



## **Verified Scalability Guide for Cisco APIC, Release 3.1(1i) and Cisco Nexus 9000 Series ACI-Mode Switches, Release 13.1(1i)**

### **Overview 2**

#### **General Scalability Limits 2**

#### **Multiple Fabric Options Scalability Limits 4**

#### **Cisco ACI Multi-Site 5**

#### **Fabric Topology, SPAN, Tenants, Contexts (VRFs), External EPGs, Bridge Domains, Endpoints, and Contracts Scalability Limits 6**

#### **VMM Scalability Limits 17**

#### **Layer 4 - Layer 7 Scalability Limits 19**

#### **AD, TACACS, RBAC Scalability Limits 19**

#### **QoS Scalability Limits 20**

# Overview

This guide contains the maximum verified scalability limits for ACI parameters for the Cisco APIC Release 3.1.1i and Cisco Nexus 9000 Series ACI-Mode Switches, Release 13.1.1i. 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 leafs total which is the maximum number of leafs 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 30 per fabric

Feature	L2 Fabric	L3 Fabric	Large L3 Fabric
Number of APIC controllers	3	Minimum 3 (4 also supported)	5
Number of leaf switches	80	80	200

Feature	L2 Fabric	L3 Fabric	Large L3 Fabric
Number of spines	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, 320 ports/leaf, 650 per fabric	N/A
Number of tenants	N/A	1000	3000
Number of Layer 3 (L3) contexts (VRFs)	N/A	1000	3000
Number of contracts/filters	N/A	<ul style="list-style-type: none"> <li>• 10,000 contracts</li> <li>• 10,000 filters</li> </ul>	<ul style="list-style-type: none"> <li>• 10,000 contracts</li> <li>• 10,000 filters</li> </ul>
Number of endpoint groups (EPGs)	21,000 ( 500 maximum per tenant, or 4,000 in a single tenant in a fabric )	15,000 (500 maximum per tenant, or 4,000 in a single tenant in a fabric)	15,000 (500 maximum per tenant, or 4,000 in 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	8000	8000
Number of Multicast groups per VRF	N/A	8000	8000
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
Number of vCenters	N/A	<ul style="list-style-type: none"> <li>• 200 VDS</li> <li>• 50 AVS</li> <li>• 50 Cisco ACI Virtual Edge</li> </ul>	<ul style="list-style-type: none"> <li>• 200 VDS</li> <li>• 50 AVS</li> <li>• 50 Cisco ACI Virtual Edge</li> </ul>
Number of Service Chains	N/A	1000	1000
Number of L4 - L7 devices	N/A	30 physical, 1,200 virtual (1,200 maximum per fabric)	30 physical, 1,200 virtual (1,200 maximum per fabric)
Number of ESXi hosts - VDS	N/A	3200	3200
Number of ESXi hosts - AVS	N/A	3200 (Only 1 AVS instance per host)	3200 (Only 1 AVS instance per host)

Feature	L2 Fabric	L3 Fabric	Large L3 Fabric
Number of ESXi hosts - AVE	N/A	3200 (Only 1 AVE instance per host)	3200 (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"> <li>• minTx:50</li> <li>• minRx:50</li> <li>• multiplier:3</li> </ul>	256 Minimum BFD timer required to support this scale: <ul style="list-style-type: none"> <li>• minTx:50</li> <li>• minRx:50</li> <li>• multiplier:3</li> </ul>	256 Minimum BFD timer required to support this scale: <ul style="list-style-type: none"> <li>• minTx:50</li> <li>• minRx:50</li> <li>• multiplier:3</li> </ul>
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"> <li>• 5* or 6 node APIC cluster,6 pods, 200 leaf switches max per pod, 300 leaf switches max overall</li> <li>• 7 node APIC cluster,12 pods, 200 leaf switches max per pod, 400 leaf switches max overall</li> </ul>
L3 EVPN Services over Fabric WAN - GOLF (with and without OpFlex)	N/A	1000 VRFs, 60,000 routes in a fabric	1000 VRFs, 60,000 routes in a fabric
Layer 3 Multicast routes	N/A	8,000	8,000
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
<b>Stretched Fabric</b>		
Maximum number of fabrics that can be a stretched fabric	N/A	3
Maximum number of Route Reflectors	N/A	6
<b>Multi-Pod</b>		
Maximum number of PODs	N/A	12

