



RAIDIX 4.4.3

JetStor RDX16 / RDX24 / RDX48 / RDX60 / RDX90

Administrator's Guide

Document revision 1.0

Advanced Computer & Network Corp.

5001 Baum Blvd, Ste 680

Pittsburgh, PA 15213

support@acnc.com

www.acnc.com

Table of Contents

| | |
|--------------------------------------------------------------|-----------|
| INTRODUCTION..... | 6 |
| INTENDED AUDIENCE | 6 |
| HOW THE GUIDE IS ORGANIZED..... | 6 |
| GUIDE CONVENTIONS | 7 |
| CONTACT US | 7 |
| ABOUT RAIDIX 4.4.3..... | 8 |
| DUAL-CONTROLLER MODE | 11 |
| 1. QUICK START | 13 |
| 1.1. WEB INTERFACE OVERVIEW | 13 |
| 1.1.1. Software Interface Language | 15 |
| 1.1.2. Wizards | 15 |
| 1.1.3. Surrounding RAIDIX Systems | 16 |
| 1.2. DUAL-CONTROLLER MODE CONFIGURATION | 17 |
| 1.2.1. Set Up InfiniBand Synchronization | 20 |
| 1.2.2. Set Up iSCSI Synchronization | 22 |
| 1.2.3. Set Up Fibre Channel Synchronization..... | 23 |
| 1.2.4. Set Up SAS Synchronization | 24 |
| 1.3. GETTING STARTED WITH THE SOFTWARE | 24 |
| 1.3.1. Create a RAID | 33 |
| 1.3.2. Create a LUN | 24 |
| 1.3.3. Create a Random Access Optimization LUN | 26 |
| 1.3.4. Set Up InfiniBand, iSCSI, Fibre Channel, and SAS..... | 27 |
| 1.3.5. Create a NAS Share | 33 |
| 2. WORKING WITH RAID..... | 36 |
| 2.1. RAID LEVELS | 36 |
| 2.2. CREATE A RAID | 38 |
| 2.2.1. RAID States | 40 |
| 2.3. INITIALIZE A RAID | 41 |
| 2.4. RELOAD A RAID..... | 42 |
| 2.5. MIGRATE A RAID | 43 |
| 2.6. DELETE A RAID..... | 43 |
| 2.7. SET UP ADVANCED RAID RECONSTRUCTION..... | 43 |
| 2.8. SET UP SSD CACHE | 44 |

| | | |
|-----------|--------------------------------------------------------|-----------|
| 2.8.1. | The Amount of the Consumed RAM for the SSD cache | 45 |
| 2.8.2. | Enable SSD Cache..... | 46 |
| 2.8.3. | Disable SSD Cache | 47 |
| 3. | WORKING WITH LUNS | 49 |
| 3.1. | CREATE A LUN..... | 49 |
| 3.2. | RANDOM ACCESS OPTIMIZATION | 51 |
| 3.2.1. | Create a Random Access Optimization LUN | 52 |
| 3.3. | MASKING RULES | 57 |
| 3.3.1. | Target Masking Rules..... | 58 |
| 3.3.2. | Host Masking Rules..... | 61 |
| 3.4. | DELETE A LUN | 64 |
| 4. | NETWORK ATTACHED STORAGE (NAS)..... | 65 |
| 4.1. | MANAGE NAS USERS AND GROUPS | 66 |
| 4.1.1. | Create Users and Groups | 66 |
| 4.1.2. | Edit and Delete NAS Users and Groups | 68 |
| 4.2. | CREATE A NAS SHARE..... | 69 |
| 4.2.1. | Create an SMB Share | 71 |
| 4.2.1.3. | SET UP TRUSTS BETWEEN ACTIVE DIRECTORY DOMAINS..... | 75 |
| 4.3. | EDIT A NAS SHARE | 76 |
| 4.4. | DELETE A NAS SHARE | 76 |
| 4.5. | CREATE, EDIT, AND DELETE A QUOTA | 76 |
| 4.6. | NAS IN DUAL-CONTROLLER MODE | 78 |
| 5. | WORKING WITH HOSTS AND INTERFACES | 80 |
| 5.1. | MANAGE HOST ALIASES | 80 |
| 5.1.1. | Create a Host Alias | 80 |
| 5.1.2. | Delete a Host Alias | 83 |
| 5.1.3. | View the Established Sessions | 84 |
| 5.2. | WORKING WITH ADAPTERS..... | 84 |
| 5.2.1. | Information on iSCSI Ports | 85 |
| 5.2.2. | Information on InfiniBand Adapters | 86 |
| 5.2.3. | Information on Fibre Channel Adapters | 87 |
| 5.2.4. | Information on SAS Adapters | 88 |
| 5.3. | SETTING UP iSCSI | 89 |
| 5.3.1. | Enable/Disable iSCSI Support..... | 89 |
| 5.3.2. | Disabling of iSCSI Initiator with MPIO | 89 |

| | | |
|-----------|--------------------------------------------------|------------|
| 5.3.3. | Create an iSCSI Target | 90 |
| 5.3.4. | Set the Advanced iSCSI Target Parameters | 92 |
| 5.3.5. | Delete an iSCSI target | 94 |
| 5.4. | QoSMIC SET UP..... | 94 |
| 5.4.1. | QoSmic Learning | 95 |
| 5.4.2. | Application Types..... | 98 |
| 5.4.3. | Import of Signatures | 99 |
| 5.4.4. | QoSmic Recognizing..... | 100 |
| 6. | SETTING UP ADVANCED RAID PARAMETERS..... | 102 |
| 6.1. | WORKING WITH ADVANCED RAID PARAMETERS | 105 |
| 6.2. | WORKING WITH RAID PROFILES | 106 |
| 6.2.1. | Create a RAID Profile | 107 |
| 6.2.2. | Modify a Profile | 109 |
| 6.2.3. | Delete a Profile | 109 |
| 7. | WORKING WITH DRIVES | 110 |
| 7.1. | SET UP HOT SPARE..... | 112 |
| 7.2. | SMART-DIAGNOSTICS | 114 |
| 7.3. | DRIVES SCAN | 117 |
| 7.4. | REPLACE A DRIVE | 119 |
| 8. | SETTING UP SYSTEM PARAMETERS | 123 |
| 8.1. | WORKING WITH NODES | 123 |
| 8.1.1. | View the Nodes Current Status..... | 123 |
| 8.1.2. | Automatic Failover..... | 126 |
| 8.1.3. | Manual Failover | 126 |
| 8.1.4. | Restarting and Switching Off the System | 128 |
| 8.1.5. | Disabling Dual-Controller Mode | 128 |
| 8.1.6. | Uninterruptible Power Source Configuration | 129 |
| 8.2. | SETTING UP NETWORK PARAMETERS..... | 130 |
| 8.2.1. | General Networking Settings | 131 |
| 8.2.2. | Set Up Virtual IP-address | 132 |
| 8.2.3. | Set Up Network Interfaces Parameters | 132 |
| 8.2.4. | Set Up Channel Bonding | 133 |
| 8.3. | WORKING WITH A LICENSE..... | 135 |
| 8.3.1. | View the Product License Information | 135 |
| 8.3.2. | Updating the Licensing Parameters | 136 |

| | |
|--------------------------------------------------------------------------------------------|------------|
| 8.4. SETTING UP NOTIFICATIONS..... | 137 |
| 8.4.1. Set Up Server Parameters..... | 137 |
| 8.4.2. Set Up Subscriber Parameters | 138 |
| 8.5. CONFIGURATION RECOVERY | 139 |
| 8.5.1. Save Configuration File | 140 |
| 8.5.2. Download Configuration Files..... | 141 |
| 8.5.3. Upload Configuration Files | 141 |
| 8.5.4. Apply Saved Configuration File | 141 |
| 8.5.5. Delete Configuration File | 141 |
| 8.5.6. Advanced Configuration of the System | 142 |
| 9. SYSTEM MONITORING..... | 147 |
| 9.1. INFORMATION ON SENSORS | 147 |
| 9.1.1. Information on Controller Sensors | 147 |
| 9.1.2. Information on Enclosure Sensors | 148 |
| 9.2. INFORMATION ON PERFORMANCE..... | 152 |
| 9.2.1. Data Rate | 152 |
| 9.2.2. Trace Statistics | 157 |
| 9.3. DOWNLOADING SYSTEM LOGS | 165 |
| 9.4. COMMAND CONSOLE INTERFACE | 165 |
| 9.5. SYSTEM STATUS ASSESSMENT | 166 |
| 10. TROUBLESHOOTING..... | 169 |
| GLOSSARY | 174 |
| APPENDIX A COMMAND CONSOLE INTERFACE FEATURES AND SYNTAX..... | 178 |
| APPENDIX B SETTING UP FIBRE CHANNEL CONNECTION ON GNU/LINUX | 256 |
| APPENDIX C SETTING UP INFINIBAND CONNECTION ON WINDOWS SERVER 2008 R2..... | 257 |
| APPENDIX D SETTING UP INFINIBAND CONNECTION ON SUSE LINUX ENTERPRISE SERVER OS..... | 260 |
| APPENDIX E MPIO SETUP ON SUSE LINUX ENTERPRISE SERVER OPERATING SYSTEM | 261 |
| APPENDIX F MPIO SETUP ON VMWARE ESX 4.1, 5.1, 5.5, 6.0 | 264 |
| APPENDIX H MPIO SETUP ON WINDOWS SERVER 2008 R2 | 273 |

Introduction

Intended Audience

This Guide is intended to help administrators that operate storage systems (from now on, storage) based on the RAIDIX 4.4.3 software.

The Guide provides information on how to configure and administrate the storage system based on the RAIDIX 4.4.3 software.

How the Guide is Organized

This Guide includes the following sections:

1. [Quick Start](#): provides a general description of RAIDIX 4.4.3 web interface, description of quick dual-controller mode configuration, first system start procedure.
2. [Working with RAID](#): provides a detailed description of operations with RAID: the creation of arrays, RAID parameters setup, RAID reload and migration.
3. [Working with LUNs](#): provides a detailed description of operations with LUNs and creation of masking rules.
4. [Working with NAS](#): provides the detailed description of the configuration and work with Network Attached Storage, implemented in RAIDIX 4.4.3.
5. [Working with Hosts and Interfaces](#): provides the detailed description of the managing hosts' aliases and current session's procedure and the procedure of setting up iSCSI that provides detailed information on adapters.
6. [Setting Up Advanced RAID Parameters](#): provides the detailed information about additional RAID parameters that improve system performance and reliability.
7. [Setting Up System Parameters](#): provides the detailed information about nodes, license, sensors, network set up parameters, and so on.
8. [Working with Drives](#): describes the procedure hot spare setup, the procedure of scanning the drive, information about SMART diagnostics, and so on.

9. [System Monitoring](#): provides detailed information about nodes, sensors' states and system performance.
10. [Troubleshooting](#) provides detailed descriptions of errors that may occur in the process of using RAIDIX 4.4.3 and explains how to troubleshoot and fix these errors.

Guide Conventions

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

| Convention | Meaning |
|----------------|-----------------------------------------|
| Bold | Screen names, file names |
| <i>Italics</i> | Term references, additional information |
| Courier New | Commands |

Some of the text elements contain information that needs special attention. These text elements are marked with the following symbols:



Note — a note, which provides valuable information.



Warning — directions, which you must follow to guarantee the proper work of the software.

Contact Us

If you have any comments or suggestions on the improvement of the RAIDIX administrator documentation, please send them to the doc.comments@raidix.com. Specify the document's title and revision number in your letter. Refer to the specific pages and sections whenever possible.

About RAIDIX 4.4.3

The RAIDIX 4.4.3 software is a unique development of RAIDIX for the creation of high-performance data storage systems.

The RAIDIX 4.4.3 software is perfectly suited for tasks with high-performance requirements, reliability, and operational continuity through the use of the parallel computing and unique mathematical algorithms of our design.

RAIDIX 4.4.3 supports a **single-controller** mode when one node is active (Figure 1) and a **dual-controller Active-Active mode**, when two nodes are active and have access to the same set of drives (Figure 2),

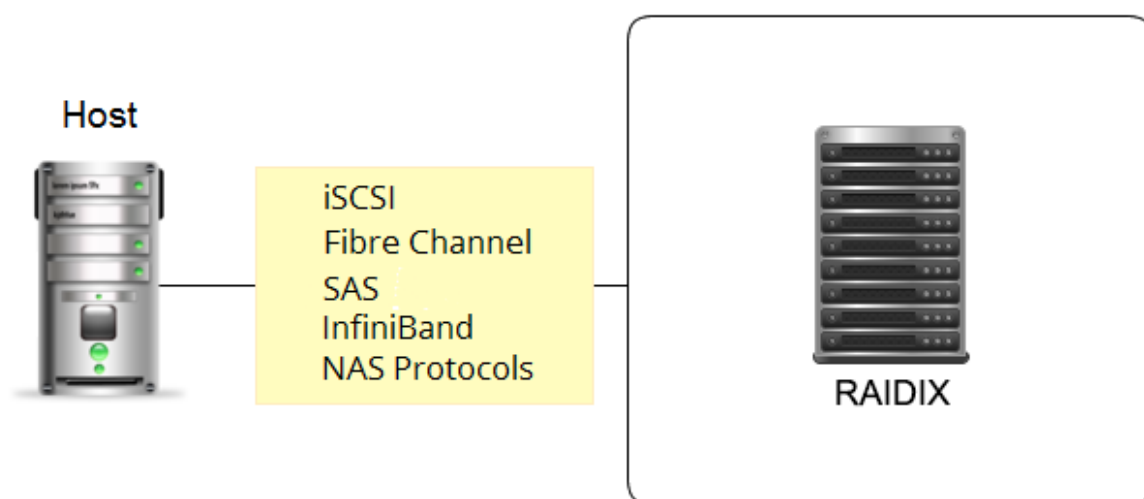


Figure 1. RAIDIX 4.4.3 Single-controller mode

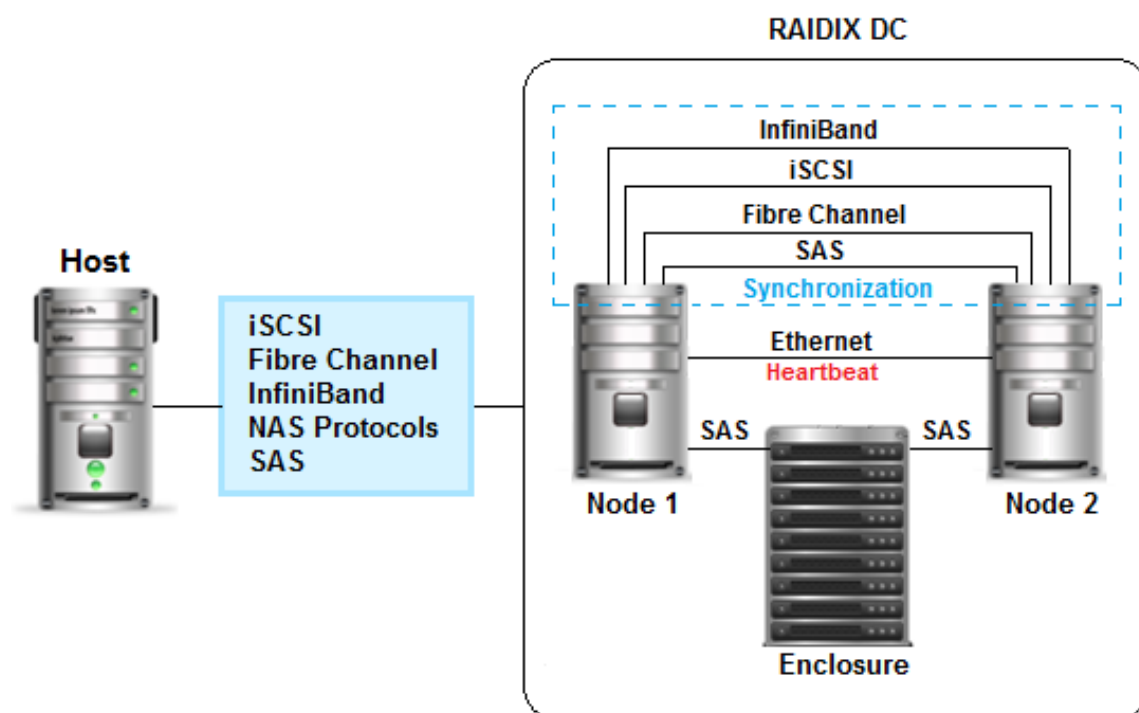



Figure 2. RAIDIX 4.4.3 Dual-controller mode

The RAIDIX 4.4.3 functionality of Network Attached Storage (NAS) supports both: single and dual-controller modes and operates in the Active-Passive mode. For detailed information about working with RAIDIX Network Attached Storage, see [section 5](#) of this document.

You can manage the system from RAIDIX 4.4.3 web interface (Figure 4) or the command console, using CLI commands (detailed description is given in [Appendix A](#)).

RAIDIX 4.4.3 Features

| | |
|--------------------------------------------|-----------------------------------------------------|
| RAID levels | RAID 0, RAID 5, RAID 6, RAID 7,3, RAID 10, RAID N+M |
| Maximum number of drives in one RAID group | 64 |
| Maximum number of drives in the system | 1200 |
| Number of LUNs | 487 |

| | |
|--------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Number of hosts supported (direct connection) | 32 |
| Browsers support | <ul style="list-style-type: none"> • Mozilla Firefox - from 17.0; • Opera - from 11.6; • Google Chrome - from 24.0; • Apple Safari - from 6.0; • Internet Explorer - from 10 (Browser mode: IE10). |
| Virtualization platforms support | <ul style="list-style-type: none"> • VMware ESX Server 4.1 / 5.1 / 5.5; • VMware ESXi Server 6.0; • KVM (Kernel-based Virtual Machine); • RHEV (Red Hat Enterprise Virtualization). |
| Client OS support | <ul style="list-style-type: none"> • MAC OS X (10.6, 10.7, 10.8, 10.9, 10.10), • Windows Server 2008 (32-, 64-bit), • Windows Server 2008 R2, • Windows Server 2012, • Windows XP (32-, 64-bit), • Windows 7 (32-, 64-bit), • Windows 8 (32-, 64-bit), • Solaris 10, • Linux (including but not limited): Red Hat Linux, SuSE, ALT-Linux, CentOS Linux, Ubuntu Linux, Solaris 10. |
| MPIO | <ul style="list-style-type: none"> • MAC OS X 10.7, 10.8, 10.9, 10.10 with X-San component, • Windows Server 2008 (32-, 64-bit), • Windows Server 2008 R2, • Windows Server 2012, • Windows 7 (64-bit), • Windows 8 (64-bit), • VMware ESX Server 4.1/ 5.1/ 5.5, • Linux (including but not limited): Red Hat Linux, SuSE, ALT-Linux, CentOS Linux, Ubuntu Linux, Solaris 10 <div>  <p>Windows 7 and Windows 8 do not support MPIO by default, but you can setup this functionality if needed (for details, see "RAIDIX 4.4 MPIO Setup on Windows 7 and Windows 8 Recommendations")</p> </div> |
| High-speed communication channels support | <ul style="list-style-type: none"> • InfiniBand (FDR, QDR, DDR), • Fibre Channel 8, 16 Gbit/s – in both directions – maximum amount of ports - 24, • Backward compatible with 4Gb and 2Gb Fibre Channel devices, • SAS 12 Gbit/s, • iSCSI (1,10 Gbit/s). |
| NAS protocols support | <ul style="list-style-type: none"> • SMB, • NFS, |

- FTP,
- AFP.

Dual-Controller Mode

The RAIDIX 4.4.3 Dual-Controller mode ensures continuity of data access.

The RAIDIX 4.4.3 software achieve the failure-resistance through the use of two cluster nodes that operate simultaneously in the **Active-Active** mode (Dual-controller mode) and have access to the standard set of drives. Nodes interact over InfiniBand, iSCSI (over Ethernet), and Fibre Channel interfaces which enable synchronization of data and caches state.

Due to the standby 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 transparently takes the entire workload, allowing the administrator to fix errors without stopping the system.



Make sure the cluster nodes have the same hardware configuration before starting to work with the system in dual-controller mode (types and models of processors and adapters, as well as RAM of the nodes in the cluster, must be the same).

In RAIDIX 4.4.3 you can place arrays asymmetrically on the nodes, each RAID will be active on one of the nodes through which access to the RAID resources will be provided. You can specify the array preference parameter (**Affinity**) for every created RAID. It means that you can define on which node RAID will be **Active** (available for performing read and write operations) and on which it will be **Passive** (not available for read and write operations).

The RAIDIX 4.4.3 dual-controller mode allows 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. For more details, see section [2.5 Migrate a RAID](#).

Each node has its IP address (for details, see RAIDIX 4.4.3 Installation and Configuration Guide). To configure dual-controller mode, open the web interface of

each node, using the appropriate IP address, and set up Heartbeat connection parameters separately on every node. For more information, see section 1.2 [Dual-Controller Mode Configuration](#).

Duplication of hardware components and interfaces provides the protection against the following failures:

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

If any of these failures occur, the system will perform automatic nodes switching (**failover**). For more information, see section [7.1.2 Automatic Failover](#). Also, RAIDIX 4.4.3 allows you to perform manual failover (we recommend to use this option during the maintenance and updates installation).

1. Quick Start

This section provides the simple instructions on how to manage RAIDIX 4.4.3 software:

1. [Web Interface Overview.](#)
2. [Dual-Controller Mode Configuration.](#)
3. [Create a RAID.](#)
4. [Create a LUN.](#)
5. [Set Up InfiniBand.](#)
6. [Set Up iSCSI.](#)
7. [Set Up Fibre Channel.](#)
8. [Set Up SAS.](#)
9. [Create a NAS share.](#)

1.1. Web Interface Overview



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



To start working with the system, complete the installation and license configuration. For details, see *RAIDIX 4.4 Installation and Configuration Guide*.

After you installed the system, open the web interface and do the following:

1. Launch any supported web browser. In the address bar, type the IP address of the system using the following format:

`http://{IP address}`



A system administrator must assign the IP address of the system (for details, see *RAIDIX 4.4 Installation and Configuration Guide*)

2. The RAIDIX login dialog will appear. (Figure 3).



Username

Password

Login

Figure 3. RAIDIX login window

3. In the **Username** field type **admin**, in the **Password** field type **raidix**, and click **Login**.

The RAIDIX 4.4.3 dashboard will open. The main page includes information about Volumes, Drives, RAIDs, SparePools and general system status (Figure 4).

RAIDIX Storage NAS Hosts & Interfaces System Monitoring

raidix1 raidix2 ? raidix

Dashboard RAIDs SparePools Drives Drives Scan RAID Profiles RAID Import

Volumes

LUN1, 3 723 GiB Raw Space, 10 073 GiB

Drives

LSI CORP SAS2X36 24 Drives 20 973 GB

7 working 7 TB 11 available 10,07 TB 6 reserved 3,9 TB

RAIDs

| RAID Name | Status | RAID Level | SparePool | Drives | Drives Status |
|-----------|---------------|------------|------------|----------|---------------|
| RAID1 | Active Online | raid7.3 | SparePool1 | 7 Drives | ✓ All |

SparePools

| SparePool Name | Drives | Free Drives | Max Free Size | Linked RAIDs | Status |
|----------------|----------|-------------|---------------|--------------|--------|
| SparePool1 | 6 Drives | 7-12 | 1 000 GB | RAID1 | ✓ |

Figure 4. RAIDIX 4.4.3 Dashboard: dual-controller mode

1.1.1. Software Interface Language

To select interface language in the upper right corner of the RAIDIX 4.4.3 UI, click on the username and choose a language from the drop-down list. The **English**, **Russian**, **Chinese**, and **Korean** languages are available (Figure 5).

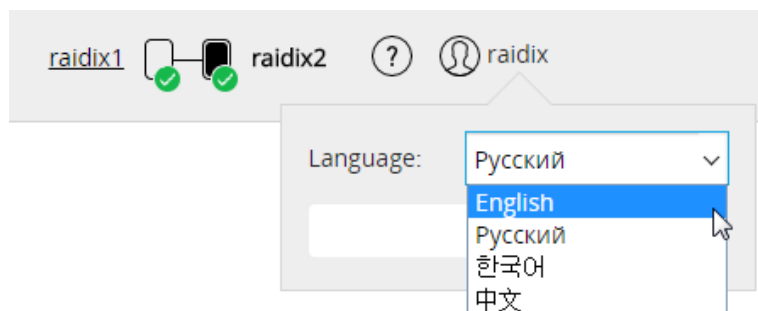


Figure 5. RAIDIX 4.4.3 interface language

1.1.2. Wizards

You can quickly and easily perform initial setup with the wizards available in the RAIDIX 4.4.3 software. The following operations are assisted:

- **RAID Creation.** The wizard is located in the **Storage** menu on the **Dashboard** page (Figure 6).
- **LUN Creation.** The wizard is located in the **Storage** menu on the **Dashboard** page (Figure 6).
- **SparePool Creation.** The wizard is located in the **Storage** menu on the **Dashboard** page (Figure 6).
- **Share Creation.** The wizard is located in the **NAS** menu on the **Shares** page.
- **Configure DC.** The wizard is located in the **System** menu on the **Nodes** page.

The screenshot shows the RAIDIX 4.4.3 Administrator's Guide interface. The top navigation bar includes links for Storage, NAS, Hosts & Interfaces, System, and Monitoring. The main dashboard area shows the following sections:


- Volumes:** A green bar indicating L0, 1 GiB.
- Drives:** A section showing Direct Attached Drives (100 drives, 0 GB) and a summary of 50 working and 50 available drives.
- RAIDs:** A table showing RAID configurations. The table has columns for RAID Name, Status, RAID Level, SparePool, Number of Drives, and Drives Status. The first row shows a RAID named R0, which is Online, RAID0, with a SparePool of —, 50 drives, and All drives are healthy.
- Buttons:** There are buttons for 'Create a RAID', 'Create a LUN', and 'Create a SparePool'.

At the bottom right, the software version is 4.4.3, and it is powered by RAIDIX.

Figure 6 An example of wizards presented on the Dashboard page

1.1.3. Surrounding RAIDIX Systems

The RAIDIX 4.4.3 software implements an option that allows you to discover 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 7).

To access the **Surrounding RAIDIX List** page, in the upper left corner of the menu, click the list icon .

RAIDIX

Storage

NAS

Hosts & Interfaces

System

Monitoring

Francesco

Laura

admin

Surrounding RAIDIX List

Interface

Broadcast Address

Timeout

Discover

| Type | Node Name | Alerts | Cluster Alerts |
|--------------------------------------|-----------------------------------------------|----------------------------------------------------|------------------------------------------------------------------------|
| <div> <div></div> <div></div> </div> | <div> <div>woof</div> <div>bark</div> </div> | <div> <div>1 RAID</div> <div>1 RAID</div> </div> | <div> <div>1 RAID is not synchronized</div> <div>Failover</div> </div> |
| <div> <div></div> </div> | <div> <div>localhost.samba.local</div> </div> | <div> <div>—</div> </div> | <div> <div>—</div> </div> |
| <div> <div></div> </div> | <div> <div>mymds</div> </div> | <div> <div>—</div> </div> | <div> <div>—</div> </div> |
| <div> <div></div> </div> | <div> <div>raidix</div> </div> | <div> <div>1 RAID</div> <div>6 drives</div> </div> | <div> <div>—</div> </div> |
| <div> <div></div> </div> | <div> <div>raidix1.raidix.com</div> </div> | <div> <div>—</div> </div> | <div> <div>—</div> </div> |
| <div> <div></div> </div> | <div> <div>raidix</div> </div> | <div> <div>1 drive</div> </div> | <div> <div>—</div> </div> |

Figure 7 Surrounding RAIDIX List page

To find RAIDIX systems in your network, specify **Interface**, **Broadcast Address**, and **Timeout** (waiting period for RAIDIX systems response) parameters values or just use default values and click **Discover**.

1.2. Dual-Controller Mode Configuration

After you connected the nodes to the one enclosure, installed the system and configured license parameters (for details, see *RAIDIX 4.4 Installation and Configuration Guide*), it is necessary to configure dual-controller mode, [setting Heartbeat connection](#) and [nodes synchronization](#).



After configuring dual-controller mode, you can manage both nodes from the web interface of any of them, by clicking the links to the remote node.

To set up **Heartbeat** connection, do the following:

1. In the main menu click **System**, and move to the **Network** page.
2. In the Network Interface Parameters, select the line of the port, that you want to use for the Heartbit, and click . The pane to set up the network interfaces parameters will appear (Figure 8).

| Port Name | Status | Link Speed | Enabled | DHCP | IP Address | Netmask | MTU |
|-----------|--------|-------------|----------------------------------------|------------------------------|----------------------|----------------------|-----------------------------------|
| eth0 | Up | 100 Mb/s | Yes | No | 172.16.21.172 | 255.255.254.0 | 1500 |
| eth1 | Up | — | <input checked="" type="checkbox"/> On | <input type="checkbox"/> Off | <input type="text"/> | <input type="text"/> | <input type="text" value="1500"/> |
| eth2 | Down | — | Yes | No | — | — | 1500 |
| eth3 | Down | 10 000 Mb/s | Yes | No | 10.0.0.1 | 255.255.254.0 | 1500 |

Figure 8. Setting up Heartbeat connection parameters

- Specify **Enabled** parameter as **On**, specify the **IP Address** and **Netmask** parameters values in the appropriate fields (for more information, see [7.2.3. Set Up Network Interfaces Parameters](#)).



Heartbeat IP-address on both nodes must be specified from one subnet

- Click **Apply** to save settings.
- Repeat steps 1-4 for the second node.
- In the main menu click **System**, and move to the **Nodes** page (Figure 9).

RAIDIX

Storage

NAS

Hosts & Interfaces

System

Monitoring

Nodes

Network

License

Notifications

Configuration Recovery

Nodes

| ID | Name | Heartbeat | RAIDs | Alerts | Node Status | Failover | Power Control |
|----|---------|-----------|-------|---------|-------------|----------|---------------|
| — | raidix1 | — | — | Healthy | — | — | Power On |

[Configure DC](#)

Figure 9. Nodes page before configuring dual-controller mode

- Click **Configure DC**. The pane to add the remote node settings will appear. (Figure 10).

Nodes

| ID | Name | Heartbeat | RAIDs | Alerts | Node Status | Failover | Power Control |
|----|---------|-----------|-------|---------|-------------|----------|---------------|
| — | raidix1 | — | — | Healthy | — | — | Power On |

Remote IP:

[Configure](#) [Cancel](#)

Figure 10. Dual-controller mode configuration

8. In the **Remote IP** field, specify the **Heartbeat** IP-address of the remote node (specified in step 3) and click **Configure**.
9. Parameters for the second node will appear in the table. Both nodes will have automatically assigned **ID** in the **Heartbeat** column (Figure 11).



If after Heartbeat connection is established, the warning that PR (Persistent Reservation) synchronization is broken occur, reset the node (hard reset).

Nodes

| ID | Name | Heartbeat | RAIDs | Alerts | Node Status | Failover | Power Control |
|----|-------|-----------|-------|---------|-------------|----------|---------------|
| 0 | Pinky | 10.0.0.2 | — | Healthy | Healthy | — | Power On |
| 1 | Brain | 10.0.0.1 | — | Healthy | Healthy | — | Power On |

[Disable DC Mode](#)

UPS: ☐ Off
Status: Not connected
[UPS Configuration](#)

Enable Write Through Without Synchronization: ☐ Off

Single Mode: ☐ Off

Figure 11. Nodes page. Dual-controller mode

For the system to work properly in the dual-controller mode, it is necessary to select targets for the cache synchronization over the Fibre Channel, InfiniBand or iSCSI and start OpenSM service (for synchronization over InfiniBand).

1.2.1. Set Up InfiniBand Synchronization

To set up caches' synchronization over the InfiniBand, do the following:

1. In the main menu click **Hosts & Interfaces** and move to the **Adapters** page. The table, where each target corresponds to the one InfiniBand adapter will appear (Figure 12).

The screenshot shows the RAIDIX web interface with the 'Hosts & Interfaces' menu selected. The 'Adapters' sub-tab is active, displaying four tables of network adapter information. On the right, there are buttons for 'Synchronization Targets' and 'OpenSM'.

ISCSI Ports

| Port Name | MAC | IP Address | Netmask | Status | Policy | Link Speed |
|-----------|-------------------|---------------|---------------|--------|--------|------------|
| eth0 | 00:25:90:7b:25:c4 | 172.16.21.101 | 255.255.254.0 | Up | — | 100 Mb/s |
| eth1 | 00:25:90:7b:25:c5 | — | — | Down | — | — |

InfiniBand Adapters

| Device/Port | Port GUID | Node GUID | State | Link Status | Link Speed |
|-------------|---------------------|---------------------|--------|-------------|------------|
| mlx4_0:1 | 0025:90ff:ff17:7799 | 0025:90ff:ff17:7798 | Active | LinkUp | 40 Gb/s |

Fibre Channel Adapters

| Model | WWNN | WWPN | Link Status | Link Speed |
|-----------------------|-------------------------|-------------------------|-------------|------------|
| ATTO Celerity FC-162E | 20:00:00:10:86:03:95:4a | 10:00:00:10:86:03:95:4a | Up | 16 Gb/s |
| ATTO Celerity FC-162E | 20:00:00:10:86:03:95:4b | 10:00:00:10:86:03:95:4b | Up | 16 Gb/s |

SAS Adapters

| State | Link Speed | Phy Number | SAS Address | Port | Connector Number |
|-------|------------|------------|-------------------------|------|------------------|
| Down | Unknown | 0 | 50:06:05:b0:08:91:88:40 | — | 0 |
| Down | Unknown | 1 | 50:06:05:b0:08:91:88:40 | — | 0 |
| Down | Unknown | 2 | 50:06:05:b0:08:91:88:40 | — | 0 |
| Down | Unknown | 3 | 50:06:05:b0:08:91:88:40 | — | 0 |

Figure 12. Adapters page

2. Click **Synchronization Targets** to open a setup window (Figure 13).

Select Targets for Synchronization

×

| Type | WWN/GUID/IQN |
|-------------------------------------------|-------------------------|
| <input type="checkbox"/> IB | 0025:90ff:ff17:7798 |
| <input type="checkbox"/> FC | 21:00:00:10:86:02:94:3e |
| <input type="checkbox"/> FC | 21:00:00:10:86:02:94:3c |
| <input type="checkbox"/> SAS | 500605b0090b24c0 |
| <input type="checkbox"/> SAS | 500605b008918840 |
| <input type="checkbox"/> SAS | 500605b0090b24c4 |
| <input type="checkbox"/> SAS | 500605b008918844 |
| <input checked="" type="checkbox"/> ISCSI | — |

Remote ISCSI IP: 172.16.21.102

Select

Cancel

Figure 13. Selecting Targets for Synchronization

3. Select GUID of the corresponding InfiniBand adapters for nodes cache synchronization and click **Select**.
4. Repeat steps 1-3 on the second node, selecting the appropriate GUID adapters for synchronization.



The best way to set up the sync is the following: leave active InfiniBand connections only for synchronization and specify adapters with LinkUp status as targets for synchronization in the appropriate window (Figure 1).

To start **OpenSM** service and set up InfiniBand cache nodes synchronization, do the following:



Only one subnet manager can run on each InfiniBand fabric, so the configuration steps should be performed only on one node.

1. In the main menu click **Hosts & Interfaces** and move to the **Adapters** page (Figure 12).
2. Click **OpenSM**. A window to specify **OpenSM** service parameters appears (Figure 14).

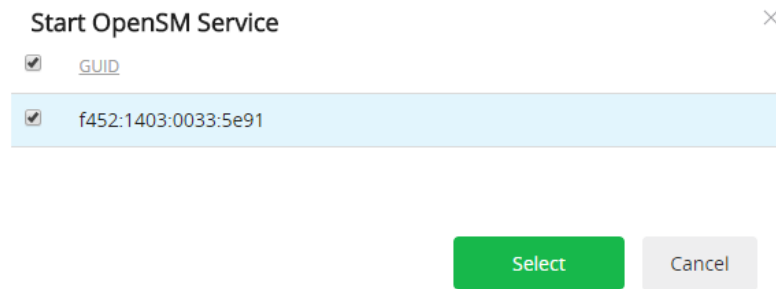


Figure 14. Start OpenSM Service window

3. Select ports for synchronization and click **Select**.

If synchronization of caches is configured correctly, after creating of a RAID array, its synchronization status will change to **On**.

1.2.2. Set Up iSCSI Synchronization

To set up caches' synchronization over iSCSI, do the following:

1. In the main menu click **Hosts & Interfaces** and move to the **iSCSI Targets** page.
2. In the widget, select **On** for the **Enable iSCSI** to turn on **iSCSI:** parameter (Figure 15).

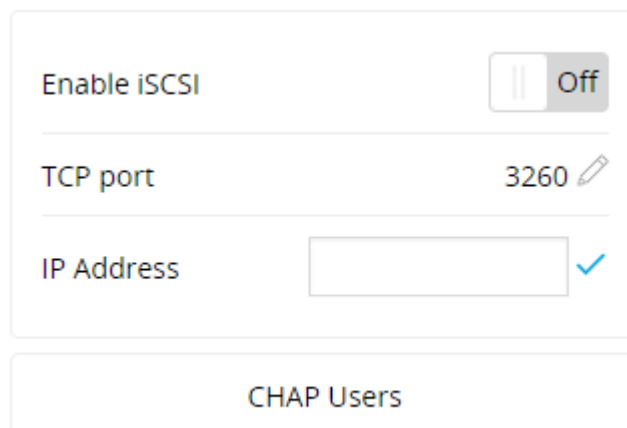
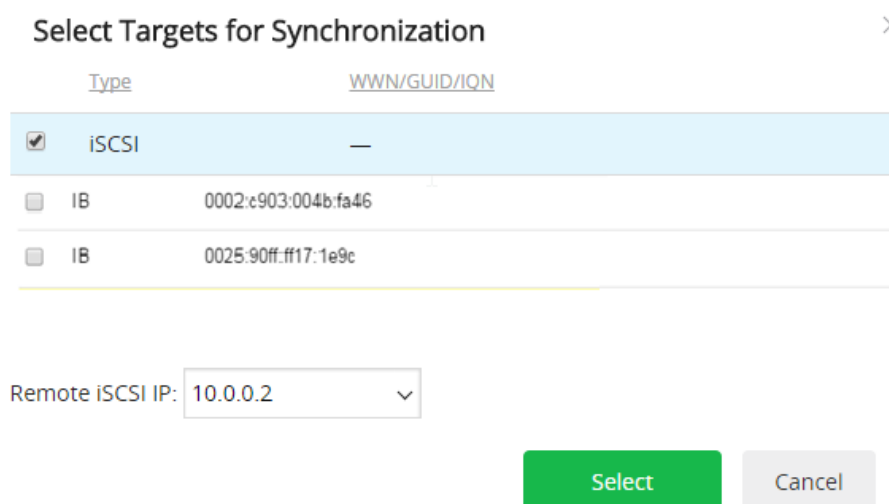


Figure 15. Enabling iSCSI

3. In the main menu, click **Hosts & Interfaces** and move to the **Adapters** page.
4. Click **Synchronize targets** to open a setup window (Figure 16).



| Type | WWN/GUID/IQN |
|-------------------------------------------|---------------------|
| <input checked="" type="checkbox"/> iSCSI | — |
| <input type="checkbox"/> IB | 0002:c903:004b:fa46 |
| <input type="checkbox"/> IB | 0025:90ff:ff17:1e9c |

Remote iSCSI IP: 10.0.0.2

Select Cancel

Figure 16. Targets for synchronization window

1. Select the **iSCSI target type** and **Remote iSCSI IP** for nodes cache synchronization and click **Select**.

If everything was configured correctly, after the creation of a RAID array, its synchronization status would change to **On**.



If desired IP address of the port is not on the list, verify that [network settings of the remote node](#) are correct. To check the network configuration, in the main menu click **System** and move to the **Network** page. Network settings are displayed in the remote node interface.

1.2.3. Set Up Fibre Channel Synchronization



Qlogic Fibre Channel adapters do not support the synchronization of caches in the DC mode.

To set up caches synchronization over the Fibre Channel, do the following:

1. In the main menu click **Hosts & Interfaces**, and move to the **Adapters** page. The table, where each target corresponds to one Fibre Channel port, appears (Figure 12).
2. Click **Synchronize Targets** to open a setup window (Figure 13).

3. Select WWN of the corresponding Fibre Channel ports for nodes cache synchronization and click **Select**.
- .
4. Repeat steps 1-3 on the second node, selecting the appropriate ports for synchronization.



You can set up synchronization by doing the follow: make connections through Fibre Channel active only for synchronization, define adapters that have Up status for the port connection and select them as targets for synchronization in the corresponding window. (Figure 14).

1.2.4. Set Up SAS Synchronization


To set up caches' synchronization over SAS, do the following:

1. In the main menu, go to the **Hosts & Interfaces** and move to the **Adapters** page. The table, where each target corresponds to one SAS ports, appears (Figure 13).
2. Click **Synchronize Targets** to open a setup window.
3. Select **SAS Address** of the corresponding SAS ports for nodes cache synchronization and click **Select**.
4. Repeat steps 1-3 on the second node.

1.3. Getting Started with the Software

1.3.1. Create a LUN


To create a LUN, do the following:

1. In the main menu click **Storage** and move to the **Dashboard** page.
2. In the **RAID List** click  to expand the LUN list of the selected array.
3. Click **Create a LUN**, the **Create a LUN** window will appear (Figure 17).

Create a LUN



RAID "r2"



LUN Parameters

| | | | |
|----------------------|----------|------------|------------------------------|
| LUN Name | LUN Size | Block Size | RAO |
| <input type="text"/> | 1842 GiB | 512 B | <input type="checkbox"/> Off |

LUN is visible to initiators ☒

Create

Cancel

Figure 17. Create a LUN window

- In the **LUN Name** field, specify the LUN name and enable **NAS**, if you are going to create a NAS Share and click **Create**.

A new LUN will be set up with the default parameters (maximum available size and **Block Size** equal to **512 B** (for details about changing LUN settings, see 3.1 "[Create a LUN](#)") (Figure 18).

RAID List

| | RAID Name | RAID Size | Free Size | RAID Level | Drives | Cache | SparePool | Stripe | |
|--|---------------|-----------|-----------|------------|--------|-------|-----------|--------|--|
| | raid0 | 1 862 GiB | 0 GiB | RAID0 | 0,1 | 1 GiB | — | 128 KB | |
| | L1 | 1 862 GiB | | | | | | | |
| | Create a LUN | | | | | | | | |
| | Create a RAID | | | | | | | | |


Figure 18. RAID List section. LUN is created

1.3.2. Create a Random Access Optimization LUN

To create a LUN with [random access optimization](#) (RAO), do the following:



On one RAID you can create only one LUN with RAO. The LUN with RAO must be one and only on the RAID, you must not add any other LUNs or LUNs with RAO on the same RAID.

1. In the main menu click **Storage** and move to the **Dashboard** page.
2. In the **RAID List** click  to expand the LUN list of the selected array.
3. Click **Create a LUN**, the **Create a LUN** window will appear (Figure 17).
4. Enable RAO. You will see the **Advanced Random Access Optimization Parameters** panel (Figure 19).
5. In the **LUN Name** field, specify the LUN name and enable **NAS**, if you are going to create a NAS Share and click **Create**.

A new LUN will be created with the default parameters (For more information, see 3.2.1 «[Create a Random Access Optimization LUN](#)»).

Advanced Random Access Optimization Parameters

Sparse ?

On

Memory ?

0.25

Size ?

BIO Rotation Interval ?

64

Slab Size ?

2

Enable 512e ?

Off

Ack Threads ?

1

BIO Threads ?

4

Block Map Cache Size ?

128

Block Map Period ?

16380

CPU Threads ?

2

Deduplication ?

On

Logical Threads ?

1

Physical Threads ?

1

Read Cache Size ?

0

Write Policy ?

Auto

Apply

Cancel

Figure 19 RAO LUN settings panel

1.3.3. Set Up InfiniBand, iSCSI, Fibre Channel, and SAS

The way the connection is initiated through InfiniBand, iSCSI, Fibre Channel, and SAS depends on the selected mode of access of the initiators to the LUNs (the type of host access to LUN) and is called masking policy. By default, the system works in **Standard** mode: after LUN is created, it is available to all initiators. The system administrator can change the masking policy and select **Enterprise** mode: after LUN is created, it is not available to all hosts. To check or change the current access type, see section 3.2 [“Masking Rules.”](#)

1.3.3.1. Set Up InfiniBand

To set up the InfiniBand connectivity, perform the following steps:



To set up the InfiniBand connectivity for **Enterprise** policy, refer to steps 1-11.

To set up the InfiniBand connectivity for **Standard** policy, refer to steps 1-3.

1. Connect initiator to a storage system.
2. In the main menu click **Host & Interfaces** and move to the **Adapters** page.
3. Click the **OpenSM**, the window to start OpenSM (InfiniBand Subnet Manager) will open (Figure 20).




Figure 20. Start OpenSM Service window

4. Select ports for connection with the host and click **Select**.
5. Open the Hosts & Interfaces menu.
6. In the **Hosts Aliases** section click **Create a Host Alias**. The pane to create a new host will open. (Figure 21).

Hosts Aliases

| Host Alias | Host (WWN/GUID/IQN) | Realtime | LUN List |
|------------|---------------------|----------|----------|
| All | ✓ All Hosts | Off | — |

Alias Name: Host (WWN/GUID/IQN): Realtime: ☐ Off [Select](#)

Figure 21. Hosts page. Host Alias creation

7. In the **Alias Name** field type the host alias name.
8. In the **Host (WWN/GUID/IQN)** field type host GUID — a unique InfiniBand ID.
9. Assign priority to host alias (**Realtime** parameter) (for more information, see [section 4.1](#)).
10. In the **LUN** field click **Add a LUN** to add a LUN to the appropriate host alias.
11. Assign LUNs to an appropriate host alias by checking the corresponding checkbox and click **Select**. The selected LUNs appear in the host alias pane.
12. Click **Create**. The Host Alias appears (Figure 22).

Hosts Aliases

| Host Alias | Host (WWN/GUID/IQN) | Realtime | LUN List |
|------------|--------------------------------------|----------|----------------------------|
| Alias1 | ✓ 0x005214030022bca00002c903004bf2f2 | Off | lun1 (R/W) |
| All | ✓ All Hosts | Off | — |

[+ Create a Host Alias](#)

Figure 22. Host Alias for InfiniBand is created

1.3.3.2. Set Up iSCSI

To set up the iSCSI connectivity, do the following:



To set up the iSCSI connectivity for **Enterprise** policy, refer to steps 1-13.

To set up the iSCSI connectivity for **Standard** policy, refer to steps 1-5.

To set up the iSCSI connectivity, perform the following steps:

1. Connect Initiator to a storage system.
2. Open the **Hosts & Interfaces** menu and move to the **iSCSI Targets** page.
3. To turn on iSCSI, select **On** for **Enable iSCSI** parameter in the wizard (Figure 23).

A screenshot of the iSCSI configuration wizard. It features three main sections: 'Enable iSCSI' with a toggle switch currently set to 'Off'; 'TCP port' with the value '3260' and a pencil icon for editing; and 'IP Address' with an empty text input field and a blue checkmark icon. Below these sections is a button labeled 'CHAP Users'.

Figure 23. Enabling iSCSI

4. If necessary, in the **IP Address** field, enter the IP address of the interface to detect the iSCSI target.

If you leave this field blank, the iSCSI targets will be available for all interfaces.



