Rappels théoriques concernant les architectures de stockage traditionnel

Plus de post-scriptum dans mes inscriptions aux ANF ?

[…]
Ayant vécu sur la liste ASR des moments intenses de surprise à la lecture des choix de collègues sur des RAID ou des médias de stockage HDD/SSD, je ferai volontiers une intervention
[...]

Qui suis-je ?

Pourquoi le stockage ?

Histoire d'un labo

- 2004 100Go
- 2008 1To
- 2012 10To
- 2016 100To

« le disque remplace la bande, et la mémoire Flash remplace le disque »

Jim Gray,
Microsoft Research Dec.2006

Le couple disque/bande reste encore très « compétitif »

Gilles Requilé LMGC

RESINFO Journée Sauvegarde 23/3/06

Supports et architectures de sauvegarde de données

Le support magnétique

Camembert rotatif

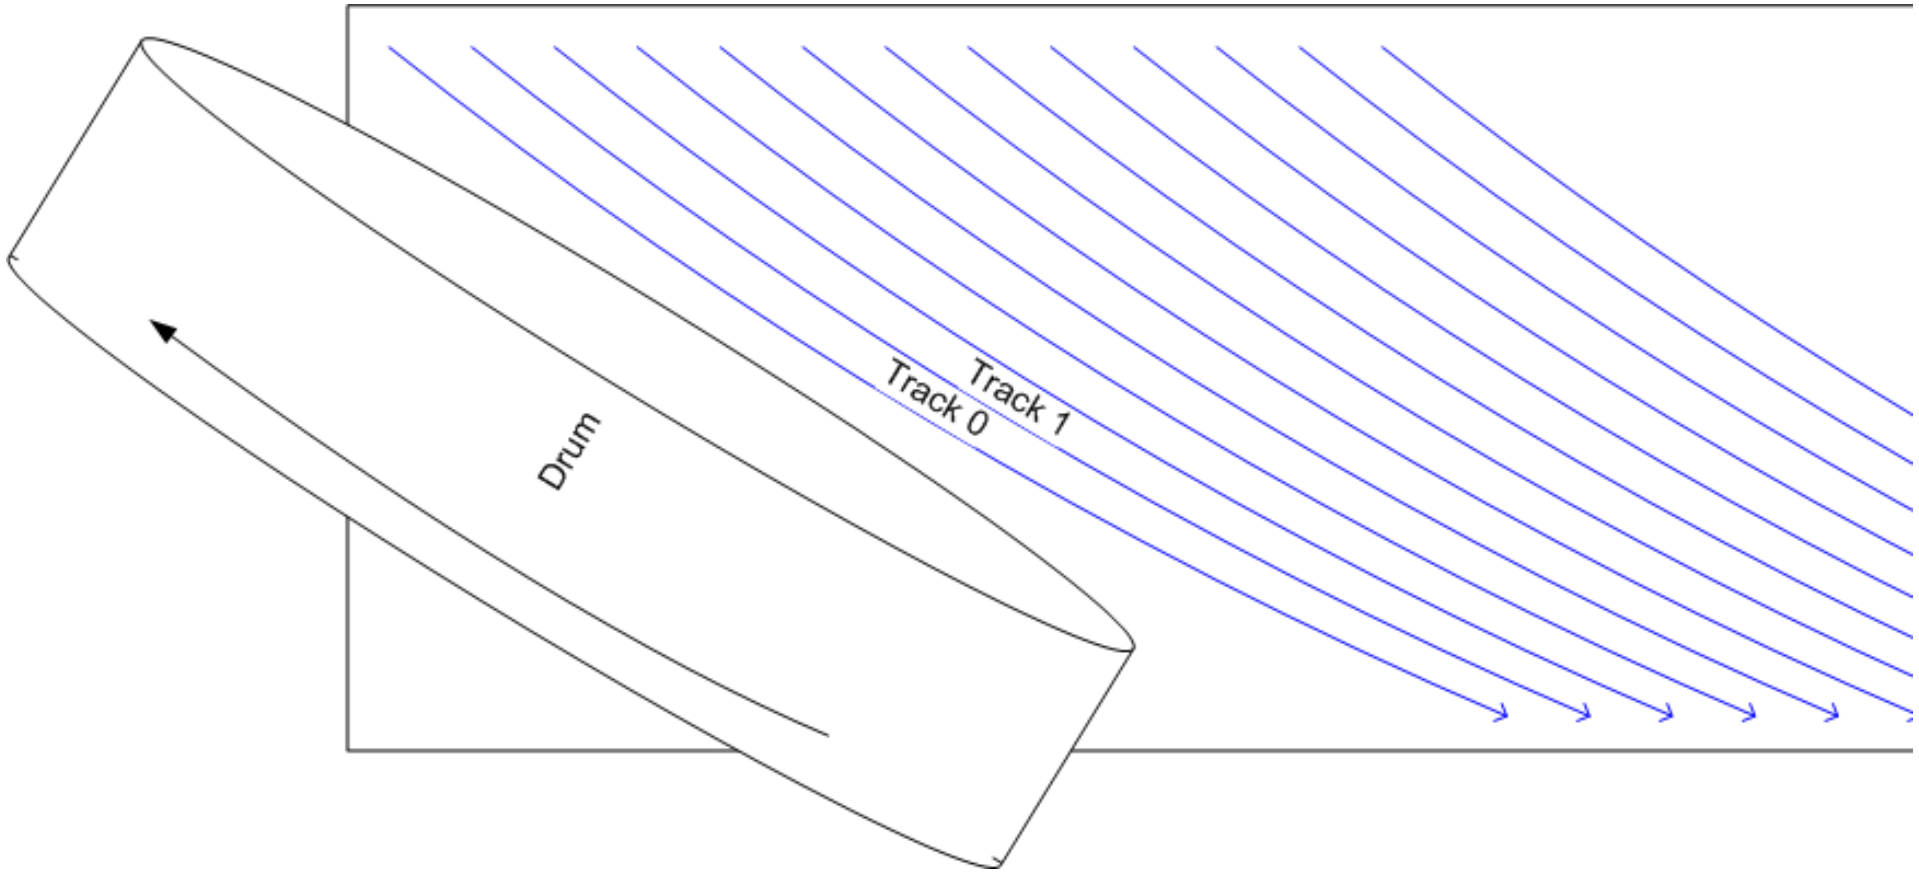

La gigue du bras

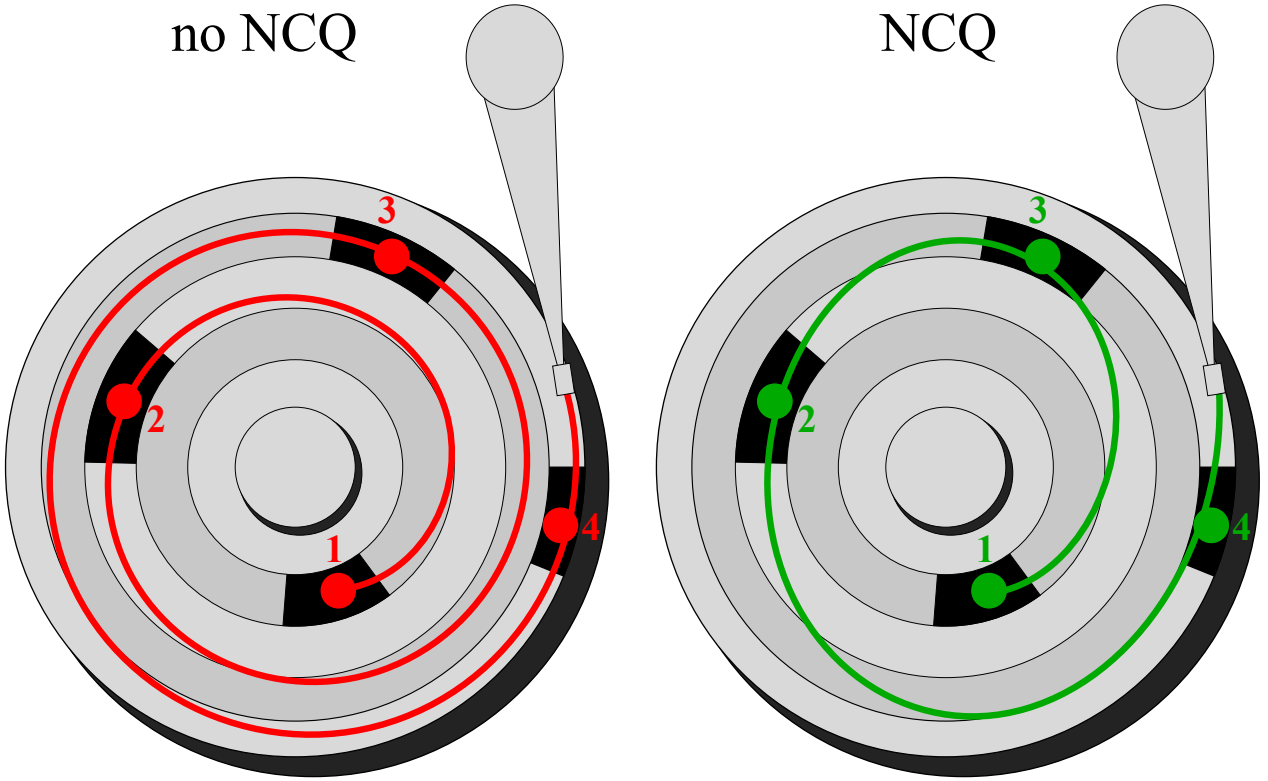

Fini les équations de Maxwell's

millipede

?

Non-volatile random-access memory (NVRAM) is random-access memory that retains its information when power is turned off (non-volatile)

