



Ncore Ceph RBD Gateway

Quick Start Guide

Product Version 26.1 | Guide Revision 1.1 | August 2026

Table of Contents


Introduction	3
1. CPU Core Isolation	4
2. Quick Start	5
3. Connection Settings	6
4. Dashboard Overview	7
5. Attach Ceph RBD Images to Ncore Gateway	8
6. iSCSI Targets [how to export block devices via iSCSI]	9
7. CHAP Management - User	10
7a. CHAP Management - Mutual	11
8. NVMe-oF Subsystems [how to export block devices via NVMe-oF]	12
9. Basic Troubleshooting	13

Welcome to Ncore Ceph RBD Gateway

Ncore Ceph RBD Gateway provides a unified management platform to discover, attach, and expose Ceph RBD images through industry-standard block-storage protocols.

With Ncore Ceph RBD Gateway, you can:

- Expose Ceph RBD images through iSCSI targets.
- Expose Ceph RBD images through NVMe-oF using TCP or RDMA.
- Provide block-storage access to Windows, Linux, and enterprise applications.
- Manage RBD images, iSCSI targets, and NVMe-oF subsystems from one desktop interface.



Ncore Ceph RBD Gateway architecture overview

Operational model: The Windows management application connects securely to the Ncore Gateway service. The gateway attaches Ceph RBD images and presents them to remote clients through iSCSI or NVMe-oF.

Capability	Purpose
Unified management	Manage Ceph RBD images and storage exports from one interface.
Enterprise-ready operation	Use authenticated management, TLS, and protocol-level host controls.
High-performance data path	Use dedicated CPU processing groups and accelerated storage transports.
Observability	Review attachment state, export assignment, active connections, and gateway status.

1. Ncore Gateway - CPU Core Isolation

Each exported Ceph RBD is assigned a dedicated CPU core. When the configured exports exceed the available cores, Ncore distributes additional exports across the existing processing cores.




Figure 1 - CPU core and export mapping

Benefit	Description
Parallel processing	Each export is processed independently when a dedicated core is available.
Minimal contention	Dedicated processing groups provide predictable workload isolation.
Optimized mapping	The gateway pins protocol processing groups to configured CPU cores.

Live Connection Mapping						
Each active iSCSI connection is assigned to a dedicated processing group on its designated CPU core.						
Target Configurations Portal Groups Initiator Groups CHAP Active Connections						
Active Connections						
Connection	State	Address	Target Node	Login Phase	Thread	TSIH
ID: 1 (CID: 1)	Running	192.168.100.33 → 192.168.100.2	db01	full feature phase	iscsi_poll_group_81	2
ID: 2 (CID: 1)	Running	192.168.100.33 → 192.168.100.2	vm01	full feature phase	iscsi_poll_group_82	3
ID: 3 (CID: 1)	Running	192.168.100.33 → 192.168.100.2	db92	full feature phase	iscsi_poll_group_83	4

Figure 2 - Live connection mapping

2. Quick Start

Prepare the Linux host by configuring HugePages, optional RDMA modules, the container runtime, and the Ncore Gateway service.

Component	Recommended
Operating system	64-bit Linux distribution, such as Ubuntu, Red Hat Enterprise Linux, or Rocky Linux
CPU	Multi-core x86_64 processor
Memory	16 GB RAM minimum
HugePages	8 x 1 GB HugePages or higher minimum
Network	10 GbE or faster
Container runtime	Docker Engine or Podman

Configure HugePages

Ncore SAN supports 1 GB HugePages.

1 GB HugePages

8 pages x 1 GB = 8 GB

```
default_hugepagesz=1G hugepagesz=1G hugepages=20
```

Ubuntu or Debian

Add the selected options to GRUB_CMDLINE_LINUX_DEFAULT in /etc/default/grub, preserving existing options. Then run:

```
sudo update-grub
sudo reboot
```

Red Hat Enterprise Linux or Rocky Linux

Add the selected options to all installed kernel entries, then restart the host:

```
sudo grubby --update-kernel=ALL --args="<selected HugePage options>"
sudo reboot
```

Verify HugePages

```
grep -E '^HugePages_|^Hugepagesize|^Hugetlb' /proc/meminfo
```

Expected: 1 GB pages = HugePages_Total 8 and Hugepagesize 1048576 kB; 2 MB pages = HugePages_Total 4096 and Hugepagesize 2048 kB.

Configure optional RDMA modules

For NVMe-oF over RDMA, load the modules applicable to the installed adapter: **rdma_cm**, **rdma_ucm**, **ib_uverbs**, **ib_core**, and **iw_cm**. RDMA modules are not required for iSCSI or NVMe-oF over TCP.

Verify the container and gateway

- Confirm Docker Engine or Podman is installed and enabled.
- Confirm the Ncore SAN container is running.
- Confirm a valid Ncore SAN license is installed.
- Launch the Windows management application and connect to the configured gateway URL.

3. Connection Settings

Configure the management application to connect to the Ncore Gateway, authenticate, install the TLS certificate, and verify connectivity.

