

An Introduction to EMC CLARiiON CX4 Disk-Drive Spin Down Technology

Applied Technology

Abstract

This white paper describes EMC® CLARiiON® disk-drive Spin Down technology available on the CX4 series. It describes the implementation of the disk-drive Spin Down feature and provides details for using this feature through Navisphere® Manager and CLI. Some best practices are also provided, along with charts showing how much power can be saved with this technology.

October 2009

Copyright © 2009 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS IS.” EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see EMC Corporation Trademarks on EMC.com

All other trademarks used herein are the property of their respective owners.

Part Number h6632

Table of Contents

Executive summary	4
Introduction	4
Audience	4
Terminology	4
Power savings with disk-drive Spin Down	5
Criteria to enter standby mode	5
Criteria to exit standby mode	5
Drive spin-up algorithm	6
Different ways to manage disk-drive Spin Down	6
Managing disk-drive Spin Down at the storage-system level	6
Using Navisphere Manager to manage disk-drive Spin Down.....	6
Using Navisphere CLI commands to manage disk-drive Spin Down.....	7
Managing disk-drive Spin Down at the RAID group level	8
Using Navisphere Manager to manage disk-drive Spin Down at the RAID group level	8
Using Navisphere CLI to manage disk-drive Spin Down at the RAID group level.....	10
Reading disk-drive Spin Down information at the drive level	11
Using Navisphere Manager to display disk-drive Spin Down information at the disk level ...	11
Using Navisphere CLI to display disk-drive Spin Down information at the disk level	12
Power Saving Statistics Logging	12
Navisphere CLI support.....	13
Tree view changes in Navisphere Manager GUI	13
Reliability	14
How EMC tested this technology	14
Ongoing processes that ensure reliability	15
Sniffing – a background health-check service	15
Error handling	16
Configuration rules	16
Best practices	17
Tested use cases	17
Conclusion	18

Executive summary

Over the past several years, there have been rapid and drastic increases in energy costs, and it appears that this trend will continue for the foreseeable future. Many organizations are having difficulty meeting the growing demand for power in their data centers, either because power companies cannot generate enough power or the necessary infrastructure simply does not exist.

To help solve these critical problems, EMC continues to develop and implement advanced energy-efficient technologies, such as the disk-drive Spin Down feature now available on EMC® CLARiiON® CX4. This feature is included in FLARE® release 29. It reduces power consumption by automatically *spinning down drives* when they are not accessed for a certain period of time, and spinning the drives up when they are needed for an I/O operation. CLARiiON background services make sure that the drive remains in a healthy condition by periodically spinning up the drives for a health check. No separate enabler is required for using this feature.

Introduction

This white paper discusses CLARiiON's new disk-drive Spin Down technology, including how it works, how to enable it, and when and how to use it.

When this feature is enabled, the CLARiiON storage system keeps track of how long disks are idle. If a drive is idle for more than 30 minutes, CLARiiON will *spin down* the drive's motor so that it is no longer spinning. The electronics for the drive are still powered up. If there is an I/O request for that drive, CLARiiON spins the drive up to its normal speed.

FLARE automatically senses if a drive is *qualified* for spin down. (Certain drives have been qualified by EMC to work with the feature.) For a drive to enter spin-down mode, it must be a qualified drive and the spin-down feature must be enabled on the storage system. The criteria that FLARE uses to decide when to enter *standby* mode is discussed in the "Criteria to enter standby mode" section on page 5.

For flexibility, CLARiiON allows you to manage spin down at different levels; for example you can manage this feature at the storage-system level and the RAID-group level. You can also monitor this feature at the disk level. This is discussed in the "Different ways to manage disk-drive Spin Down" section on page 6. The paper also discusses the power savings that can be achieved, and presents tested use cases with this technology.

Audience

This white paper is intended for EMC customers, partners, and employees who are considering the use of the disk-drive Spin Down feature. It is assumed that the reader is familiar with CLARiiON storage systems and its management software. A good introductory paper is *EMC CLARiiON Fibre Channel Storage Fundamentals* available on Powerlink®.

Terminology

Spin-down capable RAID group – A RAID group that consists of drives that EMC has qualified for the spin-down feature. Note that disks in this group cannot be located in a vault slot (first five drive slots in Bus 0 DAE 0).

Spin-down eligible drives – These drives must be spin-down qualified drives, and they should satisfy the criteria described in the "Configuration rules" section on page 16.

Spin-down qualified drives – Drives that EMC has qualified for the spin-down feature in the CLARiiON storage system. EMC tests these drives to make sure that they can withstand multiple spin-up and spin-down cycles during the drive's lifetime.

Idle time requirement – When the idle time for a drive exceeds this parameter, the drive goes into standby mode.

Standby mode – When a drive goes into standby mode, CLARiiON spins down the drive so that it is no longer spinning.

Power savings with disk-drive Spin Down

The amount of power that the spin-down feature saves is proportional to how much time a disk spends in idle mode. The disk's drive motor does not run in idle mode; the only power required in this mode is the power used by the electronics. Table 1 shows how much power drives use in different modes, and how much power is saved. Note that Table 1 only compares power consumed by the same drive in different modes; it does not include the power consumed by CLARiiON's other components such as the storage processor enclosure, disk array enclosures, and so on.

