SDC 2015 - Host Managed SMR

Albert Chen • Jim Malina • TK Kato

absolutely™ WD

Shingled Magnetic Recording

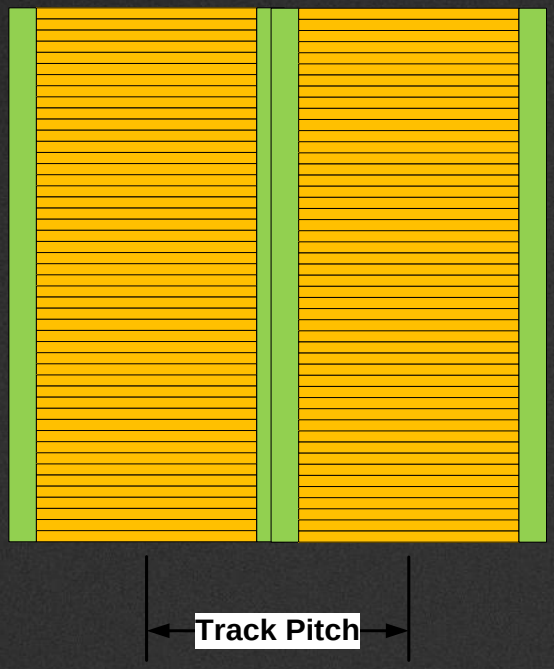

Random Write
Random Read

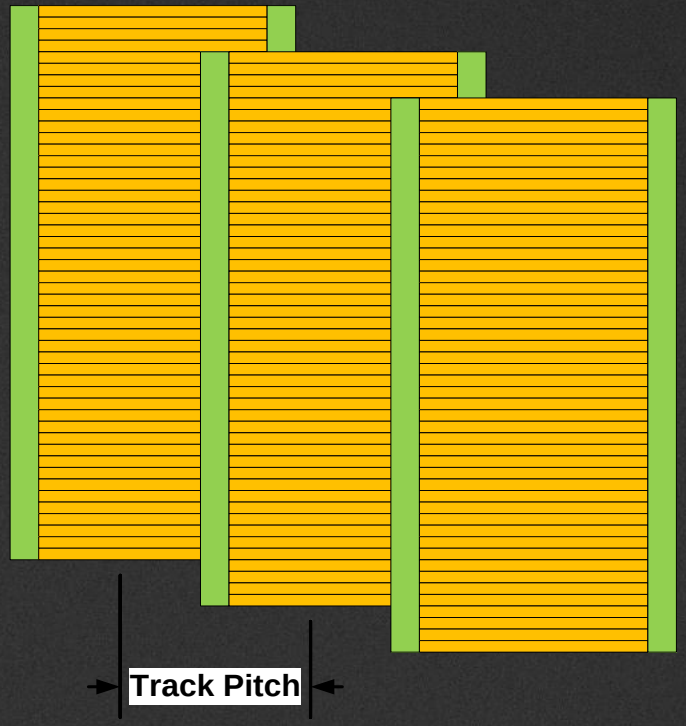

Sequential Write
Random Read

Abstraction Layers

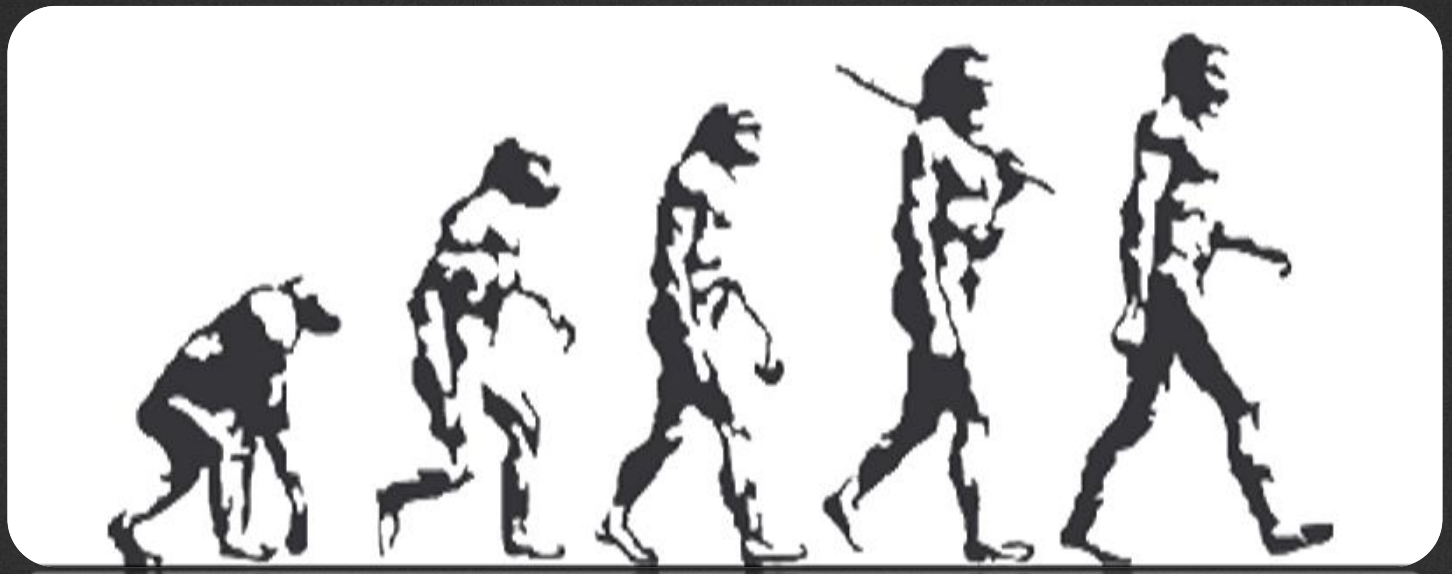

