

**XINNOR**

**Xinnor xiRAID 4.0.4  
Administrator's Guide**

# Contents

<b>Xinnor xiRAID 4.0.4 Administrator's Guide.....</b>	<b>4</b>
Introduction.....	4
About Xinnor xiRAID 4.0.4.....	5
Using xicli.....	6
Command-Line Interface (CLI) Overview.....	6
Accepting EULA.....	7
License.....	8
Using RAID.....	10
Creating the RAID.....	10
RAID Initialization.....	14
Showing RAID State.....	15
Deleting the RAID.....	19
Managing RAID configuration.....	20
Unloading the RAID.....	20
RAID Reconstruction.....	21
Increasing Size and Changing Level of RAID.....	22
Changing RAID Parameters.....	28
Drives.....	31
Manual Drive Replacement or Excluding.....	31
Automatic Drive Replacement.....	32
Drive I/O Error Counter.....	36
Removing Drive Metadata and Resetting Current Error Count.....	39
Configuration Files and Metadata.....	40
xiRAID RAID Configuration.....	40
Importing RAID.....	42
Configuration File Recovery.....	45
Restoring the RAID.....	48
Scanning RAIDs and Drives, LED Indication.....	49
Update Check Service.....	51

- Notifications..... 52
  - System Log..... 52
  - Setting up Email Notifications..... 54
- Number of CPU Threads of the xiraid Module..... 58
- General Configuration Recommendations..... 60
  - RAID Creation..... 60
  - Using NVMe-oF devices to create a RAID..... 61
  - RAID and System Setup Recommendations..... 61
  - File System Mounting Examples..... 67
- DKMS Specifics when Updating Linux Kernel..... 69
- Authenticating gRPC Client..... 70
- Troubleshooting..... 72

# Xinnor xiRAID 4.0.4 Administrator's Guide

Instructions, practical guides, and tips for administering Xinnor xiRAID 4.0.4

## Introduction

### Intended Audience

This guide is intended for administrators and users of RAIDs based on the Xinnor xiRAID 4.0.4 software.


The guide contains instructions on how to configure and manage RAIDs in Xinnor xiRAID 4.0.4.


### Guide Conventions

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

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

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

 Tip - a note, which provides valuable information.

 Warning - binding instructions to guarantee the proper work of the software.

## About Xinnor xiRAID 4.0.4

Xinnor xiRAID 4.0.4 is high-performance software RAID developed specifically for NVMe storage devices and new types of SAN networks. Xinnor xiRAID 4.0.4 technologies use high potential of Flash devices (NVMe, SAS, SATA) to create a fast fault-tolerant RAID available as a local block device with opportunity of export via network by using auxiliary software.

Xinnor xiRAID 4.0.4 is a Linux kernel module and a management utility, which is built and configured for the most popular distributions (see the Xinnor xiRAID 4.0.4 System Requirements document). The software is installed on servers with slots for Flash memory devices or with connected JBODs. Xinnor xiRAID 4.0.4 enables you to combine drives into high-performance fault-tolerant RAIDs.

## Xinnor xiRAID 4.0.4 Specifications

Supported RAID levels	<ul style="list-style-type: none"><li>• RAID 0</li><li>• RAID 1</li><li>• RAID 10</li><li>• RAID 5</li><li>• RAID 6</li><li>• RAID 7.3</li><li>• RAID 50</li><li>• RAID 60</li><li>• RAID 70</li><li>• RAID N+M</li></ul>
Maximum number of drives in a RAID	64.
Maximum number of drives in the system	Depends on hardware configuration.
Maximum number of RAIDs	128.
Maximum RAID size	Defined by drive sizes.

Space for RAID metadata storage

The system reserves the first 100 MiB and the last 100 MiB of each drive in a RAID.

## Using xicli

Manage your software Xinnor xiRAID in Linux by using the `xicli` program.

Most of the commands listed in this document require superuser privileges. Please log in as an administrator or root to run these. However, the following commands can be run without superuser privileges: all commands with the `show` subcommand (raid show, config show, drive faulty-count show, settings eula show, license show etc), setting eula modify, and any command with the `--help` parameter.

## Command-Line Interface (CLI) Overview

### Conventions on CLI command syntax

Item format	Description
item	A required item (command, subcommand, argument, option).
<item>	A placeholder variable.
[item]	An optional item.

In the CLI, enter commands in the following format:

```
# xicli <command> <subcommand> <required_args> [optional_args]
```

To show the full list of commands, run

```
# xicli -h
```

To show the `xicli` version, run

```
# xicli -v
```

CLI syntax specifics:

1. Type the arguments of the subcommands in one line.
2. Subcommand arguments are separated by spaces.
3. Use short or long forms of subcommand argument options.

4. To get the list of all subcommands and arguments, add the `-h` option::

```
# xicli <command> <subcommand> -h
```

### The list of available commands <command>

config	Operations with the configuration file.
drive	Operations with the drives.
license	Operations with the license.
log	Operations with the event log.
mail	Operations with the mail notifications.
pool	Operations with the spare pools.
raid	Operations with the RAIDs.
settings	Operations with the additional settings of the <code>xicli</code> program.
update	Operations with the Update Check service.

A detailed description of the commands and subcommands is presented in the corresponding sections of the document.

## Accepting EULA

The first time after the installation you run any `xicli` command (except `settings eula modify`, `settings eula show`), you will be prompted to accept the EULA.

After accepting the EULA, the ran command executes, and you can use Xinnor xiRAID 4.0.4.



Commands related to the EULA don't require running `xiraid.target` service.

To change the acceptance status of the EULA, run

```
# xicli settings eula modify
```

## Argument for the `eula modify` subcommand

### Required argument

---

<code>-s</code>	<code>--status</code>	The status of the EULA acceptance. Possible values: <b>accepted</b> , <b>not_accepted</b> .
-----------------	-----------------------	--

To view the acceptance status of the EULA, run


```
# xicli settings eula show
```

## Argument for the `eula show` subcommand

### Optional argument

---

<code>-f</code>	<code>--format</code>	Output format:  <ul style="list-style-type: none"> <li>• <b>table</b>;</li> <li>• <b>json</b>;</li> <li>• <b>prettyjson</b> – human-readable json.</li> </ul> <p>The default: <b>table</b>.</p>
-----------------	-----------------------	---

 Accepting the EULA is not required after updating Xinnor xiRAID.

## License

You can manage the license with the command

```
# xicli license <subcommand>
```

Subcommands for the `license` command:

<code>delete</code>	Delete the current license.
<code>show</code>	Show info on the current license.

---



update

Update the current license.

To start working with the system, add the valid license file on each node. To do so, you need the hardware key (hwkey) which can be found by running the command:

```
# xicli license show
```

Command output example when no license was added:

```
Kernel version: 3.10.0-957.27.2.el7.x86_64
hwkey: 2EE10A9F12DD2662
license_key: null
version: 0
crypto_version: 0
created: 0-0-0
expired: 0-0-0
disks: 0
levels: 0
type: nvme
disks_in_use: 0
status: expired
```

Command output example when a license was added:

```
Kernel version: 3.10.0-957.27.2.el7.x86_64
hwkey: 2EE10A9F12DD2662
license_key: E3C6C0C0EEABE4274DBACACAE8D31E
76C2FB30CF18A3EB832BC007BC0B32DC4CDF9D4C135
D78344C19DF3E0D9AA995F64C4E0AFA9A441A51292D
version: 1
crypto_version: 0
created: 2019-4-14
expired: 2020-12-31
disks: 128
levels: 60
type: nvme
disks_in_use: 0
status: valid
```

### Description of the `license` command output

Kernel version	Kernel version.
hwkey	Hardware key.
license_key	License key.
version	Software version.

## Description of the `license` command output (continued)

<code>crypto_version</code>	Version of crypto-API for the license generator.
<code>created</code>	The date when the license was created.
<code>expired</code>	License expiration date.
<code>disks</code>	Maximum number of drives.
<code>levels</code>	Maximum RAID level. RAID levels from minimal to maximal: 0, 1, 10, 5, 6, 7 (stands for 7.3), 50, 60, 70 (includes N+M).
<code>type</code>	Drive type.
<code>disks in use</code>	Number of used drives in the system.
<code>status</code>	License state.

You can save the command output as a text file by running the command:

```
# xicli license show > license_request.txt
```

To get your license key, send your hardware key to the Xinnor support team at [support@xinnor.io](mailto:support@xinnor.io).

After you get your license file, copy it to the system, and apply the license key by running the command:

```
# xicli license update -p </path/to/>license.txt
```

To check the applied license, run:

```
# xicli license show
```

## Using RAID

In this chapter you will learn how to create and delete xiRAID RAID objects.

### Creating the RAID

You can create the RAID with the command

```
# xicli raid create <args> [optional_args]
```

For the argument descriptions, see the table below. For recommendations on configuring RAID settings, see the [General Configuration Recommendations](#) chapter.



Creating xiRAID RAID over xiRAID RAID devices is not allowed. To pool a large number of drives into a single address area, use RAIDs 10, 50, 60, 70.

Minimum number of drives required to create a RAID:

- of levels 5, 6 – at least 4 drives;
- of level 7.3 - at least 6 drives;
- of level 10 – at least 2 drives (the number of drives must be even);
- of level 0 – at least 1 drive;
- of level 1 – at least 2 drives;
- of levels 50, 60 – at least 8 drives (make sure the total drives number is multiple of the `--group-size` parameter value);
- of level 70 - at least 12 drives (make sure the total drives number is multiple of the `--group-size` parameter value);
- of level N+M – at least 8 drives.

Creating xiRAID RAIDs requires at least 1024 MiB of RAM.

## Arguments for the `create` subcommand

### Required arguments

<code>-n</code>	<code>--name</code>	The name of the RAID.  The maximum RAID name length is <b>28</b> characters.
<code>-l</code>	<code>--level</code>	The level of the RAID: <b>0, 1, 5, 6, 7, 10, 50, 60, 70, or nm.</b>
<code>-d</code>	<code>--drives</code>	The list of block devices ( <code>/dev/sd*</code> , <code>/dev/mapper/mpath*</code> , <code>/dev/nvme*</code> , <code>/dev/dm-*</code> ) separated by a space.
<code>-gs</code>	<code>--group_size</code>	<b>Only for RAIDs 50, 60, or 70.</b>  The number of drives for one RAID group of level 5, 6, or 7.3 of the appropriate RAID 50, 60, or 70.

