Linaro ltd.

Working with cheap flash drives

Arnd Bergmann

Performance footprint of an SD card

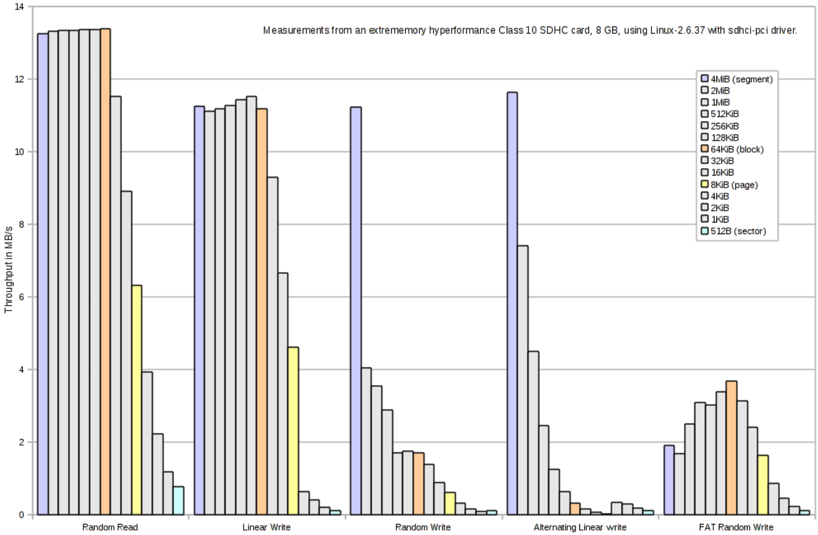

Measurements from an extrememory hyperformance Class 10 SDHC card, 8 GB, using Linux-2.6.37 with sdhci-pci driver.

