



## **Verified Scalability Guide for Cisco APIC, Release 2.3(1e) and Cisco Nexus 9000 Series ACI-Mode Switches, Release 12.3(1e)**

### **Overview 2**

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

Fabric Topology, SPAN, Tenants, Contexts, External EPGs, Bridge Domains, Endpoints, and Contracts

#### **Scalability Limits 4**

#### **Multiple Fabric Options Scalability Limits 11**

#### **VMM Scalability Limits 12**

#### **Layer 4 - Layer 7 Scalability Limits 14**

#### **AD, TACACS, RBAC Scalability Limits 14**

# Overview

This guide contains the maximum verified scalability limits for ACI parameters for the Cisco APIC Release 2.3.1e and Cisco Nexus 9000 Series ACI-Mode Switches, Release 12.3.1e. 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.

**NOTE:** The maximum number of leaf switches overall is 400 per fabric scale and the maximum number of physical ports is 19,200 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
Number of spines	Maximum spines per pod: 6. Total spines 16.	Maximum spines per pod: 6. Total spines 16.	Maximum spines per pod: 6. Total spines 16.

Feature	L2 Fabric	L3 Fabric	Large L3 Fabric
Number of FEXs	N/A	20 FEXes per leaf, 320 FEX ports/leaf, 650 FEXes per fabric.	N/A
Number of tenants	N/A	1000	3000
Number of Layer 3 (L3) contexts	N/A	1000	3000
Number of contracts/filters	N/A	<ul style="list-style-type: none"> <li>• 2000 contracts</li> <li>• 10,000 filters</li> </ul>	<ul style="list-style-type: none"> <li>• 2000 contracts</li> <li>• 10,000 filters</li> </ul>
Number of endpoint groups (EPGs)	21,000 (500 maximum per tenant)	15,000 (500 maximum per tenant)	15,000 (500 maximum per tenant)
Number of Isolation enabled EPGs	250	250	250
Number of endpoints (EPs)	180,000	180,000	180,000
Number of bridge domains (BDs)	21,000	15,000	15,000
Number of IP longest prefix matches (IP LPMs) (for external connection) <b>Note</b> This limit exists across all protocols/transit scenarios	N/A	<ul style="list-style-type: none"> <li>• IPv4; 40,000</li> <li>• IPv6; 20,000</li> <li>• Per Leaf IPv4 and IPv6; 10,000 total</li> </ul>	<ul style="list-style-type: none"> <li>• IPv4; 40,000</li> <li>• IPv6; 20,000</li> <li>• Per Leaf IPv4 and IPv6; 10,000 total</li> </ul>
Number of BGP + number of OSPF sessions + EIGRP (for external connection)	N/A	1,200	1,200
Number of Multicast groups	N/A	8000	8000
Number of Multicast groups per VRF	N/A	8000	8000
Number of vCenters	N/A	<ul style="list-style-type: none"> <li>• 200</li> <li>• 50 AVS</li> </ul>	<ul style="list-style-type: none"> <li>• 200</li> <li>• 50 AVS</li> </ul>
Number of Service Chains	N/A	1000	1000
Number of L4 - L7 devices	N/A	30 physical, 1,200 virtual (1200 maximum per fabric)	30 physical, 1,200 virtual (1500 maximum per fabric)
Number of ESXi hosts - VDS	N/A	3200	3200