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 a Target**. The panel to create an iSCSI will open (Figure 24).

iSCSI Targets are not created

| Target Name | CHAP Mode | IP Filter | LUN List | Login | Password |
|--------------------------------|-----------|---------------------------|---------------------------|-------|----------|
| iqn.2015-04.com.raidix:target0 | None | Add an IP | + Add LUN | | |

[Create](#) [Cancel](#)

Figure 24. iSCSI page. Target creation

6. Define the following parameters and click **Create**:

- **Target Name:** type IQN – a unique iSCSI ID,
- **CHAP Mode:** specify the authorization mode used by host and target to authenticate each other via the exchange of a secret code or password: **None**, **Unidirectional**, **Bidirectional** (for details refer to section 6.7.2 “[Create an iSCSI Target](#)”),
- **IP Filter:** assign a host (hosts) by adding its IP-address to the masking rule.
- **LUN List:** select the LUN, which will be available for the target.

7. Created target appears on the **iSCSI** page (Figure 25).

iSCSI Targets

| Target Name | CHAP Mode | IP Filter | LUN List | Login | Password | |
|--------------------------------|---------------|-----------|-------------------------------------------------------------------|---------|----------|-------------------------------------|
| iqn.2015-04.com.raidix:target0 | None | All | — | — | — | ✎ ✕ |
| iqn.2015-04.com.raidix:target1 | Bidirectional | All | ada ata lsi | timur88 | ***** | ✎ ✕ |

[+ Create a Target](#)

Figure 25. 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**.

8. Open the **Hosts & Interfaces** menu.

9. In the **Hosts** section click **Create a Host Alias**. The pane to create a new host will appear (Figure 21).
10. In the **Host Alias** field type the host alias name.
11. In the **Host (WWN/GUID/IQN)** field type IQN – a unique iSCSI ID of the host.
12. Assign priority to host alias (**Realtime** parameter) (for more information, see section 5.1).
13. In the **LUN** field click **Add a LUN** to add a LUN to the appropriate host alias. The **Select LUN** window will appear.
14. Assign LUNs to a suitable host alias by checking the corresponding checkbox and click **Select**. The selected LUNs will appear in the host alias pane.
15. Click **Create**. The Host Alias will appear (Figure 26).




| RAIDIX | | | | |
|------------------------------------------------------|------------------------------------|----------|----------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Storage NAS Hosts & Interfaces System Monitoring (1) | | | | |
| Hosts Adapters iSCSI Targets | | | | |
| Hosts Aliases | | | | |
| Host Alias | Host (WWN/GUID/IQN) | Realtime | LUN List | |
| Alias1 | iqn.1994-05.com.redhat:a9b5d49372a | On | <div> <div>11 (R/W)</div> <div>14 (R/W)</div> <div>13 (R/W)</div> <div>12 (R/W)</div> </div> |   |
| All | All Hosts | Off | — |  |
| + Create a Host Alias | | | | |

Figure 26. Host Alias for iSCSI is created

1.3.3.3. Set Up Fibre Channel

To set up a Fibre Channel connectivity for **Enterprise** policy (when after creation, LUNs are not available to any host), perform the following procedure:



In the **Standard** mode, a Fibre Channel connection will be established automatically

1. Connect host to the storage.
2. Open the **Hosts & Interfaces** menu and move to the **Hosts Aliases** page.
3. Click **Create a Host Alias**. The pane to create a new host alias will open (Figure 21).
4. In the **Host Alias** field type the host alias name.

5. In the **Host (WWN/GUID/IQN)** field type WWN – a unique Fibre Channel host ID.
6. Assign priority to host alias (**Realtime** parameter) (for more information, see [section 5.1](#)).
7. In the **LUN list** field click **Add a LUN** to add a LUN to an appropriate host alias.
8. Assign LUNs to a suitable host alias by checking the corresponding checkbox and click **Select**. The selected LUNs appears in the host alias pane.
9. In the drop-down list of the corresponding LUN define the host access policy to the LUN and click **Create**. The host alias will be created.

| RAIDIX | | | | |
|-------------------------------------------------------------|---------------------|----------|----------------------------------------------|-----|
| Storage NAS Hosts & Interfaces System Monitoring | | | | |
| Hosts Adapters iSCSI Targets | | | | |
| Hosts Aliases | | | | |
| Host Alias | Host (WWN/GUID/IQN) | Realtime | LUN List | |
| Alias1 | ✓ 100000108603a1ea | On | 11 (R/W) 14 (R/W) 13 (R/W) 12 (R/W) | ✎ ✕ |
| All | ✓ All Hosts | Off | — | ✎ |
| ⊕ Create a Host Alias | | | | |

Figure 27. Host Alias for Fibre Channel is created

1.3.3.4. Set Up SAS



In RAIDIX 4.4.3 you can enable the SAS connection only in the **single-controller** mode

After performing initial SAS configuration including wide ports and SAS modes configuration (for details refer to the *RAIDIX 4.4 Installation and Configuration Guide*), continue configuration by setting up SAS connection for **Enterprise** policy (note, that after you created LUNs, they will be unavailable to any host):



In the **Standard** mode, a connection over SAS will be established automatically

1. Connect host to the storage.
2. Open the **Hosts & Interfaces** menu and move to the **Hosts Aliases** page.
3. Click **Create a Host Alias**. The pane to create a new host alias will open (Figure 21).

4. In the **Host Alias** field type the host alias name.
5. In the Host (WWN/GUID/IQN) field type host's SAS address.
6. Assign priority to host alias (**Realtime** parameter) (for more information, see [section 5.1](#)).
7. In the **LUN list** field click **Add LUN** to add a LUN to an appropriate host alias.
8. Assign LUNs to a suitable host alias by checking the corresponding checkbox and click **Select**. The selected LUNs will appear in the host alias pane.
9. In the drop-down list of the corresponding LUN define the host access policy to the LUN and click **Create**. The host alias will be created.

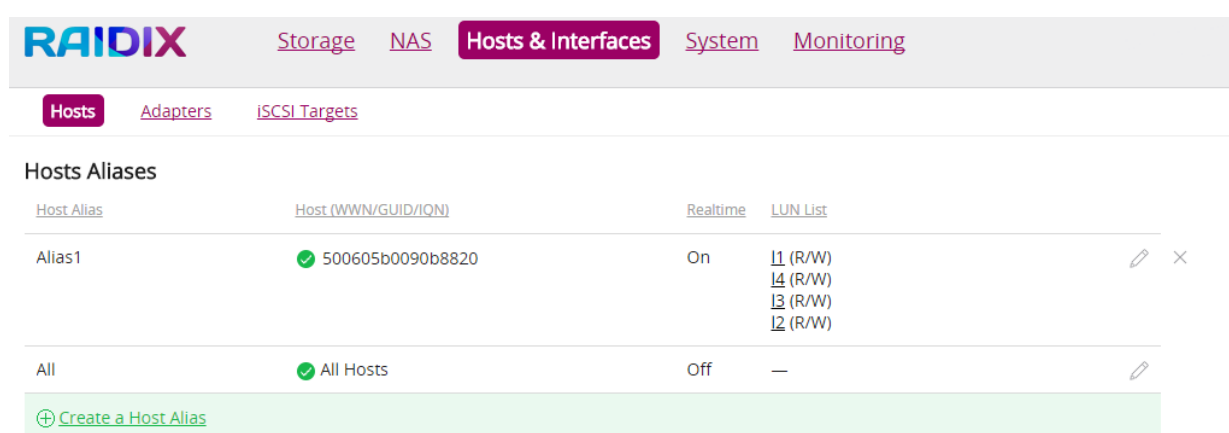


Figure 28. Host Alias for SAS is created

1.3.4. Create a RAID

To create a RAID array, do the following:

1. In the main menu click **Storage** move to the **Dashboard** page.
2. Click **Create a RAID**. The panel to create an array will open (Figure 29).

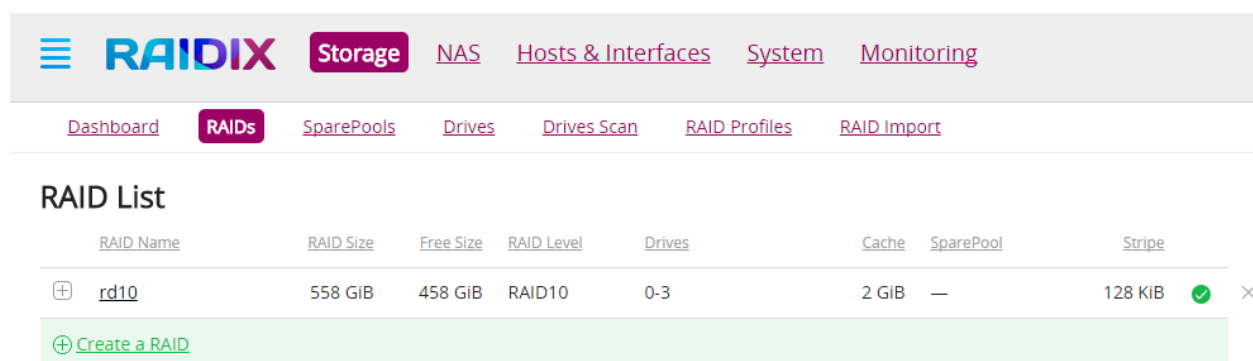


Figure 29. The Storage page. Creating a RAID. Dual-Controller Mode

3. In the pane, define the following parameters:

- In the **RAID Name** field type the array name.
- In the **RAID Level** field, select the level of the array (for details, see 2.1 “[RAID Levels](#)”).
- In the **Drives** field, click **Select** and select the drives from the list.
- Click **OK**.
- In the **Cache** field, define cache size.

Click **Create**. In the **RAID List** the created array will be displayed. (Figure 30).




Figure 30. An example of creating RAID N+M in Dual-Controller mode

1.3.5. Create a NAS Share

To create a NAS share, do the following actions:

1. Create a **RAID** and a **LUN** on it with **enabled NAS** parameter.
2. Open the **NAS** menu and move to **Users and Groups** page. Add information about NAS users (for more details, see section 4.1.1 “[Create Users and Groups](#)”).
3. Open the **Shares** page and click **Format the LUN**.
4. In the opened window select a file system and click **Select**.
5. In the **Shares List** section click **Create a Share**. A pane to create a new share will appear (Figure 31).

RAIDIX
Storage
NAS
Hosts & Interfaces
System
Monitoring

Shares
Users & Groups
Quotas

Shares List

L1
300 GiB
Free Size: —
RAID "rd1"
File System: EXT4
Format the LUN
[Create a Quota](#)

| Share Name | Path | Protocol | Squash | Browsable | WORM | Guest | Users & Groups | Hosts |
|-----------------------------------------|---------------------------------------------------------------|----------|-----------------------------|-----------|------|-------|----------------|-------------------------|
| <input type="text" value="Share Name"/> | <input type="text" value="/"/> Select Path | NFS | <input type="checkbox"/> No | | | | | 0 Hosts |
| <div> Create Cancel </div> | | | | | | | | |

Figure 31. Creating a share on the NAS page

- Specify the **Share Name** in the appropriate text field.
- Specify other parameters values (for more, see section 4.2 "[Create a NAS share](#)") and click **Create**.
- In **Shares List** section the created share with specified parameters will appear (Figure 32).

RAIDIX
Storage
NAS
Hosts & Interfaces
System
Monitoring
[Lambri](#)

Shares
Users & Groups
Quotas

Shares List

L1
136 GiB
Free Size: —
RAID "R1"
File System: XFS
Format LUNs
[Create a Quota](#)

| Share Name | Path | Protocol | Squash | Browsable | WORM | Guest | Users & Groups | Hosts |
|----------------------------------|------|----------|--------|-----------|------|----------|----------------|--------|
| Share1 | / | SMB | — | Yes | No | Disabled | 1 0 | 1 Host |
| + Create a Share | | | | | | | | |

Figure 32. Created Share is displayed

2. Working with RAID

This section of the document describes [RAID levels](#) and provides the following information on operations with arrays:

1. [RAID Levels](#)
2. [Initialize a RAID](#)
3. [Reload a RAID](#)
4. [Migrate a RAID](#)
5. [Delete a RAID](#)
6. [Set Up Hot Spare](#)
7. [Estimate Drives State](#)
8. [Replace a Drive](#)
9. [Set Up Advanced RAID Reconstruction](#)
10. [Set Up SSD Cache](#)

2.1. RAID Levels

There are several methods of combining hard drives, and each of them has its pros and cons. These methods are called RAID levels. Fault tolerance and performance of the system will be defined by the RAID level.

RAIDIX 4.4.3 software allows creating the following RAID levels: RAID 0, RAID 5, RAID 6, RAID 7.3, RAID 10 and RAID N+M. Each RAID scheme offers a different balance of performance, data protection and storage efficiency.

RAID 0 — the level of interleaving blocks without mirroring. The data blocks are distributed across several drives. Data are in parallel access mode that provides high performance. However, due to the lack of redundancy, RAID 0 does not provide data reliability, the failure of one drive in RAID leads to the whole array degradation.

RAID 5 — the level of interleaving blocks with distributed parity. RAID 5 requires at least five disks. RAID 5 can sustain 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. Both data and redundancy information are duplicated. RAID 6 is characterized by 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 disks for checksums allocation. RAID N+M requires at least eight disks and can sustain complete failure up to **32** drives in the same group (depending on the number of parity disks).

RAID 10 — the architecture of this mirrored array represents RAID 0, the segments of which are RAID 1 arrays, instead of separate drives. Accordingly, this array level must contain at least four drives. Data integrity is maintained in case of a half drives failure, the irreversible array destruction occurs when two drives of one mirrored pair fail.

Main RAID characteristics are summarized in the table below, where **Storage efficiency** is the ability to store and manage data that consumes the least amount of space with little to no impact on performance; resulting in a lower total operational cost.

| RAID Level | Storage Efficiency | Read Performance | Write Performance | Data Protection | Number of Drives |
|------------|--------------------|------------------|-------------------|-----------------|------------------|
| RAID 0 | 100% | High | Very high | No | 1-64 |
| RAID 5 | 80-98% | Very high | Very high | Yes | 5-64 |
| RAID 6 | 60-97% | Very high | Very high | Yes | 5-64 |
| RAID 7.3 | 40-95% | Very high | Very high | Yes | 5-64 |

| RAID Level | Storage Efficiency | Read Performance | Write Performance | Data Protection | Number of Drives |
|------------|--------------------|------------------|-------------------|-----------------|-----------------------------------|
| RAID N+M | 50-94% | Medium - High | Medium - High | Yes | 8-64 |
| RAID 10 | 50% | High | High | Yes | 2-64 (only even number of drives) |

Besides that, RAIDIX allows creating arrays of RAID 5 Initialized (RAID5i), RAID6 Initialized (RAID6i), RAID 7.3 Initialized (RAID 7.3i) and RAID N+M Initialized (RAID N+Mi).

Distinctive features of initialized RAID are:



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

For detailed information see section 5.4. "[Setting Up Advanced RAID Parameters.](#)"

2.2. Create a RAID


To create a RAID, perform the following procedure:

1. Open the **Storage** menu and move to the **RAIDs** page.
2. In the **RAID List** section click **Create a RAID**. The pane to create a RAID appears
3. (Figure 33).

The screenshot shows the RAIDIX web interface. At the top, there's a navigation bar with 'Storage' selected. Below it, a sub-navigation bar shows 'RAIDs' as the active tab. The main content area is titled 'RAID List is empty'. Below this, there's a table with columns: RAID Name, RAID Size, Free Size, Level, Drives, Cache, SparePool, Stripe, Sync, and Affinity/Status. Below the table, there's a form to create a new RAID. The form has a 'RAID Name' field, a 'RAID Size' dropdown (set to 1 GiB), a 'Level' dropdown (set to RAID0), a 'Drives' dropdown (set to 128 KB), a 'Cache' dropdown (set to 128 KB), a 'Stripe' dropdown (set to 128 KB), a 'Sync' dropdown (set to 128 KB), and an 'Affinity/Status' dropdown (set to 128 KB). There's also a 'Use Profile' dropdown (set to default). At the bottom right of the form, there are 'Create' and 'Cancel' buttons.

Figure 33. RAID creation pane. Dual-controller mode

- Learn the meaning of the following RAID parameters on the **RAID List** and specify the required values:

| Parameter | Description |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| RAID Name (required) | A unique array name |
| RAID Size | Array size (GiB) |
| Free Size | Free space on RAID (not occupied by LUNs) (GiB) |
| RAID Level (required) | <p>Possible values: RAID 0, RAID 5, RAID 5i, RAID 6, RAID 6i, RAID 7.3, RAID 7.3i, RAID 10, RAID N+M, and RAID N+Mi (for details see 2.1 “RAID Levels”).</p> <p>If you create initialized RAID, the initialization process starts; its completion is displayed in percent. Working with an array is impossible until the initialization is complete.</p> |
| Drives (required) | <p>Drives list.</p> <p>Type the numbers of drives separated by commas or dash or click Select to check the numbers of drives in any order.</p> |
| <div>  It is recommended to include drives of the same size in one RAID </div> | |
| Cache | Cache size (GiB). The system automatically sets 1 GiB by default |
| SparePool (optional) | <p>Name of a set of drives, which will be used for hot spare in case of RAID drive failure (not allowed for RAID0).</p> <p>You can assign a SparePool to an array from page SparePools.</p> |
| Stripe | The minimum size of the data stripe (in kilobytes) to be recorded on a drive in the array. The system automatically sets 128KiB by default |
| Sync | <p>Cache synchronization status. The option is available only in dual-controller mode. The possible values are:</p> <ul style="list-style-type: none"> On – caches of the specified RAID on both nodes are synchronized; Off – caches of the specified RAID on both nodes are not synchronized; Error – an error occurred during the array synchronization. |
| Affinity/Status (required) | <p>ID of the node, from which it will be possible to perform read and write operations with RAID (node has preferable access to RAID).</p> <p>The option is available only in dual-controller mode. Nodes identifiers are displayed on RAIDIX 4.4.3 main page.</p> |

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|



By default, **Affinity** corresponds to the node ID, on which the array is created. You can also change **Affinity** value by performing RAID migration to the specified node. See more information in section 2.5 ["Migrate a RAID."](#)

| | |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Use a Profile (optional) | Profile of RAID parameters that will be applied to the array (settings for cache, SDC, etc.). The system uses the Default profile settings automatically (for more information about managing RAID parameters refer to Setting Up Advanced RAID Parameters). |
|--------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

5. Click **Create**. The array appears in the **RAID List** section (Figure 34).

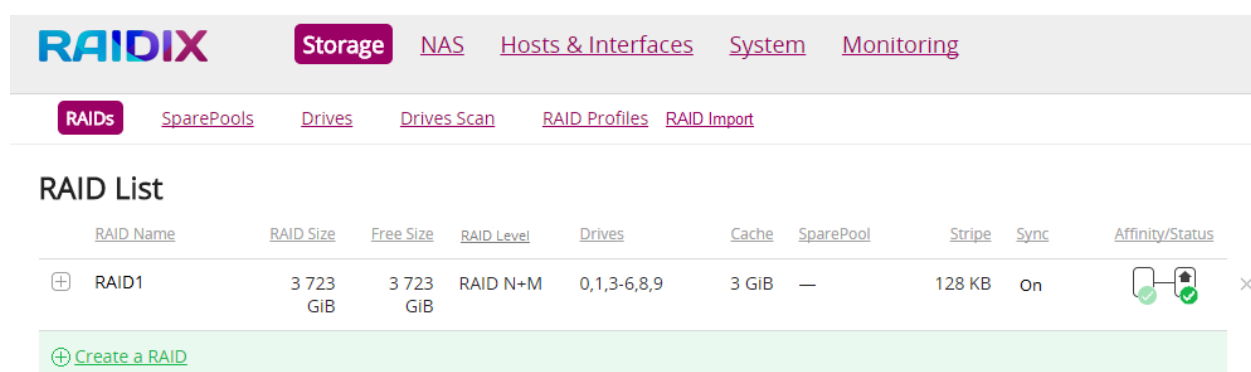


Figure 34. An example of creating a RAID N+M in dual-controller mode

2.2.1. RAID States

After creating a RAID, the current array status will be displayed (translucent filling of the icon corresponds to passive RAID states). Possible variants are:



Online is displayed in the following cases:

- **Active** – the array is active on the node, specified in its Affinity. All array drives work correctly.
- **Passive** – the array is passive on the node, which is not specified in its Affinity. All array drives function properly.



Warning displays in the following states:

- **Degraded** – the array is available and operating, but some of the drives are removed from the system or faulty;
- **Failed over** – the array that was active on the node, specified in its Affinity, was failed over (transferred) to the second node. As a result, the array became temporarily active on the second node and passive on the first;

- **Initializing** – an array is being initialized. During the procedure, you cannot work with the array. The completion status is displayed in percent.
- **Passive** – the array is passive on the node, specified in its Affinity;
- **Reconstructing** – one or several drives are currently reconstructing, but you can operate with the array;
- **SDC** – the array contains corrupted data;
- **Transitioning** – the array is switching to the passive mode.



Bad displays in the following states:

- **Offline** – the array is offline (not available);
- **Not Loaded** – the configuration exists, but fails to be loaded. To load the array configuration, use one of the following ways, described in section [Troubleshooting](#).

2.3. Initialize a RAID

The initialization procedure is started automatically when creating RAID 5i, RAID 6i, RAID 7.3i, and RAID N+Mi (Figure 35).

| RAID Name | RAID Size | Free Size | RAID Level | Drives | Cache | SparePool | |
|-----------------|-----------|-----------|------------|-----------|-------|-----------|--------------------------------|
| ⊕ r1 | 2 792 GiB | 0 GiB | RAID0 | 0,1,3 | 1 GiB | — | 128 KB Online Initializing: 1% |
| ⊕ r2 | 1 116 GiB | 1 116 GiB | RAID5i | 2,4,5,7,8 | 1 GiB | — | 128 KB |
| ⊕ Create a RAID | | | | | | | |

Figure 35. The initialization process has started

When initialization is completed, **Random Writes** and **Silent Data Corruption** parameters will become available for change from the **Advanced RAID Parameters** window.

To stop initialization process, you can use the following CLI command:

```
rdcli raid init stop -n <raid_name_value>
```

To resume initialization, use the following CLI command:

```
rdcli raid init start -n <raid_name_value>
```

2.4. Reload a RAID

In RAIDIX 4.4.3 software RAID reload must be performed to apply cache size and static array parameters changes.

Also, if the array changed its status to **Not loaded**, a RAID reload can help to bring it back to a normal state. See more details in section [8. «Troubleshooting»](#).

You can reload the RAID from the system web interface, as well as from CLI. Detailed commands description is given in [Appendix A](#).



It is not recommended to reload a RAID array while the system is working under a load

To reload the RAID array from web interface, do the following:

1. In the **Storage** menu, select the **RAIDs** tab. In the **RAID List** click the corresponding array name. The RAID page opens (Figure 36).

RAID "sdf"

luna, 400 GiB Free, 3 323 GiB

LUN List

| LUN Name | LUN Size | Block Size | NAS | ID | Type | Serial Number |
|----------|----------|------------|-----|----|------|---------------------------|
| luna | 400 GiB | 512 B | No | 0 | — | 0000000000000000096df878a |

[Add LUN](#)

Drives List

| UID | Capacity | Vendor | Model | Serial | ID |
|-----|----------|---------|------------|--------------|-----|
| 5 | 1 000 GB | TOSHIBA | MK1001TRKB | 3200A06DFRC5 | sdx |
| 9 | 1 000 GB | TOSHIBA | MK1001TRKB | Z150A0BTFM16 | sdv |
| 11 | 1 000 GB | TOSHIBA | MK1001TRKB | 3200A07ZFRC5 | sdn |
| 14 | 1 000 GB | TOSHIBA | MK1001TRKB | Z150A0F4FM16 | sdl |

"sdf" Info

Online, Passive
 Sync: On
 Level: RAID0
 Cache: 1 GiB
 No SparePool
 Random Reads: On
 Random Writes: Off
 SSD Cache Drive: —

[Show RAID parameters](#)

Select Drives for SSD caching

Reload the RAID

Migrate the RAID

Delete the RAID

[Drives Scan](#)

Figure 36. RAID reload

2. Click **Reload the RAID**. Confirm your decision in the appeared dialog box by clicking **Yes**.

2.5. Migrate a RAID

For nodes workload balancing, perform RAID migration from any node of the cluster. The **Affinity** parameter value changes to the opposite during the RAID migration procedure. To migrate the array, do the following:

1. In the **Storage** menu, select the **RAIDs** tab. In the **RAID List** click the corresponding array name. The RAID page opens (Figure 36).
2. Click **Migrate the RAID**. Confirm your decision in the appeared dialog box by clicking **Yes**.

2.6. Delete a RAID

To delete a RAID, do the following:

1. In the **Storage** menu, select the **RAIDs** tab. In the **RAID List** click the corresponding array name. The RAID page opens (Figure 36).
2. Click **Delete the RAID**. Confirm your decision in the appeared dialog box by clicking **Yes**.

2.7. Set Up Advanced RAID Reconstruction

In RAIDIX 4.4.3 software also implemented the ability to define advanced RAID reconstruction parameters.

To configure the advanced RAID reconstruction parameters, open the **Storage** menu and move to the **RAIDs** page. In the right part of the page click **Driver Parameters**; the widget expands (Figure 37):

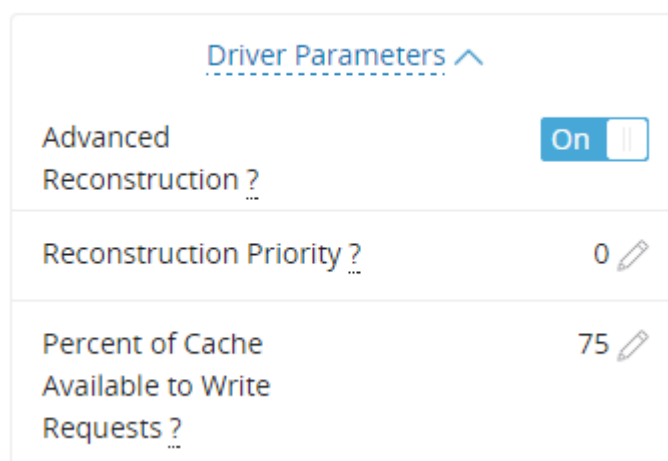


Figure 37. Driver Parameters widget

RAIDIX 4.4.3 provides the following modes of **Advanced Reconstruction**:

- **On:** advanced reconstruction is enabled. Parameter allows optimizing reading speed by eliminating from the process up to two (in the case of RAID 7.3 – up to three) drives with lower reading speed or which are in a process of their software data restore.
- **Off:** function is disabled.



It is not recommended to enable **Advanced Reconstruction** mode if the system contains a large number of drives and RAIDs



Enabling **Random Write** and **Silent Data Corruption** functions stops **Advanced Reconstruction** function

2.8. Set Up SSD Cache

Cache is a component placed in front of a primary device to transparently store data so that future data requests can be served faster. The more requests that can be served from cache, the faster is the overall system performance. As **SSDs are adopted for random read workload** and have sufficient IO speeds, it allows a host to read the cached data much faster than if it had to read the same data directly from the hard drive.

The SSD cache is configured to function as a secondary I/O cache. The increase in performance varies with workload and the configured cache capacity. Systems with **random read** workload benefit the most.

2.8.1. The Amount of the Consumed RAM for the SSD cache

The amount of the consumed by the SSD cache RAM memory, depends on the size of the SSD drive.

To calculate the amount of the RAM memory for the SSD cache, the RAIDIX 4.4.3 system uses the following formula:

$$[\text{The number of megabytes in gigabytes}] * [\text{The number of gigabytes on your disk}] / 450$$

Where 450 is the coefficient of the memory that the system allocates for the SSD cache hash table. The size of consumed memory cannot be less than 150 MB.

For example, for the 130 GB SSD disk, the amount of the consumed memory will be $1024 * 130 / 450 = 295$ MB.

2.8.2. Enable SSD Cache

RAIDIX 4.4.3 software implements **SSD Cache** to speed up access to data. For SSD cache can be used:

- an SSD drive;
- MD RAID composed of SSD drives.

To enable SSD cache using an SSD drive, perform the following:



Only one SSD drive can be used as SSD cache for a RAID group.

Each RAID group should have its SSD drive.

SSD drive that is already utilized in a RAID group cannot be assigned for caching.

In DC mode: If SSD drive is used on one node, it cannot be used on another node.

1. Open the **Storage** menu and move to the **RAIDs** page.
2. Click on the name of the appropriate RAID to access to its page. In the right part of the page, click **Select Drives for SSD Cache** (Figure 38).

The screenshot shows the RAIDIX interface for RAID 'r1'. The top navigation bar includes 'Storage', 'NAS', 'Hosts & Interfaces', 'System', and 'Monitoring'. The left sidebar shows 'Dashboard', 'RAIDs', 'SparePools', 'Drives', 'Drives Scan', 'RAID Profiles', and 'RAID Import'. The main content area displays RAID 'r1' with three LUNs: lun2 (900 GiB), lun1 (900 GiB), and lun3 (992 GiB). Below the LUN list is a table with columns: LUN Name, LUN Size, Block Size, NAS, ID, Type, and Serial Number. The table lists lun1, lun2, and lun3 with their respective details. To the right of the RAID configuration, there is a 'r1' Info panel showing RAID Level: RAID0, Cache: 1 GiB, No SparePool, Random Reads: Off, Random Writes: Off, and SSD Cache Drive: —. Below the info panel, there is a red box highlighting the 'Select Drives for SSD Caching' button. Other buttons visible include 'Reload the RAID', 'Delete the RAID', and 'Drives Scan'.

Figure 38. RAID page. Selecting drives for SSD Cache

3. In the opened window select an SSD drive to be used as SSD cache, and click **Select** (Figure 39).




Figure 39. Selecting Drives for SSD Cache window



In DC mode: Select other SSD drive for SSD cache on the second node.

To enable SSD cache using MD RAID composed of SSD drives, perform the following:

1. Create an MD RAID from SSD drives, using the following command:

```
mdadm --create /dev/md0 --level=0 --chunk=4 --raid-devices=2  
/dev/sdc /dev/sdd
```

where chunk size is MD RAID stripe size, in KiB;

level is MD RAID level;

raid-devices – number of SSD drives, included in RAID.



We recommend you to create an MD RAID with chunk = **4** for better performance and more efficient SSD drives use.

2. Open the **Storage** menu and move to the **RAIDs** page.
3. Click on the name of the appropriate RAID to access to its page. In the right part of the page click **Select Drives for SSD Cache** (Figure 38).
4. In the opened window select the MD RAID to be used as SSD cache, and click **Select**.



In DC mode: Create another MD RAID with identical parameters on the second node and select it for SSD cache.

2.8.3. Disable SSD Cache

To disable SSD caching, perform the following:

1. Open the **Storage** menu and move to the **RAIDs** page.
2. Click on the name of the appropriate RAID to access to its page. In the right part of the page click **Select Drives for SSD Cache** (Figure 38).
3. In the opened window click **Remove SSD Cache** (Figure 39).

3. Working with LUNs

To work with system, you need to create at least one LUN partition after creating RAID.

You can manage LUNs in the **Storage** menu or on the **RAID** page. The following operations with LUNs are available:

1. [Create a LUN](#)
2. [Masking Rules](#)
3. [Delete a LUN](#)

3.1. Create a LUN

In RAIDIX 4.4.3 a LUN can be created in the following ways:

- from the Storage tab;
- from the RAID page.

To create a LUN, perform the following procedure:

1. Open the **Storage** menu and move to the **RAIDs** page (Figure 40) or click on the name of the array to open the appropriate **RAID** page (Figure 41).



If the **Storage** page contains the message **RAID List is Empty**, you need to create a RAID following the procedure described in 2.2 "[Create a RAID](#)".

The screenshot shows the RAIDIX interface with the **Storage** menu selected. The **RAIDs** sub-menu is active, displaying a table titled "RAID List". The table has columns for RAID Name, RAID Size, Free Size, RAID Level, Drives, Cache, SparePool, and Stripe. A single RAID entry is shown: R1, 1 567 GiB, 1 567 GiB, RAID0, 0-22, 1 GiB, —, 128 KB. Below the table, there are two buttons: "Create a LUN" and "Create a RAID".

| RAID Name | RAID Size | Free Size | RAID Level | Drives | Cache | SparePool | Stripe |
|-----------|-----------|-----------|------------|--------|-------|-----------|--------|
| R1 | 1 567 GiB | 1 567 GiB | RAID0 | 0-22 | 1 GiB | — | 128 KB |

Buttons: [Create a LUN](#), [Create a RAID](#)

Figure 40. RAID List section. LUN list

RAIDIX

Storage NAS Hosts & Interfaces System Monitoring

Dashboard RAIDs SparePools Drives Drives Scan RAID Profiles RAID import

< Back to RAIDs

RAID "R1"

Free, 1 567 GiB

LUN List is empty

| LUN Name | LUN Size | Block Size | NAS | ID | Type | Serial Number |
|---------------------------------------|----------|------------|-----|----|------|---------------|
| <input type="text" value="LUN Name"/> | 1567 GiB | 512 B | Off | - | - | - |

Create Cancel

Figure 41. RAID page. LUN list section

- Click **Create a LUN**. Fill in the required fields:

| Parameter | Description |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LUN Name (required) | A unique LUN name |
| LUN Size (required) | LUN size (GiB/TiB) The maximal possible value of the parameter is set by default. |
| Block Size (required) | Block size: 512 B, 4096 B |
| LUN is visible to initiators (optional) | Parameter allows the LUN to be used for NAS and automatically enables or disables it for all targets |
| ID | Parameter is displayed only on the RAID page. LUN number used for addressing purposes. The parameter can be used to diagnose the issues of the storage operation. |
| Type | The system automatically specifies the default host type, providing LUN interaction with hosts of any supported operating systems. |
| Serial Number | Parameter is displayed only on RAID page. LUN serial number. Parameter can be used to diagnose the issues of the storage operation |

- Click **Create**. The created LUN appears (Figure 42).

RAIDIX

Storage

[NAS](#)

[Hosts & Interfaces](#)

[System](#)

[Monitoring](#)

RAIDs

[SparePools](#)

[Drives](#)

[Drives Scan](#)

[RAID Profiles](#)

RAID List

| | RAID Name | RAID Size | Free Size | RAID Level | Drives | Cache | SparePool | Stripe | |
|---------------|---------------------------|---------------------------|---------------------------|----------------------------|------------------------|-----------------------|---------------------------|------------------------|--|
| | R1 | 1 567 GiB | 0 GiB | RAID0 | 0-22 | 1 GiB | — | 128 KB | |
| | L1 | 1 567 GiB | | | | | | | |
| Create a LUN | | | | | | | | | |
| Create a RAID | | | | | | | | | |

Figure 42. Created LUN

3.2. Random Access Optimization

Random access optimization (RAO) – is a block virtualization technology that allows you to create pools of deduplicated data for block-based storages. RAO uses deduplication to reduce storage consumption by eliminating duplicated blocks of information. Instead of writing the same data more than once, each duplicated block of data is registered as the link to the source (original, recorded earlier) block of data.

LUN with RAO includes several main parts (Figure 43):

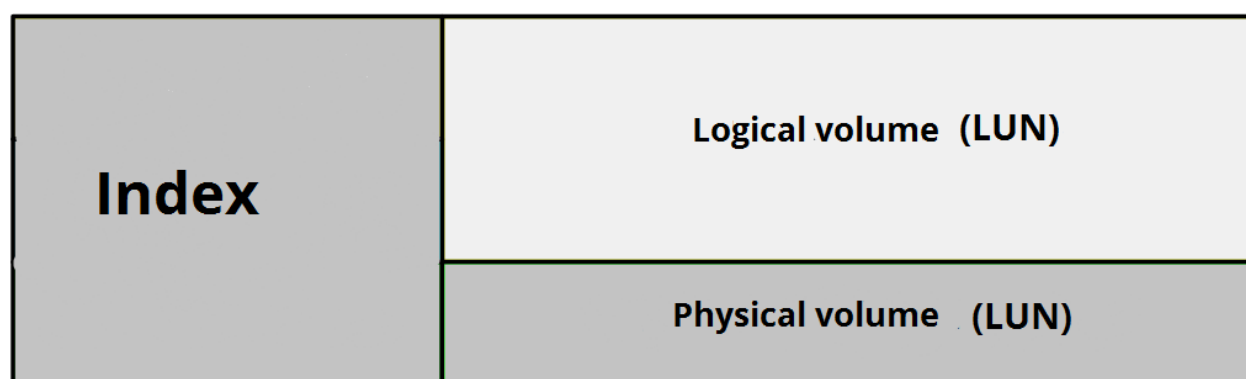


Figure 43 Schema of LUN with RAO

- **Index** – reads information from the logical volume, recognizes the repeating blocks of data (duplicates), writes the original data blocks on the physical volume, records and stores links to original data blocks, and provides the interaction between the logical volume, physical volume, and index.

- **Logical volume** – is a virtual block device that user uses to directly write/ read the information. If the size of the logical volume exceeds the size of the physical volume, the system operates in the thin provisioning mode.
- **Physical volume** – is divided into the slabs of the same size. In this slabs, physical volume stores the original data blocks. One volume with RAO can have up to 8069 slabs. The default value of one slab is 2 GB. You can use volume with the default value in simple systems with the low workload since with the 2 GB slab, the maximum use of the physical storage space will not exceed 16 GB. The recommended value of the slab is 32 GB, since with 32 GB slab size, the maximum use of the physical storage space can be up to 256 TB.

The storage will use at least one slab for the RAO metadata, and, therefore, will not use it to save data.



After you created the volume with the RAO, you will not be able to change the index parameters, physical and logical volumes parameters.

The LUN with RAO reads data from a logical volume, analyzes it, deduplicates and writes data on the physical volume. If space on the physical volume comes to an end, the system warns you about this:

- If the physical volume is filled by 70% – 90%, the system will display a warning message (Warning).
- If the physical volume is filled by 100%, the system will display an error (Error).

In both cases, the system will send you an e-mail notification with the warning.

3.2.1. Create a Random Access Optimization LUN

To create LUN with RAO, do the following:

Creation of more LUNs on the RAID, containing a LUN with RAO, may lead to poor performance.



We strongly recommend creating only one LUN with RAO on the RAID. Do not add more LUNs or LUNs with RAO on the same RAID.

In the same vein, do not create LUNs with RAO on RAID levels RAID0 and RAID 10.




If you use virtual machines on RAO LUN, you have to enable [Enable 512](#).

1. In the main menu, click **Storage** and move to the **Dashboard** page.



If the **Storage** page contains the message **RAID List is Empty**, you need to create a RAID following the procedure described in 2.2 "[Create a RAID](#)".

2. In the **RAID List** click  to expand the LUN list of the selected array.
3. Click **Create a LUN**, the **Create a LUN** window will appear (Figure 17).
4. Enable RAO. You will see the **Advanced Random Access Optimization Parameters** panel (Figure 19).
5. Click **Create a LUN**. Fill in the required fields.

If you want to create a big RAO LUN, we recommend you to use the following parameters:



| LUN Size (up to _ Tb) | Cache size (not less than _ Gb) | Slab (not less than _ GB) |
|-----------------------|---------------------------------|---------------------------|
| 20 | 18 | 8 |
| 40 | 24 | 8 |
| 50 | 26 | 8 |
| 60 | 30 | 8 |
| 70 | 32 | 16 |
| 100 | 32 | 16 |

The RAO LUN size depends on the extent of the RAM. For every 10 Tb of RAO LUN, you need 3 GB of RAM.

| Parameter | Description |
|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LUN Name (required) | A unique LUN name |
| LUN Size (required) | <p>LUN size (GiB/TiB)</p> <p>The maximal possible value of the parameter is set by default (equals the size of the RAID).</p> <p>The recommended minimal size for RAO LUN is 100 Gb.</p> |
| Block Size (required) | Block size: 512 B, 4096 B |
| LUN is visible to initiators (optional) | Parameter allows the LUN to be used for NAS and automatically enables or disables it for all targets |
| RAO (required) | Enables random access optimization |
| Logical Size | <p>The logical size of the LUN. By default is equal to the difference between the physical size of the volume and the size of the index.</p> <p>If the size of the logical volume exceeds the size of the physical volume, the system will work in the thin provisioning mode.</p> |
| Sparse | <p>Enables Random Access sparse indexing.</p> <p>The sparse indexing guarantees the largest amount of data stored in the index memory. Sparse indexing allows you to store the amount of data that ten times exceeds the amount of data stored with conventional deduplication.</p> <p>With the enabled sparse indexing, the amount of deduplicated data will be 10 TB for 1 GB of RAM. The amount of 1 GB of the sparse index is enough to deduplicate and store 40 TB of data (from the physical storage).</p> <ul style="list-style-type: none"> • On (default) – sparse indexing enabled. • Off – sparse indexing disabled. <p>Enabled sparse indexing requires a lot of hardware resources.</p> |
| Memory | <p>Specifies the amount of server memory in gigabytes; the default size is 0.25 GB.</p> <p>Specifies the amount of server memory in gigabytes; the default size is 0.25 GB.</p> <p>When you create the RAO LUN, the value for the Memory parameter of the LUN Size that is bigger than 1Tb should be at least 1 GB. For other Memory values, you can use the default LUN Size value (0.25 GB).</p> |

Parameter**Description**

When you create the RAO LUN in DC, it is necessary to consider that the total amount of RAM allocated for the RAO LUN must not exceed the amount of memory on each node.

Size

Specifies the index size in gigabytes. The recommended value is $\frac{1}{4}$ of the LUN Size. If not specified, default is calculated based on the memory allocated to the server through the Memory option.

The default value of Size is 1 (200 GB).

The Memory default value 0.25 is equal to 50 GB of index size.



It is impossible to create the RAO LUN with index size larger than 4 TB.


Creating the RAO LUN with big index size takes time.

The process of RAO LUN creating is resource-intensive and can affect overall system performance. Therefore, when you create the RAO LUN, you must disable/ stop all processes of writing or reading from other LUNs.

With processes disabled, the approximate time of creating of RAO LUN is the following:

| Index Size (Gb) | Time (min) |
|-----------------|------------|
| 1000 | 5 |
| 1200 | 6 |
| 1400 | 8 |
| 1600 | 9 |
| 1800 | 10 |
| 2000 | 10 |
| 2200 | 12 |
| 2400 | 13 |
| 2600 | 14 |
| 2800 | 15 |
| 3000 | 16 |

| Parameter | Description |
|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| BIO Rotation Interval | Specifies the number of I/O operations to enqueue for each block I/ O submission thread before directing work to the next thread. |
| Slab Size | Specifies the size of the increment by which a Random Access Optimization volume can grow in gigabytes. Using a smaller size than the default maximum constrains the total maximum physical volume size. Must be a power of two between 1 and 32 gigabytes. |
| Enable 512e | Enables 512-byte block device emulation mode for Random Access Optimization. |
| Ack Threads | Specifies the number of threads to use for acknowledging completion of requested Random Access Optimization I/O operations. |
| BIO Threads | Specifies the number of threads to use for submitting I/O operations to the storage device. Each additional thread after the first will use an additional 18 MB of RAM. |
| Block Map Cache Size | Specifies the amount of memory allocated for cached block map pages in megabytes. Must be a multiple of 4 KB. The default is 128, which is the required minimum. |
| Block Map Period | A value between 1 and 16380 which determines the number of block map updates which may accumulate before cache pages are flushed to disk: the default value is 16380. |
| CPU Threads | The number of threads used in costly operations such as hashing. |
| Deduplication | Enables deduplication on Random Access Optimization volume. Disabled deduplication will result in increased performance and amount of place that data occupies on the drive. <ul style="list-style-type: none">• On (default) – deduplication is enabled.• Off – deduplication is disabled. |
| Logical Threads | Specifies the number of threads across which subdivide parts of the Random Access Optimization processing based on logical block addresses. |
| Physical Threads | Specifies the number of threads across which to divide parts of the Random Access Optimization processing based on physical block addresses. |

| Parameter | Description |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Read Cache Size | Specifies the extra Random Access Optimization volume read cache size in megabytes. This space is in addition to a system-defined minimum. This defaults to 0. |
| Writing Policy | <p>Specifies the write policy: sync, async.</p> <p>Auto. – the default option that corresponds to the writing policy of the RAID</p> <p>Sync. – writes are acknowledged only after data is stably written.</p> <p>Async. – writes are acknowledged after data has been cached for writing to stable storage. Data, which has not been flushed is not guaranteed to persist in this mode.</p> <div><p>The creation of the RAO LUN with enabled <code>Write Policy = Sync (Sync.)</code> parameter can take a long time.</p><p><code>Write Policy = Sync (Sync.)</code> parameter is automatically turned on, if <code>Write Through for RAID</code> is enabled.</p><p>Use the option <code>Write Policy = Sync</code> only on SSD drives</p></div> |

3.3. Masking Rules

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

RAIDIX 4.4.3 software enables the following masking rules:

Target Masking Rules

These rules imply specifying a target; using this target a particular LUN is available for hosts.



In the case of using Fibre Channel, InfiniBand and SAS adapters, a target is meant as a port on an adapter, using this port the corresponding LUN will be available to hosts.

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

In RAIDIX 4.4.3 software the target masking rules creation is carried out on the LUN page (for details, see section 3.3.1.1 "[Create a Target Masking Rule](#)").

Hosts Masking Rules

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

RAIDIX 4.4.3 software enables the following masking policy support:

- **Standard:** after creation, LUNs are available to all hosts.
- **Enterprise:** after creation, LUNs are not available to any host.



Standard mode is set by default. You can change masking policy mode with [CLI commands](#).

In RAIDIX 4.4.3 software the hosts masking rules creation is carried out on the LUN page (for details, see section 3.3.2.1 "[Create a Host Masking Rule](#)" and section 5.1.1 "[Create a Host Alias](#)")

3.3.1. Target Masking Rules

Target masking rule creation enables additional features on organizing hosts access to LUN. Besides, the following advantages come to life:

- **Performance Increase**

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.

- **Fault Tolerance Increase**

RAIDIX 4.4.3 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.

- **Managing 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.

3.3.1.1. Create a Target Masking Rule

To create a masking rule, perform the following:

1. Open the **Storage** menu.
2. In the **RAID List** section click  to expand the LUN list of an appropriate array.
3. Click on the name of an appropriate LUN. The **LUN** page opens (Figure 44).

RAIDIX
Storage
NAS
Hosts & Interfaces
System
Monitoring

Dashboard
RAIDs
SparePools
Drives
Drives Scan
RAID Profiles
RAID import

[Back to RAID "R1"](#)

Host Masking Rules for LUN "L1" are not created

[+ Create a Masking Rule](#)

Target List for LUN "L1"




| Type | Target Name | Remote Node Target | IP Filter | Available for LUN | |
|------|-------------------------|--------------------|-----------|-------------------|--|
| FC | 21:00:00:10:86:02:94:3e | — | | Yes | |
| FC | 21:00:00:10:86:02:94:3c | — | | Yes | |
| FC | 21:00:00:10:86:02:94:3f | — | | Yes | |
| FC | 21:00:00:10:86:02:94:3d | — | | Yes | |

[+ Create a Target \(iSCSI only\)](#)


Figure 44. The LUN page

6. The **Target List** section allows assigning a target to a LUN, and it contains the following user interface elements:

| 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 execute 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. SAS — a particular port assigned for LUN. The value is set automatically. |
| Target Name | Target name |
| Remote Node Target | <p>Name of the target on the remote node that is linked with the target on a local node. The option is available only in dual-controller mode and only for InfiniBand, SAS or Fibre Channel targets.</p> <p>The option allows you to combine two one-type targets on different cluster nodes in one; that simplifies further masking rules set up.</p> |

| Parameter | Description |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP Filter | <p>IP Filter allows assigning the masking rule to a particular IP address or several IP addresses.</p> <p>Click All to select all available IP-addresses.</p> <p> This option is enabled only for iSCSI targets</p> |
| Available for LUN | <p>The rule allows/denies access to a particular LUN:</p> <p> Yes – the rule allows access to a LUN for a particular target.</p> <p> No – the rule denies access to a LUN for a particular target.</p> |

4. To set up InfiniBand, Fibre Channel, and SAS access policy:

- Working in dual-controller mode, click the  element to enable editing.
Click **Add Link** in the **Remote Node Target** field and select the target that will be linked with the current target on the local node in the appeared window.
- In the **Available for LUN** field select the appropriate access policy. Click **OK** in the confirmation window.

