



Verified Scalability Guide for Cisco APIC, Release 4.1(1), Multi-Site, Release 2.1(1), and Cisco Nexus 9000 Series ACI-Mode Switches, Release 14.1(1)

Overview 2

General Scalability Limits 2

Multiple Fabric Options Scalability Limits 5

Cisco ACI vPod Scalability Limits 6

Cisco ACI Multi-Site Scalability Limits 6

Fabric Topology, SPAN, Tenants, Contexts (VRFs), External EPGs, Bridge Domains, Endpoints, and Contracts Scalability Limits 8

VMM Scalability Limits 26

Layer 4 - Layer 7 Scalability Limits 28

AD, TACACS, RBAC Scalability Limits 29

Cisco Mini ACI Fabric and Virtual APICs Scalability Limits 29

Cisco Cloud APIC Scalability Limits 30

Cisco ACI and UCSM Scalability 30

QoS Scalability Limits 31

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 4.1(1)
- Cisco ACI Multi-Site, Release 2.1(1)
- Cisco Nexus 9000 Series ACI-Mode Switches, Release 14.1(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.

General Scalability Limits

- **L2 Fabric:** In Legacy mode there is no routing, L3 context, nor contract enabled in the L2 fabric profile. A tenant in this profile does not need to be mapped to one dedicated ACI tenant. A tenant can be represented by a set of EPGs instead. To improve the load sharing among APIC controller nodes, you must distribute EPGs and BDs across different ACI tenants.
- **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 200 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 overall is 400 per fabric scale. The maximum number of physical ports is 19,200 per fabric. The maximum number of remote leaf switches is 40 per fabric, with total number of unique BDs deployed on all remote leaf switches in a POD not exceeding 1,000.

General Scalability Limits

Configurable Options	L2 Fabric	L3 Fabric	Large L3 Fabric
Number of APIC controllers Note * denotes preferred cluster size. While the higher number of controllers is supported, the preferred size is based on the number of leaf switches in the environment.	3* or 4 node APIC cluster	3* or 4 node APIC cluster	5*, 6, or 7 node APIC cluster
Number of leaf switches	80	80 for 3-node cluster 200 for 4-node cluster	300 for 5- or 6-node cluster 400 for 7-node cluster
Number of spine switches	Maximum spines per pod: 6. Total spines per fabric: 24.	Maximum spines per pod: 6. Total spines per fabric: 24.	Maximum spines per pod: 6. Total spines per fabric: 24.
Number of FEXs	20 FEXs per leaf switch 576 ports per leaf switch 650 FEXs per fabric	20 FEXs per leaf switch 576 ports per leaf switch 650 FEXs per fabric	20 FEXs per leaf switch 576 ports per leaf switch 650 FEXs per fabric
Number of tenants	1,000	1,000	3,000
Number of Layer 3 (L3) contexts (VRFs)	N/A	1,000	3,000
Number of contracts/filters	N/A	• 10,000 contracts • 10,000 filters	• 10,000 contracts • 10,000 filters
Number of endpoint groups (EPGs)	For a fabric with a single Tenant: 4,000 For a fabric with multiple Tenants: 500 per Tenant, up to 21,000 total across all Tenants	For a fabric with a single Tenant: 4,000 For a fabric with multiple Tenants: 500 per Tenant, up to 15,000 total across all Tenants	For a fabric with a single Tenant: 4,000 For a fabric with multiple Tenants: 500 per Tenant, up to 15,000 total across all Tenants
Number of Isolation enabled EPGs	400	400	400
Number of bridge domains (BDs)	21,000	15,000	15,000

