



Verified Scalability Guide for Cisco APIC, Release 3.2(6), Multi-Site, Release 1.2(5), and Cisco Nexus 9000 Series ACI-Mode Switches, Release 13.2(6)

[Overview](#) 2

[General Scalability Limits](#) 2

[Multiple Fabric Options Scalability Limits](#) 5

[Cisco ACI Multi-Site Scalability Limits](#) 5

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

[VMM Scalability Limits](#) 21

[Layer 4 - Layer 7 Scalability Limits](#) 24

[AD, TACACS, RBAC Scalability Limits](#) 24

[QoS Scalability Limits](#) 25

Overview

This guide contains the maximum verified scalability limits for ACI parameters for the Cisco APIC Release 3.2.(6), Cisco Multi-Site Release 1.2(5), and Cisco Nexus 9000 Series ACI-Mode Switches, Release 13.2.(6). 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 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 an ACI tenant.
- **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:** Multipod enables provisioning a more fault-tolerant fabric comprised of multiple pods with isolated control plane protocols. Also, multipod 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, multipod enables provisioning multiple pods per floor or building and providing connectivity between pods through spine switches.

Multipod 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) each deployed as a single Pod in this release and managed by an independent APIC controller cluster. An ACI Multi-Site policy manager 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.

NOTE: 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

Feature	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	300 for 5- or 6-node cluster 400 for 7-node cluster
Number of spine switches	Maximum spines per pod: 6. Total spines 24.	Maximum spines per pod: 6. Total spines 24.	Maximum spines per pod: 6. Total spines 24.
Number of FEXs	N/A	20 per leaf switch, 576 ports/leaf switch, 650 per fabric	N/A
Number of tenants	N/A	1,000	3,000
Number of Layer 3 (L3) contexts (VRFs)	N/A	1,000	3,000
Number of contracts/filters	N/A	<ul style="list-style-type: none"> • 10,000 contracts • 10,000 filters 	<ul style="list-style-type: none"> • 10,000 contracts • 10,000 filters
Number of endpoint groups (EPGs)	21,000 (500 maximum per tenant, or 4,000 if there is a single tenant in a fabric)	15,000 (500 maximum per tenant, or 4,000 if there is a single tenant in a fabric)	15,000 (500 maximum per tenant, or 4,000 if there is a single tenant in a fabric)
Number of Isolation enabled EPGs	400	400	400
Number of bridge domains (BDs)	21,000	15,000	15,000
Number of BGP + number of OSPF sessions + EIGRP (for external connection)	N/A	3,000	3,000
Number of Multicast groups	N/A	8,000	8,000
Number of Multicast groups per VRF	N/A	8,000	8,000
Number of static routes to a single SVI/VRF	N/A	5,000	10,000
Number of static routes on a single leaf switch	N/A	5,000	10,000

Feature	L2 Fabric	L3 Fabric	Large L3 Fabric
Number of vCenters	N/A	<ul style="list-style-type: none"> • 200 VDS • 50 AVS • 50 Cisco ACI Virtual Edge 	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • minTx:50 • minRx:50 • multiplier:3 	256 Minimum BFD timer required to support this scale: <ul style="list-style-type: none"> • minTx:50 • minRx:50 • multiplier:3 	256 Minimum BFD timer required to support this scale: <ul style="list-style-type: none"> • minTx:50 • minRx:50 • multiplier:3
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 leaf switches overall	<ul style="list-style-type: none"> • 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	8,000	8,000

Feature	L2 Fabric	L3 Fabric	Large L3 Fabric
Number of Routes in Overlay-1 VRF	1,000	1,000	1,000

Multiple Fabric Options Scalability Limits

Configurable Options	Per Leaf Scale	Per Fabric Scale
Stretched Fabric		
Maximum number of fabrics that can be a stretched fabric	N/A	3
Maximum number of Route Reflectors	N/A	6
Multi-Pod		
Maximum number of PODs	N/A	12
Maximum number of leaf switches per POD	N/A	200
Maximum number of leaf switches overall	N/A	400

Cisco ACI Multi-Site Scalability Limits

Stretched vs. non Stretched—Stretched in multi-site means the Multi-Site fabric has stretched objects such as EPG, BD, VRF or Subnet across multiple sites or has cross site contracts between EPGs. The scale parameters for this scenario is described in the "Stretched (Multi-Site)" column. 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 scale parameters for this scenario is described in the "Non-Stretched (APIC)" column in the table below:



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

Multi-Site General Scalability Limits

Configurable Options	Scale
Sites	10
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	60
Number of Templates	300
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

Scaling Item	Stretched (Multi-Site)	Non-Stretched (APIC)
Tenants	300	2,500
VRFs	800	3,000
BDs	3,000	10,000
Contracts	3,000	2,000
End Points	100,000	150,000 including: <ul style="list-style-type: none"> • 100,000 - stretched from other sites • 50,000 - locally learned in site-local
EPGs	3,000	10,000
Isolated EPGs	400	400
MicroSegment EPGs	400	400
IGMP Snooping	8,000	8,000
L3Out external EPGs	500	2,400
Subnets	8,000	10,000
Number of L4-L7 Logical devices	400	Refer to L4-L7 Site Local Scalability Numbers
Number of graph instances	250	Refer to L4-L7 Site Local Scalability Numbers
Number of device clusters per tenant	10	Refer to L4-L7 Site Local Scalability Numbers

Scaling Item	Stretched (Multi-Site)	Non-Stretched (APIC)
Number of interfaces per device cluster	Any	Refer to L4-L7 Site Local Scalability Numbers
Number of graph instances per device cluster	125	Refer to L4-L7 Site Local Scalability Numbers

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

The table below 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 switch is listed as LSE for scalability limits purpose only. The switch is capable of supporting LSE2 platform features, consult specific feature documentation for full list of supported devices.

ALE/LSE Type	ACI-Supported TORs
ALE v1	<ul style="list-style-type: none"> • N9K-C9396PX + N9K-M12PQ • N9K-C93128TX + N9K-M12PQ • N9K-C9396TX + N9K-M12PQ
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	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.

Configurable Options	Per Leaf Scale	Per Fabric Scale
Fabric Topology		
Number of PCs, vPCs	320 (with FEX HIF)	N/A
Number of encaps per access port, PC, vPC (non FEX HIF)	3,000	N/A
Number of encaps 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
Number of static port bindings	For ALE v1 and v2: 30,000 For LSE and LSE2: 60,000	400,000
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 2,000 logical ports (port x VLAN) per leaf	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of endpoints (EPs)	Default profile (Dual stack)— <ul style="list-style-type: none"> • For ALE v1 and v2: <ul style="list-style-type: none"> • MAC: 12,000 • IPv4: 12,000 or • IPv6: 6,000 or • IPv4: 4,000 • IPv6: 4,000 Default Profile or High LPM Profile— <ul style="list-style-type: none"> • For LSE or LSE2: <ul style="list-style-type: none"> • MAC: 24,000 • IPv4: 24,000 • IPv6: 12,000 IPv4 Scale profile— <ul style="list-style-type: none"> • For LSE and LSE2: <ul style="list-style-type: none"> • MAC: 48,000 • IPv4: 48,000 • IPv6: Not supported • For ALE v1 and v2: Not supported High Dual Stack Scale profile— <ul style="list-style-type: none"> • For LSE: <ul style="list-style-type: none"> • MAC: 64,000 • IPv4: 64,000 • IPv6: 24,000 • For LSE2: <ul style="list-style-type: none"> • MAC: 64,000 • IPv4: 64,000 • IPv6: 48,000 	16-slot and 8-slot modular spine switches: Max. 450,000 Proxy Database Entries in the fabric, which can be translated into any one of the following: <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) 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.
	IPv4 Scale profile— <ul style="list-style-type: none"> • For LSE and LSE2: <ul style="list-style-type: none"> • MAC: 48,000 • IPv4: 48,000 • IPv6: Not supported • For ALE v1 and v2: Not supported High Dual Stack Scale profile— <ul style="list-style-type: none"> • For LSE: <ul style="list-style-type: none"> • MAC: 64,000 • IPv4: 64,000 • IPv6: 24,000 • For LSE2: <ul style="list-style-type: none"> • MAC: 64,000 • IPv4: 64,000 • IPv6: 48,000 	4-slot modular spine switches: Max. 360,000 Proxy Database Entries in the fabric, which can be translated into any one of the following: <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) 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.
	<ul style="list-style-type: none"> • For ALE v1 and v2: Not supported 	

