IBM Systems Storage

Storage Futures
Technology Outlook

Dr. Axel Koester
Enterprise Storage Consultant
axel.koester@de.ibm.com

Advances in Caching

Flash Memory

SAN Volume Controller & Flash

Innovations from the Labs

Hard Disk Performance *per GB* drops alarmingly

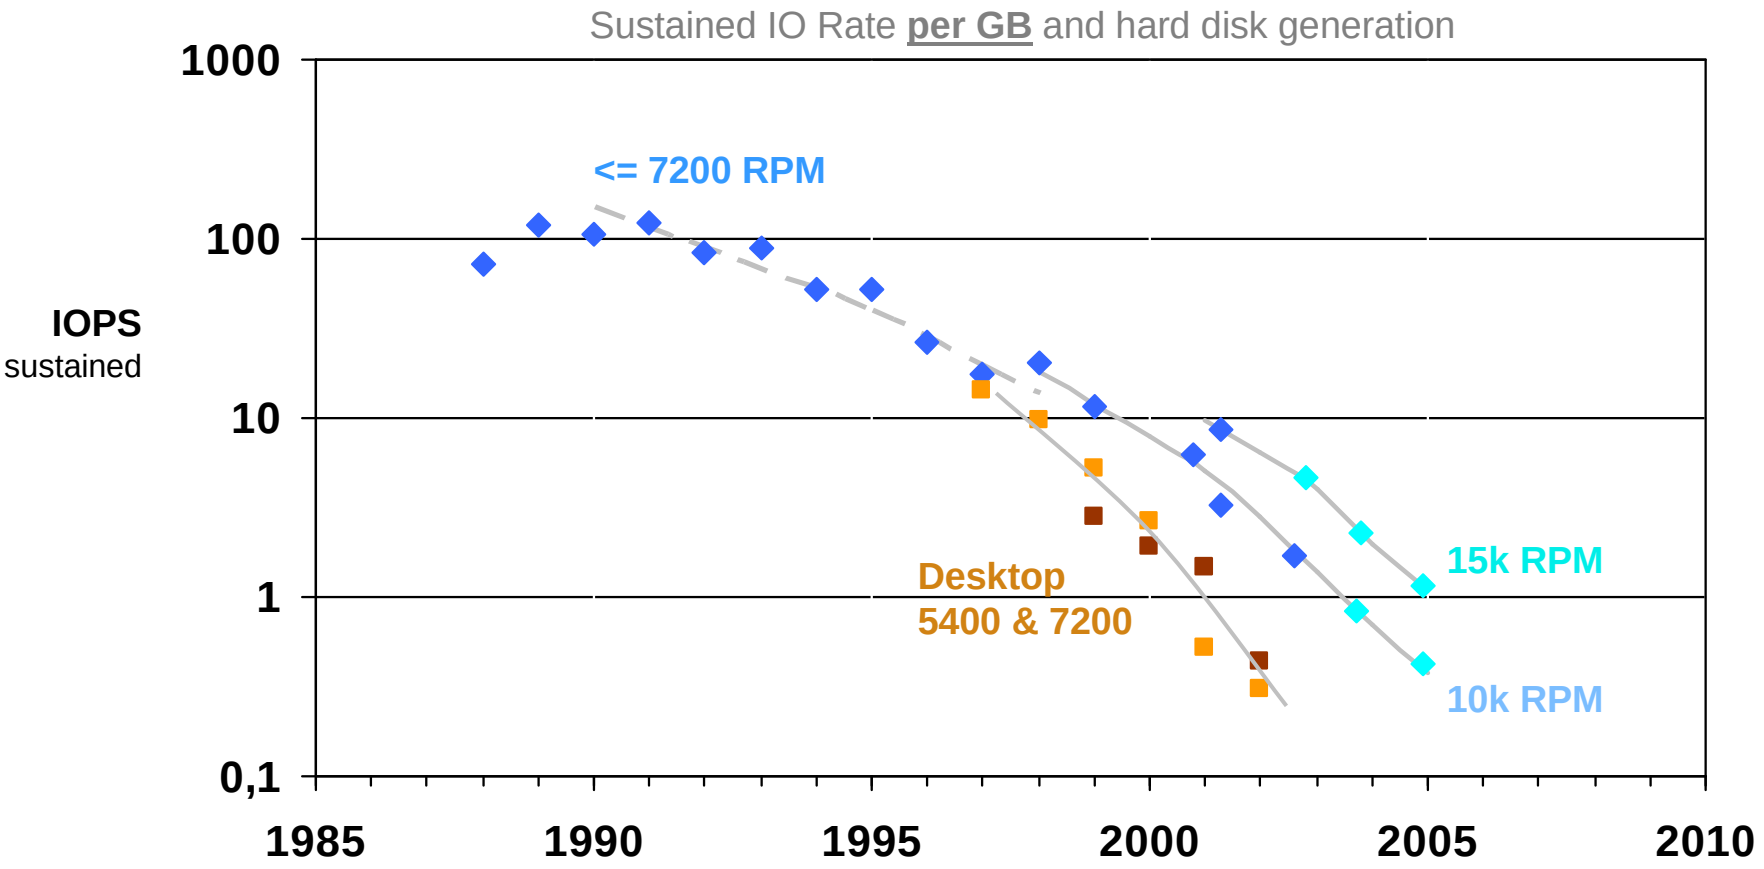

Quicker Access? 20k RPM?

- 20.000 RPM disks run **hot**
- **RPM ×2 = Power Consumption up ×8**
- => smaller platters, less capacity

Western Digital ®
VelociRaptor 2.5"
20kRPM Prototype

(*) Air Friction Loss ~ {RPM} ^{3}

Cache Innovations for larger Disks

Adaptive Multistream Prefetch Caching

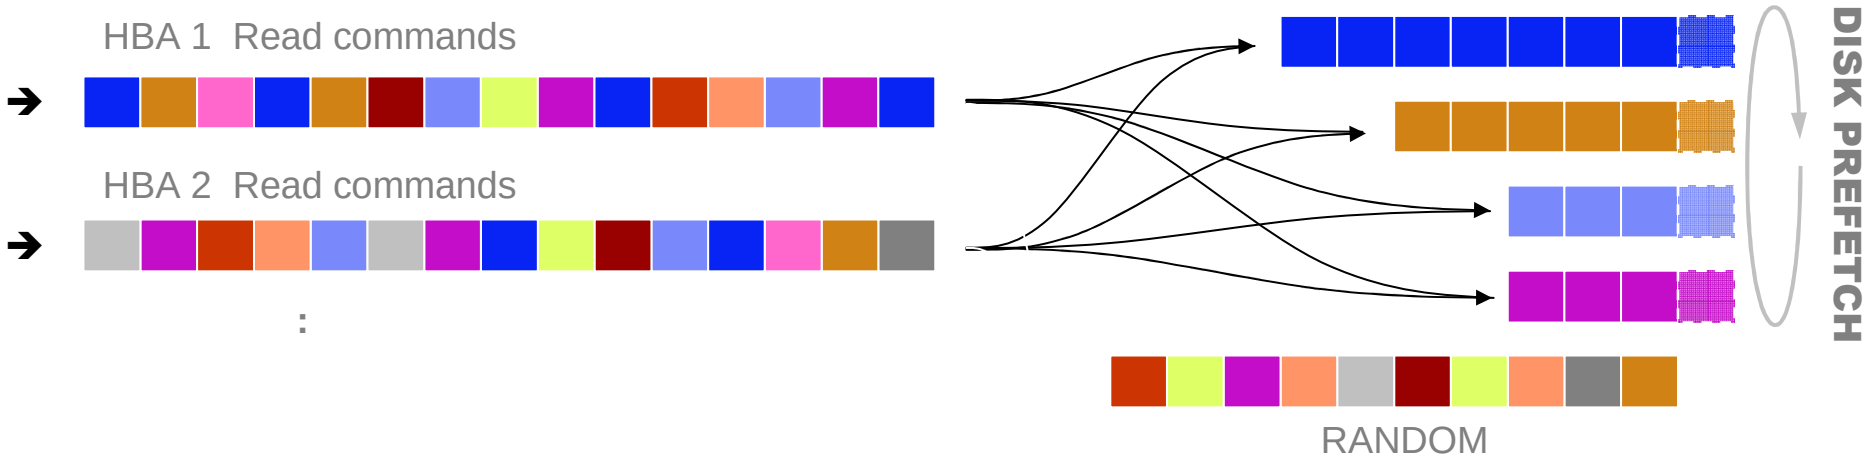

- Finds sequential data patterns in chaotic access streams
- …cross all adapters and IO clients
- …in realtime at > 120.000 IO/s

Caching Innovation – DB2 Real World Example

DB2 *Table Scans*, achieving cache speed with SARC & AMP
(= optimal prefetch prediction)