Read performance scatter graph

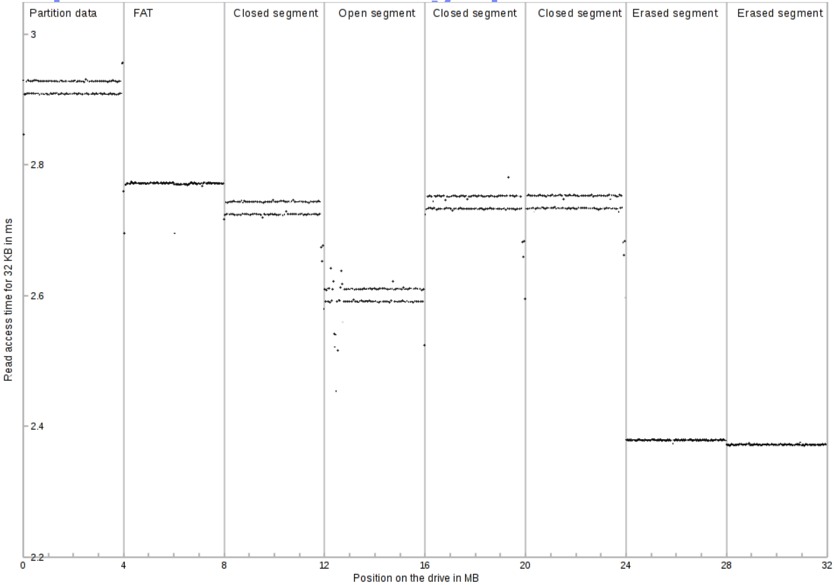

Static wear leveling over erase blocks

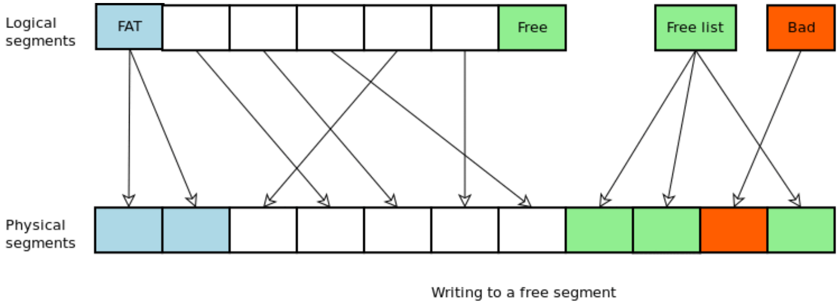

Static wear leveling over erase blocks

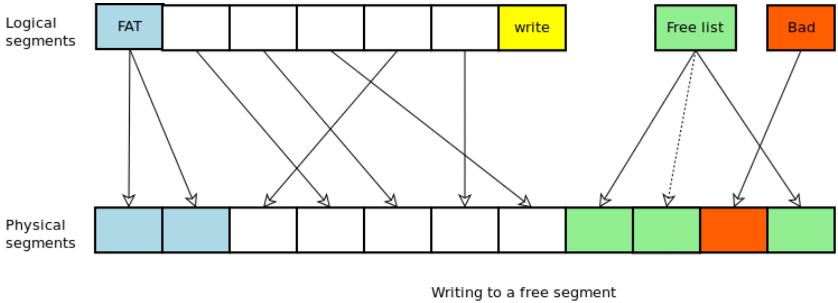

Static wear leveling over erase blocks

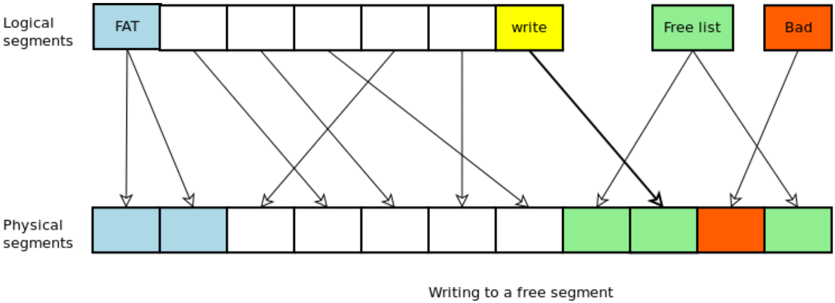

Static wear leveling over erase blocks

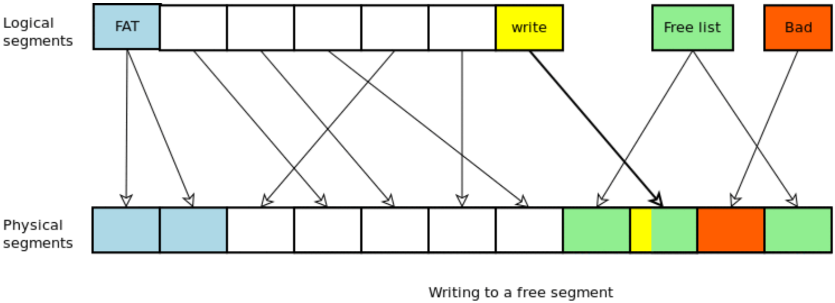

Static wear leveling over erase blocks

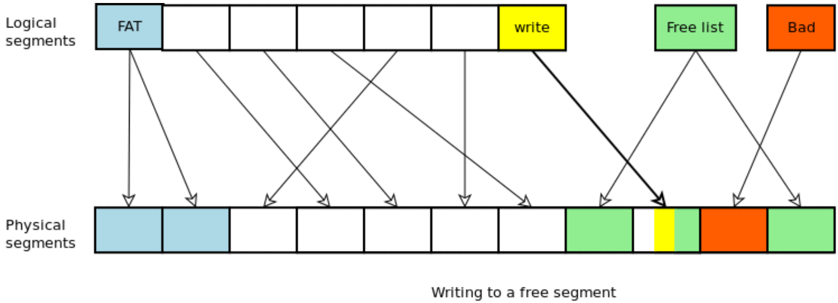

Static wear leveling over erase blocks

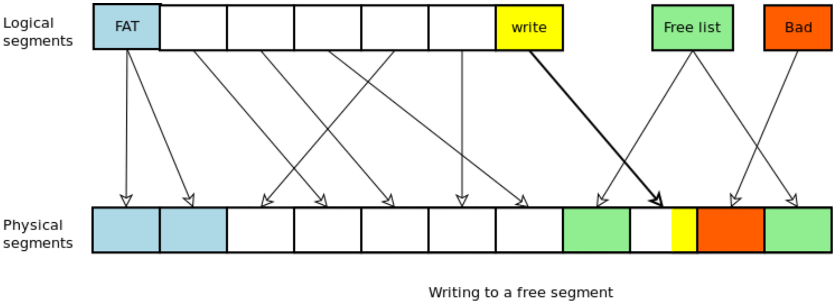

Static wear leveling over erase blocks

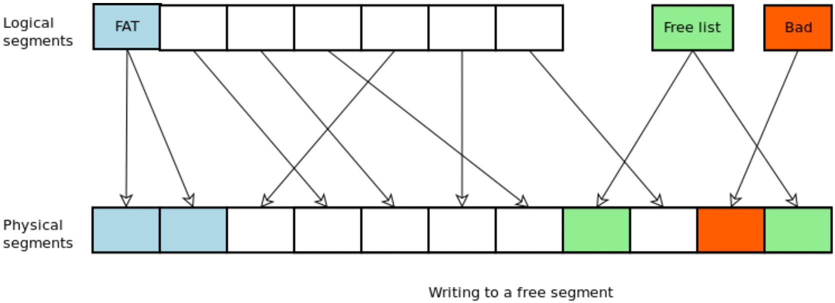

Static wear leveling over erase blocks

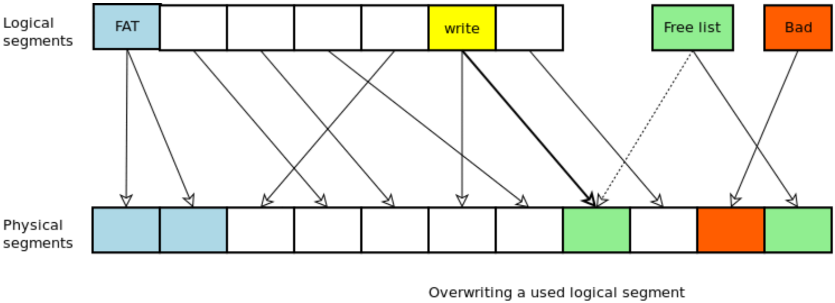

Static wear leveling over erase blocks

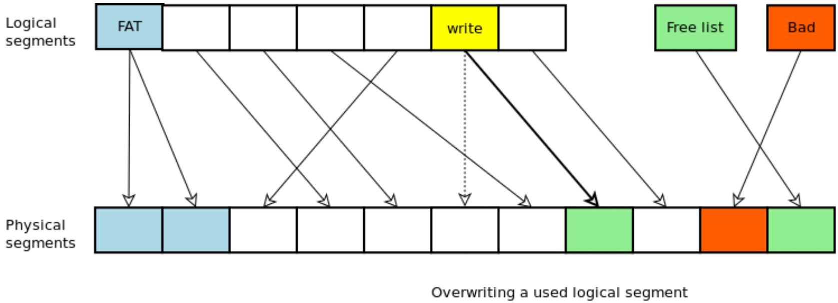

Static wear leveling over erase blocks

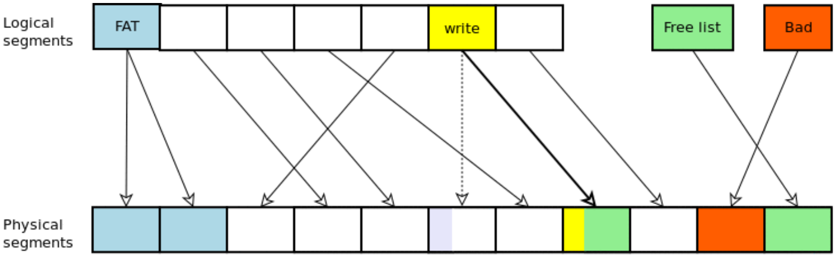

Overwriting a used logical segment

Static wear leveling over erase blocks

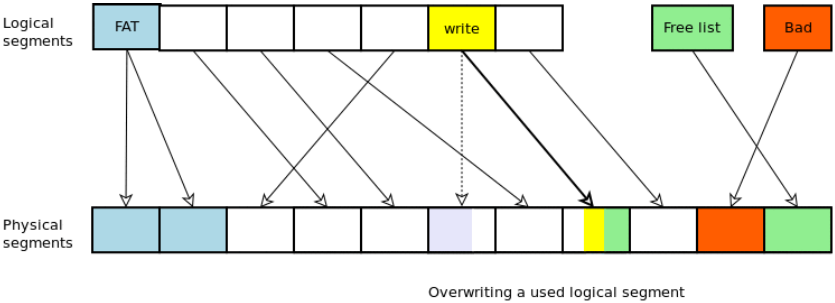

Static wear leveling over erase blocks

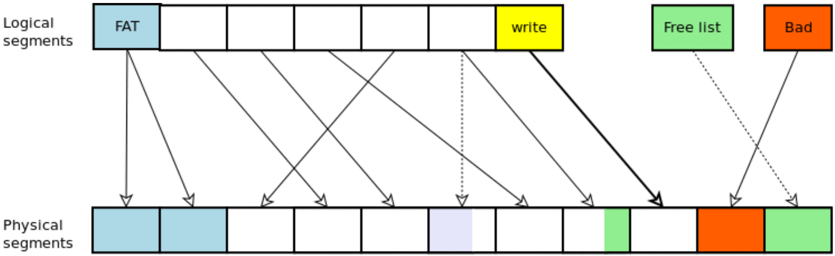

Writing to multiple segments

Static wear leveling over erase blocks

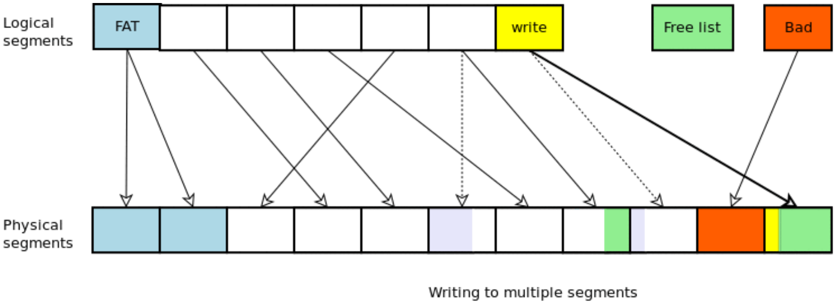

Static wear leveling over erase blocks

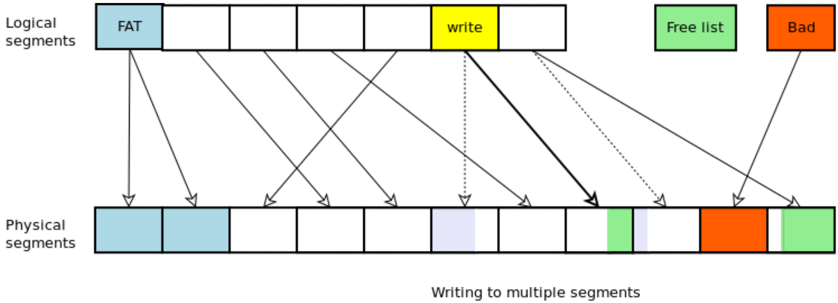

Writing to multiple segments

Static wear leveling over erase blocks

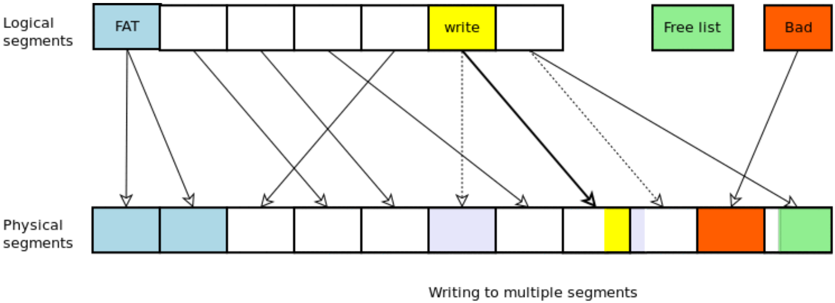

Static wear leveling over erase blocks

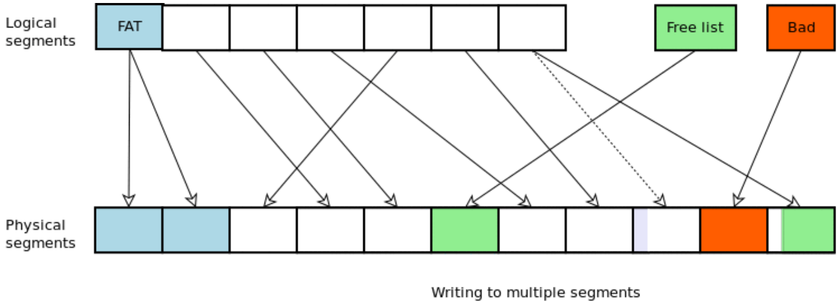

Static wear leveling over erase blocks

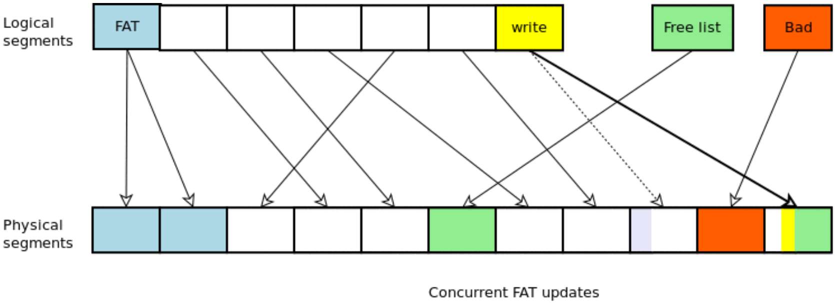

Static wear leveling over erase blocks

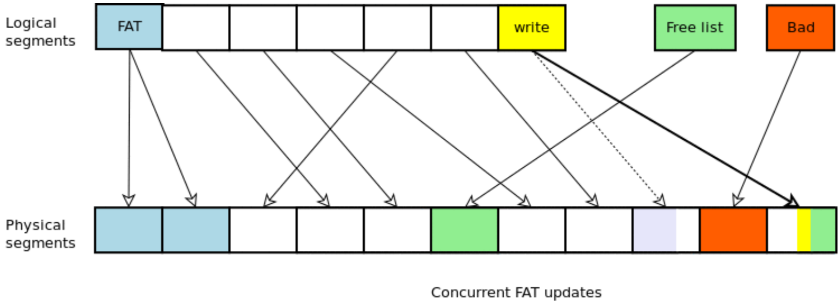