Configurable Options	Per Leaf 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: <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) The formula to calculate in mixed mode is as follows: $\#MAC + \#IPv4 + \#IPv6 \leq 180,000$
Number of MAC EPGs	N/A	125
Number of Multicast Groups	Default or High LPM profile (Dual stack): 8,000 IPv4 Scale profile: 8,000 High Dual stack Scale profile: 512	8,000
Number of Multicast Groups per VRF	8,000	8,000
Number of IPs per MAC	4,096	4,096
SPAN	ALE based ToRs: <ul style="list-style-type: none"> • 4 uni-directional or 2 bi-directional access/tenant sessions • 4 uni-directional or 2 bi-directional fabric sessions LSE based ToRs: <ul style="list-style-type: none"> • 8 uni-directional or 4 bi-directional sessions (fabric, access, or tenant) 	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of ports per SPAN session	ALE-based TORs: <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 TORs: <ul style="list-style-type: none"> • 30 – total number of unique ports (fabric + access) across all types of SPAN sessions 	N/A
Number of source EPGs in tenant SPAN sessions Note The numbers listed in this row assume that only tenant SPAN is configured. If both, Access and Tenant SPAN are configured, the following formula applies for both ingress and egress SPAN: $E + P + E * P + EPP + v6FePP + 0.5 * v4FePP \leq 230$ Where: <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 	ALE-based ToR switches: <ul style="list-style-type: none"> • 230 ingress direction + 50 egress direction LSE-based ToR switches: <ul style="list-style-type: none"> • 230 bidirectional • 460 unidirectional (230 ingress + 230 egress) 	N/A
Common pervasive gateway	256 virtual IPs per Bridge Domain	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of Data Plane policers at interface level	ALE: <ul style="list-style-type: none"> • 64 ingress policers • 64 egress policers LSE and LSE2: <ul style="list-style-type: none"> • 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 pppi copp tcam-usage'</code> command to check on LSE/LSE2 platforms	N/A
Maximum number of SNMP trap receivers	10	10
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		
Number of Contexts (VRFs) per tenant	ALE: 50 LSE: 128	ALE: 50 LSE: 128
Number of application profiles per tenant (or per Context (VRF))	N/A	N/A
Contexts (All numbers 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	N/A	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 service graphs per device cluster	N/A	500
L3 Out per context (VRF)	--	400
Maximum number of Routed, Routed Sub-interface, or SVIs per L3 Out	<ul style="list-style-type: none"> • 8 for Routed and Routed sub-interface • 1000 for SVI 	<ul style="list-style-type: none"> • 12 for Routed and Routed sub-interface • 1000 for SVI
Maximum number of BGP neighbors	400	3000
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
<p>Maximum 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 profile (Dual stack) -</p> <ul style="list-style-type: none"> • For ALE v1 and v2: <ul style="list-style-type: none"> • IPv4: 10,000 or • IPv6: 6,000 or • IPv4: 4,000, IPv6: 4,000 • IPv6 wide prefixes (> /64): 1,000 • For LSE or LSE2: <ul style="list-style-type: none"> • IPv4: 20,000 or • IPv6: 10,000 • IPv6 wide prefixes (>= /84): 1,000 <p>NOTE: For LSE2 and FX2 models there's no restriction on wide prefixes</p> <p>IP4 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 • For ALE v1 and v2: Not supported <p>High Dual Stack Scale Profile –</p> <ul style="list-style-type: none"> • For LSE or LSE2: <ul style="list-style-type: none"> • IPv4: 38,000 or • IPv6: 19,000 • IPv6 wide prefixes (>= /84): 1,000 <p>NOTE: For LSE2 and FX2 models there's no restriction on wide prefixes</p> • For ALE v1 and v2: Not supported 	<p>N/A</p>

