THE POSSIBILITIES ARE INFINITE FUJITSU

MAID for Green

Energy Conservation with Fujitsu ETERNUS Storage Systems

Version 1.1

November 2007

Fujitsu

Table of Contents

1. Introduction

This document describes best practices on how customers can conserve energy using ECO mode on Fujitsu ETERNUS storage systems.

1-1. Scope

Today, because of strong demand for storage capacity, customers are faced with serious issues of power and cooling in their data centers. The density of today's storage systems, the insatiable demand for storage and the rising cost of energy are forerunners of an impending energy crisis in the data center.

Analysts predict that data centers will soon reach their capacity limits in terms of power and air conditioning. This may impose tight restrictions on companies who want to expand their business but also need to substantially increase IT support in growing market segments, if there are no changes in the energy efficiency of IT systems and their operation. Gartner predicts that within the next several years, half of the world's data centers will become obsolete because of power and space restrictions and that energy costs will eat up to one-third of IT budgets. Gartner forecasts that in 2008, 50% of all data centers will no longer be able to meet the power and cooling requirements of their IT systems[1].

Apart from higher energy costs, key causes are the ongoing demand for greater performance and the trend to more densely packaged systems (for storage TB/square feet). The growing use of Internet media and the rise in electronic communication via the Internet are, for many experts, the areas with the greatest future IT support needs. But also the accelerating digitalization of business processes and other influences such as national legislation (e.g. Sarbanes-Oxley Act, Basel II) or disaster recovery measures are dominant driving factors. The introduction of newer and even more powerful applications in these usage scenarios often forces IT departments to opt for the latest hardware technology. Storage systems and also servers of the latest generation almost invariably offer greater performance and capacity in less space than comparable predecessors.

Fujitsu has recognized the impending need for "Green" technology by imbedding the vision of environmental awareness with ecologically responsible programs spanning the last decade. The result from these programs from Fujitsu has produced ecologically responsible design and technology leadership into their products[2]. As a leader in delivering RoHS compliant storage Fujitsu recognized the need to address the problem of toxic e-waste.

Fujitsu's ETERNUS storage systems deliver features designed to address the power and ecological demands of our customers. Using the virtualization features of ETERNUS storage systems, such as RAID migration, RAID expansion and LUN concatenation, ETERNUS storage systems deliver the capability to improve capacity utilization. Improved capacity utilization means reduced power demand which results in cost savings. All of the ETERNUS storage systems support ECO mode based on MAID (Massive Array of Idle Disks) technology that spins down infrequently accessed RAID groups. With effective use of ECO mode customers may realize power saving of 20% or more from their ETERNUS storage system.

In order to manage power supply issue, choosing a system that consumes less power would be the primary agenda, but there are other ways to save power consumption and the ECO mode that is available on Fujitsu's ETERNUS storage systems is worth considering as a way to conserve energy.

1-2. Fujitsu ETERNUS Storage System Products

The current offerings of ETERNUS storage systems are Fujitsu's fourth generation of its disk array product series and meet all the requirements in mission-critical environment. ETERNUS8000 is highly scalable in both vertical and horizontal dimensions, with high performance and large capacity disk array; while, ETERNUS4000 is a flexible mid-range disk array. ETERNUS2000 is an entry level disk array. In all models, all the major components (fans, power supplies, etc.) are redundantly equipped and are hot-swappable to

[1] "Meeting the Data Center Power and Cooling Challenge", Michael A. Bell, Gartner Research ID Number G00144986, 12/19/06.

[2] Fujitsu's environmental policy: http://www.fujitsu.com/global/about/environment/

ensure business continuity with high availability.

1-3. Fujitsu ETERNUS Storage System Reliability

1-3.1. Hot Spare Disks

Hot spare disks are used as a pool of spare disk drives when a disk drive in a RAID group has failed or is in error status.