Table 1. Power saving with disk-drive Spin Down

Type of Drive	Power Consumption in Active state (W)	Power Consumption in Standby state (W)	Power Saving (%)
5400 RPM SATA II	8.87	4.06	54.23
7200 RPM SATA II	12.15	4.77	60.74

How much power this feature saves depends on the application that uses the disk and how often the application accesses the data on the disk. For example, in a multi-purpose fully loaded CLARiiON CX4-960 environment, if more than 50 percent of the data is archived (infrequently accessed), the spin-down feature can save up to 30 percent in power usage.

Criteria to enter standby mode

FLARE monitors the I/O activity on each drive in a disk-spin-down capable RAID group. This monitoring process runs independently on each storage processor in the CLARiiON storage system. The drives are allowed to enter standby mode when both storage processors report that the drives have not been accessed for 30 minutes. (The idle-time requirement is set to 30 minutes; this is a fixed value.)

It is important to note that FLARE monitors the I/O activity at the RAID group level and *not* at the LUN level. So, it is critical to design the LUN layout based on the *access profiles* of the LUNs when using the spin-down feature. For example, if two LUNs are placed in a RAID group – one that is accessed very frequently and another that is accessed very rarely – the drives will probably never enter standby mode since they will not be idle long enough to satisfy the idle time requirement.

These drives enter the standby mode (within one to two minutes) after the spin-down feature is enabled:

- Drives that are not in a RAID group.
- Drives that are in a RAID group but do not have LUNs created on them.

Drives that are configured as hot spares enter standby mode after being idle for 30 minutes.

Criteria to exit standby mode

When FLARE receives an I/O operation for a drive that is in standby mode, one of the following is possible:

- **If the I/O is from an external host (on the server side)** — The I/O is kept in a *pending* state in the FLARE's host-side layer while a request is sent to the drive to exit standby mode. If the host-side

driver receives more I/O requests, it saves them until all drives in the RAID group are in a *ready* state, then the host completes the I/O operations.

- **If the I/O is an internal operation originating within FLARE** (to Create or Destroy LUNs or RAID groups; Rebuild; Equalize; Background Zero; Background Verify; Cache Flush; Cache Prefetch; or RAID group expansion) — The I/O is kept in a *pending* state while a request is sent to the drive to exit standby mode. When all the drives enter the ready state, the I/O operation is completed.
- **If the I/O operation is a drive health check** (Sniff Verify background service) — The drive may or may not exit standby mode. This is explained in more detail in the “Sniffing – a background health-check service” section on page 15.

Drives that are configured as hot spares exit standby mode when they need to swap in for any other drive in the storage system.

Drive spin-up algorithm

FLARE limits the maximum number of drives that can be spinning up in one enclosure at any point of time. This is done to make sure that each enclosure remains within its power budget.

In the worse-case scenario, where all 15 drives in an enclosure are in standby mode and need to spin up at the same time:

- FLARE allows up to four drives to spin up at the same time.
- FLARE waits 12 seconds before starting to spin up the next set of four drives in the enclosure.

These steps are taken when an enclosure filled with multiple drives is powered up. Internal lab testing shows that the worse-case time for the last drive in a disk array enclosure (DAE) to become ready for I/O is well within the latency toleration limits of most operating systems. This worse-case analysis assumes the following are true:

- The DAE is filled with 15 drives.
- All 15 drives are in standby mode.
- All 15 drives need to exit the standby mode at the same time.
- Data needs to be accessed from the last drive in the DAE.
- All drives are in a “healthy” state and there are no delays due to issues with the drive hardware.

Different ways to manage disk-drive Spin Down

To provide flexibility, FLARE provides different ways to manage the spin-down feature.

Managing disk-drive Spin Down at the storage-system level

You can enable or disable the spin-down feature at the storage system level. When you enable this feature at the storage-system level, drives that are not bound in a RAID group immediately transition to standby mode. RAID group drives follow the spin-down setting at the RAID group level. (This is explained later in the white paper.) When you disable this feature at this level, all drives in standby mode immediately transition out of standby mode.

By default, the spin-down feature is disabled on the storage system and needs to be enabled to allow the spin-down capable drives to go into standby mode. If the feature is enabled at the storage-system level, it is important to make sure that it is also enabled at the RAID group level. The enable/disable setting for the spin-down feature maintains its setting after storage-system and storage-processor reboots.

Using Navisphere Manager to manage disk-drive Spin Down

To manage this feature in Navisphere Manager, right-click the storage system’s serial number and select **Power Savings...** to open the **Storage System Power Savings** dialog box shown in Figure 1.