Static wear leveling over erase blocks

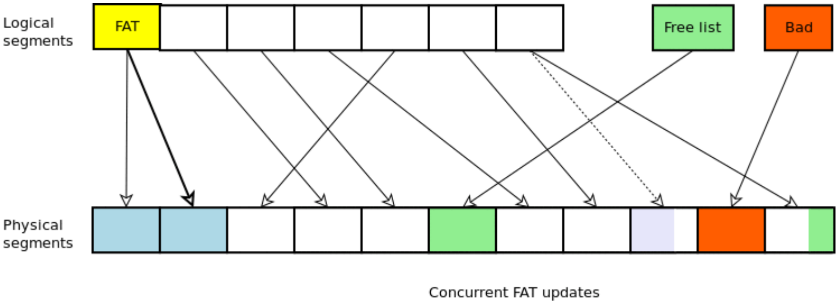

Static wear leveling over erase blocks

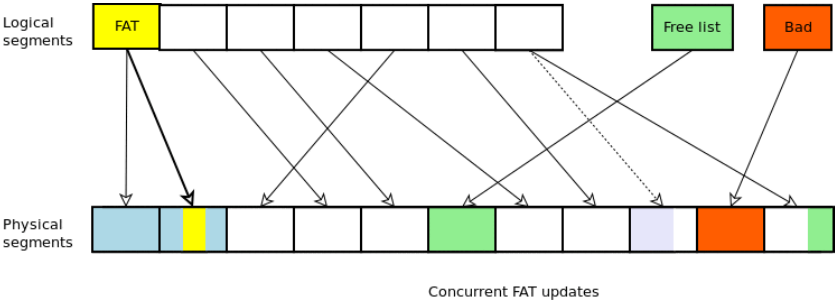

Static wear leveling over erase blocks

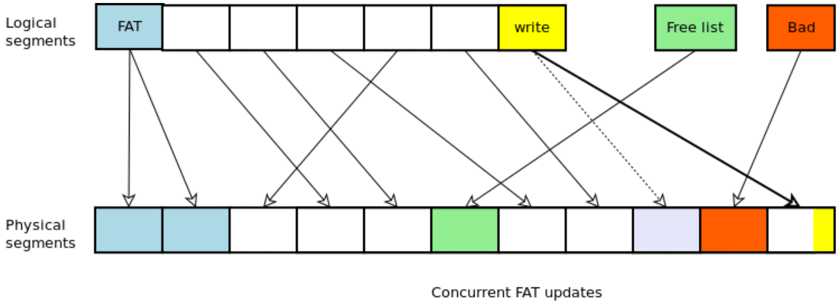

Static wear leveling over erase blocks

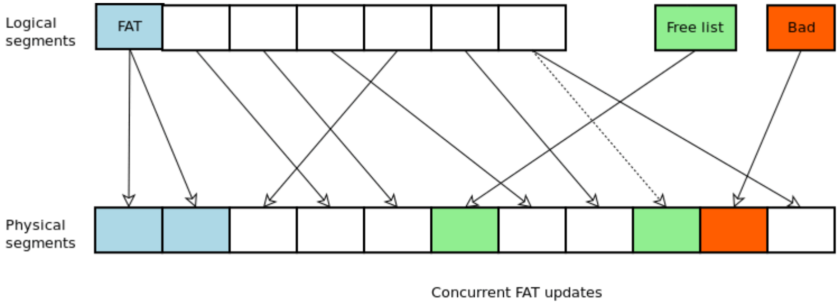

Static wear leveling over erase blocks

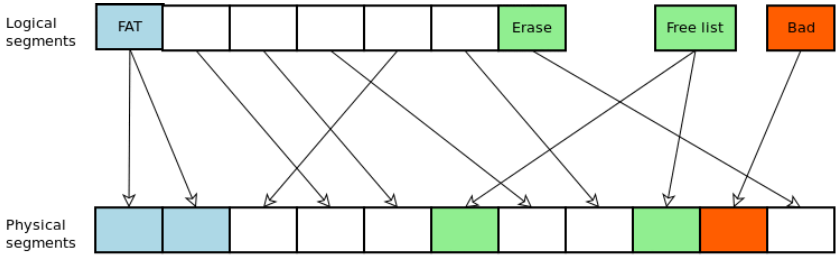

Erasing a segment

Static wear leveling over erase blocks

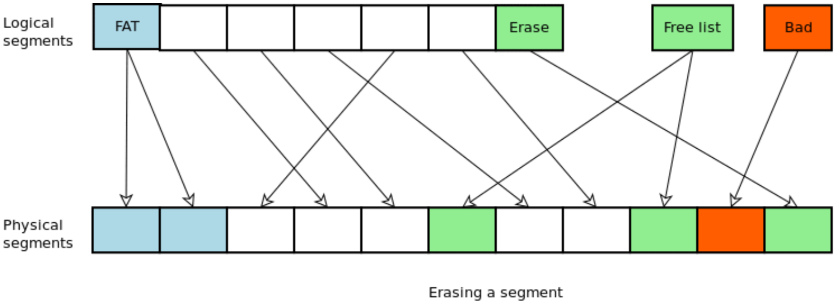

Static wear leveling over erase blocks

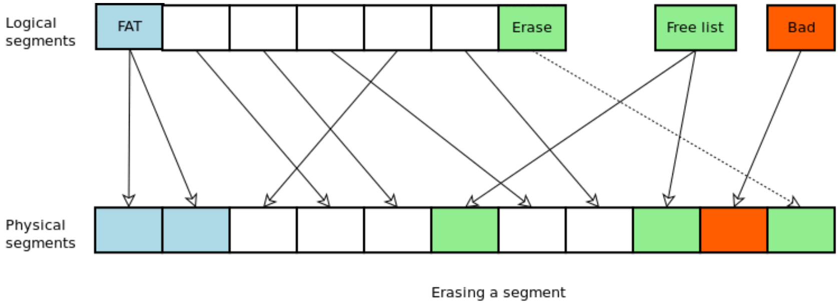

Static wear leveling over erase blocks

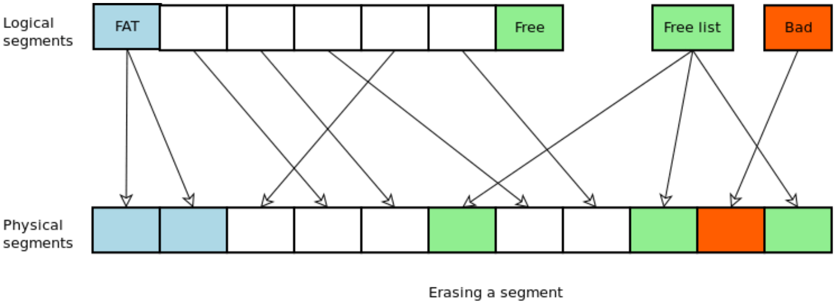

Static wear leveling over erase blocks

Linear access optimized garbage collection

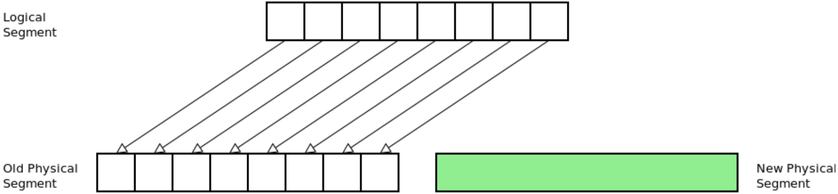

Linear access optimized garbage collection

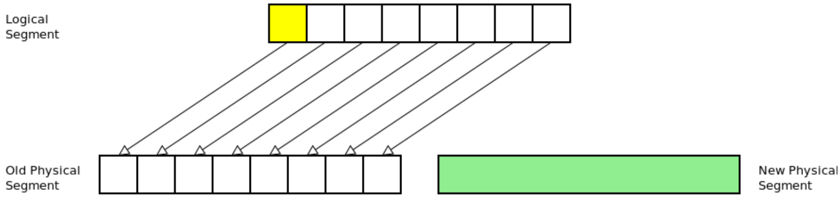

Linear access optimized garbage collection

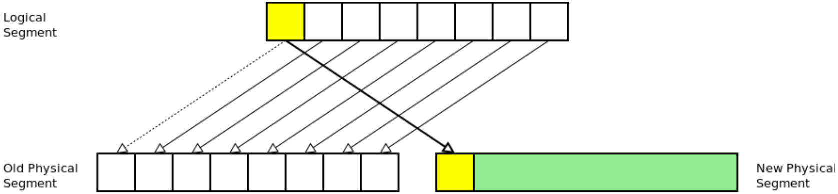

Linear access optimized garbage collection

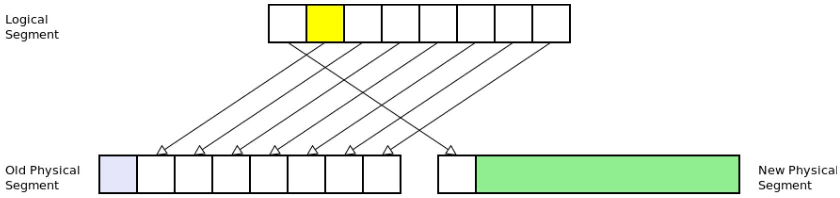

Linear access optimized garbage collection

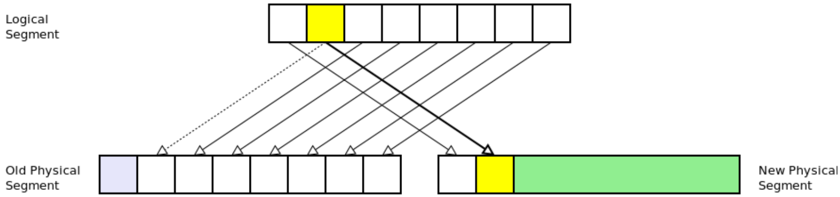

Linear access optimized garbage collection

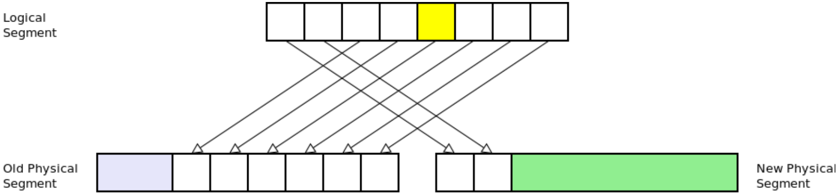

Linear access optimized garbage collection

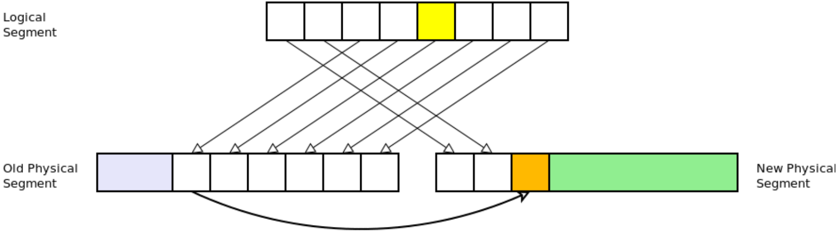

Linear access optimized garbage collection

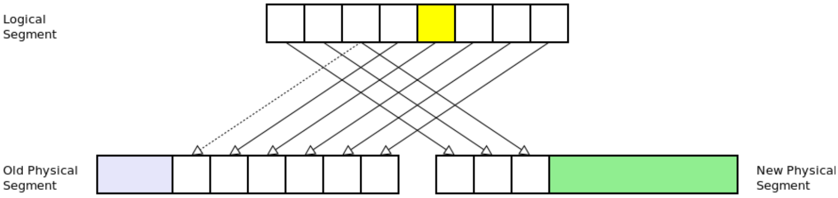

Linear access optimized garbage collection

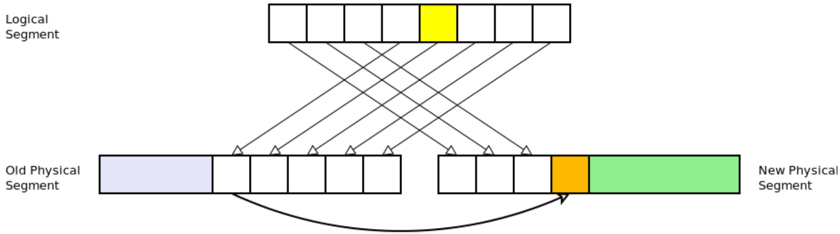

Linear access optimized garbage collection

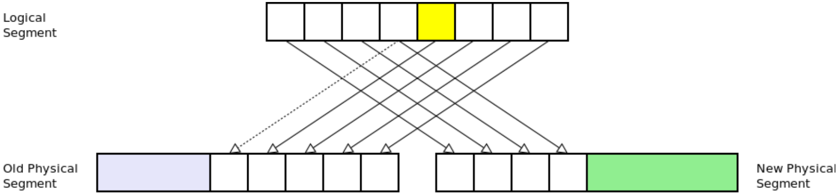

Linear access optimized garbage collection