**Arguments for the create subcommand (continued)**

		Possible values are integers from <b>4</b> to <b>32</b> .
-sc	--synd_cnt	<b>Only for RAID<sub>s</sub> N+M.</b> The number of syndromes M. Possible values are integers from <b>4</b> to <b>32</b> . Additional conditions: $N+M \leq 64$ and $M \leq N$ .
Optional arguments		
-bs	--block_size	RAID block size: <b>512</b> or <b>4096</b> bytes. The default: <b>4096</b> .
-inp	--init_prio	<b>Except RAID 0.</b> Initialization priority in %. Possible values are from <b>0</b> to <b>100</b> (maximum rate of initialization). The default: <b>100</b> .
-mwe	--merge_write_enabled	<b>Except RAID<sub>s</sub> 0, 1, 10.</b> Enable ( <b>1</b> ) or disable ( <b>0</b> ) the Merge function for write operations. The default: <b>0</b> .
-mre	--merge_read_enabled	<b>Except RAID<sub>s</sub> 0, 1, 10.</b> Enable ( <b>1</b> ) or disable ( <b>0</b> ) the Merge function for read operations. The default: <b>0</b> .
-ml	--memory_limit	RAM usage limit in MiB. Possible values: <b>0</b> and integers from <b>1024</b> to <b>1048576</b> . The <b>0</b> value sets unlimited RAM usage. The default: <b>0</b> .
-mm	--merge_max	<b>Except RAID<sub>s</sub> 0, 1, 10.</b> Maximum wait time (in microseconds) for stripe accumulation for the Merge functions. Possible values: integers from <b>1</b> to <b>100000</b> .

**Arguments for the `create` subcommand (continued)**

		The default: <b>1000</b> .
<code>-mw</code>	<code>--merge_wait</code>	<p><b>Except RAIDs 0, 1, 10.</b></p> <p>Wait time (in microseconds) between requests for the Merge functions.</p> <p>Possible values: integers from <b>1</b> to <b>100000</b>.</p> <p>The value must be less than the <code>merge_max</code> value.</p> <p>The default: <b>300</b>.</p>
<code>-rcp</code>	<code>--recon_prio</code>	<p><b>Except RAID 0.</b></p> <p>Reconstruction priority in %.</p> <p>Possible values are from <b>0</b> to <b>100</b> (maximum rate of reconstruction).</p> <p>The default: <b>100</b>.</p>
<code>-re</code>	<code>--resync_enabled</code>	<p><b>Except RAIDs 0, 1, 10.</b></p> <p>Enable (<b>1</b>) or disable (<b>0</b>) the Resync function.</p> <p>The default: <b>1</b>.</p>
<code>-rl</code>	<code>--request_limit</code>	<p>Number of simultaneous I/O requests on RAID.</p> <p>Possible values: from <b>0</b> (unlimited) to <b>4294967295</b>.</p> <p>The <b>0</b> value disables the restriction.</p> <p>The default: <b>0</b>.</p>
<code>-rsp</code>	<code>--restripe_prio</code>	<p>Restripping priority in %.</p> <p>Possible values are from <b>0</b> to <b>100</b> (maximum rate of restripping).</p> <p>The default: <b>100</b>.</p>
<code>-se</code>	<code>--sched_enabled</code>	<p>Enable (<b>1</b>) or disable (<b>0</b>) the scheduling function.</p> <p>The default: <b>0</b>.</p>
<code>-sp</code>	<code>--sparepool</code>	Name of the spare pool to assign to the RAID.

## Arguments for the `create` subcommand (continued)

`-ss`      `--strip_size`      Strip size in KiB.  
 Possible values: **16**, **32**, **64**, **128**, or **256**.  
 The default: **16**.

Example: Create the RAID 5 named "media5" over 4 NVMe drives — "nvme0n1", "nvme1n1", "nvme2n1", "nvme3n1", with strip size equal to 64 KiB and enabled Merge function for write operations.

```
# xicli raid create -n media5 -l 5 -d /dev/nvme0n1 /dev/nvme1n1 /
dev/nvme2n1 /dev/nvme3n1 -ss 64 -mwe 1
```

## RAID Initialization

Initialization will start automatically after a RAID (except RAID 0) is created.



To improve the system performance under the load, try **decreasing initialization priority** by changing the corresponding RAID parameter.



The random write performance is higher on an initialized RAID.

To start or continue the RAID initialization, run

```
# xicli raid init start <arg>
```

### Argument for the `init start` subcommand

Required argument

---

<code>-n</code>	<code>--name</code>	The name of the RAID.
-----------------	---------------------	-----------------------

To stop the RAID initialization, run

```
# xicli raid init stop <arg>
```

### Argument for the `init stop` subcommand

Required argument

## Argument for the `init stop` subcommand (continued)

`-n`                                    `--name`                                    The name of the RAID.

Example: Start initialization of the RAID "media5":

```
# xicli raid init start -n media5
```

## Showing RAID State

You can view info about the RAID with the command

```
# xicli raid show [optional_args]
```

### Arguments for the `show` subcommand

Optional arguments

<code>-n</code>	<code>--name</code>	The name of the RAID. Without the argument, show info on all xiRAID RAIDs.
<code>-o</code>	<code>--online</code>	Only show RAIDs that are in the "online" state (RAIDs that were not unloaded by the <code>raid unload</code> command). The argument takes no value.
<code>-u</code>	<code>--units</code>	Dimension: <ul style="list-style-type: none"> <li>• <b>s</b> – sectors (1 sector=512 bytes);</li> <li>• <b>k</b> – kilobytes;</li> <li>• <b>m</b> – megabytes;</li> <li>• <b>g</b> – gigabytes.</li> </ul> The default: <b>g</b> .
<code>-f</code>	<code>--format</code>	Output format: <ul style="list-style-type: none"> <li>• <b>table</b>;</li> <li>• <b>json</b>;</li> <li>• <b>prettyjson</b> – human-readable json.</li> </ul>

## Arguments for the `show` subcommand (continued)

The default: **table**.

`-e`      `--extended`      Show extended output.

The argument takes no value.

Example: Show information on the "media5" RAID:

```
# xicli raid show -n media5 -e
```

RAIDs name	static	state	devices	health	wear	serials	params	info
media5	size: 4 GiB level: 1 strip_size: 16 block_size: 4096 sparepool: - active: True config: True	online initialized	0 /dev/sdh online 1 /dev/sdd online	100% 100%	N/A N/A	drive-scsi7 drive-scsi3	init_prio : 100 init_depth : 128 recon_prio : 100 recon_depth : 64 memory_limit_mb : 0 sched_enabled : 0 request_limit : 0 restripe_prio : 100	memory_usage_mb : -

## Description of the `show` subcommand output

### Row

### Description

`name`      RAID name.

`static`      Static RAID parameters:

- `size`.
- `level`.
- `synd_cnt` – only for RAID5s N+M – number of syndromes.
- `block_size` – RAID block size.
- `group_size` – only for RAID50, 60, and 70 – size of the corresponding RAID group.
- `strip_size`.
- `sparepool` – name of the assigned spare pool.
- `active`:
  - True, if the RAID's block device is in the system.
  - False, if:



## Description of the `show` subcommand output (continued)

Row	Description
	<ul style="list-style-type: none"> <li>▪ The RAID was not loaded after reboot.</li> <li>▪ The RAID is unloaded.</li> </ul> <ul style="list-style-type: none"> <li>• config:               <ul style="list-style-type: none"> <li>◦ True, if the RAID is in the configuration file.</li> <li>◦ False, if the RAID is missing.</li> </ul> </li> </ul>
state	RAID state: <ul style="list-style-type: none"> <li>• online – the RAID is available and ready to work.</li> <li>• initialized – initialization is finished.</li> <li>• initing – the RAID is initializing.</li> <li>• degraded – the RAID is available and ready for work but some drives are missing or failed.</li> <li>• reconstructing – the RAID is reconstructing.</li> <li>• offline – the RAID is unavailable.</li> <li>• need_recon – the RAID needs reconstruction.</li> <li>• need_init – the RAID needs initialization.</li> <li>• read_only – the license has expired. The RAID is read-only.</li> <li>• unrecovered – RAID can't complete reconstruction because of unrecoverable sections.</li> <li>• none – RAID was unloaded via the <code>unload</code> command or was not restored after reboot.</li> <li>• restriping – RAID is restriping.</li> <li>• need_resize – restriping was finished, the RAID size increase is available.</li> <li>• need_restripe – restriping was stopped and not finished.</li> </ul>
devices	The list of devices included in the RAID, and their current states:

**Description of the `show` subcommand output (continued)**

Row	Description
	<ul style="list-style-type: none"> <li>• online – the drive is active.</li> <li>• offline – the drive is missing or unavailable.</li> <li>• reconstructing – the drive is in process of reconstruction.</li> <li>• need_recon – the drive needs reconstruction.</li> </ul>
health	<p>To show, use the command with the <code>-e</code> parameter.</p> <p>Percent of valid drive data.</p> <p>When health is 100% – no reconstruction required.</p>
wear	<p>To show, use the command with the <code>-e</code> parameter.</p> <p>The wear percentage of the SSD or NVMe drive.</p> <p>When the drive reaches the 90% threshold, the system sends an error message to the mail.</p> <p>The S.M.A.R.T. values "Percentage used endurance indicator" and "Percentage Used" are used to check SSD and NVMe drives respectively.</p>
serials	<p>To show, use the command with the <code>-e</code> parameter.</p> <p>Serial numbers of drives in RAID.</p>
params	<p>To show, use the command with the <code>-e</code> parameter.</p> <p>Editable RAID parameters:</p> <ul style="list-style-type: none"> <li>• <code>init_prio</code> – (except RAID 0) initialization priority: from 0% to 100%.</li> <li>• <code>init_depth</code> – the queue depth of initialization requests.</li> <li>• <code>recon_prio</code> – (except RAID 0) reconstruction priority: from 0% to 100%.</li> <li>• <code>recon_depth</code> – the queue depth of reconstruction requests.</li> <li>• <code>memory_limit_mb</code> – the value limited RAM usage, in megabytes.</li> </ul>

## Description of the `show` subcommand output (continued)

Row	Description
	<ul style="list-style-type: none"> <li>• <code>merge_write_enabled</code> – is the function Merge enabled (1) or disabled (0) for write operations.</li> <li>• <code>merge_read_enabled</code> – is the function Merge enabled (1) or disabled (0) for read operations.</li> <li>• <code>resync_enabled</code> – is the function Resync enabled (1) or disabled (0).</li> <li>• <code>sched_enabled</code> – is the function Scheduling enabled (1) or disabled (0).</li> <li>• <code>request_limit</code> – number of simultaneous I/O requests on RAID (0 for no limit).</li> <li>• <code>restripe_prio</code> – priority for restriping: from 0% to 100%.</li> <li>• <code>merge_wait_usecs</code> – waiting time between requests when Merge is enabled.</li> <li>• <code>merge_max_usecs</code> – maximum time to wait for stripe accumulation with Merge enabled.</li> </ul>
info	<p data-bbox="560 1137 895 1169">Dynamic RAID values:</p> <ul style="list-style-type: none"> <li>• <code>init_progress</code> – initialization progress: from 0% to 100%.</li> <li>• <code>recon_progress</code> – reconstruction progress: from 0% to 100%.</li> <li>• <code>memory_usage_mb</code> – amount of RAM usage; if <code>memory_limit_mb = 0</code> (not limited), then <code>memory_usage_mb</code> is not displayed.</li> <li>• <code>restripe_progress</code> – restriping progress: from 0% to 100%.</li> </ul>

## Deleting the RAID



Warning! The result of the command is irreversible. Read the description carefully.

