

Cisco Nexus Dashboard Verified Scalability Guide, Release 4.1.1

New and Changed Information 2

About supported node types and features 2

Scale for LAN Deployments 4

Scale for SAN Deployments 19

Trademarks ?

Revised: August 7, 2026

New and Changed Information

The table provides an overview of the significant changes to this guide for this current release. The table does not provide an exhaustive list of all changes made to the guide or of the new features in this release.

Feature	Description	Where Documented
Unification of scale limits in one guide	Beginning with Nexus Dashboard release 4.1.1, verified scale limits are combined in one guide.	*Cisco Nexus Dashboard Verified Scalability Guidelines, Release 4.1.1* (this document)
Support for creating a Data Center VXLAN and an IPFM fabric on the same Nexus Dashboard cluster	With this release, Nexus Dashboard added support for creating a Data Center VXLAN EVPN and an IPFM fabric on the same cluster. The fabrics do not share data.	• For configuration information, see Creating LAN and ACI Fabrics and Fabric Groups • For scale information, see *Cisco Nexus Dashboard Verified Scalability Guidelines, Release 4.1.1* (this document)

Overview

This document provides the maximum verified platform scalability limits for Nexus Dashboard.

The values listed below are explicitly tested and validated but may not represent the maximum theoretical system limits for Nexus Dashboard. For more information, see the specific tables for the functionality that you are using for specific scalability information.



Note Unless explicitly called out, the listed scale numbers apply to all form factors.

This release supports the following cluster form factors:

- Physical appliance (`.iso`) – This form factor refers to the Cisco UCS physical appliance hardware with the Nexus Dashboard software stack pre-installed on it.
- Virtual Appliance – The virtual form factor allows you to deploy a Nexus Dashboard cluster using VMware ESX (`.ova`) or RHEL KVM (`.qcow2`).

  The virtual form factor supports the following two profiles:

  - Data node – This profile with higher system requirements is designed for higher scale and/or unified deployment.
  - App node – This profile with lower system requirements can be deployed as secondary nodes. Can also be deployed as primary nodes but does not support unified deployment.

  In addition, beginning with Nexus Dashboard release 4.1(1), support is available for running a virtual Nexus Dashboard (vND) on the AWS public cloud. See "Deploying vNDs in Amazon Web Services" in the *Cisco Nexus Dashboard Deployment and Upgrade Guide* for more information.

About supported node types and features

These node types have been available for releases prior to Nexus Dashboard release 4.1.1.

- SE-NODE-G2 (UCS-C220-M5). The product ID of the 3-node cluster is SE-CL-L3.
- ND-NODE-L4 (UCS-C225-M6). The product ID of the 3-node cluster is ND-CLUSTER-L4.

Beginning with Nexus Dashboard release 4.1.1, this node type is now also available.

- ND-NODE-G5S (UCS-C225-M8). The product ID of the 3-node cluster is ND-CLUSTERG5S.

In addition, in LAN deployments, these are the available features that you can leverage.

- **Controller**: Also referred to as Fabric Management. This feature is used to manage NX-OS and non-NX-OS switches (such as Catalyst, ASR, and so on). This includes creating any non-ACI fabric types, as well as performing software upgrades and creating new configurations on those fabrics.
- **Telemetry**: This feature provides telemetry functionality, similar to the functionality provided by Nexus Dashboard Insights in releases prior to Nexus Dashboard release 4.1.1. You can enable and use the **Telemetry** feature when you create or edit a fabric through **Manage** > **Fabrics**.
- **Orchestration**: You can use the **Orchestration** feature through Nexus Dashboard to connect multiple ACI fabrics together, and consolidate and deploy tenants, along with network and policy configurations, across multiple ACI fabrics. You can enable and use the **Orchestration** feature when you add an ACI through **Admin** > **System Settings** > **Multi-cluster connectivity** > **Connect Cluster**.

You can enable these features independently or, in some cases, as one of these combined feature sets.