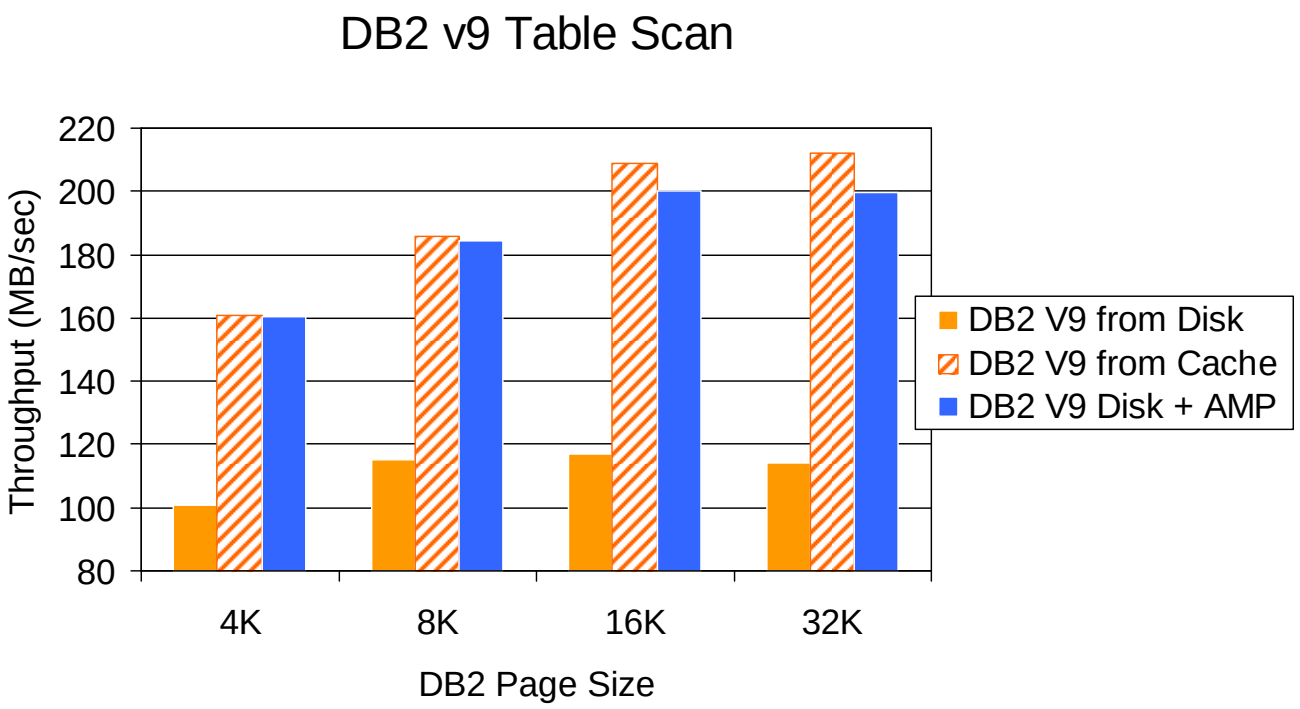

DB2 Table Scan with AMP is equivalent to Table Scan in Cache

Wise Order Writes (upcoming in DS8000)

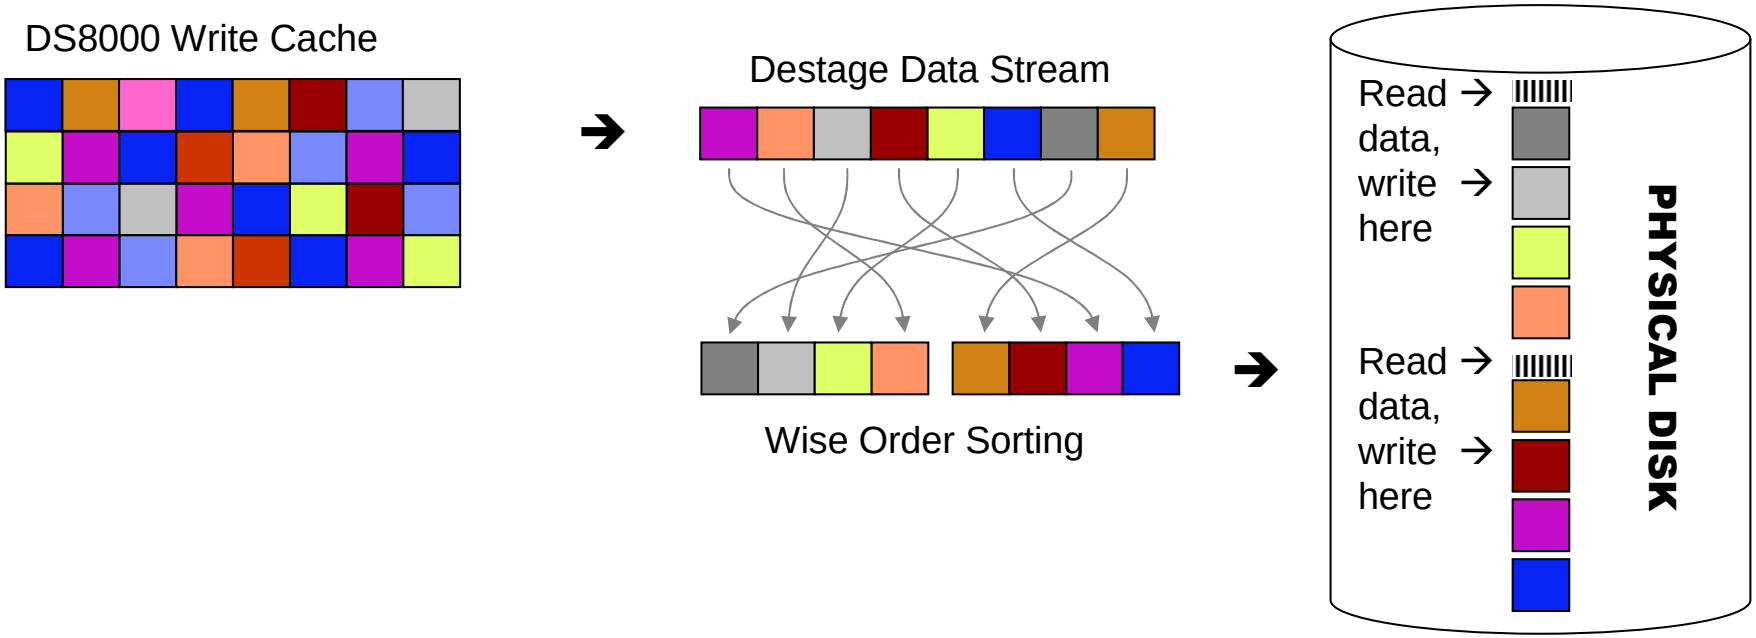

- Optimize head movements for low-latency read access
- Delay writes in cache until head is in proximity

➔ much better mixed workload behavior (close to 100% read)

Wise Order Writes compared with "classics"