You can delete the RAID without possibility to restore the RAID and data on it with the command

```
# xicli raid destroy <arg> [optional_args]
```

### Arguments for the `destroy` subcommand

Mutually exclusive required arguments

<code>-n</code>	<code>--name</code>	The name of the RAID.
<code>-a</code>	<code>--all</code>	Delete all the xiRAID RAIDs. The argument takes no value.

Optional arguments

	<code>--force</code>	Force the command execution.
	<code>--config_only</code>	Remove RAID only from config.

Example: Deleting the RAID "media5":

```
# xicli raid destroy -n media5
```

## Managing RAID configuration

In this chapter you will learn how to configure and manage xiRAID RAID objects.

### Unloading the RAID

You can remove (unload) the RAID with possibility to restore the RAID and save data on it with the command

```
# xicli raid unload <arg>
```

### Arguments for the `unload` subcommand

Mutually exclusive required arguments

<code>-n</code>	<code>--name</code>	The name of the RAID.
<code>-a</code>	<code>--all</code>	Unload all available xiRAID RAIDs.

## Arguments for the `unload` subcommand (continued)

The argument takes no value.

Example: Unloading the RAID "media5":

```
# xicli raid unload -n media5
```

To restore unloaded RAIDs, run:

```
# xicli raid restore {-n <raid_name>|-a}
```

## RAID Reconstruction

Reconstruction of a RAID (except RAID 0) starts automatically after a drive has been replaced in the RAID. While a RAID is being reconstructed, the functions initialization and restripping are paused.



To improve the system performance under the workload, try **decreasing reconstruction priority** by changing the corresponding RAID parameter.

To start the RAID reconstruction, run

```
# xicli raid recon start <arg>
```

### Argument for the `recon start` subcommand

Required argument

<code>-n</code>	<code>--name</code>	The name of the RAID.
-----------------	---------------------	-----------------------

To stop the RAID reconstruction, run

```
# xicli raid recon stop <arg>
```

### Argument for the `recon stop` subcommand

Required argument

<code>-n</code>	<code>--name</code>	The name of the RAID.
-----------------	---------------------	-----------------------

Example: Start RAID "media5" reconstruction:

```
# xicli raid recon start -n media5
```

## Increasing Size and Changing Level of RAID

In this chapter, you will learn about the following RAID operations:

- Changing the RAID level with the addition of new disks.
- Increasing the size of a RAID by adding new disks.
- ~~Increasing the size of a RAID by replacing its disks with larger disks (*vertical scaling* RAID operation).~~



The operation of increasing the RAID level by replacing the disks with larger ones is not supported in Xinnor xiRAID 4.0.4. Please refer to the "Xinnor xiRAID 4.0.4 Known Issues" for more information.

RAID operations that add new disks to a RAID consist of two steps: *restripping* (`raid restripe` command), which selects the disks to be added and the RAID level, and *resizing* (`raid resize` subcommand), which applies changes to the RAID (which is in `need_resize` status).

## Requirements and Specifics of Restripping and Resizing

### Requirements and specifics:

- Except RAID N+M.
- Only one RAID can be restriped at a time.
- To improve the performance of your system under workload, try to **change the priority of restripping** by changing the corresponding RAID parameter.
- RAID state must not be one of the following:
  - `offline`;
  - `need_restripe`;
  - `restripping`;
  - `degraded`.



The RAID level can only be changed if the number of added drives is sufficient to fit all the data stored prior to the level change. The number of data drives in a group must not decrease. Note that for RAID levels 50, 60, and 70, the total number of drives must be a multiple of the group size. Please keep in mind that RAID 50 includes one syndrome drive per group, while RAID 60 has two and RAID 70 has three. These drives do not store information but are necessary for data recovery in case of loss.

Example: Adding 1 drive to a RAID 60 configuration with 14 disks and 2 groups is not feasible. The only possible configuration for 15 drives is 3 groups of 5 drives, with only 3 of them containing data. Performing this operation would result in a new RAID with fewer data drives, hence it cannot be done. However, it is possible to add 2 drives to create a new RAID configuration consisting of 2 groups, each containing 8 drives (2 syndrome and 6 data drives). This will increase the data storage space.

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

Current level	New level	Requirements	Minimal number of drives you should add
RAID 0	RAID 0		1
	RAID 1	RAID 0 must contain only 1 drive.	1
	RAID 10	RAID 0 contains only 1 drive.	3
		RAID 0 contains more than 1 drive	The number of drives to be added must be equal to the number of drives in the RAID 0.
RAID 1	RAID 1		2
	RAID 10		2
	RAID 5		2
RAID 10	RAID 10		4
	RAID 5		1
	RAID 50	The number of RAID 50 disks must be a multiple of its group size.	1

## The available options for RAID level changes and the minimum required number of drives (continued)

<b>Current level</b>	<b>New level</b>	<b>Requirements</b>	<b>Minimal number of drives you should add</b>
RAID 5	RAID 5		1
	RAID 6		1
	RAID 10	The number of RAID 10 disks must be an even number.  The number of information disks in a group must not decrease.	1
	RAID 50	The number of disks in RAID 50 must be at least 8.  The number of disks in RAID 50 must be a multiple of its group size.  The number of information disks in a group must not decrease.	1
RAID 6	RAID 6		1
	RAID 7.3		1
	RAID 60	The number of disks in RAID 60 must be at least 8.  The number of disks in RAID 60 must be a multiple of its group size.  The number of information disks in a group must not decrease.	1
RAID 7.3	RAID 7.3		1
	RAID 70	The number of disks in RAID 70 must be at least 8.  The number of disks in RAID 70 must be a multiple of its group size.	1



## The available options for RAID level changes and the minimum required number of drives (continued)

Current level	New level	Requirements	Minimal number of drives you should add
		The number of information disks in a group must not decrease.	
RAID 50	RAID 50	The number of disks in RAID 50 must be a multiple of its group size.	1
		The number of information disks in a group must not decrease.	
	RAID 60	The number of disks in RAID 60 must be at least 8.	1
		The number of disks in RAID 60 must be a multiple of its group size.	
		The number of information disks in a group must not decrease.	
RAID 60	RAID 60	The number of disks in RAID 60 must be a multiple of its group size.	1
		The number of information disks in a group must not decrease.	
	RAID 70	The number of disks in RAID 70 must be at least 12.	1
		The number of disks in RAID 70 must be a multiple of its group size.	
		The number of information disks in a group must not decrease.	
RAID 70	RAID 70	The number of disks in RAID 70 must be a multiple of its group size.	1
		The number of information disks in a group must not decrease.	

## raid restripe

The chapter describes the available subcommands for the restripe operation.

Restripping refers to any change in RAID configuration, such as the position of checksums or data drives, with the aim of changing the RAID level or size.

To start RAID restripping, run

```
# xicli raid restripe start <args>
```

### Arguments for the `restripe start` subcommand

Required arguments

-n	--name	The name of the RAID.
-l	--level	The new level for the RAID.  If you are only increasing the RAID size, enter the current RAID level for this argument.
-gs	--group_size	<b>Only for RAIDs 50, 60, and 70.</b>  The new group size for the RAID.  Possible values: integers from <b>4</b> to <b>32</b> .
-d	--drives	The list of block devices (/dev/sd*, /dev/mapper/mpath*, /dev/nvme*, /dev/dm-*) separated by a space to add to the RAID.

To pause RAID restripping, run

```
# xicli raid restripe stop <arg>
```

### Argument for the `restripe stop` subcommand

Required argument

-n	--name	The name of the RAID.
----	--------	-----------------------

To continue RAID restripping, run

```
# xicli raid restripe continue <arg>
```

## Argument for the `restripe continue` subcommand

### Required argument

---

<code>-n</code>	<code>--name</code>	The name of the RAID.
-----------------	---------------------	-----------------------

## raid resize

The chapter describes the command for the resizing operation.

Following the restriping process, the RAID goes into the 'need\_resize' state. We recommend starting the resizing operation in order to restore optimal system speed.



The `raid resize` command should be executed when there is no workload on the RAID.

```
# xicli raid resize <arg>
```

## Argument for the `resize` subcommand

### Required argument

---

<code>-n</code>	<code>--name</code>	The name of the RAID.
-----------------	---------------------	-----------------------

## Examples of Restriping and Resizing

Example: Restripe the RAID "media5" by adding a new drive `/dev/sdi` without changing the RAID level (increasing the RAID size):

```
# xicli raid restripe start -n media5 -l 5 -d /dev/sdi
```

Example: Increasing the RAID "media5" size by replacing the drives {3 GB each} with drives of larger size {5 GB each: `/dev/sde`, `/dev/sdf`, `/dev/sdg`, `/dev/sdh`}.

In this example, the RAID size will be increased from 9 GB to 15 GB.



The operation of increasing the RAID level by replacing the disks with larger ones is not supported in Xinnor xiRAID 4.0.4. Please refer to the "Xinnor xiRAID 4.0.4 Known Issues" for more information.

Example: Restripping of the “media5” RAID with adding new drives /dev/sdf /dev/sdg /dev/sdh and RAID level changing from 5 to 6:

```
# xicli raid restripe start -n media5 -l 6 -d /dev/sdf /dev/sdg /dev/sdh
```

After restripping is finished, the RAID state is *need\_resize* until you run

```
# xicli raid resize -n <raid_name>
```

## Changing RAID Parameters

To improve the system performance under the workload, try decreasing initialization, reconstruction, or restripping priorities.

See recommendations on configuring RAID parameters in the chapter [RAID and System Setup Recommendations](#).

