

Dell PowerVault MD3860f

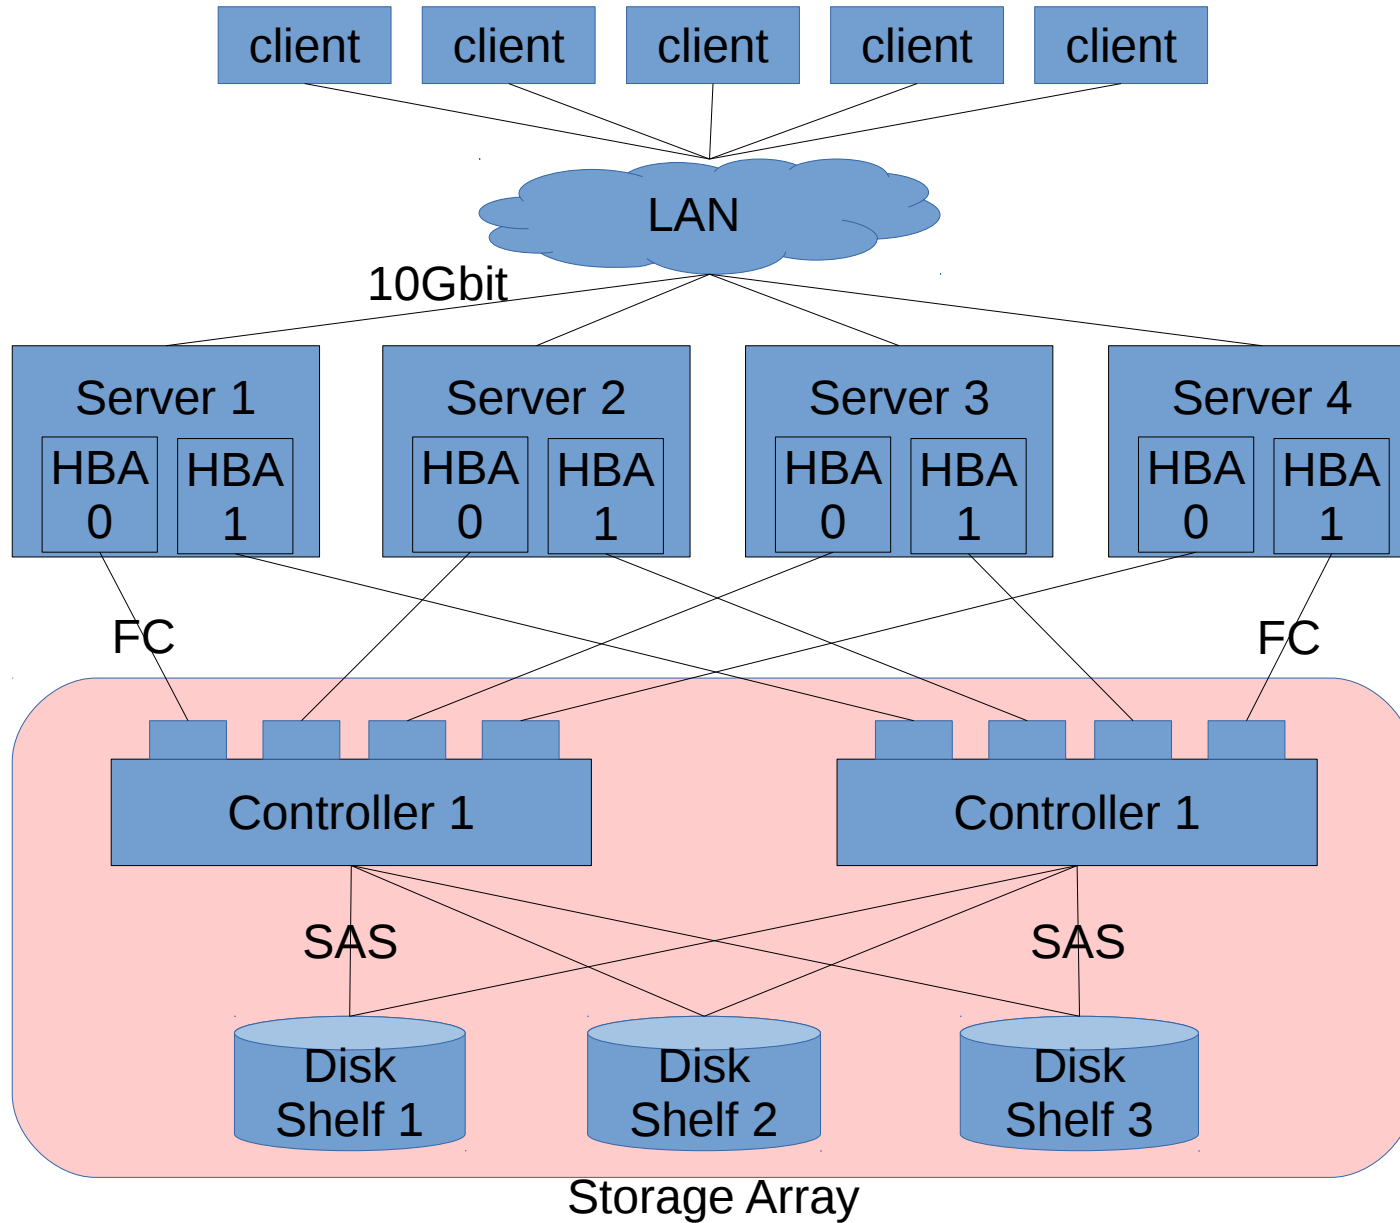
Test di pre-collaudo

MD3860f HW

- Dual controller (active-active)
- Front-end: 4xFC16 ports per controller
 - 16GB/s total
- Back-end: 3x6Gbps SAS bus per controller (1 int + 2 ext)
 - 4.5GB/s total
- Disks: 180 NLSAS 4TB
- Base with 60 HDDs
- Two expansion box with 60 HDDs each



Testing environment



RAID Configuration



- Traditional RAID-6:
 - 18 disk pools of 10 HDDs (8+2)
- Dynamic Disk Pool:
 - 6 disk pools of 30 HDDs
 - In each disk pool has 8TB (= 2 HDD) of “reservation capacity”
 - 1 disk pool of 180 HDDs
 - 8 TB of “reservation capacity”

File system configuration

- GPFS v3.5.0-19
- Block size = 4MB
- Block Allocation type: scattered
- pagepool=16G
- maxMBps=6400

Dynamic Disk Pool

- DDP – distributed RAID6
- D-stripe – fixed size=4GB
-




Dynamic Disk Pools Technical Report:
http://i.dell.com/sites/doccontent/shared-content/data-sheets/en/Documents/Dynamic_Disk_Pooling_Technical_Report.pdf

