

Thomas-Krenn.AG[®]

The server experts



*SSD Caches
Festplatten-RAIDs mit Flashspeicher
beschleunigen*

TK Roadshow 2012

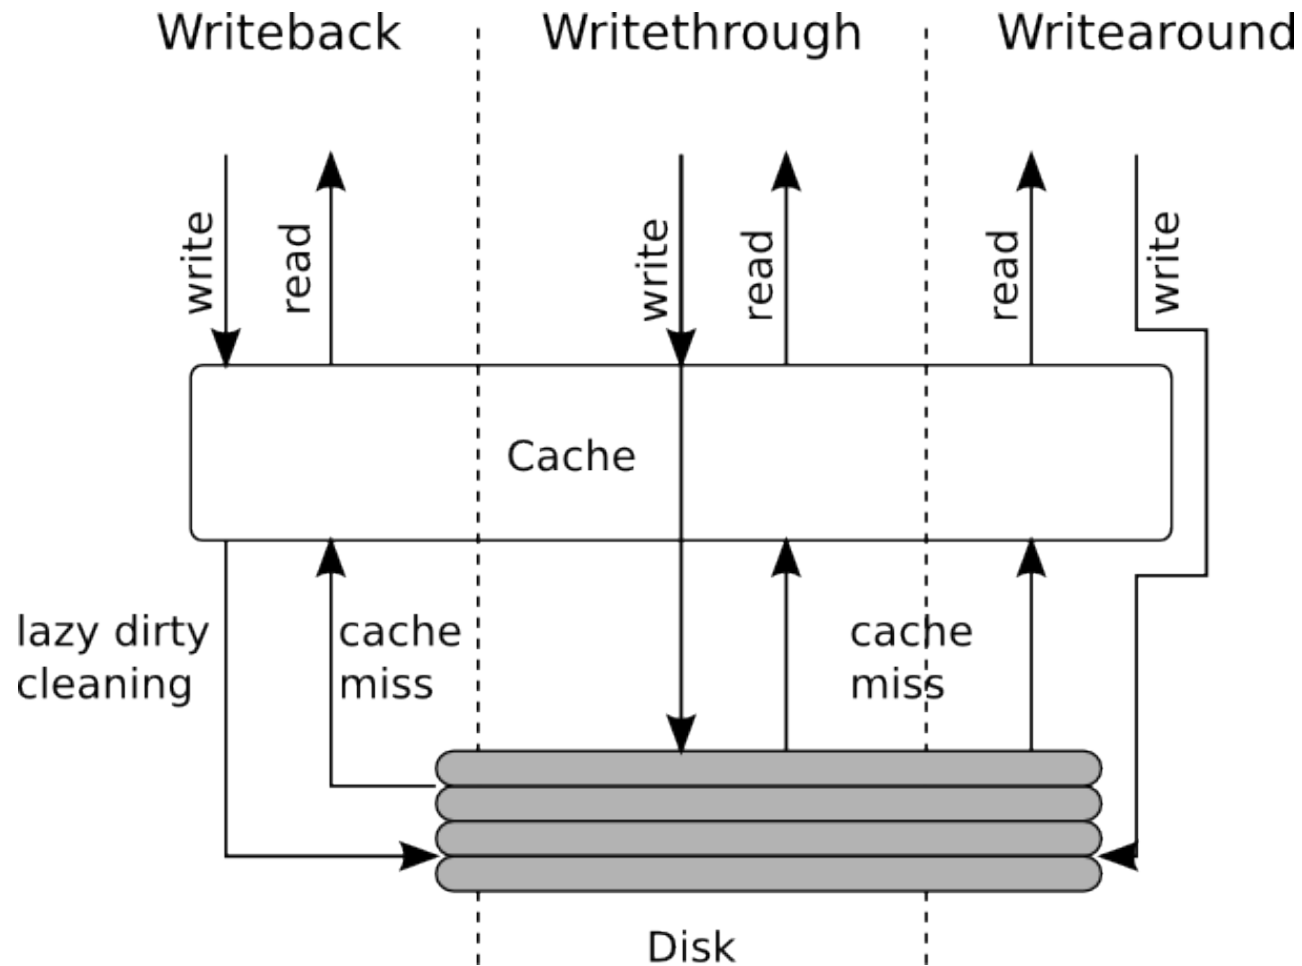
Agenda

- 1) Wann cachem?
- 2) Definitionen
- 3) Feature-Vergleich
- 4) Funktionsweise
- 5) Performance Tests
- 6) Fazit