To change the RAID dynamic parameters, run:

```
# xicli raid modify <arg> [optional_args]
```

### Arguments for the `modify` subcommand

Required argument		
<code>-n</code>	<code>--name</code>	The name of the RAID.  The maximum RAID name length is <b>28</b> characters.
Optional arguments		
<code>-inp</code>	<code>--init_prio</code>	<b>Except RAID 0.</b>  Initialization priority in %.  Possible values are from <b>0</b> to <b>100</b> (maximum rate of initialization).  The default: <b>100</b> .
<code>-mwe</code>	<code>--merge_write_enabled</code>	<b>Except RAIDs 0, 1, 10.</b>  Enable ( <b>1</b> ) or disable ( <b>0</b> ) the Merge function for write operations.  The default: <b>0</b> .
<code>-mre</code>	<code>--merge_read_enabled</code>	<b>Except RAIDs 0, 1, 10.</b>

**Arguments for the `modify` subcommand (continued)**

		Enable (1) or disable (0) the Merge function for read operations. The default: 0.
-ml	--memory_limit	RAM usage limit in MiB. Possible values: 0 and integers from 1024 to 1048576. The 0 value sets unlimited RAM usage. The default: 0 (unlimited).
-mm	--merge_max	<b>Except RAID 0, 1, 10.</b> Maximum wait time (in microseconds) for stripe accumulation for the Merge functions. Possible values: integers from 1 to 100000. The default: 1000.
-mw	--merge_wait	<b>Except RAID 0, 1, 10.</b> Wait time (in microseconds) between requests for the Merge functions. Possible values: integers from 1 to 100000. The value must be less than the merge_max value. The default: 300.
-rcp	--recon_prio	<b>Except RAID 0.</b> Reconstruction priority in %. Possible values: from 0 to 100 (maximum rate of reconstruction). The default: 100.
-re	--resync_enabled	<b>Except RAID 0, 1, 10.</b> Enable (1) or disable (0) the resync function.

**Arguments for the `modify` subcommand (continued)**

		The default: <b>1</b> .
<code>-rl</code>	<code>--request_limit</code>	Number of simultaneous I/O requests on RAID.  Possible values: integers from <b>0</b> to <b>4294967295</b> .  The <b>0</b> value disables the restriction.  The default: <b>0</b> .
<code>-rsp</code>	<code>--restripe_prio</code>	Restripping priority in %.  Possible values are from <b>0</b> to <b>100</b> (maximum rate of restripping).  The default: <b>100</b> .
<code>-se</code>	<code>--sched_enabled</code>	Enable ( <b>1</b> ) or disable ( <b>0</b> ) the scheduling function.  The default: <b>0</b> .
<code>-sp</code>	<code>--sparepool</code>	Name of the spare pool to assign to the RAID.  The <b>null</b> value removes the spare pool from the RAID.  Spare pool can not be assigned to RAID 0.
	<code>--force_online</code>	Change RAID state to online if the RAID has unrecoverable sections.  I/O operations on unrecoverable sections may lead to data corruption.  The argument takes no value.
	<code>--force_resync</code>	Force RAID re-initialization.  The argument takes no value.

Example: Setting reconstruction priority for the RAID "media5" equal to 50%:

```
# xicli raid modify -n media5 -rcp 50
```

# Drives

## Manual Drive Replacement or Excluding

To exclude or replace a drive in a RAID, run:

```
# xicli raid replace <args>
```



If you manually replace a drive that is a part of a spare pool, the drive excludes from the spare pool.

### Arguments for the `replace` subcommand

#### Required arguments

<code>-n</code>	<code>--name</code>	The name of the RAID.
<code>-no</code>	<code>--number</code>	The number of the drive. To find out the number of the drive, use <pre># xicli raid show</pre>
<code>-d</code>	<code>--drive</code>	The new block device. To remove the drive (to mark it as missing) set the <b>null</b> value.

Example: In the RAID "media5", replacing the drive "0" with the drive "nvme4n1":

1. Mark the drive "0" as "missing":

```
# xicli raid replace -n media5 -no 0 -d null
```

2. Replace the drive "0" with the drive "nvme4n1":

```
# xicli raid replace -n media5 -no 0 -d /dev/nvme4n1
```

## Automatic Drive Replacement

A drive can be automatically replaced after it

- was physically removed from a RAID;
- exceeded the threshold value of wear (90%);
- exceeded the threshold value of I/O errors (3).

To automatically replace drives on a RAID, create a spare pool, then assign the created spare pool to the RAID. You can only assign one spare pool to each RAID. We recommend creating a sparepool with storage devices of the same type.

If the system has a spare pool, you can assign it to an existing RAID or when creating a new RAID.

### Commands for managing spare pools

To add drive(s) to the spare pool, run

```
# xicli pool add <args>
```

#### Arguments for the `add` subcommand

Required arguments

<code>-n</code>	<code>--name</code>	The name of the spare pool.
<code>-d</code>	<code>--drives</code>	The list of block devices ( <code>/dev/sd*</code> , <code>/dev/mapper/mpath*</code> , <code>/dev/nvme*</code> , <code>/dev/dm-*</code> ) separated by a space.

To create the spare pool, run

```
# xicli pool create <args>
```

#### Arguments for the `create` subcommand

Required arguments

<code>-n</code>	<code>--name</code>	The name for the spare pool.
-----------------	---------------------	------------------------------



## Arguments for the `create` subcommand (continued)

<code>-d</code>	<code>--drives</code>	The list of block devices ( <code>/dev/sd*</code> , <code>/dev/mapper/mpath*</code> , <code>/dev/nvme*</code> , <code>/dev/dm-*</code> ) separated by a space.
-----------------	-----------------------	--

To delete the spare pool, run

```
# xicli pool delete <arg>
```

## Argument for the `delete` subcommand

Required argument

<code>-n</code>	<code>--name</code>	The name of the spare pool.
-----------------	---------------------	-----------------------------

To remove drive(s) from the spare pool, run

```
# xicli pool remove <args>
```

## Arguments for the `remove` subcommand

Required arguments

<code>-n</code>	<code>--name</code>	The name of the spare pool.
<code>-d</code>	<code>--drives</code>	The list of block devices ( <code>/dev/sd*</code> , <code>/dev/mapper/mpath*</code> , <code>/dev/nvme*</code> , <code>/dev/dm-*</code> ) separated by a space.

To show info on the spare pool, run

```
# xicli pool show [optional_args]
```

## Arguments for the `show` subcommand

Optional arguments

<code>-n</code>	<code>--name</code>	The name of the spare pool. Without the argument, show info on all spare pools.
<code>-f</code>	<code>--format</code>	Output format:

## Arguments for the `show` subcommand (continued)

- **table**;
- **json**;
- **prettyjson** – human-readable json.

The default: **table**.

`-u`

`--units`

Size units:

- **s** – sectors (1 sector=512 bytes);
- **k** – kilobytes;
- **m** – megabytes;
- **g** – gigabytes.

The default: **g**.

SparePools			
name	devices	serials	sizes
pool1	0 /dev/sda <b>ready</b>	drive-scsi0	5 GiB
	1 /dev/sdb <b>ready</b>	drive-scsi1	5 GiB

Possible drive states:

- **ready** – the drive is able for replacement;
- **absent** – drive is missing in the system;
- **failed** – attempt to replace with this drive from the spare pool failed, the drive will not be used for replacement.

To manage delay timer for the drive replacement from the spare pools, run

```
# xicli settings pool modify <arg>
```



When you change any parameter of the `xicli settings pool modify` command, the `xiraid-scanner.service` restarts.

## Argument for the `pool modify` subcommand

### Required argument

---

<code>-rd</code>	<code>--replace_delay</code>	<p>Delay time (in seconds) for the drive replacement from the spare pools.</p> <p>Only one delay time is used for all the spare pools.</p> <p>Possible values: integers from <b>1</b> to <b>3600</b>.</p> <p>The default: <b>180</b>.</p>
------------------	------------------------------	---

To show delay time used for the drive replacement from the spare pools, run

```
# xicli settings pool show
```

## Argument for the `pool show` subcommand

### Optional argument

---

<code>-f</code>	<code>--format</code>	<p>Output format:</p> <ul style="list-style-type: none"> <li>• <b>table</b>;</li> <li>• <b>json</b>;</li> <li>• <b>prettyjson</b> – human-readable json.</li> </ul> <p>The default: <b>table</b>.</p>
-----------------	-----------------------	---

Example: Creating a sparepool “pool1” and assigning it to the RAID “media5”:

1. Create a sparepool:

```
# xicli pool create -n pool1 -d /dev/sda /dev/sdb
```

2. Assign the created sparepool to the RAID:

```
# xicli raid modify -n media5 -sp pool1
```

Example: Setting the replacement timer for the sparepools to 60 seconds:

```
# xicli settings pool modify -rd 60
```

## Drive I/O Error Counter

You can keep track of drives where I/O errors (faults) have started to appear so that you can replace such drives with healthy ones in a timely manner.



We recommend setting up email notifications (to learn more, see the [Setting up Email Notifications](#)) chapter to trace drives with I/O errors.

*Fault threshold* is the common number of faults for each drive, above which the drive will be removed from the RAID or replaced with a suitable drive from the spare pool. You can set the fault threshold value in the range from 1 to 1000. If you change the fault threshold value, the current number of faults on the drives is reset.

When a drive is removed from a RAID *because* the fault threshold is exceeded:

- if the RAID has a SparePool with the suitable drive, *the removed drive* will be replaced and then the RAID reconstruction will start;
- if *the removed drive* has not been replaced in the RAID (automatically or manually), the drive will return in the RAID after resetting the current number of faults on that drive;
- the `drive clean` command applied to *the removed drive* resets the current number of faults and does not remove metadata from the drive.

To manage the threshold value of I/O errors for all drives, run

```
# xicli settings faulty-count modify <arg>
```



When you change any parameter of the `xicli settings faulty-count modify` command, the `xiraid-scanner.service`, `xiraid-version_lock.service`, `xiraid-mail.service` restart.

### Argument for the `faulty-count modify` subcommand

Required argument

---

<code>-t</code>	<code>--threshold</code>	The threshold value for all drives.
-----------------	--------------------------	-------------------------------------

## Argument for the `faulty-count modify` subcommand (continued)

If you set a new fault threshold value, the current numbers of faults are reset for all the drives.

Possible values: integers from **1** to **1000**.

The default: **3**.

Example: Set the drive fault threshold value to 10:

```
# xicli settings faulty-count modify -t 10
```

To show the threshold value of I/O errors, run

```
# xicli settings faulty-count show
```

## Argument for the `faulty-count show` subcommand

Optional argument

`-f`

`--format`

Output format:

- **table**;
- **json**;
- **prettyjson** – human-readable json.

