



Verified Scalability Guide for Cisco APIC, Release 3.0(1k) and Cisco Nexus 9000 Series ACI-Mode Switches, Release 13.0(1k)

Overview 2

General Scalability Limits 2

Fabric Topology, SPAN, Tenants, Contexts (VRFs), External EPGs, Bridge Domains, Endpoints, and Contracts Scalability Limits 4

Multiple Fabric Options Scalability Limits 12

VMM Scalability Limits 12

Layer 4 - Layer 7 Scalability Limits 14

AD, TACACS, RBAC Scalability Limits 14

Cisco ACI Multi-Site Scalability Limits, Release 1.0(1x) 15

QoS Scalability Limits 15

Overview

This guide contains the maximum verified scalability limits for ACI parameters for the Cisco APIC Release 3.0.1k and Cisco Nexus 9000 Series ACI-Mode Switches, Release 13.0.1k. 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 and 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

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 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 (VRFs)	N/A	1000	3000
Number of contracts/filters	N/A	<ul style="list-style-type: none"> • 2000 contracts • 10,000 filters 	<ul style="list-style-type: none"> • 2000 contracts • 10,000 filters
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 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"> • 200 • 50 AVS 	<ul style="list-style-type: none"> • 200 • 50 AVS
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 (1200 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)
Number of VMs	N/A	Depends upon server scale	Depends upon server scale
Number of configuration zones per fabric	30	30	30

Feature	L2 Fabric	L3 Fabric	Large L3 Fabric
Number of BFD sessions	<ul style="list-style-type: none"> • 256 (physical interface, VRF) • 512 per Leaf 	<ul style="list-style-type: none"> • 256 (physical interface, VRF) • 512 per Leaf 	<ul style="list-style-type: none"> • 256 (physical interface, VRF) • 512 per Leaf
Multi-Pod NOTE: * = 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	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 (VRFs), External EPGs, Bridge Domains, Endpoints, and Contracts Scalability Limits

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)	1750	N/A
Number of encaps per FEX HIF, PC, vPC	20	N/A
Number of member links per PC, vPC	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

Configurable Options	Per Leaf Scale	Per Fabric Scale
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)	Default profile (Dual stack) For ALE v1 and v2: <ul style="list-style-type: none"> • MAC: 12,000 • IPv4: 12,000 or • IPv6: 6000 or • IPv4: 4000, IPv6: 4000 For LSE and LSE2: <ul style="list-style-type: none"> • MAC: 24,000 • IPv4: 24,000 • IPv6: 12,000 IPv4 Scale profile— For ALE v1 and v2: Not supported For LSE and LSE2: <ul style="list-style-type: none"> • MAC: 48,000 • IPv4: 48,000 • IPv6: Not supported 	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

Configurable Options	Per Leaf Scale	Per Fabric Scale
SPAN	<p>ALE based ToRs:</p> <ul style="list-style-type: none"> • 4 uni-directional or 2 bi-directional infra/tenant sessions • 4 uni-directional or 2 bi-directional fabric sessions <p>LSE based ToRs:</p> <ul style="list-style-type: none"> • 8 uni-directional or 4 bi-directional sessions (fabric, infra, or tenant) 	N/A
Number of ports per SPAN session	<ul style="list-style-type: none"> • All leaf access ports could be in one session • All leaf fabric ports could be in one session 	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"> • 230 ingress direction + 50 egress direction <p>LSE based TORs:</p> <ul style="list-style-type: none"> • 230 bi-directional • 460 uni-directional 	N/A
Common pervasive gateway	256 virtual IPs per Bridge Domain	N/A
Maximum number of Data Plane policers at interface level	<ul style="list-style-type: none"> • 64 ingress policers • 64 egress policers <p>For LSE and LSE2:</p> <ul style="list-style-type: none"> • 7 ingress policers • 7 egress policers 	N/A
Maximum number of Data Plane policers at EPG and interface level	128 ingress policers	N/A
Maximum number of SNMP trap receivers	10	10
Tenants		
Number of Contexts (VRFs) per tenant	50	50
Number of application profiles per tenant (or per Context (VRF))	N/A	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
Contexts (All numbers applicable to dual stack unless explicitly called out)		
Maximum number of Contexts (VRFs)	400	N/A
Maximum number of BDs per Context (VRF)	256	N/A
Maximum ECMP (equal cost multi-path) for BGP best path	16	N/A
Number of VRFs per tenant	N/A	50
Number of BDs per VRF	N/A	1750
Number of isolated EPGs	N/A	250
Border Leafs per L3 Out	N/A	8
Maximum number of vzAny Provided Contracts	16 per Context (VRF)	N/A
Maximum number of vzAny Consumed Contracts	16 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"> • 8 for Routed and Routed sub-interface • 1000 for SVI
Maximum number of Dynamic Routing protocol peers for BGP	400	3000
Maximum number of Dynamic Routing protocol peers for OSPF	300	N/A
Maximum number of Dynamic Routing protocol peers for EIGRP	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> <p>For ALE v1 and v2:</p> <ul style="list-style-type: none"> • IPv4: 10,000 or • IPv6: 6000 or • IPv4: 4000, IPv6: 4000 • IPv6 wide prefixes (> /64): 1000 <p>For LSE and LSE2:</p> <ul style="list-style-type: none"> • IPv4: 20,000 or • IPv6: 10,000 • IPv6 wide prefixes (>= /84): 1000 <p>IPv4 Scale profile</p> <p>For ALE v1 and v2: Not supported</p> <p>For LSE and LSE2:</p> <ul style="list-style-type: none"> • IPv4: 38,000 • IPv6: Not supported 	<ul style="list-style-type: none"> • IPv4; 40,000 or • IPv6; 20,000 or • IPv4; 10,000, IPv6; 10,000
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"> • 200 for SVI • 32 for subinterfaces 	<ul style="list-style-type: none"> • 400 for SVI • 32 for subinterfaces
Maximum number of ARP entries for L3 Outs	7500	N/A
Shared L3 Out	<ul style="list-style-type: none"> • IPv4: 2000 or • IPv6: 1000 	<ul style="list-style-type: none"> • IPv4: 6000 or • IPv6: 3000
Configurable Options	per Leaf scale	per Fabric scale
Maximum number of L3 Outs	400 (per leaf scale)	2400 (single stack)
External EPGs		
Number of External EPGs	600	2400 on ALE 4000 on LSE (single stack)
Number of External EPGs per L3 Out	250	400