1) Wann cachchen?

```
if(dataSet == known &&  
    ramAvailable != enough &&  
    appAccess == analyzed &&  
    perfTests == available)  
    CheckForCacheTechnologies();  
else  
    NeedMoreInfo();
```

2) Definitionen



3) Feature-Vergleich

	WB		WT	WA/read-only
FlashCache				
CacheCade		1		
MaxCache		2		

¹ Including ForcedWB

! Achtung: Derzeit kein redundanter Cache mit mehreren SSDs möglich

² Including InstantWB

3) Feature-Vergleich

		Skip sequential I/O
FlashCache	1	<input checked="" type="checkbox"/>
CacheCade	2	<input type="checkbox"/>
MaxCache	3	<input checked="" type="checkbox"/>

- Das Festplatten-RAID kann für sequentiellen I/O schneller sein als die SSD

¹ Konfigurierbar via sysctl (threshold)

² Keine weiteren Details bekannt

³ Nicht konfigurierbar, immer ausgelassen

3) Feature-Vergleich

	Cache persistence	
FlashCache	<input checked="" type="checkbox"/> !	1
CacheCade	<input checked="" type="checkbox"/>	2
MaxCache	<input checked="" type="checkbox"/>	3



- Das Festplatten-RAID kann für sequentiellen I/O schneller sein als die SSD

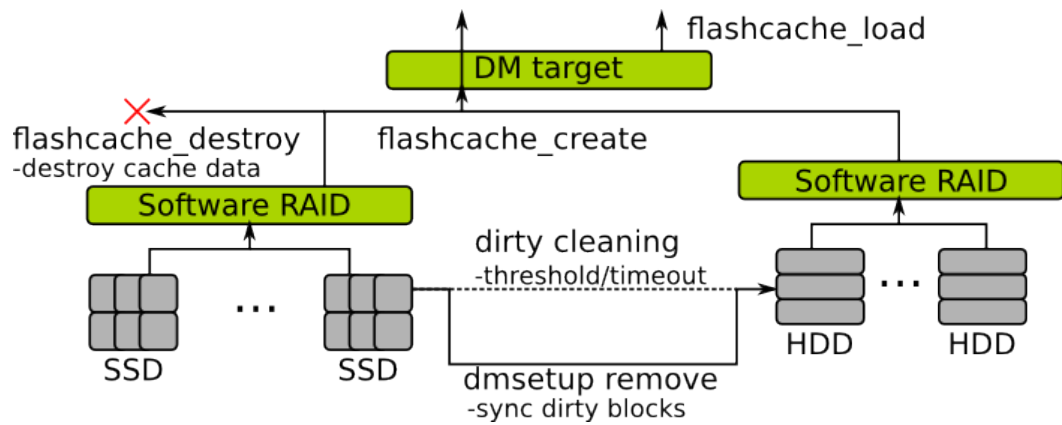
¹ Nur WB Cache ist persistent

² Persistent in alle Cache Modus

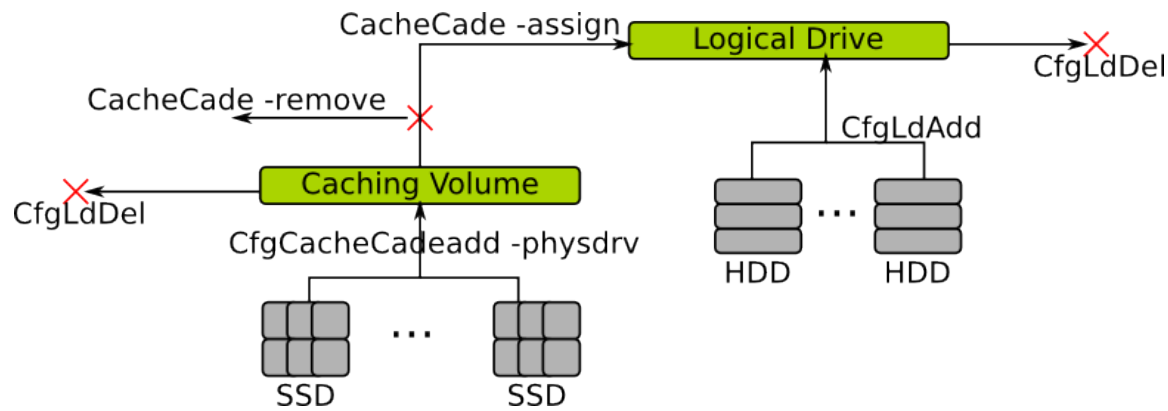
³ Dirty Shutdown: WT – discarded, WB – LV failed

