Smart Array technology: advantages of battery-backed cache

technology brief

Abstract

Many factors affect overall system performance. This paper explains the importance of using battery-backed cache (referred to as the Array Accelerator on HP Smart Array Controllers) to maximize performance of the storage system and to prevent loss or corruption of business-critical data. It also identifies technology introduced in HP Smart Array Controllers that provides exceptional data protection while minimizing the long-term costs of that protection. It is assumed that readers have at least a basic understanding of RAID storage technology

Introduction

Many factors affect overall system performance. This paper examines the importance of using battery-backed cache on Smart Array controllers to maximize performance of the storage system and to preserve posted-write data in cache.

Role of cache in the storage system

The process of transferring data to and from disk storage includes storing the data temporarily in cache memory located on the RAID (redundant array of independent disks) array controller that is managing the data transfer. Cache is comprised of high-speed silicon memory DRAM chips. The access time for writing data to or reading data from DRAM (dynamic random access memory) memory is roughly 106, or a million times faster than the typical access time for writing directly to or reading directly from a set of disk drives. In a posted-write operation, as soon as the host computer writes data to the cache, the write operation is completed; and the host is freed up immediately to perform another operation. The host does not have to wait for the write data to be transferred to disk. Therefore, use of cache memory on HP RAID array controllers significantly speeds up write operations and increases overall system performance.

Posted-write operations also provide other benefits. If the host computer reads from data still residing in write cache, the RAID controller will quickly deliver this read data from its cached contents. If the host writes new data to the same location, the controller will over-write the cached contents, thus eliminating an extra disk write operation.

Write cache can improve the performance of any RAID level by combining two or more adjacent write requests into a single request for a disk drive. Such an operation reduces disk rotational latency and processing overhead. Write cache can significantly improve RAID5 performance by accumulating enough write requests to perform a full-stripe write to the disk drives. This eliminates the need to read-modify-write parity data for each write access, since all members of the stripe needed to update the parity are available in the write cache.

Need for a battery-backed cache

Once a data write is completed, the cache must be able to protect this data in the event of a power loss before the write data is moved to the drives. If external power to the system is interrupted, if the host computer fails, or if the RAID controller fails, battery power maintains data in the cache long enough for the user to recover the data and preserve the integrity of business-critical information. Depending on the capacity of the batteries when external power is lost, battery-backed cache is typically capable of retaining the integrity of data in the cache for several days. The batteries used in HP's battery-backed cache are rechargeable, allowing them to maintain full charge even after a power failure. HP protects cached data by not allowing posted-write operations unless the batteries are at an acceptable charge level to back up data for a predetermined duration.

With battery-backed cache available, the RAID controller can also speed up other operations such as capacity expansion and RAID migration by temporarily storing data in battery-backed cache DRAM instead of on the drives.

CAUTION:
To avoid risking the integrity of business-critical data, users should always properly shut down the operating system; this guarantees that all posted-write data in the cache will be transferred to the disk drives before power is removed.

Recovering data from battery-backed cache

If an unexpected power loss occurs while data is held in the posted-write cache, the Smart Array controller will automatically write this data to disk when power is restored within a few days. If power is not restored within a few days, the write-cache batteries will be drained and posted-write data in the cache will be lost. Once system power is restored, the rechargeable batteries will recharge within a few hours.

In the event of a server power supply failure, the Smart Array controller card and all of the drives can be moved to another server to allow the data in the write cache to be written to the drives. In the event of a controller failure, the cache module containing posted-write data can be moved to a new Smart Array controller. That new Smart Array controller containing the posted-write data in its cache must be attached to the original drives for which the posted-write data is intended.

There is a special concern when using an embedded RAID controller with battery-backed cache. Since the RAID controller is embedded on the server board, if the server board fails, the same model server board must replace it, and the battery-backed cache should be installed on the new server board.

Selection criteria for battery-backed cache

RAID controller vendors offer products using a variety of battery configurations. Implementations of battery-backed cache vary by type of battery used, usable battery capacity, and the simplicity of data recovery in the event of external power loss.