Configurable Options	L2 Fabric	L3 Fabric	Large L3 Fabric
Number of BGP + number of OSPF sessions + EIGRP (for external connection)	N/A	3,000	3,000
Number of Multicast routes	N/A	32,000	32,000
Number of Multicast routes per VRF	N/A	32,000	32,000
Number of static routes to a single SVI/VRF	N/A	5,000	5,000
Number of static routes on a single leaf switch	N/A	10,000	10,000
Number of vCenters	N/A	• 200 VDS • 50 AVS • 50 Cisco ACI Virtual Edge	• 200 VDS • 50 AVS • 50 Cisco ACI Virtual Edge
Number of Service Chains	N/A	1,000	1,000
Number of L4 - L7 devices	N/A	30 managed or 50 unmanaged physical HA pairs, 1,200 virtual HA pairs (1,200 maximum per fabric)	30 managed or 50 unmanaged physical HA pairs, 1,200 virtual HA pairs (1,200 maximum per fabric)
Number of ESXi hosts - VDS	N/A	3,200	3,200
Number of ESXi hosts - AVS	N/A	3,200 (Only 1 AVS instance per host)	3,200 (Only 1 AVS instance per host)
Number of ESXi hosts - AVE	N/A	3,200 (Only 1 AVE instance per host)	3,200 (Only 1 AVE instance per host)
Number of VMs	N/A	Depends upon server scale	Depends upon server scale
Number of configuration zones per fabric	30	30	30
Number of BFD sessions per leaf switch	256 Minimum BFD timer required to support this scale: • minTx:50 • minRx:50 • multiplier:3	256 Minimum BFD timer required to support this scale: • minTx:50 • minRx:50 • multiplier:3	256 Minimum BFD timer required to support this scale: • minTx:50 • minRx:50 • multiplier:3

Configurable Options	L2 Fabric	L3 Fabric	Large L3 Fabric
Multi-Pod NOTE: * is preferred cluster size	• 3* or 4 node APIC cluster • 6 pods • 80 leaf switches overall	• 3* or 4 node APIC cluster • 6 pods • 80 for 3-node cluster 200 for 4-node cluster	• 5* or 6 node APIC cluster, 6 pods, 200 leaf switches max per pod, 300 leaf switches max overall • 7 node APIC cluster, 12 pods, 200 leaf switches max per pod, 400 leaf switches max overall
L3 EVPN Services over Fabric WAN - GOLF (with and without OpFlex)	N/A	1,000 VRFs, 60,000 routes in a fabric	1,000 VRFs, 60,000 routes in a fabric
Layer 3 Multicast routes	N/A	32,000	32,000
Number of Routes in Overlay-1 VRF	1,000	1,000	1,000

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	200
Maximum number of leaf switches overall	400
Maximum number of Route Reflectors for L3Out	24

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

Cisco 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 ACI Multi-Site Scalability Limits

Stretched Vs. Non-Stretched

Stretched in Multi-Site means that the fabric has stretched objects such as EPGs, BDs, VRFs, or subnets across multiple sites or has cross-site contracts between EPGs.

Non-Stretched in Multi-Site means all objects such as EPG, contract, and BD are local to a site only and do not cross the local-site boundary.

The total number of stretched and non-stretched objects must not exceed the maximum verified scalability limit for that object, which are listed in their respective sections in this guide.



Note For maximum scale Multi-Site configurations with many features enabled simultaneously, it is recommended that those configurations be tested in a lab before deployment.

Multi-Site General Scalability Limits

Configurable Options	Scale
Sites	12
Pods per site	12
Leaf switches per site	200

Multi-Site Object Scale

Configurable Options	Scale
Policy Objects per Schema	500
Templates per Schema	5
Application Profiles per Schema	200
Number of Schemas	80
Number of Templates	400
Multi-Site Orchestrator Users (nonparallel*) *Multi-Site Orchestrator processes requests sequentially from multiple users even if they are deploying different schemas.	50

Multi-Site Scalability Limits for Stretched Objects

Configurable Options	Scale (Stretched)
Tenants	400
VRFs	1,000
BDs	4,000
Contracts	4,000
Endpoints	150,000 including: • 100,000 - learned from other sites • 50,000 - locally learned in site-local

Configurable Options	Scale (Stretched)
EPGs	4,000
Isolated EPGs	400
Microsegment EPGs	400
IGMP Snooping	8,000
Layer-3 multicast routes	8,000
L3Out external EPGs	500
Subnets	8,000
Contract Preferred Group (BD/EPG combinations)	250* *EPGs must be added to the preferred group gradually
Number of L4-L7 logical devices	400
Number of graph instances	250
Number of device clusters per tenant	10
Number of interfaces per device cluster	Any
Number of graph instances per device cluster	125

Multi-Site VRF/BD VNID Translation Scale

Configurable Options	Scale
Fixed spines	21,000
Modular spines	42,000

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 leaf switches. This information is helpful to determine which leaf 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 switch is listed as LSE for scalability limits purposes only; the switch supports LSE2 platform features. Consult specific feature documentation for the full list of supported devices.

