



## **xiRAID Classic 4.3.0 Release Notes**

# Contents

<b>Release Notes.....</b>	<b>3</b>
Changes to RAID Reinitialization.....	3
Support Added for New Versions of RHEL and Proxmox.....	3
Changes to Trial Licenses.....	3
Performing TRIM on RAID Drives.....	4
Spare Pools in High Availability Clusters.....	4
Sending Test Email Notifications.....	4
Initialization and Reconstruction Priority.....	4
Scanning for RAID Silent Data Corruption .....	5
Known Issues.....	5

# Release Notes

## Changes to RAID Reinitialization

The logic for automatic RAID reinitialization has changed. Previously, when activating a RAID, if we found that the RAID was not correctly unloaded, we would perform a full recalculation and rewrite the syndromes in the RAID. In the current version, we check the syndromes in the stripes and only rewrite them if they are incorrect in a specific stripe, logging this event. This change speeds up the check (since reading from drives is generally faster than writing) and reduces drive wear.

Additionally, the `--force_resync` flag for the `raid modify` command is now deprecated and will be removed in a future release. To reset the initialization counter, set the state of the RAID to `need_init`, and start RAID reinitialization, use the `raid init reset` command.

## Support Added for New Versions of RHEL and Proxmox

The following Linux distributions are now supported:

- RHEL, Rocky Linux, and Alma Linux versions 9.6 and 10
- Proxmox versions 8.3 and 8.4

## Changes to Trial Licenses

xiRAID Classic will no longer include a trial license with the installation. To request a trial license, please do not hesitate to contact us at [support@xinnor.io](mailto:support@xinnor.io).

## Performing TRIM on RAID Drives

xiRAID Classic can now perform the TRIM procedure on drives during the RAID creation process. The TRIM procedure will automatically run on RAID drives if the following conditions are met:

- All drives support Deterministic Read Zero after TRIM (RZAT).
- None of the drives contain any metadata. This check ensures that no data is accidentally deleted. If a drive contains metadata, performing the TRIM procedure will make it impossible to restore the data from that drive. This condition helps prevent unintentional data loss.
- The `--no_trim` option is not specified.

For more details, refer to [Creating a RAID](#).

## Spare Pools in High Availability Clusters

Spare pools are now supported in configurations where xiRAID Classic is integrated into a Pacemaker cluster. For more details, refer to [Managing Spare Pools in a Cluster](#).

## Sending Test Email Notifications

Administrators can now send a test email notification to configured notification recipients via the `mail send` command:

```
xicli mail send
```

This may be helpful after configuring a Mail Transfer Agent to ensure the MTA is functioning correctly and that all recipients are receiving messages as expected.

## Initialization and Reconstruction Priority

The default values for initialization priority (`--init_prio`) and reconstruction priority (`--recon_prio`) have been changed from 100% to 50%. This change will not have a noticeable impact on the performance of initialization and reconstruction.

For instructions on how to change these values, refer to Changing RAID Parameters.

## Scanning for RAID Silent Data Corruption

Users can now scan xiRAID Classic RAIDs for Silent Data Corruption (SDC). Running an SDC scan is recommended when you suspect or want to prevent data integrity issues that might be caused by undetected corruption in the storage array. Silent data corruption can occur without any obvious signs or warnings, meaning the data is corrupted but the system or xiRAID Classic itself doesn't show errors, potentially leading to serious issues when the data is accessed later. For instructions on how to run an SDC scan, see [#unique\\_18](#).

## Known Issues

### 1. `xicli raid show`, `xicli raid unload` and `xicli raid destroy` behavior after destroying or unloading a RAID

When the commands `xicli raid show`, `xicli raid unload`, or `xicli raid destroy` are issued shortly after a `xicli raid destroy` or `xicli raid unload` command, they may fail because the initial destroy/unload process is still completing. This does not lead to data corruption or any other issues. Simply repeating the commands will ensure they are successful.

### 2. Disabling CPU Cores

Do not disable CPU cores on a system with xiRAID Classic as this not supported.

### 3. GPG errors after running apt update

On Ubuntu and Proxmox, you may encounter errors related to the expiry of the `pkg.xinnor.io` GPG keys (the next expiry is on November 26, 2029) when running `apt update` or updating from xiRAID Classic 4.2.0. For instructions on how to resolve these errors, refer to [#unique\\_20](#).

### 4. Information about device wear cannot be displayed for SATA and SAS drives

The wear value in the output of the `xicli raid show -e` command will likely show N/A for SATA and SAS devices. At the moment, information about device wear is available only for NVMe drives.

## 5. Unavailable drives in spare pools

If xiRAID Classic attempts to use a drive from a spare pool and the drive is unavailable, it will be permanently removed from the spare pool. The drive will not be automatically added back when it becomes available again. To use the drive in the spare pool again, you must manually add it back.

## 6. Issues with kernel updates on Proxmox

A bug in DKMS versions 3.0.9 and 3.0.10 prevents the proper installation of new kernel packages on Proxmox systems with xiRAID Classic. This issue occurs if your system has DKMS modules that use the `BUILD_EXCLUSIVE_KERNEL` parameter.

To work around this issue, users should disable DKMS autoinstall globally on their system by creating a special marker file:

```
touch /etc/dkms/no-autoinstall
```

This will cause future kernel package installations to skip the process of building DKMS modules, including the xiRAID Classic kernel module. As a result, to keep xiRAID Classic operational on new kernels, the DKMS autoinstall procedure must be run manually each time a new kernel is installed. For example:

```
dkms autoinstall --kernelver 6.8.12-7-pve
```

In the example above, `6.8.12-7` is the version of the recently installed kernel update.

## 7. Increasing the size of a RAID by replacing its drives with larger ones can introduce defects that may affect the system's functionality

A critical defect has been discovered in all xiRAID Classic 4.0.x versions related to the feature of increasing the size of a RAID by replacing its drives with larger ones (vertical scaling RAID operation). It is strongly advised not to use this feature. This bug will be addressed and fixed in the upcoming release.

This defect does not affect the following possibilities:

- Changing the RAID level with the addition of new drives.
- Increasing the size of a RAID by adding new drives.