- Controller and Telemetry
- Orchestration and Telemetry
- Controller, Telemetry, and Orchestration (not supported on an App node cluster or in a cluster with SE-NODE-G2 nodes)

Guidelines and limitations

- For Nexus Dashboard release 4.1.1, you cannot mix the newer ND-NODE-G5S (UCS-C225-M8) nodes in a cluster with the older SE-NODE-G2 (UCS-C220-M5) and ND-NODE-L4 (UCS-C225-M6) nodes.
- A 6-node physical appliance cluster is primarily designed for extended scale NX-OS or ACI fabrics with the Telemetry feature enabled and is not recommended for non-Telemetry deployments.
- The virtual form factor does not support all features in many cluster sizes and types, as described in this document.

Scale for LAN Deployments

Form factors, cluster sizes, and general scale

Form factors and cluster sizes for LAN deployments

Note The following tables contains supported switch scale based on the cluster form factor and size. For complete feature-specific scale, see one of the following tables in this document.

For supported form factors and cluster sizes for IPFM and SAN fabrics, see the later sections of this document.

If you are mixing NX-OS and ACI fabrics in the same cluster, then the lesser of the two listed scale numbers applies.

These updates are now available beginning with Nexus Dashboard release 4.1.1, as described in the following tables.

- `ND-NODE-G5S` (UCS-C225-M8) is now available as a new physical node type
- 1-node and 3-node virtual clusters (data) are now available as new virtual cluster types

In addition:

- The `5-node virtual cluster (app)` is not supported as a greenfield deployment, and is only supported when upgrading existing clusters from Nexus Dashboard release 3.2.x to Nexus Dashboard release 4.1.1. The cluster must be moved to a supported form factor by taking a back up and restore. A `3-node virtual cluster (data)` (Controller only) would be the equivalent for the `5-node virtual cluster (app)` (Controller only).
- For the virtual appliance, the App profile with 1.5TB disk is not supported as a greenfield deployment, but an upgrade from Nexus Dashboard release 3.2.x to 4.1.1 is supported with this configuration.

Table 1: Form factors and cluster sizes for LAN deployments: 1-node physical clusters

Node type	Enabled Features	Supported NX-OS scale	Supported ACI Scale
SE-NODE-G2 and ND-NODE-L4	Controller only	50 switches (managed) 100 switches (monitored)	N/A
	Controller and Telemetry	2 fabrics 25 switches Flow Telemetry: 500 flows/second Traffic Analytics: 10,000 conversations/minute, 1 concurrent troubleshoot job For full Telemetry scale, see Telemetry scale limits, on page 13.	

Node type	Enabled Features	Supported NX-OS scale	Supported ACI Scale
ND-NODE-G5S	Controller and Telemetry	2 fabrics 50 switches Flow Telemetry: 1,000 flows/second Traffic Analytics: 10,000 conversations/minute, 1 concurrent troubleshoot job For full Telemetry scale, see Telemetry scale limits, on page 13.	
	Controller, Telemetry, and Orchestration **Note** For non-production use only		

Table 2: Form factors and cluster sizes for LAN deployments: 3-node physical clusters

Node type	Enabled Features	Supported NX-OS scale	Supported ACI Scale	Notes
SE-NODE-G2 or ND-NODE-L4	Controller only	500 switches (managed) 1000 switches (monitored)	N/A	
ND-NODE-L4	Telemetry only	12 fabrics 400 switches Flow Telemetry: 10,000 flows/second Traffic Analytics: 100,000 conversations/minute, 5 concurrent troubleshoot jobs For full Telemetry scale, see Telemetry scale limits, on page 13.		Supports mixed ACI and NX-OS deployment with 50/50 split
SE-NODE-G2	Controller and Telemetry or Telemetry only	8 fabrics 250 switches Flow Telemetry: 10,000 flows/second Traffic Analytics: 100,000 conversation/minute, 5 concurrent troubleshoot jobs For full Telemetry scale, see Telemetry scale limits, on page 13.		Supports mixed ACI and NX-OS deployment with 50/50 split