ALE/LSE type	ACI-supported leaf switches
ALE v1	• N9K-C9396PX + N9K-M12PQ • N9K-C93128TX + N9K-M12PQ • N9K-C9396TX + N9K-M12PQ
ALE v2	• N9K-C9396TX + N9K-M6PQ • N9K-C93128TX + N9K-M6PQ • N9K-C9396PX + N9K-M6PQ • N9K-C9372TX 64K • N9K-C9332PQ • N9K-C9372PX
LSE	N9K-C93108TC-EX + N9K-C93180YC-EX + N9K-C9336C-FX2
LSE2	N9K-C93108TC-FX + N9K-C93180YC-FX + N9K-C9348GC-FXP



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.

Fabric Topology

Configurable Options	Per Leaf Switch 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	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)	For ALE v1 and v2: 9,000 For LSE and LSE2: 10,000	N/A

Configurable Options	Per Leaf Switch Scale	Per Fabric Scale
Number of static port bindings	For ALE v1 and v2: 30,000 For LSE and LSE2: 60,000	400,000
STP	All VLANs	N/A
Mis-Cabling Protocol (MCP)	256 VLANs per interface 2,000 logical ports (port x VLAN) per leaf	N/A

Configurable Options	Per Leaf Switch Scale	Per Fabric Scale
Maximum number of endpoints (EPs)	Default profile (Dual stack)— • For ALE v1 and v2: • MAC: 12,000 • IPv4: 12,000 or • IPv6: 6,000 or • IPv4: 4,000 - IPv6: 4,000 Default Profile or High LPM Profile— • For LSE or LSE2: • MAC: 24,000 • IPv4: 24,000 • IPv6: 12,000 IPv4 Scale profile— • For LSE and LSE2: • MAC: 48,000 • IPv4: 48,000 • IPv6: Not supported • For ALE v1 and v2: Not supported High Dual Stack Scale profile— • For LSE: • MAC: 64,000 • IPv4: 64,000 • IPv6: 24,000 • For LSE2: • MAC: 64,000 • IPv4: 64,000 • IPv6: 48,000 • For ALE v1 and v2: Not supported	16-slot and 8-slot modular spine switches: Max. 450,000 Proxy Database Entries in the fabric, which can be translated into any one of the following: • 450,000 MAC-only EPs (each EP with one MAC only) • 225,000 IPv4 EPs (each EP with one MAC and one IPv4) • 150,000 dual-stack EPs (each EP with one MAC, one IPv4, and one IPv6) The formula to calculate in mixed mode is as follows: $\#MAC + \#IPv4 + \#IPv6 \leq 450,000$ NOTE: Four fabric modules are required on all spines in the fabric to support above scale.

Configurable Options	Per Leaf Switch Scale	Per Fabric Scale
		4-slot modular spine switches: Max. 360,000 Proxy Database Entries in the fabric, which can be translated into any one of the following: • 360,000 MAC-only EPs (each EP with one MAC only) • 180,000 IPv4 EPs (each EP with one MAC and one IPv4) • 120,000 dual-stack EPs (each EP with one MAC, one IPv4, and one IPv6) The formula to calculate in mixed mode is as follows: $\#MAC + \#IPv4 + \#IPv6 \leq 360,000$ NOTE: Four fabric modules are required on all spines in the fabric to support above scale.