RAID configurations

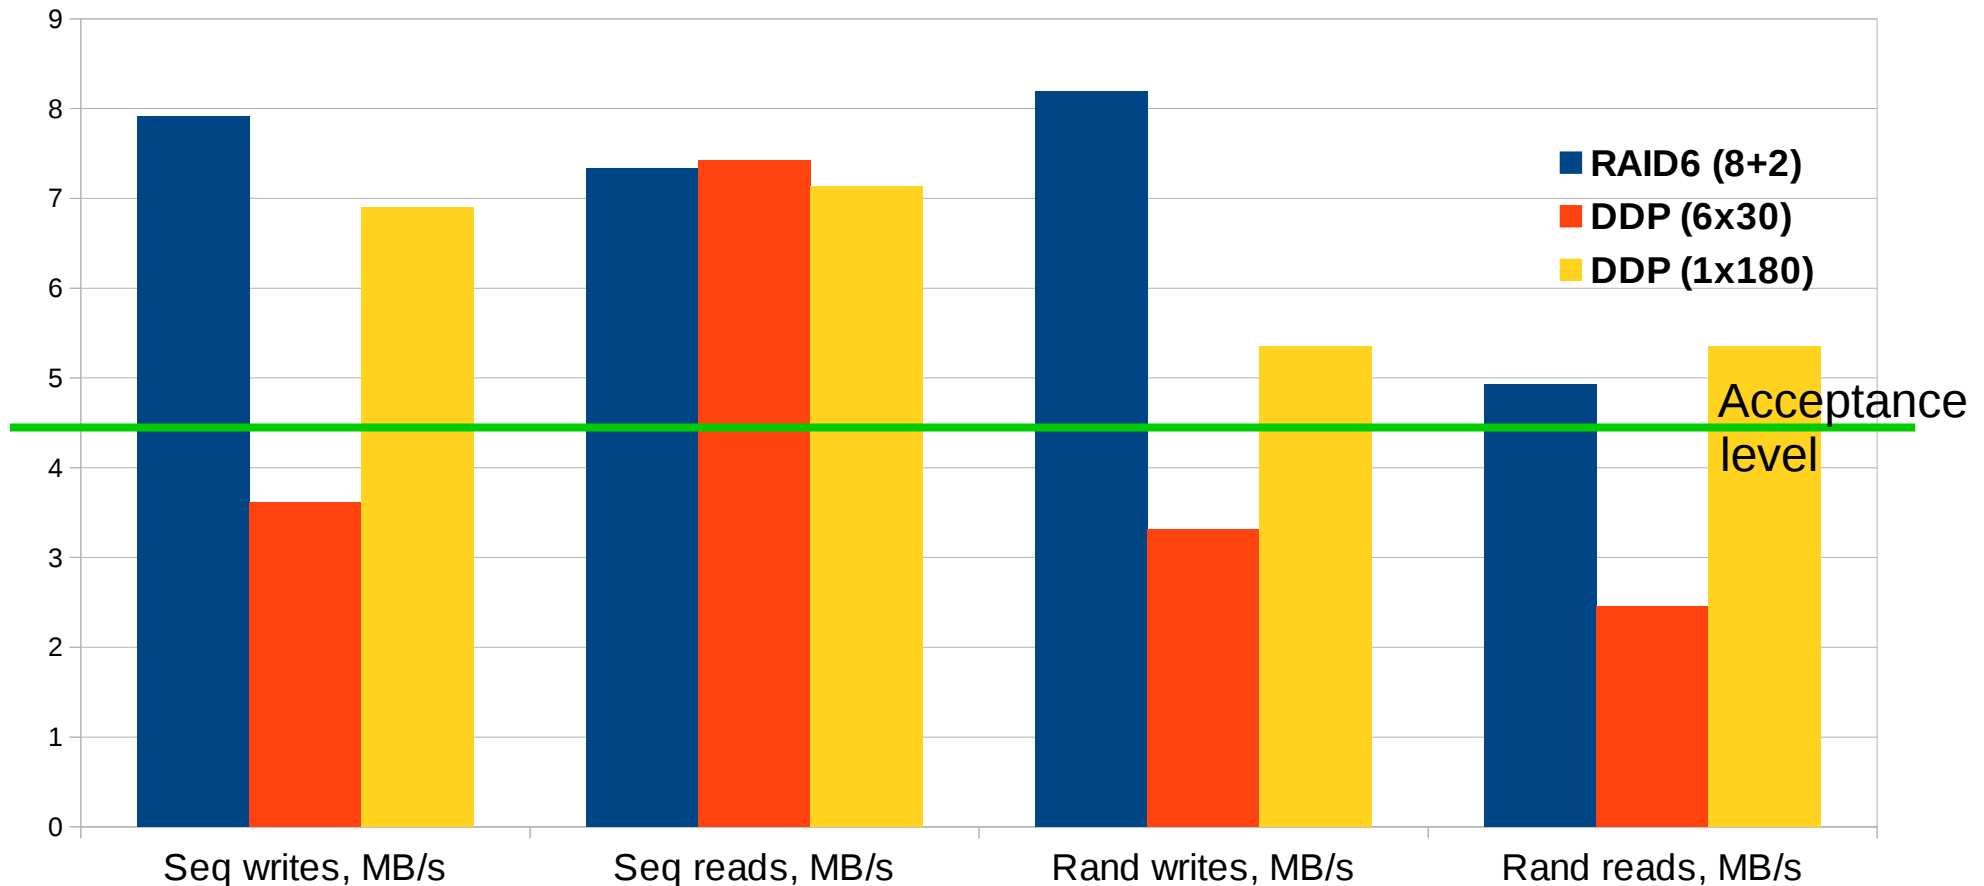
RAID configuration	RAID6 (8+2)	DDP (6x30)	DDP (1x180)
N. of pools	18	6	1
N. disks per pool	10	30	180
% used for parity	20	20	20
N. of reserved capacity disks	0	2x6	2
N. of LUNs	18	12	16
Usable space, TB	576	530	562
Critical conditions (N. Of failed disks)	2 in 1 pool	2 in 1 pool	2 in 1 pool
Data loss probability (3 disks failure)	P*10	P*30	P*180
Recovery Time from critical, hours	50 (rebuild of 1 disk)	<3 (estimate)	<0.5 (estimate)