Le match des années 2000

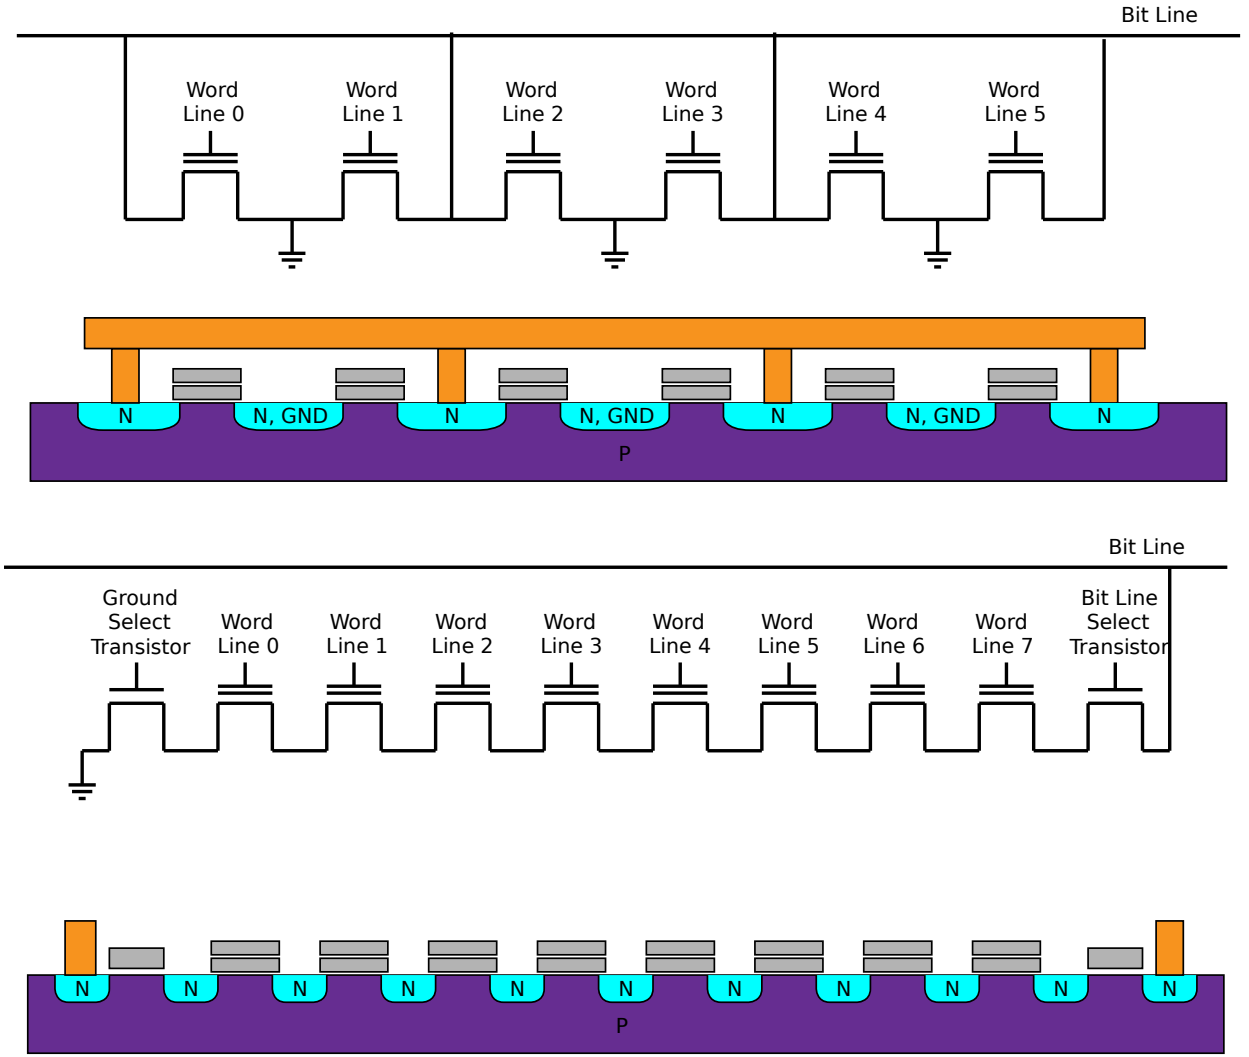

10000 écritures et moins ?

Pour modifier 1 bit, en écrire 512 ?

La répartition de l'usure.

