4 unidirectional or 2 bidirectional access/tenant sessions • 4 unidirectional or 2 bidirectional fabric sessions LSE-based ToR switches: <ul style="list-style-type: none"> • 32 unidirectional or 16 bidirectional sessions (fabric, access, or tenant) 	N/A
Number of ports per SPAN session	ALE-based ToR switches: <ul style="list-style-type: none"> • All leaf access ports could be in one session. • All leaf fabric ports could be in one session. LSE/LSE2-based ToR switches: <ul style="list-style-type: none"> • 63 – total number of unique ports (fabric + access) across all types of span sessions 	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
<p>Number of source EPGs in tenant SPAN sessions</p> <p>Note The numbers listed in this row assume that only tenant SPAN is configured.</p> <p>If both, Access and Tenant SPAN are configured, the following formula applies for both ingress and egress SPAN:</p> $E + P + E * P + EPP + v4FePP + 0.5 * v4FePP \leq 230$ <p>Where:</p> <ul style="list-style-type: none"> • E— Number of source EPGs in Tenant SPAN • P—Number of source Ports in access SPAN without any filters • EPP—Number of (Epg,Port) Pairs in access SPAN with EPG filter only (no filter group) • v4FePP—Number of (v4 filter entry, Port) Pairs in access SPAN with filter group • v6FePP—Number of (v6 Filter entry, Port) Pairs in access SPAN with filter group 	<p>ALE-based ToR switches:</p> <ul style="list-style-type: none"> • 230 ingress direction + 50 egress direction <p>LSE-based ToR switches:</p> <ul style="list-style-type: none"> • 230 bidirectional • 460 unidirectional (230 ingress + 230 egress) 	<p>N/A</p>
<p>Number of SPAN ACL filter TCAM entries</p> <p>SPAN filters are supported on -EX, -FX, and -FX2 TORs only.</p> <p>SPAN filters are not supported in the following:</p> <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 <p>Total number of TCAM entries is calculated using the following formula:</p> $\begin{aligned} & (\text{IPv4-filters}) * \\ & (\text{IPv4-filter-source-groups}) + 2 * \\ & (\text{IPv6-filters}) * \\ & (\text{IPv6-filter-source-groups}) + 2 * \\ & (\text{no-filter-source-groups}) \end{aligned}$	<p>N/A</p>

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of L4 Port Ranges	<p>16 (8 source and 8 destination)</p> <p>First 16 port ranges consume a TCAM entry per range.</p> <p>Each additional port range beyond the first 16 consumes a TCAM entry per port in the port range.</p> <p>Filters with distinct source port range and destination port range count as 2 port ranges.</p> <p>You cannot add more than 16 port ranges at once.</p>	N/A
Common pervasive gateway	256 virtual IPs per Bridge Domain	N/A
Number of Data Plane policers at the interface level	<p>ALE:</p> <ul style="list-style-type: none"> • 64 ingress policers • 64 egress policers <p>LSE and LSE2:</p> <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	<p>256</p> <p>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</p>	N/A
Number of SNMP trap receivers	10	10
<p>IP SLA probes*</p> <p>*With 1 second probe time and 3 seconds of timeout</p>	100	1500

Configurable Options	Per Leaf Scale	Per Fabric Scale
First Hop Security (FHS)* With any combination of BDs/EPGs/EPs within the supported limit	2000 endpoints 1000 bridge domains	N/A
Number of Q-in-Q tunnels (both QinQ core and edge combined)	1980	N/A
Number of TEP-to-TEP atomic counters	N/A	1600

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	N/A	1,200
Total routes* * Including both, remote leaf and multi-pod	N/A	60,000
External EPGs	N/A	2000 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	ALE: 50 LSE: 128	ALE: 50 LSE: 128

VRFs (Contexts)