Node type	Enabled Features	Supported NX-OS scale	Supported ACI Scale	Notes
SE-NODE-G2 / ND-NODE-L4 [1]	Telemetry and Orchestration	N/A	• **SE-NODE-G2**: 12 fabrics 100 switches Flow Telemetry: 10,000 flows/second Traffic Analytics: 100,000 conversations/minute, 5 concurrent troubleshoot jobs Split supported is ACI (75) and NX-OS (25) for IPN/ISN use case • **ND-NODE-L4**: 12 fabrics 400 switches Flow Telemetry: 10,000 flows/second Traffic Analytics: 100,000 conversations/minute, 5 concurrent troubleshoot jobs Split supported is ACI (375) and NX-OS (25) for IPN/ISN use case For full Telemetry scale, see Telemetry scale limits, on page 13.	NX-OS switches supported only in **External and inter-fabric connectivity** fabric types
ND-NODE-L4	Controller, Telemetry, and Orchestration	20 fabrics 250 switches Flow Telemetry: 10,000 flows/second Traffic Analytics: 100,000 conversations/minute, 5 concurrent troubleshoot jobs For full Telemetry scale, see Telemetry scale limits, on page 13. For full Orchestration scale, see Orchestration scale limits, on page 14.		

Node type	Enabled Features	Supported NX-OS scale	Supported ACI Scale	Notes
ND-NODE-G5S	Controller, Telemetry, and Orchestration	500 switches Flow Telemetry: 10,000 flows/second Traffic Analytics: 100,000 conversations/minute, 5 concurrent troubleshoot jobs For full Telemetry scale, see Telemetry scale limits, on page 13. For full Orchestration scale, see Orchestration scale limits, on page 14.		

[1] When mixing SE-NODE-G2 and ND-NODE-L4 nodes in the same cluster, the lower scale value applies.

Table 3: Form factors and cluster sizes for LAN deployments: 6-node physical clusters

Node type	Enabled Features	Supported NX-OS scale	Supported ACI Scale	Notes
SE-NODE-G2 and ND-NODE-L4	Telemetry only	• **SE-NODE-G2**: 32 fabrics 500 switches • **ND-NODE-L4**: 32 fabrics 750 switches Flow Telemetry: 20,000 flows/second Traffic Analytics: 200,000 conversations/minute, 8 concurrent troubleshoot jobs For full Telemetry scale, see Telemetry scale limits, on page 13.	• **SE-NODE-G2**: 20 fabrics 750 switches • **ND-NODE-L4**: 20 fabrics 1000 switches Flow Telemetry: 20,000 flows/second Traffic Analytics: 200,000 conversations/minute, 8 concurrent troubleshoot jobs For full Telemetry scale, see Telemetry scale limits, on page 13.	For a mix of ACI and NX-OS deployment with a 50/50 split, the scale will be 750 for a ND-NODE-L4-only cluster and 500 for a SE-NODE-G2-only cluster or a SE-NODE-G2/ND-NODE-L4 mixed cluster
SE-NODE-G2 and ND-NODE-L4	Telemetry and Orchestration [2]	N/A	20 fabrics 500 switches Flow Telemetry: 20,000 flows/second Traffic Analytics: 200,000 conversation/minute, 8 concurrent troubleshoot job For full Telemetry scale, see Telemetry scale limits, on page 13.	Split supported is ACI (450) and NX-OS (50 for IPN/ISN)

[2] To enable Telemetry and Orchestration on a 6-node physical cluster, when the fabric is added through **Onboard existing LAN fabric** under **Select a category**, choose **External and Inter-fabric connectivity** in the **Select a type** step during the Create/Onboard Fabric workflow.

