FB2HB2MB6

Ari Mujunen, Aalto University
Metsähovi Radio Observatory

FlexBuffTwoHarroBoxTooMarkBoxSix

Mark 6 hardware

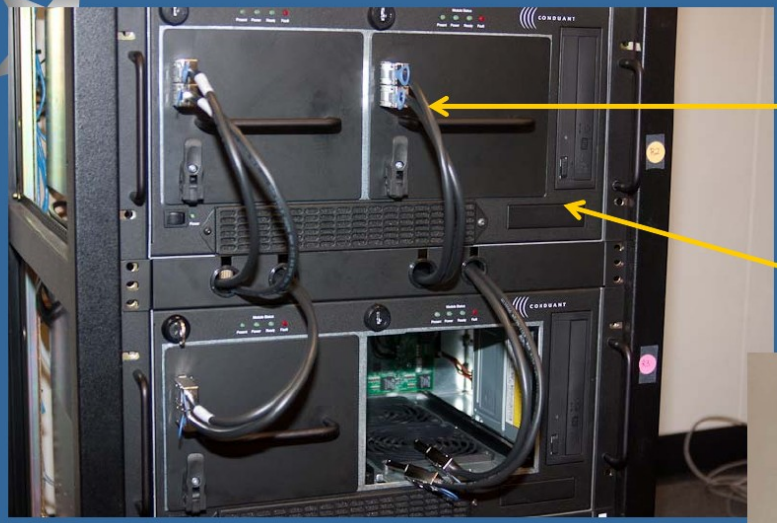

High-speed data connections to module front-panel via two standard SAS cables

Existing Mark 5 chassis is upgradeable to Mark 6

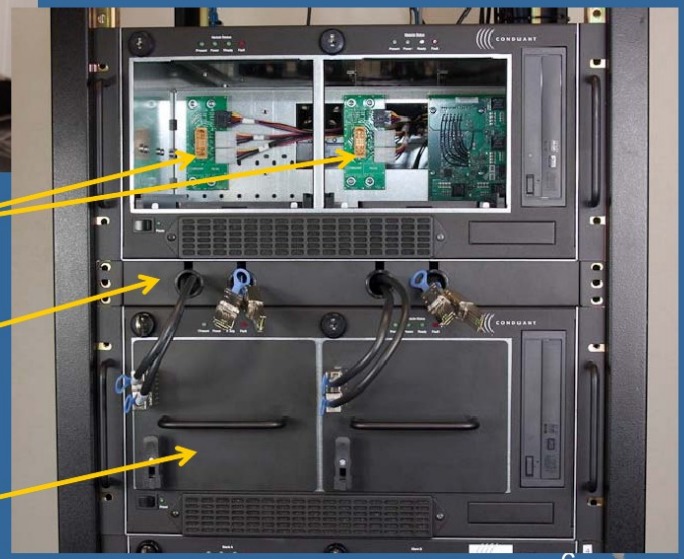

New chassis backplanes for disk power management

Cable-management panel (unused cables retract into panel)

Existing Mark 5 SATA disk modules are upgradeable to Mark 6 (new backplane and front panel)

USD 495 USD 300

Mark 5 SATA Drive Module Upgrade to Mark 6

Cost:
USD 62 per each drive +labor...

New Front Panel

Connectors for two eSATA cables

8x LED (1 per drive)

