



RAIDIX 5.2 ADMINISTRATOR'S GUIDE

Document version 1.23

CONTENTS

Introduction.....	9
Intended Audience.....	9
How the Guide is Organized.....	9
Guide Conventions.....	9
Contact Us.....	10
Third-party Software Use Limitation	10
About RAIDIX 5.2.....	11
Dual-Controller Mode	13
1. Web Interface Overview	15
1.1 Web Interface Language.....	16
1.2 Wizards.....	17
1.3 Surrounding RAIDIX List.....	17
2. Quick Start	19
2.1 RAID Creation.....	19
2.2 LUN Creation	21
2.3 Setting up Host Connection	22
2.3.1 InfiniBand Setup	22
2.3.2 iSER Setup	23
2.3.3 iSCSI Setup.....	25
2.3.4 Fibre Channel Setup.....	27
2.4 NAS Share Creation	27
3. Working with RAID	29
3.1 RAID Levels.....	29
3.2 RAID Engines.....	31
3.2.1 Generic	32
3.2.2 ERA.....	32
3.3 RAID Creation.....	33
3.4 RAID States.....	38
3.5 RAID Initialization	39
3.6 Scheduled Background Integrity Check	41
3.6.1 Features of Background Integrity Check	41
3.6.2 Configuration of Background Integrity Check.....	42
3.6.3 Background Integrity Check in DC	46
3.7 RAID Reload.....	46

3.8	RAID Migration.....	47
3.9	RAID Restriping.....	48
3.10	RAID Deletion	51
4.	Advanced RAID Parameters	52
4.1	Working with Advanced RAID Parameters	57
4.1.1	Advanced Reconstruction Setup.....	59
4.2	RAID Profiles	60
4.2.1	Creating a RAID Profile.....	61
4.2.2	Modifying a RAID Profile.....	62
4.2.3	Downloading a RAID Profile	62
4.2.4	Deleting a RAID Profile	62
5.	Drives.....	63
5.1	Information about drives.....	63
5.2	Drive Health	64
5.2.1	Error Counter.....	65
5.2.2	Health Tracking	66
5.2.3	S.M.A.R.T. Diagnostics.....	67
5.3	Drive Performance	70
5.4	Drive Replacement in the RAID.....	72
5.4.1	Automatic Drive Replacement.....	73
5.4.2	Manual Drive Replacement	75
5.4.2.1	NVMe SSD Hot Plugging	77
6.	LUN	79
6.1	Creating the LUN	79
6.2	LUN Extension.....	83
6.2.1	LUN Extension over the RAID.....	83
6.2.2	LUN Extension over Multiple RAIDs.....	84
6.3	Masking Rules.....	85
6.3.1	Target Masking Rules.....	86
6.3.1.1	Create a Target Masking Rule	86
6.3.2	Host Masking Rules.....	88
6.3.2.1	Create a Host Masking Rule	89
6.3.2.2	Change a Default Masking Rule.....	90
6.3.2.3	Delete a Host Masking Rule.....	90
6.3.3	Unmasking LUN.....	91
6.4	SSD Cache Setup.....	92

6.4.1	Configuring SSD Cache.....	93
6.4.2	Disabling SSD Cache	96
6.5	Replication	96
6.5.1	Overview of the Replication.....	97
6.5.2	Creating the LUN with Replication Metadata.....	97
6.5.3	Replication Settings	99
6.5.3.1	Replication Settings for SC Systems	100
6.5.3.2	Replication Settings for DC system.....	105
6.5.3.3	Changing Replication Nodes Addresses	110
6.5.4	Managing Replication.....	110
6.5.4.1	Suspension and Resumption of Replication.....	110
6.5.4.2	Consistency Check	111
6.5.4.3	Replacing Data with Replica	111
6.5.4.4	Changing Replication Roles	112
6.5.4.5	Changing Replication Protocols.....	112
6.5.4.6	Disabling of Replication	112
6.5.5	Disabling the LUN with Replication Metadata	113
6.6	LUN Deletion.....	113
7.	Network Attached Storage (NAS).....	114
7.1	Creating a NAS Share	114
7.2	Managing NAS Users and Groups	119
7.2.1	Create Users and Groups	119
7.2.2	Edit and Delete NAS Users and Groups	120
7.2.3	Active Directory Connection.....	122
7.2.3.1	Setting up Active Directory Connection.....	122
7.2.3.2	Set up Trusts Between Active Directory Domains.....	123
7.2.4	LDAP Connection	124
7.2.4.1	Connecting to LDAP Server.....	124
7.2.5	Assigning User and Group for a Share.....	126
7.3	Managing Quotas	127
7.4	Editing a NAS Share	128
7.5	Delete a NAS Share.....	128
8.	Working with Hosts and Interfaces.....	130
8.1	Working with Host in DC	130
8.2	Managing Host Aliases	130
8.2.1	Creating a Host Alias	130
8.2.1.1	Set up Advanced Prioritization Parameters.....	132

8.2.2	Deleting a Host Alias.....	133
8.2.3	Viewing Established Sessions	133
8.3	Working with Adapters	133
8.3.1	Information on iSCSI Ports.....	134
8.3.2	Information on InfiniBand Adapters	135
8.3.3	Information on Fibre Channel Adapters.....	135
8.3.3.1	Modifying the Fibre Channel parameters	136
8.3.4	Information on SAS Adapters	136
8.4	Setting up iSCSI	137
8.4.1	Enabling and Disabling iSCSI Support.....	137
8.4.2	Disconnecting an iSCSI Initiator on Windows Server	138
8.4.3	Creating an iSCSI Target.....	138
8.4.3.1	Creating CHAP users.....	140
8.4.4	Setting up the Advanced iSCSI Target Parameters	141
8.4.5	Deleting an iSCSI Target	142
8.5	Setting up QoSmic	142
8.5.1	QoSmic Learning	143
8.5.2	Starting QoSmic Learning.....	143
8.5.2.1	Stopping QoSmic Learning	145
8.5.3	Application Types.....	145
8.5.4	Import of Signatures.....	146
8.5.5	QoSmic Recognizing	146
8.6	NVMe-oF	147
8.6.1	Configuring Connection	148
8.6.2	ANA Configuration	150
8.6.3	Deleting Ports and Subsystems	150
9.	Working with SAN Optimizer	151
9.1	SAN Connection	151
9.1.1	Configuring FC Connection	151
9.1.2	Configure iSCSI Connection	154
9.1.3	Bonding of SAN Drives in DC	154
9.1.4	Creating Proxied LUN	156
9.1.5	Deleting SAN Connection Points.....	156
10.	Setting up System Parameters	157
10.1	Working with Nodes.....	157
10.1.1	Dual-Controller System Configuration	157
10.1.1.1	Setting up Heartbeat	158

10.1.1.2	Setting up Cache Synchronization of Nodes	160
10.1.2	Viewing Node Status	164
10.1.3	Failover and Migration of RAID	167
10.1.3.1	Manual Failover	169
10.1.3.2	Automatic Failover	169
10.1.3.3	Failback	170
10.1.3.4	Migration of RAID	170
10.1.4	Disabling Dual-Controller Mode.....	172
10.1.5	Rebooting, Shutting down, and Powering on DC System.....	173
10.1.6	Uninterruptible Power Source Configuration	174
10.2	Setting up Network Parameters	176
10.2.1	General Network Settings.....	176
10.2.2	Setting up SSL Certificate.....	177
10.2.3	Setting up Virtual IP Address	179
10.2.4	Setting up Network Interface Parameters.....	180
10.2.5	Setting up Channel Bonding.....	181
10.3	Working with License.....	182
10.3.1	Viewing Product License Information	183
10.3.2	License Management.....	185
10.4	Setting up Notifications.....	186
10.4.1	Setting up SMTP Server Parameters.....	187
10.4.1.1	Configuring SMTP Recipients	189
10.4.2	Setting up SNMP Server Parameters	190
10.4.2.1	Configuring SNMP Recipients	194
10.5	Configuration Recovery	195
10.5.1	Saving Current System Configuration.....	195
10.5.2	Downloading Configuration File	196
10.5.3	Uploading Configuration File	196
10.5.4	Applying System Configuration.....	196
10.5.5	Deleting Configuration File	196
10.5.6	Advanced Configuration of the System	196
10.5.6.1	Cloning System Configuration	196
10.5.6.2	Import of a RAID	197
10.6	Time Settings	199
10.7	Users.....	201
10.7.1	Managing System Users	202
10.7.2	Password Policy	203

10.8	System Drive	204
11.	System Monitoring	206
11.1	Sensors Info	206
11.1.1	Information on Controller Sensors	206
11.1.2	Information on Enclosure Sensors	207
11.2	Performance Info	209
11.2.1	Data Rate	210
11.2.2	IOPS Charts	212
11.2.3	Latency Charts	212
11.2.4	Transfer Size Charts	213
11.2.5	Trace Statistics	214
11.3	Console	220
11.4	Downloading System Logs	221
11.4.1	Downloading System Logs Using Web Interface	221
11.4.2	Downloading System Logs Using CLI	222
11.5	System Status Assessment	223
11.6	Command Queue	225
12.	Troubleshooting	227
	Glossary	237
	Appendix A. CLI Features and Syntax	241
	adapter	243
	dc	244
	drive	247
	eraraid	250
	fc	256
	host	259
	ib	260
	iscsi	262
	lun	268
	mask	274
	metadata	276
	nas	278
	network	298
	nodes	318
	notify mail	319

notify snmp	327
nvme fabrics	337
param	344
profile raid	355
qosmic	361
raid	363
replication	373
sas target	377
sensor	378
session	379
sparepool	380
ssdcache	381
system	383
log_collector	392
Appendix B. MPIO Setup on Linux Enterprise Server	393
Appendix C. Sending Logs to a Remote Linux System Through Rsyslog Service	396

INTRODUCTION

Intended Audience

This Guide is intended to help administrators that operate storage systems based on the RAIDIX 5.2 software (from now on, the system).

The Guide provides information on how to configure and maintain the storage system based on the RAIDIX 5.2 software.

How the Guide is Organized

You can find information on specific issues in the following chapters:

1. Web Interface Overview: navigation, language, wizards, surrounding RAIDIX systems.
2. Quick Start: creation of a functional storage system with default settings.
3. Working with RAID: properties, creation and deletion, integrity check, configuring base parameters, reload, migration.
4. Advanced RAID Parameters: configuration of RAID performance and reliability.
5. Working with Drives: SparePool, drive replacement, drive scan, S.M.A.R.T.
6. Working with LUN: creation and deletion, LUN extension, masking rules, replication, SSD cache.
7. Working with NAS: a share, SMB, NFS, FTP, AFP, assignment of access privileges, quota.
8. Working with Hosts and Interfaces: host alias, connected sessions, InfiniBand, Fibre Channel, iSCSI, iSER, QoS, NVMe-oF.
9. Working with SAN Optimizer: connection and disconnection, management of SAN drives in the DC system, proxied LUN.
10. Setting Up System Parameters: nodes, system modes (Single-Controller and Dual-Controller), Failover and Failback, license, network, system time, saving and applying of system configuration, sending system notifications.
11. System Monitoring: sensors, system performance, Command Line Interface (CLI), alert history, system logs.
12. Troubleshooting: description of possible troubles and solutions.


Guide Conventions

The Guide uses the typefaces and formatting to specify different names and terms:

Convention	Uses
Bold	GUI controls, option value, minor titles.
<i>Italic</i>	Emphasis, term references, documentation titles, section titles, file paths.
Text color	Instructions below are only for specific situations and configurations.
<code>Monospace</code>	Commands, command utilities, and console-driven text.

Text paragraphs that need your special attention are marked with the following frame:

 *Note* – a note, which provides valuable information.

 *Warning* – binding instructions to guarantee the proper work of the software.

Contact Us

If you have any comments or suggestions on the improvement of the RAIDIX end-user documentation, send them at doc.comments@raidix.com. Specify a document title, a chapter number, and a page. Refer to the specific pages and sections whenever possible.

Third-party Software Use Limitation

Perform the upgrade of the RAIDIX software only by update packages provided by the RAIDIX company. Don't use third-party updates and package installation tools (`yum`, `rpm`) because it may result in the system malfunction and loss of warranty.

ABOUT RAIDIX 5.2

The RAIDIX 5.2 software is a development of RAIDIX LLC that implements software-based RAID over operating system kernel. RAIDIX 5.2 enables you to create high-performance fault-tolerant storage systems. The software is distributed as an ISO image of an operating system with features available under a license.

RAIDIX 5.2 supports a *Single-controller* mode (SC mode) when one node is active (Figure 1) and a *Dual-controller* active-active mode (DC mode), when both nodes are active and have access to the same set of drives (enclosure) (Figure 2). To learn more, see [Dual-Controller Mode](#).

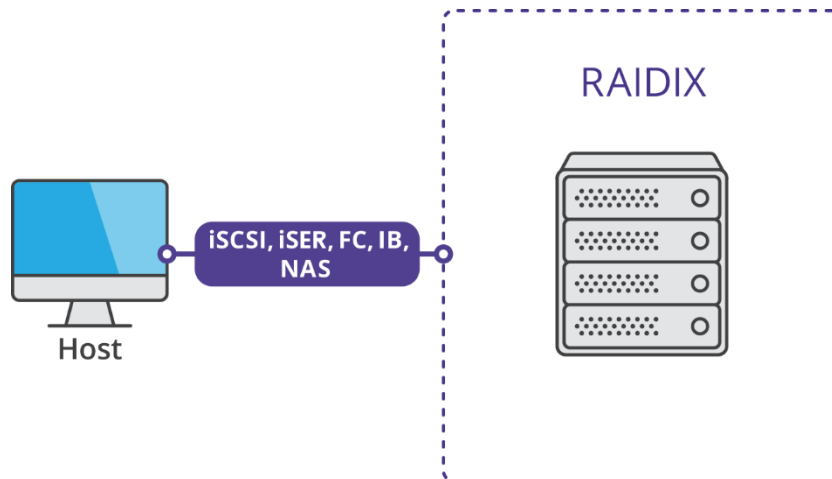


Figure 1. Workflow of RAIDIX 5.2 in Single-controller mode

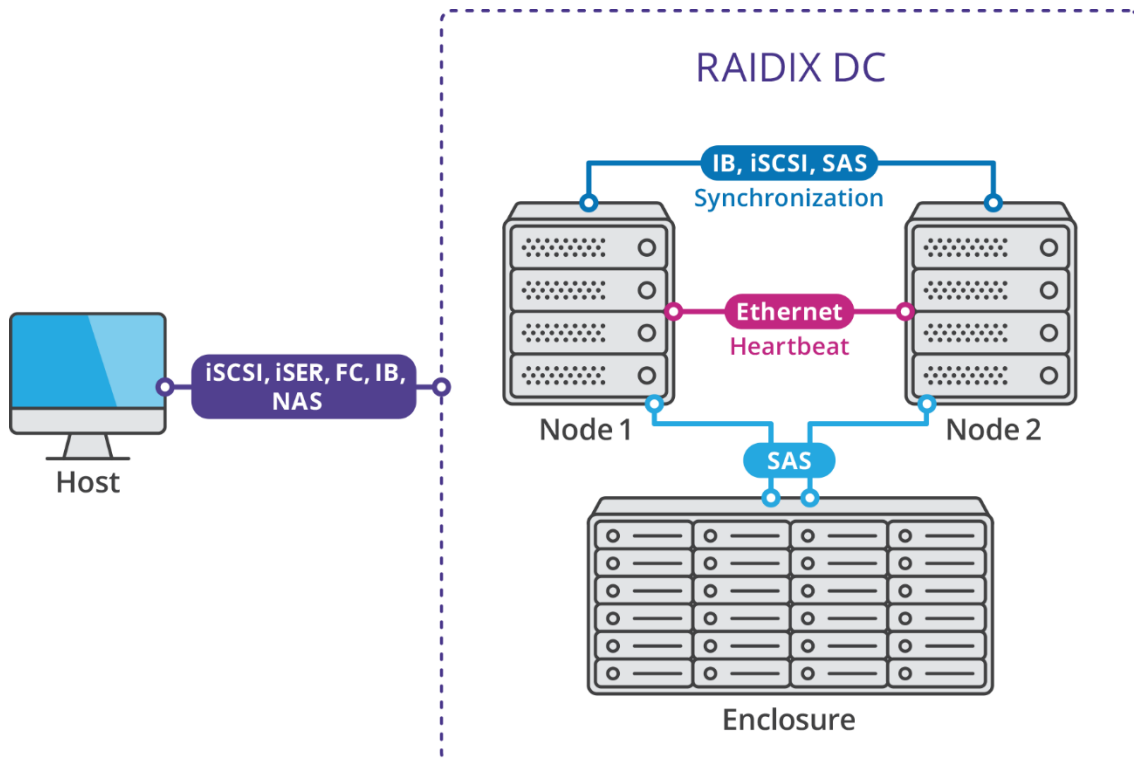


Figure 2. Workflow of RAIDIX 5.2 in Dual-controller mode

RAIDIX 5.2 enables you to build both Network Attached Storage (NAS) and Storage Area Network (SAN) solutions.

The system can be operated through web interface (hereafter – GUI) and command line interface (CLI). To learn more about CLI commands, see [Appendix A](#).

RAIDIX 5.2 Features

RAID levels	Generic RAID	ERA RAID
	<ul style="list-style-type: none"> • 0; • 10; • 5; • 6; • 7.3; • N+M. 	<ul style="list-style-type: none"> • 0; • 1; • 10; • 50; • 60; • 70; • 5; • 6; • 7.3; • N+M.
Maximum number of drives in one RAID	64.	
Maximum number of RAIDs	64.	
Maximum number of drives in the system	600.	
Number of LUNs	447.	
Number of hosts supported (direct connection)	32.	
Browsers support	Latest versions of the browsers: <ul style="list-style-type: none"> • Mozilla Firefox; • Opera; • Google Chrome; • Apple Safari; • Microsoft Edge (with Microsoft EdgeHTML 15 or higher). 	
Virtualization platforms support	<ul style="list-style-type: none"> • VMware ESXi Server 6.5, 6.7, 7.0; • RHEV (Red Hat Enterprise Virtualization) 4.4; • Microsoft Hyper-V Server 2019; • Proxmox VE 6.x, 7.x. 	
Client/initiator OS support	<ul style="list-style-type: none"> • macOS 10.14, 10.15; • Microsoft Windows Server 2016, 2019, 2022; • Microsoft Windows 10, 11; • Linux (including but not limited to): <ul style="list-style-type: none"> ○ RHEL (CentOS, Alma, Rocky) 7, 8; ○ Ubuntu 18.04, 20.04, 22.04; ○ SLES 12 SP5, 15 SP3; ○ ALT 8, 9. 	

Client/initiator OS compatibility	<ul style="list-style-type: none"> • macOS 11.
High-speed communication channel support	<ul style="list-style-type: none"> • InfiniBand (EDR, FDR, and QDR generations); • Fibre Channel 8 and 16 Gbit/s; • iSCSI (1, 10, 40, 100 Gbit/s); • iSER.
NVMe-oF	Over Infiniband in the ib or eth mode.
Supported NAS protocols	<ul style="list-style-type: none"> • SMBv3 and SMBv2; • NFSv4 and below; • FTPv1 and below; • AFPv3 and below.

Dual-Controller Mode

Dual-Controller mode ensures continuity of data access while one node is down. The RAIDIX 5.2 software achieves fault-tolerance by using two cluster nodes operating simultaneously in the *Active-Active* mode and having access to one set of drives.

Duplication of hardware components and interfaces protects against the following failures:

- failure of one of the hardware components (CPU, motherboard, power supply unit, controller, system drive);
- failure of an enclosure connection interface (SAS cable, I/O device);
- the shutdown of one node;
- software errors or system failure on one node.

If a failure occurs, the system performs automatic *failover*, which allows an administrator to fix errors without stopping the system. To learn more, see [Automatic failover](#).

Each node obtains metadata of the other node via heartbeat (over Ethernet channel).

Nodes interact over InfiniBand, Ethernet (iSCSI), and SAS interfaces, which enables synchronization of data and caches.

Due to the cache synchronization in both directions, a remote node always contains an up-to-date copy of data in a local node cache. Thus, if one of the nodes fails, the other node translucently takes the entire workload.



Before enabling DC, make sure both nodes have the same hardware configuration. Types and models of processors and adapters, and RAM sizes must be the same.

In RAIDIX 5.2, you can place RAIDs on the nodes asymmetrically. Each RAID will be Active on one node, which will control access to the RAID. You can specify the RAID preference parameter *Affinity* for each created RAID, so you can define on which node RAID will be *Active* (available for I/O operations) and on which it will be *Passive* (not available for I/O operations).

Dual-controller mode enables you to perform *RAID migration* from any node of the cluster to balance the nodes workload. This leads to the *Affinity* parameter change: RAID becomes active on another node. To learn more, see [Migrate a RAID](#).

Each node has an IP address. To learn more, see *RAIDIX 5.2 Software Installation and Configuration Guide*.

You can configure Dual-controller mode (set up heartbeat) on any of two nodes. The configuration requires the IP address of the second node. To learn more, see [Dual-Controller mode configuration](#).

1. WEB INTERFACE OVERVIEW

! Make sure JavaScript and cookies are enabled in your web browser.

! To start working with the system, complete the installation and the license configuration. To learn more, see *RAIDIX 5.2 Software Installation and Configuration Guide*.

To open the web interface:

1. Open any supported web browser. In the address bar, enter an IP address of the system using the following format:

`http://<IP_address>`

i A system administrator assigns an IP address of the system. To learn more, see *RAIDIX 5.2 Software Installation and Configuration Guide*.

! The system can also use HTTPS. Since RAIDIX uses a self-signed security certificate by default, the browser will generate the warning about insecure connection while using HTTPS.

To proceed with HTTPS:

- In web browser, open advanced settings and add an exception;
- or [set up SSL certificate](#).

2. In the opened window, enter a username and a password.

By default, to log in GUI, use the username **admin** and the password **raidix-era**.

3. Click **LOG IN**.

The opened **STORAGE | DASHBOARD** page (Figure 3) includes information on volumes, drives, RAIDs, SparePools, and general system status.

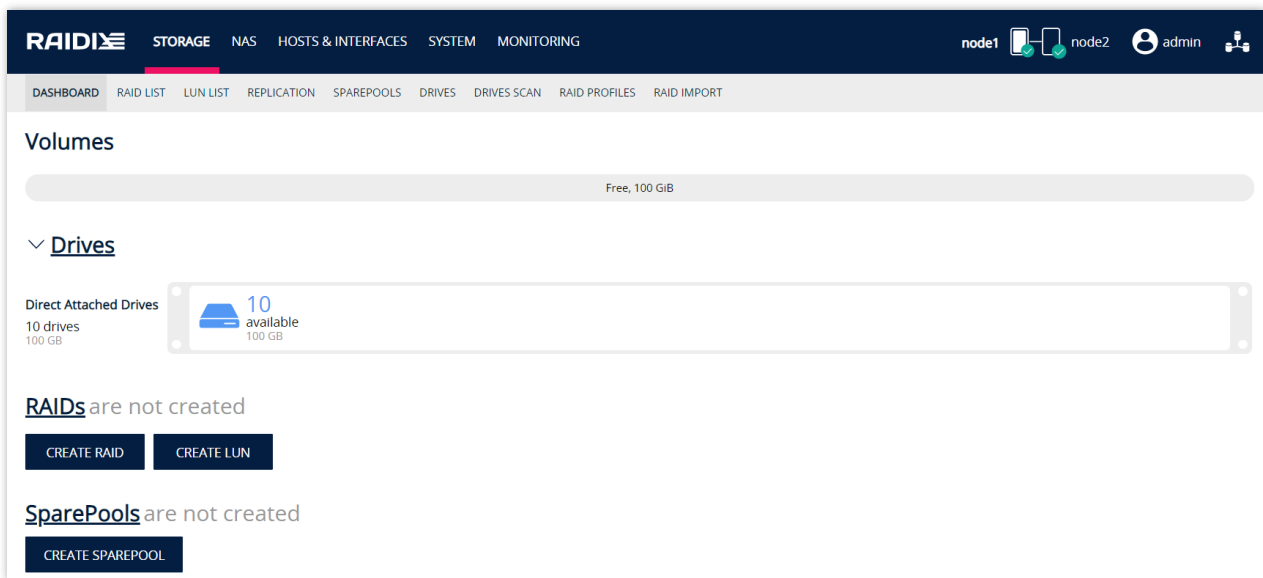


Figure 3. DASHBOARD page. The system is in DC mode

1.1 Web Interface Language

Available languages are:

- English;
- Russian;
- Japanese.

To select the web interface language, in the upper right corner of the window, hover the pointer on the username and select a language from the **Language** field (Figure 4).

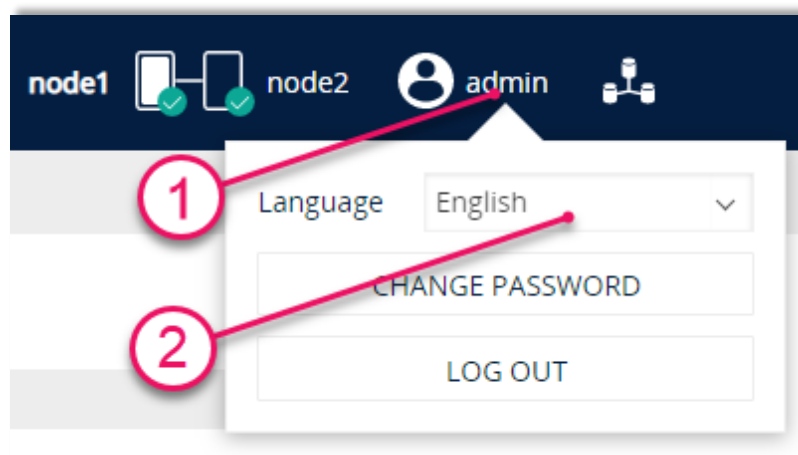


Figure 4. Selecting an GUI language

1.2 Wizards

The following Wizards are available:

1. Create RAID.

The wizard is located on the **STORAGE | DASHBOARD** page.

2. Create LUN.

The wizard is located on the **STORAGE | DASHBOARD** page.

3. Create SparePool.

The wizard is located on the **STORAGE | DASHBOARD** page.

4. Create Share.

The wizard is located on the **NAS | SHARES** page (Figure 71).


5. Configure DC.

The wizard is located on the **SYSTEM | NODES** page (Figure 111).

1.3 Surrounding RAIDIX List

The list of surrounding RAIDIX systems is made up of all RAIDIX systems that are running on the local network. You can gather information on nodes (types, names, alerts, cluster alerts) and failover statuses (Figure 5).

To view the Surrounding RAIDIX List page:

1. In the upper right corner of the menu, click  (Figure 5).
2. To find RAIDIX systems in your network:
 - with default search, click **DISCOVER**;
 - with specific search, specify **Interface**, **Broadcast Address**, and **Timeout** (waiting period for RAIDIX systems response) values, then click **DISCOVER**.

RAIDIX

STORAGE NAS HOSTS & INTERFACES SYSTEM MONITORING

node1 node2

admin

Surrounding RAIDIX List

Interface

Broadcast Address

Timeout

ens18

DISCOVER

Polled: 17 out of 63

Click to open the Surrounding RAIDIX List page

Figure 5. "Surrounding RAIDIX List" page

2. QUICK START

This chapter provides information about an overall picture of the functioning and management of the system. The chapter describes the step-by-step creation of the RAIDIX storage system with a minimal set of required parameters while additional settings remain default. You can change the additional parameters later optimizing the system for different workload patterns.

To find out about managing the DC mode, see the chapter [Dual-Controller Mode Configuration](#).

2.1 RAID Creation

To create a RAID:

1. Select **STORAGE > RAID LIST**.
2. In the *RAID List* section, click **Create RAID**. The widget of RAID creation opens (Figure 6).

Figure 6. RAID LIST page. RAID creation widget

3. In the **RAID Name** field, type a RAID name.

! You can't create an ERA RAID named *power* or *uevent*.

4. In the **RAID Engine** field, select a RAID Engine value (to learn more, see [RAID Engines](#)). Depending on selected RAID Engine, other RAID parameters will differ:

- **For the Generic RAID Engine, define the following parameters:**

- In the **RAID Level** field, select a level of the RAID (to learn more, see [RAID Levels](#)).
- In the **Drives** field, click **Select** and select drives from the list. When done, click **SELECT**.

i During creation of a Generic RAID, you can select both SSDs and HDDs. If a RAID contains both SSD and HDD, the SSD will be under specific consumption.

- In the **RAM cache size** field, select a RAM cache size (the default is minimal available value).

- In case of complex write patterns, turn on **4 KiB blocks write mode**. It may improve RAID performance.
- Manage **Write Barriers Support** for increasing either data reliability or system performance.
- In the **Stripe** field, select a stripe size (the default is **128 KiB**).
- In the **SparePool** field, select a SparePool of the replacing drives (the default is no value).

- For the ERA RAID Engine, define values for the following parameters:



ERA RAIDs are available only on systems with CPUs, which support AVX. ERA RAID N+M requires minimum AVX2.

- In the **RAID Level** field, select the level of the RAID. To learn more, see [RAID Levels](#).
- Next to the **Drives** field, click **Select** and select the drives from the list. To confirm, click **SELECT**.
- In the **Affinity** field, select the active node for the RAID (the default is node on which the RAID is creating).
- In the **Stripe** field, select the stripe size (the default is **Auto**).

Following parameters are only for ERA RAIDs all levels but 0.

- In the **SparePool** field, select available SparePool unit (the default is no value).
- In the **Memory limit** field, define the RAM limit size. The default is no value, which means the system automatically calculates the value based on Stripe size:
 - 16 KiB stripe – 2048 MiB RAM;
 - 32 KiB stripe – 2048 MiB RAM;
 - 64 KiB stripe – 4096 MiB RAM;
 - 128 KiB stripe – 8192 MiB RAM;
 - 256 KiB stripe – 16384 MiB RAM.

5. Click **CREATE**. In the *RAID List* section, the created RAID displays (Figure 7).

RAIDIX

STORAGE


NAS

HOSTS & INTERFACES

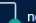
SYSTEM


MONITORING


node1



node2



 admin



DASHBOARD

RAID LIST

LUN LIST

REPLICATION

SPAREPOOLS

DRIVES

DRIVES SCAN

RAID PROFILES

RAID IMPORT

RAID List

>>

RAID Name

RAID Size

Free Size

RAID Level

Drives

RAID Engine

SparePool

Sync

Affinity/Status

>

[GEN_RAIDS](#)

39 GiB

39 GiB


RAID 5


0,2,5-7

Generic

—

On





Create RAID

Figure 7. Created Generic RAID

2.2 LUN Creation

To create a LUN:

1. Select **STORAGE > RAID LIST**.
2. Click the name of a RAID. The RAID page opens (Figure 8).

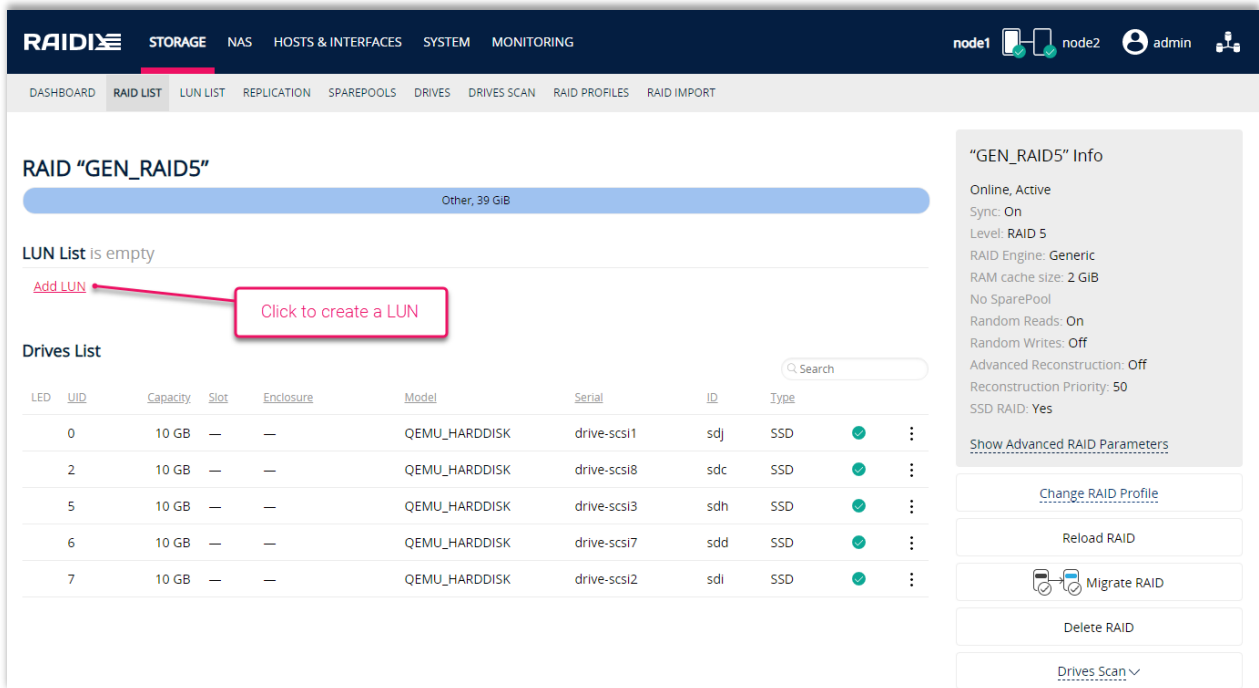


Figure 8. RAID page

3. In the *LUN List* section, click **Add LUN**. The widget to create a LUN shows (Figure 9).

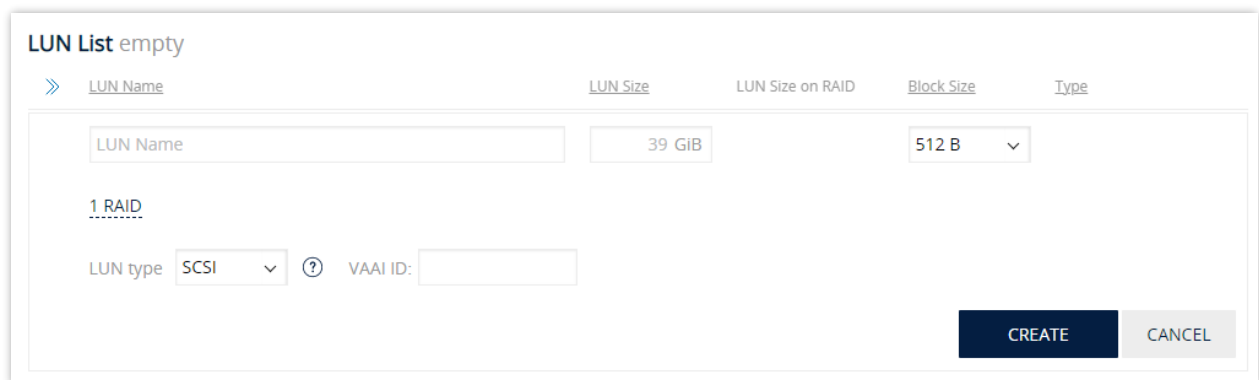


Figure 9. RAID LIST page. The widget for creating a LUN

4. In the appropriate fields, specify a LUN name, size, type, and Block Size.
If you are going to create a NAS share, select the type **NAS**.
5. Click **CREATE**.

The created LUN will show on the RAID page in the *LUN List* section (Figure 10), and on the **STORAGE | DASHBOARD** page in the graphic representation of the corresponding RAID.

The system sets some parameters of created LUN by default. To learn how to manage LUN parameters, see [Create a LUN](#).

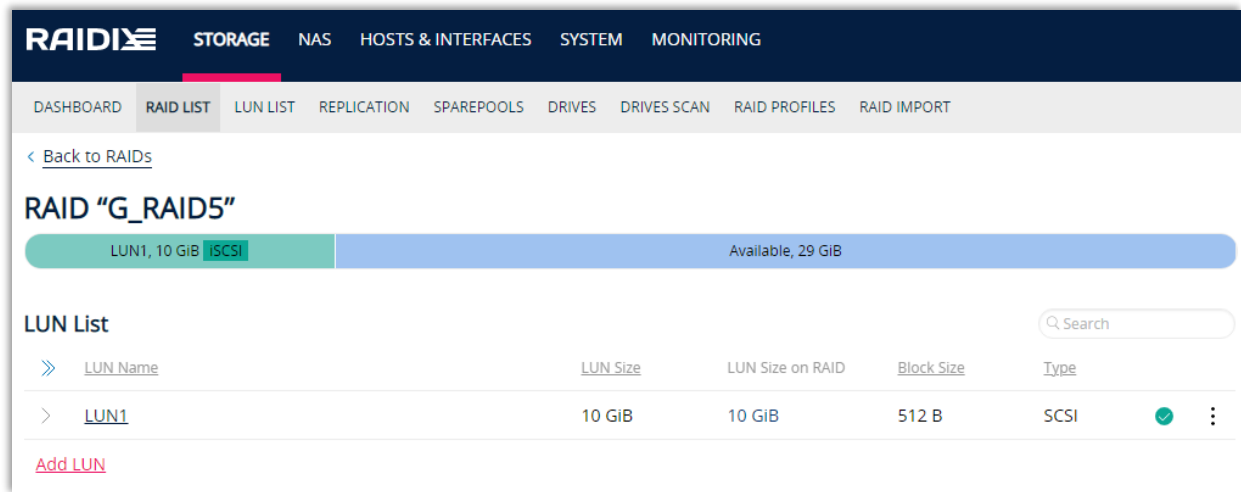


Figure 10. RAID page. LUN List section with created LUN

2.3 Setting up Host Connection

To set up a connection to a host:

1. Setup a network between a host and the system:
 - 1.1. Physically connect an initiator and a target.
 - 1.2. Additionally:
 - for InfiniBand, set up OpenSM;
 - for iSCSI, enable the iSCSI function and create an iSCSI target;
 - for iSER, enable iSCSI and iSER functions, and create an iSCSI target;
 - for FC, no need for additional setup.
2. Create a LUN, which will be available for a host as a block device:
 - 2.1. Create a LUN.
 - 2.2. Manage LUN access rules.

2.3.1 InfiniBand Setup

To set up the InfiniBand SRP connection:

1. Connect a host to the system.
2. Select **HOSTS & INTERFACES > ADAPTERS**.
3. Click the **OpenSM** button.
4. In the opened *Start OpenSM Service* window, select ports for connection with the host, and then click **SELECT** (Figure 11).

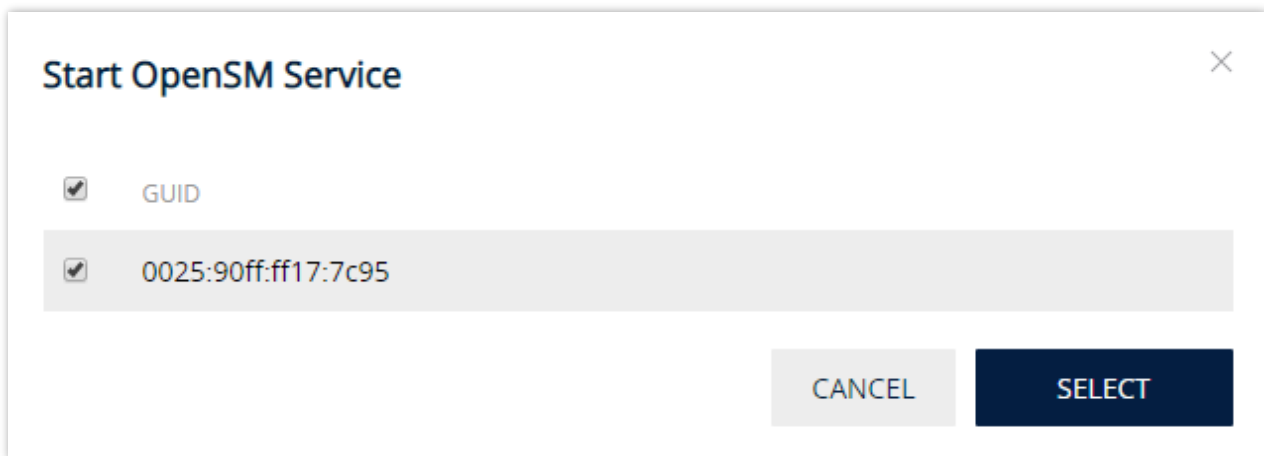


Figure 11. "Start OpenSM Service" window

! If the InfiniBand network has no running OpenSM service, for the DC system, start OpenSM on both nodes.

5. For a LUN with of the NAS type the setup is complete.

To set up LUN access, see [Masking rules](#).

2.3.2 iSER Setup

! When using iSER as a transport, the configuration in which a host has ConnectX-4 adapter and a target has ConnectX-3 adapter is not supported.
For correct work, the version of the adapter on the target must be the same or higher than the version of the adapter on the host.

To set up connection to a host via iSER:

1. On both the host and the target, switch the mode of the InfiniBand adapter to Ethernet.
To learn more about switching adapter modes, see the [Mellanox website](#).
2. Select **HOSTS & INTERFACES > ISCSI TARGETS**.
3. In the upper right of the page, on the pane, turn on iSCSI using the **Enable iSCSI** switch (Figure 12).
4. In the upper right of the page, on the pane, turn on iSER using the **Enable iSER** switch (Figure 12).

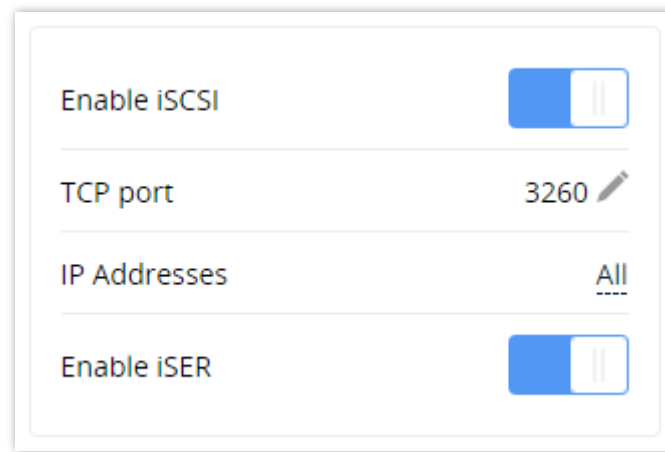


Figure 12. Pane on the iSCSI TARGETS page. “Enable iSCSI” and “Enable iSER” switches

- By default, iSER is enabled on all available network interfaces. If iSER is required only on particular interfaces, in the *IP Addresses* string, click **All**, then specify IP addresses of the interfaces.

i *IP address* and *TCP port* parameters are for the local node only, therefore when you change these settings in the DC mode, you must implement all the changes on the remote node.

- Click **Create Target**. The widget to create an iSER target shows (Figure 13).

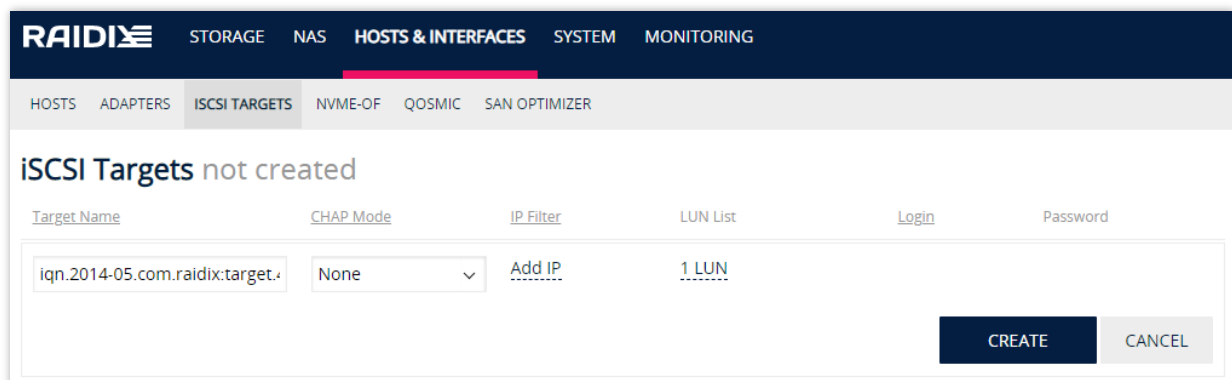
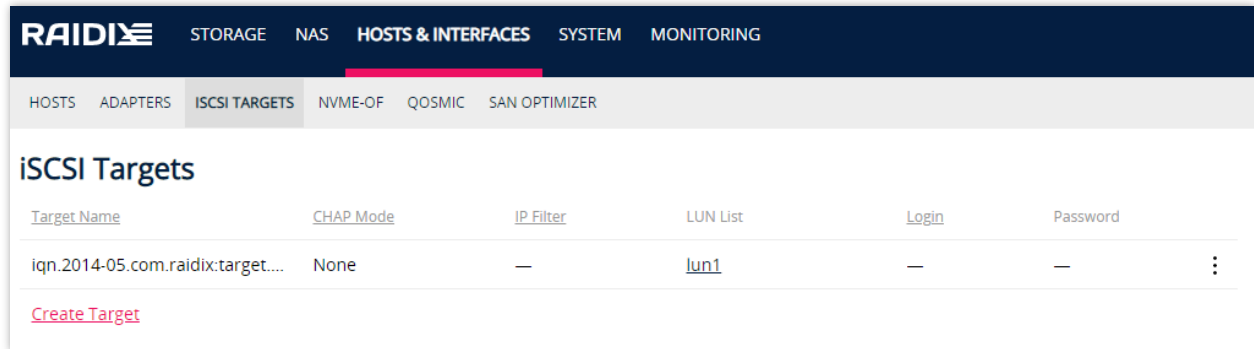


Figure 13. Widget for iSER target creation

- Define the following parameters and click **CREATE**:
 - Target name** – unique iSER ID (IQN);
 - CHAP mode** – specify the authorization mode used by host and target to authenticate each other via the exchange of a secret code or a password (to learn more, see [Create an iSCSI Target](#)):
 - None;
 - Unidirectional;
 - Bidirectional.
 - IP filter** – assign a host (or hosts) by adding its IP address to the masking rule;
 - LUN list** – select LUNs, which will be available for the target;

- **Login** – CHAP login. The field is editable if CHAP mode is *bidirectional*,
 - **Password** – CHAP password. The field is editable if CHAP mode is *bidirectional*.
8. Created target shows in the *iSCSI Targets* section (Figure 14).



RAIDIX STORAGE NAS HOSTS & INTERFACES SYSTEM MONITORING						
HOSTS ADAPTERS ISCSI TARGETS NVME-OF QOSMIC SAN OPTIMIZER						
iSCSI Targets						
Target Name	CHAP Mode	IP Filter	LUN List	Login	Password	
iqn.2014-05.com.raidix:target....	None	—	lun1	—	—	⋮
Create Target						

Figure 14. Created iSER target

i If you have created the iSER target with *Unidirectional* or *Bidirectional* CHAP mode, it is necessary [to create a CHAP user](#) by clicking the **CHAP Users**.

9. For a LUN of the NAS type the setup is complete.

To set up LUN access, see [Masking rules](#).

2.3.3 iSCSI Setup

To set up iSCSI connection:

- i**
- To set up the iSCSI connection for *Enterprise* policy, refer to steps 1-11.
 - To set up the iSCSI connection for *Standard* policy, refer to steps 1-3.

1. Physically connect an initiator to the storage system.
2. Select **HOSTS & INTERFACES > ISCSI TARGETS**.
3. In the upper right of the page, on the pane, turn on iSCSI using the **Enable iSCSI** switch (Figure 15).

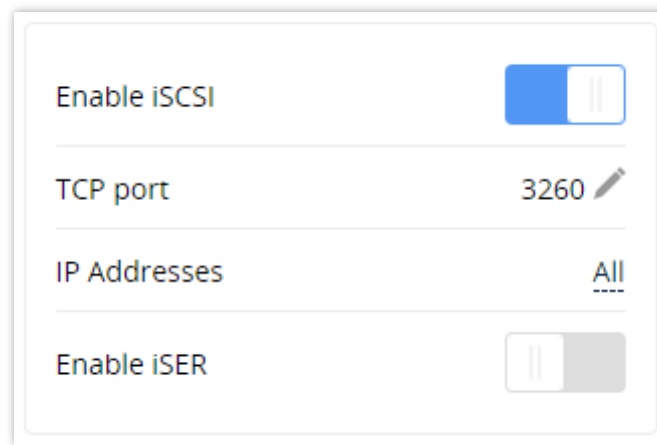


Figure 15. iSCSI TARGETS page. "Enable iSCSI" switch

4. By default, iSCSI enables on all available network interfaces. If iSCSI is required only on particular interfaces, in the *IP Addresses* string, click **All**, then specify IP addresses of the interfaces.

i *IP address* and *TCP port* parameters are for the local node only, therefore when you change these settings in the dual controller mode, you must implement all the changes on the remote node.

5. Click **Create Target**. The panel to create an iSCSI shows (Figure 16).

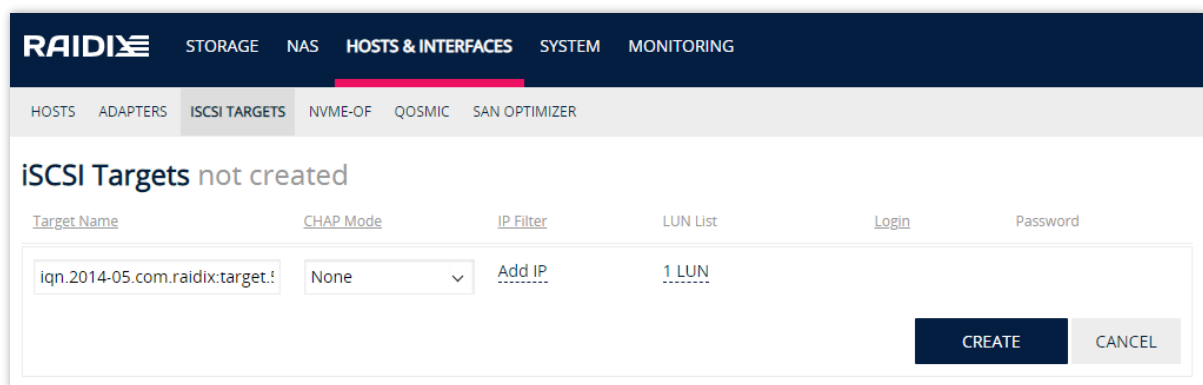
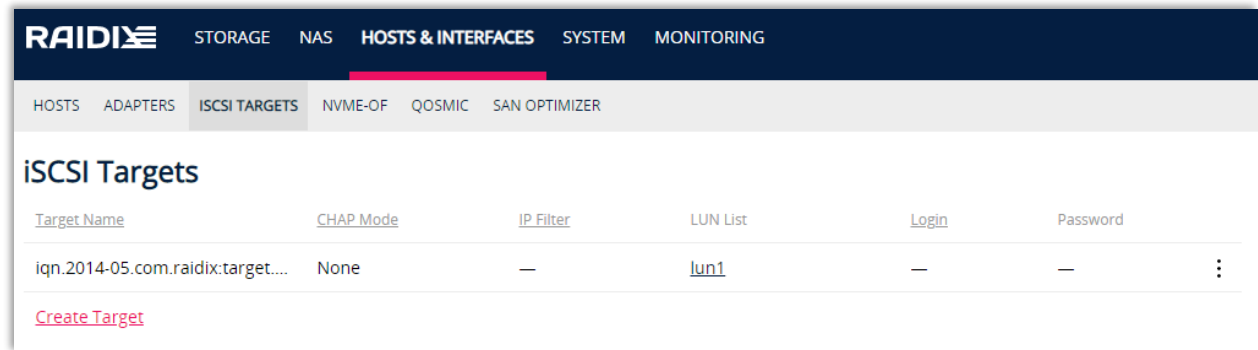


Figure 16. Widget for creation of a target

6. Define the following parameters:
- **Target name** – unique iSCSI ID (IQN);
 - **CHAP mode** – specify the authorization mode used by host and target to authenticate each other via the exchange of a secret code or password. To learn more, see [Create an iSCSI Target](#)):
 - None;
 - Unidirectional;
 - Bidirectional.
 - **IP filter** – assign a host (hosts) by adding its IP address to the masking rule;
 - **LUN list** – select LUNs, which will be available for the target;

- **Login** – CHAP login. The field is editable if CHAP mode is *bidirectional*.
 - **Password** – CHAP password. The field is editable if CHAP mode is *bidirectional*.
7. Click **CREATE**. Created target will show in the *iSCSI Targets* section (Figure 17).



RAIDIX STORAGE NAS HOSTS & INTERFACES SYSTEM MONITORING					
HOSTS ADAPTERS ISCSI TARGETS NVME-OF QOSMIC SAN OPTIMIZER					
Target Name	CHAP Mode	IP Filter	LUN List	Login	Password
iqn.2014-05.com.raidix:target....	None	—	lun1	—	—
Create Target					

Figure 17. Created iSCSI target



If you have created the iSCSI target with *Unidirectional* or *Bidirectional* CHAP mode, it is necessary [to create a CHAP user](#) by clicking the **CHAP Users** button.

To set up LUN access, see [Masking rules](#).

2.3.4 Fibre Channel Setup

For a LUN of the SCSI type, Fibre Channel connection establishes automatically, no additional settings required.

To set up LUN access, see [Masking rules](#).

2.4 NAS Share Creation

For simplified NAS creation, you may create a share with enabled guest access (available for SMB). In this case you may skip the step 2.

To create a NAS share:

1. Create a RAID with a LUN of the NAS type.
2. Select **NAS > USERS AND GROUPS**. Add NAS user(s). To learn more, see [Create Users and Groups](#).
3. Open the **NAS | SHARES** page.
4. Click **Format LUN**. In the opened window, select a file system and click **SELECT**.
5. In the *Shares List* section, click **Create Share**.
6. In the opened widget (Figure 18) specify share parameters (to learn more, see [Create a NAS share](#)).

For simplified NAS creation without additional access settings, set for the **Guest** parameter the value **Read / Write**.

Figure 18. Creating a share with NFS protocol

7. Click **CREATE**.

i When creating a share, a short-time disconnect of other shares may occur due to the reloading of the NAS services.

In the Shares List section, the created share with specified parameters displays (Figure 19).

Share Name	Path	Protocol	Guest	Users & Groups	Hosts
docs	/	NFS	—	—	0 Clients

Figure 19. Created NFS share

3. WORKING WITH RAID

This section describes [RAID levels](#), [RAID engines](#), and provides the information about consistency parameters.

See about the management of performance parameters in [Advanced RAID Parameters](#).

3.1 RAID Levels

There are several methods of combining hard drives called RAID levels. Each level has its pros and cons and offers a different balance of performance, data protection, and storage efficiency.

RAIDIX 5.2 enables you creating RAID levels: 0, 1, 5, 6, 7.3, 10, N+M, 50, 60, and 70.

RAID 0 – the level of interleaving blocks without mirroring. The data blocks are distributed across several drives. Data is in parallel access mode that provides high performance. Due to the lack of redundancy, RAID 0 doesn't provide data reliability – the failure of one drive in RAID leads to the whole RAID degradation.

RAID 1 – architecture of mirroring without parity or striping. The data is mirrored on all drives of the RAID, and the RAID can only be of the smallest member drive size. Random read performance of a RAID 1 may equal up to the sum of each member's performance, while the write performance remains at the level of a single slowest drive.

RAID 5 – the level of interleaving blocks with distributed parity. RAID 5 requires at least five drives. RAID 5 sustains the complete failure of *one* drive in the same group and provides a minimal degree of reliability.

RAID 6 – the level of interleaving blocks with double parity distribution. RAID 6 has improved performance, since each drive processes the I/O requests (entries) independently, allowing parallel access to the data. RAID 6 can sustain the complete failure of *two* drives in the same group. Redundant parity information provides additional time to restore redundancy without loss of information.

RAID 7.3 – the level of interleaving blocks with triple parity distribution. The RAID 6 analog, but has a higher degree of reliability: three checksums are calculated using different algorithms, the capacity of three drives is allocated for checksums. Thus, the RAID 7.3 can sustain the complete failure of *three* drives in the same group.

RAID N+M – the level of interleaving blocks with M checksums. RAID N+M allows a user to choose the number of drives for checksums allocation. RAID N+M requires at least eight drives and can sustain complete failure up to 32 drives in the same group (segments of which are RAID depending on the number of parity drives).



Speaking of RAID levels 10, 50, 60, and 70, a "group" is a RAID, which is the segment of the RAID 0 striping, and a "group size" is a number of drives in the group.

RAID 10 – the architecture of this mirrored RAID represents RAID 0, which components are RAID 1 instead of separate drives. Each RAID 1 consists of two drives, the minimum number of RAID 1 in a mirrored array is 2. Thus, in RAID 10 the minimum number of drives is 4. Data integrity is maintained in case of failure of half drives; the irreversible RAID destruction occurs when two drives of one mirrored pair already fail.

RAID 50 – RAID 0 striping combination across multiple RAID 5. With such a combination, RAID 50 may show better performance with reduced latency.

Requirements:

- At least 8 drives, the drive number must be a multiple of the group size.
For example, in case of 16 drives and the group size of 8, the RAID will be of 14 drives size from 2 groups of 8 drives; in case of the group size 4, the RAID will be of 12 drives size from 4 groups of 4 drives.
- The group size is at least 4 drives.
- RAID type: ERA.

Features:

- The amount of RAID 50 is calculated as $(N - N / gs) * size$, where N – total number of drives, gs – group size, size – the size of one drive.
- Recoverable from 1 failure in each stripe.
- RAID initialization is applicable.

RAID 60 – RAID 0 striping combination across multiple RAID 6. RAID 60 is the equivalent of RAID 50 with a higher level of fault tolerance.

Requirements:

- At least 8 drives, the drive number must be a multiple of the group size.
For example, in case of 24 drives and the group size of 12, the RAID will be of 20 drives size from 2 groups of 12 drives; in case of the group size 8, the RAID will be of 18 drives size from 3 groups of 8 drives.
- The group size is at least 4 drives.
- RAID type: ERA.

Features:

- The amount of RAID 60 is calculated as $(N - (2 * N / gs)) * size$, where N – total number of drives, gs – group size, size – the size of one drive.
- Recoverable from 2 failures in each stripe.
- RAID initialization is applicable.

RAID 70 – RAID 0 striping combination across multiple RAID 7. RAID 70 is the equivalent of RAID 60 with a higher level of fault tolerance.

Requirements:

- At least 8 drives, the drive number must be a multiple of the group size.

For example, in case of 32 drives and the group size of 16, the RAID will be of 26 drives size from 2 groups of 16 drives; in case of the group size 8, the RAID will be of 20 drives size from 4 groups of 8 drives.

- The group size is at least 4 drives.
- RAID type: ERA.

Features:

- The amount of RAID 70 is calculated as $(N - (3 * N / gs)) * size$, where N – total number of drives, gs – group size, size – the size of one drive.
- Recoverable from 3 failures in each stripe.
- RAID initialization is applicable.

Initialized RAIDs

The RAIDIX software enables you to create initialized RAIDs 5i, 6i, 7.3i, and N+Mi.

Features of an initialized RAID are:

- Ability to perform an additional setup of Random Writes parameter to increase the random operations rate.
- Opportunity to use the Silent Data Corruption (SDC) detection, which enables correcting hidden drive errors in real time.

3.2 RAID Engines

RAIDIX 5.2 has two different Engines for RAIDs: Generic and ERA, each of which has a different scope.

You can select a RAID Engine during the RAID creation.

Comparison of RAID Engine Functionality

Feature	Generic Engine	ERA Engine
NAS	Yes	Yes
Active Directory	Yes	Yes
NFS+LDAP	Yes	Yes
DC	Yes	Yes
RAID N+M	Yes	Yes
RAID 0, 1, 10, 5, 6, 7.3	Yes	Yes
RAID 50, 60, 70	No	Yes
Reconstruction per Drive	Yes	Yes
Replication	Yes	No
SAN Optimizer	Yes	No

Feature	Generic Engine	ERA Engine
QoSmic	Yes	Yes ⁱ
Automatic Write Through	Yes	No
SDC Correction	Yes	No
Partial Reconstruction	Yes	Yes
Advanced Reconstruction	Yes	No
Target Masking	Yes	Yes
Host Masking	Yes	Yes
SNMP	Yes	Yes
SMTP	Yes	Yes
Background Integrity Check	Yes	Yes
SparePool	Yes	Yes
LUN Extension	Yes	Yes ⁱⁱ
RAID Performance >300K IOPS	No	Yes
LUN Block Size 512 B (support of VMWare volumes)	Yes	Yes
SSD cache	Yes	No

3.2.1 Generic

We recommend choosing Generic Engine to create HDD RAIDs.

Features of Generic RAID

- Flexible settings for increased performance under different loads, including RAM cache settings.



To protect the information in the RAM cache from loss of power, we recommend using a UPS.

- c RAID levels are 0, 1, 10, 5, 6, 7.3, N+M, as well as initialized 5i, 6i, 7.3i, and N+Mi.
The initialization blocks I/O operations for the RAID.
- SDC detection and correction features.

3.2.2 ERA

We recommend choosing ERA Engine to create SSD RAIDs.

ⁱ For HDD RAIDs with the workload less than 100K IOPS

ⁱⁱ Within a single ERA RAID

Features of ERA RAID

- ERA algorithms use the potential of flash drives to increase SSD RAID performance.
 - Available interfaces: NVMe SSD, SAS SSD, SATA SSD.
- Available RAID levels: 0, 1, 5, 6, 10, 50, 60, 70, 7.3, and N+M.
- RAID initialization (except ERA RAID 0) starts automatically and works in background keeping the RAID available for I/O operations.
- Fast SSDs are accessed directly without the need for a RAM cache.

Limitations of ERA RAID

- ERA available only on processors that support AVX. ERA RAID N+M requires minimum AVX2.
- SDC detection and correction features are unavailable.
 - Integrity check.
 - Scheduled Background Integrity Check.
- The extension of LUN size is limited to the free size of one RAID (on which the LUN is located).
- The SATA interface for SSD is available only for SC mode.
- The RAID Profile feature is unavailable.



If you don't have the "Flash" option in your license, you can create ERA RAIDs, but only on HDDs.

Note the following features and limitations when using ERA RAIDs on HDDs (also applies to ERA RAIDs on SSDs):

- RAIDs 1, 50, 60, and 70 are available.
- No SSD cache.
- No RAM cache.
- No Advanced Reconstruction.

3.3 RAID Creation

In RAIDIX 5.2, you can create a RAID using:

- the *widget* **Create RAID** on the page **STORAGE | RAID LIST**;
- the *wizard* **CREATE RAID** on the page **STORAGE | DASHBOARD**;
- the *wizard* **CREATE LUN** on the page **STORAGE | DASHBOARD**.

To create a RAID using the *widget*:

1. Select **STORAGE > RAID LIST**.
2. In the *RAID List* section, click **Create RAID**.
3. In the opened widget (Figure 20), fill in the fields.

RAID List empty

RAID Name

RAID Size

Free Size

RAID Level

Drives

RAID Engine

SparePool

RAID Name

RAID Engine
Generic

RAID Level
0

RAM cache size
4 GiB

Stripe
128 KiB

Use Profile
default

Drives

Select

4 KiB blocks write mode

Write Barriers Support

CREATE


CANCEL


Figure 20. RAID creation widget. The system in DC mode

Parameter	Description
RAID Name	RAID name. <div> ! Don't name ERA RAID <i>power</i> or <i>uevent</i>. </div>
RAID Engine	RAID Engine (to learn more, see RAID Engines). Possible values: <ul style="list-style-type: none"> Generic – preferable for HDDs; ERA – preferable for SSDs. <div> ! ERA Engine works only on systems with CPUs that support AVX. ERA RAID N+M requires minimum AVX2. </div>

Parameter	Description																						
RAID Level	<p>RAID level.</p> <p>To learn more, see RAID Levels.</p> <p>Possible values:</p> <table> <thead> <tr> <th>Generic RAID</th><th>ERA RAID</th></tr> </thead> <tbody> <tr><td>• 0</td><td>• 0</td></tr> <tr><td>• 5</td><td>• 1</td></tr> <tr><td>• 5i</td><td>• 5</td></tr> <tr><td>• 6</td><td>• 6</td></tr> <tr><td>• 6i</td><td>• 7.3</td></tr> <tr><td>• 7.3</td><td>• 10</td></tr> <tr><td>• 7.3i</td><td>• 50</td></tr> <tr><td>• 10 (1ⁱⁱⁱ)</td><td>• 60</td></tr> <tr><td>• N+M</td><td>• 70</td></tr> <tr><td>• N+Mi</td><td>• N+M</td></tr> </tbody> </table> <div> <p>i For Generic RAID:</p> <p>If you create an <i>initialized</i> RAID, the initialization process starts; its progress displays in percentages.</p> <p>The working with initializing RAID is impossible until the initialization is complete (to learn more, see Generic RAID Initialization).</p> </div> <div> <p>i For ERA RAID:</p> <p>If you create an ERA RAID (except ERA RAID 0), the initialization process will start in background; its progress displays in percentages.</p> <p>You can work with initializing RAID and configuring the initialization priority to manage drives workload (to learn more, see ERA Raid Initialization).</p> </div>	Generic RAID	ERA RAID	• 0	• 0	• 5	• 1	• 5i	• 5	• 6	• 6	• 6i	• 7.3	• 7.3	• 10	• 7.3i	• 50	• 10 (1 ⁱⁱⁱ)	• 60	• N+M	• 70	• N+Mi	• N+M
Generic RAID	ERA RAID																						
• 0	• 0																						
• 5	• 1																						
• 5i	• 5																						
• 6	• 6																						
• 6i	• 7.3																						
• 7.3	• 10																						
• 7.3i	• 50																						
• 10 (1 ⁱⁱⁱ)	• 60																						
• N+M	• 70																						
• N+Mi	• N+M																						
Use Profile	<p>Only for Generic RAIDs. Optional.</p> <p>Select a saved Profile with RAID parameters.</p> <p>A Profile is a set of additional RAID parameters, which you can apply to a RAID during creation.</p> <p>If there are no saved Profiles, the system uses the default Profile. To learn more about managing RAID profiles, see the RAID Profiles chapter).</p>																						

ⁱⁱⁱ To create RAID 1, select RAID 10 with only 2 drives.

Parameter	Description
Drives	<p>Drives list.</p> <p>Add drives with the same sizes into a RAID.</p> <p>You can select drives in two ways:</p> <ul style="list-style-type: none"> Under the Drives field, click Select, tick required drives and confirm by clicking SELECT. Type numbers (UIDs) of drives separated by commas or a number range separated by a hyphen (<i>for example</i>: 0-3,5,7). <div>  To create a RAID, you can select both HDDs and SSDs. Note the specific SSD wear in RAIDs containing both SSDs and HDDs. </div>
RAM cache size	<p>Only for Generic RAIDs.</p> <p>Cache size (in GiB).</p> <p>The default is 4 GiB.</p> <p>We recommend selecting the following RAM size depending on the number of drives in the RAID:</p> <ul style="list-style-type: none"> up to 12 drives – from 4 GiB RAM; from 12 to 16 drives – from 8 GiB RAM; from 16 to 24 drives – from 12 GiB RAM; from 24 to 32 drives – from 16 GiB RAM; from 32 drives – from 24 GiB RAM.
4 KiB blocks write mode	<p>Only for Generic RAID 0 and RAID 10.</p> <p>Enable or disable 4 KiB blocks write mode.</p> <p>This mode might improve performance while working with complex write patterns.</p>
Write Barriers Support	<p>Only for Generic RAIDs.</p> <p>Saves data from the RAM write buffer to drives.</p> <p>It increases data integrity, but may decrease system performance.</p>
Block Size	<p>Only for ERA RAIDs.</p> <p>Sector partitioning format of drives.</p> <p>Possible values:</p> <ul style="list-style-type: none"> 512 B 4096 B (by default)
SparePool	<p>Not allowed for RAID 0. Optional.</p> <p>Name of a set of drives that you can assign for hot spare in case of RAID drive failure. To learn more, see the Automatic Drive Replacement chapter.</p>

Parameter	Description																
Stripe	<p>The minimum size of the data stripe (in KiB) to be recorded on drives in the RAID.</p> <p>Possible values:</p> <table> <thead> <tr> <th>Generic RAID</th><th>ERA RAID</th></tr> </thead> <tbody> <tr> <td>• 16 KiB</td><td>• Auto (by default)</td></tr> <tr> <td>• 32 KiB</td><td>• 16 KiB</td></tr> <tr> <td>• 64 KiB</td><td>• 32 KiB</td></tr> <tr> <td>• 128 KiB (by default)</td><td>• 64 KiB</td></tr> <tr> <td>• 256 KiB</td><td>• 128 KiB</td></tr> <tr> <td>• 512 KiB</td><td>• 256 KiB</td></tr> <tr> <td>• 1 MiB</td><td></td></tr> </tbody> </table> <p>The Auto value sets the stripe size depending on a drive type (HDD: 64 KiB, SSD: 32 KiB, NVMe: 16 KiB).</p>	Generic RAID	ERA RAID	• 16 KiB	• Auto (by default)	• 32 KiB	• 16 KiB	• 64 KiB	• 32 KiB	• 128 KiB (by default)	• 64 KiB	• 256 KiB	• 128 KiB	• 512 KiB	• 256 KiB	• 1 MiB	
Generic RAID	ERA RAID																
• 16 KiB	• Auto (by default)																
• 32 KiB	• 16 KiB																
• 64 KiB	• 32 KiB																
• 128 KiB (by default)	• 64 KiB																
• 256 KiB	• 128 KiB																
• 512 KiB	• 256 KiB																
• 1 MiB																	
Affinity/Status	<p>Only in DC mode.</p> <p>ID of the node, from which it will be possible to perform I/O on the RAID (node has preferable access to the RAID).</p> <p>Nodes identifiers are displayed on the STORAGE DASHBOARD page.</p> <div>  By default, Affinity corresponds to the node ID, on which a RAID was created. You can also change Affinity value by performing RAID Migration. To learn more, see Migrate a RAID. </div>																
Memory limit	<p>Only for ERA RAIDs but RAID 0.</p> <p>RAID RAM memory limit, in MiB. The Limit determines the available amount of RAM for the RAID.</p> <p>Possible values: integers from 1024 to 1048576.</p> <p>If the limit value is not specified, the value sets automatically based on Stripe size:</p> <ul style="list-style-type: none"> • 16 KiB stripe – 2048 MiB RAM; • 32 KiB stripe – 2048 MiB RAM; • 64 KiB stripe – 4096 MiB RAM; • 128 KiB stripe – 8192 MiB RAM; • 256 KiB stripe – 16384 MiB RAM. 																

- Click **CREATE**. The created RAID (Figure 21) shows in the *RAID List* section.

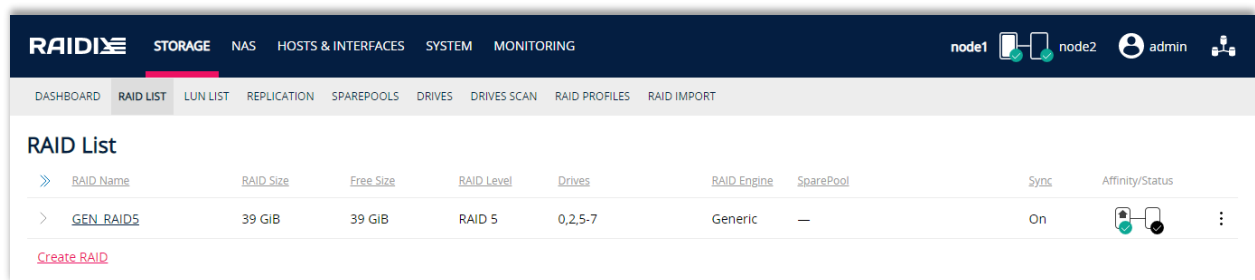


Figure 21. RAID LIST page. Created Generic RAID


In the line with a RAID object, additional values show:

RAID Size	RAID size (in GiB).
Free Size	Free space on RAID (not occupied by LUNs) (GiB).
Sync	<p>Only in DC mode for Generic RAID.</p> <p>Cache synchronization status.</p> <p>For ERA RAID, synchronization is inapplicable, so the value is “—”.</p> <p>The possible values:</p> <ul style="list-style-type: none"> • <i>On</i> – caches of the specified RAID on both nodes are synchronized; • <i>Off</i> – caches of the specified RAID on both nodes are not synchronized; • <i>Error</i> – an error occurred during the RAID synchronization.


i If synchronization is broken, Write Through (to learn more about write through, see the [Advanced RAID Parameters](#) chapter) will be enabled on RAID. Working in Write Through mode ensures better data integrity but might negatively affect performance.



3.4 RAID States

After creating a RAID, the current RAID status shows:

- in DC mode – in the RAID line in the *Affinity/Status* field;
- in SC mode – in the RAID line to the left of .

Possible RAID statuses

	Online shows in the following cases:
<i>Active</i>	The RAID is active on the node, specified in its Affinity. All RAID drives work correctly.
<i>Initializing: finished</i>	The RAID initialization completed successfully.
<i>Passive</i>	<p>Only in DC system.</p> <p>The RAID is passive on the node, which is not specified in its Affinity. All RAID drives work correctly.</p>

<i>Hold</i>	<p>Only for ERA RAIDs in DC system.</p> <p>The ERA RAID is passive on the node that is not specified in its Affinity. All RAID drives work correctly.</p> <p>The node with the passive ERA RAID provides limited information about that RAID.</p>
<p> Warning shows in the following states:</p>	
<i>Degraded</i>	A RAID is available and operating, but some of the drives are removed from the system or faulty.
<i>Failed over</i>	<p>Only in DC system.</p> <p>A RAID that was active on the node, specified in its Affinity, was failed over (transferred) to the second node. As a result, the RAID became temporarily active on the second node and passive on the first.</p>
<i>Initializing</i>	A RAID is being initialized. During the initialization, you cannot work with the Generic RAID, but you can work with ERA RAID. The completion status is displayed in percent.
<i>Passive</i>	<p>Only in DC system.</p> <p>A RAID is passive on the node specified in its Affinity.</p>
<i>Reconstructing</i>	One or several drives are currently reconstructing, but you can operate with the RAID.
<i>SDC</i>	<p>A RAID contains corrupted data.</p> <p>While a RAID has the SDC status, an error will occur during the reading from the damaged stripe.</p>
<i>Transitioning</i>	A RAID is switching to the passive mode.
<p> Bad shows in the following states:</p>	
<i>Offline</i>	<p>A RAID is offline.</p> <p>The RAID lacks too many drives and is not available.</p>
<i>Not Loaded</i>	<p>A RAID configuration exists, but fails to be loaded.</p> <p>To load the RAID configuration, use one of the following ways described in Troubleshooting.</p>
<i>Caches are not synchronized</i>	<p>Only in DC system.</p> <p>A synchronization network channel between nodes does not exist.</p>

3.5 RAID Initialization

Features of an initialized RAID are:

- Ability to perform an additional setup of the **Random Writes** parameter to increase the random operations rate.
- Opportunity to use the **Silent Data Corruption (SDC)** detection, which enables correcting hidden drive errors in real time.

To learn about Random Writes and SDC functions, see [Advanced RAID Parameters](#).

Depending on RAID Engine, Initialization differs:

- Generic RAID Initialization:
 - The Initialization starts automatically after a RAID 5i, 6i, 7.3i, or N+Mi is created.
 - During the initialization process, RAID is disabled for I/O.
 - To [stop](#) and [start](#) initialization process, use CLI commands.
- ERA RAID Initialization:
 - The Initialization starts automatically after you create ERA RAID (except RAID 0).
 - You can work with the RAID during the Initialization.
 - The default value of Initialization priority is **5%**.
 - In order to control the workload on RAID drives and the utilization of system resources, manage the **Initialization priority** parameter (Figure 22).

Setting the parameter to **0** provides the lowest priority initialization providing computing resources for other system processes, while setting the parameter to **100** provides the initialization process the highest priority among the processes.

RAID Size	Free Size	Level	Drives	RAID Engine	SparePool	Stripe	Sync	Affinity/Status
837 GiB	837 GiB	RAID 5	5-8	ERA	—	16 KiB		

Dynamic Parameters

Enable merge optimization	<input type="checkbox"/>	?
Low thread optimization	<input type="checkbox"/>	?
Reconstruction Priority	5	?
Initialization priority	5	?
Memory limit	2048	?

Figure 22. "Advanced RAID Parameters" window. The Initialization Priority parameter

You can check initialization progress:

- in the upper-left corner of GUI, by hovering a node name or icon (Figure 23);
- on the **STORAGE | RAID LIST** page, by hovering the *Affinity/Status* icon (Figure 24).

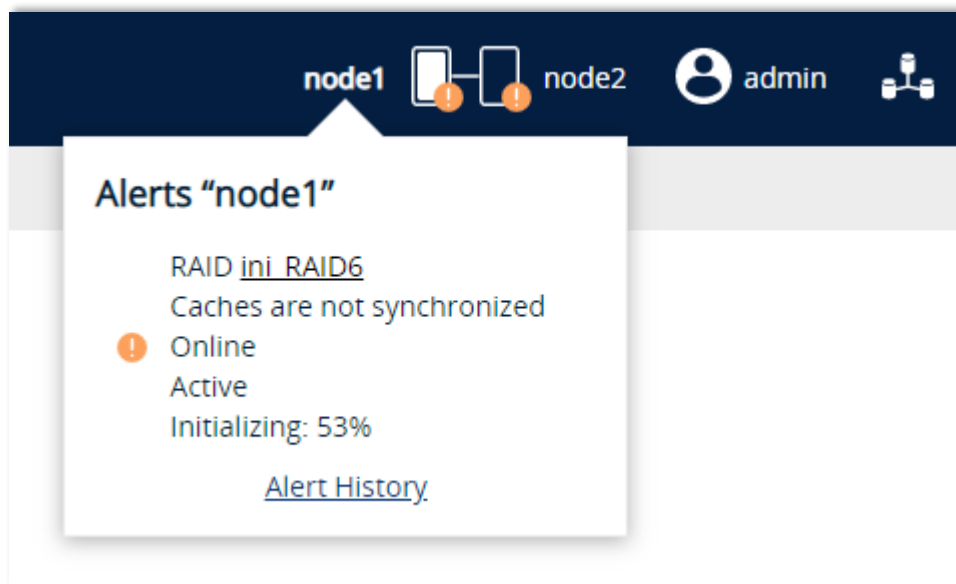


Figure 23. Progress of the Initialization process in the GUI menu

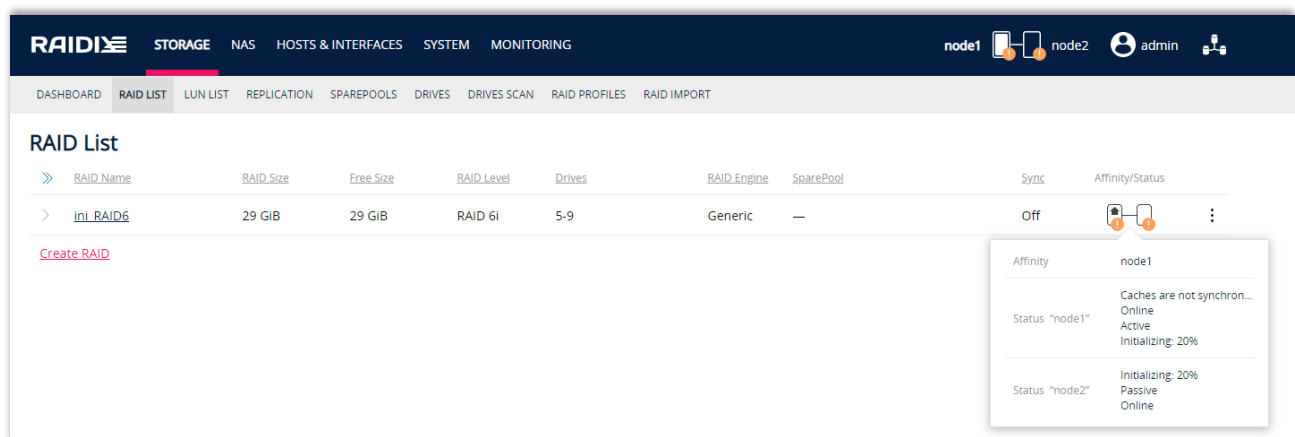


Figure 24. Progress of the Initialization process on the RAID List page

3.6 Scheduled Background Integrity Check

Scheduled Background Integrity Check (in this chapter, called SBIC) runs automatically in accordance with a schedule set by an administrator. SBIC enables the system to find and correct data corruptions.

3.6.1 Features of Background Integrity Check

Before configuring Scheduled Background Integrity Check, note the following specifics of this feature:


- SBIC doesn't block the system while working.
- You can set a separate schedule of SBIC for each RAID.
- Before enabling SBIC, set its schedule. The schedule defines periods when SBIC is running and periods during which it suspends. If SBIC is enabled and periods of suspending are not specified, then the SBIC will run continuously.

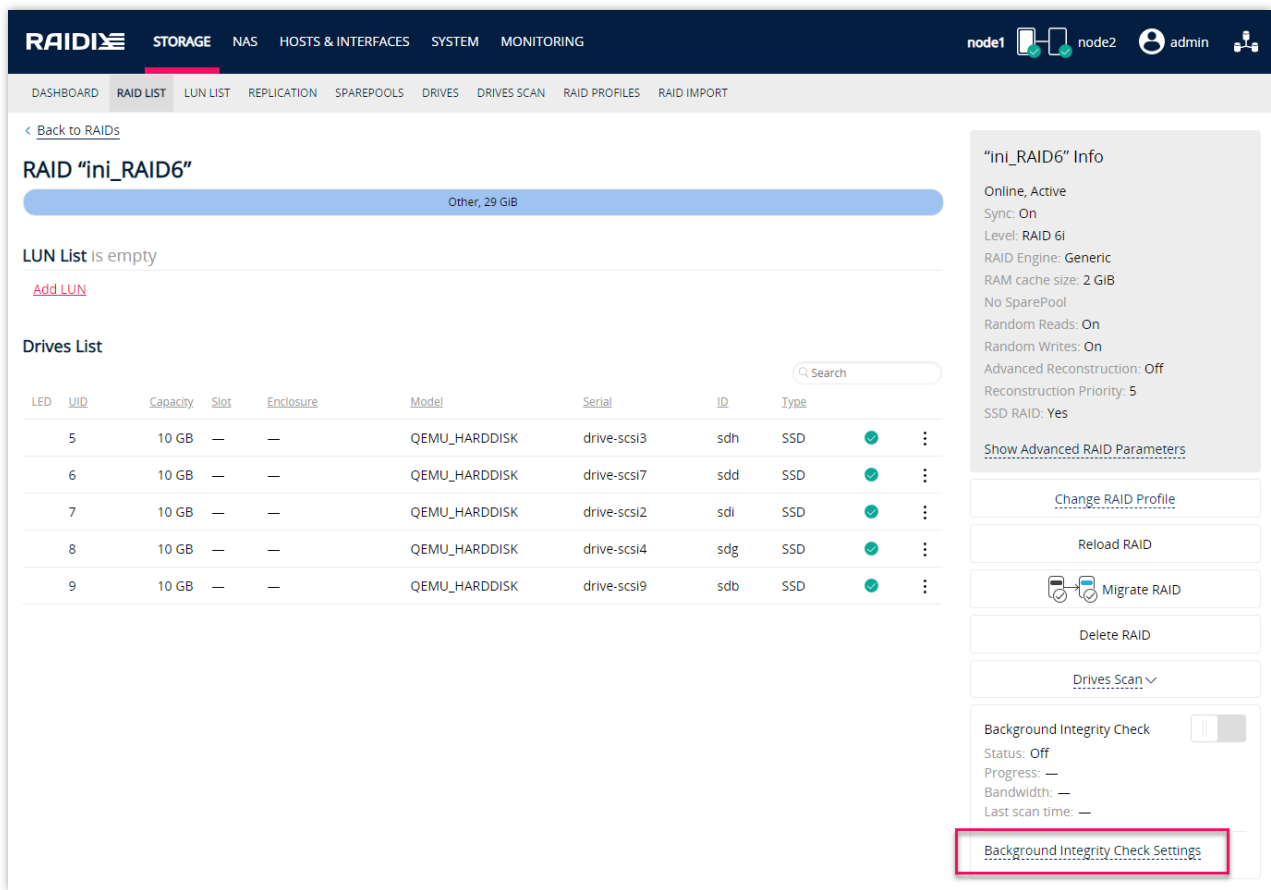
- During periods when SBIC is running, a new check starts in 30 seconds after the previous was complete. SBIC is running until the period of suspension starts. If the SBIC wasn't complete until then, it restarts in the next period.
- On initialized generic RAIDs, SBIC runs in the Silent Data Corruption (SDC) detection mode. The system corrects the error after the detection of silent data corruption.
- If during SBIC of ERA RAIDs or non-initialized Generic RAIDs the system detects a drive with faults over [Threshold fault count](#), the drive excludes from the RAID.
- If a RAID is reloaded during SBIC, the check suspends and will resume after the RAID is back online.
- If a RAID goes *Offline* during SBIC, the check suspends and will resume after the RAID is back *Online*.
- When deleting a RAID, its SBIC stops. The SBIC settings and schedule will be deleted along with the RAID.

3.6.2 Configuration of Background Integrity Check

You can manage SBIC from the *Background Integrity Check Settings* window.

To open the *Background Integrity Check Settings* window (Figure 27), use any of the following ways:

- On the RAID page, click **Background Integrity Check Settings** (Figure 25).
- On the **STORAGE | RAID LIST** page, click  in the line of the corresponding RAID, then click **Background Integrity Check Settings** (Figure 26).