Configurable Options	Per Leaf Switch Scale	Per Fabric Scale
		Fixed spine switches: Max. 180,000 Proxy Database Entries in the fabric, which can be translated into any one of the following: • 180,000 MAC-only EPs (each EP with one MAC only) • 90,000 IPv4 EPs (each EP with one MAC and one IPv4) • 60,000 dual-stack EPs (each EP with one MAC, one IPv4, and one IPv6) The formula to calculate in mixed mode is as follows: $\#MAC + \#IPv4 + \#IPv6 \leq 180,000$
Number of Multicast Routes	Default (dual stack), IPv4, or High LPM scale profile: 8,000 with (S,G) scale not exceeding 4,000 High Dual stack Scale profile: • LSE: 512 • LSE2: 32,000 with (S,G) scale not exceeding 16,000	32,000
Number of Multicast Routes per VRF	Default (dual stack), IPv4, or High LPM scale profile: 8,000 with (S,G) scale not exceeding 4,000 High Dual stack Scale profile: • LSE: 512 • LSE2: 32,000 with (S,G) scale not exceeding 16,000	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

Configurable Options	Per Leaf Switch Scale	Per Fabric Scale
SPAN	ALE-based leaf switches: • 4 unidirectional or 2 bidirectional access/tenant sessions • 4 unidirectional or 2 bidirectional fabric sessions LSE-based leaf switches: • 8 unidirectional or 4 bidirectional sessions (fabric, access, or tenant)	N/A
Number of ports per SPAN session	ALE-based leaf switches: • All leaf access ports could be in one session. • All leaf fabric ports could be in one session. LSE/LSE2-based leaf switches: • 30 – total number of unique ports (fabric + access) across all types of SPAN sessions	N/A

Configurable Options	Per Leaf Switch Scale	Per Fabric Scale
Number of SPAN sources in each direction	ALE-based leaf switches: • $2 * (V + FP + AP1 + AP2 + (V * AP1)) \leq 460$ (ingress direction) • $2 * (V + FP + AP1 + AP2 + (V * AP1)) \leq 100$ (egress direction) LSE-based leaf switches: • $2 * (V + FP + AP1 + AP2 + (V * AP1) + AP3_v6) + AP3_v4 \leq 480$ (for each direction) Where: • V: Number of source VLANs in Tenant SPAN. Each source EPG may contain multiple VLANs. • FP: Number of source ports in Fabric SPAN • AP1: Number of source ports in Access SPAN without any filters • AP2: Number of (VLAN, Port) pairs in Access SPAN with EPG/L3Out filters. Each EPG/L3Out 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
Maximum 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 Switch Scale	Per Fabric Scale
Maximum number of Data Plane policers at the interface level	ALE: • 64 ingress policers • 64 egress policers LSE and LSE2: • 7 ingress policers • 3 egress policers	N/A
Maximum number of Data Plane policers at EPG and interface level	128 ingress policers	N/A
Maximum number of interfaces with Per-Protocol Per-Interface (PPPI) CoPP	63	N/A
Maximum 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 ppqi copp tcam-usage'</code> command to check on LSE/LSE2 platforms	N/A
Maximum number of SNMP trap receivers	10	10
IP SLA probes* *With 1 second probe time and 3 seconds of timeout	100	200
First Hop Security (FHS)* With any combination of BDs/EPGs/EPs within the supported limit	2,000 endpoints 1,000 bridge domains	N/A
Maximum number of Q-in-Q tunnels (both QinQ core and edge combined)	1,980	N/A
Maximum number of TEP-to-TEP atomic counters (tracked by 'dbgAcPathA' object)	N/A	1,600

Tenants

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of Contexts (VRFs) per tenant	ALE: 50 LSE: 128	ALE: 50 LSE: 128