Re-use Handle from old Module

pre-installed

Front Panel

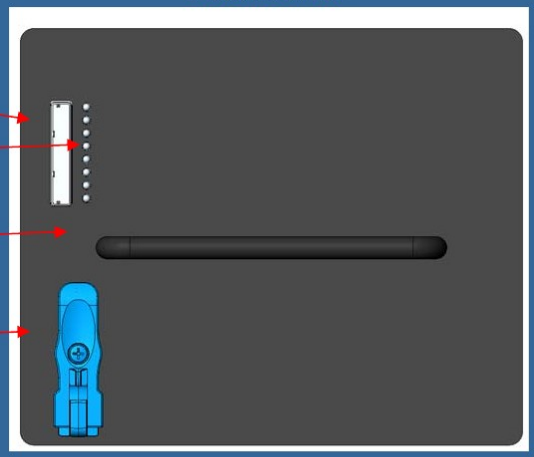

Cooling slots

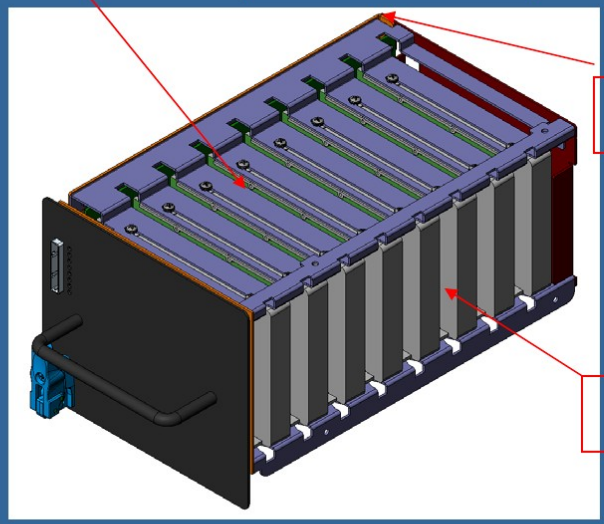

New PCB and power connector

Easily removable disks

Rear Panel

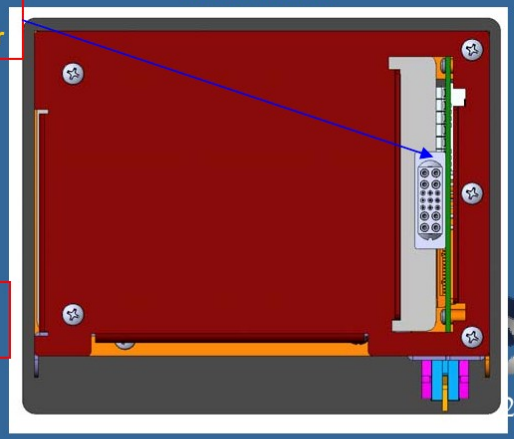

Mark 6 hardw

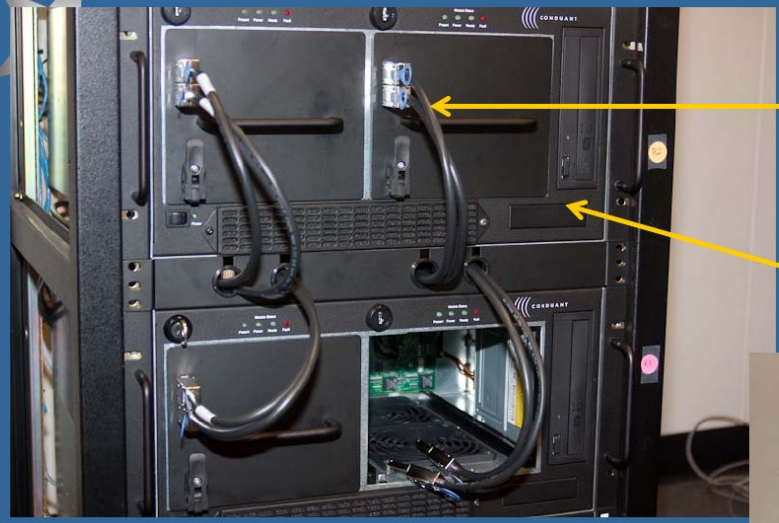

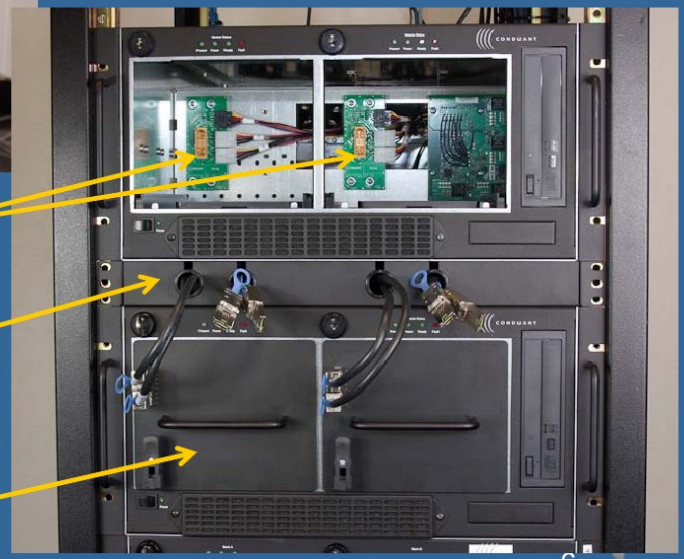

New chassis backplanes for disk power management

Cable-management panel (unused cables retract into panel)

Existing Mark 5 SATA disk modules are upgradeable to Mark 6 (new backplane and front panel)