3) Feature-Vergleich

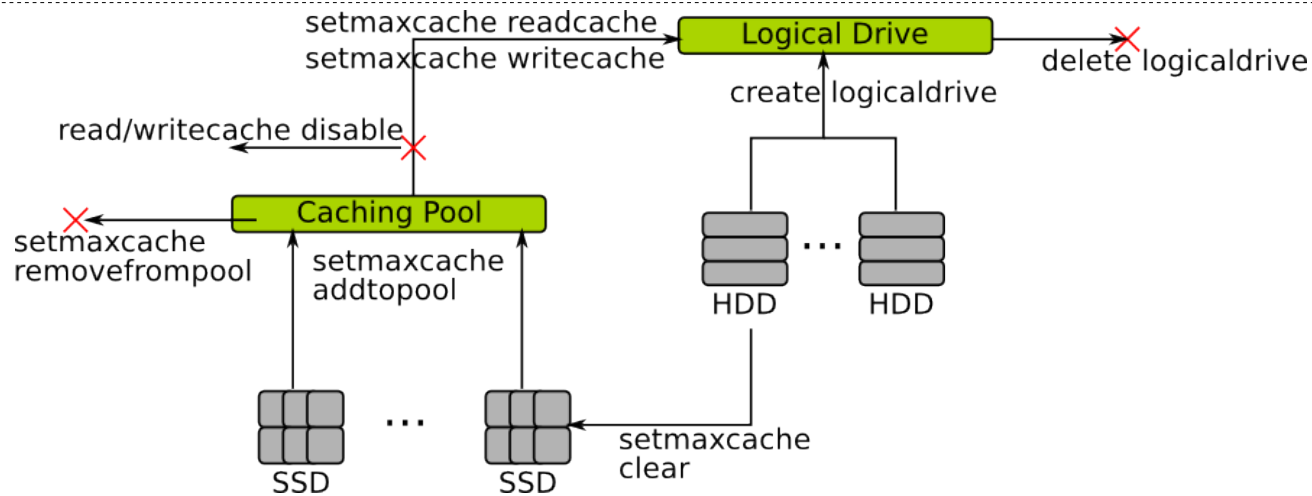
	Hot spot detection
FlashCache	
CacheCade	
MaxCache	