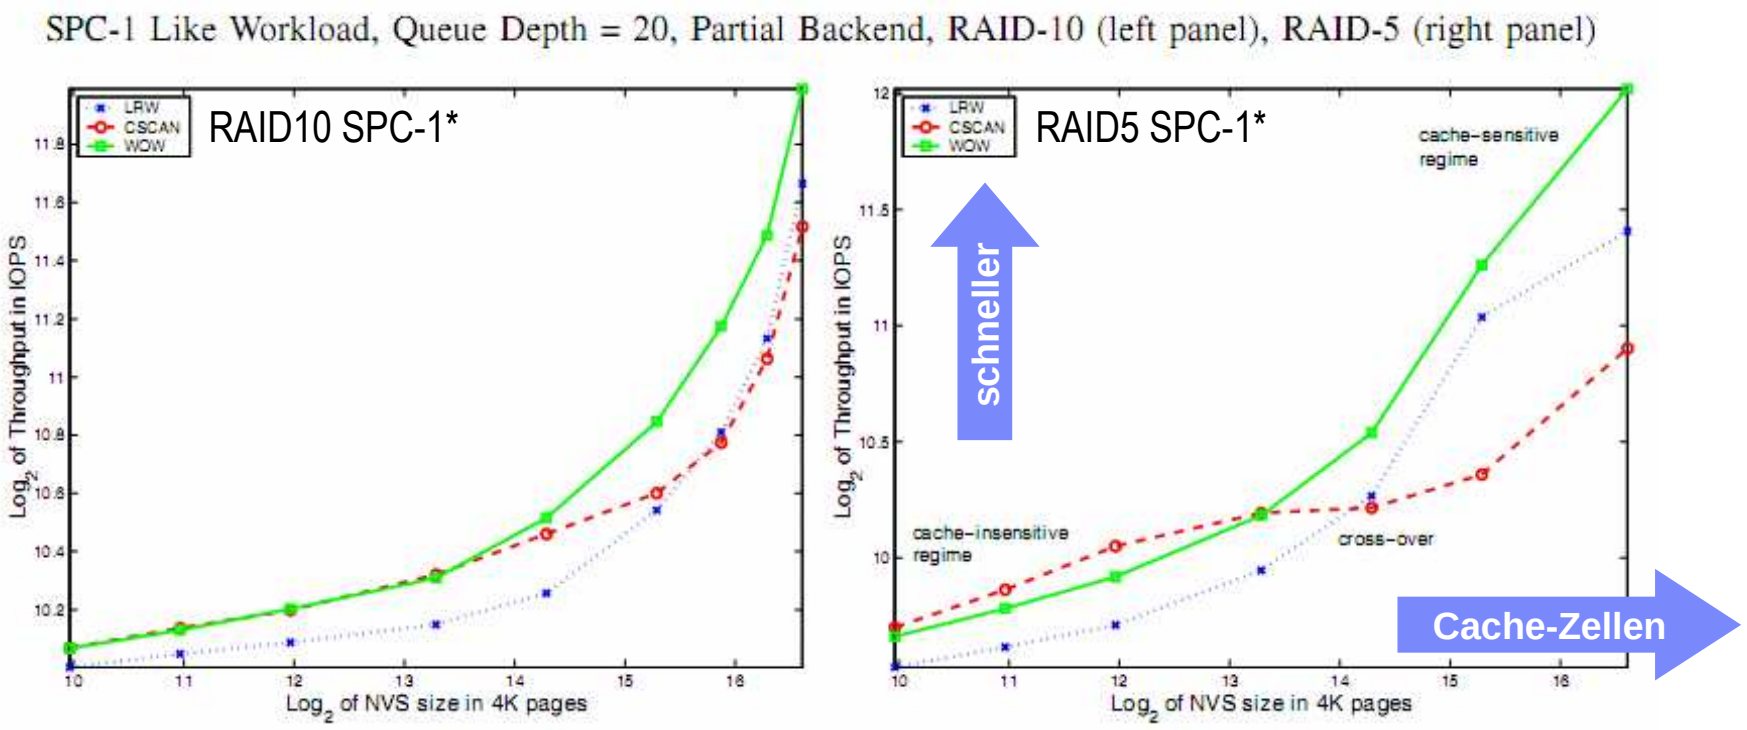

- **WOW** *(wise order writes)* versus "second best" **CSCAN** *(cyclical scan)* and classic **LRW** *(least recently written)* under SPC-1-like workload