You must register the hot spare disk in the ETERNUS storage system. If the hot spare disk has been registered, when one of the disk drives in a RAID group has a problem, data from this drive is automatically rebuilt onto the hot spare disk. Note: ECO mode does not apply to hot spare disks (more details in section 3-4).

1-3.2. Pre-failure Detection

ETERNUS storage systems utilize S.M.A.R.T. (Self-Monitoring Analysis and Reporting Technology) in the disk drive to patrol and collect error statistics that make it possible to identify failing disk drives before data redundancy in the RAID group is lost. This feature provides disk drive exchange and maintains full redundancy during copy.

The ETERNUS storage system diagnoses disk drives to detect a possible failure. When the disk drive requires preventive maintenance, an ETERNUS storage system automatically rebuilds the data of the RAID group to a hot spare to ensure data redundancy.

Disk drives are not detached immediately, only after the data backup copy has completed. With the redundant copy function generating the data onto the hot spare disk with full redundancy, it is possible to recover from other disk drive errors during the recovery process, further improving availability.

This self-monitoring analysis and reporting technology improves reliability in the RAID groups including drives used with ECO mode. Note: ECO mode is disabled for a RAID group: When it contains a logical volume that is undergoing rebuild/copy back (more details in section 3-4).

The following diagram shows how this feature is implemented.

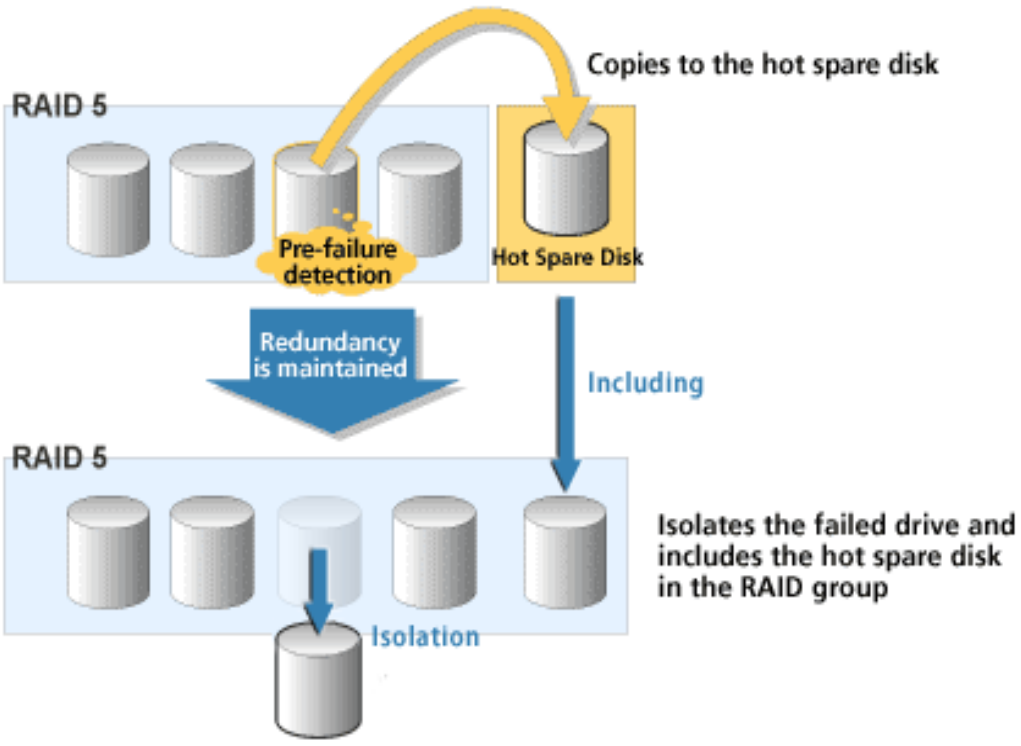

2. Fujitsu's ETERNUS Storage System Energy Conservation

2-1. Definitions for Data Storage Tiers

"Online" is used to describe highly available transactional data, typical of bank-transactions, financial applications and other mission critical systems. It requires the high performance and reliability of high speed fibre channel (FC) or Serial Attached SCSI (SAS) disk drives.