Contexts

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

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of Contexts (VRFs)	ALE: 400 LSE and LSE2: 800	3,000
Maximum ECMP (equal cost multipath) for BGP best path	16	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
Maximum number of vzAny Provided Contracts	Shared services: Not supported Non-shared services: 70 per Context (VRF)	N/A
Maximum 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
Maximum number of BGP neighbors	400	3,000
Maximum number of OSPF neighbors	300 (Maximum number of VRFs with an l3out where OSPF is the only routing protocol enabled, cannot exceed 142)	N/A
Maximum number of EIGRP neighbors	16	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of IP Longest Prefix Matches (LPM) entries Note The total of (# of IPv4 prefixes) + 2*(# of IPv6 prefixes) must not exceed the scale listed for IPv4 alone	Default profile (Dual stack) - • For ALE v1 and v2: • IPv4: 10,000 or • IPv6: 6,000 or • IPv4: 4,000, IPv6: 4,000 • IPv6 wide prefixes (> /64): 1,000 • For LSE or LSE2: • IPv4: 20,000 or • IPv6: 10,000 • IPv6 wide prefixes (>= /84): 1,000 NOTE: For LSE2 and FX2 models there's no restriction on wide prefixes. IPv4 Scale Profile – • For LSE or LSE2: • IPv4: 38,000 • IPv6: Not supported • For ALE v1 and v2: Not supported High Dual Stack Scale Profile – • For LSE or LSE2: • IPv4: 38,000 or • IPv6: 19,000 • IPv6 wide prefixes (>= /84): 1,000 NOTE: For LSE2 and FX2 models there's no restriction on wide prefixes. • For ALE v1 and v2: Not supported	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum 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	High LPM Scale profile – • For LSE or LSE2: • IPv4: 128,000 or • IPv6: 64,000 • IPv6 wide prefixes ($\geq /84$): 1,000 - NOTE: For LSE2 and FX2 models there's no restriction on wide prefixes. • For ALE v1 and v2: Not supported	N/A
Maximum number of Secondary addresses per logical interface	1	1
Maximum number of L3 interfaces per Context	• 1,000 SVIs • 8 Routed interfaces • 100 sub-interfaces with or without port-channel	N/A
Maximum number of L3 interfaces	• 1,000 SVIs • 8 Routed interfaces • 1,000 sub-interfaces with or without port-channel	N/A
Maximum number of ARP entries for L3 Outs	7,500	N/A
Shared L3 Out	• IPv4 Prefixes: 2,000 or • IPv6 Prefixes: 1,000	• IPv4 Prefixes: 6,000 or • IPv6 Prefixes: 3,000
Maximum number of L3 Outs	400 For LSE and LSE2: 800	2,400 (single-stack) 1,800 (dual-stack)

External EPGs

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of External EPGs	800	ALE: 2,400 LSE: 4,000 The listed scale is calculated as a product of (Number of external EPGs)*(Number of border leaf switches for the L3Out) For example, the following combination adds up to a total of 2,000 external EPGs in the fabric (250 external EPGs * 2 border leaf switches * 4 L3Outs): • 250 External EPGs in L3Out1 on leaf1 and leaf2 • 250 External EPGs in L3Out2 on leaf1 and leaf2. • 250 External EPGs in L3Out3 on leaf3 and leaf4 • 250 External EPGs in L3Out4 on leaf3 and leaf4
Number of External EPGs per L3Out	250	600 The listed scale is calculated as a product of (Number of external EPGs per L3Out)*(Number of border leaf switches for the L3Out) For examples, 150 external EPGs on L3Out1 that is deployed on leaf1, leaf2, leaf3, and leaf4 adds up to a total of 600
Maximum 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	ALE: 1,000 IPv4 LSE: refer to LPM scale section.	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum 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	ALE: 1,000 LSE and LSE2: • Default Profile: • IPv4 (/32): 16,000 • IPv6 (/128): 12,000 Combined number of host prefixes and endpoints can't exceed 12,000. • IPv4 Profile: • IPv4 (/32): 16,000 Combined number of host prefixes, mcast routes, and endpoints can't exceed 56,000. • IPv6 (/128): 0 • High Dual Stack Profile: • IPv4 (/32): 64,000 Combined number of host prefixes, mcast routes, and endpoints can't exceed 64,000. • IPv6 (/128): 24,000 (LSE) Combined number of host prefixes and endpoints can't exceed 24,000. • IPv6 (/128): 48,000 (LSE2 only) Combined number of host prefixes and endpoints can't exceed 48,000. • High LPM Profile: • IPv4 (/32): 24,000 Combined number of host prefixes, mcast routes, and endpoints can't exceed 24,000. • IPv6 (/128): 12,000 Combined number of host prefixes and endpoints can't exceed 12,000.	N/A