Les caractéristiques

L'entrée/sortie

Et l'art d'attendre

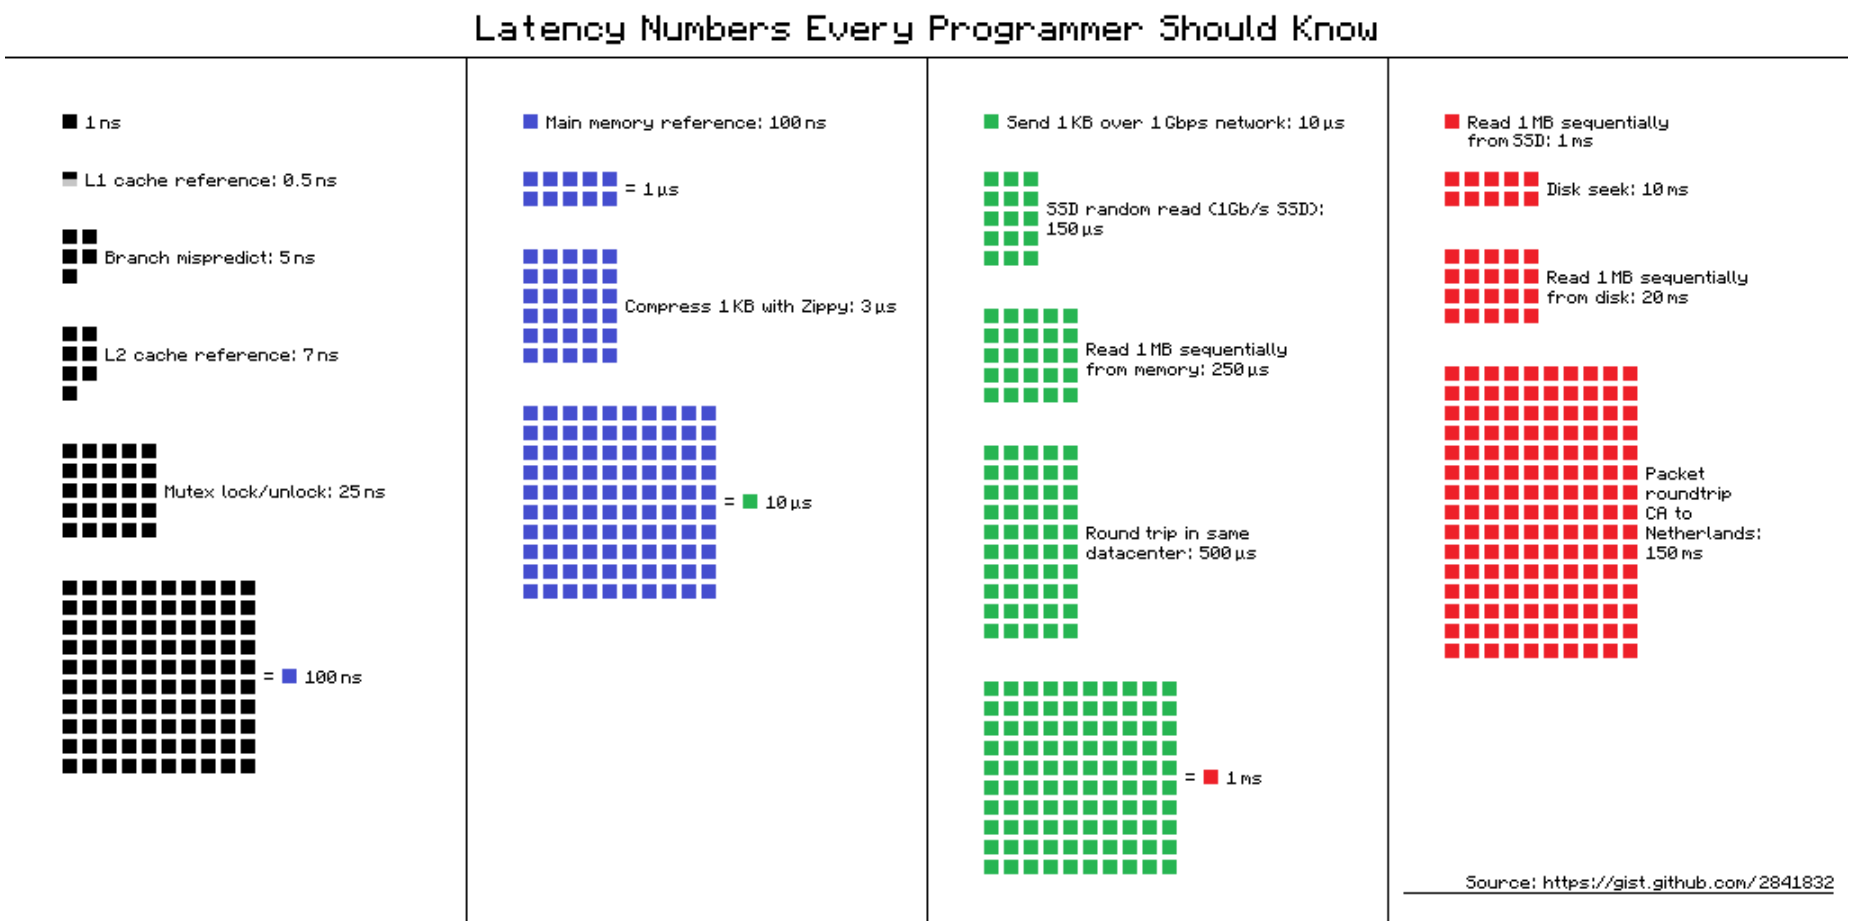

https://people.eecs.berkeley.edu/~rcs/research/interactive_latency.html

Les caractéristiques

La bande passante

Les caractéristiques

La durée de vie

Les caractéristiques

Le taux d'erreur

Les caractéristiques

Le taux de panne

Intel® Solid-State Drive DC S3710 Series

Product Specification

- Capacity:
  - 200GB, 400GB, 800GB, 1.2TB
- Components:
  - Intel® 20nm NAND Flash Memory
  - High Endurance Technology (HET) Multi-Level Cell (MLC)
- Form Factor: 2.5–inch
- Read and Write IOPS[1,2] (Full LBA Range, IOMeter* Queue Depth 32)
  - Random 4KB[3] Reads: Up to 85,000 IOPS
  - Random 4KB Writes: Up to 45,000 IOPS
  - Random 8KB[3] Reads: Up to 52,000 IOPS
  - Random 8KB Writes: Up to 21,000 IOPS
- Bandwidth Performance[1]
  - Sustained Sequential Read: Up to 550 MB/s[4]
  - Sustained Sequential Write: Up to 520 MB/s
- Endurance: 10 drive writes per day[5] for 5 years
  - 200GB: 3.6PB 400GB: 8.3PB
  - 800GB: 16.9PB 1.2TB: 24.3PB
- Latency (average sequential)
  - Read: 55 µs (TYP)
  - Write: 66 µs (TYP)
- Quality of Service[6,8]
  - Read/Write: 500 µs / 5 ms (99.9%)
- Performance Consistency[7,8]
  - Read/Write: Up to 90%/90% (99.9%)
- AES 256-bit Encryption
- Altitude[9]
  - Operating: -1,000 to 10,000 ft
  - Operating[10]: 10,000 to 15,000 ft
  - Non-operating: -1,000 to 40,000 ft
- Product Ecological Compliance
  - RoHS*
- Compliance
  - SATA Revision 3.0; compatible with SATA 6Gb/s, 3Gb/s and 1.5Gb/s interface rates