"Nearline" describes data less frequently accessed, which must still be readily available on demand. This includes medical records, images, sound files and e-mail archives. Here the use of lower-cost, high capacity yet reliable Nearline disk drives is a practical solution. The storage space for disk-to-disk copy can be also defined as Nearline.

"Offline" describes data rarely if ever accessed but which must be retained for long term data security, regulatory requirement or as a multi-generational archive of data snapshots where high capacity tape drives are the most suitable medium.

2-1.1. Tiered Storage with ETERNUS Storage Systems

Customers can choose and mix any type of disk drives in ETERNUS storage systems, ranging from 73GB to 300GB / 15,000 rpm (15K rpm) fibre channel disk drives, 300GB / 10,000 rpm (10K rpm) fibre channel disk drives (ETERNUS4000 and ETERNUS8000 only) and 500GB and 750GB / 7,200 rpm (7.2K rpm) Nearline disk drives. Information on the various disk drives available for the ETERNUS storage systems is located in section 2.3. ETERNUS storage systems do not require any additional hardware logic to use different types or sizes of disk drives, so customers have the freedom of configuring disk drives to maximize the usage and efficiency.

The description of 3 tiers of storage within ETERNUS storage systems:
Tier 1: High performance (15K rpm), high availability
Tier 2: Nearline drives deliver good performance (7.2K rpm) at a lower price point
Tier 3: ECO mode - MAID delivers new levels of economy

2-2. Virtualization in ETERNUS Storage System

Fujitsu's ETERNUS storage systems provide two features to manage the logical volumes. Logical volume migration changes the RAID group where a logical volume is located. The entire logical volume is migrated from one RAID group to another in a manner that is transparent to servers and application software.

One way is to migrate logical volume(s) from one RAID group type to another RAID group type in place. In other words, the disk drives used in the current RAID group will be used in the destination RAID group. This feature is useful when customers want to change the RAID-level and keep using the same disk drives. For example, RAID-5 (3+1) can be formed into RAID-10 (2+2) without requiring any additional disk drives to change the RAID group type.

Another feature is to prepare two separate RAID groups and migrate the logical volume(s) from one RAID group to another. This is useful when logical volume need to be moved to a RAID group using different type of disk drives. For example, logical volumes can be migrated from RAID-5 using 10,000 rpm disk drives to RAID-1 using 15,000 rpm disk drives.

In both features, the accesses to logical volumes are automatically switched within ETERNUS storage systems when the logical volume migration is complete.

By combining this feature with multi-tiered storage and ECO mode, customers can easily and efficiently manage the storage capacity.

2-2.1. Information Lifecycle Management (ILM)

The ability to mix different type of disk drives allows the creation of a system whose flexibility matches the value and frequency of the use of data.

A cost-effective 24/7/365 data storage system can be achieved with the ETERNUS storage system by storing frequently used data on high performance (15K rpm) disk drives, while using Nearline disk drives for data backup and archiving.

