

Verified Scalability Guide for Cisco Nexus Dashboard Fabric Controller, Release 12.2.3

New and Changed Information 2

Verified Scale Limits for Release 12.2.3 2

Full Cisco Trademarks with Software License ?

Revised: February 27, 2026

New and Changed Information

The following 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.

Date	Changes
July 16, 2025	First release of this document.

Verified Scale Limits for Release 12.2.3

This section provides verified scalability values for various deployment types for Cisco Nexus Dashboard Fabric Controller, Release 12.2.3.

The values are validated on testbeds that are enabled with a reasonable number of features and aren't theoretical system limits for Cisco Nexus Dashboard Fabric Controller software or Cisco Nexus/MDS switch hardware and software. When you try to achieve maximum scalability by scaling multiple features at the same time, results might differ from the values that are listed here.

Nexus Dashboard System Resources

The following table provides information about Server Resource Requirements to run NDFC on top of Nexus Dashboard. Refer to Nexus Dashboard Capacity Planning to determine the number of switches supported for each deployment.

Cisco Nexus Dashboard can be deployed using number of different form factors. NDFC can be deployed on the following form factors:

- pND - Physical Nexus Dashboard
- vND - Virtual Nexus Dashboard

Table 1: Server Resource Requirements to run NDFC on top of Nexus Dashboard

Deployment Type	Node Type	CPUs	Memory	Storage (Throughput: 40-50 MB/s)
Fabric Discovery	Virtual Node (vND) – app node	16 vCPUs	64 GB	550 GB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2 x 10-core 2.2GHz Intel Xeon Silver CPU	256 GB of RAM	4 x 2.4 TB HDDs 400 GB SSD 1.2 TB NVME drive
	Physical Node (pND) (PID: ND-NODE-L4)	2.8GHz AMD CPU	256 GB of RAM	4 x 2.4 TB HDDs 960 GB SSD 1.6 TB NVME drive

Deployment Type	Node Type	CPUs	Memory	Storage (Throughput: 40-50 MB/s)
Fabric Controller	Virtual Node (vND) – app node	16 vCPUs	64 GB	550 GB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2 x 10-core 2.2GHz Intel Xeon Silver CPU	256 GB of RAM	4 x 2.4 TB HDDs 400 GB SSD 1.2 TB NVME drive
	Physical Node (pND) (PID: ND-NODE-L4)	2.8GHz AMD CPU	256 GB of RAM	4 x 2.4 TB HDDs 960 GB SSD 1.6 TB NVME drive
SAN Controller	Virtual Node (vND) – app node (with SAN Insights)	16 vCPUs (with physical reservation)	64 GB (with physical reservation)	550 GB SSD
	Data Node (vND) – Data node (with SAN Insights)	32 vCPUs (with physical reservation)	128GB (with physical reservation)	3 TB SSD
	Physical Node (pND) (PID: SE-NODE-G2)	2 x 10-core 2.2GHz Intel Xeon Silver CPU	256 GB of RAM	4 x 2.4 TB HDDs 400 GB SSD 1.2 TB NVME drive
	Physical Node (pND) (PID: ND-NODE-L4)	2.8GHz AMD CPU	256 GB of RAM	4 x 2.4 TB HDDs 960 GB SSD 1.6 TB NVME drive

Scale Limits for NDFC Fabric Discovery

Table 2: Scale Limits for Fabric Discovery Persona and Nexus Dashboard

Profile	Deployment Type	Verified Limit
Fabric Discovery	1-Node vND (app node)	100 switches
Fabric Discovery	3-Node vND (app node)	200 switches
Fabric Discovery	5-Node vND (app node)	1000 switches
Fabric Discovery	1-Node pND	100 switches
Fabric Discovery	3-Node pND	1000 switches

Scale Limits for NDFC Fabric Controller

Table 3: Scale Limits for Fabric Controller Persona and Nexus Dashboard

Profile	Deployment Type	Verified Limit
Fabric Controller	1-Node vND (app node)	50 switches
Fabric Controller	3-Node vND (app node)	100 switches
Fabric Controller	5-Node vND (app node)	400 switches for Easy Fabrics[1]
Fabric Controller	5-Node vND (app node)	1000 switches for External Fabrics[2]
Fabric Controller	1-Node pND	50 switches
Fabric Controller	3-Node pND	500 switches for Easy Fabrics[1]
Fabric Controller	3-Node pND	1000 switches for External Fabrics[2]

[1] Easy Fabrics include Data Center VXLAN EVPN fabrics and BGP fabrics.

[2]External Fabrics include Flexible Network fabrics, Classic LAN fabrics, External Connectivity Network fabrics, and Multi-Site Interconnect Network fabrics. Both managed and monitored mode are supported.

Table 4: Scale Limits for Switches and Fabrics in Fabric Controller

Description	Verified Limit
Switches per fabric	200
Physical Interfaces per NDFC instance[1]	30000

[1] Supported scale for 1-node vND is 2500 physical interfaces.

Table 5: Scale Limits For Provisioning New Data Center VXLAN EVPN Fabrics (also referred to as "Greenfield" Deployment)

Description	Verified Limit
Fabric Underlay Overlay	
Switches per fabric	200
Overlay Scale for VRFs and Networks[1]	500 VRFs, 2000 Layer-3 Networks or 2500 Layer-2 Networks
VRF instances for external connectivity	500
IPAM Integrator application	150 networks with a total of 4K IP allocations on the Infoblox server

Description	Verified Limit
ToR and Leaf devices	A Data Center VXLAN EVPN fabric can manage both Layer-2 ToR switches and leaf switches. Maximum scale for this sort of fabric is 40 leaf switches and 320 ToR switches. Maximum of 32 ToR switches (or 16 vPC-ToR pairs) can be connected per leaf-vPC pair.
Endpoint Locator[2]	
Endpoints	100000
VXLAN EVPN Multi-Site Domain	
Sites	30
Virtual Machine Manager (VMM)[3]	
Virtual Machines (VMs)	5500
VMware Center Servers	4
Kubernetes Visualizer application	Maximum of 160 namespaces with maximum of 1002 pods

[1] Supported scale for 1-node vND is 250 VRFS and 1000 networks.