5. To set up iSCSI access policy, perform the following:

- Create a target iSCSI (for details on how to create an iSCSI-target refer to 5.3.2 "[Create an iSCSI Target](#)").
- In the **IP Filter** field select an IP-address or several IP-addresses.
- In the **Available for LUN** field select the appropriate access policy. Click **OK** in the confirmation window.

3.3.2. 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.

The following levels of host access to LUNs are supported in RAIDIX 4.4.3:

- Read Only,**

- **Read/Write,**
- **Exclude.**

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.

If you add at least one host with **Read / Write** access level to the list, all other users will not have access to the LUN. Accordingly, to obtain permission, it is necessary to add them to the list.




Only one host can have both read and write access. If several hosts have simultaneous read and write access, a conflict leading to data destruction can occur.

Thus, all hosts except for one must have only read access. If the access is denied, the host is not deleted from the list, and you can assign it a different access level if required.

3.3.2.1. Create a Host Masking Rule

To create a host masking rule, perform the following procedure:

1. Open the **Storage** menu and move to the **RAIDs** page.
2. In the **RAID List** section click  to expand the LUN list.
3. Click on the appropriate LUN name. The **LUN** page opens.
4. In the **Host Masking Rules** section click **Create a Masking Rule**. A pane to create a host masking rule appears (Figure 45).

RAIDIX Storage NAS Hosts & Interfaces System Monitoring

Dashboard RAIDs SparePools Drives Drives Scan RAID Profiles

< Back to RAID "R1"

Host Masking Rules for LUN "L1" are not created

Host Alias Access Host (WWN/GUID/IQN)

All Read Only

[Add a New Host Alias](#)

Create Cancel

Figure 45. LUN page. A Host Masking Rule creation

5. In the **Host Alias** field select previously created host alias from the drop-down list (see section 4.1 "[Create a Host Alias](#)").
6. If the host alias is not created yet, click **Add a New Host Alias**. The window allows you to select a host by the session (Figure 46).

Add a New Host Alias

Alias Name

Select Hosts

| <input type="checkbox"/> Host (WWN/GUID/IQN) | Target Port | Node |
|-------------------------------------------------------------|--------------------------------|------------------|
| <input type="checkbox"/> localhost | LocalRDtarget LocalRDtarget | Kabal Stryker |
| <input type="checkbox"/> 2100001086026dc5 | 21:00:00:10:86:02:7d:6d | Kabal |
| <input type="checkbox"/> 0x691e17ffff902500002590ffff1720b1 | 0025:90ff:ff17:1e68 | Stryker |
| <input type="checkbox"/> 2100001086026dc4 | 21:00:00:10:86:02:7d:6c | Kabal |
| <input type="checkbox"/> 0xb12017ffff902500002590ffff171e69 | 0025:90ff:ff17:20b0 | Kabal |

Save Cancel

Figure 46. Add a New Host Alias window

7. Check appropriate hosts by their sessions.
8. In the **Host Alias Name** field type the name of host alias and click **Save**.
9. Specify access level of host aliases to LUNs in the field **Access: Read Only, Read/Write, Exclude** and click **Create**. A new masking rule appears in the list.

3.3.2.2. Change a Default Masking Rule

To check current host masking rules policy, use the following command:

```
rdcli param system show
```



- value <1> corresponds to the **Standard** mode (after creation, LUNs are available to all hosts),
- value <0> corresponds to the **Enterprise** mode (after creation, LUNs are not available to any host).

To change masking policy, execute the command:

```
rdcli param system modify --luns_visible_by_default 0|1
```



3.3.2.3. Delete a Host Masking Rule

To delete a host masking rule from the LUN page, perform the following:

1. Open the **Storage** tab.
2. In the **RAID List** section click  in the appropriate array line. The **LUN list** will open.
3. Click the appropriate LUN name to open the **LUN** page.
4. In the **Host Masking Rules for LUN** section click  in the appropriate rule line.
 7. Click **Yes** in the appeared confirmation window.

3.4. Delete a LUN

To delete a LUN, perform the following procedure:

1. Open the **Storage** menu.
2. In the **RAID List** section click  to expand the LUN list.
3. Click  of the corresponding LUN. A confirmation message appears. In the confirmation message click **OK**.

4. Network Attached Storage (NAS)

In RAIDIX 4.4.3 software provided 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 4.4.3 platform include:

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

RAIDIX 4.4.3 NAS functionality includes the ability to create and edit public folders – Shares with adjustable parameters (path, protocol, visibility, host selection, visibility).

Moreover, in RAIDIX 4.4.3 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.

Also [quoting](#) support is added: ability to specify the amount of LUN size, available for the particular share users or user groups.

To manage NAS settings, open the **NAS** tab (Figure 47).

The screenshot shows the RAIDIX web interface. At the top, there's a navigation bar with the RAIDIX logo and several menu items: Storage, NAS (highlighted in a dark purple box), Hosts & Interfaces, System, and Monitoring. Below this, there's a sub-navigation bar with Shares (highlighted in a dark purple box), Users & Groups, and Quotas. The main content area is titled 'Shares List'. It shows details for a share named 'lun1' with a size of 20 GiB, free size of 19,96 GiB, RAID 'r1', and File System: XFS. There are links for 'Format the LUN' and 'Create a Quota'. Below this, there's a table with columns: Share Name, Path, Protocol, Squash, Browsible, WORM, Guest, Users & Groups, and Hosts. The table contains one row for 'Share1' with path '/', protocol 'NFS', and '1 Host'. At the bottom of the table, there's a green button labeled '+ Create a Share'.

| Share Name | Path | Protocol | Squash | Browsible | WORM | Guest | Users & Groups | Hosts |
|------------|------|----------|--------|-----------|------|-------|----------------|--------|
| Share1 | / | NFS | No | — | — | — | — | 1 Host |

Figure 47. NAS page

The section provides the following detailed description of available operations with NAS:

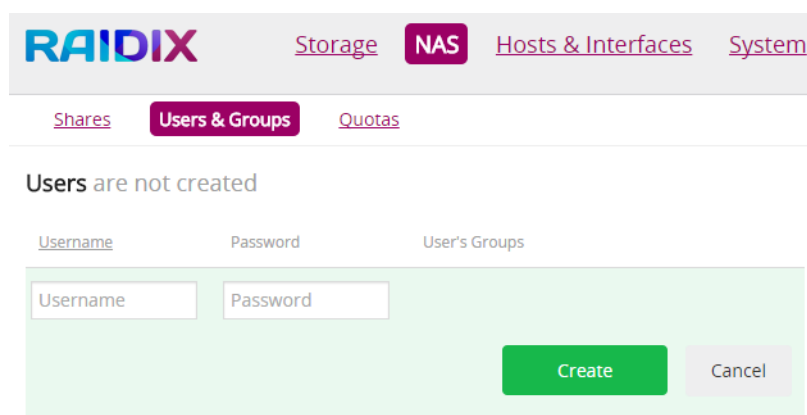
1. [Manage NAS users and groups.](#)
2. [Create a share.](#)
3. [Edit a share.](#)
4. [Delete a share.](#)
5. [Quoting set up.](#)

4.1. Manage NAS Users and Groups

4.1.1. Create Users and Groups

To create new NAS users or groups, perform the following steps:

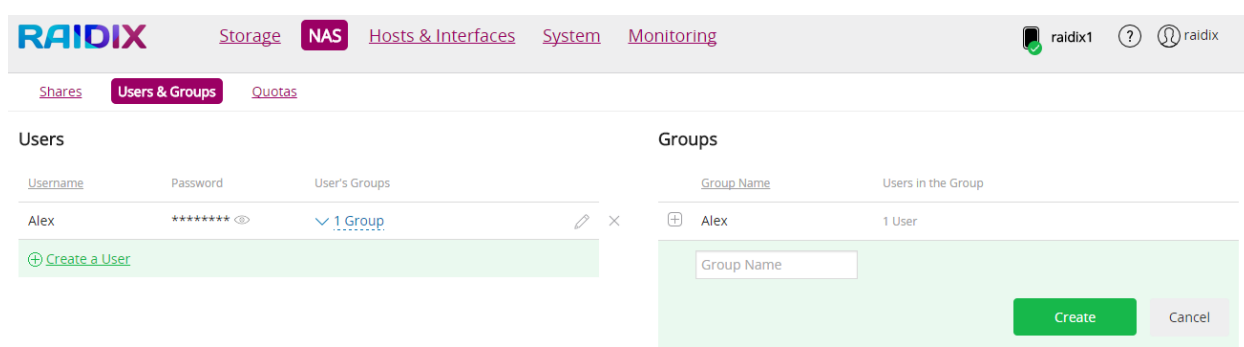
1. Select the **NAS** menu item.
2. In the **Users & Groups** section, click **Create a User**. The pane for a new user or group creation appears (Figure 48).



The screenshot shows the RAIDIX web interface with the 'NAS' tab selected. Under the 'Users & Groups' sub-tab, a message states 'Users are not created'. Below this, there is a form with three input fields: 'Username', 'Password', and 'User's Groups'. The 'Username' and 'Password' fields are currently empty. To the right of the 'User's Groups' field, there is a 'Create' button and a 'Cancel' button.


Figure 48. NAS users creation

3. Specify **Username** (in Latin letters) and **Password**. Click **Create**; the user will be created. At the same time, a group with the same name for this user will be created. This group will be assigned as primary for this user by default (Figure 49).
4. To create a group, select an item **Group** and specify its name (in Latin letters). Click **Create**.



The screenshot shows the RAIDIX web interface with the 'NAS' tab selected. Under the 'Users & Groups' sub-tab, there are two main sections: 'Users' and 'Groups'. The 'Users' section shows a table with one user, 'Alex', whose password is masked with asterisks. Below the table is a '+ Create a User' button. The 'Groups' section shows a table with one group, 'Alex', which contains '1 User'. Below the table is a form with a 'Group Name' input field and 'Create' and 'Cancel' buttons.

Figure 49. NAS group creation

5. To view or to add users to a group, click . Click **Add a User to the Group**; the window with available users appears. Choose users from the list and click **Select** to add them to the group (Figure 50).

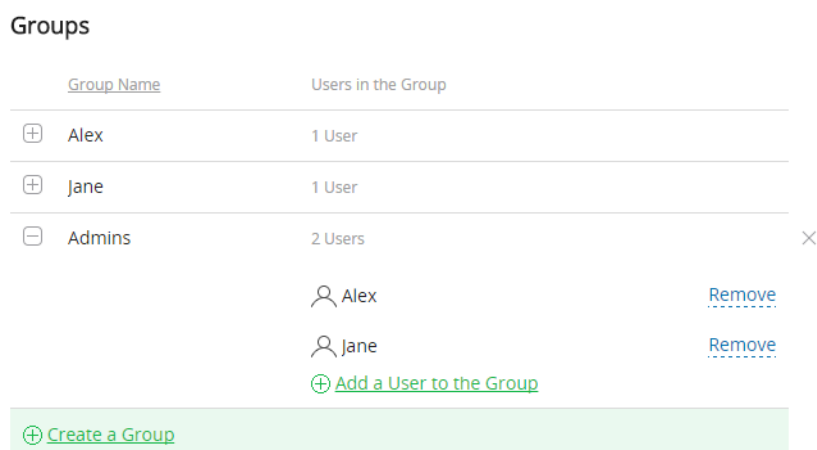







Figure 50. Adding users to a group

4.1.2. Edit and Delete NAS Users and Groups

6. To delete a user, to change its password or a list of groups, in the **Users & Groups** page of the **NAS** menu item do the following:

- To change the password, in the users section click  to enable editing. Enter the new one in the **Password** field and click **Apply**.
Click  to view the corresponding user password.
- To delete a user or a group, click  in the corresponding line. In the confirmation window click **Yes**.

NAS users can be included into several groups simultaneously; each group can be assigned as a primary.

- To add a user to other groups, check them in the group list and click **Apply**.
- To change the group's priority, click  **N Groups** in the **Users** section. In the appeared list check a group to be primary (Figure 51).
- To delete a non-primary group, in the **Users and Groups** section, click  in the corresponding group line and confirm your selection by clicking **Yes**.

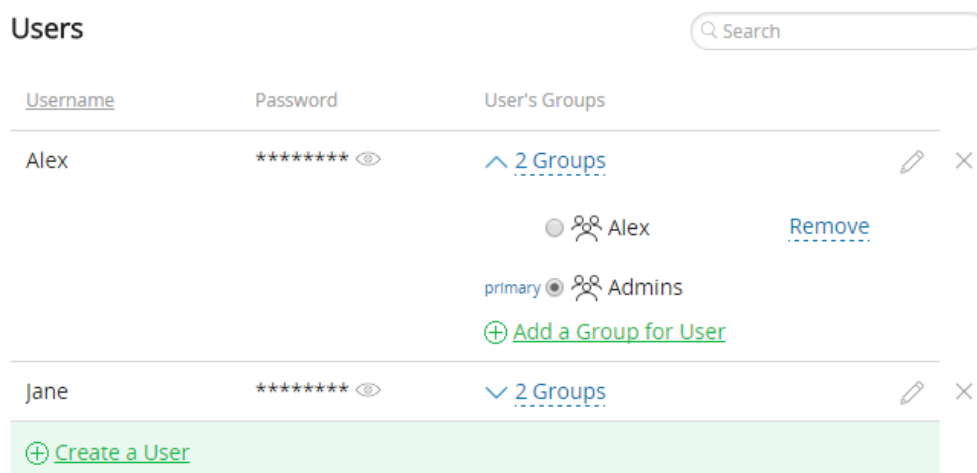


Figure 51. Users section. Setting priority to a group



To remove a Primary group, it is necessary first to delete all users from it by assigning them another Primary group.

4.2. Create a NAS Share

The basic elements of working with NAS in RAIDIX 4.4.3 software are data folders, available for a custom list of users – in other words – shared folders (**Shares**). Access to shares can be organized via SMB, NFS, FTP, and AFP protocols. In the case of using SMB protocol, setting up the Active Directory users parameters is available.

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

To create a new share, perform the following steps:

1. Create a RAID (for more information, see [2.2 «Create a RAID»](#)) and LUN partitions on it with **enabled NAS** (for more information, see [3.1 «Create a LUN»](#)).
2. Open the **NAS** menu item. The **Shares** page opens.
3. Click **Format LUNs** and select the file system to format with:
 - XFS;
 - EXT4.
4. Click **Create a Share**. A pane for share creation appears (Figure 52).

The screenshot shows the RAIDIX web interface. At the top, there's a navigation bar with 'RAIDIX' logo and tabs: Storage, NAS (selected), Hosts & Interfaces, System, and Monitoring. Below this, there's a sub-navigation bar with 'Shares' (selected), Users & Groups, and Quotas. The main section is titled 'Shares List'. It shows a summary for 'lun1' with a size of 20 GiB, free size of —, RAID 'r1', and File System: XFS. There are links for 'Format the LUN' and 'Create a Quota'. Below this is a table with columns: Share Name, Path, Protocol, Squash, Browseable, WORM, Guest, Users & Groups, and Hosts. The 'Create' form is overlaid on the table. It has a 'Share Name' input field, a 'Path' input field with a 'Select Path' link, a 'Protocol' dropdown menu set to 'NFS', a 'Squash' toggle set to 'No', and a '0 Hosts' link. At the bottom right of the form are 'Create' and 'Cancel' buttons.

Figure 52. Shares section. Creating a share

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

| Parameter | Description |
|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Share Name (required) | Name of the share, available for a particular list of users |
| Path (required) | The path to the share |
| Protocol (required) | Type of the protocol, via which access to a folder on the NAS-server is provided. Possible values are: <ul style="list-style-type: none"> NFS – Network File System; SMB – Server Message Block; FTP – File Transfer Protocol; AFP – Apple Filing Protocol. |
| Squash (NFS only) | Enabling the parameter allows users with read and write permissions to modify or delete files, which they do not own. Possible values are: <ul style="list-style-type: none"> Yes – enabled; No – disabled. |
| Browseable (SMB only) | The option allows customizing the share visibility to the host in the list of available shares. Possible values are: <ul style="list-style-type: none"> Yes – shared folder will be displayed; No – shared folder will not be displayed. |
| WORM (Write Once Read Many) (SMB only) | Enabling this option allows prohibiting the user to modify or delete the file after its creation for extended period. Files become unavailable for modification/ deletion after 1 second (grace period) |

| Parameter | Description |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | parameter) after the moment of their creation and within five years. You can modify this option from CLI . |
| Guest (SMB, FTP, and AFP only) | Parameter specifies guest access to the shared folder. Possible values are: <ul style="list-style-type: none"> • Disabled - no access; • Read Only; • Read/Write. |
| Users & Groups | Option allows selecting share users and groups and setting their permissions. |
| Hosts (SMB, FTP, and AFP only) | Option allows specifying a list of hosts that will have access to the shared folder. Hosts are added by IP-address or name with the following access types: Read/Write or Read Only . By selecting All you can provide access for all hosts. |



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



You must not create multiple shared folders with one and the same way.

For example, for the LUN with the name «Lun1» and shared samba folder with the name «Share1», and the path parameter – «path1», the full path to the mount point will be: «/mnt/nas/Lun1/path1».

You must not create the shared folder with the same path to the mount point.

It means, you must not create the shared samba folder with the name «Share2» for the LUN with the name «Lun1», and the path parameter – «path1», as the full path to the mount point will be: «/mnt/nas/Lun1/path1».

4.2.1. Create an SMB Share

While [working with SMB shares](#), [configuration of the authorization parameters for Active Directory](#) is available. It is recommended to perform Active Directory settings configuration before the share creation.

4.2.1.1. Integration with Active Directory

Active Directory allows several domain users to work with shares, assigning to the users different access permissions to the folder.

Besides that, while working with Active Directory in RAIDIX 4.4.3, an ability *to assign an administrator of root directory* of the share and its files is implemented.

To configure Active Directory parameters perform the following steps:

1. Open the NAS menu item.
2. Expand the **Active Directory** widget on the **Shares** page (Figure 53).

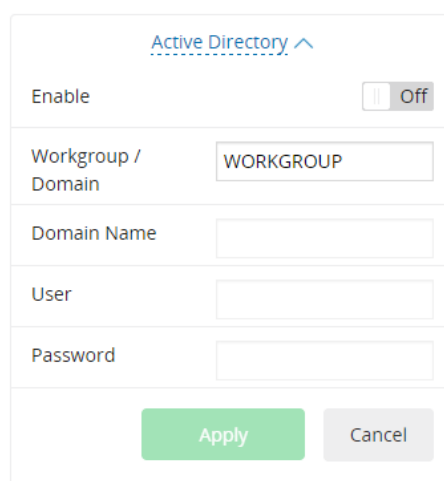


Figure 53. Active Directory settings window

3. Enable the Active Directory and fill in the following fields.

| Parameter | Description |
|--------------------|-------------------------------------------------------------------------------|
| Enable | Enabling/ Disabling Active Directory |
| Workgroup / Domain | Name of user group or domain, for which the connection settings are performed |
| Domain Name | Active Directory domain controller |
| User | Name of system administrator, 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.

- Click **Apply**. The connection will be implemented.

4.2.1.2. Configure SMB Share Parameters

When selecting SMB protocol, specify the following share parameters (Figure 54):

Figure 54. A set of fields to create an SMB share

- Specify the **Share Name**.
- Select or specify **Path** to the share.
- Select **SMB** Protocol.
- Set **Browseable** parameter value. The option allows customizing the share visibility to the host in the list of available shares.
- Specify **WORM** (Write once read many) parameters:
 - Yes** – parameter is enabled: files in the folder are read-only and cannot be modified. By default, files become unavailable for modification/deletion after 1 second (grace period) after the moment of their creation and within 60 months. Specify the **R/O Period** (in months) and **Grace Period** (in seconds) to change it.
 - No** – parameter is disabled.



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 grace period to perform modifications.



If a user mounts to a root directory, WORM does not work for its subdirectory. It is recommended to mount a folder with enabled WORM directly to the root directory.

6. Specify guest access permissions.
7. Define access permissions for share users, listed in the **Users & Groups** section. To do this, click on the value of the parameter; the window to add users will appear (Figure 55).

Users & Groups for Share

Load All

Users

Groups

| Name | Read Only | Read / Write | Admins |
|------|--------------------------|--------------------------|--------------------------|
| Alex | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |
| Jane | <input type="checkbox"/> | <input type="checkbox"/> | <input type="checkbox"/> |

Apply

Cancel

Figure 55. Adding share users

7. If [Active Directory](#) connection is enabled, setting will be performed only for the domain users.
8. You can assign share root directory **administrator** from the list of Active Directory users. The administrator will be able to manage all share files and assign ACL permissions to other users.
9. In the **Users & Groups** field, assign the appropriate permissions to users and groups by selecting them in the list. Share root directory administrator should be assigned **Read/Write** permissions.
10. If there is no user in the list, it must be created in the **Users & Groups** page of the **NAS** menu item. For details, see section 5.1.1 "[Creating New Users and Groups.](#)"

11. In the **Hosts** field, set the list of hosts that will have access to the created folder. To do this, in the **Host** field type IP-address or host name and select the access type: **Allow** or **Deny** (Figure 56). Click **Apply**, and the rule appears in the list.

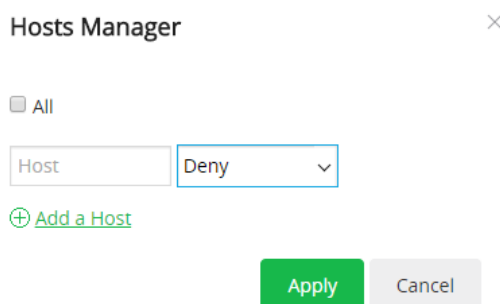


Figure 56. Hosts Manager window for SMB protocol

12. Click **Create**; the created folder appears in the list (Figure 57).

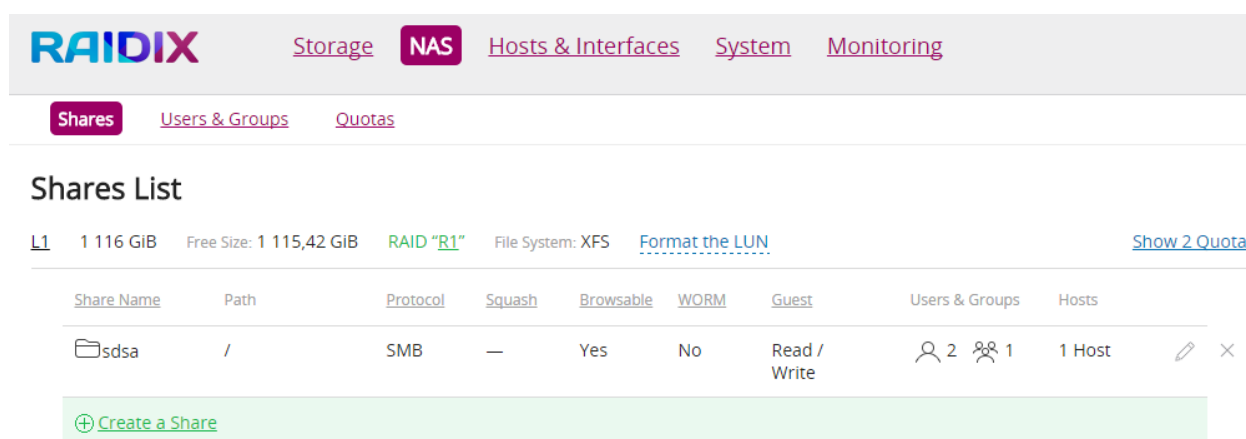


Figure 57. Shares List. Created SMB share

4.2.1.3. Set Up Trusts Between Active Directory Domains

In RAIDIX 4.4.3 software allows a user to set up trusts between Active Directory domains. Thus, users in one domain can also be authenticated and authorized to use resources of another domain.

To set up trusts between domains execute the command, specify the values as the domains' names:

```
rdcli nas samba modify -t|--trusts value -u value -p value
```


Where the value should include:

- -t domain name (or several domains), with a trust relationship.
- -u – username for the specified domain
- -p – password

Domain names with the trust relationship.

4.3. Edit a NAS Share


To edit the created share, perform the following actions:

1. Open the **NAS** menu item. The **Shares** page opens.
2. To edit a NAS share, click  of the corresponding folder name. A panel for editing share parameters opens. Fields **Share name**, **Protocol**, and **Path** remain unchanged.
3. Define new share parameters and click **Apply**.

4.4. Delete a NAS Share



Before deleting a NAS share, verify that it is not used by clients

To delete a share, click  in the corresponding share line in section **Shares** of the **NAS** tab. Confirm your choice by clicking **Yes** in the dialog box. The shared folder will be deleted.

4.5. Create, Edit, and Delete a Quota

Regarding RAIDIX 4.4.3, the **quota** is the limitation of LUN space, available for users of the share, created on this LUN.



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



To set up quoting, perform the following steps:

1. Open the **NAS** menu item and move to the **Quotas** page.
2. Click **Create Quota** of the appropriate LUN. The pane for adding a quota opens (Figure 58).

Figure 58. A pane to add a quota

3. Fill in the following fields and click **Create**.

| Parameter | Description |
|----------------------|------------------------------------------------------------------------|
| Users & Groups Names | Name of users and groups which the restriction applies to |
| Quota | LUN size, available for using, in MiB |
| Used | Parameter displays the size of the used LUN space of the allowed limit |

4. Information about the created quota appears. If several users or groups were selected, quotas of the same size would be set up for all of them.
5. To edit the size of the quota, click  in the appropriate quota line, enter new parameter value and click **Apply**.
6. Click , if you will need to delete a quota.

4.6. NAS in Dual-Controller Mode

Work with NAS in dual-controller mode is performed from the node, where RAIDs with LUNs on them are **active**.

To start working with NAS in dual-controller mode, an administrator should assign 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 create a new virtual IP, perform the following steps:

1. Open the **System** menu item and move to the **Network** page.
2. In the **Virtual IP Settings** section click **Create Virtual IP-address**. A pane for creation a new IP opens (Figure 59).



Virtual IP Settings

| Virtual IP-address | Port Name | State on Local Node | Netmask | Node |
|----------------------|-----------------------------------|-------------------------------------|----------------------|----------------------|
| <input type="text"/> | <input type="text" value="eth0"/> | <input type="text" value="Active"/> | <input type="text"/> | <input type="text"/> |

Figure 59. A pane to create a new Virtual IP

3. Define the following parameters of NAS in dual-controller 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 variants: <ul style="list-style-type: none"> • Active – work with NAS resources from this node is possible • Passive – work with NAS resources from this node will become possible only after nodes failover. |
| Netmask | Netmask that will be assigned to the virtual interface |

4. Click **Create** to apply settings.
5. To edit NAS virtual IP settings, click  in the corresponding IP address line.
6. To delete IP-address, click .
7. Repeat steps 1-6 on the second node.

5. Working with Hosts and Interfaces

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

5.1. Manage Host Aliases

Host Alias Name is an easy-to-remember name assigned to the host. Several hosts can be combined by one host alias.

Host page elements let you perform the following operations with hosts:

1. [Create Host Aliases and prioritize them.](#)
2. [Delete Host Aliases.](#)
3. [View information on all the Established Host Sessions.](#)

5.1.1. Create a Host Alias

To create a host alias, perform the following procedure:

1. Open the Hosts & Interfaces menu.
2. In the **Host** page click **Create a Host Alias**. The pane to create the host alias appears (Figure 60).

The screenshot displays the RAIDIX web interface for the 'Hosts & Interfaces' section. The 'Hosts' sub-tab is active. Below the navigation bar, the 'Hosts Aliases' section is shown. It contains a table with the following data:

| Host Alias | Host (WWN/GUID/IQN) | Realtime | QoSmic | LUN List |
|------------|---------------------|----------|--------|----------|
| All | ✓ All Hosts | Off | — | — |

Below the table, there is a form to create a new host alias. The form includes the following fields and controls:

- Alias Name**: A text input field.
- Host (WWN/GUID/IQN)**: A text input field containing 'WWW/GUID/IQN'.
- Realtime**: A toggle switch set to 'Off'.
- QoSmic**: A toggle switch set to 'Off'.
- Add a LUN**: A green button with a plus icon.
- Create**: A green button.
- Cancel**: A grey button.

Figure 60. 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**,
 - for SAS – **SAS Address**.
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.



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.

7. Turn on QoSMic service if you are planning to use it (for more information, see [QoSMic Set Up](#))
8. To set up host access level to LUN (to create a masking rule), click **Add a LUN** and select the corresponding LUNs, access to which you are going to set up. The defined LUN (LUNs) appears in the **LUN List** field (Figure 61).



If you have even one Host Alias with added LUN, you must add Host Alias with added LUN for every initiator.

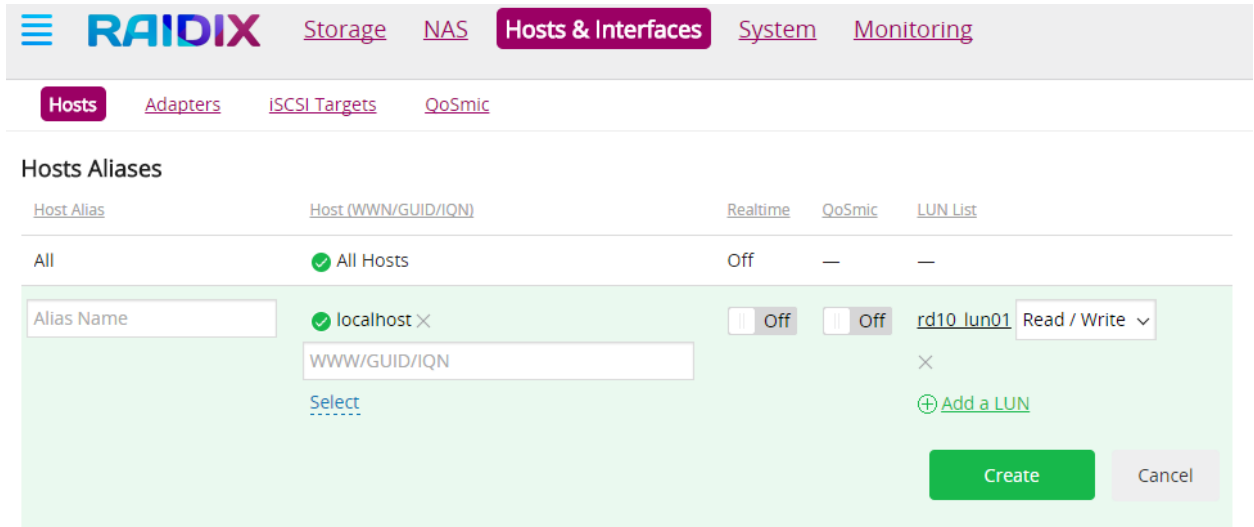


Figure 61. Host Aliases section. A LUN added

9. Specify the host access level to the LUN, selecting the appropriate item from the drop-down list:
- Read Only,
 - Read/Write,
 - Exclude.
10. Click **Create**. A new host alias will be created.



In RAIDIX 4.4.3 a host alias creation and specifying the host access level can also be done from the LUN page, for more details see section [3.2.2 «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**

5.1.1.1. Set Up Advanced Prioritization Parameters

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

| Parameter | Value |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Max Data Rate of Non- | Parameter allows you to limit the maximum data rate for requests from hosts without the highest priority, in MB. Analogue of attribute <code>request_saturation</code> while setting from the command console. A |

| Parameter | Value |
|--------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Real-Time Requests | nonnegative integer, the default value – 0 MB (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 | <p>Parameter allows you to limit the maximum number of simultaneous requests from the hosts which do not have the highest priority.</p> <p>Analogue of attribute <code>command_saturation</code> while setting from the command console. The default value is 4 MB.</p> <p>Setting 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.</p> <p>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.</p> <p>If sessions with hosts which do not have the highest priority are not created, commands in all sessions will have equal priority.</p> <p>Setting the <code>request_saturation</code> parameter as 0 disables the <code>command_saturation</code> parameter.</p> |

The parameters can be set from the widget **Driver Parameters** on the **Hosts** page (Figure 62) or using CLI commands (more information see [in Appendix A](#)).





Figure 62. Driver Parameters widget

5.1.2. Delete a Host Alias

To delete a host alias perform the following procedure:

1. Open the **Hosts & Interfaces** menu item.
2. In the **Host** section click  of the host alias you want to delete, a confirmation message appears.
3. Click **Yes**. The host alias will be deleted.

5.1.3. View the Established Sessions

The **Session** is an established connection between RAID 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) using the **Other Sessions** section of the **Hosts** page (Figure 63).

| Other Sessions | | | <input type="text" value="Search"/> |
|---------------------|-------------------------|-------|-------------------------------------|
| Host (WWN/GUID/IQN) | Target Port | Node | |
| 210000108602689a | 21:00:00:10:86:03:57:ec | Node1 | |
| 21000010860357ec | 21:00:00:10:86:02:68:9a | Node2 | |

Figure 63. Other Sessions section

| Parameter | Description |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Host(WWN/GUID/IQN) | A unique host identifier: <ul style="list-style-type: none">• WWN – a unique Fibre Channel ID,• GUID – a unique InfiniBand ID,• IQN – a unique iSCSI-host ID.• SAS Address— a unique SAS address. |
| Target Port | A unique target port which is used for interaction between the host and RAID |
| Node | Node name |

5.2. Working with Adapters

To work with adapters open the **Hosts & Interfaces** menu item and move to the **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).

5.2.1. Information on iSCSI Ports

Information on iSCSI ports is available in the **iSCSI Ports** section on the **Adapters** page (Figure 64)

RAIDIX

[Storage](#)

[NAS](#)

[Hosts & Interfaces](#)

[System](#)

[Monitoring](#)

[Monitor](#)

[Hosts](#)

[Adapters](#)

[iSCSI Targets](#)

[QoSmic](#)

iSCSI Ports

| Port Name | MAC | IP Address | Netmask | Status | Policy | Link Speed |
|-----------|-------------------|---------------|---------------|--------|--------|-------------|
| eth2 | d4:85:64:cc:5e:90 | 172.16.21.174 | 255.255.254.0 | Up | — | 1 000 Mb/s |
| eth3 | d4:85:64:cc:5e:91 | 192.168.1.174 | 255.255.255.0 | Up | — | 1 000 Mb/s |
| eth4 | 00:1b:21:c1:3d:1c | 10.5.5.174 | 255.255.255.0 | Up | — | 10 000 Mb/s |
| eth5 | 00:1b:21:c1:3d:1d | 10.4.4.174 | 255.255.255.0 | Up | — | 10 000 Mb/s |

Figure 64. The Adapters Page. 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 |

| Parameter | Description |
|------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status | <p>Connection status of a corresponding Ethernet port. The parameter can have the following values:</p> <ul style="list-style-type: none"> • Up – a connection is established • Down – there is no connection. Possible reasons: the cable is switched incorrectly, or the driver is missing • Unknown – 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>For detailed description of Policy option refer to 7.2.3 “Set Up Network Interfaces Parameters”.</p> |
| Link Speed | Link speed in Mb/sec |

5.2.2. Information on InfiniBand Adapters

The information on the used InfiniBand adapters is available in the section **InfiniBand Adapters** of **Adapters** page (Figure 65).

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 65. The InfiniBand Adapters section of the Adapters page

| 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 | <p>A physical state of the InfiniBand adapter port:</p> <ul style="list-style-type: none"> • Init – the port is initialized, but not active • Active – physical connection is established, the port is active • Down – there is no physical connection |

| Parameter | Description |
|-------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Link Status | Status of connection through the InfiniBand adapter: <ul style="list-style-type: none"> LinkUp – connection is established. This state appears only if Active connection is established. Polling – there is no connection. Possible reasons: incorrect cable connection or no driver loaded Down – there is no connection with the port. |
| Link Speed | Connection speed (Gb/sec) |

5.2.3. Information on Fibre Channel Adapters

Information on Fibre Channel Adapters is available in the **Fibre Channel Adapters** section in the **Adapters** page (Figure 66).

Fibre Channel Adapters

| Model | WWNN | WWPN | Link Status | Link Speed |
|-----------------------|-------------------------|------------------|-------------|------------|
| ATTO Celerity FC-82EN | 20:00:00:10:86:03:68:71 | 2100001086036871 | Up | 8 Gb/s |
| ATTO Celerity FC-82EN | 20:00:00:10:86:03:68:70 | 2100001086036870 | Up | 8 Gb/s |

Figure 66. The 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"> Up – the connection is established. Down – there is no connection. Possible reasons: incorrect cable connection or no driver loaded. Unknown - connection status is unknown. |
| Link Speed | Connection speed (Gb/s) |

5.2.3.1. Modifying the Fibre Channel Parameters

In RAIDIX 4.4.3 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 --help
```

5.2.4. Information on SAS Adapters

Information on SAS adapters is available in the **SAS Adapters** section on the **Adapters** page (Figure 67).

| SAS Adapters | | | | | |
|--------------|------------|------------|-------------------------|------|------------------|
| State | Link Speed | Phy Number | SAS Address | Port | Connector Number |
| Down | Unknown | 0 | 50:06:05:b0:08:91:88:40 | — | 0 |
| Down | Unknown | 1 | 50:06:05:b0:08:91:88:40 | — | 0 |
| Down | Unknown | 2 | 50:06:05:b0:08:91:88:40 | — | 0 |
| Down | Unknown | 3 | 50:06:05:b0:08:91:88:40 | — | 0 |
| Up | 12.0 Gbit | 4 | 50:06:05:b0:08:91:88:40 | 1 | 1 |
| Up | 12.0 Gbit | 5 | 50:06:05:b0:08:91:88:40 | 1 | 1 |
| Up | 12.0 Gbit | 6 | 50:06:05:b0:08:91:88:40 | 1 | 1 |
| Up | 12.0 Gbit | 7 | 50:06:05:b0:08:91:88:40 | 1 | 1 |
| Up | 6.0 Gbit | 0 | 50:06:05:b0:09:0b:24:c0 | — | 0 |
| Up | 6.0 Gbit | 1 | 50:06:05:b0:09:0b:24:c0 | — | 0 |

Figure 67. SAS Adapters section

| Parameter | Description |
|------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| State | <p>State of connection through the SAS adapter:</p> <p>Up – connection is established.</p> <p>Down – there is no connection with the port. Possible reasons: incorrect cable connection or no driver loaded.</p> |
| Link Speed | Connection speed (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 |

5.3. Setting Up iSCSI

iSCSI architecture in RAIDIX 4.4.3 software is based on a “client-server” model.

- **Client** (host) is a device sending the reading/writing requests to the “server”.
- **Server** is the storage target processing requests from hosts.

5.3.1. Enable/Disable iSCSI Support

To enable/disable iSCSI support, specify **On** for **Enable iSCSI** parameter (Figure 68).

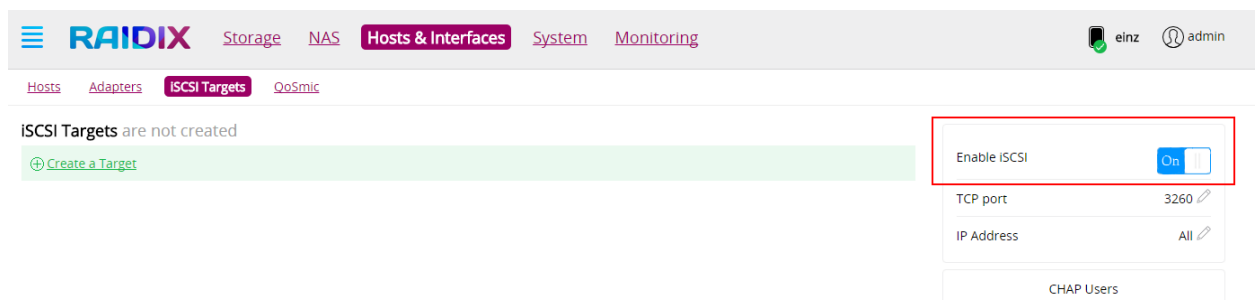


Figure 68. iSCSI page. Enabling iSCSI

5.3.2. Disabling of iSCSI Initiator with MPIO

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

For Windows Server 2008 R2:

1. Right-click on **Computer**, select **Manage** -> **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**.

For Windows Server 2008 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.

5.3.3. Create an iSCSI Target

To create an iSCSI target, perform the following procedure:

1. In the Hosts & Interfaces menu open the iSCSI Targets page.
2. Click **Create a Target**. The pane to set up the iSCSI target appears (Figure 69).

The screenshot shows the RAIDIX web interface for creating an iSCSI target. The top navigation bar includes links for Storage, NAS, Hosts & Interfaces (selected), System, and Monitoring. Below this, there are sub-links for Hosts, Adapters, iSCSI Targets (selected), and QoSmic. The main content area is titled 'iSCSI Targets are not created'. It contains a form with the following fields:

- Target Name:** iqn.2015-04.com.raidix:targett
- CHAP Mode:** None (dropdown menu)
- IP Filter:** All (dropdown menu)
- LUN List:** A list of LUNs with delete icons: L2, L1, L2, and L3.
- Login:** (empty field)
- Password:** (empty field)

At the bottom right of the form, there are two buttons: 'Create' (green) and 'Cancel' (grey). A green '+ Add a LUN' button is also present in the LUN List area.



Figure 69. iSCSI target creation



When creating a new iSCSI-target verify 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. |

| Parameter | Description |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| CHAP Mode | <p>The CHAP authorization mode:</p> <ul style="list-style-type: none"> • None – there is no authorization, • Bidirectional mode – mutual target-hosts authentication. Host sends login and password to the target. After verification target sends login and password to the host. • Unidirectional mode (One-way authentication) – only the host introduces itself to the target. <p> For Bidirectional and Unidirectional modes of CHAP authorization it is necessary to create CHAP users</p> |
| 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. <p> By default all iSCSI targets are not available for all IP-addresses</p> |
| LUN List | Option allows assigning a LUN (LUNs) to an iSCSI target |
| Login | Option is available only in Bidirectional CHAP mode (Duplex authorization) |
| Password | Option is available only in Bidirectional CHAP mode (Duplex authorization) |

5.3.3.1. Create CHAP users

In case of selection **Bidirectional** and **Unidirectional** modes of CHAP authorization, create CHAP users, performing the following steps:

1. Open the **Hosts & Interfaces** menu item and move to the **iSCSI Targets** page.
2. Click **CHAP Users**. A window with logins and passwords of users for CHAP-authorization appears (Figure 70).

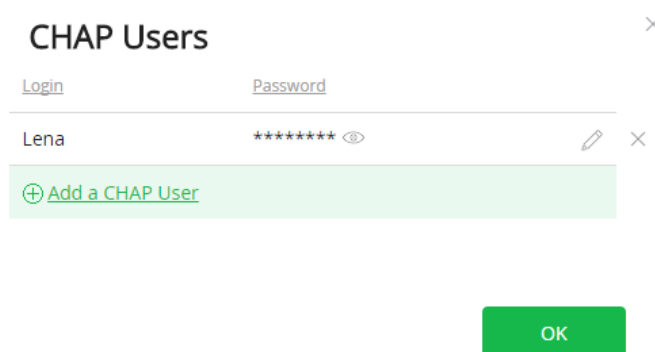





Figure 70. CHAP Users window

3. Verify that user is on the list. If not, click **Add a 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 appropriate user line.

5.3.4. Set the Advanced iSCSI Target Parameters

To setup the advanced iSCSI target parameters, perform the following:

1. Open the **Hosts & Interfaces** menu item and move to the **iSCSI Targets** page.
2. Click  of the corresponding target and in the appeared pane click **Parameters**. The **Target Parameters** window opens (Figure 71):

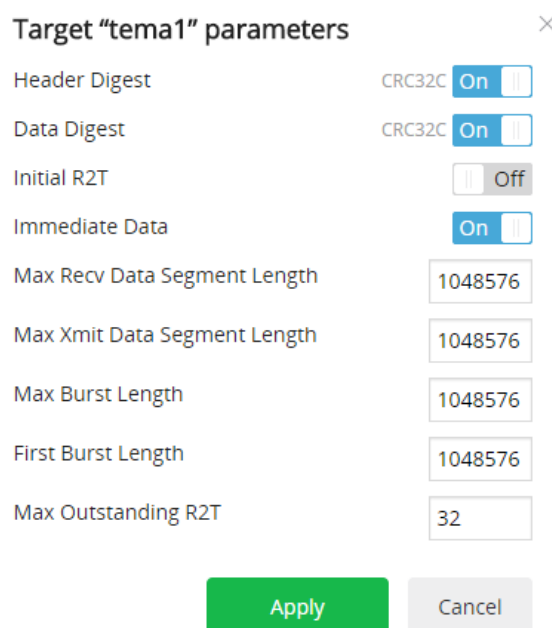


Figure 71. Target Parameters window


| Parameter | Description |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Header Digest | <p>Option provides the header integrity during the interaction with the specified iSCSI-target</p> <ul style="list-style-type: none">• CRC32C – option is enabled.• Off – option is disabled (set by default). |
| Data Digest | <p>Option provides the data integrity during the interaction with the specified iSCSI-target</p> <ul style="list-style-type: none">• CRC32C – option is enabled.• Off – option is disabled (set by default). |
| Initial R2T | <p>Option is used to turn off the default use of R2T for unidirectional and the output part of bidirectional commands.</p> <p>The possible parameter values are:</p> <ul style="list-style-type: none">• On – option is enabled.• Off – option is disabled (set by default). |
| Immediate Data | <p>Option allows including PDU as a part of initiator commands.</p> <p>The possible parameter values are:</p> <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); set by default.• Off – option is disabled. |
| Max Recv Data Segment Length | <p>Option is used to specify the maximum data segment length in bytes that can be received in one iSCSI PDU.</p> <p>Minimum value is 512.</p> <p>Maximum value is 1048576 (set by default).</p> |
| Max Xmit Data Segment Length | <p>Option is used to specify the maximum data segment length in bytes that could be sent in one iSCSI PDU.</p> <p>Minimum value is 512.</p> <p>Maximum value is 1048576 (set by default).</p> |
| Max Burst Length | <p>Option is used to specify the maximum SCSI data payload in bytes in a Data-In or a solicited Data-Out iSCSI sequence.</p> <p>Minimum value is 512.</p> <p>Maximum value is 1048576 (set by default).</p> |
| First Burst Length | <p>Option is used to specify the maximum amount in bytes of unsolicited data an iSCSI host may send to the target during the execution of a single SCSI command.</p> <p>Minimum value is 512.</p> |

| Parameter | Description |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | 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. By default, the value is 32 . |

3. After you select the corresponding option, click **Apply**.

5.3.5. Delete an iSCSI target

To delete an iSCSI target, perform the following procedure:

1. Open the **Hosts & Interfaces** menu item and move to the **iSCSI Targets** page.
2. Click  of the corresponding target to delete; a confirmation message will appear.
3. Click **Yes**. The iSCSI target will be deleted.

5.4. QoSmic Set Up

QoSmic service, implemented in RAIDIX 4.4.3, 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 quickly 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 execute Learning and Recognizing for applications simultaneously.

5.4.1. 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.



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 RAID is located on the different nodes (one at a time on each).



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

5.4.1.1. To Start Learning

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

To start QoSmic learning from web interface, do the following:

1. Open the **Host & Interfaces** menu and move to the **Hosts** page (Figure 70).

2. In the **Host**, page select the host that you want to use for QoSmic and click. Parameters panel will open.
3. In the parameters panel select **Host**, create **Alias**, and add **LUN** (or use parameters you created earlier).



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



You can change the list of initiators in the process of learning.

4. Click **Apply**.

| Host Alias | Host (WWN/GUID/IQN) | Realtime | QoSmic | LUN List |
|------------|---------------------------------------------------------------------------------------------------------------------|-----------------------------|-----------------------------|--------------------------------------------------------|
| All | ✓ All Hosts | Off | — | — |
| FC8_2 | ✓ 21000010860357ea × ✓ 21000010860357eb × <input type="text" value="WWW/GUID/IQN"/> Select | On <input type="checkbox"/> | On <input type="checkbox"/> | L1_FC8 Read / Write ▾ × + Add a LUN |
| IB_Host | ✓ 0x000000000000000000000000000000002c903004bf2f3 ✓ 0x000000000000000000000000000000002c903004bf2f4 | Off | On | L2_IB (R/W) |

Figure 72 Parameters panel

5. Open the **QoSmic** page (see Fig. 71).

The **QoSmic** page will display the list of initiators in the top half. These are initiators for which you **can Start Recognition**.

6. In the right part of the page, click **Start Learning**.

The learning will start.

QoSmic Hosts Aliases

| Host Alias | Realtime | Running Application |
|------------|----------|---------------------|
| FC8_2 | On | — |
| IB_Host | Off | — |

Applications

| Name | Warnings | TBD: Learning | Type |
|------------------|-----------------------------------------------------|---------------|----------|
| Io_1M_seq_w | Less than 90% of signatures have the same size, ... | 8 % | Unwanted |
| SQL | — | 1 % | Favored |
| Frametest_4k_w_4 | Less than 90% of signatures have the same size, ... | 3 % | Favored |

TBD: Upload a New Application

Figure 71. 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.



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



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.

To start learning from CLI, do the following:

1. To start learning:

```
rdcli qosmic teacher start -n <name> -hn <host name>
```

2. To stop the learning:

```
rdcli qosmic teacher stop
```

3. To look through the learning:

```
rdcli qosmic teacher show
```

4. To restart the learning:

```
rdcli qosmic teacher resume -n <name> -hn <host name>
```

5.4.1.2. Learning Will Not Start

The learning will not start in the following cases:


- You can start learning only for one application.
- If you activated the process of recognizing.
- The button **Start Learning** will be blocked if there are no hosts in the service.

5.4.1.3. Interruption of Learning

You can interrupt the process of learning by clicking “Stop Learning”. If the process of learning is interrupted, the service will add the application to the Applications list on the QoSmic page.



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**.

5.4.2. Application Types

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

- **low_intensity**: Low intensity on > 50% of intervals, the signature may be invalid.
- **same_size**: Less than 90% of signatures have the same size, the signature may be invalid.
- **short_learning_time**

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 will open.
2. In the parameters panel choose the type of the application from the drop-down list:
 - **Favored** means the selected application has high priority.
 - **Unwanted** means the selected application will be ignored.
 - **Disabled** means the application does not provide signatures for QoSmic.

To manage applications from CLI, do the following:

1. To create an application:

```
rdcli qosmic apps create -f <filepath>.
```
2. To clean application from warnings:

```
rdcli qosmic apps clean -n <name>.
```
3. To delete the application:

```
rdcli qosmic apps delete -n <name> -f
```
4. To change the type of application:

```
rdcli qosmic apps modify -n <name> -t <type>
```
5. To rename the application:

```
rdcli qosmic apps rename -n <name> -nn <nname>
```
6. To show the application:

```
rdcli qosmic apps show
```

5.4.3. Import of Signatures

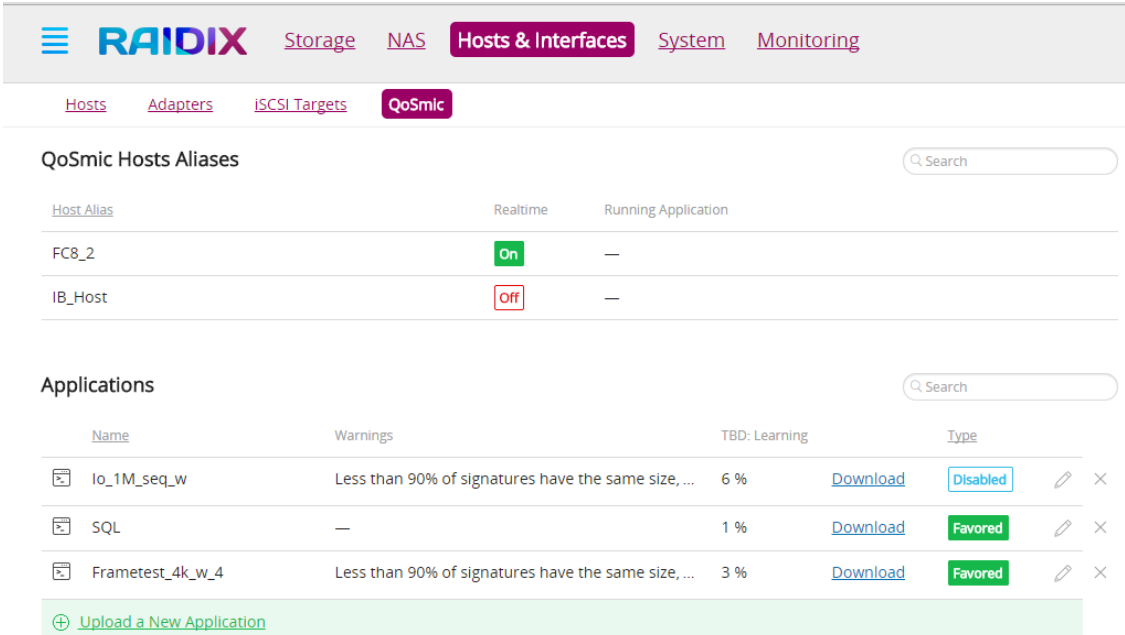
You can download collected signatures for the further use. For example, you can download the application (signatures) and upload them to another server.

To download signatures:

- On the QoSmic page, in the Applications section, select the application which signatures you want to download and click **Download** (See Fig.72).

To upload signatures:

1. On the **QoSmic** page, in the **Applications** sections, under the list of applications click **Upload a New Application**.
2. Select a file of the application you want to upload.
3. Click Open.



QoSmic Hosts Aliases

| Host Alias | Realtime | Running Application |
|------------|----------|---------------------|
| FC8_2 | On | — |
| IB_Host | Off | — |

Applications

| Name | Warnings | TBD: Learning | Type |
|------------------|-----------------------------------------------------|---------------|-------------------|
| Io_1M_seq_w | Less than 90% of signatures have the same size, ... | 6 % | Download Disabled |
| SQL | — | 1 % | Download Favored |
| Frametest_4k_w_4 | Less than 90% of signatures have the same size, ... | 3 % | Download Favored |

[Upload a New Application](#)

Figure 72. Applications download

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

5.4.4. QoSmic Recognizing

After you performed 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 from the web interface, do the following:

1. Click **Start Recognizer** on the left side of the **QoSmic** page
2. Recognizing will follow the rules:
 - a. When **Realtime** and **QoSmic** are off on the **Hosts** page, any type of application will not affect **Realtime** (it means that **Realtime** will be off).
 - b. When **Realtime** is on and **QoSmic** is off on the **Hosts** page, any type of application will not affect **Realtime**.
 - c. When **Realtime** is off and **QoSmic** is on, than **Realtime** will be on for **Favored** applications and off for any other type of applications.
 - d. When both **Realtime** and **QoSmic** are turned on, than **Realtime** will be off for **Unwanted** applications and on for any other type of applications.

To start recognizing from CLI:

1. Open the **Monitoring** menu and move to the **Support** page.
2. Switch on **Console**.
3. To perform recognizing, run the following commands:

- a. To start the recognizing:

```
rdcli qosmic recognizer start
```

- b. To stop the recognizing:

```
rdcli qosmic recognizer stop
```

- c. To show recognizing:

```
rdcli qosmic recognizer show
```




6. Setting Up Advanced RAID Parameters

RAIDIX 4.4.3. Software implements advanced RAID configurations that help **to speed up your system's performance** and increase its reliability (by enabling Silent Data Corruption function). Basically, advanced RAID configurations provides you with cache settings that makes **RAIDIX cache highly configurable**. For most workloads the default settings fit very well, but if you are not satisfied with the given performance or you want to achieve the highest level of performance, system can be tuned through the adjustment of several parameters. These parameters can be saved as a **RAID Profile**, and then, it can be applied for several RAIDs and can be changed for each array separately.

RAID parameters 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 (for detail refer to 6.1 [Working with RAID Profiles](#)).

| Parameter | Description |
|----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Clean Segments Replacement | Clean cache segments contain the data that have been already flushed to the disks. In RAIDIX 4.4.3 clean segments are replaced by the LRU (Least Recently Used) algorithm. By the LRU algorithm the least recently used segments are flushed on the drives at first. |
| Dirty Segments Replacement | Dirty cache segments contain the data that have not been flushed to the disks yet. Dirty segments are replaced with LRU algorithm. RAIDIX 4.4.3 software implements the following replacement algorithms: LRU (Last Recently Used) – for the most efficient use this algorithm with the sequential workload. Cycle – for the most efficient use this algorithm with the random workload. |
| Max Dirty Segments for LRU | Maximum number of dirty segments to be flushed at once on the disks with LRU algorithm. Possible values: from 1 to 64. The default value is 24 . |

| | |
|------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Max Dirty Segments for Cycle | <p>Maximum number of dirty segments to be flushed at once on the disks with LRU algorithm. Possible values: from 1 to 64.</p> <p>The default value is 24.</p> |
| Read Ahead Past | <p>Minimum number of read data in the cache behind, required to start read ahead. Possible values: from 1 to 300 MiB. Default value is 64.</p> <p>The default value means that the client computer must request 64 MiB of sequential data to activate the cache (to run read ahead). Setting value too low would lead to recache. Setting value too high will not start read ahead.</p> |
| Read Ahead Distance | <p>Maximum interval of read data in cache looked forward for data presence, used to decide whether a read ahead is needed. Possible values: from 1 to 300 MiB.</p> <p>Setting up this parameter allows optimizing cache during the sequential read from drives. Thus, to make reading process smoother, the system can be set to read more data, than the client computer is currently requesting and save data in cache. This is called Read Ahead.</p> <p>The default value is 64.</p> <p>If number of Read Ahead Distance is too high, it will cause pauses in reading, setting the parameter too low would render the cache not as effective.</p> |
| Read Ahead At Once | <p>Maximum size of read ahead request, in stripes. Possible values: from 0 to 64. Default value is 24.</p> <p>If the request size from client computers is not limited, the requests might be uneven, possibly interrupting playback for other clients. Setting the value of this parameter allows to control (to limit) requests size, providing smooth reading from the array.</p> |
| Alt Buffers Size | <p>Size of alternative cache buffer (MiB). Possible values: from 200 to 4096. Default value is 512.</p> |
| Write Back at Once | <p>Maximum size of a single write back request, in stripes.</p> <p>Possible values: from 1 to 64.</p> <p>The parameter allows you to limit amount of cache, used for each write command from a client. Default value is 24.</p> |
| Max Write Back Sequential | <p>Maximum number of concurrent sequential requests to be written back. Possible values: from 1 to 1024. The default value is 8.</p> |
| Max Write Back Random | <p>Maximum number of random requests to be written back. Possible values: from 1 to 1024. The default value is 1024.</p> |

| | |
|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Max Write Data | Maximum amount of data to be written on the disks simultaneously (MiB). Possible values: from 1 to 2048 . Default value is 256 . |
| Number of Threads | Number of processing threads. Possible values: from 1 to 32 . Default value is 16 . |
| Random Reads | Option allows optimizing operation with array for random read. Possible values are: <ul style="list-style-type: none"> • Enabled, • Disabled. |
| Random Writes | Option allows optimizing the operation with array for random write. Possible values are: <ul style="list-style-type: none"> • Enabled, • Disabled. <div>  This parameter can be specified only for initialized arrays of RAID5, RAID6, RAID 7.3 and RAID N+M levels </div> |
| Silent Data Corruption | Option allows detecting and recovering corrupted RAID array data. By default it is Disabled . Option may operate in the following modes: <ul style="list-style-type: none"> • Detection – drive errors detection is running; • Correction – errors detection and, if possible, data recovery. <div>  This parameter can be specified only for initialized arrays of RAID5, RAID6, RAID 7.3 and RAID N+M levels </div> <div>  Silent Data Corruption suspends the following functions: Random Reads, Random Writes, and Advanced Reconstruction </div> |
| Verify Cache Sum | Enabling this option lets the system compare the cache checksums during the synchronization process (only in Dual-Controller mode) |
| Write Through | Option allows performing write operations directly on the disks and on the cache. Only if the data is written on the disks, the host will receive the confirmation report. Write through practically minimizes the risk of data loss. By default, Write Back mechanism is used. |
| Reconstruct at Once | Maximum size of a reconstruction request (MiB). Possible values: from 1 to 64 . Default value is 24 . |

Moreover, the following cache settings are not included in the profile and should be applied separately:

| | |
|-------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enable Write Through Without Synchronization | Enable this option to run write through mechanism when caches are not synchronized. Available only in dual-controller mode. This option is available under the UPS configuration widget on the Nodes page (System Nodes) |
| Percent of Cache Available to Write Requests (0-100%) | Maximum percent of total cache volume, occupied by dirty segments (0-100%) Parameter is aimed to guarantee part of cache for read operations when intensive write is running. Analogue of attribute cache_saturation while setting from CLI. Default value is 75 . You can set up this parameter in the Driver Parameters widget (Storage RAIDs). |

6.1. Working with Advanced RAID Parameters

While creating a RAID, **Advanced RAID Parameters** are applied by default, you can change Advanced RAID Parameters by [creating a RAID Profile](#) and applying it to the RAIDs or modify the default advanced RAID parameters of the appropriate RAID.

To modify **Advanced RAID Parameters**, perform the following procedure:

1. Open the **Storage** menu and move to the **RAIDs** page.
2. In the **RAID List** section click on the name of a corresponding array. The **RAID** page opens.
3. Click Show RAID Parameters. The Advanced RAID Parameters window opens.
4. Click **Advanced Settings** to expand the additional configuration parameters list (Figure 73).

Advanced RAID "Raid1" Parameters

RAID Size

Free Size

Level

Drives

Cache

SparePool

Stripe

Sync

Affinity/Status

4 652 GiB

4 652 GiB

RAID0

0-4

1 GiB

—

256 KiB

Off

Static Parameters

Alt Buffers Size ?

Write Back at Once ?

Max Write Data ?

Number of Threads ?

Verify Cache Sum ?

Reconstruct At Once ?

512

24

256

16

Off

24

Dynamic Parameters

Clean Segments Replacement ?

Dirty Segments Replacement ?

Max Dirty Segments for LRU ?

Max Dirty Segments for Cycle ?

Read Ahead Past ?

Read Ahead Distance ?

Read Ahead at Once ?

Max Write Back Sequential ?

Max Write Back Random ?

Random Reads ?

Random Writes ?

Silent Data Corruption ?

Write Through ?

LRU

LRU

24

24

64

64

24

8

1024

On

Off

Disabled

Off

Save as a RAID Profile...

Apply

Cancel

Figure 73 RAID Page. Advanced RAID Parameters window



In RAIDIX 4.4.3 **dynamic RAID parameters** are parameters that are applied dynamically without requiring RAID reload. While modifying **static RAID parameters, the array must be reloaded**

- Specify RAID parameters and click **Apply** to save the settings.
- You can click **Save as a RAID Profile** and apply the parameters for other RAIDs.

6.2. Working with RAID Profiles

RAIDIX 4.4.3 allows you to perform the following operations with profiles:

- [Create a RAID Profile](#) that sets the advanced RAID parameters.
- [Modify a RAID Profile](#) and apply it for a RAID group,
- [Delete a RAID Profile.](#)

6.2.1. Create a RAID Profile



If no parameter profiles created by administrator, system will apply default settings (**Default** RAID profile). It is not allowed to change the default RAID profile parameters.

To create a new **RAID profile**, do the following:

1. In the **Storage** menu open the **RAID Profiles** page.
2. In the **RAID Profiles** list click **Create a Profile**. The RAID parameters are able to set up (Figure 74).

RAIDIX

Storage

NAS

Hosts & Interfaces

System

Monitoring

Francesco

Laura

Dashboard

RAIDs

SparePools

Drives

Drives Scan

RAID Profiles

RAID Profiles

default

local_xfs_4k_video

nfs_4k_video

Create a Profile

Create a Profile

Profile Name

Clean Segments Replacement ?

Dirty Segments Replacement ?

Max Dirty Segments for LRU ?

Max Dirty Segments for Cycle ?

Read Ahead Past ?

Read Ahead Distance ?

Read Ahead at Once ?

Alt Buffers Size ?

Write Back at Once ?

Max Write Back Sequential ?

Max Write Back Random ?

Max Write Data ?

Number of Threads ?

Random Reads ?

Random Writes ?

Silent Data Corruption ?

Verify Cache Sum ?

Write Through ?

Reconstruct At Once ?

LRU

LRU

24

24

64

64

24

512

24

8

1024

256

16

On

Off

Disabled

Off

Off

24

Save

Figure 74 RAID Profiles page. Creating a profile

- Enter the **Profile Name**, modify the parameters, and click **Save**. The new RAID Profile appears in the **RAID Profiles** list.
- The new saved profile will be available to select from the drop-down list on the **RAIDs** page when you start to create a RAID (Use Profile field), on the **RAID Profile** page, and on the **Advanced RAID Parameters** window of the appropriate RAID (Figure 75).
- Settings of one profile can be applied for any amount of arrays; however, settings should be applied to each array separately.

RAID List

| RAID Name | RAID Size | Free Size | Level | Drives | Cache | SparePool | Stripe | |
|-----------|-----------|-----------|--------|-----------|-------|-----------|--------|-----|
| ⊕ r1 | 2 792 GiB | 0 GiB | RAID0 | 0,1,3 | 1 GiB | — | 128 KB | ✓ × |
| ⊕ r2 | 1 116 GiB | 1 116 GiB | RAID5i | 2,4,5,7,8 | 1 GiB | — | 128 KB | ✓ × |

RAID Name: — RAID Level: RAID0 Drives: — Cache: 1 GiB SparePool: — Stripe: 128 KB

Use Profile: default
 default
 nas

Create Cancel

Figure 75 Created profile appears in the Use Profile list

6.2.2. Modify a Profile

To modify a profile, do the following:

1. In the **Storage** menu open the **RAID Profiles** tab.
2. In the **RAID Profiles** list click the name of the appropriate RAID profile.
3. Perform the required modifications and click **Apply**.

6.2.3. Delete a Profile

To delete a **RAID profile**, do the following:

1. In the **Storage** menu open the **RAID Profiles** tab.
2. In the **RAID Profiles** list click **×**. In the confirmation window click **Yes**.

7. Working with Drives




To view information on drives open the **Storage** menu and move to the **Drives** page (Figure 76).


The screenshot shows the RAIDIX web interface. The top navigation bar includes 'Storage', 'NAS', 'Hosts & Interfaces', 'System', and 'Monitoring'. The 'Storage' menu is expanded, showing 'Dashboard', 'RAIDs', 'SparePools', 'Drives', 'Drives Scan', and 'RAID Profiles'. The 'Drives' page is active, displaying a 'Drives Scan' button and a status indicator '24 working 1.75 TB'. Below this, a table lists the drives with the following columns: LED, UID, Capacity, Slot, Model, RAID Name, SparePool, Serial, and ID. The table contains 8 rows of drive information.

| LED | UID | Capacity | Slot | Model | RAID Name | SparePool | Serial | ID |
|-----|-------|----------|-------------|-------|-----------|----------------------|--------|-----|
| 0 | 73 GB | 0 | DH0072FAQRD | r1 | — | 3TA06ZTC00009946S2TR | sdc | ... |
| 1 | 73 GB | 11 | EH0072FAWJA | r1 | — | 3TA14VRK00009040Q75N | sdn | ... |
| 2 | 73 GB | 4 | EH0072FARWC | r1 | — | PKVE2XTE | sdg | ... |
| 3 | 73 GB | 13 | EH0072FARWC | r1 | — | PKVGMDE | sdp | ... |
| 4 | 73 GB | 8 | EH0072FARWC | r1 | — | PKVGN00E | sdh | ... |
| 5 | 73 GB | 5 | EH0072FAWJA | r1 | — | 3TA0Y6M600009028HTZL | sdh | ... |
| 6 | 73 GB | 7 | EH0072FARWC | r1 | — | PKVGN0HE | sdj | ... |
| 7 | 73 GB | 1 | EH0072FARWC | r1 | — | PKV71STE | sdd | ... |

Figure 76 Drives page

| Parameter | Description |
|-----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| LED | Drive indicator button. After pushing the button, a drive indicator in the enclosure turns on. Allows associating the real drive with its prototype in RAIDIX software. |
| UID | Unique identifier assigned to the drive in enclosure |
| Capacity | Hard drive capacity, GB |
| Slot | Slot number (numbering starts from zero). |
| Model | Drive model ID |
| RAID Name | RAID name |
| SparePool | SparePool name, into which the drive is included |
| Serial | Device name according to the serial number of the drive |
| ID | The device name in DevFS |

| Parameter | Description |
|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Local / Remote | <p>Drive status on the corresponding node:</p> <ul style="list-style-type: none"> OK – works correctly This status appears in the following situations:<ul style="list-style-type: none">Dirty – a drive contains partitions.Faulty warning – drive is in the RAID, despite one or two I/O errors have occurred. <ul style="list-style-type: none">Needs reconstructing – a drive needs to be reconstructed.Reconstructing – a drive is reconstructing.Error – a failure occurred.Faulty error – more than 3 I/O errors occurred on the drive; this drive cannot be used in RAID or in SparePool.It is recommended to replace the faulty drive, but you can also click Clean Drive to return drive to a normal status.No drive – there is no drive in the system. |

Drive with **Dirty** status can be cleaned right from the RAIDIX 4.4.3 GUI. To do this, open the **Drives** page, move to the  icon and click the appeared **Clean Drive** element in the appropriate drive line (Figure 77). Click **Yes** in the confirmation message appeared.

RAIDIX
Storage
NAS
Hosts & Interfaces
System
Monitoring

Dashboard
RAIDs
SparePools
Drives
Drives Scan
RAID Profiles
RAID Import

Drives

LSI CORP SAS2X36 Revision: 0718 ID: 500304800119597f 25 Drives Size: 21 973 GB [Hide](#)

1 error 1 000 GB

1 warning 1 000 GB

6 working 4,6 TB

16 available 14,37 TB

1 dirty 1 000 GB


| LED | UID | Capacity | Slot | Model | RAID Name | SparePool | Serial | ID | |
|-----|----------|----------|-------------|-------|-----------|----------------------|--------|----|-----|
| 0 | 1 000 GB | 25 | MK1001TRKB | r2 | — | 3200A06DFRC5 | sdv | ✓ | ... |
| 1 | 1 000 GB | 27 | MK1001TRKB | — | — | 3200A066FRC5 | sdx | ✓ | ... |
| 2 | 300 GB | 10 | ST3300657SS | r2 | — | 3SJ1TT5P00009102EMSD | sdi | ✓ | ... |
| 3 | 1 000 GB | 3 | MK1001TRKB | r2 | — | 3200A07UFRC5 | sdd | | |
| 4 | 1 000 GB | 26 | MK1001TRKB | — | — | Z150A0BTFM16 | sdw | | |
| 5 | 1 000 GB | 14 | MK1001TRKB | — | — | 3200A07ZFRC5 | sdl | | |

Dirty

Clean Drive

Figure 77 Removing partitions from the Dirty drive

You can also clean the **Faulty error** drive (reset error counter) from the RAIDIX 4.4.3 GUI.

To do this, open the **Drives** page, move to the  icon and click the appeared **Clean Drive** element in the appropriate drive line. Click **Yes** in the confirmation message appeared.



While cleaning drive with **Faulty error** or **Faulty warning** statuses please 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.

7.1. Set Up Hot Spare

RAIDIX 4.4.3 implements a **SparePool** functionality that allows performing an automatic replacement (hot spare) of a drive in a RAID in case of this drive disconnection or failure. In terms of RAIDIX 4.4.3, SparePool is a set of drives, which will be used for a hot spare of a failed drive in the array. One SparePool can be assigned for several RAIDs.



After drive replacement in RAID 5, RAID6, RAID7.3, RAID10 and RAID N+M arrays reconstruction process will start.

To create a SparePool, perform the following procedure:

1. Open the **Storage** menu and move to the **SparePools** page.
2. Click **Create a SparePool**. A pane for SparePool creation will appear (Figure 78).

Figure 78 SparePools page

3. Fill in the required fields:

| Parameter | Description |
|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| SparePool Name (required) | Name of the SparePool |
| Drives (required) | <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> <p>One drive can be included only in one SparePool. Drives, already included in an array, cannot be included in SparePool.</p> </div> |
| Free Drives | SparePool drives, available to be included in RAIDs |
| Max Free Size | Maximum size of the drive included in the SparePool. |
| RAIDs | <p>Arrays, associated with the SparePool (an automatic replacement (hot spare) of the specified array drives will be performed with drives from the specified SparePool).</p> <p>Click Select RAIDs link, select one or several arrays and click Select.</p> |

| Parameter | Description |
|-----------|-------------|
|-----------|-------------|

From the right of the table current status of the corresponding SparePool is displayed. Possible variants are:



OK – SparePool is available for use;



Warning – there are no drives of suitable size in the SparePool to replace a failed drive in the associated array, or there are no free drives (available for replacement) in the SparePool.

4. Click **Create**. The created SparePool appears on the list.

| SparePool Name | Drives | Free Drives | Max Free Size | RAID |
|----------------|--------|-------------|---------------|------|
| SP1 | 2-4 | 2-4 | 1 000 GB | — |

[Create a SparePool](#)

Figure 79. Created SparePool



After selecting drives to be included in a SparePool, the system will perform test write on them. If any of the drives cannot pass the automatic write test, the system will return the error: such drives cannot be included in SparePools. Please, select other drives.

7.2. SMART-diagnostics

Contemporary hard drives contain sensors, which help to reveal and store the logs on discovered defects that can lead to the device failure. The hard drives can also use the [SMART](#) technology for self-check. The results obtained during the self-check are stored in logs. SATA and SAS drives have different SMART values.

To display the SMART information for the particular drive click **SMART**. The SMART window opens (Figure 80).

SMART

Device: TOSHIBA MK1001TRKB
Version: 0105

Serial number: 3200A08BFRC5

Device type: disk

Transport protocol: SAS

Local Time is: Mon Apr 13 16:27:28 2015 MSK

Device supports SMART and is Disabled

Temperature Warning Disabled or Not Supported

SMART Health Status: OK

Current Drive Temperature: 36 C

Drive Trip Temperature: 65 C

Manufactured in week 12 of year 2012

Specified cycle count over device lifetime: 50000

Accumulated start-stop cycles: 341

Specified load-unload count over device lifetime: 600000

Accumulated load-unload cycles: 17787

Elements in grown defect list: 2

Error counter log:

| | Errors Corrected by | | | Total | Correction | Gigabytes | Total |
|--------|---------------------|---------|----------|-----------|-------------|--------------|-------------|
| | ECC | | rereads/ | errors | algorithm | processed | uncorrected |
| | fast | delayed | rewrites | corrected | invocations | [10^9 bytes] | errors |
| read: | 0 | 12 | 1 | 1 | 0 | 775171.148 | 0 |
| write: | 0 | 0 | 0 | 0 | 0 | 43885.986 | 0 |

Non-medium error count: 1628

No self-tests have been logged

Long (extended) Self Test duration: 10899 seconds [181.7 minutes]

Cancel

Figure 80 SMART window

The table below contains the description of 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 |

| Parameter | Description |
|-------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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">• Manufacturing Defect List – list of defects discovered during the test phase supplied by the drive vendor. This list remains unchanged• GDL (Grown Defect List) – a list of defects, which appear during the drive use. The list is regularly updated |
| Vendor Cache Information | <p>The parameter provides the following information:</p> <ul style="list-style-type: none">• Blocks Sent to Initiator: if the SAS drive is used, host-adaptor is the initiator and the drive itself is a target. The parameter specifies the number of data blocks sent to Initiator. These are not necessarily 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• Blocks Received from Initiator: number of drive sectors where the write operation has ever been performed• Blocks Read from Cache and sent to the Initiator – 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• Number of Read or Write Commands whose size <= Segment Size: 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• Number of Read or Write Commands whose size > Segment Size: number of commands or data which had to be split into several drive to transfer to the drive or from the drive |

| Parameter | Description |
|-----------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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">• Quick test 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, <p>Complete drive scan that can be initiated by user only.</p> |

7.3. Drives Scan

In RAIDIX 4.4.3 software an ability to scan drives of an array 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.

The scanning can be enabled from the **Drives** page, from the **RAID** page and the **Drives Scan** page.

To run drive scan, specify the scanning **Type** (Read Scan/ Write Scan), **RAID Name** (if it is required) and the **Size** of data to be read/written on the drives (1 GB, 10 GB, 100 GB, All RAID) and click **Scan**.

When the scan starts, the system will automatically open the scan results page – the **Drive Scan** section (Figure 81).

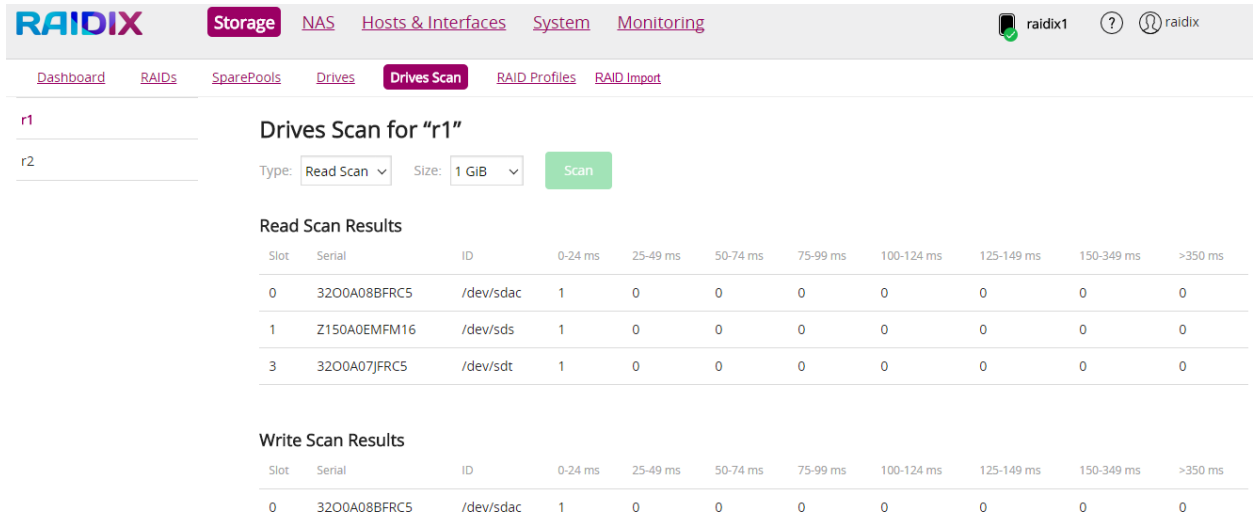


Figure 81 **Drives Scan** page. **Read Scan Results**

| Parameter | Description |
|----------------------------------|------------------------------------|
| Slot | Slot number in drive enclosure |
| Serial | Serial number of drive in the slot |
| ID | Drive ID in the system |
| Time intervals (in milliseconds) | Command execution time intervals |



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



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

Recommendations on interpreting 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 **Drive Scan** page. For example, the results of a **Read Scan** test on Figure 81 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.

It is commended that the rightmost columns of the tables should contain zero values. If a scan revealed drives with values uncommon for the overall statistics, they must be replaced with new ones, as these drives performance can have negative impact on the entire array performance. For example, if a full RAID scan was performed and non-zero data is only in the rightmost column >350 ms, this may indicate a disk failure.

7.4. Replace a Drive

RAIDIX 4.4.3 provides several ways of drive replacement. Each way of drive replacement initiates the reconstruction process (RAID recovery). You can set the reconstruction priority in the web-interface (for details see [Set Up Reconstruction Priority](#)).

RAIDIX 4.4.3 provides several ways to replace a drive:

- **Automatically:** a drive is replaced from the SparePool that is assigned to the RAID. Replacement of the failed drive occurs automatically. For more information see section 7.1 [«Set Up a Hot Spare»](#).
- **Manually:** replace a drive using one of the following ways:
 1. Physically replace a drive in the enclosure. To do this:
 - Define a failed drive (when a failure occurs, a drive RED LED will blink);
 - remove the failed drive;
 - insert a fault-free drive.

A new drive must meet the following requirements:



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

To remove metadata from a disk use the following CLI command:

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

where sdX is the name of the block device.

2. Manually replace the drive via web-interface.

- Manual drive replacement allows administrator to replace failed drive in array with any free (not used in RAID) drive in the system (no matter whether this drive is included in **SparePool**, assigned to this RAID or not).

The drive is not removed from the SparePool, into which it is included.



The option of manual drive replacement in web-interface will be unavailable if there is no drive of suitable size in the system.

7.4.1.1. Manually Replace a Drive

To replace a drive manually, perform the following steps:

- In the **Storage** tab move to the **Drives** page (Figure 82).
- Click **Replace** in the failed drive line.

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 shows 'Dashboard', 'RAIDs', 'SparePools', 'Drives', 'Drives Scan', and 'RAID Profiles'. The 'Drives' page is active, displaying a summary for 'HP D2700 SAS AJ941A' with 24 drives and a size of 1.752 GB. A table lists the drives with columns: LED, UID, Capacity, Slot, Model, RAID Name, SparePool, Serial, and ID. The first drive (UID 0) is highlighted, and a 'Replace' button is visible next to its status indicator.

| LED | UID | Capacity | Slot | Model | RAID Name | SparePool | Serial | ID |
|-----|-------|----------|-------------|-------|-----------|----------------------|--------|----|
| 0 | 73 GB | 0 | DH0072FAQRD | r1 | — | 3TA06ZTC00009946S2TR | sdb | ✓ |
| 1 | 73 GB | 11 | EH0072FAWJA | r1 | — | 3TA14VRK00009040Q75N | sdm | ✓ |
| 2 | 73 GB | 4 | EH0072FARWC | r1 | — | PKVE2XTE | sdf | ✓ |
| 3 | 73 GB | 13 | EH0072FARWC | r1 | — | PKVGMDE | sdo | ✓ |
| 4 | 73 GB | 8 | EH0072FARWC | r1 | — | PKVGN00E | sdj | ✓ |

Figure 82 Drive replacement from the Drives page

- A window with information about all drives, available to replace the failed drive, will appear (Figure 83). Select the corresponding drive and click **Select**.

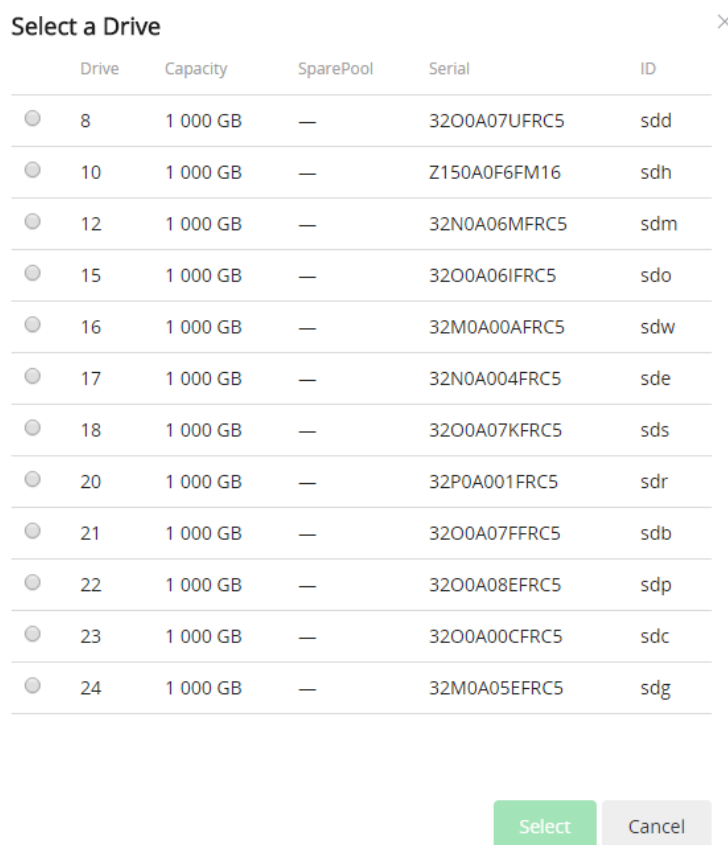


Figure 83 Selecting a drive to replace

After inserting a new drive, that meets the requirements listed above, the array reconstruction procedure runs automatically (Figure).

| | | | | | | | |
|----|----------|---------|------------|--------------|-----|--------------------|---------|
| 6 | 1 000 GB | TOSHIBA | MK1001TRKB | 32M0A04FFRC5 | sdg | Reconstructing: 2% | Replace |
| 11 | 1 000 GB | TOSHIBA | MK1001TRKB | 32O0A078FRC5 | sdi | | Replace |

Figure 84. Drives list section on the RAID page



In RAIDIX software, a mechanism of RAID **Partial reconstruction** is implemented, allowing restoring only a particular area on a hard drive, containing corrupted data. Therefore, in case you reinstall the already used drive, the area on which data was recording during the absence of the drive will be restored.


If you want to **stop reconstruction** procedure on the drive (for example, in case if you need to replace the reconstructing drive), use the following command in console:

```
rdcli raid reconstruct stop -n raid_name
```

If you want to **resume reconstruction**, type:

```
rdcli raid reconstruct start -n raid_name
```

7.4.1.2. Set Up Reconstruction Priority

To do so, open the **Storage** menu and move to the **RAIDs** page. In the right part of the page click **Driver Parameters**; the widget expands (Figure 84). For the **Reconstruction Priority** parameter click the  icon to enable editing. The reconstruction priority can be set from 0-100, wherein:

- **0** – the reconstruction stops completely if any other activity takes place;
- **100** – reconstruction mechanism has the highest priority.

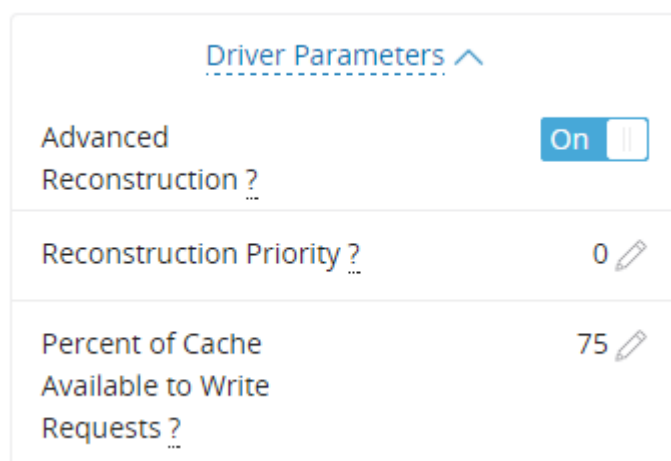


Figure 84 Driver Parameters widget

Specify the value and click the  icon.

8. Setting Up System Parameters

To set up the system parameters, use the **System** menu. This section provides a detailed description of RAIDIX 4.4.3 tasks:

1. [Working with Nodes](#)
2. [Setting Up Network Parameters](#)
3. [Working with a License](#)
4. [Setting Up Notifications](#)
5. [Configuration Recovery](#)

8.1. Working with Nodes

You can manage the nodes on the page **Nodes** of the **System** menu. The **Nodes** page elements let you perform the following operations with hosts:




1. [View the Nodes Current Status.](#)
2. [Automatic Nodes Switching](#)
3. [Manual Nodes Switching.](#)
4. [Restarting and Switching Off the System.](#)
5. [Configuring Dual-Controller Mode.](#)
6. [Disabling Dual-controller Mode.](#)
7. [Configuring the Connection of Uninterruptible Power Source \(UPS\).](#)






8.1.1. View the Nodes Current Status

To see the information on nodes current statuses, open the **System** menu and stay on the **Nodes** page (Figure 85):

Figure 85 Nodes page. Single-controller mode

Figure 86 Nodes page. Dual-controller mode

| Parameter | Description |
|-----------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ID | Node identifier |
| Name | <p>Node name. To edit the title of a node, click .</p> <p>Click  to save changes.</p> <p>The name of the local (current) node is displayed without the underscore. The name of the second (remote) node is displayed with underscore and is a link to switch to remote node interface.</p> |
| Heartbeat | IP address for connection to the second node in cluster |
| RAIDs | Displays the amount of Active, Passive, Active / Passive Failed Over RAIDs on every node. |
| | <p> 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.</p> |
| Alerts | The general node status: license, drives, adapters, arrays. Possible variants are available: |

| Parameter | Description |
|---------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| |  Healthy – a node runs correctly; |
| |  Warning – 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 Degraded, Reconstructing, Initializing, SDC, or Passive Failed Over; Adapter is Down (no connection); License status is Warning, because a hardware modification was made. Drive state is Reconstructing, Dirty or 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). |
| |  Error – 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, or more than one hardware modification was made Incorrect license data RAID is Offline or Not Loaded Drive in array is Failure, Faulty error, Error or No disk Heartbeat connection interface is unavailable (Down) |
| Node Status | <p>The node status in dual-controller mode. Possible variants:</p>  Node is OK – no migrated passive RAID with Affinity , corresponding to the node ID. |
| |  Failback – a node Failover has occurred: the system has detected passive RAID with Affinity , corresponding to the node ID. To return to the initial arrays state, click Failback . |
| Failover | <p>An option allows the nodes to change their cluster states. For details, see section 7.1.3 Manual Nodes Switching.</p> |
| Power Control | <p>The element is intended to reload and turn off the node:</p> <ul style="list-style-type: none"> Power On; Restart; Power Off. |

8.1.2. Automatic Failover

Distinctive feature of RAIDIX 4.4.3 is duplication of hardware components and interfaces in dual-controller system mode. This contributes to the smooth operation of the system, for example, in case of one of the hardware components failure, drive shelves` connection interface failure, or power-off of one of the nodes.

If any of these failures occurs, the system will perform automatic **failover** (changing the mode of the node). Thus, if one node fails, the other node transparently takes the entire workload: all arrays, that were **Active** on the failed node will become **Active** on the second node (available for read and write operations), allowing the administrator to fix appeared errors without stopping the system.

After nodes failover, the system switches to **Manual Mode**. To return the initial arrays state, click the **Failback** element on any page of the web interface of any node (Figure 87).

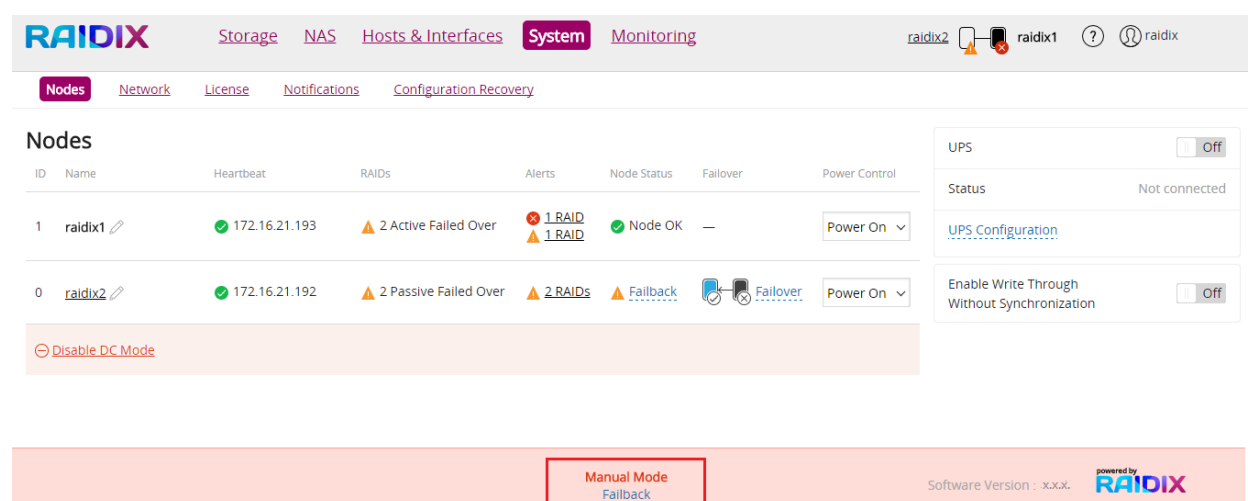


Figure 87 Performing nodes failback

8.1.3. Manual Failover

In addition to automatic nodes switching, it is possible to switch the modes of the operating nodes manually in dual-controller mode. The service allows you to change the current cluster state to the opposite one by migrating RAIDs with **Affinity** of one node to the second node. We recommend using the option during the maintenance (including updates installation) on the first node.



Nodes switching must be executed from the node to which you plan to failover (transfer) the arrays.

1. In RAIDIX 4.4.3 manual nodes failover can be performed in the following ways:
 - By clicking **Failover** on page **Storage** (Figure 88).




Figure 88 Performing failover from the Storage page

- By clicking **Failover** on page **Nodes** of the **System** tab in the appropriate node line (Figure 89);

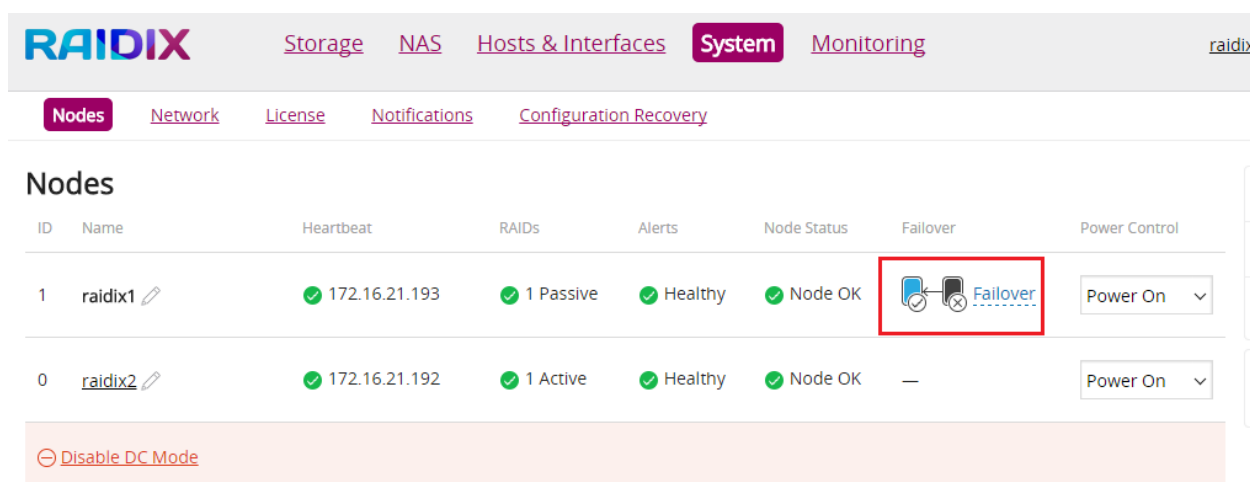


Figure 89 Performing failover from the Nodes page

2. Click **Yes** in the Confirmation window.

As a result, all arrays that were **Active** on the remote node and which have ID of the remote node in their **Affinity** will be transferred (failed over) to the local node. After the transfer (failover), these arrays will become **Active** on the local node and **Passive** on the remote node and will change their status to **Failed Over** (Figure 90). After that, you can continue to work with LUNs on the failed over RAIDs in a regular mode.

RAIDIX

Storage

NAS

Hosts & Interfaces

System

Monitoring

Nodes

Network

License

Notifications

Configuration Recovery

Nodes

| ID | Name | Heartbeat | RAIDs | Alerts | Node Status | Failover | Power Control |
|----|---------|-------------------------------------|---------------------------------------------|------------------------------|--------------------------------|-------------------------------------------|---------------|
| 1 | raidix2 | <div><div></div>172.16.21.193</div> | <div><div></div>1 Active Failed Over</div> | <div><div></div>1 RAID</div> | <div><div></div>Node OK</div> | — | Power On |
| 0 | raidix1 | <div><div></div>172.16.21.192</div> | <div><div></div>1 Passive Failed Over</div> | <div><div></div>1 RAID</div> | <div><div></div>Failback</div> | <div><div></div><div></div>Failover</div> | Power On |

Disable DC Mode

Figure 90 Display information about failed over RAIDs

Remote node will become ready for the maintenance.

8.1.4. Restarting and Switching Off the System

To restart the storage system, perform the following procedure:

1. In the **System** menu, open the **Nodes** page (Figure 85).
2. In the **Power Control** field from the drop-down list select **Restart** or **Power Off**.
3. Confirm you want to restart or switch off the system.

8.1.5. Disabling Dual-Controller Mode



After disabling dual-controller mode, an array will remain active on the node, on which it was active before the disabling. You can operate with each node separately in single-controller mode.

To disable dual-controller mode, from the web interface of any node open the **System** menu item and in the **Nodes** page click **Disable DC Mode** (Figure 91).

The screenshot shows the RAIDIX web interface. At the top, there's a navigation bar with tabs: Storage, NAS, Hosts & Interfaces, System (selected), and Monitoring. Below this, there's a sub-navigation bar with links: Nodes (selected), Network, License, Notifications, and Configuration Recovery. The main content area is titled 'Nodes' and contains a table with columns: ID, Name, Heartbeat, RAIDs, Alerts, Node Status, Failover, and Power Control. There are two rows of nodes: 'raidix1' and 'raidix2'. Below the table, there's a red box highlighting a button labeled 'Disable DC Mode'.

| ID | Name | Heartbeat | RAIDs | Alerts | Node Status | Failover | Power Control |
|----|---------|-----------------|-------------|--------|-------------|----------|---------------|
| 1 | raidix1 | ✓ 172.16.21.193 | ✓ 2 Passive | — | ✓ Node OK | Failover | Power On |
| 0 | raidix2 | ✓ 172.16.21.192 | ✓ 2 Active | — | ✓ Node OK | — | Power On |

⊖ Disable DC Mode

Figure 91 Disabling DC mode

8.1.6. Uninterruptible Power Source Configuration

In RAIDIX 4.4.3, functionality to connect the system to an uninterruptible power source (UPS) is implemented. The option is disabled by default.

To enable UPS monitoring, perform the following steps:

1. Connect UPS to the system with RAIDIX 4.4.3 installed, following the procedure, described in the appropriate UPS documentation. RAIDIX 4.4.3 supports the following types of UPS connection: through COM-port and USB.
2. In the **System** menu item move to the **UPS** widget on the **Nodes** page (Figure 92).




Figure 92 UPS configuration widget

3. Click **UPS Configuration**. The UPS window opens (Figure 93).

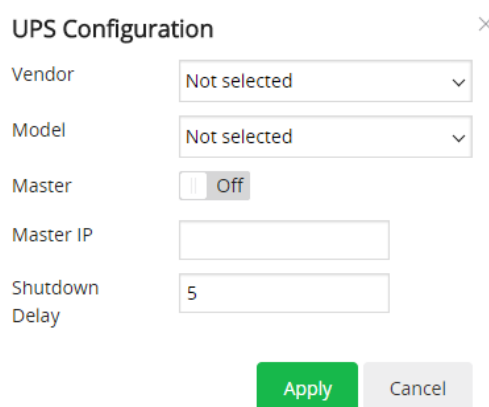


Figure 93 Configuring UPS

4. Select **Vendor** of the UPS, clicking the drop-down list in the corresponding field.
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, select **On** in the appropriate field.
 - If current system is **Slave**, specify the IP address of the **Master** system.
7. If you 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.

You can also organize interaction of RAIDIX 4.4.3 with UPS by uploading the necessary driver with the following CLI command (for details, see [Appendix A](#)):