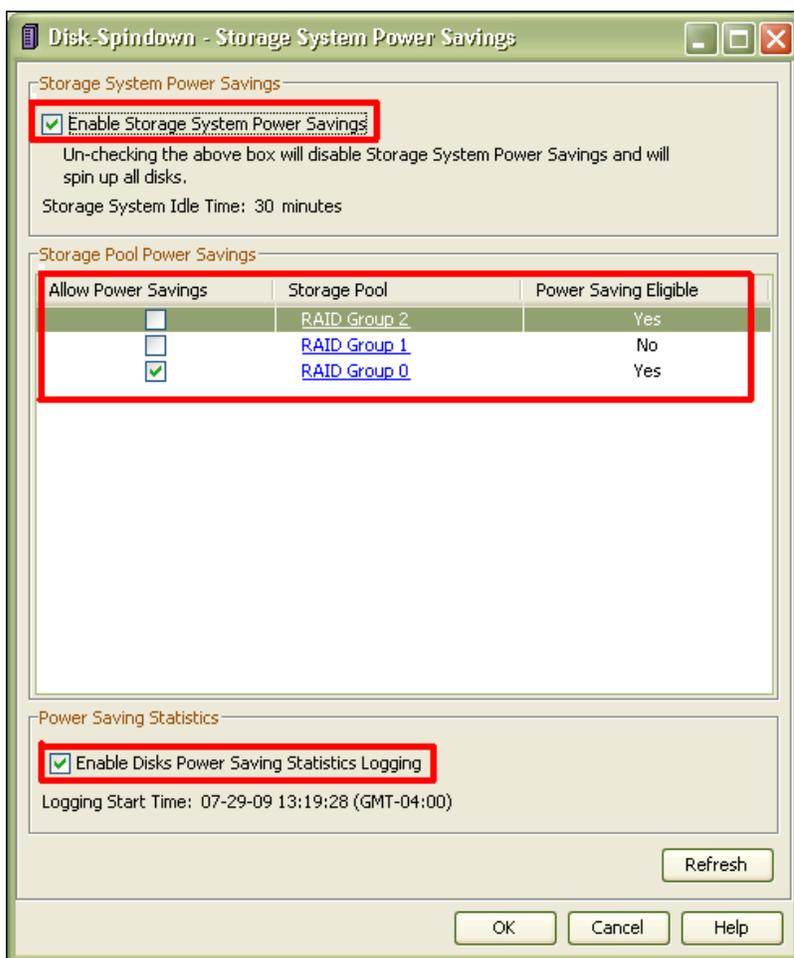


Figure 1. Storage System Power Savings dialog box

In this dialog box you can:

- Select **Enable Storage System Power Savings** to enable the spin-down feature for this storage system.
- Select **Storage Pool Power Savings** to enable the spin-down feature for a specific RAID group. In this example, RAID groups 0 and 2 are eligible for spin down in this example.
- Click a RAID group link (in the Storage Pool column) to open the **RAID Group Power Savings** dialog box if you wish to configure settings for that RAID group. This dialog box is shown in Figure 5 on page 10.
- Select **Enable Disks Power Saving Statistics Logging** to enable statistics logging.

Using Navisphere CLI commands to manage disk-drive Spin Down

This CLI command enables or disables the spin-down feature at the storage system level:

```
naviseccli -h IP address powersaving -globalsettings <ON|OFF>
```

This command shows the setting for the spin-down setting:

```
naviseccli -h IP address powersaving -info <-systemidletime> <-globalpowersavingsettings>
```

Managing disk-drive Spin Down at the RAID group level

You can also enable or disable spin down at the RAID group level. By default, the spin-down feature is disabled for all RAID groups. The setting for this feature will be the same after storage-system and storage-processor reboots.

Using Navisphere Manager to manage disk-drive Spin Down at the RAID group level

The Navisphere Manager GUI displays the RAID groups that are eligible for disk spin down by appending the words **Power Savings** to other RAID group information. This is discussed in the “Tree view changes in Navisphere Manager GUI” section on page 13.

Creating RAID groups

You can also select spin-down-eligible drives in Navisphere Manager when you create RAID groups in the **Disk-Spindown – Create Storage Pool** dialog box shown in Figure 2.

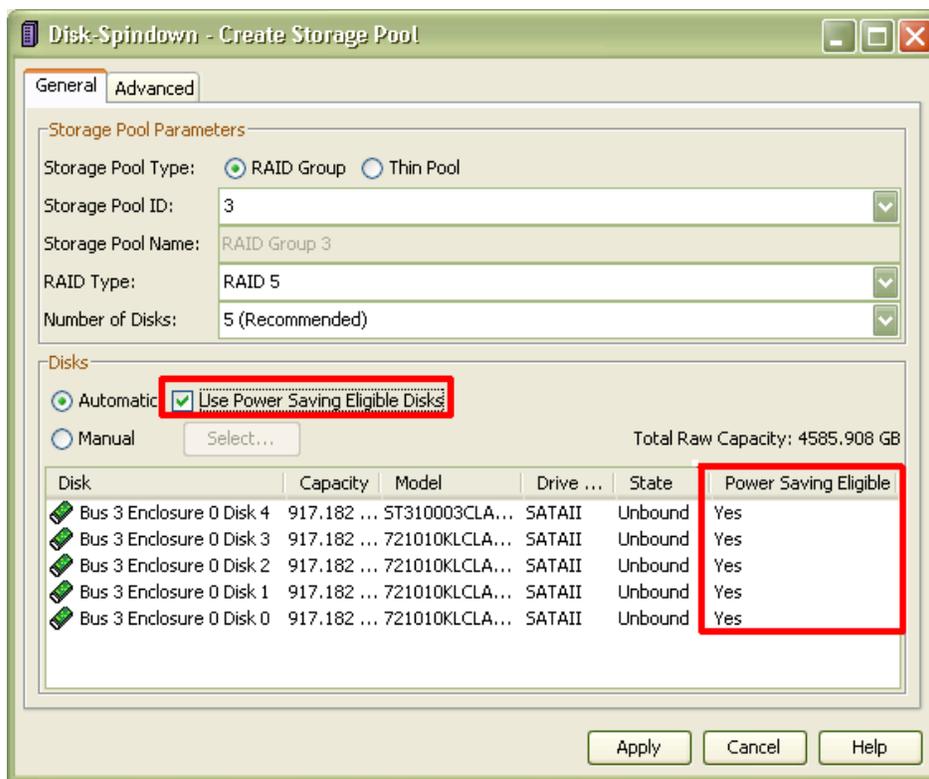


Figure 2 Disk-Spindown - Create Storage Pool dialog box: Selecting drives automatically

If you wish to create a RAID group of power-saving-eligible drives, select **Use Power Saving Eligible Disks** as shown in the example. Also note that a new column called Power Saving Eligible has been added. In this example, all of the listed disks are power saving eligible.