Feature	L2 Fabric	L3 Fabric	Large L3 Fabric
Number of ESXi hosts - AVS	N/A	3200 (Only 1 AVS instance per host)	3200 (Only 1 AVS 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	<ul style="list-style-type: none"> <li>• 256 (physical interface, VRF)</li> <li>• 512 per Leaf</li> </ul>	<ul style="list-style-type: none"> <li>• 256 (physical interface, VRF)</li> <li>• 512 per Leaf</li> </ul>	<ul style="list-style-type: none"> <li>• 256 (physical interface, VRF)</li> <li>• 512 per Leaf</li> </ul>
Multi-Pod	3 node APIC cluster, 6 pods, 80 leaf switches overall	3 node APIC cluster, 6 pods, 80 leaf switches overall	<ul style="list-style-type: none"> <li>• 5 node APIC cluster, 6 pods, 200 leaf switches max per pod, 300 leaf switches max overall</li> <li>• 7 node APIC cluster, 10 pods, 200 leaf switches max per pod, 400 leaf switches max overall</li> </ul>
L3 eVPN Services over Fabric WAN (with and without OpFlex)	N/A	1000 VRFs, 1000 L3outs, 60,000 routes in a fabric	1000 VRFs, 60,000 routes in a fabric
Layer 3 Multicast	N/A	8000 multicast routes	8000 multicast routes

## Fabric Topology, SPAN, Tenants, Contexts, External EPGs, Bridge Domains, Endpoints, and Contracts Scalability Limits

Configurable Options	Per Leaf Scale	Per Fabric Scale
<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	8	N/A

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
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 ports x VLANS (static port bindings)	For ALE v1 and v2: 30,000 For LSE and LSE2: 60,000	400,000
STP	All VLANs	N/A
Maximum number of endpoints (EPs)	For ALE v1 and v2: <ul style="list-style-type: none"> <li>• IPv4: 12,000 or</li> <li>• IPv6: 6000 or</li> <li>• IPv4: 4000, IPv6; 4000</li> </ul> For LSE and LSE2: <ul style="list-style-type: none"> <li>• MAC: 24,000</li> <li>• IPv4: 24,000</li> <li>• IPv6: 12,000</li> </ul>	180,000
Number of MAC EPGs	N/A	125
Number of Multicast Groups	8000	8000
Number of Multicast Groups per VRF	8000	8000
Number of IPs per MAC	1024	1024
SPAN	<ul style="list-style-type: none"> <li>• 4 uni-directional or 2 bi-directional infra/tenant sessions</li> <li>• 4 uni-directional or 2 bi-directional fabric sessions (VXLAN is carried in the SPAN)</li> </ul>	N/A

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
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>	N/A
Number of source EPGs in tenant sessions (Note: Number of source EPGs above presumes that only tenant span is configured)	<p>ALE based TORs:</p> <ul style="list-style-type: none"> <li>• 230 ingress direction + 50 egress direction</li> </ul> <p>LSE based TORs:</p> <ul style="list-style-type: none"> <li>• 230 bi-directional</li> <li>• 460 uni-directional</li> </ul>	N/A
Common pervasive gateway	256 virtual IPs per Bridge Domain	N/A
Maximum number of Data Plane policers	<ul style="list-style-type: none"> <li>• 64 ingress policers</li> <li>• 64 egress policers</li> </ul> <p>For all EX series switches:</p> <ul style="list-style-type: none"> <li>• 7 ingress policers</li> <li>• 7 egress policers</li> </ul>	N/A
Maximum number of SNMP trap receivers	10	10
<b>Tenants</b>		
Number of Contexts per tenant	50	50
Number of application profiles per tenant (or per Context)	N/A	N/A
<b>Contexts (All numbers applicable to dual stack unless explicitly called out)</b>		
Maximum number of Context	400	N/A
Maximum number of BDs per Context	256	N/A
Number of VRFs per tenant	N/A	50
Number of BDs per VRF	N/A	1750

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
Number of isolated EPGs	N/A	250
Border Leafs per L3 Out	N/A	8
Maximum number of LPM Prefixes for External EPG Classification	1000 IPv4	N/A
Maximum number of vzAny Provided Contracts	16 per Ctx	N/A
Maximum number of vzAny Consumed Contracts	16 per Ctx	N/A
Number of service graphs per device cluster	N/A	500
L3 Out per context	--	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 Dynamic Routing protocol peers for BGP	400	2400
Maximum number of Dynamic Routing protocol peers for OSPF	300	N/A
Maximum number of Dynamic Routing protocol peers for EIGRP	16	N/A
Maximum number of Static Routes	<ul style="list-style-type: none"> <li>• IPv4; 10,000 or</li> <li>• IPv6; 6000 or</li> <li>• IPv4; 4000, IPv6; 4000</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 External Routes	<p>For ALE v1 and v2:</p> <ul style="list-style-type: none"> <li>• IPv4; 10,000 or</li> <li>• IPv6; 5000 or</li> </ul> <p>For LSE and LSE2:</p> <ul style="list-style-type: none"> <li>• IPv4; 20,000 or</li> <li>• IPv6; 10,000</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>