```
rdcli param ups modify -d value
```

8.2. Setting Up Network Parameters

To set up the network parameters, open the **System** menu and move to the **Network** page (Figure 94).

The elements of the **Network** page allow you to perform the following:

1. [Setting up the General Network Settings.](#)
2. [Setting up Virtual IP-address.](#)

3. [Setting up Network Interfaces Parameters.](#)
4. [Setting Up Channel Bonding.](#)

RAIDIX Storage NAS Hosts & Interfaces **System** Monitoring raidix1 raidix2 ? RAIDIX

Nodes **Network** License Notifications Configuration Recovery

General Networking Settings

| Primary Interface | IP Gateway | Primary DNS | Secondary DNS | Tertiary DNS |
|-------------------|---------------|--------------|---------------|--------------|
| eth0 | 172.16.21.254 | 192.168.60.1 | — | — |

Virtual IP Settings

[+ Create a Virtual IP address](#)

Network Interfaces Parameters

| Port Name | Status | Link Speed | Enabled | DHCP | IP Address | Netmask | MTU |
|-----------|--------|------------|---------|------|---------------|---------------|------|
| eth0 | Up | 1 000 Mb/s | Yes | No | 172.16.21.193 | 255.255.254.0 | 1500 |
| eth1 | Down | — | No | No | — | — | 1500 |

Bonding Settings


| Bond Name | Policy | Interfaces |
|------------|--------|-------------------------------|
| Balance-rr | | + Add a Slave |

Create **Cancel**

Figure 94 Network page. Dual-controller mode

8.2.1. General Networking Settings

To set up the general network parameters, perform the following procedure:

1. Open the **System** menu and move to the **Network** page.
2. In the **General Networking Settings** section click , the pane to set up the network parameters opens (Figure 95).

General Networking Settings

| Primary Interface | IP Gateway | Primary DNS | Secondary DNS | Tertiary DNS |
|-------------------|---------------|--------------|---------------|--------------|
| eth0 | 172.16.21.254 | 192.168.60.1 | | |

Apply **Cancel**

Figure 95 Network page. General Networking Settings section

3. 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 |

8.2.2. Set Up Virtual IP-address

Setting up virtual IP-address is necessary to work with Network Attached Storage (NAS) in dual-controller mode. See more information about creation a new virtual IP in section [4.6 NAS in Dual-Controller Mode](#).

8.2.3. Set Up Network Interfaces Parameters

To set up the network interfaces parameters, perform the following:

1. Open the **System** menu and move to the **Network** page.
2. Click the  icon of the corresponding Ethernet port. The pane to set up the network interfaces parameters appears (Figure 96). You can set up the following parameters:

Network Interfaces Parameters

| Port Name | Status | Link Speed | Enabled | DHCP | IP Address | Netmask | MTU |
|-----------|--------|------------|----------------------------------------|------------------------------|--------------------------------------------|--------------------------------------------|-----------------------------------|
| eth0 | Up | 1 000 Mb/s | <input checked="" type="checkbox"/> On | <input type="checkbox"/> Off | <input type="text" value="172.16.21.193"/> | <input type="text" value="255.255.254.0"/> | <input type="text" value="1500"/> |
| eth1 | Down | — | No | No | — | — | 1500 |

Apply Cancel

Figure 96. Network page. Network Interfaces Parameters section

3. Fill in the required fields and click **Apply**.

| Parameter | Description |
|-----------|--------------------|
| Port Name | Ethernet port name |

| Parameter | Description |
|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Status | A connection status of the appropriate Ethernet port. The parameter can have the following values: <ul style="list-style-type: none">• Up – a connection is established.• Down – there is no connection. Possible reasons: the cable is switched incorrectly or the driver is missing.• Unknown – link status is not defined. |
| Link Speed | Link Speed, in Mbps |
| Enabled | A parameter stating whether the port is enabled. <ul style="list-style-type: none">• On – the port is enabled.• Off – the port is disabled. |
| 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. |

8.2.4. Set Up Channel Bonding

Channel bonding enables aggregation of multiple network interfaces into a single logical bonded interface as well as to distribute the load between these channels, increasing capacity in both directions. Bonding may improve system performance.

To configure channel bonding, perform the following:

1. Open the **System** menu and move to the **Network** page (Figure 96).
2. Click **Create a Bond**. The pane for bonding appears (Figure 97).






Figure 97 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. |
| Active-backup (Active-backup policy) | Only one slave interface in the bond is active. A different slave becomes active if, and only if, the active slave fails. 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, specify the Ethernet interfaces to be bound to as a slave.
5. Click **Create**.
6. To unbound the interfaces, click . In the confirmation window click **Yes** to confirm the operation.
7. To edit the bonding parameters, click .

8.3. Working with a License

8.3.1. View the Product License Information

To view the information on the license, open the **System** menu and move to the **License** page (Figure 98).

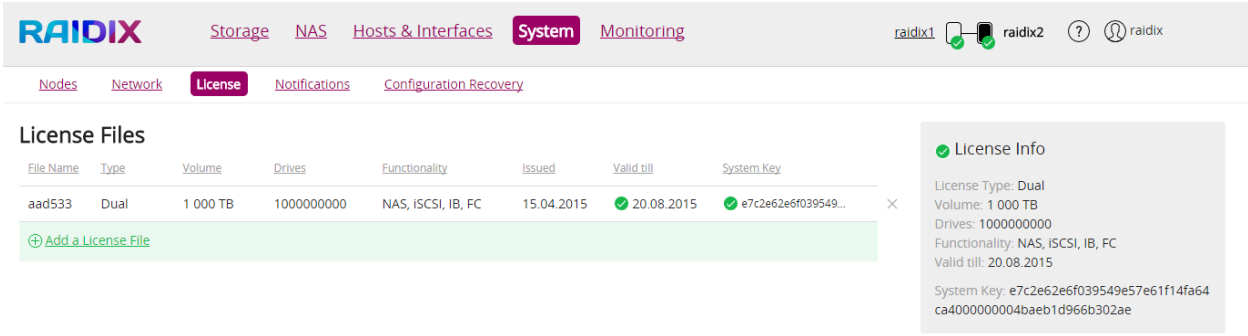







Figure 98 License page

To start working with the system, it is necessary to add a valid license file. For more details, refer to *RAIDIX 4.4 Installation and Configuration Guide*.

The table below provides detailed information on each downloaded license file:

| Parameter | Description |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| File name | The name of a downloaded license file |
| 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 |

| Parameter | Description |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Drives | Number of drives available for RAID creation |
| Functionality | Available functionality for current license |
| Issued | License issue date |
| Valid till | License expiry date.  OK – the license is valid.  Bad – the license has expired. |
| System Key | License request-key.  – the key is valid; no hardware modifications were made.  – the key is valid; one hardware modification was made.  – 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. |

8.3.2. Updating the Licensing Parameters



You need to update and prolong the licensing parameters when updating the hardware components or when license period has expired.

To update the license parameters, perform the following:

1. Copy the current system license key (**Key** parameter of page **License**).
2. Send the license key of your system to the RAIDIX technical support.
3. Download the received license file by clicking **Add License File** on the **License** page.

8.4. Setting Up Notifications

RAIDIX 4.4.3 enables sending system state messages using e-mail. The notifications can contain the information on the following subjects: RAID, iSCSI, Drive, Network, SparePool, License, Controller Sensors, and Enclosure Sensors. The types of notification messages can be selected by the user: **Info**, **Warning**, or **Error**.

8.4.1. Set Up Server Parameters

To set up server e-mail notification parameters, perform the following procedure:



You can configure server parameters once and use it in the administration of all email alerts

1. In the **System** menu, open the **Notifications** page.
2. In the **E-mail Notification** section click **Server Parameters**. A window **Server E-mail Notification Parameters** appears (Figure 99).

| Server Parameters ^ | |
|----------------------------------------------------------------------------|---------------------------------|
| IP/DNS Name | <input type="text"/> |
| Port | <input type="text" value="25"/> |
| Type of Authorization | <input type="text" value="On"/> |
| Encryption | <input type="checkbox"/> Off |
| Sender E-mail | <input type="text"/> |
| User Name (Optional) | <input type="text"/> |
| Password (Optional) | <input type="text"/> |
| <input type="button" value="Apply"/> <input type="button" value="Cancel"/> | |

Figure 99 Server E-mail Notification Parameters window

3. Define the following server parameters:

| Parameter | Description |
|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| IP/DNS Name | IP address or DNS name of the system from which notifications are sent. |
| Port | Port number |
| 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 |
| Encryption | Encryption status. Possible values: <ul style="list-style-type: none">• Off• On |
| Sender E-mail | An e-mail address of the notification sender. |
| User Name (Optional) | A user sends notifications. The parameter can be not specified. |
| Password (Optional) | Password for e-mail server access. The parameter can be not specified. If the server security policy allows unauthorized access, you can leave the field blank. |

4. Click **Apply**.

8.4.2. Set Up Subscriber Parameters

To set up subscriber parameters, perform the following procedure:

1. In the **System** menu, open the **Notifications** page and perform [setting up server parameters](#).
2. In the **E-mail Notification** section, click **Add a Recipient**. The pane to create e-mail notification appears (Figure 100).

Figure 100 Notifications page. Adding a recipient

3. In the **Recipient** field, type an e-mail address of the notification receiver.
4. In the **Notifications** field, click **Select** and check the types of notifications to be sent.
5. Click **Create**. A new subscriber e-mail address will appear in the subscribers list (Figure 101).

Figure 101. Recipient added

6. To verify that set up was correct, click **Test** in the appropriate user line. A test e-mail should be sent to the specified user e-mail address.

8.5. 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 will be applied immediately.

Use the elements of the **Configuration Recovery** page of the **System** menu to manage the system configuration files (Figure 102).

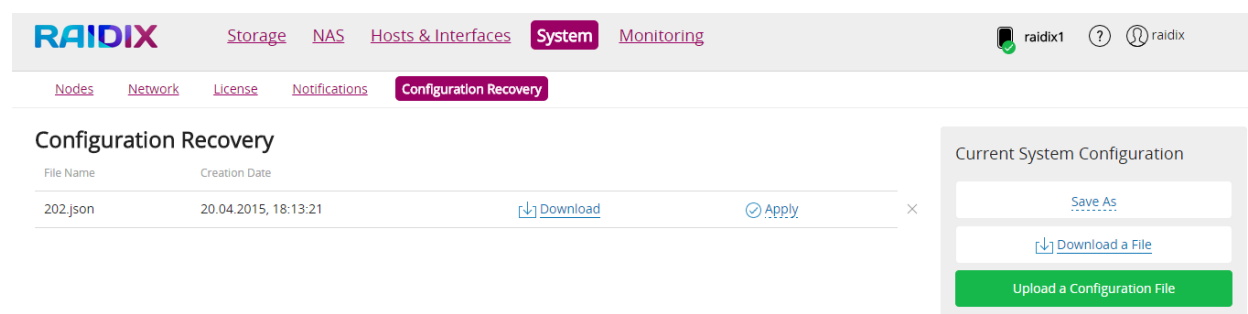


Figure 102. Configuration Recovery page

The tools on the **Configuration Recovery** page enable the following operations:

1. [Saving Configuration Files.](#)
2. [Downloading Configuration Files.](#)
3. [Uploading Configuration Files.](#)
4. [Applying Saved Configuration Files.](#)
5. [Deleting Configuration Files.](#)



Do not use configuration recovery after changing hardware system configuration or creating an array of a different level.

Do not initiate system configuration recovery until the RAID array initialization is completed.

8.5.1. Save Configuration File



To be able to restore current system settings after modifying them, it is recommended to create a backup copy of system configuration file regularly.

To save the current system configuration, perform the following procedure:

1. In the **System** menu, move to the **Configuration Recovery** page.

2. Click **Save As**, the **Save Current Configuration as** window appears (Figure 103):




Figure 103 Saving configuration

3. Specify the file name and click **Save**.

8.5.2. Download Configuration Files

To download a configuration file to your computer, click the **Download a File** on the **Configuration Recovery** page. The configuration file will be saved.

8.5.3. Upload Configuration Files

1. To upload a configuration file from your computer, click **Upload a Configuration File** on the **Configuration Recovery** page.
2. Select uploaded configuration file and click **Open**. The configuration file will be uploaded.


8.5.4. Apply Saved Configuration File

To apply a configuration file from the list of saved files, click **Apply** in the corresponding file line on the **Configuration Recovery** page. The configuration file will be applied.



If you have rearranged the drives or enclosure to a workable system, set up the **network**, **target** and **masking** settings.

8.5.5. Delete Configuration File

To delete a configuration file, click  of the corresponding configuration file on the **Configuration Recovery** page. Click **Yes** in the confirmation message.

8.5.6. Advanced Configuration of the System

8.5.6.1. System Configuration Cloning

In RAIDIX 4.4.3 software 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 **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 output of the `rdclone -h` will look like this:

```
[root@Trump utils]# rdclone -h
Tool for configuration cloning.
Usage:  rdclone  [-h]  [-t]  [-f <configuration_file_path>]  [-o
<output_file_path>]
    -f  Path to the configuration file.
        If not defined - current configuration will be used.
    -F  Read configuration from STDIN
    -h  Show this help text and exit.
    -o  The path to the file where set of rdcli commands will be
written.
        If not defined - STDOUT will be used.
```

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 similar hardware, on which you want to recreate configuration, should not contain any other objects in the configuration file.

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 the `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 "??".

8.5.6.2. Connecting Data Drives of Another System to the Working System

In RAIDIX 4.4.3 you can import another 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 was correctly restored (Online). If RAID was 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.

To connect data drives of another system to the working system through the web-interface, do the following:

1. In the main menu select **Storage** and move to the **RAID Import** page.
2. In the **RAID Import** list, select RAID for import.
3. In the drop-down list, select the cache size for the RAID you want to import (Figure 104)

| RAID Name | RAID Level | Drives | Number of drives | Status | Cache |
|-----------|--------------|--------------|------------------|----------|--------------|
| R2 | RAID0 | 5 | 2 | ✗ | |
| R1 | RAID6 | 11 | 7 | ✗ | |
| R3 | RAID0 | 12,13,15,16 | 5 | ✗ | |
| R6i | RAID6 | 14 | 6 | ✗ | |
| R4 | RAID0 | 17-21 | 5 | ✓ | 1 GiB |
| R5 | RAID10 | 22-25 | 4 | ✓ | |
| Rnmi | RAID N+M | 26 | 13 | ✗ | |
| Rnumi | RAID N+M | 27 | 10 | ✗ | |

Figure 104 RAID Import

4. Click **Get Volumes**.

It is not necessary for the import if your system does not have RAIDs or LUNs with the same names or does not have RAIDs at all.

5. 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/or volumes. To rename RAID and/or volumes, do the following:
 - a. Click **Rename drives metadata**.

- b. In the appeared dialog, select volumes that belong to the RAID for import, specify **New RAID name** and click **Select**.
6. Click **Import**.

To connect data drives of another system to the working system through the CLI, do the following:

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 was not specified, the search will be based on all available drives.

The `rdctl` utility will find and display information about RAIDs that were 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 needed for the disambiguation when several RAIDs have the same name, and you need to understand 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 -n|--name <raid_name>
```

This command will upload the 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 -n|--name <raid_name>
```

9. System Monitoring

To control the system's state and to operate the system via [command console](#), use the **Monitoring** menu. 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](#).

9.1. Information on Sensors

To work with sensors open the **Monitoring** menu item and move to the **Sensors** page.

The **Sensors** page provides information on the controller sensors and sensors of the connected drive enclosures.




9.1.1. Information on Controller Sensors

To access to the controller sensors, open the **Monitoring** menu item and move to the **Sensors** page. 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 top table (Figure 105) on the **Sensors** page.

| Name | Value | Status |
|--------------------------------------------------|---------|--------|
| IOH Therm Margin | -47.000 | ✓ |
| BB -12.0V | -11.959 | ✓ |
| PSGet SDR 004b command failed: Unspecified error | --- | ✓ |
| PS1 Curr Out % | 20.000 | ✓ |
| PS1 Power In | 184.000 | ✓ |
| P1 Therm Margin | -72.000 | ✓ |
| BB +5.0V STBY | 5.107 | ✓ |
| BB +3.3V | 3.353 | ✓ |

Figure 105 Sensors page. Information on nodes sensors

| Parameter | Description |
|-----------|-------------|
| Name | Sensor name |

| Parameter | Description |
|-----------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Value | Current sensor value |
| Status | Current sensor status: <ul style="list-style-type: none"> OK – works correctly; Warning – sensor is not defined; Error – the critical error occurred, or the measured value exceeds the allowed limit. |

If the sensor results are outside the allowed value 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 e-mail notifications (if the [appropriate settings](#) were performed)

9.1.2. Information on Enclosure Sensors

To access to the enclosure sensors, open the **Monitoring** menu item and move to the **Sensors** page. In the left part of the page select the **Enclosure Controllers**, information about the power supply, cooling mechanism and sensors of the connected enclosure is displayed (Figure 106).

RAIDIX

Storage

NAS

Hosts & Interfaces

System

Monitoring

RAIDIX

?

raidix

Sensors

Performance

Support

History

Controller Sensors

Enclosure Sensors

Enclosure Sensors

| Enclosure Logical ID | Enclosure Vendor ID | Product ID | Product Revision Level |
|----------------------|---------------------|------------------|------------------------|
| 50014380052bf340 | HP | D2700 SAS AJ941A | 0038 |

Power Supply




| DC OverVoltage | DC UnderVoltage | DC OverCurrent | Failure | Off | OverTemp Failure | OverTemp Warning | AC Failure | DC Failure | Disabled | Status |
|----------------|-----------------|----------------|---------|-----|------------------|------------------|------------|------------|----------|--------|
| No | No | No | — | No | No | No | No | No | — | ✓ |
| No | No | No | — | No | No | No | No | No | — | ✓ |

Temperature

| Temperature | OverTemp Failure | OverTemp Warning | UnderTemp Failure | UnderTemp Warning | Disabled | Status |
|-------------|------------------|------------------|-------------------|-------------------|----------|--------|
| 32 °C | No | No | No | No | — | ✓ |
| 23 °C | No | No | No | No | — | ✓ |







Figure 106 Sensors page. Information on enclosure sensors.

| Parameter | Description |
|------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Enclosure Logical ID | Unique enclosure ID |
| Enclosure Vendor ID | Enclosure Manufacturer ID |
| Product ID | Enclosure ID in INQUIRY format |
| Product Revision Level | Enclosure firmware version |
| Power Supply | |
| DC OverVoltage | Parameter indicates whether an overvoltage condition has been detected at the power supply output. Possible variants: <ul style="list-style-type: none"> Yes No |
| DC UnderVoltage | Parameter indicates whether an undervoltage condition has been detected at the power supply output. Possible variants: <ul style="list-style-type: none"> Yes No |
| DC OverCurrent | Parameter indicates whether an overcurrent condition has been detected at the power supply output. Possible variants: <ul style="list-style-type: none"> Yes No |

| Parameter | Description |
|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Failure | <p>Parameter returns Information about the possible failure in the sensor work (the system cannot read its state). Possible values:</p> <ul style="list-style-type: none"> Yes – a problem has occurred, impossible to read sensor state. No – sensor works correctly. |
| Off | <p>Parameter returns information about the elements state.</p> <ul style="list-style-type: none"> Yes – element (power supply or cooler) is disabled No - element (power supply or cooler) is disabled but sensor is working |
| OverTemp Failure | <p>Parameter indicates the power supply has detected a temperature above the safe operating temperature range. The power supply may shut down. Possible variants:</p> <ul style="list-style-type: none"> Yes No |
| OverTemp Warning | <p>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.</p> |
| AC Failure | <p>Bit set to one indicates that the power supply is not receiving the specified A.C. power.</p> |
| DC Failure | <p>Bit set to one indicates that the power supply is unable to supply the specified D.C. power.</p> |
| Disabled | <p>If sensors were disabled by external command, parameter indicates the value Yes. If no disabling was performed, parameter indicates the value No.</p> |
| Status | <div>  OK – works correctly; </div> <div>  Warning – sensor is not defined; </div> <div>  Error – the critical error occurred, or the measured value exceeds the allowed limit. </div> |

Cooling

| | |
|------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Actual Fan Speed | Actual fan speed (in rotations per minute). |
| Failure | <p>Parameter returns Information about the possible failure in the sensor work (the system cannot read its state). Possible values:</p> <ul style="list-style-type: none"> Yes – a problem has occurred, impossible to read sensor state. |

| Parameter | Description |
|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | <ul style="list-style-type: none"> No – sensor works correctly. |
| Off | <p>Parameter returns information about the elements state.</p> <ul style="list-style-type: none"> Yes – element (power supply or cooler) is disabled No - element (power supply or cooler) is disabled but sensor is working |
| Actual Speed Code | Current cooling mode |
| Disabled | If sensors were disabled by external command, parameter indicates the value 1 . If no disabling was performed, parameter indicates the value 0 . |
| Status | <p> OK – works correctly;</p> <p> Warning – sensor is not defined;</p> <p> Error – the critical error occurred, or the measured value exceeds the allowed limit.</p> |
| Temperature | |
| Temperature | Sensor temperature, °C. |
| UnderTemp Failure | Bit set to one indicates that the temperature is below the safe operating temperature range or lower than the value indicated by the critical low threshold. |
| UnderTemp Warning | A under temperature warning bit set to one indicates that the temperature is below the normal operating temperature range or lower than the value. |
| Disabled | If sensors were disabled by external command, parameter indicates the value Yes . If no disabling was performed, parameter indicates the value No . |
| Status | <p> OK – works correctly;</p> <p> Warning – sensor is not defined;</p> <p> Error – the critical error occurred, or the measured value exceeds the allowed limit.</p> |

9.2. Information on Performance

Performance monitoring module helps to map up configuration, estimate and increase the performance both inside of the storage system itself and during the data transfer process. To access the system performance data, open the **Monitoring** menu and move to the **Performance** page (Figure 107).

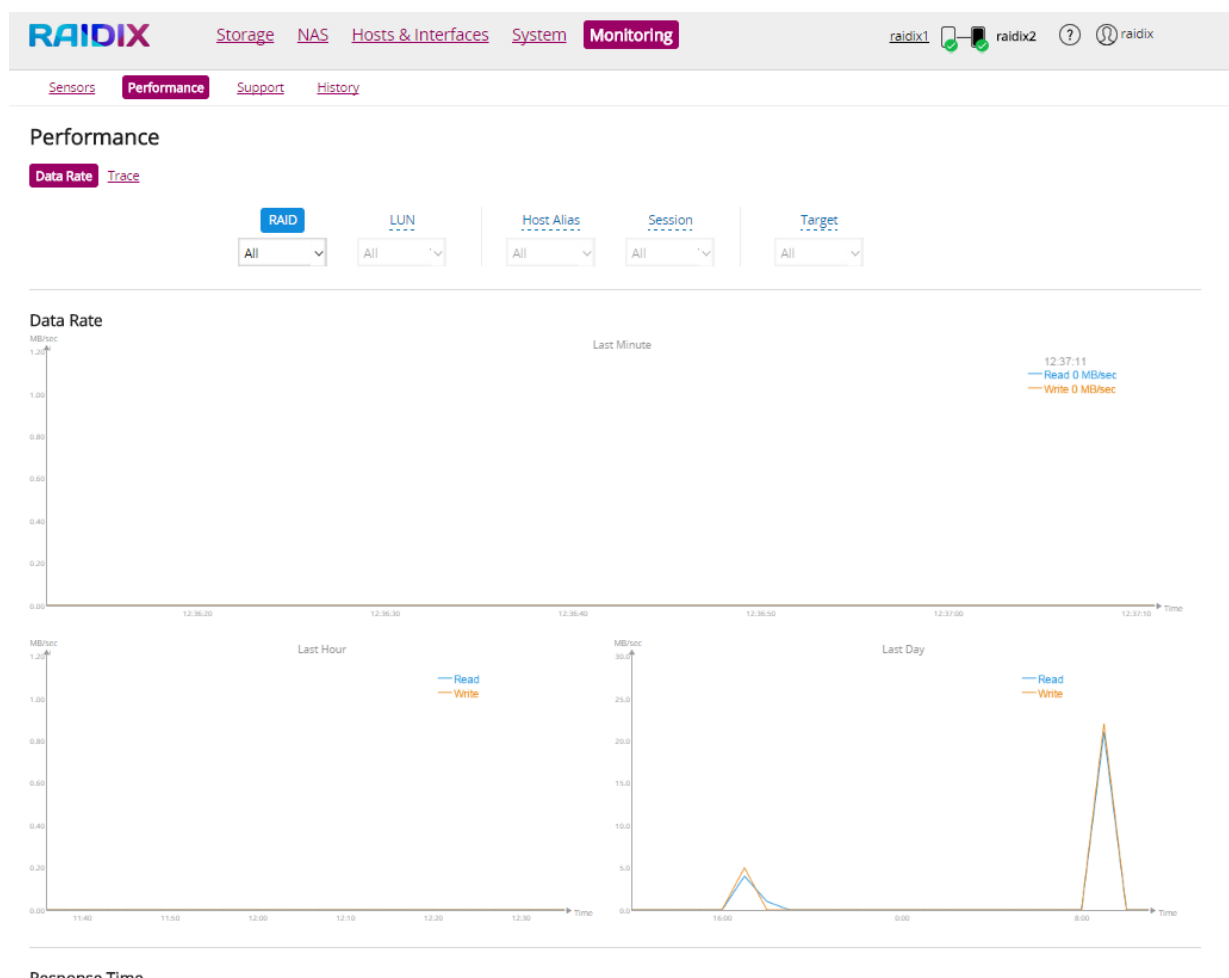


Figure 107 Performance monitoring module. Data Rate section fragment

The module consists of two functional sections: [Data Rate](#) and [Trace](#).

9.2.1. Data Rate

Data Rate (Monitoring | Performance | Data Rate) section provides detailed information on the system performance and allows making assessment of the data exchange rate using the following charts:

- Data Rate;
- Response Time;
- Transfer Size.

You can set up the chart display parameters using the following filters on the top of the **Data Rate** page (Figure 108):

- **RAID** – filter provides statistics for all/ specified RAIDs;
- **LUN** - filter provides statistics for all/ specified LUNs;
- **Host Alias** – filter provides statistics for all/ specified created Host Aliases;
- **Session** - filter provides statistics for all/ specified established host sessions;
- **Target** – filter provides statistics for all/ specified existing targets.

Performance



Figure 108 Data Rate page filters

To set up the displayed parameters, click the appropriate filter name and select the information to display in the drop-down list appeared.

9.2.1.1. Data Rate Charts

Data Rate charts (Figure 109) 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 109 Data Rate charts

The vertical axis of the charts shows the operations speed (in MB/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 110).




Figure 110 Last Hour chart data detalization

Today chart displays the last day data rate. To view the detailed information to an hour, point your mouse to the graph.

9.2.1.2. Response Time Charts

Response Time charts allow evaluating the system activity of command processing at different time intervals (Figure 111). 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).

Response time statistics are displayed on the following charts:

- Last Minute Summary,
- This Minute Summary,
- Last Hour Summary,
- This Hour Summary,
- Last Day Summary,
- Today Summary.

Response Time

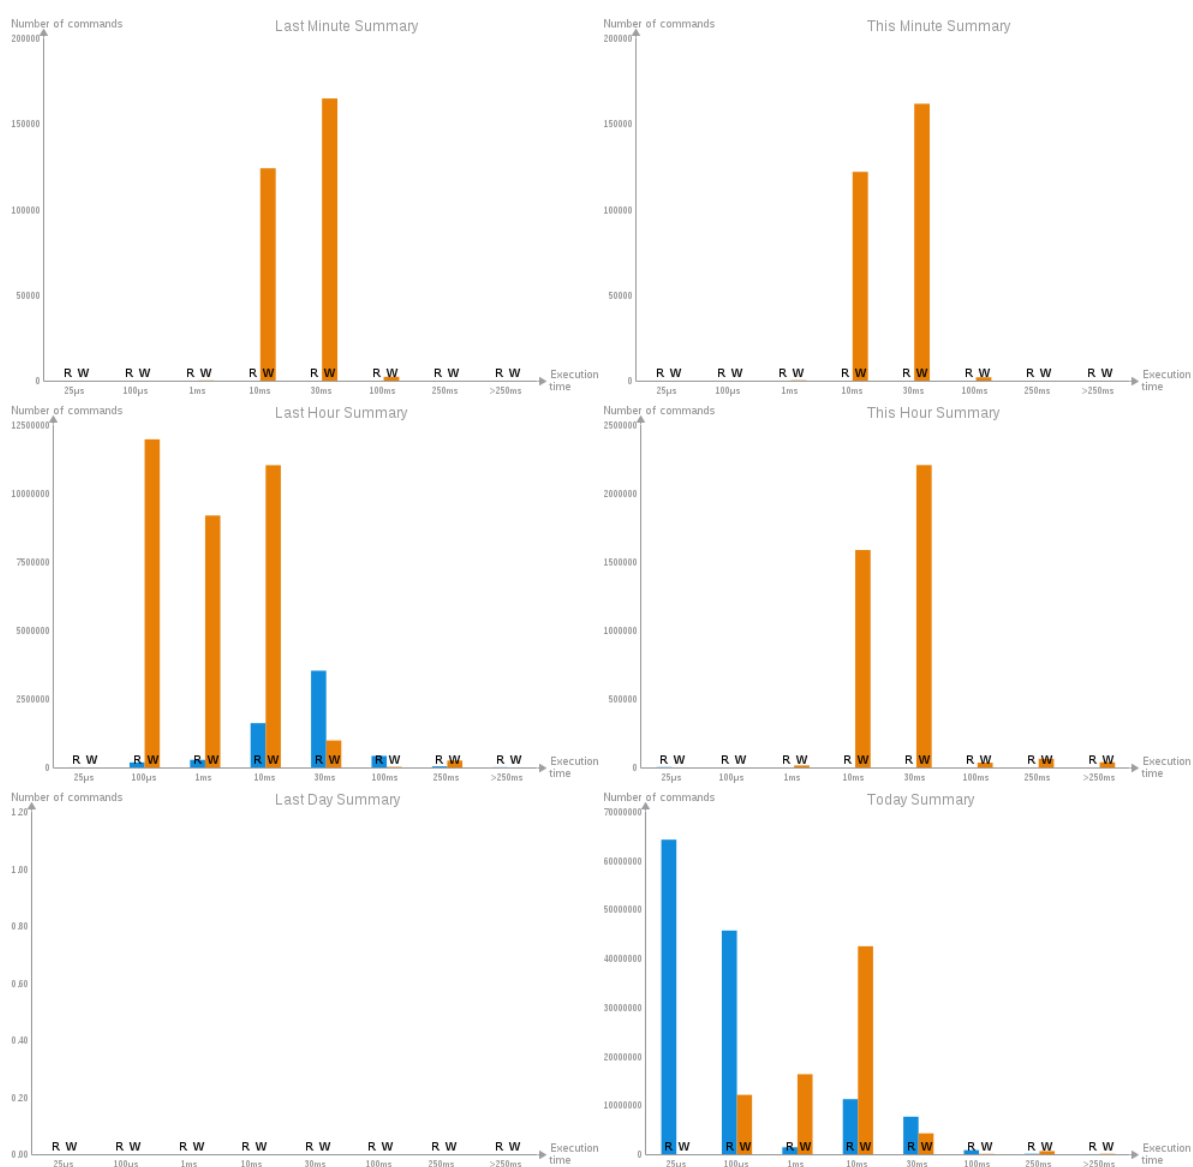


Figure 111 Response Time charts (fragment)

9.2.1.3. Transfer Size Charts

Transfer Size charts (Figure 112) 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.

Transfer Size

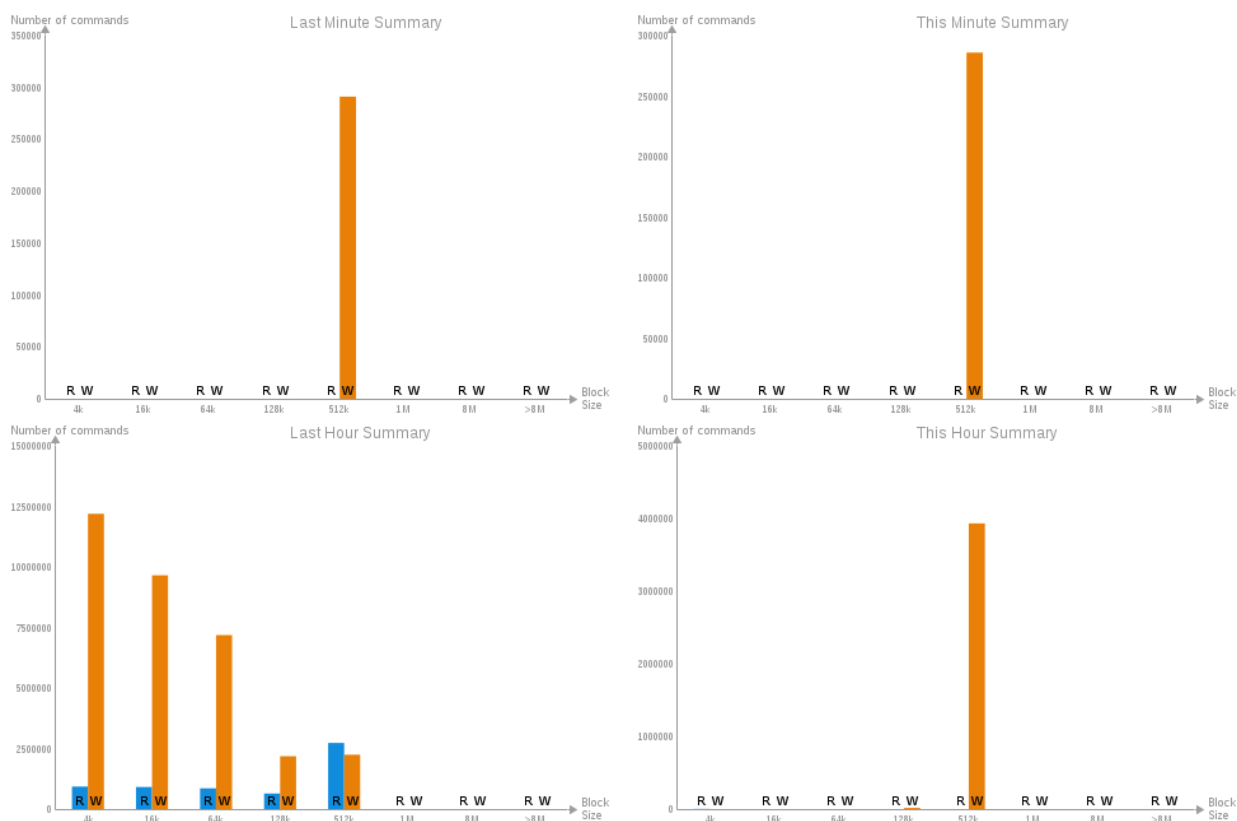


Figure 112 Transfer Size charts (fragment)

9.2.2. Trace Statistics

Trace section (System | Statistics | Trace) is aimed at engineers, performing the primary selection of storage configuration and commissioning operations (Figure 113).

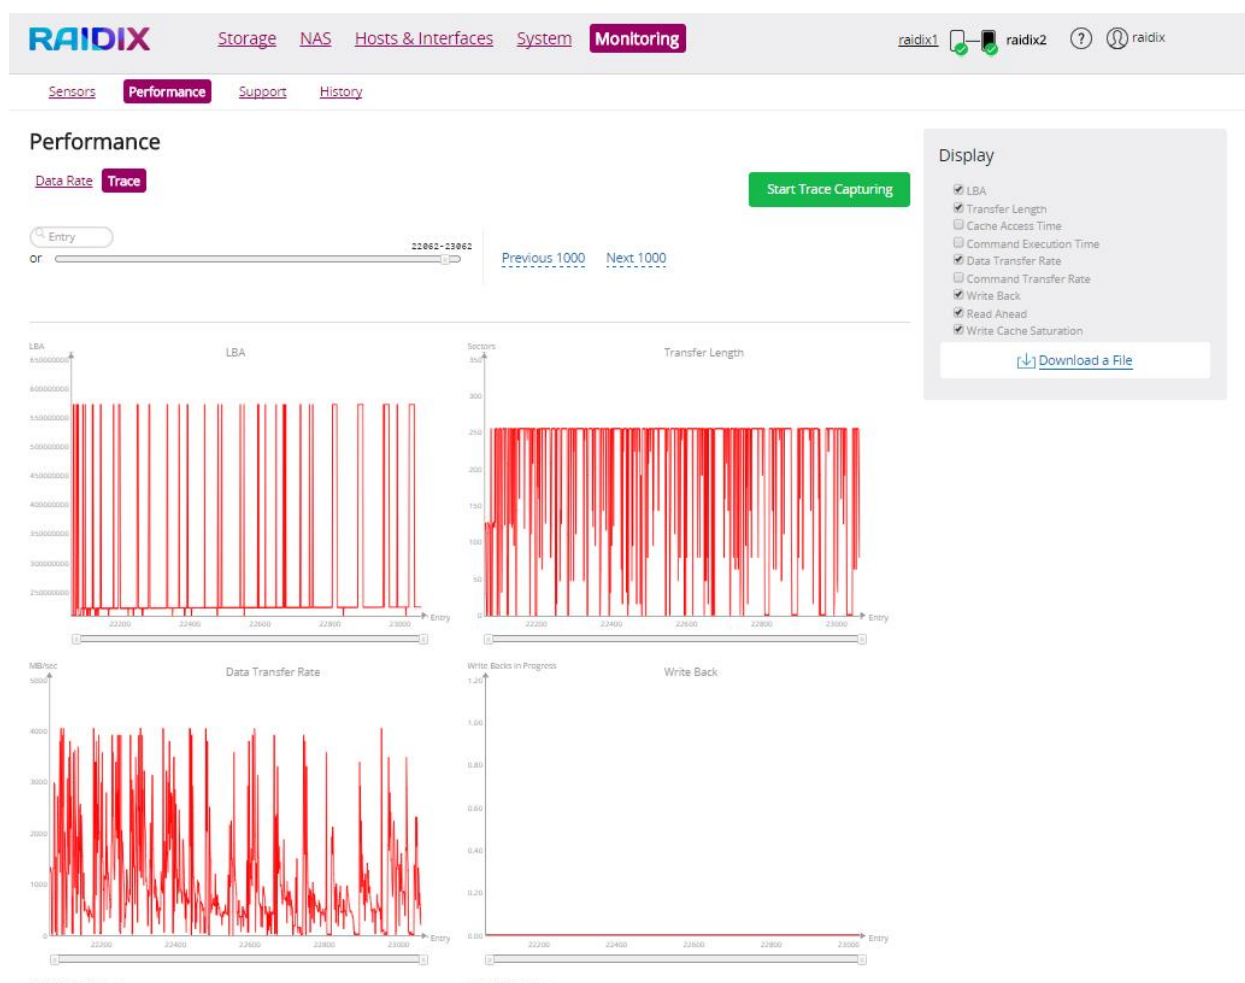


Figure 113 Performance monitoring module. Trace section fragment

The displayed information can be set using the control panel filters at the top of the page (Figure 114). 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);
- 1000 previous or next entries (click the corresponding icon in GUI).

To start statistics capturing, click the **Start Trace Capturing** (Figure 114). Collecting information about processing SCSI-commands by the system will start. Trace capturing and processing of the results will take not more than 30 seconds. Then the results will be displayed on the charts. To stop trace capturing earlier click **Stop Trace Capturing**.

Performance

[Data Rate](#) **Trace**[Start Trace Capturing](#)

Figure 114. Trace page control panel filters

In the right side of the control panel **Display**, the filter is implemented. To display a chart, check an appropriate item. The following charts are available in **Trace** section (6 of 10 are shown by default) (Figure 115):

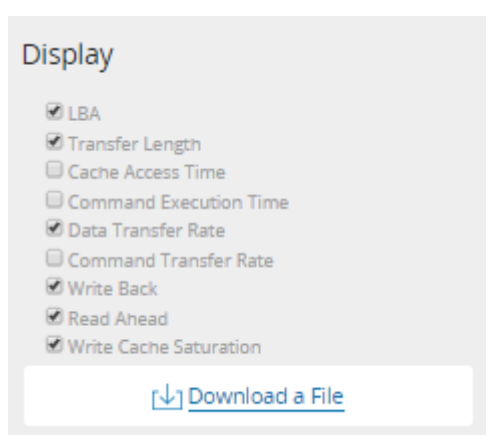


Figure 115 Display filter

You can download the traces in csv format by clicking the **Download a File** (Figure 115).

You can also scale the graph, selecting an appropriate area on it with the mouse (scaled area will be marked as the light-blue color of the horizontal axis (Figure 116)).




Figure 116 Scaling example

1. **LBA Graph** shows the logical blocks address from which the data transfer starts and the received commands (Figure 117):




Figure 117 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.

2. **Transfer Length** graph (Figure 118) shows the volume of data, transferred by commands; information on the each command size is represented in sectors.




Figure 118 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.

3. **Cache Access Time** graph shows the time during which commands were in the cache queue waiting to be run (Figure 119).




Figure 119 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).

4. **Command Execution Time** shows time devoted to commands execution (Figure 120).




Figure 120 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).

5. **Data Transfer Rate** shows the appropriate command execution data transfer rate (Figure 121).




Figure 121 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 MB/sec).

6. **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 122).




Figure 122 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 MB/sec).

7. **Write Back** graph (from cache to disk data transfer statistics) shows the number of Write Back commands during a particular command execution (Figure 123).




Figure 123 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.

8. **Read Ahead** graph shows read ahead operations selection (Figure 124).




Figure 124 Read Ahead graph example

The horizontal axis is an internal entry number of the appropriate command execution (read or write) and the vertical is the number of read-ahead requests.

9. **Write Cache Saturation** graph shows the value of cache saturation with dirty data (non-synchronized with disk cache data) (Figure 125).




Figure 125 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. As higher the Dirty cache segments amount is the quicker the commands get access to the cache.

9.3. Downloading System Logs

In RAIDIX 4.4.3 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 from the **Support** page of **Monitoring** menu, perform the following:

1. Click **Collect System Logs** (Figure 126).

System Logs

[Collect System Logs](#)

Figure 126 Collecting System Logs

2. Wait until the direct link to download logs appears and click the link (Figure 127). System logfile will be downloaded on your computer.

System Logs

[Collect System Logs](#)

[Download System Logs](#)

Figure 127 Downloading System Logs

9.4. Command Console Interface

This section contains information on the console command interface.

Command Console lets the administrators perform more functions in addition to those operations available through the web UI. The command line has a number of requirements for the typed commands. For details on syntax and the list of available commands, see [Appendix A](#).

To access the console, open the **Monitoring** menu item and move to the **Support** page. Click **On** to display console (Figure 128).

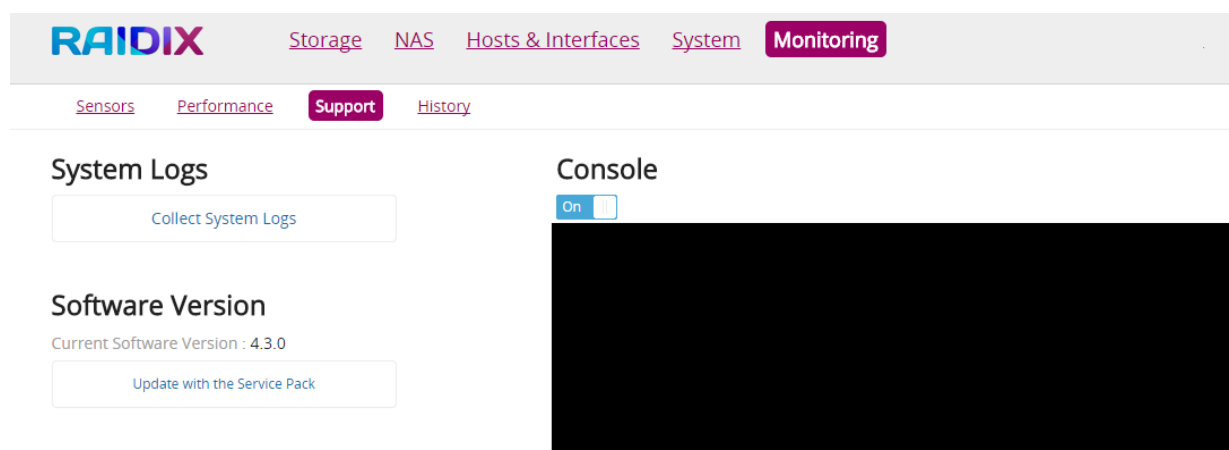


Figure 128 RAIDIX Command Console

9.5. System Status Assessment

To view detailed information on system alerts, open the **Monitoring** menu and move to the **History** page (Figure 129). Specify the number of alert items in the **Number of Alerts** field and click **Reload**.

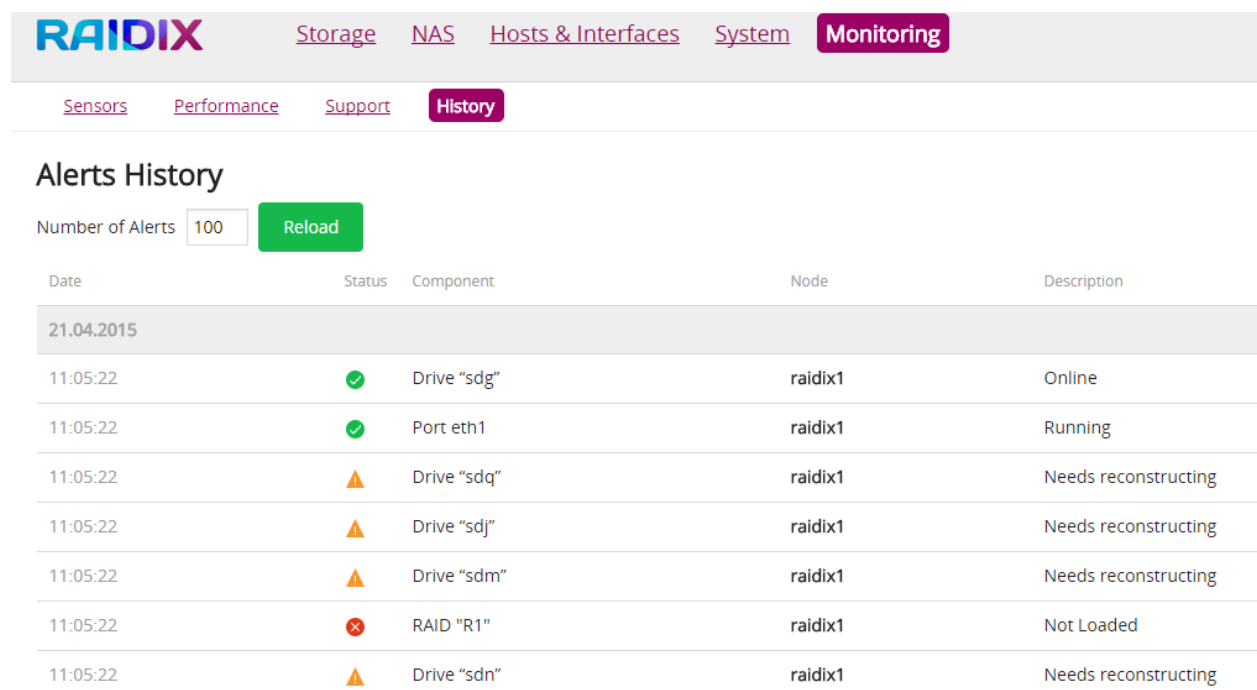


Figure 129 Alerts History page. Information on system's status

To find an **alert**, type any date, component or description in the search field (Figure 130).

Alerts History

Number of Alerts [Reload](#)

22.04

previous 1 2 3 next

| Date | Status | Component | Node | Description |
|------------|--------|--------------|------|-------------|
| 22.04.2015 | | | | |
| 8:17:54 | ✓ | Port "eth0" | Tony | Running |
| 8:13:23 | ✓ | Port "bond0" | Tony | Running |

Figure 130 Alerts searching

The following information on alerts is presented:

| Parameter | Description |
|-----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Date | Date and time then alert occurred |
| Status | <p>Parameter defines node devices state:</p> <ul style="list-style-type: none"> Healthy – works correctly; Info – information. The current status is displayed in case the Ethernet port not involved in the data transmission and unavailable. Warning – a non-critical failure has occurred: <ul style="list-style-type: none"> RAID status is Degraded, Reconstructing, Initializing, SDC, or array is active on non-preferable node (opposite to its affinity); Adapter status is Down (there is no physical connection); Drive status is Dirty, Faulty Warning, Needs Reconstructing, Reconstructing; License status is Warning because a hardware modification was made. Bad – 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 Faulty error, No drive or Error; RAID is Offline, Not Loaded, or an error during the initialization has occurred. |
| Component | <p>Alert component. Possible values:</p> <ul style="list-style-type: none"> Drive Port RAID License |


| Parameter | Description |
|-------------|--------------------------------------------|
| Node | Node name |
| Description | A description of the problem device status |

10. Troubleshooting

This section contains descriptions of possible errors and steps required to solve them.

| Nº | Problem description | Recommendations |
|----|--------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 1 | GUI is not available | <ul style="list-style-type: none"> Probable reason is that the Ethernet interface of the node is not available. <p>Verify nodes availability using the appropriate IP-address.</p> <p>If the node is unavailable, it is most likely that the Ethernet interface is not operational.</p> <p>If the node is not operational, check if there is enough of free memory space on the drive (<code>df -h</code>) and enough of the storage memory (<code>free -h</code>).</p> <div> <p>Turn off the node with failed Ethernet interface manually:</p> <ul style="list-style-type: none"> In single-controller mode, the node can be turned off only in case of total absence of workload In dual-controller mode before turning off the failed node, transfer (failover) all arrays to the healthy node (by clicking Failover) </div> |
| 2 | The drives are unavailable in GUI after the new drive enclosure connection or after the reconnection of already used enclosure | <ul style="list-style-type: none"> Reboot the system. If system reboot did not help, reboot the bin and reboot the system again. |
| 3 | Current RAID state is Not Loaded | <p>The state means that the configuration exists but fails to be loaded. To load the array, use one of the following ways:</p> <ul style="list-style-type: none"> Verify the specified parameters value (including an increase of the cache size). <p>Enter the following CLI commands to restore RAID and LUN on it:</p> <pre>rdcli raid restore -n <RAID_name> and rdcli lun restore -n <LUN_name></pre> |

| № | Problem description | Recommendations |
|----|-----------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | <p>Reboot the system.</p> <p>Reload the array.</p> |
| 4 | Current LUN, Share, Mask state is Not Loaded | <p>Perform the following commands:</p> <pre>rdcli raid restore -n <имя_RAID> rdcli lun restore -n <имя_LUN></pre> |
| 5 | Synchronization does not start automatically | <ul style="list-style-type: none"> • Check that all connection status of all ports for synchronization is Up; • Check that targets for synchronization were selected correctly; • Verify OpenSM service is running on the first node. On the proper port (if InfiniBand is used). • When synchronizing over iSCSI verify that iSCSI is enabled, and IP-address of the remote controller, used for synchronization, is available. |
| 6 | If you renamed RAIDs and volumes, there is a possibility that LUNs will not display the file system after the import. | <ul style="list-style-type: none"> • In CLI, after the import, run the following command: <pre>rdcli lun modify --name <lun_name> --formatted yes</pre> |
| 7 | Errors occurring during upgrading of the firmware on the SAS target. | <ul style="list-style-type: none"> • You must upgrade the firmware on the SAS target by using Isiutil (not through the SAS Flash). |
| 8 | Errors occur due to the firmware of the SAS adapters. | <ul style="list-style-type: none"> • When you are working with adapters that are based on the LSI SAS2308 6 GB (9201 / 9205 / 9207 / 9211) chip, make sure, that the version of the firmware is P19. |
| 9 | Errors occurring in the process of configuration or after the configuration of the SAS adapters' ports. | <ul style="list-style-type: none"> • You must set up adapters' ports by using Isiutil and through dividing them into groups (see, <i>RAIDIX 4.4 Software Installation and Configuration Guide</i>). |
| 10 | If the Linux initiator with the ATTO SAS was loaded before the LSI SAS target, the initiator will not define LUNs. | <ul style="list-style-type: none"> • We recommend you to install the LSI SAS card on the Linux initiator. • If the Linux initiator with the ATTO SAS card does not identify LUNs of the LSI SAS target, you must disconnect the SAS cable from either the initiator or the target and connect it back. After that action, all LUNs will be defined. |

| Nº | Problem description | Recommendations |
|----|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 11 | <p>The Isscsi initiator displays the following record before adding the LUNs:</p> <pre>disk Raidix N/A 4.4. - disk Raidix N/A 4.4. -</pre> <p>Or the following:</p> <pre>disk mpt3sas mpt 3sas target 4161 -</pre> | <ul style="list-style-type: none"> Remove devices by running the following command: <pre>echo 1 > /sys/class/scsi_device/1\:0\:19\:0/dev ice/delete</pre> <p>And perform <code>rescan devices</code>.</p> After you performed <code>rescan devices</code>, find all LUNs by performing the following command: <pre>echo '- - -' > /sys/class/scsi_host/host1/scan</pre> Perform the <code>lsscsi</code>, the result should be similar to the following: <pre>[root@raidix ~]# lsscsi disk Raidix LunRDX 4.4. /dev/sdb disk Raidix LunRDX 4.4. /dev/sdc disk ATA WDC WD5003ABYX-0 01.0 /dev/sdq disk ATA GB0750EAFJK HPG1 /dev/sdr</pre> <div>  <ul style="list-style-type: none"> Due to the nature of the operating system, it is not possible to perform LUN rescan on MAC OS. </div> |
| 12 | <p>When you write a large file on the LUN with the help of the MAC OS file manager, the auxiliary 4K file appears. You will not be able to delete this file if you configured WORM with a small <code>grace_period</code>.</p> | <ul style="list-style-type: none"> Try to establish a longer <code>grace_period</code>, to have time to delete the file. Log in to the RAIDIX through <code>ssh</code> and delete the file manually. |
| 13 | <p>The maximum number of LUNs that you can create is 487. OS Windows</p> | <ul style="list-style-type: none"> If you need to use all 487 LUNs, we advise you to use Linux, as it defines all the LUNs |

| Nº | Problem description | Recommendations |
|----|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | defines the limited amount of LUNs: from 0 to 230. | <ul style="list-style-type: none"> If you use Windows, we recommend to create not more than 231 LUN so the system will see all of them. |
| 14 | Due to the prohibition on the creation of the LUNs, it is not possible to upgrade DC and save the RAIDs when you are upgrading from the RAIDIX 3.x version of the system. | <ul style="list-style-type: none"> Do not upgrade until the issue is resolved in future versions. Transfer all data from the RAID to the temporary storage and then delete the RAID. Upgrade the system to the RAIDIX 4.4.3 version, configure RAID and upload all data from the temporary storage on this RAID (or, after the upgrade and RAIDs reconstruction, create new RAID on which you will transfer data from the temporary storage and delete the old RAID). Upgrade the system and use it in the single controller mode. |
| 15 | It is impossible to run the initialization of the RAID in the DC mode after the use of the Single Mode. | <ul style="list-style-type: none"> Do not create initialized RAIDs in the Single Mode. |
| 16 | After upgrading from RAIDIX 4.3.1 to RAIDIX 4.4.3, users with the administrator rights may not be displayed for the SMB protocol. | <ul style="list-style-type: none"> After upgrading, you must add users with administrator rights to the SMB either through the web interface or the CLI (see, Manage NAS Users and Groups). |
| 17 | When you delete the host alias in the process of QoSmic learning, the Stop Learning button becomes inactive. | <ul style="list-style-type: none"> In the Console, run the following command: <pre>rdcli qosmic teacher stop</pre> This command will stop the process of learning. |
| 18 | After the rebooting of the both nodes in the Dual Controller mode in ESXi 6.0, reboot/ crash of the LUN adapters are not displayed. | <ul style="list-style-type: none"> You need to scan both adapters twice, |
| 19 | Thin provisioning disabled by default. | The thing provisioning mode is disabled by default, because when a logical size of RAO LUN is larger than physical, creating a file system on initiator can take a long |

| № | Problem description | Recommendations |
|----|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | time. The larger the LUN, the longer it will take to create the file system. |
| 20 | Windows 2008R2 initiator that is connected through InfiniBand in the single controller mode (SRP OFED driver) can work only through one route (cable). | The limitation is due to the specifics of work of many ports cards with OFED initiators. |
| 21 | <p>When you try to perform writing of the new, not deduplicated data on the RAO LUN with enabled thin provisioning, the writing will not start if the amount of data exceeds the RAO LUN drive volume.</p> <p>When you try to perform writing of the data, that is partially identical to the previously recorded data, the writing starts, but will be stopped in case RAO LUN runs out of space. In both cases, the system will display an error.</p> | <ul style="list-style-type: none"> • Delete the new data that you wrote on the RAO LUN. • Use this LUN only in read-only mode. |
| 22 | LUN blocks for ESXi | When you are using ESXi, the block size for the LUN can be only 512 bytes, since ESXi does not support LUNs with a block size of 4KB. |

Glossary

| Term | Definition |
|-----------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| AFP (Apple Filing Protocol) | Network protocol that provides access to files in Mac OS X. |
| Cache | <p>A part of RAM of the array, in which the data for lowering the time required to serve the hosts (initiators) requests, is temporarily stored.</p> <p>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.</p> <p>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</p> |
| 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 | <p>A specific computer program enabling the operating system access to a target device. Usually supplied by a device vendor.</p> <p>In Driver Parameters widget, RAIDIX kernel module raidix_mod is understood under the term 'driver', and prioritization and reconstruction parameters can be set.</p> |
| DSM (Device-Specific Module) | A particular software component supplied by a storage vendor; built into the Microsoft MPIO architecture and supports interaction with the particular hardware model |
| FC (Fibre Channel) | High-speed data transfer interface used for storage networking such as workstations, mainframes, supercomputers, and storage devices |
| 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 |

| | |
|---------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| GUID (Globally Unique Identifier) | Globally unique 128-bit identifier In RAIDIX 4.4.3 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 |
| Host (Initiator) | Initiator is a SCSI-device that starts a data request. In RAIDIX 4.4.3 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 |
| 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 an array. 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 array |
| 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. The service is called <i>multipath</i> |
| 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 |
| 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 RAIDIX 4.4.3 software. |
| PDU (Protocol Data Unit) | Protocol data exchange unit |
| RAID (Redundant Array of Independent Disks) | An array of hardware drives managed by a controller |
| SAN (Storage Area Network) | Storage area network is a dedicated storage network that provides access to consolidated block-level storage (such as disk arrays, 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 (SAS) is a point-to-point protocol used to transfer data between servers and SAS storage devices. |
| Sensor | <p>An element for measuring the state of the disk array components, for example, measuring voltage, fan speed or temperature. The sensor converts the monitored value into a signal, which is useful for measuring, transmission, transfer, storage, and registration of monitored object state information.</p> <p>As a rule, sensor is a circuit, an optical device, a switch or a specialized resistor inside the array</p> |
| 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 |
| SMART (Self-Monitoring Analyzing and Reporting Technology) | <p>A monitoring technology for the hardware assessment using built-in self-diagnosis hardware and a mechanism of forecasting disk failures.</p> <p>A special software performing SMART parameters monitoring and warning users in case of device pre-failure state</p> |
| SMB (Server Message Block) | Network application protocol for remote access to files, printers, and other network resources, as well as for inter-process communication. |
| SparePool | Set of drives, which will be used for hot spare of a failed drive in the array. |

| | |
|------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 |
| 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

Command Console Interface Features and Syntax

The following format is used to type the commands in the command console:

```
rdcli {object} {method} {parameters}
```

Possible values of the {object} parameter:

| | | | |
|---------|---------|--------------|------------|
| adapter | iscsi | notify | sas target |
| dc | lun | param | session |
| | | profile raid | sparepool |
| drive | mask | qosmic | |
| fc | nas | raid | system |
| host | network | | |
| ib | nodes | sensor | |

| Object Name | Functionality |
|------------------------------|---------------------------------------------------------------|
| adapter | Adapters management |
| dc | Cluster management |
| drive | Drives 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 | E-mail notifications settings management |
| param | System parameters management |
| profile raid | RAID profiles parameters management |
| qosmic | Qosmic management |

| | |
|---------------------------|-----------------------------------------------------------------|
| raid | Array parameters management |
| sas | SAS settings management |
| sensor | Sensors control |
| session | Displaying information about all sessions running in the system |
| sparepool | SparePools parameters management |
| system | System settings configuration |

To get the full list of commands, use the command

```
rdcli --fullhelp
```

Command console syntax usage:

1. The command parameters are typed in one line.
2. The parameter in braces is mandatory for typing command ({object}, {method}).
3. The parameter in square brackets is a clarifying one.
4. Command parameters are separated with spaces.
5. Both short and long forms of command attributes can be used. For example:

```
rdcli raid create -n [--name] [type parameter value - raid name] -l
[or --level] [specify array level] -d [--drives] [specify numbers of
drives in the array]
```

6. To get the list of all methods and objects of lower levels for every object, specify attribute **-h** after the object:

```
rdcli {object} -h
```

7. To get the list of all methods and objects of lower levels for every method, specify attribute **-h** after the method:

```
rdcli {object}{method} -h
```

8. Presence of element <value> in the command description indicates the need to enter the parameter value after the corresponding attribute. Parameter value is separated from the parameter with a space, without quotas or brackets.

The table below contains a detailed description of commands.

| Object | Method | Command Syntax | Functionality |
|---------|--------|---------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| adapter | show | <code>rdcli adapter show</code> | Command shows list of all adapters in the system. |
| dc | create | <code>rdcli dc create</code> | <p>Command is used to set up Heartbeat connection to configure DC (Dual Controller) mode.</p> <p>Mandatory attribute:</p> <p><code>-a --ipaddr <value></code> – IP address of the remote node.</p> |
| | delete | <code>rdcli dc delete</code> | <p>Command disables Dual Controller mode.</p> <p>Optional attribute:</p> <p><code>-f --force</code> – forces command execution.</p> |
| | modify | <code>rdcli dc modify</code> | <p>Command allows changing cluster parameters.</p> <p>Optional attributes:</p> <p><code>-t --targets <value></code> – targets for synchronization (comma-separated);</p> <p><code>-a --iscsi_ip <value></code> – IP address for iSCSI synchronization on the remote node;</p> <p><code>-sm --single_mode <value></code> – single node cluster;</p> <p><code>-wws --wt_without_sync <value></code> – enables write through on every RAID without synchronization;</p> <p><code>-pr --pr_sync_enabled <value></code> – enables or disables PR (Persistent Reservation) synchronization. Possible values: 1 (on) – enabled; 0 (off) – disabled.</p> |

| Object | Method | Command Syntax | Functionality |
|--------|----------|--------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | failover | rdcli dc failover | Command performs nodes failover. Optional attribute: -f --force – forces command execution. |
| | 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 | show | rdcli drive show | Command displays info about all drives used in the system. Optional attributes: -f --free – displays all drives which are not in RAIDs; -u --uid <value> – displays drive with specified UID; -r --raid <value> – displays drives in the specified RAID; -s --smart – displays SMART-diagnostics results; -sp --sparepool <value> – displays drives, included in the specified SparePool; -sl --slot – displays the full drives list, sorted by their slot numbers; -ss --ssd – shows SSD drives only. |
| | locate | rdcli drive locate | Command is used to enable LED-indicators. Mandatory attribute: -l --led_locate <value> – enables or disables LED-indicators. |

| Object | Method | Command Syntax | Functionality |
|-------------|--------|------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>Possible values: on (enabled); off (disabled).</p> <p>Optional attributes:</p> <p>-a --all – applies to all drives;</p> <p>-u --uid <value> – enables or disables LED-indicator on the specified drive.</p> |
| | clean | rdcli drive clean | <p>Command removes partitions from the drives in the Dirty state or resets error counter on the Faulty drive.</p> <p>Optional attributes:</p> <p>-u --uid <value> – specifies drive UID;</p> <p>-a --all – applies to all appropriate drives.</p> |
| fc target | show | rdcli fc target show | Command displays the list of Fibre Channel targets used in system. |
| | modify | rdcli fc target modify | <p>Command allows associating the appropriate Fibre Channel targets on the local and remote nodes.</p> <p>Mandatory attributes:</p> <p>-n --target_name <value> – local target name;</p> <p>-r --remote_target <value> – remote target name.</p> <p>Optional attribute:</p> <p>-f --force – forces command execution.</p> |
| fc celerity | show | rdcli fc celerity show | Command displays the list of all Fibre Channel ATTO targets used in system. |

| Object | Method | Command Syntax | Functionality |
|--------|--------|---------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | modify | <code>rdcli fc celerity modify</code> | <p>Command displays Fibre Channel adapters parameters (only for ATTO Celerity), which can be modified.</p> <p>Optional attributes:</p> <p><code>-s --speed_override <value></code> – changes default speed of Fibre Channel card. Possible values: 2, 4, 8, 16 (for 16 Gbit/s card). To use the default speed, set the value: -1;</p> <p><code>-t --topology_override <value></code> – changes default topology of Fibre Channel card. Possible values: 0 (no_override); 1 (for_loop); 2 (for_ptp); 3 (for_loop_preferred); 4 (for_ptp_preferred); -1 (to use default topology);</p> <p><code>--num_w_q <value></code> – number of work queues. Default value is 4;</p> <p><code>--can_queue <value></code> - maximum number of commands per adapter. Default value is 512;</p> <p><code>--event_log_mask <value></code> – specifies a bit mask of events to report to the system log. Default is 1 (critical events only);</p> <p><code>--sgl_page_size <value></code> – scatter/gather list (SGL) page size in number of S/G entries. If your application is doing a lot of very large transfers, you may want to increase the SGL page size. Default value is 128;</p> <p><code>--force_low_dma <value></code> – forces all DMA mappings to below 4 GB. This may severely impact performance and should only be used on systems with 512 GB or more of physical memory layouts.</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>Default is 0;</p> <p>--target_link_delay_mask <value> – enables target mode link delay on a per-port basis. Each bit in the mask represents a physical port. Default is 0x00000000 (all off);</p> <p>--addr_override <value> – address override;</p> <p>--num_r_q <value> – number of receive queue pairs. Default is 2;</p> <p>--num_r_q_128 <value> - number of 128 byte receive buffers. Default value is 256;</p> <p>--old_device_reset <value> – use the old device reset method. Possible values: 1 (on) – use the old device reset method; 0 (off) – use new method. Default is 0;</p> <p>--port_retry_count <value> – maximum number of retries allowed for a port which is logged out. Default is 0;</p> <p>--num_xcb <value> – number of Exchange Control Blocks (XCBs) available for Target Mode requests. Each request requires one XCB. Incoming commands above the set limit are dropped. This setting is only valid when Target Mode is enabled. Default value is 512;</p> <p>--io_time_out <value> – time before an I/O command is timed out by the driver, in seconds. For no timeout set: 0. Default value is 30;</p> <p>--private_els_timeout <value> – timeout, in seconds, for ELS commands when on a private loop. Default value is 2;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>--interrupt_mode <value> – defines the interrupt mode to use. Possible values: 0 (legacy), 1 (MSI), 2 (MSI-X). Default is 2;</p> <p>--max_logins <value> – defines the maximum number of ports that can be logged in at once. Default value is 256;</p> <p>--rport_timeout <value> – time before an FC rport device is timed out by the OS, in seconds. Default is 30;</p> <p>--num_sg_lists <value> – number of SGL pages. Default value is 256;</p> <p>--target_mode_link_delay <value> – delays link initialization until a target mode driver (TMD) connects to the HBA driver. Default value is 0 (do not delay);</p> <p>--target_mode <value> – enables target mode function. Possible values: 1 (on), 2 (off). Default is 1 (on);</p> <p>--per_cpu_msix_cnt <value> – creates one MSIX vector per online CPU. Default is 1;</p> <p>--max_e_q <value> – maximum number of event queues;</p> <p>--initiator_mode_mask <value> – enables initiator mode on a per-port basis. Each bit in the mask represents a physical port. Default is 0xffffffff (all on);</p> <p>--cpu_affinity <value> – enables CPU affinity between I/O request and response. Possible values: 1 (enabled), 0 (disabled). Default is 0;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>--im --initiator_mode <value> – enables initiator mode function. Possible values: 1 (on), 2 (off). Default is 1;</p> <p>--reg_slg_pcmt <value> – percentage of total exchanges with registered SGLs. Default is 30;</p> <p>--quick_init <value> – enables quick link initialization. Possible values: 1 (on), 2 (off). Default is 0.</p> <p>--num_w_q_pgs <value> – number of pages per work queue. Default value is 1, which corresponds to 64 entries;</p> <p>--port_retry_delay <value> – amount of time to wait for a port to log back in after being lost, in ms. Default is 0, which allows the NVRAM Port Down Timeout value to be used;</p> <p>--eq_delay <value> – maximum interrupts per second delay. Default value: -1;</p> <p>--num_rpi <value> – number of RPIs. Default value is 2048;</p> <p>--use_transport_layer <value> – attaches to the SCSI transport layer. Possible values: 1 (on), 2 (off). Default is 1;</p> <p>-num_r_q_2k <value> – number of 2k receive buffers. Default is 32;</p> <p>--num_r_q_pgs <value> – number of pages per receive queue. Default value is 1, which corresponds to 512 entries;</p> <p>--atto_max_sectors <value> – maximum number of disk sectors in a single data transfer. Default value is 65535 (largest possible setting);</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|-------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>--sg_tablesize <value> – maximum number of entries in a scatter/gather table. Default value is 255, maximum is 256;</p> <p>--num_scb <value> – number of Sequence Control Blocks (SCBs), available for Target Mode requests. This setting is only valid when Target Mode is enabled. Default is 1024;</p> <p>--num_m_q_pgs <value> - number of pages per mailbox queue. Default is 1, which corresponds to 16 entries.</p> <p>--cmd_per_lun <value> - maximum number of commands per LUN. Default value is 16;</p> <p>--t10_dif <value> - enable T10-DIF functionality. Possible values: 1 (enabled), 0 (disabled). Default is 0;</p> <p>--cmd_retry_count <value> - maximum number of retries allowed for a command. Default value is 20.</p> <p>--target_mode_mask <value> - enables target mode on a per-port basis. Each bit in the mask represents a physical port. Default is 0x00000000 (all off).</p> |
| host | create | rdcli host create | <p>Command assigns the alias name to one or several hosts.</p> <p>Mandatory attributes:</p> <p>-n --name <value> – alias name;</p> <p>-l --initiators <value> – comma-separated list of hosts, which will be merged into one alias name.</p> <p>Optional attributes:</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-rt --realtime <value> – assigns priority to the created alias name of initiator/initiators;</p> <p>-qs --qosmic <value> – enable or disable QoSmic.</p> |
| | delete | rdcli host delete | <p>Command deletes the specified host alias.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – host alias.</p> |
| | modify | rdcli host modify | <p>Command allows you to modify the list of hosts, merged into the alias with the specified name.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – host alias.</p> <p>Optional attributes:</p> <p>-l --initiators <value>– list of comma-separated IQN/WWN/GUID of hosts;</p> <p>-r --newname <value> – new name of host alias;</p> <p>-rt --realtime <value> – assigns a priority (Realtime) to hosts, merged into the specified host alias;</p> <p>-qs --qosmic <value> – enables QoSmic.</p> |
| | show | rdcli host show | <p>Command shows all host aliases in the system.</p> <p>Optional attribute:</p> <p>-n --name <value> – displays the information on the specified host alias.</p> |

| Object | Method | Command Syntax | Functionality |
|--------|---------------|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ib | target show | <code>rdcli ib target show</code> | Command shows list of targets used in the system. |
| | target modify | <code>rdcli ib target modify</code> | <p>Command allows associating the appropriate InfiniBand targets on the local and remote nodes.</p> <p>Mandatory attributes:</p> <ul style="list-style-type: none"> -n --target_name <value> – local target name; -r --remote_target <value> – remote target name. <p>Optional attribute:</p> <ul style="list-style-type: none"> -f --force – forces command execution. |
| | opensm start | <code>rdcli ib opensm start</code> | Command starts OpenSM service. |
| | opensm stop | <code>rdcli ib opensm stop</code> | Command stops OpenSM service. |
| | opensm add | <code>rdcli ib opensm add</code> | <p>Command allows adding port GUID to start OpenSM.</p> <p>Mandatory attribute:</p> <ul style="list-style-type: none"> -g --guid <value> – specifies port GUID for OpenSM service. <p>Optional attribute:</p> <ul style="list-style-type: none"> -f --force – forces command execution and restarts OpenSM service if it is already launched. |
| | opensm del | <code>rdcli ib opensm del</code> | <p>Command deletes port GUID, for which OpenSM service will not be launched.</p> <p>Mandatory attribute:</p> |

| Object | Method | Command Syntax | Functionality |
|--------|---------------|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-g --guid <value> – port GUID.</p> <p>Optional attribute:</p> <p>-f --force – forces command execution and restarts OpenSM service if it is already launched.</p> |
| | opensm show | rdcli ib opensm show | <p>Command displays OpenSM service state. Possible values: 1 (service is started); 0 (service is stopped). In case service is started, the command also shows list of ports GUID, used for OpenSM service operation.</p> |
| iscsi | modify | rdcli iscsi modify | <p>Command allows to manage iSCSI support.</p> <p>Optional attributes:</p> <p>-e --enable <value> – enables or disables iSCSI. Possible values: 1 (enabled), 0 (disabled);</p> <p>-p --port <value> – specifies the port for iSCSI connection on the local node;</p> <p>-a --address – IP address of the interface to detect the iSCSI target on the local node;</p> <p>-f --force – forces command execution.</p> |
| | show | rdcli iscsi show | <p>Command shows iSCSI state.</p> |
| | target create | rdcli iscsi target create | <p>Command creates iSCSI target with specified name and authentication type.</p> |

| Object | Method | Command Syntax | Functionality |
|--------|---------------|---------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>Mandatory attributes:</p> <ul style="list-style-type: none"> -n --name <value> – iSCSI tagret name; -m --mode <value> – authentication type. Possible values: bidir (bidirectional); unidir (unidirectional); none (no authentication). <p>Optional attributes:</p> <ul style="list-style-type: none"> -l --login <value> – login for bidirectional authentication; -p --password <value> – password for bidirectional authentication. |
| | target delete | rdcli iscsi target delete | <p>Command deletes the specified iSCSI target.</p> <p>Mandatory attribute:</p> <ul style="list-style-type: none"> -n --name <value> – iSCSI target name. |
| | target modify | rdcli iscsi target modify | <p>Command allows changing parameters of the specified iSCSI target.</p> <p>Mandatory attribute:</p> <ul style="list-style-type: none"> -n --name <value> – iSCSI target name. <p>Optional attributes:</p> <ul style="list-style-type: none"> -m --mode <value> – authentication type; -l --login <value> – login for bidirectional authentication; -p --password <value> – password for bidirectional authentication; -hd --HeaderDigest <value> – HeaderDigest; |

| Object | Method | Command Syntax | Functionality |
|--------|-------------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-dd --DataDigest <value> – DataDigest;</p> <p>-rt --InitialR2T <value> – InitialR2T;</p> <p>-id --ImmediateData <value> – ImmediateData;</p> <p>-rd --MaxRecvDataSegmentLength <value> – MaxRecvDataSegmentLength;</p> <p>-xd --MaxXmitDataSegmentLength <value> – MaxXmitDataSegmentLength;</p> <p>-mb --MaxBurstLength <value> – MaxBurstLength;</p> <p>-fb --FirstBurstLength <value> – FirstBurstLength;</p> <p>-mr --MaxOutstandingR2T <value> – MaxOutstandingR2T.</p> |
| | target show | rdcli iscsi target show | <p>Command displays all created iSCSI targets.</p> <p>Optional attribute:</p> <p>-n --name <value> – displays the information on the specified iSCSI target.</p> |
| | chap create | rdcli iscsi chap create | <p>Command creates CHAP user with the specified login and password.</p> <p>Mandatory attributes:</p> <p>-l --login <value> – CHAP login;</p> <p>-p --password <value> – CHAP password.</p> |

| Object | Method | Command Syntax | Functionality |
|--------|---------------|----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | chap delete | <code>rdcli iscsi chap delete</code> | <p>Command deletes the specified CHAP user.</p> <p>Mandatory attribute:</p> <p>-l --login <value> – CHAP login.</p> |
| | chap modify | <code>rdcli iscsi chap modify</code> | <p>Command allows changing the CHAP user login and password.</p> <p>Mandatory attributes:</p> <p>-l --login <value> – CHAP login;</p> <p>-p --password <value> – CHAP password.</p> |
| | chap show | <code>rdcli iscsi chap show</code> | <p>Command shows all the created CHAP users.</p> <p>Optional attribute:</p> <p>-l --login <value> – shows information on the specified CHAP user.</p> |
| | iprule create | <code>rdcli iscsi iprule create</code> | <p>Command creates an allow rule for the specified target with the specified hostname and IP address.</p> <p>Mandatory attributes:</p> <p>-t --target <value> – target name;</p> <p>-a --ip <value> – specifies host IP address.</p> |
| | iprule delete | <code>rdcli iscsi iprule delete</code> | <p>Command deletes allow rule for the specified target with the specified hostname and IP.</p> <p>Mandatory attributes:</p> <p>-t --target <value> – target name;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|-------------|--------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | iprule show | <code>rdcli iscsi iprule show</code> | <p><code>-a --ip <value></code> – specifies host IP address.</p> <p>Command displays the list of all IP rules.</p> <p>Optional attribute:</p> <p><code>-tn --target <value></code> – displays information on the IP rules for the specified target.</p> |
| lun | create | <code>rdcli lun create</code> | <p>Command creates a LUN with the specified name on the specified RAID.</p> <p>Mandatory attributes:</p> <p><code>-n --name <value></code> - LUN name;</p> <p><code>-r --raid <value></code> - name of RAID, on which the LUN is created.</p> <p>Optional attributes:</p> <p><code>-s --size <value></code> – LUN size, GiB. Default value is 32;</p> <p><code>-o --offset <value></code> – offset value. Default value is 128;</p> <p><code>-b --blocksize <value></code> – block size, B. Possible values: 512, 4196. Default value is 512;</p> <p><code>-no --nas_only <value></code> – create LUN only for NAS;</p> <p><code>-ra --ra</code> –enables RAO;</p> <p><code>-am --alb_mem</code> –specifies the amount of RAO server memory in GiB; the default size is 1. Possible values: special decimal values: 0.25, 0.5, 0.75; or any positive integer;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-as --alb_size – specifies the RAO server index size, in GiB. If not specified, a default is calculated based on the memory allocated to the RAO server based on the alb_mem option;</p> <p>-aspl --alb_sparse – enables Random Access sparse indexing.</p> <p>The sparse indexing guarantees the largest amount of data stored in the index memory. Sparse indexing allows you to store the amount of data that 10 times exceeds the amount of data stored with conventional deduplication. Enabled sparse indexing requires a lot of hardware resources. Possible values: 1 (On) – enabled; 0 (Off) – disabled. Default is 1;</p> <p>-rabril --ra_bio_rotation_interval - specifies the number of I/O operations to enqueue for each block I/ O submission thread before directing work to the next thread;</p> <p>-rass --ra_slab_size – specifies the size of the increment by which a Random Access Optimization volume can grow in GiB. Using a smaller size than the default maximum constrains the total maximum physical volume size. Possible values: powers of two between 1 and 32 GiB;</p> <p>-rabmcs --ra_block_map_cache_size – specifies the amount of memory allocated for cached block map pages, in MiB. The value</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>must be a multiple of 4 KiB. The default is 128, which is the required minimum;</p> <p>-rabmp --ra_block_map_period – determines the number of block map updates which may accumulate before cache pages are flushed to disk. Possible values: from 1 to 16380. Default value is 16380;</p> <p>rae512 --ra_enable512e – enables 512-byte block device emulation mode for Random Access Optimization;</p> <p>-raat --ra_ack_threads – specifies the number of threads to use for acknowledging completion of requested Random Access Optimization I/O operations. Possible values: 1, 2;</p> <p>-rabt --ra_bio_threads – specifies the number of threads to use for submitting I/O operations to the storage device. Possible values: 1, 2. The second thread will use additional 18 MB of RAM;</p> <p>-ract --ra_cpu_threads <value> – specifies the number of threads to be used for costly operations such as hashing. Possible values: 1, 2;</p> <p>-rals --ra_logical_size – the logical size of the LUN. By default, is equal to the difference between the physical size of the volume and the size of the index. If the size of the logical volume exceeds</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>the size of the physical volume, the system will work in the thin provisioning mode;</p> <p>-ralt --ra_logical_threads – specifies the number of threads across which subdivide parts of the Random Access Optimization processing based on logical block addresses;</p> <p>-rapt --ra_physical_threads – specifies the number of threads across which parts of the Random Access Optimization processing are subdivided, based on physical block addresses. Possible values: 1, 2;</p> <p>-rarcs --ra_read_cache_size – specifies the extra Random Access Optimization volume read cache size in MiB. This space is in addition to a system-defined minimum. Default value is 0;</p> <p>-rawpl --ra_write_policy – specifies the write policy: synchronous (writes are acknowledged only after data is stably written) or asynchronous (writes are acknowledged after data has been cached for writing to stable storage). Possible values: sync, async.</p> <p>-rad --ra_deduplication – enables deduplication for RAO.</p> |
| | delete | rdcli lun delete | Command deletes the specified LUN. |

| Object | Method | Command Syntax | Functionality |
|--------|--------|------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>Mandatory attribute:</p> <p>-n --name <value> - LUN name.</p> <p>Optional attribute:</p> <p>-f --force – forces command execution.</p> |
| | modify | rdcli lun modify | <p>Command allows changing parameters of the specified LUN.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – LUN name.</p> <p>Optional attributes:</p> <p>-fs --fs_type <value> – format the LUN with the specified file system. Possible values: xfs, ext4;</p> <p>-f --force – forces command execution;</p> <p>-raat --ra_ack_threads – specifies the number of threads to use for acknowledging completion of requested Random Access Optimization I/O operations. Possible values: 1, 2;</p> <p>-rabt --ra_bio_threads – specifies the number of threads to use for submitting I/O operations to the storage device. Possible values: 1, 2. The second thread will use additional 18 MB of RAM;</p> <p>-ract --ra_cpu_threads <value> – specifies the number of threads to be used for costly operations such as hashing. Possible values: 1, 2;</p> <p>-ralt --ra_logical_threads – specifies the number of threads across</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>which subdivide parts of the Random Access Optimization processing based on logical block addresses;</p> <p>-rapt --ra_physical_threads – specifies the number of threads across which parts of the Random Access Optimization processing are subdivided, based on physical block addresses. Possible values: 1, 2;</p> <p>-rarcs --ra_read_cache_size – specifies the extra Random Access Optimization volume read cache size in MiB. This space is in addition to a system-defined minimum. Default value is 0;</p> <p>-rawp --ra_write_policy – specifies the write policy: synchronous (writes are acknowledged only after data is stably written) or asynchronous (writes are acknowledged after data has been cached for writing to stable storage). Possible values: sync, async;</p> <p>-rad --ra_deduplication – enables deduplication for RAO;</p> <p>-rabmcs --ra_block_map_cache_size – specifies the amount of memory allocated for cached block map pages, in MiB. The value must be a multiple of 4 KiB. The default is 128, which is the required minimum;</p> <p>-rabmp --ra_block_map_period – determines the number of block map updates which may accumulate before cache pages are</p> |

| Object | Method | Command Syntax | Functionality |
|--------------------|---------|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | flushed to disk. Possible values: from 1 to 16380 . Default value is 16380 ; |
| | show | <code>rdcli lun show</code> | <p>Command shows all the created LUNs.</p> <p>Optional attributes:</p> <p>-n --name <value> – displays the information on the specified LUN;</p> <p>-f --free – displays all the free LUNs.</p> |
| | restore | <code>rdcli lun restore</code> | <p>Command restores the LUN.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – LUN name.</p> |
| mask target | create | <code>rdcli mask target create</code> | <p>Command creates a masking rule for LUN and target.</p> <p>Mandatory attributes:</p> <p>-l --lunname <value> – specifies the name of LUN, for which masking rule is created;</p> <p>-t --targetname <value> - specifies the target, for which masking rule is created.</p> |
| | delete | <code>rdcli mask target delete</code> | <p>Command deletes masking rule for targets.</p> <p>Optional attributes:</p> <p>-l --lunname <value> – specifies the name of LUN, for which masking rule is created;</p> |

| Object | Method | Command Syntax | Functionality |
|-----------|--------|------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-tn --targetname <value> - specifies the target, for which masking rule is created;</p> <p>-il--id <value> - masking rule ID.</p> |
| | show | rdcli mask target show | <p>Command shows the list of all masking rules for targets.</p> <p>Optional attributes:</p> <p>-l --lunname <value> – displays all masking rules for the specified LUN;</p> <p>-tn --targetname <value> – displays all masking rules for the specified target;</p> <p>-il--id <value> – displays masking rule with the specified ID.</p> |
| mask host | create | rdcli mask host create | <p>Command creates a masking rule for host.</p> <p>Mandatory attributes:</p> <p>-l --lunname <value> – specifies the name of LUN, for which masking rule is created;</p> <p>-hn --hostname <value> – host name;</p> <p>-p --permissions <value> – defines access permissions. Possible values: r (Read Only), w (Read/Write).</p> |
| | modify | rdcli mask host modify | <p>Commands allows changing parameters of masking rule for host.</p> <p>Mandatory attributes:</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-l --lunname <value> – specifies the name of LUN, for which masking rule is created;</p> <p>-h --hostname <value> – host name;</p> <p>-p --permissions <value> – defines access permissions. Possible values: r (Read Only), w (Read/Write).</p> |
| | delete | <code>rdcli mask host delete</code> | <p>Command deletes masking rule for host.</p> <p>Optional attributes:</p> <p>-l --lunname <value> – deletes masking rule for host for the specified LUN;</p> <p>-h --hostname <value> – deletes masking rule for the specified host;</p> <p>-i --id <value> – deletes the specified masking rule.</p> |
| | show | <code>rdcli mask host show</code> | <p>Command shows the list of all masking rules for hosts.</p> <p>Optional attributes:</p> <p>-l --lunname <value> – displays masking rules for host for the specified LUN;</p> <p>-h --hostname <value> – displays masking rules for the specified host;</p> <p>-i --id <value> – displays the specified masking rule.</p> |

| Object | Method | Command Syntax | Functionality |
|--------------------|--------|-------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| metadata raid | scan | <code>rdcli metadata raid scan</code> | Command searches for RAID on all free drives. Optional attribute <code>-d --drives <value></code> – specifies drives for scan. |
| | rename | <code>rdcli metadata raid rename</code> | Command renames imported RAID or volumes in case in case of name conflict. Optional attributes: <code>-d --drives <value></code> – specifies volumes for renaming; <code>-on --old_name <value></code> – RAID old name in drives metadata; <code>-n --name<value></code> – assigns the specified new name to the RAID or volume. |
| | import | <code>rdcli metadata raid import</code> | Command imports RAID into the current configuration. Mandatory attribute: <code>-r --raid <value></code> – RAID name. Optional attribute: <code>-cs --cache_size <value></code> – specifies cache size for the RAID. |
| metadata volume | scan | <code>rdcli metadata volume scan</code> | Command shows list of volumes for the specified RAID. Mandatory attribute: <code>-r --raid <value></code> – RAID name. Optional attribute: <code>-cs --cache_size <value></code> – cache size for the RAID. |
| | rename | <code>rdcli metadata volume rename</code> | Command renames volumes of the specified RAID. Mandatory attributes: |

| Object | Method | Command Syntax | Functionality |
|----------|--------|-----------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-r --raid <value> – name of RAID, which includes the volume to rename;</p> <p>-n --name <value> – new name of the volume.</p> <p>Optional attribute:</p> <p>-cs --cache_size <value> – cache size for the RAID.</p> |
| nas user | create | rdcli nas user create | <p>Command creates a NAS user with specified name and password.</p> <p>Mandatory attributes:</p> <p>-n --name <value> – username;</p> <p>-p --password <value> – user password.</p> <p>Optional attributes:</p> <p>-g --groups <value> – groups for the user;</p> <p>-pg --primary_group <value> – primary group for the user.</p> |
| | modify | rdcli nas user modify | <p>Command allows changing NAS user parameters.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – username.</p> <p>Optional attributes:</p> <p>-g --groups <value> – list of supplementary groups in which user is included;</p> <p>-p --password <value> – changes user password;</p> <p>-pg --primary_group <value> – changes primary group for the user.</p> |

| Object | Method | Command Syntax | Functionality |
|-----------|--------|-------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | delete | <code>rdcli nas user delete</code> | <p>Command deletes the specified NAS user.</p> <p>Mandatory attribute:</p> <p><code>-n --name <value></code> – username.</p> |
| | show | <code>rdcli nas user show</code> | <p>Command displays list of all NAS users.</p> <p>Optional attributes:</p> <p><code>-n --name <value></code> – displays information on the specified NAS user;</p> <p><code>-si --show_id</code> - displays NAS users UIDs;</p> <p><code>-a --ads</code> - displays list of Active Directory users when SMB share via Active Directory connection is established;</p> <p><code>--all</code> – displays list of all users, including Active Directory users.</p> |
| nas group | create | <code>rdcli nas group create</code> | <p>Command creates a NAS user group.</p> <p>Mandatory attribute:</p> <p><code>-n --name <value></code> – user group name.</p> |
| | show | <code>rdcli nas group show</code> | <p>The command displays list of NAS user groups.</p> <p>Optional attributes:</p> <p><code>-n --name <value></code> – displays information about NAS user group with specified name;</p> <p><code>-si --show_id</code> – displays user groups IDs;</p> |

| Object | Method | Command Syntax | Functionality |
|-----------|------------|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-a --ads – displays list of Active Directory user groups when SMB share via Active Directory connection is established;</p> <p>--all – displays list of all users, including Active Directory users.</p> |
| | delete | rdcli nas group delete | <p>Command deletes NAS user group.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – user group name.</p> <p>Optional attribute:</p> <p>-a --ads – removes Active Directory user groups.</p> |
| nas share | show | rdcli nas share show | <p>Command displays list of all created shares, regardless of their type.</p> <p>Optional attribute:</p> <p>-l --lun <value> – displays list of all shares, created on the specified LUN.</p> |
| | nfs create | rdcli nas share nfs create | <p>Command creates a NFS share on the specified LUN.</p> <p>Mandatory attributes:</p> <p>-n --name <value> – NFS share name;</p> <p>-l --lun <value> – LUN name.</p> <p>Optional attributes:</p> <p>-p --path <value> – path to NAS share on LUN;</p> <p>-s --squash <value> – enables or disables Squash option;</p> |

| Object | Method | Command Syntax | Functionality |
|------------|--------|-----------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-rs --root_squash <value> – enables or disables squash option for a root user;</p> <p>-rh --ro_hosts <value> – list of hosts with Read Only permissions;</p> <p>-wh --rw_hosts <value> – list of hosts with Read/Write permissions.</p> |
| nfs modify | | <code>rdcli nas share nfs modify</code> | <p>Command allows changing NFS share parameters.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – NFS share name.</p> <p>Optional attributes:</p> <p>-s --squash <value> – enables or disables Squash option;</p> <p>-rs --root_squash <value> – enables or disables squash option for a root user;</p> <p>-rh --ro_hosts <value> – list of hosts with Read Only permissions;</p> <p>-wh --rw_hosts <value> – list of hosts with Read/Write permissions.</p> |
| nfs delete | | <code>rdcli nas share nfs delete</code> | <p>Command deletes NFS share.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – NFS share name.</p> <p>Optional attribute:</p> |

| Object | Method | Command Syntax | Functionality |
|--------|------------|----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-f --force – forces command execution and restarts NFS service. In case this is the last share on LUN, this LUN should be unmounted, which requires restarting the NFS service.</p> |
| | nfs show | rdcli nas share nfs show | <p>Command displays list of all NFS shares.</p> <p>Optional attributes:</p> <p>-n --name <value> – displays information about the specified NFS share;</p> <p>-l --lun <value> – displays information about NFS shares, which are created on the specified LUN.</p> |
| | afp create | rdcli nas share afp create | <p>Command creates an AFP share on the specified LUN.</p> <p>Mandatory attributes:</p> <p>-n --name <value> – AFP share name;</p> <p>-l --lun <value> – LUN name.</p> <p>Optional attributes:</p> <p>-p --path <value> – path to AFP share on LUN;</p> <p>-g --guest <value> – guest access type. Possible values: ro (Read Only), rw (Read/Write), off (guest access is disabled);</p> <p>-ru --ro_users <value> – list of users with Read Only permissions;</p> <p>-wu --rw_users <value> – list of users with Read/Write permissions;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|------------|----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-rgl--ro_groups <value> – list of groups of users with Read Only permissions;</p> <p>-wgl--rw_groups <value> – list of groups of users with Read/Write permissions;</p> <p>-ah --allow_hosts <value> – list of hosts, which have access to the share;</p> <p>-dh --deny_hosts <value> – list of hosts, for which access to the share is denied.</p> |
| | afp modify | rdcli nas share afp modify | <p>Command allows changing AFP share parameters.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – AFP share name.</p> <p>Optional attributes:</p> <p>-gl--guest <value> – guest access type. Possible values: ro (Read Only); rw (Read/Write); off (guest access is disabled);</p> <p>-ru --ro_users <value> – list of users with Read Only permissions;</p> <p>-wu --rw_users <value> – list of users with Read/Write permissions;</p> <p>-rgl--ro_groups <value> – list of groups of users with Read Only permissions;</p> <p>-wgl--rw_groups <value> – list of groups of users with Read/Write permissions;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|------------|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-ah --allow_hosts <value> – list of hosts, which have access to the share;</p> <p>-dh --deny_hosts <value> – list of hosts, for which access to the share is denied.</p> |
| | afp delete | rdcli nas share afp delete | <p>Command deletes AFP share.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – AFP share name.</p> <p>Optional attribute:</p> <p>-f --force – forces command execution and restarts AFP service.</p> <p>In case this is the last share on LUN, this LUN should be unmounted, which requires restarting the AFP service.</p> |
| | afp show | rdcli nas share afp show | <p>Command displays list of all AFP shares.</p> <p>Optional attributes:</p> <p>-n --name <value> – displays the information about specified AFP share;</p> <p>-l --lun <value> – displays information about AFP shares, which are created on the specified LUN.</p> |
| | ftp create | rdcli nas share ftp create | <p>Command creates an FTP share on the specified LUN.</p> <p>Mandatory attributes:</p> <p>-n --name <value> – FTP share name;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|------------|----------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-l --lun <value> – LUN name.</p> <p>Optional attributes:</p> <p>-p --path <value> – path to the share on LUN;</p> <p>-g --guest <value> – guest access type: ro (Read Only); rw (Read/Write); off (guest access is disabled);</p> <p>-ru --ro_users <value> – list of groups of users with Read Only permissions;</p> <p>-wu --rw_users <value> – list of groups of users with Read/Write permissions;</p> <p>-rg --ro_groups <value> – list of groups of users with Read Only permissions;</p> <p>-wg --rw_groups <value> – list of groups of users with Read/Write permissions.</p> |
| | ftp modify | rdcli nas share ftp modify | <p>Command allows changing FTP share parameters.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – FTP share name;</p> <p>Optional attributes:</p> <p>-g --guest <value> – guest access type: ro (Read Only); rw (Read/Write); off (guest access is disabled);</p> |

| Object | Method | Command Syntax | Functionality |
|------------|--------|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-ru --ro_users <value> – list of groups of users with Read Only permissions;</p> <p>-wu --rw_users <value> – list of groups of users with Read/Write permissions;</p> <p>-rg --ro_groups <value> – list of groups of users with Read Only permissions;</p> <p>-wg --rw_groups <value> – list of groups of users with Read/Write permissions.</p> |
| ftp delete | | rdcli nas share ftp delete | <p>Command deletes FTP share.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – AFP share name.</p> <p>Optional attribute:</p> <p>-f --force – forces command execution and restarts FTP service.</p> <p>In case this is the last share on LUN, this LUN should be unmounted, which requires restarting the FTP service.</p> |
| ftp show | | rdcli nas share ftp show | <p>Command displays list of all FTP shares.</p> <p>Optional attributes:</p> <p>-n --name <value> – displays the information about specified FTP share;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------------|------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | -l --lun <value> – displays information about FTP shares, which are created on the specified LUN. |
| | samba create | rdcli nas share samba create | <p>Command creates an SMB share on the specified LUN.</p> <p>Обязательные параметры:</p> <p>Mandatory attributes:</p> <p>-n --name <value> – SMB share name;</p> <p>-l --lun <value> – LUN name.</p> <p>When adding attributes, you can specify the following parameters:</p> <p>-p --path <value> – path to the share on LUN;</p> <p>-g --guest <value> – guest access type: ro (Read Only); rw (Read/Write); off (guest access is disabled);</p> <p>-ru --ro_users <value> – list of groups of users with Read Only permissions;</p> <p>-wu --rw_users <value> – list of groups of users with Read/Write permissions;</p> <p>-rg --ro_groups <value> – list of groups of users with Read Only permissions;</p> <p>-wg --rw_groups <value> – list of groups of users with Read/Write permissions;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-ah --allow_hosts <value> – list of hosts, which have access to the share;</p> <p>-dh --deny_hosts <value> – list of hosts, for which access to the share is denied;</p> <p>-a --admins <value> – list of root directory administrators;</p> <p>-b --browseable <value> – visibility of SMB share in the list of available network resources. Possible values: 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> <p>-worm --worm <value> –enables or disables WORM (Write Once Read Many) option. Possible values: 0 (disabled); 1 (enabled);</p> <p>-gpl --grace_period <value> – period of time, during which the share is available for modification/ deletion, in seconds. The default value is 1 (minimum possible value);</p> <p>-rpl --ro_period <value> - period of time, during which the share is not available for modification while WORM parameter is enabled, in months. The default value is 60 (5 years);</p> <p>-c --comment <value> – comment to the share.</p> |


| Object | Method | Command Syntax | Functionality |
|--------|--------------|-------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | samba modify | <code>rdcli nas share samba modify</code> | <p>Command allows changing SMB share parameters.</p> <p>Mandatory attributes:</p> <p><code>-n --name <value></code> – SMB share name.</p> <p>Optional attributes:</p> <p><code>-a --admins <value></code> – edits the list of root directory administrators;</p> <p><code>-ah --allow_hosts <value></code> – list of hosts, which have access to the share;</p> <p><code>-b --browseable <value></code> – visibility of SMB share in the list of available network resources. Possible values: 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> <p><code>-dh --deny_hosts <value></code> – list of hosts, for which access to the share is denied;</p> <p><code>-gl--guest <value></code> – guest access type: ro (Read Only); rw (Read/Write); off (guest access is disabled);</p> <p><code>-gpl--grace_period <value></code> – period of time, during which the share is available for modification/ deletion, in seconds. The default value is 1 (minimum possible value);</p> |

| Object | Method | Command Syntax | Functionality |
|--------------|--------|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-ru --ro_users <value> – list of groups of users with Read Only permissions;</p> <p>-rg --ro_groups <value> – list of groups of users with Read Only permissions;</p> <p>-rp --ro_period <value> - period of time, during which the share is not available for modification while WORM parameter is enabled, in months. The default value is 60 (5 years);</p> <p>-wg --rw_groups <value> – list of groups of users with Read/Write permissions;</p> <p>-wu --rw_users <value> – list of groups of users with Read/Write permissions;</p> <p>-worm --worm <value> –enables or disables WORM (Write Once Read Many) option. Possible values: 0 (disabled); 1 (enabled);</p> <p>-c --comment <value> – comment to the share.</p> |
| samba delete | | rdcli nas share samba delete | <p>Command deletes SMB share.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – SMB share name.</p> <p>Optional attribute:</p> |


| Object | Method | Command Syntax | Functionality |
|-----------|------------|----------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-f --force – forces command execution and restarts SMB service.</p> <p>In case this is the last share on LUN, this LUN should be unmounted, which requires restarting the SMB service.</p> |
| | samba show | rdcli nas share samba show | <p>Command returns a list of all SMB shares.</p> <p>Optional attributes:</p> <p>-n --name <value> – displays the information about specified SMB share;</p> <p>-l --lun <value> – displays information about SBM shares, which are created on the specified LUN.</p> |
| nas samba | show | rdcli nas samba show | <p>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.</p> |
| | modify | rdcli nas samba modify | <p>Command allows connecting (disconnecting) Samba to (from) Active Directory and changing settings of Microsoft Windows workgroup.</p> <p>Optional attributes:</p> <p>-a --ads <value> – changes Active Directory connection state.</p> <p>Possible values: 1 – try to connect to Active Directory; 0 – connection will not be disabled, but winbind service will be</p> |


| Object | Method | Command Syntax | Functionality |
|-----------|--------|------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>stopped and information about connection will be removed from samba and kerberos configuration files;</p> <p>-w --workgroup <value> – name of Windows workgroup (if ads = 0) or Active Directory domain (if ads = 1);</p> <p>-r --realm <value> – domain controller name. If ads = 1, this attribute is mandatory; if ads = 0, this attribute is not used;</p> <p>-u --user <value> – name of user with permission to connect the work station to the domain. If ads = 1, this attribute is mandatory; if ads = 0, this attribute is not used;</p> <p>-p --password <value> – user password. If ads = 1, this attribute is mandatory; if ads = 0, this attribute is not used;</p> <p>-t --trusts <value> – sets trusts between Active Directory domains.</p> |
| nas quota | show | rdcli nas quota show | <p>Command displays information about the created quotas.</p> <p>Optional attributes:</p> <p>-u --user <value> – displays quotas set for the specified user;</p> <p>-g --group <value> – displays quotas set for the specified group;</p> <p>-l --lun <value> – displays quotas set for the specified LUN.</p> |
| | create | rdcli nas quota create | <p>Command allows creating quota of the on the specified LUN.</p> <p>Mandatory attribute:</p> <p>-l --lun <value> – LUN name.</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>Optional attributes:</p> <ul style="list-style-type: none"> -s --size <value> – quota size, in MiB; -u --users <value> – users, for which quota is set; -g --groups <value> – groups, for which quota is set; -gpl--grace_period <value> – sets a grace period (in minutes), after which the share is unavailable for modifications. |
| | modify | <code>rdcli nas quota modify</code> | <p>Command allows editing quota parameters.</p> <p>Mandatory attribute:</p> <ul style="list-style-type: none"> -l --lun <value> – LUN name. <p>Optional attributes:</p> <ul style="list-style-type: none"> -s --size <value> – quota size, in MiB; -u --users <value> – users, for which the quota is set; -g --groups <value> – groups, for which the quota is set; -gpl--grace_period <value> – sets a grace period (in minutes), after which the share is unavailable for modifications. |
| | delete | <code>rdcli nas quota delete</code> | <p>Command deletes quotas on the specified LUN.</p> <p>Mandatory attribute:</p> <ul style="list-style-type: none"> -l --lun <value> – LUN name. <p>Optional attributes:</p> <ul style="list-style-type: none"> -u --users <value> – deletes quotas, set for the specified users; |

| Object | Method | Command Syntax | Functionality |
|---------|----------|----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | -gl --groups <value> – deletes quotas, set for the specified groups. |
| network | show | rdcli network show | Command displays general networking settings: dns, gateway, primary interface. |
| | modify | rdcli network modify | <p>Command allows changing general network parameters.</p> <p>Optional attributes:</p> <p>-hn --hostname <value> – host name;</p> <p>-if --primary_interface <value> – primary interface;</p> <p>-dns1 --dns1 <value> – primary DNS;</p> <p>-dns2 --dns2 <value> – secondary DNS;</p> <p>-dns3 --dns3 <value> – tertiary DNS;</p> <p>-gw --gateway <value> – gateway.</p> <p> If the main interface has been assigned IP-address via DHCP, it is prohibited to change DNS and Gateway parameters when performing this command.</p> |
| nodes | discover | rdcli nodes discover | <p>Command initiates system scanning to discover new RAIDIX systems in the local network.</p> <p>Optional attributes:</p> <p>-if --interface <value> - interface name in the system;</p> <p>-tl --timeout <value> – waiting time for response;</p> |

| Object | Method | Command Syntax | Functionality |
|-------------------|--------|--------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | -ba --broadcast_address <value> – broadcast address. |
| | show | rdcli nodes show | Command displays information on RAIDIX systems in the local network. |
| network interface | show | rdcli network interface show | Displays the network interface parameters. Optional attributes: -n --name <value> – interface name in the system; -b --bonding <value> – bonding name. |
| | modify | rdcli network interface modify | Command allows modifying network interface parameters. Mandatory attribute: -n --name <value> – interface name. Optional attributes: -a --ipaddr <value> – interface IP address; -d --dhcp <value> – specifies whether the port address is defined by DHCP. Possible values: 1 (DHCP is used); 0 (DHCP is not used); -m --netmask <value> – net mask. -M --mtu <value> – MTU (maximum transmission unit) value; -u --up <value> – port status. Possible values: 1 (port is enabled), 0 (port is disabled). |

| Object | Method | Command Syntax | Functionality |
|--------------|--------|----------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| network bond | create | <code>rdcli network bond create</code> | <p>Command allows combining specified interfaces, defining bonding policy type.</p> <p>Mandatory attributes:</p> <p>-p --policy <value> – bonding mode. Possible values: 0, 1, 2, 3, 4, 5, 6;</p> <p>-s --slaves <value> – interface list (comma-separated with no space between).</p> <p> Primary interface cannot participate in bonding</p> |
| | delete | <code>rdcli network bond delete</code> | <p>Command deletes the specified channel bonding.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – bonding name.</p> |
| | show | <code>rdcli network bond show</code> | <p>Command displays all configured channel bondings.</p> <p>Optional attributes:</p> <p>-n --name <value> – displays info about the specified bonding.</p> |
| | modify | <code>rdcli network bond modify</code> | <p>Command allows modifying the specified channel bonding parameters:</p> <p>Mandatory attributes:</p> <p>-n --name <value> – bonding name.</p> <p>Optional attributes:</p> |

| Object | Method | Command Syntax | Functionality |
|-------------|--------|--------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-p --policy <value> – bonding mode. Possible values: 0, 1, 2, 3, 4, 5, 6;</p> <p>-s --slaves <value> – interface list (comma-separated with no space between).</p> |
| network vip | show | rdcli network vip show | Command displays list of all virtual IP addresses created in the system. |
| | create | rdcli network vip create | <p>Command allows assigning virtual IP-address to the specified interface.</p> <p>Mandatory attributes:</p> <p>-i --interface <value> – network interface;</p> <p>-a --ipaddr <value> – IP address.</p> <p>Optional attributes:</p> <p>-n --netmask <value> – netmask;</p> <p>-s -- spare <value> – specifies the node, on which virtual IP will be Active in regular system mode. Possible values: 0 (Active) – this node is primary for this IP (the option is specified by default); 1 (Passive) – the local node is not primary for this VIP.</p> <p> Value 0 (Active) should be set on the node from which work with shares is performed.</p> |

| Object | Method | Command Syntax | Functionality |
|----------------------------|--------|-----------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | delete | <code>rdcli network vip delete</code> | <p>Command deletes the specified virtual IP address.</p> <p>Mandatory attribute:</p> <p><code>-na --name <value></code> – virtual IP address.</p> |
| | modify | <code>rdcli network vip modify</code> | <p>Command changes VIP settings.</p> <p>Mandatory attribute:</p> <p><code>-na --name <value></code> – virtual IP address.</p> <p>Optional attributes:</p> <p><code>-a --ipaddr <value></code> – interface IP address;</p> <p><code>-n --netmask <value></code> – netmask;</p> <p><code>-s --spare <value></code> – specify the primary node for IP address (the node, on which the virtual IP will be Active in regular system mode);</p> <p>Possible values: 0 (Active) – the local node is primary for this IP (sets by default after VIP creation); 1 (Passive) – the local node is not primary for this virtual IP.</p> |
| notify mail profile | modify | <code>rdcli notify mail profile modify</code> | <p>Command allows configuring server e-mail notification parameters.</p> <p><code>-a --auth <value></code> – authentication type. Possible values: on, off, plain, scram-sha-1, cram-md5, gssapi, digest-md5, login, ntlm;</p> |

| Object | Method | Command Syntax | Functionality |
|--------------------------|--------|---------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-e --encryption <value> – TLS encryption. Possible values: 0 – disabled, 1 – enabled;</p> <p>-p --password <value> – sender password;</p> <p>-pr --port <value> – port;</p> <p>-s --server <value> – IP address of server – sender of notifications;</p> <p>-sn --sender <value> – sender name;</p> <p>-u --user <value> – sender e-mail;</p> <p>-fm --format <value> – format of notification messages. Possible values: text, html.</p> |
| | show | rdcli notify mail profile show | Command displays server e-mail notification parameters. |
| notify mail recipient | create | rdcli notify mail recipient create | <p>Command creates e-mail notifications recipient.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – notification recipient e-mail.</p> <p>Optional attributes:</p> <p>-c --controller_sensors <value> – controller sensors notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-d --drive <value> – volumes state notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-e --enclosure_sensors <value> – enclosure sensors notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-s --iscsi <value> – iSCSI- adapters notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-l --license <value> – license notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-ln --lun <value> – LUN state notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-il --network <value> – network state notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-r --raid <value> – RAID state notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-pl--sparepool <value> –SparePool state notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-sn --system_notification <value> - system notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only.</p> |
| | modify | rdcli notify mail recipient modify | <p>Command allows changing e-mail notification parameters.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – notification recipient e-mail.</p> <p>Optional attributes:</p> <p>-sn --system_notification <value> – enables or disables e-mail notification sending. Possible values: 1 (enabled), 0 (disabled);</p> <p>-r --raid <value> – sets RAID state notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-ln --lun <value> – sets LUN state notification types. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-d --drive <value> – sets volumes state notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-i --network <value> – sets network state notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-l --license <value> – sets license state notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-p --sparepool <value> – sets SparePool state notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-c --controller_sensors <value> – sets controller sensors notifications type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-e --enclosure_sensors <value> – sets enclosure sensors notifications type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only;</p> <p>-s --iscsi <value> – sets iSCSI- adapters notification type. Possible values: ok – all notifications, warning – warnings and errors, error – errors only.</p> |

| Object | Method | Command Syntax | Functionality |
|--------------------|--------|-------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | delete | <code>rdcli notify mail recipient delete</code> | Command deletes e-mail notification recipient Mandatory attribute: -n --name <value> – notifications recipient e-mail. |
| | show | <code>rdcli notify mail recipient show</code> | Command displays list of all e-mail notification recipients. Optional attribute: -n --name <value> – notification recipient e-mail. |
| notify mail | test | <code>rdcli notify mail test</code> | Command sends test e-mail to the recipients. Optional attribute: -r --recipient <value> – notification recipient e-mail. |
| param raid | show | <code>rdcli param raid show</code> | Command displays list of RAID parameters. Optional attribute: -d --default – displays list of default RAID parameters. |
| | modify | <code>rdcli param raid modify</code> | Command allows changing RAID parameters. Optional attributes: -cs --cache_size <value> – cache size, in MiB. Default value: 4096 ; -ss --stripe_size <value> – sets stripe size, in MiB; -acc --alg_compat_code <value> – changes the checksum calculation algorithm; -d --default – applies default settings. |

| Object | Method | Command Syntax | Functionality |
|--------------|--------|----------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| param driver | show | <code>rdcli param driver show</code> | <p>Command displays list of driver parameter values.</p> <p>Optional attribute:</p> <p><code>-d --default</code> – displays default settings.</p> |
| | modify | <code>rdcli param driver modify</code> | <p>Command allows changing driver parameters.</p> <p>Optional attribute:</p> <p><code>-cc --command_saturation <value></code> - maximum number of non-real-time commands. Possible values: from 0 to (2³²-1). Default value: 4;</p> <p><code>-ft --flush_threshold <value></code> - minimum cache saturation percentage enabling cache flushing. Possible values: from 0 to 100. Default value: 25;</p> <p><code>-il --io_timeout <value></code> – disk I/O timeout, in ms. Possible values: from 0 to (2³²-1). Default value: 8000;</p> <p><code>-lj --log_level <value></code> – logging level. Possible values: from 0 to (2³²-1). Default value: 256;</p> <p><code>-pfl --predictive_failover <value></code> – enables or disables data reconstruction when most drives are complete. Possible values: 1 (enabled), 0 (disabled). Default value: 0;</p> <p><code>-rl --recon_min <value></code> – reconstruction priority. Possible values: from 0 to 100. Default value: 0;</p> |

| Object | Method | Command Syntax | Functionality |
|-----------------|--------|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-rt --request_sync_timeout <value> – request synchronization timeout, in ms. Possible values: from 0 to (2³²-1). Default value: 10000;</p> <p>-rs --request_saturation <value> - Maximum RAID data rate level for non-real-time requests, MB/s. Default value: 0;</p> <p>-rc --recon_commands <value> – Number of outstanding reconstruction requests. Possible values: from 0 to (2³²-1). Default value: 8;</p> <p>-s --cache_saturation <value> - Maximum cache saturation level for non-real-time requests, in percent. Possible values: from 0 to 100. Default value: 75;</p> <p>-ct --calc_thread_num <value> - Number of calculation threads. Possible values: from 0 to 32. Default value: 0 (sets automatically; depending on number of cores);</p> <p>-d --default – applies default settings.</p> |
| param system | show | rdcli param system show | Command displays current masking policy. |
| | modify | rdcli param system modify | <p>Command changes masking policy.</p> <p>Optional attribute:</p> <p>-lv --luns_visible_by_default <value> – masking policy. Possible</p> |

| Object | Method | Command Syntax | Functionality |
|------------------|--------|-------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | values: 0 - after creation LUNs are not available to any host (Enterprise mode); 1 - after creation LUNs are available to all hosts (Standard mode). |
| param ups | show | <code>rdcli param ups show</code> | Command displays current UPS settings. |
| | modify | <code>rdcli param ups modify</code> | <p>Command allows modifying settings of specified UPS driver.</p> <p>Optional attributes:</p> <p>-d --driver <value> – driver name. Possible values: 'apcsmart', 'bcmxcp', 'belkin', 'belkinunv', 'bestfcom', 'bestfortress', 'bestuferrups', 'bestups', 'blazer_ser', 'etapro', 'everups', 'gamatronic', 'genericups', 'isbmex', 'ivtscd', 'liebert', 'liebert-esp2', 'metasys', 'mge-shut', 'mge-utalk', 'microdowell', 'oldmge-shut', 'oneac', 'optiups', 'powercom', 'powerman-pdu', 'powerpanel', 'rhino', 'safenet', 'solis', 'tripplite', 'tripplitesu', 'upscod2', 'victronups', 'bcmxcp_usb', 'blazer_ser', 'blazer_usb', 'richcomm_usb', 'tripplite_usb', 'usbhid-ups';</p> <p>-p --port <value> – port (for COM);</p> <p>-s --shutdown <value> – shutdown timeout;</p> <p>-e --enable <value> – enables or disabled UPS. Possible values: yes, no;</p> |

| Object | Method | Command Syntax | Functionality |
|--------------|--------|---------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-m --master <value> – sets UPS type as master/slave. Possible values: yes, no;</p> <p>-a --ipaddr <value> – master-server IP address for slave.</p> |
| profile raid | create | rdcli profile raid create | <p>Command creates RAID static parameters profile with the specified name.</p> <p>Mandatory attribute:</p> <p>-n --profile_name <value> – profile name.</p> <p>The following attributes can be used with the command to set the corresponding parameters:</p> <p>-rao --read_ahead_at_once <value> – maximum size of read ahead request, in stripes. Possible values: from 1 to 64. Default value is 24;</p> <p>-wbo --write_back_at_once <value> – maximum number of stripes in single write back request. Possible values: from 1 to 64. Default value is 24;</p> <p>-mad --max_alt_data <value> – the size of alternative cache buffer, in MiB. Possible values: from 200 to 4096. Default value is 512;</p> <p>-rad --read_ahead_distance <value> – number of cache segments looked forward for data presence, used to decide whether read</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>ahead is needed. Possible values: from 1 to 300. Default value is 64;</p> <p>-rap --read_ahead_past <value> – minimum number of stripes in cache behind, required to start read ahead. Possible values: from 1 to 300. Default value is 24;</p> <p>-cac --cache_alg_clean <value> – displays type of cache algorithm for dirty segments replacement (LRU);</p> <p>-cad --cache_alg_dirty <value> – type of cache algorithm for dirty segments replacement: LRU or Cycle;</p> <p>-lwo --lru_wb_at_once <value> – specifies the maximum number of dirty segments to be flushed at once on the disks with LRU algorithm. Possible values: from 1 to 64, but not greater than write_back_at_once (wbo) value. The default value is 24;</p> <p>-cwo --cycle_wb_at_once <value> – specifies maximum number of dirty segments to be flushed at once on the disks with Cycle algorithm. Possible values: from 1 to 64, but not greater than write_back_at_once (wbo) value. The default value is 24;</p> <p>-mwbs --max_write_back_seq <value> - maximum number of concurrent sequential requests to be written back. Possible values: from 1 to 1024. Default value is 8;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p><code>-mwbr --max_write_back_rand <value></code> - maximum number of concurrent random requests to be written back. Possible values: from 1 to 1024. Default value is 8;</p> <p><code>-mwd --max_write_data <value></code> - total amount of data simultaneously written on the RAID, in MiB. Possible values: from 64 to 2048. Default value is 256;</p> <p><code>-rc --reconstruct_at_once <value></code> – maximum size of a reconstruction request, in MiB. Possible values: from 1 to 64. Default value is 24;</p> <p><code>-r --random_read <value></code> – enables or disables random read optimization. Possible values: 0 (disabled), 1 (enabled);</p> <p><code>-w --random_write <value></code> – enables or disables random write optimization. Possible values: 0 (disabled), 1 (enabled);</p> <p><code>-tn --thread_num <value></code> – the number of processing threads. Possible values: from 1 to 32. Default value is 16;</p> <p><code>-vcs --verify_cache_sum <value></code> – enables or disables verification of cache dirty segments match on the both controllers;</p> <p><code>-sl --verify_synd <value></code> – enables or disables Silent Data Corruption detection. Possible values: 0 (disabled), 1 (detection is enabled), 2 (detection and correction are enabled);</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|---------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-wt --write_through –performs write operations directly on the disks. Possible values: 0 (disabled); 1 (enabled); 2 (auto – enabled only for sequential write).</p> |
| | modify | rdcli profile raid modify | <p>Command allows editing the specified static RAID parameters profile.</p> <p>Mandatory attribute:</p> <p>-n --profile_name <value> – profile name.</p> <p>The following attributes can be used with the command:</p> <p>-cac --cache_alg_clean <value> – displays type of cache algorithm for dirty segments replacement (LRU);</p> <p>-cad --cache_alg_dirty <value> – type of cache algorithm for dirty segments replacement: LRU or Cycle;</p> <p>-cwo --cycle_wb_at_once <value> – specifies maximum number of dirty segments to be flushed at once on the disks with Cycle algorithm. Possible values: from 1 to 64, but not greater than write_back_at_once (wbo) value. The default value is 24;</p> <p>-lwo --lru_wb_at_once <value> – specifies the maximum number of dirty segments to be flushed at once on the disks with LRU algorithm. Possible values: from 1 to 64, but not greater than write_back_at_once (wbo) value. The default value is 24;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-mad --max_alt_data <value> – the size of alternative cache buffer, in MiB. Possible values: from 200 to 4096. Default value is 512;</p> <p>-mwbr --max_write_back_rand <value> - maximum number of concurrent random requests to be written back. Possible values: from 1 to 1024. Default value is 8;</p> <p>-mwbs --max_write_back_seq <value> - maximum number of concurrent sequential requests to be written back. Possible values: from 1 to 1024. Default value is 8;</p> <p>-mwd --max_write_data <value> - total amount of data simultaneously written on the RAID, in MiB. Possible values: from 64 to 2048. Default value is 256;</p> <p>-rad --read_ahead_distance <value> – number of cache segments looked forward for data presence, used to decide whether read ahead is needed. Possible values: from 1 to 300. Default value is 64;</p> <p>-rao --read_ahead_at_once <value> – maximum size of read ahead request, in stripes. Possible values: from 1 to 64. Default value is 24;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p><code>-rap --read_ahead_past <value></code> – minimum number of stripes in cache behind, required to start read ahead. Possible values: from 1 to 300. Default value is 24;</p> <p><code>-rcol --reconstruct_at_once <value></code> – maximum size of a reconstruction request, in MiB. Possible values: from 1 to 64. Default value is 24;</p> <p><code>-r --random_read <value></code> – enables or disables random read optimization. Possible values: 0 (disabled), 1 (enabled);</p> <p><code>-s --verify_synd <value></code> – enables or disables Silent Data Corruption detection. Possible values: 0 (disabled), 1 (detection is enabled), 2 (detection and correction are enabled);</p> <p><code>-tnl --thread_num <value></code> – the number of processing threads. Possible values: from 1 to 32. Default value is 16;</p> <p><code>-vcs --verify_cache_sum <value></code> – enables or disables verification of cache dirty segments match on the both controllers;</p> <p><code>-wbo --write_back_at_once <value></code> – maximum number of stripes in single write back request. Possible values: from 1 to 64. Default value is 24;</p> |

| Object | Method | Command Syntax | Functionality |
|-------------|--------|---------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-wt --write_through <value> – performs write operations directly on the disks. Possible values: 0 (disabled); 1 (enabled); 2 (auto – enabled only for sequential write);</p> <p>-w --random_write <value> – enables or disables random write optimization. Possible values: 0 (disabled), 1 (enabled).</p> |
| | delete | rdcli profile raid delete | <p>Command deletes the specified static RAID parameters profile.</p> <p>Mandatory attribute:</p> <p>-n --profile_name <value> – profile name.</p> <p>Optional attribute:</p> <p>-f --force – forces command execution.</p> |
| | show | rdcli profile raid show | <p>Command shows all the created static RAID parameters profiles.</p> <p>Optional attributes:</p> <p>-n --profile_name <value> – displays specified profile;</p> <p>-r --raid – displays list of RAIDs, parameters of which correspond to the specified profile.</p> |
| qosmic apps | show | rdcli qosmic apps show | Command shows full list of applications. |
| | create | rdcli qosmic apps create | <p>Command adds new application from the file with signatures.</p> <p>Mandatory attribute:</p> <p>-f --filepath <value> – path to the application.</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|--------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | Optional attribute: -i --icon <value> – path to the application icon. |
| | delete | rdcli qosmic apps delete | Deletes data on application (or file with signatures). Mandatory attribute: -n --name <value> – specifies name of application to be deleted and rebuilds model; Optional attribute: -f --force – forces command execution, while learning is in progress. |
| | modify | rdcli qosmic apps modify | Command changes settings of the RT (priorities). Mandatory attribute: -n --name <value> – application name. Optional attribute: -i --icon <value> – path to the application icon; -t --type <value> – application types. Possible values: avored (critically important); unwanted (non-priority), disabled (priority is not defined). |
| | rename | rdcli qosmic apps rename | Command renames application. Mandatory attributes: |

| Object | Method | Command Syntax | Functionality |
|-------------------|--------|------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| qosmic teacher | | | <p>-n --name <value> – current application name;</p> <p>-nn --nname <value> – new application name.</p> |
| | clean | <code>rdcli qosmic apps clean</code> | <p>Command deletes all warnings from the application.</p> <p>Mandatory attributes:</p> <p>-n --name <value> – current application name.</p> |
| | show | <code>rdcli qosmic teacher show</code> | <p>Command shows all applications and their settings.</p> |
| | start | <code>rdcli qosmic teacher start</code> | <p>Command starts process of learning.</p> <p>Mandatory attributes:</p> <p>-hn --hosts <value> – host names (comma-separated);</p> <p>-n --name <value> – application name.</p> <p>Optional attribute:</p> <p>-i --icon <value> – path to the application icon.</p> |
| | stop | <code>rdcli qosmic teacher stop</code> | <p>Command stops process of learning</p> |
| | resume | <code>rdcli qosmic teacher resume</code> | <p>Command resumes the process of learning</p> <p>Mandatory attributes:</p> <p>-hn --hosts <value> – host names (comma-separated);</p> <p>-n --name <value> – application name.</p> |

| Object | Method | Command Syntax | Functionality |
|--------------------------|--------|--------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | Optional attribute: -i --icon <value> – path to the application icon. |
| qosmic recognizer | show | <code>rdcli qosmic recognizer show</code> | Shows the process of recognizing |
| | start | <code>rdcli qosmic recognizer start</code> | Starts the process of recognizing |
| | stop | <code>rdcli qosmic recognizer stop</code> | Stops the process of recognizing |
| raid | show | <code>rdcli raid show</code> | Command displays list of all created RAIDs. Optional attributes: -n --name <value> – RAID name; -s --raid_status <value> – displays all RAIDs with the specified status; -si --sdcinfo <value> – displays information about silent data corruptions on RAID volumes; -l --legacy <value> – displays legacy RAIDs (RAIDs with 3.x topology). |
| | create | <code>rdcli raid create</code> | Command creates a RAID with the specified name, level and number of drives. Mandatory attributes: |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-n --name <value> - RAID name;</p> <p>-l --level <value> - RAID level;</p> <p>-d --drives <value> - RAID drives.</p> <p>Optional attributes:</p> <p>-cs --cache_size <value> – specifies cache size, in MiB;</p> <p>-ss --stripe_size <value> – specifies allows stripe size, in MiB;</p> <p>-rd --redundant_drives <value> – specifies number of redundant drives to create RAID N+M;</p> <p>-p --profile <value> - uses RAID parameters from the specified profile;</p> <p>-a --affinity <value> – specifies Affinity;</p> <p>-sp --sparepool <value> – assigns a SparePool to the array.</p> <p>-wt --write_through <>—performs write operations directly on the disks. Possible values: 0 (disabled); 1 (enabled); 2 (auto – enabled only for sequential write).</p> |
| | | | <p>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> |


| Object | Method | Command Syntax | Functionality |
|--------|--------|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | modify | <code>rdcli raid modify</code> | <p>Command modifies RAID parameters.</p> <p>Mandatory attribute:</p> <p>-n --name <value> - RAID name.</p> <p>Optional attributes:</p> <p>-rao --read_ahead_at_once <value> – maximum size of read ahead request, in stripes. Possible values: from 1 to 64. Default value is 24;</p> <p>-wbo --write_back_at_once <value> – maximum number of stripes in single write back request. Possible values: from 1 to 64. Default value is 24;</p> <p>-mad --max_alt_data <value> – the size of alternative cache buffer, in MiB. Possible values: from 200 to 4096. Default value is 512;</p> <p>-rad --read_ahead_distance <value> – number of cache segments looked forward for data presence, used to decide whether read ahead is needed. Possible values: from 1 to 300. Default value is 64;</p> <p>-rap --read_ahead_past <value> – minimum number of stripes in cache behind, required to start read ahead. Possible values: from 1 to 300. Default value is 24;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-cac --cache_alg_clean <value> – displays type of cache algorithm for dirty segments replacement (LRU);</p> <p>-cad --cache_alg_dirty <value> – type of cache algorithm for dirty segments replacement: LRU or Cycle;</p> <p>-lwo --lru_wb_at_once <value> – specifies the maximum number of dirty segments to be flushed at once on the disks with LRU algorithm. Possible values: from 1 to 64, but not greater than write_back_at_once (wbo) value. The default value is 24;</p> <p>-cwo --cycle_wb_at_once <value> – specifies maximum number of dirty segments to be flushed at once on the disks with Cycle algorithm. Possible values: from 1 to 64, but not greater than write_back_at_once (wbo) value. The default value is 24;</p> <p>-mwbs --max_write_back_seq <value> – maximum number of concurrent sequential requests to be written back. Possible values: from 1 to 1024. Default value is 8;</p> <p>-mwbr --max_write_back_rand <value> – maximum number of concurrent random requests to be written back. Possible values: from 1 to 1024. Default value is 8;</p> <p>-mwd --max_write_data <value> – total amount of data simultaneously written on the RAID, in MiB. Possible values: from 64 to 2048. Default value is 256;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|--------|----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-rcol --reconstruct_at_once <value> – maximum size of a reconstruction request, in MiB. Possible values: from 1 to 64. Default value is 24;</p> <p>-rl --random_read <value> – enables or disables random read optimization. Possible values: 0 (disabled), 1 (enabled);</p> <p>-wl --random_write <value> – enables or disables random write optimization. Possible values: 0 (disabled), 1 (enabled);</p> <p>-tnl --thread_num <value> – the number of processing threads. Possible values: from 1 to 32. Default value is 16;</p> <p>-vcs --verify_cache_sum <value> – enables or disables verification of cache dirty segments match on the both controllers;</p> <p>-sl --verify_synd <value> – enables or disables Silent Data Corruption detection. Possible values: 0 (disabled), 1 (detection is enabled), 2 (detection and correction are enabled);</p> <p>-wt --write_through <value> – performs write operations directly on the disks. Possible values: 0 (disabled); 1 (enabled); 2 (auto – enabled only for sequential write);</p> <p>-cs --cache_size <value> – cache size, in MiB;</p> <p>-p --profile <value> – changes RAID parameters to the parameters of the specified RAID profile;</p> |

| Object | Method | Command Syntax | Functionality |
|--------|------------|------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-sp --sparepool <value> – assigning a SparePool to the RAID;</p> <p>-ssd --ssdcache <value> – UID of SSD to be assigned as SSD cache.</p> |
| | migrate | <code>rdcli raid migrate</code> | <p>Command changes Affinity of the specified RAID (performs migration of RAID on the other node).</p> <p>Mandatory attribute:</p> <p>-n --name <value> – RAID name.</p> |
| | delete | <code>rdcli raid delete</code> | <p>Command deletes RAID (including from configuration file).</p> <p>Mandatory attribute:</p> <p>-n --name <value> – RAID name.</p> <p>Optional attribute:</p> <p>-f --force – forces command execution.</p> |
| | restore | <code>rdcli raid restore</code> | <p>Command restores RAID.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – RAID name.</p> |
| | init show | <code>rdcli raid init show</code> | <p>Command shows RAID initialization status.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – RAID name.</p> |
| | init start | <code>rdcli raid init start</code> | <p>Command starts RAID initialization.</p> <p>Mandatory attribute:</p> |

| Object | Method | Command Syntax | Functionality |
|--------|-------------------|------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-n --name <value> – RAID name.</p> <p>Optional attribute:</p> <p>-f --force – forces command execution.</p> |
| | init stop | rdcli raid init stop | <p>Command stops RAID initialization.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – RAID name.</p> <p>Optional attribute:</p> <p>-f --force – forces command execution.</p> |
| | reconstruct start | rdcli raid reconstruct start | <p>Command starts RAID reconstruction.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – RAID name.</p> |
| | reconstruct stop | rdcli raid reconstruct stop | <p>Command stops RAID reconstruction.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – RAID name.</p> |
| | replace | rdcli raid replace | <p>Command allows replacing drive with the specified UID (attribute --old) in the RAID with a drive which is not used in other RAIDs.</p> <p>If new drive is not specified, system automatically selects the first suitable drive.</p> <p>Mandatory attribute:</p> |

| Object | Method | Command Syntax | Functionality |
|-------------------|--------|-------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>-n --name <value> – RAID name;</p> <p>-ou --old <value> – old drive UID.</p> <p>Optional attributes:</p> <p>-nu --new <value> – new drive UID;</p> <p>-f --force – forces command execution.</p> |
| | reload | rdcli raid reload | <p>Command reloads the specified RAID.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – RAID name.</p> <p>Optional attribute:</p> <p>-f --force – forces command execution.</p> |
| | | |  It is not recommended to reload a RAID while the system is under high load. |
| sas target | show | rdcli sas target show | Command displays list of SAS targets, used in the system. |
| | modify | rdcli sas target modify | <p>Command allows associating SAS targets on the local and remote nodes.</p> <p>Mandatory attributes:</p> <p>-n --target_name <value> – local target name;</p> <p>-r --remote_target <value> – remote target name.</p> <p>Optional attribute:</p> |

| Object | Method | Command Syntax | Functionality |
|------------------|--------|-------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | -f --force – forces command execution. |
| sensor | show | <code>rdcli sensor show</code> | <p>Command displays list of all sensors in the system.</p> <p>Optional attribute:</p> <p>e --enclosure – displays information about sensors of all connected enclosures.</p> |
| session | show | <code>rdcli session show</code> | Command displays all established sessions. |
| sparepool | create | <code>rdcli sparepool create</code> | <p>Command creates new SparePool with the specified name and list of drives.</p> <p>Mandatory attributes:</p> <p>-n --name <value> – SparePool name;</p> <p>-d --drives <value> – drive UUIDs.</p> |
| | modify | <code>rdcli sparepool modify</code> | <p>Command modifies configuration of the specified SparePool.</p> <p>Mandatory attributes:</p> <p>-n --name <value> – SparePool name;</p> <p>-d --drives <value> – drive UUIDs.</p> |
| | delete | <code>rdcli sparepool delete</code> | <p>Command deletes the specified SparePool.</p> <p>Mandatory attribute:</p> <p>-n --name <value> – SparePool name;</p> <p>Optional attribute:</p> |

| Object | Method | Command Syntax | Functionality |
|--------|----------------|-----------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | -f --force – forces command execution. |
| | show | rdcli sparepool show | <p>Command displays all the created SparePools.</p> <p>Optional attribute:</p> <p>-n --name <value> – displays the information on the specified SparePool.</p> |
| system | show | rdcli system show | Command displays information about system: free RAM volume and RAIDIX software version. |
| | restart | rdcli system restart | Command restarts RAIDIX software. |
| | 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>Optional attributes:</p> <p>-hw --hardware_key – displays hardware key;</p> <p>-s --summary – displays information only about period of validity, functionality, license type and volume.</p> |
| | license create | rdcli system license create | <p>Command adds license file.</p> <p>Mandatory attribute:</p> |

| Object | Method | Command Syntax | Functionality |
|--------|----------------|-----------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | -f --filename <value> - path to the license file. |
| | license delete | rdcli system license delete | <p>Command deletes the specified license file.</p> <p>Mandatory attribute:</p> <p>-f --filename <value> - path to the license file.</p> |
| | config reread | rdcli system config reread | Command rereads the configuration file. |
| | config restore | rdcli system config restore | <p>Command restores the configuration file.</p> <p>Optional attributes:</p> <p>-a --auto – performs automatic recovery;</p> <p>-d --drive <value> – performs recovery from manually specified block device (for example: /dev/sda) (the option blocks the previous one);</p> <p>-f --file <value> – specifies full path to the file;</p> <p>-o --stdout <value> – shows configuration file, which will replace the current configuration file (option will work only if at least one of two previous attributes is specified);</p> <p>-l --legacy – restores RAIDs from the file with 3.x topology;</p> <p>--remote-node – performs recovery using configuration file from the remote node.</p> |

| Object | Method | Command Syntax | Functionality |
|--------|-----------------|-------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | settings show | <code>rdcli system settings show</code> | Command displays current system settings. |
| | settings modify | <code>rdcli system settings modify</code> | <p>Command allows changing system settings.</p> <p>Optional attributes:</p> <ul style="list-style-type: none"> -d --debug <value> – logging level; -ct --cmd_timeout <value> – command daemon request execution timeout; -rt --restore_timeout <value> – system restoration timeout; -et --exec_timeout <value> – maximum system command execution timeout. |
| | reboot | <code>rdcli system reboot</code> | Command reboots the system. |
| | shutdown | <code>rdcli system shutdown</code> | Command shuts down the system. |
| | user show | <code>rdcli system user show</code> | <p>Command displays list of all RAIDIX web interface users.</p> <p>Optional attribute:</p> <ul style="list-style-type: none"> -l --login <value> – displays a user with the specified login. |
| | user create | <code>rdcli system user create</code> | <p>Command creates a user with the specified name and password.</p> <p>Mandatory attributes:</p> <ul style="list-style-type: none"> -l --login <value> – user login; -p --password <value> – user password. |

| Object | Method | Command Syntax | Functionality |
|--------|-------------|--------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| | | | <p>Optional attributes:</p> <p>-la --language <value> – selects the interface language. Possible values: en (English); ru (Russian); ko (Korean); zh (Chinese).</p> |
| | user modify | rdcli system user modify | <p>Command changes user login and language.</p> <p>Mandatory attributes:</p> <p>-l --login <value> – user login;</p> <p>-la --language <value> – selects the interface language. Possible values: en (English); ru (Russian); ko (Korean); zh (Chinese).</p> |
| | user delete | rdcli user delete | <p>Command deletes the specified user.</p> <p>Mandatory attribute:</p> <p>-l --login <value> – user login.</p> |

Appendix B

Setting Up Fibre Channel connection on GNU/Linux

To set up Fibre Channel connection on GNU/Linux, perform the following procedure:

If you use 8 Gb adapter from ATTO Technology, load the `celernity8fc.ko` driver:

```
insmod celerity8fc.ko initiator_mode=1
```

If you use 16 Gb adapter from ATTO Technology, load the `celernity16fc.ko` driver:

```
insmod celerity16fc.ko initiator_mode=1
```

If you use a slow speed adapter from ATTO Technology, load the `celernityfc.ko` driver:

```
insmod celerityfc.ko initiator_mode=1
```

Appendix C

Setting Up InfiniBand Connection on Windows Server 2008 R2

To set up InfiniBand connection on Windows Server 2008 R2, perform the following procedure:

1. Download an appropriate InfiniBand HCA driver:
 9. <https://www.openfabrics.org/downloads/Windows/>
2. During the installation, select “Will be installed on local hard drive” in the SRP drop-down list to add SRP functionality (Figure 131):

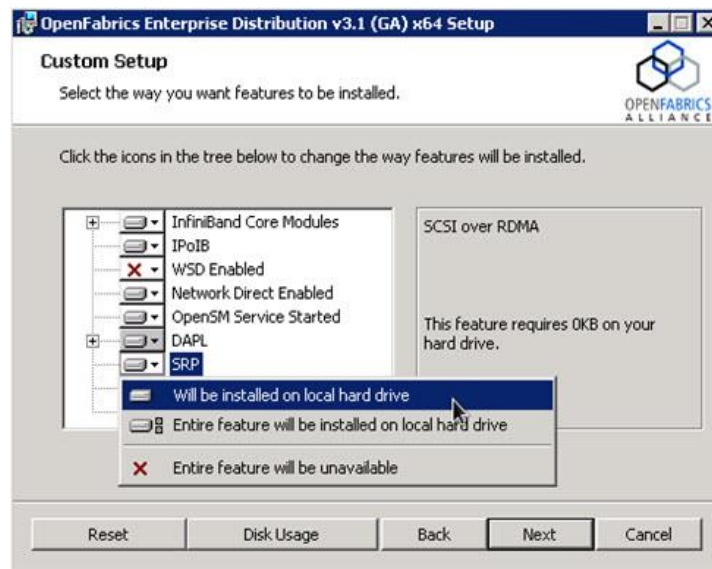


Figure 131. The InfiniBand driver installation window

3. Reload the system.
4. Launch the **InfiniBand Subnet Manager** (OpenSM Service):



Only one Subnet Manager can be active on each InfiniBand network: from the host side, from the storage system or from InfiniBand switch

5. From the Host Side: select **Auto** for **Startup Type** (Server Manager | Configurations | Services | InfiniBand Configuration and Management | Startup Type) and click **Start**.
6. From the Storage System Side: select **Active** or **Active/Passive** for **Start OpenSM Service** (Configuration Control | System Configuration | Details | InfiniBand Settings | Start OpenSM Service) and click Update.

7. From InfiniBand switch: for detailed information refer to manufacturer's manual.
8. To improve performance of the system, change in registry the following parameter:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\ibsrp\Parameters
/ModeFlags from 0 to 2.

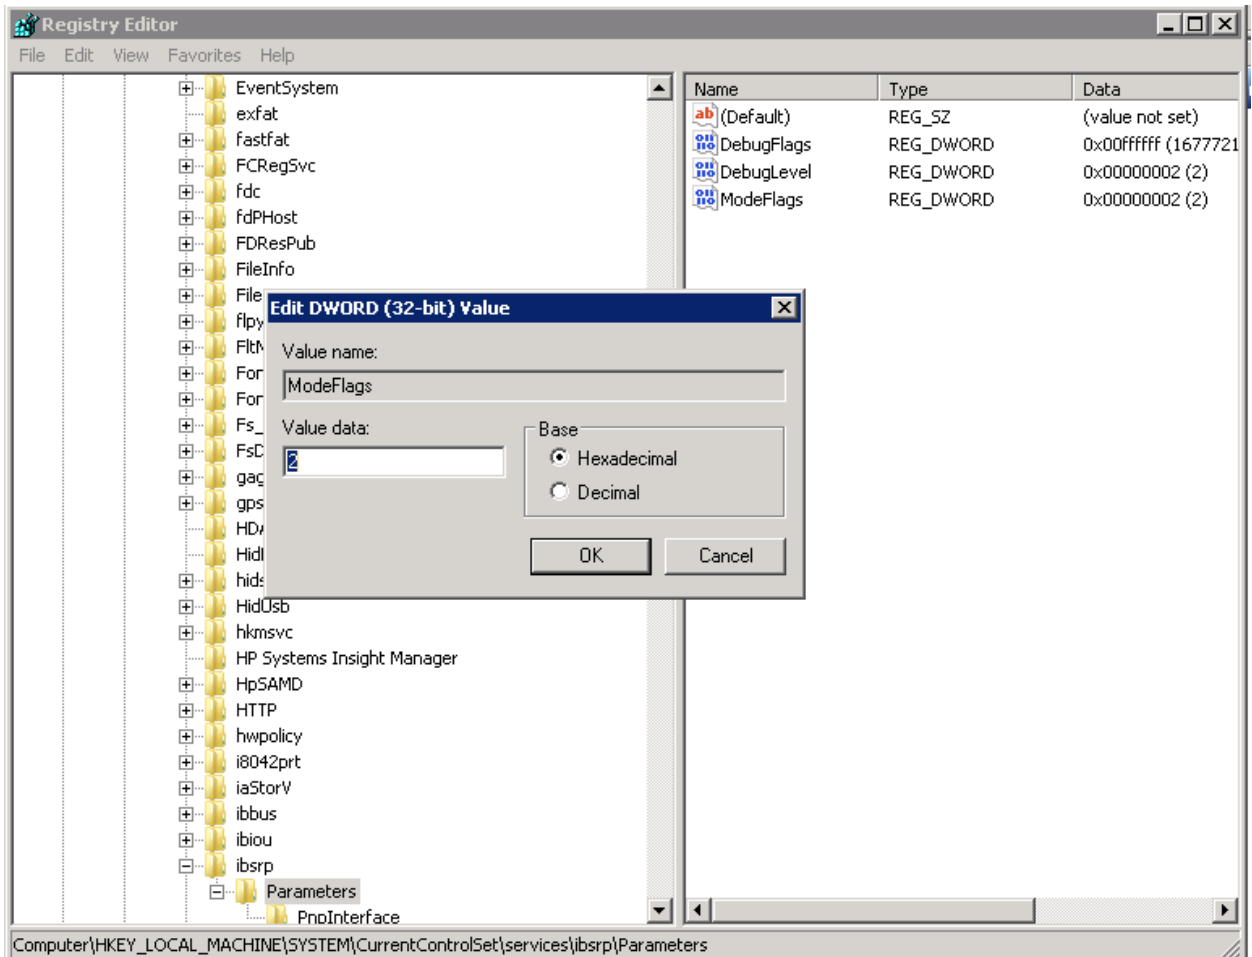


Figure 132. InfiniBand performance improvement

9. In **Computer Management Console**, select **Device Manager** and update the corresponding configuration. A new disk device appears in the list (Figure 133):

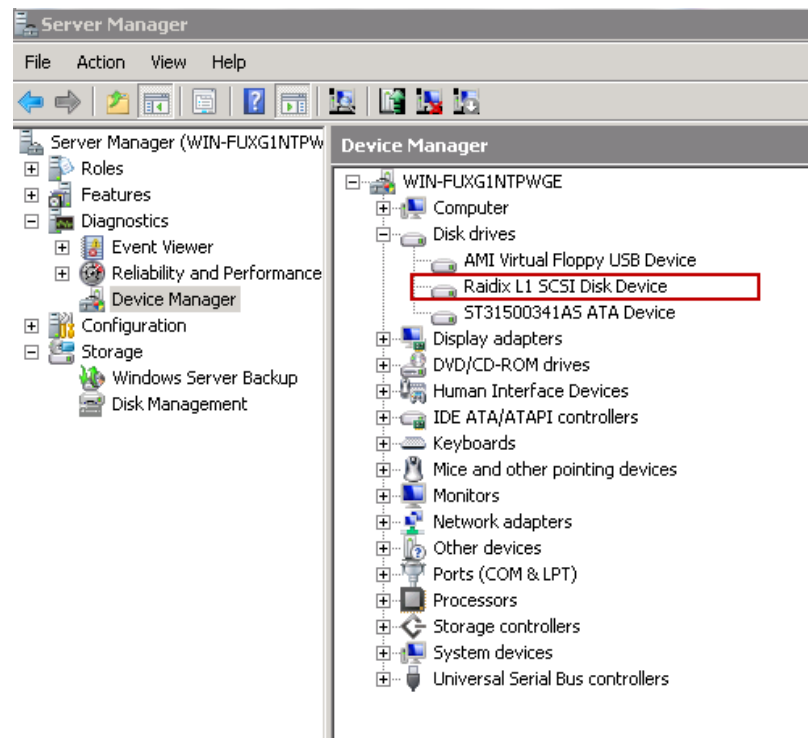


Figure 133. A new disk device displaying

Appendix D

Setting Up InfiniBand Connection on Suse Linux Enterprise Server OS

To set up InfiniBand connection on Suse Linux Enterprise Server Operating System, perform the following procedure:

1. Load the following driver:

```
usr/share/raidix/ib_srpt.ko
```

2. Start up the OpenSM service:

```
service opensmd start
```

3. Add a target to a corresponding host:

```
echo `ibsrpdm -c | tail -1` > /sys/class/infiniband_srpt/srpt-<device>-<port>/add_target
```

4. To improve the performance over InfiniBand, run the `ib_srp` module with the following parameter value `srp_sg_tablesize=255`:

```
rmmod ib_srp  
modprobe ib_srp srp_sg_tablesize=255
```

5. Execute the command to check the availability of the new Raidix LUN disk device:

```
lsscsi
```


Appendix E

MPIO Setup on Suse Linux Enterprise Server Operating System



To setup MPIO, make sure you have administrative permissions for Suse Linux Enterprise Server on the server connected to the storage system with RAIDIX 4.4.3 installed.

To setup MPIO on Suse Linux Enterprise Server Operating System, perform the following procedure:

1. Install the `multipath-tools` package.
2. Create the `/etc/multipath.conf` file containing the following:

```
defaults {
    user_friendly_names    yes
    path_checker            tur
    prio                   alua
    no_path_retry           10
        fast_io_fail_tmo   5
    polling_interval       5
    features                "0"
}

devices {
    device {
        vendor              "Raidix"
        product              "*"
        path_grouping_policy "group_by_prio"
        path_selector        "round-robin 0"
        path_checker         "tur"
        prio                 "alua"
        failback             immediate
        rr_min_io            100
        rr_weight            "priorities"
        no_path_retry        12
    }
}
```

```

        features      "1 queue_if_no_path"
        # For initiators with scsi_dh_alua
        #hardware_handler"1 alua"
    }
}

```

3. If CentOS 7.0, Red Hat 7 was installed on initiator, or scsi_dh_alua module was uploaded on the initiator, then you must uncomment (remove #) the following line:

```
#hardware_handler
```

4. Load the device-mapper multipath target module:

```
modprobe dm_multipath
```

5. Specify to run the multipath service during the load:

```
chkconfig multipathd on
```

6. Run the multipath service:

```
service multipathd start
```

7. To check the MPIO-device status, run the following command:

```
multipath -l
```

8. A list of MPIO device parameters appears:

```

mpathb (0Raidixl1 52b96d00a66d55359652b2a8) dm-0 Raidix,l1
mpathbo (0Raidixpetropavlovka 512b520b129adb1d07b223ce) dm-3
Raidix,petropavlovka
size=110G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='round-robin 0' prio=-1 status=active
|  `-- 4:0:0:3 sde 8:64 active undef running
`+- policy='round-robin 0' prio=-1 status=enabled
   `-- 5:0:0:3 sdi 8:128 active undef running
mpathbn (0Raidixkreml 512b520b129adb1d0125e389) dm-2
Raidix,kreml
size=100G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='round-robin 0' prio=-1 status=active
|  `-- 4:0:0:1 sdd 8:48 active undef running
`+- policy='round-robin 0' prio=-1 status=enabled
   `-- 5:0:0:1 sdh 8:112 active undef running
mpathbm (0Raidixnevsky 512b520b129adb1d898750e4) dm-1
Raidix,nevsky
size=110G features='1 queue_if_no_path' hwhandler='0' wp=rw

```

```
|+- policy='round-robin 0' prio=-1 status=active
|  `-- 4:0:0:2 sdc 8:32  active undef running
`+- policy='round-robin 0' prio=-1 status=enabled
   `-- 5:0:0:2 sdg 8:96  active undef running

mpathbl  (0Raidixtverskaya                    512b520b129adb1d0399ab37)  dm-0
Raidix,tverskaya

size=100G features='1 queue_if_no_path' hwhandler='0' wp=rw
|+- policy='round-robin 0' prio=-1 status=active
|  `-- 4:0:0:0 sdb 8:16  active undef running
`+- policy='round-robin 0' prio=-1 status=enabled
   `-- 5:0:0:0 sdf 8:80  active undef running
```

9. There are two ways of using the MPIIO devices:

- Direct device use,
- Using LVM2 (Logical Volume Manager).

10. Creating LVM2 volume using MPIIO:

- Make sure you have the **add /dev/dm-.*** rule specified for the required MPIIO device in the `/etc/lvm/lvm.conf` file;
- Prepare the `pvcreate /dev/dm-<N>` partition;
- Create the volume group using the `vgcreate vg /dev/dm-<N>` command;
- Create a logical volume named `<test>` sized 10GBs using the `lvcreate -L10G -ntest vg` command;
- Create a file system on the logical volume:

```
mkfs.ext3 /dev/vg/test
```

The `/dev/vg/test` logical volume is connected.

Appendix F

MPIO Setup on VMware ESX 4.1, 5.1, 5.5, 6.0



To setup MPIO, make sure you have administrative permissions for VMware ESX 4.1/5.1/5.5/6.0 on the server connected to the storage system with RAIDIX 4.4.3 installed.

To setup MPIO on VMware ESX 4.1, 5.1, 5.5 perform the following procedure:

Storage Setup

1. Create a RAID-array of the required size (for details, see section «[Creating a RAID](#)»).
2. Create a LUN with the 512 byte block size (for details, see section «[Creating a LUN](#)»).
3. Make sure that masking rules are specified to the LUN (for details, see section «[Masking Rules](#)»).

VMware ESX Setup

1. Connect the server with VMware ESX to the storage system with RAIDIX 4.4.3 software installed and reload ESX.
2. Install the FC driver if required.
3. Navigate to the ESX – **vSphere Client** management console and open the **Inventory** (Home | Inventory) window (Figure 134)

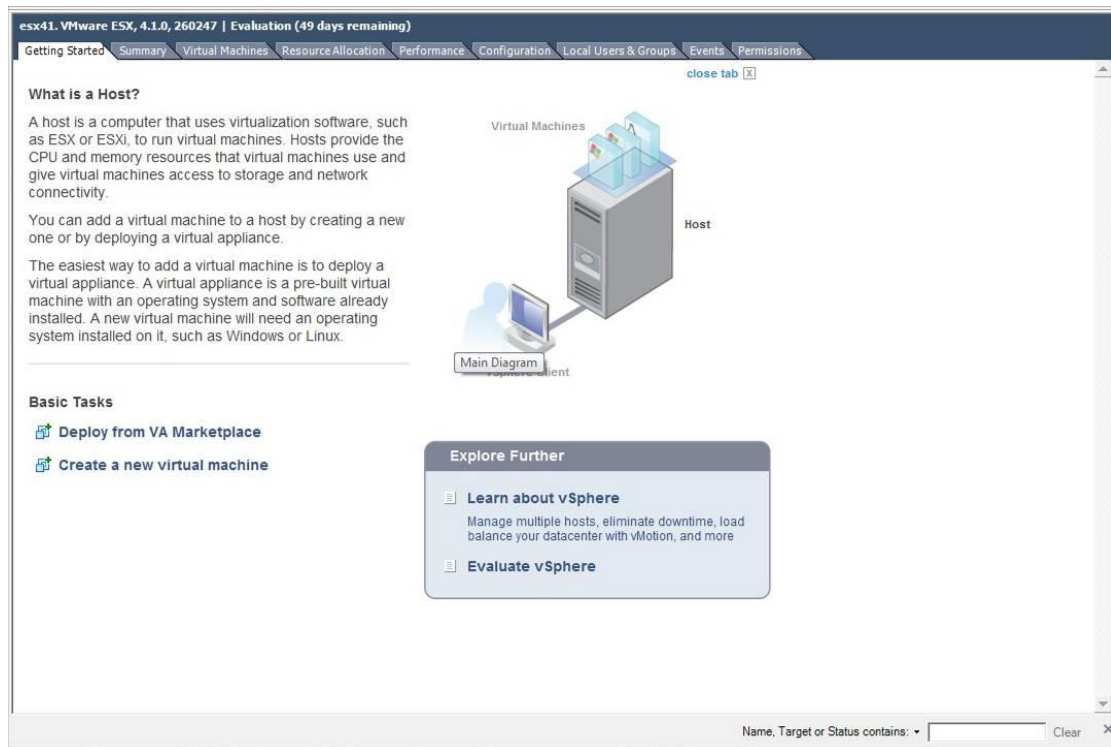


Figure 134. Inventory vSphere client

4. Select the **Configuration | Storage Adapters** tab (Figure 135)

| Name | Runtime Name | Operational State | LUN | Type |
|------------------------------------------------------------------|-----------------|-------------------|-----|------|
| Raidix Fibre Channel Disk (naa.60041560c4ff49d23d3b2a5f22136b22) | vmhba2:C0:T0:L0 | Mounted | 0 | disk |
| Raidix Fibre Channel Disk (naa.60041560c4ff49d23d3b2a5f4909bcd7) | vmhba2:C0:T0:L1 | Mounted | 1 | disk |
| Raidix Fibre Channel Disk (naa.60041560c4ff49d23d3b2a5fad1b8429) | vmhba2:C0:T0:L2 | Mounted | 2 | disk |

Figure 135. Storage Adapters page

5. Select the **Paths** tab and make sure that the statuses of paths for connection with active/passive controller are displayed correctly. Check the selected MPIO policy (Figure 136 and Figure 137).

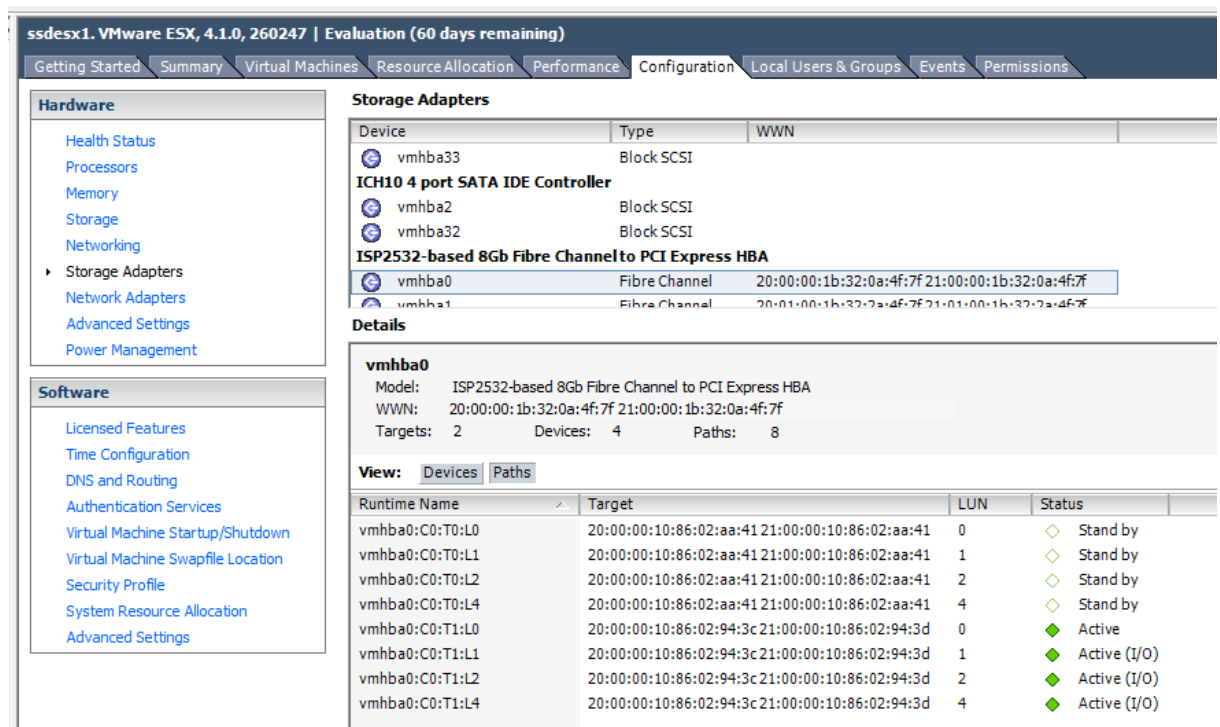


Figure 136. Configuration | Paths

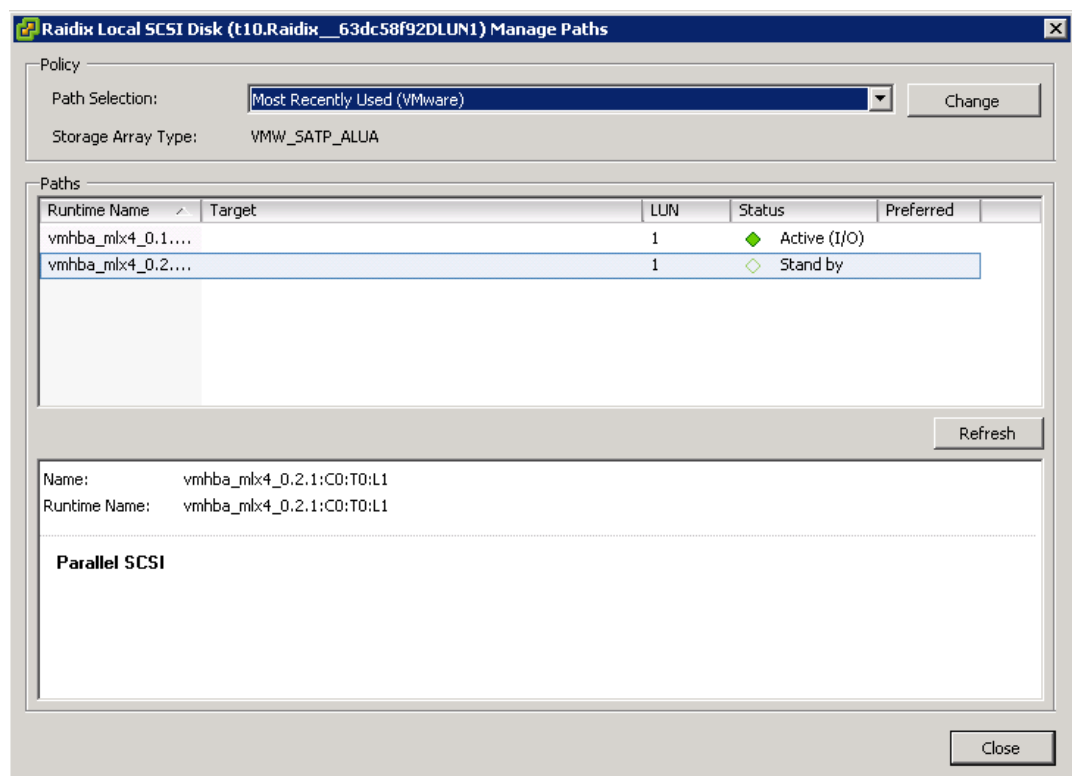


Figure 137 Selecting of MPIO policy

- Make sure that newly created LUN is displayed on the Fibre Channel ports connected to the storage.

7. If the LUN information is not displayed, run the **Rescan All** command located in the upper right part of the window and then click **Refresh**.
8. Select the **Storage** from the menu (Figure 138).

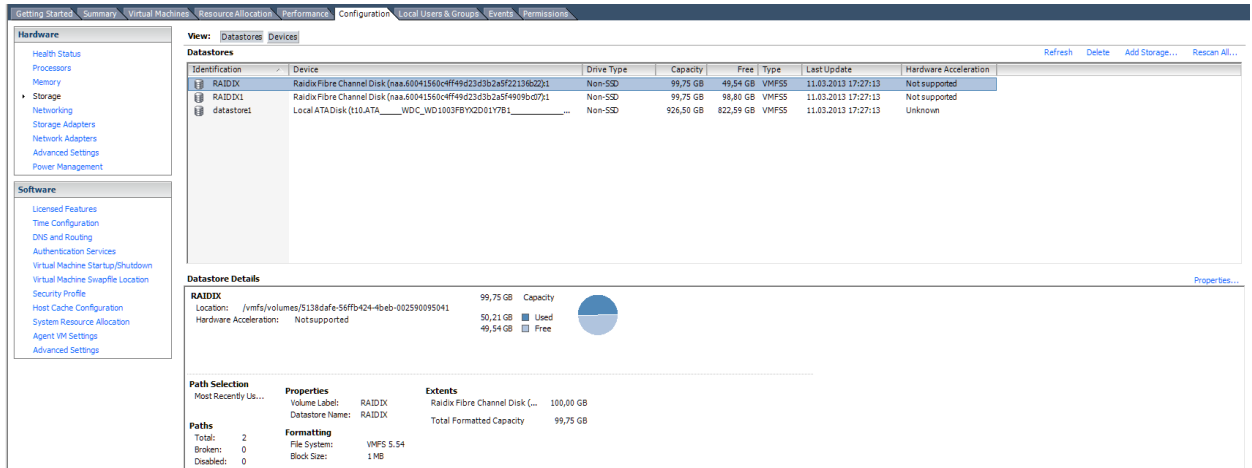


Figure 138. Storage menu

9. Select the **Add Storage** in the **Datastores** window (Figure 139). Select a check box for the **Storage Type** parameter to create a LUN (Add Storage | Storage Type | Disk/LUN) and click **Next** (Figure 139).

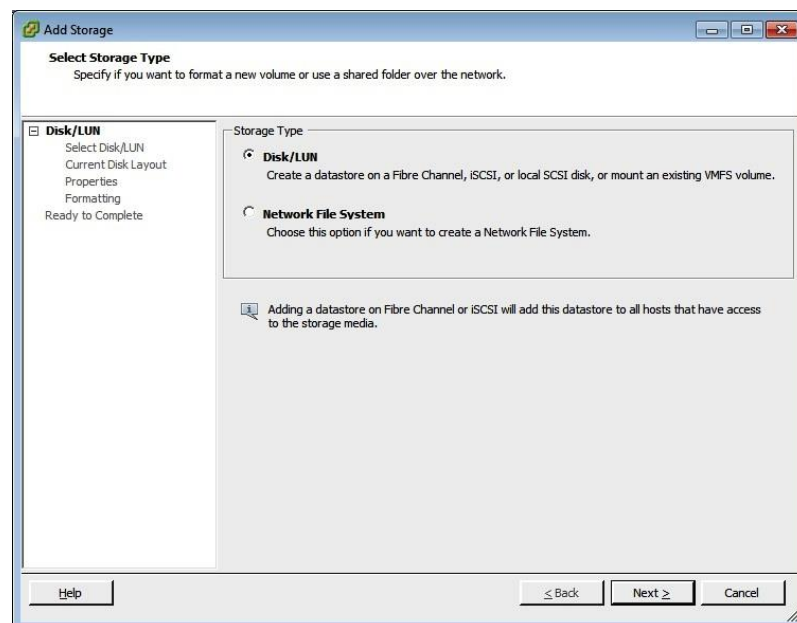


Figure 139. Add Storage window

10. Follow the Windows Wizard on adding storages. When you complete the procedure using the Wizard, a new storage appears in the **Datastores** window.

Appendix G

MPIO Setup on Windows Server 2008



To setup MPIO, make sure you have administrator permissions for Windows Server 2008 on the server connected with the storage system with RAIDIX 4.4.3 installed.

To setup MPIO on Windows Server 2008, perform the following procedure:

Storage Setup

1. Create a RAID-array of the required size (for details, see section «[Creating a RAID](#)»).
2. Create a LUN (for details, see section «[Creating a LUN](#)»).
3. Make sure that masking rules are specified to the LUN (for details, see section «[Masking Rules](#)»).

Windows Server 2008 Setup

1. Connect the server with Windows Server 2008 to the storage with RAIDIX 4.4.3 software installed.
2. Make sure you have installed the Fibre Channel adapter driver.
3. Install Windows MPIO. To do so:
 - Open the **Server Manager** (Figure 140):

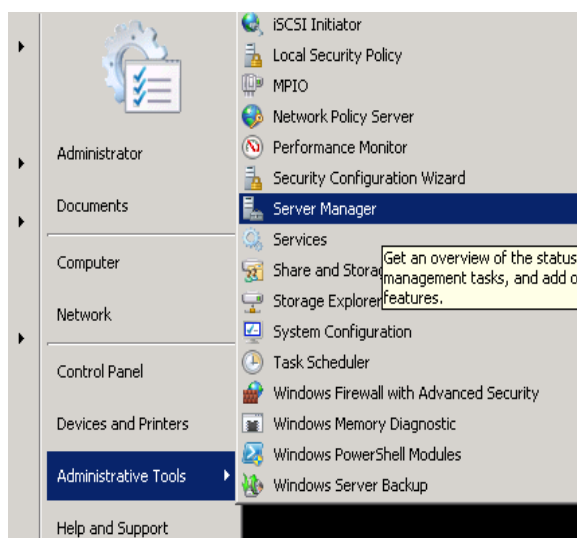


Figure 140. Starting the installation Windows MPIO

- Select **Features** (Figure 141):




Figure 141. Server Manager

- Click the Add **Features** link, select the **Multipathing I/O** component in the window, which opens and click **Install**.
4. Add drives and setup MPIO balancing mode:
- Navigate to the device MPIO device management panel (Start | Administrative Tools | MPIO);
 - Before adding new drives, verify, that previously created unused MPIO-devices (not related to any of existing LUNs) are deleted. To do it, open **MPIO Devices** section and delete all previously created Raidix devices (example on Figure 142) if there are so. Reload the system

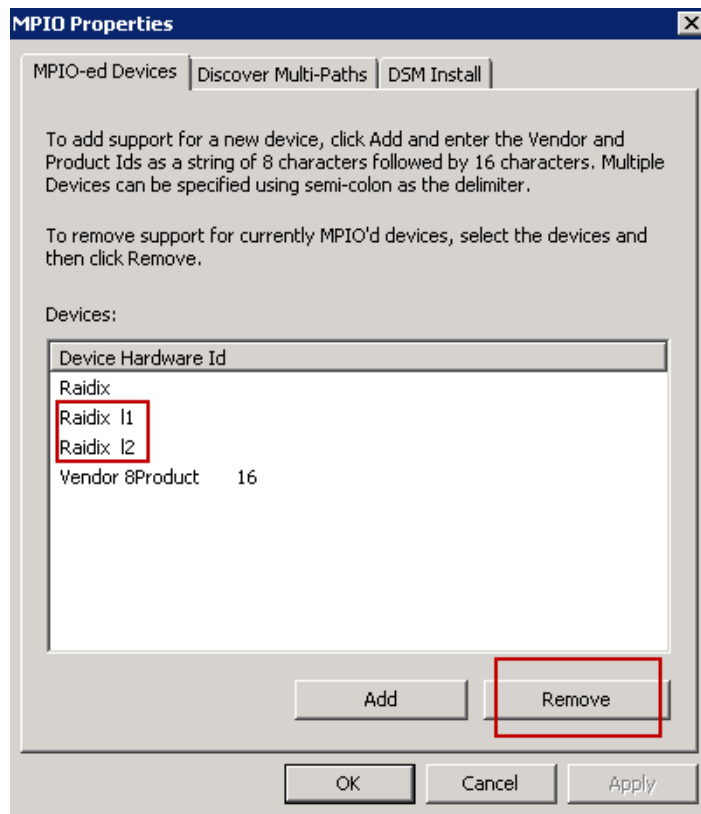


Figure 142. Delete unused MPIO-devices

- Open the Discover Multi-Paths tab;
- In the **SPC-3 compliant** section, select the device to add and click **Add** (Figure 143):




Figure 143. MPIO. Adding device

- Open the Device Manager window, navigate to the Disk Drives node, right-click and select Properties for the corresponding devices.
- Select the MPIO tab in the window, which opens.
- Specify the **Load Balance Policy** and select **Fail Over Only** policy (Figure 144Ошибка! Источник ссылки не найден.).

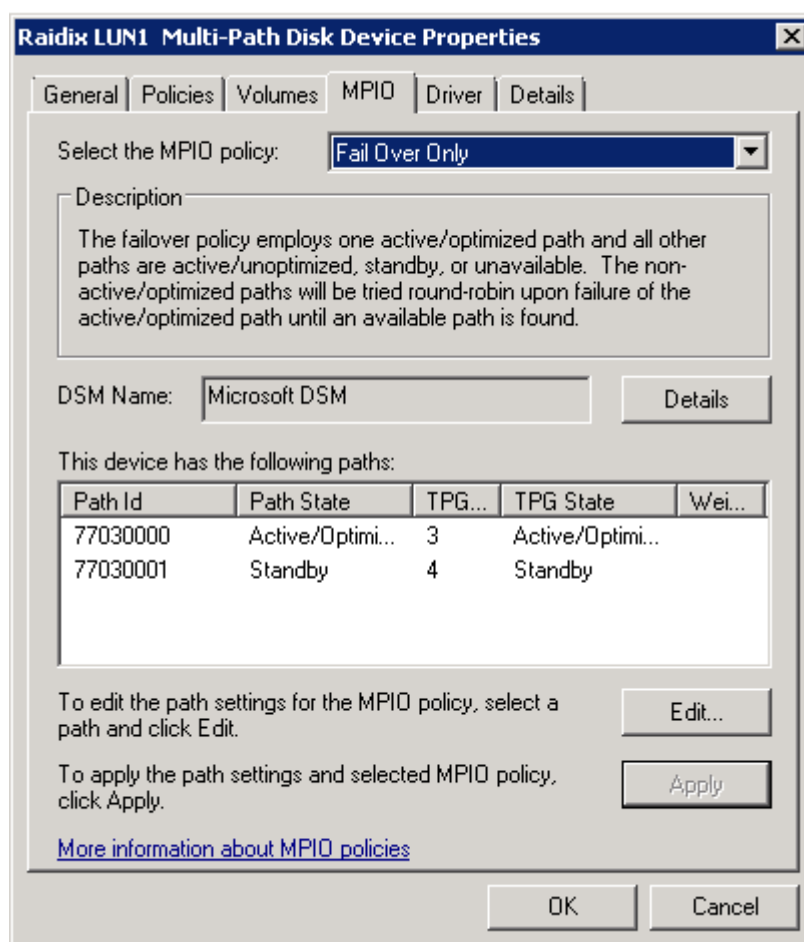


Figure 144. MPIO Policy: Fail Over Only



If the **Fail Over Only** policy is used for all paths except for one, select the **Standby** state

Fail Over Only is a mode when no load balancing is performed. An “active path” corresponds to the connection with active controller and standby path corresponds to the connection with passive controller. All operations with the storage are performed in accordance with initial path. In case of failure (the data for this path is unavailable), the system starts using one of the backup paths.

In the Microsoft DSM parameters of the any RAIDIX Multipath device, enable the Path Verify Parameter, by doing the following:

1. In the **DSM Name: Microsoft DSM** line, click **Details**.
2. Mark the **Path Verify Enable** parameter. (Figure 145).

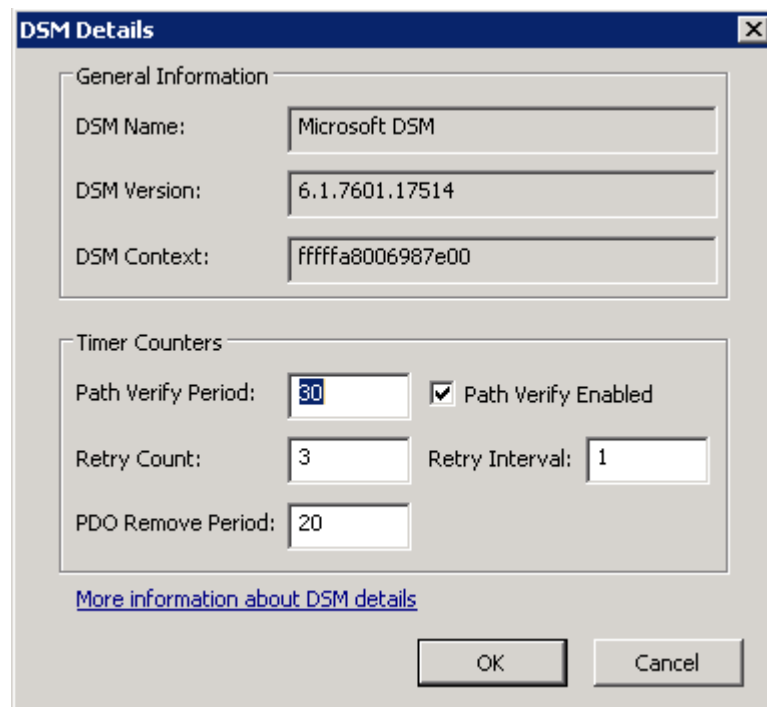


Figure 145 Path Verify Enabled

After you enable the option on one device, the option will be automatically applied to all other devices.

Appendix H

MPIO Setup on Windows Server 2008 R2



To setup MPIO, make sure you have administrator permissions for Windows Server 2008 R2 on the server connected with the storage system with RAIDIX 4.4.3 installed.

To setup MPIO on Windows Server 2008 R2, perform the following procedure:

Storage Setup

1. Create a RAID-array of the required size.
2. Create a LUN.
3. Make sure that masking rules are specified to the LUN.

Windows Server 2008 R2 Setup

1. Connect the server with Windows Server 2008 R2 to the storage with RAIDIX 4.4.3 software installed.
2. Make sure you have installed the Fibre Channel adapter driver.



A connection via QLogic 16Gb Fibre Channel adapter can cause a 15-minutes delay of LUNs appearance on the storage and on the initiator side

3. Install Windows MPIO. To do so:
 - Open the **Server Manager** window (Figure 146):

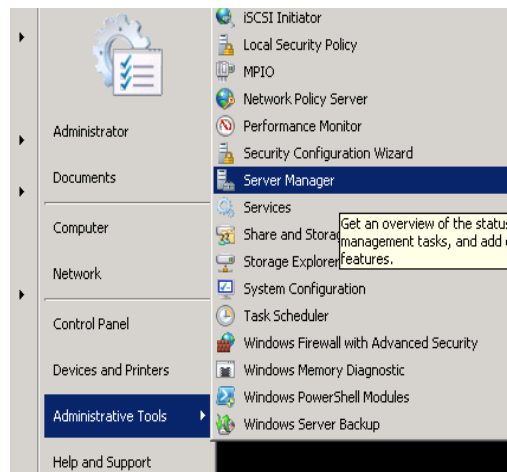


Figure 146. Windows MPIO

- Select **Features** (Figure 147):




Figure 147. Server Manager

4. Click the **Add Features** link. In the window that opens, select **Multipathing I/O** component and click **Install**.
5. Add drives and setup MPIO balancing mode. To do so, perform the following:
 - Navigate to the device MPIO Device Management panel (Start | Administrative Tools | MPIO);
 - Before adding new drives, verify, that previously created unused MPIO-devices (not related to any of existing LUNs) are deleted. To do it, open **MPIO Devices** section and delete all previously created Raidix devices (see example on Figure 148) if there are so. Reload the system;

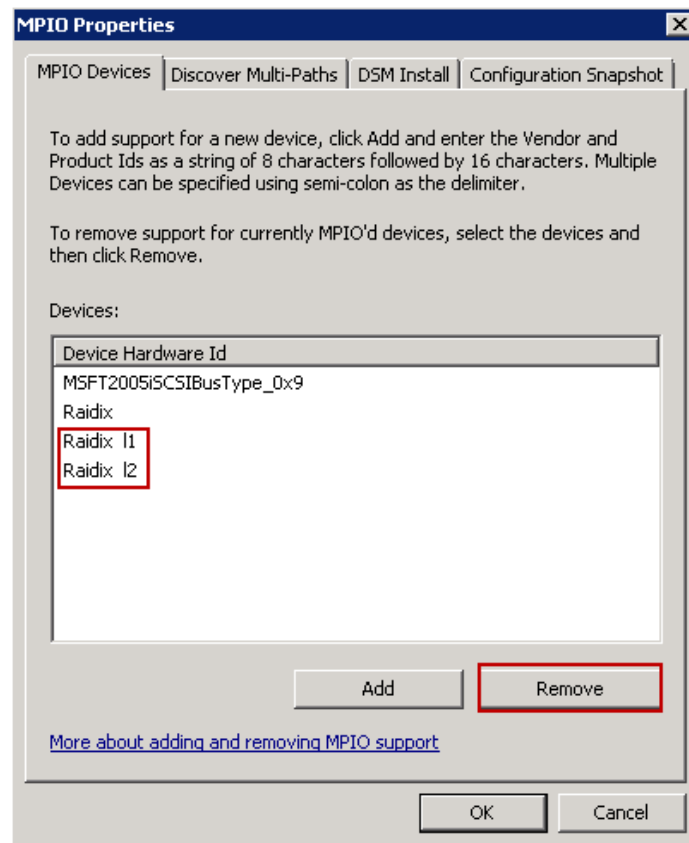


Figure 148. Delete the unused LUNs

- Open the Discover Multi-Paths tab;
- In the **SPC-3 compliant** section, select the device to add and click Add (Figure 149).

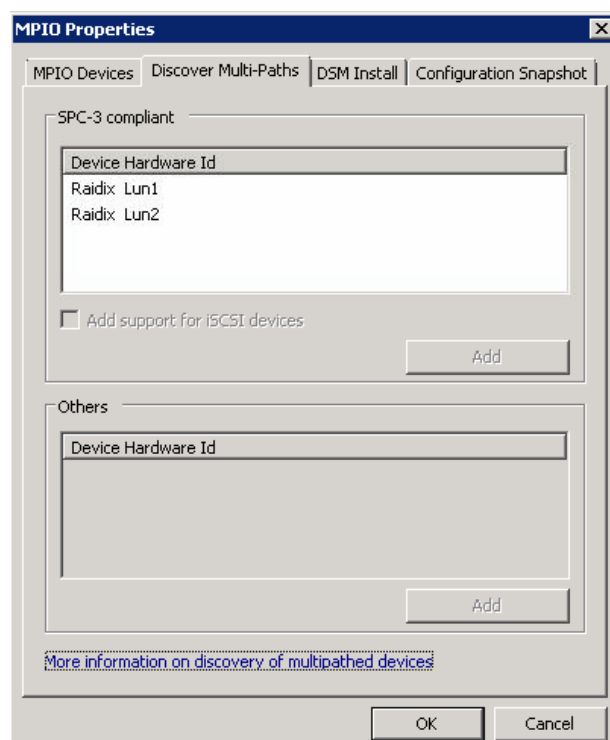


Figure 149. MPIO. Add device

- Open the **Device Manager** window, navigate to the **Disk Drives** node, right-click and select **Properties** for the corresponding devices.
- Select the **MPIO** tab in the window, which opens.
- Specify the **Load Balance Policy** and select **FailOver Only** policy

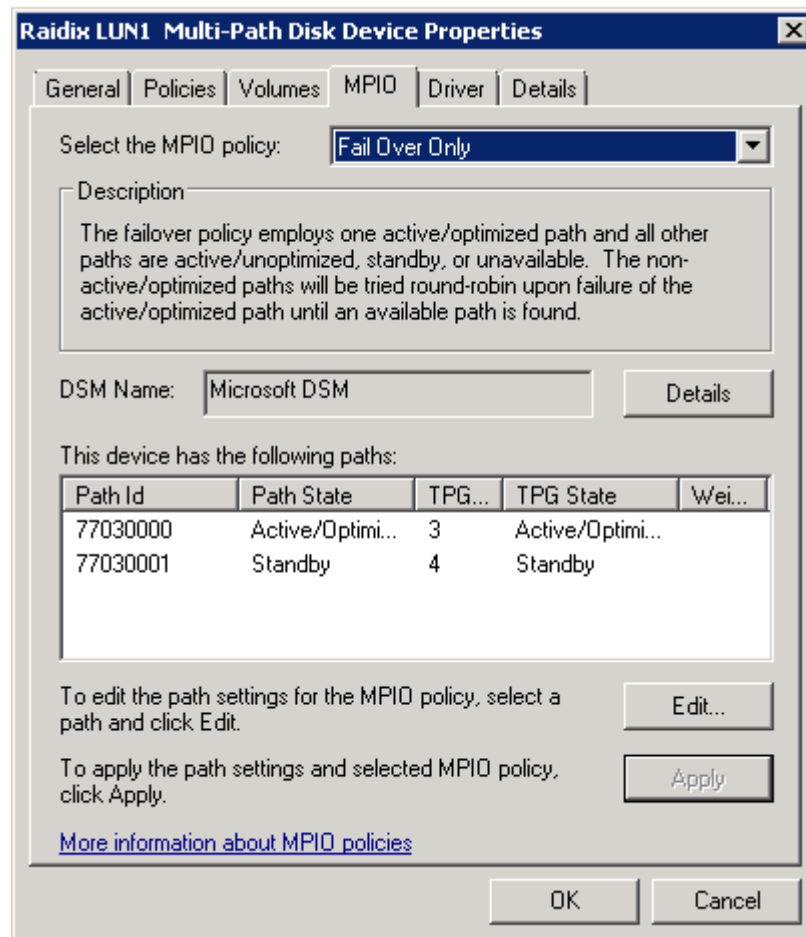


Figure 150. MPIO Policy: Fail Over Only



If the **Fail Over Only** policy is used for all paths except for one, select the **Standby** state

Fail Over Only is a mode when no load balancing is performed. An “active path” corresponds to the connection with active controller and standby path corresponds to the connection with passive controller. All operations with the storage are performed in accordance with initial path. In case of failure (the data for this path is unavailable), the system starts using one of the backup paths.

In the Microsoft DSM parameters of the any RAIDIX Multipath device, enable the Path Verify Parameter, by doing the following:

1. In the **DSM Name: Microsoft DSM** line, click **Details**.

Mark the **Path Verify Enable** parameter. (Figure 145).

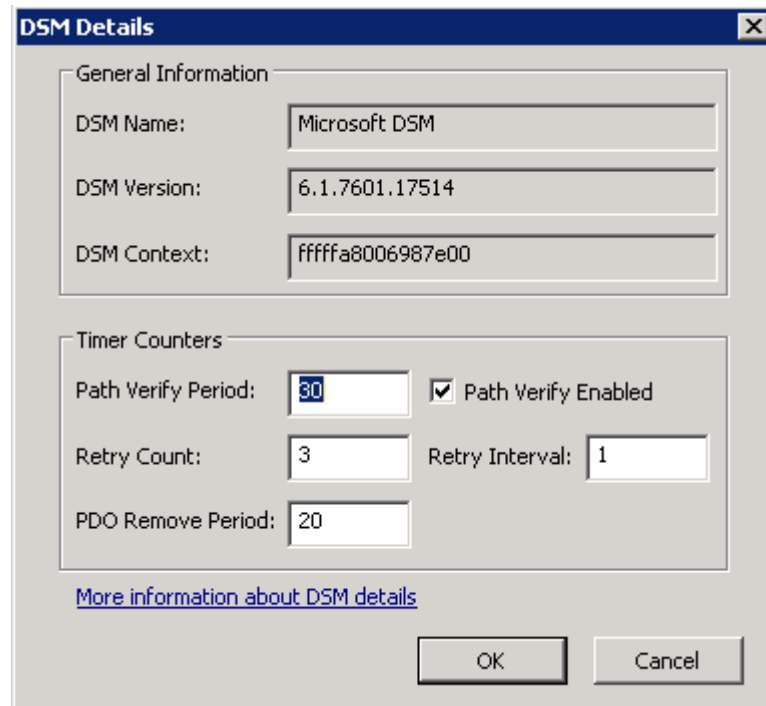


Figure 151 Path Verify Enabled

After you enable the option on one device, the option will be automatically applied to all other devices.