Bridge Domain

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of BDs	1,980 Legacy mode: 3,500 On ALE leaf switches with multicast optimized mode: 50	15,000
Maximum number of BDs with Unicast Routing per Context (VRF)	ALE: 256 LSE: 1,000	1,750
Maximum number of subnets per BD	1,000, cannot be for all BDs.	1,000 per BD
Maximum number of EPGs per BD	3,960	4,000
Number of L2 Outs per BD	1	1
Number of BDs with Custom MAC Address	1,000 On ALE leaf switches with multicast optimized mode: 50	1,000 On ALE leaf switches with multicast optimized mode: 50
Maximum number of EPGs + L3 Outs per Multicast Group	128	128
Maximum number of BDs with L3 Multicast enabled	1,750	1,750
Maximum number of VRFs with L3 Multicast enabled	64	64
Maximum 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
Maximum number of PIM Neighbors	1,000	1,000
Maximum number of PIM Neighbors per VRF	64	64
Maximum number of L3Out physical interfaces with PIM enabled	32	N/A

Endpoint Groups (Under App Profiles)

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum 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
Maximum number of domains (physical, L2, L3)	100	N/A
Maximum number of VMM domains	N/A	• 200 VDS • 50 AVS • 50 Cisco ACI Virtual Edge
Maximum number of native encapsulations	• One per port, if a VLAN is used as a native VLAN. • Total number of ports, if there is a different native VLAN per port.	Applicable to each leaf independently
Maximum number of 802.1p encapsulations	• 1, if path binding then equals the number of ports. • If there is a different native VLAN per port, then it equals the number of ports.	Applicable to each leaf independently
Can encapsulation be tagged and untagged?	No	N/A
Maximum number of Static endpoints per EPG	Maximum endpoints	N/A
Maximum number of Subnets for inter-context access per tenant	4,000	N/A
Maximum number of Taboo Contracts per EPG	2	N/A
IP-based EPG (bare metal)	4,000	N/A
MAC-based EPG (bare metal)	4,000	N/A

Contracts

Cisco ACI supports two types of compression for policy CAM (content-addressable memory):

- **Bidirectional compression** ensures that bidirectional rules consume a single entry in the policy CAM and is supported starting with Cisco APIC release 3.2(1).
- **Policy TCAM indirection compression** enables multiple contracts to refer to the same filter rules and is supported starting with Cisco APIC release 4.0(1).

If you enable compression in release 4.0(1) or later, APIC will use either or both optimizations depending on the configuration. When enabling compression on -EX switches, APIC will apply bidirectional compression. The policy TCAM compression feature requires -FX leaf switches or newer.

Configurable Options	Per Leaf Scale	Per Fabric Scale
Security TCAM size	Default Scale profile • For ALE v1: 4,000 • For ALE v2: 40,000 • For LSE and LSE2: 64,000 IPv4 Scale profile • For LSE and LSE2: 64,000 • For ALE v1/v2: N/A High Dual Stack Scale profile • For LSE: 8,000 • For LSE2: 128,000 • For ALE v1/v2: N/A High LPM Scale profile • For LSE and LSE2: 8,000 • For ALE v1/v2: N/A	N/A
Software policy scale with Policy Table Compression enabled (Number of <code>actrlRule</code> Managed Objects)	Dual stack profile: • LSE (N9K-C9336C-FX2 only): 80,000 • LSE2 (N9K-C93180YC-FX only): 80,000 High Dual Stack profile: • LSE2 (N9K-C93180YC-FX and N9K-C93600CD-GX only) : 140,000	N/A
Approximate TCAM calculator given contracts and their use by EPGs	Number of entries in a contract X Number of Consumer EPGs X Number of Provider EPGs X 2	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of consumers (or providers) of a contract that has more than 1 provider (or consumer)	100	100
Number of consumers (or providers) of a contract that has a single provider (or consumer)	1,000	1,000
Scale guideline for the number of Consumers and Providers for the same contract	N/A	Number of consumer EPGs * number of provider EPGs * number of filters in the contract <= 50,000