  - ATA/ATAPI Command Set – 2 (ACS-2 Rev 7); includes SCT (Smart Command Transport) and device statistics log support
  - Enhanced SMART ATA feature set
  - Native Command Queuing (NCQ) command set
  - Data set management Trim command
- Power Management
  - 5V or 5V+12V SATA Supply Rail[11]
  - SATA Interface Power Management
  - OS-aware hot plug/removal
  - Enhanced power-loss data protection feature
- Power[12]
  - Active: Up to 6.9 W (TYP)[8]
  - Idle: 600 mW
- Weight:
  - 200GB: 82 grams ± 2 grams
  - 400GB: 82 grams ± 2 grams
  - 800GB: 88 grams ± 2 grams
  - 1.2TB: 94 grams ± 2 grams
- Temperature
  - Operating: 0° C to 70° C
  - Non-Operating[13]: -55° C to 95° C
  - Temperature monitoring and logging
  - Thermal throttling
- Shock (operating and non-operating): 1,000 G/0.5 ms
- Vibration
  - Operating: 2.17 G_{RMS} (5-700 Hz)
  - Non-Operating: 3.13 G_{RMS} (5-800 Hz)
- Reliability
  - Uncorrectable Bit Error Rate (UBER): 1 sector per $10^{\wedge 17}$ bits read
  - Mean Time Between Failures (MTBF): 2 million hours
  - End-to-End data protection
- Certifications and Declarations
  - UL*, CE*, C-Tick*, BSMI*, KCC*, Microsoft* WHCK, VCCI*, SATA-IO*
- Compatibility
  - Windows 7* and Windows 8*, and Windows 8.1*
  - Windows Server 2012* R2
  - Windows Server 2012*
  - Windows Server 2008* Enterprise 32/64bit SP2
  - Windows Server 2008* R2 SP1
  - Windows Server 2003* Enterprise R2 64bit SP2
  - VMWare* 5.1, 5.5
  - Red Hat* Enterprise Linux* 5.5, 5.6, 6.1, 6.3, 7.0
  - SUSE* Linux* Enterprise Server 10, 11 SP1
  - CentOS* 64bit 5.7, 6.3
  - Intel® SSD Toolbox with Intel® SSD Optimizer