Table 4: Form factors and cluster sizes for LAN deployments: Virtual clusters

Node type	Enabled Features	Supported NX-OS scale	Supported ACI Scale
Virtual Clusters: 1-node virtual			
App ESXi/KVM	Controller only	50 switches (Managed) 100 switches (Monitored)	N/A
Data ESXi/KVM	Controller and Telemetry	2 fabrics 25 switches Flow Telemetry: 500 flows/second Traffic Analytics: 3000 conversation/minute, 1 concurrent troubleshoot job For full Telemetry scale, see Telemetry scale limits, on page 13.	
	Controller, Telemetry, and Orchestration **Note** For non-production use only		
Virtual Clusters: 3-node virtual			
App ESXi/KVM	Controller only	Managed: 100 switches Monitored: 200 switches	N/A
App ESXi only	Orchestration only	N/A	Same verified scalability as provided in *Nexus Dashboard Orchestrator, release 4.2(3)*
Data ESXi/KVM	Controller only	Managed: 400 switches Monitored: 1000 switches	N/A
Data ESXi/KVM	Controller, Telemetry, and Orchestration	100 switches Flow Telemetry: 2500 flows/second Traffic Analytics: 5000 conversation/minute, 2 concurrent troubleshoot jobs 1 max online assurance job across all fabrics For full Telemetry scale, see Telemetry scale limits, on page 13.	
Virtual Clusters: 6-node virtual			

Node type	Enabled Features	Supported NX-OS scale	Supported ACI Scale
3 Data and 3 App nodes ESXi **Note** Data nodes must be `primary`.	Controller, Telemetry, and Orchestration	200 switches Flow Telemetry: 5000 flows/second Traffic Analytics: 10,000 conversation/minute, 5 concurrent troubleshoot job For full Telemetry scale, see Telemetry scale limits, on page 13.	

General scale limits

Table 5: General scale limits

Category	Verified Scale Limit
Standby nodes in a cluster	Up to 2 standby nodes. • Physical node clusters: Standby nodes are supported for most 3-node or larger clusters • Virtual node clusters: Standby nodes are supported only with a 3-node vND (app) profile for a Controller-only or Orchestration-only deployment.
Users configured on the cluster	1000
API requests rate	1000 requests in 6 seconds
Login domains	8
Clusters connected using multi-cluster connectivity	10
Fabrics across all clusters connected using multi-cluster connectivity	100
Switches across all clusters connected using multi-cluster connectivity	3500
Maximum latency between nodes for intra-cluster connectivity	50 ms
Maximum latency between cluster nodes and switches	150 ms
Maximum latency between nodes for inter-cluster connectivity	500 ms

Controller scale limits

This section provides verified scalability values for LAN fabrics.

Table 6: Fabric underlay and overlay scale limits

Category	Verified Scale Limit
Fabrics in a VXLAN EVPN fabric group	50
Switches per fabric	200
Physical interfaces **Note** Physical interfaces are for brownfield deployments.	11,500
Overlay scale for VRFs and networks	1000 VRFs, 3000 Layer 2 or Layer 3 networks Supported scale for 1-node virtual Nexus Dashboard is 250 VRFS and 1000 networks.
Overlay associations	This category defines the number of networks per switch and attached to the interfaces. The supported total number of networks x switches x interfaces associations are 5 million per cluster, where: • You have 3000 or fewer networks, and • You have 200 or fewer switches For example: • **Supported**: 3000 networks x 100 switches x 10 interfaces = 3 million (supported because this total is less than 5 million) • **Not supported**: 3000 networks x 200 switches x 10 interfaces = 6 million (unsupported because this total exceeds 5 million)
VRF instances for external connectivity	1000
IPAM integrator application	150 networks with a total of 4k IP allocations on the Infoblox server
ToR devices **Note** There is no support for ToR devices in a brownfield deployment.	40 leaf switches with 320 ToR switches A Data Center VXLAN EVPN fabric can manage both Layer 2 ToR switches and leaf switches. 32 ToR switches (or 16 vPC-ToR pairs) can be connected per leaf-vPC pair.