2-3. ETERNUS Disk Drive Support

2-3.1. ETERNUS2000 Disk Drive Support

SAS Disk Drives
(15,000 rpm)
SAS 3Gbps

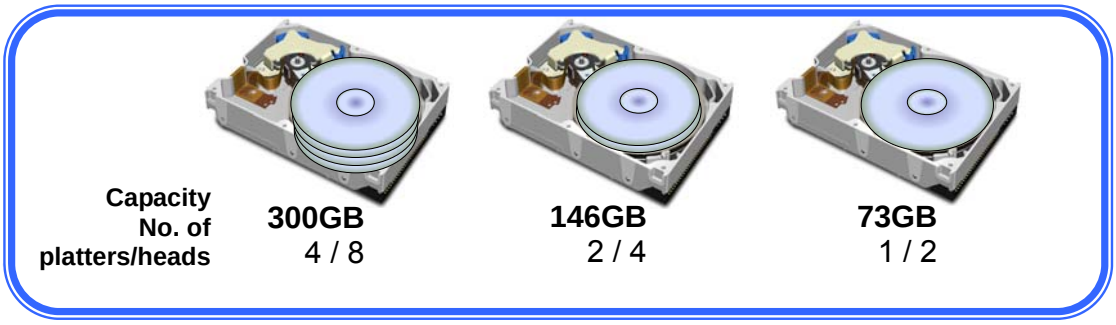

Nearline SATA Disk Drives
(7,200 rpm)
SATA 3Gbps

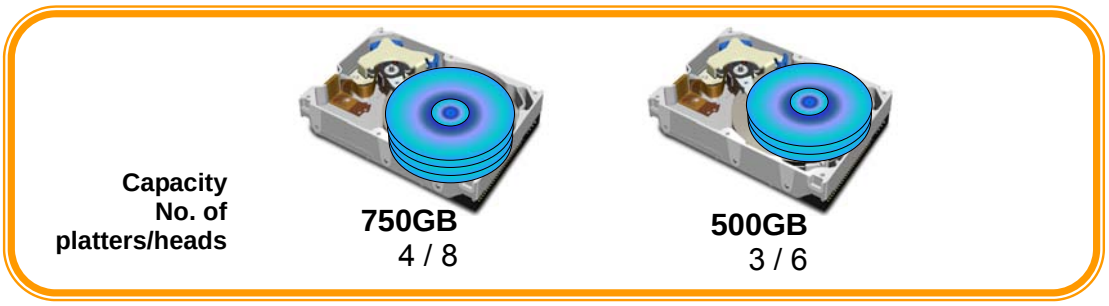

2-3.2. ETERNUS4000 and ETERNUS8000 Disk Drive Support

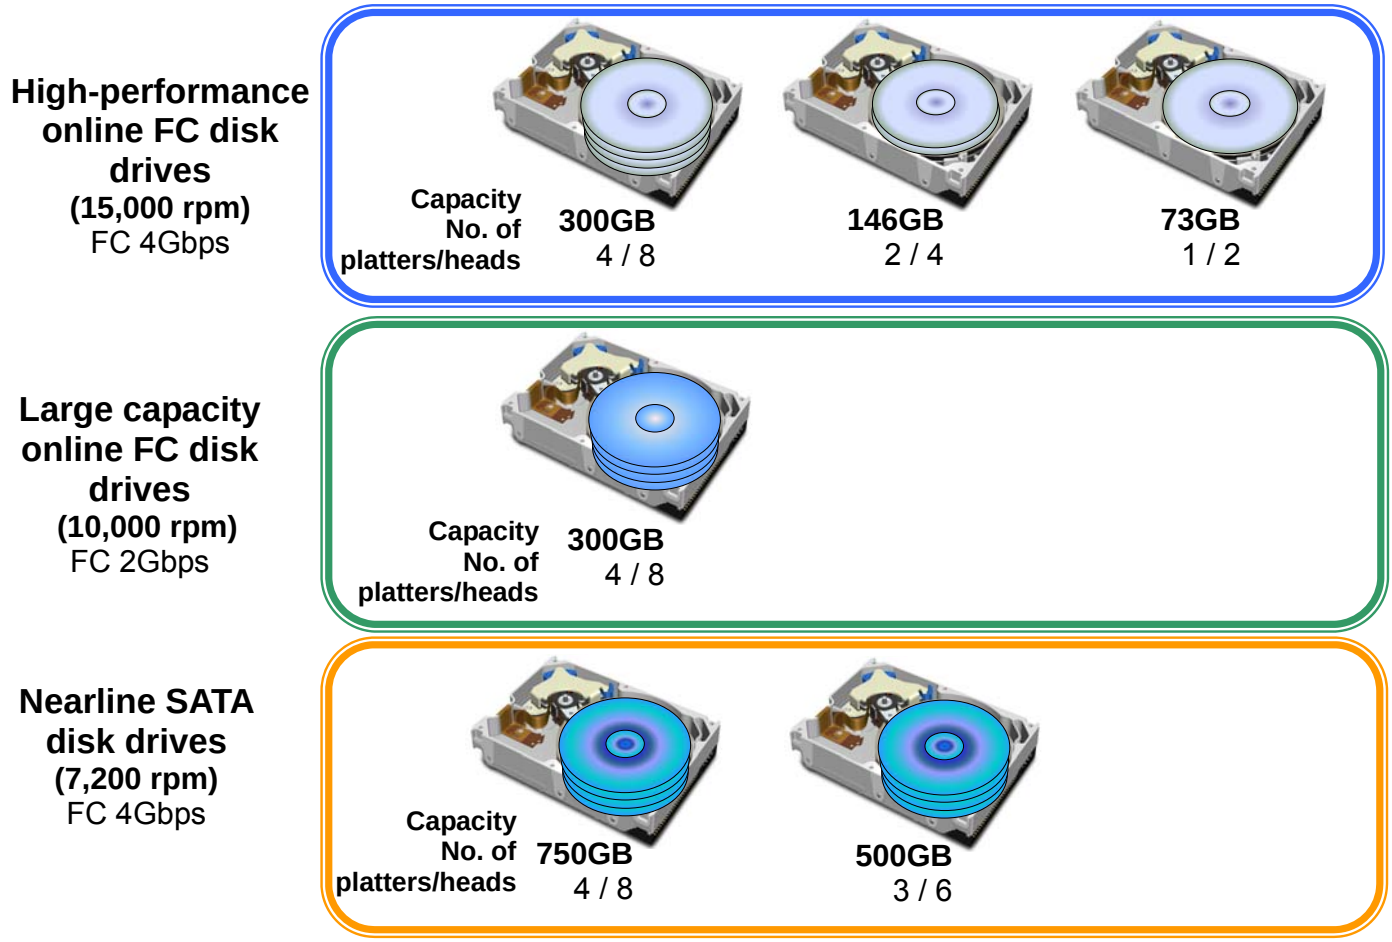

The Nearline 500GB/750GB disk drives enable more cost-effective backup and archiving operations. The Nearline disk drives operate at 7,200 rpm and achieve highly reliable system status through use of redundant interfaces. The Nearline disk drive enables significant cost reductions compared to use of existing standard FC disk drives. In general FC disk drives should be used for on-line storage operations with mission critical systems; which need high transaction performance and high reliability. Nearline disk drives should be used for less critical storage operations needing high-capacity at a lower price point, such as low access storage, backup and archiving operations.

With the variety of RAID levels 1, 10, 5 and 6, customers can create tiers of storage space, such as a RAID-10 using 146GB / 15,000 rpm fibre channel disk drives for high performance transaction operations, RAID-5 using 300GB / 10,000 rpm fibre channel disk drives (ETERNUS4000 and ETERNUS8000 only) for backup and a RAID-5 using 500GB / 7,200 rpm Nearline disk drives for archiving.

2-4. ECO mode – MAID Technology

ETERNUS storage systems are also equipped with the ECO mode based on the MAID (Massive Array of Idle Disks) technology. This feature manages the on/off state of the disk drives' spindle motors based on policy settings as well as usage patterns. This feature enables significant power consumption reductions compared to non-ECO mode environments. For example, customers can specify to spin up disk drive motors when these disk drives are used in a target volume for backup operation and spin down the motors of the disk drives when the copy operation is completed. By using this feature, customers can save energy consumption of 20% or more.

3. Best Practices

3-1. Features for ECO mode

ETERNUS8000, ETERNUS4000 and ETERNUS2000 storage systems are equipped with ECO mode. There is no license fee to enable this feature.