(*) **simulated** SPC-1® OLTP Workload; R/W = 40/60, each 60% random

Cache Segmentation : DS8000 / SVC and competitors

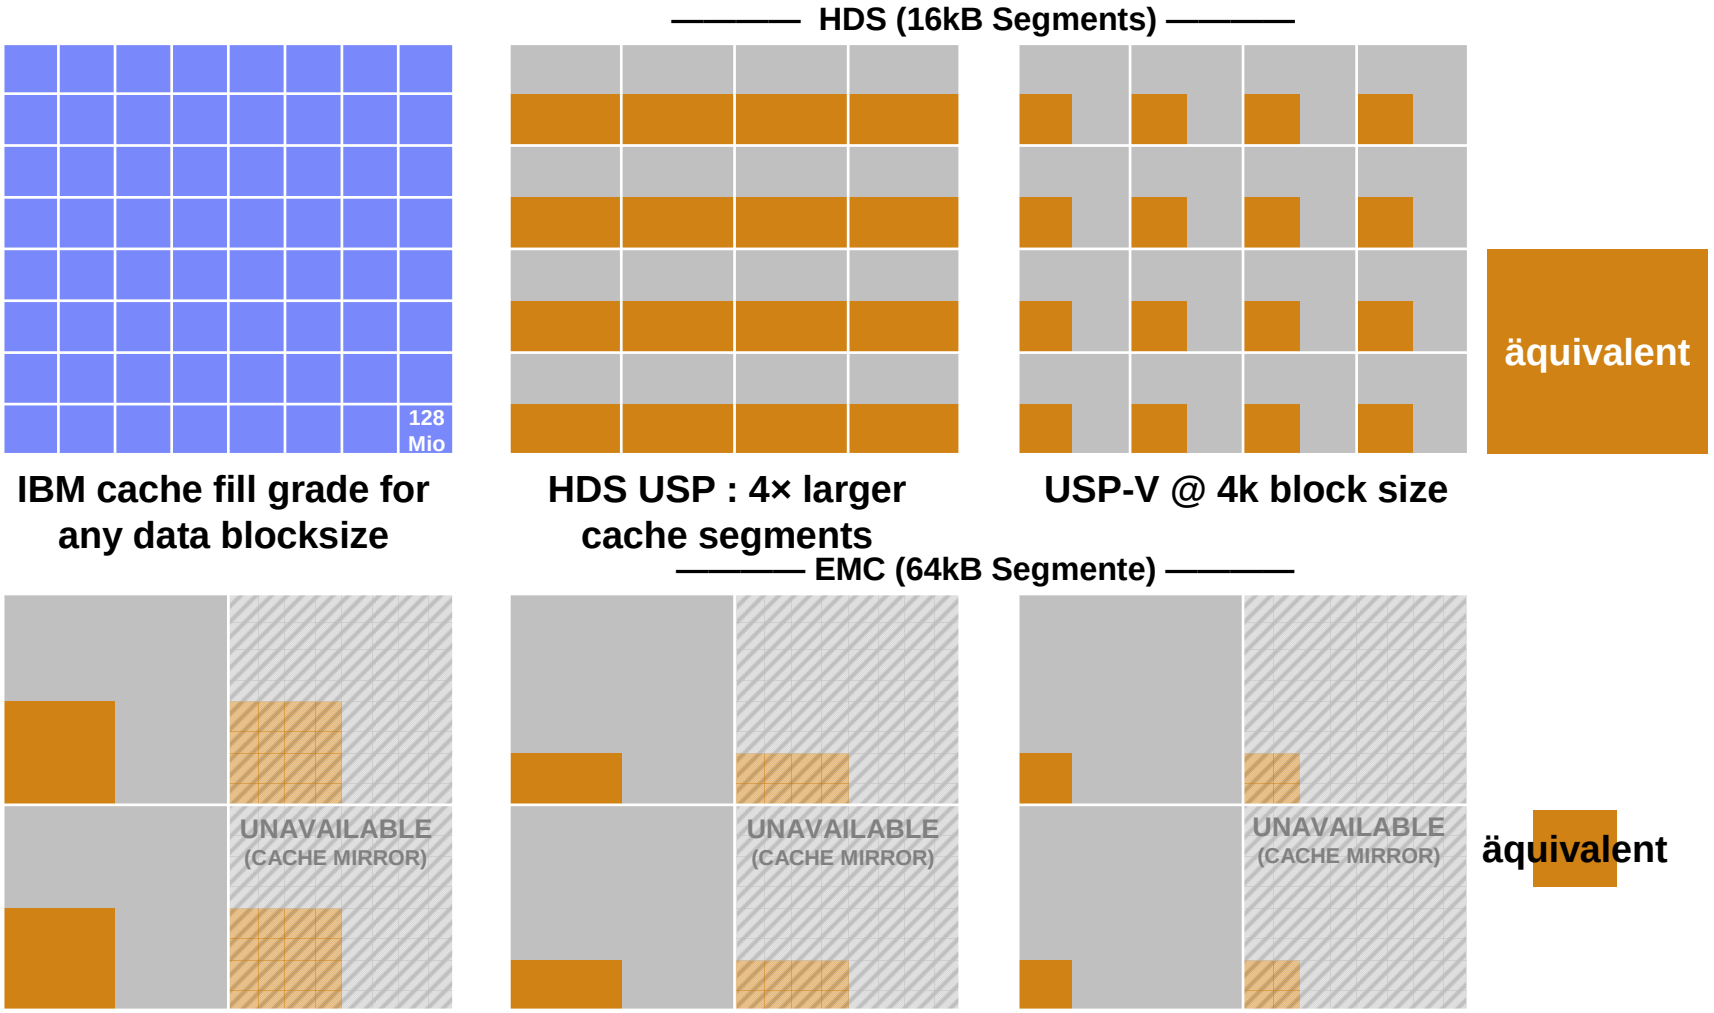

Flash Memory versus (?) Disks

Internal Flash Performance is non-trivial

Read

Very fast

Sequential Write

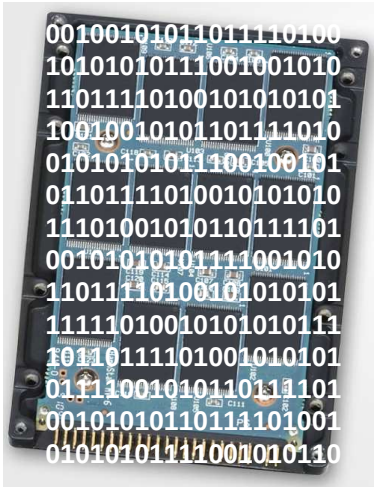

Fast

Random Write

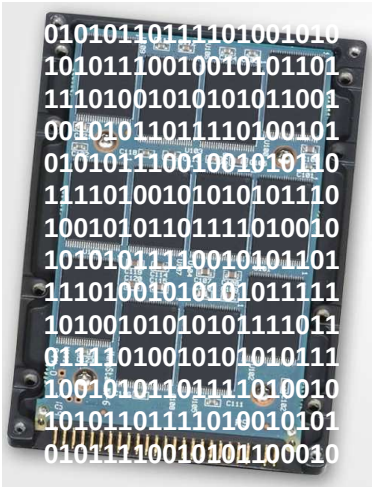

Slow

How Flash Storage Cells work

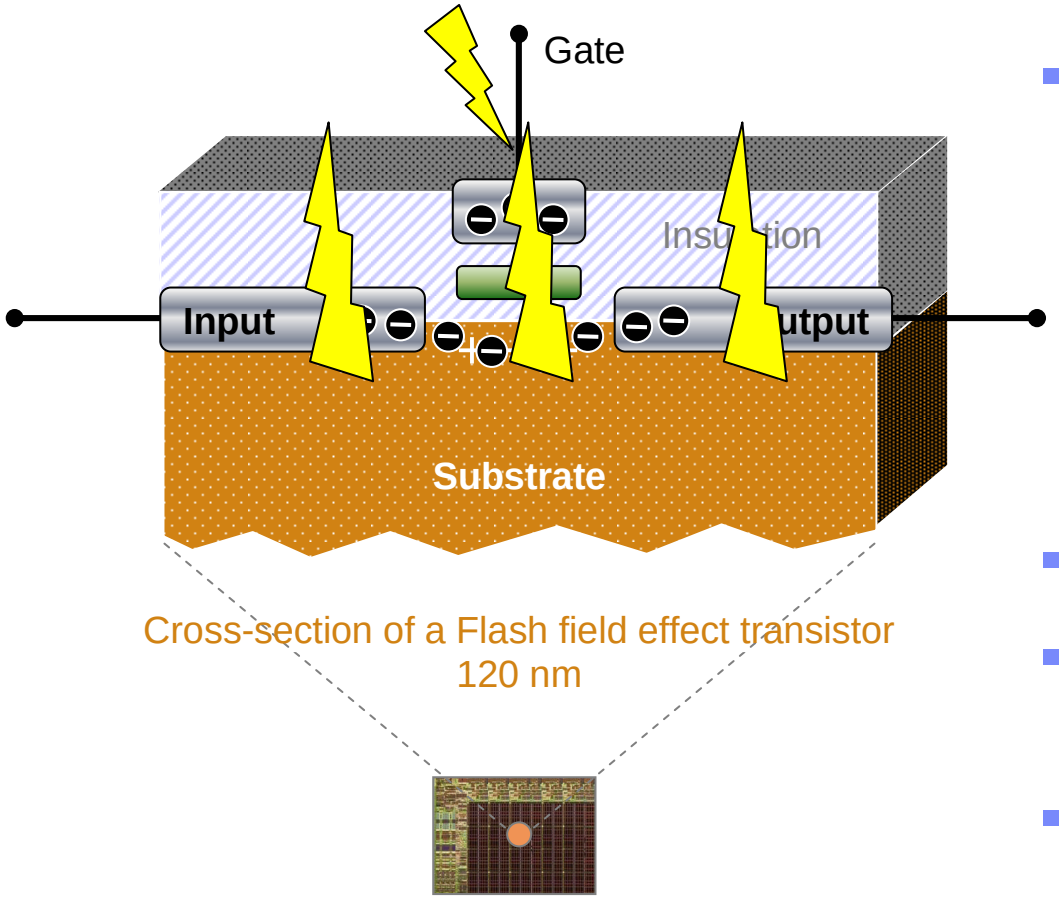

Cross-section of a Flash field effect transistor
120 nm

- **Floating** (insulated) **Gate**

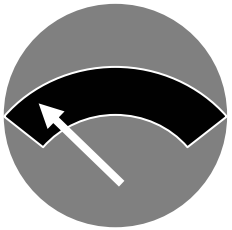

- Only able to write ONEs
- Cannot delete ONEs individually, only block-wise
- Deletion wear after 100.000×

Overwriting & Deleting Flash Data

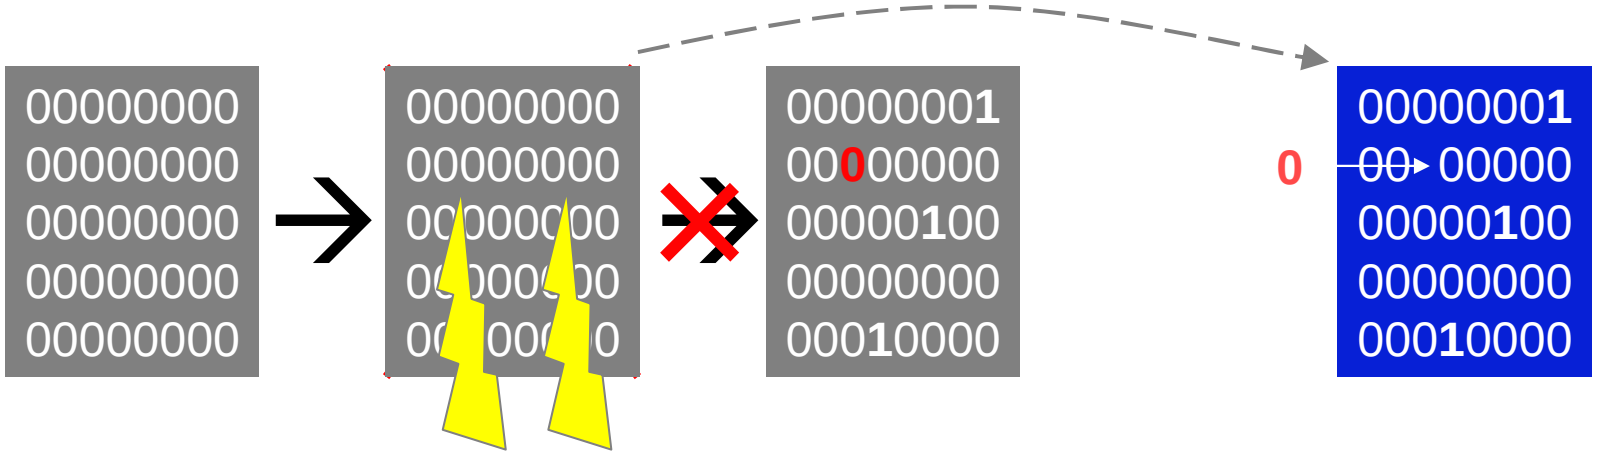

- Random Write is *not* the optimal workload for Flash
- **Delete block** sizes are much larger than typical IO size
- Constant relocation of often-used blocks helps leveling wear

