



Verified Scalability Guide for Cisco APIC, Release 2.0(1m) and Cisco Nexus 9000 Series ACI-Mode Switches, Release 12.0(1)

[Overview](#) 2

Overview

This guide contains the maximum verified scalability limits for ACI parameters for the Cisco APIC Release 2.0(1m) and Cisco Nexus 9000 Series ACI-Mode Switches, Release 12.0(1). These values are based on a profile where each feature was scaled to the numbers specified in the tables. These numbers do not represent the theoretically possible 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.

Feature	L2 Fabric	L3 Fabric	Large L3 Fabric
Number of APIC controllers	3	3	5
Number of leaf switches	80	80	200
Number of spines	6	6	6
Number of FEXs	N/A	12 per leaf, 120 per fabric	N/A
Number of tenants	N/A	1,000	3,000
Number of Layer 3 (L3) contexts	N/A	1,000	3,000
Number of contracts/filters	N/A	<ul style="list-style-type: none"> • 1,000 contracts • 10,000 filters 	<ul style="list-style-type: none"> • 1,000 contracts • 10,000 filters

Feature	L2 Fabric	L3 Fabric	Large L3 Fabric
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	150	150	150
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) Note This limit exists across all protocols/transit scenarios	N/A	<ul style="list-style-type: none"> • IPv4; 40,000 • IPv6; 20,000 • Per Leaf IPv4 and IPv6; 10,000 total 	<ul style="list-style-type: none"> • IPv4; 40,000 • IPv6; 20,000 • Per Leaf IPv4 and IPv6; 10,000 total
Number of BGP + number of OSPF sessions + EIGRP (for external connection)	N/A	1,200	1,200
Number of Multicast groups	N/A	8,000	8,000
Number of Multicast groups per VRF	N/A	8,000	8,000
Number of vCenters	N/A	<ul style="list-style-type: none"> • 50 vDS • 50 AVS 	<ul style="list-style-type: none"> • 50 vDS • 50 AVS
Number of Service Chains	N/A	1,000	1,000
Number of L4 - L7 devices	N/A	30 physical, 1,200 virtual (1,200 maximum per fabric)	30 physical, 1,200 virtual (1,500 maximum per fabric)
Number of ESX hosts - VDS	N/A	3,200	3,200
Number of ESX hosts - AVS	N/A	3,200 (Only 1 AVS instance per host)	3,200 (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"> • 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	N/A	6 pods, 80 nodes overall	4 pods, 200 nodes max per pod, 300 leaf switches overall
L3 eVPN Services over Fabric WAN	N/A	1,000 VRFs, 60,000 routes in a fabric	1,000 VRFs, 60,000 routes in a fabric

Feature	L2 Fabric	L3 Fabric	Large L3 Fabric
Layer 3 Multicast	N/A	8,000 multicast routes	8,000 multicast routes

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

Configurable Options	Per Leaf Scale	Per Fabric Scale
Fabric Topology		
Maximum number of vPCs	320 (hif vPC with FEX)	(Number of leafs /2) X 48
Maximum number of encaps per vPC	1,750 (ports X encap < 64,000) If the BD is in classic mode, 48 X 3,500 = 168,000 port-VLAN combination is supported in the L2 Fabric mode. Note With EPG deployed on FEX : port-VLAN combination is restricted to 10,000.	N/A
Maximum number of Member Links per vPC	8	N/A
Maximum number of PCs	48	(Number of leafs) X 48
Maximum number of encaps per PC	1,750 (ports X encap < 64,000)	N/A
Maximum number of Member Links per PC	8	N/A
Maximum number of PCs, access ports	48	(Number of leafs) X 48
Maximum number of encaps per access port	1,750 (ports X encap < 64,000)	N/A
Number of ports x VLANs (global scope) (No FEX, no local scope VLANS)	64,000	N/A
Number of ports x VLANs (local scope) (FEX and/or local scope VLANS)	10,000	N/A
STP	All VLANs	N/A
Maximum number of endpoints (EPs)	<ul style="list-style-type: none"> • IPv4; 12,000 or • IPv6; 6,000 or • IPv4; 4,000, IPv6; 4,000 	180,000
Number of Multicast Groups	8,000	8,000
Number of Multicast Groups per VRF	8,000	8,000