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
Maximum number of Secondary (VIP) addresses per logical interface	1	1
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>	<ul style="list-style-type: none"> <li>• 400 for SVI</li> <li>• 32 for subinterfaces</li> </ul>
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>
<b>Configurable Options</b>	<b>per Leaf scale</b>	<b>per Fabric scale</b>
Maximum number of L3 Outs	400 (per leaf scale)	2400 (single stack)
<b>External EPGs</b>		
Number of External EPGs	600	2400 (single stack)
Number of External EPGs per L3 out	250	400
<b>Bridge Domain</b>		
Maximum amount of BDs	1750; if legacy mode, 3,500; if Multicast optimized mode then 50	15,000
Maximum number of subnets per BD	512 (cannot be for all BDs)	512 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 Multicast optimized mode is used, then 50	1750; if legacy mode, 3500; If Multicast optimized mode is used, then 50
Number of Multicast groups	8000	8000



<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
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	4	N/A
Number of DHCP relay labels per BD	2	2
DHCP relay for secondary subnets in a BD	No	No
Number of external EPGs per L2 out	1	1
Maximum number of PIM Neighbors	1000	1000
Maximum number of PIM Neighbors per VRF	64	64
<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> </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

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
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 encaps 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
<b>Contracts</b>		
Security TCAM size	<ul style="list-style-type: none"> <li>• 4000 (for ALE v1)</li> <li>• 40,000 (for ALE v2)</li> <li>• 61,000 (for LSE and LSE2)</li> </ul> <p><b>Note</b> For TOR to ALE mapping, see the reference table below.</p>	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	25	25
Maximum number of EPGs consuming the same contract	25	25
<b>FEX VPC</b>		
Maximum EPGs behind FEX VPC port	20	N/A
<b>FCoE</b>		
Maximum number of VSAN	32	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of VFC	151 <b>Note</b> This number includes VFCs configured on switch ports and FEX ports.	N/A
Maximum number of FDISC per port	96	N/A
Maximum number of FDISC per SB	96	N/A

ALE 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

## 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>		

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of PODs	N/A	10
Maximum number of nodes per POD	N/A	200
Maximum number of leaf switches overall	N/A	400

## 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
Datacenters in a vCenter	N/A	2
Total number of (VMM domain, VMM controller (vCenter/vShield)) instances	N/A	<ul style="list-style-type: none"> <li>• 200 vDS</li> <li>• 50 AVS</li> </ul>
Number of ESX hosts per AVS	240	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 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
Support RBAC for AVS	N/A	Yes
Support RBAC for vDS	N/A	Yes
<b>Microsegmentation/DFW with AVS</b>		

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
Number of ESX hosts per AVS	100	N/A
Number of Microsegment EPGs	1,000	N/A
Number of DFW flows per vEth	10,000	N/A
Number of DFW denied and permitted flows per ESX host	250,000	N/A
Number of VMM domains per Microsegment EPG	N/A	10
<b>Microsoft</b>		
Number of SCVMMs	N/A	5
VMM domains for Microsoft (in addition to that of VMware)	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
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

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
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	1,500
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)
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)

<b>Configurable Options</b>	<b>Per Leaf Scale</b>	<b>Per Fabric Scale</b>
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







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

**Asia Pacific Headquarters**  
Cisco Systems (USA) Pte. Ltd.  
Singapore

**Europe Headquarters**  
Cisco Systems International BV  
Amsterdam, The Netherlands

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).