1. Performance values vary by capacity
2. Performance specifications apply to both compressible and incompressible data
3. 4KB = 4,096 bytes; 8KB = 8,192 bytes.
4. MB/s = 1,000,000 bytes/second
5. Based on JESD218 standard.
6. Based on Random 4KB QD=1 workload, measured as the time taken for 99.9 percentile of commands to finish the round-trip from host to drive and back to host
7. Based on Random 4KB QD=32 workload, measured as the (IOPS in the 99.9th percentile slowest 1-second interval)/(average IOPS during the test)
8. Measurement taken once the workload has reached steady state but including all background activities required for normal operation and data reliability
9. Altitude pressure is simulated in a test chamber; excludes soft error
10. Extended operation at a higher altitude might impact reliability.
11. If both 12V and 5V power supplies are present, defaults to 5V+ 12V power supplies. Does not support 12 volt only.
12. Based on 5Vpower supply
13. Please contact your Intel representative for details on the non-operating temperature range

SEAGATE

Enterprise Capacity 3.5 HDD

VERSION 5

Caractéristiques	4Kn SAS			
Capacité	8 To	6 To	4 To	2 To
Modèle de base	ST8000NM0065	ST6000NM0105	ST4000NM0095	ST2000NM0115
Modèle doté de PowerBalance™	—	—	—	—
Modèle Seagate Secure®	ST8000NM0095	ST6000NM0205	ST4000NM0075	—
Modèle Seagate Secure avec autochiffrement FIPS[1,2]	ST8000NM0125	ST6000NM0255	—	—
Fonctions spécifiques				
Protection des informations (DIF T10)	Oui	Oui	Oui	Oui
Capteur d'humidité	Oui	Oui	Oui	Oui
Super parité	Oui	Oui	Oui	Oui
Sans halogène	Oui	Oui	Oui	Oui
Technologie PowerChoice™	Oui	Oui	Oui	Oui
Technologie PowerBalance™	Oui	Oui	Oui	Oui
Mémoire cache	256 Mo	256 Mo	128 Mo	128 Mo
Mise en cache des écritures avancée (Flash NOR interne)	Oui	Oui	Oui	Oui
Fiabilité/Intégrité des données				
Vibrations, hors fonctionnement : 10 Hz – 500 Hz (G)	5	5	5	5
Temps moyen entre deux pannes	2 000 000 heure	2 000 000 heure	2 000 000 heure	2 000 000 heure
Taux de panne annualisé pour un fonctionnement 24h/24 et 7j/7 (AFR)	0,44 %	0,44 %	0,44 %	0,44 %
Erreurs de lecture irrécupérables par bit lu, max.	1 secteur par 10E15	1 secteur par 10E15	1 secteur par 10E15	1 secteur par 10E15
Durée de fonctionnement (heures)	8 760	8 760	8 760	8 760
Octets par secteur	4 096, 4 160, 4 224	4 096, 4 160, 4 224	4 096, 4 160, 4 224	4 096, 4 160, 4 224
Durée de la garantie limitée (années)	5	5	5	5
Performances				
Vitesse de rotation (tr/min)	7 200	7 200	7 200	7 200
Vitesse d'accès interface (Gbits/s)	12,0, 6,0, 3,0	12,0, 6,0, 3,0	12,0, 6,0, 3,0	12,0, 6,0, 3,0
Taux de transfert en continu max.	249 Mo/s	226 Mo/s	226 Mo/s	226 Mo/s
Latence moyenne (ms)	4,16	4,16	4,16	4,16
Ports d'interface	Double	Double	Double	Double
Vibrations rotationnelles à 1 500 Hz (rad/s²)	12,5	12,5	12,5	12,5
Consommation				
Consommation au repos, moyenne (W)	8,5 W	8,37 W	6,52 W	4,57 W
Consommation moyenne en fonctionnement	12 W	9,2 W	7,74 W	6,34 W
Alimentation requise	+12 V et +5 V	+12 V et +5 V	+12 V et +5 V	+12 V et +5 V
Caractéristiques environnementales/Température				
Température en fonctionnement (°C)	5 °C – 60 °C	5 °C – 60 °C	5 °C – 60 °C	5 °C – 60 °C
Résistance aux chocs en fonctionnement (lecture/écriture) (G)	70/40 G	70/40 G	70/40 G	70/40 G
Résistance aux chocs, hors fonctionnement, 1 ms/2 ms (G)	150/250	150/250	150/300	150/300
Caractéristiques physiques				
Hauteur (mm/po)[3]	26,1 mm/1,028 po.	26,1 mm/1,028 po.	26,1 mm/1,028 po.	26,1 mm/1,028 po.
Largeur (mm/po, maximum)[3]	101,85 mm/4,01 po.	101,85 mm/4,01 po.	101,85 mm/4,01 po.	101,85 mm/4,01 po.
Profondeur (mm/po, maximum)[3]	147 mm/5,787 po.	147 mm/5,787 po.	147 mm/5,787 po.	147 mm/5,787 po.
Poids (g/lb)	780 g/1,72 lb	705 g/1,55 lb	680 g/1,5 lb	610 g/1,34 lb
Unités par carton	20	20	20	20
Cartons par palette/Cartons par niveau	40/8	40/8	40/8	40/8