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of leaf switches per POD	N/A	200
Maximum number of leaf switches overall	N/A	400

## Cisco ACI Multi-Site

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 (ie. EPG, contract, 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	8
Leaf switches per site	200

### Multi-Site Object Scale

Configurable Options	Scale
Policy Objects per Schema	500
Templates per Schema	5
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	200	2500
VRFs	400	3000
BDs	2000	10,000
Contracts	2000	2000

Scaling Item	Stretched (Multi-Site)	Non-Stretched (APIC)
End Points	50,000	100,000 including: <ul style="list-style-type: none"> <li>• 50,000 - stretched from other sites</li> <li>• 50,000 - locally learned in site-local</li> </ul>
EPGs	2000	10,000
Isolated EPGs	400	400
MicroSegment EPGs	400	400
IGMP Snooping	8000	8000
L3Out external EPGs	500	2400
Subnets	2000	10,000

## 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.

ALE/LSE Type	ACI-Supported TORs
ALE v1	<ul style="list-style-type: none"> <li>• N9K-C9396PX + N9K-M12PQ</li> <li>• N9K-C93128TX + N9K-M12PQ</li> <li>• N9K-C9396TX + N9K-M12PQ</li> </ul>
ALE v2	<ul style="list-style-type: none"> <li>• N9K-C9396TX + N9K-M6PQ</li> <li>• N9K-C93128TX + N9K-M6PQ</li> <li>• N9K-C9396PX + N9K-M6PQ</li> <li>• N9K-C9372TX 64K</li> <li>• N9K-C9332PQ</li> <li>• N9K-C9372PX</li> </ul>
LSE	N9K-C93108TC-EX + N9K-C93180YC-EX
LSE2	N9K-C93108TC-FX + N9K-C93180YC-FX

NOTE: Unless explicitly called out, LSE represents both LSE and LSE2 and ALE represents both ALE v1 and ALE v2 in table below:

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
<b>Fabric Topology</b>		
Number of PCs, vPCs	320 (with FEX HIF)	N/A
Number of encaps per access port, PC, vPC (non FEX HIF)	1750	N/A
Number of encaps per FEX HIF, PC, vPC	20	N/A
Number of member links per PC, vPC* *vPC total ports = 16, 8 per leaf	8	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
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
Maximum number of endpoints (EPs)	<p>Default profile (Dual stack) -</p> <ul style="list-style-type: none"> <li>For ALE v1 and v2: <ul style="list-style-type: none"> <li>MAC: 12,000</li> <li>IPv4: 12,000 or</li> <li>IPv6: 6000 or</li> <li>IPv4: 4000, IPv6: 4000</li> </ul> </li> <li>For LSE or LSE2: <ul style="list-style-type: none"> <li>MAC: 24,000</li> <li>IPv4: 24,000</li> <li>IPv6: 12,000</li> </ul> </li> </ul> <p>IPv4 Scale profile—</p> <ul style="list-style-type: none"> <li>For LSE and LSE2: <ul style="list-style-type: none"> <li>MAC: 48,000</li> <li>IPv4: 48,000</li> <li>IPv6: Not supported</li> </ul> </li> <li>For ALE v1 and v2: Not supported</li> </ul> <p>High Dual Stack Scale profile—</p> <ul style="list-style-type: none"> <li>For LSE and LSE2: <ul style="list-style-type: none"> <li>MAC: 64,000</li> <li>IPv4: 64,000</li> <li>IPv6: 24,000</li> </ul> </li> <li>For ALE v1 and v2: Not supported</li> </ul>	<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"> <li>360,000 MAC-only EPs(each EP with one MAC only)</li> <li>180,000 IPv4 EPs(each EP with one MAC and one IPv4)</li> <li>120,000 dual-stack EPs(each EP with one MAC, one IPv4 and one IPv6)</li> </ul> <p>The formula to calculate in mixed mode is as follows: #MAC + #IPv4 + #IPv6 &lt;= 360,000</p> <p>NOTE: Four fabric modules are required on all spines in the fabric to support above scale.</p> <hr/> <p>Fixed spine switches (N9K-C9364C):</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"> <li>180,000 MAC-only EPs (each EP with one MAC only)</li> <li>90,000 IPv4 EPs (each EP with one MAC and one IPv4)</li> <li>60,000 dual-stack EPs (each EP with one MAC, one IPv4, and one IPv6)</li> </ul> <p>The formula to calculate in mixed mode is as follows: #MAC + #IPv4 + #IPv6 &lt;= 180,000</p>
Number of MAC EPGs	N/A	125
Number of Multicast Groups	<p>Default profile (Dual stack): 8,000</p> <p>IPv4 Scale profile: 8,000</p> <p>High Dual stack Scale profile: 0</p>	8000
Number of Multicast Groups per VRF	8000	8000
Number of IPs per MAC	1024	1024