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

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of Contexts (VRFs)	ALE: 400 LSE and LSE2: 800	3000
Maximum ECMP (equal cost multipath) for BGP best path	64	N/A
Maximum ECMP (equal cost multipath) for OSPF best path	64	N/A
Maximum ECMP (equal cost multipath) for Static Route best path	64	N/A
Number of isolated EPGs	400	400
Border Leafs per L3 Out	N/A	12
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
L3 Out per context (VRF)	N/A	400
Number of BFD neighbors	<ul style="list-style-type: none"> • Up to 256 sessions using the following minimum BFD timers: • minTx:50 • minRx:50 • multiplier:3 	N/A
Number of BGP neighbors	400	10,000
Number of OSPF neighbors	300	N/A
Number of EIGRP neighbors	16	N/A
Number of static routes to a single SVI/VRF	5000	N/A
Number of static routes on a single leaf switch	10,000	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
<p>Number of IP Longest Prefix Matches (LPM) entries</p> <p>Note 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"> • ALE v2: <ul style="list-style-type: none"> • IPv4: 10,000 or • IPv6: 6000 or • IPv4: 4000, IPv6: 4000 • IPv6 wide prefixes (> /64): 1000 • For LSE or LSE2: <ul style="list-style-type: none"> • IPv4: 20,000 or • IPv6: 10,000 • IPv6 wide prefixes (>= /84): 1000 <p>NOTE: For LSE2 and FX2 models there's no restriction on wide prefixes.</p> <p>IPv4 scale profile:</p> <ul style="list-style-type: none"> • For LSE or LSE2: <ul style="list-style-type: none"> • IPv4: 38,000 • IPv6: Not supported • ALE v2: Not supported <p>High Dual Stack scale profile:</p> <ul style="list-style-type: none"> • LSE or LSE2: <ul style="list-style-type: none"> • IPv4: 38,000 or • IPv6: 19,000 • IPv6 wide prefixes (>= /84): 1000 <p>NOTE: For LSE2 and FX2 models there's no restriction on wide prefixes.</p> • ALE v2: Not supported 	<p>N/A</p>

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of IP Longest Prefix Matches (LPM) entries (Continued) Note The total of (# of IPv4 prefixes) + 2*(# of IPv6 prefixes) must not exceed the scale listed for IPv4 alone		N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
	<p>High LPM Scale profile –</p> <ul style="list-style-type: none"> • LSE or LSE2: <ul style="list-style-type: none"> • IPv4: 128,000 or • IPv6: 64,000 • IPv6 wide prefixes ($\geq /84$): 1000 <p>NOTE: For LSE2 and FX2 models there's no restriction on wide prefixes.</p> <ul style="list-style-type: none"> • ALE v2: Not supported <p>High Policy profile:</p> <ul style="list-style-type: none"> • LSE2 (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C9364C-GX, C93180YC-FX3, and C93108TC-FX3P switches with 32GB of RAM only): <ul style="list-style-type: none"> • IPv4: 20,000 or • IPv6: 10,000 • LSE (N9K-C9336C-FX2 and N9K-C93180YC-EX): <ul style="list-style-type: none"> • IPv4: 8000 • IPv6: 4000 <p>High IPv4 EP Scale profile:</p> <ul style="list-style-type: none"> • LSE2 (N9K-C93180YC-FX, N9K-C93600CD-GX, and N9K-C93180YC-FX3 switches with 32GB of RAM only): <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 (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C93180YC-FX3, and N9K-C93108TC-FX3P switches with 	

Configurable Options	Per Leaf Scale	Per Fabric Scale
	32GB of RAM only): <ul style="list-style-type: none"> • IPv4: 20,000 • IPv6: 10,000 • LSE: Not supported 	
Number of Secondary addresses per logical interface	1	1
Number of L3 interfaces per Context	<ul style="list-style-type: none"> • 1000 SVIs • 16 Routed interfaces • 100 sub-interfaces with or without port-channel 	N/A
Number of L3 interfaces	<ul style="list-style-type: none"> • 1000 SVIs • 16 Routed interfaces • 1000 sub-interfaces with or without port-channel 	N/A
Number of ARP entries for L3 Outs	7,500	N/A
Shared L3 Out	<ul style="list-style-type: none"> • IPv4 Prefixes: 2000 or • IPv6 Prefixes: 1000 	<ul style="list-style-type: none"> • IPv4 Prefixes: 6000 or • IPv6 Prefixes: 3000
Number of L3 Outs	400 For LSE and LSE2: 800	2400 (single stack)