RAID "ini_RAID6"

Other, 29 GiB

LUN List is empty

[Add LUN](#)

Drives List

LED	UUID	Capacity	Slot	Enclosure	Model	Serial	ID	Type	
	5	10 GB	—	—	QEMU_HARDDISK	drive-scsi3	sdh	SSD	✓
	6	10 GB	—	—	QEMU_HARDDISK	drive-scsi7	sdd	SSD	✓
	7	10 GB	—	—	QEMU_HARDDISK	drive-scsi2	sdl	SSD	✓
	8	10 GB	—	—	QEMU_HARDDISK	drive-scsi4	sdg	SSD	✓
	9	10 GB	—	—	QEMU_HARDDISK	drive-scsi9	sdb	SSD	✓

"ini_RAID6" Info

Online, Active
 Sync: On
 Level: RAID 6i
 RAID Engine: Generic
 RAM cache size: 2 GiB
 No SparePool
 Random Reads: On
 Random Writes: On
 Advanced Reconstruction: Off
 Reconstruction Priority: 5
 SSD RAID: Yes

[Show Advanced RAID Parameters](#)

[Change RAID Profile](#)

[Reload RAID](#)

[Migrate RAID](#)

[Delete RAID](#)

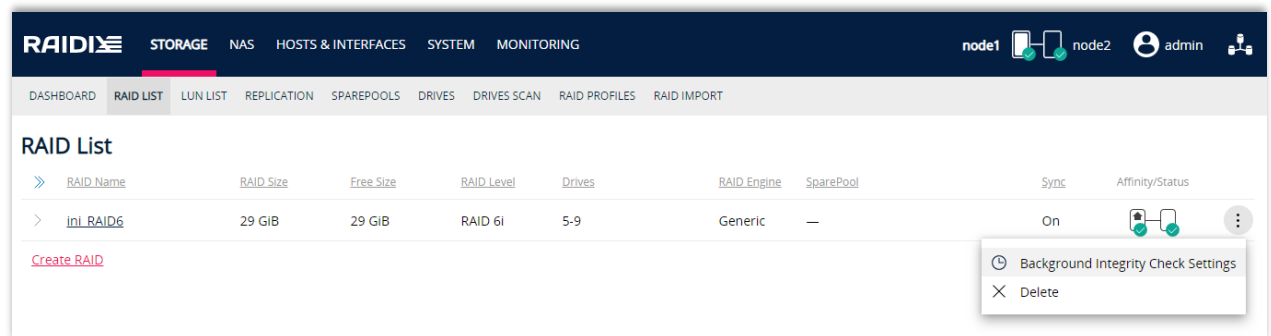
[Drives Scan](#)

Background Integrity Check

Status: Off
 Progress: —
 Bandwidth: —
 Last scan time: —

[Background Integrity Check Settings](#)

Figure 25. Background Integrity Check Settings on the RAID page



RAID List

RAID Name	RAID Size	Free Size	RAID Level	Drives	RAID Engine	SparePool	Sync	Affinity/Status
ini_RAID6	29 GiB	29 GiB	RAID 6i	5-9	Generic	—	On	✓

[Create RAID](#)

Background Integrity Check Settings
 Delete

Figure 26. Background Integrity Check Settings on the RAID LIST page

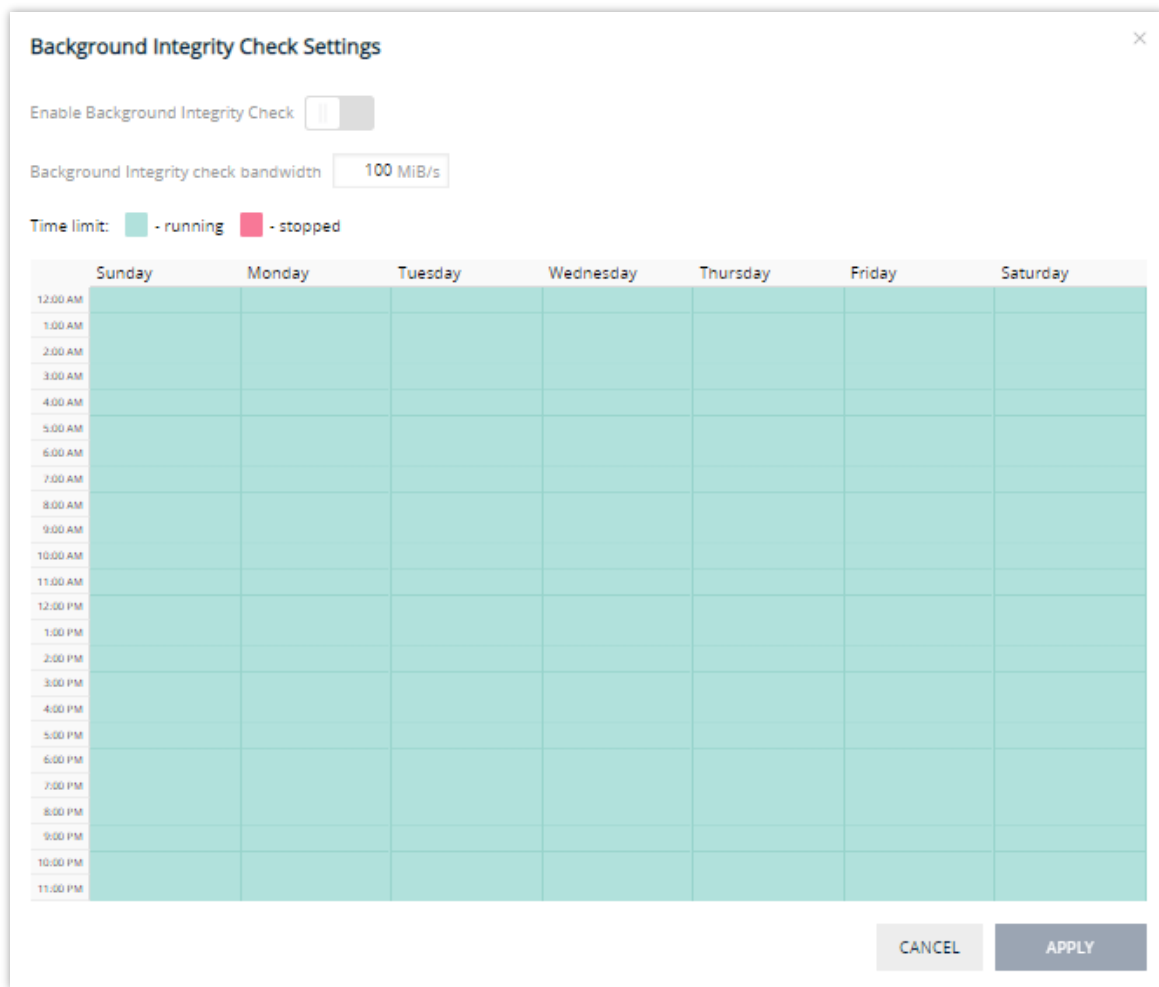



Figure 27. "Background Integrity Check Settings" window

To configure SBIC:

- In the *Background Integrity Check Settings* window, configure the following settings:
 - Enable **Background Integrity Check** by using the corresponding switch.
 - In the field **Background Integrity Check Bandwidth**, specify the scanning rate, in MiB/s.
The default value is **100 MiB/s**.
 - Configure the schedule of checks. Periods of time, during which background integrity check runs, are shown with green. New check starts in 30 seconds after the previous check was complete.

To add suspension periods:

- In the green area, select a suspension interval with the pointer. The selected interval will change its color to pink and the selected time range will be shown in the upper-right corner of the interval.
 - A click on the green area creates a 1-hour suspension interval.
 - A suspension interval can be changed by stretching it. Time step is 15 minutes.
 - You can also delete a suspension interval by clicking .
- After the schedule is set, click **APPLY** (Figure 28).

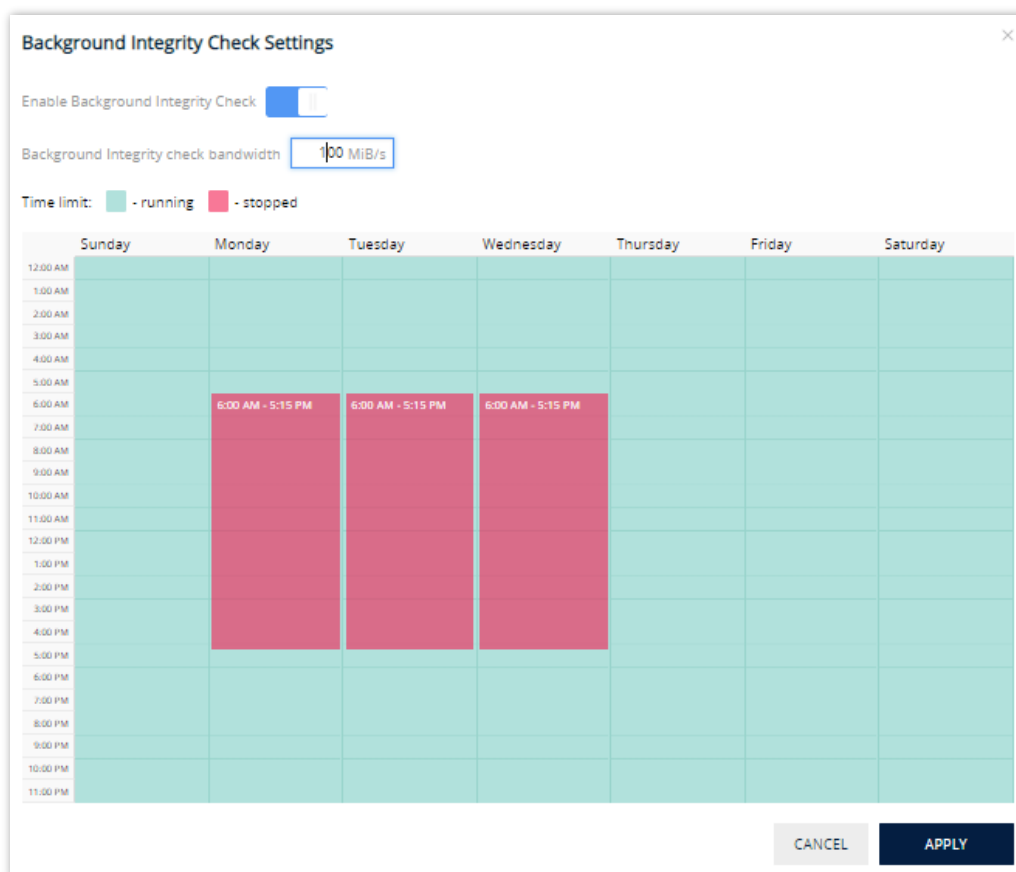




Figure 28. “Background Integrity Check Settings” window. Suspension periods

After background integrity check is configured, the icon  displays in the line of configured RAID on the **RAID LIST** page.

Background integrity check will start when the first interval when the check is allowed starts.

Status of the Integrity Check

The state and progress of Background Integrity Check is shown:

- In the *Background Integrity Check* pane on the RAID page (Figure 29).
Information in the widget updates every 30 seconds.
- In the **STORAGE > RAID LIST** page by hovering the pointer over the  icon in a RAID line (Figure 30).

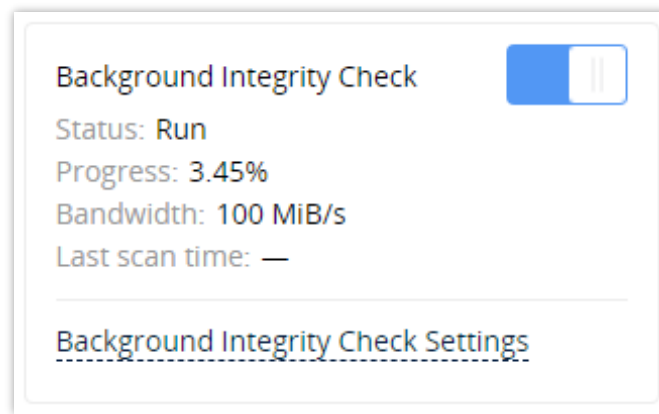


Figure 29. Background Integrity Check state on the RAID page

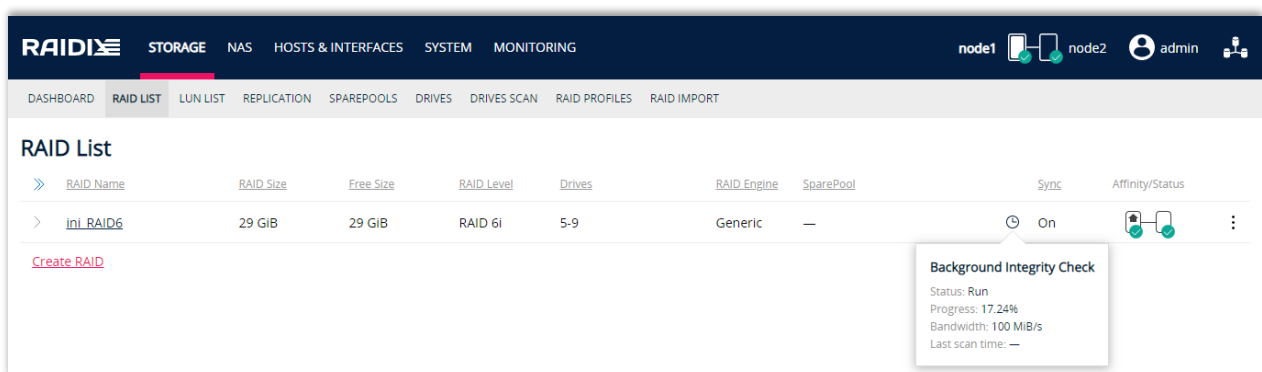


Figure 30. Background Integrity Check state on the RAID List page

3.6.3 Background Integrity Check in DC

For DC systems, SBIC has the following specific features:

- You can configure SBIC on any node in accordance with the section [Configure Background Integrity Check](#).
- If Failover occurred during SBIC, the SBIC suspends and will continue after Failback.
 - If Failback occurs during the suspension interval, the SBIC will continue in the next allowed period.
- If during SBIC a RAID was migrated to the second node, the SBIC will continue on the node that has become Active when the Migration is complete.
- If the active node reloads while SBIC is running, the SBIC will suspend and will continue after the system is booted and Failbacked.

3.7 RAID Reload

i Only for Generic RAID.

! Do not reload RAID's while the system is under a load.

RAID Reload is performed for:

- applying cache size and static RAID parameters changes;
- if the RAID status has changed to *Not loaded*, the RAID reload can help to bring it back to a normal state. To learn more, see [Troubleshooting](#).

To reload a RAID:

1. Select **STORAGE > RAID LIST**.
2. Select a RAID.
3. On the opened page, click **Reload RAID** (Figure 31).
4. To confirm, click **RELOAD**.

The screenshot shows the RAIDIX web interface. The top navigation bar includes 'STORAGE', 'NAS', 'HOSTS & INTERFACES', 'SYSTEM', and 'MONITORING'. The 'STORAGE' section is active, showing 'RAID LIST', 'LUN LIST', 'REPLICATION', 'SPAREPOOLS', 'DRIVES', 'DRIVES SCAN', 'RAID PROFILES', and 'RAID IMPORT'. The main content area displays 'RAID "ini_RAID6"' with a status bar indicating 'Other, 29 GiB'. Below this, the 'LUN List' is empty, and the 'Drives List' contains five entries (LED 5-9) with columns for Capacity, Slot, Enclosure, Model, Serial, ID, and Type. On the right, the 'ini_RAID6 Info' panel shows details like 'Online, Active', 'Sync: On', 'Level: RAID 6i', 'RAID Engine: Generic', 'RAM cache size: 2 GiB', 'No SparePool', 'Random Reads: On', 'Random Writes: On', 'Advanced Reconstruction: Off', 'Reconstruction Priority: 5', and 'SSD RAID: Yes'. At the bottom of this panel, the 'Reload RAID' button is highlighted with a red rectangle. Other buttons include 'Change RAID Profile', 'Migrate RAID', 'Delete RAID', and 'Drives Scan'.

Figure 31. Reload RAID button on the RAID page

3.8 RAID Migration

Migration of a RAID is changing the RAID Affinity parameter (Active or Passive) for a node in the DC system.

Use the Migration functionality to manage workload balancing between DC system nodes.

You can perform Migration of a RAID on any node and for any RAID.

To migrate a RAID:

1. Select **STORAGE > RAID LIST**.
2. Select a RAID you want to migrate.
3. On the opened page, on the right, click **Migrate RAID**. To confirm, click **PERFORM**.

3.9 RAID Restripping

i Only for ERA RAID.

Restripping allows the next modifications of created RAIDs:

- Change a RAID level.
- Increase RAID size by adding new drives.

Requirements and options:

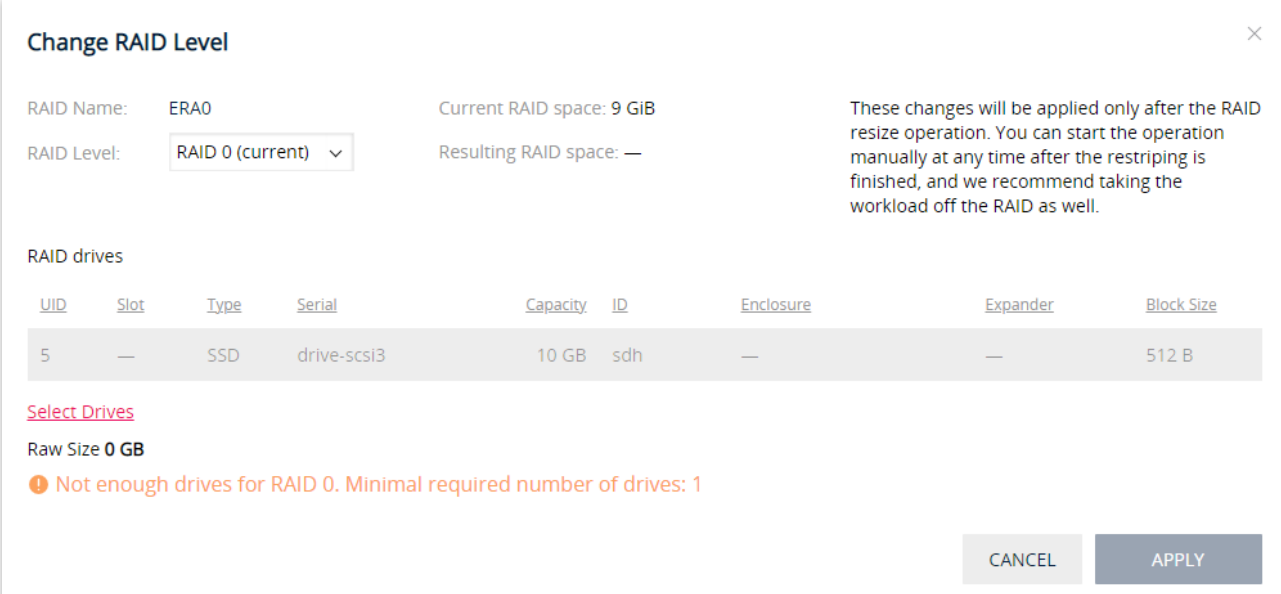
- Initialization of a RAID must be finished.
- Only one RAID can be restriped at a time.
- While restripping, reconstruction is unavailable.
- RAID state mustn't be "Degraded" and "Need recon".
- Restripping is performed in real time.

The available options for RAID level changes and the minimum required number of drives

RAID level change	Minimal number of drives you should add
RAID 0 to RAID 1	<p>i RAID 0 should contain only 1 drive.</p> <p>1</p>
RAID 0 to RAID 10	<ul style="list-style-type: none"> • If a RAID 0 contains only 1 drive: 3 • If a RAID 0 contains more than 1 drive: The number of drives to be added must be equal to the number of drives in RAID 0.
RAID 1 to RAID 10	2
RAID 1 to RAID 5	2
RAID 10 to RAID 5	2
RAID 5 to RAID 6	1
RAID 6 to RAID 7.3	1

To change a RAID level or increase a RAID size:

1. Restripe a RAID:
 - 1.1. Select **STORAGE > RAID LIST**.
 - 1.2. Select a RAID.
 - 1.3. Click **Change RAID level** at the right of the page.
 - 1.4. Depending on a value in the **RAID Level** field (Figure 32):
 - To increase RAID size, keep a value unchanged (current).
 - To change a RAID level, select a new value.



The 'Change RAID Level' window displays the current RAID configuration and options for changing the level. It includes fields for RAID Name, RAID Level, Current RAID space, and Resulting RAID space. A table lists the RAID drives with columns for UID, Slot, Type, Serial, Capacity, ID, Enclosure, Expander, and Block Size. A warning message indicates that there are not enough drives for RAID 0.

Change RAID Level [X]

RAID Name: ERA0 Current RAID space: 9 GiB

RAID Level: RAID 0 (current) ▼ Resulting RAID space: —

These changes will be applied only after the RAID resize operation. You can start the operation manually at any time after the restriping is finished, and we recommend taking the workload off the RAID as well.

RAID drives

UID	Slot	Type	Serial	Capacity	ID	Enclosure	Expander	Block Size
5	—	SSD	drive-scsi3	10 GB	sdh	—	—	512 B

[Select Drives](#)

Raw Size 0 GB

⚠ Not enough drives for RAID 0. Minimal required number of drives: 1

CANCEL APPLY

Figure 32. Change RAID Level window

- 1.5. Add drives.
- 1.6. Click **APPLY**.
- 1.7. Wait until restriping done.

Restriping status displays at the right of the RAID page on the *Change RAID Level* pane (Figure 33).

You can pause and resume restriping using the corresponding button on the Restriping status pane or the context menu in the RAID line on the **RAID LIST** page.

"ERA0" Info

Restripping, Active

Sync:

Level: RAID 1

RAID Engine: ERA

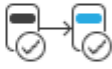
Block Size: 4,096 B

No SparePool

Reconstruction Priority: 5

SSD RAID: Yes

[Show Advanced RAID Parameters](#)

 Migrate RAID

Delete RAID

[Drives Scan](#)

Change RAID Level

Status: Restripping

Progress: 1%


Level: RAID 0 > RAID 1

Size: 9 GiB > 10 GiB

[Stop changing the RAID level](#)

Figure 33. Restripping status pane

2. Perform RAID resize:

 We recommend to stop workload on the RAID for the duration of resizing.

2.1. On the RAID page, click **Start resizing**.

2.2. To confirm, click **RESIZE**.

Wait for RAID initialization to complete. The initialization progress is displayed in the RAID line in the **RAID LIST** window and in system alerts.

3.10 RAID Deletion

To delete a RAID:

1. Select **STORAGE > RAID LIST**.
2. Select a RAID. The RAID page will open (Figure 31).
3. Click **Delete RAID**.
4. Click **DELETE** to confirm.



If you try to delete a RAID containing an SSD-cached LUN or a caching LUN, the deletion might take a long time due to flushing of cached data.

There is an option of quick deletion of such RAIDs without cache flushing. However, data integrity is not guaranteed because the deleted data will be lost.

4. ADVANCED RAID PARAMETERS

Advanced RAID Parameters influence a performance and a fault tolerance. You can set up Advanced Parameters during RAID creation (see RAID Profiles) and manage the parameters of an already created RAID.

Some Advanced Parameters require reboot of the RAID after applying. Those parameters called *static*. Advanced parameters that don't require a reboot called *dynamic*.

Parameters of a Generic RAID can be set and saved as a RAID profile on the page of the appropriate RAID and on the **RAID PROFILES** page, so that it can be applied to any RAID group later (refer to [Working with RAID Profiles](#) for details). RAID profile function is inapplicable to ERA RAID.

Advanced Parameters depend on the level and Engine include:



- cache setup;
- functions that enhance random I/O performance;
- reconstruction setup;
- SDC detection;
- optimization of a small-block sequential write with a large write depth;
- optimization of a load with a low number of threads.

Parameters for Generic RAID

Parameter	Description
Static Parameters	
4 KiB blocks write mode	Only for RAID 0 and RAID 10. Write mode with the minimum block size of 4 KiB.
Alt Buffers Size	Size of alternative cache buffer (MiB). Possible values: from 200 to 4096 . The default: 512 .
Write Back at Once	Maximum size of a single write back request, in stripes. The parameter enables you to limit amount of cache used for each write command from a host. Possible values: from 1 to 128 . The default: 24 .
Max Write Data	Maximum size of write buffer, in MiB. The buffer contains the blocks currently written to the RAID. Specify larger values if the block size is large or to write from many hosts simultaneously. Possible values: from 64 to 2048 . The default: 256 .
Number of Threads	Number of processing threads. Possible values: from 0 to 256 . The default: 16 .

Parameter	Description
Verify Cache Sum	Only in DC mode. Enabling comparison cache checksums during the synchronization process.
Reconstruct at Once	Maximum size of a reconstruction request (MiB). Possible values: from 1 to 64 . The default: 24 .
Dynamic Parameters	
Clean Segments Replacement	The parameter determines the mechanism for clean segments replacement from the cache into the corresponding areas of the main memory. <i>Clean cache segments</i> contain data that has been already flushed to the drives. In RAIDIX 5.2, clean segments are replaced with LRU (Least Recently Used) algorithm. LRU algorithm implies that the least recently used segments are flushed on the drives at first.
Dirty Segments Replacement	The parameter determines the mechanism for dirty segments replacement from the cache into the corresponding areas of the main memory. <i>Dirty cache segments</i> contain the data not flushed to the drives yet. The following replacement algorithms available: <ul style="list-style-type: none"> • LRU (Last Recently Used) –algorithm shows the best efficiency during <i>sequential</i> workload; • Cycle – algorithm shows the best efficiency during <i>random</i> workload.
Max Dirty Segments for LRU	Maximum number of dirty segments to be flushed at once on the drives with LRU algorithm. Possible values: from 1 to 128 . The default: 24 .
Max Dirty Segments for Cycle	Maximum number of dirty segments to be flushed at once on the drives with Cycle algorithm. Possible values: from 1 to 128 . The default: 24 .
Max Write Back Sequential	Maximum number of concurrent sequential requests to be written back. Possible values: from 1 to 1024 . The default value 8 is optimal for most cases.
Max Write Back Random	Maximum number of random requests to write back. Possible values: from 1 to 1024 . The default: 1024 .
Write Barriers Support	Saves data from the RAM write buffer to Drives. It increases data integrity, but may decrease system performance. The default: Enabled .
Random Reads	Option allows optimizing operation with RAID for random read. The default: Enabled .
Random Writes	Only for initialized RAIDs. Option allows optimizing the operation with RAID for random write. The default: Enabled .

Parameter	Description
Silent Data Corruption	<p>Only for initialized RAID.</p> <p>Option enables system detecting and recovering corrupted RAID data.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • Detection – drive errors detection only is running; • Correction – errors detection and, if possible, data recovery; • Disabled – the SDC feature is turned off (set by default). <div> <p>i Silent Data Corruption suspends the following functions:</p> <ul style="list-style-type: none"> • Random Reads; • Random Writes; • Advanced Reconstruction. </div> <div> <p>! While a RAID is in SDC state, there will be an error when trying to read from the corrupted stripe.</p> </div>
Write Through	<p>Write Through mechanism means that the data is written to the cache and the RAID, and the write finishes when both have finished.</p> <p>In RAIDIX, writes update the cache and then the RAID.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • Enabled – data is written to the drives. Once the data is written to the drives, initiator receives the confirmation report. • Disabled – Write Through is disabled, and Write Back mechanism is used instead. <p>A write-back scheme handles writes by updating values only to the block in the cache, then writing the modified block to the RAID when the block is replaced.</p> <ul style="list-style-type: none"> • Auto – this option enables the Automatic Write Through feature. The feature automatically selects between Write Back or Write Through modes based on information from sequence detector. <ul style="list-style-type: none"> ○ If the pattern of write is determined as random, Write Back mechanism is used. ○ If the pattern of write is determined as sequential, Write Through mechanism is used. <div> <p>i In DC mode, you can enable Automatic Write Through only on a RAID with at least one LUN.</p> </div> <div> <p>i Using Automatic Write Through may improve performance for mixed write pattern in DC mode in case synchronization is a bottleneck.</p> <p>The settings of Write Through parameter specified for the RAID are applied for all LUNs on this RAID.</p> </div> <p>The default: Disabled.</p>

Parameter	Description
Reconstruction Priority	<p>Priority of data reconstruction process, compared with other processes in the system.</p> <p>Possible values are integers from 0 to 100, where:</p> <ul style="list-style-type: none"> 0 – reconstruction stops if any other activity takes place; <div style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <p> When Background integrity check is enabled and Reconstruction priority is set to 0, reconstruction process will only run while integrity check pause.</p> </div> <ul style="list-style-type: none"> 100 – reconstruction mechanism has the highest priority. <p>The default: 5.</p>
Max Number of Reconstruction Requests	<p>Maximum number of concurrent requests for reconstruction.</p> <p>Possible values: integers from 1 to 64.</p> <p>The default: 8.</p>
Percent of Cache Available for Write Requests	<p>Cache percentage used for dirty segments.</p> <p>The specified cache percent is reserved for write operations, the rest is for read operations even while the intensive write is running.</p> <p>Possible values: from 1 to 100%.</p> <p>The default: 75.</p>
Advanced Reconstruction	<p>Advanced Reconstruction is mechanism optimizing the read operation performance by excluding the drives with the lowest read rates.</p> <p>Possible values:</p> <ul style="list-style-type: none"> Enabled; Disabled. <p>The default: Disabled.</p>
Moreover, the following cache setting is not included in the profile and applied for nodes:	
Write Through Without Synchronization	<p>Only in DC mode.</p> <p>Enable this option to run Write Through when caches are not synchronized.</p> <p>This option is available on the SYSTEM NODES page under the <i>UPS configuration</i> pane.</p> <div style="border: 1px solid red; padding: 10px; margin: 10px 0;"> <p> If you have created a RAID in DC with the Write Through Without Synchronization parameter enabled while <i>cache synchronization isn't configured</i>, the Write Through parameter for the RAID will remain enabled after cache synchronization is configured.</p> <p>In this case, you can change the Write Through settings manually.</p> </div>

ERA RAID parameters (all parameters are dynamic)

Parameter	Description
Enable merge optimization	<p>Except RAIDs 0, 1, and 10.</p> <p>Optimization for sequential loads of the low block size and the large queue length.</p> <p>The function only works when the condition is met:</p> $\text{data_drives} * \text{stripe_size} \leq 1024$ <p>where</p> <p>“data_drives” is a number of drives in the RAID (for RAIDs 5, 6, or 7.3) or in one RAID group (for RAIDs 50, 60, or 70) that are <i>dedicated for data</i>,</p> <p>“stripe_size” is a selected stripe size for the RAID (Stripe value in GUI, stripe_size in CLI) in KiB.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • Enabled; • Disabled. <p>The default: Disabled.</p>
Low thread optimization	<p>Optimization for loads produced by applications of low number of threads. Effective for a syndrome RAID.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • Enabled; • Disabled. <p>The default: Disabled.</p>
Reconstruction priority	<p>Priority of data reconstruction process, compared with other processes in the system.</p> <p>Possible values are integers from 0 to 100, where:</p> <ul style="list-style-type: none"> • 0 – reconstruction stops if any other activity takes place; <div style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <p>i When Background integrity check is enabled and Reconstruction priority is set to 0, reconstruction process will only run while integrity check pause.</p> </div> <ul style="list-style-type: none"> • 100 – reconstruction mechanism has the highest priority. <p>The default: 5.</p>
Initialization priority	<p>Priority of drives initialization process.</p> <div style="border: 1px solid black; padding: 10px; margin: 10px 0;"> <p>i Initialization can be demanding on the system computational resources if high workload comes on the RAID being initialized.</p> </div> <p>Possible values are integers from 0 to 100, where:</p> <ul style="list-style-type: none"> • 0 – initialization stops if any other activity takes place; • 100 – initialization mechanism has the highest priority. <p>The default: 5.</p>

Parameter	Description
Memory limit	RAM memory limit, defined in MiB. Possible values: integers from 1024 to 1048576 . By default, the system calculates this value based on the Stripe size.

4.1 Working with Advanced RAID Parameters

While creating a RAID, Advanced RAID Parameters are applied by default. For Generic RAID, the parameters can be changed in addition by [creating a RAID profile](#) and applying it to the RAIDs. The default Advanced RAID Parameters can be modified for each RAID.

To modify Advanced RAID Parameters:

1. Select **STORAGE > RAID LIST**.
2. In the *RAID List* section, click on the name of a corresponding RAID. The page of the selected RAID opens.
3. Click **Show advanced RAID parameters**. The Advanced RAID parameters window opens (Generic RAID – Figure 34; ERA RAID – Figure 35).

Advanced RAID "ini_RAID6" Parameters

RAID Size	Free Size	Level	Drives	RAID Engine	RAM cache size	SparePool	Stripe	Sync	Affinity/Status
29 GiB	29 GiB	RAID 6i	0-4	Generic	2 GiB	—	128 KiB	On	

Static Parameters

4 KiB blocks write mode	On	?
Alt Buffers Size	512	?
Write Back at Once	24	?
Max Write Data	256	?
Number of Threads	4	?
Verify Cache Sum	<input type="checkbox"/>	?
Reconstruct At Once	24	?

Dynamic Parameters

Clean Segments Replacement	LRU	?
Dirty Segments Replacement	LRU	?
Max Dirty Segments for LRU	24	?
Max Dirty Segments for Cycle	24	?
Max Write Back Sequential	8	?
Max Write Back Random	1024	?
Write Barriers Support	<input checked="" type="checkbox"/>	?
Random Reads	<input checked="" type="checkbox"/>	?
Random Writes	<input checked="" type="checkbox"/>	?
Silent Data Corruption	Disabled	?
Write Through	Disabled	?
Reconstruction Priority	5	?
Max Number of Reconstruction Requests	8	?
Percent of Cache Available to Write Requests	75	?
Advanced Reconstruction	<input type="checkbox"/>	?

Save as RAID Profile...

CANCEL APPLY

Figure 34. "Advanced RAID Parameters" window for a Generic RAID

Advanced RAID "ERA_RAID5" Parameters

RAID Size	Free Size	Level	Drives	RAID Engine	SparePool	Stripe	Sync	Affinity/Status
837 GiB	837 GiB	RAID 5	5-8	ERA	—	16 KiB		

Dynamic Parameters

Enable merge optimization	<input type="checkbox"/>	?
Low thread optimization	<input type="checkbox"/>	?
Reconstruction Priority	5	?
Initialization priority	5	?
Memory limit	2048	?

CANCEL APPLY

Figure 35. "Advanced RAID Parameters" window for an ERA RAID

- Specify RAID parameters and click **APPLY** to save the settings.

For Generic RAID, you can click **Save as RAID Profile** and apply the parameters for other RAID.

4.1.1 Advanced Reconstruction Setup

 Only for Generic RAID.

The Advanced Reconstruction feature optimizes the performance of reading operations for syndrome RAID. Instead of reading data from "slow" drives, the system "recovers" data due to data redundancy of a RAID.

Advanced Reconstruction is available for RAID 5, 6, 7.3, and N+M.

To configure the advanced RAID reconstruction parameters:

1. Select **STORAGE > RAID LIST**.
2. Select a RAID.
3. On the *Info* pane of the opened page, click **Show Advanced RAID Parameters**.
4. In the opened *Advanced RAID Parameters* window, use the **Advanced Reconstruction** switch (Figure 36).

Dynamic Parameters		
Clean Segments Replacement	LRU	?
Dirty Segments Replacement	LRU	?
Max Dirty Segments for LRU	24	?
Max Dirty Segments for Cycle	24	?
Max Write Back Sequential	8	?
Max Write Back Random	1024	?
Write Barriers Support	<input checked="" type="checkbox"/>	?
Random Reads	<input checked="" type="checkbox"/>	?
Random Writes	<input checked="" type="checkbox"/>	?
Silent Data Corruption	Disabled	?
Write Through	Disabled	?
Reconstruction Priority	5	?
Max Number of Reconstruction Requests	8	?
Percent of Cache Available to Write Requests	75	?
Advanced Reconstruction	<input type="checkbox"/>	?

Figure 36. "Advanced RAID Parameters" window. The Advanced Reconstruction switch



We don't recommend enabling Advanced reconstruction if the system has big number of drives and RAIDs.

4.2 RAID Profiles



Only for Generic RAID.

During RAID creation, you can apply Advanced Parameters through RAID Profile. The default RAID Profile is suitable for all workloads. To enhance performance, set up Advanced Parameters following your specific workload. RAID Profile is applicable separately and for any number of RAIDs.

4.2.1 Creating a RAID Profile

i If you didn't create a RAID Profile, the system applies default settings (Default RAID profile).
You can't change parameters of the Default RAID profile.

To create a RAID Profile:

1. Select **STORAGE > RAID PROFILES**.
2. In the *RAID Profiles* section, click **Create Profile** (Figure 37).

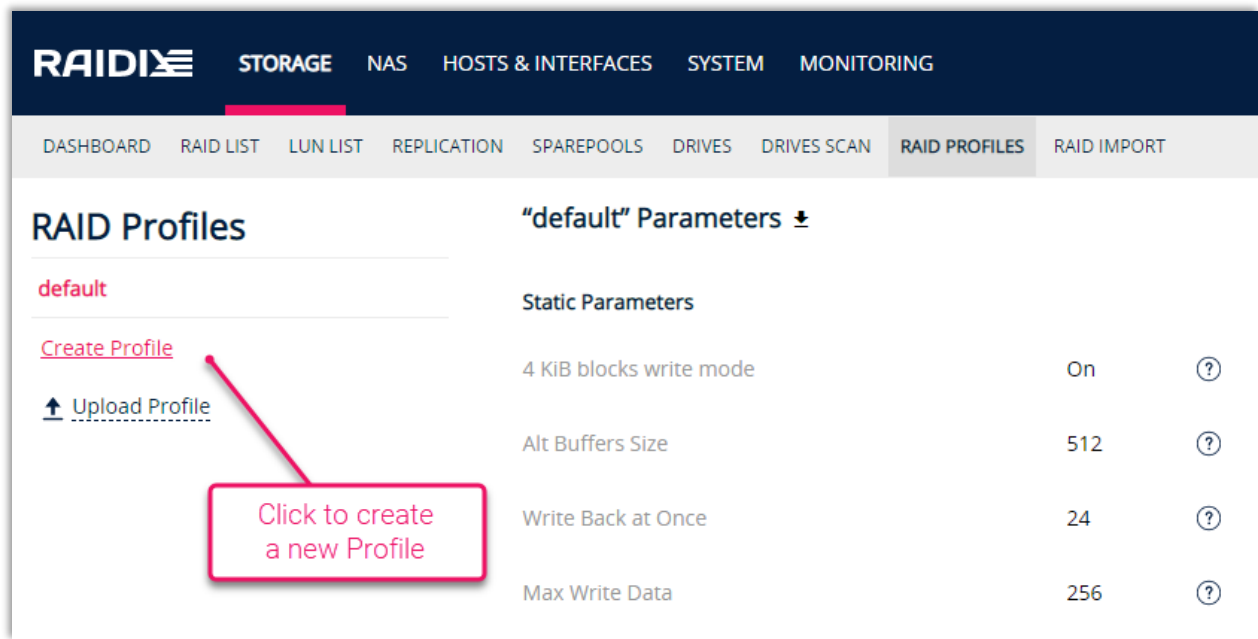


Figure 37. Create Profile control on the RAID Profiles page

3. Fill in the **Profile Name** field and modify required parameters.
4. Click **SAVE**.

You can apply a RAID Profile during the RAID creation (Figure 38).

RAID List is empty

>>	RAID Name	RAID Size	Free Size	RAID Level	Drives
<div>RAID Name <input type="text"/></div> <div>RAID Engine Generic ▾ ⓘ RAID Level 0 ▾</div> <div>Use Profile <div><div>default ▾</div><div>default</div><div>RAID_Profile1</div><div>RAID_Profile2</div></div> Drives <div><div>Drives</div><div>Select</div></div></div>					

Figure 38. List of RAID Profiles in the RAID creation widget


4.2.2 Modifying a RAID Profile

1. Select **STORAGE > RAID PROFILES**.
2. In the *RAID Profiles* section, click the name of an appropriate RAID profile.
3. Perform the required modifications and click **SAVE**.

4.2.3 Downloading a RAID Profile

1. Select **STORAGE > RAID PROFILES**.
2. In the *RAID Profiles* section, click the name of an appropriate RAID profile.
3. Click **DOWNLOAD PROFILE**.

4.2.4 Deleting a RAID Profile

1. Select **STORAGE > RAID PROFILES**.
2. In the *RAID Profiles* section, click  by the name of the profile to be deleted.
3. To confirm, click **DELETE**.

5. DRIVES

This chapter provides information about the drives in the system that are intended to create a RAID or already are part of a RAID.

To learn more about the drives on which the system is installed, see the [System Drive](#) chapter.

5.1 Information about drives

To see information about drives in the system, open the **STORAGE | DRIVES** page (Figure 39).

Drives

INTEL SR2612UR Revision: I106 ID: 500019b9ff5ff000 9 drives Size: 8,300 GB

5 working 5 TB, 2 available 1.3 TB, 2 reserved 2 TB

LED	UID	Capacity	Slot	Model	Type	RAID Name	SparePool	ID	Local	Remote
>	0	1,000 GB	0	MK1001TRKB	HDD	GEN_RAID5	—	sdb		
>	1	1,000 GB	3	MK1001TRKB	HDD	GEN_RAID5	—	sdc		
>	2	1,000 GB	6	MK1001TRKB	HDD	GEN_RAID5	—	sdf		
>	3	1,000 GB	9	MK1001TRKB	HDD	GEN_RAID5	—	sdh		
>	4	300 GB	11	ST3300657SS	HDD	—	—	sdj		
>	5	1,000 GB	7	MK1001TRKB	HDD	GEN_RAID5	—	sdg		
>	6	1,000 GB	8	MK1001TRKB	HDD	—	Sparepool1	sda		
>	7	1,000 GB	10	MK1001TRKB	HDD	—	Sparepool1	sdi		
>	8	1,000 GB	5	MK1001TRKB	HDD	—	—	sde		

Drive Parameters

Threshold Fault Count: 3

Health Monitoring: ☐

Drives Scan

Type: Read Scan

RAID Name: GEN_RAID5

Size: 1 GiB


Offset: 0%




Integrity Check: ☐

SCAN

Figure 39. DRIVES page

Drive parameters on the **STORAGE | DRIVES** page.

Parameter	Description
LED	Drive indicator button. After clicking  , a drive indicator on the enclosure turns on. It allows to associate the real drive with its prototype in the RAIDIX software.
UID	Unique identifier of the drive in the system.
Capacity	Hard drive capacity, in GB.
Slot	Slot number. Numbering starts from "0".
Model	ID of the drive model.
Type	The type of the drive: HDD or SSD.
RAID Name	Name of the RAID containing the drive.
SparePool	SparePool name, into which the drive is added.

Parameter	Description
ID	A device name in the DevFS file system.
Status	<p>Drive status on the corresponding node:</p> <ul style="list-style-type: none">  <i>OK</i> – works correctly.  The <i>Warning</i> status appears in the following situations: <ul style="list-style-type: none"> <i>dirty</i> – drive contains partitions. <i>faulty warning</i> – drive have had the number of I/O errors less than the value of the <i>Threshold fault count</i> parameter. <i>needs reconstructing</i> – drive reconstruction is required. <i>reconstructing</i> – drive is reconstructing.  The <i>Bad</i> status appears in the following situations: <ul style="list-style-type: none"> <i>error</i> – a failure occurred. <i>faulty error</i> – the number of I/O errors occurred on the drive is equal or more than the number of <i>Threshold fault count</i>, this drive cannot be used in RAID or in SparePool. <i>no drive</i> – there is no drive in the system. <p>DC-systems display the status of drives connected via an external enclosure on the remote node. If the drives are connected directly to a node, only the local status shows.</p>

5.2 Drive Health

The health of a drive in RAIDIX is determined by the number of drive I/O errors when the system was unable to complete a read or write operation on the RAID, and by drive internal errors diagnosed by the drives themselves via S.M.A.R.T.

Drive internal errors do not always lead immediately to I/O errors, but a certain amount of such errors or combinations of such errors suggests that the drive is becoming failed and is recommended for replacement.

I/O errors are not always due to drive problems. The cause may be, for example, a bad cable connection to the drive.

The main causes of errors on the drive:

- Problems with the drive: mechanical damage or wear;
- Hardware problems: cables, adapters, or SAS Expanders;
- Software problems: driver or firmware errors.

The following drive health monitoring features are available to you in RAIDIX:

- Drive Error counter.
- Drive Health status.

- S.M.A.R.T. diagnostics.

For details about each function, see the corresponding subchapters below.

5.2.1 Error Counter

Drive Error Counter is used to keep track of drives that *are part of a RAID* where I/O errors have started to occur. A *drive I/O error* (later on "*drive error*") means that the system was unable to complete a read/write operation to the drive. If the number of drive errors is zero, such drive is considered *healthy*, if the number is more than zero but less than the Threshold fault count value – *unhealthy*, if the number is equal to the Threshold fault count value – *failed*.

The following options are available to you relating to the Error Counter:

- defining the value of the threshold fault count parameter;
- clearing the current value of the number of drive errors.



Threshold fault count

The *Threshold fault count* parameter indicates the number of errors on the drive above which the drive is marked as "faulty error" and excluded from its RAID.

If the number of occurred errors is less than the *Threshold fault count* value, the drive is marked as "faulty warning" but not excluded from its RAID.

To change the *Threshold fault count* value:

i After you change the parameter value, read/write fault count will reset for all drives.

1. Open the **STORAGE | DRIVES** page.
2. Expand **Drive Parameters** on the right of the **DRIVES** page (Figure 40).
3. In the line of the *Threshold Fault Count* parameter, click  and specify a new parameter value. Confirm by clicking .

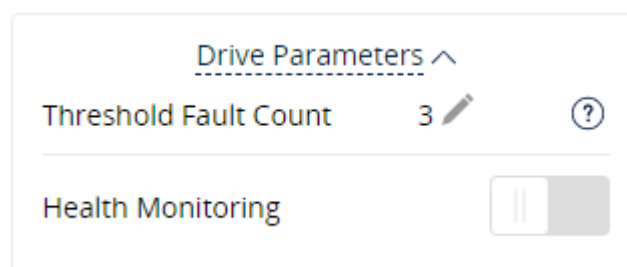


Figure 40. Drive Parameters pane

Cleaning a drive

Depending on whether a drive is in a RAID or not, the drive cleanup enables you to:


- Reset the drive faulty counter if the drive is in a RAID (drive is in the “Faulty” state).



While cleaning a drive with Faulty error or Faulty warning statuses, note that these states were assigned to drive because some I/O errors occurred on it. With high possibility these errors can occur again, therefore, when such statuses appear, it is recommended to replace the failed drive with a new one as soon as possible.

- Remove all data on the drive if the drive is not in a RAID (drive is in the “Dirty” state).

To clean a drive:

1. Open the **STORAGE | DRIVES** page.
2. Click  against switching drive and select **Clean**.
3. To confirm, click **CLEAN**.

5.2.2 Health Tracking



Only for SAS HDDs and SAS SSDs.

The drive automatic health tracking feature tracks S.M.A.R.T. attributes to *alert* you in advance (before I/O errors occur) about faulty drives as well as to enable the function of *proactive automatic replacement* of problematic drives (“advance automatic replacement”, see [Automatic Drive Replacement](#) for details).

In contrast to the Error Counter, Health tracking does not require drives to belong to a RAID.

You can enable the function on the **STORAGE | DRIVES** page with the **Health Monitoring** switch on the *Drive Parameters* pane.

The following S.M.A.R.T. attributes are used to check drive health:

- Unrecovered read errors.
- Unrecovered verify errors.
- Unrecovered write errors.

By default, if at least *one* of these errors affects a drive, the system will display a notification that the drive health is not OK.

You can configure the parameters via the CLI (see `rdcli param drive`).

5.2.3 S.M.A.R.T. Diagnostics

Contemporary drives contain sensors, which help to reveal and store the logs on discovered defects that can lead to the device failure. The drives can also use the S.M.A.R.T. technology for self-check. The results obtained during the self-check are stored in logs. SATA and SAS drives have different S.M.A.R.T. values.


S.M.A.R.T. diagnostics provides information on drives used by the system administrator to manually interpret and make decisions about replacing the drives.

For each S.M.A.R.T.-enabled drive, you have access to:

- Full S.M.A.R.T. information.
- General health notifications.
- Notifications about SSD wear percentage.

Full S.M.A.R.T. information

To display S.M.A.R.T. information for the drive:

1. Open the **STORAGE | DRIVES** page.
2. Click  in the line of the required drive, and then click **SMART**. The *SMART* window will open (Figure 41).

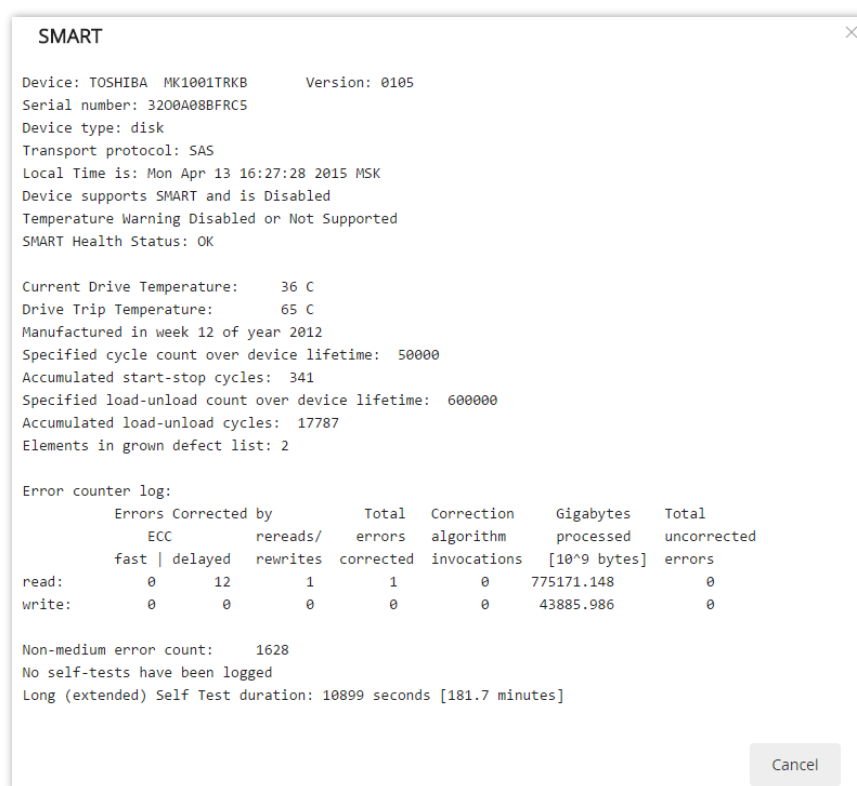


Figure 41. SMART window

Some parameters on the *SMART* window

Parameter	Description
Device	Information on drive vendor, model number, and firmware version.
Serial Number	Drive serial number. The last 8 symbols on the right are actually the drive serial number. All other numbers are unique identifier of the drive vendor.
Device Type	The device type.
Transport protocol	Drive connection protocol (interface) (SAS or SATA).
Local Time	Local time of the system.
Device supports SMART and is Enabled	The parameter specifies whether the drive supports SMART function and the function is activated.
Temperature Warning	Parameter specifies whether the temperature rise alert is turned on.
SMART Health Status	The current overall drive health.
Current Drive Temperature	The current drive temperature (in Celsius).
Drive Trip Temperature	Maximum internal drive temperature which has ever been logged.
Elements in Grown Defect List	<p>The hard drive defines and supports the list of sectors where write cannot be performed. Such sectors are called "drive surface defects." There are two types of defects:</p> <ul style="list-style-type: none">• <i>Manufacturing Defect List</i> – list of defects discovered during the test phase supplied by the drive vendor. This list remains unchanged;• <i>GDL</i> (Grown Defect List) – a list of defects, which appear during the drive use. The list is regularly updated.

Parameter	Description
Vendor Cache Information	<p>The parameter provides the following information:</p> <ul style="list-style-type: none"> • <i>Blocks Sent to Initiator</i> – if the SAS drive is used, host adapter is the initiator and the drive itself is a target. The parameter specifies the number of data blocks sent to Initiator. It is not only the drive data that can be SMART information. In most cases this is sector information, thus we can state that the parameter specifies the number of drive sectors where the read operation has ever been performed. • <i>Blocks Received from Initiator</i> – number of drive sectors where the write operation has ever been performed. • <i>Blocks Read from Cache and sent to the Initiator</i> – caching performance effectiveness level. If the initiator requests the same block twice and the data is in the drive cache, there is no need to read it again. Thus, the value is always equal or higher than Blocks Sent to the Initiator value. The higher this value is the more intensively the drive heads must operate. • <i>Number of Read or Write Commands whose size <= Segment Size</i> – The hard drive directs the data to the computer using the blocks sequence through the cache area named the cache segment. This value is a number of commands with the size less or equal to the cache segment. The size of most commands exceeds the segment size. • <i>Number of Read or Write Commands whose size > Segment Size</i> – number of commands or data, which had to be split into several drive to transfer to the drive or from the drive.
Vendor (Factory) Information	
Number of Hours Powered Up	The number of hours during which the drive power was turned on regardless of whether read or write operations took place during this time. When calculating the powered up hours, idle time is also counted. If the powered up drive was in sleep mode, sleep time is also counted in.
Number of Minutes until next SMART test	<p>The drive can be verified using two types of tests within the SMART technology:</p> <ul style="list-style-type: none"> • <i>Quick test</i> that takes several seconds and is initiated by the drive itself but can also be run manually. The information is updated after testing is complete; • <i>The complete drive scan</i> that can be initiated only by user.

Notifications about general health

The notifications are based on S.M.A.R.T. drive data and provide binary health information – OK or not OK.

The system scans the S.M.A.R.T. data and shows notifications every 24 hours.

You can set up receiving notifications by SMTP and SNMP, see [Setting up Notifications](#) for details.

Notifications about SSD wear percentage

The notifications are based on S.M.A.R.T. SSD data and warn of drive worn out when the 70% and 90% thresholds are reached.


The system scans the S.M.A.R.T. data and shows notifications every 24 hours.

You can set up receiving notifications by SMTP and SNMP, see [Setting up Notifications](#) for details.

5.3 Drive Performance

 Only for Generic RAID.


In RAIDIX 5.2, ability to scan drives of a RAID on performing read/write operations is implemented. Scan results show the number of read/write commands, conducted by the system in different time intervals. Analysis of the results allows indicating drives with low performance.

 Drive Scan shouldn't be run during Advanced Reconstruction because performance parameters will not be objective due to data calculation instead of data read. Besides, under these conditions, some drive may be marked as Slow due to simultaneous read from RAID and drives during the scanning.

To run drive scan:

 After starting a new test, all previous tests results will be removed.

1. Open one of the pages:
 - **STORAGE | DRIVES.**
 - **STORAGE | DRIVES SCAN.**
 - RAID page.
2. Configure Drive Scan using the following parameters:
 - **RAID name;**
 - the scanning **Type** (**Read Scan** or **Write Scan**);

 When starting a new write scan test, all data on the drives will be deleted.

 To run a write test, the RAID must not contain LUNs.

- the **Size** of data to be read/written on the drives (**1 GiB, 10 GiB, 100 GiB, All RAID**);
- **Offset** (**0%, 25%, 50%, 75%**) between the beginning of the drive data area and the start point of the scan;
- **Integrity Check** function during the scan (available only for initialized RAIDs).

**Limitations:**

- Integrity Check available only for initialized RAID5;
- SDS detection unavailable for RAID5 and 6 in degraded mode;
- SDS detection available for RAID5 with *one* missing drive.

The check is running with the read scan detecting SDCs. If the function detects an SDC, the system will show an alert after the scan finished:

- in the RAID status field;
- in the node Alerts section;
- on the **MONITORING | ALERTS** page.



Drive Scan Integrity Check function doesn't correct SDC.

3. Click **SCAN**.

After the scan finished, the system shows the tooltip *Scan completed*.

When you start the scan from the **DRIVES** page, the system will automatically open the **DRIVE SCAN** page with scan results (Figure 42).

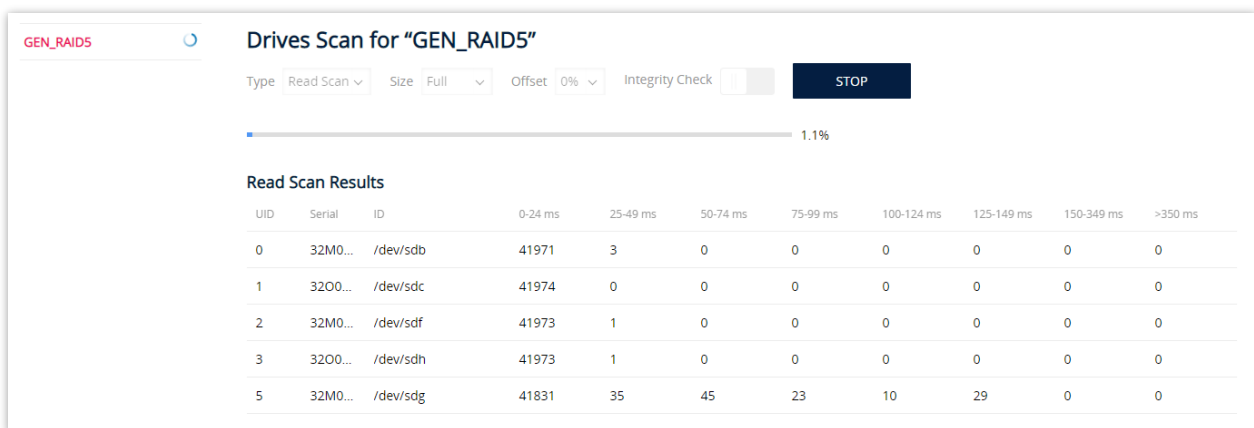


Figure 42. Read Scan in progress. The results are displaying dynamically

Parameter	Description
UID	Unique identifier of the drive in the system.
Serial	Serial number of drive in the slot.
ID	A device name in the DevFS file system.
Time intervals	Command execution time intervals (in milliseconds).

Interpretation of Read Scan and Write Scan results

When the scanning option is used, the information about the number of commands operated in different time intervals will display in the appropriate sections of the **DRIVES SCAN** page. For example,

the results of a Read Scan test on Figure 42 show that the drive in a slot number 3 has the best speed characteristics, as most of requests (33) were performed in the shortest interval (0-24 milliseconds). Zero value of the last columns (time range >100 milliseconds) for the drive tells that all requests during the read operation were quick.

The best case is when the rightmost columns of the tables contain zero values.

If the scan reveals drives with values uncommon for the overall statistics, these drives may negatively affect the whole RAID performance. For example, if a full RAID showed non-zero values only in the rightest column >150 ms, this indicates a drive failure. Consider replacing slow drives with new ones

5.4 Drive Replacement in the RAID

If a RAID drive has disconnected from the system, has failed, or may fail (based on S.M.A.R.T. diagnostic data), you can manually replace such a drive or set up an automatic replacement.

Ways to replace a drive in the RAID:

- Manual replacement:
 - Physical replacement of a drive in the enclosure.
 - Replacing a drive with another drive that already is in the system.
- Automatic replacement:
 - Automatic replacement of a faulty or disconnected drive.
 - Proactive automatic replacement of a potentially malfunctioning drive.

After the drive replacement is finished, the reconstruction of the RAID drives will start.

Reconstruction is a process of data restoring. You can set up the reconstruction priority depending on host workload (to learn more about the priority, see [Advanced RAID Parameters](#)).



For Generic RAID, the reconstruction may take more time after a failover or reboot of a RAID.



To stop the reconstructing for a Generic RAID, use [rdcli raid reconstruct stop](#).

To start the reconstructing, use [rdcli raid reconstruct start](#).

The *Advanced Reconstruction* and *Reconstruction per drive* features are used for the reconstruction process, improving RAID performance.

Improvements of reconstruction:

- Improved RAID fault tolerance to the single I/O failures.

If several drives in a RAID require reconstruction at the same time, the system will reconstruct the drive with the smallest number of inconsistent data first (the *Advanced Reconstruction* feature).
- Increased the RAID reconstruction speed.

When several drives need reconstruction, only the inconsistent data on each drive will be reconstructed (the *Reconstruction per drive* feature).

The *Reconstruction per drive* feature has the following specifics depending on the RAID type:

- ERA RAIDs are more fault-tolerant than Generic RAIDs:
Data on ERA RAIDs will be recovered if the number of corrupted strips in each stripe does not exceed the number of parities in the RAID; in Generic, if the number of drives with corrupted data does not exceed the number of parities in the RAID.
- Displaying the reconstruction:
For ERA RAIDs, the system displays the percentage of data that do not need reconstruction (the percentage named "RAID health percentage").
For Generic RAIDs, the system displays the arithmetic average of the percentages of the reconstruction for all drives that need to be reconstructed.

5.4.1 Automatic Drive Replacement

The following methods of automatic drive replacement are available to you:

- Automatic replacement of a faulty or disconnected drive.
- Proactive automatic replacement of a potentially malfunctioning drive.

To enable automatic replacement of disconnected or faulty RAID drives, create a SparePool and assign it to the RAID.

You can optionally enable proactive automatic replacement after you have assigned a SparePool for the RAID. The proactive replacement mechanism uses S.M.A.R.T. data to replace drives in the RAID before they fail.

i After replacing a drive in Generic RAID (except RAID 0) or ERA RAID (except RAID 0), the system will start data reconstruction on the RAID drives.

! After auto-replacement of a drive from the SparePool, the drive will be displayed in the SparePool. Manually unselect such drive from the SparePool and add a new drive to the SparePool if possible. Otherwise, information about the number of free drives in the SparePool may be incorrect.

By default, RAID auto-replacement starts with some time delay to prevent accidental short-term drive disconnections. You can view or change the time delay by using the CLI commands:

```
$ rdcli param system show
```

```
$ rdcli param system modify
```

the parameter `--hot_replace_timeout` (to learn more, see [Appendix A](#)).

You can create the SparePool through:

- the *widget* *Create a SparePool* on the page **STORAGE | SPAREPOOLS**;
- the *wizard* **CREATE SPAREPOOL** on the page **STORAGE | DASHBOARD**.

To create the SparePool via the *widget*:

i Consider that specific SSD wearing out occurs in RAIDs that contain both an SSD and an HDD.

- Select **STORAGE > SPAREPOOLS**.
- Click **Create SparePool**. The widget for creating the SparePool opens (Figure 43).

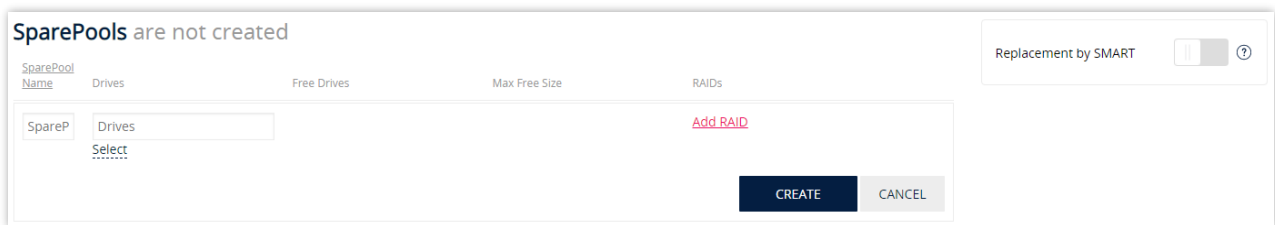




Figure 43. SparePool creation widget

- Fill in required fields:

Parameter	Description
Editable fields	
SparePool Name	Name of the SparePool.
Drives	<p>UIDs of drives included in the SparePool.</p> <p>In the Drives field, type numbers of drives to be included in the SparePool or click the Select link, check the appropriate drives, and click SELECT.</p> <div> i One drive can be included only in one SparePool. Drives, already included in an array, cannot be included in SparePool. </div>
RAIDs	<p>Optional.</p> <p>RAID, associated with the SparePool (an automatic replacement (hot spare) of the specified RAID drives will be performed with drives from the specified SparePool).</p> <p>Click the Select RAIDs link, select one or several RAID and click Select.</p> <p>From the right of the table status of the corresponding SparePool is displayed. Possible variants are:</p> <div> <div>  OK – SparePool is available for use; </div> <div>  Warning – there are no drives of suitable size in the SparePool to replace a failed drive in the associated RAID, or there are no free drives (available for replacement) in the SparePool. </div> </div>

Parameter	Description
Information fields	
Free Drives	SparePool drives available to be included in RAIDs.
Max Free Size	Maximum size of the drive included in the SparePool.

4. Click **CREATE**.

i After selecting drives is included in the SparePool, the system will perform a test write on them. If any of these drives cannot pass the automatic write test, the system will return an error notification that the drives cannot be included in SparePools, and you should select other drives.

To enable proactive auto-replacement of drives for the RAID:

1. Make sure the SparePool is assigned for the RAID.
2. Turn on **Health Monitoring** (to learn more, see [Health Monitoring](#)).
3. Open **STORAGE | SPAREPOOLS** and turn on **Replacement by health status** (Figure 44).

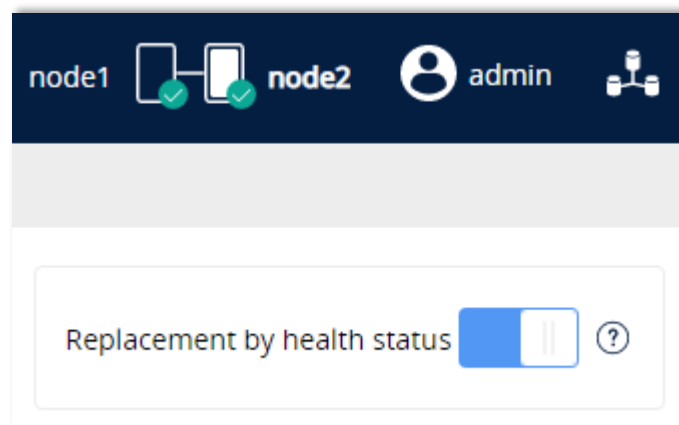


Figure 44. The switch "Replacement by health status"

5.4.2 Manual Drive Replacement

i Consider that specific SSD wearing out occurs in RAIDs that contain both an SSD and an HDD.

The following methods of manual drive replacement are available to you:

- Physical replacement of a drive in the enclosure.
- Replacing a drive with another drive that already is in the system.

Requirements for new drives used in the replacement:

- The physical replacement:

- the new drive capacity must not be less than the capacity of the failed drive;
- the new drive interface must not differ from the capacity of the failed drive;
- if the new drive was used before, you have to remove metadata from it.

To remove metadata from a drive, use the following CLI command:

```
$ dd if=/dev/zero of=/dev/<sdX> bs=1M count=1000 oflag=direct
```

where <sdX> is a name of a block device.

- The replacement with the drive that already is in the system:
 - the new replacing drive is not used in any RAIDs;
 - the new replacing drive is not used in any SparePools of another RAID.

If the replacing drive is included in the SparePool assigned to this RAID, the drive stays in the SparePool.

To physically replace a drive in the drive enclosure:


1. Define a failed drive (when a failure occurs, the drive LED stops blinking).
2. Remove the failed drive.
3. Insert a healthy drive instead of the removed one.

After replacing a drive in the Generic RAID (except RAID 0) or the ERA RAID (except RAID 0), the system will start data reconstruction on the RAID drives.

To replace a drive manually with a drive that already is in the system:



The option of manual drive replacement in GUI will be available if there is a drive with suitable size in the system.

1. Select **STORAGE > DRIVES**.
2. In the menu of the failed drive, click  and select **Replace** (Figure 45).

The screenshot shows the RAIDIX web interface. The top navigation bar includes: RAIDIX, STORAGE, NAS, HOSTS & INTERFACES, SYSTEM, and MONITORING. Below this is a sub-navigation bar with: DASHBOARD, RAID LIST, LUN LIST, REPLICATION, SPAREPOOLS, DRIVES (selected), DRIVES SCAN, RAID PROFILES, and RAID IMPORT.

The main section is titled "Drives". It shows a summary for "INTEL SR2612UR" with Revision: I106, ID: 500019b9ff5ff000, 9 drives, and Size: 8,300 GB. Below this is a visual representation of drive status: 5 working (5 TB), 2 available (1.3 TB), and 2 reserved (2 TB).

A table lists the drives with columns: LED, UID, Capacity, Slot, Model, Type, RAID Name, SparePool, ID, Local, and Remote. A context menu is open for the drive with UID 2, showing options: SMART and Replace.

LED	UID	Capacity	Slot	Model	Type	RAID Name	SparePool	ID	Local	Remote
>	0	1,000 GB	0	MK1001TRKB	HDD	GEN_RAID5	—	sdb	✓	✓
>	1	1,000 GB	3	MK1001TRKB	HDD	GEN_RAID5	—	sdc	✓	?
>	2	1,000 GB	6	MK1001TRKB	HDD	GEN_RAID5	—	sdf	✓	↔
>	3	1,000 GB	9	MK1001TRKB	HDD	GEN_RAID5	—	sdh	✓	✓

Figure 45. Drive replacement from the Drives page

- In the opened Select a Drive for Replacement window, select a drive and click **REPLACE** (Figure 46).

After replacing a drive in Generic RAID (except RAID 0) or ERA RAID (except RAID 0), the system will start data reconstruction on the RAID drives.

The screenshot shows a window titled "Select a Drive for Replacement" with a search bar. It contains a table with columns: UID, Capacity, Slot, Enclosure, SparePool, Serial, ID, and Type. Three drives are listed, with the third drive (UID 8) selected. At the bottom right are "CANCEL" and "REPLACE" buttons.

UID	Capacity	Slot	Enclosure	SparePool	Serial	ID	Type
6	1,000 GB	8	500019b9ff5ff000	Sparepool1	3200A078FRC5	sda	HDD
7	1,000 GB	10	500019b9ff5ff000	Sparepool1	32M0A04FFRC5	sdi	HDD
8	1,000 GB	5	500019b9ff5ff000	—	32M0A05SFRC5	sde	HDD

Figure 46. Selecting a drive, which will be used as a new one

5.4.2.1 NVMe SSD Hot Plugging

General hot-plug features:

- All U.2 2.5" form factor NVMe SSDs are hot pluggable, but only one should be hot plugged at a time.
- Hot-plug operations only supported after the operating system is fully loaded.
- PCIe Adapter Card NVMe devices do not support hot-plug operations.

- Complete hot-plug operations within 1 second.

To learn about hot-swap NVMe SSD capabilities and restrictions for the hardware platforms you are using, refer to the appropriate service manuals.

6. LUN

To enable block or file device for a host, create a LUN on a RAID. LUN properties depend on tasks to be solved.



RAIDIX 5.2 provides LUN interaction with hosts of any supported operating systems.

On the **STORAGE | LUN LIST**, you can:

- view, create, delete, edit, and format LUNs;
- view and manage [Proxied LUNs](#);
- view and manage [SSD Cache](#).

On the **STORAGE | REPLICATION**, you can:

- view and manage [Replicated LUNs](#);
- view and manage [Replication Metadata LUN](#).

6.1 Creating the LUN

In RAIDIX 5.2, you can create a LUN through:

- the *widget* **Create LUN** on the **STORAGE | RAID LIST** page;
- the *widget* **Add LUN** on a RAID page;
- the *widget* **Create LUN** on the **STORAGE | LUN LIST** page;

During the LUN creation on the **LUN LIST** page, you should manually select a RAID for the LUN.

- the *wizard* **CREATE LUN** on the page **STORAGE | DASHBOARD**.

To create a LUN via *the widget* on the RAID page:

1. Select **STORAGE > RAID LIST**.
2. Select a RAID.
3. On the opened page, click **Add LUN** (Figure 47).

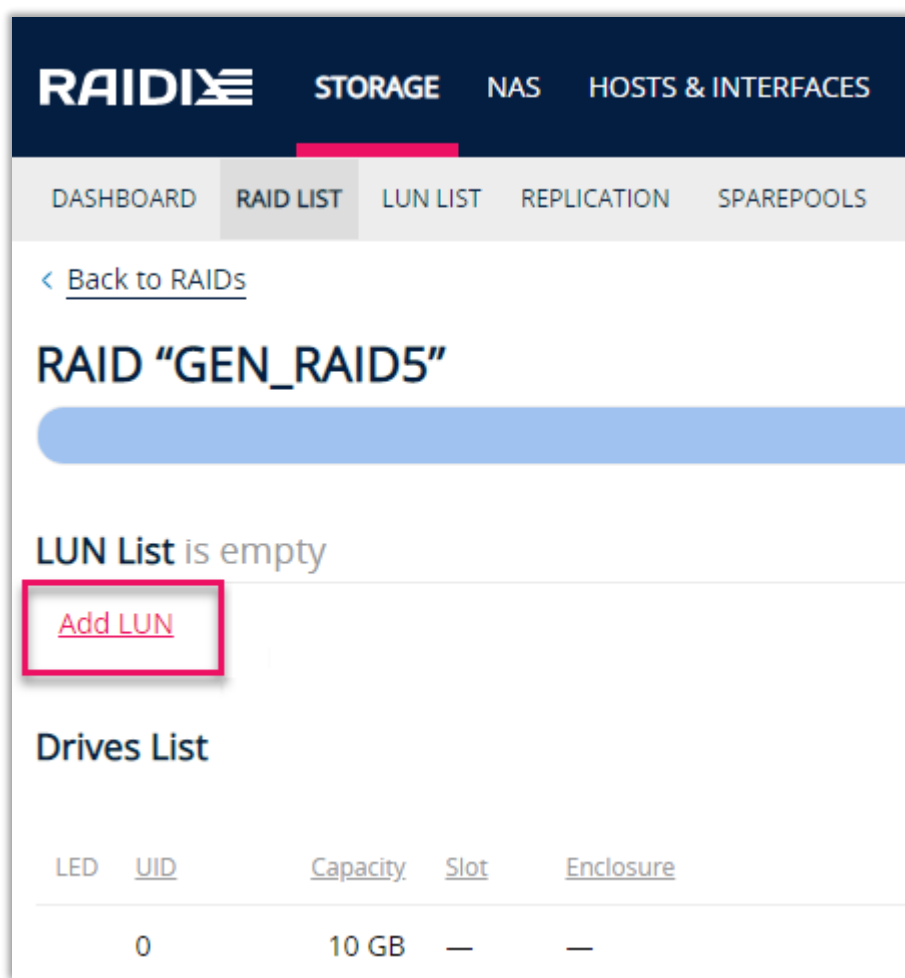



Figure 47. RAID page: Add LUN link

- Set up parameters (depends on RAID Engine).


LUN parameters of Generic RAID

Parameter	Description
Widget parameters	
LUN Name	A unique LUN name. Possible values: any letters, numbers, and underscore, from 1 to 15 characters.
LUN Size	LUN size (in GiB or TiB). The maximal value of the parameter is set by default.
Block Size	Block size: 512 B or 4096 B .
LUN type	Area of use of the LUN. Possible values: <ul style="list-style-type: none"> NAS – you can create a share or ssd cache. SCSI – block access by any of SCSI transport protocols. NVMe-oF – block access by any of NVMe-oF transport protocols. The default: SCSI .

Parameter	Description
VAAI ID	<p>Only for VAAI (vStorage API for Array Integration).</p> <p>A value of the prod_id parameter used for VAAI.</p> <p>The value must be the same for all LUNs on a single RAID (or RAIDs when extending a LUN) participating in VAAI.</p> <div>  Changing the parameter may cause the LUN on the host to be lost. We recommend reconnecting the LUN on the host after each change of this parameter. </div> <p>You can change the parameter in the window <i>LUN Parameters</i>.</p>
Next parameters show only after a LUN is created	
ID	<p>LUN number used for addressing purposes.</p> <p>Parameter can be used to identify a port number for Replication and to diagnose storage issues.</p>
LUN size on RAID	<p>Space which LUN takes up on the RAID. If the LUN is extended on several RAIDs, this value can be less than the LUN size.</p>
RAIDs List of LUN	<p>Displays on the LUN page and in the LUN Parameters window.</p> <p>List of RAIDs, on which LUN is extended. See LUN extension.</p>
Serial Number	<p>Displays on the LUN page and in the LUN Parameters window.</p> <p>Parameter can be used to diagnose storage issues.</p>
Sequential read parameters in the window LUN Parameters	
Read Ahead Past	<p>Minimum size of the sequential interval of read, in MiB.</p> <p>Possible values: from 1 to 2048.</p> <p>The default: 64.</p>
Read Ahead Distance	<p>Distance (in MiB) from the end of an interval, at which Read Ahead is performed.</p> <p>Possible values: from 1 to 2048.</p> <p>The default: 64.</p>
Read Ahead at Once	<p>Size of Read Ahead request, in MiB.</p> <p>Possible values: from 1 to 2048.</p> <p>The default: 24.</p>
Minimum Size of Sequential Pattern Request	<p>Minimum size of request (in sectors) which can be recognized as a part of sequential pattern.</p> <p>Possible values: from 1 to 2048.</p> <p>The default: 1.</p>
Range Timeout	<p>Maximum lifetime of interval, in seconds.</p> <p>Possible values: from 1 to 2048.</p> <p>The default: 10.</p>
Max Read Ahead Distance	<p>Maximum size from the end of the interval, at which read ahead is performed, in MiB.</p> <p>Possible values: from 1 to 2048.</p> <p>The default: 512.</p>

Parameter	Description
Max Read Ahead at Once	Maximum size of each read ahead request, in MiB. Possible values: from 1 to 2048 . The default: 24 .
Sequential Write Parameters on the page LUN Parameters	
Minimum Size of Sequential Interval	Minimum size of sequential interval of write, in MiB. Possible values: from 1 to 2048 . The default: 64 .
Minimum Size of Sequential Pattern Request	Minimum size of request (in sectors), which can be recognized as a part of sequential pattern. Possible values: from 1 to 2048 . The default: 1 .
Range Timeout	Maximum lifetime of interval, in seconds. Possible values: from 1 to 2048 . The default: 10 .

LUN parameters of ERA RAID

Parameter	Description
LUN Name	LUN name. Possible values: any letters, numbers, and underscore, from 1 to 15 characters.
LUN Size	LUN size (in GiB or TiB). The parameter is set by default to a maximum value.
Block Size	Block size. Possible values: <ul style="list-style-type: none"> • 512 B (available only for RAIDs with 512 B block size) • 4096 B
LUN type	Area of use of the LUN. Possible values: <ul style="list-style-type: none"> • NAS – you can create a share or ssd cache. • SCSI – block access by any of SCSI transport protocols. • NVMe-oF – block access by any of NVMe-oF transport protocols. The default: SCSI .
VAAI ID	<p>Only for ESXi VAAI Clone Blocks.</p> <p>A value of the prod_id parameter.</p> <div>  Each time you change this parameter, the LUN will be re-added to all targets, which will lead to a short-term loss of this LUN on the hosts. </div> <p>You can change the parameter in the window <i>LUN Parameters</i>.</p>

Parameter	Description
Next parameters shown on the page LUN Parameters only after creation of a LUN	
ID	Address LUN number. The parameter can be used to diagnose storage problems.
Serial Number	Displays on the LUN page and in the LUN Parameters window. LUN serial number. The parameter can be used to diagnose storage problems.

- Click **CREATE**.

6.2 LUN Extension

! During the extension of a LUN, its shares are unavailable to users.

i Before extending a LUN on an ERA RAID, we recommend waiting until the RAID initialization and reconstruction is complete if these processes are running. Alternatively, you can set the initialization priority to **0** during LUN extension.

You can extend the size of a LUN:

- with free size of the RAID that contains the LUN;
- with free size of additional Generic RAIDs, which located on one node, that you added to the LUN.

- i**
- You can add neither Generic nor ERA RAID to a LUN on an ERA RAID.
 - You can extend the size of a *LUN that is under a file system* by a maximum of 200 TiB at a time.

6.2.1 LUN Extension over the RAID

To extend a LUN:

- On the LUN page, click **Show LUN Parameters**.
- In the opened window, specify the LUN size in the appropriate field and click **APPLY**.
Maximum available size shown below the field.
- In the confirmation window, click **EXTEND**. Wait until the process of LUN extension is complete.

6.2.2 LUN Extension over Multiple RAID5s

i We recommend adding RAID5s with the same level as a RAID at which a LUN was created.

To add a RAID to an existing LUN:

1. On the page of a LUN, click **Show LUN Parameters**. The LUN Parameters window will open (Figure 48).

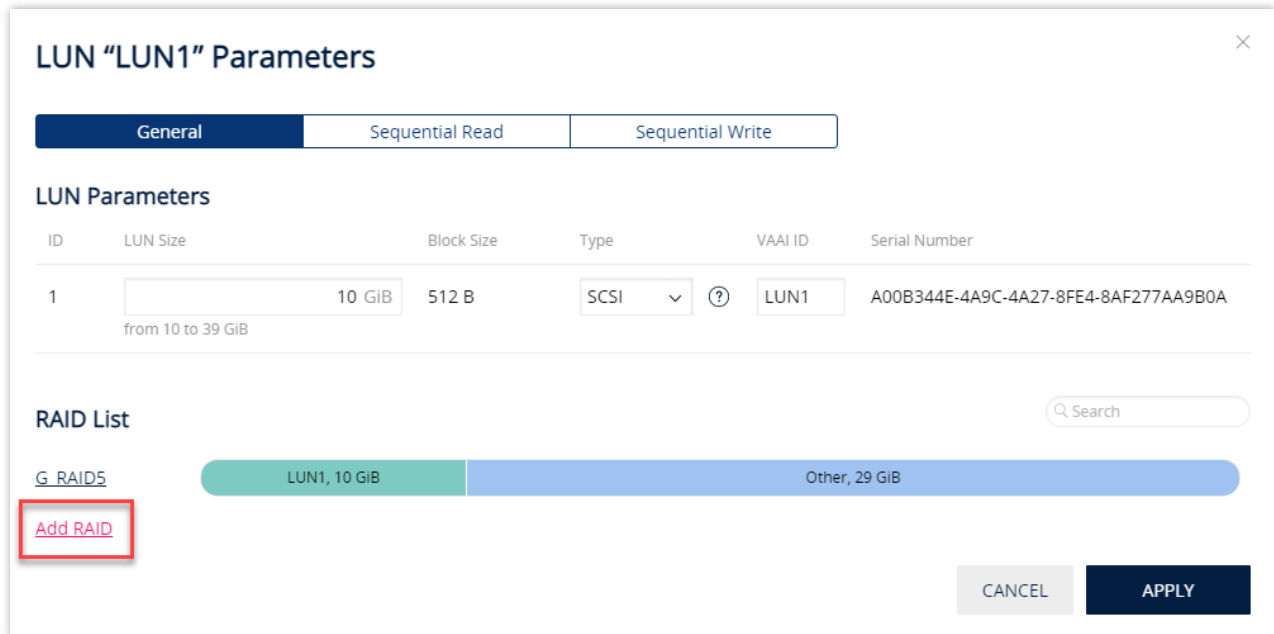
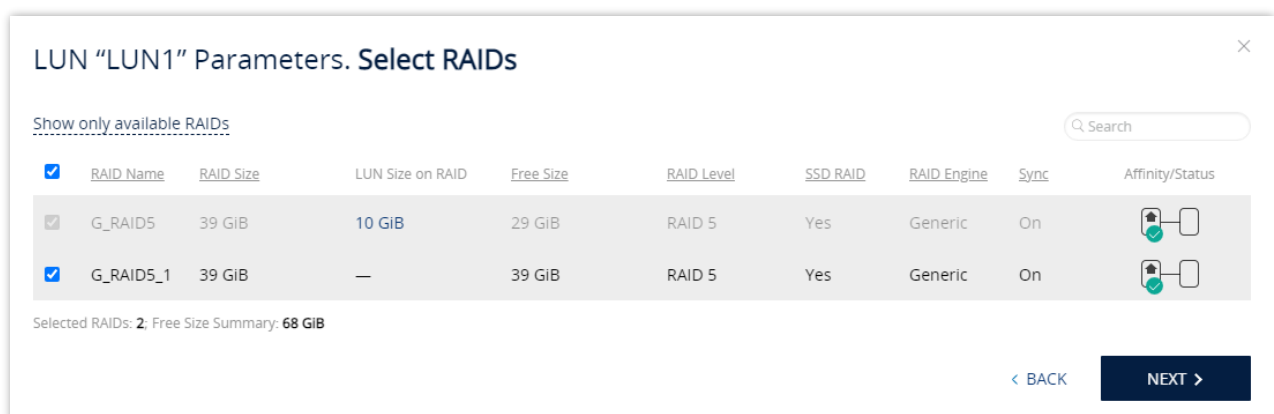


Figure 48. LUN page. The LUN Parameters pane

2. Below the *RAIDs List of LUN* section header, click **Add RAID**. Select RAID5s to be added to LUN and click **NEXT** (Figure 49).



<input checked="" type="checkbox"/>	RAID Name	RAID Size	LUN Size on RAID	Free Size	RAID Level	SSD RAID	RAID Engine	Sync	Affinity/Status
<input checked="" type="checkbox"/>	G_RAID5	39 GiB	10 GiB	29 GiB	RAID 5	Yes	Generic	On	
<input checked="" type="checkbox"/>	G_RAID5_1	39 GiB	—	39 GiB	RAID 5	Yes	Generic	On	

Selected RAID5s: 2, Free Size Summary: 68 GiB

Figure 49. Selecting RAID5s to be added to a LUN

3. To confirm, click **EXTEND**.

The name of the new RAID shows on the RAID5s list for the LUN (Figure 50).

- Specify new LUN size and click **APPLY**.

Figure 50. LUN Parameters after a Generic RAID was added

6.3 Masking Rules

i Only for LUNs of the SCSI type.

Masking implies setting up the rules for controlling access to the storage.

RAIDIX 5.2 enables the following masking rules.

Target masking rules

These rules imply specifying a target, from which a particular LUN will be available for hosts.

i In case of using Fibre Channel or InfiniBand adapters, a target is meant as a port on an adapter, from which the corresponding LUN will be available to hosts.

In case of using iSCSI, a target is a virtual device, which is created to run iSCSI requests and provide hosts access to LUNs over IP network.

The target masking improves system performance and fault tolerance.

In RAIDIX 5.2, the target masking rules creation is carried out on the LUN page. To learn more, see [Create a target masking rule](#).

Hosts masking rules

Hosts masking rules let you specify the access level of a host to a particular LUN or manage a host access to all LUNs simultaneously.

RAIDIX 5.2 software enables the following masking policy support:

In RAIDIX 5.2, the hosts masking rules creation is carried out on the LUN page. To learn more, see [Create a host masking rule](#) and Create a Host Alias.

6.3.1 Target Masking Rules

Target masking rule creation enables additional features on organizing hosts access to LUN.

Besides, the following advantages come to life:

- Increasing performance

You can assign a particular target to each LUN; using this target, LUN will interact with a particular host. As every flow is directed by a particular connection with similar speed, the effectiveness of data exchange with hosts increases and so does the overall performance of the system.

- Increasing fault tolerance

RAIDIX 5.2 supports MPIO (Multi-Path Input-Output) technology, which provides hosts access for LUNs via multiple data paths simultaneously.

Thus, you can assign several targets to each LUN. The software for LUN access management on several ports must be enabled on client computers. The connection will not be broken even if a problem occurs with any port of the target device or client computer. For data exchange, another designated target will be activated.

- Limiting access to LUNs

Limited access to LUNs is implemented in the following way: LUN is assigned to an individual target to be used by a particular user. Thus, this LUN is unavailable from other targets.

- Manageable hosts

By assigning several ports, you can manage hosts without the hardware switch. If no port is assigned to a host, it can access all LUNs through any ports by parameters specified when you assigned the hosts permissions.

6.3.1.1 Create a Target Masking Rule

To create a masking rule:

1. Select **STORAGE > RAID LIST**.
2. Click the name of the required RAID.
3. In the *LUN List* section, click the name of the required LUN. The LUN page opens (Figure 51).

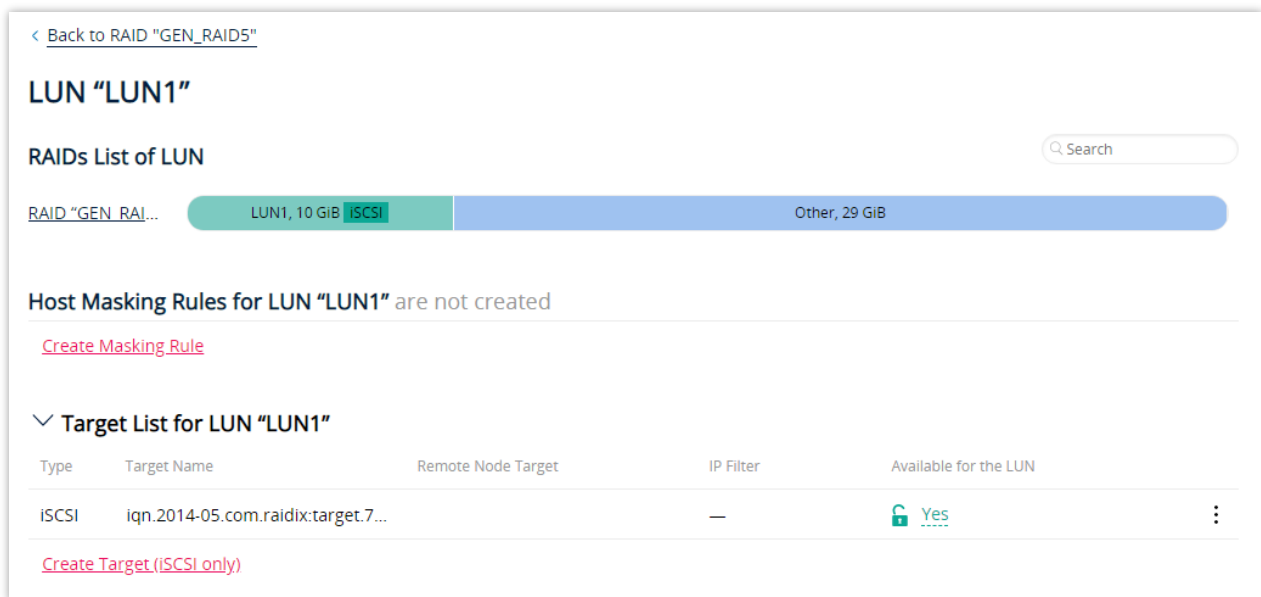






Figure 51. Masking Rules control and Target List section on the LUN page

4. The *Target List* section enables to assign a target to a LUN, and it contains the following parameters:

Parameter	Description
Type	<p>Target type:</p> <ul style="list-style-type: none"> FC (Fibre Channel) – a particular port assigned for LUN. The value defined automatically. iSCSI – a virtual device which is created to run iSCSI queries and to provide access to the logical device LUN over IP. IB (InfiniBand) – a particular port assigned for LUN. The value is defined automatically.
Target name	Target name.
Remote node target	<p>Only in DC mode. Only for InfiniBand, and Fibre Channel targets.</p> <p>Name of the target on the remote node that is linked with the target on the local node.</p> <p>The option enables you to combine two one-type targets on different nodes into one; that simplifies further masking rules setup.</p>
IP filter	<p>IP Filter enables a masking rule to be assigned a particular IP address or several IP addresses.</p> <p>Click All to select all available IP addresses.</p> <div style="border: 1px solid black; padding: 10px; margin-top: 10px;"> <p> This option is enabled <i>only</i> for iSCSI targets.</p> </div>
Available for LUN	The rule allows or denies an access to a particular LUN:

Parameter	Description
 Yes	– the rule enables access to a LUN for a particular target.
 No	– the rule denies access to a LUN for a particular target.