How Solid State Drives (SSDs) handle block deletion

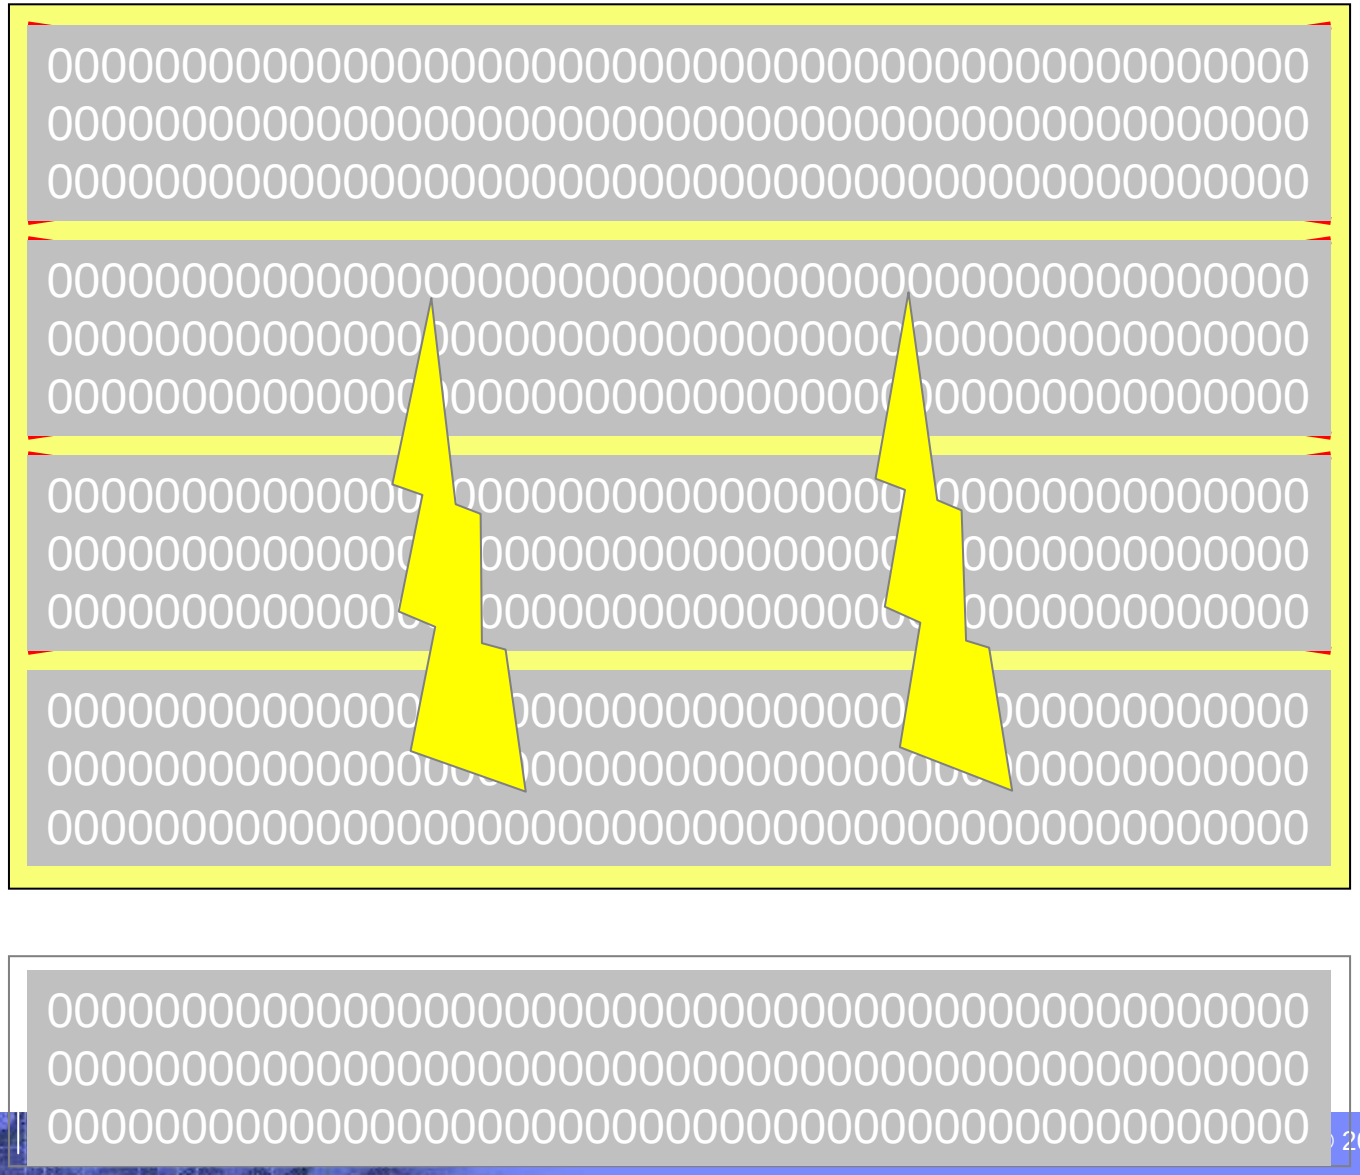

Cost-effective Approach to Random IO Performance

1. Use **Flash PCI-Memory** instead of **SSDs**
2. **Serialize** all random IOs (= less deletions)
3. Don't overwrite in-place
4. Optimize IO patterns at the system level, upfront of the drives

ADTRON® SSD 2,5" 160GB
$80-$115 *per gigabyte*

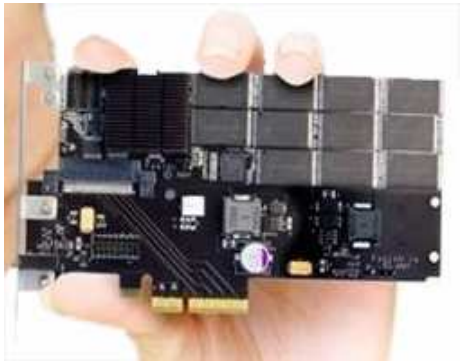

PCI-based Flash Memory
$30 *per gigabyte*

Random IO Serialization

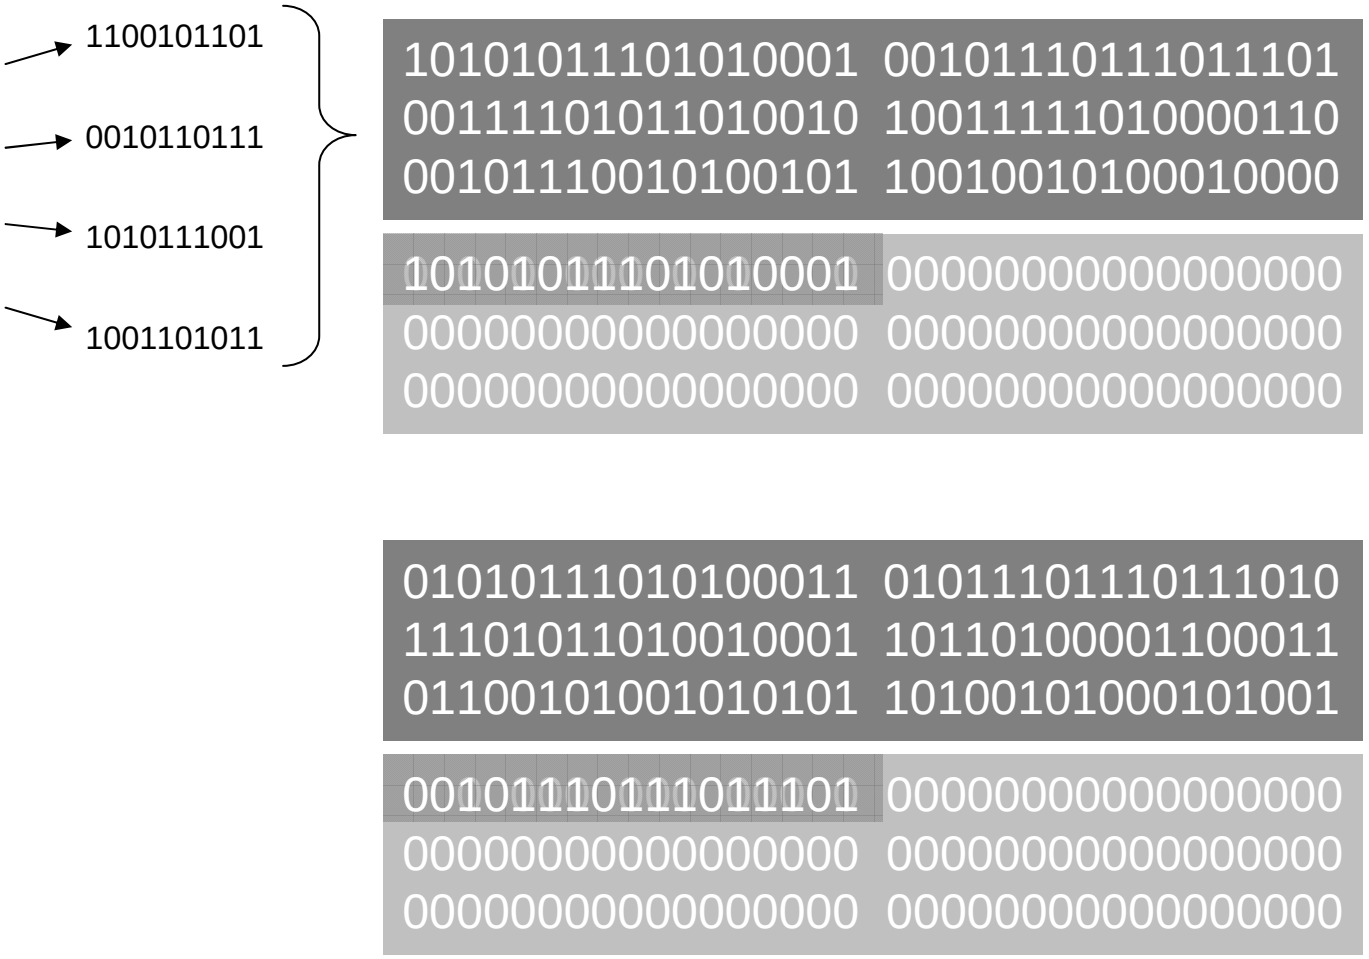

IBM Quicksilver : 1 Mio IOPS World Record

- Technology Demonstrator: IBM SAN Volume Controller + Flash
- Database workload 70/30, 0% cache, running for 2 hours, delivered **1 Mio IOPS** at **700µs response time** (peak 1.1M)
- Flash-optimizing SVC controller →