If you wish to manually select disks for a RAID group, select the **Manual** radio button instead of **Automatic**. When you select **Manual**, the **Disk-Spindown – Disk Selection** dialog box shown in Figure 3 opens.

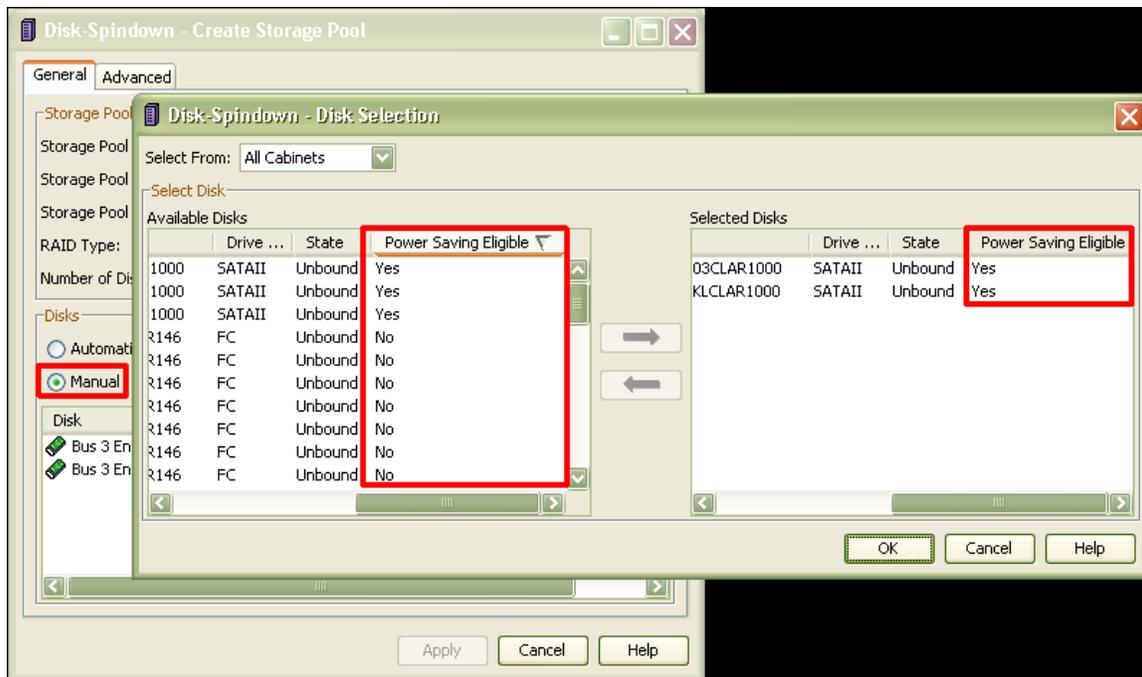


Figure 3. Disk-Spin-down – Disk Selection dialog box: Selecting disks manually

Enabling RAID groups

As mentioned earlier, the spin-down feature has to be enabled at both the storage-system level *and* the RAID group level to allow the drives to transition to standby mode. To enable this feature when creating a new RAID group, select the **Advanced** tab in the **Disk-Spin-down - Create Storage Pool** dialog box as shown in Figure 4.

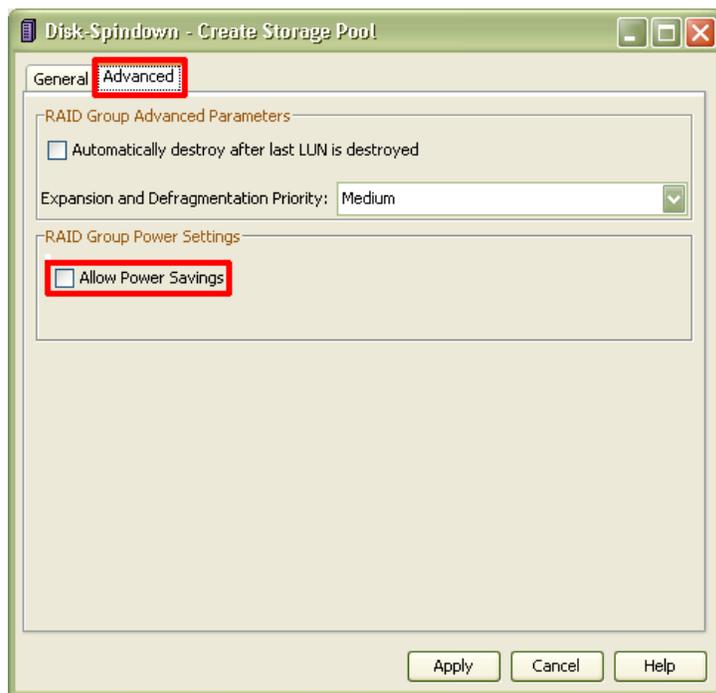


Figure 4. Enabling spin down on new RAID groups

To view and configure spin-down settings for an already existing RAID group, right-click the RAID group name in the Navisphere tree view (shown in Figure 8 on page 14), and select **Power Savings...** in the right-click menu to open the **Disk-Spindown – RAID Group Power Savings** dialog box, shown in Figure 5. This dialog box allows you to:

- Enable the spin-down feature on the RAID group by selecting **Allow Power Savings**.
- Check the spin-down setting at the storage system level. There is a link that opens up the Storage System Power Savings window shown in Figure 1.
- Check if the RAID group is eligible for disk spin down or not.
- Check whether each drive in the RAID group is qualified and eligible for disk spin down, and if the drives are in the standby state or not. This is done from the Disks tab.