Configurable Options	Per Leaf Scale	Per Fabric Scale
Number of IPs per MAC	256	256
SPAN	<ul style="list-style-type: none"> • 4 infra/tenant sessions • 4 fabric sessions (VXLAN is carried in the SPAN) 	8 fabric sessions per fabric
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 EPG/BDs per SPAN session	280	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"> • 64 ingress policers • 64 egress policers 	N/A
Maximum number of SNMP trap receivers	10	10
Tenants		
Number of Contexts per tenant	8	8
Number of application profiles per tenant (or per Context)	N/A	N/A
Contexts (All numbers applicable to dual stack unless explicitly called out)		
Maximum number of Context	200	N/A
Maximum number of BDs per Context	256	N/A
Border Leafs per L3 Out	N/A	4
Maximum number of LPM Prefixes for External EPG Classification	1,000 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
Maximum number of L3 Outs	400	400
L3 Out per context	--	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 • 200 for SVI 	<ul style="list-style-type: none"> • 8 for Routed and Routed sub-interface • 400 for SVI

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of Dynamic Routing protocol peers for BGP	300	N/A
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"> • IPv4; 10,000 or • IPv6; 6,000 or • IPv4; 4,000, IPv6; 4,000 	<ul style="list-style-type: none"> • IPv4; 40,000 or • IPv6; 20,000 or • IPv4; 10,000, IPv6; 10,000
Maximum number of External Routes	<ul style="list-style-type: none"> • IPv4; 10,000 or • IPv6; 6,000 or • IPv4; 4,000, IPv6; 4,000 	<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 L3 out	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: 2,000 or • IPv6: 1,000 	<ul style="list-style-type: none"> • IPv4: 6,000 or • IPv6: 3,000
External EPGs		
Number of External EPGs per L3 out	16	400
Number of L3 out	N/A	400
Bridge Domain		
Maximum amount of BDs	1,750; if legacy mode, 3,500; if Multicast optimized mode then 50	15,000
Maximum number of subnets per BD	16 (cannot be for all BDs)	16 per BD

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of EPGs per BD	3,499 (cannot exceed 3,500 total) 3,499 is supported in hardware but please refer to the per fabric scale for the effective software support for this release.	3,499
Number of L2 Outs per BD	1	1
Number of BDs with Custom MAC Address	1,750; if legacy mode, 3500; If Multicast optimized mode is used, then 50	1,750; if legacy mode, 3,500; If Multicast optimized mode is used, then 50
Number of Multicast groups	8,000	8,000
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	1,000	1,000
Maximum number of PIM Neighbors per VRF	64	64
Endpoint Groups (Under App Profiles)		
Maximum amount of EPGs	Normally 1,750; 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

Configurable Options	Per Leaf Scale	Per Fabric Scale
Maximum number of VMM domains	<ul style="list-style-type: none"> • 50 vDS • 5 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 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	50	N/A
Maximum number of Taboo Contracts per EPG	2	N/A
IP-based EPG	4,000	N/A
Contracts		
Security TCAM size	<ul style="list-style-type: none"> • 4,000 (for ALE v1) • 32,000 (for ALE v2) <p>Note 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	10	10
Maximum number of EPGs consuming the same contract	10	10
FEX VPC		
Maximum EPGs behind FEX VPC port	20	N/A

Configurable Options	Per Leaf Scale	Per Fabric Scale
FCoE		
Maximum number of VSAN	16	N/A
Maximum number of VFC	16	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

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	4
Maximum number of nodes per POD	N/A	200
Maximum number of leaves overall	N/A	300

VMM Scalability Limits

Configurable Options	Per Leaf Scale	Per Fabric Scale
VMware		
Number of vCenters (vDS)	N/A	50 (Verified with a load of 10 events/minute for each vCenter)
Number of vCenters (AVS)	N/A	10
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"> • 50 vDS • 10 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
Support RBAC for AVS	N/A	Yes
Support RBAC for vDS	N/A	Yes
Microsegmentation/DFW with AVS		
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
Microsoft		
Number of controllers per SCVMM domain	N/A	5
Number of SCVMM domains	N/A	5

Configurable Options	Per Leaf Scale	Per Fabric Scale
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
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	1,500
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	100
Deployment scenario for ASA (transparent or routed)	N/A	Yes

Configurable Options (L4-L7 Configurations)	Per Leaf Scale	Per Fabric Scale
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

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
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

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any Internet Protocol (IP) addresses and phone numbers used in this document are not intended to be actual addresses and phone numbers. Any examples, command display output, network topology diagrams, and other figures included in the document are shown for illustrative purposes only. Any use of actual IP addresses or phone numbers in illustrative content is unintentional and coincidental.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com).

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<https://www.openssl.org>)

This product includes software written by Tim Hudson (tjh@cryptsoft.com).

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)

© 2016 Cisco Systems, Inc. All rights reserved.



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.