1 Les disques avec autochiffrement et les disques avec autochiffrement certifié FIPS 140-2 Validated ne sont pas disponibles dans tous les modèles ni dans tous les pays. Certains modèles requièrent la prise en charge d'un contrôleur ou d'un hôte compatible avec le TCG.
2 Certification FIPS 140-2 de niveau 2 en cours.
3 Ces dimensions du boîtier de base sont conformes à la norme Small Form Factor (SFF-8201) relative à la compacité de l'appareil que vous pouvez consulter à la page www.sffcommittee.org (en anglais uniquement). Pour obtenir des informations sur les dimensions des connecteurs, référez-vous à la norme SFF-8223.

		SSD	HDD
Latence	:	66µs	4.16ms
BW	:	8k*21.10^3=172Mo/s	1/(4.16*10^-3)*4096=1Mo/s
AFR	:	0.44 %/an	0.44 %/an
UBER	:	10^17/4k = $2.4x10^{13}$	10^15/4k = $2.4x10^{11}$

MTBF(hrs) = 876,000 / AFR(%)

La vie réelle

Hard Drive Failure Stats through 3/31/2016

Cumulative from 4/2013 through period indicated

			3/31/2014 (1 year)		3/31/2015 (2 years)		03/31/2016 (3 years)	
MFG	Model	Drive Size	Drive Count	Annualized Failure Rate	Drive Count	Annualized Failure Rate	Drive Count	Annualized Failure Rate
HGST	HDS5C3030ALA630	3TB	4,591	0.85%	4,596	0.74%	4,552	0.81%
HGST	HDS5C4040ALE630	4TB	2,582	1.33%	2,653	1.16%	2,706	1.03%
HGST	HDS722020ALA330	2TB	4,713	1.08%	4,664	1.15%	4,264	1.57%
HGST	HDS723030ALA640	3TB	1,020	1.54%	1,013	1.83%	998	1.71%
HGST	HMS5C4040ALE640	4TB	47	2.67%	7,026	1.18%	7,075	0.79%
HGST	HMS5C4040BLE640	4TB	494	20.29%	3,100	0.48%	3,091	0.38%
HGST	HUH728080ALE600	8TB	—	—	—	—	45	3.84%
Seagate	ST3000DM001	3TB	4,074	13.92%	485	28.26%	—	—
Seagate	ST31500341AS	1.5TB	404	22.27%	259	24.12%	—	—
Seagate	ST31500541AS	1.5TB	1,746	9.87%	1,485	10.18%	45	10.12%
Seagate	ST32000542AS	2TB	211	8.03%	81	9.93%	—	—
Seagate	ST33000651AS	3TB	287	6.53%	234	5.27%	—	—
Seagate	ST4000DM000	4TB	8,800	3.83%	14,803	2.83%	34,729	2.90%
Seagate	ST4000DX000	4TB	179	0.75%	175	1.61%	207	2.95%
Seagate	ST6000DX000	6TB	—	—	495	1.70%	1,882	1.42%
Toshiba	DT01ACA300	3TB	58	4.63%	47	4.23%	47	4.22%
Toshiba	MD04ABA400V	4TB	—	—	—	—	146	2.21%
Toshiba	MD04ABA500V	5TB	—	—	—	—	45	2.05%
WDC	WD20EFRX	2TB	—	—	—	—	133	10.56%
WDC	WD30EFRX	3TB	578	8.78%	1,045	7.90%	1,054	6.74%
WDC	WD40EFRX	4TB	—	—	45	9.01%	46	2.14%
WDC	WD60EFRX	6TB	—	—	450	6.64%	458	5.71%

BACKBLAZE

D.A.Patterson,

G.Gibson,

et

R.H.Katz,

1988

3 A Solution: Arrays of Inexpensive Disks

Rapid improvements in capacity of large disks have not been the only target of disk designers, since personal computers have created a market for inexpensive magnetic disks These lower cost disks have lower performance as well as less capacity Table I below compares the top-of-the-line IBM 3380 model AK4 mainframe disk, Fujitsu M2361A "Super Eagle" minicomputer disk, and the Conner Peripherals CP 3100 personal computer disk

Characteristics	*IBM 3380*	*Fujitsu M2361A*	*Conners CP3100*	*3380 v 3100 (>1 means 3100 is better)*	*2361 v 3100*
Disk diameter (inches)	14	10 5	3 5	4	3
Formatted Data Capacity (MB)	7500	600	100	01	2
Price/MB(controller incl)	$18-$10	$20-$17	$10-$7	1-2 5	1 7-3
MTTF Rated (hours)	30,000	20,000	30,000	1	1 5
MTTF in practice (hours)	100,000	?	?	?	?
No Actuators	4	1	1	2	1
Maximum I/O's/second/Actuator	50	40	30	6	8
Typical I/O's/second/Actuator	30	24	20	7	8
Maximum I/O's/second/box	200	40	30	2	8
Typical I/O's/second/box	120	24	20	2	8
Transfer Rate (MB/sec)	3	2 5	1	3	4
Power/box (W)	6,600	640	10	660	64
Volume (cu ft)	24	3 4	03	800	110

Table I *Comparison of IBM 3380 disk model AK4 for mainframe computers, the Fujitsu M2361A "Super Eagle" disk for minicomputers, and the Conners Peripherals CP 3100 disk for personal computers By "Maximum I/O's/second" we mean the maximum number of average seeks and average rotates for a single sector access Cost and reliability information on the 3380 comes from widespread experience [IBM 87] [Gawlick87] and the information on the Fujitsu from the manual [Fujitsu 87], while some numbers on the new CP3100 are based on speculation The price per megabyte is given as a range to allow for different prices for volume discount and different mark-up practices of the vendors (The 8 watt maximum power of the CP3100 was increased to 10 watts to allow for the inefficiency of an external power supply, since the other drives contain their own power supplies)*

Select Mean Time Between Failures (MTBF): Pessimistic Estimate (36.5K)

Nonrecoverable Error Rate: 10^11

Drive Capacity: 8 TB

Sector Size: 4 KB

Quantity of Disks: 8

Volumes: 1

Volume Rebuild Speed (MB/s): 1

Submit

Faut calculer !