5. To set up InfiniBand or Fibre Channel access policy:

- While working in DC mode, to enable editing, click the  element.
- Click **Add link** in the **Remote node target** field and select the target to be linked with the current target on the local node in the showed window.
- In the **Available for the LUN** field, select the appropriate access policy. To confirm, click **Yes**.

6. To set up iSCSI access policy:

- Create a target iSCSI (for details on how to create an iSCSI-target refer to the section [Create an iSCSI target](#)).
- In the **IP filter** field, select an IP address or several IP addresses.
- In the **Available for the LUN** field, select the appropriate access policy. To confirm, click **OK**.

6.3.2 Host Masking Rules

Host masking rules let you specify the access level of the host to a particular LUN or manage the host access to all LUNs simultaneously.



We recommend you always configuring host masking rules. In some cases, if you do not restrict access to LUNs, MPIO may not work correctly.



If a LUN is formatted into a journaling file system, we recommend that you configure masking rules before mounting the LUN.

The following levels of host access to LUNs are supported in RAIDIX 5.2:

- Read Only;
- Read/Write.

Host masking rules depend on the mode, in which you are working by default:

After creation, LUNs are available to all hosts: all LUN users are automatically accessed to read and write if the access level for any host is not specified.

After addition of, at least, one host with Read/Write access level to the list, other users will not have access to the LUN. Thus, to obtain permission, it is necessary to add them to the list.

**For LUNs formatted not into a cluster file system (XFS, ext4, NTFS, and others):**

Only one host can have both read and write access to a LUN. If several hosts have the read and write access simultaneously, data corruption can occur.

If you disable the access to a LUN for a host, the host remains in the list, so you can assign it a different access level if required.

6.3.2.1 Create a Host Masking Rule

To create a host masking rule:

1. Select **STORAGE > RAID LIST**.
2. Select a LUN.
3. In the *Host Masking Rules for LUN* section, click **Create Masking Rule**. A widget to create a host masking rule opens (Figure 52).

The screenshot shows the RAIDIX web interface. The top navigation bar includes 'STORAGE', 'NAS', 'HOSTS & INTERFACES', 'SYSTEM', and 'MONITORING'. The 'STORAGE' tab is active, and the 'RAID LIST' sub-tab is selected. Below the navigation bar, there is a breadcrumb trail: 'DASHBOARD > RAID LIST > SPAREPOOLS > DRIVES > DRIVES SCAN > RAID PROFILES > RAID IMPORT > PROXIED LUNS'. The main content area is titled 'LUN "common"'. Below the title, it says 'Host Masking Rules for LUN "common" are not created'. There are three columns: 'Host Alias', 'Access', and 'Host (WWN/GUID/IQN)'. Under 'Host Alias', there is a dropdown menu with 'alias1' selected and a link 'Add New Host Alias'. Under 'Access', there is a dropdown menu with 'Read Only' selected. At the bottom right, there are 'CREATE' and 'CANCEL' buttons.

Figure 52. LUN page. A Host Masking Rule creation

4. In the **Host Alias** field, select previously created host alias from the drop-down list. To learn more, see [Create a host alias](#).
5. If the host alias is not created yet, click **Add New Host Alias**.

The opened window allows you to select a host by the session (Figure 53).

Add New Host Alias

Alias Name

Select Hosts

Host (WWN/GUID/IQN)	Target Port	Node
<input type="checkbox"/> scst_local_host	scst_local_tgt	node1
	scst_local_tgt	node2

Figure 53. Add a New Host Alias window

- 5.1. Check appropriate hosts by their sessions.
- 5.2. In the **Alias Name** field, type the name of host alias.
- 5.3. Click **SAVE**.
6. Specify access level of host aliases to LUNs in the field **Access: Read Only** or **Read/Write**.
7. Click **CREATE**.



After creation or changing of LUN masking rules, it is necessary to rescan devices on initiator. Manually change LUN status to Offline and then back to Online.

6.3.2.2 Change a Default Masking Rule

The changing is available through CLI (to learn more about CLI syntax, see [Appendix A](#))

To check current host masking policy, run:

```
$ rdcli param system show
```


where a value of the `luns_visible_by_default` parameter defines default policy:

- the value **1** means that after creation, LUNs are available to all hosts;
- the value **0** means that after creation, LUNs are not available to any host.

To change masking policy, run:

```
$ rdcli param system modify -lv {0|1}
```

6.3.2.3 Delete a Host Masking Rule

1. Select **STORAGE > LUN LIST**.
2. Select a LUN.
3. In the *Host masking rules for LUN* section, click  in the string of the appropriate rule and select **Delete**.

4. To confirm, click **DELETE**.

6.3.3 Unmasking LUN

To unmask a LUN for all hosts:

1. Select **STORAGE > LUN LIST**.
2. Select a LUN.
3. On the LUN page:
 - If all targets have masks, click **Unmask LUN on all targets** (Figure 54).
 - If targets have different masks, on the **Mask LUN on all targets** button, click **:** and select **Unmask LUN on all targets** (Figure 55).

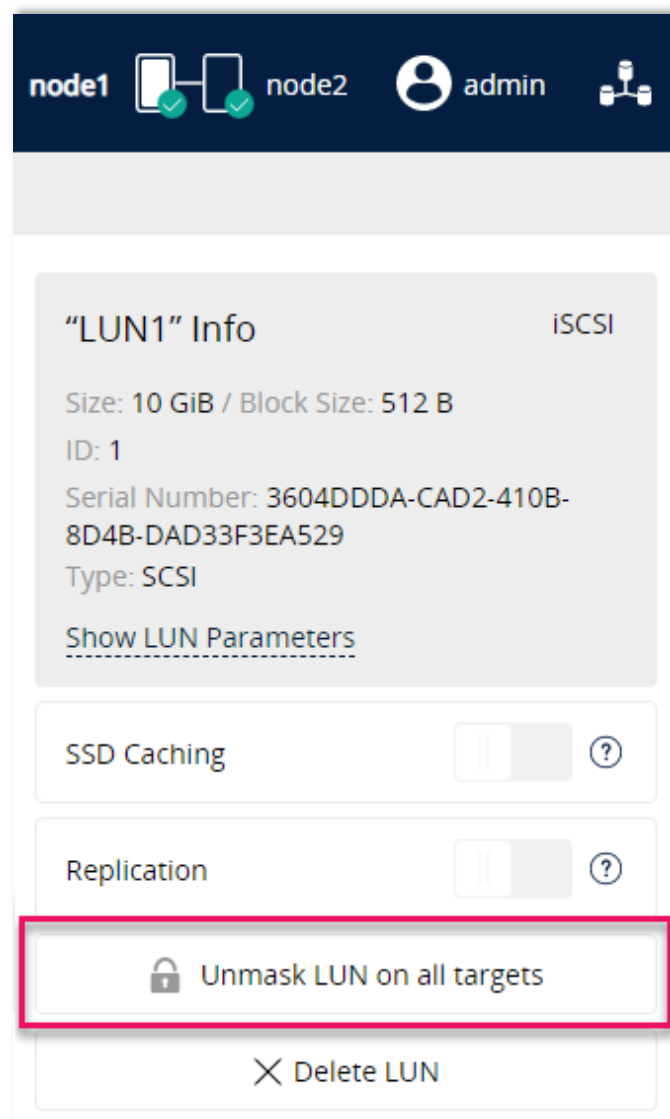


Figure 54. "Unmask LUN on all targets" button

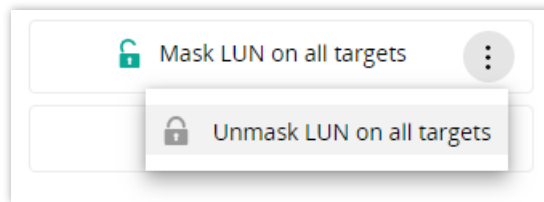


Figure 55. “Unmask LUN on all targets” button in case of different masks on targets

4. In the opened window, click **UNMASK** (Figure 56). LUN becomes invisible on hosts.

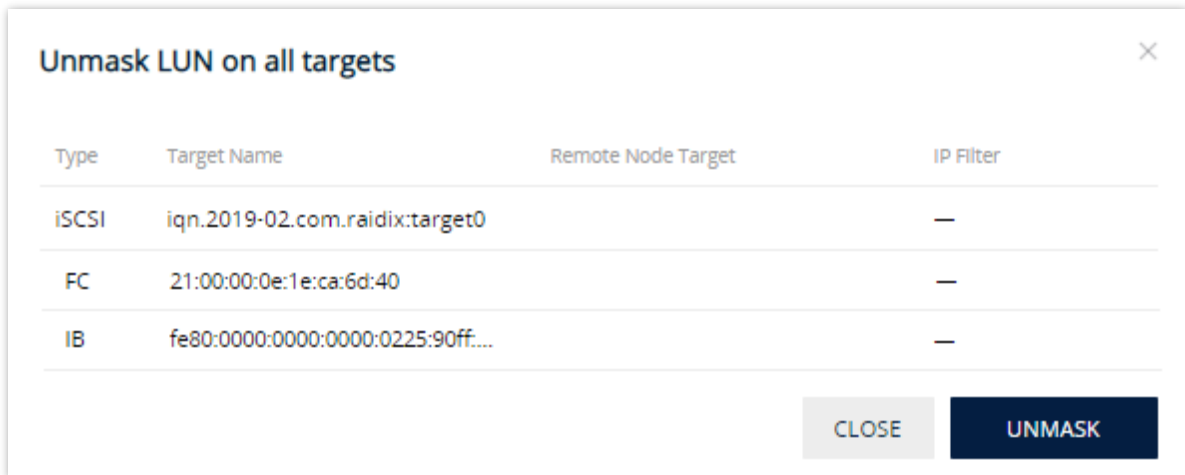


Figure 56. Confirmation of LUN unmasking

6.4 SSD Cache Setup

i Only for Generic RAID.

SSD cache is a high-performance buffer memory of limited size on SSDs. In comparison to HDDs, SSDs are better for random read and write workload, SSD cache provides a higher rate of access to the most often used data (“hot” data) on the hard drives as long as the workload is random. SSD cache works simultaneously with main cache.

In RAIDIX 5.2, SSD cache is a specially configured LUN (caching LUN). After the configuration, you can enable the caching LUN as SSD cache for the data LUNs (primary LUNs).

To use SSD caching:


1. Create a RAID and a caching LUN on it.
2. Set up SSD cache (caching LUN).
3. Create a RAID and a primary LUN on it.
4. Enable SSD caching for the primary LUN.

Parameters of the SSD caching mechanism

1. Storage system:

- To create a caching LUN, required [some amount](#) of RAM, depending on the LUN size.
 - To enable caching LUN as an SSD cache to a primary LUN, required *260 MiB* RAM.
2. Caching LUN:
- RAID parameters:
 - SSD RAID.
 - Engine – Generic RAID.
 - RAID 0 is only for read operations.
 - LUN parameters:
 - The LUN type is **NAS**.
 - The LUN doesn't have shares.
3. Primary LUN:
- RAID parameters:
 - Engine – Generic RAID.
 - LUN parameters:
 - To enable SSD cache, the value of the **SSD Caching** parameter is **on**.


Features of the SSD cache managing

 You should not use the SSD write cache (RWC) along with replication, as it will significantly reduce system performance.

- You can replicate a primary LUN.
 - Both a caching LUN and a replication metadata LUN can inherit to the same RAID. However, this configuration may reduce performance of the RAID with the primary LUN in case of sequential big-sized block I/O operations.
- During a reboot of a RAID with SSD cache, all primary LUNs with that SSD cache will be unavailable.
- For the DC Active/Active system configuration, create a RAID with caching LUN on both nodes.
- You can disable SSD cache only from a node, on which a RAID with a caching LUN is active.
- The Disabling of SSD cache or SSD caching may take a long time in case SSD cache contains data not flushed to primary LUN(s).
 - You can speed up the disabling, but data stored on SSD cache will be lost.

6.4.1 Configuring SSD Cache

You can use a LUN, created on SSD RAID (the RAID, which includes only SSDs) as SSD cache. The LUN must not have neither shares nor target masks.

 The system requires **260** MiB of RAM to add SSD cache to one LUN.

To configure SSD cache:

1. Select **STORAGE > LUN LIST**.
2. In the *SSD Cache* section, click **SSD Cache Settings** (Figure 57).

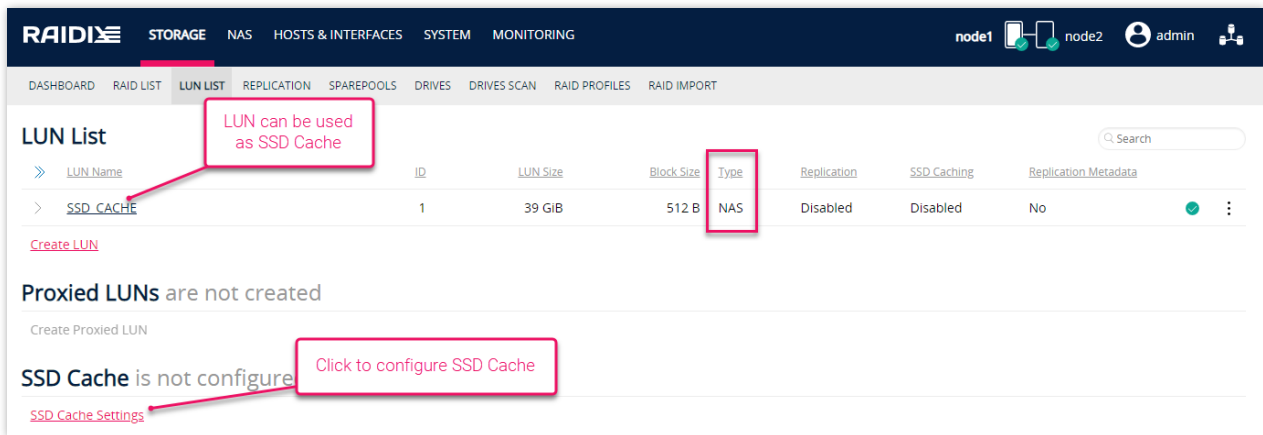


Figure 57. SSD Cache Settings control

3. In the opened window, select a LUN to be used as SSD cache and click **NEXT**.

i If a LUN created on RAID 0 is selected to be used as SSD cache, then SSD cache will be used for only read operations to ensure data integrity.

4. In the **Memory Size** field, set the RAM size to be used by SSD cache.

i *Memory size* parameter defines the maximum amount of memory to be used by SSD cache. Values of *Maximum amount of memory* and *Minimum amount of memory* parameters are calculated as follows:

Maximum amount of memory (in MiB) is calculated according to the formula:

$$\text{Maximum size} = 42 \cdot \text{size} / 1,024 + 260 + (60 \cdot \text{size} / 1,024) / 100$$

where *size* stands for LUN size in MiB.

Minimum amount of memory (in MiB) is calculated according to the formula:

$$\text{Minimum size} = 260 + 0.1 \cdot [42 \cdot \text{size} / 1,024 + (60 \cdot \text{size} / 1,024) / 100]$$

The default value, which is set automatically in the **memory size** field, is calculated according to the formula:

$$\text{Default size} = 260 + 0.25 \cdot [42 \cdot \text{size} / 1,024 + (60 \cdot \text{size} / 1,024) / 100]$$

If you choose a parameter value, which is smaller than the specified maximum amount, then the caching LUN will be partially filled with data. The smaller is the memory size value, the less LUN space will be used.

5. If you use the caching LUN only for reading operations, toggle the **Read Only** switch to the **On** position. Click **NEXT**.
6. Select LUNs for which you want to enable caching.
7. Click **COMPLETE**.

In the *SSD Cache* section, after cache configuration, the following information available:

Parameter	Description
LUN	Name of the caching LUN.
Node Name	Name of the node on which a RAID with the caching LUN is active.
RAM Usage	Max RAM size available for SSD Cache, in MiB.
MD Restoring	The value 1 stands for that SSD Cache Metadata is restoring on the remote node. Possible values: 0, 1 .
Read Only	SSD Cache is used only for read operations. Possible values: Yes, No .
Allocated	Total amount of used SSD Cache space.
RRC Allocation	Read-data percent in SSD Cache, in %.
RWC Allocation	Write-data percent in SSD Cache, in %.

To enable SSD Cache for a LUN:



In the DC system, RAID's of both the SSD Cache and the LUN must be active on the node from which enabling is performed.

Herewith, on the node that is passive for the RAID with LUN, the control that enables SSD cache will be inactive.

1. Select **STORAGE > LUN LIST**.
2. Select a LUN for which you are going to enable SSD Cache.
3. On the LUN page, enable SSD Cache by the switch **SSD Caching** (Figure 58).

To disable SSD Cache for a LUN:

1. Select **STORAGE > LUN LIST**.
2. Select a LUN with enabled SSD Cache.
3. On the LUN page, disable SSD Cache by the switch **SSD caching** (Figure 58).



SSD cache disabling may take a long time if it contains data not flushed to the primary LUN.

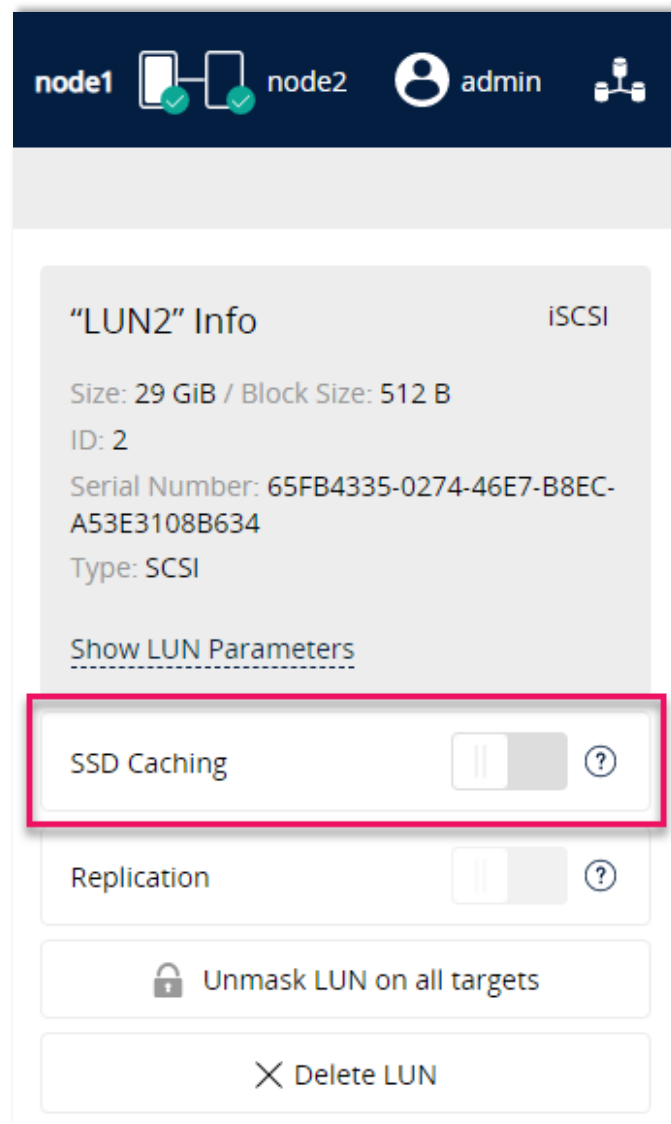


Figure 58. Enabling SSD Caching for the LUN

6.4.2 Disabling SSD Cache

1. Select **STORAGE > LUN LIST**.
2. In the *SSD Cache* section, click **DISABLE SSD CACHE**.
3. To confirm, click **DISABLE**.

i If SSD cache contains data not flushed to the RAID before disabling, then its deletion may take a long time due to flushing of cached data.

i If more than one LUN (main volume) is connected to the SSD cache, the time to disconnect one LUN from the SSD cache may be extended by flushing the data from the cache to each connected LUN.

6.5 Replication

Replication is the duplication of data on two RAIDIX systems.

Replication meets the following requirements:


- Business continuity.
- Increased data storage reliability.
- High data integrity without performance degradation.

6.5.1 Overview of the Replication

RAIDIX 5.2 supports replication of the One-to-One topology on the LUN level. Each LUN on one RAIDIX 5.2 system (in SC or DC mode) pairs with another LUN on the other RAIDIX 5.2 system (in SC or DC mode). Replication takes place via TCP/IP v4.

 Replication is only available for LUNs providing block access (SCSI type).

 We recommend using a dedicated channel for replication.

 You should not use the SSD write cache (RWC) along with replication, as it will significantly reduce system performance.

For replication, the administrator sets LUN on one storage system to the Primary role and LUN on the second storage system to the Secondary role. The roles can be changed by the administrator if necessary.

One node can have up to 32 main LUNs and up to 32 additional LUNs used in replication at the same time.

Up-to-date data is stored on the primary LUN, which collects all data and transfer the difference to the secondary LUN. The secondary LUN is not available for hosts.

Available replication protocols:

- *Synchronous* – initiator receives write confirmation only after data is safely replicated to the secondary LUN. Use of synchronous protocol ensures data integrity but may lead to degradation of performance. This type of replication is demanding in terms of channel capacity and network latency.
- *Asynchronous* – write is acknowledged on the initiator after data is written on the primary LUN. The primary storage tracks data storage and replicates in to the secondary LUN periodically. Asynchronous protocol helps to avoid performance deterioration with replication enabled.

6.5.2 Creating the LUN with Replication Metadata

Replication metadata is stored on a specifically dedicated LUN – metadata LUN. Create metadata LUN on both storage systems, participating in replication.

Any LUN created on RAID with redundant drives (any but RAID 0) and unmasked on all targets can be used as a metadata LUN. Metadata LUN is not available for users.



You can use replication on the systems with [SSD cache](#).

You can create the caching LUN and the LUN with replication metadata on one SSD RAID. The performance of this configuration can be lower in case of sequential read/write and large blocks than that with the caching LUN and the LUN with replication metadata are on different RAID.



For [asynchronous replication protocol](#), we recommend creating the replicated LUN and replication metadata LUN on different RAID.

To create a replication metadata LUN:

1. Create a LUN with the NAS type in accordance with the "Create a LUN" chapter.



Metadata LUN cannot be created on RAID 0.

Minimal size of Metadata LUN is 1 GiB. This size supports up to 32638 GiB of replication volumes size. As replication volumes size increases, you can extend Metadata LUN accordingly.

Create only one Metadata LUN on each storage, regardless of the number of LUN participating in replication.

2. Select **STORAGE > REPLICATION**.
3. In the *Replication Metadata* section, click the **Select LUN for Replication Metadata** control (Figure 59).

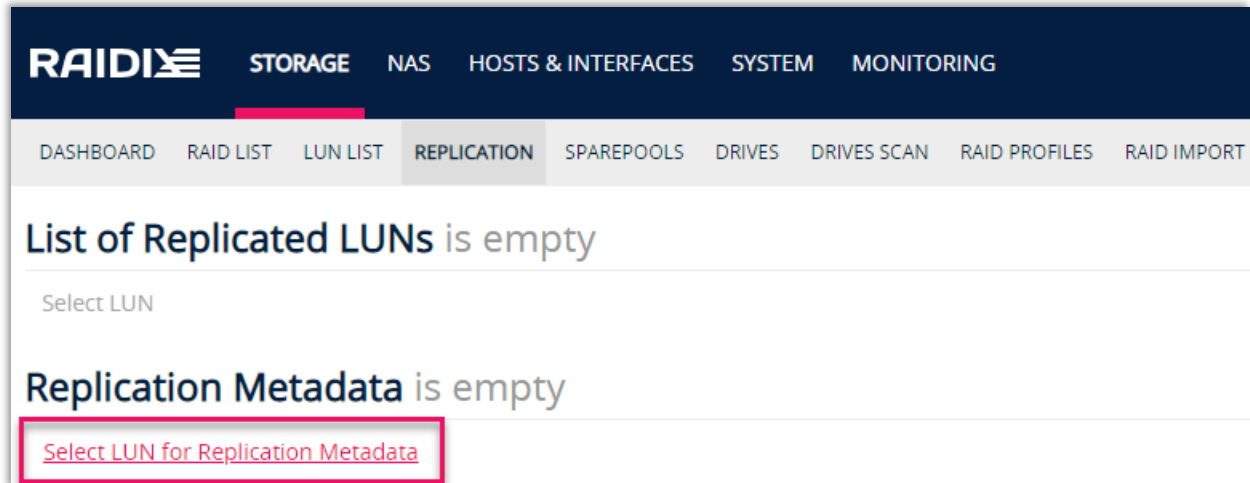


Figure 59. Selecting LUN for replication metadata

4. In the opened window, select LUN for replication metadata and click **SELECT**.

Repeat steps 1-4 on the second storage. After that, you can set up the Replication.

In DC system, in *Replication Metadata* section, you can see metadata LUNs from each node (Figure 60).

Replication Metadata				
LUN Name	Node Name	ID	Size	
m1	node1	2	2 GiB	✓
m2	node2	3	2 GiB	✓
Disable Replication Metadata				

Figure 60. List of metadata LUNs in DC system

6.5.3 Replication Settings

After the Metadata LUN is created, configure replication parameters. Replication is performed for two LUNs on different storages or nodes of one DC system. You can create LUNs participated in replication on any-level RAID.

! During replication setup or reconfiguration, the access of both LUNs involved in replication will be "read-only". Once you set the "Primary" role for one of the LUNs, the access will be resumed.

! If you change the roles of LUNs, change the LUN masking rules as well.

i For **asynchronous replication protocol**, we recommend creating the replicated LUN and replication metadata LUN on different RAIDs.

Requirements and restrictions to replication LUNs:

- LUN type must be SCSI.
- LUNs participating in Replication must be of the same size.
- LUNs participating in Replication must not have the name *all*.
- Secondary LUN must be unavailable for hosts (set up masking policy).
- Each node can contain up to 32 primary LUNs and up to 32 secondary LUNs simultaneously.
- You can create replication LUNs and the metadata LUN on the same RAID.

You can setup, manage, and monitoring the replication:

- on the LUN page;
- on the **STORAGE | REPLICATION** page.

6.5.3.1 Replication Settings for SC Systems

To configure Replication:

Given instruction uses **REPLICATION** page.

1. Create *identical-sized* LUNs on each storage participating in Replication.
2. Enable and set up Replication on both systems:
 - 2.1. Select **STORAGE > REPLICATION**.
 - 2.2. In the List of Replicated *LUNs* section, click **Select LUN** (Figure 61).

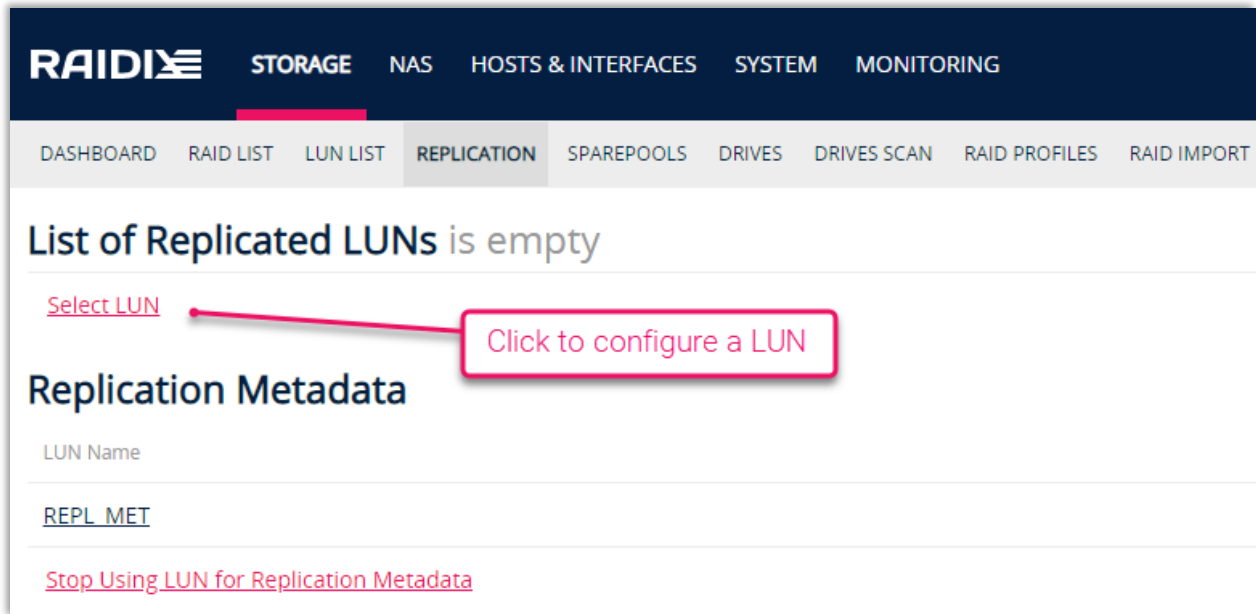


Figure 61. A control to start the configuration of Replication on the REPLICATION page

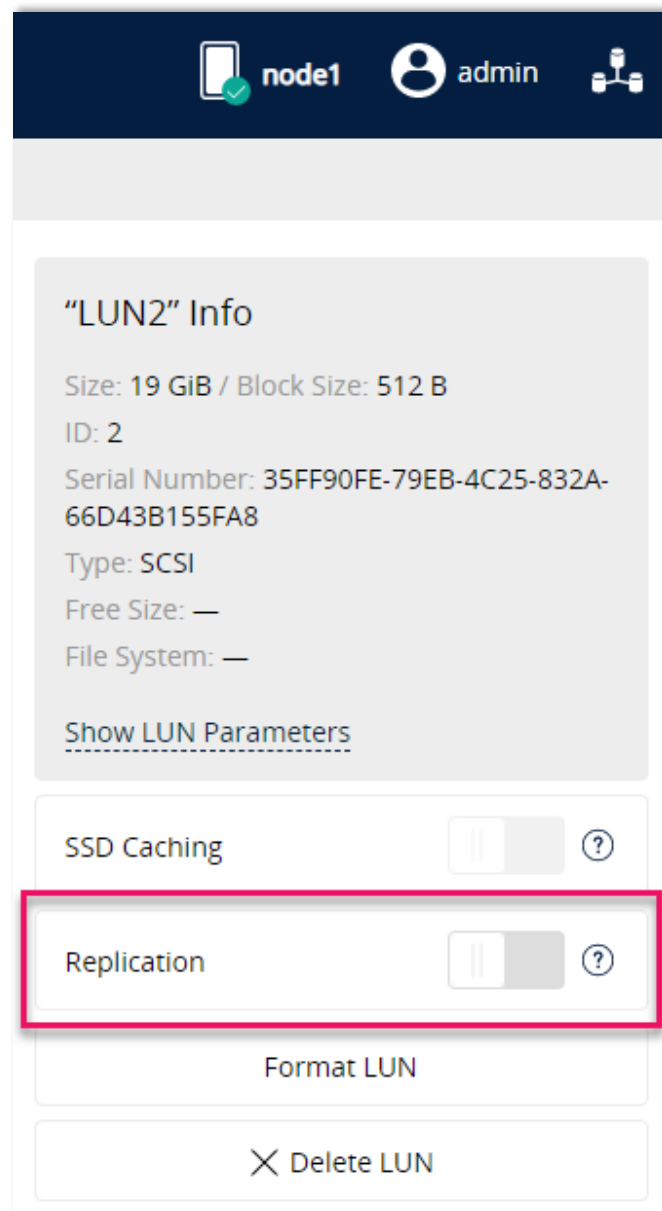


Figure 62. Enabling the Replication on the LUN page

- 2.3. Select a LUN available for Replication and click **NEXT**.
- 2.4. In the opened window, enter the following parameters and click **APPLY**:
 - **Local IP** – IP address of the local storage;
 - **Remote IP** – IP address of the remote storage (the second site of replication);
 - **Remote LUN name** – the name of the LUN on the remote storage;
 - **Protocol** – synchronous (Sync) or asynchronous (Async) replication.

i Performance increase from using the asynchronous protocol may not be sufficient. We recommend you to use the synchronous replication mode, which ensures data integrity in case a node with the primary LUN fails.

- 2.5. Click **COMPLETE**.

After you configured first LUN, information about Replication with the state Inconsistent/DUnknown will shows (Figure 64).

Replicated LUN List										
LUN Name	ID	Size	Local IP	Remote IP	Remote LUN Name	Protocol	Role	Data status	Replica Match	State
LUN2	2	19 GiB	172.16.22.23	172.16.22.57	LUN2	Sync	Secondary	Inconsistent/DUnknown	0%	Off

Figure 63. Replication info after the first LUN was configured on the REPLICATION page

The *State* field shows replication process state. Possible values and their meanings are given in the table:

State	Meaning
Off	The LUN is not replicated over this connection, since the connection is not "Connected".
Established	All writes to that LUN are replicated online. This is the normal state.
StartingSyncS	Full synchronization, initiated by the administrator, is just starting. The next possible states are: "SyncSource" or "PausedSyncS".
StartingSyncT	Full synchronization, initiated by the administrator, is just starting. Next state: "WFSyncUUID".
WFSyncUUID	Synchronization is about to begin. Next possible states: "SyncTarget" or "PausedSyncT".
SyncSource	Synchronization is currently running, with the local node being the source of synchronization.
SyncTarget	Synchronization is currently running, with the local node being the target of synchronization.
VerifyS	The local node is the source of an ongoing synchronization, but synchronization is currently paused. This may be due to a dependency on the completion of another synchronization process, or due to synchronization having been manually interrupted by "drbdadm pause-sync".
VerifyT	The local node is the target of an ongoing synchronization, but synchronization is currently paused. This may be due to a dependency on the completion of another synchronization process, or due to synchronization having been manually interrupted by "drbdadm pause-sync".
PausedSyncS	On-line device verification is currently running, with the local node being the source of verification.

State	Meaning
PausedSyncT	On-line device verification is currently running, with the local node being the target of verification.
Ahead	Data replication was suspended, since the link cannot cope with the load. This state is enabled by the configuration "on-congestion" option.
Behind	Data replication was suspended by the node with the secondary LUN, since the link cannot cope with the load. This state is enabled by the configuration "on-congestion" option on the peer node.

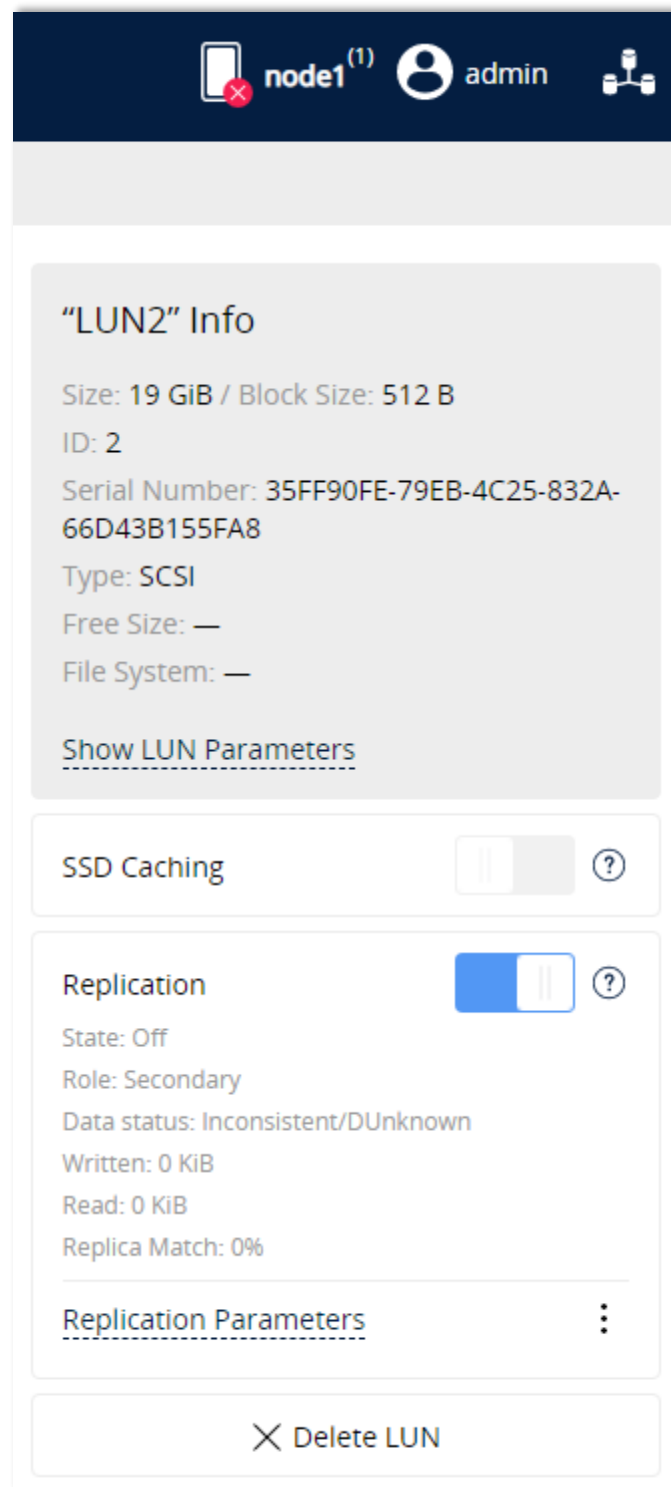
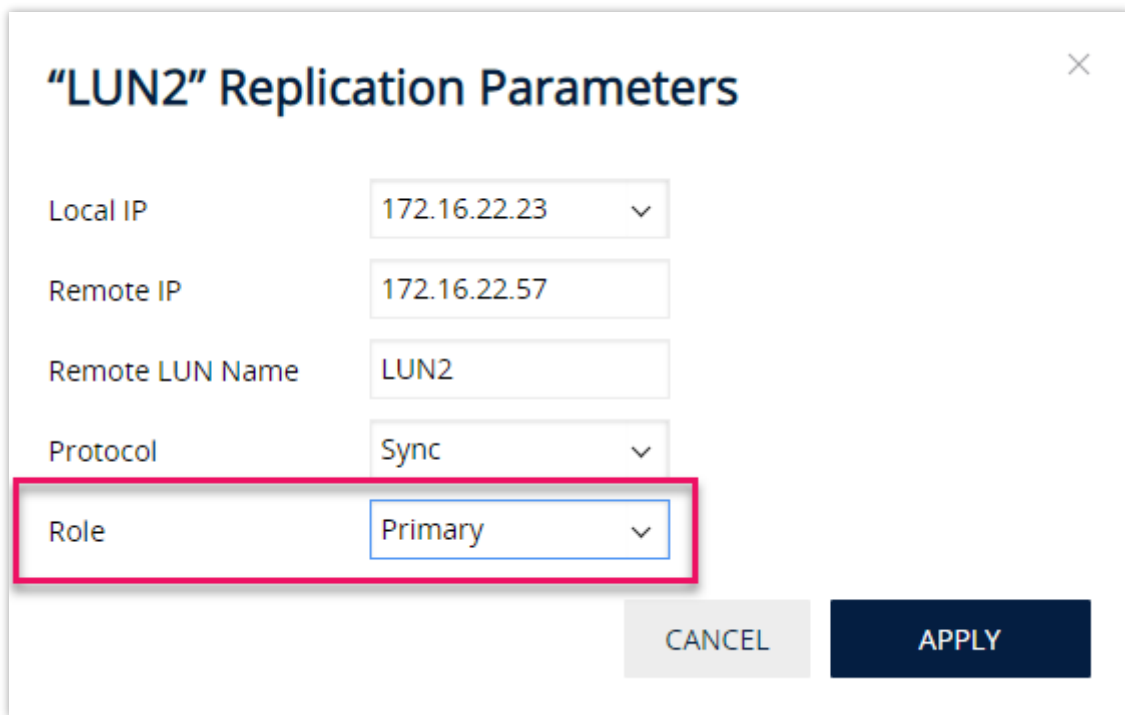


Figure 64. Replication info after the first LUN was configured on the LUN page

3. Repeat the previous step on the second storage.
4. Set the **Primary** role to one of the LUNs:
 - 4.1. On the storage intended for primary LUN, click **Replication Parameters** on the LUN page. The window *Replication Parameters* opens (Figure 65).
 - 4.2. In the **Role** field, select the value **Primary**.
 - 4.3. Click **APPLY**.



"LUN2" Replication Parameters

Local IP: 172.16.22.23

Remote IP: 172.16.22.57

Remote LUN Name: LUN2

Protocol: Sync

Role: Primary

CANCEL APPLY

Figure 65. Selecting replication role for the LUN

After you selected the primary LUN, data synchronization will start. The progress of synchronization is displayed on the LUN page.

- i** During the initial synchronization, we don't recommend to perform:
- failover and failback;
 - suspend and resume replication.
- It will cause starting over of the synchronization.

- i** Replication will start over after:
- migration of a RAID with LUN that participate in replication;
 - disabling of replication and enabling it again (even with metadata LUN kept untouched).

For using replication in isolate networks, you should allow data transfer through replication ports.

Port's number calculates as "11000 + LUN ID"

To get LUN ID, see the */D* parameter on the **STORAGE | REPLICATION** page.

6.5.3.2 Replication Settings for DC system

For DC systems, configure replication as for SCs with the following specifics:

- !** When the system in DC mode with replication, you can't Migrate a RAID with replication LUN or metadata LUN.

- On each node of the DC-system, which has a primary or secondary LUN, configure one or more virtual IP addresses.
- When you are setting up replication, for the **Local IP** parameter, enter a VIP that is active on the local node of the DC system.
- The systems for replication must be connected through a switch.

Combinations of systems for replication

We recommend the following combinations of systems:

- From *DC Active/Passive* to *SC*.
Replicated LUNs are on RAIDs that are active on only one node.
- From *DC Active/Active* to *SC*.
Replicated LUNs are on RAIDs that are active on different nodes.
- From *DC Active/Active* to *DC Active/Active*.
The main and additional LUNs are on each node of each DC-system.

Configuring the storage for replication from the DC Active/Passive system to the SC system

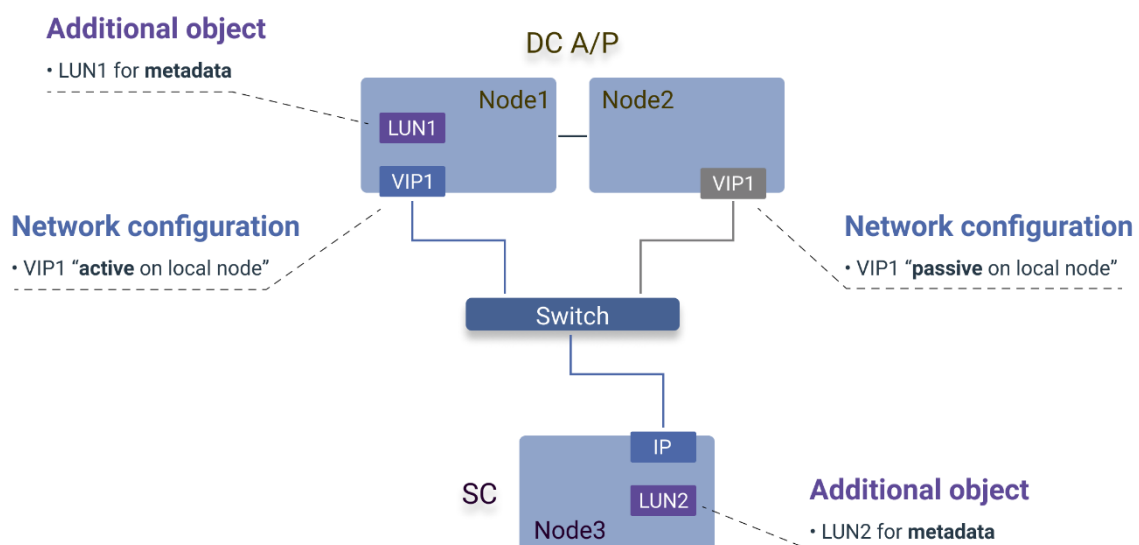


Figure 66. Scheme of a storage for replication from DC Active/Passive to SC

Configuring the network and creating additional storage objects for replication from DC Active/Passive to SC

Active DC node (Node1)	
LUN for replication metadata (LUN1)	On a RAID that is active for this node, create a LUN for metadata.
VIP "active on local node" (VIP1)	Create a VIP with the value Active for the parameter State on Local Node .

Passive DC node (Node2)

VIP "passive on local node" (VIP1)	Create a VIP with the value Passive for the parameter State on Local Node . Values Virtual IP address and Netmask must match with the values of the created VIP on Node1.
------------------------------------	--

SC node (Node3)

LUN for replication metadata (LUN2)	Create a LUN for metadata on this node.
-------------------------------------	---

Specifics of replication configuration

- On Node1, for the parameter **Local IP**, enter the value VIP1.
- On Node3, for the parameter **Remote IP** enter the value VIP1.

Configuring the storage for replication from the DC Active/Active system to the SC system

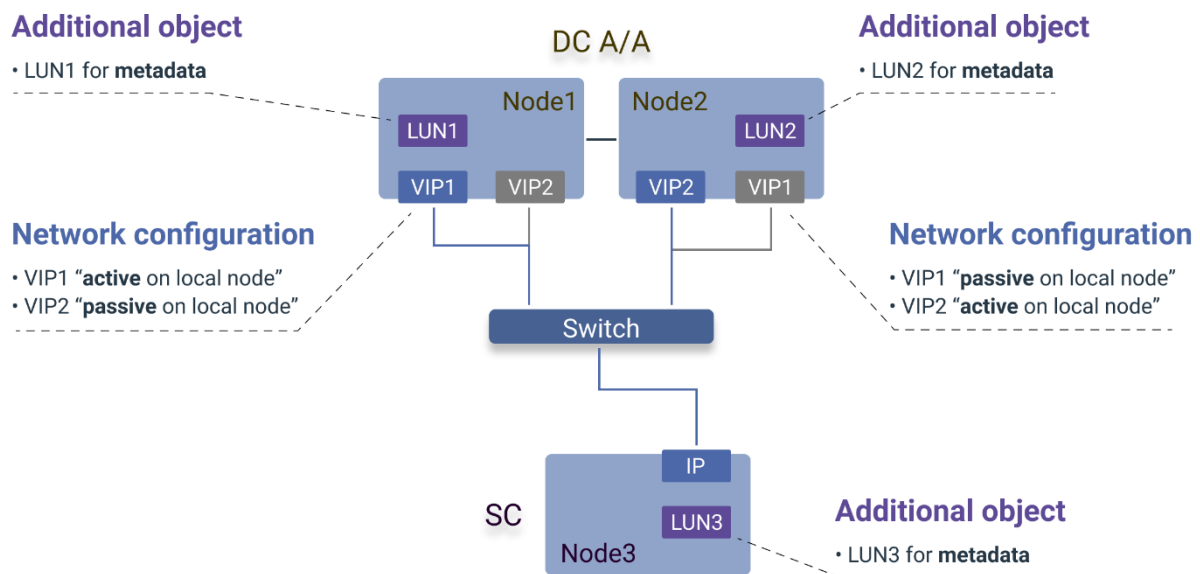


Figure 67. Scheme of a storage for replication from DC Active/Active to SC

Configuring the network and creating additional storage objects for replication from DC Active/Active to SC

Active DC node (Node1)

LUN for replication metadata (LUN1)	On a RAID that is active for this node, create a LUN for metadata.
VIP "active on local node" (VIP1)	Create a VIP with the value Active for the parameter State on Local Node .
VIP "active on local node" (VIP2)	Create a VIP with the value Passive for the parameter State on Local Node . Values Virtual IP address and Netmask must match with the values of the created VIP2 on Node2.

Active DC node (Node2)

LUN for replication metadata (LUN2)	On a RAID that is active for this node, create a LUN for metadata.
VIP "active on local node" (VIP1)	Create a VIP with the value Passive for the parameter State on Local Node . Values Virtual IP address and Netmask must match with the values of the created VIP1 on Node1.
VIP "active on local node" (VIP2)	Create a VIP with the value Active for the parameter State on Local Node .

SC node (Node3)

LUN for replication metadata (LUN2)	Create a LUN for metadata on this node.
-------------------------------------	---

Specifics of replication configuration

- When configuring *a replication* on Node1, for the parameter **Local IP**, enter the value VIP1.
- When configuring *a replication* on Node2, for the parameter **Local IP**, enter the value VIP2.
- When configuring *a replica* on Node3, for the parameter **Remote IP**, enter the value of the active VIP of the corresponding node.

Configuring the storage for replication between DC Active/Active systems

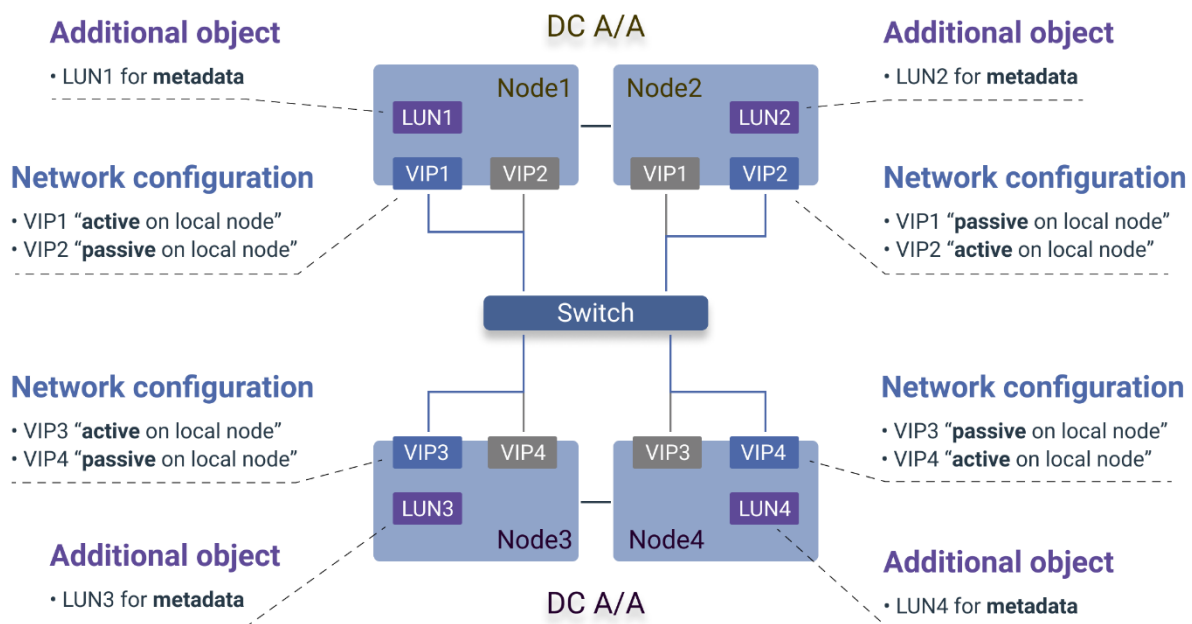


Figure 68. Scheme of a storage for replication from DC Active/Active to DC Active/Active

Configuring the network and creating additional storage objects for replication from DC Active/Active to DC Active/Active

Active DC node (Node1)

LUN for replication metadata (LUN1)	On a RAID that is active for this node, create a LUN for metadata.
VIP "active on local node" (VIP1)	Create a VIP with the value Active for the parameter State on Local Node .
VIP "active on local node" (VIP2)	Create a VIP with the value Passive for the parameter State on Local Node . Values Virtual IP address and Netmask must match with the values of the created VIP2 on Node2.

Active DC node (Node2)

LUN for replication metadata (LUN2)	On a RAID that is active for this node, create a LUN for metadata.
VIP "active on local node" (VIP1)	Create a VIP with the value Passive for the parameter State on Local Node . Values Virtual IP address and Netmask must match with the values of the created VIP1 on Node1.
VIP "active on local node" (VIP2)	Create a VIP with the value Active for the parameter State on Local Node .

Active DC node (Node3)

LUN for replication metadata (LUN3)	On a RAID that is active for this node, create a LUN for metadata.
VIP "active on local node" (VIP3)	Create a VIP with the value Active for the parameter State on Local Node .
VIP "active on local node" (VIP4)	Create a VIP with the value Passive for the parameter State on Local Node . Values Virtual IP address and Netmask must match with the values of the created VIP4 on Node4.

Active DC node (Node4)


LUN for replication metadata (LUN4)	On a RAID that is active for this node, create a LUN for metadata.
VIP "active on local node" (VIP3)	Create a VIP with the value Passive for the parameter State on Local Node . Values Virtual IP address and Netmask must match with the values of the created VIP3 on Node3.
VIP "active on local node" (VIP4)	Create a VIP with the value Active for the parameter State on Local Node .

Specifics of replication configuration

- When configuring *a replication* on a node, for the parameter **Local IP**, use the value of an active VIP of this node.
- When configuring *a replica* on a node, for the parameter **Remote IP**, use the value of the active VIP of the node that has a primary LUN.

6.5.3.3 Changing Replication Nodes Addresses

If after replication is configured, replication channel has changed, and IP addresses of replication sites have changed, use the following order to reconfigure replication nodes addresses:

1. On one storage, select **STORAGE > REPLICATION**.
2. In a string with LUN for replication, click  and select **Replication Parameters**.
3. In the opened window, specify **Remote IP** and click **APPLY**.
4. Move to the second storage.
5. Select **STORAGE > REPLICATION**.
6. In the *Replication* section, click **Replication Parameters**. The *Replication Parameters* page opens.
7. In the opened window, specify **Local IP** and **Remote ID** and click **APPLY**.
8. Move back to the first storage. Specify **Local IP** in the *Replication Parameters* and click **APPLY**.

If the IP address of only one of the storages has changed within the same subnetwork, at first specify the new **Remote IP** in *Replication Parameters* on the remote storage. Then specify **Local IP** on the local storage.

6.5.4 Managing Replication

For replication management, the following actions available:

- Suspend and resume replication.
- Start consistency check.
- Replace data with a replica.
- Change VIPs (when DC is the part of replication).



We recommend the following procedure to change a VIP:

1. Create a new VIP.
2. Change an old VIP to the new one in replication settings.
3. Delete the old VIP.

6.5.4.1 Suspension and Resumption of Replication

As synchronization of data on the primary and the secondary nodes starts right after replication is configured, so it may affect performance. You can suspend the replication data synchronization during intense I/O operations and resume at an appropriate time.



Always suspend replication on the secondary storage before rebooting or turning off the power. After each failure of the secondary node, start the consistency check. If some data remains unsynchronized, see [Troubleshooting](#).

To suspend data synchronization of the Replication:

1. On the primary storage, select **STORAGE > REPLICATION**.
2. In the *List of Replicated LUNs* section, click  and select **Suspend Replication** (Figure 69).

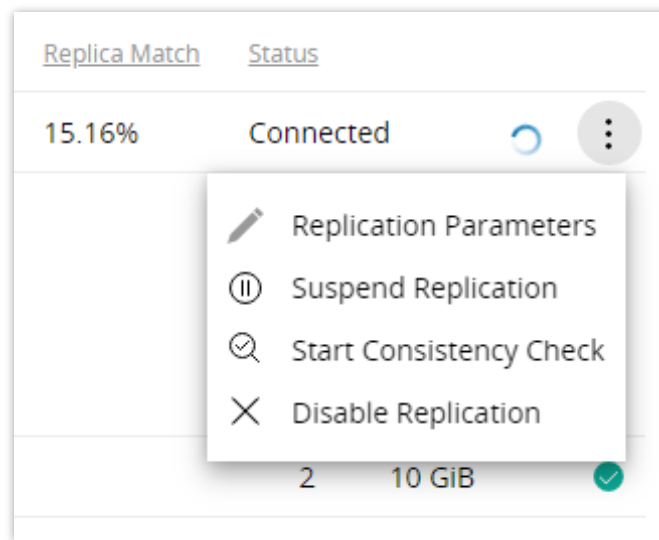



Figure 69. Manage options of the Replication on the REPLICATION page

To resume data replication synchronization:


1. On the primary storage, select **STORAGE > REPLICATION**.
2. In the *List of Replicated LUNs* section, click  and select **Resume Replication**.

6.5.4.2 Consistency Check

Consistency check enables comparison of data in the primary and the secondary LUNs and detection of possible data integrity failures.

You can suspend and resume consistency check at any time.


To start consistency check:

1. On the main storage, select **STORAGE > REPLICATION**.
2. In the *List of Replicated LUNs* section, click  and select **Start Consistency Check**.

6.5.4.3 Replacing Data with Replica

If one of the storages contains inconsistent data, replace it with replica from the second replication site. Perform this action on the primary or the secondary LUN.

To replace data with its replica:

1. On the storage with inconsistent data, select **STORAGE > REPLICATION**.
2. In the *List of Replicated LUNs* section, click  and select **Replace Data with Replica**.
3. To confirm, click **REPLACE**.

6.5.4.4 Changing Replication Roles

Administrator can change roles in a replication pair so that the primary LUN becomes secondary, and vice versa. Role changing is available after data synchronization is complete.

To change replication roles:

1. On the storage with primary LUN, open **Replication Parameters** of the LUN:
 - 1.1. Select **STORAGE > REPLICATION**.
 - 1.2. In a string with replicated LUN, click **:** and select **Replication Parameters**.
2. In the opened *Replication Parameters* window, change the **Role** parameter value to **Secondary** and click **APPLY**.

The system will show attention message.

3. On the second storage, open replication parameters of the LUN.
 - 3.1. Select **STORAGE > REPLICATION**.
 - 3.2. In a string with replicated LUN, click **:** and select **Replication Parameters**.
4. In the opened *Replication Parameters* window, change the **Role** parameter value to **Primary** and click **APPLY**.

Data synchronization between replicated LUNs starts.

6.5.4.5 Changing Replication Protocols

RAIDIX 5.2 supports synchronous and asynchronous replication protocols.

To change the replication protocol:

1. On the storage with primary LUN, open *Replication Parameters* of the LUN:
 - 1.1. Select **STORAGE > REPLICATION**.
 - 1.2. In a string with replicated LUN, click **:** and select **Replication Parameters**.
2. In the opened *Replication Parameters* window, change the **Protocol** parameter value and click **APPLY**.
3. On the second storage, open replication parameters of the LUN.
 - 3.1. Select **STORAGE > REPLICATION**.
 - 3.2. In a string with replicated LUN, click **:** and select **Replication Parameters**.
4. In the opened *Replication Parameters* window, change the **Protocol** parameter value and click **APPLY**.

6.5.4.6 Disabling of Replication

1. On any storage participated in the Replication, select **STORAGE > REPLICATION**.
2. In the *List of Replicated LUNs* section, in the replicated LUN line, click **:** and select **Disable Replication**.
3. To confirm, click **DISABLE**.

6.5.5 Disabling the LUN with Replication Metadata

If no LUNs participate in replication, it is possible to disable the Metadata LUN.

To disable Metadata LUN:

1. Select **STORAGE > REPLICATION**.
2. In the *Replication Metadata* section, click **Stop Using LUN for Replication Metadata** (Figure 70).
3. To confirm, click **DISABLE**.

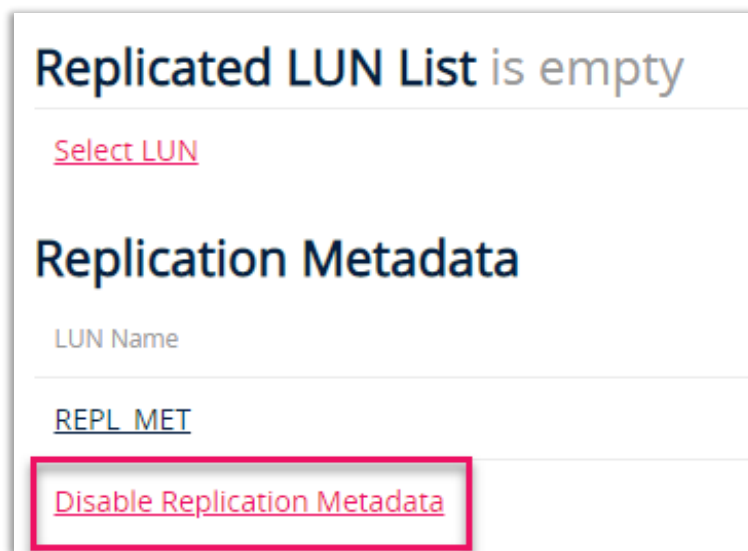


Figure 70. Disabling LUN from replication metadata

6.6 LUN Deletion

You can delete a LUN from:

- the LUN page;
- the RAID page;
- the **STORAGE | LUN LIST** page.

To delete a LUN from the **LUN LIST** page:

1. Select **STORAGE > LUN LIST**.
2. In the *LUN List* section, in the line of the LUN to be deleted, click **:** and select **Delete**.
3. To confirm, click **DELETE**.



If a LUN you try to delete is an SSD-cached LUN or a caching LUN, the deletion might take a long time because of the flushing of cached data.

There is an option of quick deletion of such LUNs without cache flushing. However, data integrity is not guaranteed because the deleted data will be lost.

7. NETWORK ATTACHED STORAGE (NAS)

RAIDIX 5.2 software provides an ability to work with Network Attached Storage – NAS. Network system architecture represents a NAS server combined with storage system on RAIDIX platform and interacts with client computers via SMB/CIFS, NFS, FTP, or AFP protocols.

Main advantages of working with NAS on RAIDIX 5.2 platform include:

- lower exploitation cost compared to SAN;
- work over the local Ethernet network;
- possibility of file sharing – the simultaneous user access to large volumes of data.

NAS functionality includes:

- An ability to create and edit public folders – Shares with adjustable parameters (path, protocol, visibility, host selection, visibility).
- An ability to set the WORM (write once read many) parameter and ACL (Access Control List) support for SMB shares is implemented. Setting up WORM allows an administrator to prohibit file modification after its creation for an extended period. Thus, it is impossible to delete the file, to modify its content, and to change its metadata.
- An ability to specify the amount of LUN size, available for the particular share users or user groups ([quotas](#)).



To start working with NAS in DC mode, assign a virtual IP address to a network storage before creating shares. In the case of nodes failover, RAIDs migrate to another node together with a virtual IP address, so the share is always available for a user.

To manage NAS settings, open pages under the **NAS** menu (Figure 71).

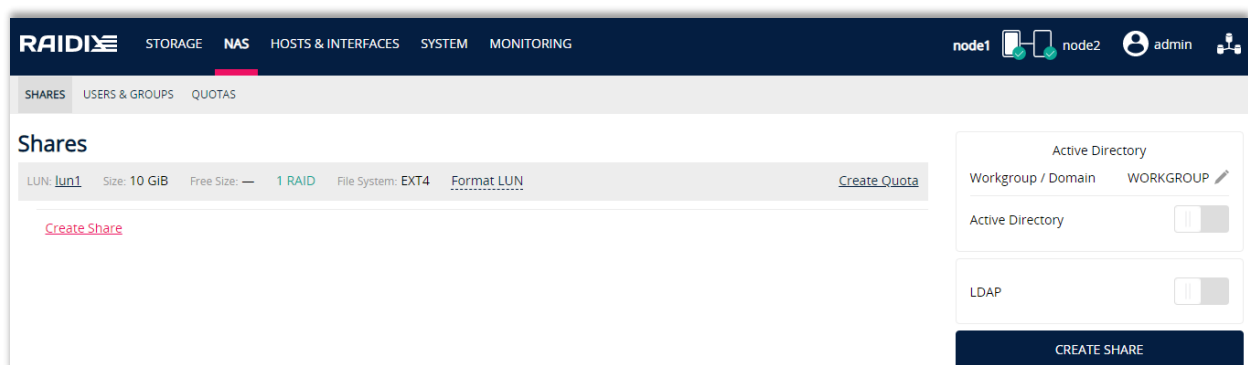


Figure 71. SHARES page

7.1 Creating a NAS Share

The basic elements of working with NAS in RAIDIX 5.2 are data folders, available for a custom list of users, in other words, shared folders (Shares). Organize access to shares via SMB, NFS, FTP, and AFP

protocols. In the case of using SMB protocol, setting up the Active Directory user parameters is available.

Creating the share provides users a file access to storage data.

i You can manage access to one directory over only one NAS protocol.

You can create a NAS share through:

- the *widget* **Create Share** on the page **NAS | SHARES**;
- the *wizard* **CREATE SHARE** on the page **NAS | SHARES**;

To create a share via the *widget*:

1. Create a RAID and a LUN of the NAS type. To learn more, see [Create a RAID](#) and [Create a LUN](#).
2. Select **NAS > SHARES**.
3. Click **Format LUN** and select the file system to format with:
 - XFS;
 - EXT4.

! To format a LUN to the EXT4 file system, the LUN size must not exceed **260,096** GiB.

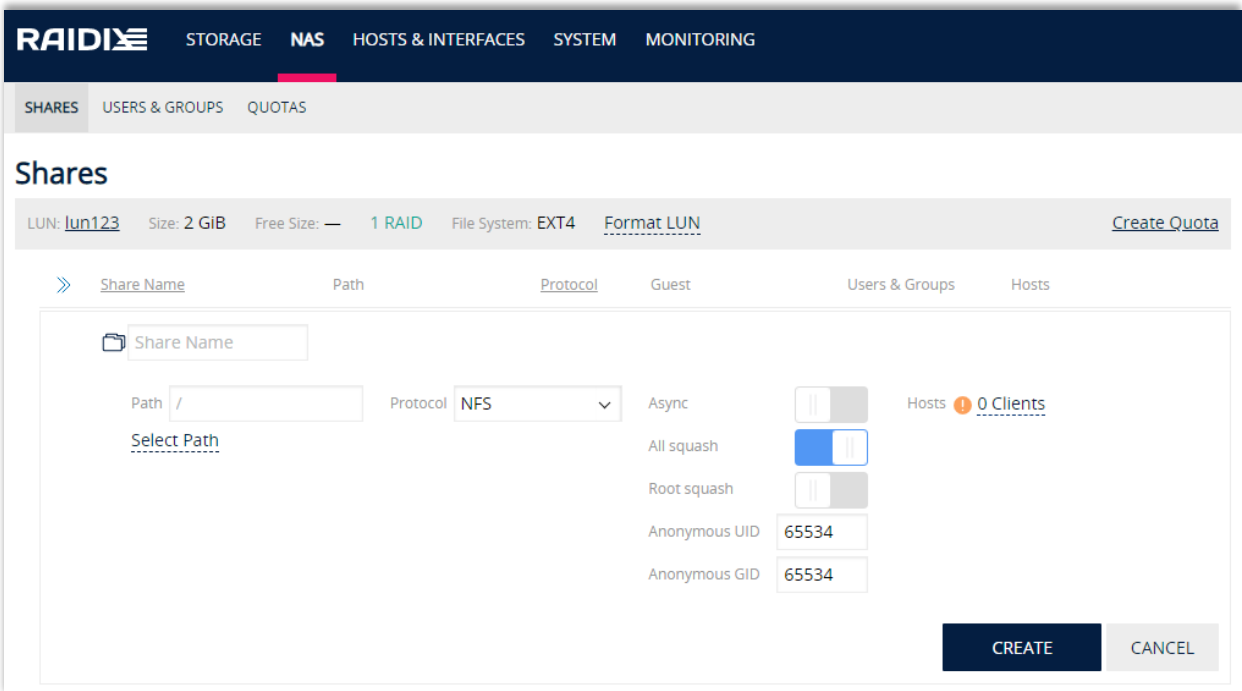
! XFS may shut down if an XFS share has no free space.
When using XFS, leave at least 5% free space.

i Formatting of large LUNs might take a long time. To accelerate the process, try changing the value of the parameter *Max write back sequential* from **8** to **64**.

i When you need [to format a large LUN with XFS on an HDD ERA RAID](#), to reduce the formatting time, you may wait until the RAID initialization is complete or set initialization priority to **0** while formatting.

4. Click **Create Share**. The widget for share creation opens (Figure 72).



i When a share is just created, NAS services restart, which may cause a short access loss to other shares.





The screenshot shows the RAIDIX web interface. The top navigation bar includes 'STORAGE', 'NAS' (selected), 'HOSTS & INTERFACES', 'SYSTEM', and 'MONITORING'. Below this, the 'SHARES' tab is active. The 'Shares' section has a sub-header with 'LUN: lun123', 'Size: 2 GiB', 'Free Size: —', '1 RAID', 'File System: EXT4', and 'Format LUN'. A 'Create Quota' link is on the right. The main form area has tabs for 'Share Name', 'Path', 'Protocol', 'Guest', 'Users & Groups', and 'Hosts'. The 'Path' tab is selected, showing a 'Share Name' input field, a 'Path' input field with a 'Select Path' link, a 'Protocol' dropdown set to 'NFS', and several toggle switches for 'Async', 'All squash', and 'Root squash'. Below these are input fields for 'Anonymous UID' and 'Anonymous GID', both set to '65534'. A 'Hosts' section shows '0 Clients'. At the bottom right are 'CREATE' and 'CANCEL' buttons.

Figure 72. Widget for creating a share with selected NFS protocol

5. Fill in the following information and click **CREATE**.

Parameter	Description
Share Name	Name of the share.
Path	<p>A path to the share on the system.</p> <div>  <p>For all shares <i>on one LUN</i>, the Path parameters must be different, and the Path of one share must not be a subfolder of the Path of another share.</p> </div>
Protocol	<p>Type of the protocol, by which access to a folder on the NAS server is provided.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • NFS – Network File System; • SMB – Server Message Block; • FTP – File Transfer Protocol; • AFP – Apple Filing Protocol. <div>  <p>Do note that if the host already has an SMB share mounted, after creating a new SMB share you will be able to mount it on the host in about 5 minutes. To learn more, see Troubleshooting.</p> </div>

Parameter	Description
Async	<p>Only for NFS.</p> <p>Sync or async writing mode for the NFS share.</p> <p>In sync mode, all changes to the filesystem are immediately flushed to drive; the respective write operations are pending.</p> <p>In async mode, the system buffers write operations and optimizes the actual writes; write operations are running continuously.</p> <p>The default: disabled (sync).</p>
All Squash	<p>Only for NFS.</p> <p>All connections (including root) to the shared folder are made from an anonymous user.</p> <p>The default: enabled.</p>
Root Squash	<p>Only for NFS.</p> <p>When enabled, the root user on the client machine connects to the shared folder as an anonymous user.</p> <p>The default: disabled.</p> <p>When disabled, gives the root user on the client machine access to the shared folder if the All squash parameter is disabled.</p>
Anonymous UID	<p>Only for NFS.</p> <p>UID of the anonymous user from whom the connections to the shared folder will go when All squash or Root squash is enabled.</p> <p>Possible values: from 1 to 4294967294.</p> <p>The default: 65534.</p>
Anonymous GID	<p>Only for NFS.</p> <p>GID of the anonymous group from whom the connections to the shared folder will go when All squash or Root squash is enabled.</p> <p>Possible values: from 1 to 4294967294.</p> <p>The default: 65534.</p>
Browseable	<p>Only for SMB.</p> <p>The option allows customizing the share visibility to the host in the list of available shares.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • Yes – display the shared folder; • No – not to display the shared folder.

Parameter	Description
WORM (Write Once Read Many)	<p>Only for SMB.</p> <p>Enabling this option allows prohibiting the user to modify or delete the file after its creation for extended period.</p> <ul style="list-style-type: none"> Yes – enable the parameter Files in the folder are read-only and cannot be modified. By default, files become unavailable for modification/deletion after 1 second (the WORM Grace period) after the moment of their creation and within 60 months (R/O period). Specify the R/O period (in months) and WORM Grace period (in seconds) to change these values. No – parameter is disabled. <div style="border: 1px solid #f00; padding: 5px; margin-top: 10px;"> <p> WORM Grace period starts from the moment when file is created, and if you are going to copy a big file, be sure that you will have enough WORM Grace period to perform modifications.</p> </div> <div style="border: 1px solid #000; padding: 5px; margin-top: 10px;"> <p> The WORM attribute of a share doesn't inherit for nested shares.</p> </div>
Guest	<p>Except NFS.</p> <p>Parameter specifies guest user access to the shared folder.</p> <p>Possible values:</p> <ul style="list-style-type: none"> Disabled – no access; Read Only; Read/Write.
Users & Groups	<p>Except NFS.</p> <p>Option allows selecting share users and groups and setting their permissions.</p>
Hosts	<p>Except FTP.</p> <p>Option allows specifying a list of hosts that will have access to the shared folder.</p> <p>Hosts are added by IP address or name with the following access types:</p> <ul style="list-style-type: none"> For all the protocols: <ul style="list-style-type: none"> All – provides access for all hosts. For NFS: <ul style="list-style-type: none"> Read/Write; Read Only. For SMB and AFP: <ul style="list-style-type: none"> Deny; Allow.

6. Click **CREATE** (Figure 73).

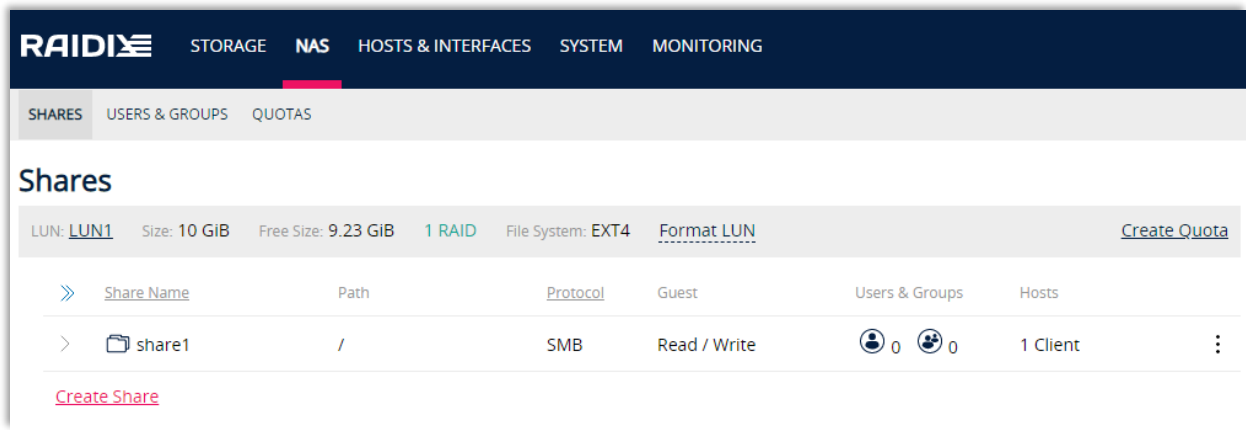


Figure 73. Shares List - created SMB share



In the list of LUNs for shares creation, only LUNs that are not available for iSCSI, InfiniBand, Fibre Channel targets are displayed. Before creation a share on a particular LUN, verify the LUN is not available through any of these protocols.

7.2 Managing NAS Users and Groups

7.2.1 Create Users and Groups

To create a NAS user:

1. Select **NAS > USERS & GROUPS**.
2. Click **Create User**. The pane for a new user or group creation opens (Figure 74).

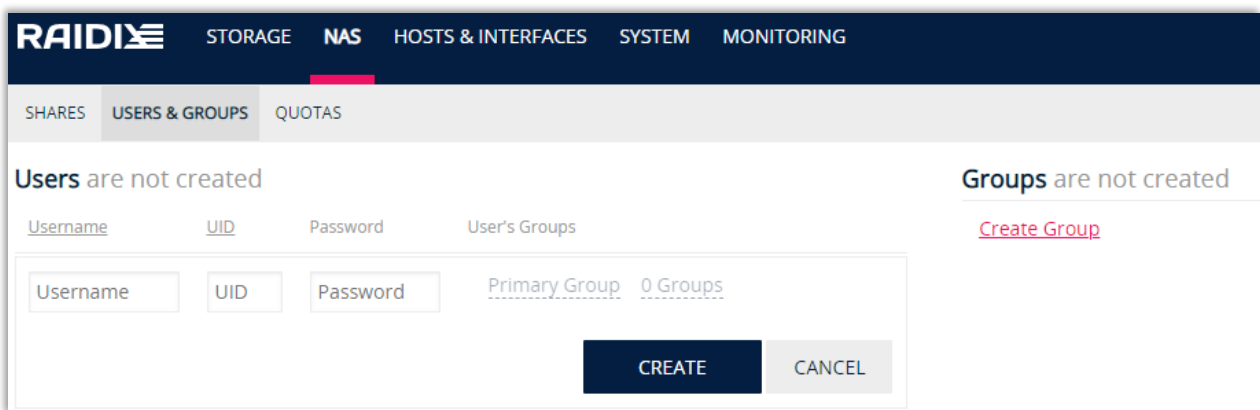


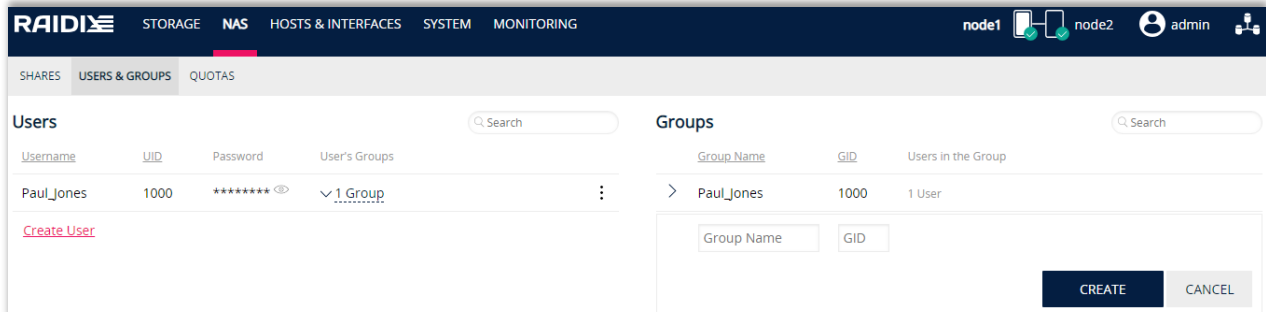
Figure 74. Widget of NAS user creation

3. Specify Username (in Latin letters) and Password.
4. Click **CREATE**.

At the same time, a group with the same name for this user creates. This group will be assigned as a primary for this user by default.

To create a NAS group:

1. Select **NAS > USERS & GROUPS**.
2. Click **Create Group**.
3. Specify a **Group Name** (Figure 75).
4. Click **CREATE**.



Username	UID	Password	User's Groups
Paul_Jones	1000	*****	1 Group

Group Name	GID	Users in the Group
Paul_Jones	1000	1 User

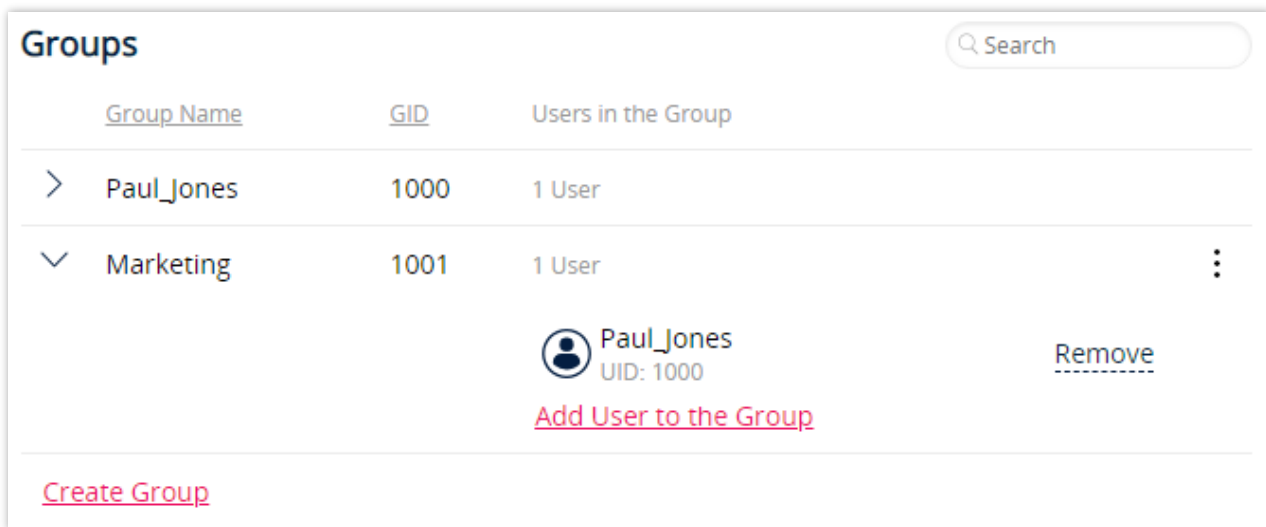
Group Name: GID:

CREATE **CANCEL**

Figure 75. Widget of NAS group creation

To view or add users to a group:

1. Select **NAS > USERS & GROUPS**.
2. Expand the desired group by clicking **>** (Figure 76).
3. Click **Add User to the Group**; the window with available users opens.
4. Choose users from the list and click **SELECT**.



Group Name	GID	Users in the Group
> Paul_Jones	1000	1 User
Marketing	1001	1 User

Paul_Jones
UID: 1000
Add User to the Group

Remove

Create Group

Figure 76. Expanded NAS group


7.2.2 Edit and Delete NAS Users and Groups

To change a user password:

1. Select **NAS > USERS & GROUPS**.
2. Click **:** in the string of the appropriate user, then click **Edit**.

3. Enter a new password and click **APPLY**.

To delete a user:

1. Select **NAS > USERS & GROUPS**.
2. Click  in the string of the appropriate user, and then click **Delete**.
3. In the confirmation window, click **DELETE**.

To change a list of groups for a user:

1. Select **NAS > USERS & GROUPS**.
2. In the column *User's Groups*, click the number of groups, in which the user is included. The list of groups opens (Figure 77).
3. NAS users can be included into several groups simultaneously; each group can be assigned as a primary. Select the group to be primary for the user.
4. To include the user to new groups, click **Add Group for User**. Select the required groups in the opened window and click **SELECT**.
5. To remove the user from a group, next to the group, click **Remove**. Click **REMOVE** in the confirmation window. You can also remove a user from the group in the *Groups* section.

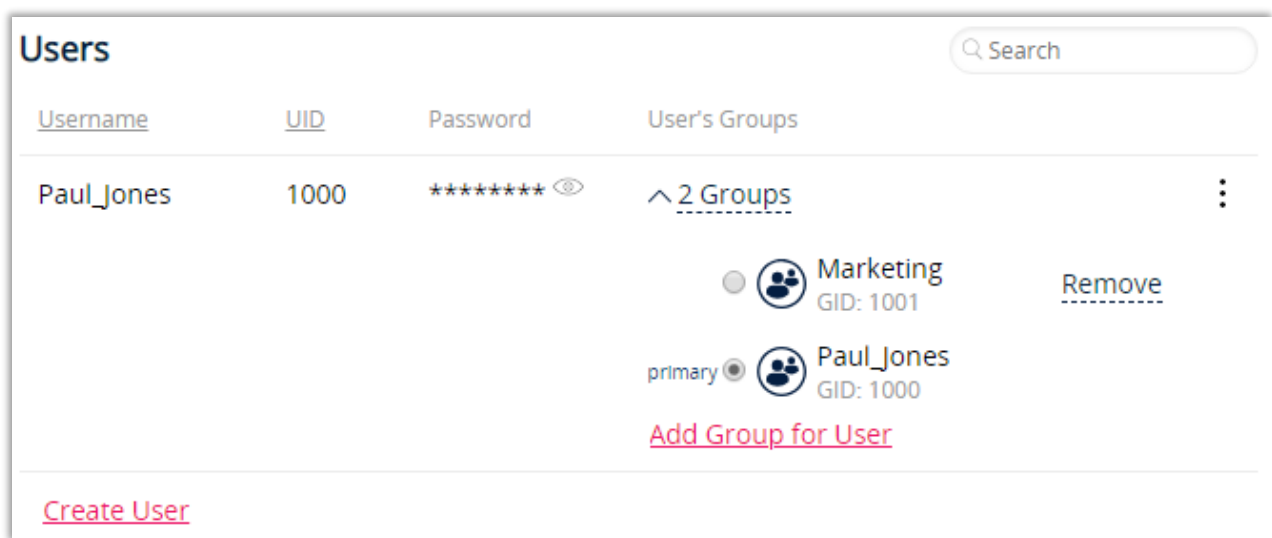



Figure 77. Users section. Setting the priority to a user's NAS group

To remove a group that is not primary for any user:

1. Select **NAS > USERS & GROUPS**.
2. In the *Groups* section, click  and select **Delete**.
3. To confirm, click **DELETE**.



Before removing a primary group, delete all users from it by assigning them to another primary group.

7.2.3 Active Directory Connection

You can change the workgroup name or set up AD at the **NAS | SHARES** page on the Active Directory pane (Figure 78).

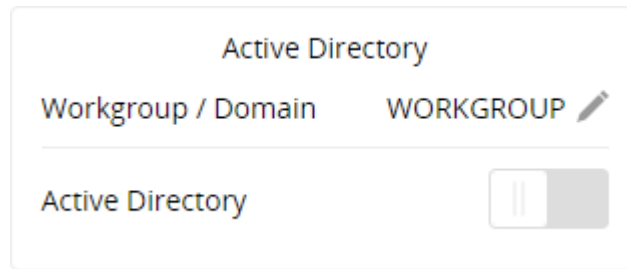


Figure 78. Active Directory pane

7.2.3.1 Setting up Active Directory Connection

We recommend configuring Active Directory (AD) connection *before* creating a share.

Active Directory allows several domain users to work with shares by assigning to a share different user access permission.

Besides that, while working with Active Directory in RAIDIX 5.2, an ability to assign an administrator of root directory of the share and its files is implemented.

To configure Active Directory parameters:

1. Open the **NAS | SHARES** page.
2. Toggle the **Active Directory** switch to the **on** position.
3. On the opened window (Figure 79), manage the connection.

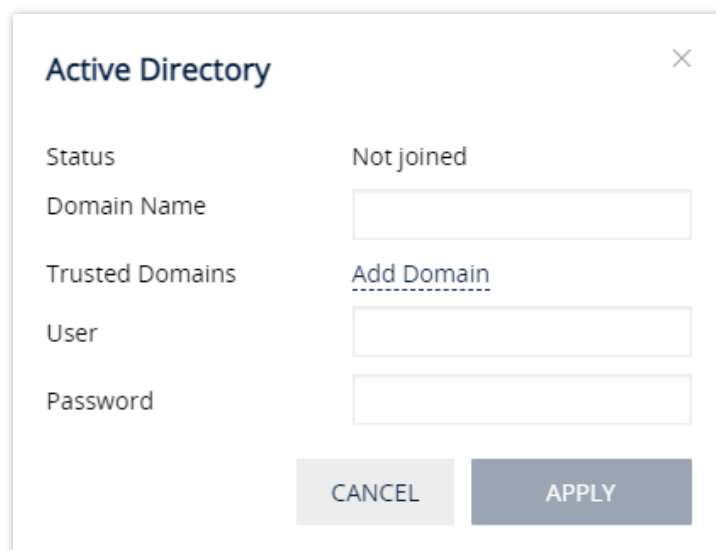


Figure 79. Active Directory settings

Parameter	Description
Active Directory	Enable or disable Active Directory.
Domain name	Active Directory domain <i>name</i> .
Trusted domains	Active Directory trusted domain <i>names</i> .
User	Name of a system administrator that is establishing the connection. The username is required only during initial connection to the domain and not saved in the system.
Password	User password. The password is required only during initial connection to the domain and not saved in the system.

4. Click **APPLY**.



After AD was connected to the system for the first time, you must reboot all nodes in the system. Subsequent connections do not require a reboot.

7.2.3.2 Set up Trusts Between Active Directory Domains

RAIDIX 5.2 allows a user to set up trusts between Active Directory domains. Thus, users in one domain can also be authorized for resources of another domain.

To add trusted domains:

1. Open the **NAS | SHARES** page.
2. Toggle the **Active Directory** switch to the **on** position.

If you already have a connection to AD, you need to reconnect it.

3. Click **Add Domain** in the **Trusted Domains** field. The window *Trusted Domains* (Figure 80) will show.

Figure 80. Adding trusted domains

- Enter the name of a trusted domain in the field **Domain name** and click **Add Domain**. After you added all trusted domains, click **APPLY**.

7.2.4 LDAP Connection

Connecting to LDAP enables LDAP users to work with NFS shares. You can configure different file and directory access rights for users.

LDAP (Lightweight Directory Access Protocol) is an application protocol intended for accessing the directory information services based on X.500. The information model of LDAP is based on entries. Each entry having a distinguished name, (DN) is a collection of attributes. Each DN is a global unique identifier for the catalogue and unambiguously specifies the entry. Each attribute has its own type and one or more values. Types are mnemonic strings, containing an attribute destination – i.e., "cn" stands for common name. Value syntax depends on the attribute type.