Mark 6 16Gbps demonstration system

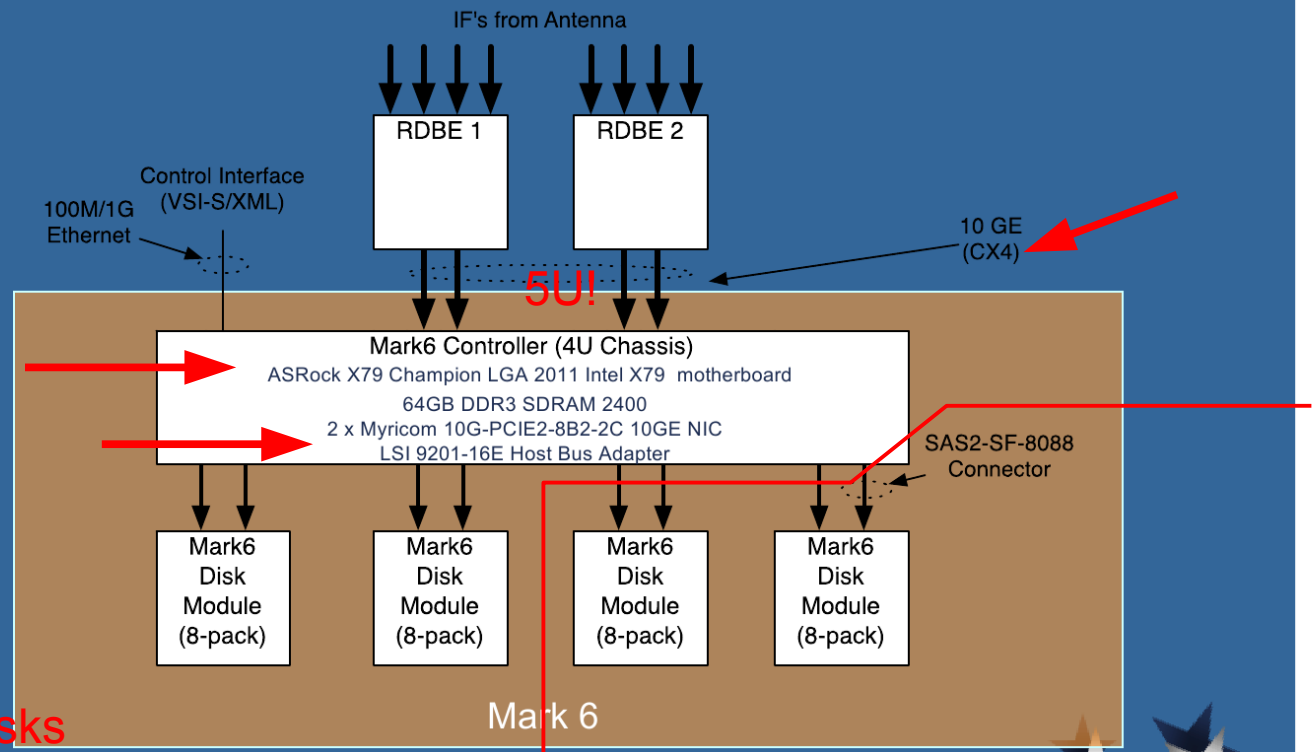

Only
32 disks
in total

7k€ / 3k€

HW+SW Interchangeability

- Ran this ---> on FB HW
- jive5ab with FB SW should run on Mark6 HW just fine

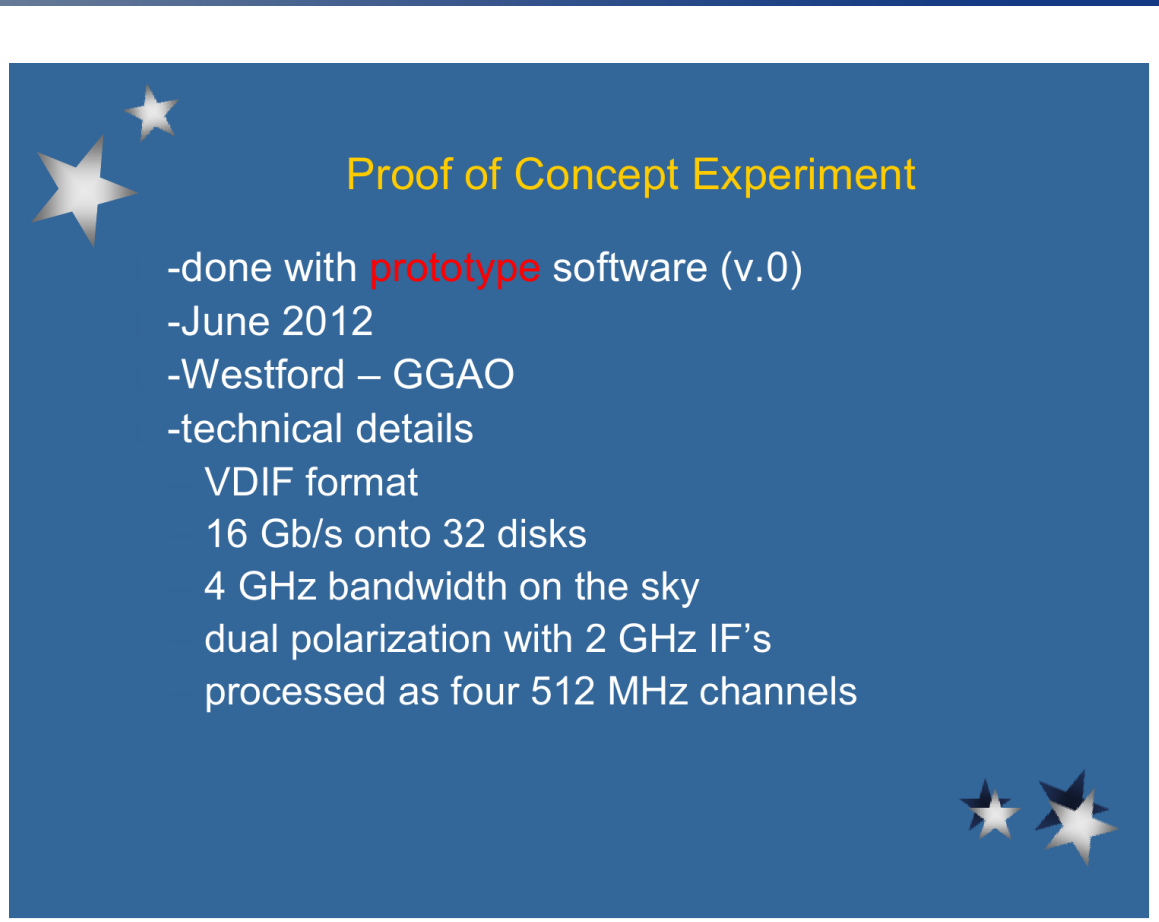

FB2HB2MB6 Mark6 16Gbps demonstration system

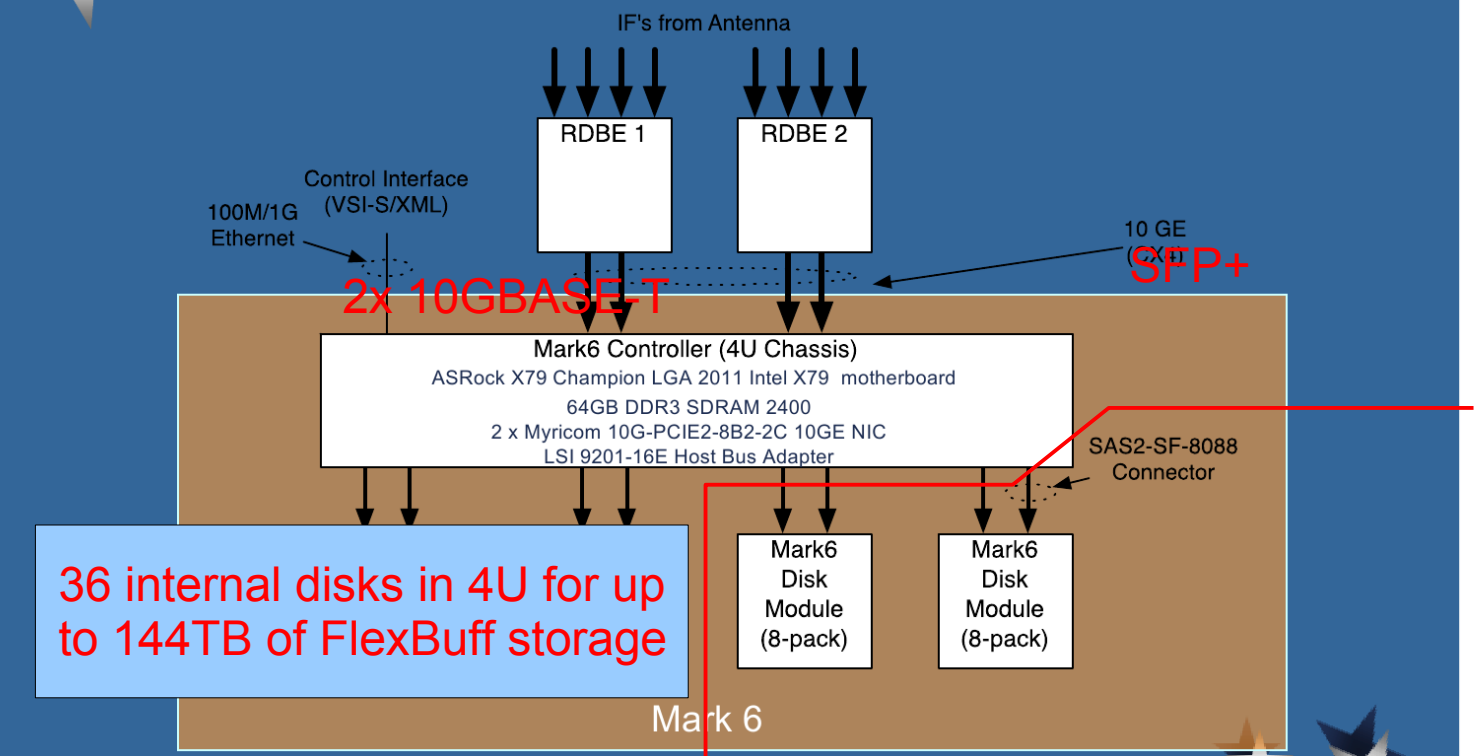

Motherboard

- 2x LGA2011 Xeon
- 7 slots PCIe x8

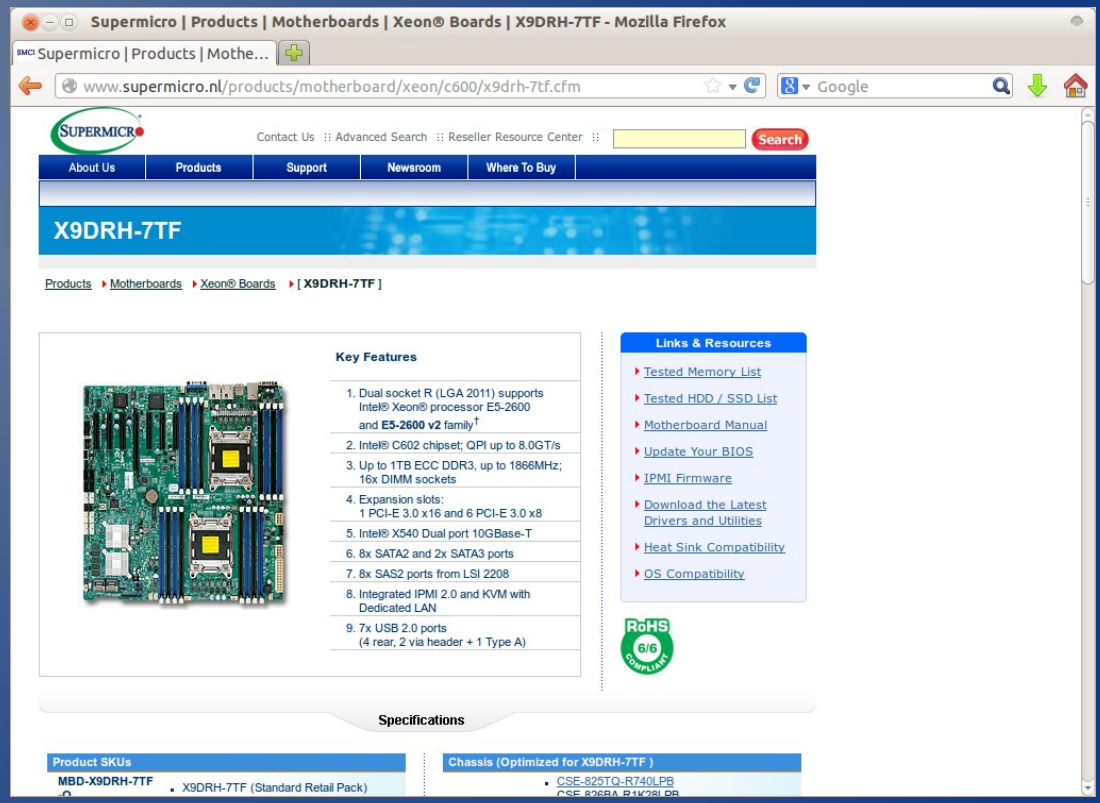

Disk Controllers

- SATA III 6Gbps Generation
- LSI3008 chips
- One controller per 8 disks

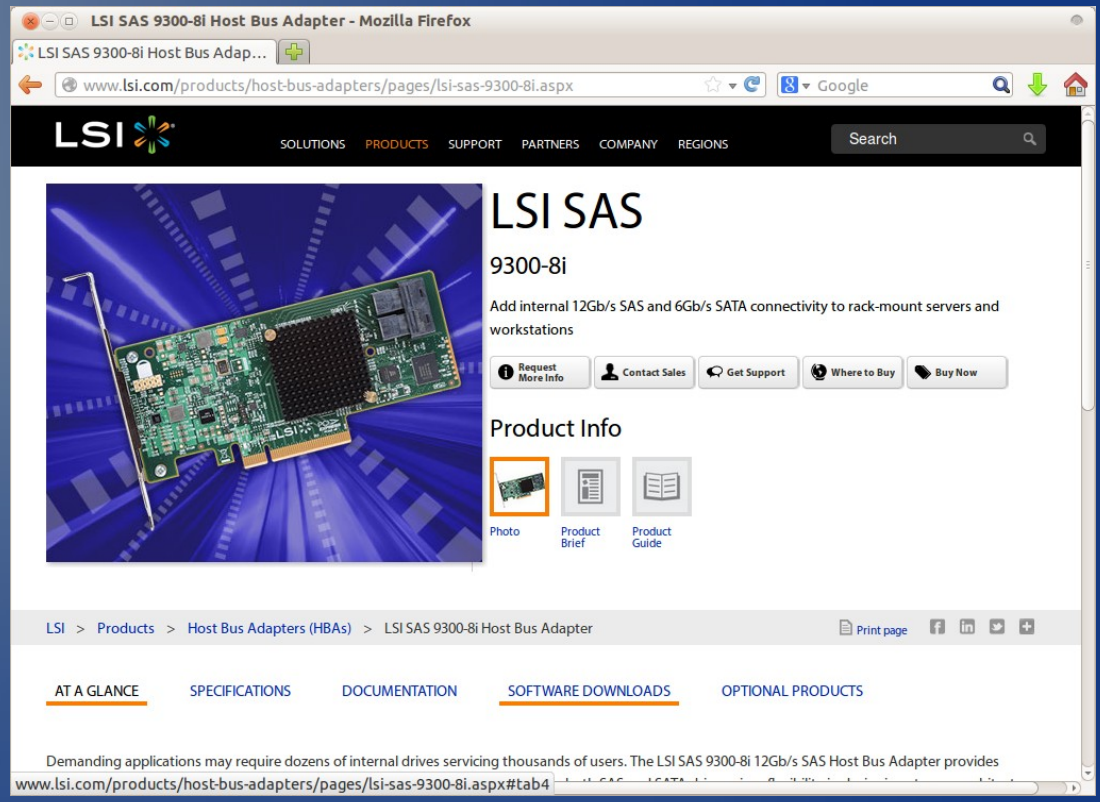

10GE Adapters

- Dual SFP+

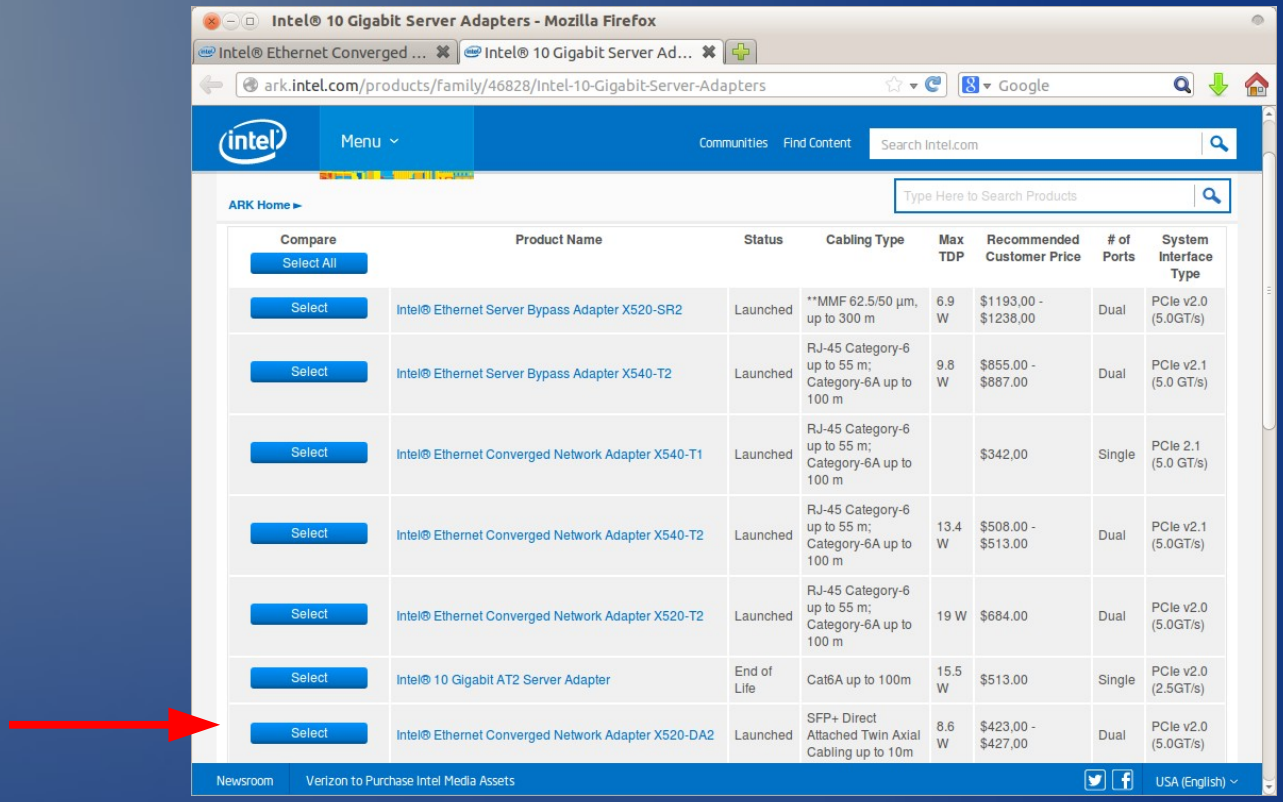

10G Optical Modules

- Reasonable cost
- Takes ~1 month
  - Just like branded modules...

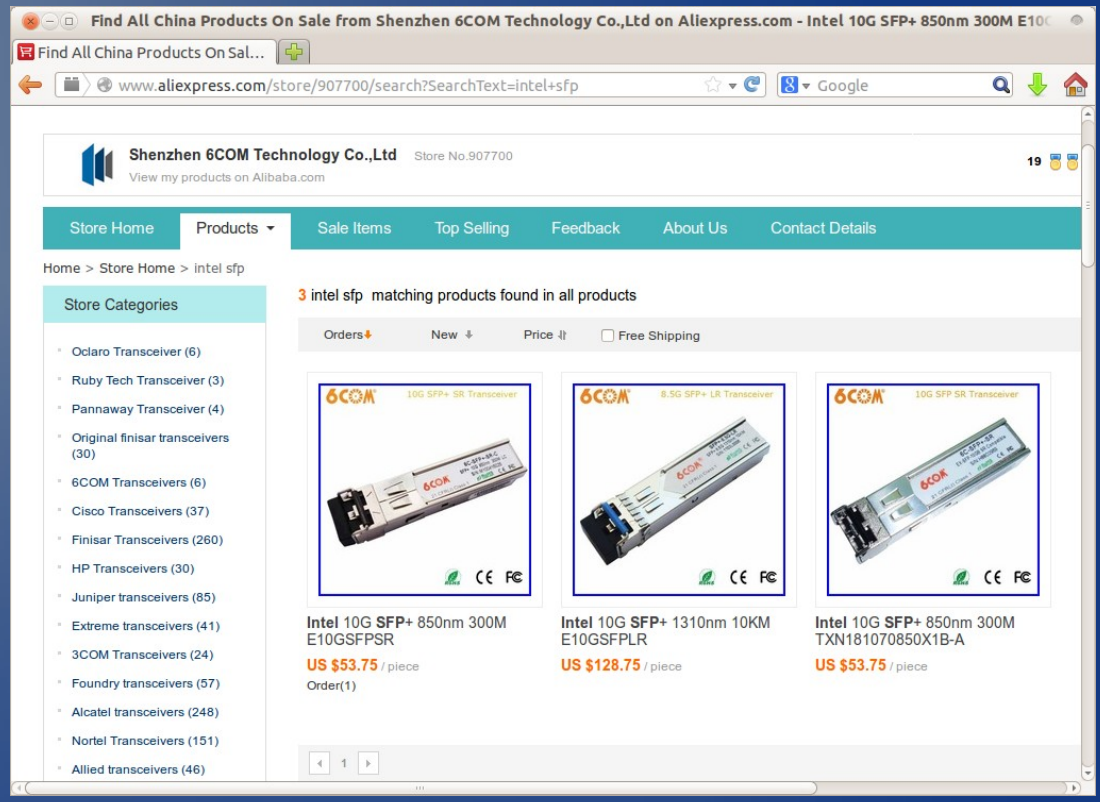

Parts List

- 6304e
  - From www.damicon.fi published no-VAT prices
  - 4U 36-disk, 2x6-core Xeon, 64GB RAM, 4x10GE (2x-T, 2xSFP+), 2xMark6