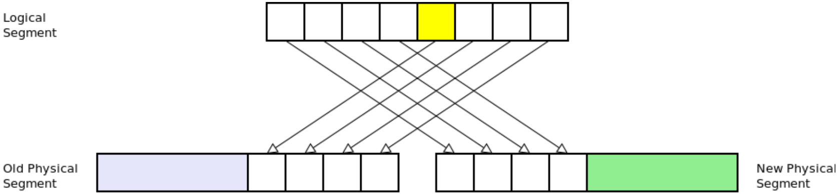

Linear access optimized garbage collection

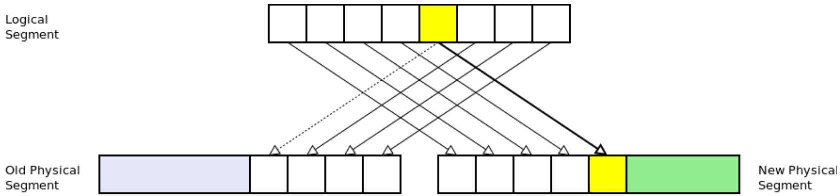

Linear access optimized garbage collection

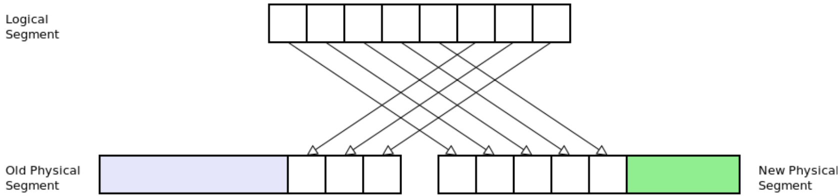

Linear access optimized garbage collection

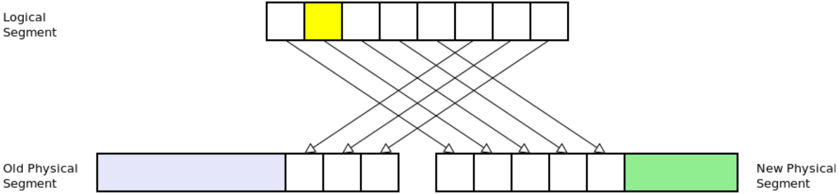

Linear access optimized garbage collection

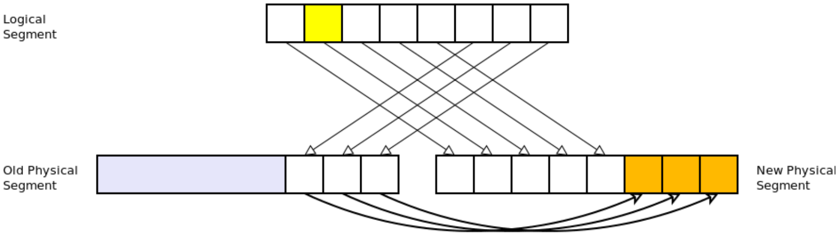

Linear access optimized garbage collection

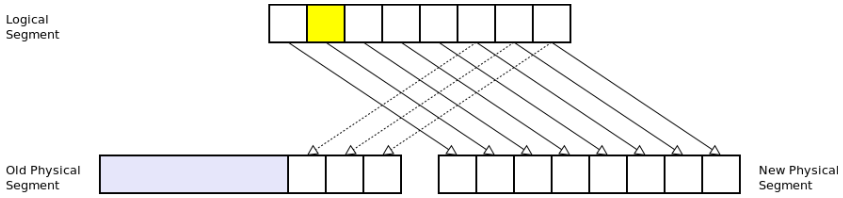

Linear access optimized garbage collection

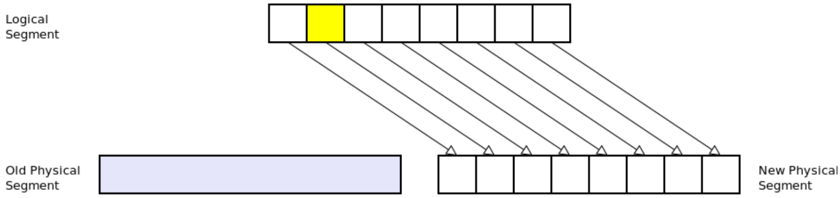

Linear access optimized garbage collection

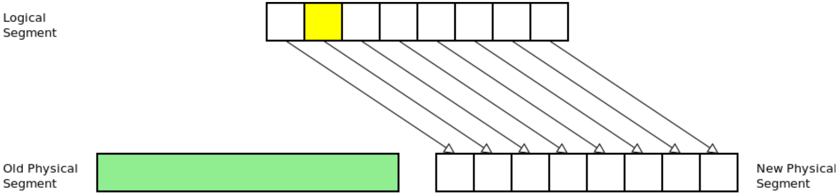

Linear access optimized garbage collection

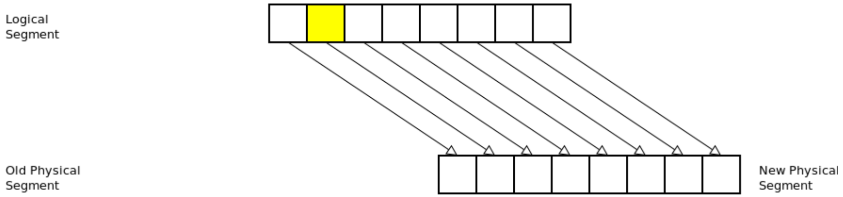

Linear access optimized garbage collection

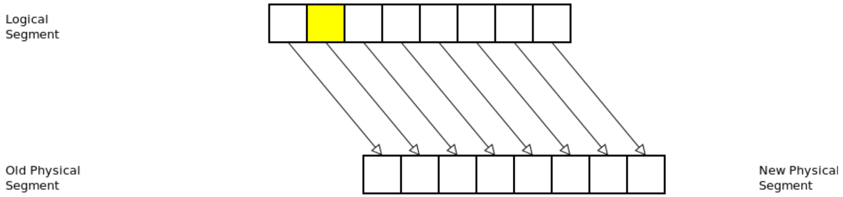

Linear access optimized garbage collection

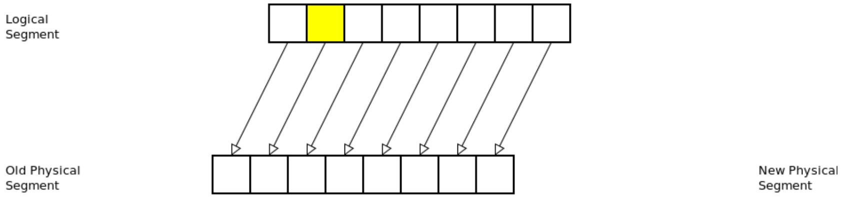

Linear access optimized garbage collection

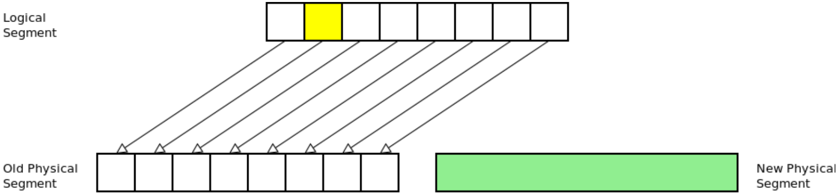

Linear access optimized garbage collection

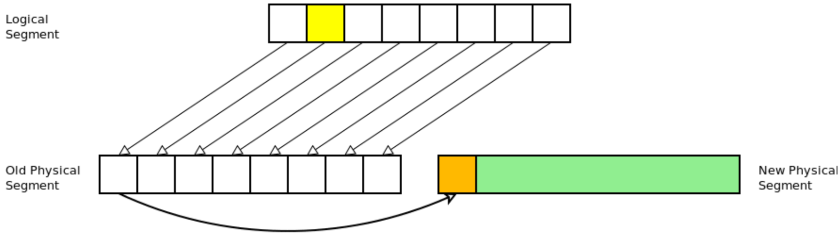

Linear access optimized garbage collection

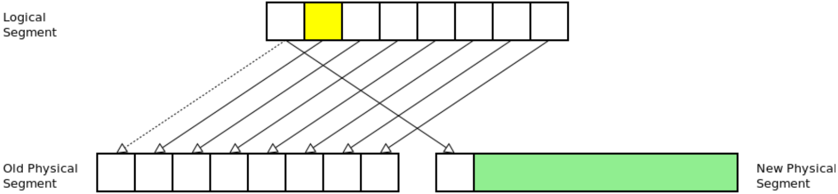

Linear access optimized garbage collection

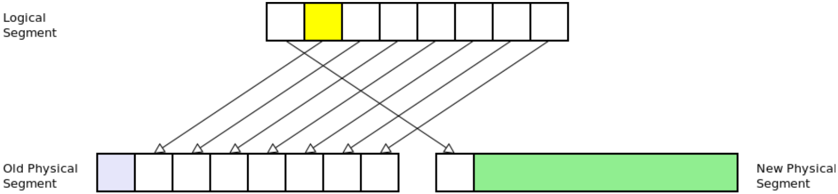

Linear access optimized garbage collection

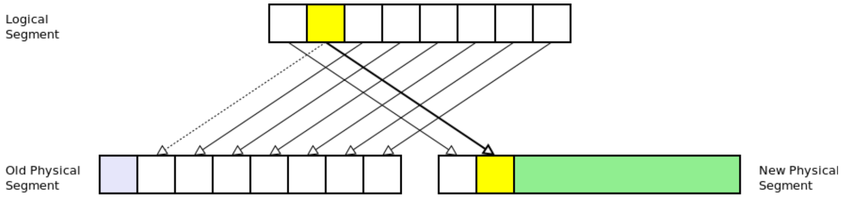

Linear access optimized garbage collection

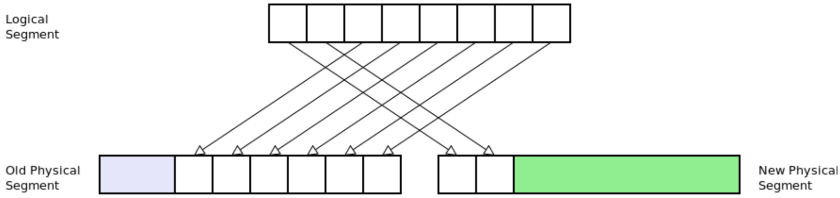

Linear access optimized garbage collection

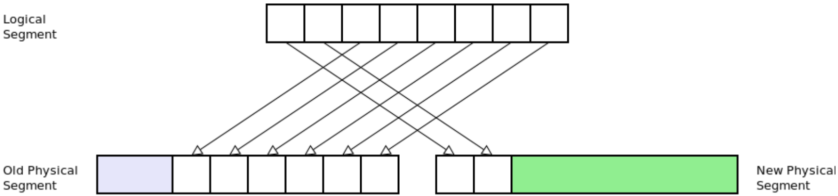

Linear access optimized garbage collection

Garbage collection with block remapping

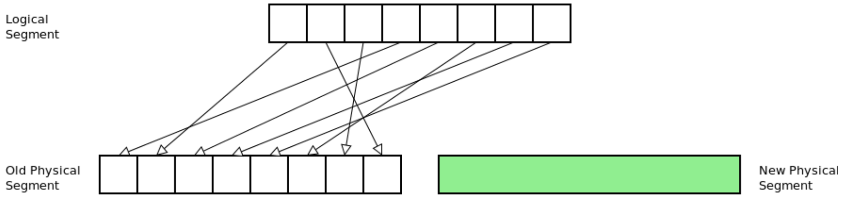

Garbage collection with block remapping