LDAP catalogue entries form a hierarchic tree-like structure. The structure reflects the geographic and/or organizational arrangement of stored data. Entries representing countries appear at the top of the tree structure. Below them, there are entries representing country areas and organizations. Below, there are entries representing organization departments, people, printers, documents, etc.

An entry is accessed by its unique name, consisting of the entry name (Relative Distinguished Name, RDN) and its ancestor entries' names.

7.2.4.1 Connecting to LDAP Server

To connect to a LDAP server:

- Open the **NAS | SHARES** page.
- In the pane on the right, toggle the **LDAP** switch to the **on** position (Figure 81).

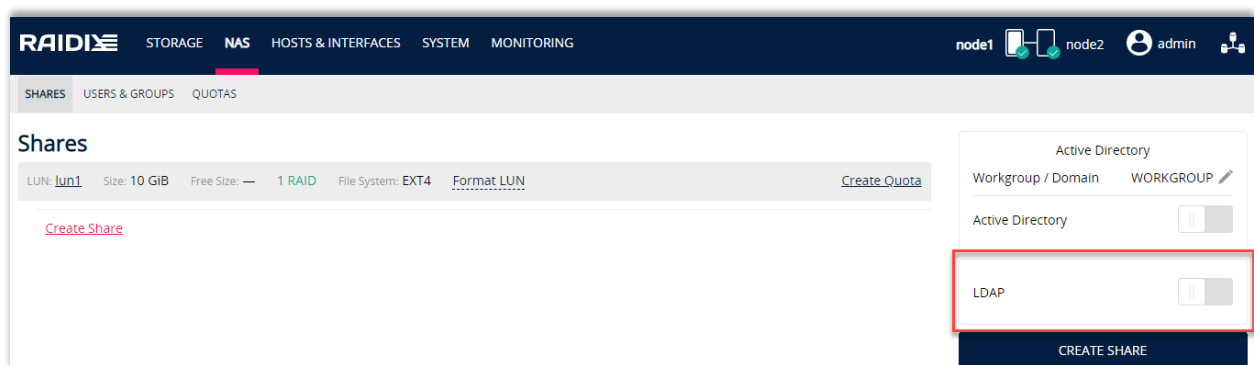


Figure 81. Enabling LDAP

- The *LDAP Settings* window opens (Figure 82).

LDAP Settings

Bind DN

cn=admin,dn=com

?

Password

.....

?

Servers

ldap://testldap1.rd.com

×

?

Server Address

Add Server

Search Base


dc=raidix,dc=com

?

CANCEL

APPLY

Figure 82. LDAP Settings widget

Parameter	Destination
Bind DN	Path to LDAP, where the information for authorization during execution of queries to server is placed.
Password	Password for LDAP authorization, used in combination with Bind DN.
Servers	List of LDAP servers. The servers must be from the same cluster. <div>  Before adding an LDAP server, check its availability. </div>
Search Base	Path to LDAP, from which all requests start.

- Specify the required settings in accordance with the table above and click **APPLY**.
While working in DC mode, repeat the settings on the second node.

Viewing LDAP users

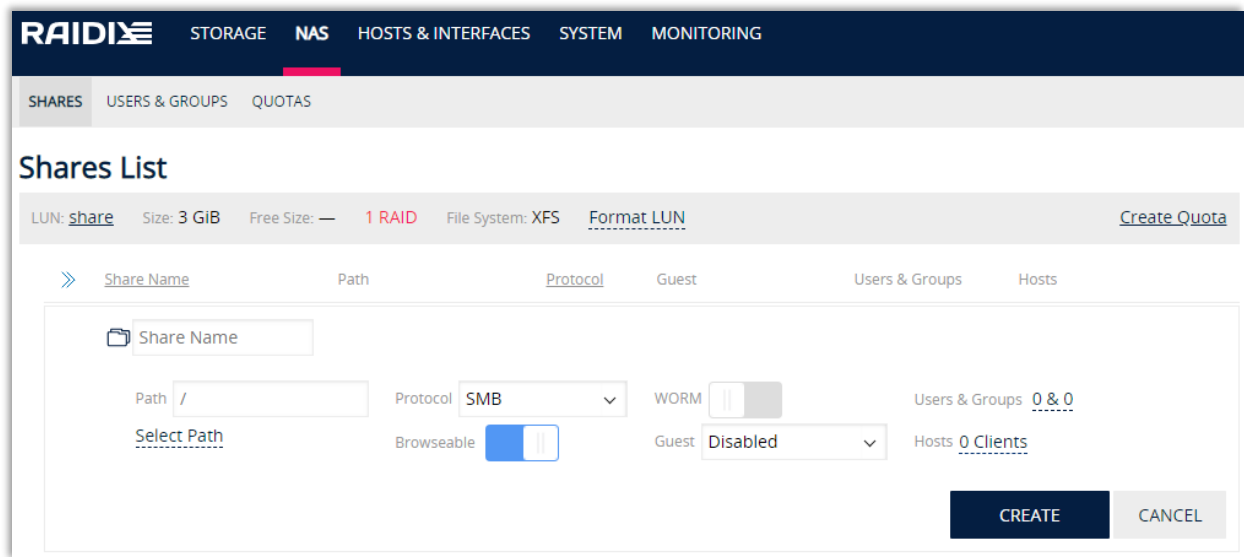
To view LDAP users, in CLI run

```
$ rdcli nas user show -t ldap
```

To learn more, see [Appendix A](#).

7.2.5 Assigning User and Group for a Share

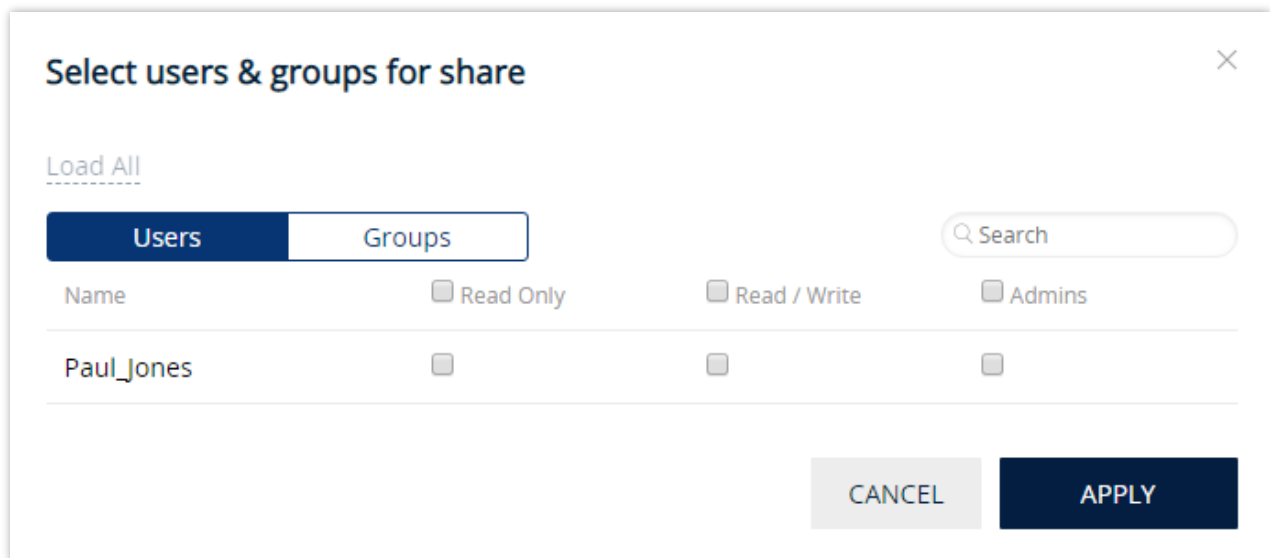
During creation or editing SMB, FTP, or AFP share, click a value of the *Users & Groups* parameter (Figure 83):



The screenshot shows the RAIDIX web interface. The top navigation bar includes 'STORAGE', 'NAS', 'HOSTS & INTERFACES', 'SYSTEM', and 'MONITORING'. Below this, a sub-navigation bar has 'SHARES', 'USERS & GROUPS', and 'QUOTAS'. The main section is titled 'Shares List'. It features a header with 'LUN: share', 'Size: 3 GiB', 'Free Size: —', '1 RAID', 'File System: XFS', 'Format LUN', and a 'Create Quota' link. Below the header is a table with columns: 'Share Name', 'Path', 'Protocol', 'Guest', 'Users & Groups', and 'Hosts'. The 'Users & Groups' column for the 'share' entry is highlighted, showing a dropdown menu with '0 & 0' selected. Other fields visible include 'Path: /', 'Protocol: SMB', 'WORM' (disabled), 'Browseable' (checked), 'Guest: Disabled', and 'Hosts: 0 Clients'. At the bottom right are 'CREATE' and 'CANCEL' buttons.

Figure 83. Widget to create an SMB share

The window to assign users and groups shows (Figure 84).



The screenshot shows a dialog window titled 'Select users & groups for share'. It has a close button (X) in the top right corner. Below the title is a 'Load All' link. There are two tabs: 'Users' (selected) and 'Groups'. To the right of the tabs is a search bar labeled 'Search'. Below the tabs is a table with columns: 'Name', 'Read Only', 'Read / Write', and 'Admins'. The 'Name' column contains the entry 'Paul Jones'. The 'Read Only', 'Read / Write', and 'Admins' columns each have a checkbox. At the bottom right are 'CANCEL' and 'APPLY' buttons.

Figure 84. Adding share users

If [Active Directory](#) connection is enabled, the setting will be applied only for the domain users.

By the list of Active Directory users, you can assign an administrator to a share root directory. The administrator will be able to manage all share files and assign ACL permissions to other users.

In the **Users & Groups** field, assign the appropriate permissions to users and groups by selecting them in the list. Administrator of a share root directory should have the Read/Write permission.

If there is no user in the list, on the **NAS | USERS & GROUPS** page, create a user. To learn more, see [Create Users and Groups](#).

7.3 Managing Quotas

In RAIDIX 5.2, the quota is the limitation of LUN space, available for users of the share, created on this LUN.

i To create quotas for using the LUN space, at least one share should be set up on this LUN.

To create quota:

1. Select **NAS > QUOTAS**.
2. In the section of the appropriate LUN, click **Create Quota**. The pane for adding a quota opens (Figure 85).

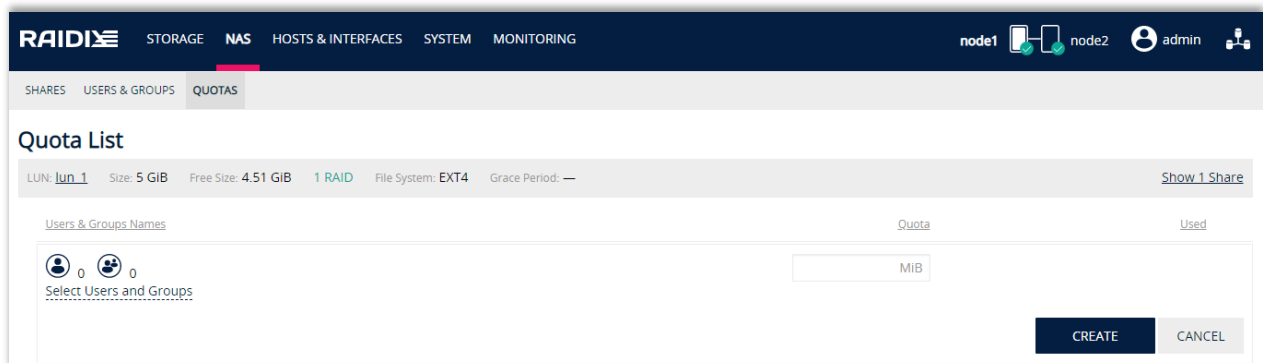


Figure 85. Pane to add a quota

3. Click **Select Users and Groups**. In the opened window, select required users and/or groups and click **SELECT**.
4. Enter a value for quota.
5. Click **CREATE**.

Information about the created quotas for users and groups appears in the table. In the last column, the size of the used space is displayed.

To edit grace period:

i **Except for Windows hosts.**

Quota must be created.

Period (in minutes) that starts after overrunning the quota, during which the user can proceed the writing process on the share. After Grace Period expired, the writing process stops.

The default value is **15** minutes.

1. Select **NAS > QUOTAS**.

- For a chosen LUN, click  next to the *Grace Period* parameter (Figure 86), and enter a new value (in minutes).

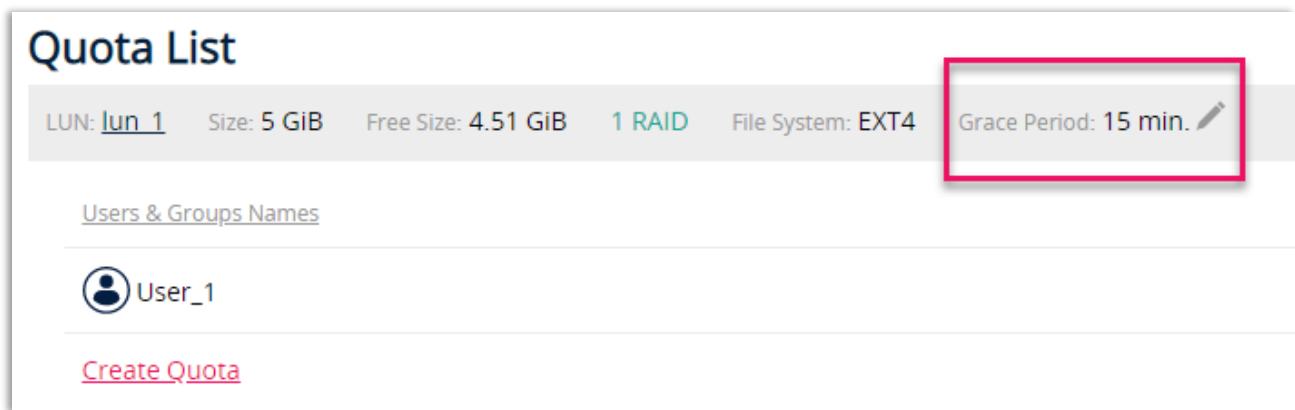




Figure 86. Grace Period parameter


To change quota:

- Select **NAS > QUOTAS**.
- Click  and select **Edit**.
- After the changes, click **APPLY**.

To delete quota:

- Select **NAS > QUOTAS**.
- Click  and select **Delete**.
- To confirm, click **DELETE**.

7.4 Editing a NAS Share

- Select **NAS > SHARES**.
- In the line of the required share, click , then click **Edit**. The pane of share settings opens. Fields *Share name*, *Path*, and *Protocol* are not available for editing.
- Define new share parameters and click **APPLY**.

For the NFS protocol, you can enable RDMA support. Support can only be managed through the CLI (see [rdcli nas nfs modify](#)). The option affects all NFS shares created on the system. Your network adapter must support RDMA.

7.5 Delete a NAS Share



Before you delete a NAS share, make sure clients do not use the share.



While a share is being deleted, access to other shares may be lost for a short time due to the restart of the NAS service.

1. Select **NAS > SHARES**.
2. In the corresponding share line, click and select **Delete**.
3. To confirm, click **DELETE**.

8. WORKING WITH HOSTS AND INTERFACES

This chapter explains how to manage hosts aliases, view established sessions, adapters' info, set up OpenSM service (for InfiniBand), and create iSCSI targets.

8.1 Working with Host in DC

If the system works in DC, set up Multipath on hosts. Setting up Multipath for different hosts described in the *RAIDIX 5.2 Appendixes*.

8.2 Managing Host Aliases

i Only for LUNs of the SCSI type.

Assign easy-to-remember name – host alias name – to your hosts. Combine several hosts into one host alias.

HOST & INTERFACES page elements let you to:

1. [Create and prioritize host aliases.](#)
2. [Delete host aliases.](#)
3. [View information on all the established host sessions.](#)

8.2.1 Creating a Host Alias

1. Select **HOSTS & INTERFACES > HOSTS**.
2. Click **Create Host Alias**. The widget to create the host alias opens (Figure 87).

The screenshot shows the RAIDIX management interface. The top navigation bar includes 'STORAGE', 'NAS', 'HOSTS & INTERFACES' (which is highlighted), 'SYSTEM', and 'MONITORING'. Below this, a sub-navigation bar shows 'HOSTS', 'ADAPTERS', 'ISCSI TARGETS', 'NVME-OF', 'QOSMIC', and 'SAN OPTIMIZER'. The 'HOSTS' tab is selected. The main content area is titled 'Host Aliases'. It contains a form with the following fields and controls:

- Host Alias**: A label for the section.
- Host (WWN/GUID/IQN)**: A text input field with a 'Select' button below it.
- Realtime**: A toggle switch.
- QoSmic**: A toggle switch.
- LUN List**: A section with an 'Add LUN' button and a message 'No LUNs available'.
- Alias Name**: A text input field.
- CREATE** and **CANCEL** buttons at the bottom right.

Figure 87. Hosts Aliases section

3. In the **Alias Name** field, type the host alias name.
4. If the session is set up, click **Select** and choose an appropriate session. Click **Select**.
5. If the session is not set up, type a unique identifier for host:

- for Fibre Channel host – *WWN*;
 - for InfiniBand host – *GUID*;
 - for iSCSI host – *IQN*;
6. Specify the Realtime parameter that sets priority for a host:
- **On** – requests from a particular host are processed before any others regardless of the load from the other hosts.
 - **Off** – requests from a particular host are processed secondarily, with a lower priority.
- i** Host prioritization means giving the highest priority to the requests from a corresponding host. These requests are executed at first, irrespective of the load from the hosts with a normal priority.
- i** At presence of at least one Host Alias with added LUN, it is necessary to add Host Alias with added LUN for every initiator.
7. Turn on QoSMic service if you are planning to use it. To learn more, see [QoSmic set up](#).
8. To set up host access level to LUN (to create a masking rule), click **Add LUN** and select (Figure 88) the corresponding LUNs, access to which you are going to set up. The defined LUN (LUNs) displays in the *LUN List* section.

Select LUNs
×

<input type="checkbox"/>	LUN Name	Access	LUN Size	Block Size	Serial	
<input type="checkbox"/>	I5	Read / Write	1 GiB	512 B	D5C0AF7A-BE7F-4455-A424-...	✓
<input type="checkbox"/>	I4	Read / Write	1 GiB	512 B	E34A7EF9-3F8F-479E-807A-7...	✓
<input type="checkbox"/>	I3	Read / Write	1 GiB	512 B	5B16E7A0-4904-466B-99C8-...	✓
<input type="checkbox"/>	I2	Read / Write	1 GiB	512 B	47FA7CDF-6DFD-4844-A157-...	✓
<input type="checkbox"/>	I1	Read / Write	1 GiB	512 B	8A5024C1-48D9-4CFE-B82F-...	✓

CANCEL
SELECT

Figure 88. Selecting LUNs to add them to the alias

- 8.1. Specify the host access level to LUNs, selecting the appropriate item from the field:
- **Read Only**;
 - **Read / Write**.
9. Click **CREATE**.

! After creation or change of masking rules, it is necessary to rescan drives (on Windows hosts) or SCSI bus (on Linux hosts).

i In RAIDIX 5.2, host alias creation and specifying the host access level can also be done from the LUN page, for more details see section [Host Masking Rules](#).

If a Host Alias and a Masking rule were created on the LUN page, information would also be displayed on page **HOSTS** in section *Hosts Aliases*.

8.2.1.1 Set up Advanced Prioritization Parameters

If you assigned the highest priority for processing requests from a host alias (on the page **HOSTS**, the Realtime parameter value for this host alias was set as **On**), modify the following parameters to optimize the system work:

Parameter	Description
Max data rate of non-real-time requests	Parameter allows you to <i>limit the maximum data rate for requests</i> from hosts without the highest priority, in MB. It is analog of attribute <code>--request_saturation</code> while setting from the command console. A nonnegative integer, the default value – 0 MiB (no limits are set; commands in all sessions have equal priority). When the parameter value is set more than 0 , sessions without the highest priority are guaranteed maximum total bandwidth, equal to the value of this parameter. Sessions with the highest priority are guaranteed total bandwidth, equal to the difference between the total capacity of the system and the value of this parameter.
Max number of non-real-time requests	Parameter allows you to limit <i>the maximum number of simultaneous requests from the hosts</i> , which do not have the highest priority. Analogue of attribute <code>--command_saturation</code> while setting from the command console. The default value is 4 MiB. Setting the parameter <code>--command_saturation</code> as 0 means no limits are set for a total number of simultaneously processing requests in sessions, which do not have the highest priority. Setting the parameter value higher than 0 means that sessions, which do not have the highest priority, are guaranteed simultaneous processing of a total number of requests, equal to the parameter value.

i If sessions with hosts not having the highest priority are not created, commands in all sessions will have equal priority.

Setting the `--request_saturation` parameter as **0** disables the `--command_saturation` parameter.

You can set the parameters on the **HOSTS & INTERFACES | HOSTS** page in the pane *Real-Time Settings* (Figure 89) or using CLI commands. To learn more, see [Appendix A](#).

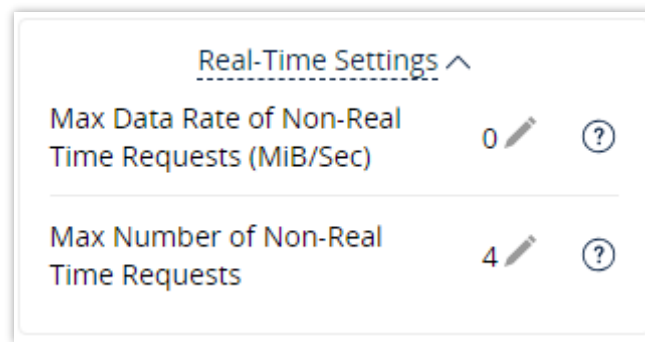


Figure 89. Real-Time Settings pane

8.2.2 Deleting a Host Alias

1. Select **HOSTS & INTERFACES > HOSTS**.
2. Click and click **Delete**.
3. Click **DELETE** in the confirmation window.

8.2.3 Viewing Established Sessions

The session is an established connection between RAIDIX and a client computer. The client system is visible from the storage and is in the list of the established sessions.

You can view the information on all sessions with the system nodes (with hosts, not combined in aliases) on the **HOSTS & INTERFACES | HOSTS** page in the *Other Sessions* section (Figure 90).

Other Sessions Search		
Host (WWN/GUID/IQN)	Target Port	Node
scst_local_host	scst_local_tgt	node1
	scst_local_tgt	node2

Figure 90. Other Sessions pane

Parameter	Description
Host (WWN/GUID/IQN)	A unique host identifier: <ul style="list-style-type: none"> • <i>WWN</i> – a unique Fibre Channel ID; • <i>GUID</i> – a unique InfiniBand ID; • <i>IQN</i> – a unique iSCSI-host ID;
Target Port	A unique target port used for interaction between the host and RAIDIX.
Node	Node name.

8.3 Working with Adapters

To start working with adapters, open the **HOSTS & INTERFACES | ADAPTERS** page.

The **ADAPTERS** page contains detailed information on all adapters and ports, used in the system:

1. [iSCSI ports.](#)
2. [InfiniBand adapters.](#)
3. [Fibre Channel adapters.](#)
4. [SAS adapters.](#)

Besides that, on the **ADAPTERS** page, you can set up targets for synchronization (while configuring dual-controller mode after the system installation) and start OpenSM service (in the case of setting nodes synchronization or connection to the host over InfiniBand).

8.3.1 Information on iSCSI Ports

Information on iSCSI ports is available on the **HOSTS & INTERFACES | ADAPTERS** page in the *iSCSI Ports* section (Figure 91)

Port Name	MAC	IP Address	Netmask	Status	Policy	Link Speed
ens18	aa:99:cd:a6:78:4a	172.16.22.45	255.255.252.0	Up	—	1,000 Mb/s

Figure 91. iSCSI Ports section

Parameter	Description
Port name	Ethernet port name.
MAC	MAC address.
IP Address	IP address which is used to access the particular interface.
Netmask	Netmask.
Status	<p>Connection status of a corresponding Ethernet port. The parameter can have the following values:</p> <ul style="list-style-type: none"> • <i>Up</i> – a connection is established; • <i>Down</i> – there is no connection. Possible reasons: the cable is switched incorrectly, or the driver is missing; • <i>Unknown</i> – the port is not defined.
Policy	<p>Mode of aggregating multiple networks interfaces.</p> <p>Policy defines the mode of load balancing between network interfaces combined into a single logical bonded interface.</p> <p>To learn more, see Set Up Network Interfaces Parameters.</p>
Link speed	Link speed, in MiB/s.

8.3.2 Information on InfiniBand Adapters



While using two-port Mellanox VPI adapters, avoid the configuration with one port in Ethernet mode and another port in InfiniBand mode.

The information on the used InfiniBand adapters is available on the **HOSTS & INTERFACES | ADAPTERS** page in the section *InfiniBand Adapters* (Figure 92).

InfiniBand Adapters					
Device:Port	Port GUID	Node GUID	State	Link Status	Link Speed
mlx4_0:1	f452:1403:0015:6351	f452:1403:0015:6350	Down	Disabled	40 Gb/s
mlx4_0:2	f452:1403:0015:6352	f452:1403:0015:6350	Active	LinkUP	40 Gb/s

Figure 92. ADAPTERS page: InfiniBand Adapters section

Parameter	Description
Device:Port	Port number.
Port GUID (Globally Unique Identifier)	A unique name of the InfiniBand adapter port.
Node GUID	A unique name of the InfiniBand adapter.
State	A physical state of the InfiniBand adapter port: <ul style="list-style-type: none"> • <i>Init</i> – the port is initialized, but not active; • <i>Active</i> – physical connection is established, the port is active; • <i>Down</i> – there is no physical connection.
Link status	Status of connection through the InfiniBand adapter: <ul style="list-style-type: none"> • <i>LinkUp</i> – connection is established. This state appears only if Active connection is established; • <i>Polling</i> – there is no connection. Possible reasons: incorrect cable connection or no driver loaded; • <i>Down</i> – there is no connection with the port.
Link speed	Connection speed (Gb/sec).

8.3.3 Information on Fibre Channel Adapters

Information on Fibre Channel Adapters is available on the **HOSTS & INTERFACES | ADAPTERS** page in the *Fibre Channel Adapters* section (

Figure 93).

Fibre Channel Adapters				
Model	WWNN	WWPN	Link Status	Link Speed
QLE2562	20:00:00:24:ff:5c:f2:e2	21:00:00:24:ff:5c:f2:e2	Down	Unknown
QLE2562	20:00:00:24:ff:5c:f2:e3	21:00:00:24:ff:5c:f2:e3	Down	Unknown
ATTO Celerity FC-82EN	20:00:00:10:86:aa:bb:cc	21:00:00:10:86:03:66:e0	Up	8 Gb/s
ATTO Celerity FC-82EN	20:00:00:10:86:aa:bb:cc	21:00:00:10:86:03:66:e1	Up	8 Gb/s

Figure 93. ADAPTERS page: the Fibre Channel Adapters section

Parameter	Description
Model	Adapter model.
WWNN (World Wide Node Name)	A unique node name of the Fibre Channel device.
WWPN (World Wide Port Name)	A unique name of the adapter port.
Link status	Status of connection through the Fibre Channel adapter: <ul style="list-style-type: none">• <i>Up</i> – the connection is established;• <i>Down</i> – there is no connection. Possible reasons: incorrect cable connection or no driver loaded;• <i>Unknown</i> – connection status is unknown.
Link speed	Connection speed (Gb/s).

8.3.3.1 Modifying the Fibre Channel parameters

In RAIDIX 5.2, a number of Fibre Channel adapter parameters are available for modification (only for ATTO).

To change the parameters, use the following CLI command that lists the adapter parameters and the instructions on how to modify it:

```
$ rdcli fc celerity modify {-h|--help}
```

8.3.4 Information on SAS Adapters



Only cache synchronization and drive enclosure connection are supported.

Information on SAS adapters is available on the **HOSTS & INTERFACES | ADAPTERS** page in the *SAS Adapters* section (Figure 94).



Use an adapter ATTO SAS only to connect the enclosure.

SAS Adapters					
State	Link Speed	PHY Number	SAS Address	Port	Connector Number
Up	12 Gbit	0	50:06:05:B0:09:0B:75:B0	0	0
Up	12 Gbit	1	50:06:05:B0:09:0B:75:B1	0	0
Up	12 Gbit	2	50:06:05:B0:09:0B:75:B2	0	0
Up	12 Gbit	3	50:06:05:B0:09:0B:75:B3	0	0
Up	12 Gbit	4	50:06:05:B0:09:0B:75:B4	1	1
Up	12 Gbit	5	50:06:05:B0:09:0B:75:B5	1	1
Up	12 Gbit	6	50:06:05:B0:09:0B:75:B6	1	1
Up	12 Gbit	7	50:06:05:B0:09:0B:75:B7	1	1

Figure 94. SAS Adapters section

Parameter	Description
State	State of connection through the SAS adapter: <ul style="list-style-type: none"> <i>Up</i> – connection is established; <i>Down</i> – there is no connection with the port. Possible reasons: incorrect cable connection or no driver loaded. <i>Unknown</i> – the state is unknown.
Link speed	Connection speed (in Gb/s).
Phy number	SAS physical links.
SAS address	A unique SAS worldwide name.
Port	Logical port's number.
Connector number	Physical port's number.

8.4 Setting up iSCSI

8.4.1 Enabling and Disabling iSCSI Support

To enable or disable iSCSI support, use the **Enable iSCSI** switch (Figure 95).

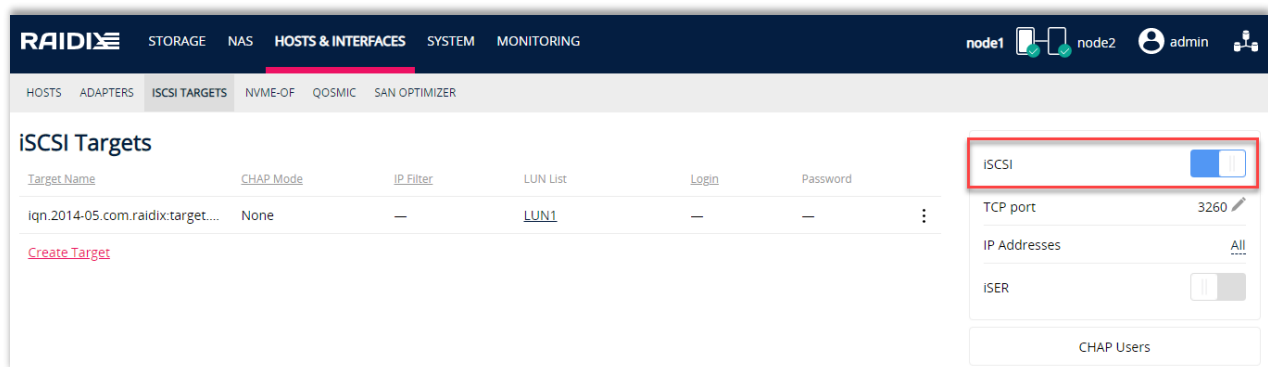


Figure 95. Enable iSCSI switch in the on position

8.4.2 Disconnecting an iSCSI Initiator on Windows Server

Due to some specifics in the process of Windows Server work, to disconnect an iSCSI target properly, do the following:

For Windows Server 2012 R2:

1. Open the **Start** menu and go to **Disk Management**.
2. Right click on **Disk X** and select **Offline**.
Disk X, in this case, is RAIDIX LUN.
3. Repeat step 2 for each LUN.
4. Open the **iSCSI Initiator Properties**, select the target, and click **Disconnect**.

If the LUN is under the load, then after you click **Disconnect**, the LUN will not disconnect immediately, but as soon as writing from Windows cache to LUN will be over.

8.4.3 Creating an iSCSI Target

To create an iSCSI target:

1. Select **HOSTS & INTERFACES > iSCSI TARGETS**.
2. Click **Create Target**. The pane to set up the iSCSI target appears (Figure 96).

The screenshot shows the RAIDIX web interface with the 'HOSTS & INTERFACES' tab selected. Under 'iSCSI TARGETS', there is a form titled 'iSCSI Targets not created'. The form has six fields: 'Target Name' (containing 'iqn.2014-05.com.raidix:target:'), 'CHAP Mode' (a dropdown menu set to 'None'), 'IP Filter' (containing 'Add IP'), 'LUN List' (containing '1 LUN'), 'Login', and 'Password'. At the bottom right of the form are two buttons: 'CREATE' and 'CANCEL'.

Figure 96. iSCSI target creation

i While creating a new iSCSI-target, make sure that iSCSI-target IQN is unique within your network.

3. Fill in the following fields and click **CREATE**:

Parameter	Description
Target name	The full name of iSCSI target.
CHAP mode	<p>The CHAP authorization mode:</p> <ul style="list-style-type: none"> • None – no authorization; • Unidirectional mode (One-way authentication) – only the host introduces itself to the target; • Bidirectional mode – mutual target-hosts authentication. Host sends login and password to the target. After verification target sends login and password to the host. <div> <p>i For Bidirectional and Unidirectional modes of CHAP authorization, it is necessary to create CHAP users.</p> </div>
IP Filter	<p>The IP addressing rule defines the access to target from a particular IP address or from all available IP addresses.</p> <p>The possible parameter's values are:</p> <ul style="list-style-type: none"> • IP address – the rule defines the access to target from a particular IP address; • All – the rule defines the access to target from all available IP addresses. <div> <p>i By default, all iSCSI targets are not available for all IP addresses.</p> </div>

Parameter	Description
LUN List	Option allows assigning a LUN (LUNs) to an iSCSI target. <div> <i>i</i> While working with Windows initiator, in case of necessity of adding a LUN, changing its parameters or deleting, break the session with the initiator, perform required actions and then reestablish the connection with the initiator. </div>
Login	Option is available only in Bidirectional CHAP mode (Duplex authorization).
Password	Option is available only in Bidirectional CHAP mode (Duplex authorization).

8.4.3.1 Creating CHAP users

In case of selection Bidirectional and Unidirectional modes of CHAP authorization, create CHAP users

To create a CHAP user:

1. Select **HOSTS & INTERFACES > ISCSI TARGETS**.
2. Click **CHAP Users**. A window with logins and passwords of users for CHAP authorization appears (Figure 97).

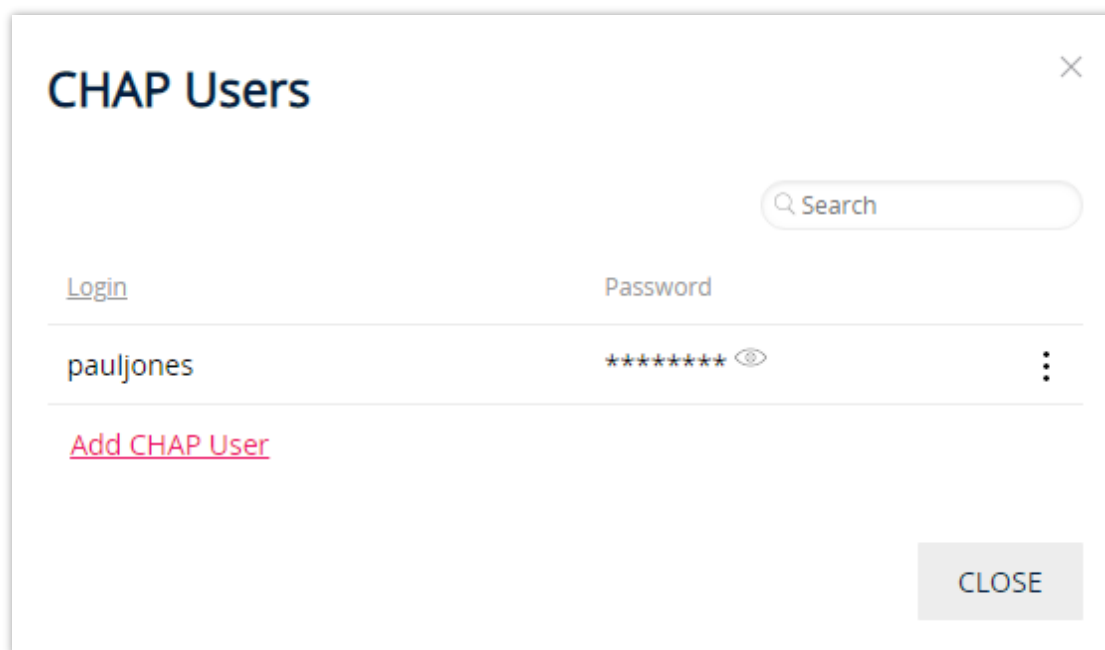




Figure 97. CHAP Users pane

3. Verify that user is in the list. To add a new user, click **Add CHAP User** and type its login and password in the appropriate fields. Click **CREATE**.
4. To view the password, click  in the appropriate user line.
5. To edit the password, click  in the required CHAP user line, and then click **Edit**. Change the password and click **APPLY**.

8.4.4 Setting up the Advanced iSCSI Target Parameters

1. Select **HOSTS & INTERFACES > ISCSI TARGETS**.
2. In the row of the required iSCSI target, click **:** and select **Edit**.
3. In the opened widget, click **Parameters**. The *Target Parameters* window opens (Figure 98):

Target "iqn.2014-05.com.raidix:target.73431" parameters

Header Digest	<input type="checkbox"/>	?
Data Digest	<input type="checkbox"/>	?
Initial R2T	<input type="checkbox"/>	?
Immediate Data	<input checked="" type="checkbox"/>	?
Max Recv Data Segment Length	1048576	?
Max Xmit Data Segment Length	1048576	?
Max Burst Length	1048576	?
First Burst Length	1048576	?
Max Outstanding R2T	32	?

CANCEL **APPLY**


Figure 98. Target Parameters window

Parameter	Description
Header Digest	Option provides the header integrity during the interaction with the specified iSCSI-target: <ul style="list-style-type: none"> • CRC32C – option is enabled; • Off – option is disabled (set by default).
Data Digest	Option provides the data integrity during the interaction with the specified iSCSI-target: <ul style="list-style-type: none"> • CRC32C – option is enabled; • Off – option is disabled (set by default).

Parameter	Description
Initial R2T	Option is used to turn off the default use of R2T for unidirectional and the output part of bidirectional commands. The possible parameter values are: <ul style="list-style-type: none"> On – option is enabled; Off – option is disabled (set by default).
Immediate Data	Option allows including PDU as a part of initiator commands. The possible parameter values are: <ul style="list-style-type: none"> On – option is enabled (If Immediate Data is set to Yes and InitialR2T is set to Yes (default), then only immediate data is accepted in the first burst); Off – option is disabled. The default: On .
Max Recv Data Segment Length	Option is used to specify the maximum data segment length in bytes that can be received in one iSCSI PDU. Minimum value is 512 . Maximum value is 1048576 (set by default).
Max Xmit Data Segment Length	Option is used to specify the maximum data segment length in bytes that could be sent in one iSCSI PDU. Minimum value is 512 . Maximum value is 1048576 (set by default).
Max Burst Length	Option is used to specify the maximum SCSI data payload in bytes in a Data-In or a solicited Data-Out iSCSI sequence. Minimum value is 512 . Maximum value is 1048576 (set by default).
First Burst Length	Option is used to specify the maximum amount in bytes of unsolicited data an iSCSI host might send to the target during the execution of a single SCSI command. Minimum value is 512 . Maximum value is 1048576 (set by default).
Max Outstanding R2T	Option is used to specify the maximum number of outstanding R2Ts per task, excluding any implied initial R2T that might be part of that task. Minimum value is 1 ; maximum value is 65535 . The default: 32 .

- After you select the corresponding options, click **APPLY**.

8.4.5 Deleting an iSCSI Target

- Select **HOSTS & INTERFACES > ISCSI TARGETS**.
- In the line of an iSCSI target, click  and select **Delete**.
- To confirm, click **DELETE**.

8.5 Setting up QoS mic



Only for LUNs of the SCSI type.

QoSmic service, implemented in RAIDIX 5.2, allows assigning different priority for hosts in the system (Realtime parameter) without Administrator intervention. QoSmic automatically prioritizes applications requests according to the type of application (critical, non-critical) running on the host. This service responds to the storage system load changes and maintains required level of performance for critical and important applications.

QoSmic work includes two stages:

- Learning – launch the application on the host and switch on learning. While learning QoSmic reads signatures (specific metrics of the application), this process can take up to 8 hours. After learning is finished, assign a priority to the application and QoSmic saves all data (signatures and type of applications) to the QoSmic database.
- Recognizing – QoSmic recognizes applications by the signatures database and distributes the load on the system, adjusting the Realtime option. Thus, when the critical or important application (Favored) is run on a host, QoSmic service will automatically prioritize this host by processing its requests in the first place (Realtime – On). All requests from hosts with less important applications (Unwanted) will be handled with a small delay.
- Stages of QoSmic work always run separately. It is impossible to run Learning and Recognizing for applications simultaneously.

8.5.1 QoSmic Learning

The main purpose of QoSmic service is to turn on and off the Realtime parameter for the hosts recognized as hosts with important or critical applications. To recognize important/ critical applications and turn on the Realtime parameter, QoSmic needs to collect signatures of these applications first.

QoSmic collects signatures of the application through learning or you can upload signatures previously collected by QoSmic.

i When you run learning on a single controller, QoSmic will learn only one application at a time. In the dual controller mode, you can run learning of different applications on every controller simultaneously, if RAIDs are located on different nodes.

! In the DC mode, if failover happened during learning, the number of signatures will stop to grow, thus learning occurs only on the active node.

8.5.2 Starting QoSmic Learning

You can start learning from the web interface of the service or from CLI.

To start QoSmic learning:

1. Select **HOSTS & INTERFACES > HOSTS**.
2. Create a host alias included the **Qosmic** switch in the **On** position and added LUN (Figure 99).

i You can change the list of hosts in the process of learning.

! If there were no Host or Host Alias created, the learning will not start.

3. Click **CREATE**.

Figure 99. Creating a host alias: parameter panel

4. Open the **HOSTS & INTERFACES | QOSMIC** page (Figure 100). The **QOSMIC** page displays the list of hosts.
5. In the right part of the page, expand the *Teacher* pane and click **LEARN NEW APPLICATION**. The learning starts.

Figure 100. QOSMIC page

To collect the application signatures successfully, the host with the application should work under the high load in the process of learning. Run the application on the selected host to start learning.

i The full duration of one learning cycle is **8** hours or **40 000** signatures.

i The information on the learning process that list of applications shows may be different from the information that displays the Learning Progress section. In the case of any discrepancy between the information on training, believe the data under the Learning Progress section. The actual information about learning will be displayed on the list of applications after the learning will be stopped/completed.


! Learning is impossible to launch in the following cases:

- Learning is currently in progress for another application.
- Recognition is currently in progress.
- There are no hosts for learning (in this case, the button **Start Learning** becomes inactive).
- Learning has been already performed for the selected application.

8.5.2.1 Stopping QoSmic Learning

You can stop the process of learning by clicking **Stop Learning**. If the process of learning was stopped, the service will add the application to the *Applications list* on the **QOSMIC** page.

i The application will be saved only after the system will collect at least one signature and will run at least 1% of the learning time.

You can proceed the learning. To proceed the learning, click  next to the appropriate application, and then **Continue Learning**.


8.5.3 Application Types

After learning is finished, the following warnings can be displayed:

- *low_intensity* – low intensity on more than 50% of intervals, the signature might be invalid.
- *same_size* – less than 90% of signatures have the same size, the signature might be invalid.
- *short_learning_time* – not enough time was spent on learning, signatures may be invalid.

After learning is completed, set up the application type.

To setup the application type, do the following:

1. On the **QOSMIC** page, in the Applications list select the application the type of which you want to setup and click . The parameters panel opens.
2. In the parameters panel choose the type of the application from the field:
 - **Favored** – the selected application has high priority.
 - **Unwanted** – the selected application will be ignored.
 - **Disabled** – the application does not provide signatures for QoSmic.

8.5.4 Import of Signatures

You can download collected signatures for the further use. For example, you can download the application (signatures) and upload them to another server.

To download signatures, on the **QOSMIC** page, in the *Applications* section, select the application whose signatures you want to download and click **Download** (Figure 101).

To upload signatures, do the following:

1. Select **HOSTS & INTERFACES > QOSMIC**.
2. Click **Upload New Application** under the list of applications.
3. Select the application you want to upload.
4. Click **Open**.

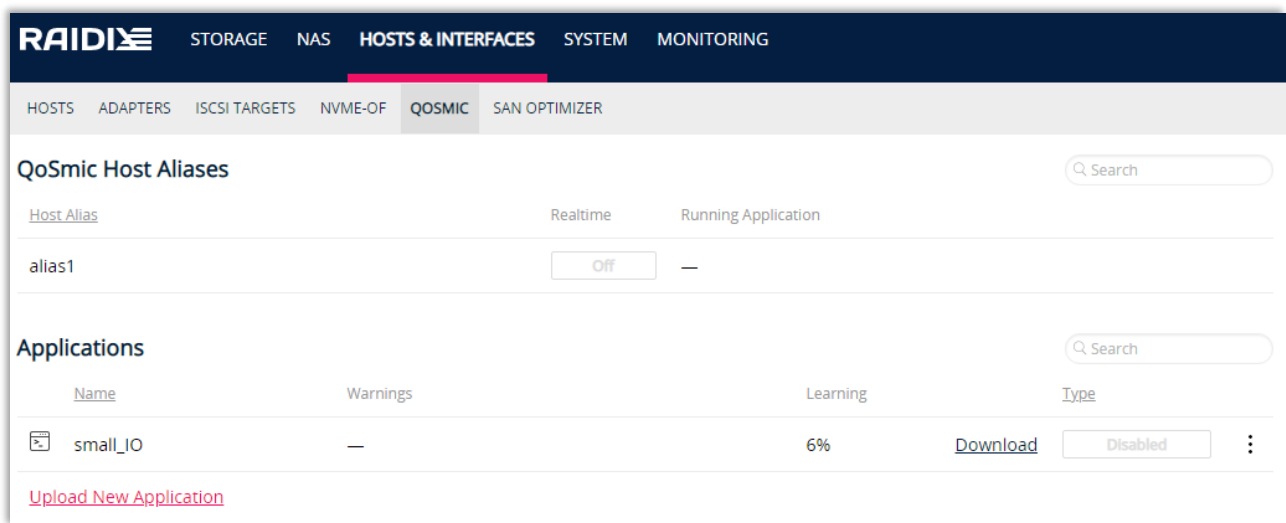


Figure 101. New application added



If the file for import has the same name as the already uploaded file, the system will add (1) to the uploaded file.

8.5.5 QoSmic Recognizing


After you perform learning of the application and set up the type, you can start applications recognition with QoSmic. QoSmic will recognize applications with *Favored* and *Unwanted* types only. You can start recognizing from the web interface of the service or CLI.


To start recognizing:


1. Click **Start recognizer** on the left side of the **QOSMIC** page
2. Recognizing follows the rules:
 - On the **HOSTS** page, when Realtime and QoSmic are **Off**, any type of application will not affect Realtime (it means that Realtime will be off).

- On the **HOSTS** page, when Realtime is **On** and QoSmic is **Off**, any type of application will not affect Realtime.
- On the **HOSTS** page, when Realtime is **Off** and QoSmic is **On**, then Realtime will be On for *Favored* applications and Off for any other type of applications.
- On the **HOSTS** page, when both Realtime and QoSmic are **On**, then Realtime will be Off for *Unwanted* applications and On for any other type of applications.

8.6 NVMe-oF

 Only for ERA RAID.

 At least NVMe-oF version 1.4 is required to work with the DC system.

 Use the ElRepo OS kernels for proper operation of CentOS hosts with NVMe-oF.

NVMe-oF is a specification-defined extension to NVMe that enables NVMe-based communication over interconnects other than PCIe. This interface makes it possible to connect “external” storage enclosures to a server, either directly or through a switch, while still using NVMe as the fundamental communication mechanism.

As an interconnection for NVMe-oF, you can use

- over RDMA: InfiniBand HBAs in any mode (ib or eth);
- over TCP: Ethernet HBAs and InfiniBand HBAs in any mode (ib or eth).

Definitions

Namespace is defined as a quantity of nonvolatile memory that can be formatted into logical blocks. This definition is virtually identical to the SCSI concept of a logical unit.

NVMe Qualified Name (NQN) is an identifier for a remote NVMe target. Similar with iSCSI Qualified Name (IQN).

Subsystem is a logical entity includes masking rules for namespaces and a list of ports to connect with hosts. Similar with

Referral – list of NVMe-oF ports to search for namespaces.

ANA

Asymmetric Namespace Access (ANA) is an NVMe standard that was implemented as a way for the target to inform an initiator of the most optimal way to access a given namespace.

8.6.1 Configuring Connection

Specifics:

- To work with the NVMe-oF protocol, create LUNs of the NVMe-oF type.
- The host masking rules are configured via subsystems.
- The QoSmic feature is unavailable.
- Configuring the level of namespace access (read only or read and write) for initiators is not supported.
- Configuring load prioritization for subsystems is not supported.



We recommend you to allocate no more than 20 LUNs (namespaces) per initiator. See the [recommendations](#) for details.

To set up NVMe-oF:

1. Configure a subsystem:
 - 1.1. Open **HOSTS AND INTERFACES | NVME-OF**.
 - 1.2. Click **Create Subsystem**.
 - 1.3. In the opened widget, set the parameters of the subsystem (Figure 102).

Parameters in the widget of the subsystem creation

Parameter	Description
Name	Name of the subsystem.
NQN	NQN identifier of the subsystem. If you leave the field empty, the system will automatically insert a value.
NVMe-oF version	Version of the NVMe-oF protocol. The default: 1.4 .
Hosts	The list of NQNs of hosts, for which access to namespaces will be allowed.
Namespaces (LUNs)	List of namespaces to which access will be allowed from the selected hosts.

Subsystems empty

Name	NQN	NVMe-oF version	Hosts	Namespaces (LUNs)
<input type="text" value="rdx.nvmeof.sub0"/>	<input type="text" value="nqn.2021-10.com.raidix"/>	<input type="text" value="1.4"/>	Add host	Add LUN

Figure 102. NVMe-oF subsystem creation widget

- 1.4. Click **CREATE**.
2. Configure ports:
 - 2.1. Click **Create Port**.
 - 2.2. In the opened widget, set the parameters of a port (Figure 103).

Parameters in the widget of the NVMe-oF port creation

Parameter	Description
Name	Name of the NVMe-oF port.
Type	Type of the port. Possible values: TCP , RDMA . The default: TCP .
Family	The parameter is not changed. Address Family of the port – ipv4 .
Address	IP address of the port.
Port	Number of the port. The default: 4420 .
Port ID	–
Security	Security type of the network connection. Possible values: <ul style="list-style-type: none"> • Not specified. • Required. • Not required. The default: Not required .
Subsystems	The list of subsystems that will be visible on this port.
Inline data size	Inline data size (in bytes) for the NVMe-oF port. For TCP: <ul style="list-style-type: none"> • Possible values: integers from 2^{10} to $2^{31}-1$ and -1. • The default: -1. For RDMA: <ul style="list-style-type: none"> • Possible values: integers from -1 to 2^{14}. • The default: 0.

Ports empty

Name	Type	Family	Address	Port	Port ID	Security	Subsystems	Inline data size
tcp:172.16.22	TCP ▼	ipv4	172.16.22.1 ▼	4420	—	Not specil ▼	Add Port Subsystems	-1

CREATE **CANCEL**

Figure 103. NVMe-oF port creation widget

2.3. Click **CREATE**.

3. **For DC system:** repeat steps 1-2 on the second node and configure ANA.

Subsystems created on one node should have an identical analogue (name, namespaces, and masking rules) on the second node.

8.6.2 ANA Configuration



Synchronization between nodes must be configured for the proper operation of ANA.

1. On each node, create an additional NVMe-oF port and add subsystems to them.
2. On each node, on the **HOSTS & INTERFACES | NVME-OF** page, click **Port Synchronization** and select the port for the synchronization.

8.6.3 Deleting Ports and Subsystems

You can delete ports and subsystems at the **HOSTS & INTERFACES | NVME-OF** page by clicking in the line of an appropriate object.

Before deleting a port, make sure that it has no subsystems.

To remove subsystems from a port:

1. Open the **HOSTS & INTERFACES | NVME-OF** page.
2. In a port string, click and select **Edit**.
3. In the *Subsystem* column, remove all subsystems by clicking .
4. Click **SAVE**.

After that, you can delete the port as usual.

9. WORKING WITH SAN OPTIMIZER

SAN Optimizer is a technology of virtualization, which enables to organize several independent storages into a joint virtual storage. SAN optimizer allows to connect targets on external storages to RAIDIX and to work with them in the same way as with local drives; for instance, to create RAIDs and LUNs. SAN Optimizer enables virtualization (forwarding) of a drive as a block device without being modified.

SAN Optimizer has the following advantages:

- High performance due to caching in RAIDIX;
- Enhanced reliability due to creation of RAIDs;
- Budget solution of storage with no additional purchases of expensive equipment.

Before starting to work with SAN Optimizer, apply a valid license with support of this function.

9.1 SAN Connection

You can connect external storages to RAIDIX via iSCSI or FC interface.

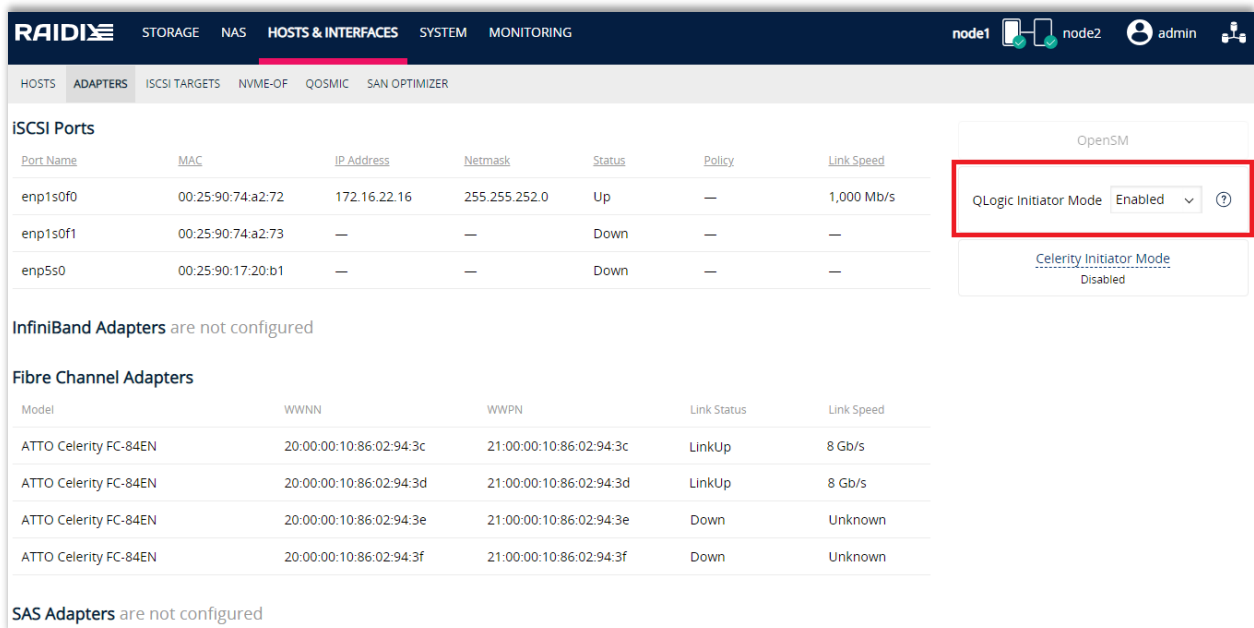


We recommend configuring host aliases before using adapters.

9.1.1 Configuring FC Connection

To configure the SAN connection via FC:

1. Select **HOSTS & INTERFACES > ADAPTERS**.
2. Turn on **QLogic Initiator Mode** (recommended value is **Enabled**) (Figure 104):
 - **Disabled** – initiator mode will never be enabled.
 - **Exclusive** – initiator mode will be enabled on load, disabled on enabling target mode, and then on disabling target mode enabled back.
 - **Dual** – Initiator Modes will be enabled. Target Mode can be activated when ready.
 - **Enabled** – initiator mode will always stay enabled.



The screenshot shows the RAIDIX web interface with the 'ADAPTERS' tab selected. The 'ISCSI Ports' table lists three ports: enp1s0f0 (Up, 1,000 Mb/s), enp1s0f1 (Down), and enp5s0 (Down). The 'OpenSM' settings panel on the right shows 'QLogic Initiator Mode' set to 'Enabled' (highlighted with a red box) and 'Celerity Initiator Mode' set to 'Disabled'.

Port Name	MAC	IP Address	Netmask	Status	Policy	Link Speed
enp1s0f0	00:25:90:74:a2:72	172.16.22.16	255.255.252.0	Up	—	1,000 Mb/s
enp1s0f1	00:25:90:74:a2:73	—	—	Down	—	—
enp5s0	00:25:90:17:20:b1	—	—	Down	—	—

InfiniBand Adapters are not configured

Fibre Channel Adapters

Model	WWNN	WWPN	Link Status	Link Speed
ATTO Celerity FC-84EN	20:00:00:10:86:02:94:3c	21:00:00:10:86:02:94:3c	LinkUp	8 Gb/s
ATTO Celerity FC-84EN	20:00:00:10:86:02:94:3d	21:00:00:10:86:02:94:3d	LinkUp	8 Gb/s
ATTO Celerity FC-84EN	20:00:00:10:86:02:94:3e	21:00:00:10:86:02:94:3e	Down	Unknown
ATTO Celerity FC-84EN	20:00:00:10:86:02:94:3f	21:00:00:10:86:02:94:3f	Down	Unknown

SAS Adapters are not configured

Figure 104. Enabling QLogic Initiator Mode

i Adapters FC ATTO Celerity, provided that initiator mode is enabled, can be used both as a target and as an initiator. The required mode can be set in CLI.

i Adapters FC Qlogic cannot be used as a target if the initiator mode is enabled.

- Click **Celerity Initiator Mode**. The window *Celerity Initiator Mode Settings* opens.
- Toggle the switch **Celerity Initiator Mode** to the **On**. The list of available ATTO Celerity adapters opens (Figure 105).

Celerity Initiator Mode Settings



Celerity Initiator Mode
☒
?

Enable initiator mode on all adapters
☐

<input type="checkbox"/>	Model	WWNN	WWPN	Link Status
<input type="checkbox"/>	ATTO Celerity FC-84EN	20:00:00:10:86:02:94:3c	21:00:00:10:86:02:94:3c	Up
<input type="checkbox"/>	ATTO Celerity FC-84EN	20:00:00:10:86:02:94:3d	21:00:00:10:86:02:94:3d	Up
<input type="checkbox"/>	ATTO Celerity FC-84EN	20:00:00:10:86:02:94:3e	21:00:00:10:86:02:94:3e	Down
<input type="checkbox"/>	ATTO Celerity FC-84EN	20:00:00:10:86:02:94:3f	21:00:00:10:86:02:94:3f	Down

CANCEL
APPLY

Figure 105. Enabling Celerity Initiator Mode

5. Tick the adapters for which initiator mode to be enabled. To enable initiator mode for all adapters, use the switch **Enable initiator mode on all adapters**. Click **APPLY**.
6. Select **HOSTS & INTERFACES > SAN OPTIMIZER**.
7. Select the FC connection point in the section *SAN Connection Points*. Click , then click  **Select Targets** (Figure 106).

RAIDIX

STORAGE

NAS

HOSTS & INTERFACES

SYSTEM

MONITORING

node1

node2

admin

HOSTS

ADAPTERS

ISCSI TARGETS

NVME-OF

QOSMIC


SAN OPTIMIZER

SAN Connection Points

Type	Address	Added Targets	Added Drives	RAIDs	Proxied LUNs
FC	21:00:00:10:86:02:8f:0b	2 targets	2 drives	1 RAID	—
ISCSI	255.255.255.255:5678 IQN: iqn.1994-05.com.redhat:c838604f75a5	1 target	2 drives	1 RAID	1 LUN

Select Targets

Figure 106. SAN Connection Points

8. In the opened window, select required targets and click **SELECT**. In the confirmation window, click **CHANGE**. The number of added targets will appear in the appropriate cell of the section *SAN Connection Points*.
9. To add drives, at the end of the line of appropriate target, click  and select **Select Drives**.
10. In the opened window, select the required drives and click **SELECT**.
11. The added drives will appear in the section *SAN Drives*, on the page **STORAGE | DASHBOARD** in the section *Drives*, and on the **DRIVES** page in the corresponding section.

9.1.2 Configure iSCSI Connection

! iSCSI target on which CHAP authentication enabled cannot be used as SAN target.

To configure the SAN connection via iSCSI:

1. Select **HOSTS & INTERFACES > SAN OPTIMIZER**.
2. Below the section *SAN Connection Points*, click **Create SAN Connection Point (iSCSI Only)** (Figure 107).

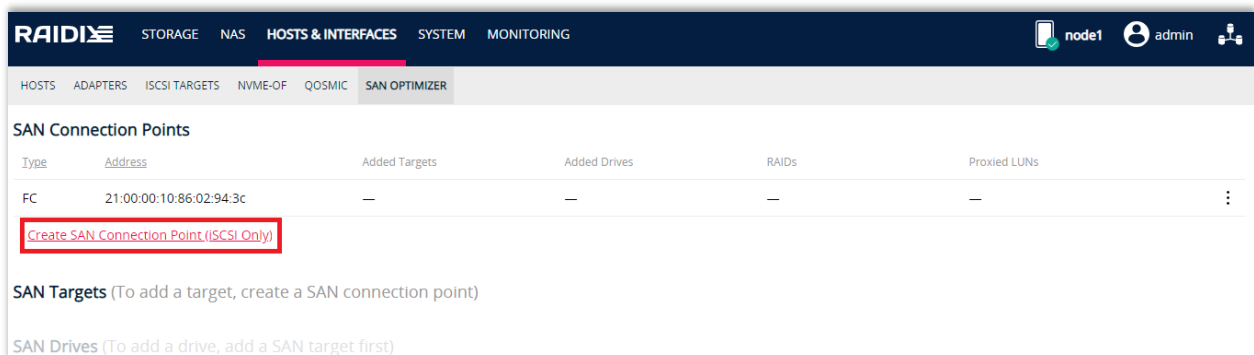


Figure 107. Create SAN Connection Point over iSCSI

3. In the opened pane, specify IP address of the iSCSI target and the port for connection and click **CREATE**. The connection point displays in the section *SAN Connection Points*.

i The default port of iSCSI connection is **3260**.

4. Click **:** and select **Select Targets**.
5. Select the required target and click **SELECT**. Click **CHANGE** in the confirmation window. The number of added targets will appear in the appropriate cell of the section *SAN Connection Points*.
6. To add drives, click **:** at the end of the string of appropriate target and click **Select Drives**.
7. The added drives will appear in the section *SAN Drives*, on the page **STORAGE | DASHBOARD** in the section *Drives*, and on the **DRIVES** page in the corresponding section.

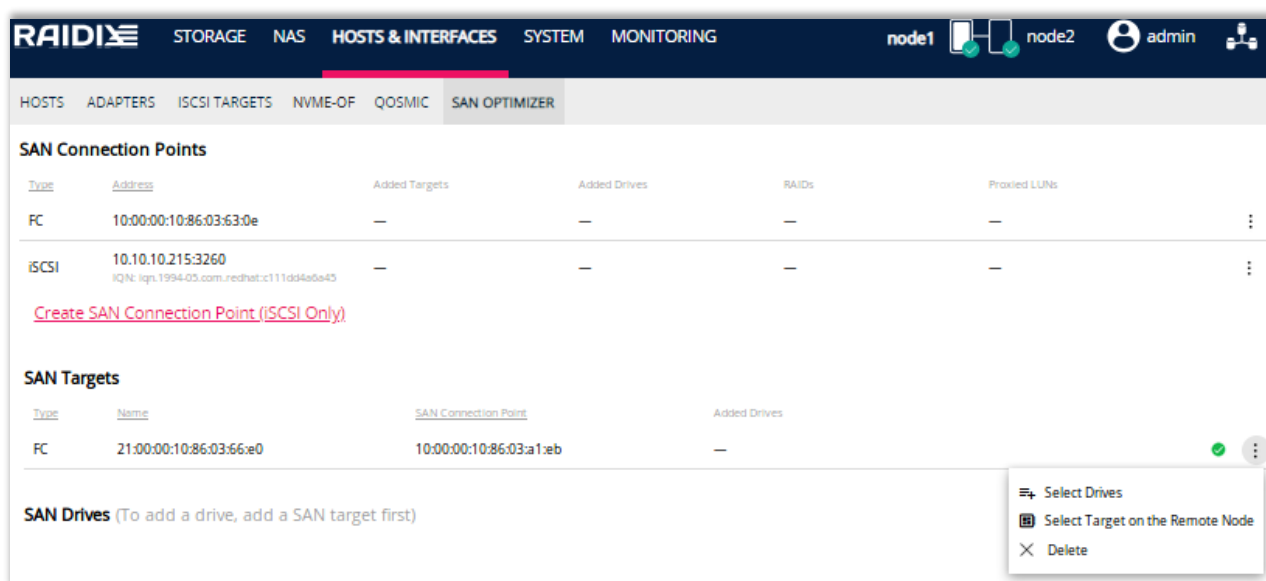
9.1.3 Bonding of SAN Drives in DC

While working in DC, it is possible to add SAN drives only on one node. Use bonding to provide access to SAN drives on the second node.

To bond SAN drives in DC system:

1. Configure an iSCSI or FC connection point, select required targets and drives on one node in accordance with subsections [Configure FC Connection](#) or [Configure iSCSI Connection](#).
2. Add SAN targets FC or iSCSI connection point and SAN targets iSCSI on the second node.

- Click  in the string of the appropriate SAN target and click **Select Target on the Remote Node** (Figure 108).



The screenshot shows the RAIDIX web interface with the 'SAN OPTIMIZER' tab selected. Under 'SAN Targets', there is a table with one entry:

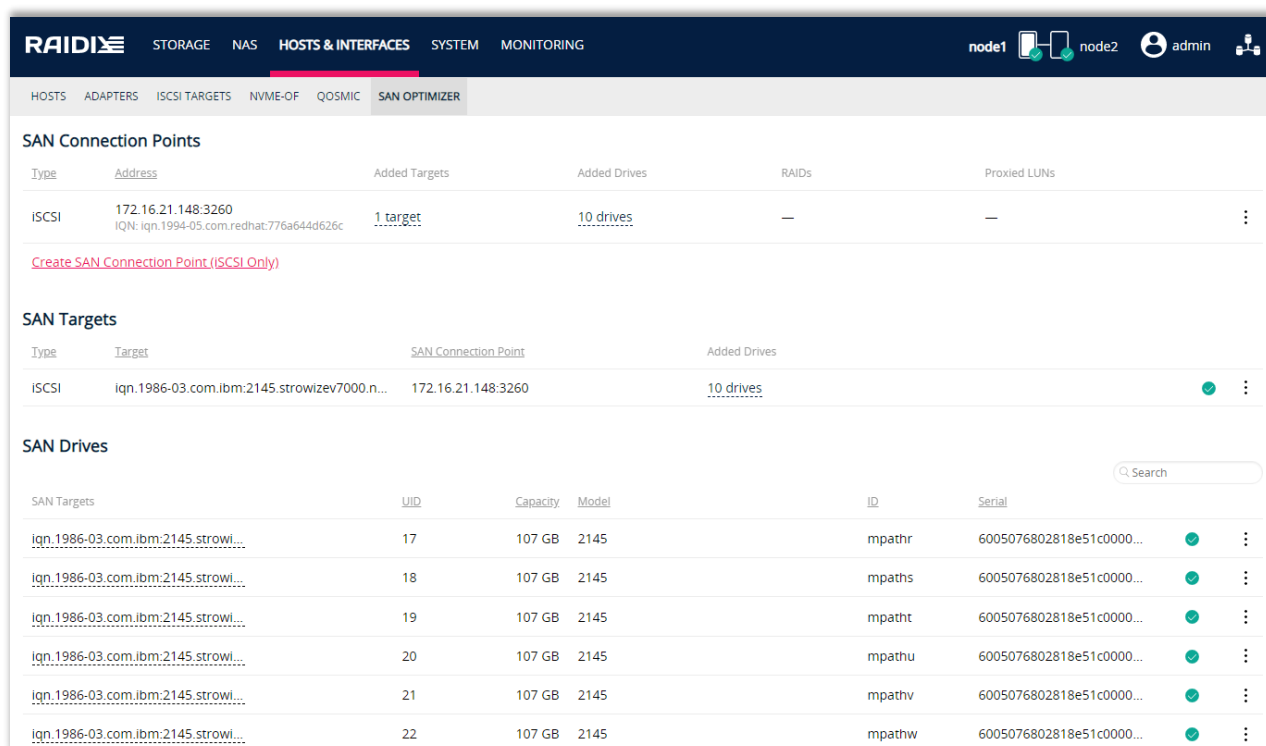
Type	Name	SAN Connection Point	Added Drives
FC	21:00:00:10:86:03:66:e0	10:00:00:10:86:03:a1:eb	—

A context menu is open over the first target, showing the following options:

- Select Drives
- Select Target on the Remote Node
- Delete

Figure 108. Selecting a target on the remote node

- In the opened window, select the appropriate target and click **Select**. After you add the targets, all SAN drives on the target become available and displayed in the section *SAN Drives* (Figure 109), on the page **STORAGE | DASHBOARD** in the section *Drives*, and on the **DRIVES** page in the corresponding section.



The screenshot shows the RAIDIX web interface with the 'SAN OPTIMIZER' tab selected. Under 'SAN Drives', there is a table listing 10 drives:

SAN Targets	UUID	Capacity	Model	ID	Serial
iqn.1986-03.com.ibm:2145.strowi...	17	107 GB	2145	mpathr	6005076802818e51c0000...
iqn.1986-03.com.ibm:2145.strowi...	18	107 GB	2145	mpaths	6005076802818e51c0000...
iqn.1986-03.com.ibm:2145.strowi...	19	107 GB	2145	mpathr	6005076802818e51c0000...
iqn.1986-03.com.ibm:2145.strowi...	20	107 GB	2145	mpathu	6005076802818e51c0000...
iqn.1986-03.com.ibm:2145.strowi...	21	107 GB	2145	mpathv	6005076802818e51c0000...
iqn.1986-03.com.ibm:2145.strowi...	22	107 GB	2145	mpathw	6005076802818e51c0000...

Figure 109. SAN drives on the second node after bonding

- Perform the bonding for every SAN target.

9.1.4 Creating Proxied LUN

It is possible to create RAIDs and LUNs using SAN drives. Furthermore, it is possible to create proxied LUNs.

i You can create only one proxied LUN on every SAN drive.

To create a proxied LUN use one of the following ways:

- From the **SYSTEM | LUN LIST** page:
 1. In the *Proxied LUNs* section, click **Create Proxied LUN**.
 1. In the opened widget, set up the parameters of the proxied LUN:
 - LUN name.
 - Block size.
 - Visibility to hosts.
 2. Click **CREATE**.
- From the page **HOSTS & INTERFACES > SAN OPTIMIZER**.
 1. In the *SAN Drives* section, in the string of the drive on which you want to create a proxied LUN, click **:** and select **Create Proxied LUN**.
 2. Follow the instructions of the wizard.

You can view and manage proxied LUNs on the page **STORAGE | LUN LIST**.

9.1.5 Deleting SAN Connection Points

To delete added SAN drives, SAN targets or SAN connection points, click **:** in the appropriate string and click **Delete**.

i While deleting SAN connection point, targets and drives related to the point will be also deleted. Proxied LUNs, RAIDs, and LUNs created on the SAN drives won't be deleted. They appear automatically when the connection point is re-added to the system.
Delete proxied LUNs on the page **STORAGE | LUN LIST**.

10. SETTING UP SYSTEM PARAMETERS

To set up the system parameters, use the **SYSTEM** menu. This section provides detailed description of RAIDIX 5.2 tasks:

1. [Working with nodes.](#)
2. [Setting up network parameters.](#)
3. [Working with a license.](#)
4. [Setting up notifications.](#)
5. [Configuration recovery.](#)
6. [Time settings.](#)
7. [Users.](#)
8. [System Drive.](#)

10.1 Working with Nodes

You can manage the nodes on the **SYSTEM | NODES** page. The page elements let you do the following operations with hosts:

1. Configuring Dual-Controller mode.
2. View the nodes current status.
3. Automatic node switching.
4. Manual node switching.
5. Disabling Dual-Controller mode.
6. Rebooting, shooting down, and powering on DC system.
7. Uninterruptible power source configuration.

10.1.1 Dual-Controller System Configuration

Before starting DC system configuration, it is necessary to:

1. Install the system.
2. Add license file(s) (for details, see *RAIDIX 5.2 Installation and Configuration Guide*).
3. Connect the nodes to enclosures and hypervisor volumes.

To start working with DC-system, [set up heartbeat connection](#) and [cache synchronization](#).



Use different paths for heartbeat and for cache synchronization.

Both connections between system nodes must be point-to-point.



After configuring dual-controller mode, you can manage both nodes from the **SYSTEM | NODES** page of any of them by clicking the links to the remote node.



If on one of the nodes a share has Active Directory users, you are unable to create a DC.
In this case, configure Active Directory on the second node as well or disable AD on the first node.

10.1.1.1 Setting up Heartbeat

You can set up heartbeat over Ethernet.

You can set up heartbeat in two ways:

- by using *widget* **Create DC** on the **SYSTEM | NODES** page;
- by using *wizard* **CONFIGURE DC** on the **SYSTEM | NODES** page.



While using two-port Mellanox VPI adapters, avoid the configuration with one port in Ethernet mode and another port in InfiniBand mode.

To set up heartbeat connection via the *widget*:

1. Set up the network interface on the first node:
 - 1.1. Select **SYSTEM > NETWORK**.
 - 1.2. In the *Network Interface Parameters* section, select the line of the port, that you want to use for heartbeat, then click and select **Edit**. The widget to set up the network interfaces parameters opens (Figure 110).

Port Name	Status	Link Speed	Enabled	DHCP	IP Address	Netmask	MTU
ens18	Up	1,000 Mb/s	<input type="checkbox"/>	<input type="checkbox"/>	172.16.22.143	255.255.252.0	1500

Figure 110. Setting up heartbeat connection parameters

- 1.3. Toggle the switch **Enabled** to the **On** position, then enter the IP address and Netmask parameters in the appropriate fields. To learn more, see [Set up network interfaces parameters](#).



Heartbeat IP address on both nodes must be from one subnet.

- 1.4. To save settings, click **APPLY**.
2. Repeat the step 1 on the second node.
3. Create DC:
 - 3.1. At any node, select **SYSTEM > NODES**.
 - 3.2. Click **Create DC**. The widget to add the remote node settings opens (Figure 111).

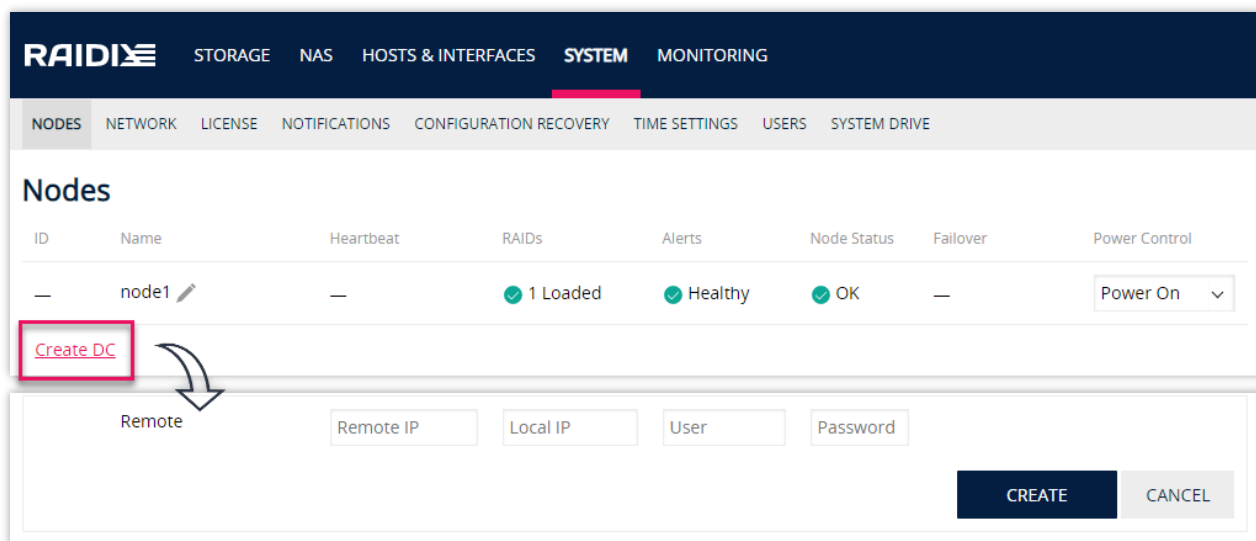


Figure 111. NODES page in SC

- 3.3. In the Remote IP field, enter the Heartbeat IP address of the remote node (specified in step 1 or 2) and click **CREATE**.

Parameters for the second node will appear in the table. Both nodes will get an automatically assigned ID shown in the ID column (Figure 119).

4. If necessary, enable synchronization of Persistent Reservations.

To learn more about this function, see the «Synchronization of Persistent Reservations» subchapter below.

Synchronization of Persistent Reservations



We recommend that you always enable synchronization of Persistent Reservations when using Windows or Hyper-V clusters as initiators.

The SCSI protocol provides a reservation mechanism defined in the SCSI-2 command set and then extended with updated functionality in SCSI-3.

SCSI Reservations (SCSI-2) is used to control access to a SCSI device (LUN) available to multiple initiators. An initiator sets a reservation on a LUN to prevent another initiator from making changes to the LUN. Normally, the same initiator would perform a release on the affected LUN.

SCSI-3 Persistent Reservations allows restoring the state of reservations even if a SCSI bus is reset (for example, when recovering from an error). This mechanism also supports multiple paths from initiator to device (LUN).

Synchronization of Persistent Reservations (hereafter – PR sync) keeps the same status of LUN reservations on the two nodes of the DC system. However, the LUNs is only accessed via Multipath Active-Passive, and ALUA Active-Active becomes unavailable.

To enable PR sync for a specific LUN:

1. Turn on PR sync support in the DC system:

- 1.1. Open the **SYSTEM | NODES** page.
- 1.2. Toggle **Persistent reservation synchronization** to the on position (Figure 112).
2. Enable PR sync for a LUN:

You can only enable PR sync with CLI.

```
$ rdcli lun modify -name <lun_name> --pr_sync 1
```

To learn more about the command, see [Appendix A](#).

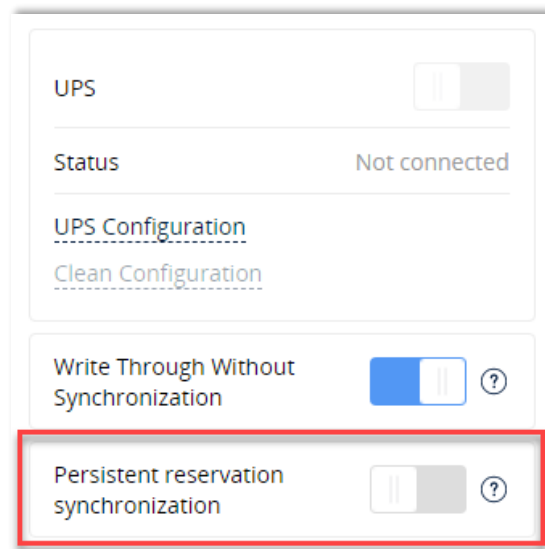


Figure 112. "Persistent reservation synchronization" switch on the NODES page

10.1.1.2 Setting up Cache Synchronization of Nodes

For the proper operation of the DC system, configure targets to synchronize node caches on each node. InfiniBand network, iSCSI, and SAS protocols are available for the synchronization.

Set up InfiniBand synchronization

! While using two-port Mellanox VPI adapters, avoid the configuration with one port in Ethernet mode and other port in InfiniBand mode.

i We recommend using a dedicated InfiniBand channel for cache synchronization.

To set up cache synchronization over InfiniBand:

1. At any node, select **HOSTS & INTERFACES > ADAPTERS**. At the InfiniBand Adapters section, each target corresponds to the one InfiniBand port (Figure 113).

RAIDIX STORAGE NAS **HOSTS & INTERFACES** SYSTEM MONITORING node1 node2 admin

HOSTS **ADAPTERS** ISCSI TARGETS NVME-OF QOSMIC SAN OPTIMIZER

iSCSI Ports

Port Name	MAC	IP Address	Netmask	Status	Policy	Link Speed
eno1	00:25:90:7b:c7:76	172.16.21.111	255.255.252.0	Up	—	1,000 Mb/s
eno2	00:25:90:7b:c7:77	12.12.12.11	255.255.255.0	Up	—	1,000 Mb/s
ib0	a0:00:02:10:fe:80:00:...	10.1.1.1	255.255.255.0	Up	—	56,000 Mb/s

InfiniBand Adapters

Device:Port	Port GUID	Node GUID	State	Link Status	Link Speed
mlx4_0:1	0025:90ff:ff17:7c95	0025:90ff:ff17:7c94	Active	LinkUp	56 Gb/s

Fibre Channel Adapters are not configured

SAS Adapters are not configured

Target Synchronization

OpenSM

QLogic Initiator Mode: Disabled

Celerity Initiator Mode: Disabled

Figure 113. ADAPTERS page

- To open a setup window (Figure 114), click **Target Synchronization**.

Select Targets for Synchronization

Type	WWN/GUID/IQN
<input type="checkbox"/> IB	fe80:0000:0000:0000:0025:90ff:ff17:7c95
<input type="checkbox"/> FC	20:00:00:10:86:02:94:3c

CANCEL SELECT

Figure 114. Select Targets for Synchronization window

- Select GUID of the corresponding InfiniBand ports for nodes' cache synchronization and click **SELECT**.
- Repeat steps 1-3 on the second node, selecting the appropriate InfiniBand GUIDs for synchronization.
- Start OpenSM service:
 - Select **HOSTS & INTERFACES > ADAPTERS**.
 - Click the **OpenSM** button. A window to specify OpenSM service parameters opens (Figure 115).

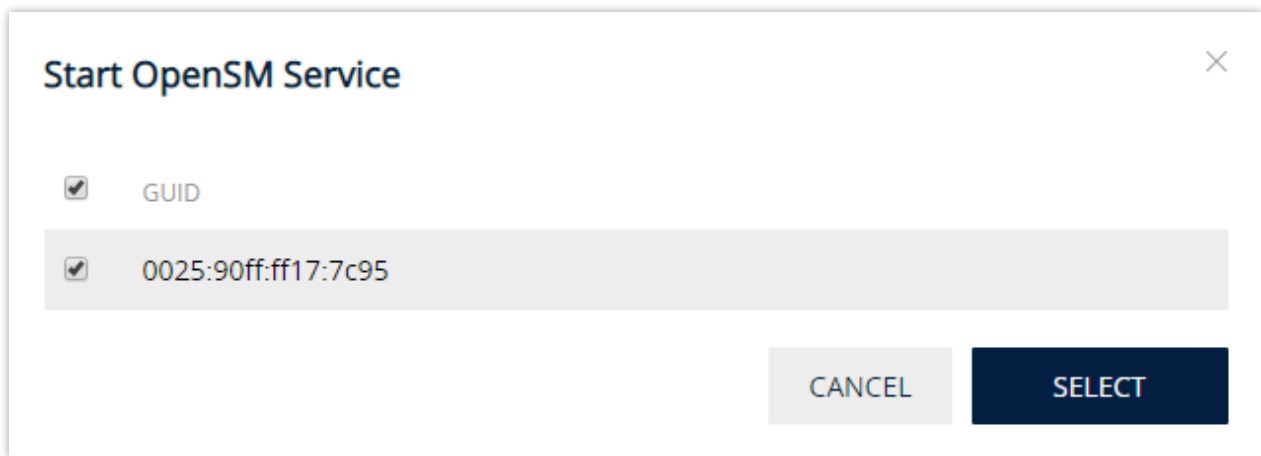


Figure 115. Start OpenSM Service window

5.3. Select ports' GUIDs for synchronization and click **SELECT**.

For cache synchronization, start OpenSM only on one node.

Set up iSCSI synchronization



Before setting up iSCSI synchronization, make sure that IP addresses of ports for synchronization are present in the list of IP addresses in iSCSI settings.



- Use different paths for heartbeat and for cache synchronization.
- Both connections between system nodes must be point-to-point.

To set up cache synchronization over iSCSI:

1. At any node, select **HOSTS & INTERFACES > ISCSI TARGETS**.
2. To turn on iSCSI, in the pane, toggle the **Enable iSCSI** switch to the **On** position (Figure 116).

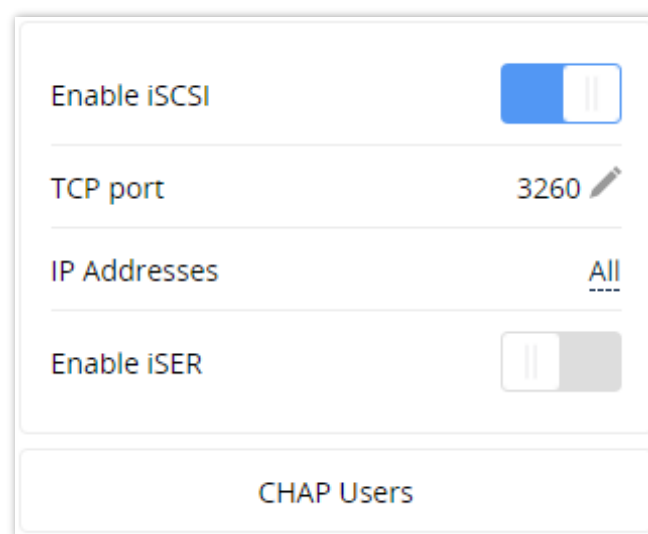


Figure 116. Enabling iSCSI

3. Select **HOSTS & INTERFACES > ADAPTERS**.
4. Click **Target Synchronization** to open a setup window (Figure 117).

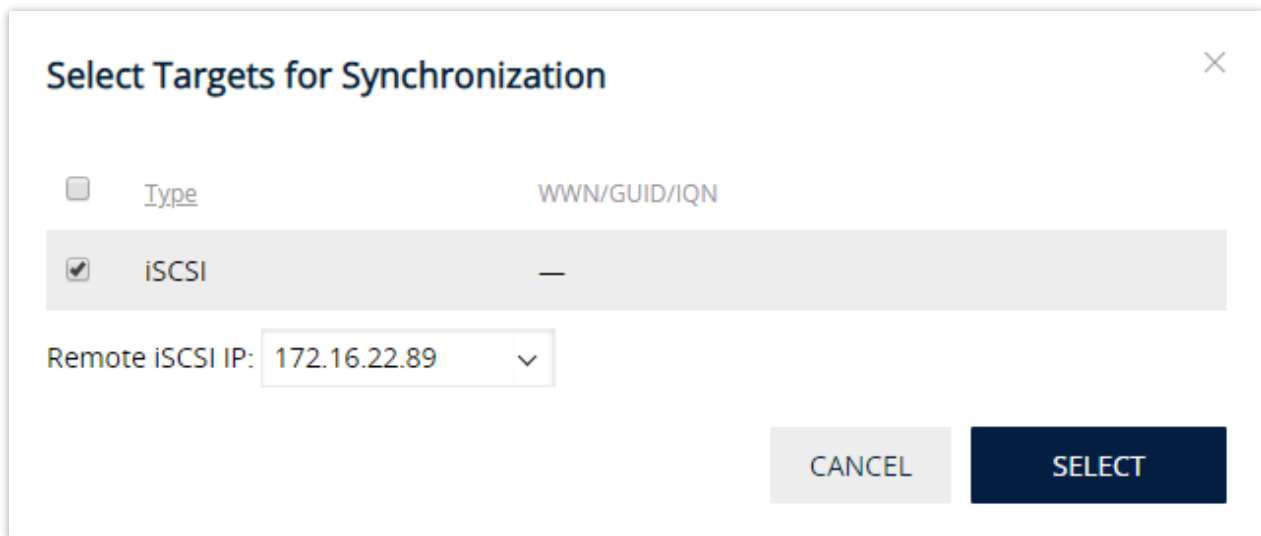


Figure 117. Targets for synchronization window

5. Select the iSCSI target type.
6. Enter Remote iSCSI IP for the nodes cache synchronization, and click **SELECT**.

! If the desired IP address of the port is off the list, verify that [network settings of the remote node](#) are correct. To check the network configuration, select **SYSTEM > NETWORK**.