ETERNUS storage systems power consumption is reduced by setting and managing scheduled periods of operation for the disk drives of each RAID group. ECO mode is set on a per RAID group basis. Setup and management of the disk operation schedule is done with ETERNUSmgr for each RAID group. ETERNUS has no technical restriction on ECO mode, so any RAID group can operate with ECO mode enabled. However, from the practical stand point, we highly recommend to assess the usage of each RAID group and enable ECO mode on the RAID groups that carry less frequently accessed data, such as archival or backup.

Customers can choose the wait time for ECO mode to start from the last access. ETERNUS offers 10 minutes to 60 minutes with 10 minute increments. The default is set to 30 minutes.

A server access and access by either Equivalent Copy (EC) or One Point Copy (OPC) will cause the motor of a stopped disk drive to be spun up, allowing the access to proceed normally within a short period.[3]

If a disk drive is activated from the stopped state more than a set amount of times[4] in a day, a state of increased access frequency is assumed and the ECO mode will cease stopping the disk drive motor.

3-2. Mode of Operations to Deploy ECO mode

There are two ways to actually deploy the ECO mode in ETERNUS storage system. One way is to set up a policy on when to enable and disable the ECO mode and another way is to set up so that ECO mode is enabled after certain period of no disk access.

(Example 1): Operation schedule is set as 9:00 to 21:00, the allowed number of activations is two times and there are no accesses outside of the scheduled period.

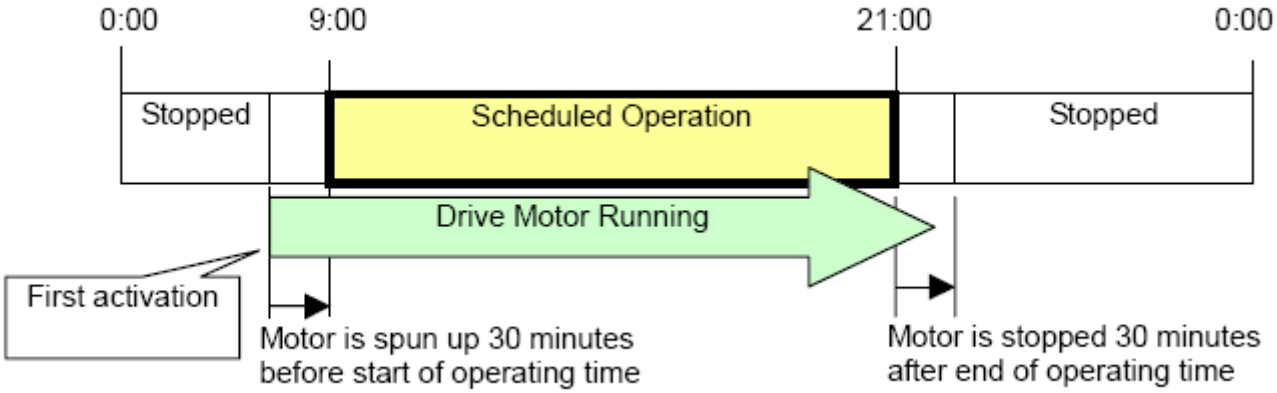

[3] From a few seconds to up to 3 minutes.
[4] The number of activations allowed may be set to between 1 and 5 times, with the default being 3 times.

(Example 2): Operation schedule is set as 9:00 to 21:00, the allowed number of activations is two times and there are accesses outside of the scheduled period.