Configurable Options	Per Leaf Scale	Per Fabric Scale
SPAN	<p>ALE based ToRs:</p> <ul style="list-style-type: none"> <li>• 4 uni-directional or 2 bi-directional access/tenant sessions</li> <li>• 4 uni-directional or 2 bi-directional fabric sessions</li> </ul> <p>LSE based ToRs:</p> <ul style="list-style-type: none"> <li>• 8 uni-directional or 4 bi-directional sessions (fabric, access, or tenant)</li> </ul>	N/A
Number of ports per SPAN session	<ul style="list-style-type: none"> <li>• All leaf access ports could be in one session</li> <li>• All leaf fabric ports could be in one session</li> </ul> <p>NOTE: For LSE/LSE2 only: 30 (total number of unique ports (fabric + access) across all types of span sessions)</p>	N/A
<p>Number of source EPGs in tenant SPAN sessions</p> <p><b>Note</b> 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> $E + P + E * P + EPP + v6FePP + 0.5 * v4FePP \leq 230$ <p>Where:</p> <ul style="list-style-type: none"> <li>• E— Number of source EPGs in Tenant SPAN</li> <li>• P—Number of source Ports in access SPAN without any filters</li> <li>• EPP—Number of (Epg,Port) Pairs in access SPAN with EPG filter only (no filter group)</li> <li>• v4FePP—Number of (v4 filter entry, Port) Pairs in access SPAN with filter group</li> <li>• v6FePP—Number of (v6 Filter entry, Port) Pairs in access SPAN with filter group</li> </ul>	<p>ALE-based ToR switches:</p> <ul style="list-style-type: none"> <li>• 230 ingress direction + 50 egress direction</li> </ul> <p>LSE-based ToR switches:</p> <ul style="list-style-type: none"> <li>• 230 bidirectional</li> <li>• 460 unidirectional (230 ingress + 230 egress)</li> </ul>	N/A

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
Common pervasive gateway	256 virtual IPs per Bridge Domain	N/A
Maximum number of Data Plane policers at interface level	ALE: <ul style="list-style-type: none"> <li>• 64 ingress policers</li> <li>• 64 egress policers</li> </ul> LSE and LSE2: <ul style="list-style-type: none"> <li>• 7 ingress policers</li> <li>• 3 egress policers</li> </ul>	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	2000 endpoints 1000 bridge domains	N/A
Maximum number of Q-in-Q tunnels (both QinQ core and edge combined)	1980	N/A
Maximum number of TEP-to-TEP atomic counters	N/A	1600
<b>Tenants</b>		
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
<b>Contexts (All numbers applicable to dual stack unless explicitly called out)</b>		