7. Repeat the steps 3-6 on the second node.

Set up SAS synchronization

i We do not recommend using one SAS adapter at the same time to synchronize and connect the enclosures.
Such configuration may increase the RAID rebuild time with the possibility of a temporary change in RAID status (online/offline) when the node is rebooted.

Configuration of SAS adapter is described in the document *RAIDIX 5.2 Software Installation and Configuration Guide*.

To set up cache synchronization over SAS:

1. On one of the nodes select **HOSTS & INTERFACES > ADAPTERS**.
2. Click **Target Synchronization** to open a setup window.
3. Select **SAS Address** of the corresponding SAS ports for nodes' cache synchronization and click **SELECT**.



The address of the first phy in the group will be the address for the synchronization port.

For example, for 2 ports with 4 phys in each, the first port address will be phy0 and the second port address will be phy4.

- Repeat steps 1-3 on the second node.

10.1.2 Viewing Node Status

To see the information on node statuses, open the **SYSTEM | NODES** page (SC mode – Figure 118, DC mode – Figure 119).

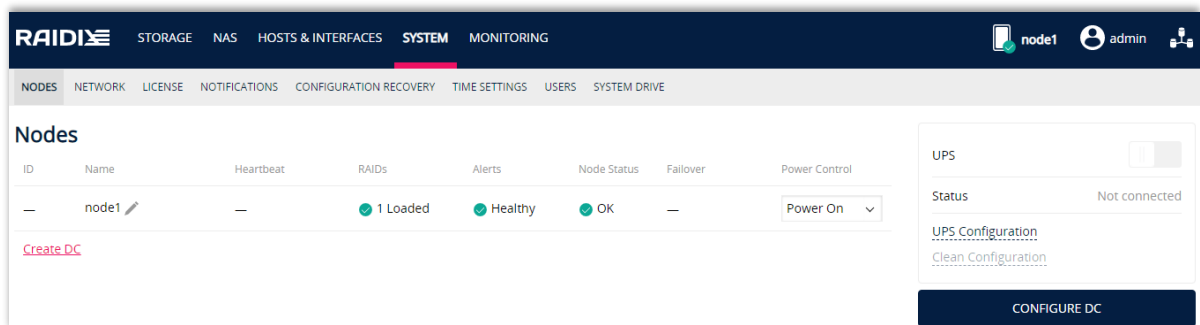


Figure 118. NODES page in SC mode

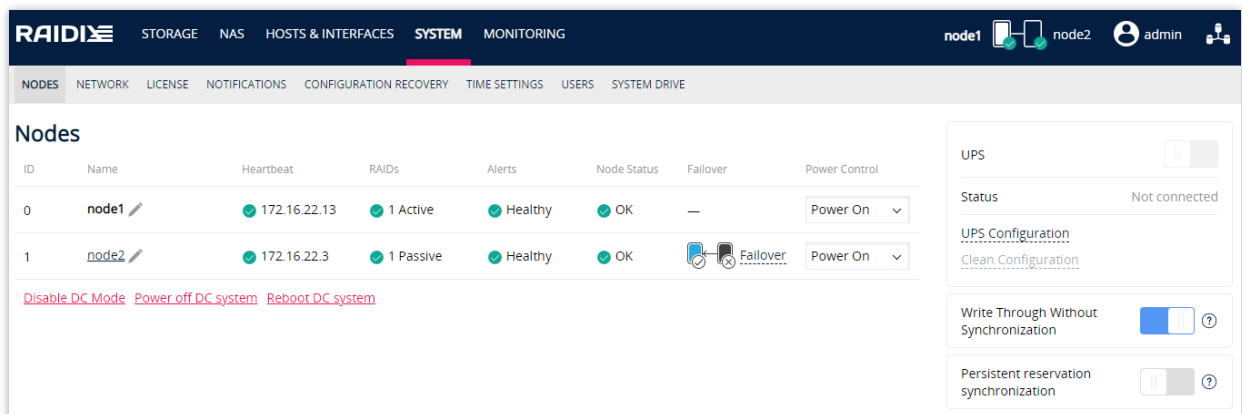








Figure 119. NODES page in DC mode

Parameter	Description
ID	Node identifier.
Name	<p>Node name.</p> <p>By default, the node is assigned a unique name consisting of six numbers and Latin letters.</p> <p>You can edit the name (only with numbers and Latin letters).</p> <p>For DC systems only: the name of the local (current) node is in bold. The name of the second (remote) node is a link to switch to remote node interface.</p>
Heartbeat	IP address for connection to the second node in cluster.

Parameter	Description
RAIDs	<p>Displays the amount of Active, Passive, Active / Passive Failed Over RAIDs on every node.</p> <div> Active/Passive Failed Over RAIDs appear after automatic or manual node failover, which causes an array becoming active on the node, opposite to its Affinity.</div>
Node Status	<p>The node status in dual-controller mode. Possible variants:</p> <div> <i>Healthy</i> – no migrated passive RAIDs with Affinity, corresponding to the node ID.</div> <div> <i>Failback</i> – a node Failover has occurred: the system has detected passive RAIDs on the node with Affinity, corresponding to the node ID. To return to the initial RAID state, click Failback.</div>

Parameter	Description
Alerts	<p>The general node status: license, drives, adapters, RAID. Possible variants are available:</p> <ul style="list-style-type: none">  <i>Healthy</i> – a node runs correctly;  <i>Warning</i> – a non-critical failure has occurred on the node. This status appears in the following situations: <ul style="list-style-type: none"> • RAID state is <ul style="list-style-type: none"> ○ Degraded; ○ Reconstructing; ○ Initializing; ○ SDC; ○ Passive Failed Over. • Adapter has the Down state (no connection). • License status is Warning because a hardware modification was made. • Drive state is <ul style="list-style-type: none"> ○ Reconstructing; ○ Dirty; ○ Faulty warning. • When you have done one change in the hardware configuration (the hardware configuration key is different from hardware license key to one element).  <i>Error</i> – a critical failure has occurred on the node. This status appears in the following situations: <ul style="list-style-type: none"> • no valid licenses; • license key is incorrect; • more than one hardware modification was made. • incorrect license data. • RAID is <ul style="list-style-type: none"> ○ Offline; ○ Not Loaded. • drive in RAID has <ul style="list-style-type: none"> ○ Failure; ○ Faulty error; ○ Error; ○ No disk. • in DC mode, Heartbeat connection interface is unavailable (Down).
Failover	<p>An option allows the nodes to change their cluster states. To learn more, see Manual Failover and RAID Migration.</p>

Parameter	Description
Power Control	The element is intended to reboot and turn off the node: <ul style="list-style-type: none">• Power On;• Restart;• Power Off.

10.1.3 Failover and Migration of RAID

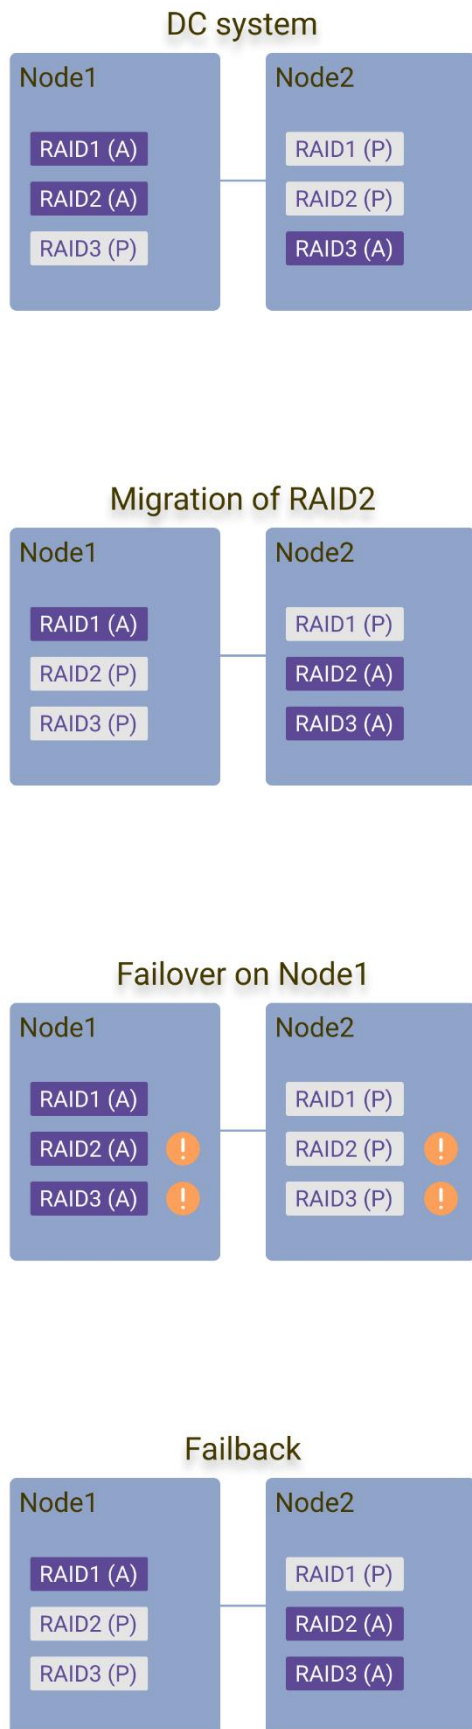
A distinctive feature of the DC system is the duplication of hardware components and interfaces, which contributes to the uninterrupted operation of the system, for example, in case of failure of one of the hardware components, failure of the interface connection to the drive shelves, or power failure of one of the nodes.

Only one node of the DC system can process I/O operations for the RAID at a time, while the second node stores a copy of that RAID. The RAID "affinity" property is responsible for which node processes the I/O and which node stores the copy. The RAID has the "Active" state on the node from which I/O operations are enabled, and has the "Passive" state if the RAID serves as a copy of the "active" RAID on that node.

To divide the RAID workload among the nodes, you can set the "affinity" property when you create the RAID, and you can change this property using the "migrate" function.

To temporarily redirect the I/O operations of all RAIDs to one node (for example, to maintain one node without stopping the entire system), use the "failover" function. The system will perform an "automatic failover" in case of a failure or abnormal termination of the previously active node or its component.

Use the "failback" function to return the I/O operations of all RAIDs to the nodes that were originally configured. This requires both nodes in the system to be operational.



Standard operation

RAID1 and RAID2 are active on Node1 (pointed with "A") and are passive on Node2 (pointed with "P"). RAID3 is active on Node2 and is passive on Node1 .

Migration

After the migration, the RAID has changed its affinity and has become active on the other node.

RAID2 has changed its affinity and has become active on Node2.

Failover

After the failover on Node1, all RAIDs that were active on Node2 have temporarily become active on Node1.

RAID 2 and RAID 3 have temporarily become active on Node1. The "Failed Over" states are displayed for these RAIDs now.

Failback

After the failback, all the RAIDs that had the "Failed Over" state have moved to their primary node.

RAID2 and RAID3 have become active on the Node2 as they were before the failover.

Figure 120. Example of migration and failover processes

10.1.3.1 Manual Failover

You can failover nodes manually if you want to perform system maintenance (including installing updates).

After switching nodes, all RAID5 become active on one of the nodes. The RAID5 whose "affinity" is not the same as the current node get the "Failed Over" status.



Perform the failover on the node on which you want to switch the workload.



Before manual failover or failback, ensure that active/optimized and active/unoptimized paths between the system and the host are operative.

However, there are situations where you can perform failover or failback with only one type of path between the system and the host:

- when the type of path or paths is the active/unoptimized;
- automatic failover.




During manual failover, the workload on the NFS share may interrupt.

To perform manual failover:

1. Select **SYSTEM > NODES**.
2. Click the **Failover** link in the line of the appropriate node (Figure 121).

On this node, all RAID5 will become active.

Nodes							
ID	Name	Heartbeat	RAID5	Alerts	Node Status	Failover	Power Control
0	node1	✓ 172.16.22.143	✓ 1 Active	✓ Healthy	✓ OK	—	Power On ▾
1	node2	✓ 172.16.22.153	✓ 1 Passive	✓ Healthy	✓ OK		Power On ▾

[Disable DC Mode](#) [Power off DC system](#) [Reboot DC system](#)

Figure 121. Performing manual Failover from the NODES page

3. To confirm, click **PERFORM**.

The other node will become ready for the maintenance.

10.1.3.2 Automatic Failover

In case of failing of one node, the system will perform the "automatic failover". The second node will take over the load of the failed node which means that all RAID5 that were active on the failed node become "active" (available for I/O operations) on the second node, and the "affinity" of those RAID5 does not change. This feature enables you to fix the problem without stopping the system.



Auto-failover will not start if the DC system is powered off or is rebooted and only one node is loaded after. To learn more, see [troubleshooting](#).



During automatic failover, the workload on the NFS share may interrupt.

10.1.3.3 Failback

After failover, the system switches to "Manual Mode".

To return the initial RAID states, click the **Failback** link on any page of the web interface of any node (Figure 122). To do so, both nodes must be operationable.



Before manual failover or failback, ensure that active/optimized and active/unoptimized paths between the system and the host are operative.

However, there are situations where you can perform failover or failback with only one type of path between the system and the host:

- when the type of path or paths is the active/unoptimized;
- automatic failover.



During failback, the workload on the NFS share may interrupt.

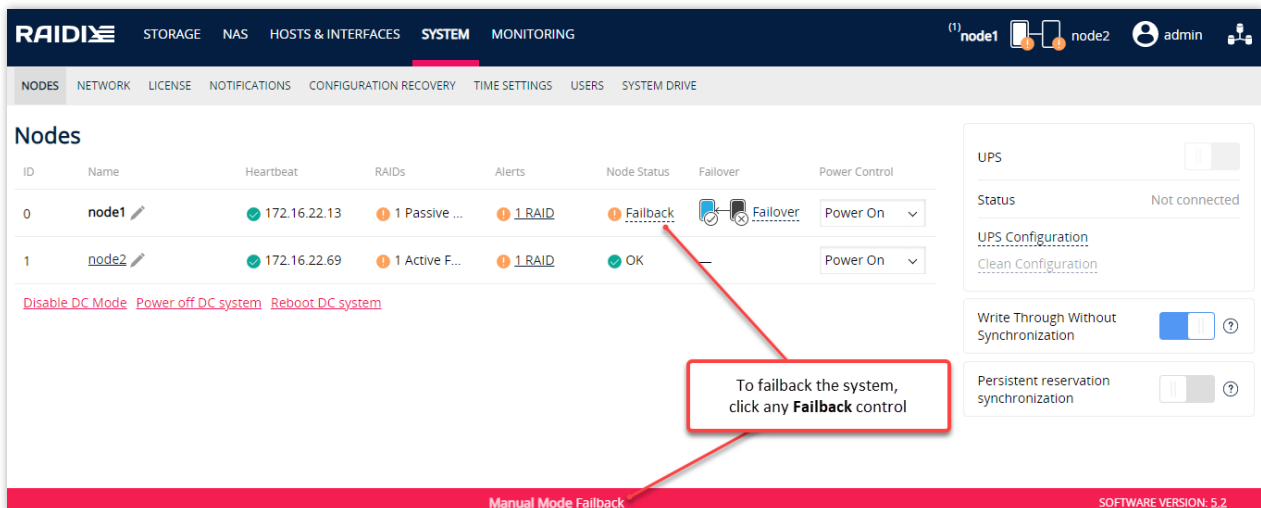


Figure 122. Performing failback

10.1.3.4 Migration of RAID

Migration changes the "affinity" of the RAID, thus changing the node that processes I/O operations on that RAID.

Specifics of the RAID migration:

- Migration of the RAID that has the LUN extended to one or several RAIDs.

When you migrate one of these RAIDs to another node, all RAIDs with that LUN will be automatically migrated.

- Migration of the RAID with the LUN(s) that is connected to the SSD cache.

To migrate the RAID, disconnect the SSD cache from the LUN(s).

- Migration of the RAID with the LUN(s) with the configured replication.

To migrate the RAID, disable the LUN(s) replication.

To migrate the RAID to the second node:

1. Select **STORAGE > RAID LIST**.
2. Click the name of the RAID to be migrated. The RAID page opens.
3. On the RAID page, click **Migrate RAID** (Figure 123).
4. In the confirmation window, click **MIGRATE**.

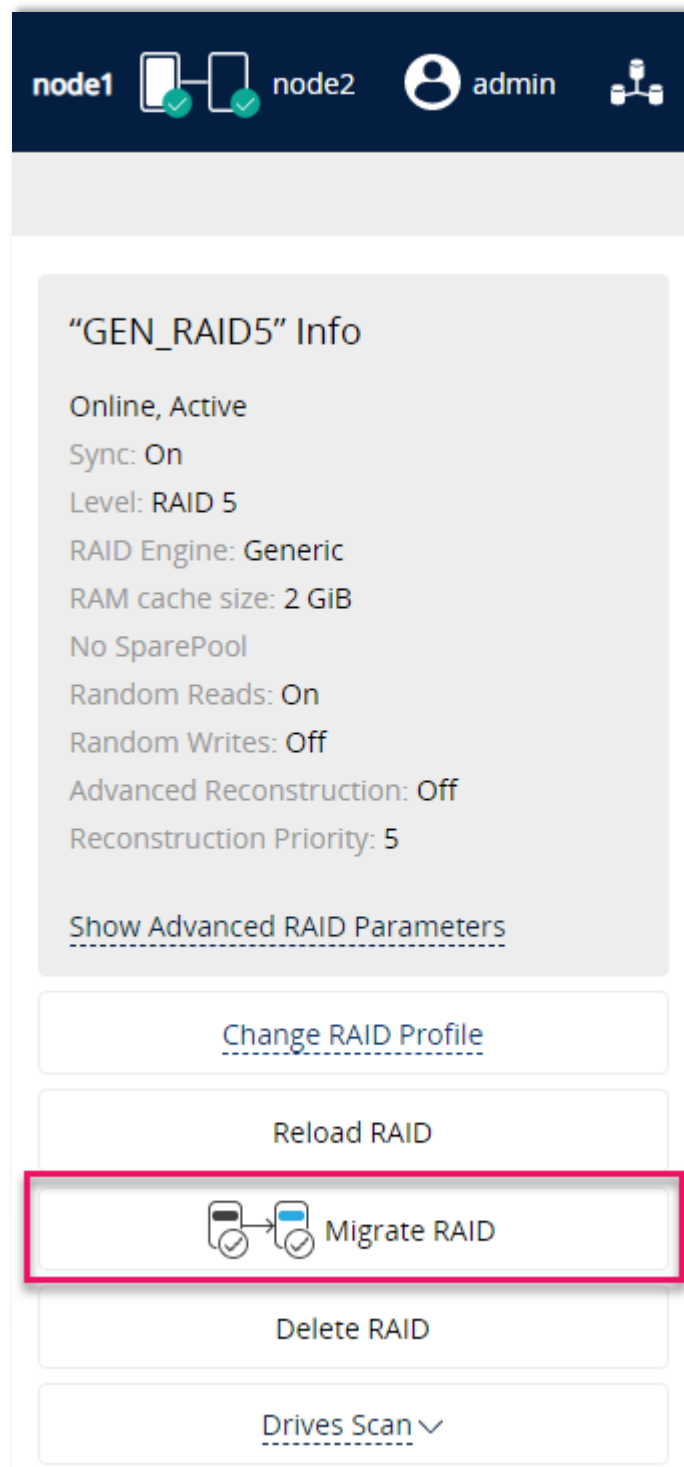


Figure 123. RAID Migration button on the RAID LIST page

10.1.4 Disabling Dual-Controller Mode

1. On any node, select **SYSTEM > NODES**.
2. Click **Disable DC Mode** (Figure 124).
3. To confirm, click **DISABLE**.



If each node in DC system has active RAID (named Active-Active DC system), you can't disable DC mode from GUI, because to reconfigure DC mode, one of the nodes must have no RAID(s). When you try disabling DC mode for Active-Active DC system, the system will show attention message:

Attention! You can't reconfigure DC mode while each node has active RAID(s). To disable Active-Active DC mode, migrate all RAID(s) on one node.

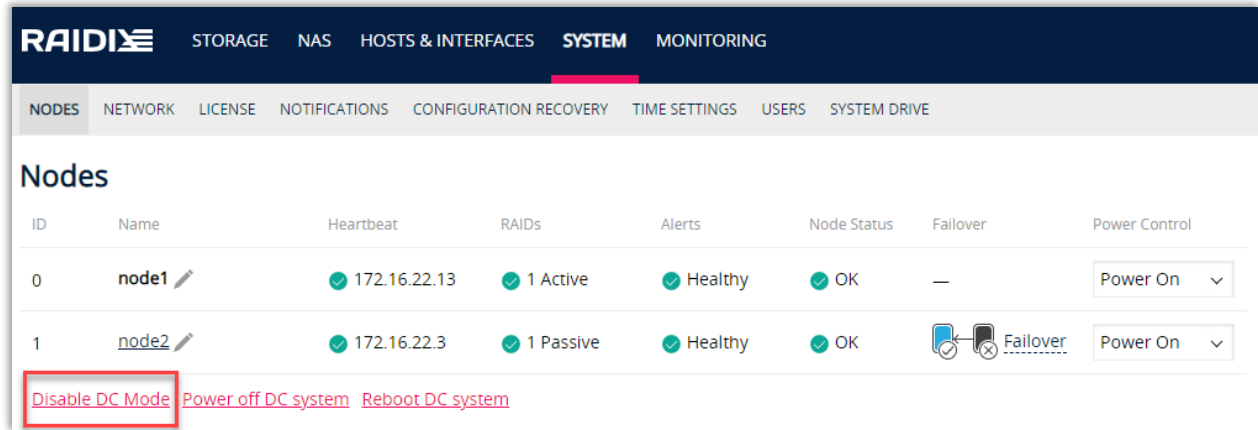


Figure 124. "Disable DC mode" control



If there are RAID(s) in the system, which are active on the different nodes, then at presence of SSD cache on these RAID(s), DC disabling after failover may take a long time due to SSD cache flushing to the drives.

10.1.5 Rebooting, Shutting down, and Powering on DC System

If you need to shut down DC system (for example, to change the server room), we recommend following the shutdown and power-on procedures described below.

Rebooting

To reboot a SC system or a node:

1. Open the **SYSTEM | NODES** page.
2. In the **Power Control** field, select **Reboot**.
3. Confirm your intention.

To reboot a DC system:

1. Open the **SYSTEM | NODES** page.
2. Click **Reboot DC system**.
3. Confirm your intension.

Shutting down

To shut down a SC system or a node:

1. Open the **SYSTEM | NODES** page.
2. In the **Power Control** field, select **Power Off**.
3. Confirm your intention.
4. If necessary, shut down the drive enclosure.

To shut down a DC system:

1. Make sure the DC system is fully operational (heartbeat is on, no failovered RAIDs).
2. Open the **SYSTEM | NODES** page.
3. Click **Power off DC system** and confirm the shutdown.
4. If necessary, shut down the drive enclosure.

Powering on

To power on a SC system or a node:

1. Make sure all necessary node interfaces are connected.
2. Power on the drive enclosure or make sure the enclosure is on.
3. Power on the node.

To power on a DC system:

1. Make sure all necessary node interfaces are connected, including heartbeat.
2. Power on the drive enclosure or make sure the enclosure is on.
3. Power on the nodes in any order.

10.1.6 Uninterruptible Power Source Configuration

The section describes how to connect the system to an uninterruptible power source (UPS). The option is disabled by default.

You can reset the configuration to the default settings by clicking **Clean Configuration**.

To configure UPS:

1. Connect UPS to the system following the procedure described in the appropriate UPS documentation. RAIDIX 5.2 supports the following types of UPS connection: through COM-port and USB.
2. Select **SYSTEM > NODES**.
3. On the *UPS* pane (Figure 125), click **UPS Configuration**. The UPS window opens (Figure 126).

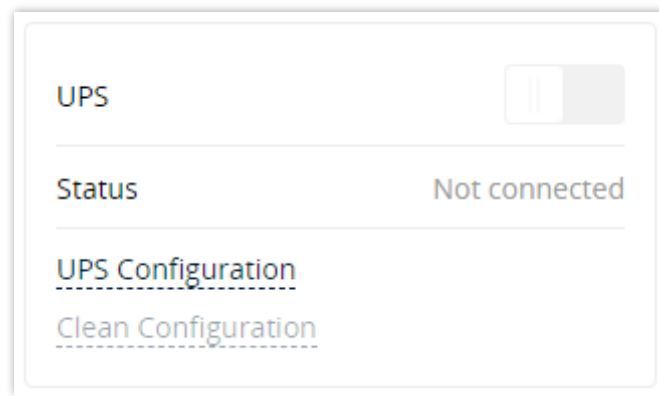


Figure 125. UPS configuration widget

Figure 126. Configuring UPS

4. In the **Vendor** field, select a vendor of the UPS.
5. In the **Model** field, select the UPS model being connected.
6. If you have a large UPS that is capable of running multiple systems simultaneously, specify Master/Slave configuration settings:
 - if the current system is Master, toggle the switch **Master** to the **On** position;
 - if the current system is Slave, in the **Master IP** field, specify the IP address of the Master system.
7. If you have selected a model with connection via COM-port, specify the port for connection in the **Port** field.
8. Wait for status *Connected* in the *Status* field: the UPS connection setup is successfully completed.

10.2 Setting up Network Parameters

To set up the network parameters, select **SYSTEM > NETWORK** (Figure 127).

The elements on the page allow you to do the following:

1. [Setting up the General Network Settings.](#)
2. [Setting up SSL certificate.](#)
3. [Setting up Virtual IP address.](#)
4. [Setting up Network Interfaces Parameters.](#)
5. [Setting Up Channel Bonding.](#)

The screenshot shows the RAIDIX web interface with the **SYSTEM > NETWORK** menu selected. The page is divided into several sections:

- General Network Settings:** A table with columns for Primary Interface, IP Gateway, Primary DNS, Secondary DNS, Tertiary DNS, and SSL Certificate. The Primary Interface is `eno1`, IP Gateway is `172.16.21.254`, Primary DNS is `172.16.20.3`, Secondary DNS is `172.16.20.4`, and Tertiary DNS is `172.16.20.252`. There is a link to [Upload SSL certificate](#).
- Virtual IP Settings:** A section with a link to [Create Virtual IP address](#).
- Network Interfaces Parameters:** A table with columns for Port Name, Status, Link Speed, Enabled, DHCP, IP Address, Netmask, MTU, and Flow Control.

Port Name	Status	Link Speed	Enabled	DHCP	IP Address	Netmask	MTU	Flow Control
eno1	Up	1,000 Mb/s	Yes	Yes	172.16.22.2	255.255.252.0	1500	Auto, Tx: Off, Rx: Off
eno2	Down	—	No	Yes	—	—	1500	Disabled
ens5	Down	—	No	No	—	—	1500	Tx: On, Rx: On
ens6	Up	100,000 Mb/s	Yes	No	11.0.0.1	255.255.255.0	9000	Tx: On, Rx: On
- Bonding Settings:** A section with a link to [Create Bond](#).

Figure 127. NETWORK page

10.2.1 General Network Settings

! The real (not virtual) IP addresses of two or more interfaces must belong to different IP networks.

To set up the general network parameters:

1. Select **SYSTEM > NETWORK**.
2. In the *General Networking Settings* section, click **:** in the appropriate interface line, and then click **Edit**. The pane to set up the network parameters opens (Figure 128).

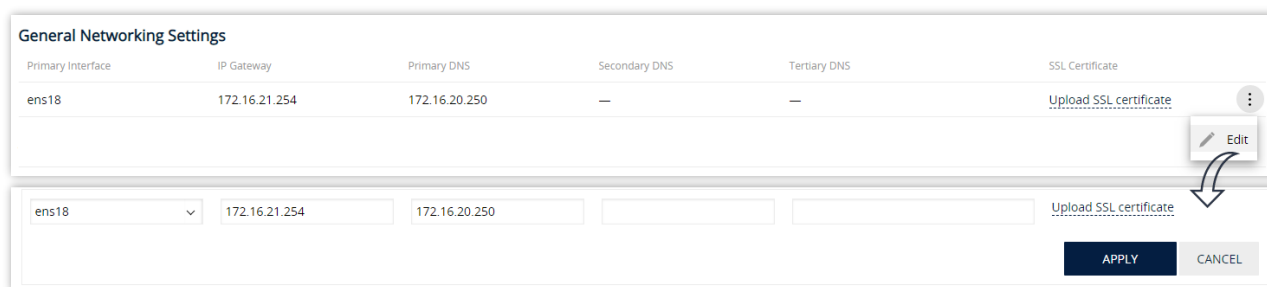


Figure 128. Network page. General networking settings

- Fill in the following fields and click **APPLY**.

Parameter	Description
Primary Interface	The main Ethernet interface of system connection. Used to perform administrative tasks.
IP Gateway	IP gateway to access external network.
Primary DNS	Primary DNS Server IP address.
Secondary DNS, Tertiary DNS	Alternative DNS Server IP addresses.

10.2.2 Setting up SSL Certificate

RAIDIX 5.2 uses self-signed SSL certificate by default. The certificate is always presented in the system and only being replaced after uploading another certificate. When uploaded certificate deletes, the system SSL certificate become active again.

To upload your own SSL certificate:

- Select **SYSTEM > NETWORK**.
- In the *General Network Settings* section, at the Primary Interface item, click **Upload SSL certificate** (Figure 129).

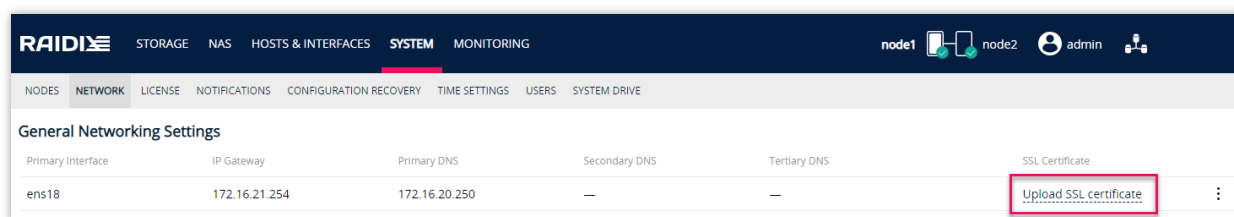


Figure 129. Upload SSL certificate

- Window to upload SSL certificate opens (Figure 130). Click **Select files** and select appropriate SSL files. You need to upload two files – *.key and *.crt.

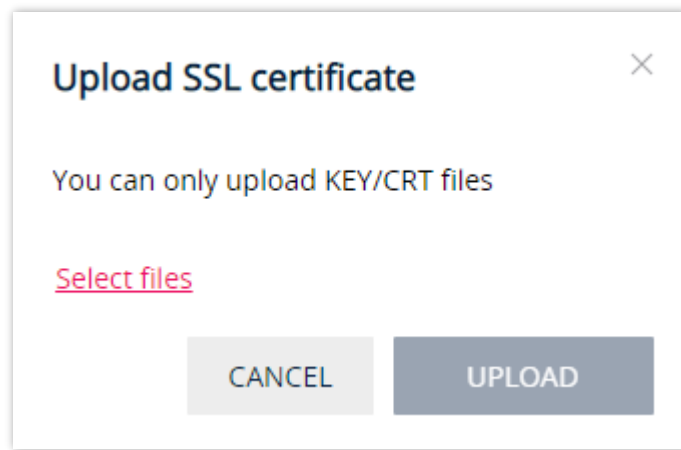


Figure 130. Uploading SSL certificate window

4. After selecting files, click **UPLOAD** (Figure 131).

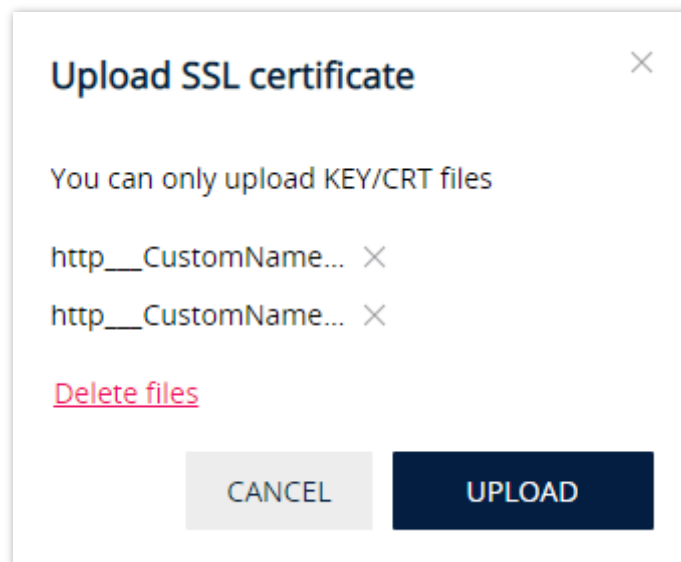


Figure 131. Uploading SSL certificate window after selecting the files

To delete an SSL certificate:

1. Select **SYSTEM > NETWORK**.
2. In the *General Network Settings* section, in the Primary Interface line, click **Remove SSL certificate** (Figure 132).

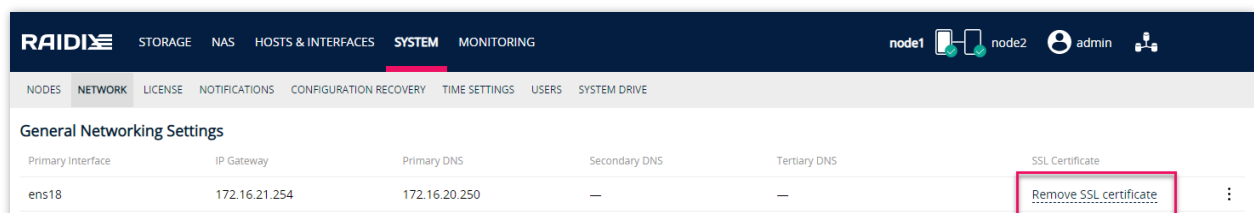


Figure 132. Remove SSL certificate

3. To confirm, click **DELETE**.

10.2.3 Setting up Virtual IP Address

Set up a virtual IP address to:

- configure multiple IP addresses on a single network interface;
- work with Network Attached Storage (NAS) in DC system;
- set up Replication in DC system.

To create a new virtual IP:

1. Select **SYSTEM > NETWORK**.
2. In the *Virtual IP Settings* section, click **Create Virtual IP-address**. A widget for creation a new IP opens (Figure 133).

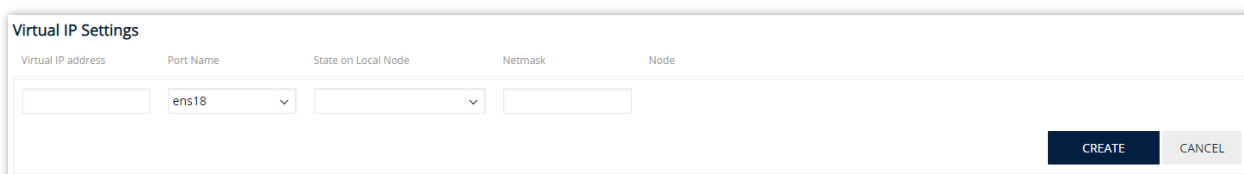


Figure 133. Widget to create a new virtual IP

3. Define the following parameters of NAS in DC mode:

Parameter	Description
Virtual IP address	Virtual IP address that will be assigned to the network storage.
Port name	A local port name for access to NAS. On this port, the NAS virtual IP in regular system mode (before nodes failover) will be set up.
State on local node	State of the NAS on the local node (analogue of the spare value attribute while configuring from CLI). Possible values: <ul style="list-style-type: none"> • <i>Active</i> – work with NAS resources from this node is possible. • <i>Passive</i> – work with NAS resources from this node will become possible only after node failover.
Netmask	Netmask to be assigned to the virtual interface.

4. Click **CREATE**.
5. Repeat steps 1-4 on the second node.

To edit a NAS virtual IP:

1. Select **SYSTEM > NETWORK**.
2. In the required virtual IP line, click **:** and select **Edit**.

3. Change required parameter values and click **APPLY**.



In DC mode, always change virtual IP on both nodes. Different virtual IPs on the nodes can lead to malfunction of RAIDIX features, including NAS and replication.

To delete a virtual IP address

1. Select **SYSTEM > NETWORK**.
2. In the line of a VIP, click and select **Delete**.
3. To confirm, click **DELETE**.

10.2.4 Setting up Network Interface Parameters

To set up the network interfaces parameters:

1. Select **SYSTEM > NETWORK**.
2. Click in the appropriate Ethernet port line, and then click **Edit**. The pane to set up the network interfaces parameters appears (Figure 134). You can set up the following parameters:

Port Name	Status	Link Speed	Enabled	DHCP	IP Address	Netmask	MTU	Flow Control
eno1	Up	1,000 Mb/s	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	172.16.22.2	255.255.252.0	1500	Auto

Figure 134. Network page. Network Interfaces parameters

3. Fill in the required fields and click **APPLY**.

Parameter	Description
Port Name	Ethernet port name.
Status	<p>A connection status of the appropriate Ethernet port.</p> <p>Possible values:</p> <ul style="list-style-type: none"> • <i>Up</i> – a connection is established. • <i>Down</i> – there is no connection. Possible reasons: the cable is switched incorrectly or the driver is missing. • <i>Unknown</i> – link status is not defined.
Link Speed	Link Speed, in Mbps.
Enabled	<p>A parameter stating whether the port is enabled.</p> <ul style="list-style-type: none"> • On – the port is enabled. • Off – the port is disabled.

Parameter	Description
DHCP (Dynamic Host Configuration Protocol)	The parameter states whether the port address is specified using DHCP protocol: <ul style="list-style-type: none"> • On – the address is received with the DHCP protocol. • Off – the address is obtained without using the DHCP protocol.
IP Address	IP address of the interface.
Netmask	Netmask.
MTU (Maximum Transmission Unit)	The maximum size of data transmission unit.
Flow Control	The flow control mode of the Ethernet interface. Possible values: <ul style="list-style-type: none"> • Rx – handle the received pause frames on this port. • Tx – send pause control frames from this port to the connected device when the received traffic increases. • Auto – standard IEEE 802.3 autonegotiation for flow control

10.2.5 Setting up Channel Bonding

Channel bonding enables aggregation of multiple network interfaces into a single logical bonded interface. This allows you to distribute the load between these channels, increasing capacity in both directions. Bonding may improve system performance.

To configure channel bonding:



1. Select **SYSTEM > NETWORK**.
2. Click **Create Bond**. The pane for bonding shows (Figure 135).

Figure 135. NETWORK page: Bonding Settings section

3. In the **Policy** field, select from the drop-down the type of bonding policy. The following types of bonding policy are available:

Type	Description
Balance-rr (Round-Robin policy)	Transmit packets in sequential order from the first available slave interface through the last. This mode provides load balancing and fault tolerance.

Type	Description
Active-backup (Active-backup policy)	Only one slave interface in the bond is active. Another slave becomes active only in case of failure of the current the active slave. The bond's MAC address is externally visible on only one port (network adapter) to avoid confusing the switch. This mode provides fault tolerance.
Balance-xor (XOR policy)	Transmit based on the selected transmit hash policy. This mode provides load balancing and fault tolerance.
Broadcast (Broadcast policy)	Transmits everything to all slave interfaces. This mode provides fault tolerance.
802.3ad IEEE 802.3ad dynamic link aggregation	Creates aggregation groups that share the same speed and duplex settings. Utilizes all slaves in the active aggregator according to the 802.3ad specification. This mode provides load balancing and fault tolerance and requires a switch with IEEE 802.3ad and LACP.
Balance-tlb (Adaptive transmit load balancing)	Channel bonding that does not require any special switch support. The outgoing traffic is distributed according to the current load (computed about the speed) on each slave interface. Incoming traffic is received by the current slave. If the receiving slave fails, another slave takes over the MAC address of the failed receiving slave.
Balance-alb	Adaptive load balancing: includes balance-tlb plus receives load balancing (rlb) for IPV4 traffic, and does not require any special switch support. The receive load balancing is achieved by ARP negotiation.

4. In the **Interfaces** field, click **Add Slave** and specify Ethernet interfaces to be bound as a slave.
5. Click **Create**.
6. To unbind the interfaces, click  in the appropriate bond line, and then click **Unbond**. Click **UNBOND** in the confirmation window.
7. To edit the bonding parameters, click  in the appropriate bond line and select **Edit**.
8. After changing the settings, click **APPLY**.



Once the bond has been created, it will additionally appear in the *Network Interfaces Parameters* section.

If you disable the bond through the corresponding object in the *Network Interfaces Parameters* section, the bond will be unbonded (bond interfaces will be removed).

10.3 Working with License

The license enables a set of features available on your system.



After installing the system, the drives that are available to the system *will not* be displayed until the license file is uploaded. To show available drives, upload the license file to the system.

The RAIDIX software with the applied license can be installed on only one hardware platform (node) and can't be reinstalled on others. Within one platform, you can change the following "insufficient" system components without having to re-issue a license:

Insufficient system components:

- network adapters;
- drives (except the system drives);
- RAM modules.

Changing the following "sufficient" components may (depending on the type of RAID) affect the need to reissue the license.

Sufficient system components:

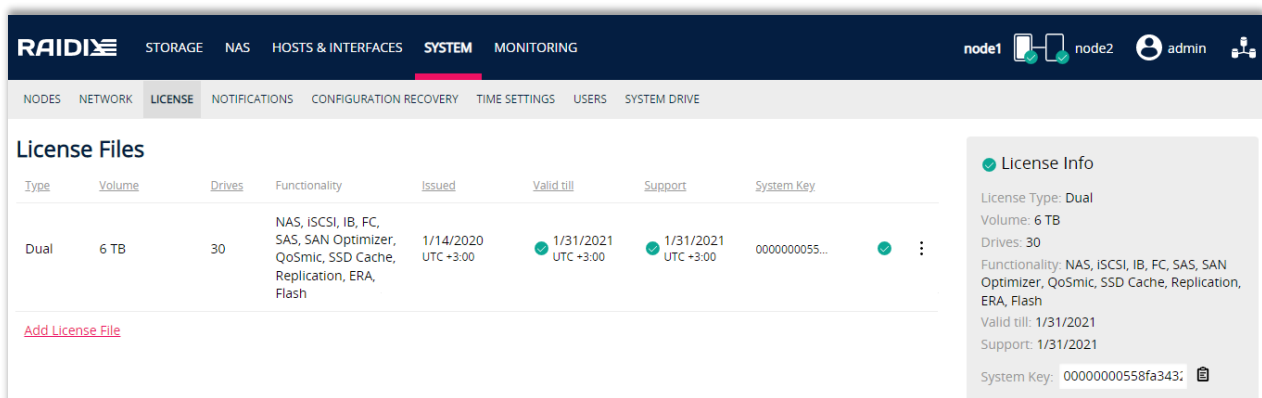
- motherboard;
- system chassis;
- CPU;
- BIOS (vendor, version, date);
- system drives.

Specifics of the license reissue depending on RAID type:

- Generic RAID: you can swap up to two sufficient hardware components without reissuing the license.
- ERA RAID: any change to a sufficient hardware component requires the license to be reissued.
- After the license has become invalid, RAID's will be in read-only mode until the license is reissued.

10.3.1 Viewing Product License Information

To view the information on the license, select **SYSTEM > LICENSE** (Figure 136).



Type	Volume	Drives	Functionality	Issued	Valid till	Support	System Key
Dual	6 TB	30	NAS, iSCSI, IB, FC, SAS, SAN Optimizer, QoSmic, SSD Cache, Replication, ERA, Flash	1/14/2020 UTC +3:00	1/31/2021 UTC +3:00	1/31/2021 UTC +3:00	0000000055...

[Add License File](#)

License Info

License Type: Dual

Volume: 6 TB

Drives: 30

Functionality: NAS, iSCSI, IB, FC, SAS, SAN Optimizer, QoSmic, SSD Cache, Replication, ERA, Flash

Valid till: 1/31/2021










Support: 1/31/2021

System Key: 00000000558fa343

Figure 136. License page with a loaded license file

The *License Files* section provides detailed information on each downloaded license file. The *License Info* panel provides combined information about the functionality and services from all the license files.

Parameter	Description
Type	License type: <ul style="list-style-type: none"> • Single – available in the single-controller configuration; • Dual – available in the dual-controller configuration.
Volume	Available volume for all LUNs.
Drives	Number of drives available for RAID creation.
Functionality	Available functionality for current license. <ul style="list-style-type: none"> • GRAID – “Generic RAID” option. See RAID Engines chapter for details. • ERA – “ERA RAID” option. See RAID Engines chapter for details. • Flash – extends “ERA” option by allowing you to create ERA RAID on SSDs. • FC – FCP/FC protocol support for SAN. • IB – support for SRP and iSER (over IB) protocols for SAN and Generic RAID cache synchronization in DC-system over iSER. • iSCSI – support of iSCSI protocol for SAN and Generic RAID cache synchronization in DC-system. • NAS – NAS file protocols support: SMB, NFS, FTP, AFP. • NVME-oF – NVMe protocols support (NVMe/TCP, NVMe/RDMA) for SAN. • VM Management – option for creation and managing virtual machines. • Replication – option for replication LUNs between multiple RAIDIX systems. • SAS – SAS protocol support for connecting disks and disk enclosures. • SAN Optimizer – “SAN Optimizer” option. See chapter Working with SAN Optimizer for more information. • SSD Cache – “SSD Cache” option for Generic RAID. • QoSmic – “QoSmic” option. See Setting up QoSmic chapter for details.
Issued	License issue date.

Parameter	Description
Valid till	<p>License expiration date and current status.</p> <p>Graphical information represents:</p> <ul style="list-style-type: none">  The license is valid.  The license will expire soon. The status shows up 31 days before the expiration.  The license has expired.
Support	<p>Customer support expiration date and current status:</p> <ul style="list-style-type: none">  The support is valid.  The support will expire soon. The status shows up 31 days before the expiration.  The support has expired.
System Key	<p>License request-key.</p> <ul style="list-style-type: none">  The key is valid; no hardware modifications were made.  The key is valid; one hardware modification was made.  <ul style="list-style-type: none"> The key is not valid; more than one hardware modifications were made. The key is not valid; the license has expired. The key is incorrect.

10.3.2 License Management

License management consists of the following actions:

- Adding, updating, and reissuing the license.
- Downloading a license file.
- Deleting a license file.

Adding, updating, and reissuing the license


You need to update the license file in order to upgrade hardware components and for renewing the license. This requires reissuing a new license file for your system.

To add a license file to the system:

1. Copy the current System key (on **LICENSE, Key** field).
2. Send the System key to RAIDIX Support department.
3. On the **LICENSE** page, apply the received license file by clicking **Add License File**.

Downloading a license file

To download an uploaded license file:

1. Select **SYSTEM > LICENSE**.
2. On the opened page, click the icon  and select **Download** (Figure 137).

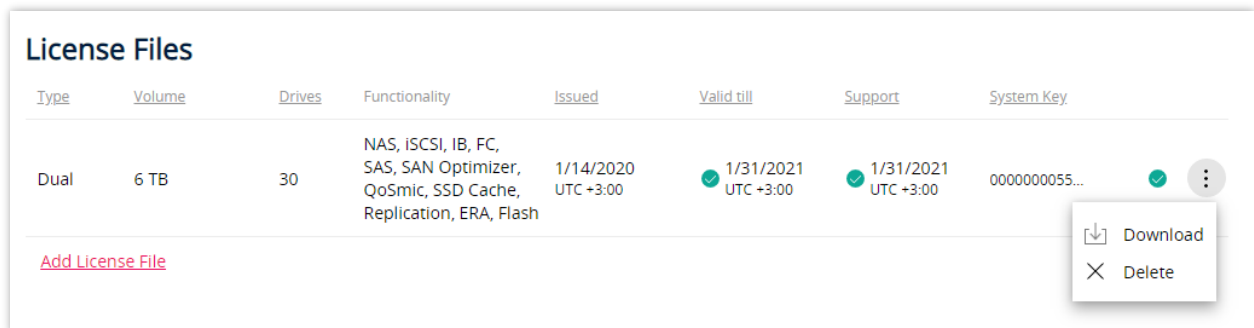



Figure 137. Download the License File

Deleting a license file

To delete a license file from the system:

1. Select **SYSTEM > LICENSE**.
2. On the opened page, click  and select **Delete** (Figure 137).

10.4 Setting up Notifications

RAIDIX 5.2 supports sending notifications on the system state messages via SMTP and SNMP.

For each RAIDIX software component, you can select the notification message types: **Info**, **Warning**, or **Error**.

List of the components is given below:

- UPS;
- RAID;
- LUN;
- Drive;
- Guest VMs;
- SparePool;
- SAN optimizer;
- iSCSI;
- Network;
- Controller Sensors;
- Enclosure Sensors;
- Samba share;

- AFP share;
- FTP share;
- NFS share;
- NVMe-oF;
- License;
- DC.

System notifications also are available for receiving.

10.4.1 Setting up SMTP Server Parameters

To set up SMTP server notification parameters:

1. Select **SYSTEM > NOTIFICATIONS**.
2. In the upper left, select **SMTP** (Figure 138).

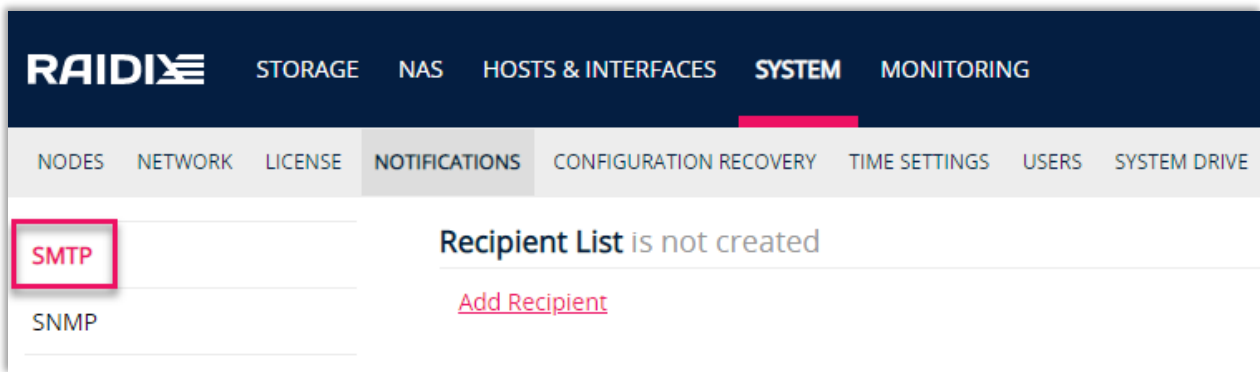
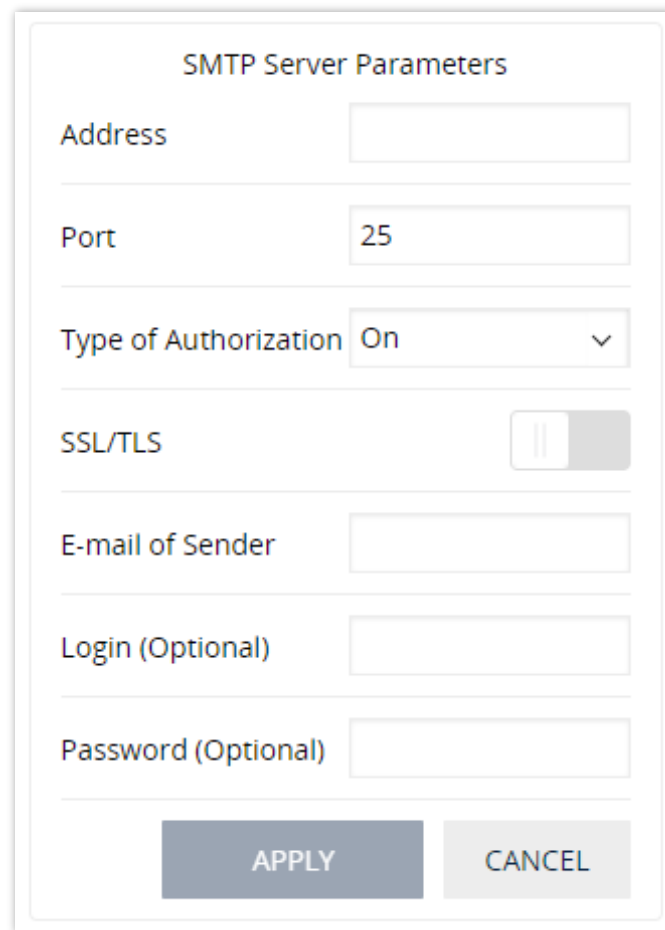


Figure 138. Configuring SMTP server

3. Move to the *SMTP Server Parameters* widget in the upper right (Figure 139).



The image shows a dialog box titled "SMTP Server Parameters". It contains several input fields and a toggle switch. The fields are: "Address" (empty), "Port" (25), "Type of Authorization" (On), "E-mail of Sender" (empty), "Login (Optional)" (empty), and "Password (Optional)" (empty). There is a toggle switch for "SSL/TLS" which is currently turned off. At the bottom, there are two buttons: "APPLY" and "CANCEL".

Figure 139. SMTP Server Parameters

4. Define the following server parameters:

Parameter	Description
Address	IP address or DNS name of the system from which notifications are sent.
Port	Port number. The default: 25 .
Type of Authorization	The following authorization types are possible: <ul style="list-style-type: none"> • On; • Off; • PLAIN; • SCRAM-SHA1; • CRAM-MD5; • GSSAPI; • DIGEST-MD5; • LOGIN; • NTLM.
SSL/TLS	Enable or disable encryption.
E-mail of Sender	An email address of the notification sender.

Parameter	Description
Login (Optional)	A user sends notifications.
Password (Optional)	Password for email server access. If the server security policy allows unauthorized access, you can leave the field blank.

- Click **APPLY**.

10.4.1.1 Configuring SMTP Recipients

To add an SMTP recipient:

- Select **SYSTEM > NOTIFICATIONS**.
- In the *Recipient List* section, click **Add Recipient**. The pane to create a recipient opens (Figure 140).

Figure 140. NOTIFICATIONS page: adding a recipient

- In the **Recipient** field, type an email address of a recipient.
- In the *Notifications* column, click **Select**, then in the opened window *Add Notification Type* select the types of notifications to be sent and click **APPLY**.
- Click **CREATE**. The recipient displays in the *Recipient List* section (Figure 141).

Figure 141. Recipient of SMTP notifications

To verify recipient's settings, click **:** in the appropriate recipient line, and then click **Test**. A test notification will be sent to the specified email.

10.4.2 Setting up SNMP Server Parameters

RAIDIX 5.2 supports three versions of SNMP protocol – v.1, v.2, and v.3.

As opposed to SNMP v.1, SNMP v.2 includes enhanced performance, security, privacy, and connection between managers. SNMP v.2 also allows clients to receive more control data in a single request. However, the new security system based on sides from SNMP v.2 failed to be widely used.

To configure SNMP v.1 or SNMP v.2:

1. Select **SYSTEM > NOTIFICATIONS**.
2. On the left, select **SNMP** (Figure 142).

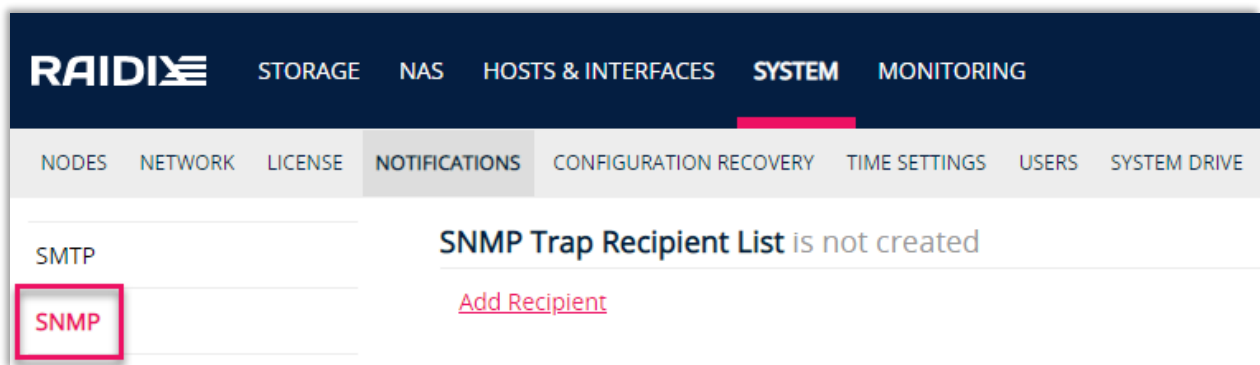


Figure 142. Configuring SNMP Server

3. Move to the *SNMP agent parameters* widget on the right (Figure 143).

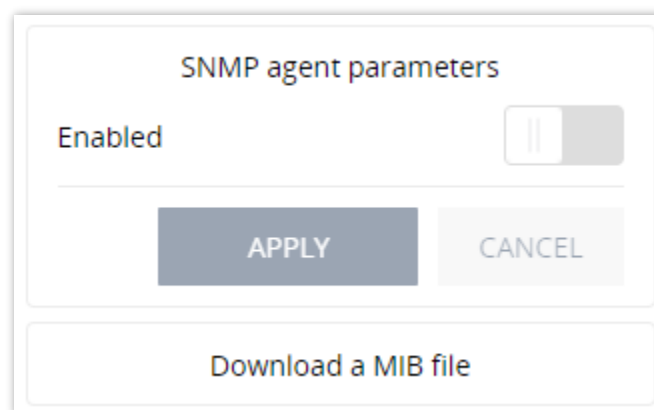


Figure 143. Configuring SNMP agent parameters

4. Configure the following SNMP agent parameters (Figure 144):

Parameter	Description
Enabled	Enable SNMP.
SNMP Version	SNMP protocol version: 1 , 2 , or 3
IP addresses	Optional. IP addresses by which a client can send requests to the server to receive information on the server state. To add another IP address, click Add IP Address and type the IP address in the appeared field.
Port	Port via which SNMP agent send requests to the server. The default port is 161 .
Community string	The string used for authentication on the server. The default value is public .

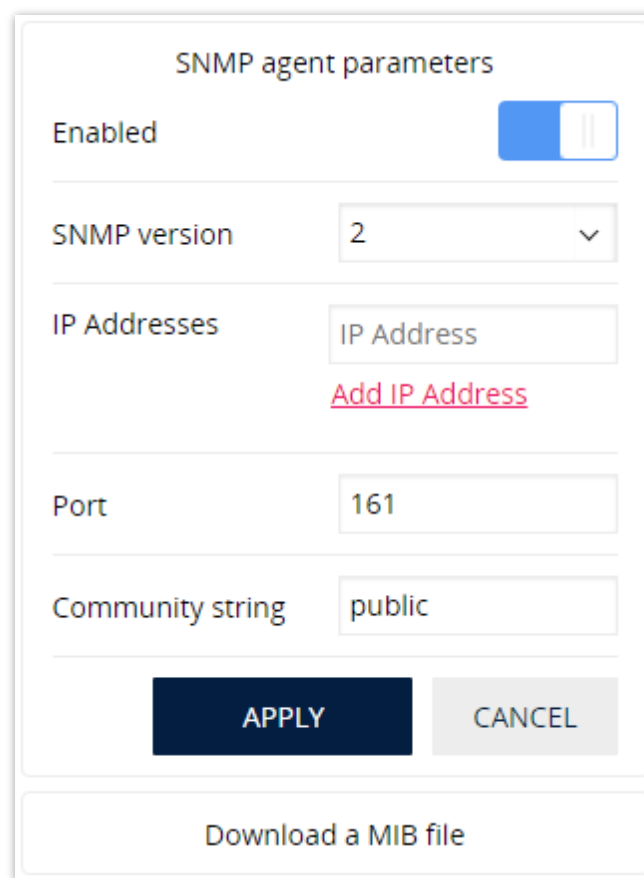


Figure 144. SNMP v2 agent parameters, same as SNMP v1

5. Click **APPLY**.

After SNMP agent is configured, you can download a MIB file – a structured file containing information on all objects of device. To do so, click **Download a MIB file**.

To configure SNMP v.3:

1. Select **SYSTEM > NOTIFICATIONS**.

2. On the left, select **SNMP** (Figure 142).
3. Move to the *SNMP agent parameters* pane on the right (Figure 143).
4. Configure the following SNMP agent parameters (Figure 145):

Parameter	Description
Enabled	Enable or disable SNMP.
SNMP Version	SNMP protocol version: 1 , 2 , or 3 .
IP addresses	<p>Optional.</p> <p>IP addresses by which a client can send requests to the server to receive information on the server state.</p> <p>To add another IP address, click Add IP Address and type the IP address in the appeared field.</p>
Port	The port number, through which SNMP agent sends requests to the server. The default: 161 .
Username	<p>Username for authentication on the server.</p> <p>Minimum length is 8 symbols.</p> <p>This parameter on SNMP server must match the username on the client side.</p>
Auth.protocol	<p>Authentication protocol: md5 or sha.</p> <p>The protocol must match the protocol selected on the client side.</p>
Auth.password	Password for authentication on the server. This parameter on SNMP server must match the password on the client side.
Auth.master key	Key used for authentication instead of the password.
Auth.localized key	Localized key used for authentication instead of the password.
Encryption Protocol	Encryption Protocol: des or aes .
Privacy password	Password used for encryption.
Privacy master key	Key used for encryption instead of the password.
Privacy localized key	Key used for encryption instead of the password.

SNMP agent parameters

Enabled

SNMP version

3

IP Addresses

IP Address

Add IP Address

Port

161

Username

Auth. protocol

md5

Auth. password

Auth. master key

Engine Id

Auth. localized key

Encryption protocol

des

Privacy password

Privacy master key

Privacy localized key

APPLY

CANCEL

Download a MIB file

Figure 145. SNMP v.3 agent parameters

- Click **APPLY**.

After SNMP agent is configured, you can download a MIB file – a structured file containing information on all objects of device. To do so, click **Download a MIB file**.

10.4.2.1 Configuring SNMP Recipients

After you have configured the SNMP settings, add the notification recipients.

To add an SNMP recipient:

- Select **SYSTEM > NOTIFICATIONS**.
- On the left, select **SNMP**.
- In the *Recipient list* section, click **Add Recipient**. The pane to create email notification opens (Figure 146).

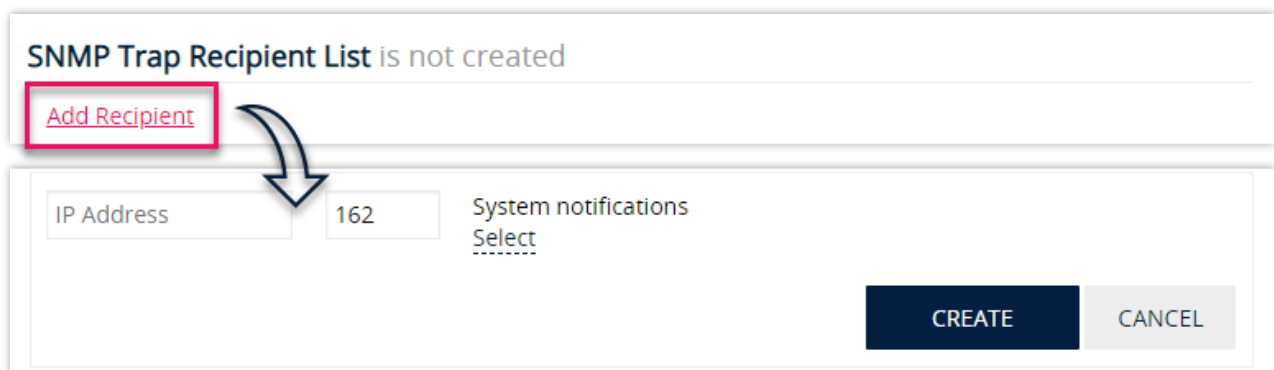
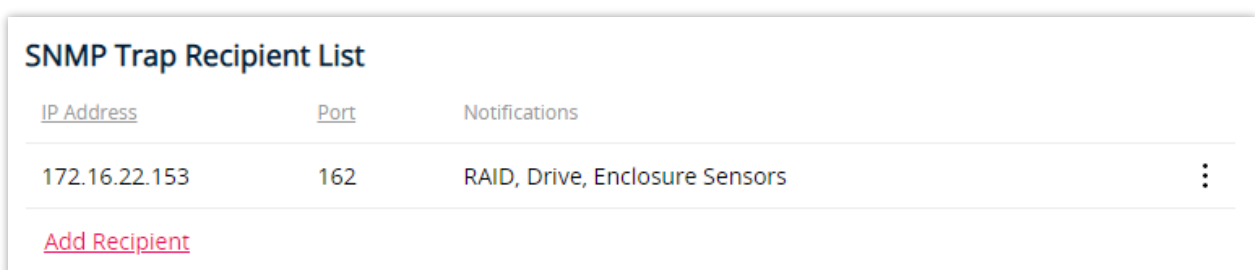


Figure 146. Adding a recipient for SNMP trap

- In the **IP address** field, type an IP address of the recipient.
- In the **Port** field, specify the port to receive SNMP notifications. The default port is **162**.
- In the *Notifications* column, click **Select**, then in the opened window *Add Notification Type* select the types of notifications to be sent and click **APPLY**.
- Click **CREATE**. The new SNMP recipient displays in the *SNMP Trap Recipient List* section (Figure 147).



IP Address	Port	Notifications
172.16.22.153	162	RAID, Drive, Enclosure Sensors

Figure 147. Recipient added

To verify the settings was correct, click **:** in the appropriate recipient line, and then click **Test**. A test SNMP Trap will be sent to the recipient.

10.5 Configuration Recovery

! While working with the system, it is important to have an ability to restore system state and settings from the system metadata. System metadata is stored on drives and contains configuration parameters of the system.

The configuration file includes information on current system configuration. Current system configuration can be recovered even if a controller or boot drive failure occurs: a user has an ability to connect drives or enclosure with data to a workable system or change the boot drive and upload configuration file on it. System configuration applies immediately.

Use the elements of the **SYSTEM | CONFIGURATION RECOVERY** page to manage the system configuration files (Figure 148).

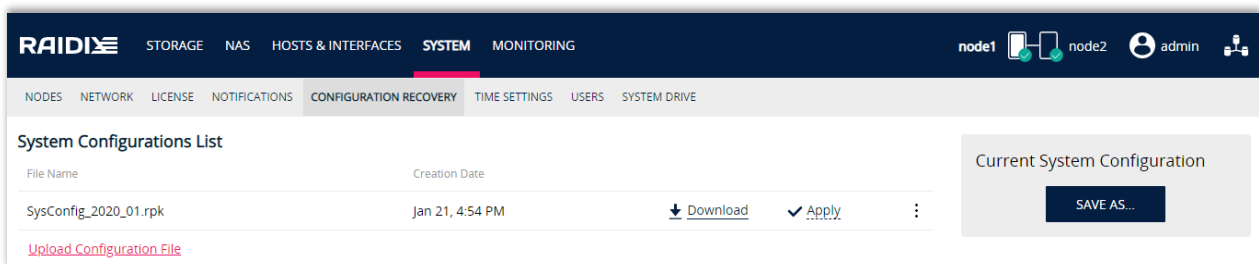


Figure 148. CONFIGURATION RECOVERY page

- !
- Don't use configuration recovery after changing hardware system configuration or creating a RAID of a different level.
 - Don't initiate system configuration recovery until the RAID initialization is completed.

10.5.1 Saving Current System Configuration

- !
- To be able to restore current system settings after modifying them, make a backup copy of system configuration file regularly.


The system creates configuration file, enables you to **Download** or **Apply** the file

To save the current system configuration:

1. Select **SYSTEM > CONFIGURATION RECOVERY**.
2. In the *Current System Configuration* pane, click **SAVE AS**.
3. In the opened window, specify a file name and click **SAVE**.

10.5.2 Downloading Configuration File

To download a configuration file:

1. Select **SYSTEM > CONFIGURATION RECOVERY**.
2. In a line of the configuration, click  **Download**.
The configuration file will be saved on your computer.

To download a file of the current configuration, first save it on the system by clicking **SAVE AS**.

10.5.3 Uploading Configuration File

To upload a configuration file:

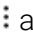
1. Select **SYSTEM > CONFIGURATION RECOVERY**.
2. In the *System Configuration List* section, click **Upload Configuration File**.
3. Select a configuration file.

10.5.4 Applying System Configuration

To apply a system configuration:

1. Select **SYSTEM > CONFIGURATION RECOVERY**.
2. In a line of the configuration, click **Apply**.
3. To confirm, click **APPLY**.

10.5.5 Deleting Configuration File

1. Select **SYSTEM > CONFIGURATION RECOVERY**.
2. In a line of the configuration, click  and select **Delete**.
3. To confirm, click **DELETE**.

10.5.6 Advanced Configuration of the System

10.5.6.1 Cloning System Configuration

In RAIDIX 5.2, you can clone the configuration of the system in the same system or another system with the similar hardware configuration.

You can perform system configuration cloning only through CLI by using the `rdclone` utility. For a list of valid commands, in the CLI run the following:

```
$ rdclone -h
```

Utility reads the current configuration from:

- actual `raidixcfg.json` file (default);
- any proper file of the configuration that was specified in the `rdclone` parameter;
- STDIN through pipe, if particular key was specified.

The utility converts the current configuration of the system into the sequence of the `rdcli` commands, which generates a list in the shell script format. The shell script list will be put into the STDOUT (by default) or to the specified in the key file.

The outcome of the `rdclone` utility is used as a `sh` file, that contains a set of commands to recreate configuration on similar hardware.

The `rdclone` utility features are:

- The `rdclone` utility generates shell script for a user or group, before NAS creation commands. The shell script searches for the user or group in the Linux system to remove them in the case of detection.
- The `rdclone` utility creates the XFS file system on a LUN, before NAS share creation commands.
- If there are drives' IDs in the RAID or SparePool creation commands, which are not in the drive section of the configuration, the `rdclone` will show a warning in STDERR. The STDERR warning will include IDs of the drives and will generate commented line of the RAID/SparePool creation, where UID of drives will be marked as "??".

10.5.6.2 Import of a RAID



Only for Generic RAID.

In RAIDIX 5.2, you can import others system RAID, which you have on the drives but do not have in the working system configuration. For example, you can use this functionality, when you want to combine several RAID, which you worked with into one.

You can connect data drives of another system to the working system through the web interface or CLI.



When you import RAID, you also import volumes that belong to this RAID. This import is possible only when RAID is correctly restored (Online). If RAID is restored incorrectly, the system will fail to read the RVM metadata, which contains information about the volumes. For the correct import, all RAID drives must be presented in the system.



RAIDs and volumes with the same names cannot exist in the same system, therefore, when you add data drives of another system into your working system, before the import, you must rename RAIDs and volumes if their names coincide with the names of the RAIDs and volumes that are already in the system.

You cannot import a RAID on which no LUNs were created.

Due to the use of statistical metadata, in the RAIDIX 5.2 systems, you are able to import only RAIDIX 4.5 RAIDs and above.

i RAID import doesn't recover mistaken deleted RAIDs and LUNs.

To import a RAID through the web-interface:

1. Select **STORAGE > RAID IMPORT**.
2. In the *RAID Import* section, click **Import** in the line of the RAID to be imported (Figure 149).
3. If names of RAIDs for import coincide with the names of the RAID that are already in the system, rename those RAIDs.

To rename a RAID:

- 3.1. Click **Rename drives metadata**.
- 3.2. In the appeared dialog window, select drives belonging to the RAID for import, specify new RAID name, and click **Select**.

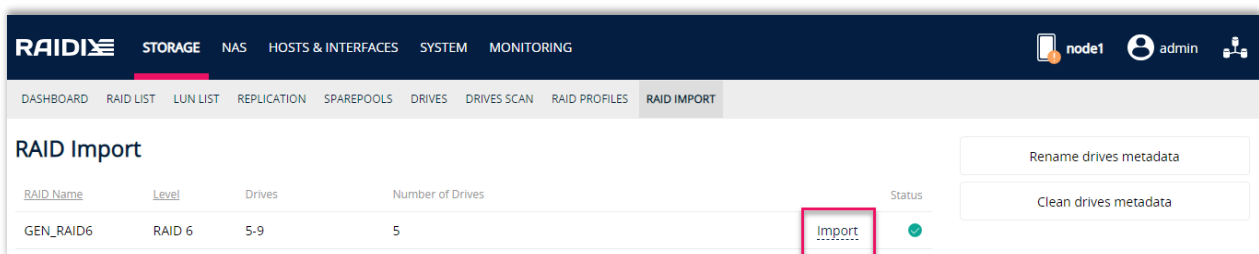




Figure 149. RAID Import

4. In the opened *RAID Import Wizard* window, select the cache size for the RAID to be imported and click **NEXT**.
5. In the appeared dialog window, check for name conflicts with already existing LUNs in the system.

If there are any conflicts, rename all such LUNs by using  icon. Click  to confirm the renaming.

6. If there are no conflicts with LUN names, click **NEXT**.
7. The summary table containing RAID parameters and the list of LUNs on the RAID opens. Check the information, then click **IMPORT**.

To import a RAID through the CLI:

1. Specify the list of drives that you want to use for the RAID search, by running the following command:

```
$ rdcli metadata raid scan [{-d|--drives} (<drive-uids>)]
```

If the list of drives is not specified, the search will be based on all available drives.

The `rdcli` utility will find and display information about RAIDs found.

2. If names of the RAID and/or volumes for import coincide with the names of the RAID and/or volumes that are already in the system, rename the RAID and volumes. To rename RAID and/or volumes, do the following:

```
$ rdcli metadata raid rename [{-on|--old_name}
<old_raid_name>] [{-d|--drives} (<drive-uids>)] {-n|--name}
<name>
```

The list of drives is optional and needed in case of ambiguity (when several RAIDs have the same name) to identify which RAID to rename.

3. Restore the RAID by its name and get the list of volumes, by running the following command:

```
$ rdcli metadata volume scan {-r|--raid} <raid_name>
```

This command will show LUNs available on RAID.

4. If necessary, rename volumes through running the following command:

```
$ rdcli metadata volume rename {-r|--raid} <raid_name>
{-n|--name} (<vols-rename-list>)
```

In this command <vols-rename-list> has the following structure:

```
(<old-name1>-<new-name1>,<old-name2>-<new-name2>)
```

5. Import RAID into the current configuration, by running the following command:

```
$ rdcli metadata raid import {-r|--raid} <raid_name>
```

If needed, specify a new cache size value for the RAID before import by using the `--cache_size` option.

10.6 Time Settings

Change system time settings on the page **SYSTEM | TIME SETTINGS** (Figure 150).

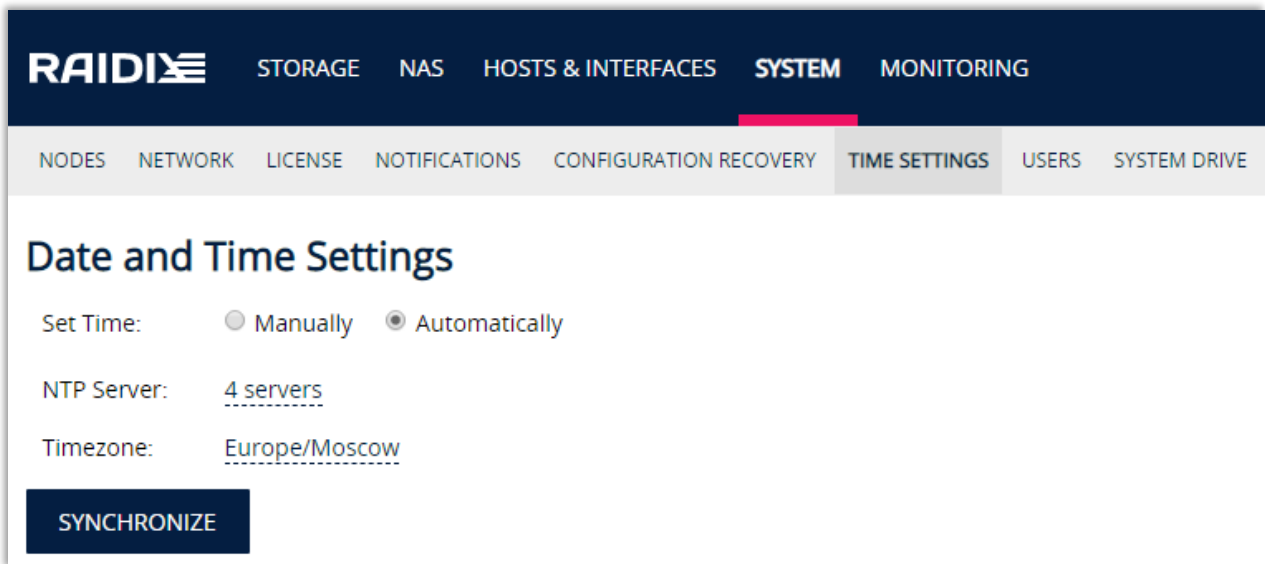


Figure 150. Date and Time Settings: Automatically

Set the system time manually or synchronize the time with NTP servers.

Current system date and time are displayed in the widget in the upper-right corner.

To set time manually, do the following:

1. Select **SYSTEM > TIME SETTINGS**.
2. In the *Set Time* section, select **Manually** (Figure 151).

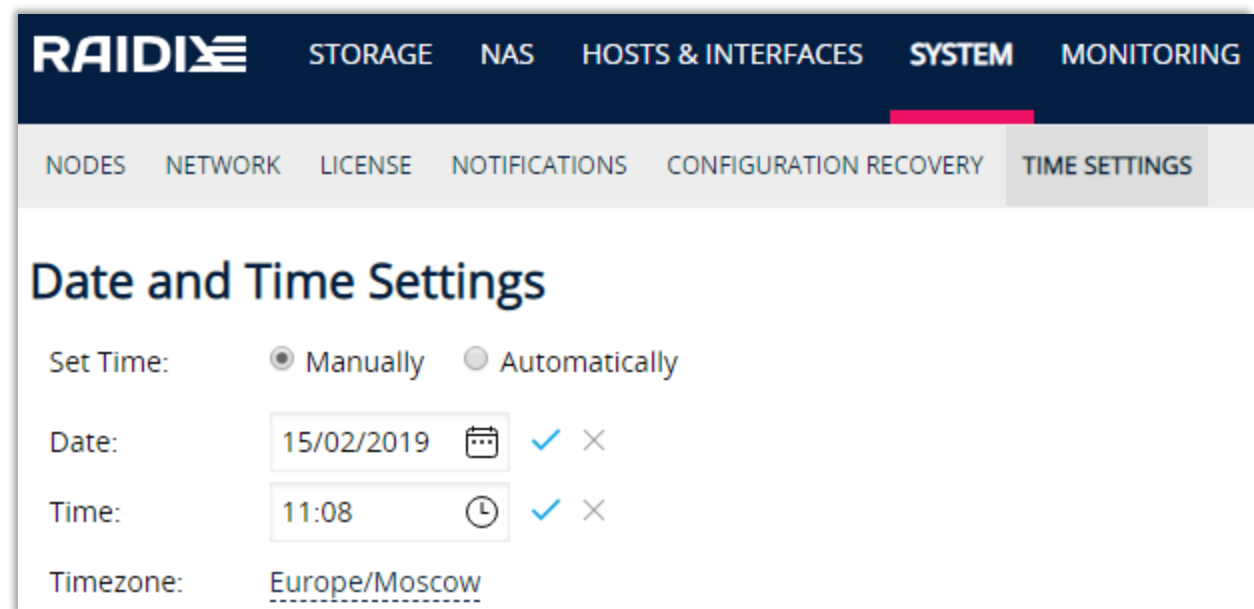




Figure 151. Date and Time Setting: Manually

3. To change date or time, click  in the appropriate line, specify the value and confirm the input by clicking .
- To discard changes, click .

4. To change timezone:
 - 4.1. Click the current value and select the correct time zone from the list.
 - 4.2. Click **SELECT**.

Time settings are immediately applied and displayed in the pane *Current System Date and Time*.

To synchronize time using NTP servers:

1. Select **SYSTEM > TIME SETTINGS**.
2. In the *Set Time* section, select **Automatically**.
3. Click the number of servers to look through NTP server list.
4. To change timezone:
 - 4.1. Click the current value and select the correct time zone from the list.
 - 4.2. Click **SELECT**.
5. Click **SYNCHRONIZE**.

10.7 Users

To divide an access to the system maintenance, RAIDIX 5.2 has 2 user roles with different possibilities:

- Administrators – all system and user operations are available, including:
 - viewing system state;
 - managing system parameters;
 - system upgrade;
 - creating system users and editing user parameters.
- Operators – operations with the system, including:
 - viewing both state and alerts of the system;
 - collecting logs of the system (see [Downloading system logs](#));
 - running S.M.A.R.T. for drives;
 - changing both personal password and system language.

The system user account has four editable parameters available form GUI:

- **Login** – a user name;
- **Password** – a password to log in the system;
- **Role** – a user role;
- **Language** – used GUI language.

In addition to the listed parameters, you can edit password policy parameters of the system user account from CLI (see [Password policy](#)).

10.7.1 Managing System Users

To create a system user:

1. Select **SYSTEM > USERS**.
2. In the *User List* section, click **Create User**.
3. Fill in the fields **Login**, **Password**, **Role**, and **Language** (Figure 152).

The screenshot shows the RAIDIX web interface with the 'SYSTEM' tab selected. Under 'SYSTEM', the 'USERS' sub-tab is active. The 'User list' section displays a table with columns: Login, Password, Role, and Language. The table contains one entry: 'admin' with a masked password, role 'Administrators', and language 'English'. A red box highlights the 'Create User' button below the table. An arrow points from this button to a form below. The form has input fields for 'Login' and 'Password', and dropdown menus for 'Operators' and 'English'. At the bottom right of the form are 'CREATE' and 'CANCEL' buttons.

Figure 152. Widget to create a system user

4. Click **CREATE**.

To change parameters of an available system user:

1. Select **SYSTEM > USERS**.
2. In the *User List* section, in the string with a user that you want to modify, click **:** and select **Edit** (Figure 153).

The screenshot shows the RAIDIX web interface with the 'SYSTEM' tab selected. Under 'SYSTEM', the 'USERS' sub-tab is active. The 'User list' section displays a table with columns: Login, Password, Role, and Language. The table contains two entries: 'admin' and 'Operator1'. For the 'Operator1' entry, a dropdown menu is open, showing 'Edit' and 'Delete' options. A red box highlights the 'Create User' button below the table.

Figure 153. Changing parameters of an available system user

3. Modify required parameters.
4. Click **APPLY**.

To delete an available system user:

1. Select **SYSTEM > USERS**.
2. In the *User List* section, in the string with a user you want to delete, click **:** and select **Delete** (Figure 153).

3. In the confirmation window, click **DELETE**.

10.7.2 Password Policy

You can set up the password policy for specific system user (individual policy) and all system users (general policy).

General password policy consists of the following rules:

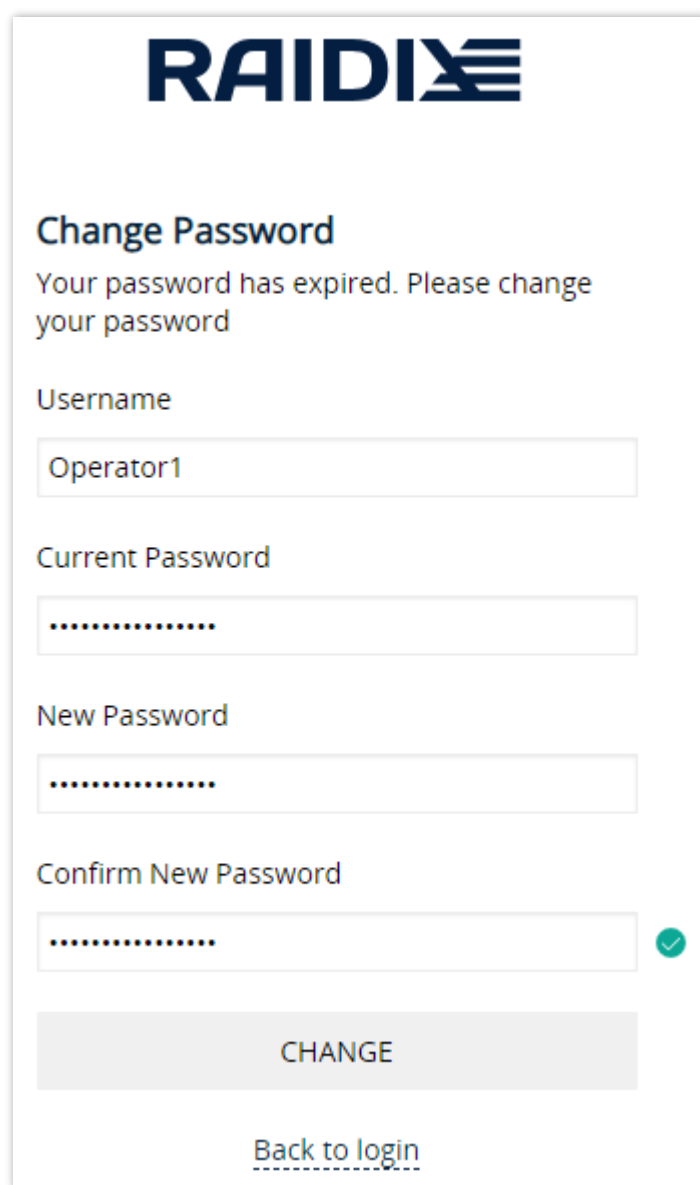
- password difficulty:
 - minimum password length;
 - minimum number of symbol classes;
 - maximum number of consecutive symbols;
 - maximum number of consecutive symbols of one class;
 - maximum number of consecutive symbols in a row;
 - maximum/minimum number of lowercase symbols;
 - maximum/minimum number of uppercase symbols;
 - maximum/minimum number of digital symbols;
 - maximum/minimum number of special characters;
 - dictionary check;
- password history:
 - length of password history;
- failed authentication:
 - amount of unsuccessful authentication attempts before user lockout;
 - time to unlock of a system user.

You can view and modify the general policy via CLI (see `system user`). By attempt of creating an improper password, the system shows the alert.

Individual password policy consists of the following rules:

- system user account lifetime:
 - lifetime of an account;
 - time to lock inactive user after password lifetime expired;
- password lifetime:
 - password lifetime;
 - time to password renew.

While a password is expired, the system shows the *Change Password* window by attempting to log in (Figure 154). To change the password, fill in the fields and click **Change**.



The image shows a web interface for changing a password. At the top is the RAIDIX logo. Below it is the title 'Change Password' and a message: 'Your password has expired. Please change your password'. There are four input fields: 'Username' with the value 'Operator1', 'Current Password' with masked characters, 'New Password' with masked characters, and 'Confirm New Password' with masked characters and a green checkmark icon to its right. Below the fields is a grey 'CHANGE' button. At the bottom is a link 'Back to login'.

Figure 154. Change Password window



A user with the Operator role can change the password through GUI only after:

- the first authentication;
- expiration of the password.

10.8 System Drive

The **SYSTEM DRIVE** page provides info about system drive partitions. To see the info, select **SYSTEM | SYSTEM DRIVE** (Figure 155).

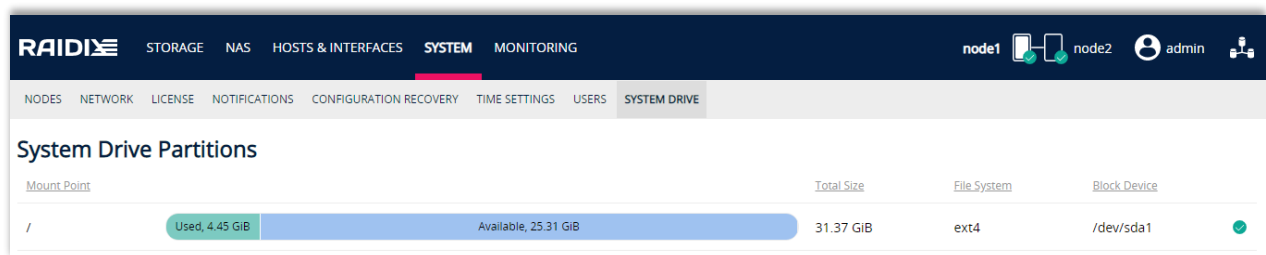



Figure 155. System Drive page

The *System Drive Partitions* section has the following fields:

Parameter	Description
Mount Point	Mount point of a partition.
Total Size	Total available size of a partition.
File System	Name of partition file system.
Block Device	Full path to the partition block device.
Status icon	<ul style="list-style-type: none"> ✓ – The partition has enough free space. ⚠ – Partition free space is between 90 – 95 %. ✗ – Partition free space is less than 95%.