The default: **table**.

To reset the current numbers of faults for drives, run

```
# xicli drive faulty-count reset <arg>
```

The RAID that contains the drive must be loaded.



When you change any parameter of the `xicli drive faulty-count reset` command, the `xiraid-scanner.service` restarts.

## Arguments for the `faulty-count reset` subcommand

Required argument

## Arguments for the `faulty-count reset` subcommand (continued)

<code>-d</code>	<code>--drives</code>	The list of block devices ( <code>/dev/sd*</code> , <code>/dev/mapper/mpath*</code> , <code>/dev/nvme*</code> , <code>/dev/dm-*</code> ) separated by a space to reset their current numbers of faults.
-----------------	-----------------------	---

Example: reset current values of fault count for drives `/dev/sda`, `/dev/sdb`, `/dev/sdd`:

```
# xicli drive faulty-count reset -d /dev/sd[a-b] /dev/sdd
```

To show the current numbers of faults for drives, run

```
# xicli drive faulty-count show [optional_args]
```

## Arguments for the `faulty-count show` subcommand

Mutually exclusive optional arguments

<code>-n</code>	<code>--name</code>	The RAID name for which drives the current number of faults will be shown.  If neither of the two arguments is specified, show the values for all drives.
<code>-d</code>	<code>--drives</code>	The list of block devices ( <code>/dev/sd*</code> , <code>/dev/mapper/mpath*</code> , <code>/dev/nvme*</code> , <code>/dev/dm-*</code> ) separated by a space to show their current numbers of faults.  If neither of the two arguments is specified, show the values for all drives.
Optional argument		
<code>-f</code>	<code>--format</code>	Output format: <ul style="list-style-type: none"> <li>• <b>table</b>;</li> <li>• <b>json</b>;</li> <li>• <b>prettyjson</b> – human-readable json.</li> </ul>


## Arguments for the `faulty-count show` subcommand (continued)

The default: `table`.

Example: show current values of fault count for drives `/dev/sda`, `/dev/sdb`, `/dev/sdd`:

```
# xicli drive faulty-count show -d /dev/sd[a-b] /dev/sdd
```

## Removing Drive Metadata and Resetting Current Error Count

 Warning! The result of the command is irreversible. Read the description carefully.

*Metadata* is Xinnor xiRAID device configuration information (to learn more, see [Configuration Files and Metadata](#)).

The `drive clean` command resets the current error counter value and/or removes metadata from selected disks depending on the status and state of those disks.

The `drive clean` command resets the current error counter and doesn't delete the metadata:

- on a disk that was removed from a RAID due to exceeding the I/O error threshold.

To remove the metadata from such disks, add a new disk to the RAID to replace the removed one.

- on a disk included in a RAID that is present in the current configuration file.

To remove metadata from the drives and/or reset their current error count, run:

```
# xicli drive clean <arg>
```

### Argument for the `clean` subcommand

Required argument

Option	Flag	Description
-d	--drives	The list of block devices ( <code>/dev/sd*</code> , <code>/dev/mapper/mpath*</code> , <code>/dev/nvme*</code> , <code>/dev/dm-*</code> ) separated by a space to

## Argument for the `clean` subcommand (continued)

reset the current fault counter and/  
or delete the metadata.

Example: Deleting metadata from drives `"/dev/nvme5n1"` and `"/dev/nvme1n1"`:

```
# xicli drive clean -d /dev/nvme1n1 /dev/nvme5n1
```

# Configuration Files and Metadata

In this chapter, you can learn:

- how the information about xiRAID RAID objects is stored;
- how to use this information to import a RAID to another system;
- how to restore the configuration if a system disk fails;
- how to restore a RAID.

## xiRAID RAID Configuration

After creating a xiRAID RAID device, the information about the created device automatically goes to the following locations in the system:

- the current configuration file `/etc/xiraid/raid.conf`;
- the metadata on the disks included in the created xiRAID RAID device.

Additionally, before changing the current configuration file, the system automatically creates a backup file of the current configuration `/etc/xiraid/raid.conf.bak`. Thus, if necessary, you can restore the previous version of the device configuration.

## The current configurations file

The current configurations file stores the most recent changes to the xiRAID RAID devices and is used when working with created xiRAID RAID devices.

The example of the configuration file:



```
{
  "raids": {
    "raid_one": {
      "name": "raid_one",
      "uuid": "C2875DCE-971E-4401-9D2C-7EBC36422B6A",
      "level": "0",
      "synd_cnt": 0,
      "strip_size": 16,
      "block_size": 4096,
      "drives": [
        "drive-scsi9"
      ],
      "size": 20578304,
      "group_size": 1
    },
    "raid_two": {
      "name": "raid_two",
      "uuid": "3967F298-893A-45E3-A79F-051FE6C499F7",
      "level": "1",
      "synd_cnt": 0,
      "strip_size": 16,
      "block_size": 4096,
      "drives": [
        "drive-scsi8",
        "drive-scsi7"
      ],
      "size": 20578304,
      "group_size": 2
    }
  },
  "drives": {
    "drive-scsi8": "1",
    "drive-scsi7": "2"
  },
  "faulty_count_threshold": 3,
  "timestamp": 1648628802.796359
}
```

The "raids" object contains the settings of the created xiRAID RAIDs.

The "drives" object contains the serial numbers of the drives with an error count greater than 0.

The "faulty\_count\_threshold" object contains the value of the error threshold value for the drives.

The "timestamp" object contains the creation date of this configuration file in timestamp format.

## Configurations metadata on disks

Since the current configurations file is stored on the system disk, to protect against system disk failure, the configuration information is also stored on the disks that belong to the xiRAID RAID devices.

Each disk contains data that enables a complete copy of the configuration file to be restored.

## Importing RAID

RAID Import is the transfer of RAID-containing disks from one system to another. Once imported, the xiRAID RAID device becomes available for management and its configuration information is added to the current configuration file of the new system. Usually, configuration information from the metadata on the drives is used for the import.

However, when importing a RAID, there may be situations where there is already a RAID with the same name on the system where the RAID is migrated, or block devices with the same serial number are used.

Below in this chapter is a list of commands that enable you to import RAID and resolve possible conflicts, and an example of the import process.

The commands of the RAID import:

- raid import apply
- raid import show

The import commands work with RAIDs that are present in the disk metadata, but are not present in the current configuration file.

See examples of the import process in the chapter [Example of Importing a RAID](#).

## raid import show

To show info about the RAID(s) that can be imported (restored) from the drives, run

```
# xicli raid import show [optional_args]
```

### Arguments for the `import show` subcommand

#### Optional arguments

-d	--drives	<p>The list of block devices (/dev/sd*, /dev/mapper/mpath*, /dev/nvme*, /dev/dm-*) separated by a space to show the info.</p> <p>Without the argument, shows the info from all drives.</p>
-f	--format	<p>Output format:</p> <ul style="list-style-type: none"> <li>• <b>table</b>;</li> <li>• <b>json</b>;</li> <li>• <b>prettyjson</b> – human-readable json.</li> </ul> <p>The default: <b>table</b>.</p>
	--offline	<p>Show non-recoverable RAID(s) in the import list.</p> <p>The argument takes no value.</p>

#### Possible conflicts:

- *name*: Conflict with *in-system RAID(s)*;
- *drives*: Conflict with *in-system RAID(s)*;
- *name*: Conflict with *import RAID(s)*;
- *drives*: Conflict with *import RAID(s)*;
- *name*: Conflict with *in-system and import RAID(s)*;
- *drives*: Conflict with *in-system and import RAID(s)*.

### Drives statuses (messages in the devices row):

- *no\_metadata* – drive has no xiRAID RAID metadata. After drive import, run drive reconstruction.
- *in\_use* – drive is in in-system RAID. After import, the drive will go to the *offline* state.
- *normal* – drive works properly (the state may be changed after import).

Example:

To show all RAIDs that are available for import, run:

```
# xicli raid import show
```

The command will find and display information about founded RAIDs on drives that can be imported. Three RAIDs available for import shown in the figure below:

RAID's uuid	info	devices	serials	import status
073DE372-94D0-43B3-9582-BFC67F3824CC	size: 9 GiB level: 0 synd_cnt: 0 strip_size: 16 block_size: 4096	0 /dev/sdl normal 1 /dev/sdm normal	drive-sca11 drive-sca12	name: OK drives: OK
61BF2210-6F1B-4F76-9F01-02D297E1C715	size: 9 GiB level: 6 synd_cnt: 2 strip_size: 16 block_size: 4096	0 /dev/sdh in_use, no_metadata 1 /dev/sdi normal 2 /dev/sdj normal 3 /dev/sdk normal	drive-sca17 drive-sca18 drive-sca19 drive-sca10	name: Conflict with in-system RAID(s) drives: Conflict with in-system RAID(s) media1
D698D77D-1B10-48D6-95FE-B037E7B0C366	size: 14 GiB level: 7 synd_cnt: 3 strip_size: 16 block_size: 4096	0 /dev/sdb normal 1 /dev/sdc normal 2 null no_metadata 3 null no_metadata 4 /dev/sdf normal 5 /dev/sdg normal	drive-sca1 drive-sca2 null null drive-sca15 drive-sca16	name: Conflict with in-system and import RAID(s) drives: OK

Conflicts are highlighted in red color. The first RAID does not conflict with any RAIDs. The second RAID conflicts with in-system and import RAIDs by the name and disks. The third RAID has a conflict with in-system RAID by the name.

## raid import apply

To import the RAID from disk metadata, run

```
# xicli raid import apply <arg> [optional_arg]
```

### Arguments for the `import apply` subcommand

Required argument

-id	--uuid	UUID of the RAID.
-----	--------	-------------------

Optional argument

-nn	--new_name	The new name for the RAID.
-----	------------	----------------------------

## Example of Importing a RAID

Example: move RAID "media5" to another system that already has a RAID with the same name:

1. Stop the workflow on the RAID.

2. Unload the RAID:

```
# xicli raid unload -n media5
```

3. Remove the RAID disks from this system and insert them into another system.

4. Check for import conflicts:

```
# xicli raid import show
```

5. Import the RAID with changing its name to "media5\_2":

```
# xicli raid import apply -id 52538D69-CF99-4471-8C85-DD42C9026A22 -nn media5_2
```

## Configuration File Recovery


If the system disk fails, you can recover xiRAID RAID objects from the metadata on the disks by using the command

```
# xicli config <subcommand> <args> [optional_args]
```

Subcommands for the `config` command:

<code>apply</code>	Apply the configuration file for all restoring RAIDs.
<code>backup</code>	Save the current configuration file (create the backup file <i>backup_raid.conf</i> at the current directory).
<code>restore</code>	Restore the configuration file from a file or from the drives.
<code>show</code>	Show configuration files stored on the drives.