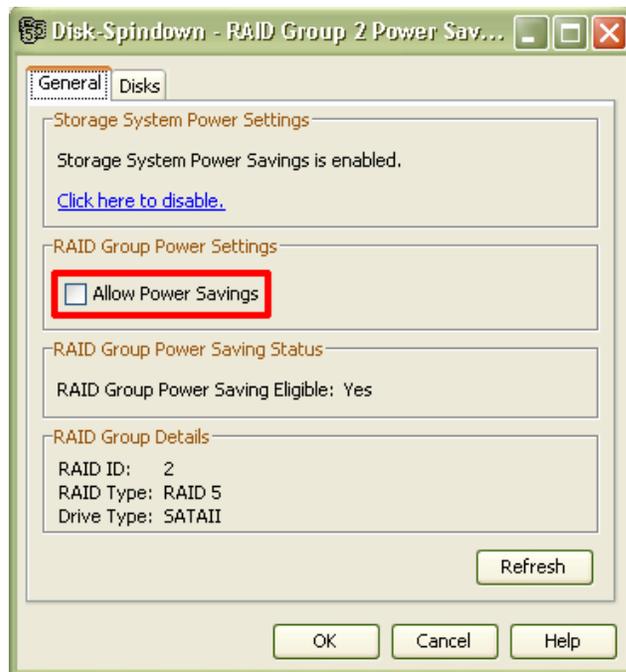


Figure 5. Disk-Spindown – RAID Group Power Savings dialog box: Enabling disk spin down on existing RAID groups

Using Navisphere CLI to manage disk-drive Spin Down at the RAID group level

Use these CLI commands to manage spin-down settings at the RAID group level:

To enable spin down when the RAID group is created:

```
naviseccli -h IP address createrg ... -powersavings ON|OFF
```

To enable or disable spin down on existing RAID groups:

```
naviseccli -h IP address chgrg ... -powersavings ON|OFF
```

To list the spin-down settings on a RAID group:

```
naviseccli -h IP address getrg ... <-powersavingssetting> <-powersavingseligible> <-isrginstandbystate> -all
```

Reading disk-drive Spin Down information at the drive level

Using Navisphere Manager to display disk-drive Spin Down information at the disk level

Navisphere allows you to read spin-down related information at the drive level; you cannot change settings at this level. To view this information, right-click the disk name in Navisphere tree view (shown in Figure 8), and select **Properties...** to open the **Bus Enclosure** dialog box. Select the **Power Settings** tab as shown in Figure 6.



Figure 6. Bus Enclosure dialog box: Displaying spin-down information at the disk level

The Bus Enclosure dialog box displays:

- If the drive is qualified and eligible for disk spin down.
- Which of the three power-saving states the drive is in:
 - Full Power – The drive is spinning at a normal rate.
 - Low Power – The drive is in standby mode; there will be delays accessing the data.
 - Reduced Power – The drive is transitioning from one state to another and cannot be accessed for data at this point.
- If **Disk Power Saving Statistics Logging** is enabled on the storage system. There is a link that opens the **Storage System Power Savings** dialog box shown in Figure 1.
- The values of the power-saving statistics counters for the drive if **Disk Power Saving Statistics Logging** is enabled on the storage system. These counters are discussed in the next section.

The current power saving state of any drive in the storage system can also be seen in the **Disk-Spindown — Disk Summary** dialog box shown in Figure 7. To open this dialog box, right-click the storage-array

name in Navisphere Manager, and select **Disk Summary...** The column labeled Power Saving State can be used to determine if the drive is in standby mode or not.