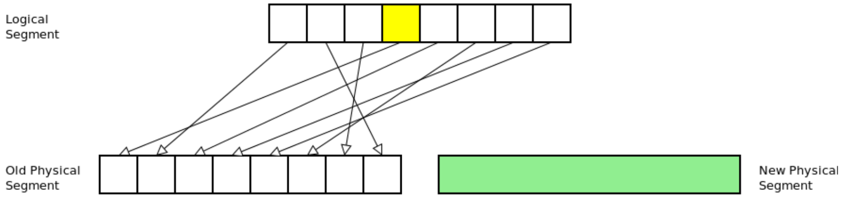

Garbage collection with block remapping

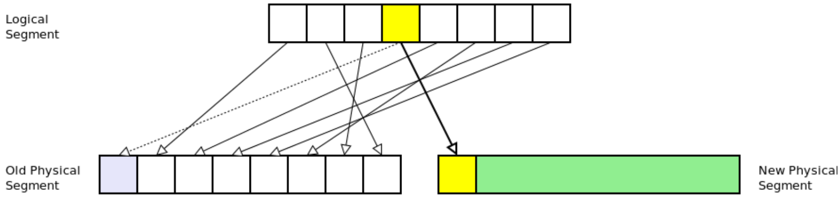

Garbage collection with block remapping

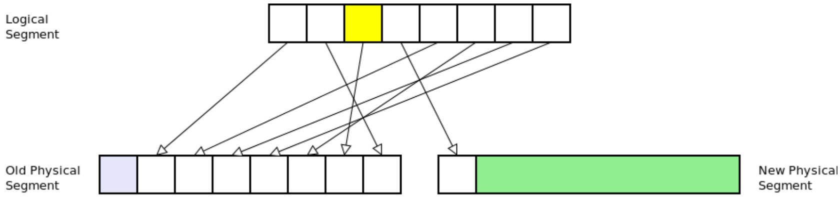

Garbage collection with block remapping

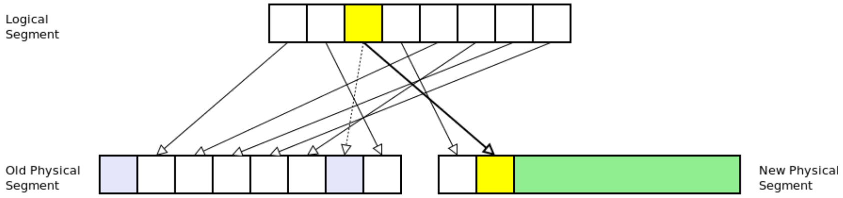

Garbage collection with block remapping

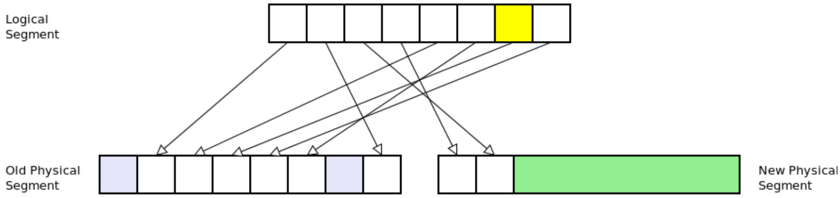

Garbage collection with block remapping

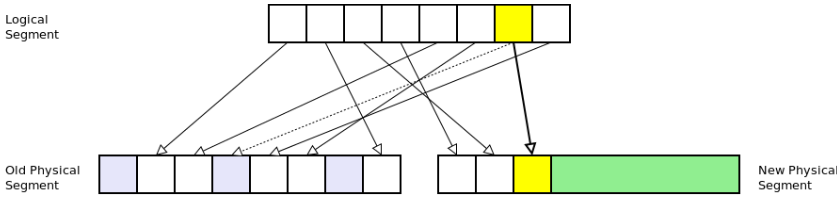

Garbage collection with block remapping

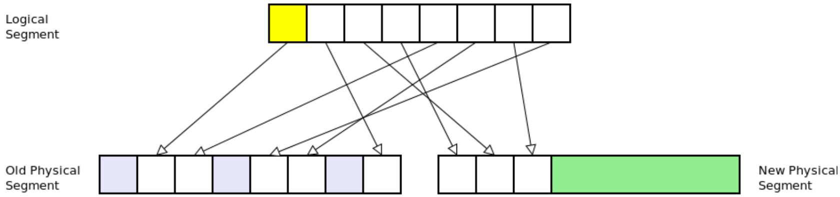

Garbage collection with block remapping

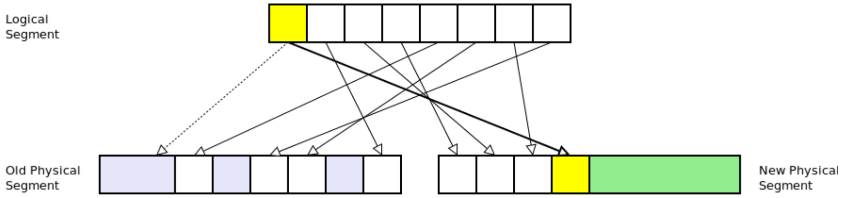

Garbage collection with block remapping

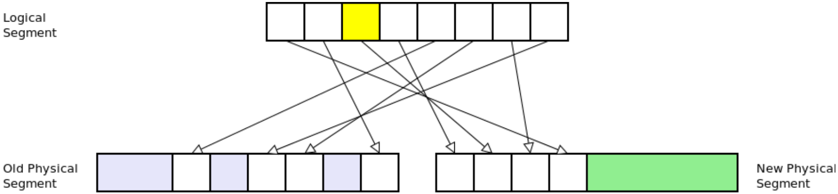

Garbage collection with block remapping

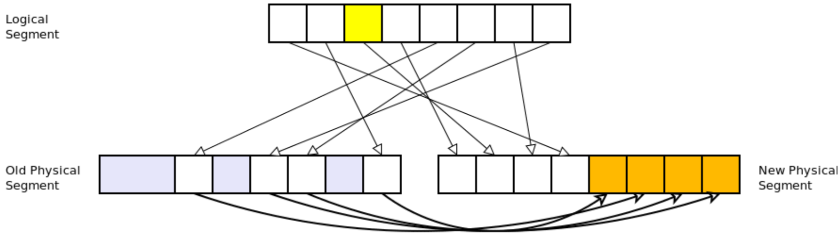

Garbage collection with block remapping

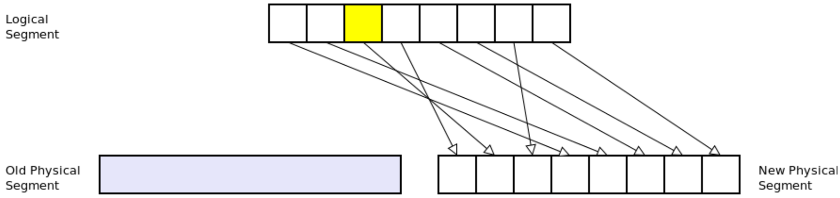

Garbage collection with block remapping

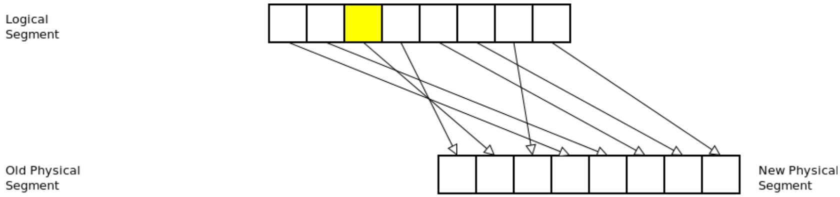

Garbage collection with block remapping

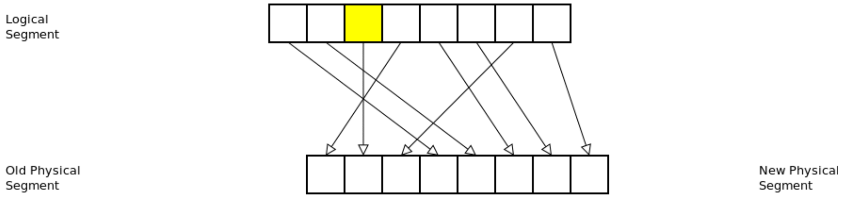

Garbage collection with block remapping

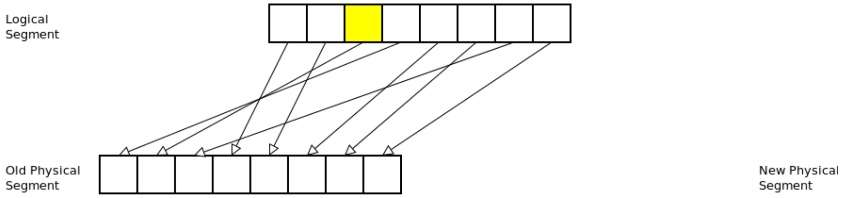

Garbage collection with block remapping

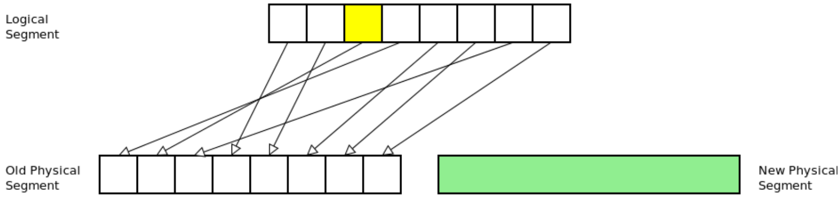

Garbage collection with block remapping

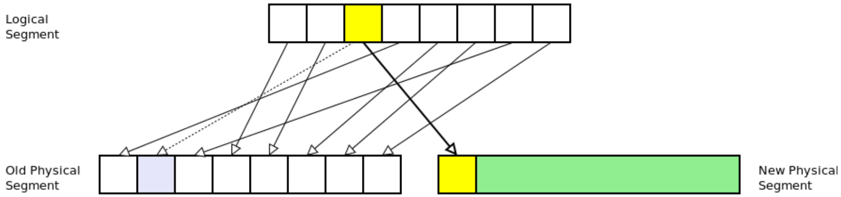

Garbage collection with block remapping

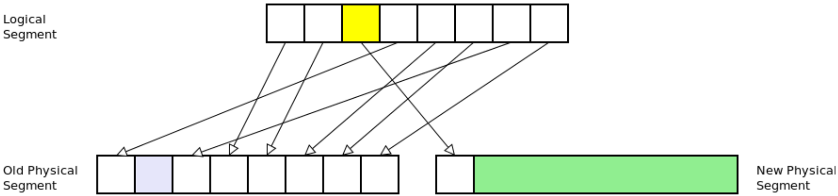

Garbage collection with block remapping

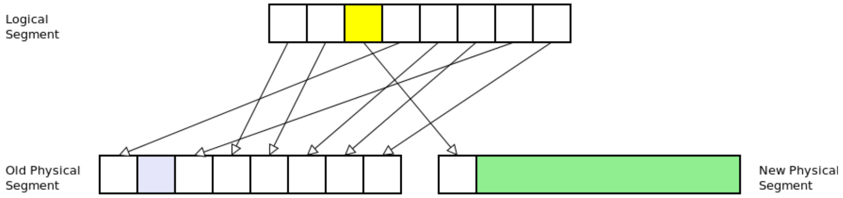

Garbage collection with block remapping

Data logging garbage collection

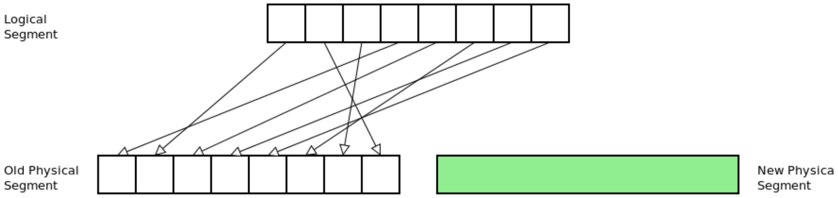

Data logging garbage collection