Table 7: Endpoint locator scale limits

Description	Verified Limit
Endpoints	100,000 **Note** For a single-node virtual cluster, the scale is reduced to 1 instance of an endpoint locator with 10,000 endpoints.
Dual-home and dual-stacked (IPv4 + IPv6) endpoints	60,000

Description	Verified Limit
BGP route reflectors or route servers	2 per fabric
Fabrics in VXLAN EVPN fabric group	30

Table 8: Virtual Machine Manager (VMM) scale limits

Description	Verified Limit
VMware Center Servers	4 **Note** For 1-node virtual Nexus Dashboard clusters, the scale is reduced to 1 VMware Center Server.
vCenter endpoints (VMs)	1-node physical cluster: 100 3-node physical cluster: 4000 6-node physical cluster: 4000 1-node virtual cluster: 100 3-node virtual cluster: 1000 6-node virtual cluster: 1000
Kubernetes clusters	4
Kubernetes Visualizer application	160 namespaces with up to 1002 pods

Table 9: Security groups scale limits

Description	Verified Limit
VXLAN security groups and selectors	4000 selectors (if 1 per security group then 4000 security groups) 12,000 bidirectional security associations
Adjust the number of security associations if the number of class maps per policy map is more than 1, so that the maximum total number of security associations is always 12,000 bidirectional when deployed on the devices.	1 class map + 1 policy map per association
Associations in a startup configuration **Note** Nexus Dashboard does not support a bootflash repartition.	12,000

IPFM fabric scale limits

This section provides verified scalability values for IPFM fabrics.

Table 10: Scale imits for IPFM fabrics based on deployment type

Deployment Type	Verified Scale Limit
1-node physical Nexus Dashboard	35 switches (2 spine switches and 33 leaf switches)
3-node physical Nexus Dashboard	120 switches (2 spine switches, 100 leaf switches, and 18 tier-2 leaf switches) - IPFM Fabric 60 switches - LAN Fabric 20 switches - VxLAN Fabric
1-node virtual Nexus Dashboard (app node)	35 switches (2 spine switches and 33 leaf switches)
3-node virtual Nexus Dashboard (app node)	120 switches (2 spine switches, 100 leaf switches, and 18 tier-2 leaf switches) - IPFM Fabric 60 switches - LAN Fabric 20 switches - VxLAN Fabric

Table 11: Scale limits for IPFM fabrics based on the type of mode

Category	Verified Scale Limit			
	NBM Active Mode Only	NBM Passive Mode Only	Mixed Mode	
			NBM Active VRF	NBM Passive VRF
Number of switches	120	32	32	32
Number of flows	32,000	32,000	32,000	32,000
Number of endpoints (discovered hosts)	5000	1500	3500	1500
VRFs	16	16	16	16
Host Policy - Sender	8000	NA	8000	NA
Host Policy - Receiver	8000	NA	8000	NA
Host Policy - PIM (Remote)	512	NA	512	NA
Flow Policy	2500	NA	2500	NA
NBM ASM group range	20	NA	20	NA
Host Alias	2500	NA	2500	NA
Flow Alias	2500	NA	2500	NA
NAT Flows	3000	3000	3000	3000
RTP Flow Monitoring	8000	8000	8000	8000

Category	Verified Scale Limit			
	NBM Active Mode Only	NBM Passive Mode Only	Mixed Mode	
			NBM Active VRF	NBM Passive VRF
PTP Monitoring	120 switches	32 switches	32 switches	32 switches

Telemetry scale limits

This section provides verified scalability values for Telemetry features of Nexus Dashboard.

Table 12: Telemetry scale limits

Description	Verified Limit
Endpoints	1-node physical cluster: 20,000 3-node physical cluster: 120,000 6-node physical cluster: 240,000 1-node virtual cluster: 20,000 3-node virtual clusters: 60,000 6-node virtual clusters: 60,000
Flow telemetry rules	500 rules per switch for both NX-OS and ACI fabrics.
Exporters for Kafka	6 exporters total for **Alerts and Events** across both NX-OS and ACI fabrics. 20 exporters for **Alerts and Events** with only Anomalies enabled without the Statistics and Advisories enabled. 6 exporters for **Usage** for ACI fabrics only. 6 email exporters. 6 syslog exporters.
Export data	5 emails per day for periodic job configurations.
Syslog	5 syslog exporter configurations across fabrics.
AppDynamics	5 apps 50 tiers 250 nodes 300 net links 1000 flow groups
DNS integration	40,000 DNS entries for physical clusters. 10,000 DNS entries for virtual clusters.