 You can't log in to GUI when system drives are full. The system will show an alert before authorization. To log in GUI, make free space on system drives.

Alerts

The changing of system drive partition status shows on the **ALERTS** page. You can [set up email and snmp notifications](#) on the **SYSTEM | NOTIFICATIONS** page.

Besides that, health notifications are also available for system drives, see [S.M.A.R.T. Diagnostics](#) for details.

11. SYSTEM MONITORING

To control the system's state and to operate the system via [command console](#), open the **MONITORING** page. System monitoring section assists a user to gather information about the [controller and enclosure sensors](#); about the [storage system's performance](#); about presence/absence of the [system faults](#) and provides an ability to [download the system logs](#).

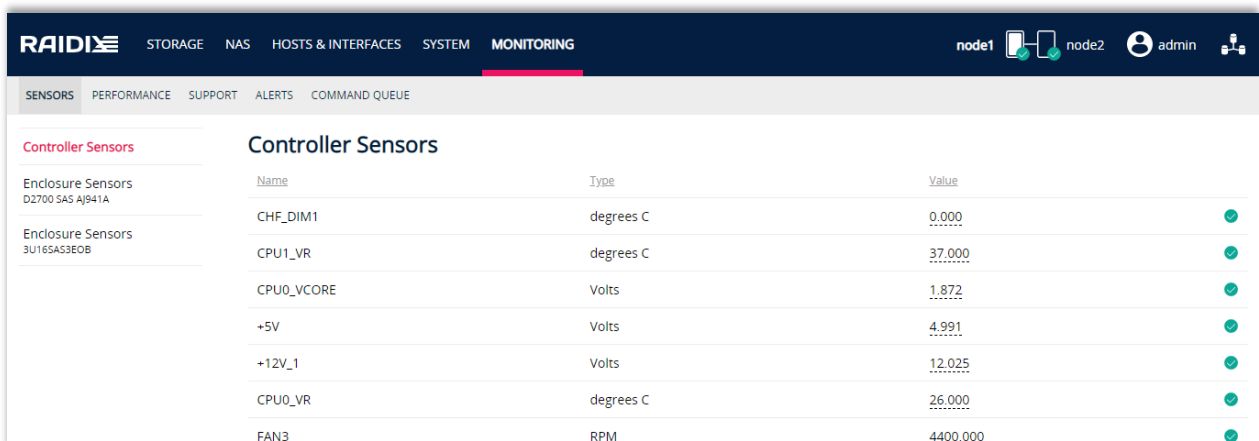
11.1 Sensors Info

To work with sensors, open the **MONITORING | SENSORS** page.

The **SENSORS** page provides information on the controller sensors and sensors of the connected drive enclosures.

11.1.1 Information on Controller Sensors




To access to the controller sensors, open the **MONITORING | SENSORS** page, and click **Controller Sensors**. Name, current value, and state indicator specifying if the value of the measured parameter is within the possible value range are listed *for every controller sensor* in the opened table (Figure 156).




RAIDIX			
STORAGE NAS HOSTS & INTERFACES SYSTEM MONITORING			
node1 node2 admin			
SENSORS PERFORMANCE SUPPORT ALERTS COMMAND QUEUE			
Controller Sensors			
Controller Sensors	Name	Type	Value
Enclosure Sensors D2700 SAS AJ941A	CHF_DIM1	degrees C	0.000
Enclosure Sensors 3U16SAS3E0B	CPU1_VR	degrees C	37.000
	CPU0_VCORE	Volts	1.872
	+5V	Volts	4.991
	+12V_1	Volts	12.025
	CPU0_VR	degrees C	26.000
	FAN3	RPM	4400.000

Figure 156. Node sensors info

Parameter	Description
Name	Sensor name.
Type	Sensor type: <ul style="list-style-type: none"> • RPM – Fan rotations per minute; • Volts; • degrees C – temperature in °C.
Value	Current sensor value.

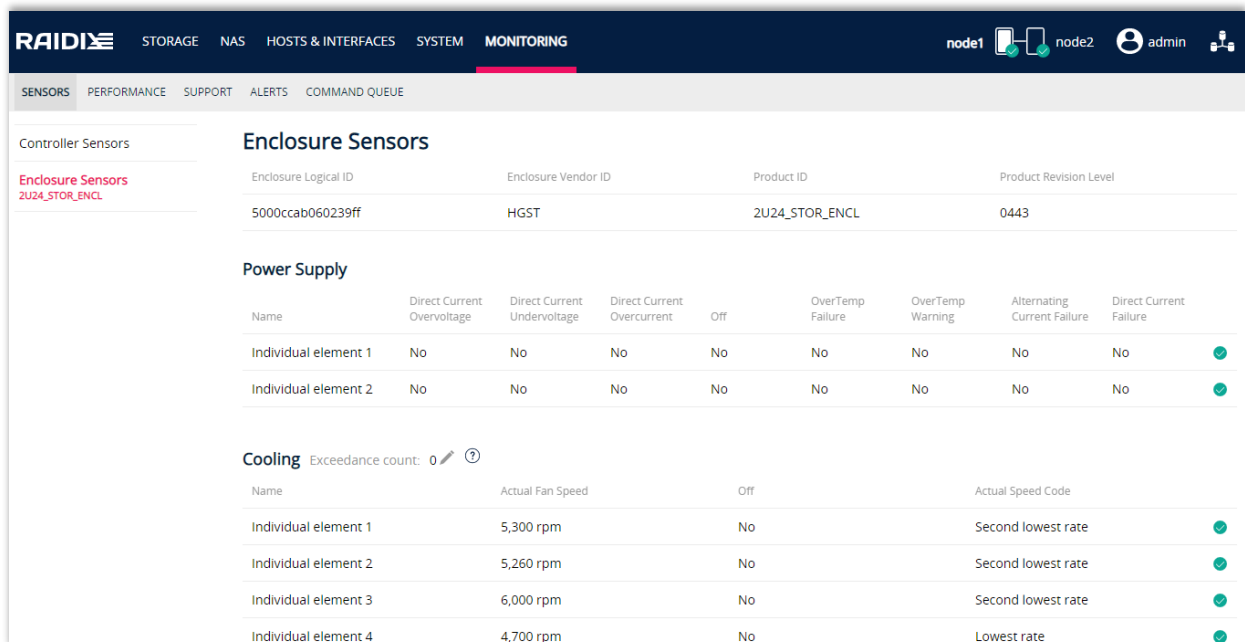
Parameter	Description
Status	Current sensor status: <ul style="list-style-type: none">  <i>OK</i> – works correctly;  <i>Warning</i> – sensor is not defined;  <i>Error</i> – the critical error occurred, or the measured value exceeds the allowed limit.

If sensor results are outside the valid range, the device scanned by the sensor is partially or completely defective.

 If readings from one or several sensors are beyond the range of allowed values, system can send email notifications (if the [appropriate settings](#) were performed).



11.1.2 Information on Enclosure Sensors



To access to the enclosure sensors, open the **MONITORING | SENSORS** page, and click **Enclosure Sensors**. Information about the power supply, cooling mechanism, and sensors of the connected enclosure is displayed (Figure 157).



The screenshot shows the RAIDIX web interface. The top navigation bar includes STORAGE, NAS, HOSTS & INTERFACES, SYSTEM, and MONITORING. The MONITORING tab is active, and the SENSORS sub-tab is selected. The left sidebar shows 'Controller Sensors' and 'Enclosure Sensors' (highlighted in red). The main content area displays 'Enclosure Sensors' for the enclosure '2U24_STOR_ENCL'. It includes a table for 'Enclosure Logical ID', 'Enclosure Vendor ID', 'Product ID', and 'Product Revision Level'. Below this, there are sections for 'Power Supply' and 'Cooling' with their respective status tables.

Enclosure Logical ID	Enclosure Vendor ID	Product ID	Product Revision Level
5000ccab060239ff	HGST	2U24_STOR_ENCL	0443

Name	Direct Current Overvoltage	Direct Current Undervoltage	Direct Current Overcurrent	Off	OverTemp Failure	OverTemp Warning	Alternating Current Failure	Direct Current Failure	Status
Individual element 1	No	No	No	No	No	No	No	No	
Individual element 2	No	No	No	No	No	No	No	No	

Cooling Exceedance count: 0  














Name	Actual Fan Speed	Off	Actual Speed Code	Status
Individual element 1	5,300 rpm	No	Second lowest rate	
Individual element 2	5,260 rpm	No	Second lowest rate	
Individual element 3	6,000 rpm	No	Second lowest rate	
Individual element 4	4,700 rpm	No	Lowest rate	

Figure 157. Enclosure sensors info

Parameter	Description
Enclosure Logical ID	Unique enclosure ID.
Enclosure Vendor ID	Enclosure Manufacturer ID.

Parameter	Description
Product ID	Enclosure ID in INQUIRY format.
Product Revision Level	Enclosure firmware version.
Power Supply	
Name	Name of a power supply sensor.
DC OverVoltage	Parameter indicates whether an overvoltage condition has been detected at the power supply output. Possible variants: <ul style="list-style-type: none"> • <i>Yes</i>; • <i>No</i>.
DC UnderVoltage	Parameter indicates whether an undervoltage condition has been detected at the power supply output. Possible variants: <ul style="list-style-type: none"> • <i>Yes</i>; • <i>No</i>.
DC OverCurrent	Parameter indicates whether an overcurrent condition has been detected at the power supply output. Possible variants: <ul style="list-style-type: none"> • <i>Yes</i>; • <i>No</i>.
Off	Parameter returns information about the elements state. <ul style="list-style-type: none"> • <i>Yes</i> – element (power supply or cooler) is disabled; • <i>No</i> – element (power supply or cooler) is disabled but sensor is working.
OverTemp Failure	Parameter indicates the power supply has detected a temperature above the safe operating temperature range. The power supply may shut down. Possible variants: <ul style="list-style-type: none"> • <i>Yes</i>; • <i>No</i>.
OverTemp Warning	Parameter indicates the power supply has detected a temperature within the safe operating temperature range, but above the normal operating temperature range. Precedes OverTemp Failure state.
AC Failure	Bit set to one indicates that the power supply is not receiving the specified A.C. power.
DC Failure	Bit set to one indicates that the power supply is unable to supply the specified D.C. power.
Status	<div>  <i>OK</i> – works correctly; </div> <div>  <i>Warning</i> – sensor is not defined; </div> <div>  <i>Error</i> – the critical error occurred, or the measured value exceeds the allowed limit. </div>
Cooling	
Name	Name of a cooling sensor.

Parameter	Description
Exceedance count	Editable parameter. Number of exceeds (situations, in which sensor readings exceed the allowed range), after which the system sends an email notification. The default: 0 (no notifications will be sent).
Actual Fan Speed	Actual fan speed (in rotations per minute).
Off	Parameter returns information about the elements state. <ul style="list-style-type: none"> <i>Yes</i> – element (a power supply or a cooler) is disabled; <i>No</i> - element (a power supply or a cooler) is disabled but sensor is working.
Actual Speed Code	Current cooling mode.
Status	<div>  <i>OK</i> – works correctly; </div> <div>  <i>Warning</i> – sensor is not defined; </div> <div>  <i>Error</i> – the critical error occurred, or the measured value exceeds the allowed limit. </div>
Temperature	
Name	Name of a temperature sensor.
Temperature	Current reading, in °C.
OverTemp Failure	The temperature is above the safe operating temperature range or higher than the value indicated by the critical upper threshold.
OverTemp Warning	The temperature is close to the critical upper threshold.
UnderTemp Failure	The temperature is below the safe operating temperature range or lower than the value indicated by the critical lower threshold.
UnderTemp Warning	The temperature is close to the critical lower threshold.
Status	<div>  <i>OK</i> – works correctly; </div> <div>  <i>Warning</i> – sensor is not defined; </div> <div>  <i>Error</i> – the critical error occurred, or the measured value exceeds the allowed limit. </div>

11.2 Performance Info

Performance monitoring module helps to map up configuration, estimate, and increase the performance both inside of the storage system and during the data transfer process.

To access the system performance data, select **MONITORING > PERFORMANCE** (Figure 158).

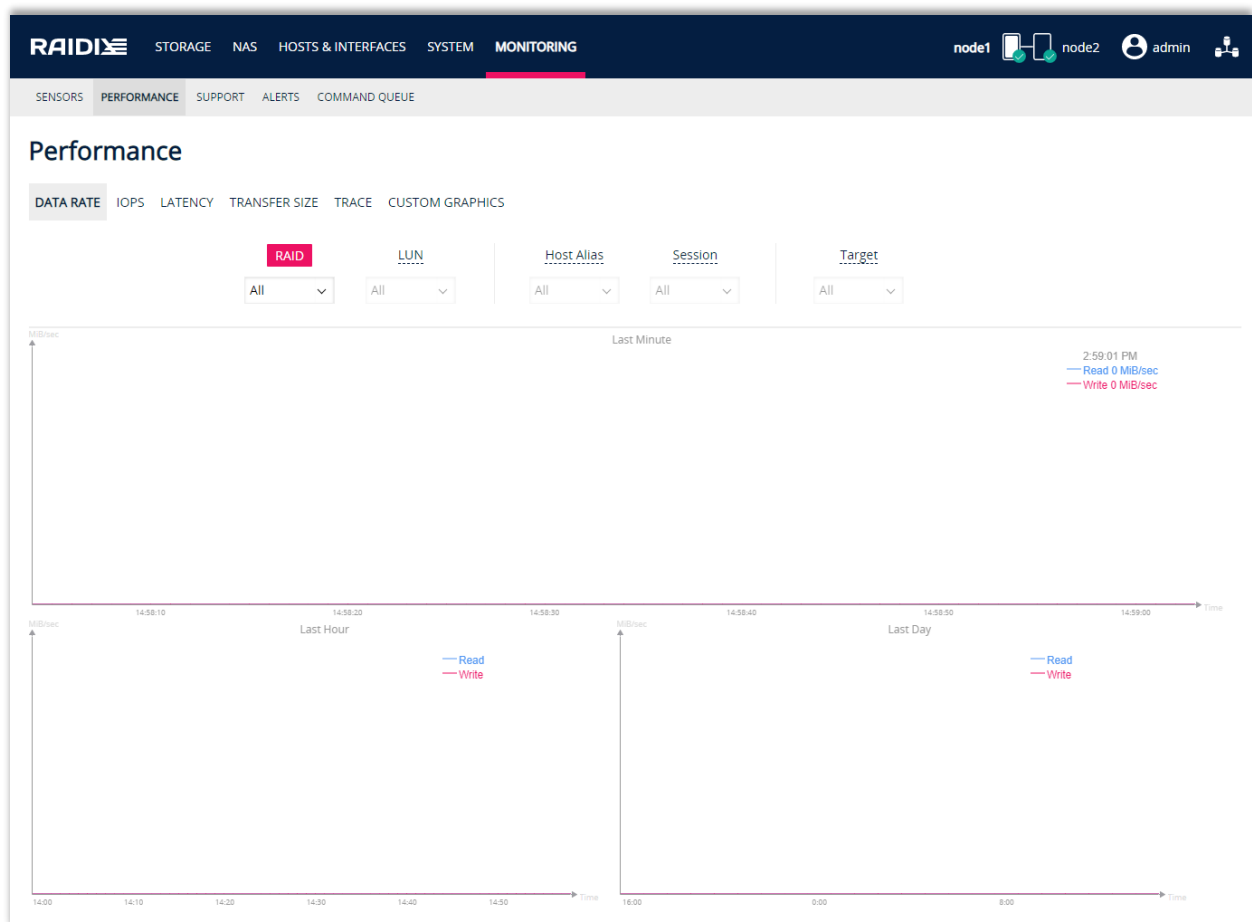


Figure 158. Data Rate section fragment

The module has the following functional sections:

- DATE RATE;
- IOPS;
- LATENCY;
- TRANSFER SIZE;
- TRACE;
- CUSTOM GRAPHICS.



After you configured replication, write processes related to data synchronization are seen only on the secondary LUN. Read processes on the primary LUN is not displayed.

11.2.1 Data Rate

DATA RATE page provides detailed information on the system performance and allows to estimate data exchange rate.

You can set up the chart display parameters using the following filters on the top of the **DATA RATE** page:

- **RAID** – filter provides statistics for all or specified RAID;

- **LUN** – filter provides statistics for all or specified LUNs;
- **Host alias** – filter provides statistics for all or specified created Host Aliases;
- **Session** – filter provides statistics for all or specified established host sessions;
- **Target** – filter provides statistics for all or specified existing targets.

To set up the displayed parameters, click the appropriate filter name and select the information to display in the drop-down list appeared.

Data Rate charts (Figure 159) are useful tools for planning and monitoring the backup procedure, allowing getting detailed real-time information on read and write operations speed at different time intervals.

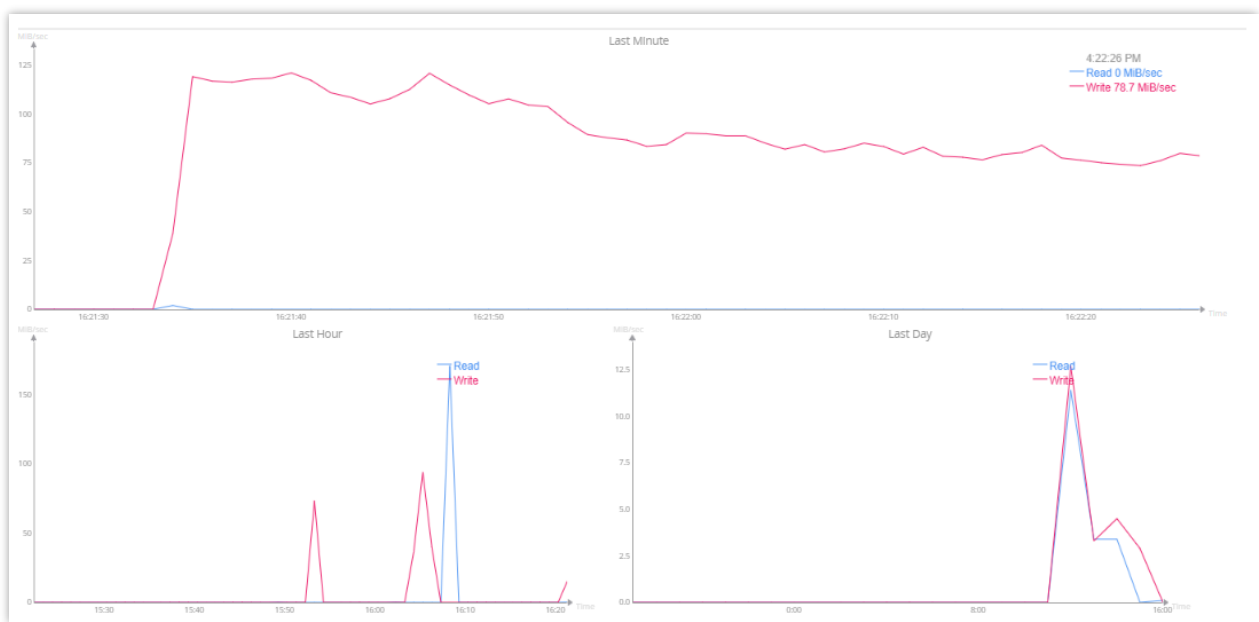


Figure 159. Data Rate charts

The vertical axis of the charts shows the operations speed (in MiB/sec). The horizontal axis displays the data transfer time interval.

Last Minute chart provides the real-time information on the data transfer rate for the last minute every second. To view the detailed information to a second, point your mouse to the graph.

Last Hour chart displays the last hour data rate; average value per minute is displayed. To view the detailed information to a minute, point your mouse to the graph (Figure 160).

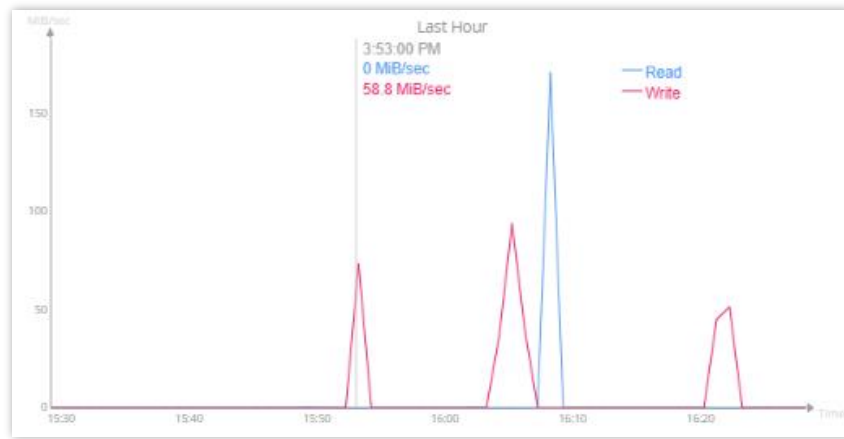


Figure 160. Last Hour chart data detalization

Last Day chart displays the last day data rate. To view the detailed information to an hour, point your mouse to the graph.

11.2.2 IOPS Charts

IOPS charts display detailed real-time information on read and write operations speed at different time intervals.

The vertical axis of the charts shows IOPS (Input/Output Operations per Second). The horizontal axis displays the data transfer time interval. Charts for reading and writing are displayed in different colors.

Last Minute chart provides the real-time information on IOPS value for the last minute every second. To view the detailed information to a second, point your mouse to the graph.

Last Hour chart displays the last hour IOPS value; average value per minute is displayed. To view the detailed information to a minute, point your mouse to the graph.

Last Day chart displays the last day IOPS value. To view the detailed information to an hour, point your mouse to the graph.

11.2.3 Latency Charts

Latency charts allow evaluating the response to commands processing at different time intervals (Figure 161). The charts vertical axis displays the number of executed commands. The horizontal axis shows where the time spent on command execution is located within a time range (in milliseconds, microseconds).

Latency statistics are displayed on the following charts:

- Last Minute Summary;
- This Minute Summary;
- Last Hour Summary;
- This Hour Summary;
- Last Day Summary;

- Today Summary.

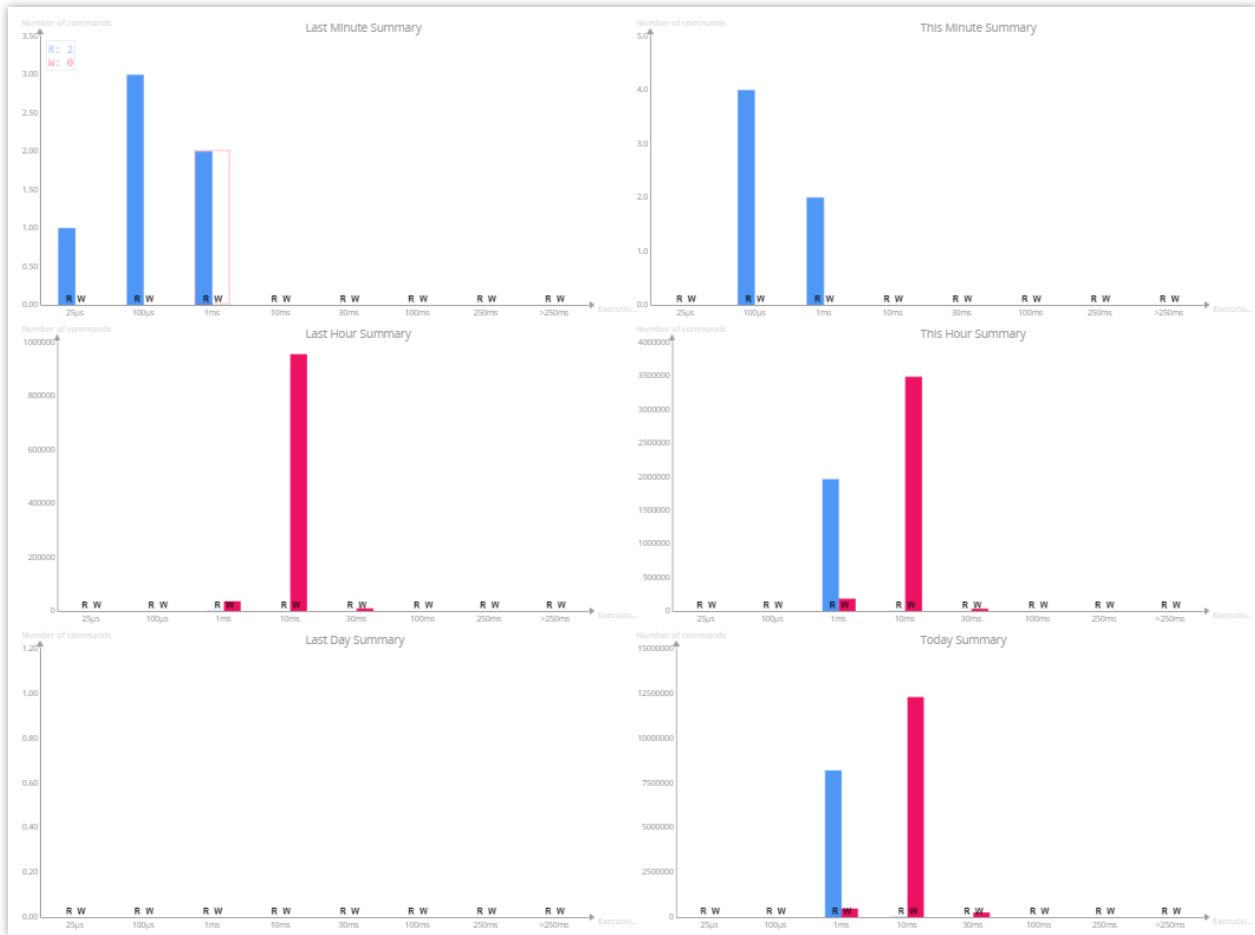


Figure 161. Latency charts (fragment)

11.2.4 Transfer Size Charts

Transfer size charts (Figure 162) show the amount of transferred data (in blocks) at different time intervals.

The following analytics time periods are available in RAIDIX software:

- Last Minute Summary;
- This Minute Summary;
- Last Hour Summary;
- This Hour Summary;
- Last Day Summary;
- Today Summary.

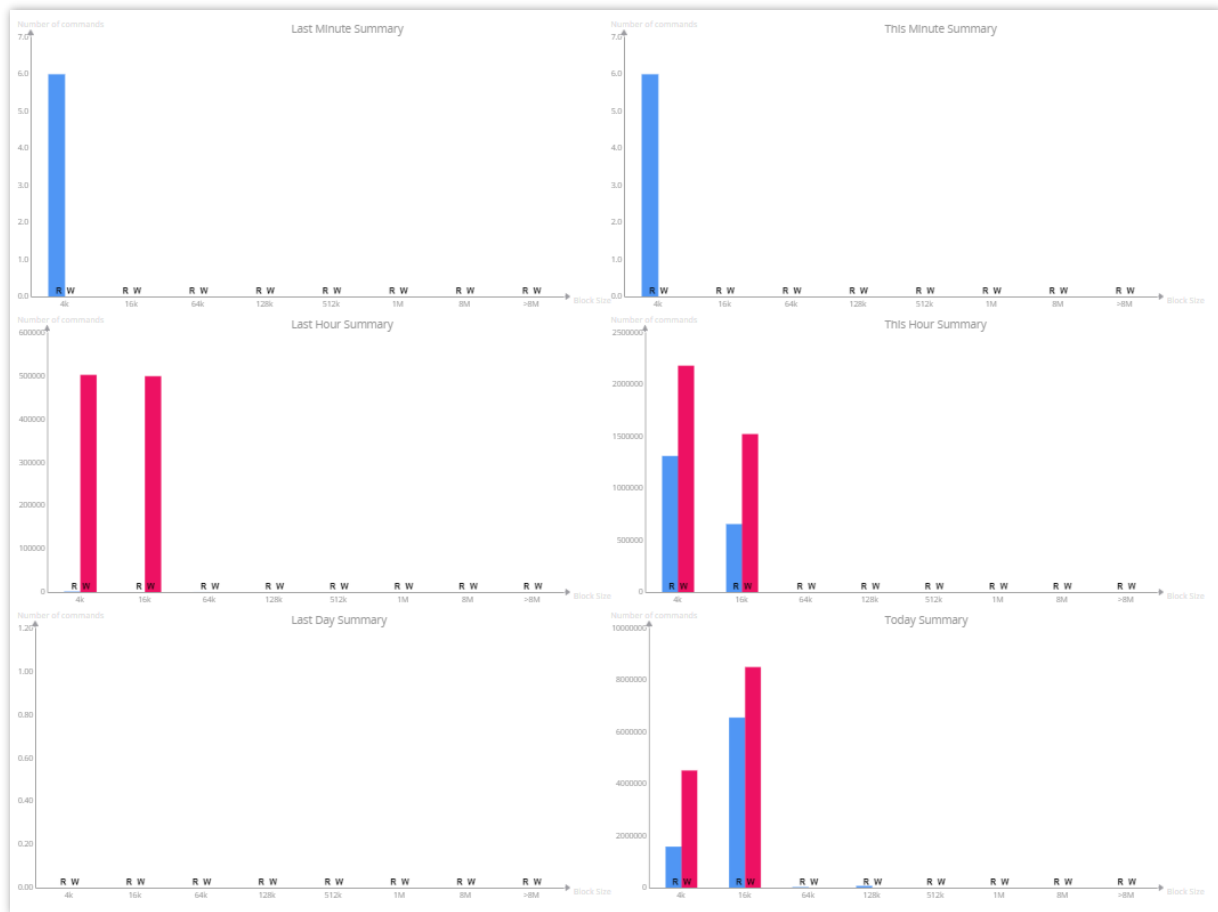


Figure 162. Transfer Size charts (fragment)

11.2.5 Trace Statistics

Trace section is intended for engineers who perform the primary selection of storage configuration and commission operations (Figure 163).



It is not recommended to trace while working on high data rates, in particular, if SSD cache is enabled since it may cause defects in the structure of trace files.

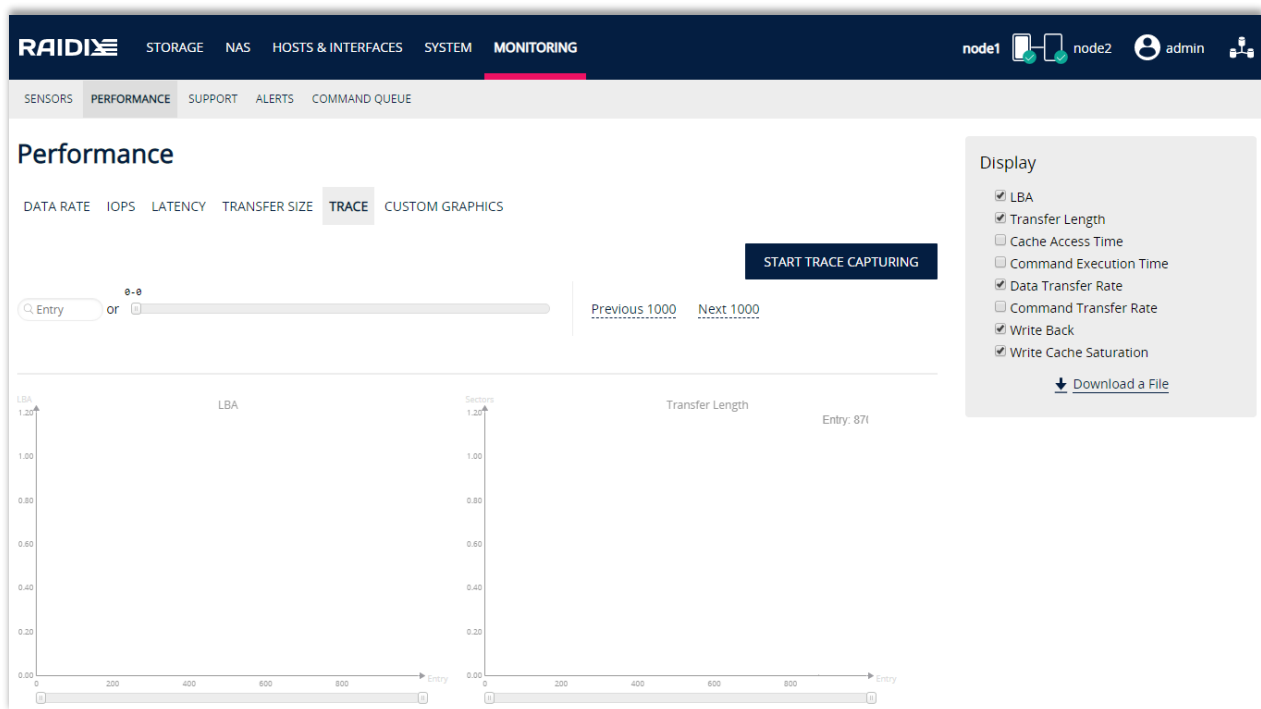


Figure 163. Performance monitoring module. Trace section fragment

The displayed information can be set using the control panel filters at the top of the page. Filters provide information about:

- specific request (entry) (to get the information, enter the internal entry number in the *Entry* window and click Enter);
- all requests history (move the cursor along the scroll bar);
- a thousand of previous or next entries (click the corresponding icon in GUI).

To start statistics capturing, click the **START TRACE CAPTURING**. The system starts collecting information about SCSI commands. Trace capturing and processing of the results will take not more than 30 seconds. Then the results display on the charts. To stop trace capturing earlier, click **STOP TRACE CAPTURING**.

You can use the pane *Display* that is provided on the right side of the window. To display a chart, check an appropriate item. The following charts are available in the *Trace* section (5 of 8 are shown by default) (Figure 164).

You can download the traces in CSV format by clicking the **Download a File** (Figure 164).

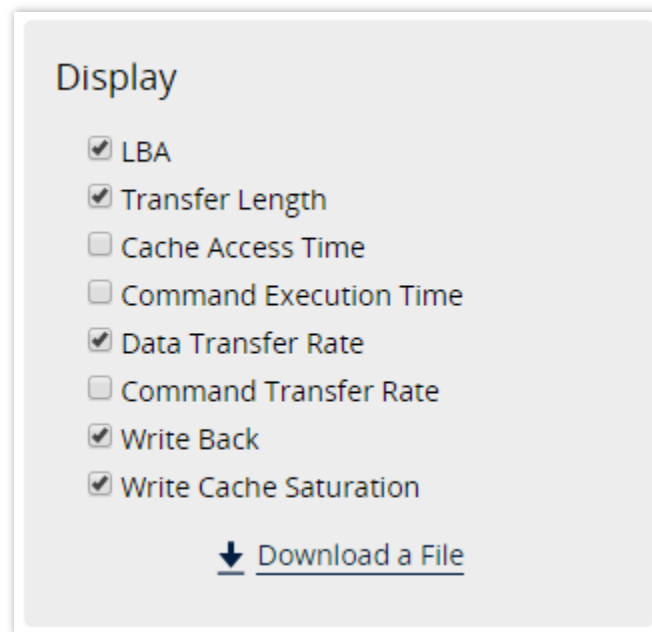


Figure 164. Display filter

To scale the graph, select an appropriate area on it with the mouse (scaled area will be marked as the light-blue color of the horizontal axis (Figure 165).



Figure 165. Scaling example

LBA graph shows the logical blocks address from which the data transfer starts and the received commands (Figure 166).

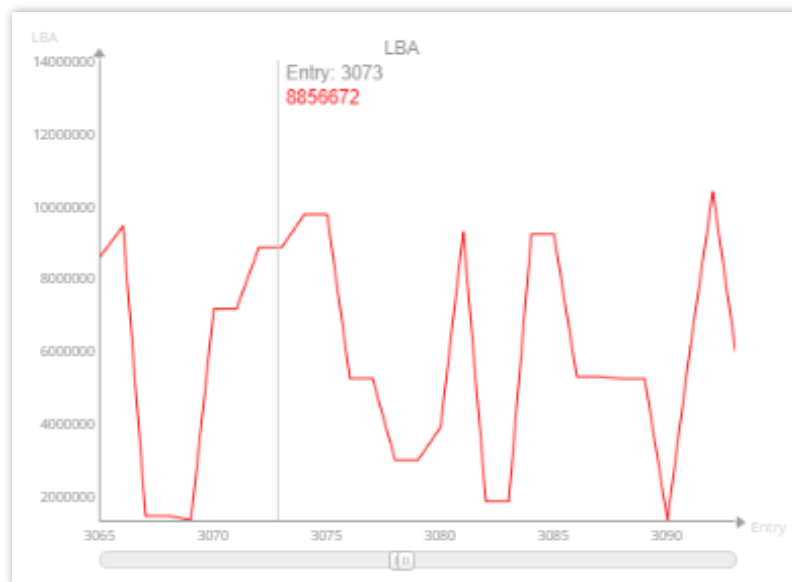


Figure 166. LBA chart example

The horizontal axis is an internal entry number of the appropriate command execution (read or write), and the vertical axis is the corresponding LBA address.

Transfer Length graph (Figure 167) shows the volume of data, transferred by commands; information on each command size is represented in sectors.

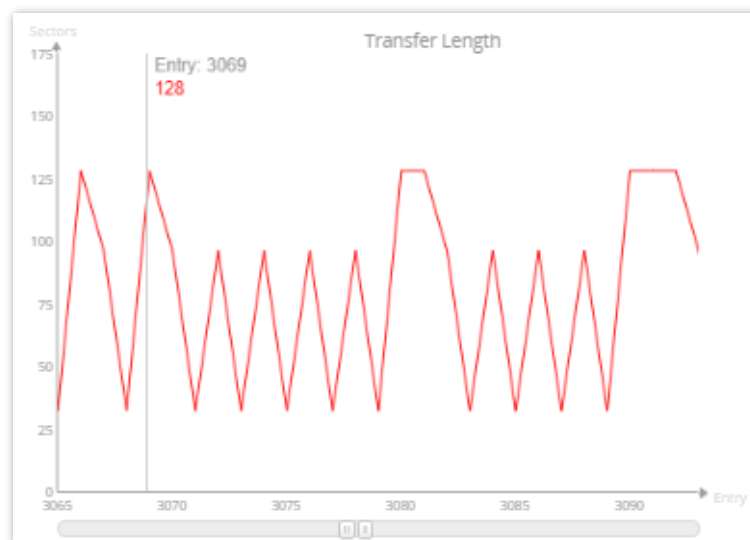


Figure 167. Transfer Length graph example

The horizontal axis is an internal number of requests (entry) on the command execution and the vertical is the corresponding request size in sectors.

Cache Access Time graph shows the time during which commands were in the cache queue waiting to be run (Figure 168).

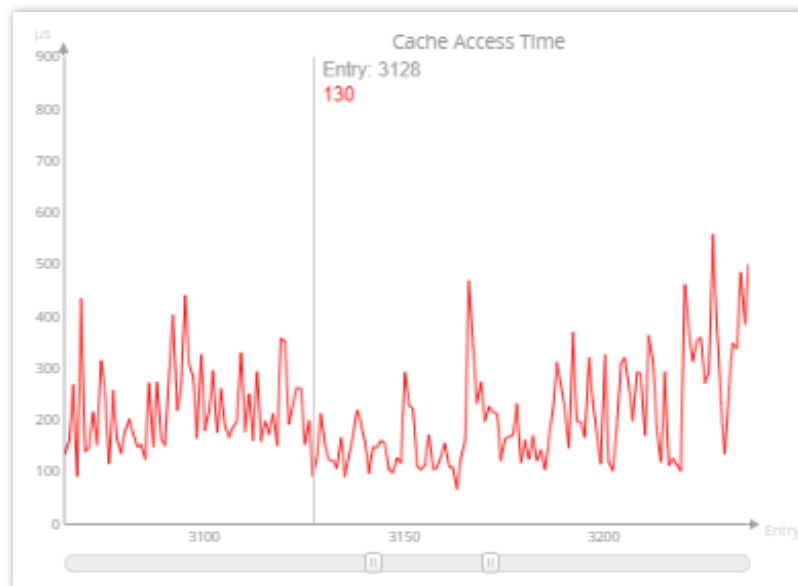


Figure 168. Cache Access Time graph example

The horizontal axis is an internal entry number of the appropriate command execution (read or write) and the vertical is time spent to get access to cache (in microseconds).

Command Execution Time graph shows time spent on commands execution (Figure 169).

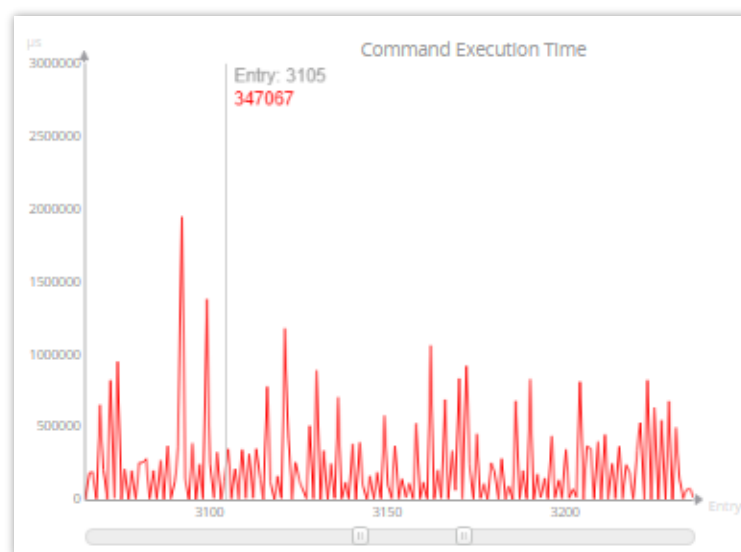


Figure 169. Command Execution Time graph example

The horizontal axis is an internal entry number of the appropriate command execution (read or write), and the vertical is time spent to run the corresponding entry since the command got access to cache (in microseconds).

Data Transfer Rate graph shows the appropriate command execution data transfer rate (Figure 170).

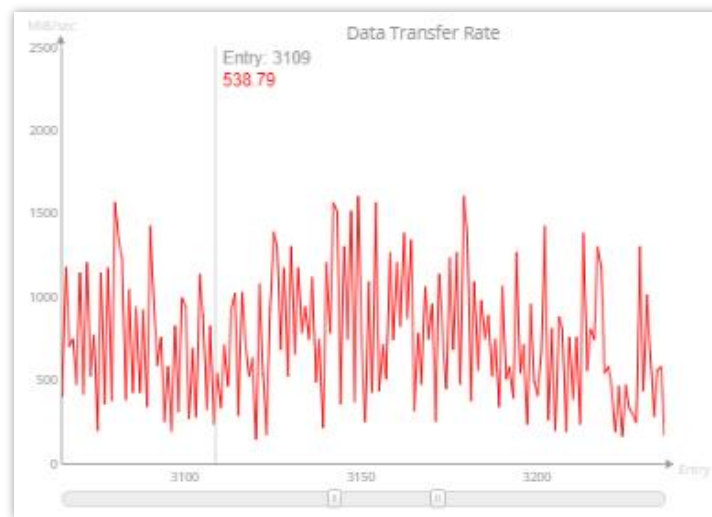


Figure 170. Data Transfer Rate graph example

The horizontal axis is an internal entry number of the appropriate command execution (read or write), and the vertical is the data transfer speed for a particular request (in MiB/sec).

Command Transfer Rate graph shows the command execution speed. Commands execution time is calculated as the sum of the time spent on processing the command and the response transfer time (Figure 171).

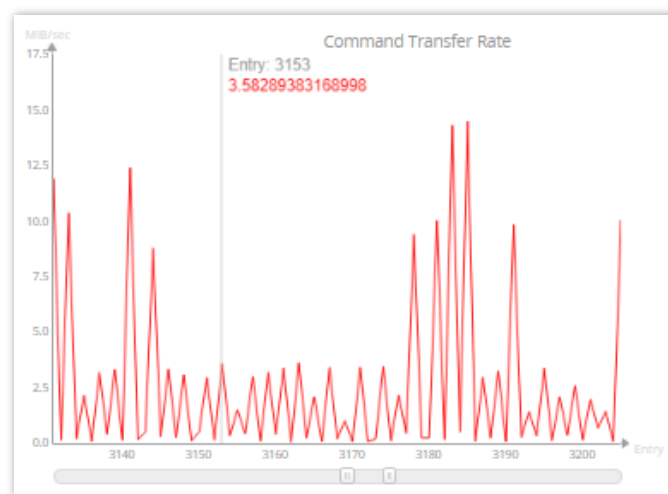


Figure 171. Command Transfer Rate graph example

The horizontal axis is an internal entry number of the appropriate command execution (read or write), and the vertical is command execution speed (in MiB/sec).

Write Back graph (from cache to disk data transfer statistics) shows the number of Write Back commands during a particular command execution (Figure 172).

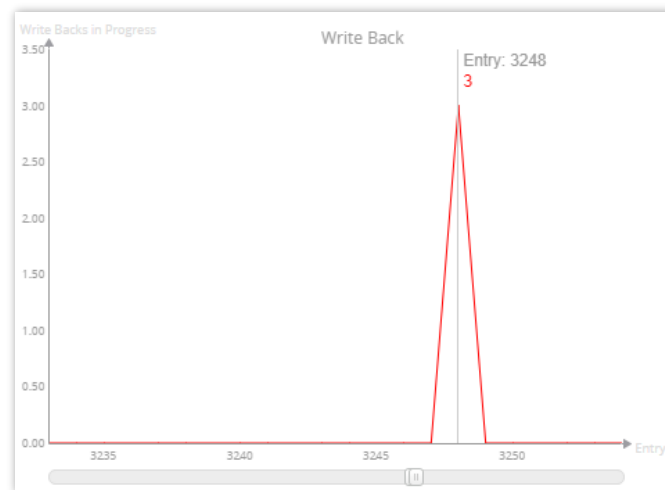


Figure 172. Write back graph example for read operation

The horizontal axis is an internal entry number of the appropriate command execution (read or write) and the vertical is the number of requests on transferring data from cache to disk when running a particular command.

Write Cache Saturation graph shows the value of cache saturation with dirty data (non-synchronized with disk cache data) (Figure 173).

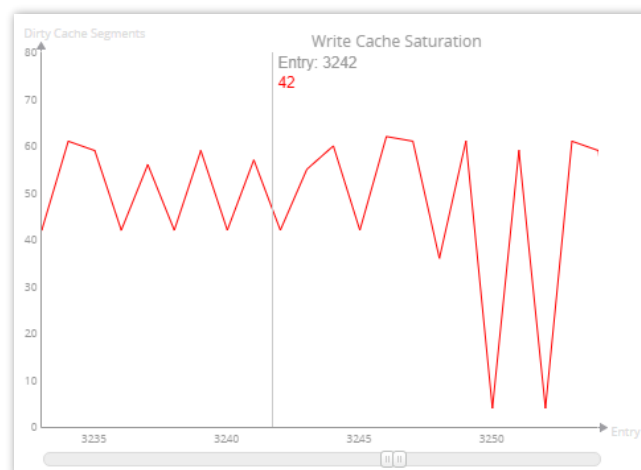


Figure 173. Write Cache Saturation graph

The horizontal axis is an internal entry number of the appropriate command execution (read or write) and the vertical is the number of Dirty cache segments. The larger the dirty cache segments amount, the quicker the commands get access to the cache.

11.3 Console

This section contains information on the console command interface.

Command console lets the administrators perform more functions in addition to those available through the web UI. The command line has a number of requirements for the typed commands. To learn more, see [Appendix A](#).

To access the console, open the **MONITORING | SUPPORT** page, and click **Show Console**. The console opens (Figure 174).

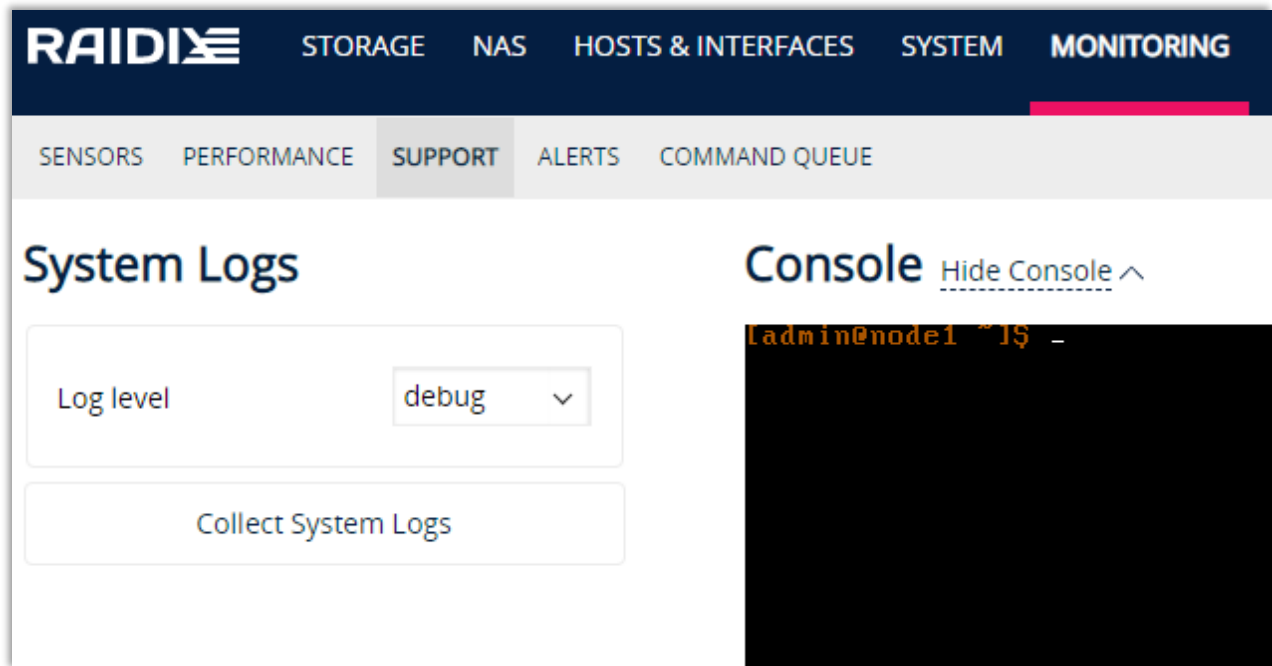


Figure 174. RAIDIX Command Console

11.4 Downloading System Logs

11.4.1 Downloading System Logs Using Web Interface

In RAIDIX 5.2, an ability to download information about the system events, that happened on the server (log file, logs), for each node is implemented.

To download system logs:

i While working in DC mode, collect system logs on each node. To do so, tick the **Collect system logs from the remote node**.

1. Open **MONITORING | SUPPORT**.
2. Select **Log level**.

In DC system, both nodes should have the same log level. Possible values: **crit**, **err**, **warn**, **notice**, **info**, **debug**. Each next level includes previous. The default is **info** – all logs are collected.

3. Click **Collect System Logs** (Figure 175).

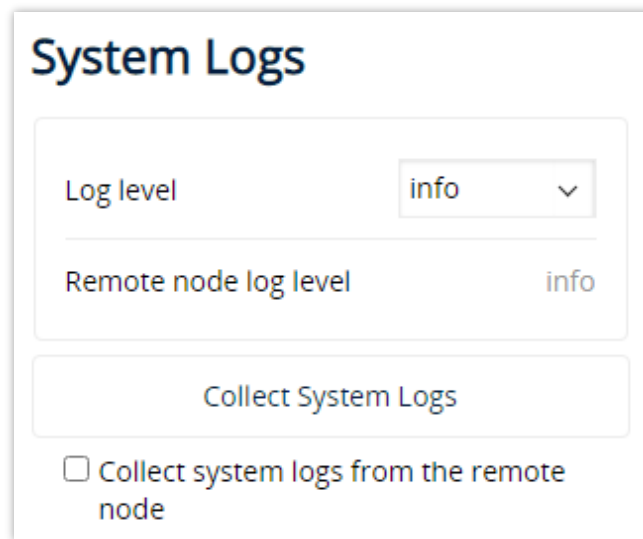


Figure 175. Collecting system logs pane in DC system

- Wait until the **Download System Logs** link displays and click the link (Figure 176). System logfile will download on your computer.

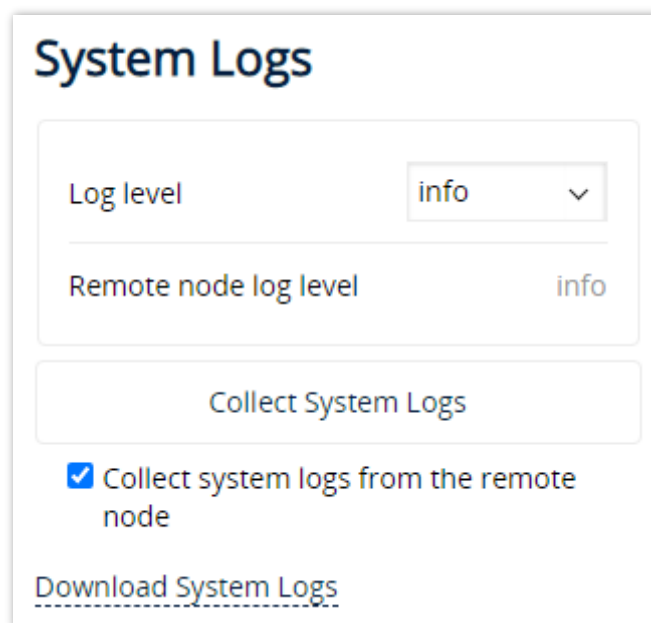


Figure 176. Download System Logs link appears after log collecting

11.4.2 Downloading System Logs Using CLI

To download system logs using CLI, do the following:


- On the **MONITORING | SUPPORT** page.
- Click **Show Console**.
- Log in as an Administrator.
- In the console, enter the command

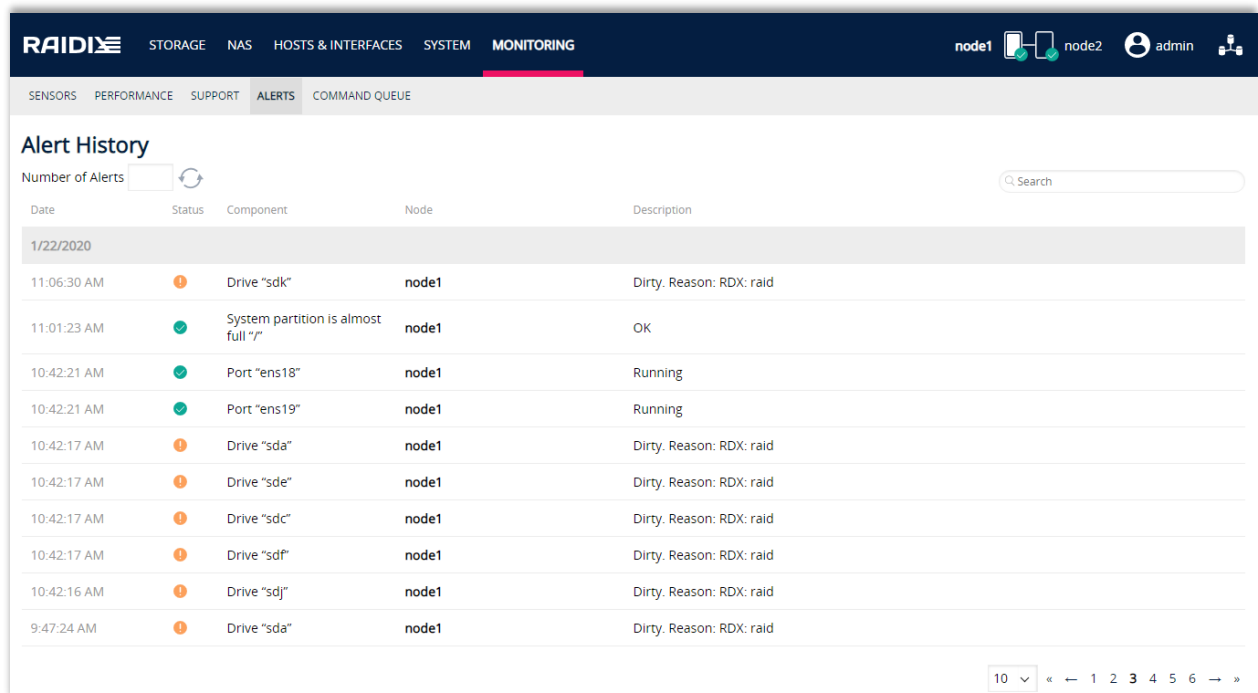
```
$ log_collector
```

The stages of log collecting will display in the console.

- After logs are collected, the system will display the name of log file (.zip archive) and file path in the `/var/spool/rdx/` directory. Download the file from the given directory.

11.5 System Status Assessment

To view detailed information on system alerts, open the **MONITORING | ALERTS** page (Figure 177). Specify the number of alert items in the **Number of Alerts** field and click the  icon.
















Date	Status	Component	Node	Description
1/22/2020				
11:06:30 AM		Drive "sdk"	node1	Dirty. Reason: RDX: raid
11:01:23 AM		System partition is almost full "r"	node1	OK
10:42:21 AM		Port "ens18"	node1	Running
10:42:21 AM		Port "ens19"	node1	Running
10:42:17 AM		Drive "sda"	node1	Dirty. Reason: RDX: raid
10:42:17 AM		Drive "sde"	node1	Dirty. Reason: RDX: raid
10:42:17 AM		Drive "sdc"	node1	Dirty. Reason: RDX: raid
10:42:17 AM		Drive "sdf"	node1	Dirty. Reason: RDX: raid
10:42:16 AM		Drive "sdj"	node1	Dirty. Reason: RDX: raid
9:47:24 AM		Drive "sda"	node1	Dirty. Reason: RDX: raid

Figure 177. Alerts History page. Information on system status

The following information on alerts is presented:

Parameter	Description
Date	Date and time then alert occurred.
Status	Parameter defines node devices state: <div>  <i>Healthy</i> – works correctly. </div>

Parameter	Description
	 <i>Warning</i> – a non-critical failure has occurred: <ul style="list-style-type: none"> RAID status is <ul style="list-style-type: none"> Degraded; Reconstructing; Initializing; SDC; Failed Over; Adapter status is Down (there is no physical connection); Drive status is <ul style="list-style-type: none"> Dirty; Faulty Warning; Reconstruction required; Reconstructing; License status is Warning because a hardware modification took place.
	 <i>Bad</i> – a critical error occurred: <ul style="list-style-type: none"> No valid licenses, license key is incorrect, or more than one hardware modification was made; Drive status is <ul style="list-style-type: none"> Faulty error; No drive; Error; RAID status is <ul style="list-style-type: none"> Offline; Not Loaded; An error has occurred during the initialization; Caches are not synchronized.
Component	Alert component. Possible values: <ul style="list-style-type: none"> Drive; Port; RAID; License; LUN; System partition; DC.
Node	Node name.
Description	A description of the problem device status.

11.6 Command Queue

While working with the system, part of the executed commands may take a long time. During the execution of such commands, in order to maintain data integrity and system operability, the system blocks using most of its functionality. Herewith, in GUI, the system disables controls (Figure 178).

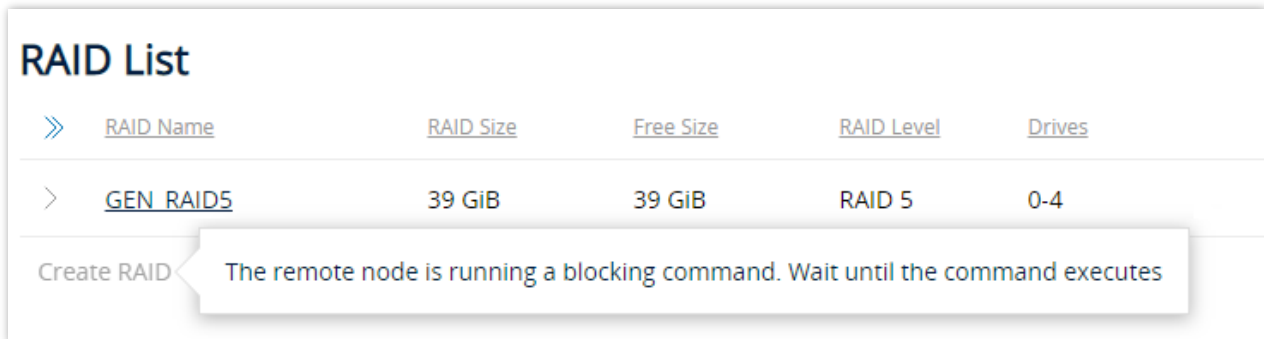


Figure 178. Disabled control

Only the show commands available for execution on the local node, regardless of whether the blocker command is currently running or not.

To get information about running commands, use the *Command queue* functionality.

To open the *Command queue* page, select **MONITORING | COMMAND QUEUE** (Figure 179).

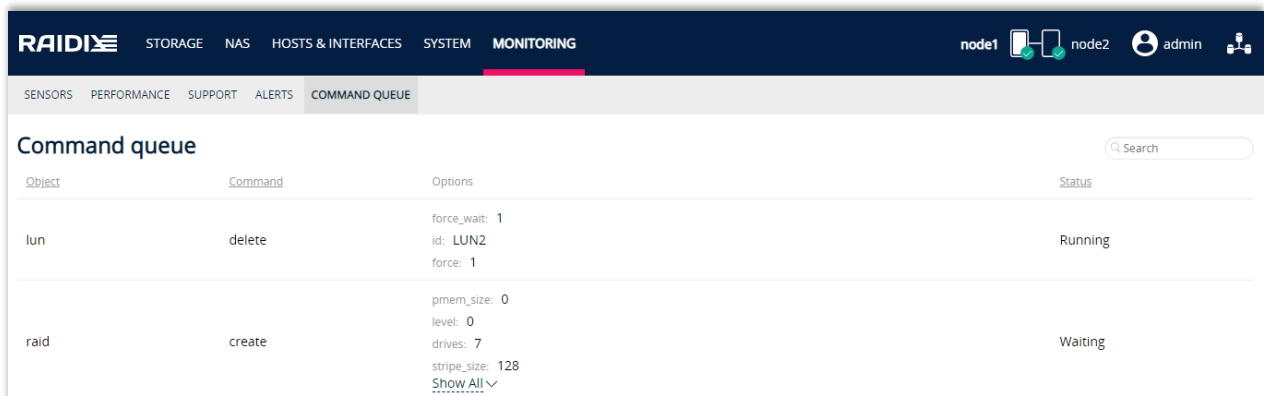


Figure 179. Command Queue page

Fields and descriptions in the *Command queue* section

Field	Description
Object	A system object on which the command was initiated.
Command	An action performed by the command.
Options	Options (parameters) with which the command was initiated.
Status	Status of the command. Possible values: <ul style="list-style-type: none"> Running – currently executed command; Waiting – the command is in the waiting state.

To view a command queue on the remote node of the DC system, run

```
$ system command show -rm <remote_ip>
```

12. TROUBLESHOOTING

This chapter contains descriptions of possible errors and steps required to solve them.

1. GUI is not available

The possible reason is that the Ethernet interface of the node is not available.

1. Verify nodes availability using the appropriate IP address. If the node is unavailable, it is most likely that the Ethernet interface is down.
2. If the node is up, check if there is enough of free memory space on the drive and storage memory using the following commands:

```
$ df -h
```

```
$ free -h
```



Turn off the node with failed Ethernet interface manually:

- In single-controller mode, the node can be turned off only in case of a total absence of workload.
- In DC mode, before turning off the failed node, transfer all arrays to the healthy node (by clicking **Failover**).

2. The drives are unavailable in GUI after a new drive enclosure is connected or after the an already used enclosure is reconnected

Reboot the system.

If the system reboot has not helped, reboot the enclosure, and then reboot the system again.

3. Current RAID state is *Not Loaded*

To load the RAID, try one of the following methods:

- Check if the status of every drive is *Healthy*.
- Verify the specified parameter values and try to increase the cache size.
- Enter the following CLI commands to restore the RAID and the LUN:

```
$ rdcli raid restore -n <RAID_name>
```

```
$ rdcli lun restore -n <LUN_name>
```

- Reboot the system.
- Reload the RAID.

4. Current LUN, Share, or Mask state is *Not Loaded*

Run the following commands:

```
$ rdcli raid restore -n <RAID_name>
```

```
$ rdcli lun restore -n <LUN_name>
```

5. Synchronization does not start automatically

- Make sure connection status of every port for synchronization is *Up*.
- Make sure targets for synchronization are selected correctly.
- In case of synchronization via InfiniBand, make sure OpenSM service is running on one node on the proper port.
- When synchronizing over iSCSI, make sure iSCSI is enabled, and IP address of the remote controller, used for synchronization, is available.
- Check out enabled IP addresses for iSCSI at **HOSTS AND INTERFACES | ISCSI TARGETS** on an upper-right panel.

6. If you renamed RAIDs and volumes, there is a possibility that LUNs will not display the file system after the import

After the import is complete, run the following command:

```
$ rdcli lun modify --name <LUN_name> --formatted yes
```

7. Errors occurring during upgrading of the firmware on the SAS target

Upgrade the firmware on the SAS target by using `managesas/srctnyccli` (not SAS Flash).

8. Errors occur due to the firmware of the SAS adapters

While working with adapters based on the LSI SAS2308 6 GB (9201 / 9205 / 9207 / 9211) chip, make sure the firmware version is P19.

9. Errors occurring in the process of configuration or after the configuration of the SAS adapters' ports

Set up adapters' ports by using `managesas` and through dividing them into groups (see *RAIDIX 5.2 Software Installation and Configuration Guide*).

10. The lsscsi on Linux initiator displays the following entry before adding the LUNs:

```
disk Raidix N/A 4.4. -
```

```
disk Raidix N/A 4.4. -
```

or

```
disk mpt3sas mpt3sas target 4161 -
```

1. Remove devices by running the following command:

```
$ echo 1 > /sys/class/scsi_device/1\:0\:19\:0/device/delete
```

2. Rescan devices.
3. After you rescan devices, find all LUNs by performing the following command:

```
$ echo '- - -' > /sys/class/scsi_host/host5/scan
```

4. Run the `lsscsi` command, the result should be similar to the following:

```
lsscsi

[0:0:1:0]    disk      IBM-ESXS  MBE2073RC          SC19  /dev/sd1
[1:0:1:0]    disk      ATA       GB0250EAFYK        HPG0  /dev/sdm
[5:0:0:1]    disk      Raidix    001                0004  /dev/sdab
[5:0:0:2]    disk      Raidix    002                0004  /dev/sdx
```



Due to operating system specifics, it is not possible to perform LUN rescan on MAC OS.

11. When you write a large file on the LUN with the help of the MAC OS file manager, the additional 4K file appears. You will not be able to delete this file if you configured WORM with a small grace period

- Try to set up the longer WORM grace period to delete the file.
- Log in to the RAIDIX through ssh and delete the file manually.

12. The maximum number of LUNs that you can create is 447. OS Windows defines the limited number of LUNs: from 0 to 230

Linux detects all LUNs, but we don't recommend creating more than 231 LUNs in Windows.

13. After upgrading from RAIDIX 4.3.1 to RAIDIX 5.2, users with the administrator rights do not appear for the SMB protocol

After upgrading, add users with administrator rights to the SMB using the web interface or CLI. To learn more, see [Manage NAS Users and Groups](#).

14. HBAs are not displayed on ESXi 6.0 initiator after reboot of the nodes in DC

Rescan the adapters twice.

15. iSCSI Initiator Small Tree software (2.0.0 – 2.1.2) will not connect to RAIDIX

The problem occurs while running CRC32C checks on Mac OS X (10.9 and 10.10).

16. Impossible to change block size for LUN while using ESXi

While using ESXi, the block size for the LUN can be only 512 bytes, since ESXi does not support LUNs with a block size of 4KiB.

17. If connection with enclosure is broken, read and write operations terminate with an error

Wait until failover occurs. I/O errors occur on hosts, as all RAIDs are on the active node but contain no drives.

18. Data corruption occurs when the connection between the active node and the enclosure is broken

The problem occurs in dual-controller-mode, if the enclosure is connected to each node with one cable and synchronization of caches is configured through the enclosure.

19. In DC system, write to an SMB share, using a virtual IP, stops after failover and does not renew after failback

The problem occurs while writing from Windows Server 2008 R2 initiator. For Windows Server 2012 R2 initiator, write renews automatically after failback as soon as the active node is restored.

20. Delays in displaying objects of web interface; impossibility to move to other pages of web interface; impossibility to add commands to queue while working in web interface

The problem occurs while some command is being executed through CLI or web interface. Wait until the command execution is complete.

Check out running commands at **MONITORING | COMMAND QUEUE**.

21. Synchronization of Persistent Reservations in DC is disabled by default

Enable Synchronization of Persistent Reservations while configuring DC if you use Windows Failover Cluster or Hyper-V.

22. After update to RAIDIX 5.2, SSD cache, formerly created, is absent in the system

After update to RAIDIX 5.2, due to architectural features, SSD cache is required to be recreated. After update in dual-controller mode, use only SAS SSD for SSD cache. SATA and PCI Express SSD are supported only in single-controller mode.

23. After updating to RAIDIX 5.2, some RAIDs have the status *Unloaded*

For some RAID types, particularly with small stripe size, bigger cache size is required in RAIDIX 5.2. After system update to RAIDIX 5.2, such RAIDs will have the status Unloaded after migration; it is necessary to increase cache size to load them. Define the necessary cache size with `rdctl` program, which you can obtain at RAIDIX technical support team.

24. Inadequate trace results. Wrong structure of the trace files

To avoid defects in the structure of trace files and inaccurate results, do not trace while working on high data rates, particularly if SSD cache is enabled.

25. Quota for user 'nobody' for NFS shares does not work

Quota for user "nobody" for NFS shares works only for Linux users. To use quotas for user "nobody" in other OS, enable the "squash" parameter.

26. Impossible to import the RAID, on which LUN, extended on several RAIDs, is located

The problem relates to renaming of the LUN. If while importing a few RAIDs there is a necessity to rename the LUN located on these RAIDs, first import the RAID, on which the renaming was initially made.

27. During write on the active node, graphs change on the passive node as well

The problem occurs while the filter **Session** is selected.

28. No data transfer between the target FC Qlogic 8 Gb and Windows 2012 initiator FC ATTO 8 Gb

It is necessary to limit the maximum request size. To do so, set the MaxXferSize value equal to **4194304** (i.e. 4 MB) in the following Windows registry keys:

- `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\clry8fc\Parameters\Device`
- `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\clry16fc\Parameters\Device`

29. Can't create RAID of RAID0 or RAID10 level

Try recreating the RAID with the parameter *4 KiB blocks write mode* enabled.

30. Can't create a bonding. The error Interface <interface name> MTU mismatches with bond interface <bonging name> MTU appears

1. Check values of the MTU parameter for all interfaces in bonding. Make sure the values are the same. MTU is displayed on the **SYSTEM | NETWORK** page or in the output of the command:

```
$ rdcli network interface show
```

2. Run the command

```
$ rdcli network bond create -mtu 1500 -n <bonding_name> -s
(<comma-separated_list_of_interfaces>) -p rr
```

31. After a failure of the secondary replication node some data remains unsynchronized even after the failed node recovers

1. Start verification check from Web UI or by running the command:

```
$ rdcli replication verify start
```

2. If after verification check is finished, the out-of-sync parameter is not equal to zero, suspend replication on any replication node.
3. Start writing to the primary node until the out-of-sync value changes.
4. Resume replication. The data synchronization will start.

32. Errors while reading and writing to a LUN added to an ESXi virtual machine as an RDM drive in the physical compatibility mode

1. Write an udev rule that disables ALUA while creating a new block device. The rule might look like this:

```
ACTION=="add", SUBSYSTEM=="block", KERNEL=="sd*",
ATTRS{vendor}=="Raidix", ATTR{device/dh_state}="detach"
```

2. Save the rule to the directory

```
/lib/udev/rules.d/<free_number>-<name>.rules
```

33. When creating multiple ERA RAIDs with a large stripe size, the system shows an error message "Not enough RAM to create RAID"

1. Find out free RAM size by running the command

```
$ rdcli system show | grep available_memory
```

2. Find out used by RAIDs RAM size by running the command

```
$ rdcli eraraid show
```

3. If necessary, change used by RAIDs RAM size by running the command

```
$ rdcli eraraid modify -n <raid_name> -ml <ram_limit>
```

where

- <raid_name> is a RAID name for which the memory size is changing;
- <ram_limit> is the amount of RAM being set.

34. Can't create LUN with replication metadata. The warning "*Can't create meta-LUN '...' for replication, LVM volume group '...vg' exists.*" shows

If a LUN with replication metadata was deleted not through RAIDIX, metadata parts may keep on drives. To create a LUN with metadata on such drives:

- use initialized RAIDs;
- delete the metadata parts via third party utilities (*pvs*, *vgs*, *lvs*, and others).

35. The system drive is not automatically returned to MDRAID. When a drive is returned to MDRAID, the drive is displayed in the list of common drives for RAID creation

After you physically return a drive in MDRAID, run

```
$ mdadm --manage /dev/<md127> --add /dev/<sdX1>
```

where <md127> is a MDRAID, <sdX1> is the drive.

Run drive rescan:

```
$ rdcli drive rescan
```

36. After manual replacement of a multipath drive, RAID reconstruction will not start automatically

Add a new drive manually using GUI or CLI.

37. depmode shows "Required key not available" when trying to load a driver

In UEFI, turn Secure Boot off.

38. Linux host does not get LUNs provided by QLogic QLE2562 8Gb, QLogic QLE2672 16Gb, or QLogic QLE2742 32Gb

A temporary solution without rebooting the initiator:

Run:

```
# rmmod qla2xxx

# modprobe qla2xxx ql2xnvmeenable=0
```

A fixed solution with the initiator rebooting:

1. Create a file `qla2xxx.conf` with the option:

```
# echo "options qla2xxx ql2xnvmeenable=0" >
/etc/modprobe.d/qla2xxx.conf
```

2. Depending on host OS:

- RHEL:

```
# dracut -force
```

- Ubuntu:

```
# update-initramfs -u
```

3. Reboot the host:

```
# reboot
```

39. Write/read operations on the host hang when the RAID goes offline

Both I/O on the storage and local disks (using SCSI commands) can hang, if the RAID goes offline during the writing.

Workarounds:

- Disconnect the host from the target (sometimes unavailable for Windows hosts). In this case:
 - Write/read operations will fail on the RAID.
- Return the RAID to the online state. Possible results:
 - Write/read operations will resume on the RAID.

- Write/read operations will fail on the RAID.

40. RAID reload is not completed if a failure occurs on the passive for this RAID node

Run

```
$ rdcli raid reload --local -n <raid_name>
```

41. On Ubuntu 20.04 host with FC QLogic16Gb HBA one mpio path is missing after restarting a node of the DC system

Try one of the options:

- First option. On the initiator, run

1. Find out adapter ports:

```
# lspci | grep -i QLogic
```

Output example:

```
05:00.0 Fibre Channel: QLogic Corp. ISP8324-based 16Gb Fibre
Channel to PCI Express Adapter (rev 02)
```

```
05:00.1 Fibre Channel: QLogic Corp. ISP8324-based 16Gb Fibre
Channel to PCI Express Adapter (rev 02)
```

2. Remove each adapter port:

```
# echo '1' > /sys/bus/pci/devices/0000\:05\:00.0/remove
```

```
# echo '1' > /sys/bus/pci/devices/0000\:05\:00.1/remove
```

3. Check that the adapter ports have been removed (the command output should be empty):

```
# lspci | grep -i QLogic
```

4. Scan PCI devices

```
# echo '1' > /sys/bus/pci/rescan
```

- Second option. If both paths are missing, on the initiator, run

```
# rmmod qla2xxx
```

```
# modprobe qla2xxx
```

- Third option. Reboot the Ubuntu system with the initiator.

42. Crash of a host CentOS 8.2 when reading large blocks over NVMe-oF

Install the last OFED and the `kmod-mlnx-nvme-5.1-OFED.5.1.2.3.7.1.rhel8u2.x86_64` package on CentOS.

43. Workload from Ubuntu 20.04 with the initiator stops

After the "failover", "auto-failover", or "migration" of the ERA RAID, initiator on Ubuntu may temporarily "lose" the path over NVMe-oF to target. With a large number of objects or high load on the storage, the paths between the initiator and the target may take a long time to recover (about 5 minutes). If you perform "failover", "failback", or "migration" before the paths are restored, the load may be interrupted.

To check the status of the paths, run on the Ubuntu with the initiator the command

```
# multipath -ll
```

44. Workload over NVMe-oF from an initiator with a large number of namespaces interrupts

When working with the DC system, if there is a high workload on a large number of namespaces, the host may disconnect from one of the subsystems after switching nodes (failover or failback) and will try to reconnect without result. To reconnect the host to the subsystem, do the manual reconnection.

For the initiators on Linux, the disconnection will be shown by the `connecting` status after running the command

```
# nvme list-subsys
```

After reconnecting the initiator to the subsystem, the status should be `live`.

For NVMe-oF configurations, we recommend allocating no more than 20 namespaces to a single initiator. The exact number of namespaces depends on the workload and configuration of the system with the initiator.

The following configuration was used for testing:

- CPU model: Intel Xeon CPU E5-2620 v4 @ 2.10GHz
- CPU cores: 32
- RAM: 125 GB
- OS: Ubuntu 20.04

45. Can't failover or migrate passive RAIDs via CLI

If the DC system is powered off or is rebooted and then only one node is powered on, the passive RAIDs on this node will remain passive:

- auto-failover will not start;
- failover and migration of the RAIDs will be unavailable.

To set the RAIDs to the active state, run

```
$ rdcli dc failover --force
```

46. Multipath crashes on Ubuntu 20.04 OS host after failover or failback

There is a small chance that multipath may suspend on an Ubuntu 20.04 host with an *LTS kernel* after failover or failback. The suspension may be temporary or permanent with the stopped workload and path failure.

47. Failed path or wrong multipath priority on Ubuntu 20.04 host.

After the failure and subsequent recovery of the DC system node, the following problems are possible on the host:

- the status of the multipath is “failed” and the path is not automatically restored;
- the multipath priority (prio) is wrong and the priority is not automatically restored.

To fix this, rescan the block device on the host:

```
# echo 1 > /sys/block/sdX/device/rescan
```

48. The error “mount error(2): No such file or directory”

If you already have an SMB share on the host from the node, you can only mount a new SMB share 5 minutes after it is created from the same node network interface.

49. System crashes when running `sysctl -a`

The `sysctl -a` command can crash the system. We recommend limiting the output, for example by using the `--pattern` argument.

50. NAS: FTP connection over VIP in a DC system fails when failover

Different FTP clients handle loss of connection to the server differently.

For Linux clients, we recommend using `lftp`.

51. Poor NFSoRDMA performance

When mounting a shared folder on a Linux client over NFSoRDMA with the `mount` command, use the `-o rdma` option. See [Linux mount man](#) and [Linux nfs man](#) for details.

GLOSSARY

Term	Definition
Advanced Reconstruction	Feature that allows to increase reading speed from a RAID with slow drives. Instead of reading data from slow drives, the system recovers this data through RAID redundancy. Inapplicable to RAID 0.
Affinity	The RAID property that means which node handles I/O operations for that RAID while the DC system is in a normal operation state.
AFP (Apple Filing Protocol)	Network protocol that provides access to files in Mac OS X.
Cache	A part of RAM of the RAID, in which the data for lowering the time required to serve the hosts (initiators) requests, is temporarily stored. Cache is used to increase the data exchange rate as the speed of read and write operations for RAM is much higher compared with the same ones for the drives, and the host exchange speed is unpredictable. The more the cache size is, the higher the system productivity is. However, it mustn't exceed the total value of the operating and system memory.
CHAP (Challenge Handshake Authentication Protocol)	Basic authentication used by host and target to authenticate each other via the exchange of a secret code or password.
DHCP (Dynamic Host Configuration Protocol)	A network protocol letting the computers automatically obtain IP address and other parameters required for operating over the TCP/IP protocol.
Driver	A specific computer program enabling the operating system access to a target device. Usually supplied by a device vendor.
FC (Fibre Channel)	High-speed data transfer interface used for storage networking such as workstations, mainframes, supercomputers, and storage devices.
NVMe-oF	A specification-defined extension to NVMe that enables NVMe-based communication over interconnects other than PCIe. This interface makes it possible to connect "external" storage enclosures to a server, either directly or through a switch, while still using NVMe as the fundamental communication mechanism.
Namespace	A quantity of nonvolatile memory that can be formatted into logical blocks. This definition is virtually identical to the SCSI concept of a logical unit.
NQN (NVMe Qualified Name)	An identifier for a remote NVMe target. Similar with iSCSI Qualified Name (IQN).
Subsystem	A logical entity includes masking rules for namespaces and a list of ports to connect with hosts.
Referral	A list of NVMe-oF ports to search for namespaces.

Term	Definition
FTP (File Transfer Protocol)	Standard protocol for transferring files over TCP networks.
GDL (Grown Defect List)	A list of defects, which appeared during the drive usage. The list is automatically updated by the drive's S.M.A.R.T. technology.
GUID (Globally Unique Identifier)	Globally unique 128-bit identifier. In the RAIDIX 5.2 software, GUID is used to identify InfiniBand hosts.
HBA (Host Bus Adapter)	A device that host uses to connect to the network or storage devices.
Heartbeat	Ethernet data channel used to exchange metadata between storage nodes in DC.
Host (Initiator)	Initiator is a SCSI-device that starts a data request. In RAIDIX 5.2 terms, initiators are hosts.
Host Alias	An alias of the host – easy-to-remember name assigned to the host. Several hosts can be combined by one host alias.
iSCSI (Internet Small Computer System Interface)	An IP-based protocol built on SCSI. It carries block-level data over traditional IP networks.
iSER (iSCSI Extensions for RDMA)	Computer network protocol that extends the iSCSI protocol to use Remote Direct Memory Access (RDMA).
InfiniBand	A computer networking communications standard used for data interconnection.
IQN	A unique iSCSI ID.
LACP (Link Aggregation Control Protocol)	An IEEE standard for combining two or more physical data channels into one logical data channel for high availability.
LBA (Logical Block Address)	Address of data block referred to during the read-write request.
LUN (Logical Unit Number)	A virtual section of a RAID. In SCSI protocol, the term is used as a way of addressing drives within the device having one SCSI Target ID, for example, a disk RAID.
MTU (Maximum Transmission Unit)	A setting that determines the size of the largest packet that can be transmitted without data fragmentation.
Multipath I/O, MPIO (Multi-Path Input-Output)	Technology that provides the hosts with access to LUN partitions in several paths. The architecture provides fault tolerance growth and load balancing for the system.
NAS (Network Attached Storage)	Network Attached Storage. Network system architecture represents a NAS server combined with storage system on RAIDIX platform and interacting with client computers via SMB/CIFS, NFS, FTP, and AFP protocols.

Term	Definition
NFS (Network File System)	Network access protocol to file systems, allowing you to connect remote file systems over the network. Supports authentication and access control features.
Nodes	Hardware-independent components of the storage system, which have own processors, cache memory, and motherboards, united into a single high-availability cluster (for the dual-controller mode). The cluster nodes are managed by the RAIDIX 5.2 software.
OpenSM (Open Subnet Manager)	InfiniBand subnet manager for centralized management.
PDU (Protocol Data Unit)	Protocol data exchange unit.
Persistent Reservations (SCSI-3)	Persistent Reservations enables multiple clients (initiators) to communicate with a target by tracking multiple initiator-to-target relationships. These relationships are links between a specific SCSI initiator port and SCSI target port for a given LUN within the SCSI target.
RAID (Redundant Array of Independent Disks)	RAID (Redundant Array of Independent Disks) is a single usable data disk, where several physical disks are combined into an array for better speed and fault tolerance.
SAN (Storage Area Network)	Storage area network is a dedicated storage network that provides access to consolidated block-level storage (such as disk RAID, tapes, optical drives) and connects it to the computers so that the operating system considers the connected resources as the local ones.
SAS (Serial-attached SCSI)	Point-to-point protocol used to transfer data between servers and SAS storage devices.
Sensor	An element for measuring the state of the disk RAID components, for example, measuring voltage, fan speed or temperature. The sensor converts the monitored parameter into a signal, which is useful for measuring, transmission, transfer, storage, and registration of monitored object state information.
SCSI (Small Computer System Interface)	A set of standards for physically connecting and transferring data between computers and peripheral devices.
Slot	The bay where the drive is inserted in the storage system rack or drive enclosure.
S.M.A.R.T. (Self-Monitoring Analyzing and Reporting Technology)	A hard disk drive monitoring technology using built-in self-diagnosis hardware and a mechanism of predicting disk failures. A special software performing S.M.A.R.T. parameters monitoring and warning users in case of device pre-failure state.
SMB (Server Message Block)	Network application protocol for remote access to files, printers, other network resources, and inter-process communication.
SparePool	Set of drives, which will be used for hot spare of a failed drive in the RAID.

Term	Definition
Squash	A parameter allows network storage users with read and write permissions to modify or delete files, which are not their own.
System metadata	Information about the system setup and properties.
Target	<p>In terms of storage systems, targets – storage devices on a data storage system, on which clients (called initiators) sent SCSI commands.</p> <p>In general, a target may be meant as a data storage system.</p>
WWN (World Wide Name)	<p>A unique identifier assigned to Fibre Channel target device.</p> <p>WWN is a 64-bit identifier. It is vendor-supplied information that is written into the programmed memory of the device itself. WWN – is the combination of the node name (World Wide Node Name, WWNN) and the port name (World Wide Port Name, WWPN).</p>

APPENDIX A.

CLI FEATURES AND SYNTAX

By default, to authenticate in CLI under an administrator role, use the login **admin** and the password **raidix-era**.

The following format is used in common to type the commands in the CLI:

```
$ rdcli <object> <method> (<required_parameters>)
[ (<oprional_parameters>) ]
```

List of the <object> values with functionality description

Object Name	Functionality
adapter	Adapter management
dc	Cluster management
drive	Drives management
eraraid	ERA RAID parameters management
fc	Setting up Fiber Channel adapters parameters
host	Hosts management
ib	InfiniBand settings management
iscsi	iSCSI settings management
lun	LUN parameters management
mask	Masking parameters management
metadata	Metadata management
nas	Network Attached Storage settings management
network	Network settings management
nodes	Displaying information on RAIDIX systems in the local network
notify mail	Email notifications settings management
notify snmp	SNMP notifications settings management
nvme fabrics	NVMe-oF settings management.
param	System parameters management
profile raid	RAID profiles parameters management
qosmic	Qosmic management
raid	Generic RAID parameters management
replication	Replication management
sas target	SAS settings management
sensor	Sensors control
session	Displaying information about all sessions running in the system
sparepool	SparePools parameters management

Object Name	Functionality
ssdcache	SSD cache management
system	System settings configuration

The command `log_collector`

The command collects system logs into a ZIP file with specified path. The type format describes in the table [log_collector](#).

Conventions on CLI-commands syntax

Item's format	Description
item	Required item (object, method, parameter, attribute).
<item>	Placeholder variable.
[item]	Additional item.
{item1 item2}	Mutually exclusive arguments.
(items)	Multiple values for the argument.

To get a hint for object, method, or parameter, use the attribute `-h`:

```
$ rdcli <object> <method> -h
```








The table below contains a detailed description of commands. The description of available commands for each object is in the tables below. The *Functionality* column contains information about a command; required parameters and additional parameters with abbreviated and full argument entries, and descriptions; and attributes if they are available for the command.

adapter

Method	Command Syntax	Functionality
show	rdcli adapter show	Command shows list of all adapters in the system.

dc

Method	Command Syntax	Functionality		
create	rdcli dc create	Command sets up Heartbeat connection to configure DC system.		
		Required parameters		
		-a	--ipaddr	Remote IP address for Heartbeat.
		-p	--password	Administrator's password on remote node. Interactive parameter: if you do not specify a parameter in the command, the system will prompt you to enter the password. Restrictions: 50 seconds to enter a password, 5 seconds for each character. If you fail to enter or enter an incorrect password, you will see the message "Remote node error: Access denied".
		-u	--user	Administrator's login on the remote node.
		Optional parameter		
		-l	--local_ip	Local IP address for Heartbeat.
		Optional parameter (does not require an argument)		
		-f	--force	Force the creation of the DC mode with different amounts of RAM (the difference is less than 256 MB) on the nodes.
delete	rdcli dc delete	Command disables dual-controller mode.		
		Optional attributes (do not require an argument)		
		-f	--force	Force command execution.
		-fw	--force_wait	Force disabling of DC mode after failover if there is SSD cache on both nodes.
		-s	--split	Disable DC leaving RAID's on different active nodes. After such disabling, re-enable DC will be impossible without removing or deleting RAID's from one of the nodes.