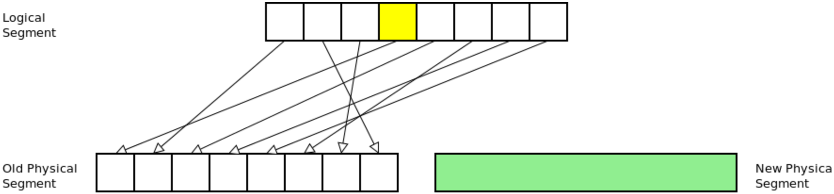

Data logging garbage collection

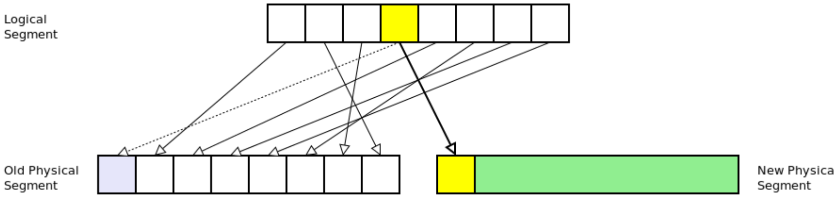

Data logging garbage collection

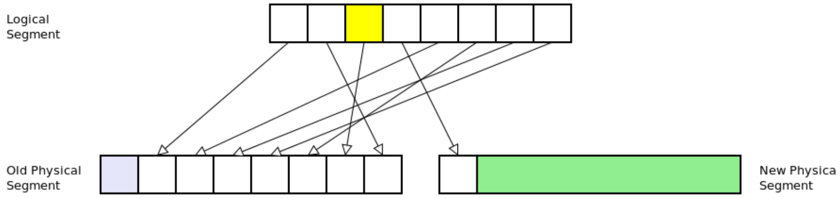

Data logging garbage collection

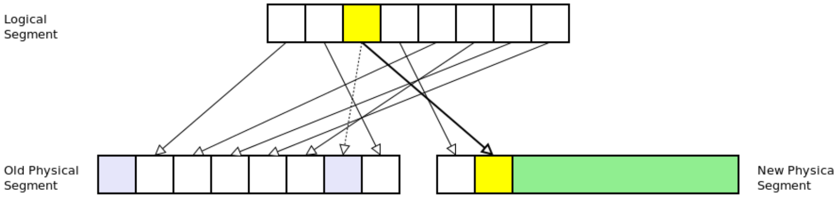

Data logging garbage collection

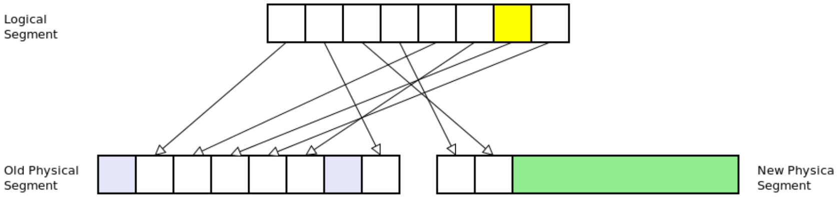

Data logging garbage collection

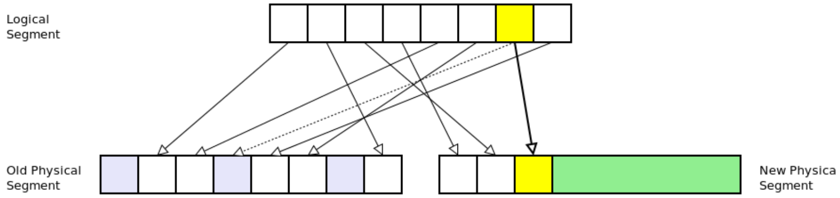

Data logging garbage collection

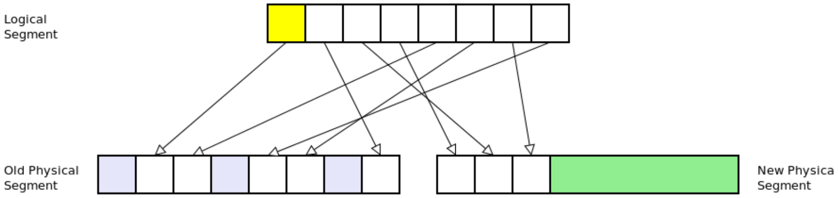

Data logging garbage collection

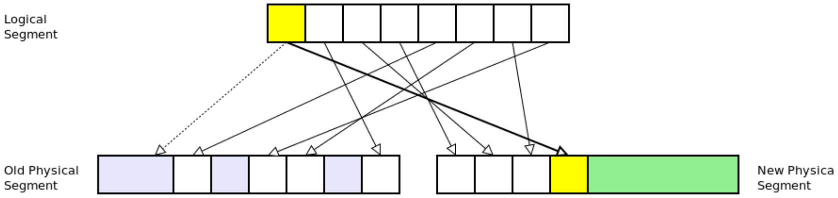

Data logging garbage collection

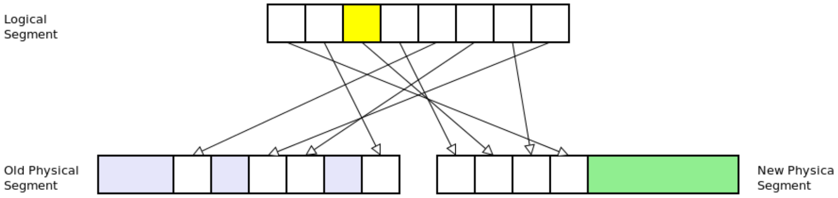

Data logging garbage collection

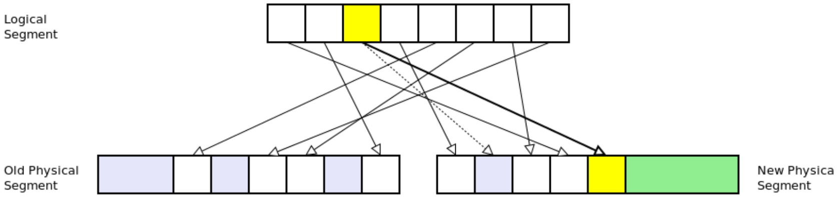

Data logging garbage collection

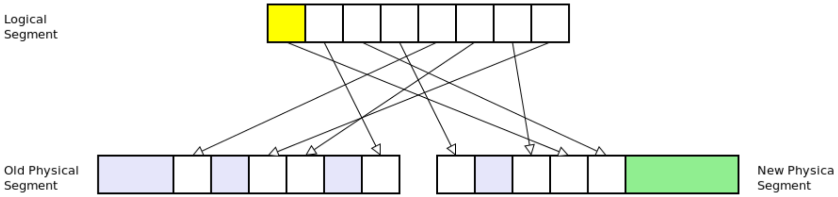

Data logging garbage collection

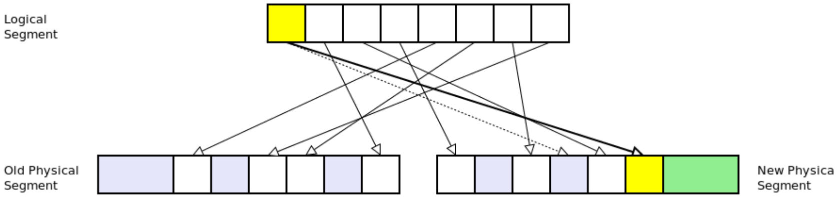

Data logging garbage collection

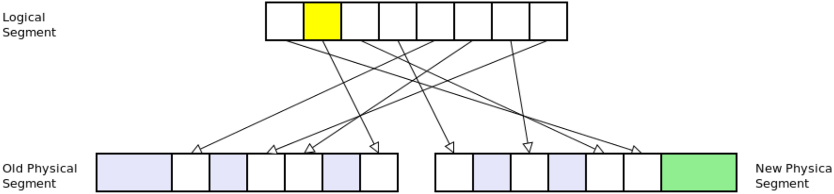

Data logging garbage collection

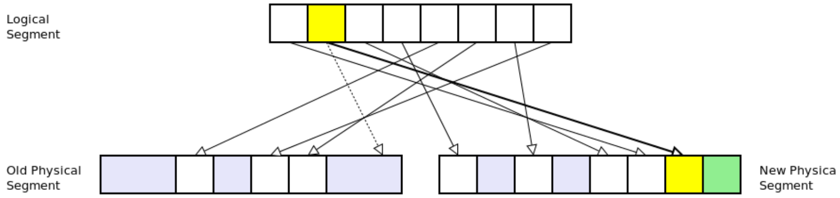

Data logging garbage collection

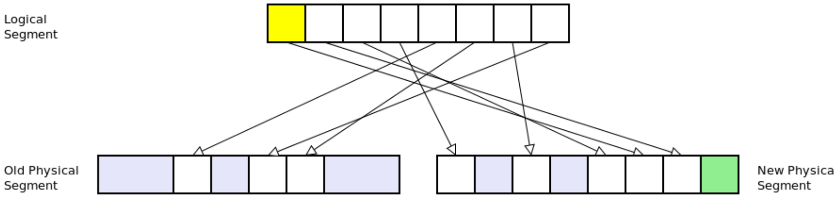

Data logging garbage collection

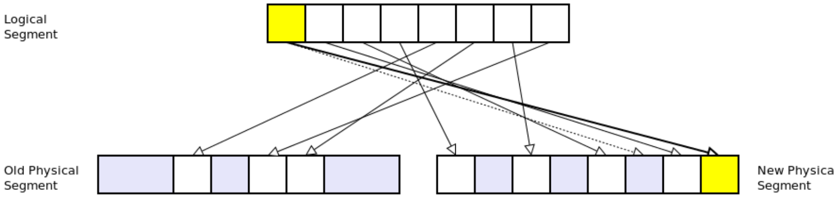

Data logging garbage collection

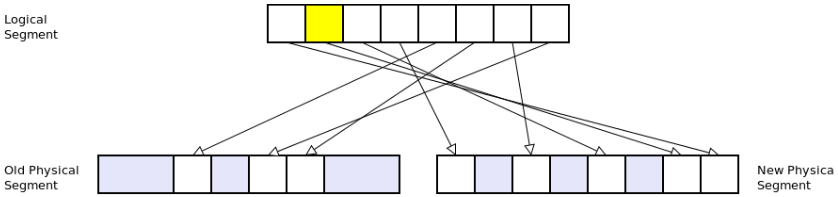

Data logging garbage collection

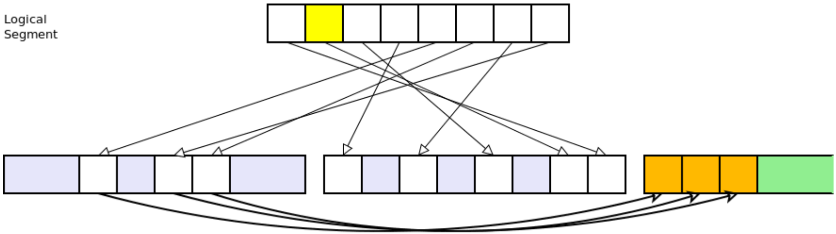

Data logging garbage collection

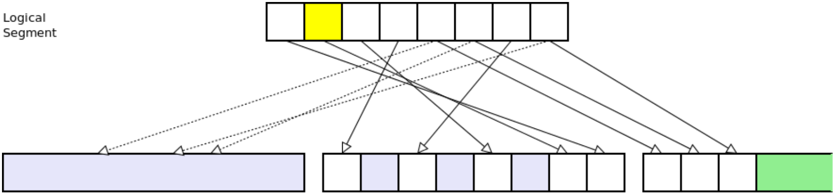

Data logging garbage collection

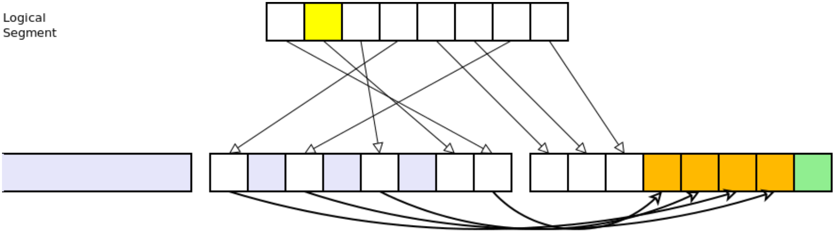

Data logging garbage collection

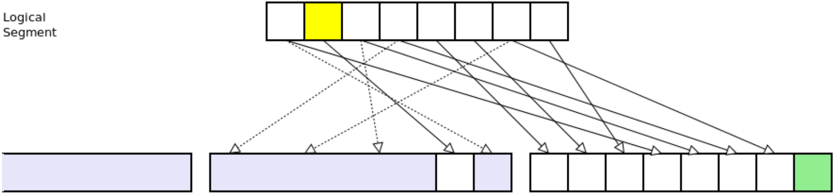

Data logging garbage collection

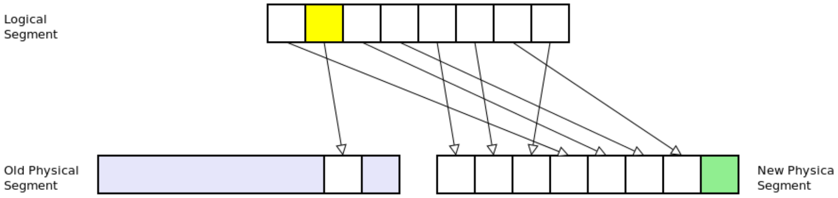

Data logging garbage collection