External EPGs

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of External EPGs	800	<p>ALE: 2400</p> <p>LSE: 4000</p> <p>The listed scale is calculated as a product of (Number of external EPGs)*(Number of border leaf switches for the L3Out)</p> <p>For example, the following 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
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>
<p>Number of LPM 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>ALE: 1000 IPv4</p> <p>LSE: refer to LPM scale section.</p>	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		N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
	<p>ALE: 1000</p> <p>LSE and LSE2:</p> <ul style="list-style-type: none"> • Default Profile: <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> • IPv4 Scale profile: <ul style="list-style-type: none"> • IPv4 (/32): 16,000 <p>Combined number of host prefixes, multicast groups, and endpoints can't exceed 56,000.</p> • IPv6 (/128): 0 <ul style="list-style-type: none"> • High Dual Stack Profile: <ul style="list-style-type: none"> • IPv4 (/32): 64,000 <p>Combined number of host prefixes, multicast groups, and endpoints can't exceed 64,000.</p> • IPv6 (/128): 24,000 (LSE) <p>Combined number of host prefixes and endpoints can't exceed 24,000.</p> <ul style="list-style-type: none"> • IPv6 (/128): 48,000 (LSE2 only) <p>Combined number of host prefixes and endpoints can't exceed 48,000.</p> <ul style="list-style-type: none"> • High LPM Profile: <ul style="list-style-type: none"> • IPv4 (/32): 24,000 <p>Combined number of host prefixes, multicast groups, and endpoints can't exceed 24,000.</p> • IPv6 (/128): 12,000 <p>Combined number of host prefixes and endpoints can't exceed 12,000.</p>	

Configurable Options	Per Leaf Scale	Per Fabric Scale
	<ul style="list-style-type: none"> • High Policy profile (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C9364C-GX, C93180YC-FX3, and C93108TC-FX3P switches with 32GB of RAM only): <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> • High IPv4 EP Scale profile (N9K-C93180YC-FX, N9K-C93600CD-GX, and N9K-C93180YC-FX3 switches with 32GB of RAM only): <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> • Multicast Heavy profile (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C93180YC-FX3, and N9K-C93108TC-FX3P switches with 32GB of RAM only): <ul style="list-style-type: none"> • IPv4 (/32): 16,000 <p>Combined number of host prefixes, multicast groups, and endpoints can't exceed 154,000.</p> • IPv6 (/128): 4000 <p>Combined number of host prefixes and endpoints can't exceed 4000.</p> 	

Bridge Domains

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of BDs	1980 Legacy mode: 3500 On ALE ToR switches with multicast optimized mode: 50	15,000
Number of BDs with Unicast Routing per Context (VRF)	ALE: 256 LSE: 1000	1750
Number of subnets per BD	1000, cannot be for all BDs.	1000 per BD
Number of EPGs per BD	3960	4000
Number of L2 Outs per BD	1	1
Number of BDs with Custom MAC Address	1750 Legacy mode: 3500 On ALE ToR switches with multicast optimized mode: 50	1750 Legacy mode: 3500 On ALE ToR switches with multicast optimized mode: 50
Number of EPGs + L3 Outs per Multicast Group	128	128
Number of BDs with L3 Multicast enabled	1750	1750
Number of VRFs with L3 Multicast enabled	64	300
Number of L3 Outs per BD	ALE: 4 LSE: 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
Number of external EPGs per L2 out	1	1
Number of PIM Neighbors	1000	1000
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
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	<ul style="list-style-type: none"> • 200 VDS • 50 AVS • 50 Cisco ACI Virtual Edge
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	4000	N/A
Number of Taboo Contracts per EPG	2	N/A
IP-based EPG (bare metal)	4000	N/A
MAC-based EPG (bare metal)	4000	N/A

Contracts

Policy TCAM compression is supported starting with Release 4.0(1). In this release, policy compression feature is supported only on LSE (N9K-C9336C-FX2) and LSE2 (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C9364C-GX, and C93180YC-FX3) switches. In the following table, "Software Policy scale" refers to the recommended maximum PE `actrlRules` deployed on any given TOR. We can retrieve the PE `actrlRule` objects deployed on TOR by querying `actrlRule` Managed Object.