How to Connect to Ncore Ceph Gateway
Follow the steps below to connect to the Ncore Ceph Gateway.

Connection Settings

- 1 Gateway URL:
- 2 Username:
- 3 Password:
- 4 TLS Certificate: ✓ Trusted

Import the Ncore Ceph Gateway self-signed certificate into the Windows Trusted Root Certification Authorities store.

Figure 3 - Connect to Ncore Ceph RBD Gateway

Setting	Description
Gateway URL	Ncore Gateway HTTPS management endpoint, for example https://192.168.100.2:9000.
Username	User account configured for gateway authentication.
Password	Password associated with the configured user account.
TLS certificate	Certificate used to establish trusted HTTPS communication.
Install certificate	Imports the certificate into the Windows Trusted Root Certification Authorities store.
Test Connection	Verifies HTTPS connectivity and authentication before saving.
Save Settings	Stores the connection configuration for future sessions.
Result: After a successful test, the application status bar displays Connected and the management functions become available.	

4. Dashboard Overview

The main dashboard lists attached Ceph RBD images and provides direct access to metrics, protocol exports, settings, search, refresh, and connection status.




Figure 4 - Ncore Ceph RBD Gateway dashboard

#	Control	Description
1	Advanced Metrics	Open the metrics and performance dashboard.
2	iSCSI	Navigation Icon, Open iSCSI management window and display device count.
3	Ceph RBD	Navigation Icon, Open RBD management window and display device count.
4	NVMe-oF	Navigation Icon, Opens NVMe-oF management window and display device count.
5	Settings	Open application settings.
6	Add RBD	Open the dialog window to attach a Ceph RBD image.
7-9	Refresh and search	Reload the list, search by name, or clear the search.
10-12	Status and assignment	Adjust zoom, review gateway status, and see where each RBD is used.

5. Attach Ceph RBD Images to Ncore Gateway

Select an available Ceph RBD image, assign a unique Name, block size, and attach it to the gateway.

The screenshot shows the 'Ncore Ceph RBD' interface. It features a table with columns: Image, Pool, Status, and Bdev. The table lists six RBD images: 'ssd-rbd1' (Ready), 'ssd-rbd2' (Attached), 'ssd-rbd3' (Ready), 'vm01' (Attached), 'vm02' (Attached), and 'db01' (Attached). Below the table, there is a 'Bdev Name' input field (labeled 2), a 'Block Size' selection (512 or 4096, with 512 selected, labeled 3), a 'Refresh' button, and an 'Attach to Gateway' button (labeled 4).

✓		Image	Pool	Status	Bdev
☐		ssd-rbd1	ssdpool	Ready	-
☐		ssd-rbd2	ssdpool	Attached	ssd-rbd2
☐		ssd-rbd3	ssdpool	Ready	-
☐		vm01	vm_pool	Attached	vm01
☐		vm02	vm_pool	Attached	vm02
☐		db01	database_pool	Attached	db01

Bdev Name:

Block Size: ☒ 512 ☐ 4096

Figure 5 - Attach a Ceph RBD image

1. Select Ceph RBD from the navigation panel and click Add RBD.
2. Select an RBD with Ready status.
3. Enter a unique name.
4. Select 512 bytes (default) or 4096 bytes for the block size.
5. Click Attach to Gateway.

Status	Meaning
Ready	The RBD image is available and can be attached.
Attached	The RBD image is already attached to the Ncore Gateway.

6. iSCSI Targets

Create and manage iSCSI targets that expose attached Ceph RBD block devices over standard IP networks.

Create iSCSI Disk

Block Device Search by name...

☒		Name	Size (GB)	Used By
☐		vm01	100.00	[iSCSI] --> vm01
☐		vm02	100.00	[iSCSI] --> vm02
☐		db01	100.00	[iSCSI] --> db01
☐		db02	100.00	
☒		archive91	100.00	

Target
Target Name:

Initiator Configuration
Initiator Network:

Portal Configuration
IP: Port:

CHAP
Select CHAP Group:

Create **Cancel**

Figure 6a - Create an iSCSI target

1. Select iSCSI from the navigation panel and click Create iSCSI Disk.
2. Select one or more unassigned Ceph RBD block devices.
3. Enter a unique target name, initiator network, and portal IP address. The default iSCSI port is 3260.
4. Optionally select a CHAP authentication group.
Review the configuration and click Create.

iSCSI DISKS

Target Configuration Portal Groups Initiator Groups CHAP Active Connections

+ Create iSCSI Disk Refresh Grid Search by name...

iSCSI Target Configurations

#		Name	LUNs	Alias	PG Tag	IG Tag	Auth	Queue Depth
1		iqn.2016-06.io.spdk:vm01	vm01	vm01_target	2	2	Disabled	64
2		iqn.2016-06.io.spdk:vm02	vm02	vm02_target	2	2	Disabled	64

Figure 6b - iSCSI target configurations

7. CHAP Management - Create CHAP User and Auth Group

Create CHAP authentication groups and user credentials for authenticated iSCSI access.