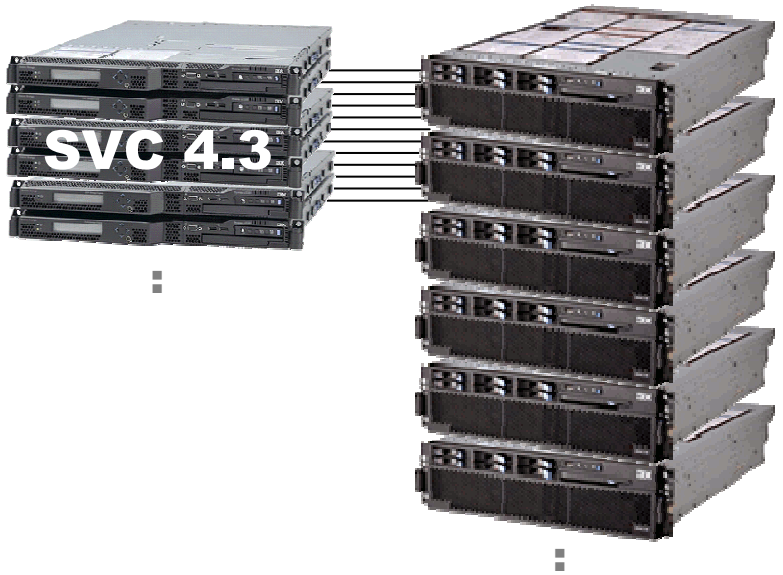

PCI Flash memory in modified SVC cages

August Press Release

IBM Quicksilver Technology

- SVC 4.3 adds fine-grained "thin provisioning" to *any* storage
- The smallest SVC grain size is 32kB
- Maximizes the use of your storage

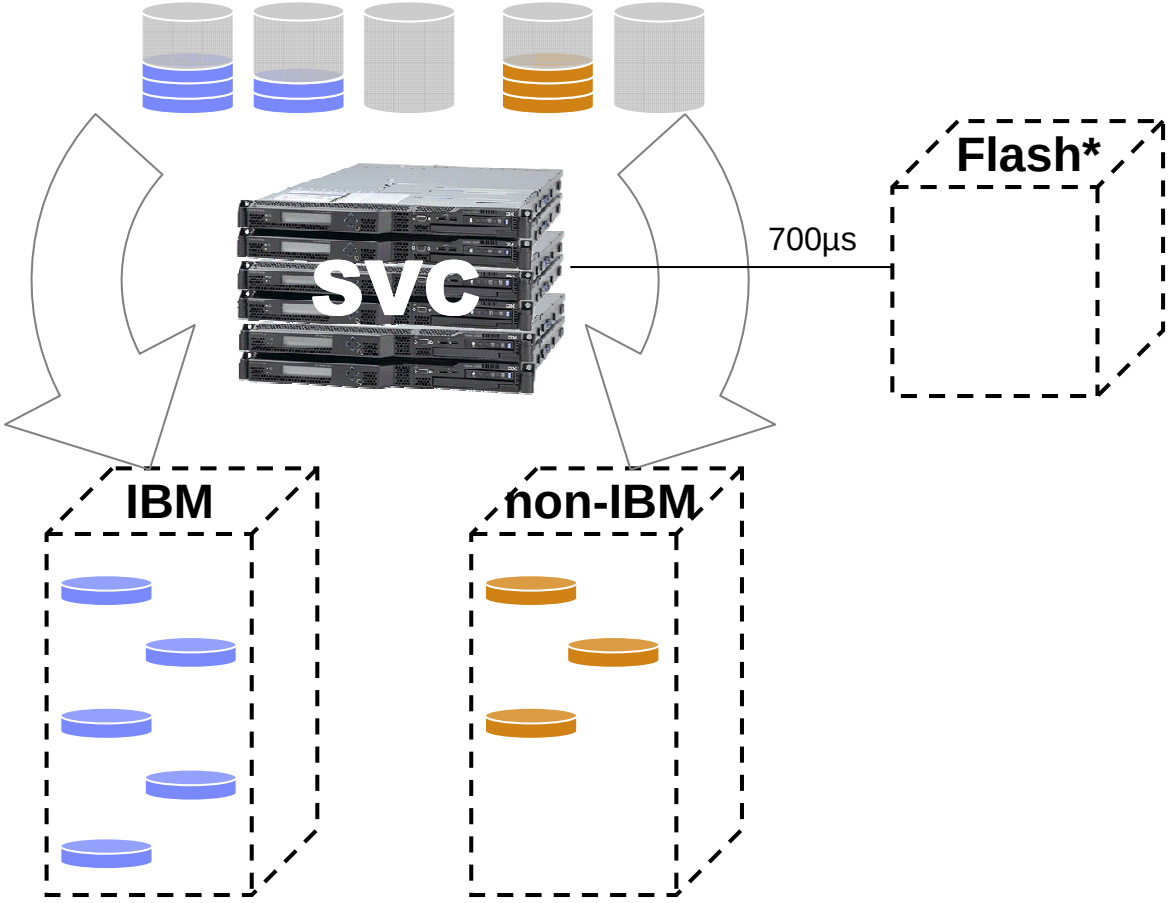

"Solid State" Alternatives

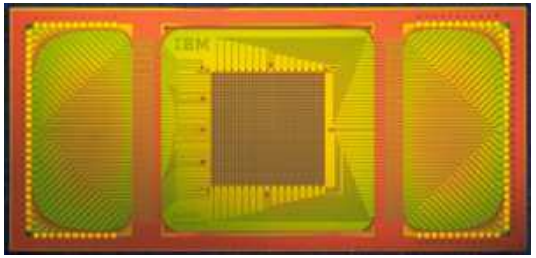

Spintronics & "RaceTrack" Memory

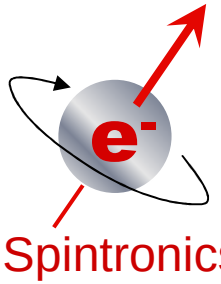

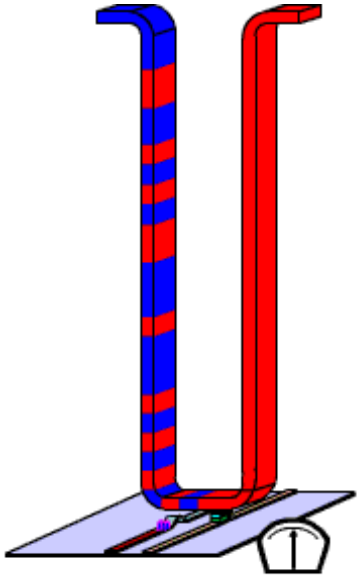

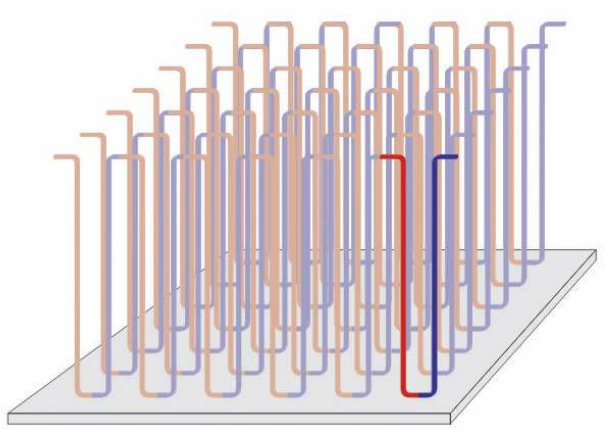

RaceTrack Storage Array

Storage in 3rd Dimension

"Large" read/write head, "small" bits on ferroelectric nano wire

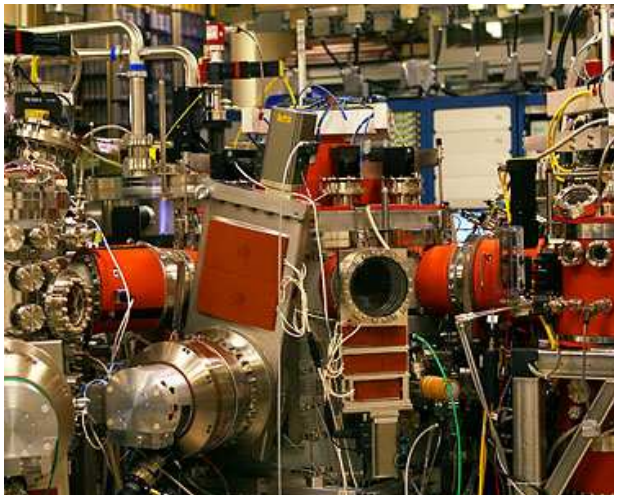

IBM Fellow Stuart Parkin, Inventor of GMR read heads, investigates "Racetrack Memory"

Magnetic Random Access Memory

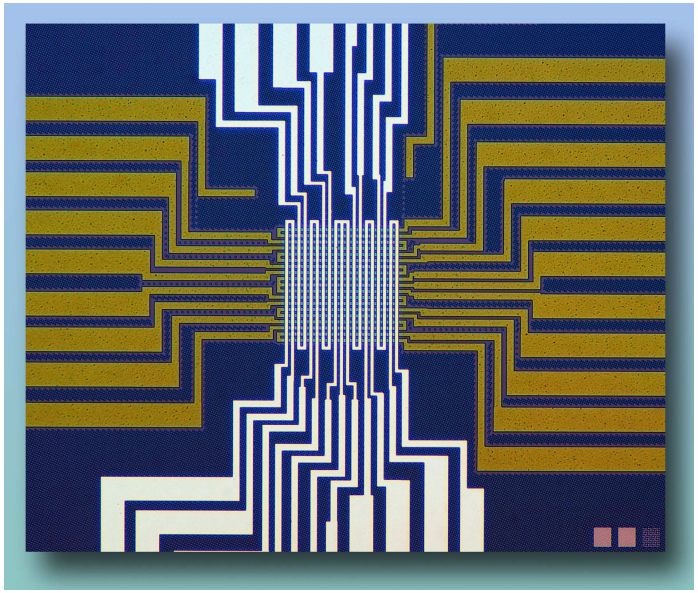

IBM Prototyp 199..