RAID Level	Formatted Capacity (GB)	Mean Time To Data Failure (MTTDF) in hours	Bit Error Rate MTTDL	Mean Time To Data Loss (MTTDL) in hours	MTTDL (Years)
RAID 0	60,293.12	4,562.50	< .01	2,281.25	0.26
RAID 1	30,146.56	155,364.04	< .01	77,682.02	8.87
RAID 10	30,146.56	155,364.04	< .01	77,682.02	8.87
RAID 5	52,756.48	1,585.35	4,562.50	3,073.92	0.35
RAID 6	45,219.84	874.75	1,849.57	1,362.16	0.16
RAID-Z3	37,683.20	857.88	13,501,875.06	6,751,366.47	770.70
RAID 50	45,219.84	17,262.67	18,250.00	17,756.34	2.03
RAID 60	30,146.56	55,109.56	25,894.01	40,501.78	4.62

RAID Level	1 Year	2 Years	3 Years	4 Years	5 Years	6 Years	7 Years	8 Years	9 Years	10 Years
RAID 0	97.86383 %	99.95437 %	99.99903 %	99.99998 %	100.00000 %	100.00000 %	100.00000 %	100.00000 %	100.00000 %	100.00000 %
RAID 1	10.66167 %	20.18663 %	28.69607 %	36.29826 %	43.08993 %	49.15749 %	54.57815 %	59.42088 %	63.74729 %	67.61244 %
RAID 10	10.66167 %	20.18663 %	28.69607 %	36.29826 %	43.08993 %	49.15749 %	54.57815 %	59.42088 %	63.74729 %	67.61244 %
RAID 5	94.25674 %	99.67015 %	99.98106 %	99.99891 %	99.99994 %	100.00000 %	100.00000 %	100.00000 %	100.00000 %	100.00000 %
RAID 6	99.80695 %	99.99963 %	100.00000 %	100.00000 %	100.00000 %	100.00000 %	100.00000 %	100.00000 %	100.00000 %	100.00000 %
RAID-Z3	0.12967 %	0.25917 %	0.38850 %	0.51766 %	0.64666 %	0.77549 %	0.90415 %	1.03265 %	1.16098 %	1.28914 %
RAID 50	38.89710 %	62.66435 %	77.18684 %	86.06050 %	91.48256 %	94.79560 %	96.81996 %	98.05690 %	98.81271 %	99.27453 %
RAID 60	19.46274 %	35.13749 %	47.76151 %	57.92855 %	66.11681 %	72.71140 %	78.02251 %	82.29993 %	85.74485 %	88.51929 %

RAID Performance Calculator

Compare two RAID configurations:

	Configuration #1	Configuration #2
RAID Type:	RAID 6	RAID 10
Drive Capacity (GB):	8000	8000
Single drive performance: ◉ **IO/s** ○ **MB/s**	170	170
Single drive cost:	410	410
Number of drives per RAID group:	5	58
Number of RAID groups:	12	1
Read operations (%):	10	10

Calculate

Results:	**Configuration #1**	**Configuration #2**
Total performance (IO/s):	**1854.55**	**5189.47**
Total usable storage capacity (TB)	**288.00**	**232.00**
RAID type:	RAID 6	RAID 10
Reads / Writes (%):	10 / 90	10 / 90
Number of RAID groups:	12	1
Number of drives per RAID group:	5	58
Total number of drives:	60	58
Single RAID group performance (IO/s):	154.55	5189.47
Capacity of a single RAID group (GB):	24000	232000
Single drive cost:	410	410
Cost per TB usable:	85.42	102.50
Total cost:	**24600.00**	**23780.00**

RAID Performance Calculator

Compare two RAID configurations:

	Configuration #1	Configuration #2
RAID Type:	RAID 6	RAID 6
Drive Capacity (GB):	8000	8000
Single drive performance: ◉ **IO/s** ○ **MB/s**	170	170
Single drive cost:	410	410
Number of drives per RAID group:	5	5
Number of RAID groups:	12	12
Read operations (%):	0	100

Calculate

Results:	**Configuration #1**	**Configuration #2**
Total performance (IO/s):	**1700**	**10200**
Total usable storage capacity (TB)	**288.00**	**288.00**
RAID type:	RAID 6	RAID 6
Reads / Writes (%):	0 / 100	100 / 0
Number of RAID groups:	12	12
Number of drives per RAID group:	5	5
Total number of drives:	60	60
Single RAID group performance (IO/s):	141.67	850
Capacity of a single RAID group (GB):	24000	24000
Single drive cost:	410	410
Cost per TB usable:	85.42	85.42
Total cost:	**24600.00**	**24600.00**

RAID Performance Calculator

Compare two RAID configurations:

	Configuration #1	Configuration #2
RAID Type:	RAID 6	RAID 6
Drive Capacity (GB):	8000	8000
Single drive performance: ○ IO/s ◉ MB/s	1	246
Single drive cost:	410	410
Number of drives per RAID group:	5	5
Number of RAID groups:	12	12
Read operations (%):	10	90

Calculate

Results:	**Configuration #1**	**Configuration #2**
Total performance (MB/s):	**10.91**	**9840**
Total usable storage capacity (TB)	**288.00**	**288.00**
RAID type:	RAID 6	RAID 6
Reads / Writes (%):	10 / 90	90 / 10
Number of RAID groups:	12	12
Number of drives per RAID group:	5	5
Total number of drives:	60	60
Single RAID group performance (MB/s):	0.91	820
Capacity of a single RAID group (GB):	24000	24000
Single drive cost:	410	410
Cost per TB usable:	85.42	85.42
Total cost:	**24600.00**	**24600.00**

SSD RAID Performance Calculator

This version of the calculator has been adapted to perform RAID performance calculation for solid state drives (SSD), devices with substantial difference in read and write performance. Examples of IOPS and throughput values for some SSD drives are provided in the table at the bottom of this page.
Supported RAID levels are RAID0, RAID1, RAID5, RAID6, RAID10 (1+0).