Method	Command Syntax	Functionality																		
		<div>  If there are RAIDds in the system, which are active on the different nodes, then at presence of SSD cache on these RAIDds, DC disabling after failover may take a long time due to SSD cache flushing to the drives. </div>																		
modify	rdcli dc modify	<p>Command configures target synchronization and changes some DC system parameters.</p> <p>Optional parameters</p> <table> <tr> <td>-t</td><td>--targets</td><td> <p>Targets for synchronization (comma-separated). Possible values:</p> <ul style="list-style-type: none"> for iSCSI: iscsi (also requires the -a parameter with the value of the remote IP address for synchronization); for InfiniBand: adapter GUID values; for SAS: adapter SAS address values. </td></tr> <tr> <td>-p</td><td>--port</td><td> <p>The NVMe-oF port for the "unoptimized" path in ANA DC system operation. You must specify the port on each DC system node. Use "+" or "-" to add or remove the port.</p> </td></tr> <tr> <td>-a</td><td>--iscsi_ip</td><td> <p>Remote node IP address for target synchronization via iSCSI. Also requires the parameter -t iscsi.</p> <div>  Do not use VIP addresses for target synchronization. </div> </td></tr> <tr> <td>-wws</td><td>--wt_without_sync</td><td> <p>Enable Write Through on every RAID without synchronization.</p> </td></tr> <tr> <td>-l</td><td>--local_ip</td><td> <p>Local IP address for Heartbeat.</p> <div>  If you change the subnet, change the IP address for Heartbeat on the second node as well. </div> </td></tr> <tr> <td>-pr</td><td>--pr_sync_enabled</td><td> <p>Enable (1) or disable (0) synchronization of Persistent Reservations for the DC system.</p> </td></tr> </table>	-t	--targets	<p>Targets for synchronization (comma-separated). Possible values:</p> <ul style="list-style-type: none"> for iSCSI: iscsi (also requires the -a parameter with the value of the remote IP address for synchronization); for InfiniBand: adapter GUID values; for SAS: adapter SAS address values. 	-p	--port	<p>The NVMe-oF port for the "unoptimized" path in ANA DC system operation. You must specify the port on each DC system node. Use "+" or "-" to add or remove the port.</p>	-a	--iscsi_ip	<p>Remote node IP address for target synchronization via iSCSI. Also requires the parameter -t iscsi.</p> <div>  Do not use VIP addresses for target synchronization. </div>	-wws	--wt_without_sync	<p>Enable Write Through on every RAID without synchronization.</p>	-l	--local_ip	<p>Local IP address for Heartbeat.</p> <div>  If you change the subnet, change the IP address for Heartbeat on the second node as well. </div>	-pr	--pr_sync_enabled	<p>Enable (1) or disable (0) synchronization of Persistent Reservations for the DC system.</p>
-t	--targets	<p>Targets for synchronization (comma-separated). Possible values:</p> <ul style="list-style-type: none"> for iSCSI: iscsi (also requires the -a parameter with the value of the remote IP address for synchronization); for InfiniBand: adapter GUID values; for SAS: adapter SAS address values. 																		
-p	--port	<p>The NVMe-oF port for the "unoptimized" path in ANA DC system operation. You must specify the port on each DC system node. Use "+" or "-" to add or remove the port.</p>																		
-a	--iscsi_ip	<p>Remote node IP address for target synchronization via iSCSI. Also requires the parameter -t iscsi.</p> <div>  Do not use VIP addresses for target synchronization. </div>																		
-wws	--wt_without_sync	<p>Enable Write Through on every RAID without synchronization.</p>																		
-l	--local_ip	<p>Local IP address for Heartbeat.</p> <div>  If you change the subnet, change the IP address for Heartbeat on the second node as well. </div>																		
-pr	--pr_sync_enabled	<p>Enable (1) or disable (0) synchronization of Persistent Reservations for the DC system.</p>																		

Method	Command Syntax	Functionality			
failover	rdcli dc failover	Command performs failover of the nodes.			
		Optional attribute (does not require an argument)			
		<table> <tr> <td>-f</td><td>--force</td><td>Force the node to failover even with no Heartbeat connection with the remote node.</td></tr> </table>	-f	--force	Force the node to failover even with no Heartbeat connection with the remote node.
-f	--force	Force the node to failover even with no Heartbeat connection with the remote node.			
failback	rdcli dc failback	Command performs failback (returns the system state before failover).			
show	rdcli dc show	Command displays cluster info (cluster ID, node ID in the cluster, nodes switching status, Heartbeat connection status, IP address of the remote node, targets for synchronization).			

drive

Method	Command Syntax	Functionality
show	rdcli drive show	Command displays info on drives used in the system.
		Optional parameters
		-u --uid Display a drive with the specified UUIDs.
		-r --raid Display drives in the specified RAID.
		-sp --sparepool Display drives included in the specified SparePool.
		-at --attribute List of drive parameters to display. Not used with the --smart parameter.
		Optional attributes (do not require an argument)
		-f --free Display all drives out of RAID's and SparePools.
		-s --smart Displays results of S.M.A.R.T. diagnostics. Not used with the --attribute parameter.
		-sl --slot Display the full drives list sorted by their slot numbers.
		-ssd --ssd Show SSD drives only.
locate	rdcli drive locate	Command manages LED-indicators.
		Mutually exclusive required parameters
		-l --led_locate Enable (1) or disable (0) LED location indicator(s). Must be used with one of the parameters: all, id, or uid.
		-lb --led_bad Enable (1) or disable (0) LED fault indicator(s). Must be used with one of the parameters: all, id, or uid.
		-lm --ledmon Enable (1) or disable (0) the ledmon service.
		Mutually exclusive optional parameters

Method	Command Syntax	Functionality		
		-u	--uid	The list of drive UUIDs (comma separated without spaces and/or hyphen for a range of disks) allowed for LED indication. To change the indication, set a new list. The parameter is used with either <code>led_locate</code> or <code>led_bad</code> parameters.
		-id	--id	The list of drive IDs (comma separated without spaces) allowed for LED indication. To change the indication, set a new list. The parameter is used with either <code>led_locate</code> or <code>led_bad</code> parameters
		-a	--all	Allow the LED indication for all drives. The parameter doesn't take arguments. The parameter is used with either <code>led_locate</code> or <code>led_bad</code> parameters
clean	rdcli drive clean	Command removes partitions from the drives in the Dirty state or resets error counter on the Faulty drive.		
Optional parameter				
		-id	--id	Drive ID.
		-u	--uid	Drive UUID.
Optional attribute (does not require an argument)				
		-a	--all	Clean all drives with appropriate statuses.
enclosure modify	rdcli drive enclosure modify	Command modifies LED statuses of the enclosures.		
Required parameter				
		-l	--led_locate	Enable (1) or disable (0) blinking.
Optional parameters				
		-e	--enclosure_id	Enclosure ID.
		-u	--uid	Drive UUID.
	rdcli drive enclosure show	Command shows enclosure LED status.		

Method	Command Syntax	Functionality
enclosure show		Optional parameters
		-e --enclosure_id Enclosure ID.
		-u --uid Drive UID.
rescan	rdcli drive rescan	Command rescans all drives.

h2>eraraid

Method	Command Syntax	Functionality																																																																																																															
create	rdcli eraraid create	<p>The command creates an ERA RAID with the specified name, level and number of drives.</p> <p>Required parameters</p> <table><tr><td>-n</td><td>--name</td><td>ERA RAID name.</td></tr><tr><td>-l</td><td>--level</td><td>ERA RAID level.</td></tr><tr><td>-d</td><td>--drives</td><td>ERA RAID drives.</td></tr><tr><td>-gs</td><td>--group_size</td><td><p>Only for ERA RAIDs 50, 60, 70.</p><p>Size of the group in RAIDs 50, 60, and 70.</p><p>Possible values: from 4 to 32. The value must be a multiple of the total number of drives in the RAID.</p><p>If not specified, the system sets a value (-gs) depending on a number of drives (D) in a RAID.</p><table><tr><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td></tr><tr><td>8</td><td>4</td><td>20</td><td>5</td><td>28</td><td>7</td><td>38</td><td>19</td><td>48</td><td>8</td><td>56</td><td>8</td></tr><tr><td>10</td><td>5</td><td>21</td><td>7</td><td>30</td><td>6</td><td>39</td><td>13</td><td>49</td><td>7</td><td>57</td><td>19</td></tr><tr><td>12</td><td>6</td><td>22</td><td>11</td><td>32</td><td>8</td><td>40</td><td>8</td><td>50</td><td>10</td><td>58</td><td>29</td></tr><tr><td>14</td><td>7</td><td>24</td><td>6</td><td>33</td><td>11</td><td>42</td><td>7</td><td>51</td><td>17</td><td>60</td><td>10</td></tr><tr><td>15</td><td>5</td><td>25</td><td>5</td><td>34</td><td>17</td><td>44</td><td>11</td><td>52</td><td>13</td><td>62</td><td>31</td></tr><tr><td>16</td><td>8</td><td>26</td><td>13</td><td>35</td><td>7</td><td>45</td><td>9</td><td>54</td><td>9</td><td>63</td><td>21</td></tr><tr><td>18</td><td>6</td><td>27</td><td>9</td><td>36</td><td>6</td><td>46</td><td>23</td><td>55</td><td>11</td><td>64</td><td>8</td></tr></table></td></tr></table> <p>Optional parameters</p> <table><tr><td>-a</td><td>--affinity</td><td><p>Only for DC systems.</p><p>ERA RAID affinity: first node (0) or second node (1).</p><p>The default: the current node.</p></td></tr></table>	-n	--name	ERA RAID name.	-l	--level	ERA RAID level.	-d	--drives	ERA RAID drives.	-gs	--group_size	<p>Only for ERA RAIDs 50, 60, 70.</p> <p>Size of the group in RAIDs 50, 60, and 70.</p> <p>Possible values: from 4 to 32. The value must be a multiple of the total number of drives in the RAID.</p> <p>If not specified, the system sets a value (-gs) depending on a number of drives (D) in a RAID.</p> <table><tr><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td></tr><tr><td>8</td><td>4</td><td>20</td><td>5</td><td>28</td><td>7</td><td>38</td><td>19</td><td>48</td><td>8</td><td>56</td><td>8</td></tr><tr><td>10</td><td>5</td><td>21</td><td>7</td><td>30</td><td>6</td><td>39</td><td>13</td><td>49</td><td>7</td><td>57</td><td>19</td></tr><tr><td>12</td><td>6</td><td>22</td><td>11</td><td>32</td><td>8</td><td>40</td><td>8</td><td>50</td><td>10</td><td>58</td><td>29</td></tr><tr><td>14</td><td>7</td><td>24</td><td>6</td><td>33</td><td>11</td><td>42</td><td>7</td><td>51</td><td>17</td><td>60</td><td>10</td></tr><tr><td>15</td><td>5</td><td>25</td><td>5</td><td>34</td><td>17</td><td>44</td><td>11</td><td>52</td><td>13</td><td>62</td><td>31</td></tr><tr><td>16</td><td>8</td><td>26</td><td>13</td><td>35</td><td>7</td><td>45</td><td>9</td><td>54</td><td>9</td><td>63</td><td>21</td></tr><tr><td>18</td><td>6</td><td>27</td><td>9</td><td>36</td><td>6</td><td>46</td><td>23</td><td>55</td><td>11</td><td>64</td><td>8</td></tr></table>	D	-gs	D	-gs	D	-gs	D	-gs	D	-gs	D	-gs	8	4	20	5	28	7	38	19	48	8	56	8	10	5	21	7	30	6	39	13	49	7	57	19	12	6	22	11	32	8	40	8	50	10	58	29	14	7	24	6	33	11	42	7	51	17	60	10	15	5	25	5	34	17	44	11	52	13	62	31	16	8	26	13	35	7	45	9	54	9	63	21	18	6	27	9	36	6	46	23	55	11	64	8	-a	--affinity	<p>Only for DC systems.</p> <p>ERA RAID affinity: first node (0) or second node (1).</p> <p>The default: the current node.</p>
-n	--name	ERA RAID name.																																																																																																															
-l	--level	ERA RAID level.																																																																																																															
-d	--drives	ERA RAID drives.																																																																																																															
-gs	--group_size	<p>Only for ERA RAIDs 50, 60, 70.</p> <p>Size of the group in RAIDs 50, 60, and 70.</p> <p>Possible values: from 4 to 32. The value must be a multiple of the total number of drives in the RAID.</p> <p>If not specified, the system sets a value (-gs) depending on a number of drives (D) in a RAID.</p> <table><tr><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td><td>D</td><td>-gs</td></tr><tr><td>8</td><td>4</td><td>20</td><td>5</td><td>28</td><td>7</td><td>38</td><td>19</td><td>48</td><td>8</td><td>56</td><td>8</td></tr><tr><td>10</td><td>5</td><td>21</td><td>7</td><td>30</td><td>6</td><td>39</td><td>13</td><td>49</td><td>7</td><td>57</td><td>19</td></tr><tr><td>12</td><td>6</td><td>22</td><td>11</td><td>32</td><td>8</td><td>40</td><td>8</td><td>50</td><td>10</td><td>58</td><td>29</td></tr><tr><td>14</td><td>7</td><td>24</td><td>6</td><td>33</td><td>11</td><td>42</td><td>7</td><td>51</td><td>17</td><td>60</td><td>10</td></tr><tr><td>15</td><td>5</td><td>25</td><td>5</td><td>34</td><td>17</td><td>44</td><td>11</td><td>52</td><td>13</td><td>62</td><td>31</td></tr><tr><td>16</td><td>8</td><td>26</td><td>13</td><td>35</td><td>7</td><td>45</td><td>9</td><td>54</td><td>9</td><td>63</td><td>21</td></tr><tr><td>18</td><td>6</td><td>27</td><td>9</td><td>36</td><td>6</td><td>46</td><td>23</td><td>55</td><td>11</td><td>64</td><td>8</td></tr></table>	D	-gs	D	-gs	D	-gs	D	-gs	D	-gs	D	-gs	8	4	20	5	28	7	38	19	48	8	56	8	10	5	21	7	30	6	39	13	49	7	57	19	12	6	22	11	32	8	40	8	50	10	58	29	14	7	24	6	33	11	42	7	51	17	60	10	15	5	25	5	34	17	44	11	52	13	62	31	16	8	26	13	35	7	45	9	54	9	63	21	18	6	27	9	36	6	46	23	55	11	64	8															
D	-gs	D	-gs	D	-gs	D	-gs	D	-gs	D	-gs																																																																																																						
8	4	20	5	28	7	38	19	48	8	56	8																																																																																																						
10	5	21	7	30	6	39	13	49	7	57	19																																																																																																						
12	6	22	11	32	8	40	8	50	10	58	29																																																																																																						
14	7	24	6	33	11	42	7	51	17	60	10																																																																																																						
15	5	25	5	34	17	44	11	52	13	62	31																																																																																																						
16	8	26	13	35	7	45	9	54	9	63	21																																																																																																						
18	6	27	9	36	6	46	23	55	11	64	8																																																																																																						
-a	--affinity	<p>Only for DC systems.</p> <p>ERA RAID affinity: first node (0) or second node (1).</p> <p>The default: the current node.</p>																																																																																																															

Method	Command Syntax	Functionality
	-bs	--block_size Drives block size in bytes. Possible values: 512, 4096 . The default: 4096 .
	-dt	--drive_type Type of a drive. Possible values: <ul style="list-style-type: none"> • slow; • middle; • fast. By default, slow – HDD and SanOp, middle – SAS SSD and SATA SSD, fast – NVMe SSD.
	-ip	--init_prio ERA RAID Initialization priority: from 0 to 100 . The default: 5 .
	-ml	--memory_limit RAM limit for ERA RAID, in MiB: from 1024 to 1048576 . By default, the system calculates this value depending on Stripe size.
	-me	--merge_enabled Enable (1) or disable (0) SCSI commands merge. Improves sequential I/O. The default: 0 .
	-mm	--merge_max Maximal request buffering time for the ERA RAID in milliseconds. Possible values: integers from 0 to 100000 . The default: 1000 .
	-mw	--merge_wait Request buffering time for the ERA RAID in milliseconds. Possible values: integers from 0 to 100000 . The value must be less than --merge_max. The default: 300 .
	-rd	--redundant_drives The number of redundant drives for ERA RAID N+M. The default: 4 .

Method	Command Syntax	Functionality		
		-rp	--recon_prio	ERA RAID Reconstruction priority: 0 to 100 . The default: 5 .
		-rl	--req_limit	Limitation on the number of requests for ERA RAID. The value 0 disables the limitation. Possible values: integers from 0 to 2147483647 . The default: 0 .
		-rp	--restripe_prio	Priority of the restripe function for the ERA RAID. Possible values: integers from 0 to 100 . The default: 100 .
		-se	--sched_enabled	Enable (1) or disable (0) optimization for low-concurrent sequential I/O. The default: 0 .
		-sp	--sparepool	Name of an existing SparePool.
		-ss	--stripe_size	ERA RAID Stripe size, in KiB: 16, 32, 64, 128, or 256 . The default: 16 .
delete	rdcli eraraid delete	The command deletes selected ERA RAID (including from configuration file).		
		Required parameter		
		-n	--name	ERA RAID name.
		Optional parameter (does not require an argument)		
		-f	--force	Force command execution.
init show	rdcli eraraid init show	The command shows ERA RAID initialization status.		
		Required parameter		
		-n	--name	ERA RAID name.
migrate	rdcli eraraid migrate	The command changes Affinity of the specified ERA RAID (performs migration of RAID on the other node).		

Method	Command Syntax	Functionality		
modify	rdcli eraraid modify	Required parameter		
		-n	--name	ERA RAID name.
		The command changes ERA RAID parameters.		
		Required parameter		
		-n	--name	ERA RAID name.
		Optional parameters		
		-ip	--init_prio	ERA RAID Initialization priority: from 0 to 100 . The default: 5 .
		-ml	--memory_limit	ERA RAID RAM limit, in MiB: from 1024 to 1048576 . By default, the system calculates this value depending on the Stripe size.
		-me	--merge_enabled	Enable (1) or disable (0) SCSI commands merge. Improves sequential I/O. The default: 0 .
		-mm	--merge_max	Maximal request buffering time for the ERA RAID in milliseconds. Possible values: integers from 0 to 100000 . The default: 1000 .
-mw	--merge_wait	Request buffering time for the ERA RAID in milliseconds. Possible values: integers from 0 to 100000 . The value must be less than --merge_max . The default: 300 .		
-rp	--recon_prio	ERA RAID Reconstruction priority: from 0 to 100 . The default: 5 .		
-rl	--req_limit	Limitation of the number of requests for the ERA RAID. The value 0 disables the limitation. Possible values: integers from 0 to 2147483647 .		

Method	Command Syntax	Functionality		
				The default: 0 .
		-rp	--restripe_prio	Priority of the restripe function for the ERA RAID. Possible values: integers from 0 to 100 . The default: 100 .
		-se	--sched_enabled	Enable (1) or disable (0) optimization for low-concurrent sequential I/O. The default: 0 .
		-sp	--sparepool	Name of an existing SparePool.
replace	rdcli eraraid replace	The command replaces the ERA RAID drive with specified UID (attribute <code>-ou</code>) with an unused by others RAID's drive.		
		Required parameters		
		-n	--name	ERA RAID name.
		-ou	--old	Removed (failed) drive UID.
		Optional parameter		
		-nu	--new	Added (fault-free) drive UID. If the added drive is not specified, the system automatically selects the first suitable drive.
resize	rdcli eraraid resize	The command changes the size of the ERA RAID.		
		Required parameter		
		-n	--name	ERA RAID name.
restripe continue	rdcli eraraid restripe continue	Command continues stopped restriping.		
		Required parameter		
		-n	--name	ERA RAID name.
restripe show	rdcli eraraid restripe show	Command shows the status of restriping.		
		Required parameter		
		-n	--name	ERA RAID name.

Method	Command Syntax	Functionality												
restripe start	rdcli eraraid restripe start	<p>Command starts restriping of the RAID.</p> <p>Required parameters</p> <table> <tr> <td>-d</td><td>--drives</td><td>List of drive UIDs to be added to the RAID. Example: 1,3,5-7.</td></tr> <tr> <td>-n</td><td>--name</td><td>ERA RAID name.</td></tr> </table> <p>Optional parameters</p> <table> <tr> <td>-l</td><td>--level</td><td> The new RAID level. Possible values: 0, 1, 10, 5, 6, 7.3. By default: keeps the current level. The available options for RAID level changes and the minimum required number of drives see in the RAID Restriping chapter. </td></tr> <tr> <td>-f</td><td>--force</td><td>Force command execution.</td></tr> </table>	-d	--drives	List of drive UIDs to be added to the RAID. Example: 1,3,5-7.	-n	--name	ERA RAID name.	-l	--level	The new RAID level. Possible values: 0, 1, 10, 5, 6, 7.3. By default: keeps the current level. The available options for RAID level changes and the minimum required number of drives see in the RAID Restriping chapter.	-f	--force	Force command execution.
-d	--drives	List of drive UIDs to be added to the RAID. Example: 1,3,5-7.												
-n	--name	ERA RAID name.												
-l	--level	The new RAID level. Possible values: 0, 1, 10, 5, 6, 7.3. By default: keeps the current level. The available options for RAID level changes and the minimum required number of drives see in the RAID Restriping chapter.												
-f	--force	Force command execution.												
restripe stop	rdcli eraraid restripe stop	<p>Command pauses restriping of the ERA RAID.</p> <p>Required parameter</p> <table> <tr> <td>-n</td><td>--name</td><td>ERA RAID name.</td></tr> </table>	-n	--name	ERA RAID name.									
-n	--name	ERA RAID name.												
show	rdcli eraraid show	<p>The command shows all existing ERA RAID.</p> <p>Optional parameters</p> <table> <tr> <td>-n</td><td>--name</td><td>ERA RAID name.</td></tr> <tr> <td>-s</td><td>--raid_status</td><td>Show ERA RAID with specified status.</td></tr> </table>	-n	--name	ERA RAID name.	-s	--raid_status	Show ERA RAID with specified status.						
-n	--name	ERA RAID name.												
-s	--raid_status	Show ERA RAID with specified status.												

fc

Method	Command Syntax	Functionality
target show	rdcli fc target show	Command displays the list of Fibre Channel targets used in the system.
target modify	rdcli fc target modify	Command allows associating the appropriate Fibre Channel targets on the local and the remote nodes.
Required parameters		
-n	--target_name	Name of the local target.
-r	--remote_target	Name of the remote target.
Optional attribute (does not require an argument)		
-f	--force	Force command execution.
celerity show	rdcli fc celerity show	Command displays parameters of celerity driver.
initiator show	rdcli fc initiator show	Command displays information on FC initiator.
Optional parameter		
-n	--name	FC initiator name.
initiator scan	rdcli fc initiator scan	Command scans drives to discover FC initiators.
Required parameter		
-n	--name	FC initiator name.
initiator target show	rdcli fc initiator target show	Command displays information on FC targets.
Optional parameters		
-n	--name	FC target name.
-i	--initiator	Initiator ID.
initiator target create	rdcli fc initiator target create	Command creates an FC target.
Required parameters		
-i	--initiator	Initiator ID.

Method	Command Syntax	Functionality
		-t --target FC target name.
		Optional parameters
		-d --drive List of IDs of block devices allowed for use through the target.
		-n --name Target name for the target identifier.
initiator target delete	rdcli fc initiator target delete	Command deletes an FC target.
		Required parameter
		-n --name Target ID.
		Optional attribute (does not require an argument)
		-f --force If not specified, then provided RAIDs and LUNs exist on the related drives, an error will occur. If specified, such RAIDs and LUNs will become unavailable.
initiator target modify	rdcli fc initiator target modify	Command changes properties of FC hosts.
		Required parameter
		-n --name Target ID.
		Optional parameters
		-t --target List of target names.
		-i --initiator ID of the initiator to which the target reassigned.
		-d --drive List of IDs of drives allowed to be used through the target.
		-r --rel_tgt_id ID of the related target on the remote node in DC.
		Optional attribute (does not require an argument)
		-f --force If not specified, then provided RAIDs and LUNs exist on the related drives, an error will occur. If specified, such RAIDs and LUNs will become unavailable.
	rdcli fc initiator target scan	Command scan drives to discover FC targets.

Method	Command Syntax	Functionality	
initiator target scan		Required parameter	
		<table> <tr> <td>-n</td><td>--name</td><td>Target ID.</td></tr> </table>	-n
-n	--name	Target ID.	
qlogic show	rdcli fc qlogic show	Command displays names of changeable parameters of qlogic driver and their current values.	
qlogic modify	rdcli fc qlogic modify	Command modifies Qlogic FC parameters.	
		Optional parameter	
		<table> <tr> <td>-i</td><td>--qlini_mode</td><td> Initiator's mode: <ul style="list-style-type: none"> • 0 – disabled; • 1 – enabled; • 2 – exclusive; • 3 – dual. The default: 0. </td></tr> </table>	-i
-i	--qlini_mode	Initiator's mode: <ul style="list-style-type: none"> • 0 – disabled; • 1 – enabled; • 2 – exclusive; • 3 – dual. The default: 0.	

h0st

Method	Command Syntax	Functionality
create	rdcli host create	Command assigns the alias name to one or several hosts.
		Required parameters
		-n --name Host alias.
		-l --initiators Comma-separated list of hosts to be merged into one alias name.
		Optional parameters
		-rt --realtime Assign the <i>Realtime</i> priority to the created alias.
		-qs --qosmic Enable (1) or disable (0) QoSmic.
delete	rdcli host delete	Command deletes the specified host alias.
		Required parameter
		-n --name Host alias.
modify	rdcli host modify	Command allows modifying the list of hosts merged into the alias.
		Required parameter
		-n --name Host alias.
		Optional parameters
		-l --initiators List of comma-separated IQN/WWN/GUID of hosts.
		-r --newname New host alias.
		-rt --realtime Assign the <i>Realtime</i> priority to the created alias.
		-qs --qosmic Enable (1) or disable (0) QoSmic.
show	rdcli host show	Command shows all host aliases in the system.
		Optional parameter
		-n --name Display the information on the specified host alias.


ib

Method	Command Syntax	Functionality									
target show	rdcli ib target show	Command shows the list of targets used in the system.									
target modify	rdcli ib target modify	<p>Command allows associating the appropriate InfiniBand targets on the local and the remote nodes.</p> <p>Required parameters</p> <table> <tr> <td>-n</td><td>--target_name</td><td>Name of the local target.</td></tr> <tr> <td>-r</td><td>--remote_target</td><td>Name of the remote target.</td></tr> </table> <p>Optional attribute (does not require an argument)</p> <table> <tr> <td>-f</td><td>--force</td><td>Force command execution.</td></tr> </table>	-n	--target_name	Name of the local target.	-r	--remote_target	Name of the remote target.	-f	--force	Force command execution.
-n	--target_name	Name of the local target.									
-r	--remote_target	Name of the remote target.									
-f	--force	Force command execution.									
opensm start	rdcli ib opensm start	Command starts OpenSM service.									
opensm stop	rdcli ib opensm stop	Command stops OpenSM service.									
opensm add	rdcli ib opensm add	<p>Command allows adding a port GUID to start OpenSM.</p> <p>Required parameter</p> <table> <tr> <td>-g</td><td>--guid</td><td>Port GUID for OpenSM service.</td></tr> </table> <p>Optional attribute (does not require an argument)</p> <table> <tr> <td>-f</td><td>--force</td><td>Force command execution and restart OpenSM service if it already started.</td></tr> </table>	-g	--guid	Port GUID for OpenSM service.	-f	--force	Force command execution and restart OpenSM service if it already started.			
-g	--guid	Port GUID for OpenSM service.									
-f	--force	Force command execution and restart OpenSM service if it already started.									
opensm del	rdcli ib opensm del	<p>Command deletes a port GUID not needed for OpenSM service.</p> <p>Required parameter</p> <table> <tr> <td>-g</td><td>--guid</td><td>Port GUID.</td></tr> </table> <p>Optional attribute (does not require an argument)</p> <table> <tr> <td>-f</td><td>--force</td><td>Force command execution and restart OpenSM service if have already started.</td></tr> </table>	-g	--guid	Port GUID.	-f	--force	Force command execution and restart OpenSM service if have already started.			
-g	--guid	Port GUID.									
-f	--force	Force command execution and restart OpenSM service if have already started.									

Method	Command Syntax	Functionality
opensm show	rdcli ib opensm show	Command displays OpenSM service state: active (1), stopped (0). If active, the command also shows a list of ports, that GUID uses for OpenSM service operation.

h2>iscsi

Method	Command Syntax	Functionality																		
modify	rdcli iscsi modify	<p>Command manages iSCSI support.</p> <p>Optional parameters</p> <table> <tr> <td>-e</td><td>--enable</td><td>Enable (1) or disable (0) iSCSI support. The default: 0.</td></tr> <tr> <td>-p</td><td>--port</td><td>Port for iSCSI connection on the local node. The default: 3260.</td></tr> <tr> <td>-a</td><td>--address</td><td>IP address of the interface to discover an iSCSI target on the local node.</td></tr> <tr> <td>-i</td><td>--iser</td><td>Enable (1) or disable (0) iSER support for targets.</td></tr> <tr> <td colspan="3">Optional attribute (does not require an argument)</td></tr> <tr> <td>-f</td><td>--force</td><td>Force command execution.</td></tr> </table>	-e	--enable	Enable (1) or disable (0) iSCSI support. The default: 0.	-p	--port	Port for iSCSI connection on the local node. The default: 3260 .	-a	--address	IP address of the interface to discover an iSCSI target on the local node.	-i	--iser	Enable (1) or disable (0) iSER support for targets.	Optional attribute (does not require an argument)			-f	--force	Force command execution.
-e	--enable	Enable (1) or disable (0) iSCSI support. The default: 0.																		
-p	--port	Port for iSCSI connection on the local node. The default: 3260 .																		
-a	--address	IP address of the interface to discover an iSCSI target on the local node.																		
-i	--iser	Enable (1) or disable (0) iSER support for targets.																		
Optional attribute (does not require an argument)																				
-f	--force	Force command execution.																		
show	rdcli iscsi show	Command shows the iSCSI state.																		
target create	rdcli iscsi target create	<p>Command creates an iSCSI target with the specified name and authentication type.</p> <p>Optional parameters</p> <table> <tr> <td>-m</td><td>--mode</td><td> Authentication type: <ul style="list-style-type: none"> • bidir – bidirectional; • unidir – unidirectional; • none – no authentication. The default: none. </td></tr> <tr> <td>-n</td><td>--name</td><td>Full iSCSI target name.</td></tr> <tr> <td>-l</td><td>--login</td><td>Login for bidirectional authentication.</td></tr> <tr> <td>-p</td><td>--password</td><td>Password for bidirectional authentication.</td></tr> <tr> <td>-ib</td><td>--ib_only</td><td>Limit (1) or don't limit (0) accessibility of the target by only InfiniBand interfaces available in the system.</td></tr> </table>	-m	--mode	Authentication type: <ul style="list-style-type: none"> • bidir – bidirectional; • unidir – unidirectional; • none – no authentication. The default: none .	-n	--name	Full iSCSI target name.	-l	--login	Login for bidirectional authentication.	-p	--password	Password for bidirectional authentication.	-ib	--ib_only	Limit (1) or don't limit (0) accessibility of the target by only InfiniBand interfaces available in the system.			
-m	--mode	Authentication type: <ul style="list-style-type: none"> • bidir – bidirectional; • unidir – unidirectional; • none – no authentication. The default: none .																		
-n	--name	Full iSCSI target name.																		
-l	--login	Login for bidirectional authentication.																		
-p	--password	Password for bidirectional authentication.																		
-ib	--ib_only	Limit (1) or don't limit (0) accessibility of the target by only InfiniBand interfaces available in the system.																		

Method	Command Syntax	Functionality															
target delete	rdcli iscsi target delete	<p>Command deletes iSCSI targets.</p> <p>Required parameter</p> <table> <tr> <td>-n</td><td>--name</td><td>ISCSI target name.</td></tr> </table>	-n	--name	ISCSI target name.												
-n	--name	ISCSI target name.															
target modify	rdcli iscsi target modify	<p>Command changes parameters of iSCSI targets.</p> <p>Required parameter</p> <table> <tr> <td>-n</td><td>--name</td><td>ISCSI target name.</td></tr> </table> <p>Optional parameters</p> <table> <tr> <td>-m</td><td>--mode</td><td>Authentication type.</td></tr> <tr> <td>-l</td><td>--login</td><td>Login for bidirectional authentication.</td></tr> <tr> <td>-p</td><td>--password</td><td>Password for bidirectional authentication.</td></tr> <tr> <td>-ib</td><td>--ib_only</td><td>Limit (1) or don't limit (0) accessibility of the target by only InfiniBand interfaces available in the system.</td></tr> </table> <div>  Advanced iSCSI target parameters are described in the section Set Up the Advanced iSCSI target parameters. </div>	-n	--name	ISCSI target name.	-m	--mode	Authentication type.	-l	--login	Login for bidirectional authentication.	-p	--password	Password for bidirectional authentication.	-ib	--ib_only	Limit (1) or don't limit (0) accessibility of the target by only InfiniBand interfaces available in the system.
-n	--name	ISCSI target name.															
-m	--mode	Authentication type.															
-l	--login	Login for bidirectional authentication.															
-p	--password	Password for bidirectional authentication.															
-ib	--ib_only	Limit (1) or don't limit (0) accessibility of the target by only InfiniBand interfaces available in the system.															
target show	rdcli iscsi target show	<p>Command displays information on iSCSI targets.</p> <p>Optional parameter</p> <table> <tr> <td>-n</td><td>--name</td><td>ISCSI target name.</td></tr> </table>	-n	--name	ISCSI target name.												
-n	--name	ISCSI target name.															
initiator show	rdcli iscsi initiator show	<p>Command displays information on iSCSI hosts.</p> <p>Optional parameter</p> <table> <tr> <td>-n</td><td>--name</td><td>Initiator name.</td></tr> </table>	-n	--name	Initiator name.												
-n	--name	Initiator name.															
initiator create	rdcli iscsi initiator create	<p>Command creates an iSCSI initiator with specified parameters.</p> <p>Required parameter</p>															

Method	Command Syntax	Functionality	
		-a	--ipaddr IP address of the remote node for iSCSI initiator.
		Optional parameters	
		-p	--port Port for iSCSI on the remote node: from 0 to 65535 . The default: 3260 .
		-sn	--session_number Number of sessions for the initiator: from 1 to 8 . The default: 2 .
		-lt	--login_timeout Login timeout, in seconds: from 3 to 120 . The default: 5 .
		-n	--name Initiator name.
initiator delete	rdcli iscsi initiator delete	Command deletes iSCSI hosts.	
		Required parameter	
		-n	--name Host name.
		Optional attribute (does not require an argument)	
		-f	--force While deleting a target, all targets with drive masks related to the target will be deleted as well. If the related drives contain RAID and LUN, they will become unavailable. If the parameter is not specified and the related drives contain RAID and/or LUN, an error occurs.
initiator scan	rdcli iscsi initiator scan	Command scans hosts to discover remote iSCSI targets.	
		At least one parameter is required	
		-n	--name Host name.
		-a	--ipaddr IP address of the remote node for iSCSI initiator.
		Optional parameter	

Method	Command Syntax	Functionality		
		-p	--port	Port for iSCSI on the remote node. The default: 3260 .
initiator target show	rdcli iscsi initiator target show	Command displays information on remote iSCSI hosts.		
		Optional parameters		
		-n	--name	Name of the remote iSCSI target.
		-i	--initiator	Initiator name.
initiator target create	rdcli iscsi initiator target create	Command creates a remote iSCSI target.		
		Required parameters		
		-i	--initiator	Initiator name.
		-t	--target	Target name.
		Optional parameters		
		-d	--drive	List of ID of block devices allowed for use through the target. By default, the list is empty.
		-n	--name	Remote iSCSI target name.
initiator target delete	rdcli iscsi initiator target delete	Command deletes iSCSI targets.		
		Required parameter		
		-n	--name	Remote iSCSI target name.
		Optional parameter		
		-f	--force	While deleting a target, all targets with drive masks related to the target will be deleted as well. If the related drives contain RAID and LUN, they will become unavailable. If the parameter is not specified and the related drives contain RAID and/or LUN, an error occurs.
	rdcli iscsi initiator target modify	Command changes parameters of iSCSI targets.		


Method	Command Syntax	Functionality
initiator target modify		Required parameters
		-n --name Remote iSCSI target name.
		Optional parameters
		-d --drive List of drive serial numbers allowed for use through the target. List is comma-separated with no space. To add or remove a serial number, use + or - characters at the beginning of the appropriate serial number.
		-i --initiator Initiator name.
		-t --target Target name.
		-r --rel_tgt_id Target name on the remote node.
		Optional attribute (does not require an argument)
		-f --force While deleting a target, all targets with drive masks related to the target will be deleted as well. If the related drives contain RAID and LUN, they will become unavailable. If the parameter is not specified and the related drives contain RAID and/or LUN, an error occurs.
initiator target scan	rdcli iscsi initiator target scan	Command scans remote iSCSI targets.
		Required parameter
		-n --name Remote iSCSI target name.
chap create	rdcli iscsi chap create	Command creates a CHAP user with specified login and password.
		Required parameters
		-l --login CHAP login.
		-p --password CHAP password.
chap delete	rdcli iscsi chap delete	Command deletes CHAP users.
		Required parameter
		-l --login CHAP login.

Method	Command Syntax	Functionality				
chap modify	rdcli iscsi chap modify	Command changes CHAP user's login and password.				
		Required parameters				
		<table> <tr> <td>-l</td><td>--login</td><td>CHAP login.</td></tr> <tr> <td>-p</td><td>--password</td><td>CHAP password.</td></tr> </table>	-l	--login	CHAP login.	-p
-l	--login	CHAP login.				
-p	--password	CHAP password.				
chap show	rdcli iscsi chap show	Command displays information on CHAP users.				
		Optional parameter				
		<table> <tr> <td>-l</td><td>--login</td><td>CHAP login.</td></tr> </table>	-l	--login	CHAP login.	
-l	--login	CHAP login.				
iprule create	rdcli iscsi iprule create	Command creates an allow rule for the specified target with the specified hostname and IP address.				
		Required parameters				
		<table> <tr> <td>-tn</td><td>--target</td><td>Target name.</td></tr> <tr> <td>-a</td><td>--ip</td><td>Host IP address.</td></tr> </table>	-tn	--target	Target name.	-a
-tn	--target	Target name.				
-a	--ip	Host IP address.				
iprule delete	rdcli iscsi iprule delete	Command deletes allow rules for the specified target with the specified hostname and IP.				
		Required parameters				
		<table> <tr> <td>-tn</td><td>--target</td><td>Target name.</td></tr> <tr> <td>-a</td><td>--ip</td><td>Host IP address.</td></tr> </table>	-tn	--target	Target name.	-a
-tn	--target	Target name.				
-a	--ip	Host IP address.				
iprule show	rdcli iscsi iprule show	Command displays the list of all IP rules.				
		Optional parameter				
		<table> <tr> <td>-tn</td><td>--target</td><td>Target name.</td></tr> </table>	-tn	--target	Target name.	
-tn	--target	Target name.				

h1>lun

Method	Command Syntax	Functionality																		
create	rdcli lun create	<p>Command creates a LUN with the specified name on the specified RAID.</p> <p>Required parameters</p> <table> <tr> <td>-n</td><td>--name</td><td> LUN name. Possible values: any letters, numbers, and underscore, from 1 to 15 characters. </td></tr> </table> <p>Mutually exclusive required parameters</p> <table> <tr> <td>-d</td><td>--drive</td><td> For SAN Optimizer. Drive UID. Incompatible with <code>--size</code>. </td></tr> <tr> <td>-r</td><td>--raid</td><td> One or several RAID names (comma separated without spaces) on which the LUN will be created. If the LUN size is smaller than the RAID size, the LUN will be created on the first RAID specified and will be incremented by the next RAID specified as the size increases. </td></tr> </table> <p>Optional parameters</p> <table> <tr> <td>-bs</td><td>--block_size</td><td> Block size, in B: 512, 4096. The default: <ul style="list-style-type: none"> for Generic RAID – 512; for ERA RAID – 4096. </td></tr> <tr> <td>-ms</td><td>--max_sectors_kb</td><td> Size of the read/write request, in KB: from 4 to 32768. The default: 1024. </td></tr> <tr> <td>-pr</td><td>--pr_sync</td><td> Enable (1) or disable (0) synchronization of Persistent Reservations for this LUN. The default: 0. </td></tr> </table>	-n	--name	LUN name. Possible values: any letters, numbers, and underscore, from 1 to 15 characters.	-d	--drive	For SAN Optimizer. Drive UID. Incompatible with <code>--size</code> .	-r	--raid	One or several RAID names (comma separated without spaces) on which the LUN will be created. If the LUN size is smaller than the RAID size, the LUN will be created on the first RAID specified and will be incremented by the next RAID specified as the size increases.	-bs	--block_size	Block size, in B: 512, 4096 . The default: <ul style="list-style-type: none"> for Generic RAID – 512; for ERA RAID – 4096. 	-ms	--max_sectors_kb	Size of the read/write request, in KB: from 4 to 32768 . The default: 1024 .	-pr	--pr_sync	Enable (1) or disable (0) synchronization of Persistent Reservations for this LUN. The default: 0 .
-n	--name	LUN name. Possible values: any letters, numbers, and underscore, from 1 to 15 characters.																		
-d	--drive	For SAN Optimizer. Drive UID. Incompatible with <code>--size</code> .																		
-r	--raid	One or several RAID names (comma separated without spaces) on which the LUN will be created. If the LUN size is smaller than the RAID size, the LUN will be created on the first RAID specified and will be incremented by the next RAID specified as the size increases.																		
-bs	--block_size	Block size, in B: 512, 4096 . The default: <ul style="list-style-type: none"> for Generic RAID – 512; for ERA RAID – 4096. 																		
-ms	--max_sectors_kb	Size of the read/write request, in KB: from 4 to 32768 . The default: 1024 .																		
-pr	--pr_sync	Enable (1) or disable (0) synchronization of Persistent Reservations for this LUN. The default: 0 .																		

Method	Command Syntax	Functionality		
		-p	--prod_id	Product identifier from the INQUIRY response. The default: the LUN name.
		-pln	--proxy_lun_no	Proxy LUN number. By default, the specified number or the number of the LUN on the target is assigned. If not specified or already taken, the first available number is assigned.
		-q	--queue_alg_flag	Enable (1) or disable (0) hosts to change type of command queue. The default: 1.
		-s	--size	LUN size, GiB. The default: all available space. Incompatible with --drive.
		-sm	--sync_mount	Enable (1) or disable (0) synchronous I/O operations on the underlying filesystem. The default for Generic RAID is 0, for ERA RAID is 1.
		-t	--type	Area of use of the LUN. Possible values: scsi , nas , nvme . The default: scsi .
		delete	rdcli lun delete	Command deletes the specified LUN.
Required parameter				
		-n	--name	LUN name.
Optional attributes (do not require an argument)				
		-f	--force	Force command execution.
		-fw	--force_wait	Wait for flush of SSD cache before deleting a LUN.

Method	Command Syntax	Functionality		
		-ff	--force_fast	Quick LUN deletion. If you use this attribute while deleting a LUN, using SSD cache, cache will not be flushed so the deletion will take less time. But data integrity is not guaranteed in this case because the deleted data can't be restored.
extend	rdcli lun extend	Command increases LUN size or extends a LUN on several RAIDs.		
		Required parameter		
		-n	--name	LUN name.
		Mutually exclusive required parameters		
		-s	--size	New LUN size. Must be larger than the current LUN size.
		-ad	--addraid	List of RAIDs to be added to the LUN.
modify	rdcli lun modify	Command changes parameters of the specified LUN.		
		Required parameter		
		-n	--name	LUN name.
		Optional parameters		
		-fs	--fs_type	Format the LUN with the specified file system: xfs , ext4 . <div> To format a LUN to EXT4 file system, the LUN size must not exceed 260096 GiB.</div>
		-ms	--max_sectors_kb	Size of the read/write request, in KiB: from 4 to 32768 . The default: 1024 .
		-pr	--pr_sync	Enable (1) or disable (0) synchronization of Persistent Reservations for this LUN. The default: 0 .

Method	Command Syntax	Functionality		
		-p	--prod_id	Product identifier from the INQUIRY response. The default: the LUN name.
		-q	--queue_alg_flag	Enable (1) or disable (0) hosts to change type of command queue. The default: 1.
		-ssd	--ssdcache	Enable SSD cache for the LUN: <ul style="list-style-type: none">• yes – enable;• no – disable.
		-sm	--sync_mount	Enable (1) or disable (0) synchronous I/O on the underlying filesystem. The default for Generic RAID is 0, for ERA RAID is 1.
		-t	--type	Area of use of the LUN. Possible values: scsi , nas , nvme . The default: scsi .
		Optional attribute (does not require an argument)		
		-f	--force	Force command execution.
		-fw	--force_wait	Wait for flush of SSD cache before deleting a LUN.
		-ff	--force_fast	Quick LUN deletion. If you use this attribute while deleting a LUN, using SSD cache, cache will not be flushed so the deletion will take less time. But data integrity is not guaranteed in this case because the deleted data can't be restored.
		show	rdcli lun show	Command shows information on all LUNs.
Optional parameters				
		-n	--name	LUN name.
		-r	--raid	RAID name.

Method	Command Syntax	Functionality		
plugin show	rdcli lun plugin show	Optional attribute (does not require an argument)		
		-f	--free	Show only LUNs with no target masks.
		Command shows parameters of the volume related to LUN.		
		Required parameter		
		-n	--name	LUN name.
plugin modify	rdcli lun plugin modify	Optional parameter		
		-p	--plugin	Plugin name.
		Command allows modification of parameters of the volume related to LUN.		
		Required parameters		
		-n	--name	LUN name.
plugin modify	rdcli lun plugin modify	-p	--plugin	Plugin whose parameters to be modified.
		-s	--set_params	List of parameters to be modified.
		Modifiable parameters of random read (range_read plugin)		
		--read_ahead_at_once	General size of read ahead request, in MiB: from 1 to 2048. The default: 24.	
		--read_ahead_distance	General size from the end of the interval, at which read ahead is performed, in MiB: from 1 to 2048. The default: 64.	
plugin modify	rdcli lun plugin modify	--read_ahead_past	Minimum number of stripes in cache behind required for read ahead: from 1 to 2048. The default: 64.	
		--range_timeout	Maximum lifetime of the interval, in seconds: from 1 to 2048. The default: 10.	

Method	Command Syntax	Functionality
	--max_read_ahead_at_once	Maximum size of each read ahead request, in MiB: from 1 to 2048 . The default: 24 .
	--max_read_ahead_distance	Maximum size from the end of the interval, at which read ahead is performed, in MiB: from 1 to 2048 . The default: 512 .
	--min_request_length	Minimum request size that can be recognized as part of a sequential pattern, in sectors: from 1 to 2048 . The default: 1 .
Modifiable parameters of random write (range_write plugin)		
	--min_request_length	Minimum request size that can be recognized as part of a sequential pattern, in sectors: from 1 to 2048 . The default: 1 .
	--min_seq_len	Minimum size of sequential write interval, in MiB: from 1 to 2048 . The default: 64 .
	--range_timeout	Maximum lifetime of the interval, in seconds: from 1 to 2048 . The default: 10 .

mask

Method	Command Syntax	Functionality		
target create	rdcli mask target create	Command creates a masking rule for a LUN and a target.		
		Required parameters		
		-l	--lunname	Name of the LUN, for which masking rule is created.
		-tn	--targetname	Target, for which masking rule is created.
target delete	rdcli mask target delete	Command deletes a masking rule for targets.		
		Required parameters		
		-l	--lunname	LUN name. Must be used with --targetname. The parameter is mutually excluded with --id.
		-tn	--targetname	Target name. Must be used with --lunname. The parameter is mutually excluded with --id.
		-i	--id	Masking rule ID. The parameter is mutually excluded with --lunname and --targetname.
target show	rdcli mask target show	Command shows the list of masking rules for targets.		
		Optional parameters		
		-l	--lunname	LUN name.
		-tn	--targetname	Target name.
		-i	--id	Masking rule ID.
host create	rdcli mask host create	Command creates a masking rule for a host.		
		Required parameters		
		-l	--lunname	LUN name.

Method	Command Syntax	Functionality		
		-hn	--hostname	Host name.
		-p	--permissions	Access permissions: <ul style="list-style-type: none">• r – read only;• w – read/write.
modify	rdcli mask host modify	Commands changes parameters of a masking rule for a host.		
Required parameters				
		-l	--lunname	LUN name.
		-hn	--hostname	Host name.
		-p	--permissions	Access permissions: <ul style="list-style-type: none">• r – read only;• w – read/write.
delete	rdcli mask host delete	Command deletes a masking rule for a host.		
Required either lunname and hostname or id				
		-l	--lunname	LUN name.
		-hn	--hostname	Host name.
		-i	--id	Masking rule ID.
show	rdcli mask host show	Command shows the list of masking rules for hosts.		
Optional parameters				
		-l	--lunname	LUN name.
		-hn	--hostname	Host name.
		-i	--id	Masking rule ID.

metadata

Method	Command Syntax	Functionality												
raid scan	rdcli metadata raid scan	<p>Command searches for RAIDS on all free drives.</p> <p>Optional parameter</p> <table> <tr> <td>-d</td><td>--drives</td><td>Drives to be scanned.</td></tr> </table>	-d	--drives	Drives to be scanned.									
-d	--drives	Drives to be scanned.												
raid rename	rdcli metadata raid rename	<p>Command renames imported RAID or volumes in case in case of name conflict.</p> <p>Required parameter</p> <table> <tr> <td>-n</td><td>--name</td><td>New name.</td></tr> </table> <p>Optional parameters</p> <table> <tr> <td>-d</td><td>--drives</td><td>Drives to be renamed.</td></tr> <tr> <td>-on</td><td>--old_name</td><td>RAID old name in metadata.</td></tr> <tr> <td>-cs</td><td>--cache_size</td><td> RAID cache size in MiB. Possible values: from 1024 to 2³²-1. The default: taken from the current RAID cache. </td></tr> </table>	-n	--name	New name.	-d	--drives	Drives to be renamed.	-on	--old_name	RAID old name in metadata.	-cs	--cache_size	RAID cache size in MiB. Possible values: from 1024 to 2³²-1 . The default: taken from the current RAID cache.
-n	--name	New name.												
-d	--drives	Drives to be renamed.												
-on	--old_name	RAID old name in metadata.												
-cs	--cache_size	RAID cache size in MiB. Possible values: from 1024 to 2³²-1 . The default: taken from the current RAID cache.												
raid import	rdcli metadata raid import	<p>Command imports a RAID into the current configuration.</p> <p>Required parameter</p> <table> <tr> <td>-r</td><td>--raid</td><td>RAID name.</td></tr> </table> <p>Optional parameter</p> <table> <tr> <td>-cs</td><td>--cache_size</td><td> RAID cache size in MiB. Possible values: from 1024 to 2³²-1. The default: taken from the current RAID cache. </td></tr> </table>	-r	--raid	RAID name.	-cs	--cache_size	RAID cache size in MiB. Possible values: from 1024 to 2³²-1 . The default: taken from the current RAID cache.						
-r	--raid	RAID name.												
-cs	--cache_size	RAID cache size in MiB. Possible values: from 1024 to 2³²-1 . The default: taken from the current RAID cache.												
volume scan	rdcli metadata volume scan	<p>Command shows list of volumes for the specified RAID.</p> <p>Required parameter</p> <table> <tr> <td>-r</td><td>--raid</td><td>RAID name.</td></tr> </table>	-r	--raid	RAID name.									
-r	--raid	RAID name.												

Method	Command Syntax	Functionality		
volume rename	rdcli metadata volume rename	Optional parameter		
		-cs	--cache_size	RAID cache size in MiB. Possible values: from 1024 to 2³²-1 . The default: taken from the current RAID cache.
		Command renames volumes of the specified RAID.		
		Required parameters		
		-r	--raid	RAID name.
		-n	--name	New name of the volume.
		Optional parameter		
		-cs	--cache_size	RAID cache size in MiB. Possible values: from 1024 to 2³²-1 . The default: taken from the current RAID cache.
drive show	rdcli metadata drive show	Command displays information on drives with metadata.		
		Optional parameter		
		-d	--drives	Drive UID.
drive clean	rdcli metadata drive clean	Command cleans drives with metadata.		
		Optional parameter		
		-d	--drives	Drive UID.

nas

Method	Command Syntax	Functionality		
user create	rdcli nas user create	Command creates a NAS user with specified name and password.		
		Required parameters		
		-n	--name	User login.
		-p	--password	User password.
		Optional parameters		
		-g	--groups	The list of groups (comma separated without spaces) for the user. By default, the user is not a member of any group other than the primary group.
		-pg	--primary_group	Primary group for the user. By default, the primary group is created with the same name as the user name.
		-id	--uid	User ID. Possible values: integers from 1000 to 4294967294 .
		Optional attribute (does not require an argument)		
		-f	--force	Recreate the existing user with the specified name.
user modify	rdcli nas user modify	Command changes NAS user parameters.		
		Required parameter		
		-n	--name	User login.
		Optional parameters		
		-g	--groups	The list of groups (comma separated without spaces) for the user. To add or remove a group, use the "+" or "-" symbols in front of the group name.
		-p	--password	User password.

Method	Command Syntax	Functionality									
		<table> <tr> <td>-pg</td><td>--primary_group</td><td>New primary group for the user.</td></tr> </table>	-pg	--primary_group	New primary group for the user.						
-pg	--primary_group	New primary group for the user.									
user delete	rdcli nas user delete	<p>Command deletes the specified NAS user.</p> <p>Required parameter</p> <table> <tr> <td>-n</td><td>--name</td><td>User login.</td></tr> </table>	-n	--name	User login.						
-n	--name	User login.									
user show	rdcli nas user show	<p>Command displays list of NAS users.</p> <p>Optional parameters</p> <table> <tr> <td>-n</td><td>--name</td><td>User login.</td></tr> <tr> <td>-t</td><td>--type</td><td> <p>Shows groups of the specified type:</p> <ul style="list-style-type: none"> • ldap – LDAP; • ads – Active Directory; • local. <p>The default: for all the types.</p> </td></tr> </table> <p>Optional attributes (do not require an argument)</p> <table> <tr> <td>-si</td><td>--show_id</td><td>UIDs of NAS users.</td></tr> </table>	-n	--name	User login.	-t	--type	<p>Shows groups of the specified type:</p> <ul style="list-style-type: none"> • ldap – LDAP; • ads – Active Directory; • local. <p>The default: for all the types.</p>	-si	--show_id	UIDs of NAS users.
-n	--name	User login.									
-t	--type	<p>Shows groups of the specified type:</p> <ul style="list-style-type: none"> • ldap – LDAP; • ads – Active Directory; • local. <p>The default: for all the types.</p>									
-si	--show_id	UIDs of NAS users.									
group create	rdcli nas group create	<p>Command creates a NAS user group.</p> <p>Required parameter</p> <table> <tr> <td>-n</td><td>--name</td><td>User login.</td></tr> </table> <p>Optional parameter</p> <table> <tr> <td>-id</td><td>--gid</td><td> <p>Group ID.</p> <p>Possible values: integers from 1000 to 4294967294.</p> </td></tr> </table> <p>Optional attribute (does not require an argument)</p> <table> <tr> <td>-f</td><td>--force</td><td>Recreate the existing group with the specified name.</td></tr> </table>	-n	--name	User login.	-id	--gid	<p>Group ID.</p> <p>Possible values: integers from 1000 to 4294967294.</p>	-f	--force	Recreate the existing group with the specified name.
-n	--name	User login.									
-id	--gid	<p>Group ID.</p> <p>Possible values: integers from 1000 to 4294967294.</p>									
-f	--force	Recreate the existing group with the specified name.									
group show	rdcli nas group show	<p>The command displays list of NAS user groups.</p> <p>Optional parameters</p>									

Method	Command Syntax	Functionality		
		-n	--name	Group name.
		-t	--type	Show groups for the specified type: <ul style="list-style-type: none"> • ldap – LDAP; • ads – Active Directory; • local. The default: for all the types.
		Optional attributes (do not require an argument)		
		-si	--show_id	Display group IDs.
group delete	rdcli nas group delete	Command deletes a NAS user group.		
		Required parameter		
		-n	--name	Group name.
		Optional parameter		
		-l	--lun	LUN name.
share show	rdcli nas share show	Command displays list of all shares.		
share nfs create	rdcli nas share nfs create	Command creates an NFS share on the specified LUN.		
		Required parameters		
		-n	--name	Share name.
		-l	--lun	LUN name.
		Optional parameters		
		-a	--async	Async (1) or sync (0) write mode for the share. The default: 0.
		-p	--path	Path to the share on the LUN.
		-s	--squash	Enable (1) or disable (0) Squash option. The default: 1.

Method	Command Syntax	Functionality		
		-rs	--root_squash	Enable (1) or disable (0) Squash option for a root user. The default: 0.
		-rh	--ro_hosts	The list of hosts (comma separated without spaces) with Read Only permissions.
		-wh	--rw_hosts	The list of hosts (comma separated without spaces) with Read/Write permissions.
		-sg	--squash_gid	GID of the anonymous group used for Squash. Possible values: from 1 to 4294967294. The default: 65534.
		-su	--squash_uid	UID of the anonymous user for Squash. Possible values: from 1 to 4294967294. The default: 65534.
share nfs modify	rdcli nas share nfs modify	Command changes NFS share parameters.		
		Required parameter		
		-n	--name	Share name.
		Optional parameters		
		-a	--async	Async (1) or sync (0) write mode for the share. The default: 0.
		-s	--squash	Enable (1) or disable (0) Squash option. The default: 1.
		-rs	--root_squash	Enable (1) or disable (0) Squash option for the root user. The default: 0.
		-rh	--ro_hosts	The list of hosts (comma separated without spaces) with Read Only permissions. To add or remove a host, use "+" or "-" before its name/address. The default: the list is empty.

Method	Command Syntax	Functionality		
		-wh	--rw_hosts	The list of hosts (comma separated without spaces) with Read/Write permissions. To add or remove a host, use "+" or "-" before its name/address. The default: the list is empty.
		-sg	--squash_gid	GID of the anonymous group used for Squash. Possible values: from 1 to 4294967294. The default: 65534.
		-su	--squash_uid	UID of the anonymous user for Squash. Possible values: from 1 to 4294967294. The default: 65534.
share nfs delete	rdcli nas share nfs delete	Command deletes an NFS share.		
		Required parameter		
		-n	--name	Share name.
		Optional attribute (does not require an argument)		
		-f	--force	Force command execution and restart NFS service. If this is the last share on LUN, this LUN should be unmounted, which requires restarting the NFS service.
share nfs show	rdcli nas share nfs show	Command displays list of all NFS shares.		
		Optional parameters		
		-n	--name	Share name.
		-l	--lun	LUN name.
share afp create	rdcli nas share afp create	Command creates an AFP share on a specified LUN.		
		Required parameters		
		-n	--name	Share name.
		-l	--lun	LUN name.

Method	Command Syntax	Functionality
Optional parameters		
-p	--path	Path to the share on the LUN.
-g	--guest	Guest access type: <ul style="list-style-type: none"> • ro – read only; • rw – read/write; • no – guest access is disabled. The default: no .
-ru	--ro_users	The list of users (comma separated without spaces) with Read Only permissions.
-wu	--rw_users	The list of users (comma separated without spaces) with Read/Write permissions.
-rg	--ro_groups	The list of user groups (comma separated without spaces) with Read Only permissions.
-wg	--rw_groups	The list of user groups (comma separated without spaces) with Read/Write permissions.
-ah	--allow_hosts	The list of hosts (comma separated without spaces) with access to the share.
-dh	--deny_hosts	The list of hosts (comma separated without spaces) with no access to the share.
share afp modify	rdcli nas share afp modify	Command changes AFP share parameters.
Required parameter		
-n	--name	Share name.
Optional parameters		

Method	Command Syntax	Functionality
	-g	--guest Guest access type: <ul style="list-style-type: none"> • ro – read only; • rw – read/write; • no – guest access is disabled. The default: no .
	-ru	--ro_users The list of users (comma separated without spaces) with Read Only permissions. To add or remove a user, use "+" or "-" before its name. The default: the list is empty.
	-wu	--rw_users The list of users (comma separated without spaces) with Read/Write permissions. To add or remove a user, use "+" or "-" before its name. The default: the list is empty.
	-rg	--ro_groups The list of user groups (comma separated without spaces) with Read Only permissions. To add or remove a group, use "+" or "-" before its name. The default: the list is empty.
	-wg	--rw_groups The list of user groups (comma separated without spaces) with Read/Write permissions. To add or remove a group, use "+" or "-" before its name. The default: the list is empty.
	-ah	--allow_hosts The list of hosts (comma separated without spaces) with access to the share. To add or remove a host, use "+" or "-" before its name/address. The default: the list is empty.

Method	Command Syntax	Functionality		
		-dh	--deny_hosts	The list of hosts (comma separated without spaces) with no access to the share. To add or remove a host, use "+" or "-" before its name/address. The default: the list is empty.
share afp delete	rdcli nas share afp delete	Command deletes an AFP share.		
		Required parameter		
		-n	--name	Share name.
		Optional attribute (does not require an argument)		
		-f	--force	Force command execution and restart AFP service. If this is the last share on LUN, this LUN should be unmounted, which requires restarting the AFP service.
share afp show	rdcli nas share afp show	Command displays list of all AFP shares.		
		Optional parameters		
		-n	--name	Share name.
		-l	--lun	LUN name.
share ftp create	rdcli nas share ftp create	Command creates an FTP share on a specified LUN.		
		Required parameters		
		-n	--name	Share name.
		-l	--lun	LUN name.
		Optional parameters		
		-p	--path	Path to the share on the LUN.

Method	Command Syntax	Functionality	
		-g	--guest Guest access type: <ul style="list-style-type: none"> • ro – read only; • rw – read/write; • no – guest access is disabled. The default: no .
		-ru	--ro_users The list of users (comma separated without spaces) with Read Only permissions.
		-wu	--rw_users The list of users (comma separated without spaces) with Read/Write permissions.
		-rg	--ro_groups The list of user groups (comma separated without spaces) with Read Only permissions.
		-wg	--rw_groups The list of user groups (comma separated without spaces) with Read/Write permissions.
share ftp modify	rdcli nas share ftp modify	Command changes FTP share parameters.	
		Required parameter	
		-n	--name Share name.
		Optional parameters	
		-g	--guest Guest access type: <ul style="list-style-type: none"> • ro – read only; • rw – read/write; • off – guest access is disabled. The default: no .
		-ru	--ro_users The list of users (comma separated without spaces) with Read Only permissions. To add or remove a user, use "+" or "-" before its name. The default: the list is empty.

Method	Command Syntax	Functionality		
		-wu	--rw_users	The list of users (comma separated without spaces) with Read/Write permissions. To add or remove a user, use "+" or "-" before its name. The default: the list is empty.
		-rg	--ro_groups	The list of user groups (comma separated without spaces) with Read Only permissions. To add or remove a group, use "+" or "-" before its name. The default: the list is empty.
		-wg	--rw_groups	The list of user groups (comma separated without spaces) with Read/Write permissions. To add or remove a group, use "+" or "-" before its name. The default: the list is empty.
share ftp delete	rdcli nas share ftp delete	Command deletes an FTP share.		
		Required parameter		
		-n	--name	Share name.
		Optional attribute (does not require an argument)		
		-f	--force	Force command execution and restart FTP service. If this is the last share on LUN, this LUN should be unmounted, which requires restarting the FTP service.
share ftp show	rdcli nas share ftp show	Command displays list of all FTP shares.		
		Optional parameters		
		-n	--name	Share name.
		-l	--lun	LUN name.
share samba create	rdcli nas share samba create	Command creates an SMB share on a specified LUN.		
		Required parameters		

Method	Command Syntax	Functionality	
	-n	--name	Share name.
	-l	--lun	LUN name.
	Optional parameters		
	-p	--path	Path to the share on the LUN.
	-g	--guest	Guest access type: <ul style="list-style-type: none">• ro – read only;• rw – read/write;• no – guest access is disabled. The default: no .
	-ru	--ro_users	The list of users (comma separated without spaces) with Read Only permissions.
	-wu	--rw_users	The list of users (comma separated without spaces) with Read/Write permissions.
	-rg	--ro_groups	The list of user groups (comma separated without spaces) with Read Only permissions.
	-wg	--rw_groups	The list of user groups (comma separated without spaces) with Read/Write permissions.
	-ah	--allow_hosts	The list of hosts (comma separated without spaces) with access to the share.
	-dh	--deny_hosts	The list of hosts (comma separated without spaces) with no access to the share.
	-a	--admins	The list of root directory administrators (comma separated without spaces).


Method	Command Syntax	Functionality	
		-b	--browsable <p>Visibility of SMB share in the list of available network resources:</p> <ul style="list-style-type: none"> 0 – share is not visible: client cannot see the share in the list of resources on the server, but will still be able to access it by specifying the path to the share with \$ character at the end; 1 – share is visible. <p>The default: 1.</p>
		-worm	--worm <p>Enable (1) or disable (0) the WORM (Write Once Read Many) option.</p> <p>The default: 0.</p>
		-gp	--grace_period <p>Period of time, during which the created file is available for modification/deletion, in seconds.</p> <p>The default: 1 (minimum possible value).</p>
		-rp	--ro_period <p>Time frame, during which the share is not available for modification while WORM parameter is enabled, in months.</p> <p>The default: 60 (5 years).</p>
		-c	--comment <p>Comment to the share.</p>
		-cs	--case_sensitive <p>Case sensitivity for the share filename.</p> <p>Possible values:</p> <ul style="list-style-type: none"> yes – case sensitive; no – not case sensitive; auto – use Samba software settings. <p>The default: auto.</p>
share samba modify	rdcli nas share samba modify	Command changes SMB share parameters. Required parameter	
		-n	--name <p>Share name.</p>
		Optional parameters	



Method	Command Syntax	Functionality
	-a	--admins The list of root directory administrators (comma separated without spaces). To add or remove an administrator, use "+" or "-" before its name. The default: the list is empty.
	-ah	--allow_hosts The list of hosts (comma separated without spaces) with access to the share. To add or remove a host, use "+" or "-" before its name/address. The default: the list is empty.
	-b	--browsable Visibility of SMB share in the list of available network resources: <ul style="list-style-type: none"> • 0 – share is not visible. A client cannot see the share in the list of resources on the server, but will still be able to access it by specifying the path to the share with \$ character at the end; • 1 – share is visible. The default: 1 .
	-dh	--deny_hosts The list of hosts (comma separated without spaces) with no access to the share To add or remove a host, use "+" or "-" before its name/address. The default: the list is empty.
	-g	--guest Guest access type: <ul style="list-style-type: none"> • ro – read only; • rw – read/write; • off – guest access is disabled. The default: no .
	-gp	--grace_period Period of time, during which the created file is available for modification/deletion, in seconds. The default: 1 (minimum possible value).

Method	Command Syntax	Functionality
	-ru	--ro_users
		The list of users (comma separated without spaces) with Read Only permissions. To add or remove a user, use "+" or "-" before its name. The default: the list is empty.
	-wu	--rw_users
		The list of users (comma separated without spaces) with Read/Write permissions. To add or remove a user, use "+" or "-" before its name. The default: the list is empty.
	-rg	--ro_groups
		The list of user groups (comma separated without spaces) with Read Only permissions. To add or remove a group, use "+" or "-" before its name. The default: the list is empty.
	-wg	--rw_groups
		The list of user groups (comma separated without spaces) with Read/Write permissions. To add or remove a group, use "+" or "-" before its name. The default: the list is empty.
	-rp	--ro_period
		Period of time, during which the share is not available for modification while WORM parameter is enabled, in months. The default: 60 (5 years).
	-worm	--worm
		Enable (1) or disable (0) WORM (Write Once Read Many) option. The default: 0 .
	-c	--comment
		Comment to the share.

Method	Command Syntax	Functionality		
		-cs	--case_sensitive	Case sensitivity for the share filename. Possible values: <ul style="list-style-type: none"> yes – case sensitive; no – not case sensitive; auto – use Samba software settings. The default: auto .
share samba delete	rdcli nas share samba delete	Command deletes an SMB share.		
		Required parameter		
		-n	--name	Share name.
		Optional attribute (does not require an argument)		
		-f	--force	Force command execution and restart SMB service. If this is the last share on LUN, this LUN should be unmounted, which requires restarting the SMB service.
share samba show	rdcli nas share samba show	Command returns list of all SMB shares.		
		Optional parameters		
		-n	--name	Share name.
		-l	--lun	LUN name.
samba show	rdcli nas samba show	Command displays NAS samba object attributes. The interface provides the ability to configure the Windows workgroup and provides the ability to connect SMB (Samba) to Active Directory.		
samba modify	rdcli nas samba modify	Command allows connecting (disconnecting) Samba to (from) Active Directory and changing settings of Microsoft Windows workgroup.		
		Optional parameters		

Method	Command Syntax	Functionality
	-a	--ads Change Active Directory connection state: <ul style="list-style-type: none"> try to connect to Active Directory (1); connection will not be disabled, but winbind service will be stopped and information about connection will be removed from samba and kerberos configuration files (0).
	-w	--workgroup <ul style="list-style-type: none"> Name of Windows workgroup (if --ads = 0); Active Directory domain (if --ads = 1).
	-r	--realm Domain controller name. <ul style="list-style-type: none"> If --ads = 1, this attribute is mandatory; if --ads = 0, this attribute is not used.
	-u	--user Name of user with permission to connect the workstation to the domain. <ul style="list-style-type: none"> If --ads = 1, this attribute is mandatory; if --ads = 0, this attribute is not used.
	-t	--trusts The list (comma separated without spaces) of Active Directory trusted domains. To change the list of the domains, specify the new list. If the uid_range parameter is not specified, when specifying the list of trusted servers, pay attention to the automatic assignment of UID ranges.
	-ur	--uid_range The comma-separated list of domains with their UID ranges. The domains are taken from the trusts parameter value. Minimal range value: 10000000, maximal range value: 19999999. The default range value: 100000. Format: domain=min-max .


Method	Command Syntax	Functionality	
		-p	<p>--password</p> <p>User password.</p> <ul style="list-style-type: none"> If <code>--ads = 1</code>, this attribute is mandatory; To prevent the entered password from remaining in the command history, enter the command without this option. The system will ask you to enter the password on a new line. if <code>--ads = 0</code>, this attribute is not used.
		Optional attribute (does not require an argument)	
		-f	<p>--force</p> <p>Force command execution.</p>
ldap show	rdcli nas ldap show	Command displays current LDAP settings.	
ldap modify	rdcli nas ldap modify	Command changes LDAP settings.	
		Optional parameters	
		-bdn	<p>--bind_dn</p> <p>Path to LDAP, where the information for authorization during execution of queries to server is placed.</p>
		-e	<p>--enable</p> <p>Enable (1) or disable (0) LDAP.</p>
		-p	<p>--password</p> <p>Password for LDAP authorization, used in combination with bind_dn.</p>
		-ma	<p>--max_uid</p> <p>Allow LDAP users whose UID is greater than this parameter.</p>
		-mi	<p>--min_uid</p> <p>Allow LDAP users whose UID is lower than this parameter.</p>
		-s	<p>--server</p> <p>List of LDAP servers. The servers must be from the same cluster</p> <div>  Before adding an LDAP server, check its availability. </div>
		-sb	<p>--search_base</p> <p>Path to LDAP, from which all requests start.</p>
nfs show	rdcli nas nfs show	Command displays NFS parameters.	
		Displayed parameters	
		nfsdcount	Number of NFS threads.

Method	Command Syntax	Functionality		
		packets_arrived	Number of NFS packets received by the server.	
		rdma	The status of the RDMA support function.	
		sockets_enqueued	Number of requests sent to the queue until the thread becomes free.	
		threads_timedout	Number of excesses of the NFS idle time (60 minutes). Under the constant load, a non-zero value might indicate an excessive number of NFS treads for the current load.	
		threads_woken	Number of «awakenings» of an idle NFS thread to process a request. The value characterizes the request processing rate. In optimal configuration, the value is approaching the value of packets_arrived.	
		<div> Restart of nfsd service leads to reset of all counters.</div>		
nfs modify	rdcli nas nfs modify	Command changes the number of NFS treads.		
Optional parameters				
		-c	--nfsdcount	Number of NFS treads. Possible values: from 1, not limited. The default: 8. <div> The NFS server will restart after changing the value of nfsdcount parameter.</div>
		-r	--rdma	Enable (1) or disable (0) the RDMA support for all NFS shares in the system. The default: 0.
quota show	rdcli nas quota show	Command displays information on quotas.		
Optional parameters				
		-u	--users	Display quotas for specified users.

Method	Command Syntax	Functionality		
quota create	rdcli nas quota create	-g	--groups	Display quotas for specified groups.
		-l	--lun	Display quotas for the specified LUN.
		Command creates a quota for the specified LUN.		
		Mutually exclusive required parameters		
		-g	--groups	Groups for which quota is set.
		-u	--users	Users for which quota is set.
		Required parameters		
		-s	--size	Quota size, in MiB.
		-l	--lun	LUN name.
		Optional parameters		
		-gp	--grace_period	Grace period (in minutes) that starts after overrunning the quota, after which the writing process stops. The default: 15 minutes.
quota modify	rdcli nas quota modify	Command allows editing quota parameters.		
		Required parameter		
		-l	--lun	LUN name.
		Optional parameters		
		-s	--size	Quota size, in MiB.
		-u	--users	Users for which quota is set.
		-g	--groups	Groups for which quota is set.
		-gp	--grace_period	Grace period (in minutes) that starts after overrunning the quota, after which the writing process stops. The default: 15 minutes.
quota refresh	rdcli nas quota refresh	Command refreshes information on quotas.		

Method	Command Syntax	Functionality
quota delete	rdcli nas quota delete	Optional parameter
		-l --lun LUN name.
		Command deletes quotas on the specified LUN.
		Required parameter
		-l --lun LUN name.
		Optional parameters
		-u --users Delete quotas for the specified users.
		-g --groups Delete quotas for the specified groups.
		-gp --grace_period Grace period (in minutes) that starts after overrunning the quota, after which the writing process stops. The default: 15 minutes.


network

Method	Command Syntax	Functionality																		
show	rdcli network show	Command displays general networking settings: <ul style="list-style-type: none"> • DNS; • gateway; • hostname; • primary interface. 																		
modify	rdcli network modify	Command changes general network parameters. <div>  If the IP address of the main interface has been assigned via DHCP, it is prohibited to change DNS and Gateway parameters when performing this command. </div> <div> Optional parameters <table> <tr> <td>-hn</td><td>--hostname</td><td>Host name.</td></tr> <tr> <td>-if</td><td>--primary_interface</td><td>Primary interface.</td></tr> <tr> <td>-dns1</td><td>--dns1</td><td>Primary DNS.</td></tr> <tr> <td>-dns2</td><td>--dns2</td><td>Secondary DNS.</td></tr> <tr> <td>-dns3</td><td>--dns3</td><td>Tertiary DNS.</td></tr> <tr> <td>-gw</td><td>--gateway</td><td>Gateway.</td></tr> </table> </div>	-hn	--hostname	Host name.	-if	--primary_interface	Primary interface.	-dns1	--dns1	Primary DNS.	-dns2	--dns2	Secondary DNS.	-dns3	--dns3	Tertiary DNS.	-gw	--gateway	Gateway.
-hn	--hostname	Host name.																		
-if	--primary_interface	Primary interface.																		
-dns1	--dns1	Primary DNS.																		
-dns2	--dns2	Secondary DNS.																		
-dns3	--dns3	Tertiary DNS.																		
-gw	--gateway	Gateway.																		
interface show	rdcli network interface show	Command displays network interface parameters. <div> Optional parameters <table> <tr> <td>-n</td><td>--name</td><td>Interface name.</td></tr> <tr> <td>-b</td><td>--bonding</td><td>Bonding name.</td></tr> </table> </div>	-n	--name	Interface name.	-b	--bonding	Bonding name.												
-n	--name	Interface name.																		
-b	--bonding	Bonding name.																		
interface modify	rdcli network interface modify	Command changes general network parameters. <div> Required parameter </div>																		

Method	Command Syntax	Functionality
		<div>-n --name Interface name.</div> <div>Optional parameters</div> <div>-a --ipaddr Interface IP address.</div> <div>-b --bridging_opts Only for bridge ports. List of keys of values of bridge parameters for files <i>/sys/class/net/<bridge device>/bridge/<key></i>.</div> <div>-d --dhcp Specify whether the port address is defined by DHCP. Possible values: <ul style="list-style-type: none"> • use DHCP (1); • do not use DHCP (0). </div> <div>-n --netmask Netmask.</div> <div>-M --mtu MTU (maximum transmission unit) value. Possible values: from 68 to 65536. The default: 1500.</div> <div>-p --pause A comma-separated list of parameters without spaces for flow control. Use "+" to add and "-" to delete the parameter. Possible values: <ul style="list-style-type: none"> • rx • tx • autoneg </div> <div>-u --up Port status: enabled (1), disabled (0).</div>
bond create	rdcli network bond create	Command allows combining specified interfaces, defining bonding policy type . <div> <i>i</i> Primary interface cannot participate in bonding. </div>

Method	Command Syntax	Functionality
Required parameters		
-p	--policy	Bonding mode. For InfiniBand, only active-backup is available. Possible values: <ul style="list-style-type: none"> • 0 or rr or balance-rr; • 1 or ab or active-backup; • 2 or xor or balance-xor; • 3 or bc or broadcast; • 4 or lagg or lacp or 802.3ad; • 5 or tlb or balance-tlb; • 6 or alb or balance-alb.
-s	--slaves	Interface list (comma-separated with no space between).
Optional parameters		
-M	--mtu	MTU size. Possible values: from 68 to 65536 . The default: 1500 .
-n	--name	Name of bonding to be created.

-bo	--bonding_opts	<p>Custom bonding options separated with commas without spaces. Specify bonding options and their values like "option=value". Use "+" to add, and "-" to remove bonding options.</p> <p>To learn more, see kernel.org.</p> <ul style="list-style-type: none"> • lp_interval The number of seconds between instances where the bonding driver sends learning packets to each slave's peer switch. Possible values: from 1 to 2³¹-1. The default: 1. • tlb_dynamic_lb Enable dynamic flow shuffling. Possible values: 0 (off), 1 (on). The default: 1. • larp_rate Only for "802.3ad" mode. LACPDU tx rate to request from 802.3ad partner. Possible values: slow, fast. The default is not set. • ad_user_port_key Only for "802.3ad" mode. In an AD system, the port-key has three parts: (00) – Duplex, (01-05) – Speed, (06-15) – User defined. This defines the upper 10 bits of the port key. Possible values: from 0 to 1023. The default: 0. • packets_per_slave Packets to send per slave in RR mode. Possible values: from 1 to 2¹⁶-1. The default: 1.
-----	----------------	--

		<ul style="list-style-type: none"> ad_actor_system In an AD system, this specifies the mac-address for the actor in protocol packet exchanges (LACPDUs). The value cannot be a multicast address. If the all-zeroes MAC is specified, bonding will internally use the MAC of the bond itself. It is preferred to have the local-admin bit set for this mac but driver does not enforce it. If the value is not given then system defaults to using the masters' mac address as actors' system address. The default: 00:00:00:00:00:00. <div>  To return this parameter to its default value, recreate the bond. </div> ad_select Only for "802.3ad" mode. 803.ad aggregation selection logic. Possible values: stable, bandwidth, count. The default: stable. ad_actor_sys_prio In an AD system, this specifies the system priority. Possible values: from 1 to 2¹⁶-1. The default: 2¹⁶-1. num_grat_arp Number of peer notifications to send on failover event. Possible values: 0 (off), 255 (maxval), 1 (default). The default: 1. updelay Delay before considering link up, in milliseconds. Possible values: from 0 (off) to 2³¹-1 (maxval).
--	--	---

		<p>The default: 0.</p> <ul style="list-style-type: none"> arp_all_targets Fail on any/all arp targets timeout. Possible values: any, all. The default: any. miimon Link check interval in milliseconds. Possible values: from 0 (off) to 2³¹-1 (maxval). The default: 0. use_carrier Use netif_carrier_ok (vs MII ioctls) in miimon. Possible values: 0 (off), 1 (on). The default: 1. primary_reselect Reselect primary slave once it comes up. Possible values: always, better, failure. The default: always. fail_over_mac For active-backup, do not set all slaves to the same MAC. Possible values: none, active, follow. The default: none. min_links Minimum number of available links before turning on carrier. Possible values: from 0 (off) to 2³¹-1 (maxval). The default: 0. downdelay Delay before considering link down, in milliseconds. Possible values: from 0 (off) to 2³¹-1 (maxval). The default: 0.
--	--	--

Method	Command Syntax	Functionality
		<ul style="list-style-type: none"> arp_validate Only for "802.3ad", "balance-tlb", "balance-alb". Validate src/dst of ARP probes. Possible values: none, active, backup, all, filter, filter_active, filter_backup. The default: none. arp_interval Only for "802.3ad", "balance-tlb", "balance-alb". Arp interval in milliseconds. Possible values: from 0 (off) to 2³¹-1 (maxval). The default: 0. xmit_hash_policy balance-xor, 802.3ad, and tlb hashing method. Possible values: layer2, layer3+4, layer2+3, encap2+3, encap3+4, vlan+srcmac. The default: layer2. all_slaves_active Keep all frames received on an interface by setting active flag for all slaves. Possible values: 0 (off), 1 (on). The default: 0. resend_igmp Number of IGMP membership reports to send on link failure. Possible values: 0 (off), 255 (maxval), 1 (default). The default: 1.
bond delete	rdcli network bond delete	Command deletes the specified channel bonding.
Required parameter		
	-n	--name
		Bonding name.

Method	Command Syntax	Functionality		
bond show	rdcli network bond show	Command displays all configured channel bondings.		
		Optional parameter		
		-n	--name	Bonding name.
bond modify	rdcli network bond modify	Command modifies channel bonding parameters.		
		Required parameter		
		-n	--name	Bonding name.
		Optional parameters		
		-s	--slaves	The list of network Interfaces (comma separated without spaces). To change the values, set the new list.
-p	--policy	Bonding mode. Possible values: <ul style="list-style-type: none">• 0 or rr or balance-rr;• 1 or ab or active-backup;• 2 or xor or balance-xor;• 3 or bc or broadcast;• 4 or lagg or lacp or 802.3ad;• 5 or tlb or balance-tlb;• 6 or alb or balance-alb.		

-bo	--bonding_opts	<p>Custom bonding options separated with commas without spaces.</p> <p>Specify bonding options and their values like "option=value". Use "+" to add, and "-" to remove bonding options.</p> <p>To learn more, see kernel.org.</p> <ul style="list-style-type: none"> • <code>lp_interval</code> The number of seconds between instances where the bonding driver sends learning packets to each slave's peer switch. Possible values: from 1 to 2³¹-1. The default: 1. • <code>tlb_dynamic_lb</code> Enable dynamic flow shuffling. Possible values: 0 (off), 1 (on). The default: 1. • <code>lacp_rate</code> Only for "802.3ad" mode. LACPDU tx rate to request from 802.3ad partner. Possible values: slow, fast. The default is not set. • <code>ad_user_port_key</code> Only for "802.3ad" mode. In an AD system, the port-key has three parts: (00) – Duplex, (01-05) – Speed, (06-15) – User defined. This defines the upper 10 bits of the port key. Possible values: from 0 to 1023. The default: 0. • <code>packets_per_slave</code> Packets to send per slave in RR mode.
-----	----------------	--

Possible values: from 1 to $2^{16}-1$.

The default: 1.

- ad_actor_system

In an AD system, this specifies the mac-address for the actor in protocol packet exchanges (LACPDUs).

The value cannot be a multicast address. If the all-zeroes MAC is specified, bonding will internally use the MAC of the bond itself.

It is preferred to have the local-admin bit set for this mac but driver does not enforce it.

If the value is not given then system defaults to using the masters' mac address as actors' system address.

The default: 00:00:00:00:00:00.



To return this parameter to its default value, recreate the bond.

- ad_select

Only for "802.3ad" mode.

803.ad aggregation selection logic.

Possible values: **stable**, **bandwidth**, **count**.

The default: **stable**.

- ad_actor_sys_prio

In an AD system, this specifies the system priority.

Possible values: from 1 to $2^{16}-1$.

The default: $2^{16}-1$.

- num_grat_arp

Number of peer notifications to send on failover event.

Possible values: 0 (off), 255 (maxval), 1 (default).

The default: 1.

- updelay

		<p>Delay before considering link up, in milliseconds. Possible values: from 0 (off) to 2³¹-1 (maxval). The default: 0.</p> <ul style="list-style-type: none"> arp_all_targets Fail on any/all arp targets timeout. Possible values: any, all. The default: any. miimon Link check interval in milliseconds. Possible values: from 0 (off) to 2³¹-1 (maxval). The default: 0. use_carrier Use netif_carrier_ok (vs MII ioctls) in miimon. Possible values: 0 (off), 1 (on). The default: 1. primary_reselect Reselect primary slave once it comes up. Possible values: always, better, failure. The default: always. fail_over_mac For active-backup, do not set all slaves to the same MAC. Possible values: none, active, follow. The default: none. min_links Minimum number of available links before turning on carrier. Possible values: from 0 (off) to 2³¹-1 (maxval). The default: 0.
--	--	--

- **downdelay**
Delay before considering link down, in milliseconds.
Possible values: from **0** (off) to **2³¹-1** (maxval).
The default: **0**.
- **arp_validate**
Only for "802.3ad", "balance-tlb", "balance-alb".
Validate src/dst of ARP probes.
Possible values: **none**, **active**, **backup**, **all**, **filter**, **filter_active**, **filter_backup**.
The default: **none**.
- **arp_interval**
Only for "802.3ad", "balance-tlb", "balance-alb".
Arp interval in milliseconds.
Possible values: from **0** (off) to **2³¹-1** (maxval).
The default: **0**.
- **xmit_hash_policy**
balance-xor, 802.3ad, and tlb hashing method.
Possible values: **layer2**, **layer3+4**, **layer2+3**, **encap2+3**, **encap3+4**, **vlan+srcmac**.
The default: **layer2**.
- **all_slaves_active**
Keep all frames received on an interface by setting active flag for all slaves.
Possible values: **0** (off), **1** (on).
The default: **0**.
- **resend_igmp**
Number of IGMP membership reports to send on link failure.
Possible values: **0** (off), **255** (maxval), **1** (default).
The default: **1**.

Method	Command Syntax	Functionality															
vip show	rdcli network vip show	Command displays list of all virtual IP addresses created in the system.															
vip create	rdcli network vip create	<p>Command allows assigning virtual IP address to the specified interface.</p> <p>Required parameters</p> <table> <tr> <td>-i</td><td>--interface</td><td>Network interface.</td></tr> <tr> <td>-a</td><td>--ipaddr</td><td>IP address.</td></tr> </table> <p>Optional parameters</p> <table> <tr> <td>-m</td><td>--netmask</td><td>Netmask.</td></tr> <tr> <td>-s</td><td>--spare</td><td> <p>Node on which virtual IP will be Active in regular system mode:</p> <ul style="list-style-type: none"> Active – this node is primary for this IP (0); Passive – the local node is not primary for this VIP (1). <p>The default: 0.</p> <div> <i>i</i> The value 0 (Active) should be set on the node from which work with shares is performed. </div> </td></tr> <tr> <td>-u</td><td>--up</td><td>Force enable (1) or disable (0) the interface. The default: determined automatically.</td></tr> </table>	-i	--interface	Network interface.	-a	--ipaddr	IP address.	-m	--netmask	Netmask.	-s	--spare	<p>Node on which virtual IP will be Active in regular system mode:</p> <ul style="list-style-type: none"> Active – this node is primary for this IP (0); Passive – the local node is not primary for this VIP (1). <p>The default: 0.</p> <div> <i>i</i> The value 0 (Active) should be set on the node from which work with shares is performed. </div>	-u	--up	Force enable (1) or disable (0) the interface. The default: determined automatically.
-i	--interface	Network interface.															
-a	--ipaddr	IP address.															
-m	--netmask	Netmask.															
-s	--spare	<p>Node on which virtual IP will be Active in regular system mode:</p> <ul style="list-style-type: none"> Active – this node is primary for this IP (0); Passive – the local node is not primary for this VIP (1). <p>The default: 0.</p> <div> <i>i</i> The value 0 (Active) should be set on the node from which work with shares is performed. </div>															
-u	--up	Force enable (1) or disable (0) the interface. The default: determined automatically.															
vip delete	rdcli network vip delete	<p>Command deletes the specified virtual IP address.</p> <p>Required parameter</p> <table> <tr> <td>-n</td><td>--name</td><td>Name of a VIP interface.</td></tr> </table>	-n	--name	Name of a VIP interface.												
-n	--name	Name of a VIP interface.															
vip modify	rdcli network vip modify	<p>Command changes VIP settings.</p> <p>Required parameter</p> <table> <tr> <td>-n</td><td>--name</td><td>Name of the VIP interface.</td></tr> </table> <p>Optional parameters</p> <table> <tr> <td>-a</td><td>--ipaddr</td><td>Interface IP address.</td></tr> </table>	-n	--name	Name of the VIP interface.	-a	--ipaddr	Interface IP address.									
-n	--name	Name of the VIP interface.															
-a	--ipaddr	Interface IP address.															