```
1  1565  SUP-CSE-847AR1400LPB
         4U ATX Supermicro SC847A-R1400LPB 36xSAS/SATA multilane 2x1400W
1   756  SUP-X9DRH-7TF
         Supermicro X9DRH-7TF EATX 2*Xeon E5 C602 16DDR3 1x16/6x8 PCIe
         6SATA 8SAS 2*10G
2   335  I-X-E5-2620V2
         Intel Xeon E5-2620v2 6-core IB-EP 2100 MHz 15 MB HT LGA2011 80 W
2    36  SUP-SNK-P0048P
         Supermicro LGA2011 X9 DP 2U passive heat sink
8    83  8GB-1600D3-ECC
         8GB 72-bit ECC DDR3 SDRAM DIMM 1600 MHz PC3-12800E CL11
4   294  SAS9300-8I
         LSI Logic SAS 9300-8i HBA 12Gb SAS 6Gb SATA 2xSFF8643 PCIe3x8 LP
8    27  CAB-8643-8087
         SAS cable SFF8643-SFF8087 internal 1m
1    25  CAB-SFF8087
         SAS cable iPass SFF8087-SFF8087 internal 50 cm
2   350? SAS9300-8E
         LSI Logic SAS 9300-8e HBA 12Gb SAS 6Gb SATA 2xSFF8644 PCIe3x8 LP
1   366  I-E10G42BTDA
         Intel E10G42BTDA 10 Gigabit X520-DA2 (SFP+) (82599EB) lp PCI-e x8 2.0
1    80  AS80
         Installation and testing
1   230  TAKUU3SERV
         Warranty extension to 3yr
```

Disks

- 5TB and 6TB disks starting to appear
  - Hitachi He-filled 6TB, Seagate Ent Cap 5TB
- 4TB around 150—300e each; 6—14k€ 144TB

```
40x disks ("two cartons" :-)
155  ST4000DM000
     4TB Seagate Barracuda 7200.14 ST4000DM000 SATA/6Gbit 64 MB NCQ
170  ST4000VN000
     4TB Seagate NAS ST4000VN000 5900 rpm SATA/6Gbit 64 MB NCQ
234  ST4000NC000
     4TB Seagate Terascale ST4000NC000 5900 rpm SATA/6Gbit 64 MB NCQ ISE
292  ST4000NM0033
     4TB Seagate Constellation ES.3 ST4000NM0033 7200 rpm SATA/6Gbit 128 MB NCQ
342  ST4000NM0023
     4TB Seagate Constellation ES.3 ST4000NM0023 7200 rpm SAS/6Gbit 128 MB
4TB disks (144TB FlexBuff capacity) 6200--13680e

3TB disks 114--283e --> 4560--11320e (108TB FlexBuff capacity),
not worth the savings
```

Enterprise Drives...?