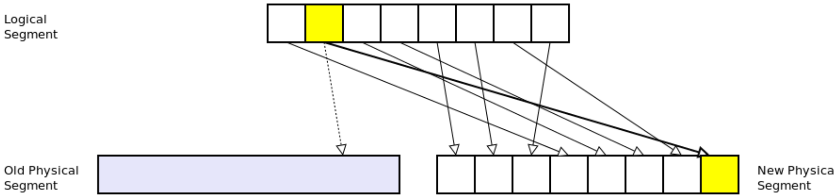

Data logging garbage collection

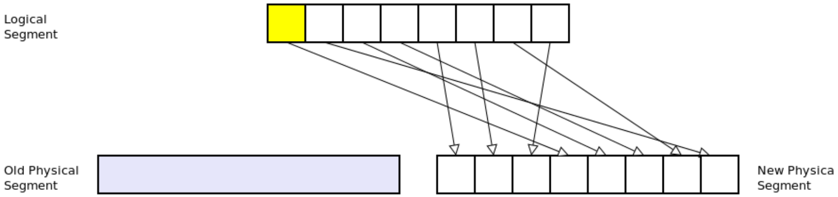

Data logging garbage collection

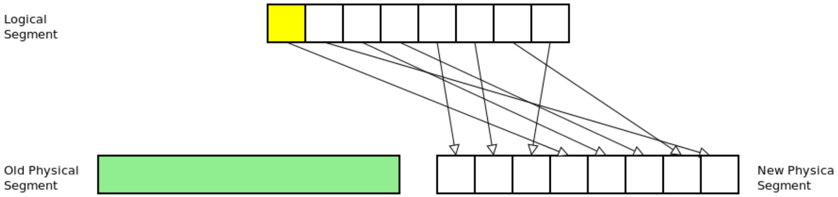

Data logging garbage collection

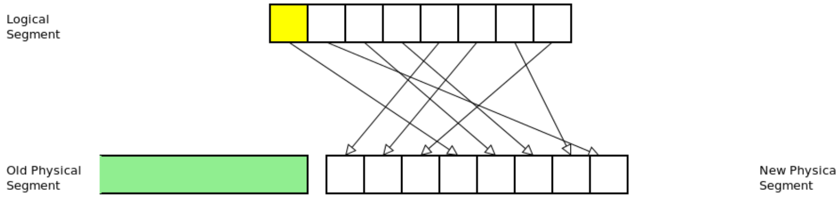

Data logging garbage collection

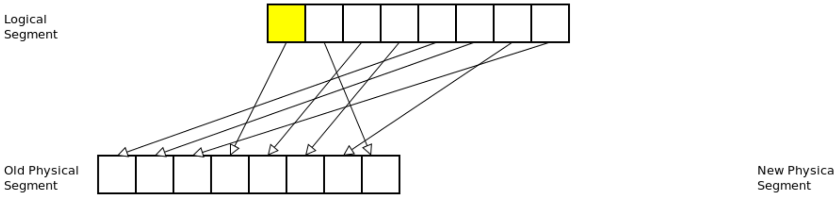

Data logging garbage collection

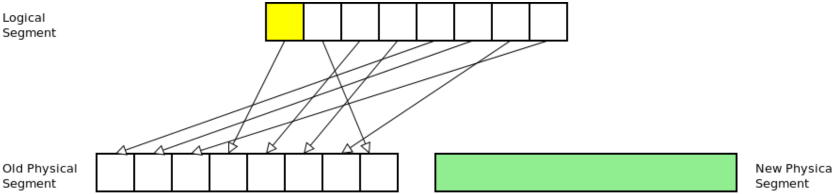

Data logging garbage collection

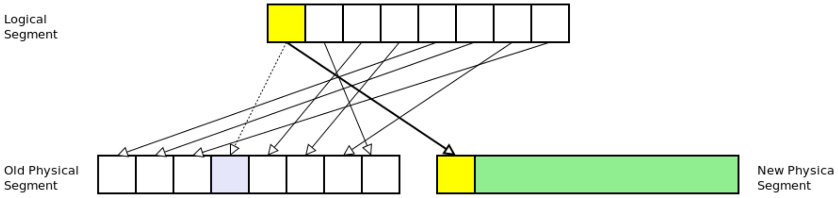

Data logging garbage collection

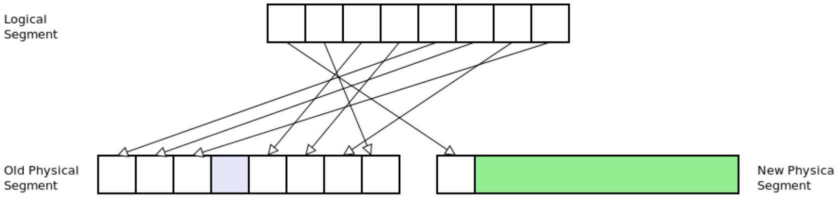

Data logging garbage collection

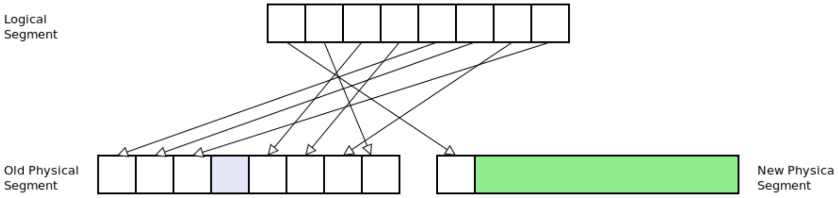

Data logging garbage collection

Running out of open erase blocks

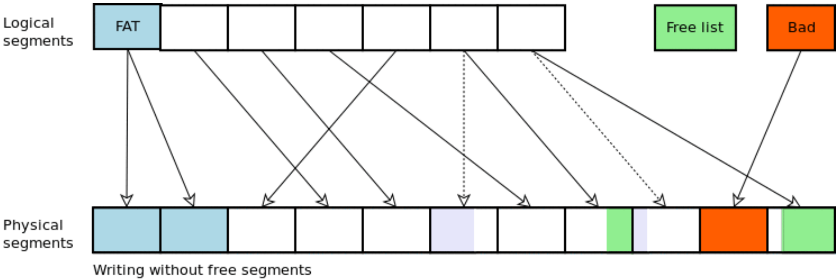

Running out of open erase blocks

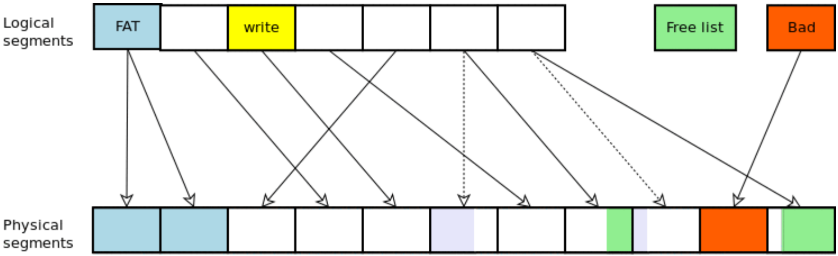

Writing without free segments: user writes to a logical segment that is not open

Running out of open erase blocks

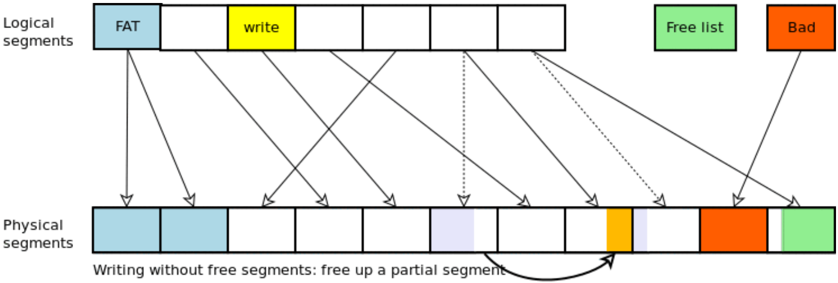

Running out of open erase blocks

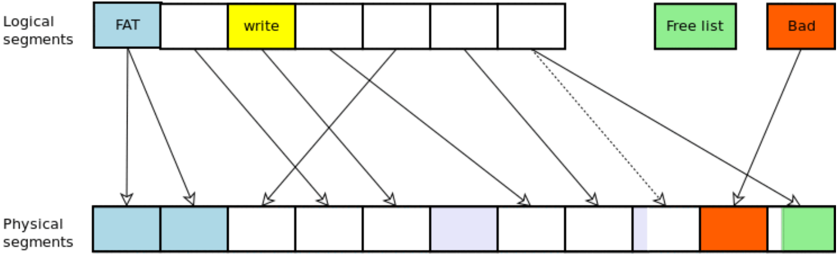

Writing without free segments: one segment is now unreferenced and can be erased

Running out of open erase blocks

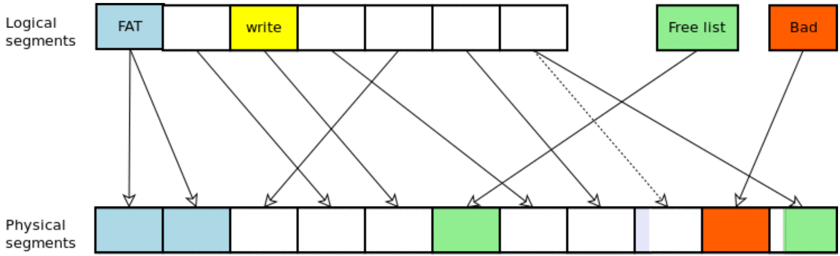

Writing without free segments: one segment is now erased and can be reused

Running out of open erase blocks

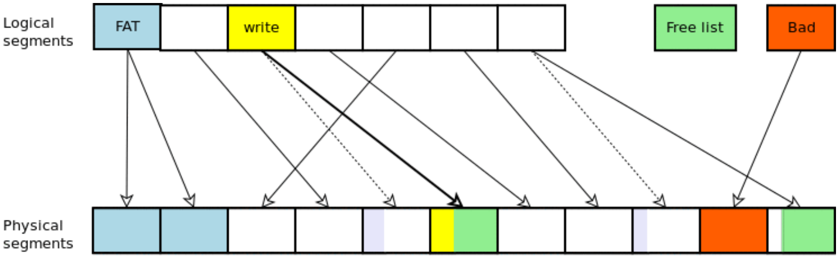

Writing without free segments: start writing to the segment just freed up

Running out of open erase blocks

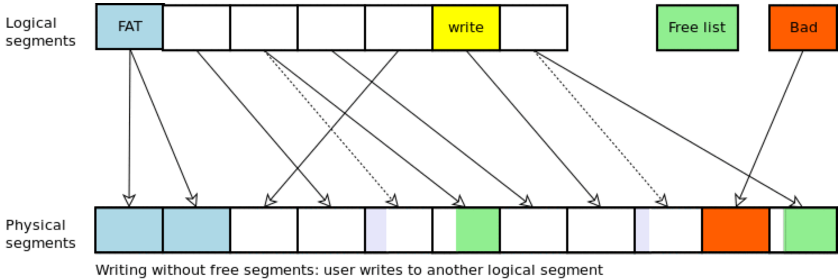
Logical segments
FAT
write
Free list
Bad
Physical segments
Writing without free segments: user writes to another logical segment