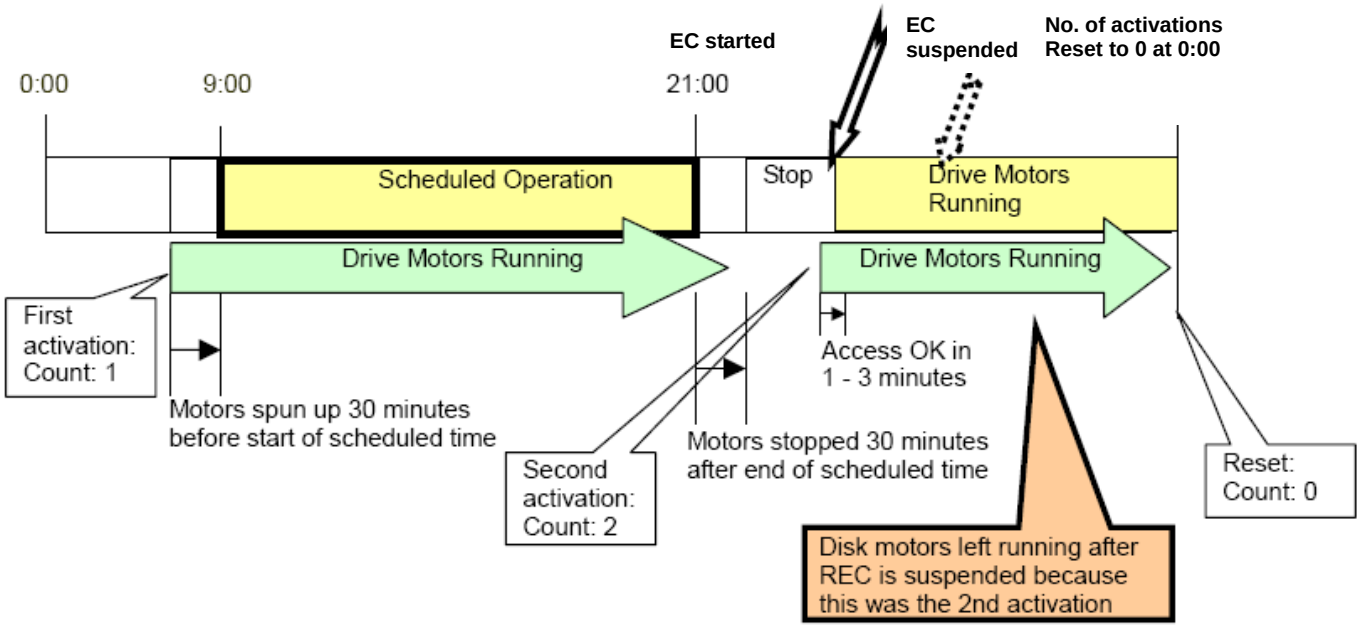

3-3. The Estimated Power Consumption

ECO mode allows customers to configure ETERNUS to spin-down selected drives by RAID group—thereby creating a third tier of addressable storage space ideal for archiving and less-frequent access requirements. ECO mode provides for increased energy conservation and natural cooling by deferring disk spindle up-time. When you measure the cost of running other storage systems against ETERNUS you will see a measurable saving which improves as the amount of data you need to store increases. ETERNUS by itself (without ECO mode) already has an efficient use of electrical power.

ECO mode is particularly useful in situations where there is a system sleep period—where data is not being accessed—the power savings can be significant.

When you want to achieve the elimination of your backup window problems and realize rapid recovery through disk-to-disk backup you can define the RAID groups for your backup target capacity and only spin them up during your backup window. This means extra power and cooling savings for this capacity on top of your economical high capacity Fujitsu Nearline disk drives.

ETERNUS ECO Mode - MAID Power Specification
Fujitsu measured 500GB/7,200 rpm Nearline disk drive
Power when spinning 19.3W
Power when not spinning 4.7W

Example for ETERNUS2000:
ETERNUS2000 Model 200 configuration: Rack mount base with 5 DEs, 72 disk drives
With 40 Nearline disk drives (8 RAID Groups) with 16 hours per day with ECO mode turned on.
25 SAS disk drives (5 RAID Groups) with no ECO mode.
Includes: 4 FC 4Gbps CA Ports, 2 CMs, 4GB Cache

Power consumption (normal) 1,504kWh @ $0.12 per kWh $180.46 for 1 month
Power consumption (ECO mode) 1,377kWh @ $0.12 per kWh $165.27 for 1 month

Estimated power cost savings of $15.19/month or $182.28/year.
Note: The estimated power cost savings excludes any additional costs for cooling that is required.