- From BackBlaze (commercial backup service provider)

Enterprise Failure Rates

- Annualized failure rate
  - 4.6% Enterprise
  - 4.2% Consumer
- Makes one suspect they are identical disks...

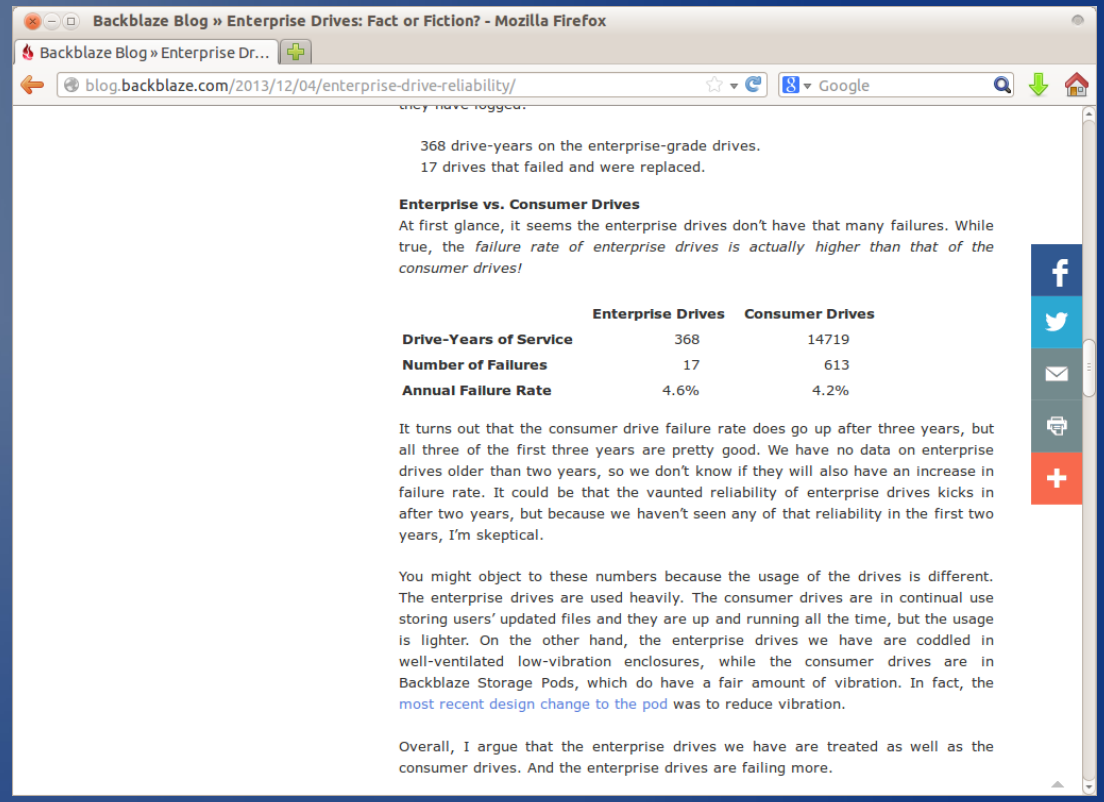

Disk Life

- Disks last maybe longer than generally expected
- Bathtub curve applies well

Disk Life

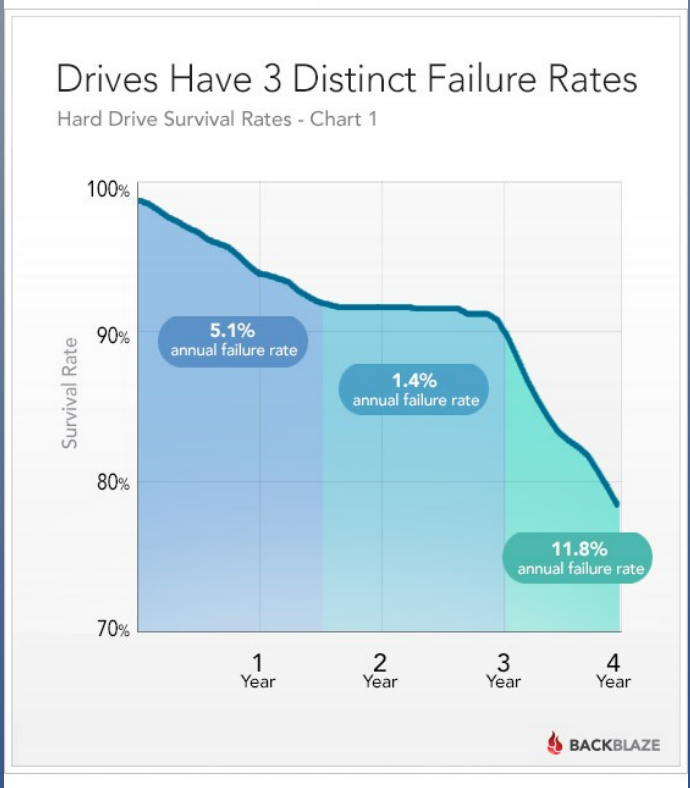
Drives Have 3 Distinct Failure Rates
Hard Drive Survival Rates - Chart 1
Survival Rate
100%
90%
80%
70%
5.1%
annual failure rate
1.4%
annual failure rate
11.8%
annual failure rate
1 Year
2 Year
3 Year
4 Year
BACKBLAZE