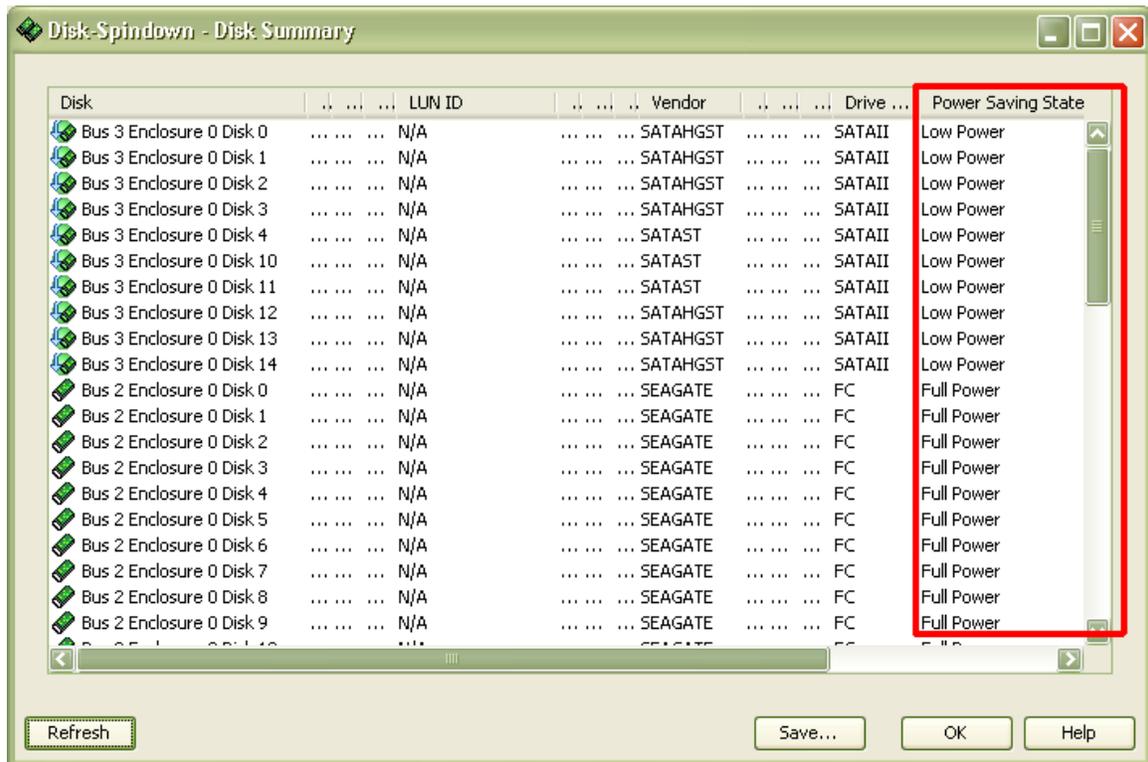


Figure 7. Disk-Spindown — Disk Summary dialog box

Using Navisphere CLI to display disk-drive Spin Down information at the disk level

The following CLI commands can be used to check the spin-down settings at the drive level:

To see information about the disk:

```
naviseccli -h IP address getdisk ... <-powersavingsdiskcapable> <-powersavingsdiskeligible> <-powersavingsstate> -all
```

This command has special switches if you want to see a specific feature on a drive.

To see disk spin-down information for all drives in the storage system:

```
naviseccli -h IP address getall -disk
```

Power Saving Statistics Logging

If **Power Saving Statistics Logging** is enabled, FLARE maintains a spin-down log on the storage system; the information in this log can help data-center managers evaluate system performance. This feature has to be enabled (or disabled) at the storage system level (shown in Figure 1). The counters are maintained and can be viewed at the drive level (shown in Figure 6). By default, this feature is *not* enabled on the storage system.

The following counters are maintained as part of this and most of them are displayed in the Bus Enclosure dialog box (Figure 6):

- Last Refreshed – last time when the statistics was retrieved by FLARE from the drives.

-
- Spinning ticks – amount of time that the drive has spent in spinning (normal) state. Each *tick* corresponds to a one-minute interval.
 - Standby ticks – amount of time that the drive has spent in standby mode. Each tick corresponds to a one minute interval.
 - Number of Spin Ups – number of times the drive has spun up from the standby state.
 - Logging start time – when Power Saving Statistics Logging was enabled on the storage system.

The first four counters are displayed in the Bus Enclosure dialog box (shown in Figure 6), and the Logging Start Time counter is displayed in the **Storage System Power Savings** dialog box (shown in Figure 1).

Please also note:

- All the Power Saving Statistics counters are reset when the Power Saving Statistics Logging feature is disabled at the storage system level or the storage system is rebooted.
- The counters are *not* reset when a single storage processor reboots.
- When a storage system reboots, the counters are reset, but the setting for the enable/disable option remains the same.
- Even if a drive is not qualified for disk-drive Spin Down, or spin down is disabled on a RAID group, these counters will be logged and displayed if Power Saving Statistics Logging is enabled at the storage system level.

Navisphere CLI support

The following Navisphere CLI commands display information related to the Power Saving Statistics Logging feature:

To see if the feature is enabled, and if so when it was enabled:

```
naviseccli -h IP address setstats -powersavings
```

To enable or disable this feature:

```
naviseccli -h IP address setstats -powersavings <-ON|OFF>
```

To display power savings statistics counters for the counters:

```
naviseccli -h IP address getdisk <-spinstats> -all
```

Tree view changes in Navisphere Manager GUI

Navisphere Manager displays disk-drive Spin Down information in the tree view shown in Figure 8. Note that:

- If all disks in a RAID group are *eligible* for disk-drive Spin Down, then Navisphere Manager appends power savings along with other metadata information, even if spin down is not enabled on this RAID group.
- There is a special icon for RAID groups that consist of drives that are eligible for disk-drive Spin Down. This new icon is displayed even if spin down is not enabled on the RAID group.
- If a drive is in standby mode, Navisphere Manager appends Low Power along with other metadata information for the drive. This metadata is not displayed when the drive is operating in the normal mode.
- A special icon is used for drives that are in standby mode. This changes to another icon when the drive is being accessed by FLARE and is not in standby mode.

