



1. Overview

The E-Disk[®] solid state disk has a feature that allows PowerGuard[®]-enabled units to IGNORE any cache disable setting. This feature forces the E-Disk[®] SSD to use the write cache for optimum performance while keeping cache data protected from power outages using PowerGuard[®].

1.1. Write Cache

Write cache is an implementation of the concept of delayed write. I/O activities are generally slow and if they are synchronous, they can significantly reduce system performance. To improve performance, the frequency of disk-writes can be specified at regular intervals. Data that needs to be written between these specified time intervals is staged in a buffer that is maintained in RAM. This write buffer is then destaged to disk at regular intervals or whenever the drive gets full.

Although write cache is a smart concept for improving performance, this technique comes with one particular risk. Cache is typically implemented using RAM, which is volatile in nature. Thus, there is a chance of losing data in case of critical power loss or system failure.

1.2. Cache Disabling

Some hard drives support the option to enable or disable the write cache. This is usually done either from the operating system's device manager or via the disk controller (HBA controller or RAID controller). Some disk or RAID controllers such as Adaptec PCI SCSI cards or RAID controllers provide the option of disabling the write cache of the attached drives while others simply default to either enabling or disabling it. Some applications may also disable the disk write cache thru the operating system. The most common reason for disabling cache is to eliminate potential data integrity issues that may arise during power loss or failure because of the volatile nature of cache. There are also other reasons such as RAID controllers with on-board cache that will inherently have better caching algorithms than those of the attached drives.

2. E-Disk[®] Features

2.1. Performance and Cache Protection

As mentioned earlier, enabling write cache provides performance benefits to most applications. Internal benchmark results generated by BitMICRO also prove that the use of the E-Disk[®] SSD's write cache can improve performance by as much as 150%. However, as mentioned earlier, using the write cache to achieve better performance comes with a tradeoff – the risk of data loss due to power loss or system failure.

To mitigate this risk of data loss while achieving optimum performance, E-Disk[®] SSDs can be equipped with PowerGuard[®] as an option. PowerGuard[®] prevents data loss due to power failure by ensuring that there is enough power in the E-Disk[®] SSD to move the contents of cache to flash memory. The Save Mode is the default mode for all types of PG protection units, including PG I, II, and III. Save Mode activation causes the storage device to automatically flush all dirty data saved in volatile memory (cache memory) to flash memory when external power falls below approximately +4.7V during operation of the device.

PowerGuard[®]-equipped E-Disk[®] SSDs can therefore provide optimum performance by making use of its write cache while offering cache protection from data loss during power loss or system failure at the same time.

2.2. Ignore Cache Disable Feature

As said earlier, some applications or disk/RAID controllers disable write cache on the attached drives. Since E-Disk[®] SSDs are drop-in replacements to hard disk drives, E-Disk[®] SSDs attached to these disk/RAID controllers can also have their write caches disabled. The E-Disk[®] cache can be disabled either by manually disabling it through the operating system's device manager or via the E-Disk[®] Analyzer. It can also be disabled by a command setting from the disk/RAID controller the drive is attached to. There may be other ways of disabling the E-Disk[®] SSD's write cache, but no matter what method was used, there is a way to achieve optimum performance by "ignoring" any disable cache settings. However, this is only possible on PowerGuard[®]-enabled E-Disk[®] SSDs since write cache protection (data integrity) can only be made via PowerGuard[®].

Essentially, ignoring any disable cache settings is making use of the write cache and therefore achieving the optimum performance from the E-Disk[®] SSD while at the same time ensuring data integrity through PowerGuard[®].

2.3. EDA Parameters

Using the E-Disk[®] Analyzer, you can ignore cache disable (on PowerGuard[®]-enabled units only) by changing the following parameters/settings:

2.3.1. PowerGuard[®] Forced Unit Access (FUA)

FUA (Forced Unit Access) provides the user the ability to require the drive to access the data in the flash media directly without using the drive's cache memory for a given command. It requires that data be read/written to the flash media before providing a command completion acknowledgment.

If PowerGuard[®] FUA is set to YES, the device server shall access the media in performing the command prior to returning completion status.

- For write operations, this causes the device server to complete the data write to the physical medium before completing the command
- For read operations, this causes the logical blocks to be retrieved from the physical medium.