Method	Command Syntax	Functionality		
		-m	--netmask	Netmask.
		-s	--spare	Node on which virtual IP will be Active in regular system mode: <ul style="list-style-type: none">Active – this node is primary for this IP (0);Passive – the local node is not primary for this VIP (1). The default: 0. <div><div><div>i</div></div>Value 0 (Active) should be set on the node from which work with shares is performed.</div>
		-u	--up	Force enable (1) or disable (0) the interface. The default: determined automatically.
bridge create	rdcli network bridge create	Command creates a network bridge.		
		Optional parameters		
		-a	--ipaddr	Interface IP address. The default: not set.
		-m	--netmask	Netmask. The default: not set.
		-M	--mtu	MTU (maximum transmission unit) value. Possible values: from 68 to 65536. The default: 1500.
		-d	--dhcp	Enable (1) or disable (0) DHCP. The default: 0.
		-u	--up	Enable (1) or disable (0) the network bridge. The default: 1.
		-p	--ports	The list of ports (comma separated without spaces) currently connected to the bridge.

Method	Command Syntax	Functionality
	-S	--stp Enable (1) or disable (0) STP support. The default: 1.
	-D	--delay Time in seconds, during which the bridge listens and examines the network available via port, before adding it. The default: 200.
	-A	--ageing Time in seconds, during which MAC is kept in the database of redirections after receiving the frame. The default: operating system value.
	-P	--prio Bridge priority in STP. This value is used to identify the root bridge in STP and to calculate the lowest cost path. The default: operating system value.

	-b	--bridging_opts	<p>String containing the key=value list (comma separated without spaces), for bridge configuration via sysfs.</p> <p>The default:</p> <ul style="list-style-type: none"> • max_age=20 • hello_time=2 • mcast_hash_elasticity=16 • mcast_startup_query_count=2 • mcast_last_member_interval=1 • mcast_igmp_version=2 • mcast_membership_interval=260 • group_fwd_mask=0 • mcast_query_interval=125 • mcast_last_member_count=2 • vlan_filtering=0 • mcast_query_use_ifaddr=0 • vlan_default_pvid=1 • nf_call_ip6tables=0 • mcast_querier=0 • group_address=0 • mcast_querier_interval=255 • mcast_router=1 • vlan_protocol=802.1Q • mcast_hash_max=512 • mcast_startup_query_interval=3125 • mcast_query_response_interval=10 • mcast_mld_version=1 • mcast_snooping=0 • vlan_stats_enabled=0
--	----	-----------------	---

Method	Command Syntax	Functionality																		
		<ul style="list-style-type: none"> mcast_stats_enabled=0 nf_call_iptables=0 nf_call_arptables=0 To learn more, see https://www.mankier.com/8/ip-link# .																		
bridge modify	rdcli network bridge modify	Command changes network bridge parameters. <div> <div>Required parameter</div> <table> <tr> <td>-n</td><td>--name</td><td>Network bridge ID.</td></tr> </table> <div>Optional parameters</div> <table> <tr> <td>-p</td><td>--ports</td><td>The list of ports (comma separated without spaces) currently connected to the bridge. Use "+" or "-" before a port name to add or remove the port.</td></tr> <tr> <td>-S</td><td>--stp</td><td>Enable (1) or disable (0) STP support. The default: 1.</td></tr> <tr> <td>-D</td><td>--delay</td><td>Period of time, during which the bridge listens and examines the network available via port, before adding it. The default: 200.</td></tr> <tr> <td>-A</td><td>--ageing</td><td>Period of time, during which MAC is kept in the database of redirections after receiving the frame, in seconds. The default: operating system value.</td></tr> <tr> <td>-P</td><td>--prio</td><td>Bridge priority in STP. This value is used to identify the root bridge in STP and to calculate the lowest cost path. The default: operating system value.</td></tr> </table> </div>	-n	--name	Network bridge ID.	-p	--ports	The list of ports (comma separated without spaces) currently connected to the bridge. Use "+" or "-" before a port name to add or remove the port.	-S	--stp	Enable (1) or disable (0) STP support. The default: 1.	-D	--delay	Period of time, during which the bridge listens and examines the network available via port, before adding it. The default: 200.	-A	--ageing	Period of time, during which MAC is kept in the database of redirections after receiving the frame, in seconds. The default: operating system value.	-P	--prio	Bridge priority in STP. This value is used to identify the root bridge in STP and to calculate the lowest cost path. The default: operating system value.
-n	--name	Network bridge ID.																		
-p	--ports	The list of ports (comma separated without spaces) currently connected to the bridge. Use "+" or "-" before a port name to add or remove the port.																		
-S	--stp	Enable (1) or disable (0) STP support. The default: 1.																		
-D	--delay	Period of time, during which the bridge listens and examines the network available via port, before adding it. The default: 200.																		
-A	--ageing	Period of time, during which MAC is kept in the database of redirections after receiving the frame, in seconds. The default: operating system value.																		
-P	--prio	Bridge priority in STP. This value is used to identify the root bridge in STP and to calculate the lowest cost path. The default: operating system value.																		

-b	--bridging_opts	<p>String containing the key=value list (comma separated without spaces) for bridge configuration via sysfs.</p> <p>To add or remove a key-value pair, use "+" or "-" before it.</p> <p>The default (depends on the processor architecture):</p> <ul style="list-style-type: none"> • max_age=20 • hello_time=2 • mcast_hash_elasticity=16 • mcast_startup_query_count=2 • mcast_last_member_interval=1 • mcast_igmp_version=2 • mcast_membership_interval=260 • group_fwd_mask=0 • mcast_query_interval=125 • mcast_last_member_count=2 • vlan_filtering=0 • mcast_query_use_ifaddr=0 • vlan_default_pvid=1 • nf_call_ip6tables=0 • mcast_querier=0 • group_address=0 • mcast_querier_interval=255 • mcast_router=1 • vlan_protocol=802.1Q • mcast_hash_max=512 • mcast_startup_query_interval=3125 • mcast_query_response_interval=1000 • mcast_mld_version=1 • mcast_snooping=0
----	-----------------	--

Method	Command Syntax	Functionality												
		<ul style="list-style-type: none"> vlan_stats_enabled=0 mcast_stats_enabled=0 nf_call_iptables=0 nf_call_arptables=0 To learn more, see https://www.mankier.com/8/ip-link# .												
bridge delete	rdcli network bridge delete	Command deletes a network bridge. Required parameter <table> <tr> <td>-n</td><td>--name</td><td>Network bridge ID.</td></tr> </table> Optional attribute (does not require an argument) <table> <tr> <td>-f</td><td>--force</td><td>Force command execution.</td></tr> </table>	-n	--name	Network bridge ID.	-f	--force	Force command execution.						
-n	--name	Network bridge ID.												
-f	--force	Force command execution.												
bridge show	rdcli network bridge show	Command displays network bridges. Optional parameter <table> <tr> <td>-n</td><td>--name</td><td>Network bridge ID.</td></tr> </table>	-n	--name	Network bridge ID.									
-n	--name	Network bridge ID.												
route ipv4 create	rdcli network route ipv4 create	Command creates a new route. Required parameters <table> <tr> <td>-d</td><td>--destination</td><td>The network address of packet destination.</td></tr> <tr> <td>-di</td><td>--destination_interface</td><td>The network interface to route packets.</td></tr> </table> Optional parameters <table> <tr> <td>-g</td><td>--gateway</td><td>The gateway address to route packets. The default: no gateway.</td></tr> <tr> <td>-pre</td><td>--preferred_source</td><td>The preferred packet source address. The default: no address.</td></tr> </table>	-d	--destination	The network address of packet destination.	-di	--destination_interface	The network interface to route packets.	-g	--gateway	The gateway address to route packets. The default: no gateway.	-pre	--preferred_source	The preferred packet source address. The default: no address.
-d	--destination	The network address of packet destination.												
-di	--destination_interface	The network interface to route packets.												
-g	--gateway	The gateway address to route packets. The default: no gateway.												
-pre	--preferred_source	The preferred packet source address. The default: no address.												

Method	Command Syntax	Functionality		
		-pri	--priority	The priority of the route. Possible values: integers from 0 (highest priority) to 4294967295 . The default: 0 .
route ipv4 delete	rdcli network route ipv4 delete	Command deletes a specified route.		
		Required parameter		
		-i	--id	The network route ID.
route ipv4 show	rdcli network route ipv4 show	Command shows existed routes.		
		Optional parameters		
		-d	--destination	The network address of packet destination.
		-di	--destination_interface	The network interface to route packets.
		-g	--gateway	The gateway address to route packets.
		-i	--id	The network route ID.
		-pre	--preferred_source	The preferred packet source address.
		-pri	--priority	The priority of the route.

nodes

Method	Command Syntax	Functionality									
discover	rdcli nodes discover	<p>Command initiates system scanning to discover new RAIDIX systems in the local network.</p> <p>Optional parameters</p> <table> <tr> <td>-if</td><td>--interface</td><td> <p>Network interface name for sending broadcast discovery packet.</p> <p>The default: the primary interface.</p> </td></tr> <tr> <td>-t</td><td>--timeout</td><td> <p>Time in seconds of waiting for answers from the nodes.</p> <p>Possible values: from 1 to 10.</p> <p>The default: 2.</p> </td></tr> <tr> <td>-ba</td><td>--broadcast_address</td><td> <p>Broadcast address that will be used for the discovery.</p> <p>The default: based on the IP address and the netmask of the network interface.</p> </td></tr> </table>	-if	--interface	<p>Network interface name for sending broadcast discovery packet.</p> <p>The default: the primary interface.</p>	-t	--timeout	<p>Time in seconds of waiting for answers from the nodes.</p> <p>Possible values: from 1 to 10.</p> <p>The default: 2.</p>	-ba	--broadcast_address	<p>Broadcast address that will be used for the discovery.</p> <p>The default: based on the IP address and the netmask of the network interface.</p>
-if	--interface	<p>Network interface name for sending broadcast discovery packet.</p> <p>The default: the primary interface.</p>									
-t	--timeout	<p>Time in seconds of waiting for answers from the nodes.</p> <p>Possible values: from 1 to 10.</p> <p>The default: 2.</p>									
-ba	--broadcast_address	<p>Broadcast address that will be used for the discovery.</p> <p>The default: based on the IP address and the netmask of the network interface.</p>									
show	rdcli nodes show	Command displays information on RAIDIX systems in the local network.									

notify mail

Method	Command Syntax	Functionality																								
profile modify	rdcli notify mail profile modify	<p>Command allows configuring email notification parameters.</p> <p>Parameters to be configured</p> <table> <tr> <td>-a</td><td>--auth</td><td> <p>Authentication type:</p> <ul style="list-style-type: none"> • on; • off; • plain; • scram-sha-1; • cram-md5; • gssapi; • digest-md5; • login; • ntlm. </td></tr> <tr> <td>-e</td><td>--encryption</td><td>Enable (1) or disable (0) TLS encryption.</td></tr> <tr> <td>-p</td><td>--password</td><td>Sender password.</td></tr> <tr> <td>-pr</td><td>--port</td><td>Port.</td></tr> <tr> <td>-s</td><td>--server</td><td>IP address of a notification sender server.</td></tr> <tr> <td>-sn</td><td>--sender</td><td>Sender email.</td></tr> <tr> <td>-u</td><td>--user</td><td>Sender name.</td></tr> <tr> <td>-fm</td><td>--format</td><td>Format of notifications: text or html.</td></tr> </table>	-a	--auth	<p>Authentication type:</p> <ul style="list-style-type: none"> • on; • off; • plain; • scram-sha-1; • cram-md5; • gssapi; • digest-md5; • login; • ntlm. 	-e	--encryption	Enable (1) or disable (0) TLS encryption.	-p	--password	Sender password.	-pr	--port	Port.	-s	--server	IP address of a notification sender server.	-sn	--sender	Sender email.	-u	--user	Sender name.	-fm	--format	Format of notifications: text or html .
-a	--auth	<p>Authentication type:</p> <ul style="list-style-type: none"> • on; • off; • plain; • scram-sha-1; • cram-md5; • gssapi; • digest-md5; • login; • ntlm. 																								
-e	--encryption	Enable (1) or disable (0) TLS encryption.																								
-p	--password	Sender password.																								
-pr	--port	Port.																								
-s	--server	IP address of a notification sender server.																								
-sn	--sender	Sender email.																								
-u	--user	Sender name.																								
-fm	--format	Format of notifications: text or html .																								
profile show	rdcli notify mail profile show	Command displays parameters of email notification server.																								
recipient create	rdcli notify mail recipient create	<p>Command creates an email notifications recipient.</p> <p>Required parameter</p> <table> <tr> <td>-n</td><td>--name</td><td>Email of notification recipient.</td></tr> </table>	-n	--name	Email of notification recipient.																					
-n	--name	Email of notification recipient.																								

Method	Command Syntax	Functionality
Optional parameters		
-a	--all	All system sensors notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
-d	--drive	Drive state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
-es	--enclosure_sensors	Enclosure sensor notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
-i	--iscsi	ISCSI adapter notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
-l	--license	License notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
-ln	--lun	LUN state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.

Method	Command Syntax	Functionality
	-r	--raid RAID state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-s	--sparepool SparePool state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-dc	--dc DC state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-it	--init_target Initiator state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-nk	--network Network notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-fsn	--ftp_share_nas FTP share state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.

Method	Command Syntax	Functionality
	-asn	--afp_share_nas AFP share state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-nsn	--nfs_share_nas NFS share state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-ssn	--samba_share_nas SMB share state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-cs	--controller_sensors Controller sensor state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-g	--guest VM notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-nf	--nvmeof NVMe-oF notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.

Method	Command Syntax	Functionality		
recipient modify	rdcli notify mail recipient modify	-u	--ups	UPS state notification type: <ul style="list-style-type: none">• ok – info notifications;• warning – warning notifications;• error – error notifications.
		-sn	--system_notification	Enable (1) or disable (0) sending of system notifications.
		Command changes email notification parameters.		
		Required parameter		
		-n	--name	Notification recipient email.
		Optional parameters		
		-a	--all	All system sensors notification type: <ul style="list-style-type: none">• ok – info notifications;• warning – warning notifications;• error – error notifications.
		-sn	--system_notification	Enable (1) or disable (0) sending of system notifications.
		-r	--raid	RAID state notification type: <ul style="list-style-type: none">• ok – info notifications;• warning – warning notifications;• error – error notifications.
		-ln	--lun	LUN state notification type: <ul style="list-style-type: none">• ok – info notifications;• warning – warning notifications;• error – error notifications.

Method	Command Syntax	Functionality
	-d	--drive Drive state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-i	--iscsi iSCSI adapter notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-l	--license License notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-s	--sparepool SparePool state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-cs	--controller_sensors Controller sensor state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-es	--enclosure_sensors Enclosure sensor notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.

Method	Command Syntax	Functionality
	-dc	--dc DC state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-it	--init_target Initiator state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-nk	--network Network notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-fsn	--ftp_share_nas FTP share state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-asn	--afp_share_nas AFP share state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.
	-nsn	--nfs_share_nas NFS share state notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; • error – error notifications.

Method	Command Syntax	Functionality		
		-g	--guest	VM notification type: <ul style="list-style-type: none">• ok – info notifications;• warning – warning notifications;• error – error notifications.
		-nf	--nvmeof	NVMe-oF notification type: <ul style="list-style-type: none">• ok – info notifications;• warning – warning notifications;• error – error notifications.
		-u	--ups	UPS state notification type: <ul style="list-style-type: none">• ok – info notifications;• warning – warning notifications;• error – error notifications.
		-ssn	--samba_share_nas	SMB share state notification type: <ul style="list-style-type: none">• ok – info notifications;• warning – warning notifications;• error – error notifications.
recipient delete	rdcli notify mail recipient delete	Command deletes an email notification recipient.		
		Required parameter		
		-n	--name	Notification recipient email.
recipient show	rdcli notify mail recipient show	Command displays list of all email notification recipients.		
		Optional parameter		
		-n	--name	Notification recipient email.
test	rdcli notify mail test	Command sends a test email to the recipients.		
		Optional parameter		
		-r	--recipient	Test letter recipient email.

notify snmp

Method	Command Syntax	Functionality
profile modify	rdcli notify snmp profile modify	Command configures SNMP agent.
Optional parameters		
-e	--enable	Enable (1) or disable (0) SNMP. The default: 0.
-sv	--snmp_version	SNMP protocol version used by SNMP agent. Possible values: 1, 2, or 3. The default: 2.
-a	--address	IP addresses used by SNMP agent.
-p	--port	SNMP agent UDP port. The default: 161.
-c	--community	For SNMP v.1 and 2 only. Community-string – name of community used for authentication. The default: public .
-secn	--security_name	For SNMP v.3 only. Username used for authentication.
-aup	--auth_proto	For SNMP v.3 only. Authentication protocol. Possible values: <ul style="list-style-type: none"> md5; sha. The default: md5 .
-ap	--auth_pass	For SNMP v.3 only. Authentication password.

Method	Command Syntax	Functionality		
		-amk	--auth_master_key	For SNMP v.3 only. Key used for authentication instead of a password.
		-alk	--auth_localized_key	For SNMP v.3 only. Localized key used for authentication instead of a password.
		-prp	--priv_proto	For SNMP v.3 only. Encryption protocol. Possible values: <ul style="list-style-type: none"> • des (set by default); • aes.
		-pp	--priv_pass	For SNMP v.3 only. Password used for encryption.
		-pmk	--priv_master_key	For SNMP v.3 only. Key used for encryption instead of a password.
		-plk	--priv_localized_key	For SNMP v.3 only. Localized key used instead of privacy passphrase.
snmp profile show	rdcli notify snmp profile show	Command displays parameters of SNMP agent.		
snmp recipient create	rdcli notify snmp recipient create	Command creates a recipient of SNMP notifications.		
		Required parameter		
		-hn	--host	Recipient IP address.
		Optional parameters		
		-po	--port	Port for receiving SNMP notifications.
		-ei	--engine_id	EngineID used for identification and encryption.

Method	Command Syntax	Functionality
	-a	--all All system sensors notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-d	--drive Drive state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-es	--enclosure_sensors Enclosure sensor notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-i	--iscsi ISCSI adapter notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-l	--license License notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-ln	--lun LUN state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.

Method	Command Syntax	Functionality
	-r	--raid RAID state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-s	--sparepool SparePool state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-dc	--dc DC state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-it	--init_target Initiator state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-nk	--network Network notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-fsn	--ftp_share_nas FTP share state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.

Method	Command Syntax	Functionality
	-asn	--afp_share_nas AFP share state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-nsn	--nfs_share_nas NFS share state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-ssn	--samba_share_nas SMB share state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-cs	--controller_sensors Controller sensor state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-g	--guest VM notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; error – error notifications.
	-nf	--nvmeof NVMe-oF notification type: <ul style="list-style-type: none"> • ok – info notifications; • warning – warning notifications; error – error notifications.

Method	Command Syntax	Functionality	
recipient modify	rdcli notify snmp recipient modify	-u	--ups UPS state notification type: <ul style="list-style-type: none"> ok – info notifications; warning – warning notifications; error – error notifications.
		-sn	--system_notification Enable (1) or disable (0) sending of system notifications.
		Command modifies parameters of recipient of SNMP notifications.	
		Required parameters	
		-n	--name This parameter is the combination of values <i>host:port</i> . <i>For example:</i> 10.10.1.1:161.
		Optional parameters	
		-hn	--host Recipient IP address.
		-po	--port Port for receiving SNMP notifications.
		-ei	--engine_id EngineID used for identification and encryption.
		-sn	--system_notification Enable (1) or disable (0) sending of system notifications.
		-a	--all All system sensors notification type: <ul style="list-style-type: none"> ok – all notifications; warning – warnings and errors; error – errors only.
		-r	--raid RAID state notification type: <ul style="list-style-type: none"> ok – all notifications; warning – warnings and errors; error – errors only.

Method	Command Syntax	Functionality
	-ln	--lun LUN state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-d	--drive Drive state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-i	--iscsi ISCSI adapter notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-l	--license License notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-s	--sparepool SparePool state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-cs	--controller_sensors Controller sensor state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.

Method	Command Syntax	Functionality
	-es	--enclosure_sensors Enclosure sensor notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-dc	--dc DC state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-it	--init_target Initiator state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-nk	--network Network notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-fsn	--ftp_share_nas FTP share state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.
	-asn	--afp_share_nas AFP share state notification type: <ul style="list-style-type: none"> • ok – all notifications; • warning – warnings and errors; • error – errors only.

Method	Command Syntax	Functionality
	-nsn	--nfs_share_nas NFS share state notification type: <ul style="list-style-type: none"> ok – all notifications; warning – warnings and errors; error – errors only.
	-g	--guest VM notification type: <ul style="list-style-type: none"> ok – info notifications; warning – warning notifications; error – error notifications.
	-nf	--nvmeof NVMe-oF notification type: <ul style="list-style-type: none"> ok – info notifications; warning – warning notifications; error – error notifications.
	-u	--ups UPS state notification type: <ul style="list-style-type: none"> ok – info notifications; warning – warning notifications; error – error notifications.
	-ssn	--samba_share_nas SMB share state notification type: <ul style="list-style-type: none"> ok – all notifications; warning – warnings and errors; error – errors only.
recipient delete	rdcli notify snmp recipient delete	Command deletes a recipient of SNMP notifications. Required parameter
	-n	--name This parameter is the combination of parameters host:port. For example: 10.10.1.1:161.
recipient show	rdcli notify snmp recipient show	Command displays information on recipients of SNMP notifications. Optional parameter

Method	Command Syntax	Functionality		
		-n	--name	This parameter is the combination of parameters host:port. For example: 10.10.1.1:161.
recipient test	rdcli notify snmp recipient test	Command sends a test notification to recipients.		
		Optional parameter		
		-n	--name	This parameter is the combination of parameters host:port. For example: 10.10.1.1:161.

nvme fabrics

Method	Command Syntax	Functionality
param modify	rdcli nvme fabrics param modify	Command changes NVMe-oF global parameters.
Required parameter		
	-l	Interval (in seconds) between Namespace extension and Namespace update. During the update, Namespaces will be disabled. The default: 15 seconds.
param show	rdcli nvme fabrics param show	Command shows global parameters of NVMe-oF.
port create	rdcli nvme fabrics port create	Command creates an NVMe-oF port.
Required parameter		
	-a	NVMe-oF port address.
Optional parameters		
	-f	Address family of the NVMe-oF port. Possible values: ipv4 .
	-in	Inline data size (in bytes) for the NVMe-oF port. For TCP: <ul style="list-style-type: none"> Possible values: integers from 2^{10} to $2^{31}-1$ and -1. The default: -1. For RDMA: <ul style="list-style-type: none"> Possible values: integers from -1 to 2^{14}. The default: 0.
	-n	Name of the NVMe-oF port. The default: <type>:<ip>:<port>.
	-p	Number of the NVMe-oF port. The default: 4420 .

Method	Command Syntax	Functionality		
		-rs	--referrals	The set of referrals (comma separated without spaces) that will be shown on the port.
		-ss	--subsystems	The set of subsystems (comma separated without spaces) that will be available on the port.
		-t	--type	Protocol type of the NVMe-oF port. Possible values: tcp , rdma . The default: tcp .
port delete	rdcli nvme fabrics port delete	Command deletes an NVMe-oF port.		
		Required parameter		
		-n	--name	Name of the NVMe-oF port.
		Optional parameter (does not require an argument)		
port modify	rdcli nvme fabrics port modify	-f	--force	Force the deletion of the NVMe-oF port.
		Command modifies the parameters of an NVMe-oF port.		
		Required parameter		
		-n	--name	Name of the NVMe-oF port.
		Optional parameters		
		-in	--inline_data_size	Inline data size (in bytes) for the NVMe-oF port. For TCP: <ul style="list-style-type: none"> Possible values: integers from 2^{10} to $2^{31}-1$ and -1. The default: -1. For RDMA: <ul style="list-style-type: none"> Possible values: integers from -1 to 2^{14}. The default: 0.

Method	Command Syntax	Functionality		
		-rs	--referrals	The set of referrals (comma separated without spaces) that will be shown on the port. To add or remove a referral, use "+" or "-" before it.
		-ss	--subsystems	The set of subsystems (comma separated without spaces) that will be available on the port. To add or remove a subsystem, use "+" or "-" before its name.
port show	rdcli nvme fabrics port show	Command shows NVMe-oF ports.		
		Optional parameters		
		-a	--address	NVMe-oF port address.
		-f	--family	Address family of the NVMe-oF port. Possible values: ipv4 .
		-n	--name	Name of the NVMe-oF port.
		-p	--port	Number of the NVMe-oF port.
		-t	--type	Protocol type of the NVMe-oF port.
referral create	rdcli nvme fabrics referral create	Command creates an NVMe-oF referral.		
		Required parameter		
		-a	--address	Address of the NVMe-oF referral.
		Optional parameters		
		-f	--family	Address family of the NVMe-oF referral. Possible values: ipv4 , ipv6 , fc .
		-n	--name	Name of the NVMe-oF referral.
		-p	--port	Port of the NVMe-oF referral.
	-s	--security	Connection security type of the NVMe-oF referral. Possible values: not specified , required , not required . The default: not specified .	

Method	Command Syntax	Functionality		
		-t	--type	Protocol type of the NVMe-oF referral. Possible values: rdma , fc , tcp . The default: rdma .
referral delete	rdcli nvme fabrics referral delete	Command deletes an NVMe-oF referral.		
		Required parameter		
		-n	--name	The name of the NVMe-oF referral.
		Optional parameter (does not require an argument)		
		-f	--force	Force deletion of the referral.
referral modify	rdcli nvme fabrics referral modify	Command changes the NVMe-oF referral parameters.		
		Required parameter		
		-n	--name	Name of the NVMe-oF referral.
		Optional parameters		
		-a	--address	Address of the NVMe-oF referral.
		-f	--family	Address family of the NVMe-oF referral. Possible values: ipv4 , ipv6 , fc .
		-p	--port	Port of the NVMe-oF referral.
		-s	--security	Connection security type of the NVMe-oF referral. Possible values: not specified , required , not required . The default: not specified .
		-t	--type	Protocol type of the NVMe-oF referral. Possible values: rdma , fc , tcp . The default: rdma .
referral show	rdcli nvme fabrics referral show	Command shows NVMe-oF referrals.		
		Optional parameters		

Method	Command Syntax	Functionality		
		-a	--address	Address of the NVMe-oF referral.
		-f	--family	Address family of the NVMe-oF referral. Possible values: ipv4 , ipv6 , fc .
		-n	--name	Name of the NVMe-oF referral.
		-p	--port	Port of the NVMe-oF referral.
		-s	--security	Connection security type of the NVMe-oF referral. Possible values: not specified , required , not required . The default: not specified .
		-t	--type	Protocol type of the NVMe-oF referral. Possible values: rdma , fc , tcp . The default: rdma .
subsystem create	rdcli nvme fabrics subsystem create	Command creates an NVMe-oF subsystem.		
		Optional parameters		
		-any	--allow_any_host	Enable (1) or disable (0) the network discovery of this subsystem for all hosts. The default: 1.
		-hs	--hosts	The set of host NQNs (comma separated without spaces) that are allowed to discover the subsystem.
		-n	--name	Name of the NVMe-oF subsystem.
		-ns	--namespaces	The set of namespaces (comma separated without spaces) that will be available on the subsystem.
		-nq	--nqn	NVMe Qualified Name (NQN) of the subsystem.
		-nv	--nvme_version	NVMe protocol version that the subsystem will support. Possible values: 1.0, 1.1, 1.2, 1.2.1, 1.3, 1.4. The default: 1.4.

Method	Command Syntax	Functionality
subsystem delete	rdcli nvme fabrics subsystem delete	Command deletes an NVMe-oF subsystem.
		Required parameter
		-n --name Name of the NVMe-oF subsystem.
		Optional parameter (does not require an argument)
subsystem modify	rdcli nvme fabrics subsystem modify	-f --force Force deletion of the subsystem.
		Command modifies the parameters of an NVMe-oF subsystem.
		Required parameter
		-n --name Name of the NVMe-oF subsystem.
		Optional parameters
		-any --allow_any_host Enable (1) or disable (0) the network discovery of this subsystem for all hosts. The default: 1.
		-hs --hosts The set of host NQNs (comma separated without spaces) that are allowed to discover the subsystem. To add or remove an NQN, use "+" or "-" before its name.
subsystem show	rdcli nvme fabrics subsystem show	-ns --namespaces The set of namespaces (comma separated without spaces) that will be available on the subsystem. To add or remove a namespace, use "+" or "-" before its name.
		-nv --nvme_version NVMe protocol version that the subsystem will support. Possible values: 1.0, 1.1, 1.2, 1.2.1, 1.3, 1.4. The default: 1.4.
		Command shows NVMe-oF subsystems.
		Optional parameters
		-any --allow_any_host Enabled (1) or disabled (0) network discovery of this subsystem for all hosts.

Method	Command Syntax	Functionality	
	-hs	--hosts	The set of host NQNs that are allowed to discover the subsystem.
	-n	--name	Name of the NVMe-oF subsystem.
	-ns	--namespaces	The set of namespaces that will be available on the subsystem.
	-nq	-nqn	NVMe Qualified Name (NQN) of the subsystem.
	-nv	--nvme_version	NVMe protocol version that the subsystem will support. Possible values: 1.0, 1.1, 1.2, 1.2.1, 1.3, 1.4.

param

Method	Command Syntax	Functionality						
plugin show	rdcli param plugin show	Command displays plugin parameters. Optional parameter						
		<table> <tr> <td>-p</td><td>--plugin</td><td>Plugin.</td></tr> </table>	-p	--plugin	Plugin.			
-p	--plugin	Plugin.						
plugin modify	rdcli param plugin modify	Command changes plugin parameters. The parameters to modify are described in rdcli lun plugin modify . Required parameters						
		<table> <tr> <td>-p</td><td>--plugin</td><td>Plugin.</td></tr> <tr> <td>-s</td><td>--set_params</td><td>Parameters to modify.</td></tr> </table>	-p	--plugin	Plugin.	-s	--set_params	Parameters to modify.
-p	--plugin	Plugin.						
-s	--set_params	Parameters to modify.						
raid show	rdcli param raid show	Command displays values of RAID parameters set while creating RAIDs. Optional attribute (does not require an argument)						
		<table> <tr> <td>-d</td><td>--default</td><td>Display default settings.</td></tr> </table>	-d	--default	Display default settings.			
-d	--default	Display default settings.						
raid modify	rdcli param raid modify	Command changes RAID parameters. Optional parameters						
		<table> <tr> <td>-cs</td><td>--cache_size</td><td>Cache size, in MiB: from 1024 to $(2^{32}-1)$. The default: 4096.</td></tr> </table>	-cs	--cache_size	Cache size, in MiB: from 1024 to $(2^{32}-1)$. The default: 4096 .			
-cs	--cache_size	Cache size, in MiB: from 1024 to $(2^{32}-1)$. The default: 4096 .						

Method	Command Syntax	Functionality	
		-ss	--stripe_size Stripe size, in KiB. Possible values: <ul style="list-style-type: none"> • 16; • 32; • 64; • 128; • 256; • 512; • 1024. The default: 128.
		-acc	--alg_compat_code Algorithm code: <ul style="list-style-type: none"> • -2 – no algorithm; • -1 – algorithm is not defined; • from 1 to 6 – other algorithms supported at hardware level. The default: 1.
		Optional attribute (does not require an argument)	
ntp show	rdcli param ntp show	-d	--default Display default settings.
		Command displays parameters of clock synchronization NTP.	
		Optional attribute (does not require an argument)	
ntp modify	rdcli param ntp modify	-f	--full Show full information.
		Command changes parameters of NTP time synchronization.	
		Optional parameters	
		-e	--enable Enable (1) or disable (0) automatic clock synchronization. The default: 1.

Method	Command Syntax	Functionality		
		-s	--servers	The list of servers (comma separated without spaces) to be used for time synchronization. As a server, you can use its IP address or domain name. To add or remove a server, use "+" or "-" before the server name.
		Optional attribute (does not require an argument)		
		-sy	--sync	Synchronize clock forcibly (once).
drive modify	rdcli param drive modify	Changes the parameters of the automatic Drive Health Monitoring.		
		Optional parameters		
		-hm	--health_monitoring	Enable (1) or disable (0) the automatic Drive Health Monitoring. The default: 0.
		-ure	--uncorrected_read_errors	Two comma-separated threshold values for the system notification (warning, error) of disk write errors. The default: 1,1.
		-uve	--uncorrected_verify_errors	Two comma-separated threshold values for the system notification (warning, error) of disk verification data errors. The default: 1,1.
		-uwe	--uncorrected_write_errors	Two comma-separated threshold values for the system notification (warning, error) of disk read errors. The default: 1,1.
		Optional attribute (does not require an argument)		
		-d	--default	Set default values for the parameters.
drive show	rdcli param drive show	Shows the parameters of the automatic Drive Health Monitoring.		
		Optional attribute (does not require an argument)		
		-d	--default	Show default values for the parameters.
driver show	rdcli param driver show	Command displays list of driver parameter values.		

Method	Command Syntax	Functionality
Optional parameter		
	-d --default	Display default settings.
driver modify	rdcli param driver modify	Command changes driver parameters.
Optional parameters		
	-cc --command_saturation	Maximum number of non-real-time commands: from 0 to $(2^{32}-1)$. The default: 4 .
	-ft --flush_threshold	Minimum cache saturation percentage enabling cache flushing: from 0 to 100 . The default: 25 .
	-i --io_timeout	Disk I/O timeout, in ms: from 0 to $(2^{32}-1)$. The default: 28000 .
	-l --log_level	Logging level: from 0 to $(2^{32}-1)$. The default: 256 .
	-rt --request_sync_timeout	Request synchronization timeout, in ms: from 0 to $(2^{32}-1)$. The default: 5000 .
	-rs --request_saturation	Maximum RAID data rate level for non-real-time requests, MB/s: from 0 to $(2^{32}-1)$. The default: 0 .
	-ct --calc_thread_num	Number of calculation threads: from 0 to 256 . The default: 0 (sets automatically depending on the number of cores).
Optional attribute (does not require an argument)		
	-d --default	Apply default settings.

Method	Command Syntax	Functionality
system show	rdcli param system show	Command displays current system settings, namely masking policy, state of delay in RAID status change, threshold of faulty counter on a drive.
system modify	rdcli param system modify	Command changes system parameters.
Modifiable parameters		
-da	--domain_name	The domain name used to form the IQN and NQN. For example, example.com .
-dad	--domain_name_date	The date (in the yyyy-mm format) used to form the IQN. This is the date when the naming authority took ownership of the domain.
-lv	--luns_visible_by_default	Masking policy: <ul style="list-style-type: none"> Enterprise mode – after creation LUNs are not available to any host (0); Standard mode – after creation LUNs are available to all hosts (1). The default: 1 .
-fct	--faulty_cnt_threshold	Threshold number of faults on a drive, at which the drive will be marked as Faulty: from 1 to 1000 . The default: 3 .
-hrbs	--hot_replace_by_smart	Enable (1) or disable (0) automatic replacement of drives from SparePools into RAIDs based on S.M.A.R.T. data. The default: 0 .
-hrt	--hot_replace_timeout	Timeout in seconds before an autoreplacement of a drive from a SparePool into a RAID. Possible values: from 0 to 1200 . The default: 180 .

Method	Command Syntax	Functionality			
		-pap	--pcie_aspm_policy	This may affect performance of PCI Express devices. Modify PCI Express Active-State Power Management policy. Possible values: default , performance , powersave . The default: default .	
ups show	rdcli param ups show	Command displays current UPS settings.			
ups modify	rdcli param ups modify	Command changes settings of the specified UPS driver.			
Optional parameters					
		-d	--driver	Driver name:	
				apcsmart	everups
				microdowell	tripplite
				bcmxcp	gamatronic
				oldmge-shut	tripplitesu
				belkin	genericups
				oneac	upscore2
				belkinunv	isbmex
				optiups	victronups
				bestfcom	ivtscd
				powercom	bcmxcp_usb
				bestfortress	liebert
				powerman-pdu	blazer_ser
				bestuferrups	liebert-esp2
				powerpanel	blazer_usb
				bestups	metasys
				rhino	richcomm_usb
				blazer_ser	mge-shut
				safenet	tripplite_usb
				etapro	mge-utalk
				solis	usbhid-ups
		-p	--port	Port (for COM). Possible values: <ul style="list-style-type: none">• dev/ttyS*;• dev/ttyUSB*;• auto. The default: auto .	

Method	Command Syntax	Functionality		
		-s	--shutdown	Shutdown timeout in seconds. The default: 5 .
		-e	--enable	Enable (1) or disable (0) UPS. The default: 1 .
		-m	--master	Set UPS type as master (1) or slave (0). The default: 1 .
		-a	--ipaddr	Master-server IP address for slave.
		Optional parameter (doesn't require arguments)		
		-c	--clean	Removes current UPS settings.
time show	rdcli param time show	Command displays the system time: date, time and time zone.		
time modify	rdcli param time modify	Command changes the system time.		
		Required at least one parameter		
		-d	--date	Date in the format YYYY-MM-DD.
		-ti	--time	Time in the format HH:MM:SS.
		-ts	--timestamp	Date and time in the format timestamp.
		-tz	--timezone	Time zone in the format tzdata.
logger show	rdcli param logger show	Command shows the current settings of the remote log function. All the logs produced by the system are sent.		
logger modify	rdcli param logger modify	Command modifies settings of the remote log function. All the logs produced by the system are sent.		
		Optional parameters		
		-ra	--remote_address	IP address of the remote server to receive logs.
		-re	--remote_enable	Enable (yes) or disable (no) log forwarding to the remote server.

Method	Command Syntax	Functionality		
		-rp	--remote_port	Port of the remote server to receive logs. By default, 514 .
sensors enclosure modify	rdcli param sensors enclosure modify	Command changes the parameters of the enclosure sensors.		
		Required parameter		
		-n	--name	The name of the enclosure sensor group.
		-t	--threshold	Trigger threshold value. Possible values: integers from 0 to 59 . The default: 0 .
sensors enclosure show	rdcli param sensors enclosure show	Command shows the current settings of the enclosure sensors.		
		Optional parameter		
		-n	--name	The name of the enclosure sensor group.
sensors modify	rdcli param sensors modify	Command modifies the sensor settings.		
		Required parameter		
		-i	--ignore	The list (comma separated without spaces) of sensor names that will be excluded from alerts. Use "+" or "-" to add or remove sensors. To add or remove a sensor whose name begins with "+" or "-", add a space at the beginning of the name and enclose the name in quotation marks, for example: ' +5v'.
sensors show	rdcli param sensors show	Command displays the current sensor settings.		
time timezone show	rdcli parameter time timezone show	Command displays the list of all time zones in the format tzdata.		
password show	rdcli param password show	Command shows the current password policy.		
		Optional parameter		
		-d	--default	Show the default password policy.

Method	Command Syntax	Functionality
password modify	rdcli param password modify	Command modifies the password policy.
		Optional parameters
		<div>-dc</div> <div>--dcredit</div> <div>Depending on the value of the parameter:</div> <ul style="list-style-type: none"> Value >0 Maximum number of credits for using numbers in the password. Value <0 Minimum number of digits. <div>Possible values: from -1000 to 1000.</div> <div>The default: 0.</div>
		<div>-dp</div> <div>--dictpath</div> <div>The path to the dictionary.</div> <div>The default: /usr/share/cracklib/empty (empty dictionary).</div>
		<div>-lc</div> <div>--lcredit</div> <div>Depending on the value of the parameter:</div> <ul style="list-style-type: none"> Value >0 Maximum number of credits for using lower case characters in the password. Value <0 Minimum number of lowercase characters in the password. <div>Possible values: from -1000 to 1000.</div> <div>The default: 0.</div>
		<div>-mcr</div> <div>--max_class_repeat</div> <div>The maximum number of consecutive symbols of one class.</div> <div>Possible values: from 0 to 1000.</div> <div>When the value is 0, the rule isn't checked.</div> <div>The default: 0.</div>


Method	Command Syntax	Functionality
	-mr	--max_repeat The maximum number of the same consecutive symbols. Possible values: from 0 to 1000 . When the value is 0, the rule isn't checked. The default: 0 .
	-ms	--max_sequence The maximum length of sequential symbols. For example, '12345' or 'fedcb'. Possible values: from 0 to 1000 . When the value is 0 , the rule isn't checked. The default: 0 .
	-mc	--minclass The minimum number of symbol classes. The classes are: <ul style="list-style-type: none"> • uppercase symbols; • lowercase symbols; • digits; • special characters. Possible values: from 0 to 4 . The default: 0 .
	-ml	--minlen The minimum password length. Possible values: from 6 to 1000 . The default: 8 .

Method	Command Syntax	Functionality
	-oc	--ocredit Depending on the value of the parameter: <ul style="list-style-type: none"> Value >0 Maximum number of credits for using special symbols in the password. Value <0 Minimum number of special symbols in the password. Possible values: from -1000 to 1000 . The default: 0 .
	-uc	--ucredit Depending on the value of the parameter: <ul style="list-style-type: none"> Value >0 Maximum number of credits for using upper case characters in the password. Value <0 Minimum number of uppercase characters in the password. Possible values: from -1000 to 1000 . The default: 0 .
	-ph	--password_history The password history length. Possible values: from 0 to 1000 . The default: 0 .
	-pwm	--password_max_attempts The login attempts number before locking of an account. Possible values: from 0 to no limit value. The default: 3 .
	-ut	--unlock_time Time to unlock a user account, in seconds. Possible values: from 0 to no limit value. The default: 300 .
Optional attribute (do not require an argument)		
	-d	--default Show the default password policy.


profile raid

Method	Command Syntax	Functionality
create	rdcli profile raid create	Command creates RAID static parameters profile with the specified name.
Required parameter		
-n	--profile_name	Profile name.
Optional parameters		
-mad	--max_alt_data	Size of alternative cache buffer, in MiB: from 200 to 4096 . The default: 512 .
-cac	--cache_alg_clean	Display type of cache algorithm for dirty segments replacement (LRU).
-cad	--cache_alg_dirty	Type of cache algorithm for dirty segments replacement: LRU or Cycle .
-lwo	--lru_wb_at_once	Maximum number of dirty segments to be flushed at once on the disks with LRU algorithm: from 1 to 128 , but not greater than {--write_back_at_once -wbo} value. The default: 24 .
-cwo	--cycle_wb_at_once	Maximum number of dirty segments to be flushed at once on the disks with Cycle algorithm: from 1 to 128 , but not greater than {--write_back_at_once -wbo} value. The default: 24 .
-mwbs	--max_write_back_seq	Maximum number of concurrent sequential requests to write back: from 1 to 1024 . The default: 8 .

Method	Command Syntax	Functionality
	-mwbr	--max_write_back_rand Maximum number of concurrent random requests to write back: from 1 to 1024 . Default value is 1024 .
	-mwd	--max_write_data Maximum write buffer size, in MiB: from 64 to 2048 . The default: 256 . The buffer contains the blocks currently written to the RAID. Specify larger values if the block size is large or if the write is performed from many hosts simultaneously.
	-rco	--reconstruct_at_once Maximum size of a reconstruction request, in MiB: from 1 to 64 . The default: 24 .
	-r	--random_read Enable (1) or disable (0) random read optimization.
	-sat	--cache_saturation Maximum percentage of cache volume available for "dirty" data. The default: 75 .
	-w	--random_write Enable (1) or disable (0) random write optimization.
	-tn	--thread_num Number of processing threads: from 0 to 256 . The value must not exceed the number of cores in the system. The default: 0 (optimal number of threads depending on number of cores).
	-vcs	--verify_cache_sum Enable (1) or disable (0) verification of cache dirty segments match on the both controllers.
	-s	--verify_synd <ul style="list-style-type: none"> 0 – disable Silent Data Corruption; 1 – enable SDC; 2 – enable the SDC detection and correction.

Method	Command Syntax	Functionality		
		-wbo	--write_back_at_once	Maximum number of stripes in single write back request: from 1 to 128 . The default: 24 .
		-wt	--write_through	<ul style="list-style-type: none">0 – disable direct drive writing;1 – enable direct drive writing;2 – enable write directly on drives automatically for sequential write (Automatic Write Through). <div> In DC mode, you can enable Automatic Write Through only on a RAID with at least one LUN.</div>
modify	rdcli profile raid modify	Command edits the specified static RAID parameter profile.		
		Required parameter		
		-n	--profile_name	Profile name.
		Optional parameters		
		-mad	--max_alt_data	Size of alternative cache buffer, in MiB: from 200 to 4096 . The default: 512 .
		-cac	--cache_alg_clean	Display type of cache algorithm for dirty segments replacement (LRU).
		-cad	--cache_alg_dirty	Type of cache algorithm for dirty segments replacement: LRU or Cycle .
		-lwo	--lru_wb_at_once	Maximum number of dirty segments to be flushed at once on the disks with LRU algorithm: from 1 to 128 , but not greater than {--write_back_at_once -wbo} value. The default: 24 .

Method	Command Syntax	Functionality
	-cwo	<p>--cycle_wb_at_once</p> <p>Maximum number of dirty segments to be flushed at once on the disks with Cycle algorithm: from 1 to 128, but not greater than {--write_back_at_once -wbo} value. The default: 24.</p>
	-mwbs	<p>--max_write_back_seq</p> <p>Maximum number of concurrent sequential requests to write back: from 1 to 1024. The default: 8.</p>
	-sat	<p>--cache_saturation</p> <p>Maximum percentage of cache volume available for "dirty" data. The default: 75.</p>
	-mwbr	<p>--max_write_back_rand</p> <p>Maximum number of concurrent random requests to write back: from 1 to 1024. The default: 1024.</p>
	-mwd	<p>--max_write_data</p> <p>Maximum write buffer size, in MiB: from 64 to 2048. The default: 256. The buffer contains the blocks currently written to the RAID. Specify larger values if the block size is large or if the write is performed from many hosts simultaneously.</p>
	-rco	<p>--reconstruct_at_once</p> <p>Maximum size of a reconstruction request, in MiB: from 1 to 64. The default: 24.</p>
	-r	<p>--random_read</p> <p>Enable (1) or disable (0) random read optimization.</p>
	-w	<p>--random_write</p> <p>Enable (1) or disable (0) random write optimization.</p>

Method	Command Syntax	Functionality		
		-tn	--thread_num	Number of processing threads: from 0 to 256 . The value must not exceed the number of cores in the system. The default: 0 (optimal number of threads depending on number of cores).
		-vcs	--verify_cache_sum	Enable (1) or disable (0) verification of cache dirty segments match on the both controllers.
		-s	--verify_synd	<ul style="list-style-type: none">0 – disable Silent Data Corruption;1 – enable SDC;2 – enable the SDC detection and correction.
		-wbo	--write_back_at_once	Maximum number of stripes in single write back request: from 1 to 128 . The default: 24 .
		-wt	--write_through	<ul style="list-style-type: none">0 – disable direct writing on the disks;1 – enable direct writing on the disks;2 – enable writing directly on drives automatically for sequential input operations (Automatic Write Through). <div> In DC mode, you can enable Automatic Write Through only on a RAID with at least one LUN.</div>
delete	rdcli profile raid delete	Command deletes a static RAID parameter profile.		
Required parameter				
	-n	--profile_name	Profile name.	
Optional attribute (does not require an argument)				
	-f	--force	Force command execution.	

Method	Command Syntax	Functionality		
show	rdcli profile raid show	Command shows all static RAID parameter profiles.		
		Optional parameters		
		-n	--profile_name	Profile name.
		Optional attribute (does not require an argument)		
		-r	--raid	Setting --profile_name is mandatory. Display list of RAIDs with the specified profile via --profile_name.
import	rdcli profile raid import	Command imports a saved RAID profile.		
		Required parameter		
		-f	--filepath	Profile name.
export	rdcli profile raid export	Command saves the current profile in the file.		
		Required parameter		
		-n	--profile_name	Profile name.


qosmic

Method	Command Syntax	Functionality									
apps show	rdcli qosmic apps show	Command shows full list of applications.									
apps create	rdcli qosmic apps create	Command adds new application from the file with signatures. Required parameter <table> <tr> <td>-f</td><td>--filepath</td><td>Path to the application.</td></tr> </table> Optional parameter <table> <tr> <td>-i</td><td>--icon</td><td>Path to the application icon.</td></tr> </table>	-f	--filepath	Path to the application.	-i	--icon	Path to the application icon.			
-f	--filepath	Path to the application.									
-i	--icon	Path to the application icon.									
apps delete	rdcli qosmic apps delete	Command deletes data on application (or file with signatures). Required parameter <table> <tr> <td>-n</td><td>--name</td><td>Name of application to delete with follow-up rebuilding.</td></tr> </table> Optional attribute (does not require an argument) <table> <tr> <td>-f</td><td>--force</td><td>Force command execution while learning is in progress.</td></tr> </table>	-n	--name	Name of application to delete with follow-up rebuilding.	-f	--force	Force command execution while learning is in progress.			
-n	--name	Name of application to delete with follow-up rebuilding.									
-f	--force	Force command execution while learning is in progress.									
apps modify	rdcli qosmic apps modify	Command changes realtime settings for applications. Required parameter <table> <tr> <td>-n</td><td>--name</td><td>Application name.</td></tr> </table> Optional parameters <table> <tr> <td>-i</td><td>--icon</td><td>Path to the application icon.</td></tr> <tr> <td>-t</td><td>--type</td><td> Application type: <ul style="list-style-type: none"> critically important (avored); non-priority (unwanted); not defined (disabled). </td></tr> </table>	-n	--name	Application name.	-i	--icon	Path to the application icon.	-t	--type	Application type: <ul style="list-style-type: none"> critically important (avored); non-priority (unwanted); not defined (disabled).
-n	--name	Application name.									
-i	--icon	Path to the application icon.									
-t	--type	Application type: <ul style="list-style-type: none"> critically important (avored); non-priority (unwanted); not defined (disabled). 									
apps rename	rdcli qosmic apps rename	Command renames application. Required parameters									

Method	Command Syntax	Functionality		
		-n	--name	Current application name.
		-nn	--nnname	New application name.
apps clean	rdcli qosmic apps clean	Command deletes all warnings from the application.		
Required parameter				
		-n	--name	Application name.
teacher show	rdcli qosmic teacher show	Command shows teaching progress for applications.		
teacher start	rdcli qosmic teacher start	Command starts learning.		
Required parameters				
		-hn	--hosts	Host IDs (comma-separated).
		-n	--name	Application name.
Optional parameter				
		-i	--icon	Path to the application icon.
teacher stop	rdcli qosmic teacher stop	Command stops process of learning.		
teacher resume	rdcli qosmic teacher resume	Command resumes learning.		
Required parameters				
		-hn	--hosts	Host IDs (comma-separated).
		-n	--name	Application name.
Optional parameter				
		-i	--icon	Path to the application icon.
recognizer show	rdcli qosmic recognizer show	Command shows the progress of recognizing.		
recognizer start	rdcli qosmic recognizer start	Command starts recognizing.		
recognizer stop	rdcli qosmic recognizer stop	Command stops recognizing.		


raid

Method	Command Syntax	Functionality												
show	rdcli raid show	<p>Command displays list of created Generic RAIDs.</p> <p>Optional parameters</p> <table> <tr> <td>-n</td><td>--name</td><td>RAID name.</td></tr> <tr> <td>-s</td><td>--raid_status</td><td>Display all RAIDs with the specified status.</td></tr> </table> <p>Optional attributes (do not require an argument)</p> <table> <tr> <td>-si</td><td>--sdcinfo</td><td>Display information on silent data corruptions on RAID volumes.</td></tr> </table>	-n	--name	RAID name.	-s	--raid_status	Display all RAIDs with the specified status.	-si	--sdcinfo	Display information on silent data corruptions on RAID volumes.			
-n	--name	RAID name.												
-s	--raid_status	Display all RAIDs with the specified status.												
-si	--sdcinfo	Display information on silent data corruptions on RAID volumes.												
create	rdcli raid create	<p>Command creates a Generic RAID with the specified name, level and number of drives.</p> <div> <p>i To create a RAID, you can select both SSDs and HDDs. Consider that specific SSD wearing out occurs in RAIDs, containing both SSD and HDD.</p> <p>i When creating several RAIDs, remember that RAIDs do not share memory size. Thus, if you use the entire amount of RAM for the first RAID, you will not be able to create the second RAID with desired cache size.</p> </div> <p>Required parameters</p> <table> <tr> <td>-n</td><td>--name</td><td>RAID name.</td></tr> <tr> <td>-l</td><td>--level</td><td>RAID level.</td></tr> <tr> <td>-d</td><td>--drives</td><td>RAID drives.</td></tr> </table> <p>Optional parameters</p> <table> <tr> <td>-cs</td><td>--cache_size</td><td>Cache buffer size, in MiB: from 1024 to (2³² - 1). The default: 4096.</td></tr> </table>	-n	--name	RAID name.	-l	--level	RAID level.	-d	--drives	RAID drives.	-cs	--cache_size	Cache buffer size, in MiB: from 1024 to (2³² - 1) . The default: 4096 .
-n	--name	RAID name.												
-l	--level	RAID level.												
-d	--drives	RAID drives.												
-cs	--cache_size	Cache buffer size, in MiB: from 1024 to (2³² - 1) . The default: 4096 .												

Method	Command Syntax	Functionality		
		-fsync	--cache_fsync_enabled	Enable (1) or disable (0) the fsync support. The default: 1.
		-ss	--stripe_size	Stripe size, in KiB.
		-m4	--mode_4k	RAID 0 and RAID 10 only. Enable (1) or disable (0) 4 KiB write.
		-rd	--redundant_drives	Number of redundant drives for RAID N+M.
		-p	--profile	RAID profile name.
		-a	--affinity	Only for DC systems. RAID affinity.
		-sp	--sparepool	SparePool name.
		-wt	--write_through	Enable (1) or disable (0) write directly on the disks; enable write directly on drives automatically for sequential input operations – Automatic Write Through (2). <div> In DC mode, you can enable Automatic Write Through only on a RAID with at least one LUN.</div>
modify	rdcli raid modify	Command changes Generic RAID parameters.		
		Required parameter		
		-n	--name	RAID name.
		Optional parameters		
		-r	--random_read	Enable (1) or disable (0) random read optimization.
		-w	--random_write	Enable (1) or disable (0) random write optimization.
		-cac	--cache_alg_clean	Display a type of cache algorithm for dirty segments replacement (LRU).

Method	Command Syntax	Functionality
	-cad	--cache_alg_dirty Type of cache algorithm for dirty segments replacement: <ul style="list-style-type: none"> • cycle is intended for random write operations; • lru is intended for sequential write operations. The default: lru .
	-cs	--cache_size Cache size in MiB. Possible values: from 1024 to 2³²-1 . The default: 1024 .
	-fsync	--cache_fsync_enabled Enable (1) or disable (0) the fsync support. The default: 1 .
	-lwo	--lru_wb_at_once Maximum number of dirty segments to be flushed at once on the disks with LRU algorithm: from 1 to 128 , but not greater than {--write_back_at_once -wbo} value. The default: 24 .
	-mad	--max_alt_data Size of alternative cache buffer, in MiB: from 200 to 4096 . The default: 512 .
	-mwd	--max_write_data Maximum write buffer size, in MiB: from 64 to 2048 . The default: 256 . The buffer contains the blocks currently written to the RAID. Specify larger values if the block size is large or if the write is performed from many hosts simultaneously.
	-mwbs	--max_write_back_seq Maximum number of concurrent sequential requests to write back: from 1 to 1024 . The default: 8 .

Method	Command Syntax	Functionality
	-mwbr	--max_write_back_rand Maximum number of concurrent random requests to write back: from 1 to 1024 . The default: 1024 .
	-p	--profile RAID profile name.
	-rcc	--recon_commands Number of outstanding reconstruction requests. Possible values: from 1 to 64 . The default: 8 .
	-sat	--cache_saturation Maximum cache saturation level for non-real-time requests, in percent: from 1 to 100 . The default: 75 .
	-s	--verify_synd Enable (1) or disable (0) Silent Data Corruption; enable Silent Data Corruption detection and the correction (2).
	-sp	--sparepool SparePool name.
	-cwo	--cycle_wb_at_once Maximum number of dirty segments to be flushed at once on the disks with Cycle algorithm: from 1 to 128 , but not greater than the {--write_back_at_once -wbo} value. The default: 24 .
	-tn	--thread_num Number of processing threads: from 0 to 256 . The value must not exceed the number of cores in the system. The default: 0 (optimal number of threads depending on number of cores).

Method	Command Syntax	Functionality		
		-wt	--write_through	<p>Enable (1) or disable (0) write directly on the disks; enable write directly on drives automatically for sequential input operations – Automatic Write Through (2).</p> <div>  In DC mode, you can enable Automatic Write Through only on a RAID with at least one LUN. </div>
		-vcs	--verify_cache_sum	Enable (1) or disable (0) verification of cache dirty segments match on the both controllers.
		-wbo	--write_back_at_once	<p>Maximum number of stripes in single write back request: from 1 to 128.</p> <p>The default: 24.</p>
		-pf	--predictive_failover	<p>Enable (1) or disable (0) data reconstruction when most drives are complete.</p> <p>The default: 0.</p>
		-rcm	--recon_min	<p>Reconstruction priority: from 0 to 100.</p> <p>The default: 5.</p>
		-rco	--reconstruct_at_once	<p>Maximum size of a single request for reconstruction: from 1 to 64.</p> <p>The default: 24.</p>
		Optional attribute (does not require an argument)		
		-f	--force	Force command execution.
verify create	rdcli raid verify create	Command allows enabling and configuring background Generic RAID integrity check.		
		Required parameters		
		-n	--name	RAID name.
		Optional parameter		


Method	Command Syntax	Functionality		
		-rr	--rate_read	Background integrity check rate, in MiB/s.
		-s	--schedule	Comma-separated list of durations, in which RAID Integrity Check stopped. Duration format: <week_day>-<hh>:<mm>=on,<week_day>-<hh>:<mm>=off Limitations: <ul style="list-style-type: none"> • one duration of Integrity Check inactivity is allowed per day • <week_day> is one of: mon, tue, wed, thu, fri, sat, or sun • <hh> must be the 24-hour clock
		-e	--enable	Enable (yes) or disable (no) the RAID integrity check. The default: no .
verify modify	rdcli raid verify modify	Command modifies the schedule and settings of background integrity check for Generic RAID.		
		Required parameter		
		-n	--name	RAID name.
		Optional parameters		

Method	Command Syntax	Functionality	
		-s	--schedule Comma-separated list of durations' modifications, in which RAID Integrity Check stopped. To add a new period: +<week_day>-<hh>:<mm>=on,+<week_day>-<hh>:<mm>=off To remove existing period: -<week_day>-<hh>:<mm>=on,-<week_day>-<hh>:<mm>=off To overwrite the whole list of durations: <week_day>-<hh>:<mm>=on,<week_day>-<hh>:<mm>=off Limitations: <ul style="list-style-type: none"> • one duration of Integrity Check inactivity is allowed per day • <week_day> is one of: mon, tue, wed, thu, fri, sat, or sun • <hh> must be the 24-hour clock
		-e	--enable Enable (1) or disable (0) background integrity check on schedule.
		-rr	--rate_read Background integrity check rate, in MiB/s.
verify show	rdcli raid verify show	Command shows information on state and progress of background integrity check for Generic RAID. Optional parameter	
		-n	--name Name of a RAID.
verify delete	rdcli raid verify delete	Command disables background integrity check and deletes its settings and schedule for Generic RAID. Required parameter	
		-n	--name Name of a RAID.
migrate	rdcli raid migrate	Command changes Affinity of the specified RAID (performs migration of RAID on the other node).	

Method	Command Syntax	Functionality
		Required parameter
		-n --name RAID name.
		Optional attribute (does not require an argument)
		-f --force Force command execution.
delete	rdcli raid delete	Command deletes Generic RAID (including from configuration file).
		Required parameter
		-n --name RAID name.
		Optional attributes (do not require an argument)
		-f --force Force command execution.
		-ff --force_fast Forced quick RAID deletion. If you use the attribute while deleting RAID, SSD cache doesn't flush and the deletion takes less time. In this case, the deleted data can't be restored.
		-fw --force_wait Wait for SSD cache to flush before deletion.
init show	rdcli raid init show	Command shows Generic RAID initialization status.
		Required parameter
		-n --name RAID name.
		Optional attribute (does not require an argument)
		-f --force Force command execution.
init start	rdcli raid init start	Command starts Generic RAID initialization.
		Required parameter
		-n --name RAID name.
		Optional attribute (does not require an argument)
		-f --force Force command execution.
init stop	rdcli raid init stop	Command stops Generic RAID initialization.
		Required parameter

Method	Command Syntax	Functionality			
		<table><tr><td>-n</td><td>--name</td><td>RAID name</td></tr></table> <p>Optional attribute (does not require an argument)</p>	-n	--name	RAID name
-n	--name	RAID name			
		<table><tr><td>-f</td><td>--force</td><td>Force command execution</td></tr></table>	-f	--force	Force command execution
-f	--force	Force command execution			
reconstruct start	rdcli raid reconstruct start	Command starts Generic RAID reconstruction.			
		Required parameter			
		<table><tr><td>-n</td><td>--name</td><td>RAID name.</td></tr></table>	-n	--name	RAID name.
-n	--name	RAID name.			
reconstruct stop	rdcli raid reconstruct stop	Command stops Generic RAID reconstruction.			
		Required parameter			
		<table><tr><td>-n</td><td>--name</td><td>RAID name.</td></tr></table>	-n	--name	RAID name.
-n	--name	RAID name.			
replace	rdcli raid replace	Command allows replacing drive with the specified UID (attribute --old) in the Generic RAID with a drive not used in other RAID. If the new drive is not specified, system automatically selects the first suitable drive.			
		Required parameters			
		<table><tr><td>-n</td><td>--name</td><td>RAID name.</td></tr></table>	-n	--name	RAID name.
-n	--name	RAID name.			
		<table><tr><td>-ou</td><td>--old</td><td>Old drive UID.</td></tr></table>	-ou	--old	Old drive UID.
-ou	--old	Old drive UID.			
		Optional parameters			
		<table><tr><td>-nu</td><td>--new</td><td>New drive UID.</td></tr></table>	-nu	--new	New drive UID.
-nu	--new	New drive UID.			
		Optional attribute (does not require an argument)			
		<table><tr><td>-f</td><td>--force</td><td>Force command execution.</td></tr></table>	-f	--force	Force command execution.
-f	--force	Force command execution.			
reload	rdcli raid reload	Command reloads the specified Generic RAID.			
		Required parameter			
		<table><tr><td>-n</td><td>--name</td><td>RAID name.</td></tr></table>	-n	--name	RAID name.
-n	--name	RAID name.			
		Optional attribute (does not require an argument)			

Method	Command Syntax	Functionality		
		-L	--local	Run the command on a local node only.
		-f	--force	Force command execution.

 Do not reload the RAID while the system is under load.

replication

Method	Command Syntax	Functionality		
create	rdcli replication create	Command creates Replication of a LUN.		
		Required parameters		
		-ipl	--local_ip	IP address of a node with a local LUN.
		-n	--name	Name of a local LUN (replicated).
		-rn	--remote_name	Name of a remote LUN (replica).
		-ipr	--remote_ip	IP address of a node with a remote LUN.
		Optional parameter		
		-p	--protocol	Replication protocol. Possible values: A (synchronous), C (asynchronous). The default: C .
delete	rdcli replication delete	Command deletes Replication of a LUN.		
		Required parameter		
		-n	--name	Name of a replicated LUN.
		Optional parameter (does not require arguments)		
		-f	--force	Force the command execution.
invalidate	rdcli replication invalidate	Command sets the status of LUN data as out-of-sync.		
		Required parameter		
		-n	--name	Name of a replicated LUN.
modify	rdcli replication modify	Command modifies settings of Replication.		
		Optional parameters		
		-ipl	--local_ip	IP address of a node with a local LUN.

Method	Command Syntax	Functionality	
		-n	--name Name of a local LUN (replicated). Interchangeable with the <code>--all</code> parameter.
		-oos	--out_of_sync_threshold Amount of data allowed to be out-of-sync, in KiB.
		-p	--protocol Replication protocol. Possible values: A (synchronous), C (asynchronous). The default: C .
		-ipr	--remote_ip IP address of a node with a remote LUN.
		-rn	--remote_name Name of a remote LUN (replica).
		-r	--role Role of the LUN participating in Replication. Possible values: <ul style="list-style-type: none"> primary To set "primary", <code>--force</code> is required. secondary
		Optional parameters (do not require arguments)	
		-a	--all Change the roles of all LUNs. Used with the <code>--role</code> parameter.
		-f	--force Required while the LUN role being changed to "primary".
suspend	rdcli replication suspend	Command suspends Replication for a LUN.	
		Required parameter	
		-n	--name Name of a replicated LUN.
		Required parameter	
resume	rdcli replication resume	Command resumes Replication for a suspended LUN.	
		Required parameter	
		-n	--name Name of a replicated LUN.
		Optional parameter (does not require arguments)	

Method	Command Syntax	Functionality		
		-d	--discard	Set the status of data on a LUN as out-of-sync (1). The default: 0.
show	rdcli replication show	Command shows settings of Replication for selected LUNs.		
		Optional parameter		
		-n	--name	Names of replicated LUNs, comma separated.
metadata clean	rdcli replication metadata clean	Command re-creates replication metadata on a LUN.		
		Required parameter		
		-n	--name	Name of a LUN with replication metadata.
		Optional parameter (does not require arguments)		
		-f	--force	Force the command execution.
metadata create	rdcli replication metadata create	Command creates a LUN with replication metadata.		
		Required parameter		
		-n	--name	Name of the replication metadata LUN.
metadata delete	rdcli replication metadata delete	Command deletes a LUN with replication metadata.		
		Required parameter		
		-n	--name	Replication metadata LUN name.
metadata show	rdcli replication metadata show	Command displays replication metadata LUNs.		
verify start	rdcli replication verify start	Command launches background check of replicated data.		
		Required parameter		
		-n	--name	LUN name.
verify stop	rdcli replication verify stop	Command stops background check of replicated data.		
		Required parameter		
		-n	--name	LUN name.

Method	Command Syntax	Functionality			
verify show	rdcli replication verify show	Command displays the progress of background check of replicated data.			
		Optional parameter			
		<table> <tr> <td>-n</td><td>--name</td><td>LUN name.</td></tr> </table>	-n	--name	LUN name.
-n	--name	LUN name.			

sas target

Method	Command Syntax	Functionality
show	rdcli sas target show	Command displays list of SAS targets used in the system.
modify	rdcli sas target modify	Command associates SAS targets on the local and remote nodes.
Required parameters		
-n	--target_name	Local target name.
-r	--remote_target	Remote target name.
Optional parameter (does not require an argument)		
-f	--force	Force command execution.

sensor

Method	Command Syntax	Functionality
show	rdcli sensor show	Command displays list of sensors in the system.
Optional parameter		
-s	--skip_status	List of statuses to skip.
Optional parameter (doesn't require arguments)		
-e	--enclosure	Display information on sensors of all connected enclosures.

session


Method	Command Syntax	Functionality
show	rdcli session show	Command displays all established sessions.

sparepool

Method	Command Syntax	Functionality		
create	rdcli sparepool create	Command creates a new SparePool with the specified name and list of drives.		
		Required parameters		
		-n	--name	SparePool name.
		-d	--drives	Drive UUIDs.
modify	rdcli sparepool modify	Command changes configuration of the specified SparePool.		
		Required parameters		
		-n	--name	SparePool name.
		-d	--drives	Drive UUIDs.
delete	rdcli sparepool delete	Command deletes a SparePool.		
		Required parameter		
		-n	--name	SparePool name.
		Optional parameter (does not require an argument)		
		-f	--force	Force command execution.
show	rdcli sparepool show	Command displays info on created SparePools.		
		Optional parameter		
		-n	--name	Name of the specific SparePool.

h2>ssdcache

Method	Command Syntax	Functionality		
create	rdcli ssdcache create	Creates SSD cache.		
		Required parameter		
		-l	--lun	Name of a LUN used as SSD cache (caching LUN).
		Optional parameters		
		-a	--affinity	Only for DC systems. Node ID at which a RAID with SSD cache is Active. Possible values: 0 , 1 .
		-fp	--flushing_prio	SSD cache flushing priority in percent. Possible values: integers from 0 (don't flush) to 100 . Default is 100 .
		-m	--memory	Amount of RAM in MiB reserved for SSD cache. Default: calculated according to the formula .
		Optional parameter (does not require an argument)		
		-ro	--readonly	"Read Only" mode for SSD cache.
delete	rdcli ssdcache delete	Deletes SSD cache.		
		Mutually exclusive required parameters		
		-a	--affinity	Only for DC systems. Node ID at which a RAID with SSD cache is Active. Possible values: 0 , 1 .
		-n	--name	Name of SSD cache (name of the caching LUN).
		Required parameters for specific situations (do not require arguments)		

Method	Command Syntax	Functionality		
		-ff	--force_fast	<p>When SSD cache in read/write mode and a RAID of caching LUN or main LUN is Offline.</p> <p>Forces the execution of a command with data loss.</p>
		-fw	--force_wait	<p>When SSD cache in read/write mode and the “--force_fast” parameter isn’t set.</p> <p>Forces the execution of the command while saving the data.</p> <div>  If SSD cache contains data, which has not been flushed to the main RAID, SSD cache disabling may take a long time due to the flush of cached data. </div>
modify	rdcli ssdcache modify	Modifies SSD cache.		
		Mutually exclusive required parameters		
		-a	--affinity	<p>Only for DC systems.</p> <p>Node ID at which a RAID with SSD cache is Active.</p> <p>Possible values: 0, 1.</p>
		-n	--name	Name of SSD cache (name of the caching LUN).
		Optional parameter		
		-fp	--flushing_prio	<p>SSD cache flushing priority in percent.</p> <p>Possible values: integers from 0 (don’t flush) to 100.</p> <p>Default is 100.</p>
show	rdcli ssdcache show	Shows SSD cache parameters.		
		Optional parameters		
		-a	--affinity	<p>Only for DC systems.</p> <p>Node ID at which a RAID with SSD cache is Active.</p> <p>Possible values: 0, 1.</p>
		-n	--name	Name of SSD cache (name of the caching LUN).

system

Method	Command Syntax	Functionality
command show	rdcli system command show	Command shows the current command queue.
Optional parameters		
-c	--command	Name of the command in the queue.
-o	--object	Name of the object to which commands refer.
-rn	--remote_node	Only show (1) commands from the remote node or only show commands from the local node (0). The default: 0.
-r	--running	Only show commands that are running now (1) or show all queue commands (0). The default: 0.
-s	--source	Show commands by the name of a source. For example: rdcli .
-u	--user	Show commands started by the specific user. For example: admin .
Optional parameter (does not require arguments)		
-a	--all	Show all commands in the queue. By default, the system shows only blocking commands.
show	rdcli system show	Command displays information about the system.
modify	rdcli system modify	Command modifies system settings.
Required parameter (does not require arguments)		
-er	--eula_revoke	Revoke user's EULA acceptance on the system. For DC system, revokes from both nodes.
restart	rdcli system restart	Command restarts RAIDIX software.

Method	Command Syntax	Functionality						
license show	rdcli system license show	<p>Command displays complete information about the following license parameters: validity, status, name of the license file, license key, number of hardware modifications, available functionality, accessible volume, license type.</p> <p>The command without parameters shows the list of licenses.</p> <p>Mutually exclusive optional parameters</p> <table> <tr> <td>-hw</td><td>--hardware_key</td><td>Display hardware key.</td></tr> <tr> <td>-s</td><td>--summary</td><td>Display information only about period of validity, functionality, license type and volume.</td></tr> </table>	-hw	--hardware_key	Display hardware key.	-s	--summary	Display information only about period of validity, functionality, license type and volume.
-hw	--hardware_key	Display hardware key.						
-s	--summary	Display information only about period of validity, functionality, license type and volume.						
license create	rdcli system license create	<p>Command adds a license file to the system.</p> <p>Required parameter</p> <table> <tr> <td>-f</td><td>--filename</td><td>The path to the license file.</td></tr> </table>	-f	--filename	The path to the license file.			
-f	--filename	The path to the license file.						
license delete	rdcli system license delete	<p>Command deletes a license file.</p> <p>Required parameter</p> <table> <tr> <td>-f</td><td>--filename</td><td>The name of a system license (filename value at system license show).</td></tr> </table>	-f	--filename	The name of a system license (filename value at system license show).			
-f	--filename	The name of a system license (filename value at system license show).						
license get	rdcli system license get	<p>Wrap a license in the file license.rpk and place it in the /tmp directory.</p> <p>Required parameter</p> <table> <tr> <td>-f</td><td>--file</td><td>The name of a system license (filename value at system license show).</td></tr> </table>	-f	--file	The name of a system license (filename value at system license show).			
-f	--file	The name of a system license (filename value at system license show).						
config restore	rdcli system config restore	<p>Command restores the configuration file.</p> <p>Mutually exclusive required parameters</p> <table> <tr> <td>-f</td><td>--file</td><td>Path to the new configuration file.</td></tr> </table>	-f	--file	Path to the new configuration file.			
-f	--file	Path to the new configuration file.						
settings show	rdcli system settings show	Command displays current system settings.						
settings modify	rdcli system settings modify	<p>Command modifies system settings.</p> <p>Optional parameters</p>						

Method	Command Syntax	Functionality		
		-ct	--cmd_timeout	Timeout for service commands, in seconds. Possible values: integers from 1 . The default: 600 .
		-co	--completion_timeout	Timeout for individual services, in seconds. Possible values: integers from 1 . The default: 500 .
		-dt	--deinit_timeout	Timeout for all services, in seconds. Possible values: integers from 1 . The default: 2500 .
		-et	--exec_timeout	Timeout for system commands, in seconds. Possible values: integers from 1 . The default: 30 .
		-l	--log_level	Type of information will be logged. Possible values: crit, err, warn, notice, info, debug . Each next level includes previous. The default: info .
		-rt	--restore_timeout	Timeout for recovery, in seconds. Possible values: integers from 1 . The default: 2400 .
		-do	--device_loss_timeout	Timeout in seconds used to reload the LUN on the initiator when the VAAI ID changes. Possible values: integers from 1 . The default: 3 .
reboot	rdcli system reboot	Command reboots the system.		
		Optional parameter (doesn't require arguments)		
		-d	--dc	Reboot the DC system entirely.

Method	Command Syntax	Functionality
shutdown	rdcli system shutdown	Command shuts down the system.
		Optional parameter (doesn't require arguments)
		-d --dc Shut down the DC system entirely.
suspend	rdcli system suspend	Restrict the system from accepting any commands other than service commands (e.g. ping). After the system is rebooted, the restriction remains. Use the <code>system resume</code> command to cancel the restriction.
resume	rdcli system resume	Cancel the <code>system suspend</code> command.
user show	rdcli system user show	Command displays list of all RAIDIX web interface users.
		Optional parameter
		-l --login User login.
user create	rdcli system user create	Command creates a user with the specified name and password.
		Required parameters
		-l --login User login.
		-p --password User password. To hide the password from the command history, use instead the parameter <code>--password_ask</code> without arguments.
		Optional parameters
		-la --language Interface language:
		<ul style="list-style-type: none"> • English (en); • Russian (ru); • Japanese (ja). The default: en .

Method	Command Syntax	Functionality
	-r	--role User roles. Possible values: <ul style="list-style-type: none"> administrators; operators. The default: operators .
	-ae	--account_expires Expiration date of the user account. Possible values: <ul style="list-style-type: none"> -1 (never); 0 (password must be changed); YYYY-MM-DD formatted date. The default: -1 .
	-id	--inactive_days Time to lock inactive user after password lifetime expired. Possible values: <ul style="list-style-type: none"> -1 (never); 0 (password must be changed); from 1 to $2^{63}-1$. The default: -1 .
	-max	--max_days The maximum number of days the password is valid since the last password change. Possible values: from -1 to $2^{63}-1$. To remove the password validation check, use the -1 value. The default: -1 . Pre-created system user(s) default: 99999 .

Method	Command Syntax	Functionality	
		-wd	<p>--warn_days</p> <p>The number of days before the password expires or is replaced, during which the user will receive warnings about the necessity of changing the password.</p> <p>To remove the notification, use the -1 value.</p> <p>Possible values: from -1 to 2⁶³-1.</p> <p>The default: -1.</p> <p>Pre-created system user(s) default: 7.</p>
		-min	<p>--min_days</p> <p>The minimum number of days the password is valid since the last password change.</p> <p>Possible values: from -1 to 2⁶³-1.</p> <p>To remove the password validation check, use the -1 value.</p> <p>The default: -1.</p> <p>Pre-created system user(s) default: 0.</p>
		Optional parameters (don't require an argument)	
		-pe	<p>--password_expires</p> <p>Force the user to change his password at the next login.</p>
user modify	rdcli system user modify	-pa	<p>--password_ask</p> <p>Flag for interactive entry of the user's password.</p> <p>With this input, the text of the password is not displayed in the command history.</p> <p>The parameter is mutually exclusive with the --password parameter.</p>
		Required parameters	
		-l	<p>--login</p> <p>User login.</p>
		Optional parameters	
		-p	<p>--password</p> <p>New user password.</p> <p>To hide the password from the command history, use instead the parameter --password_ask without arguments.</p>

Method	Command Syntax	Functionality
	-la	--language Interface language: <ul style="list-style-type: none"> English (en); Russian (ru); Japanese (ja). The default: en .
	-r	--role User roles. Possible values: <ul style="list-style-type: none"> administrators; operators. The default: operators .
	-ae	--account_expires Expiration date of the user account. Possible values: <ul style="list-style-type: none"> -1 (never); 0 (password must be changed); YYYY-MM-DD formatted date. To off the expired status, use the value -1 . The default: -1 .
	-id	--inactive_days Time to lock inactive user after password lifetime expired. Possible values: <ul style="list-style-type: none"> -1 (never); 0 (password must be changed); from 1 to 2⁶³-1. To off the lock status, use the value -1 . The default: -1 .

Method	Command Syntax	Functionality
	-max	--max_days The maximum number of days the password is valid since the last password change. Possible values: from -1 to 2⁶³-1 . To remove the password validation check, use the -1 value. The default: -1 . Pre-created system user(s) default: 99999 .
	-wd	--warn_days The number of days before the password expires or is replaced, during which the user will receive warnings about the necessity of changing the password. To remove the notification, use the -1 value. Possible values: from -1 to 2⁶³-1 . The default: -1 . Pre-created system user(s) default: 7 .
	-min	--min_days The minimum number of days the password is valid since the last password change. Possible values: from -1 to 2⁶³-1 . To remove the password validation check, use the -1 value. The default: -1 . Pre-created system user(s) default: 0 .
Optional parameters (don't require an argument)		
	-pe	--password_expires Force the user to change his password at the next login.
	-pa	--password_ask Flag for interactive entry of the user's password. With this input, the text of the password is not displayed in the command history. The parameter is mutually exclusive with the --password parameter.
user delete	rdcli system user delete	Command deletes a specified user.

Method	Command Syntax	Functionality
		Required parameter
		-l --login User login.
		Optional parameter
		-f --force Force command execution and delete user's sessions.
user unlock	rdcli system user unlock	Command unlocks a user account blocked for exceeding the limit of the incorrectly entered password.
		Required parameter
		-l --login User name.
		drive partition show
		rdcli system drive partition show
drive partition show	rdcli system drive partition show	Command shows parameters of system partitions.
		name Path to mount point.
		available_size Free space on partition, in KiB.
		bdev Name of mounted block device.
		fs_type Type of file system.
		parent Child block devices.
		total_size Total size of partition, in KiB.
		used_percent Percent of used space.
		used_size Used space on partition, in KiB.
		raid Only when MDRAID. MDRAID state info.
drive show	rdcli system drive show	Command shows parameters and allocated status of devices mounted at / and /var.

log_collector

Method	Command Syntax
log_collector	Command collects system logs and displays the path to the log file (.zip archive) in the directory /var/spool/rdx/ for downloading.

APPENDIX B. MPIO SETUP ON LINUX ENTERPRISE SERVER

i To setup MPIO, make sure you have administrative permissions for Linux Enterprise Server on the server connected to the storage system with RAIDIX 5.2 installed.

To set up MPIO on a Linux OS with an initiator:

1. Install the *multipath-tools* package.
2. Create the */etc/multipath.conf* file containing the following:

```

defaults {
    fast_io_fail_tmo          5
    features                  "0"
    no_path_retry             10
    path_checker              tur
    polling_interval          5
    prio                      alua
    user_friendly_names       yes
}

devices {
    device {
        detect_checker        no
        detect_prio           no
        failback               immediate
        no_path_retry         12
        path_grouping_policy  "group_by_prio"
        path_selector          "round-robin 0"
        path_checker           "tur"
        prio                   "alua"
        product                ".*"
        rr_min_io              100
        rr_weight              "uniform"
        vendor                 "Raidix"

        #For initiators with scsi_dh_alua
        #hardware_handler      "1 alua"
    }
    device {
        detect_checker        no
        detect_prio           no
        failback               immediate
        no_path_retry         30
        path_checker           directio
        path_grouping_policy  "group_by_prio"
        path_selector          "round-robin 0"
        prio                   ana
        product                "Raidix"
        rr_min_io              100
        rr_weight              "uniform"
        uid_attribute          ID_WWN
        vendor                 "NVME"
    }
}

```

- 2.1. If your OS is CentOS 7.0, Red Hat 7, or has the `scsi_dh_alua` module, then you must uncomment (remove the character `#`) the following line:

```
#hardware_handler
```

- 2.2. When using NVMe-oF in DC system with large number of objects (RAIDs more than 10 and LUNs more than 100), we recommend increasing the parameter `no_path_retry` (in the second "device" of the object "devices" in the *multipath.conf*).

With a large number of objects in the system, failover may take longer than 2.5 minutes (the default value is set in the *multipath.conf* configuration file; the multiplying of `polling_interval` and `no_path_retry` values). The time for the successful failover can only be determined experimentally.

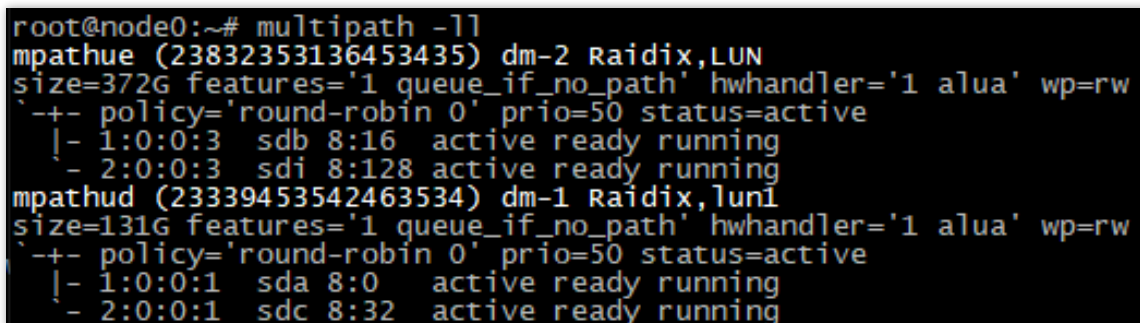
3. Run the `multipath` service:

```
# systemctl start multipathd.service
```

To check the MPIO-device status, run the following command:

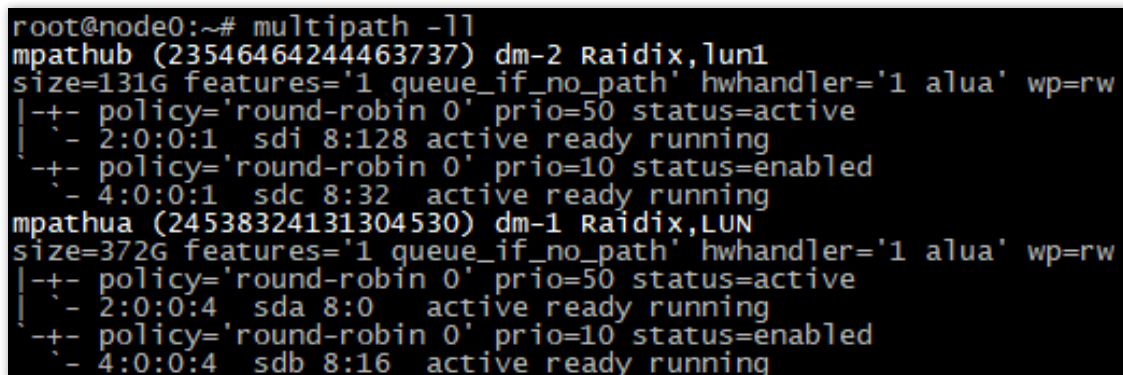
```
# multipath -ll
```

A list of MPIO device parameters appears (Figure 180, Figure 181):



```
root@node0:~# multipath -ll
mpathue (23832353136453435) dm-2 Raidix,LUN
size=372G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
+- policy='round-robin 0' prio=50 status=active
|- 1:0:0:3 sdb 8:16 active ready running
|- 2:0:0:3 sdi 8:128 active ready running
mpathud (23339453542463534) dm-1 Raidix,lun1
size=131G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
+- policy='round-robin 0' prio=50 status=active
|- 1:0:0:1 sda 8:0 active ready running
|- 2:0:0:1 sdc 8:32 active ready running
```


Figure 180. An example of MPIO device parameter output for Single mode



```
root@node0:~# multipath -ll
mpathub (23546464244463737) dm-2 Raidix,lun1
size=131G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
+- policy='round-robin 0' prio=50 status=active
|- 2:0:0:1 sdi 8:128 active ready running
+- policy='round-robin 0' prio=10 status=enabled
|- 4:0:0:1 sdc 8:32 active ready running
mpathua (24538324131304530) dm-1 Raidix,LUN
size=372G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
+- policy='round-robin 0' prio=50 status=active
|- 2:0:0:4 sda 8:0 active ready running
+- policy='round-robin 0' prio=10 status=enabled
|- 4:0:0:4 sdb 8:16 active ready running
```

Figure 181. An example of MPIO device parameter output for DC mode

APPENDIX C. SENDING LOGS TO A REMOTE LINUX SYSTEM THROUGH RSYSLOG SERVICE


 To manage the rsyslog service, you must be logged as an administrator.

Configure the node from which logs are collected

1. Run the command

```
$ rdcli param logger modify [-ra <remote_address>] [-re {yes|no}] [-rp <remote_port>]
```

where <remote_address> – an IP address of a remote system to which logs will be sent; {yes|no} – enable or disable the sending of logs; <remote_port> – a port of a remote system (the default is 514).

2. Restart the rsyslog service on each participated node:

```
# systemctl restart rsyslog.service
```

To configure the receiver node

1. Create or edit the file */etc/rsyslog.d/20-raidix-core-from-remote.conf* in the following way:

```
## Receiving logs from remote hosts

module(load="imtcp" MaxSessions="500")
input(type="imtcp" port="514" ruleset="remote")
template(name="RemoteHost" type="string" string="/var/log/remote-
%HOSTNAME%.log")
ruleset(name="remote") {
    action(type="omfile" dynaFile="RemoteHost") stop
}
```

Descriptions of the instructions

Parameter	Description
	<div>module</div> (to learn more, see rsyslog/module)
load	Module for receiving messages. Imtcp provides the ability to receive syslog messages via TCP. For more details, see rsyslog/imtcp .
MaxSessions	Sets the maximum number of sessions supported. The default is 200 .

Parameter	Description
input (to learn more, see rsyslog/input)	
type	Type of the input parameters module.
port	TCP port.
ruleset	Name of the used set of rules.
template (to learn more, see rsyslog/template)	
name	Name of the template.
type	Type of the template. string has a mandatory parameter string , which holds the template string to be applied.
string	string-based template.
ruleset (to learn more, see rsyslog/ruleset)	
name	Name of the set of rules.

- Restart the `rsyslog` service on each participated node:

```
# systemctl restart rsyslog.service
```

In this example, the messages will be logged into the files `/var/log/remote-%HOSTNAME%.log`.