Description	Verified Limit
Panduit power distribution unit (PDU) integration	1000 per Nexus Dashboard cluster 500 per fabric

Orchestration scale limits

This section provides verified scalability values for Orchestration.

Table 13: General scale limits

Category	Verified Scale Limit
Fabrics	Up to 100 fabrics total on-boarded in Nexus Dashboard. Up to 14 of those fabrics can be enabled with EVPN sessions between them. For specific details about template object scale, which depends on the type of the templates you deploy (Multi-Fabric vs Autonomous), see the tables below.
Pods per fabric	12 or 25, depending on the Cisco APIC release managing the site. For more information, see the *Cisco APIC Verified Scalability Guide* for your release.
Leaf switches per fabric	400 in a single pod 500 across all pods in multi-pod ACI fabrics The number of leaf switches supported within each fabric depends on the Cisco APIC release managing that fabric. For more information, see the *Cisco APIC Verified Scalability Guide* for your release.

Category	Verified Scale Limit
Total leaf switches across all fabrics	`(max number of fabrics) * (max number of leaf switches per fabric)`, for example: • For multi-fabric deployments, if every fabric is deployed as a multi-pod ACI fabric, then the maximum number of leaf switches is `(14 fabrics) * (500 switches) = 7000`. • For Autonomous templates, if Orchestrator is deployed in a physical Nexus Dashboard cluster, then the maximum number of leaf switches is `(100 fabrics) * (500 switches) = 50,000` • For Autonomous templates, if Orchestrator is deployed in a virtual Nexus Dashboard cluster, then the maximum number of leaf switches is `(20 fabrics) * (500 switches) = 10,000` Note that specific objects' scale (such as VRFs, BDs, EPGs, and so on) still applies, as described in the template-specific sections below.
Endpoints per fabric	The Orchestrator endpoint scale for each fabric is the same as the scale supported by the fabric's APIC. **Note** If the fabric is part of a VXLAN fabric group, the total number of endpoints is the sum of local and remote endpoints.

Templates Scale Limits



Note If a specific object's scale (such as contracts, filters, or VRFs) is not included in the following table, that object does not have a unique scale limit and the general "Policy Objects per Schema" and "Policy Objects per Template" limits apply. If any such objects were explicitly listed in previous releases, those limitations have been lifted and removed from the list.

Table 14: Application templates scale limits

Category	Verified Scale Limit
Schemas	1000
Templates per schema	30
Service graphs per schema	500
Service graph nodes per service graph	5 for Autonomous templates 2 for multi-fabric templates
Policy objects per schema	2000

Category	Verified Scale Limit
Policy objects per template	2000
Contract preferred group (BD/EPG combinations)	5000
PBR destinations per fabric (including all local and remote*) *Note that if you configure some of the new PBR use cases, such as vzAny with PBR or L3Out-to-L3Out with PBR, you may be required to implement hair-pinning of traffic across fabrics to ensure traffic can be always steered using both devices and deployed in the source and destination fabrics. As a result, the leaf nodes in a given fabric must be programmed with PBR information about the device(s) in remote fabrics as well, and those remote PBR nodes are counted toward the maximum number listed here.	1500

Table 15: Fabric resource policies scale

Category	Verified Scale Limit
Policy objects per template	2000

Table 16: Tenant policies, fabric policies, and monitoring policies templates scale

Category	Verified Scale Limit
Policy objects per template	500
Monitoring Policy Scale	
ERSPAN sessions	20 per fabric
Fabric SPAN sessions	30 per fabric

Table 17: L3Out templates scale

Category	Verified Scale Limit
IP L3Outs per template	100
SR-MPLS L3Outs per template	100
All other objects' scale	The scale for other L3Out template objects that are not explicitly listed in this table is the same as the scale supported by the fabric's APIC. For detailed information, see the Cisco APIC Verified Scalability Guide for the APIC release version managing each fabric.