FCoE NPV

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of VSANs	32	N/A
Maximum number of VFCs configured on physical ports and FEX ports	151	N/A
Maximum number of VFCs on port-channel (PC), including SAN port-channel	7	N/A
Maximum number of VFCs on virtual port-channel (vPC) interfaces, including FEX HIF vPC	151	N/A
Maximum number of FDISC per port	255	N/A
Maximum number of FDISC per leaf	1,000	N/A

FC NPV

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

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	4
Total Number of VMM domain (vCenter, Datacenter) instances.	N/A	• 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	5,000
Number of EPGs to VMware domains/vDS	N/A	5,000
Number of EPGs per vCenter/AVS	N/A	3,500
Number of EPGs to VMware domains/AVS	N/A	3,500
Number of EPGs per vCenter/Cisco ACI Virtual Edge	N/A	VLAN Mode: 1,300 VXLAN Mode: 2,000
Number of EPGs to VMware domains and Cisco ACI Virtual Edge	N/A	VLAN Mode: 1,300 VXLAN Mode: 2,000
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 (Tested with a total of 500 EPs attached to 1 vPC)	N/A
Number of Microsegment EPGs with AVS	1,000	N/A
Number of Microsegment EPGs with Cisco ACI Virtual Edge	1,000	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: 1,000

Microsoft SCVMM

Configurable Options	Per Leaf Scale (On-Demand Mode)	Per Leaf Scale (Pre-Provision Mode)	Per Fabric Scale
Number of controllers per SCVMM domain	N/A	N/A	5
Number of SCVMM domains	N/A	N/A	5
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

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	1,000	Not Supported	N/A

Microsoft Windows Azure Pack

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

Layer 4 - Layer 7 Scalability Limits

Configurable Options (L4-L7 Configurations)	Per Leaf Scale	Per Fabric Scale
Maximum number of L4-L7 logical device clusters	N/A	1,200
Maximum number of graph instances	N/A	1,000
Number of device clusters per tenant	N/A	30

Configurable Options (L4-L7 Configurations)	Per Leaf Scale	Per Fabric Scale
Number of interfaces per device cluster	N/A	Any
Number of graph instances per device cluster	N/A	500
Deployment scenario for ASA (transparent or routed)	N/A	Yes
Deployment scenario for Citrix - One arm with SNAT/etc.	N/A	Yes
Deployment scenario for F5 - One arm with SNAT/etc.	N/A	Yes

AD, TACACS, RBAC Scalability Limits

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

Cisco Mini ACI Fabric and Virtual APICs Scalability Limits

Property	Maximum Scale
Multicast Groups	200
BGP + OSPF Sessions	25
Number of Graphs Instances	20

Property	Maximum Scale
Maximum number of L4-L7 logical device clusters	3 Physical or 10 Virtual
Number of Pods	1
GOLF VRF, Route Scale	N/A
Tenants	25
Endpoints	20,000
Bridge domains (BDs)	1,000
Endpoint groups (EPGs)	1,000
VRFs	25
Number of Leafs	4
Number of Spines	2
Contracts	2,000

Cisco Cloud APIC Scalability Limits

Configurable Options	Scale
Number of Tenants	20
Number of Application Profiles	500
Number of EPGs	500
Number of Cloud End Points	1,000
Number of VRFs	20
Cloud Context Profiles	40
Number of Contracts	1,000
Number of L4-L7 Service Graphs	200
Number of L4-L7 Services Devices (AWS ALB)	100

Cisco ACI and UCSM Scalability

The following table shows verified scalability numbers for Cisco Unified Computing System with Cisco ACI [ExternalSwitch](#) app.

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

QoS Scalability Limits

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

QoS Mode	QoS Scale
Custom QoS Policy with DSCP	7
Custom QoS Policy with DSCP and Dot1P	7
Custom QoS Policy with Dot1P	38
Custom QoS Policy via a Contract	38



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.