## config apply

 Warning! The result of the command is irreversible. Read the description carefully.

To apply the configuration file for all restoring RAIDs, run

```
# xicli config apply
```

The command *applies* the current configuration file and restore all RAIDs from the file and delete all xiRAID RAIDs that are not in the file.

## config backup

To save the current configuration file (to create the backup file *backup\_raid.conf* at the current directory), run

```
# xicli config backup
```

## config restore

To restore the configuration file from a file or from the drives, run

```
# xicli config restore <arg>
```

The command *restores* (if missing) or *replaces* the current configuration file from a specified location (from a file or disk metadata), but *does not apply* it.

### Arguments for the `restore` subcommand

Mutually excluded required arguments

-f	--file	A file to restore or replace the configuration file.  If no file is specified, restore or replace from <code>/etc/xi-raid/raid.conf.bak</code> .
-d	--drives	The list of block devices ( <code>/dev/sd*</code> , <code>/dev/mapper/mpath*</code> , <code>/dev/nvme*</code> , <code>/dev/dm-*</code> ) separated by a space to

## Arguments for the `restore` subcommand (continued)

restore one newest configuration file to `/etc/xiraid.conf.drive`.

If no disks are specified, restore one newest configuration file from all disks.

## config show

To show configuration files stored on the disks, run

```
# xicli config show [optional_arg]
```

## Argument for the `show` subcommand

Optional argument

---

<code>-d</code>	<code>--drives</code>	The list of block devices ( <code>/dev/sd*</code> , <code>/dev/mapper/mpath*</code> , <code>/dev/nvme*</code> , <code>/dev/dm-*</code> ) separated by a space.  Without the argument, show from all disks.  The command also shows the newest configuration file from the drives.
-----------------	-----------------------	---

## Examples of Restoring the Current Configuration File

Example: restore the current configuration file from disk metadata.

1. Restore the most recent configuration found among all disks:

```
# xicli config restore -d
```

2. Restore the configuration file from `raid.conf.drive`:

```
# xicli config restore -f /etc/xiraid/raid.conf.drive
```

### 3. Apply the restored configuration:

```
# xicli config apply
```

Example: save a copy of the current configuration file to a flash drive.

#### 1. Change the current directory to flash drive:

```
$ cd /mnt/<device>
```

#### 2. Create a copy of the current configuration file:

```
# xicli config backup
```

## Restoring the RAID

RAIDs can be automatically restored after failures using the configuration specified in the configuration file. The timeout for all drives to be initialized is set to 10 seconds.

However, if the RAID configuration includes drives that are not available in the host system for any reason, one of the following will occur:

- If the number of remaining drives in a RAID is sufficient for its operation, it will be restored to the 'degraded' state. If the missing drives later become available in the host system, they will be added to the RAID automatically.
- If the number of remaining drives in a RAID is insufficient for its proper operation, it will be restored to the 'offline' state. In such case, contact the Xinnor support team at [support@xinnor.io](mailto:support@xinnor.io).
- If none of the drives specified in a RAID configuration are available in the host system, the RAID will not be restored.

The log message may include the following warning if not all drives included in the RAID configuration have been initialized:

```
ERROR :: For RAID 'raidname' not all devices were initialized.  
Uninitialized devices are 'drive-serial1, drive-serial2'.
```

You can restore RAID's manually from the current configuration file using

```
# xicli raid restore <arg>
```





Do not use this command if none of the drives included in the RAID configuration are available in the host system.

## Arguments for the `restore` subcommand

Mutually exclusive required arguments

<code>-n</code>	<code>--name</code>	The name of the RAID.
<code>-a</code>	<code>--all</code>	Restore all available xiRAID RAIDs. Argument takes no value.

## Scanning RAIDs and Drives, LED Indication

The *scanner* mechanism is used to monitor the statuses of RAIDs and drives:

- RAID statuses;
- RAID drives;
- automatic drive LED indication.

To manually control the LED indication, use the `drive locate` command (see description and examples below in this chapter).

### The scanner command

The service, which runs in the system, automatically manages the indication of failed and working drives and creates corresponding messages in the log file.

To manage RAIDs monitoring, the LED indication and drive SMART settings, run

```
# xicli settings scanner modify <args>
```



When you change any parameter of the `xicli settings scanner modify` command, the `xiraid-scanner.service` restarts.

## Arguments for the `scanner modify` subcommand

At least one argument is required

<code>-pi</code>	<code>--polling_interval</code>	<p>The polling interval for xiRAID RAIDs and drives in seconds.</p> <p>The parameter affects the auto-start delay for the RAID initialization, reconstruction, and restriping.</p> <p>Possible values: integers from <b>1</b> to <b>3600</b> (1 hour).</p> <p>The default: <b>1</b>.</p>
<code>-spi</code>	<code>--smart_polling_interval</code>	<p>S.M.A.R.T. drive health polling interval, in seconds.</p> <p>Possible values: integers from <b>60</b> to <b>86400</b> (24 hours).</p> <p>The default: <b>86400</b>.</p>
<code>-le</code>	<code>--led_enabled</code>	<p>Enable (<b>1</b>) or disable (<b>0</b>) the automatic LED indication of drives in the system.</p> <p>The default: <b>1</b>.</p> <p>The argument doesn't affect manual LED indication.</p>

Example: Disable automatic LED indication of drives:

```
# xicli settings scanner modify --led_enabled 0
```

To show the settings of the LED indication and drive scanner, run

```
# xicli settings scanner show
```

## Argument for the `scanner show` subcommand

Optional argument

<code>-f</code>	<code>--format</code>	Output format:
-----------------	-----------------------	----------------

## Argument for the `scanner show` subcommand (continued)

- **table**;
- **json**;
- **prettyjson** – human-readable json.

The default: **table**.

## The `locate` command

You can manually control the LED indication of the drives by using the command

```
# xicli drive locate <arg>
```

## Argument for the `locate` subcommand

Required argument

#d	##drives	
		The list of block devices ( <code>/dev/sd*</code> , <code>/dev/mapper/mpath*</code> , <code>/dev/nvme*</code> , <code>/dev/dm-*</code> ) separated by a space to switch the indication on, or switch the indication off (with the <b>null</b> value).
		The argument doesn't affect the automatic indication.

Example: Turn location indication on for drives `"/dev/nvme0n1"` and `"/dev/nvme1n1"`:

```
# xicli drive locate -d /dev/nvme0n1 /dev/nvme1n1
```

## Update Check Service

Update Check Service allows you to control the process of updating Xinnor xiRAID packages to newer versions. After installing Xinnor xiRAID 4.0.4 or updating Xinnor xiRAID products, this service locks the current version of the packages, preventing it from being automatically updated on general system update commands (`apt/yum update`).

To check for an available update, run:

```
# xicli update check
```

Disable the Update Check Service to update Xinnor xiRAID 4.0.4 to a new available version.



Please, follow the instructions provided at [xinnor.io](https://xinnor.io) to safely update your Xinnor xiRAID. Ignoring these steps may result in filesystem panick and even data loss.

To disable the Update Check Service, run:

```
# xicli update prepare
```



Please, do not run this command unless there is a new available Xinnor xiRAID version. Otherwise, the proper functioning of Xinnor xiRAID cannot be guaranteed.



The Update Check Service will inform you of any detected mounted xiRAID devices. Please, unmount the devices before continuing the update process.

We recommend setting up email notifications using [Xinnor xiRAID Administrator's Guide](#), so we could inform you about the latest Xinnor xiRAID releases. The notifications will be sent to you once every three days. The corresponding messages will be created in the log file regardless of your notifications settings.

## Notifications

### System Log

Logs contain information about the system status and system operations at specific points in time. The logs are recorded in the system log (journalctl).

You can:

- manage and view the type of system messages that will be logged;
- collect logs into a file;
- view the latest error messages on the system.

To configure the type of system messages that will be added to the system log, run

```
# xicli settings log modify <arg>
```

### Argument for the `log modify` subcommand

Required argument

Option	Argument	Description
-l	--level	The type of system messages that will be added to the system log.  Possible values: <b>error</b> , <b>warning</b> , <b>info</b> , <b>debug</b> .  Each next type includes the previous one.  The default: <b>debug</b> .

To collect all logs into a file, run

```
# xicli log collect
```

The log file will be available in the `/tmp` directory.

After collecting is complete, the message “*xiRAID logs have been collected and saved in: /tmp/xiraid\_logs\_\${hostname}\_YYYY.MM.DD\_hh-mm-ss.tar.gz*” shows.

To see the selected type of system messages for the system log, run

```
# xicli settings log show [optional_arg]
```

### Argument for the `log show` subcommand

Optional argument

Option	Argument	Description
-f	--format	Output format:

## Argument for the `log show` subcommand (continued)

- **table**;
- **json**;
- **prettyjson** – human-readable json.

The default: **table**.

To see the latest error messages on the system, run

```
# xicli log show [optional_arg]
```

## Argument for the `show` subcommand

Optional argument

<code>-l</code>	<code>--lines</code>	The number of error messages in the event log to show, starting from the last entry.  Possible values: integers from <b>1</b> to <b>1000</b> .  The default: <b>10</b> .
-----------------	----------------------	--

## Setting up Email Notifications



Make sure the system has configured MTA (Mail Transfer Agent) (for example, Postfix).

To set the receiver's email and the notification level, run

```
# xicli mail add <args>
```



When you change any parameter of the `xicli mail add` command, the `xiraid-mail.service` restarts.