Orchestrator-deployed objects scale

To better understand the scalability values captured in the following table for traditional multi-fabric deployments, it is important to clarify that there are three kinds of Orchestrator-deployed objects:

- Fabric local objects—these are the objects defined in templates associated to a single fabric, which get deployed by Orchestrator only in that specific fabric.
- Shadow objects—these are the objects deployed by Orchestrator in a fabric as a result of a contract established between fabric local and remote objects, they are the representation ("shadow)" of the remote object in the local fabric.
- Stretched objects—these are the objects defined in templates that are associated with multiple fabrics, which get deployed by Orchestrator concurrently on all those fabrics.

The table below captures the maximum number of objects that Orchestrator can deploy in a given fabric and includes the sum of all three kinds of objects described above.

For example, if you have two fabrics and you define three templates on Orchestrator—`template-1` associated to `fabric-1`, `template-2` associated to `fabric-2`, and `template-stretched` associated to both `fabric-1` and `fabric-2`—then:

- If you configure and deploy `EPG-1` in `template-1`, this will count as one EPG towards maximum allowed for `fabric-1`.
- If you configure and deploy `EPG-2` in `template-2`, this will count as one EPG towards maximum allowed for `fabric-2`.
- If you apply a contract between `EPG-1` and `EPG-2` or add both EPGs to the Preferred Group), a shadow `EPG-2` will be created in `fabric-1` and a shadow `EPG-1` in `fabric-2`. As a result, two EPGs will now be counted towards the maximum allowed in each fabric.
- Finally, if you configure and deploy `EPG-3` in `template-stretched`, it will count as another EPG in each fabric, bringing the total to 3 EPGs towards the maximum allowed scale.

It is worth adding that the maximum number of objects supported in a given fabric (and captured in the Verified Scalability Guide for Cisco APIC) must not exceed the sum of objects locally defined on APIC plus the objects pushed from Orchestrator to that fabric (Orchestrator-deployed objects).



Note For maximum scale Nexus Dashboard Orchestrator configurations with many features enabled simultaneously, we recommend that those configurations be tested in a lab before deployment.

Table 18: Orchestrator-deployed logical objects scale for multi-fabric templates

Category	Maximum number of objects per fabric for up to 4 fabrics	Maximum number of objects per fabric for 5-14 fabrics
Tenants	1000	400
VRFs	2000	1000
BDs	6000	4000
Contracts	6000	4000
EPGs	6000	4000
ESGs	5000	4000

Category	Maximum number of objects per fabric for up to 4 fabrics	Maximum number of objects per fabric for 5-14 fabrics
Isolated EPGs	500	500
Microsegment EPGs	2000	500
L3Out external EPGs	500	500
Subnets	8000	8000
L4-L7 logical devices	400	400
Graph instances	250	250
Device clusters per tenant	10	10
Number of graph instances per device cluster	125	125

Nexus Dashboard Orchestrator provides support for autonomous sites. When creating application templates, you can now choose to designate the template as Autonomous. This allows you to associate the template to one or more fabrics that are operated independently and are not connected through an Inter-Site Network (no intersite VXLAN communication).

Because autonomous sites are by definition isolated and do not have any intersite connectivity, there is no shadow object configuration across fabrics and no cross-programming of pctags or VNIDs in the spine switches for intersite traffic flow.

The autonomous templates allow for significantly higher deployment scale as shown in the following table. Since there are no stretched objects or shadow objects, the scale values shown in the table below reflect the specific fabric-local objects that Orchestrator deploys in each fabric.

Table 19: Orchestrator-deployed objects scale for autonomous templates

Category	Verified Scale Limit (per fabric)
Tenants	1000
VRFs	2000
BDs	6000
Contracts	6000
EPGs	6000
ESGs	5000
Isolated EPGs	500
Microsegment EPGs	2000
L3Out external EPGs	500
Subnets	8000
Number of L4-L7 logical devices	400

Category	Verified Scale Limit (per fabric)
Number of graph instances	250
Number of device clusters per tenant	10
Number of graph instances per device cluster	125

VRF/BD VNID translation scale

Table 20: VRF/BD VNID translation scale

Category	Verified Scale Limit
Fixed spines	21,000
Modular spines	42,000

Scale for SAN Deployments

SAN scale limits

This section provides verified scalability values for SAN deployments.