Configurable Options	Per Leaf Scale	Per Fabric Scale
Security TCAM size	<p>Default scale profile:</p> <ul style="list-style-type: none"> • ALE v2: 40,000 • LSE and LSE2: 64,000 <p>IPv4 scale profile:</p> <ul style="list-style-type: none"> • ALE v2: N/A • LSE and LSE2: 64,000 <p>High Dual Stack scale profile:</p> <ul style="list-style-type: none"> • ALE v2: N/A • LSE: 8000 • LSE2: 128,000 <p>High LPM scale profile:</p> <ul style="list-style-type: none"> • ALE v2: N/A • LSE and LSE2: 8000 <p>High Policy profile:</p> <ul style="list-style-type: none"> • LSE2 (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C9364C-GX, C93180YC-FX3, and C93108TC-FX3P switches with 32GB of RAM only): 256,000 • LSE (N9K-C9336C-FX2 and N9K-C93180YC-EX): 100,000 <p>High IPv4 EP Scale profile:</p> <ul style="list-style-type: none"> • LSE2 (N9K-C93180YC-FX, N9K-C93600CD-GX, and N9K-C93180YC-FX3 switches with 32GB of RAM only): 64,000 • Not supported on LSE <p>Multicast Heavy profile:</p> <ul style="list-style-type: none"> • LSE2 (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C93180YC-FX3, and N9K-C93108TC-FX3P switches with 32GB of RAM only): 64,000 • Not supported on LSE 	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Software policy scale with Policy Table Compression enabled (Number of <code>actrlRule</code> Managed Objects)	Dual stack profile: <ul style="list-style-type: none"> • LSE (N9K-C9336C-FX2 only): 80,000 • LSE2 (N9K-C93180YC-FX only): 80,000 High Dual Stack profile: <ul style="list-style-type: none"> • LSE2 (N9K-C93180YC-FX, N9K-C93600CD-GX, and N9K-C9364C-GX only) : 140,000 High Policy profile: <ul style="list-style-type: none"> • LSE2 (N9K-C93180YC-FX, N9K-C93600CD-GX, N9K-C9364C-GX, and C93180YC-FX3 switches with 32GB of RAM) : 256,000 • LSE (N9K-C9336C-FX2): 100,000 	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
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)	1000	1000
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
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

Configurable Options	Scale
Number of ESGs per VRF	4000
Number of ESGs per Tenant	4000
Number of L2 MAC Selectors per Leaf	5000
Number of L3 IP Selectors per Leaf	5000

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	1000

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	1000
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)
Number of vCenters (AVS)	N/A	50
Number of vCenters (Cisco ACI Virtual Edge)	N/A	50
Datacenters in a vCenter	N/A	15
Total Number of VMM domain (vCenter, Datacenter) instances.	N/A	<ul style="list-style-type: none"> • 200 VDS • 50 AVS • 50 Cisco ACI Virtual Edge
Number of ESX hosts per AVS	240	N/A
Number of ESX hosts running Cisco ACI Virtual Edge	150	N/A
Number of EPGs per vCenter/vDS	N/A	5000
Number of EPGs to VMware domains/vDS	N/A	5000
Number of EPGs per vCenter/AVS	N/A	3500
Number of EPGs to VMware domains/AVS	N/A	3500
Number of EPGs per vCenter/Cisco ACI Virtual Edge	N/A	VLAN Mode: 1300 VXLAN Mode: 2000
Number of EPGs to VMware domains and Cisco ACI Virtual Edge	N/A	VLAN Mode: 1300 VXLAN Mode: 2000
Number of endpoints (EPs) per AVS	10,000	10,000
Number of endpoints per VDS	10,000	10,000
Number of endpoints per vCenter	10,000	10,000
Number of endpoints per Cisco ACI Virtual Edge	10,000	10,000
Support RBAC for AVS	N/A	Yes
Support RBAC for VDS	N/A	Yes

Configurable Options	Per Leaf Scale	Per Fabric Scale
Support RBAC for Cisco ACI Virtual Edge	N/A	Yes
Number of Microsegment EPGs with vDS	400	N/A
Number of Microsegment EPGs with AVS	1000	N/A
Number of Microsegment EPGs with Cisco ACI Virtual Edge	1000	N/A
Number of DFW flows per vEth with AVS	10,000	N/A
Number of DFW flows per vEth with Cisco ACI Virtual Edge	10,000	N/A
Number of DFW denied and permitted flows per ESX host with AVS	250,000	N/A
Number of DFW denied and permitted flows per ESX host with Cisco ACI Virtual Edge	250,000	N/A
Number of VMM domains per EPG with AVS	N/A	10
Number of VMM domains per EPG with Cisco ACI Virtual Edge	N/A	10
Number of VM Attribute Tags per vCenter	N/A	vCenter version 6.0: 500 vCenter version 6.5: 1000