RAID Type:	RAID 6 (Stripe set with double parity)
Drive capacity (GB):	1600
Single drive performance:	(•) IO/s ( ) MB/s
Read performance:	85000
Write performance:	45000
Single drive cost:	1000
Number of drives per RAID group:	5
Number of RAID groups:	12
Read operations (%):	50

Calculate

Total Performance = 1055172.41 IO/s
Total usable capacity = 57.60 TB

RAID Type: RAID 6 (Stripe set with double parity)
Reads 50%, Writes 50%
Number of RAID groups = 12
Number of drives per RAID group = 5
Total number of drives = 60
Single RAID group performance = 87931.03 IO/s
Single drive cost = 1000
Cost per TB usable = 1041.67
Total cost = 60000.00

Notes:
Minimum number of drives per RAID6 group = 4
IO penalty (read) = 1/1 (one RAID IO per each host IO)
IO penalty (write) = 6/1 (6 RAID IOs per each host IO)
Fault tolerance = 2 disk drives per RAID group

IOPS (IO/s) and **throughput (MB/s)** are two common performance measurements for storage devices, such as disk drives and RAID sets. Most often IOPS measurement is used for random small block (4-8 KB) read and/or write operations typical for OLTP applications. MB/s is used for large block (>= 64 KB) random or sequential operations, such as file transfer and streaming. This calculator does not perform conversion between IO/s and MB/s, and vice versa. If you've switched the units, do not forget to enter applicable values in the Single drive performance fields.

Références bibliographiques

- J. S. Meena, S. M. Sze, U. Chand, and T.-Y. Tseng, “Overview of emerging nonvolatile memory technologies,” Nanoscale Research Letters, vol. 9, no. 1, p. 526, 2014.
- T. R. Albrecht et al., “Bit Patterned Magnetic Recording: Theory, Media Fabrication, and Recording Performance,” ArXiv e-prints, Mar. 2015.
- S. Bhargava and E. Yablonovitch, “HAMR Thermal Reliability via Inverse Electromagnetic Design,” ArXiv e-prints, Jul. 2014.
- “Hard Disk Drives with HAMR Technology Set to Arrive in 2018.” [Online]. Available: http://www.anandtech.com/show/9866/hard-disk-drives-with-hamr-technology-set-to-arrive-in-2018. [Accessed: 20-Nov-2016].
- M. L. Plumer, J. van Ek, and W. C. Cain, “New Paradigms in Magnetic Recording,” ArXiv e-prints, Jan. 2012.
- E. Yang et al., “Template-Assisted Direct Growth of 1 Td/in2Bit Patterned Media,” Nano Letters, vol. 16, pp. 4726–4730, Jul. 2016.
- R. F. L. Evans, R. W. Chantrell, U. Nowak, A. Lyberatos, and H.-J. Richter, “Thermally induced error: Density limit for magnetic data storage,” Applied Physics Letters, vol. 100, no. 10, p. 102402, 2012.
- J. S. Plank, “A Tutorial on Reed-Solomon Coding for Fault-Tolerance in RAID-like Systems,” Software – Practice & Experience, vol. 27, no. 9, pp. 995–1012, Sep. 1997.
- J. G. Elerath and M. Pecht, “Enhanced reliability modeling of raid storage systems,” in 37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN’07), 2007, pp. 175–184.
- C. Walter, “Kryder’s law,” Scientific American, vol. 293, no. 2, pp. 32–33, 2005.
- K. M. Greenan, J. S. Plank, and J. J. Wylie, “Mean Time to Meaningless: MTTDL, Markov Models, and Storage System Reliability.”
- D. A. Patterson, G. Gibson, and R. H. Katz, “A case for redundant arrays of inexpensive disks (RAID),” SIGMOD Rec., vol. 17, no. 3, pp. 109–116, 1988.
- “RAID Performace Calculator - WintelGuy.com.” [Online]. Available: http://wintelguy.com/raidperf2.pl. [Accessed: 11-Dec-2016].
- “RAID Reliability Calculator - Simple MTTDL Model - ServeTheHome.” [Online]. Available: https://www.servethehome.com/raid-calculator/raid-reliability-calculator-simple-mttdl-model/. [Accessed: 11-Dec-2016].
- Gray,Jim, “Tape is Dead Disk is Tape Flash is Disk RAM Locality is King,” Microsoft Research, Dec-2006.
- G. Graefe, “The five-minute rule twenty years later, and how flash memory changes the rules,” in Proceedings of the 3rd international workshop on Data management on new hardware, 2007, p. 6.
- A. Leventhal, “Triple-parity RAID and beyond,” Queue, vol. 7, no. 11, p. 30, 2009.
- “Understanding Mean Time between Failure in the Data Center - Part 2 - Schneider Electric Blog.” [Online]. Available: http://blog.schneider-electric.com/datacenter/2012/01/30/mean-time-between-failure-data-center-2/. [Accessed: 11-Dec-2016].

La patate chaude pour Benoît Parrein !