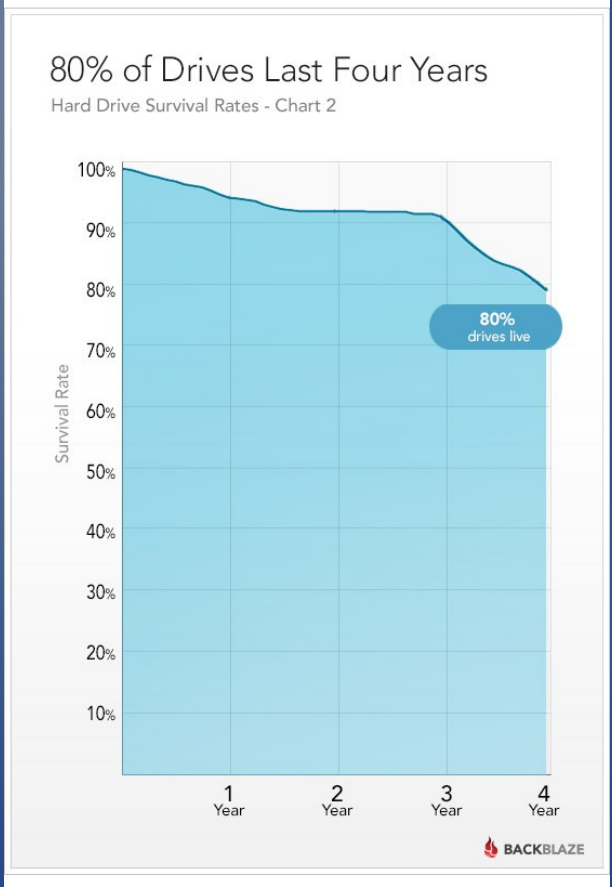
80% of Drives Last Four Years
Hard Drive Survival Rates - Chart 2
Survival Rate
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
80%
drives live
1 Year
2 Year
3 Year
4 Year
BACKBLAZE

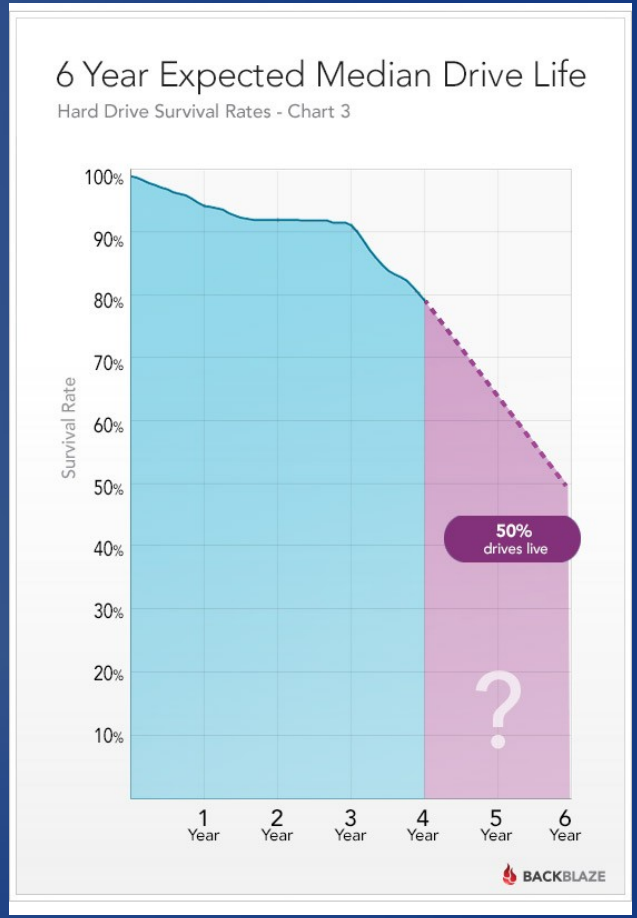
6 Year Expected Median Drive Life
Hard Drive Survival Rates - Chart 3
Survival Rate
100%
90%
80%
70%
60%
50%
40%
30%
20%
10%
50%
drives live
?
1 Year
2 Year
3 Year
4 Year
5 Year
6 Year
BACKBLAZE

Disk Brand Names..?

- On Tue 21-Jan-2014 BackBlaze published _their_ stats
- Be careful since “per-model” & “per-batch” numbers vary a lot...