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
Maximum number of Contexts (VRFs)	400	3,000
Maximum ECMP (equal cost multi-path) for BGP best path	16	N/A
Number of isolated EPGs	N/A	400
Border Leafs per L3 Out	N/A	8
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"> <li>• 8 for Routed and Routed sub-interface</li> <li>• 1000 for SVI</li> </ul>	<ul style="list-style-type: none"> <li>• 8 for Routed and Routed sub-interface</li> <li>• 1000 for SVI</li> </ul>
Maximum number of BGP neighbors	400	3000
Maximum number of BGP neighbors with authentication enabled	150	N/A
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	<p>Default profile (Dual stack) -</p> <ul style="list-style-type: none"> <li>• For ALE v1 and v2: <ul style="list-style-type: none"> <li>• IPv4: 10,000 or</li> <li>• IPv6: 6,000 or</li> <li>• IPv4: 4,000, IPv6: 4,000</li> <li>• IPv6 wide prefixes (&gt; /64): 1,000</li> </ul> </li> <li>• For LSE or LSE2: <ul style="list-style-type: none"> <li>• IPv4: 20,000 or</li> <li>• IPv6: 10,000</li> <li>• IPv6 wide prefixes (&gt;= /84): 1,000</li> </ul> </li> </ul> <p>NOTE: (only for LSE. For LSE2 there's no restriction on wide prefixes)</p> <p>IP4 Scale profile -</p> <ul style="list-style-type: none"> <li>• For LSE or LSE2: <ul style="list-style-type: none"> <li>• IPv4: 38,000</li> <li>• IPv6: Not supported</li> </ul> </li> <li>• For ALE v1 and v2: Not supported</li> </ul> <p>High Dual Stack Scale profile -</p> <ul style="list-style-type: none"> <li>• For LSE or LSE2: <ul style="list-style-type: none"> <li>• IPv4: 38,000 or</li> <li>• IPv6: 19,000</li> <li>• IPv6 wide prefixes (&gt;= /84): 1,000</li> </ul> </li> <li>NOTE: (only for LSE. For LSE2 there's no restriction on wide prefixes)</li> <li>• For ALE v1 and v2: Not supported</li> </ul>	<ul style="list-style-type: none"> <li>• IPv4; 40,000 or</li> <li>• IPv6; 20,000 or</li> <li>• IPv4; 10,000, IPv6; 10,000</li> </ul>
Maximum number of Secondary addresses per logical interface	1	1

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of L3 interfaces per Context (SVIs and sub-interfaces)	<ul style="list-style-type: none"> <li>• 200 for SVI</li> <li>• 32 for subinterfaces</li> </ul>	N/A
Maximum number of ARP entries for L3 Outs	7500	N/A
Shared L3 Out	<ul style="list-style-type: none"> <li>• IPv4: 2000 or</li> <li>• IPv6: 1000</li> </ul>	<ul style="list-style-type: none"> <li>• IPv4: 6000 or</li> <li>• IPv6: 3000</li> </ul>
Maximum number of L3 Outs	400 (per leaf scale)	2400 (single-stack) 1800 (dual-stack)
<b>External EPGs</b>		
Number of External EPGs	600	<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"> <li>• 250 External EPGs in L3Out1 on leaf1 and leaf2</li> <li>• 250 External EPGs in L3Out2 on leaf1 and leaf2.</li> <li>• 250 External EPGs in L3Out3 on leaf3 and leaf4</li> <li>• 250 External EPGs in L3Out4 on leaf3 and leaf4</li> </ul>
Number of External EPGs per L3 Out	250	<p>400</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, 100 external EPGs on L3Out1 that is deployed on leaf1, leaf2, leaf3, and leaf4 adds up to a total of 400</p>

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
Maximum number of LPM Prefixes for External EPG Classification	1000 IPv4	N/A
<b>Bridge Domain</b>		
Maximum number of BDs	1750 If legacy mode: 3,500 If 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	3499 (cannot exceed 3,500 total) 3499 is supported in hardware but please refer to the per fabric scale for the effective software support for this release.	3499
Number of L2 Outs per BD	1	1
Number of BDs with Custom MAC Address	1750 if legacy mode: 3500 If on ALE TORs with multicast optimized mode: 50	1750 if legacy mode: 3500 If on ALE TORs with multicast optimized mode: 50
Number of Multicast groups	8000	8000
Maximum number of EPGs + L3 Outs per Multicast Group	128	128
Maximum number of BDs with L3 Multicast enabled	1750	1750
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