Parameters that affect battery life and performance include:

- Age of the battery
- Temperature of the server environment
- Capacity Memory Effect (which for NiCD batteries dramatically reduces capacity unless periodic full discharges are performed)
- Number of supported deep discharge cycles before capacity reduction
- Capacity reduction caused by dendrite formations (which result from continuous trickle charging)
- Accuracy of measuring remaining capacity

Since multiple factors can reduce battery capacity over time, it is important for users to have an accurate way of determining the remaining battery capacity.

Types of batteries

Nickel Cadmium (NiCD) batteries are commonly used for battery-backed cache, but they have important limitations: They have a capacity memory effect, which reduces battery capacity; and they provide no way of accurately measuring available capacity of the battery.

For several years HP has used Lithium Manganese Dioxide (MnO2-Li) batteries for backing cache. These batteries do not have a capacity memory effect, but they support only 60 to 100 deep discharge cycles. An important point to remember is that deep discharge cycles of Lithium Manganese Dioxide batteries should be avoided because they typically reduce the remaining battery capacity

Importance of a battery maintenance schedule

Batteries used for backing cache typically have a life expectancy of three years. Therefore, HP recommends replacing these batteries within three years of use. Beyond three years, battery packs cannot be expected to retain sufficient capacity to safeguard cached data from an external power loss for the specified amount of time.

HP recommends that users of battery-backed cache implement a maintenance schedule that identifies when batteries require replacement, and then replace the batteries promptly. Battery date codes are visible on the outer edge of the battery casing in the format MMYY.

Battery replacement

For the past few years, most batteries used for backing of cache have been permanently attached, either to the RAID controller card or to a detachable cache board mounted on top of the RAID controller card. To replace the batteries, users have needed to replace each RAID controller card, or at least the detachable cache board, within three years. Unless users have implemented HP's line of Smart Array controllers, replacement of the RAID controller can be an extraordinarily expensive and tedious process involving total backup and restore of all data, application, operating system, and storage configuration information.

With the introduction of the HP Smart Array 5300 Controller series, HP introduced serviceable battery packs. This feature enables tool-free replacement of the battery packs with no need to replace either the Smart Array 5300 Controller series card or the detachable cache board.

Alternatives to battery replacement

If it is not possible or desirable to replace the batteries, users have two options for continuing to manage storage without significant risk of losing critical data.

First, management software for most RAID controllers enables storage management professionals to set how cache on the controllers will be used:

- 100 percent as read cache
- 100 percent as write cache
- Some percentage for each; for example, 50 percent read and 50 percent write or 75 percent read and 25 percent write.

Using the configuration software, users can turn off the write cache by selecting **100% read cache**. However, if the write cache is turned off, some degradation of write performance may occur. Additionally the user should no longer start operations that require the use of batteries such as RAID migration, stripe migration, and capacity expansion.

As a second option, users can replace the RAID controller in use with a newer, faster, and more powerful RAID controller. If the server itself is due for replacement by a more powerful, faster server, then redeploying the server and its RAID controller (with the write cache turned off) to a less critical application might maximize the return on investment in that equipment.

If batteries cannot be replaced and the write cache is NOT turned off, users should be prepared to experience permanent loss of data if external power or the server fails when the remaining battery capacity is insufficient to maintain data integrity.

Conclusion

Data integrity and availability are critical in today's business environment. Cache on RAID controllers increases the speed of read and write processes; but if the cache is not battery backed, write processes are subject to catastrophic data losses during server or controller failure.

Battery backing of cache is only as good as the remaining capacity of the batteries. Users of battery-backed cache on RAID controllers must be aware of battery life factors that can lead to data loss, even with the protection of battery-backed cache. A battery maintenance program should include battery replacement within three years unless other specific actions are taken.

Since data lives longer than servers, hard disk drives, or even RAID controllers, upgrading equipment while maintaining data integrity can be a significant business cost, if data compatibility and data migration are not supported by the hardware vendor. HP has addressed the needs of today's business environment by providing a family of Smart Array controllers that allow data migration and provide exceptional data protection while minimizing the long-term costs of that protection.

For more information

For more information on HP's products and solutions, see the smart array controller webpage.

Call to action

To help us better understand and meet your needs for ISS technology information, please send comments about this paper to: TechCom@HP.com.

TC020202TB, 02/2002

Printed in the US

invent