## Arguments for the `add` subcommand

### Required arguments

-a	--address	Receiver's email.
-l	--level	The notification level.  Possible values: <ul style="list-style-type: none"> <li>• <b>info</b> – Info notifications;</li> <li>• <b>warning</b> – Error and Warning notifications;</li> <li>• <b>error</b> – Error notifications.</li> </ul>

### Notification types:

Info	Warning	Error
Initialization completed on RAID (...)	Initialization not completed on RAID (...)	RAID (...) is offline now
Initialization progress on RAID (...) is (...) percent	RAID (...) is read-only now	RAID (...) is unrecovered now
Initialization started on RAID (...)	System is up after reboot/crash	After reboot/crash, RAID (...) not restored
RAID (...) is healthy now	Reconstruction not completed on RAID (...)	After reboot/crash, RAID (...) has restored in read only mode
RAID (...) is online now	SparePool (...) ran out of drives	RAID (...) is degraded now
Reconstruction completed on RAID (...)	The number of errors on the bdev (...) is increased. The current number is (...)	After reboot/crash, RAID (...) has restored in offline state
Reconstruction progress on RAID (...) is (...) percent		xiRAID license expired

**Notification types: (continued)**

Info	Warning	Error
Reconstruction started on RAID (...)		xiRAID license error: number of disks in use exceeds the allowed disks number
Drive (...) was returned to RAID (...)		Drive (...) in RAID (...) is offline now
Drive (...) from SparePool (...) was reconnected		Drive (...) from SparePool (...) was disconnected
Drive (...) in RAID (...) was automatically replaced with drive (...) from SparePool (...)		Could not automatically replace drive (...) in RAID (...) with drive (...) from SparePool (...)
Can't replace the faulty bdev (...). Replacing (...) with null		Could not automatically replace drive (...) in RAID (...) because there was no suitable drive in SparePool (...)
		The number of faults on the bdev (...) reached the fault threshold
		The bdev (...) has critical wear out (...)%

Example: Add the receiver with the "user2@email.com" email for all notification types:

```
# xicli mail add -a user2@email.com -l info
```

To remove the email from the list of email notifications, run

```
# xicli mail remove <arg>
```



When you change any parameter of the `xicli mail remove` command, the `xiraid-mail.service` restarts.

**Argument for the `remove` subcommand**

Required argument



## Argument for the `remove` subcommand (continued)

<code>-a</code>	<code>--address</code>	The email address to remove from the notifications.
-----------------	------------------------	---

To show the list of the email notifications, run

```
# xicli mail show
```

## Argument for the `mail show` subcommand

Optional argument

<code>-f</code>	<code>--format</code>	Output format:
		<ul style="list-style-type: none"> <li>• <b>table</b>;</li> <li>• <b>json</b>;</li> <li>• <b>prettyjson</b> – human-readable json.</li> </ul>

The default: **table**.

To manage email notification settings, run

```
# xicli settings mail modify <args>
```



When you change any parameter of the `xicli settings mail modify` command, the `xiraid-mail.service` restarts.

## Arguments for the `mail modify` subcommand

At least one argument is required

<code>-pi</code>	<code>--polling_interval</code>	<p>The polling interval for xiRAID RAIDs and the drives in seconds.</p> <p>Possible values: integers from <b>0</b> to <b>86400</b> (24 hours).</p> <p>The default: <b>10</b>.</p>
<code>-ppi</code>	<code>--progress_polling_interval</code>	<p>Polling interval for the progress of initialization and reconstruction, in minutes.</p>

**Arguments for the `mail modify` subcommand (continued)**

Possible values: integers from **0** to **1440** (24 hours).

The default: **10**.

To show email notification settings, run

```
# xicli settings mail show
```

**Argument for the `mail show` subcommand**

Optional argument

-f	--format	Output format:
		<ul style="list-style-type: none"> <li>• <b>table</b>;</li> <li>• <b>json</b>;</li> <li>• <b>prettyjson</b> – human-readable json.</li> </ul>

The default: **table**.

## Number of CPU Threads of the `xiraid` Module

There are two ways to control the number of CPU threads for `xiraid`:

- the `cpuignore` command

Specifics:

- possible to select specific CPUs for the limitation;
- does not require restarting the module;
- the limits set by this command are reset after the system or module restart.

- the `modprobe` command

Specifics:

- not possible to select specific CPUs. For limitation, OS starts to assign CPUs with the first CPU (ID 0);
- requires restarting the module;
- possible to configure the module to load with the restriction parameters at OS startup.

To select the CPUs that will not be used for the `xiraid` module, run

```
# xicli settings cpu-ignore modify <arg>
```

### Argument for the `cpu-ignore modify` subcommand

Required argument

`--id`

The list of CPU IDs (separated by a comma or a hyphen) that will not be used for `xiraid`.

The **null** value removes the restriction on using threads for `xiraid`.

Example: Disable CPUs with the IDs 0, 1, 2, 3, 6 for the module:

```
# xicli settings cpu-ignore modify --id 0-3,6
```

Example: Remove all restrictions on CPUs for the module:

```
# xicli settings cpu-ignore modify --id null
```

To show the list of CPUs that are not used for the `xiraid` module, run

```
# xicli settings cpu-ignore show
```

### Argument for the `cpu-ignore show` subcommand

Optional argument

`-f`

`--format`

Output format:

- **table**;
- **json**;
- **prettyjson** – human-readable json.

## Argument for the `cpu-ignore show` subcommand (continued)

The default: **table**.

To limit the number of threads for the `xiraid` module:

1. Unload the module from the kernel:

```
# rmmod xiraid
```

2. Load the module in the kernel with the `cpu_cnt` parameter:

```
# modprobe xiraid cpu_cnt=<cnt>
```

where `<cnt>` is the number of CPUs used by the module. The operating system assigns CPUs starting with the very first CPU ID.

3. Restart `xiraid` target:

```
# systemctl restart xiraid.target
```

To configure the `xiraid` module to load automatically with CPU thread parameters, in the `modprobe.conf` file (file name and location depend on your OS, see `man modprobe.conf` for details) add the line

```
options xiraid cpu_cnt=<cnt>
```

where `<cnt>` is the number of CPUs used by the module. The operating system assigns CPUs starting with the very first CPU ID.

# General Configuration Recommendations

## RAID Creation

Next recommendations are appropriate and depend on drives' parameters and vendors.

- The appropriate RAID level depends on the required availability level.

Level of availability as high as 99.999% can be achieved by using RAID 6 if the RAID consists of less than 20 drives. Use RAID 7.3 with more than 20 drives.

Level of availability as high as 99.999% can be achieved by using RAID 50 if the RAID consists of less than 16 drives. With more drives, use RAID 60 or RAID 70.

- The recommended stripe size for the xiRAID RAID is **16** KiB (set by default).

## Using NVMe-oF devices to create a RAID.

1. Xinnor xiRAID allows using NVMe-oF devices to create a RAID. Set the `--ctrl-loss-tmo` parameter to **0** to prevent command freezing because of connection loss when using these devices. It is relevant to `nvme-cli` version `>= 1.4`.

```
# nvme connect -t rdma -n nqn.Xinnor12_1 -a 10.30.0.12 -s 4420
--ctrl-loss-tmo=0
```

2. At the creation of NVMe-oF target for xiRAID RAID, you can enable Merge if the access pattern assumably will be sequential write.

Depending on the version of Linux Kernel or Mellanox drivers, NVMe-oF targets may split big requests to 32 KiB + the rest. This kind of behavior leads to constant *read-modify-writes*. For an SPDK NVMe-oF target, set the `InCapsuleDataSize` parameter denoting at by what value requests should be split.

## RAID and System Setup Recommendations

### `--init_prio`

Syndrome RAID creation starts the initialization process automatically. During it, RAID is available for reading and writing operations. Since initialization priority by default is set to **100**, you can wait until the initialization is finished, or if the access pattern is *not random write*, you can lower the initialization priority. Therefore, user I/O will be processed faster due to the reduction of initialization requests. If the initialization priority is set to **0**, initialization requests are not created during user I/O.

## --recon\_prio

The reconstruction process starts automatically. By default, reconstruction priority equals to **100**, which means reconstruction has maximum priority among other processes. Setting the priority to **0** allows the user I/O processes running before the reconstruction process.

## --restripe\_prio

The `modify` command enables to change restriping priority. If the priority value of the function is zero, restriping starts and continues only if there is no workload. By default, priority is set to **100%** that stands for the highest possible rate of the restriping process. To improve the system workload performance, try decreasing restripe priority.

## --sched\_enabled

There are 2 possible ways of handling an incoming request:

- continue execution on the current CPU;
- transfer the request to the other CPU core and continue execution. Note that it takes time for the transferring.

If the access pattern uses less than half of the system CPU, it is efficient to use the `--sched_enabled` parameter. When a lot of requests are processed by the single CPU core, enabling scheduling allows to redistribute the workload equally between all system CPUs. On multithreading access patterns, scheduling is inefficient, because useless transfer of requests from one CPU core to another wastes time.



Enable Scheduling when the access pattern is low threaded.

## --merge\_write\_enabled

The `--merge_write_enabled` improves the system workload performance when access pattern is sequential and high threaded, and the block sizes are small. This parameter sets a waiting time for all incoming requests in sequential areas.

During waiting time, requests to this area are not intentionally transferred to the drives. Instead of immediate data transfer, incoming requests are formed into a tree structure. At the end of waiting time, requests are merged together if possible. This function reduces the number of read-modify-write operations on syndrome RAIDs. Despite the extra waiting time, this function can improve the system workload performance. If the access pattern is mainly random or queue depth is small, the waiting time will not allow merging requests. In this case enabling `--merge_write_enabled` will decrease the system workload performance.



Enable merge by the `--merge_write_enabled` parameter when the access pattern is sequential and high threaded and the block sizes are small.

Since the time between incoming I/O depends on the workload intensity, size, and other parameters, it may be necessary to change `--merge_wait` and `--merge_max` parameters for better query consolidation. Usually, large I/O sizes require large values for these parameters.

The function only works when the condition is met:

```
data_drives * stripe_size ≤ 1024
```

where

- “data\_drives” is a number of drives in the RAID (for RAIDs 5, 6, 7.3 or N+M) or in one RAID group (for RAIDs 50, 60, or 70) that are *dedicated for data*;
- “stripe\_size” is a selected stripe size for the RAID (**stripe\_size** value) in KiB.

The “data\_drives” value depending on a RAID level:

RAID level	Value of data_drives
RAID 5	Number of RAID drives <i>minus</i> 1
RAID 6	Number of RAID drives <i>minus</i> 2
RAID 7.3	Number of RAID drives <i>minus</i> 3
RAID N+M	Number of RAID drives <i>minus</i> M
RAID 50	Number of drives in one RAID group <i>minus</i> 1

RAID level	Value of data_drives
RAID 60	Number of drives in one RAID group <i>minus</i> 2
RAID 70	Number of drives in one RAID group <i>minus</i> 3

Deactivate Merge when queue depth of user's workload is not enough to merge a full stripe. Activate Merge, if

```
iodepth * block_size >= data_drives * stripe_size
```

where “block\_size” is a block size of the RAID (the **block\_size** value in RAID parameters) in KiB.

## --request\_limit

This parameter limits the number of incoming requests per RAID. For example, writing files with a file system without synchronization.



To improve system workload performance, we recommend enabling the limit on the number of incoming requests by the `--request_limit` parameter when you are working with file system and the buffered writing is performed.