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 – <ul style="list-style-type: none"> • For LSE or LSE2: <ul style="list-style-type: none"> • 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 (SVIs and sub-interfaces)	<ul style="list-style-type: none"> • 1,000 for SVI • 32 for sub-interfaces with or without port-channel 	N/A
Maximum number of ARP entries for L3 Outs	7,500	N/A
Shared L3 Out	<ul style="list-style-type: none"> • IPv4: 2,000 or • IPv6: 1,000 	<ul style="list-style-type: none"> • IPv4: 6,000 or • IPv6: 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	<p>ALE: 2,400</p> <p>LSE: 4,000</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 L3 Out	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>
Maximum number of LPM Prefixes for External EPG Classification	<p>ALE: 1,000 IPv4</p> <p>LSE: refer to LPM scale section</p>	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of host prefixes for External EPG Classification	<p>ALE: 1,000</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 Profile: <ul style="list-style-type: none"> • IPv4 (/32): 16,000 <p>(combined number of host prefixes, mcast routes, and endpoints can't exceed 56,000)</p> • IPv6 (/128) : 0 • High Dual Stack Profile: <ul style="list-style-type: none"> • IPv4 (/32): 64,000 <p>(combined number of host prefixes, mcast routes, 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>	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of host prefixes for External EPG Classification (Continued)	<ul style="list-style-type: none"> • High LPM Profile: <ul style="list-style-type: none"> • 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		
Maximum number of BDs	1,980 Legacy mode: 3,500 On ALE TORs with multicast optimized mode: 50	15,000
Maximum number of BDs with Unicast Routing per Context (VRF)	ALE: 256 LSE: 1000	1750
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 If on ALE TORs with multicast optimized mode: 50	1,000 If on ALE TORs 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

Configurable Options	Per Leaf Scale	Per Fabric Scale
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)		
Maximum amount of EPGs	Normally 3,960; if legacy mode 3,500	15,000
Maximum amount of encaps 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 encaps per EPG per port	One (path or leaf binding)	N/A
Maximum number of domains (physical, L2, L3)	<ul style="list-style-type: none"> • 10 static (L2, L3, physical) • 10 dynamic 	N/A
Maximum number of VMM domains	N/A	<ul style="list-style-type: none"> • 200 vDS • 50 AVS • 50 Cisco ACI Virtual Edge
Maximum amount of native encaps	<ul style="list-style-type: none"> • 1 per port (if a VLAN is used as a native VLAN) • If there is a different native VLAN per port then it equals the number of ports 	Applicable to each leaf independently
Maximum amount of 802.1p encaps	<ul style="list-style-type: none"> • 1, if path binding then equals 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 encap 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

Configurable Options	Per Leaf Scale	Per Fabric Scale
IP-based uSeg EPG (bare metal)	4,000	N/A
MAC-based uSeg EPG (bare metal)	4,000	N/A
Contracts		
Security TCAM size	Default Scale profile <ul style="list-style-type: none"> • For ALE v1: 4,000 • For ALE v2: 40,000 • For LSE and LSE2: 64,000 IPv4 Scale profile <ul style="list-style-type: none"> • For LSE and LSE2: 64,000 • For ALE v1/v2: N/A High Dual Stack Scale profile <ul style="list-style-type: none"> • For LSE: 8,000 • For LSE2: 128,000 • For ALE v1/v2: N/A High LPM Scale profile <ul style="list-style-type: none"> • For LSE and LSE2: 8,000 • For ALE v1/v2: N/A 	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
Maximum number of EPGs providing the same contract	100	100
Maximum number of EPGs consuming the same contract	100	100
Scale guideline for # of consumers/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		
Maximum number of VSANs	32	N/A
Maximum number of VFCs configured on physical ports and FEX ports	151	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of VFCs on port-channel (PC) interfaces and virtual port-channel (vPC) interfaces	7	N/A
Maximum number of FDISC per port	96	N/A
Maximum number of FDISC per leaf	96	N/A
FC NPV		
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

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	<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	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

Configurable Options	Per Leaf Scale	Per Fabric Scale
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/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
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
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

Configurable Options	Per Leaf Scale	Per Fabric Scale
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
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 priority	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

QoS Scalability Limits

The table below shows QoS scale limits. The scale numbers depend on whether remote leafs are present in the topology as well as MPOD QoS Policy and CoS Preservation settings.

		QoS Scale with Remote Leaf in Topology	QoS Scale without Remote Leaf in Topology
MPOD QoS Policy enabled	Custom QOS Policy with DSCP	8	9
	Custom QOS Policy with DSCP and Dot1P	8	9
	Custom QOS Policy with Dot1P	43	48
	Custom QOS Policy via a Contract	43	48
CoS Preservation enabled	Custom QOS Policy with DSCP	9	9
	Custom QOS Policy with DSCP and Dot1P	9	9
	Custom QOS Policy with Dot1P	48	48
	Custom QOS Policy via a Contract	48	48



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.