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
Maximum number of PIM Neighbors	1000	1000
Maximum number of PIM Neighbors per VRF	64	64
Maximum number of L3Out physical interfaces with PIM enabled	32	N/A
<b>Endpoint Groups (Under App Profiles)</b>		
Maximum amount of EPGs	Normally 1750; if legacy mode 3500	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"> <li>• 10 static (L2, L3, physical)</li> <li>• 10 dynamic</li> </ul>	N/A
Maximum number of VMM domains	<ul style="list-style-type: none"> <li>• 200 vDS</li> <li>• 50 AVS</li> <li>• 50 Cisco ACI Virtual Edge</li> </ul>	N/A
Maximum amount of native encaps	<ul style="list-style-type: none"> <li>• 1 per port (if a VLAN is used as a native VLAN)</li> <li>• If there is a different native VLAN per port then it equals the number of ports</li> </ul>	Applicable to each leaf independently
Maximum amount of 802.1p encaps	<ul style="list-style-type: none"> <li>• 1, if path binding then equals number of ports</li> <li>• If there is a different native VLAN per port then it equals the number of ports</li> </ul>	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	4000	N/A
Maximum number of Taboo Contracts per EPG	2	N/A
IP-based EPG	4000	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
<b>Contracts</b>		
Security TCAM size	Default Scale profile <ul style="list-style-type: none"> <li>• For ALE v1: 4000</li> <li>• For ALE v2: 40,000</li> <li>• For LSE and LSE2: 61,000</li> </ul> <b>Note</b> For TOR to ALE mapping, see the reference table below. IPv4 Scale profile <ul style="list-style-type: none"> <li>• For LSE and LSE2: 61,000</li> <li>• For ALE v1/v2: N/A</li> </ul> High Dual Stack Scale profile <ul style="list-style-type: none"> <li>• For LSE and LSE2: 8,000</li> <li>• For ALE v1/v2: N/A</li> </ul>	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
<b>FEX VPC</b>		
Maximum EPGs behind FEX VPC port	20	N/A
<b>FCoE</b>		
Maximum number of VSAN	32	N/A
Maximum number of VFCs configured on physical ports and FEX interfaces.	151	N/A
Maximum number of VFCs configured on port-channel (PC) interfaces and virtual port-channel (vPC) interfaces.	7	N/A



Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of FDISC per port	96	N/A
Maximum number of FDISC per USB	96	N/A

## VMM Scalability Limits

Configurable Options	Per Leaf Scale	Per Fabric Scale
<b>VMware</b>		
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 VMware VMM Domains & Controllers (vCenter/vShield)	N/A	<ul style="list-style-type: none"> <li>• 200 vDS</li> <li>• 50 AVS</li> <li>• 50 Cisco ACI Virtual Edge</li> </ul> #VMware Domains x #Datacenters per vCenter
Number of ESX hosts per AVS	240	N/A
Number of ESX hosts per 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: 1300 VXLAN Mode: 2000
Number of EPGs to VMware domains/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

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
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 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: 1000
<b>Microsoft</b>		
Number of controllers per SCVMM domain	N/A	5
Number of SCVMM domains	N/A	5
EPGs per Microsoft VMM domain	N/A	3,000
EPGs per all Microsoft VMM domains	N/A	9,000
EP/VNICs per HyperV host	N/A	100
EP/VNICs per SCVMM	N/A	3,000
Number of logical switch per host	N/A	1
Number of uplinks per logical switch	N/A	4
Number of Windows Azure Pack subscriptions	N/A	1,000
Number of plans per Windows Azure Pack instance	N/A	150

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
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
Microsoft microsegmentation	1,000	N/A

## Layer 4 - Layer 7 Scalability Limits

<b>Configurable Options (L4-L7 Configurations)</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
Maximum number of L4-L7 logical device clusters	N/A	1,200
Maximum number of graph instances	N/A	1,000
Maximum number of VIPs per graph instance	N/A	1
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

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
Number of ACS/AD/LDAP authorization domains	N/A	4 tested (16 maximum /server type)
Number of login domains	N/A	15 (can go beyond)

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
Number of security domains/APIC	N/A	15 (can go beyond)
Number of security domains in which the tenant resides	N/A	4 (can go beyond)
Number of priority	N/A	4 tested (16 per domain)
Number of shell profiles that can be returned	N/A	4 tested (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.

		<b>QoS Scale with Remote Leaf in Topology</b>	<b>QoS Scale without Remote Leaf in Topology</b>
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](http://www.cisco.com/go/offices).