Example for ETERNUS4000:
ETERNUS4000 Model 500 configuration: Base cabinet with 28 DEs, 420 disk drives
With 275 Nearline disk drives with 16 hours per day with ECO mode turned on.
125 FC disk drives with no ECO mode.
Includes: 16 FC 4Gbps CA Ports, 2 CMs, 32GB Cache

Power consumption (normal) 8,375kWh @ $0.12 per kWh $1,005.02 for 1 month
Power consumption (ECO mode) 7,037kWh @ $0.12 per kWh $844.42 for 1 month

Estimated power cost savings of $160.60/month or $1,927.20/year.
Note: The estimated power cost savings excludes any additional costs for cooling that is required.

3-4. When ECO mode Does Not Apply

ECO mode does not apply to the following disk drives:
− System disks
− Disks in a RAID group for which an ECO mode schedule has not been set
− Hot spare disks

ECO mode will be temporarily disabled while the following failure, maintenance and use of other function situations apply:

ECO mode is disabled for entire device:
− When a controller module (CM) or disk drive access path related module contains failed parts
− When a maintenance engineer is performing a maintenance operation

ECO mode is disabled for a RAID group:
− When it contains a failed disk drive
− When it contains a logical volume that is being formatted
− When it contains a logical volume that is undergoing Logical Device Expansion (LDE)
− When it contains a logical volume that is undergoing Rebuild/Copyback
− When it contains a logical volume that is being encrypted
− When it contains a logical volume that is undergoing RAID migration
− When it contains a logical volume that includes copy data for Advanced Copy
− When an Equivalent Copy (EC) or Remote Equivalent Copy (REC) session exists (including when there is no copy data. However, disk motors may be stopped if the EC or REC session is suspended)
− When a QuickOPC (One Point Copy) or SnapOPC session exists (including when there is no copy data)

3-5. Notes on Schedule Setting and RAID Configurations

Unassigned disks: when a disk drive is added to an ETERNUS storage system that has ECO mode enabled, the disk drive is not initially registered to any RAID groups and the motor stops without condition. New RAID groups: when a RAID group is formed from unassigned disk drives, it has no ECO mode setting and so the motors of the component disk drives are always on.

ETERNUS storage system is designed to activate up to eight disk drives at once in a single Drive Enclosure (DE) in order to minimize rush current. Attempting to activate more than eight disk drives at once will result in the disk drives being activated eight at a time. When this happens, the server will see increased response times and access paths may be lost, in which case the device's RAID configuration may need to be reviewed.

When the Critical System Mode (CSM)[5] is enabled, the I/O timeout settings limit the allowed response times, so spun-off disk drives will not be activated even if a server access occurs and thus the I/O fails. In this situation, the disk operation schedule must be rigorously set to conform to the actual hours of possible business.

ETERNUS2000, ETERNUS4000 and ETERNUS8000 storage systems support the use of the Network Time

[5] For ETERNUS8000, CSM is the default setting and cannot be deactivated. For ETERNUS4000, CSM is deactivated as default and may be activated by the customer if they choose to do so.

Protocol (NTP) and daylight saving time (DST) functions for cross-site time synchronization.

4. Technical References

The specification and features of all the ETERNUS storage systems are published on Fujitsu's website at the URL listed below.

http://www.fujitsu.com/global/services/computing/storage/system/
http://www.computers.us.fujitsu.com/www/products_storage.shtml?products/storage/fujitsu/index

About This White Paper

This white paper is devoted to providing technical information and an overview of the usage of ECO mode of Fujitsu ETERNUS storage systems.

The contents of this document may be modified without any prior notice. Please contact FUJITSU LIMITED if you find any error in the descriptions. FUJITSU LIMITED is not responsible for any damage or loss that might be caused by the contents of this document.

FUJITSU LIMITED

"Fujitsu Storage System Web Site" http://www.fujitsu.com/storage/