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of LPM Prefixes for External EPG Classification	1000 IPv4	N/A
Bridge Domain		
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
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	8	N/A
Number of DHCP relay labels per BD	2	2
DHCP relay IP address will always be set to the primary SVI IP address	N/A	N/A
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
Endpoint Groups (Under App Profiles)		
Maximum amount of EPGs	Normally 1750; if legacy mode 3500	15,000

Configurable Options	Per Leaf Scale	Per Fabric Scale
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	<ul style="list-style-type: none"> • 200 vDS • 50 AVS 	N/A
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	4000	N/A
Maximum number of Taboo Contracts per EPG	2	N/A
IP-based EPG	4000	N/A
Contracts		
Security TCAM size	<ul style="list-style-type: none"> • 4000 (for ALE v1) • 40,000 (for ALE v2) • 61,000 (for LSE and LSE2) <p>Note For TOR to ALE mapping, see the reference table below.</p>	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
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

FEX VPC

Maximum EPGs behind FEX VPC port	20	N/A
----------------------------------	----	-----

FCoE

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
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"> • 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
LSE2	N9K-C93108TC-FX + N9K-C93180YC-FX

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 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
VMware		
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"> • 200 vDS • 50 AVS
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

Configurable Options	Per Leaf Scale	Per Fabric Scale
Support RBAC for AVS	N/A	Yes
Support RBAC for vDS	N/A	Yes
Number of Microsegment EPGs with AVS	1,000	N/A
Number of DFW flows per vEth with AVS	10,000	N/A
Number of DFW denied and permitted flows per ESX host with AVS	250,000	N/A
Number of VMM domains per EPG with AVS	N/A	10
Number of VM Attribute Tags per vCenter	N/A	vCenter version 6.0: 500 vCenter version 6.5: 1000
Microsoft		
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
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

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

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
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

Configurable Options	Per Leaf Scale	Per Fabric Scale
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)
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

Cisco ACI Multi-Site Scalability Limits, Release 1.0(1x)

Stretched vs. non Stretched— If you deploy separate fabrics as part of a Multi-Site architecture without stretching any object, in this release each fabric would be characterized by scalability values that are different (and lower) than the values that you would get in a “normal” fabric, that is, a standalone fabric not a part of Multi-Site.

Scaling Item	Stretched (Multi-Site)	Non-Stretched (APIC)
BDs	800	10,000
Contracts	1000	1000
End Points	50,000	100,000 including: <ul style="list-style-type: none"> • 50,000 - stretched from other sites • 50,000 - locally learned in site-local
EPGs	800	10,000
IGMP Snooping	8000	8000
L3Out external EPGs	500	2400
Leafs	50	50
Policy Objects per Schema	500	N/A
Sites	5	5
Subnets	2000	10,000
Templates per Schema	5	N/A
Tenants	100	2500
VRFs	400	3000

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
MPOD QoS Policy enabled	Custom QOS Policy with DSCP	9
	Custom QOS Policy with DSCP and Dot1P	9
	Custom QOS Policy with Dot1P	48
	Custom QOS Policy via a Contract	48

		QoS Scale
CoS Preservation enabled	Custom QoS Policy with DSCP	9
	Custom QoS Policy with DSCP and Dot1P	9
	Custom QoS Policy with Dot1P	48
	Custom QoS Policy via a Contract	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.