Running out of open erase blocks

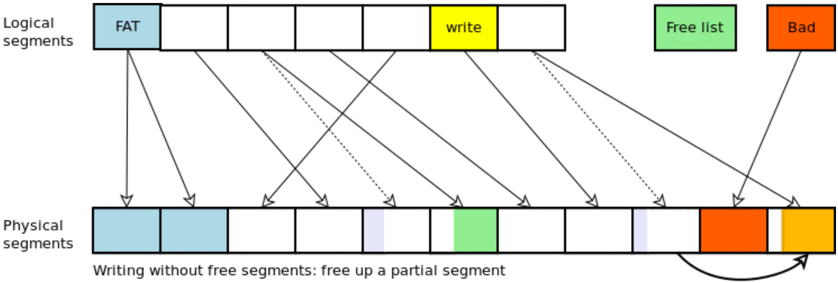

Writing without free segments: free up a partial segment

Running out of open erase blocks

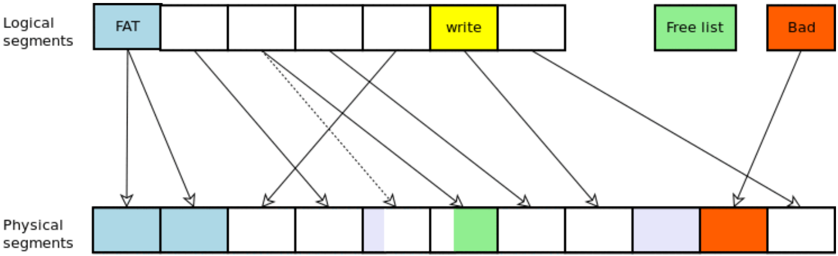

Writing without free segments: one segment is now unreferenced and can be erased

Running out of open erase blocks

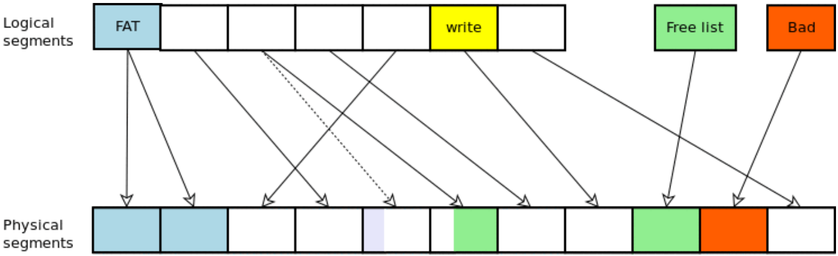

Writing without free segments: one segment is now erased and can be reused

Running out of open erase blocks

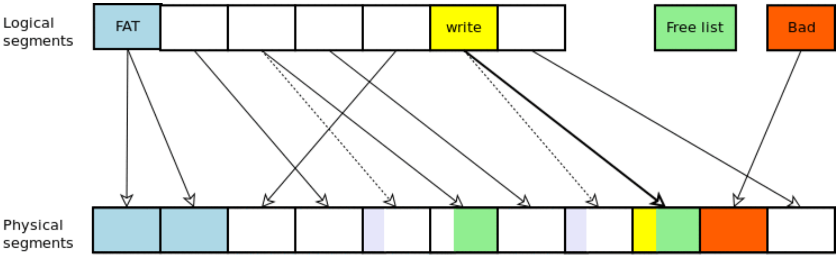

Writing without free segments: start writing to the segment just freed up

Running out of open erase blocks

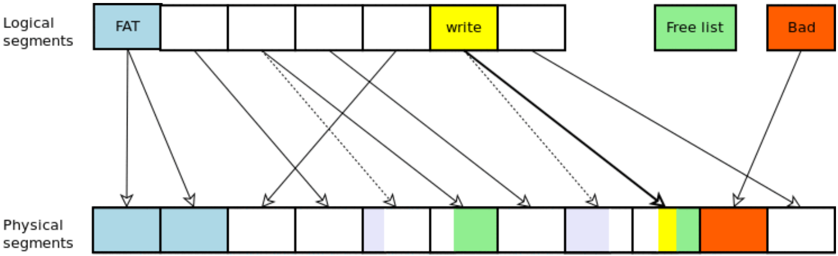

Writing without free segments: start writing to the segment just freed up

Running out of open erase blocks

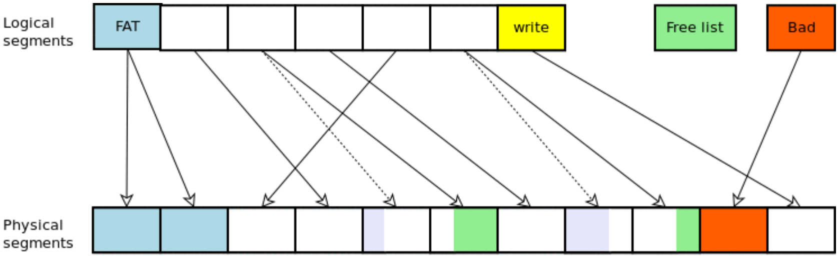

Writing without free segments: user writes to another logical segment

Running out of open erase blocks

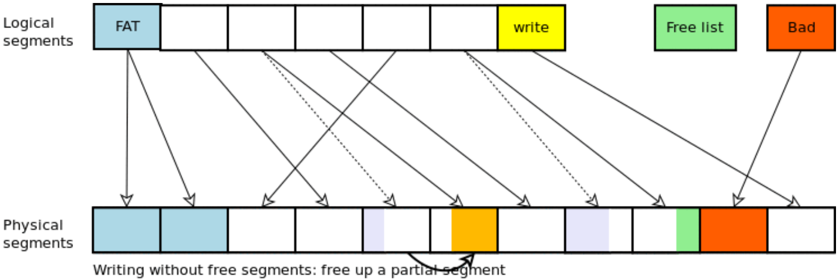

Running out of open erase blocks

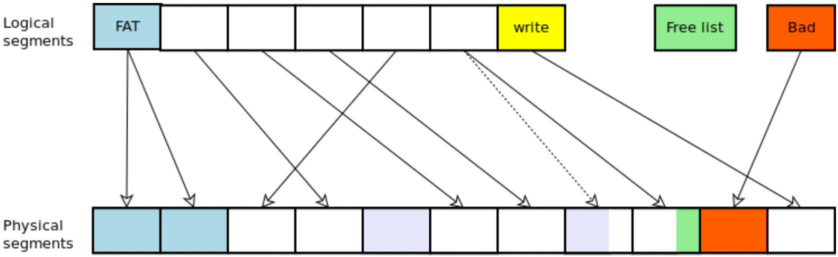

Writing without free segments: one segment is now unreferenced and can be erased

Running out of open erase blocks

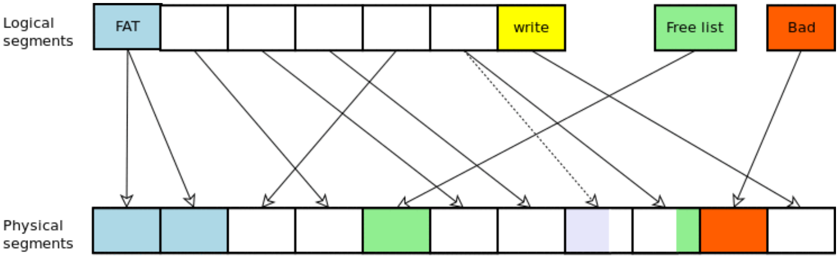

Writing without free segments: one segment is now erased and can be reused

Running out of open erase blocks

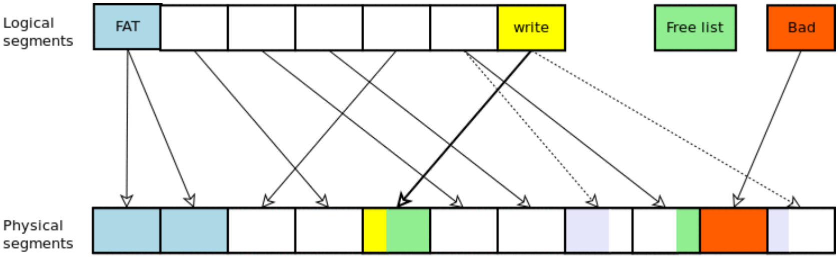

Writing without free segments: start writing to the segment just freed up

Running out of open erase blocks

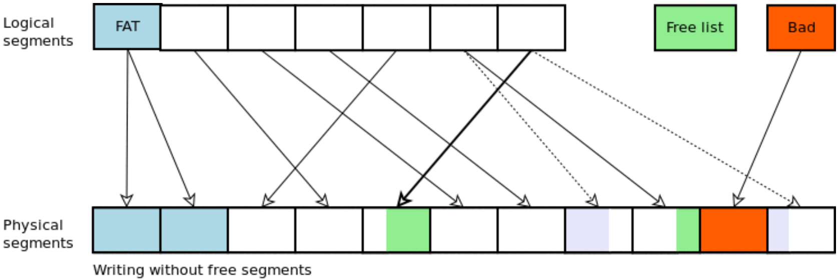

Writing without free segments

Running out of open erase blocks

Performance footprint of an SD card

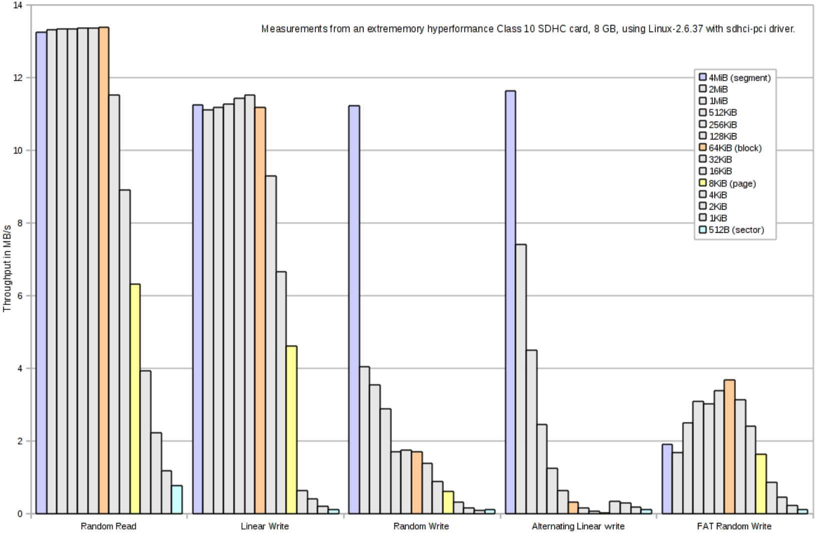

Flashbench demo

Planned optimizations

- Superpage buffer cache

Planned optimizations

- Superpage buffer cache
- Flashcache device mapper target

Planned optimizations

- Superpage buffer cache
- Flashcache device mapper target
- Erase block optimizations in elevator

Planned optimizations

- Superpage buffer cache
- Flashcache device mapper target
- Erase block optimizations in elevator
- File system block allocation

Legal Statement

This work represents the view of the author and does not necessarily represent the view of IBM.

IBM, IBM (logo), e-business (logo), pSeries, e (logo) server, and xSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States and/or other countries.

Linux is a registered trademark of Linus Torvalds.

Other company, product, and service names may be trademarks or service marks of others.

Questions?