Firmware | Driver | Storage Stack | File system | Application

absolutely™ WD

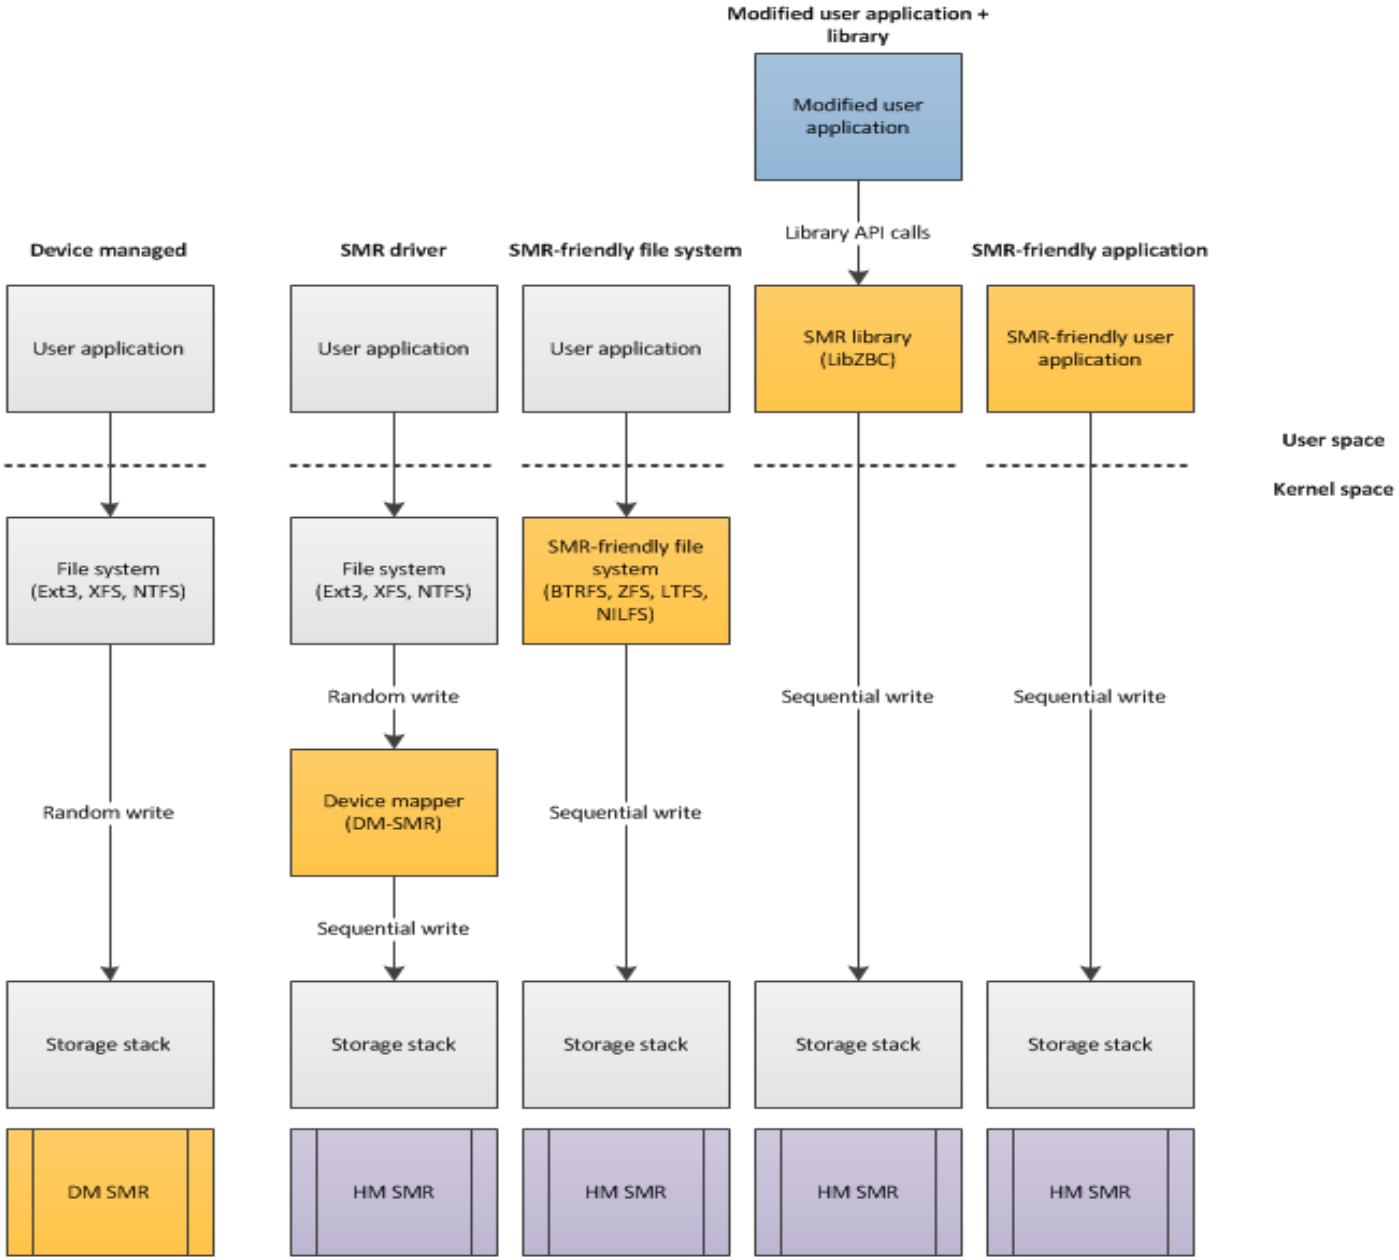

http://hselin.com/blog/2015/08/26/another-layer-of-indirection

For an in-depth presentation on how to handle host managed/aware SMR devices.

Wednesday 1:00 – 1:50
Strategies for Using Standard File systems on SMR Drives
Dr. Hannes Reinecke
SUSE

SMR Type	Drive Managed	Host Managed
Method	Handle random writes	Sequential write
Usage scenario	Client	Data center/Surveillance/24x7
Pros	Plug & play	More efficient data mgmt
Cons	Higher drive complexity Unpredictable performance	SW support required Not backwards compatible