## --force\_online

If a RAID has unrecoverable sections, then the RAID becomes unreadable (get in the offline, unrecoverable state). To try to read available data, manually turn on the online mode for the RAID by running the command

```
# xicli raid modify -n <raid_name> --force_online
```

While in the mode, I/O operations on unrecoverable sections of the RAID may lead to data corruption.

## --resync\_enabled

The function starts a RAID re-initialization after an unsafe system shutdown, thereby protecting syndromic RAIDs (all but RAID 0, RAID 1, and RAID 10) from data loss caused by a write hole.



To disable resync for all RAIDs, run

```
# modprobe xiraid resync=0
```

## Strip Size

Recommended RAID strip size is **16** KiB.

## RAM Limit

Current memory usage is being monitored and controlled to be within the limit. You can modify the `--memory-limit` parameter at any time. By default, memory usage is unlimited.

Deactivating monitoring of current memory usage and limitation control can improve system workload performance. Set `--memory-limit` to **0** to deactivate monitoring with the `modify` command.

If it is necessary to limit the use of RAM, we recommend choosing amount of RAM depending on the selected strip size for the RAD:

Strip size, in KiB	Amount of RAM, in MiB
16	2048
32	2048
64	4096
128	8192
256	16384

## NUMA

1. Create a RAID out of drives belonging to the same NUMA node, if your systems are multiprocessor.

To figure out the NUMA node drive, run:

```
# cat /sys/block/nvme0n1/device/device/numa_node
```

or via `lspci`:

```
# lspci -vvv
```

2. At creation of NVMe-oF target for xiRAID RAID, you can use network adapter of the same NUMA node as NVMe drives.

## System

1. xiRAID shows better performance with enabled hyper-threading (HT).

To find out if there is HT support on the CPU, run

```
# cat /proc/cpuinfo | grep ht
```

In the flags field, check for the ht flag.

Command output example:

```
flags : fpu vme de pse tsc msr pae mce cx8 apic sep mtrr pge mca cmov pat  
pse36 clflush mmx fxsr sse sse2 ss ht syscall nx pdpe1gb rdtscp lm con-  
stant_tsc arch_perfmon rep_good noopl xtopology cpuid tsc_known_freq pni  
pclmulqdq vmx ssse3 fma cx16 pcid sse4_1 sse4_2 x2apic movbe popcnt tsc_  
deadline_timer aes xsave avx f16c rdrand hypervisor lahf_lm abm 3dnow-  
prefetch cpuid_fault invpcid_single pti tpr_shadow vnmi flexpriority ept vpid  
ept_ad fsgsbase tsc_adjust bmi1 hle avx2 smep bmi2 erms invpcid rtm rdseed  
adx smap xsaveopt arat umip arch_capabilities
```

To check if HT is enabled, run

```
# lscpu
```

If Thread(s) per core is 1, then HT is off. HT can be enabled in BIOS/UEFI.

Command output example:

```
Architecture:                x86_64  
CPU op-mode(s):              32-bit, 64-bit  
Byte Order:                   Little Endian  
Address sizes:                40 bits physical, 48 bits virtual  
CPU(s):                       4  
On-line CPU(s) list:         0-3
```

```
Thread(s) per core:      1
...
```

2. The tuned-adm profile set to **throughput-performance** provides better performance on most of the tests:

```
# tuned-adm profile throughput-performance
```

## Workload

In Xinnor xiRAID 4.0.4, user I/O tends to be executed on the same CPU on which the user sent them. However, for some access patterns, you can transfer I/O commands to other CPUs, so the commands will not idle. You can enable I/O Scheduling to all system CPU using a parameter `--sched-enabled` (**1** – activated, **0** – deactivated).

Activating and deactivating the Scheduling function depending on the access pattern recommendations are provided in the **--shed\_enabled** section.

## Swap File

On high load servers, we recommend to disable swap file usage to increase server performance.

## File System Mounting Examples

Since the system restores xiRAID RAIDs after loading an appropriate Linux core and sending a RAID-restore command, to perform automatic mounting at system startup of file systems for these RAIDs, use one of the following instructions.



To set up automatic mounting at system startup, we recommend using `systemd.mount`.

### **systemd.mount**

Example: mounting xfs located on a RAID `/dev/xi_raidname` into `/mnt/raid/` through `systemd.mount`:

## 1. Create a file at `/etc/systemd/system/` with the mount options.

The file name must match the path of the mount directory with "/" replaced by "-" (for example, for `/mnt/raid` the file name will be "mnt-raid.mount").

*The example file `/etc/systemd/system/mnt-raid.mount`*

[Unit]

Description=Mount filesystem on Xinnor xiRAID

DefaultDependencies=no

Before=umount.target

Conflicts=umount.target

[Mount]

What=/dev/xi\_raidname

Where=/mnt/raid/

Options=defaults

Type=xf

[Install]

WantedBy=multi-user.target

## 2. Run the command

```
# systemctl daemon-reload
```

Enable automatic mounting at system startup:

```
# systemctl enable mnt-raid.mount
```

Start the service to mount the file system:

```
# systemctl start mnt-raid.mount
```

## **/etc/fstab**



Make sure to only use the block device name, such as `/dev/xi_raidname`, and not the UUID or LABEL.

Example: mounting xfs located on a RAID `/dev/xi_raidname` into `/mnt/raid/` through `/etc/fstab`:

The string from the file `/etc/fstab`

```
/dev/xi_raidname /mnt/raid/ xfs defaults 0 0
```

## DKMS Specifics when Updating Linux Kernel



If you downgrade kernel version, DKMS functionality depends on the specific distribution.

The *xiraid* kernel module uses DKMS (Dynamic Kernel Module Support) technology and is automatically built and is installed for the Linux kernel versions, listed in the document Xinnor xiRAID 4.0.4 System Requirements, of the different patch versions (without kernel API or ABI changes).

For example:

- 3.10.0-**1062**.el7.x86\_64 >> 3.10.0-**1127**.el7.x86\_64;
- 4.15.0-**112**-generic >> 4.15.0-**124**-generic.

Notice, that if you update the kernel more than the patch update (with kernel API or ABI changes), the *xiraid* kernel module will not be loaded. For example these updates will not work:

- **3**.10.0-1062.el7.x86\_64 >> **4**.18.0-193.el8.x86\_64;
- 4.**15**.0-112-generic >> 4.**18**.0-13-generic;
- **4**.15.0-112-generic >> **5**.4.0-26-generic.

To update (or change) a Linux kernel version with the installed *xiraid* module with DKMS, the OS must have a package with the header files for the kernel version to be updated:

- kernel-devel (for RHEL and RHEL-based systems);
- kernel-uek-devel (for Oracle Linux);
- linux-headers (for Ubuntu);
- pve-headers (for Proxmox).

Since some OS distributions do not have by default a package with header files (and also some repositories may not have package versions for out-of-date kernel versions), we recommend to install a package with header files for a new kernel version manually before (or simultaneously with) installing a new kernel version (see examples of commands to install packages with headers for different operating systems in the Xinnor xiRAID 4.0.4 Installation Guide.)

For example, on Ubuntu 20.04, install the `linux-image` package at the same time as the `linux-headers` package:

```
# apt install linux-image-5.15.0-27-generic linux-headers-5.15.0-27-generic
```

## Authenticating gRPC Client

To authenticate a gRPC client, set up a host and/or a port for the connection, and replace a TLS/SSL public key certificate if necessary. For gRPC connections, only server needs to provide its certificate to client (the server-side TLS connection type).

To set up client-server connection settings, run the command



When you change any parameter of the `xicli settings auth modify` command, the `xiraid-target.service` restarts. Additionally, it will cause all RAIDs to unload. Please, run this command only after stopping all mounted devices.



The command neither requires acceptance of the EULA nor running `xiraid.target` service.

```
# xicli settings auth modify <args>
```

### Arguments for the `auth modify` subcommand

At least one argument is required

---

`--host`

The host name or IP address that will be used for the connection.

After changing the host, you must regenerate and replace the certificate.

The default: **localhost**.

---

## Arguments for the `auth modify` subcommand (continued)

<code>--port</code>	The port that will be used for the connection. The default: <b>6066</b> .
---------------------	--

To view client-server connection settings, run the command


```
# xicli settings auth show
```

## Argument for the `auth show` subcommand

Optional argument

<code>-f</code>	<code>--format</code>	Output format:
		<ul style="list-style-type: none"> <li>• <b>table</b>;</li> <li>• <b>json</b>;</li> <li>• <b>prettyjson</b> – human-readable json.</li> </ul>
		The default: <b>table</b> .

To replace the certificate:

 The certificates require that the system time is not earlier than June 24, 2019.

1. copy the files to `/etc/xraid/crt/` that are strictly named

- `server-key.key` (or `server-key.pem`);
- `server-cert.crt` (or `server-cert.pem`);
- `ca-cert.crt` (or `ca-cert.pem`);

If there are `.pem` and `.key/.crt` files in `/etc/xraid/crt/` at the same time, the system will use `.key/.crt` files.

2. restart the xiraid service:

```
# systemctl restart xiraid.target
```

# Troubleshooting

**1. When attempting to use `xicli raid` commands, the following error occurs:**

**Error: Missing xiRAID RAID system module**

Possible reasons:

- After updating the OS kernel, the packages with the header files (kernel-devel, kernel-uek-devel, linux-headers, pve-headers) remain from the previous kernel version.
- Linux kernel update that is more than the patch update.

Solutions:

- Update or install the package with the kernel header files (kernel-devel, kernel-uek-devel, linux-headers, pve-headers) for the updated or installed OS kernel version. See the Xinnor xiRAID 4.0.4 Installation Guide for details.

After that, run the command

```
# dkms autoinstall
```

- Load on the Linux kernel version that was before the kernel update.

**2. When attempting to use some `xicli` commands, the following error occurs:**

**Error: failed to connect to all addresses**

Possible reason:

The client attempted to access the server using incorrect credentials, incorrect SSL certificates, or the old address that had been changed.

Solutions:



- Restart xiraid target:

```
# systemctl restart xiraid.target
```

- If the previous step did not help, update the certificates.

### 3. When attempting to use some `xicli` commands, the following error occurs in the log files:

#### **Error: Descriptors cannot not be created directly.**

Possible reason:

The Protobuf library for Python had been installed via pip prior to the installation of Xinnor xiRAID 4.0.4.

Solution:

- Uninstall Xinnor xiRAID 4.0.4 using Xinnor xiRAID 4.0.4 Installation Guide.
- Uninstall the Protobuf library for Python.
- Install Xinnor xiRAID 4.0.4.