Recovery time from HDD failure

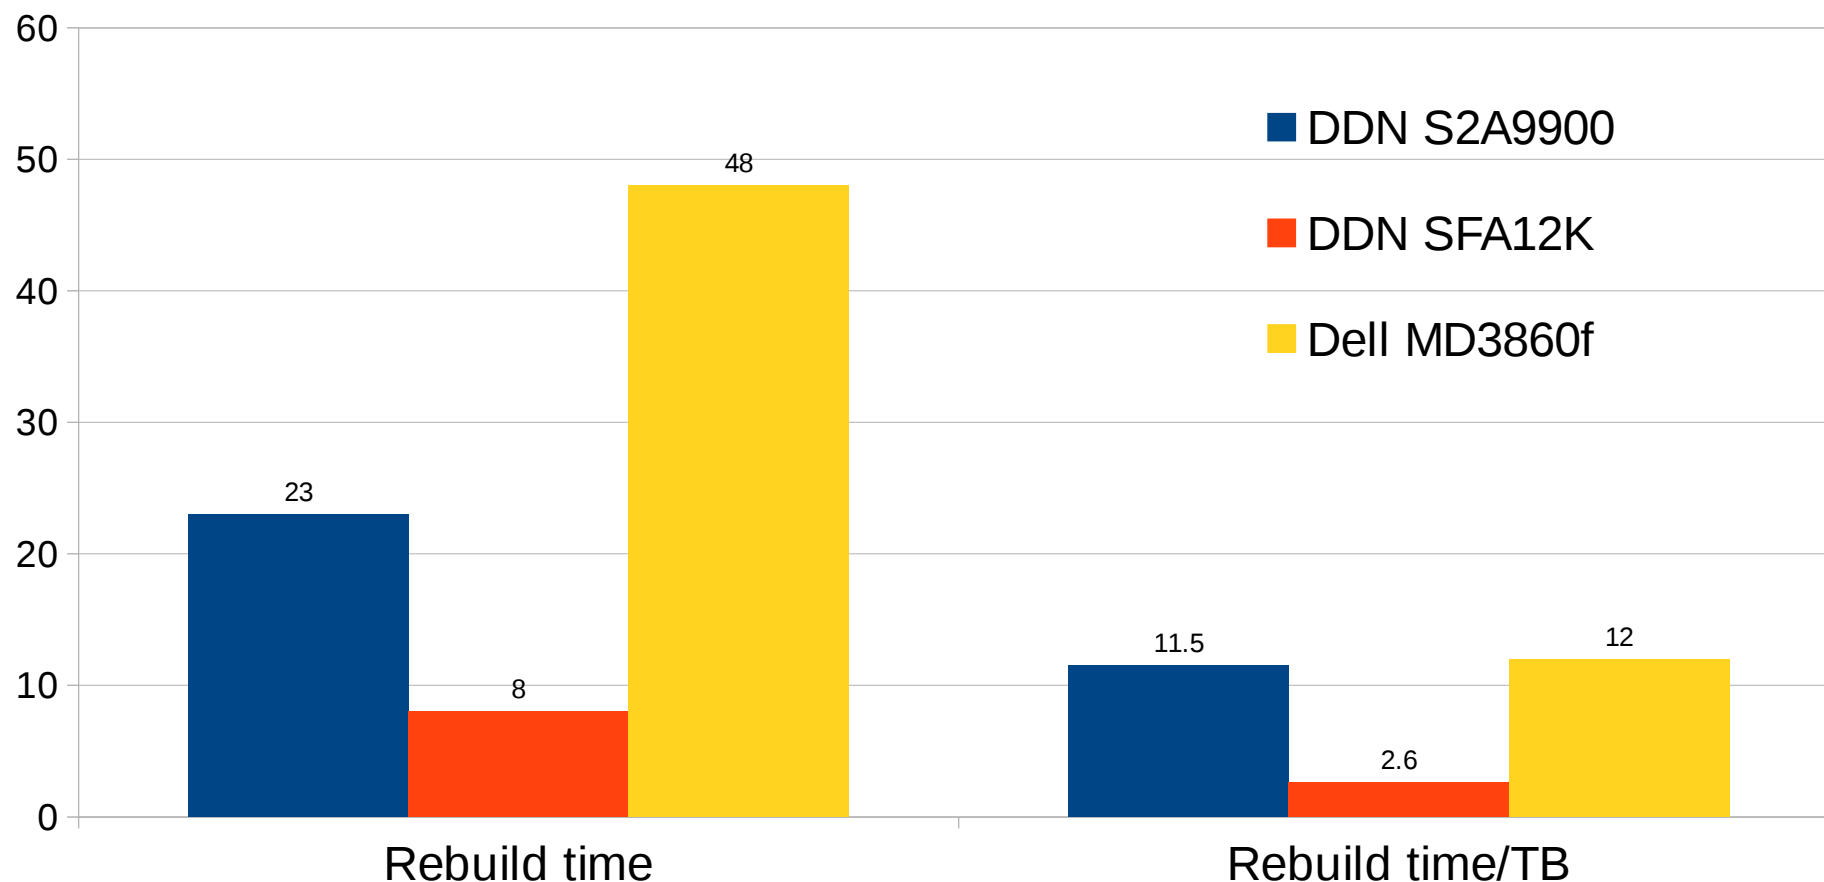
- Recovery time as reported by Dell



Performance results (MB/s/TB)




Rebuild time (hours), traditional RAID



- 4TB disks recovery (as tested at CNAF) with DDP:
 - ~ 20h under load,
 - ~9h with no load


Impact to production

- DDP 3x60 during Reconstruction: read rate at 75% respect to the optimal conditions:




Concurrent reads and writes

- DDP 3x60



- DDP 1*180



Conclusions

- Recovery time from critical conditions (2 failed disks in the same raidset) in case of traditional RAID6 is too long (>2 days in case of minimal activity)
- DDP (1x180) makes the same recovery in a few minutes
- Use of DDP is more attractive even considering lower (-20%) performance