Faster drive development

Access to application & system level semantics

Predictable drive performance

Scale with host HW

Fewer Drive Resources

Easier to manage

Zone Block Command

+ Zone ATA Command

Host Managed SMR

absolutely™ WD

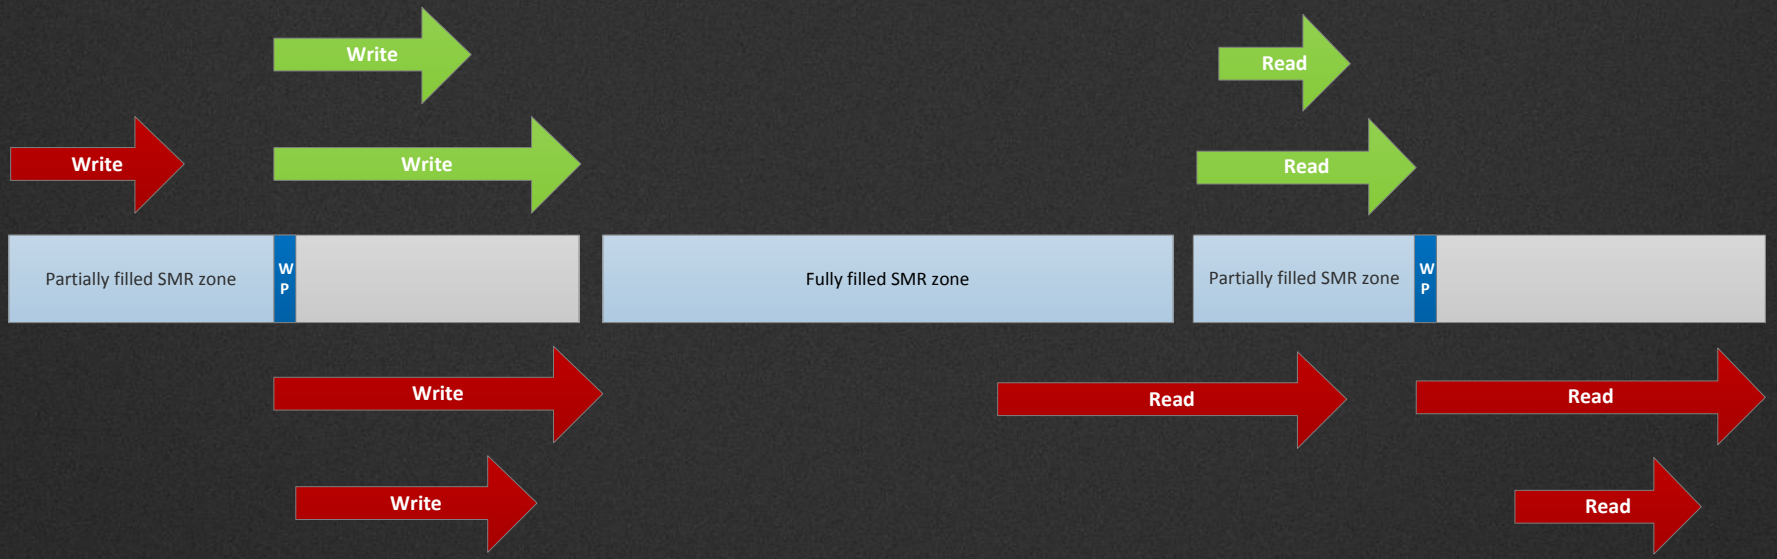

SMR Device Mapper

Acts as media-cache outside the drive with more robust resources available from the host

Write Log: “Media cache” for random writes consists of SMR and CMR zone(s)
Write Buffer: Descriptor for allocated data in Write Log
Backstore: SMR zones that makes up advertised capacity (subsumes Write Buffers during GC)