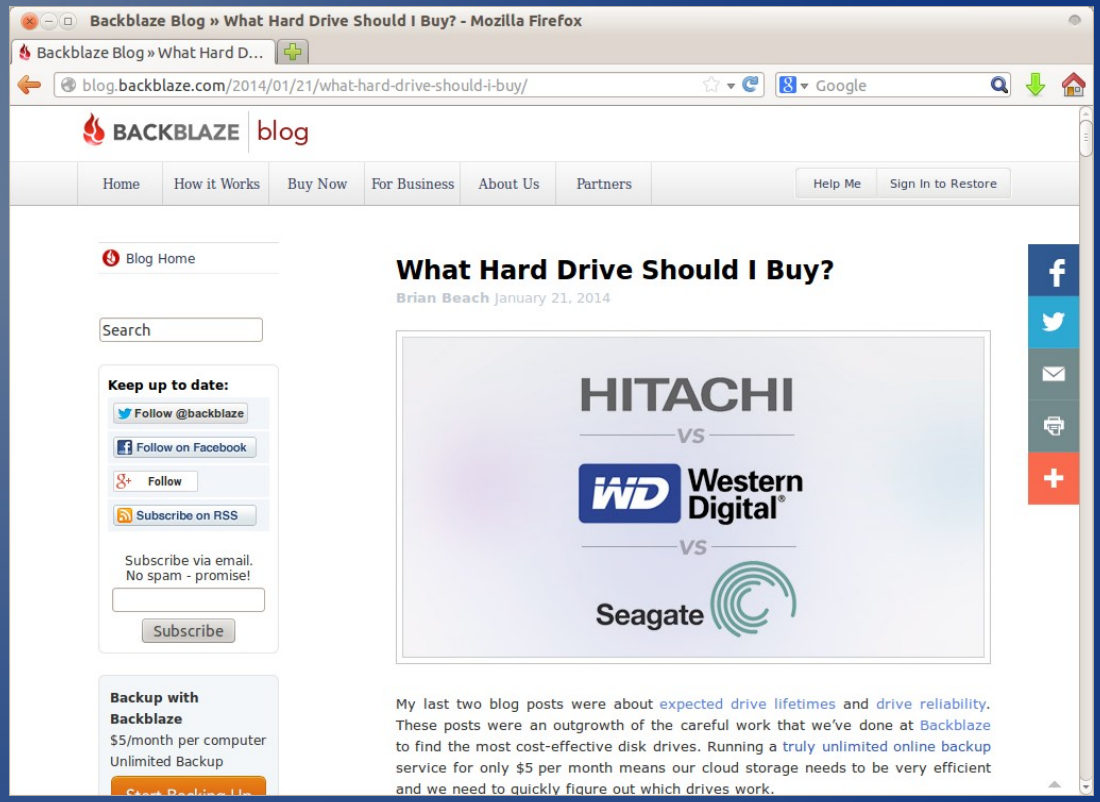

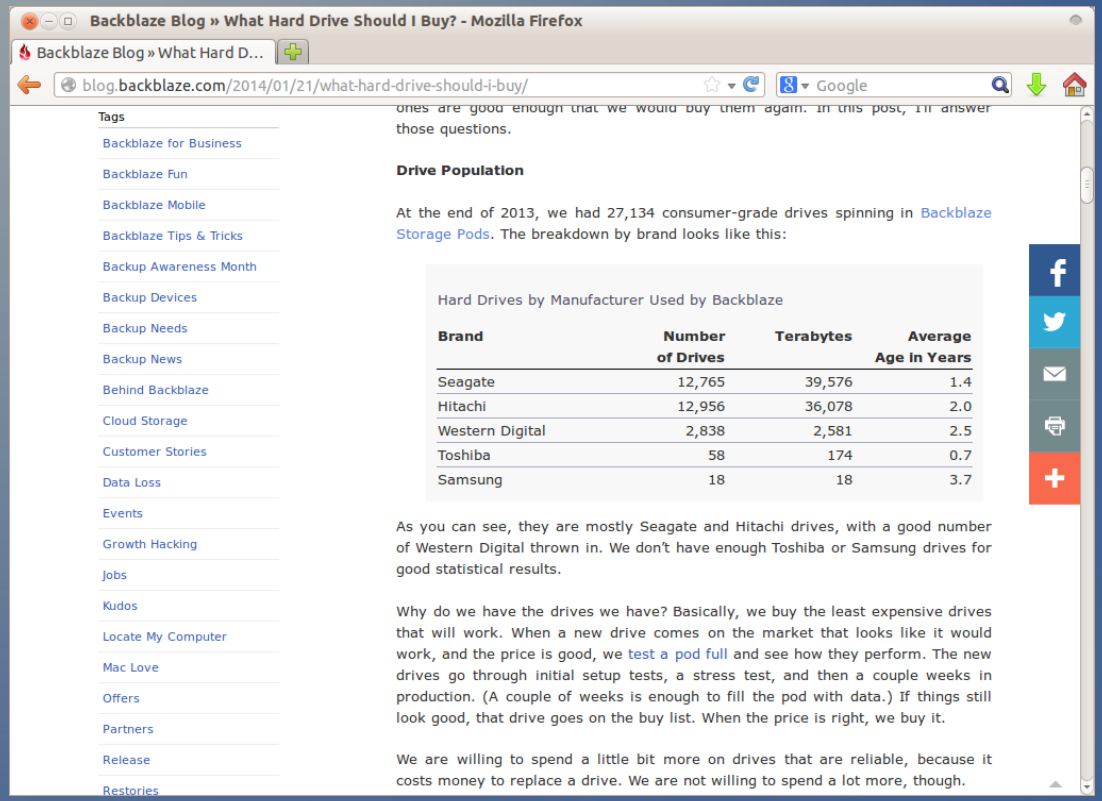

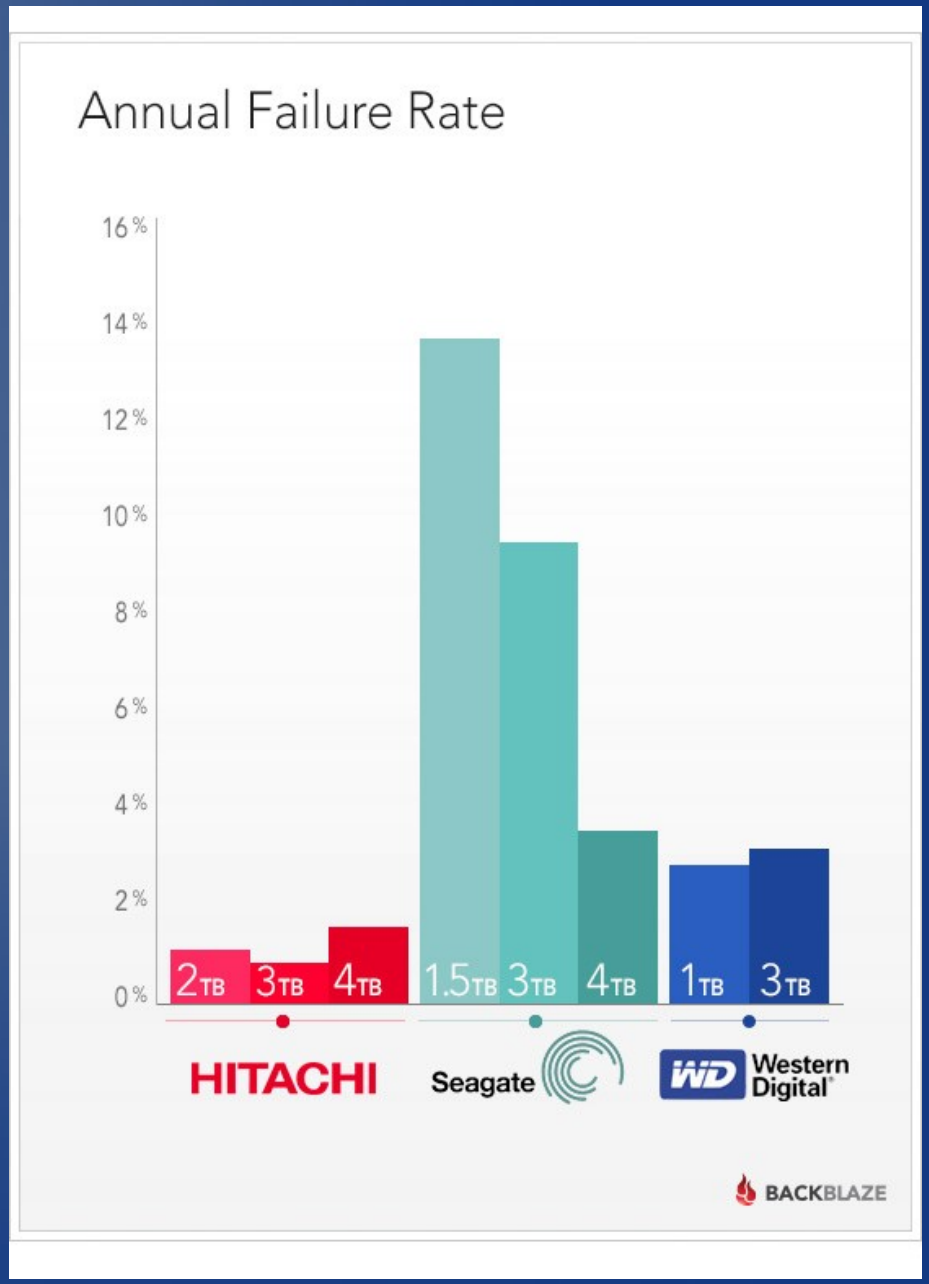

- BackBlaze buying now: Seagate Desktop HDD.15 (ST4000DM000)...!