N CHAP Management ✕

☒ CHAP ☐ Mutual CHAP

#	Auth Group	User	Secret
1	1	UserDev	*****
2	2	ncoreUser1	*****

Auth Group: 5 User: 6 Secret: 7 8

Tag: 3 4 Close

Figure 7 - CHAP management

1. Open iSCSI management and select CHAP Management.
2. Select the CHAP tab.
3. Enter a numeric Tag value (e.g. 2) if the ID don't exists, else go to Step 5.
4. click Create to create an authentication group.
5. Select the new authentication group (select a Tag ID from the drop-down list).
6. Enter the CHAP username
7. Enter user secret.
8. Click Add.

Security: Use unique CHAP credentials for each intended access scope and store secrets securely.

7a. CHAP Management - Create Mutual CHAP User

Add mutual CHAP credentials when both the initiator and target must authenticate each other.

CHAP Management

☐ CHAP
 ☒ Mutual CHAP

Auth Groups - Mutual CHAP Users

#	Auth Group	User	Secret	Mutual User	Mutual Secret
1	1	UserDev	*****		
2	2	ncoreUser1	*****	UserProd1	*****

Add Mutual CHAP User

Auth Group:
 User:
 Secret:
 Add:

MUser:
 MSecret:

Create Auth Group

Tag:
 Create:

Close

Figure 8 - Mutual CHAP management

1. Open iSCSI management and select CHAP Management.
2. Select the Mutual CHAP tab.
3. Enter new Tag ID if not already created
4. Click Create to add new Tag ID (e.g. 2) if the ID don't exist, else go to Step 5.
5. Select an authentication group (select a Tag ID from the drop-down list)
6. Enter the CHAP username
7. Enter the User secret.
8. Enter the MUser name.
9. Enter the MUser secret/password.
10. Click Add.

8. NVMe-oF Subsystems

Create NVMe-oF subsystems that expose attached Ceph RBD images to remote hosts using TCP or RDMA.

NVMe-oF Subsystem Builder — HowTo Guide
 Use this page to configure and create a new NVMe-oF subsystem.

NVMe-oF Subsystem Builder

☒ TCP
 ☐ RDMA

✓		Name	Size (GB)	Used By
☒		vm01	100.00 GB	
☐		vm02	100.00 GB	[iSCSI] --> vm02
☐		db01	100.00 GB	[iSCSI] --> db01
☐		prod01	100.00 GB	[NVMe-OF] --> prod01

Namespace ID:
 Port:

Subsystem/Export Name:

Listener Address:
 Address Family:

☒ Allow Any Host

Host NQN:

Cancel
 Create Subsystem

Figure 9a - NVMe-oF subsystem builder

1. Select NVMe-oF from the navigation panel and click Create Subsystem.
2. Choose TCP or RDMA.
3. Select an unassigned Ceph RBD block device.
4. Enter the namespace ID, subsystem NQN, listener address, port, and address family.
5. deselect "Allow Any Host" to authorized specific host NQN.
6. Review the configuration and click Create Subsystem.

NVMe-oF Subsystems							
#	NQN	Backend Device	SubType	AllowAnyHost	ListenerAddress	Hosts	Namespaces
1	nqn.2016-06.org.nvmexpress.discovery	—	Discovery	True	—	Any	None
2	nqn.2024-03.io.storage-sp-prod-rbd	rbd1	NVMe	True	TCP: 192.168.100.2:4420	Any	NS 1: rbd1
3	nqn.2016-03.io.storage-dlb-prod-cache1	rbd2	NVMe	True	TCP: 192.168.100.2:4420	Any	NS 1: rbd2
4	nqn.2024-03.io.storage-dlb-prod1	prod01	NVMe	True	TCP: 192.168.100.2:4420	Any	NS 1: prod01

Figure 9b - NVMe-oF subsystems

9. Basic Troubleshooting

Use these checks to resolve common issues in the Ncore management interface. Record the displayed error and affected object name if assistance is required.

Issue	Initial check
Cannot connect to the gateway	Verify the HTTPS gateway URL, username, password, TLS certificate, and license status. Click Test Connection.
Ceph RBD image is not listed	Click Refresh, clear the search text, and confirm that the Ceph pool and image are accessible to the gateway.
RBD cannot be attached	Confirm the RBD status is Ready and enter a unique bdev name.
RBD cannot be selected for export	Check Used By. Select an unassigned RBD or remove the existing export first.
Cannot create an iSCSI target	Verify the target name, initiator network, portal address, port 3260, selected RBD, and optional CHAP group.
iSCSI client cannot connect	Verify the portal, initiator network, firewall, and CHAP credentials.
NVMe-oF connection fails	Verify transport, listener address, port 4420, host NQN, and RBD assignment.
CHAP authentication fails	Verify the authentication group, username, secret, and mutual credentials when enabled.
Changes are not displayed	Click Refresh, clear search filters, or reopen the applicable Ncore page.
Operation remains unavailable	Check the gateway connection and license status; record the exact error before contacting support.

Still need assistance?

Provide the error message, gateway status, and affected RBD, target, or subsystem name to support@ncorestorage.com.