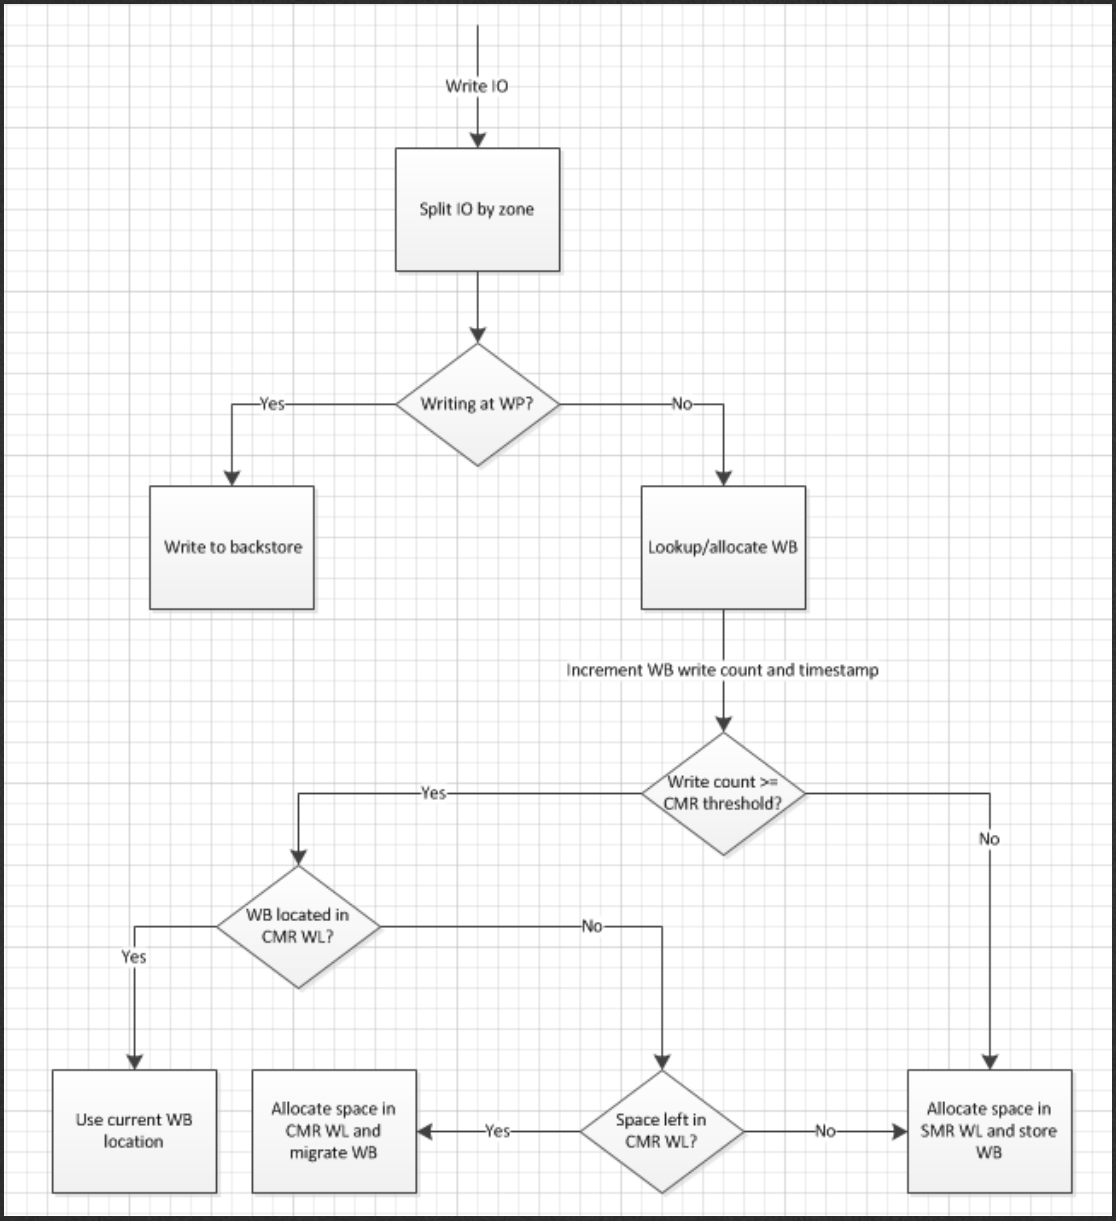

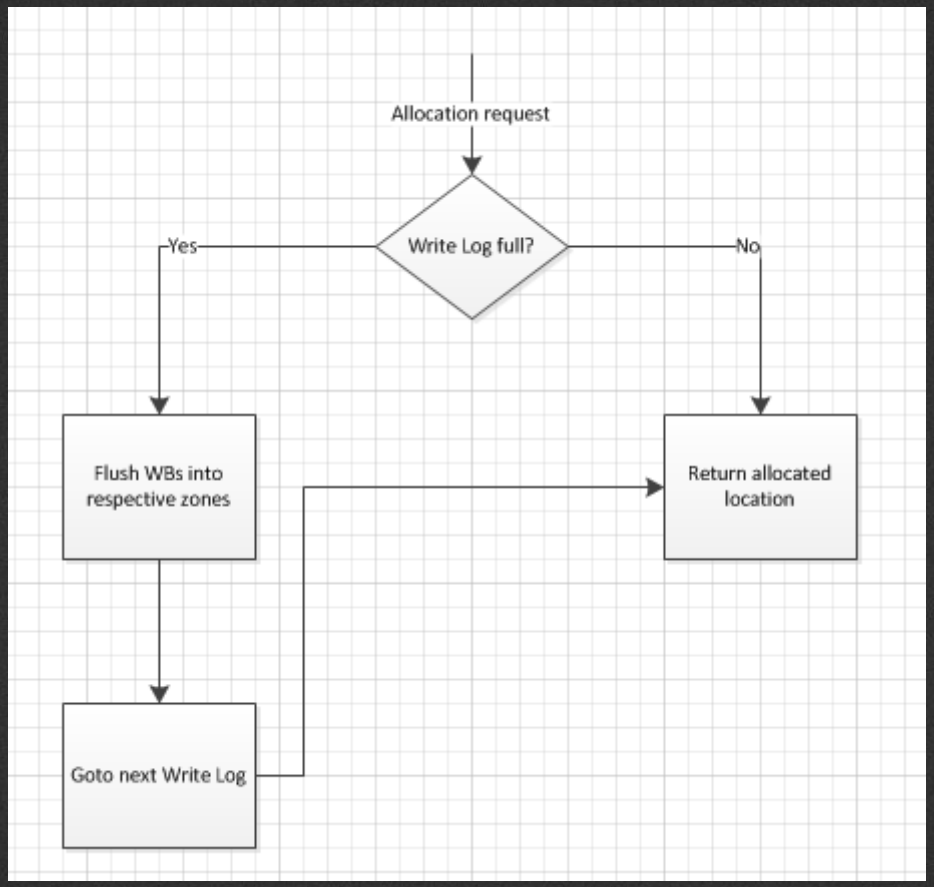

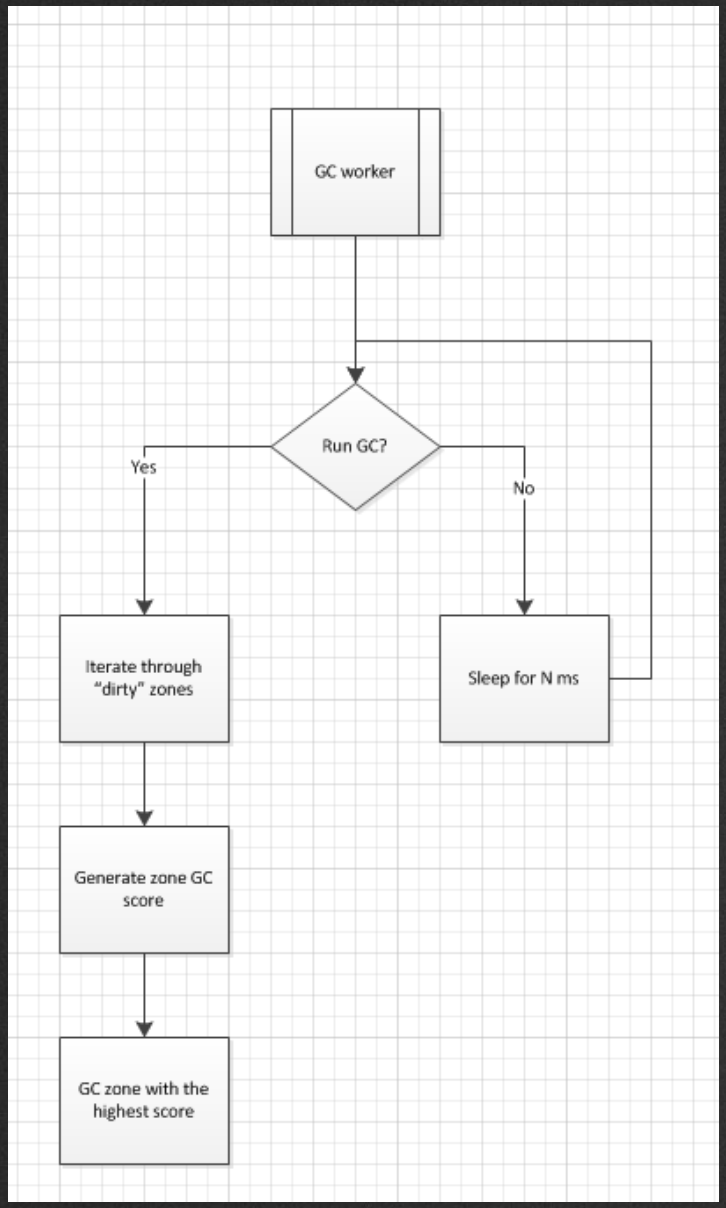

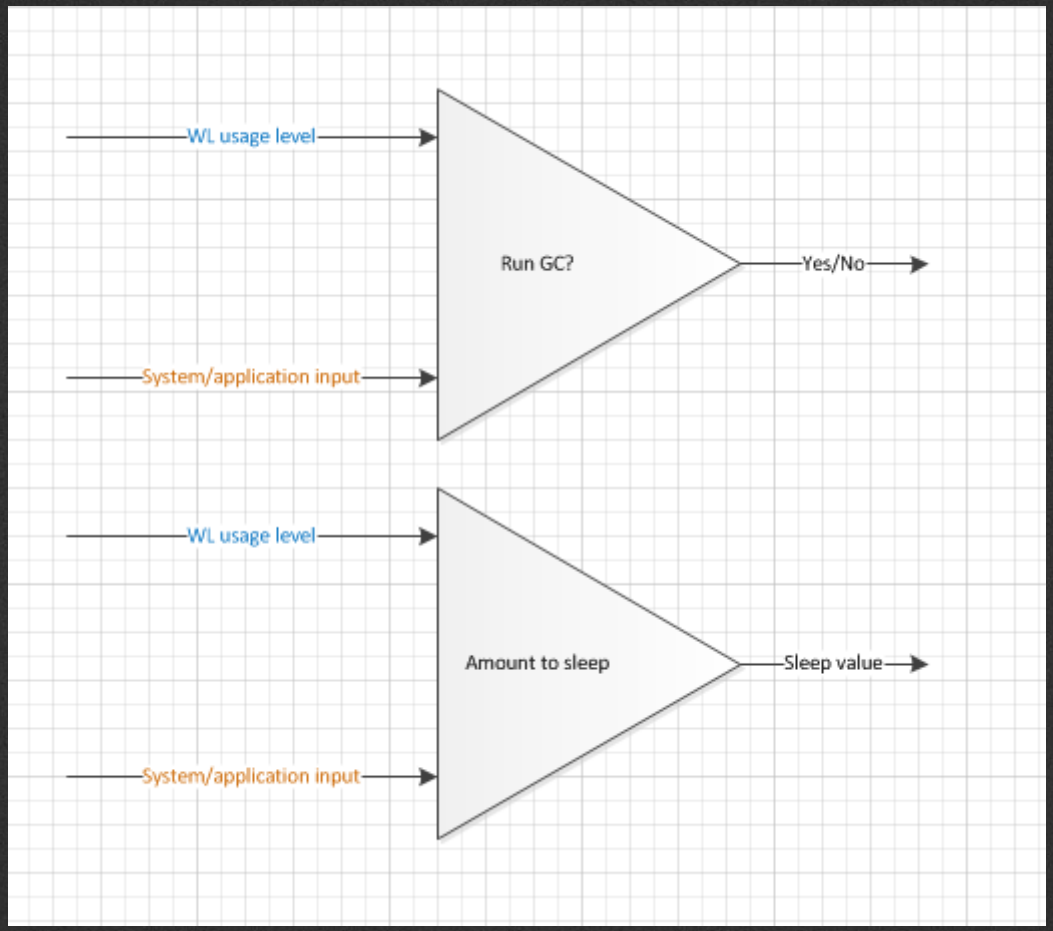

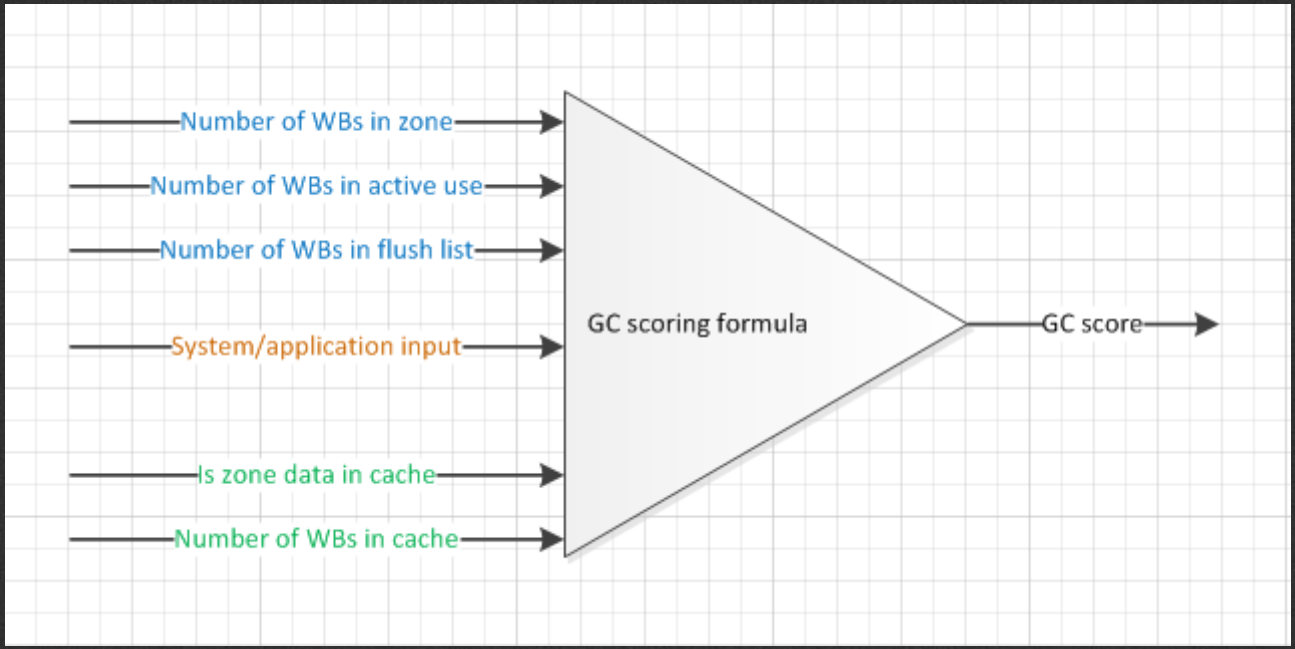

Thanks!

One more thing…

SMR Simulator

https://github.com/westerndigitalcorporation/SMR-Simulator

File system	Read beyond WP	Read spanning zones	Write not at WP	Write spanning zones	Not 4k aligned in SMR zones	Total number of violations
Ext4	5767	5687	3467	1252	0	16173
Btrfs	1231	1948	124	16	0	3319
Nilfs	0	0	0	0	0	0

Last month, WD released a SMR simulator to facilitate host-side software/file system development. SMR-Simulator is a simple tool that captures host software behavior and determines its “friendliness” to SMR technology. We hope SMR-Simulator will enable open source developers to experiment and become familiar with SMR functionalities and behaviors without the need to access real SMR (ZBC/ZAC) HW.

http://hselin.github.io/blog/2015/08/23/host-managed-smr-simulator

Contact us

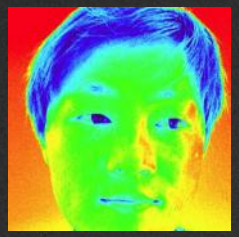

Albert Chen
Albert.Chen@wdc.com

https://www.linkedin.com/in/alberthchen

Jim Malina
Jim.Malina@wdc.com

https://www.linkedin.com/in/jimmalina

Takeaki Kato
TK.Kato@wdc.com

https://www.linkedin.com/pub/takeaki-kato/6/436/78