If PowerGuard[®] FUA is set to NO, no FUA will occur and all read/write requests go through the cache

- For read operations, any logical blocks that are contained in cache memory may be transferred to the application client directly from the cache memory.
- For write operations, logical blocks may be returned to the application client prior to writing the logical blocks to the medium.

This feature is done at command-level. Commands affected are read/write commands that are FUA. Read/Write commands are FUA if it is set on the actual command or if caching is disabled (i.e. using SCSI mode select command)

To test this feature, disable caching (Windows[®] has a way using the device manager) and observe the difference in read/write performance with PowerGuard[®] FUA settings at YES and NO

2.3.2. PowerGuard[®] SYNCH

Synchronize cache command ensures that the logical block in the cache memory, within the special range, has the most recent data value recorded on the physical medium. If a more recent data value for a logical block, within the specified range, exists in the cache

memory than on the physical medium, then the logical block from the cache memory shall be written to the physical medium. Logical blocks may not be removed from the cache memory as a result of the synchronize cache operation.

Synchronize cache command forces any pending write data in the requested set of logical blocks to be stored in the physical medium. This command may be used to ensure that the data was written and detected errors reported.

If the E-Disk[®] unit is PowerGuard[®]-enabled and PowerGuard[®] SYNCH is set to no, no synchronization of cache will occur, i.e. no forced flushing of cache

- SCSI Command affected is Synchronize Cache

2.3.3. PowerGuard[®] Write Verify

Write and verify command requests that the device server write the data transferred from the application client to the medium and then verify that the data is correctly written. The data is only transferred once from the application client to the device server.

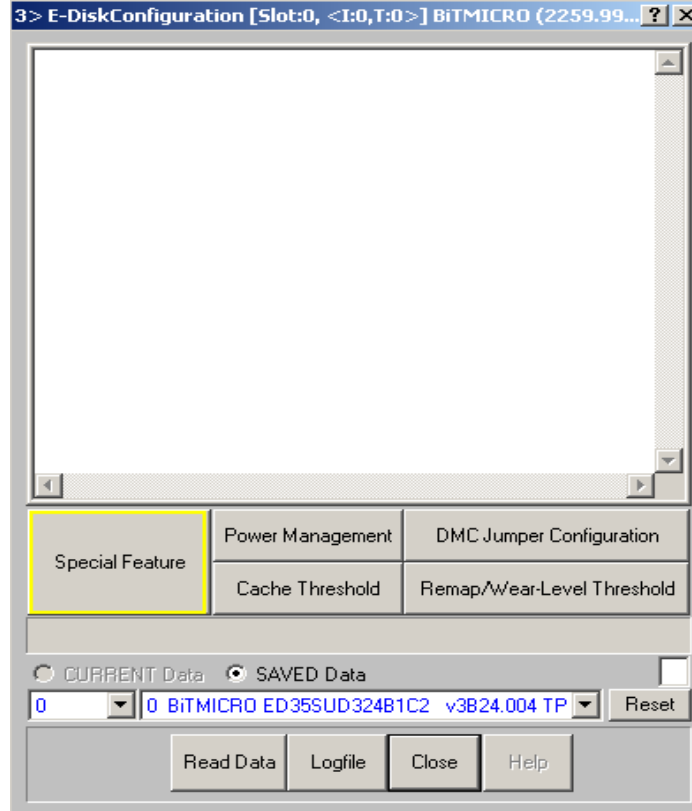
If the E-Disk[®] unit is PowerGuard[®]-enabled and PowerGuard[®] Write Verify is set to NO, write verify commands will complete as if it was only a normal write request.

- Commands affected are all SCSI write verify command variants (CDB 6,10, etc).
- As with synchronize cache, the performance effects will depend on whether the system throws this kind of command.

3. Activating Ignore Cache Disable

The following steps illustrate how to activate “ignore disable cache option”.

- 3.1. From E-Disk[®] Analyzer, go to the E-Disk[®] Configuration. Click the special feature button.



- 3.2. After the special feature button is clicked, E-Disk[®] Analyzer will display the window below. In order to activate “ignore cache disable,” you need to disable the features for PowerGuard[®] FUA, PowerGuard[®] SYNCH, PowerGuard[®] Write Verify by picking “NO” on the drop down list.