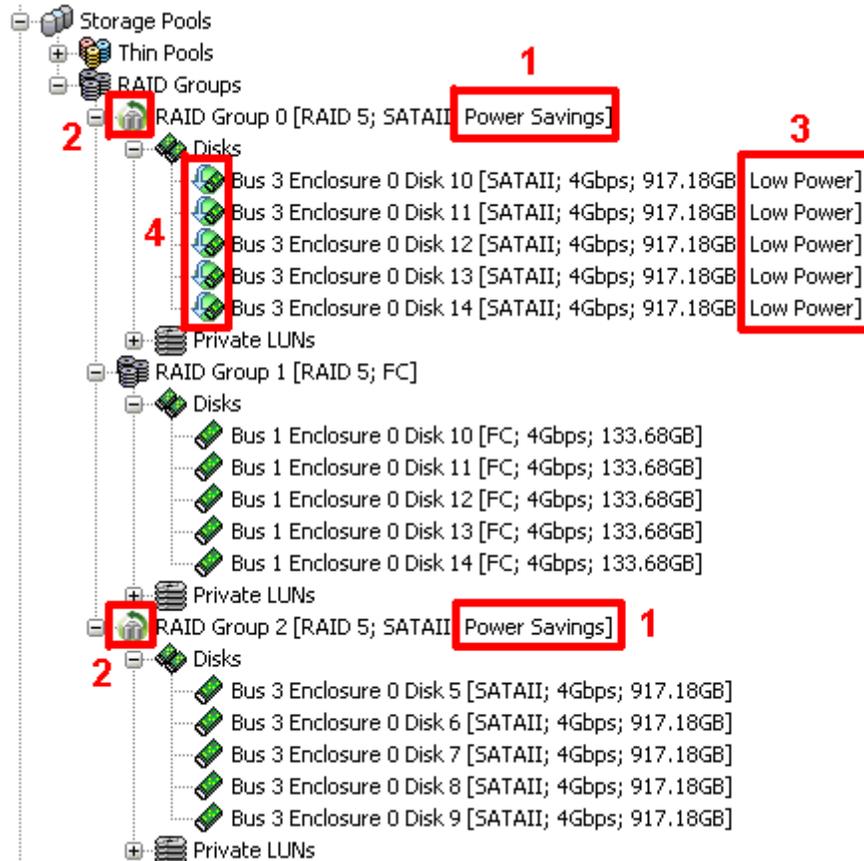


Figure 8. Navisphere Manager tree view

Reliability

Because disk-drive Spin Down technology uses disk drives in a new way – spinning them down and up repeatedly – EMC performed extensive testing to make sure that this usage model would not compromise data integrity or product reliability. EMC also created new processes and modified FLARE’s background health-check services to ensure high reliability.

How EMC tested this technology

Of course, disk suppliers perform their own start/stop tests. EMC also created and executed additional reliability tests for each drive family that is qualified for spin down. The test sequence is: spin a drive up; write to a random location; read the data just written; perform a data integrity check on the data just written against the data read back; spin the drive down; and repeat the entire test.

We ran these tests while cycling the drives between very high and very low temperatures. These tests validated the vendor’s specifications, and there was no difference between using the drives in a regular mode and using the drives in a spin-down mode.

In addition, we run this test sequence, at ambient temperature, on every drive that we ship to customers.

Ongoing processes that ensure reliability

There is an ongoing audit test at the disk suppliers and EMC manufacturing. The purpose of this process is to catch any manufacturing variations that might affect reliability. At the disk vendor's manufacturing facility, some percentage of each lot of disks is pulled aside for start/stop testing. A similar test is performed when a shipment of disks is received at EMC manufacturing.

Sniffing – a background health-check service

Sniff Verify is a FLARE service that runs in the background; this service helps detect and correct potential disk media errors before they become unrecoverable. The Sniff Verify service does this by issuing SCSI Verify commands to all disks in a LUN. It is important to note that this background service does *not* run on disks on which no LUNs are created. The SCSI Verify command causes a disk to read its media without returning any data to the storage processor on CLARiiON. If a media error is detected, the associated bad blocks on the disk are remapped. The Sniff Verify service repeats the same procedure for each LUN as it sequences through all the LUNs in a RAID group. The Sniff Verify operation makes sure that the drives will be able to service I/O when they exit the standby mode.

The algorithm for the Sniff Verify operation has been modified for disk-drive Spin Down functionality as shown in the flowchart in Figure 9. A Sniff Verify operation is not counted as an I/O when monitoring the idle time of a RAID group. When a RAID group enters standby mode and receives its first Sniff Verify request, it is rejected with a special error code. FLARE then throttles back its attempts to issue Sniff Verify requests for that RAID group until the RAID group comes out of standby mode. If the drives stay in standby mode for 24 hours, FLARE transitions all the drives to Full Power mode, which allows the next Sniff Verify request to be completed normally. This resets the rate at which Sniff Verify requests are issued for that RAID group to its normal setting. It enables the Sniff Verify operation to run for a minimum of 30 minutes or until the time the drive is free from all other I/O for 30 minutes. In this algorithm, FLARE forces the drives to exit the standby mode at least once in 24 hours to validate that the drives in standby mode can return to full power and process I/O operations.