FlashCache

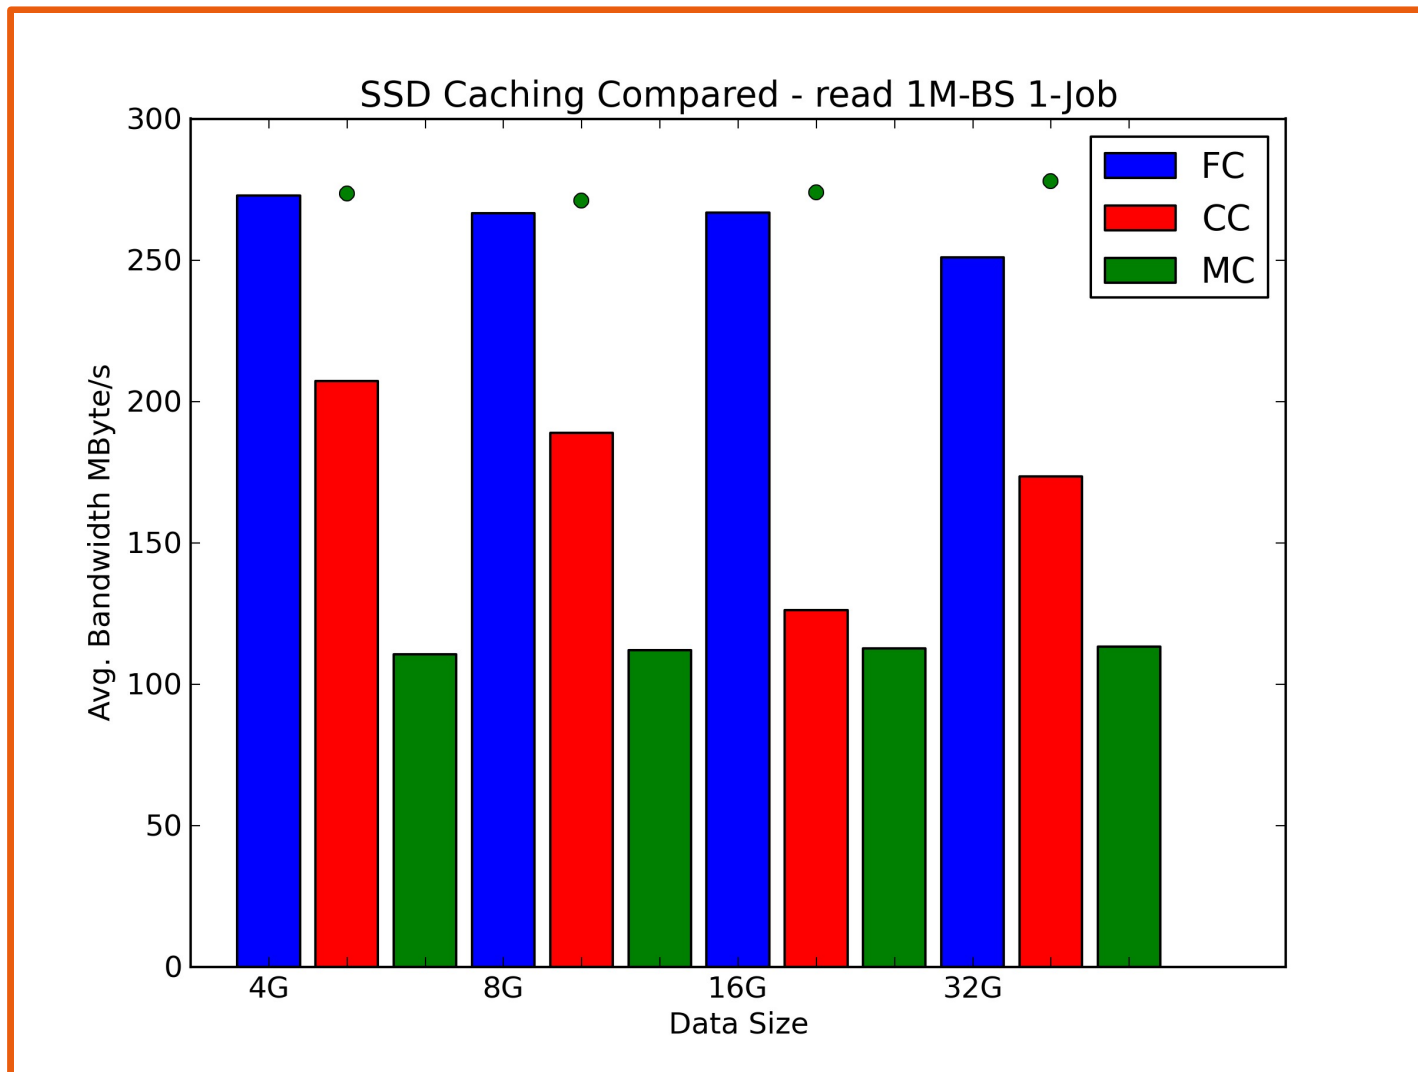


CacheCade

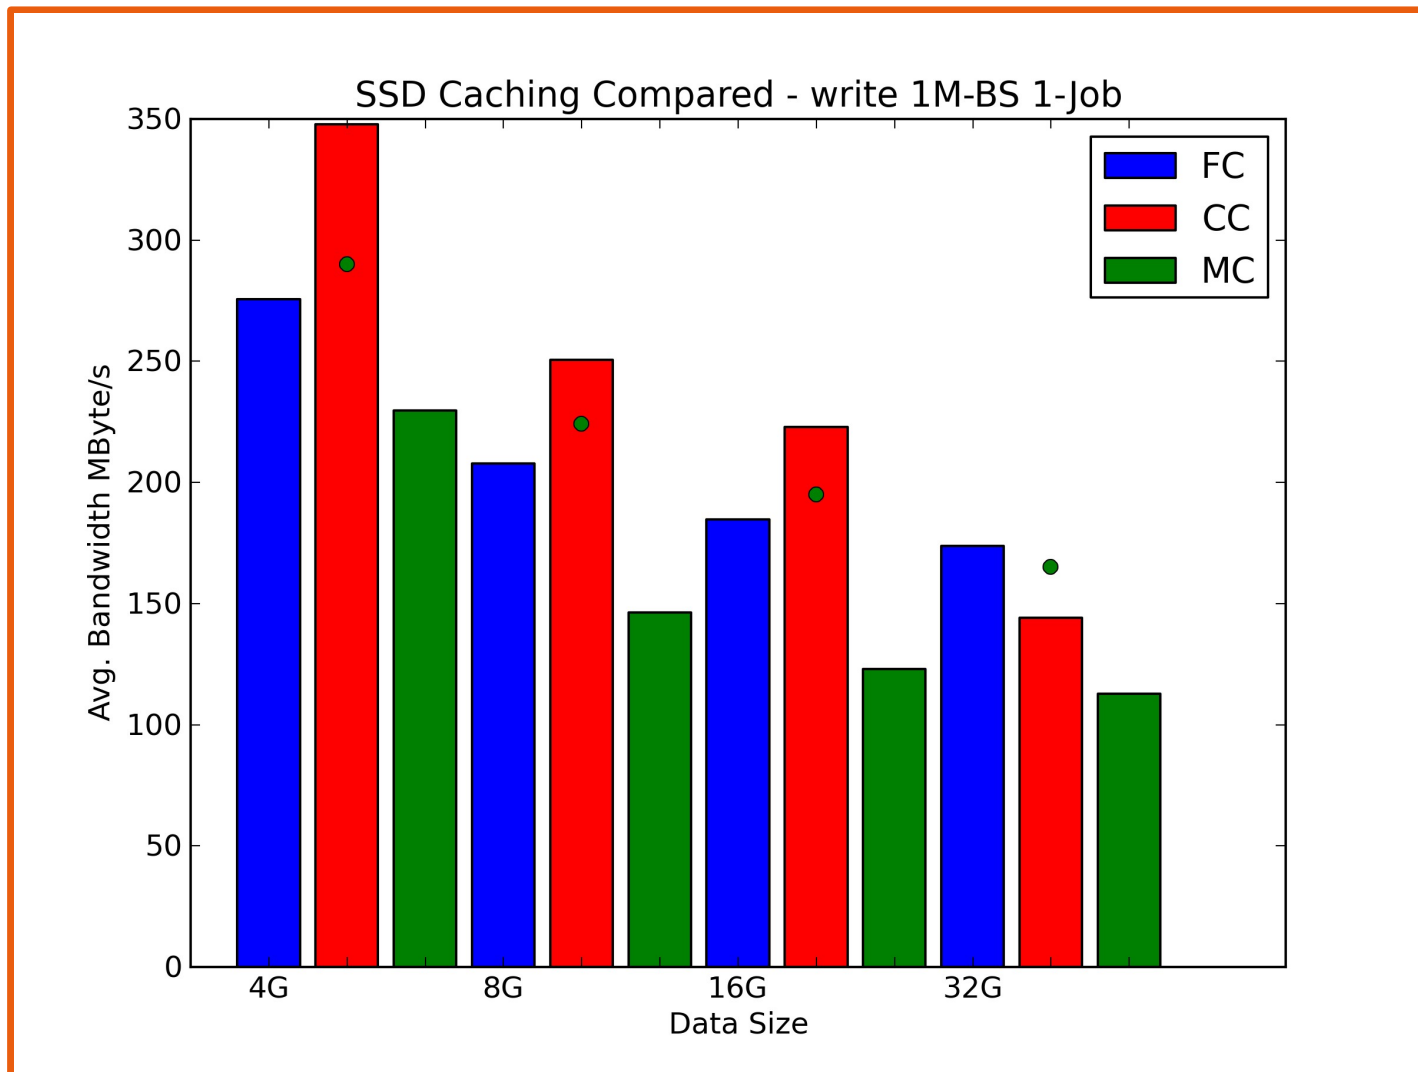


MaxCache

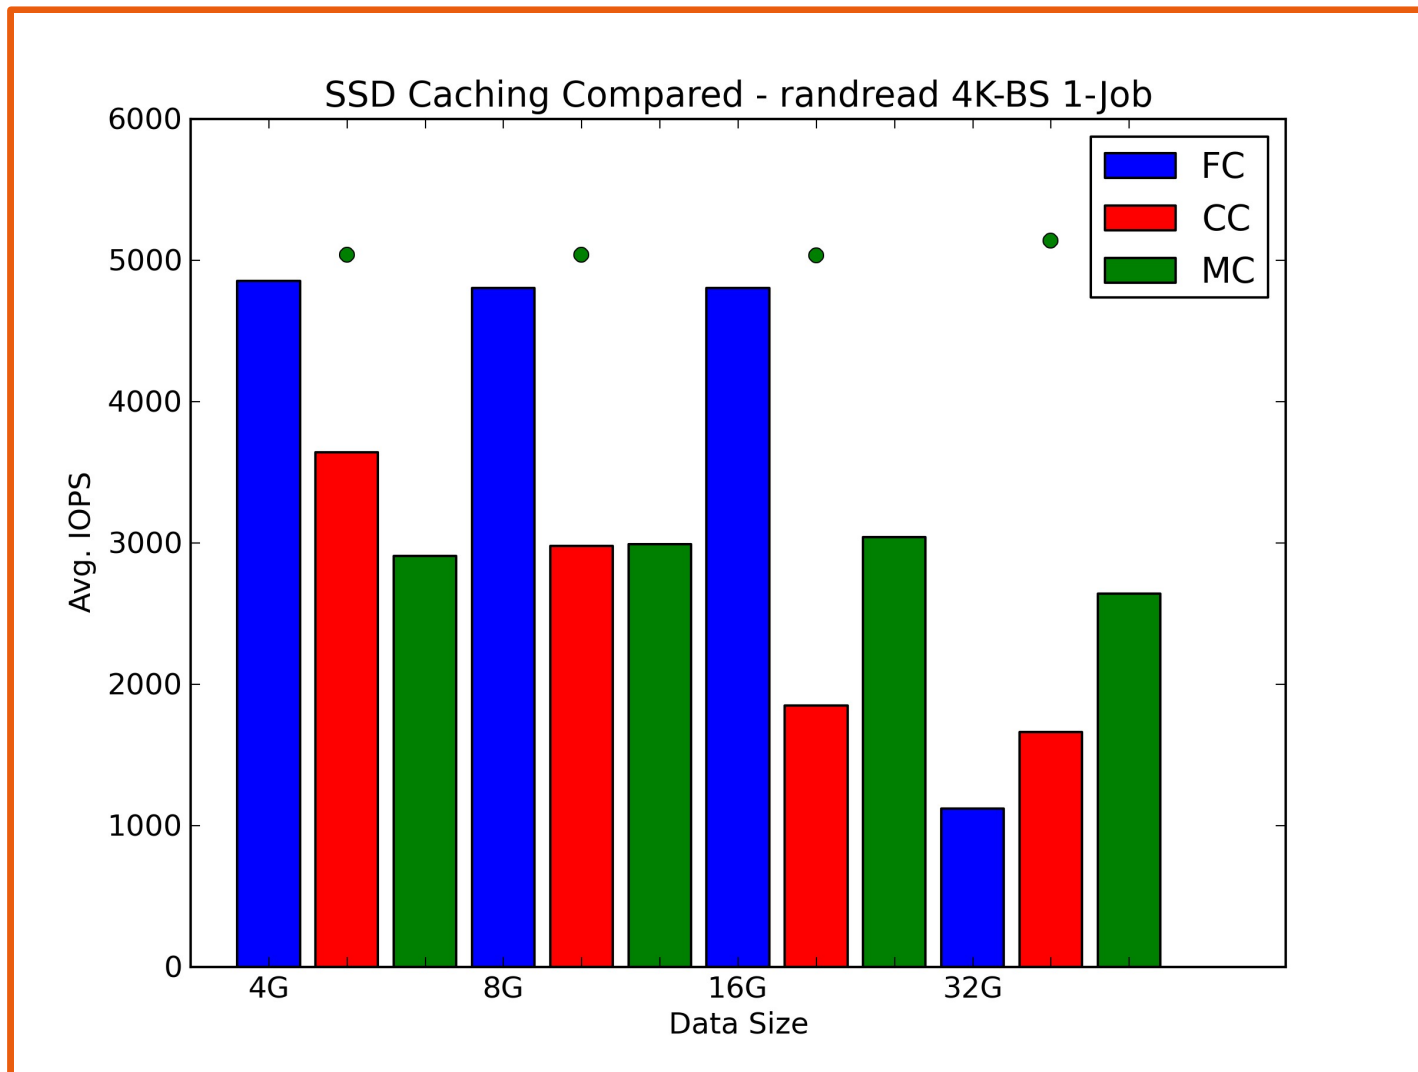
5) MBIs - read