Manufactured @Freescale Inc.
"MR2A16A" 4Mb non-volatile
35 nsec Access Time

New IBM Demonstrator:
2 nsec Access Time

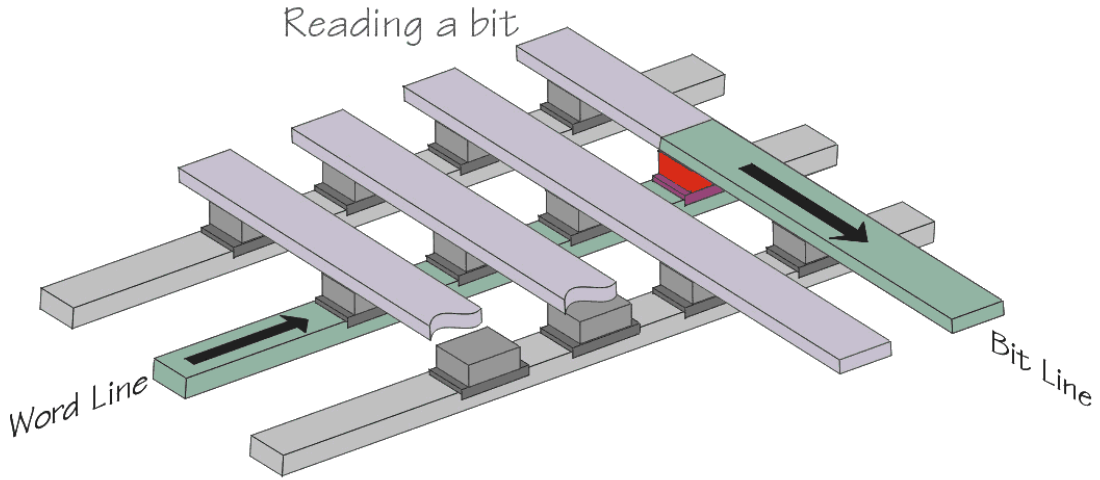

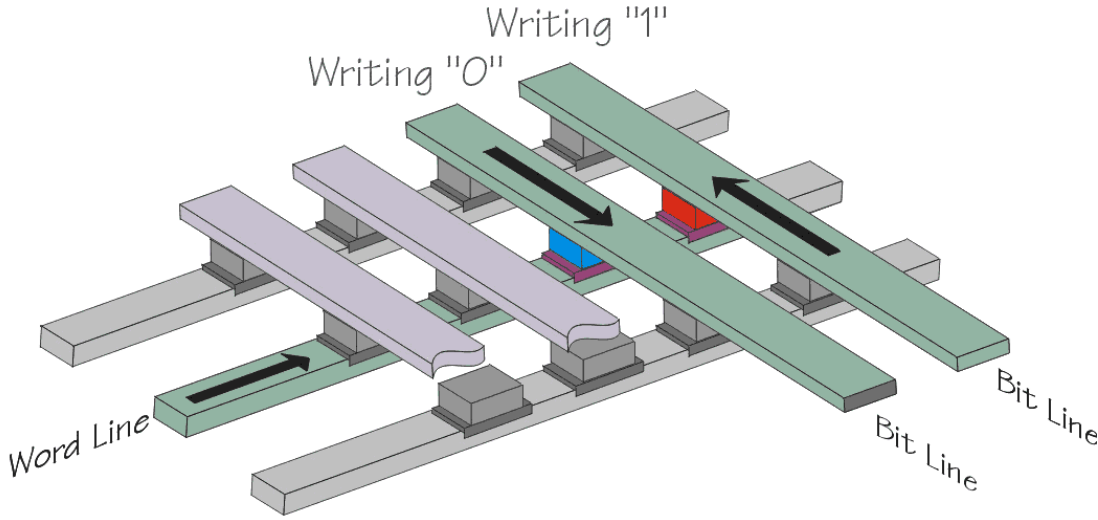

Non-volatile RAM = IT Revolution !

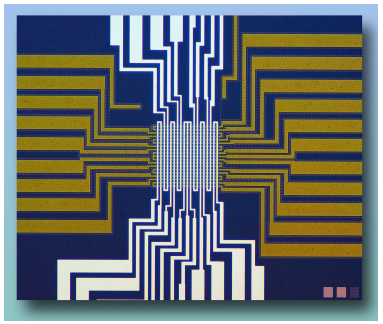
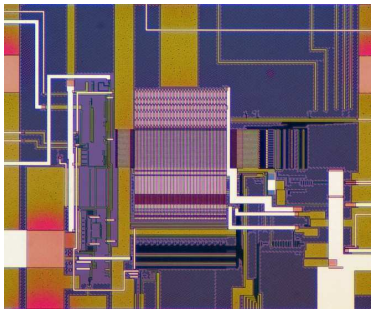
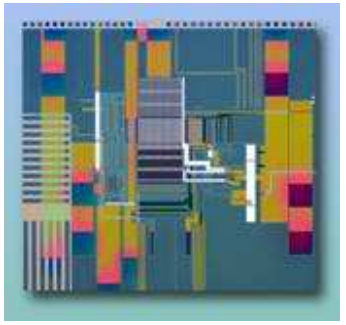

- Magnetic RAM
- Phase-Change RAM
- Ferroelectric RAM
- …

How to use non-volatile RAM

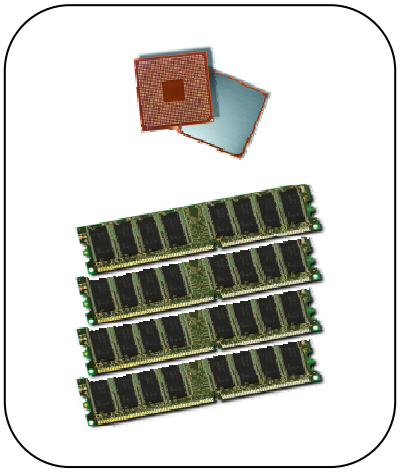

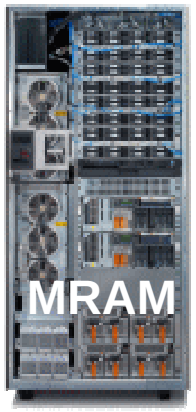

20 nsec **2 nsec**

- Online storage moves nearby the processor
- Fast (Fibrechannel-) disks disappear
- Monolithic design = higher reliability