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	3000
EPGs per all Microsoft VMM domains	N/A	N/A	9000
EP/VNICs per HyperV host	N/A	N/A	100
EP/VNICs per SCVMM	3000	10,000	10,000
Number of Hyper-V hosts	64	N/A	N/A

Configurable Options	Per Leaf Scale (On-Demand Mode)	Per Leaf Scale (Pre-Provision Mode)	Per Fabric Scale
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	1000	Not Supported	N/A

Microsoft Windows Azure Pack

Configurable Options	Per Fabric Scale
Number of Windows Azure Pack subscriptions	1000
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	3000
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 logical device clusters	1200
Number of graph instances	1000
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

Configurable Options (L4-L7 Configurations)	Per Fabric Scale
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 (can go beyond).
Number of security domains/APIC	15 (can go beyond).
Number of security domains in which the tenant resides	4 (can go beyond).
Number of priorities	4 tested (16 per domain)
Number of shell profiles that can be returned.	4 tested (32 domains total)
Number of users	8000 local / 8000 remote
Number of simultaneous logins	500 connections / NGNIX simultaneous REST logins

Cisco Mini ACI Fabric and Virtual APICs Scalability Limits

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

Property	Maximum Scale
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 Cloud APIC Scalability Limits

This section contains scalability numbers for Cisco ACI cloud deployments. The scalability limits differ based on whether it's a single cloud site or a multi-cloud deployment.

Single Cloud Site

This section contains scalability numbers for a single cloud site deployment. The same scale numbers apply to both, AWS or Azure, cloud providers.

Table 3: Single AWS Cloud Site

Configurable Options	Scale
Number of Tenants	20
Number of Application Profiles	500
Number of EPGs	500
Number of cloud Endpoints	1000
Number of VRFs	20
Cloud Context Profiles	40
Number of Contracts	1000
Number of L4-L7 Service Graphs	200
Number of L4-L7 Services Devices (AWS ALB)	100
Number of Policy Rules	70,000
Number of Regions	4
Number of CSRs per Region	2
Number of CSRs per Site	4
Number of hub networks for Transit Gateway (TGW)	2
Number of Transit Gateways per hub network	2

Configurable Options	Scale
Number of restricted domains (security domain with restricted role)	32

Table 4: Single Azure Cloud Site

Configurable Options	Scale
Number of Tenants	20
Number of Application Profiles	2000
Number of EPGs	2000
Number of cloud Endpoints	3500
Number of VRFs	60
Cloud Context Profiles	100
Number of Contracts	1000
Number of L4-L7 Service Graphs	200
Number of L4-L7 Services Devices	100
Number of Policy Rules	125,000
Number of Regions	4
Number of CSRs per Region	4
Number of CSRs per Site	16
Number of restricted domains (security domain with restricted role)	32

Multi-Cloud Deployments

This section contains scalability numbers for multi-cloud deployments. The same scale numbers apply to each cloud site (AWS or Azure) with intersite connectivity provided by the ACI Multi-Site Orchestrator. Total number of stretched and non-stretched objects must not exceed the maximum verified scalability limit for that object.

Table 5: Multi-Cloud Deployments

Configurable Options	Scale
Number of cloud sites	2
Number of managed regions per site	4
Number of CSRs per site	4
Number of CSRs per region	2

Configurable Options	Scale
Number of Tenants	5
Number of EPGs	250
Number of cloud endpoints	500
Number of VRFs	10
Cloud Context Profiles (VPC/VNET)	40
Number of Contracts	200

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	4000
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
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	128	40	N/A
Number of ACI switches connected to the same tier-1 leaf switch (multi-tier topology) with PTP globally enabled	16	16	N/A
Number of access ports with PTP enabled on a leaf switch	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	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
Number of PTP peers per leaf switch	26	26	25

NetFlow Scale

Configurable Options	Scale
Exporters per leaf switch	2
NetFlow monitor policies under BDs per leaf switch	100
NetFlow monitor policies under L3Outs per leaf switch	120
Number of records per collect interval	20,000

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: <http://www.cisco.com/go/softwareterms>. Cisco product warranty information is available at <http://www.cisco.com/go/warranty>. US Federal Communications Commission Notices are found here <http://www.cisco.com/c/en/us/products/us-fcc-notice.html>.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: [www.cisco.com go trademarks](http://www.cisco.com/go/trademarks). Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1721R)

© 2021–2022 Cisco Systems, Inc. All rights reserved.



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.