[2] Supported scale for 1-node vND is 1 instance of endpoint locator with 10000 endpoints.

[3] Supported scale for 1-node vND is 1 VMware Center Server and 1000 VMs.

Note

- Refer to the following table if you are transitioning a command line interface (CLI) configured Cisco Nexus 9000 series switches based VXLAN EVPN fabric to NDFC.
- If you are performing a Brownfield migration, where you transition an existing Data Center VXLAN EVPN fabric management to NDFC, note that the scale limts provided below no longer apply once that Brownfield migration is completed and you should be using the scale limits provided in the Greenfield table above instead.

Table 6: Scale Limits For Transitioning Existing Data Center VXLAN EVPN Fabric Management to NDFC (also referred to as "Brownfield Migration")

Description	Verified Limit
Fabric Underlay and Overlay	
Switches per fabric	200
Physical Interfaces	11500
VRF instances	400
Overlay networks	1050
VRF instances for external connectivity	400
Endpoint Locator	

Description	Verified Limit
Endpoints	50000
IPAM Integrator application	150 networks with a total of 4K IP allocations on the Infoblox server
Virtual Machine Manager (VMM)	
Virtual Machines (VMs)	5500
VMware Center Servers	4
Kubernetes Visualizer application	Maximum of 160 namespaces with maximum of 1002 pods

Scale Limits for Cohosting NDFC and Other Services

Table 7: Scale Limits for Cohosting Nexus Dashboard Insights and NDFC (NDFC 12.1.3/ND 3.0.1 and earlier)

Profile	Deployment Type	Verified Limit
Nexus Dashboard Insights and Nexus Dashboard Fabric Discovery	3-Node pND	• 50 switches • 10,000 flows/sec
Nexus Dashboard Insights and Nexus Dashboard Fabric Controller	3-Node pND	• 50 switches • 10,000 flows/sec

Table 8: Scale Limits for Cohosting Nexus Dashboard Insights and NDFC (NDFC 12.2.1/ND 3.1 and later)

Profile	Deployment Type	Verified Limit
Nexus Dashboard Insights and Nexus Dashboard Fabric Discovery (NX-OS without controller mode[1])	3-Node pND	• 250 switches • 10,000 flows/sec
Nexus Dashboard Insights and Nexus Dashboard Fabric Controller	3-Node pND	• 250 switches • 10,000 flows/sec

[1] NX-OS Discovery mode is required when you deploy Nexus Dashboard Insights for NX-OS fabrics without using NDFC.

Scale Limits for IPFM Fabrics

Table 9: Scale Limits for Nexus Dashboard and IPFM Fabrics

Profile	Deployment Type	Verified Limit
Fabric Controller	1-Node vND	35 switches (2 spine switches and 33 leaf switches)

Profile	Deployment Type	Verified Limit
Fabric Controller	3-Node vND	120 switches (2 spine switches, 100 leaf switches, and 18 Tier-2 leaf switches)
Fabric Controller	1-Node pND	35 switches (2 spine switches and 33 leaf switches)
Fabric Controller	3-Node pND	120 switches (2 spine switches, 100 leaf switches, and 18 Tier-2 leaf switches)

Table 10: Scale Limits for IPFM Fabrics

Description	Verified Limit			
	NBM Active Mode Only	NBM Passive Mode Only	Mixed Mode	
			NBM Active VRF	NBM Passive VRF
Switches	120	32	32	32
Number of flows	32000	32000	32000	32000
Number of End Points (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
PTP Monitoring	120 switches	32 switches	32 switches	32 switches

Scale Limits for NDFC SAN Controller

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

- For 1-node clusters, **80** is the maximum number of SAN fabrics that can be managed by NDFC
- For 3-node clusters, **160** is the maximum number of SAN fabrics that can be managed by NDFC

Table 11: Scale Limits for SAN Zones

Description	**Verified Limits**
Zone sets	1000
Zone	16000

Table 12: Scale Limits for Nexus Dashboard and SAN Controller Persona

Profile	**Deployment Type**	**Verified Limit**	
		Without SAN Insights	**With SAN Insights**
SAN Controller	1-Node vND (app node)[1]	80 switches, 20K ports	40 switches, 10K ports, and 40K ITs
	1-Node vND (data node)	80 switches, 20K ports	80 switches, 20K ports, and 1M ITLs/ITNs [2]
	1-Node pND (SE)	80 switches, 20K ports	80 switches, 20K ports, and 120K ITLs/ITNs
SAN Controller	3-Node vND (app node)	160 switches, 40K ports	80 switches, 20K ports, and 100K ITs
	3-Node vND (data node)	160 switches, 40K ports	160 switches, 40K ports, and 240K ITLs/ITNs
	3-Node pND	160 switches, 40K ports	160 switches, 40K ports, and 500K ITLs/ITNs

[1] App 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. NDFC 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

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 NONINFRINGEMENT 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.

All printed copies and duplicate soft copies of this document are considered uncontrolled. See the current online version for the latest version.

Cisco has more than 200 offices worldwide. Addresses and phone numbers are listed on the Cisco website at www.cisco.com/go/offices.

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.