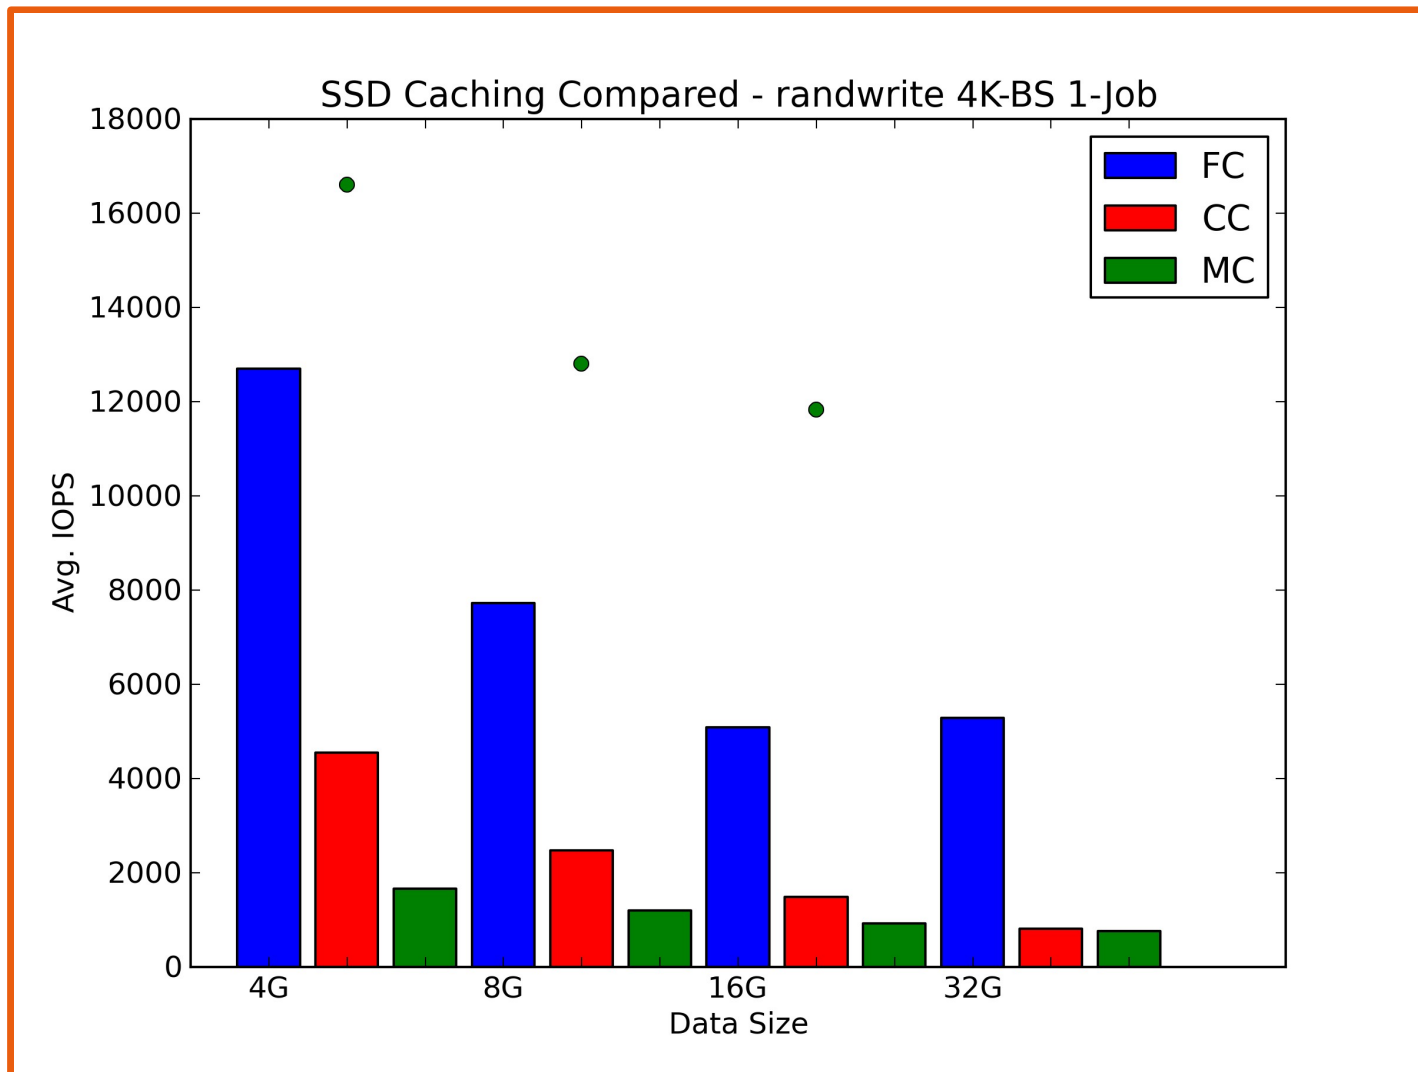
5) MBIs - write



5) IOPS - randread



5) IOPS - randwrite

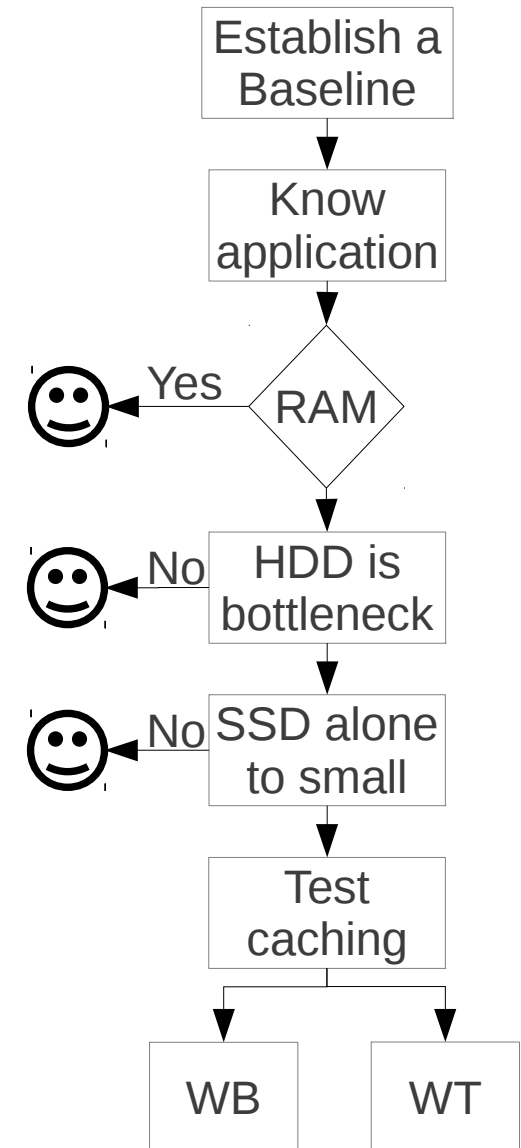


6) Fazit

1 Applikation testen

2 2x Ergebnisse überprüfen

3 Erkenntnisse umsetzen



Backup: Test system

- SSDs
 - Intel Series 320 160GB
 - Via HPA reduced to 32GB
- RAID Controller
 - LSI MegaRAID SAS 9260-4i
 - Adaptec 6805Q
- Software
 - Fio 2.0.7
 - Ubuntu 12.04
 - Updates from Release Day

Backup: Test script

- flashcache_create
 - Create a WB caching device
- cache_all=0
 - Don't cache ext4 initialization
- mkfs.ext4 -q -E lazy_itable_init=0,
lazy_journal_init=0 /dev/mapper/fc-root
- mount /dev/mapper/fc-root
- cache_all=1
- Call fio
- umount /dev/mapper/fc-root
- dmsetup remove
- flashcache_destroy /dev/sdd