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 scale.

For SAN deployments, you can set the maximum number of SAN switches and fabrics that can be managed by Nexus Dashboard by navigating to **Admin** > **System Settings** > **Fabric management** > **Discovery** and entering the necessary value in the **Maximum number of fabrics managed** field. The maximum number of SAN switches and fabrics that can be managed by Nexus Dashboard depends on the number of nodes in your cluster:

- For 1-node clusters, **24** is the default maximum number of SAN fabrics that can be managed by Nexus Dashboard. This value can be increased up to **160**. However, the number of switches across all fabrics are within the supported verified scale of **80** switches.
- For 3-node clusters, **24** is the default maximum number of SAN fabrics that can be managed by Nexus Dashboard. This value can be increased up to **160**. However, the number of switches across all fabrics are within the supported verified scale of **160** switches.

Scale limits for SAN deployments

Table 21: Scale limits for SAN deployments

Deployment Type	Verified Limit	
	Without SAN Insights	With SAN Insights
1-node virtual Nexus Dashboard (App node)[1]	80 switches, 20k ports	40 switches, 10k ports, and 40k ITs

Deployment Type	Verified Limit	
	Without SAN Insights	**With SAN Insights**
1-node virtual Nexus Dashboard (Data node)	80 switches, 20k ports	80 switches, 20k ports, and 1M ITLs/ITNs [2]
1-node physical Nexus Dashboard (SE)	80 switches, 20k ports	80 switches, 20k ports, and 120k ITLs/ITNs
3-node virtual Nexus Dashboard (App node)	160 switches, 40k ports	80 switches, 20k ports, and 100k ITs
3-node virtual Nexus Dashboard (Data node)	160 switches, 40k ports	160 switches, 40k ports, and 240k ITLs/ITNs
3-node physical Nexus Dashboard	160 switches, 40k ports	160 switches, 40k ports, and 500k ITLs/ITNs

[1] Application nodes have fewer features than data nodes. For example, the `lun` and `fc-scsi.scsi_initiator_itl_flow` features are not supported in the app ova, whereas those features are supported in the data ova. Therefore, you would have to install the data ova in order to use the `lun` or `fc-scsi.scsi_initiator_itl_flow` features.

[2] 1 million flows is the maximum number supported. If other features are enabled that consume resources, 1 million flows will not be stable in all situations. Nexus Dashboard consumes more resources per flow when processing telemetry from a larger number of devices. Watch flow counts and node memory usage (1 minute averages above ~105GB starts to show instability).

Note ITLs - Initiator-Target-LUNs

ITNs - Initiator-Target-Namespace ID

ITs - Initiator-Targets

Table 22: Scale limits for SAN zones

Description	Verified Limits
Zone sets	1000
Zone	16,000

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS REFERENCED IN THIS DOCUMENTATION ARE SUBJECT TO CHANGE WITHOUT NOTICE. EXCEPT AS MAY OTHERWISE BE AGREED BY CISCO IN WRITING, ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS DOCUMENTATION ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED.

The Cisco End User License Agreement and any supplemental license terms govern your use of any Cisco software, including this product documentation, and are located at: https://www.cisco.com/c/en/us/about/legal/cloud-and-software/software-terms.html. Cisco product warranty information is available at https://www.cisco.com/c/en/us/products/warranty-listing.html. US Federal Communications Commission Notices are found here https://www.cisco.com/c/en/us/products/us-fcc-notice.html.

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 products and features described herein as in development or available at a future date remain in varying stages of development and will be offered on a when-and if-available basis. Any such product or feature roadmaps are subject to change at the sole discretion of Cisco and Cisco will have no liability for delay in the delivery or failure to deliver any products or feature roadmap items that may be set forth in this document.

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.

The documentation set for this product strives to use bias-free language. For the purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on RFP documentation, or language that is used by a referenced third-party product.

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/c/en/us/about/legal/trademarks.html. 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. (1721R)

© 2025 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.