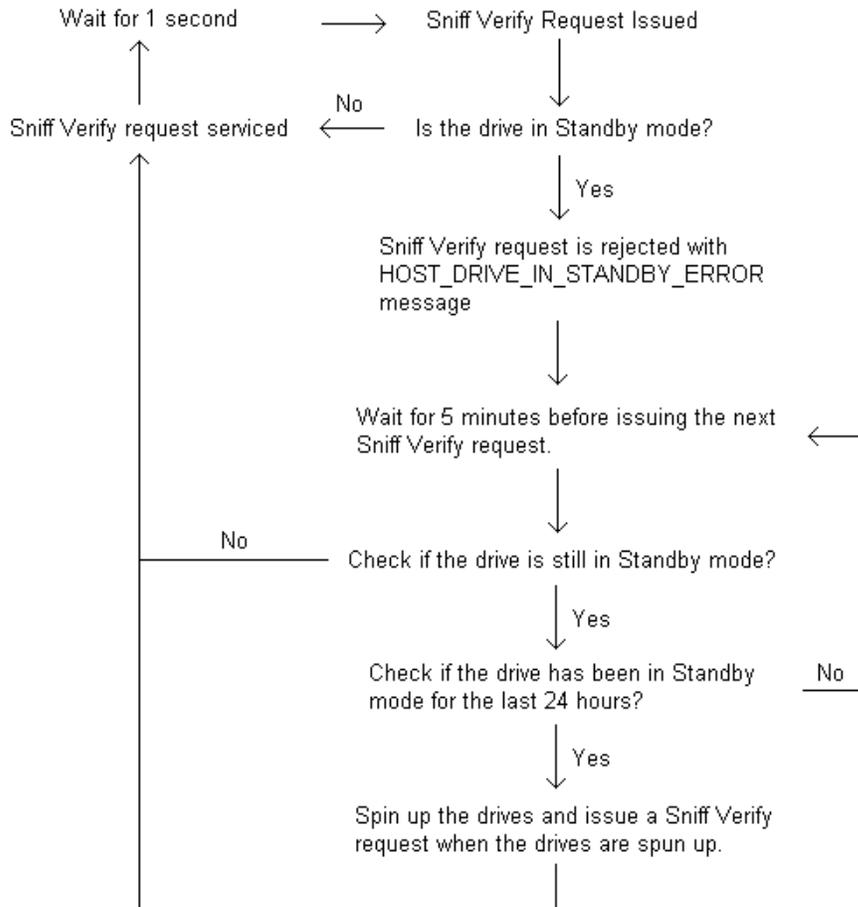


Figure 9. Flowchart for the Sniff Verify background health-check service

Error handling

FLARE handles errors the same way for disks that are in standby mode and disks that are not in standby mode. Any drive that encounters failures while entering or exiting standby mode will follow the same error processing procedure that is followed for any other CLARiiON drive. As in the case with other drives, a hot spare drive will be used if it is configured. If a drive fails when it is in standby mode, FLARE will be able to recognize that drive state and mark the drive slot as **Removed**.

Configuration rules

The following configuration rules exist for disk-drive Spin Down functionality:

- Disk-drive Spin Down is only supported on CX4 storage systems.
- Drives in vault locations (slots 0-4 on Bus 0 Enclosure 0) are not eligible for spin down.
- Unless all drives in a RAID group are eligible for spin down and satisfy the idle time criteria, none of the drives will enter the standby mode. For example, in a 4+1 RAID 5 group, FLARE will either allow all five drives to enter the standby mode at the same time, or none of the drives will enter the standby mode – depending on whether the criteria to enter standby mode is satisfied or not.
- A RAID group that is used for a layered application (such as MirrorView™, SnapView™, MetaLUNs, or Virtual Provisioning™) cannot use the spin-down feature.
- In FLARE release 29, 1 TB SATA II drives are qualified for spin down.

-
- A disk-drive Spin Down eligible drive will enter the standby mode no matter what disk array enclosure it is in.

Disk-drive Spin Down is not a licensed product and is available along with FLARE release 29. No separate enabler is required for using this feature.

Best practices

It is important to consider these items if you plan to use the disk-drive Spin Down feature:

- Since the feature is enabled at the RAID group level and not at the LUN level, it is important to carefully plan the LUN layout within different RAID groups. LUNs that are frequently accessed should not be mixed with LUNs that are rarely accessed in the same RAID group, otherwise the frequently accessed LUN may prevent the RAID group from entering standby mode.
- EMC CLARiiON Virtual LUN technology can be used to move LUNs to the “appropriate” RAID groups. Virtual LUN technology provides an online process that you can use to move LUNs from one RAID group to another RAID group.
- If possible, applications should be designed so that they access storage devices as little as possible to maximize disk idle time. For example, it is better to read and write large chunks of data a few times than to read and write smaller chunks of data many times.

Tested use cases

EMC tested some internal applications using CLARiiON’s disk-drive Spin Down feature to ensure that the applications work properly when the disks supporting them enter standby mode. The following applications were tested with disk spin down:

- EMC Documentum® Content Server Developer Edition 6.5 SP2
- EMC DiskXtender® with UNIX (version DXUL FSM 3.4)
- EMC DiskXtender with Windows (versions 6.3 and 6.4)
- EMC NetWorker® 7.5.1
- General purpose Windows NT file system
- Windows Server presented CIFS/SMB file shares

As mentioned earlier, the spin-down feature has been tested with most operating systems that CLARiiON storage systems work with. These tests showed that the operating systems were able to tolerate the latency introduced when the drives were in a spun-down state.

Refer to application-specific best practices available on EMC Powerlink® for more details.

Conclusion

For many CLARiiON implementations, the disk-drive Spin Down feature reduces power consumption. Disk-drive Spin Down provides the flexibility to set power savings policies at different granularities within the storage system. Considering the drive spin-up latency inherent in this technology, this feature is a natural fit for backup and archiving solutions. There are certain other use cases wherein this can prove very beneficial.

However, disk-drive Spin Down alone is not the “silver bullet,” and for more sizable reductions in power consumption, this technology should be just one of the tools in the toolbox. Feature and functions such as consolidation, tiered storage, virtual provisioning, and server virtualization should all be considered. Additionally, the cooling in the data center should be optimized. When all these techniques are used in concert, then the savings can really start to accumulate and have an effect on a company’s bottom line.