The SAN of the Future

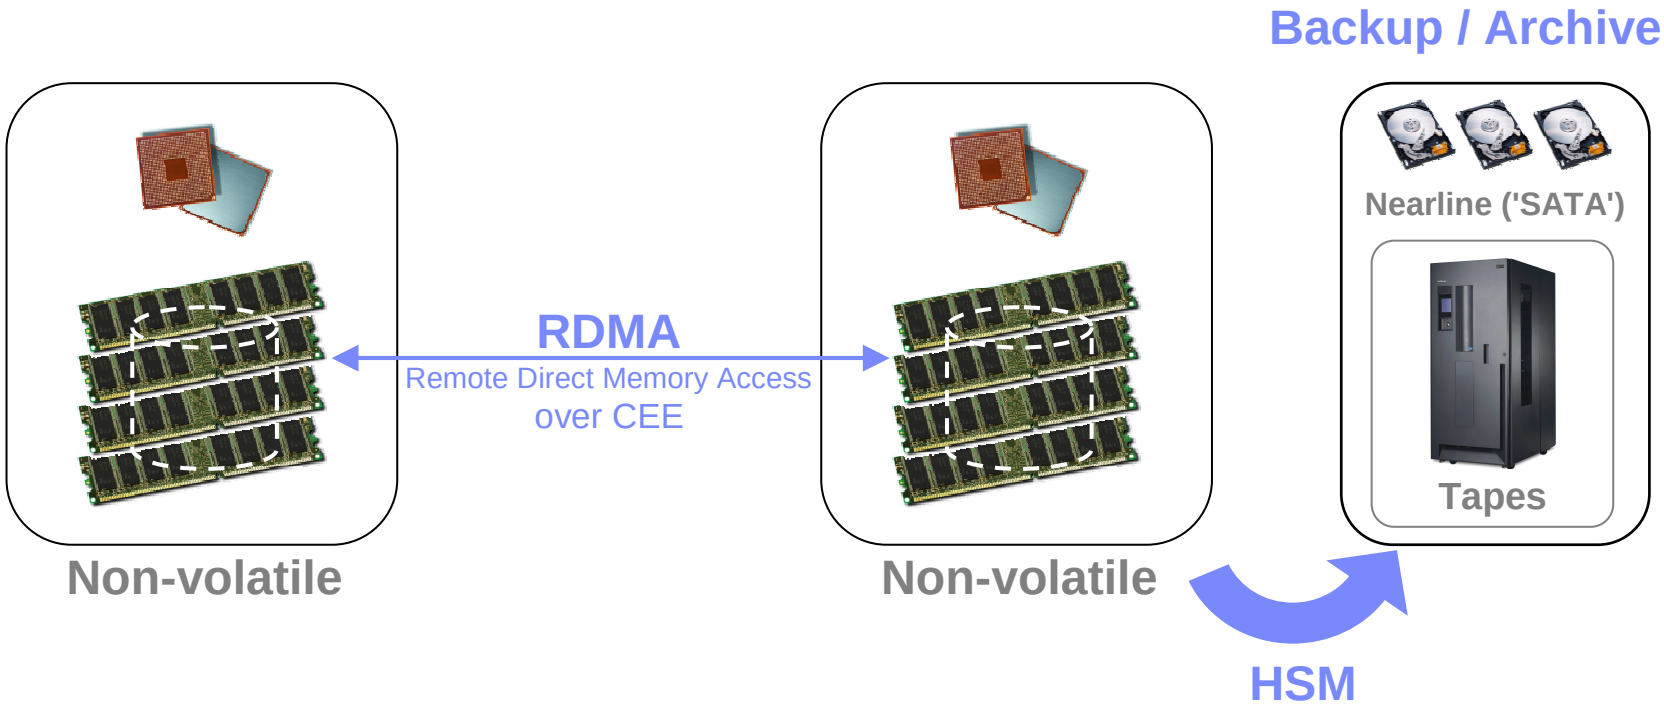

- Memory-to-memory SAN *(RDMA over Converged Enhanced Ethernet)*
- Disks → RAM
- Tapes → Disks
- Paging → HSM

Today: Consistent Caching without non-volatile RAM

Cache as a Hard Disk Memory Extension in a NAS Grid

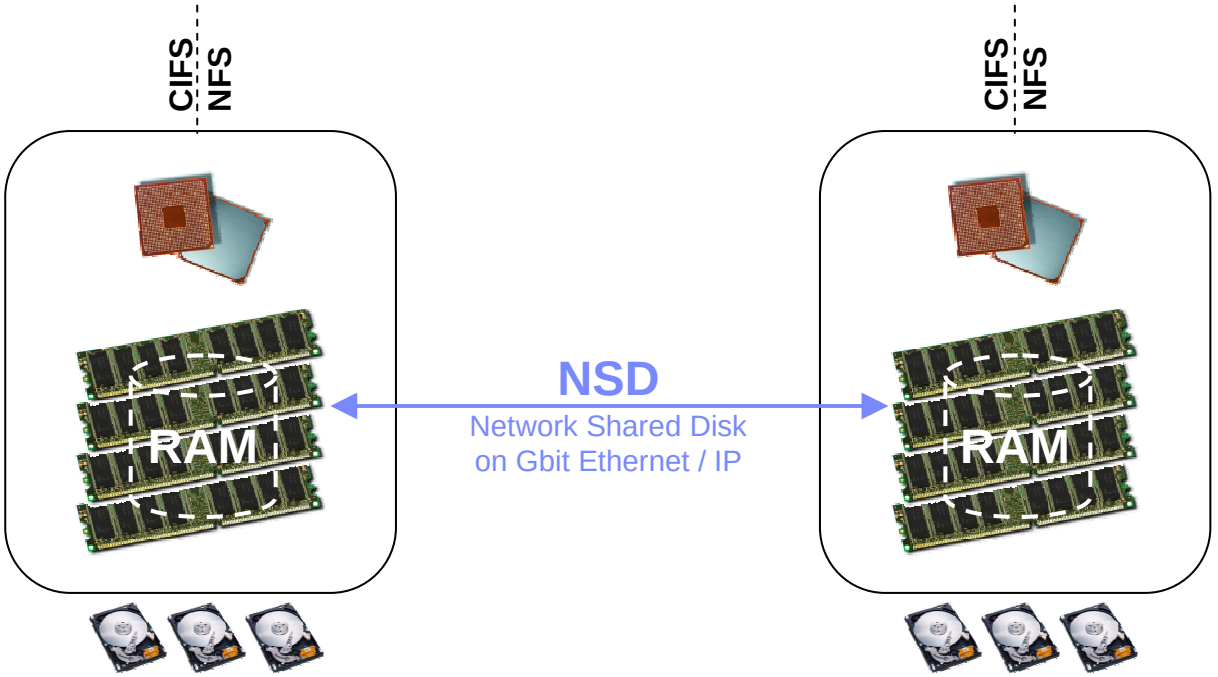

Scale-out File Services

As close as possible to "full parallel" in RAM

- Native Samba on a Cluster File System = NO NO !
- For **parallel NAS for Windows**, use CTDB !

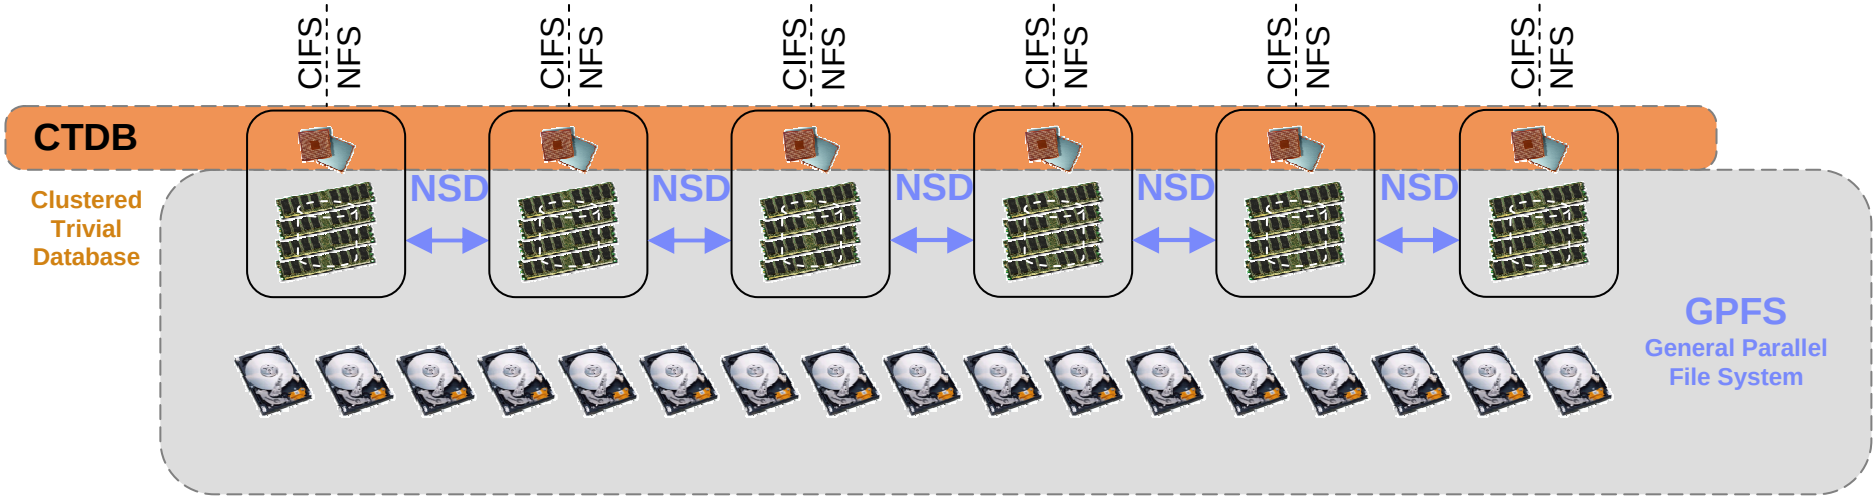

- Ultra-fast System Backup: TSM Scan **@1 Mio Files / sec**
- Ultra-fast "virtual" Full Restore

Most Recent Labs News

Image © DOCB

Slow Light

- Nano-structured silicon with **refraction index 300**
- Light is 300 times slower (can be influenced)

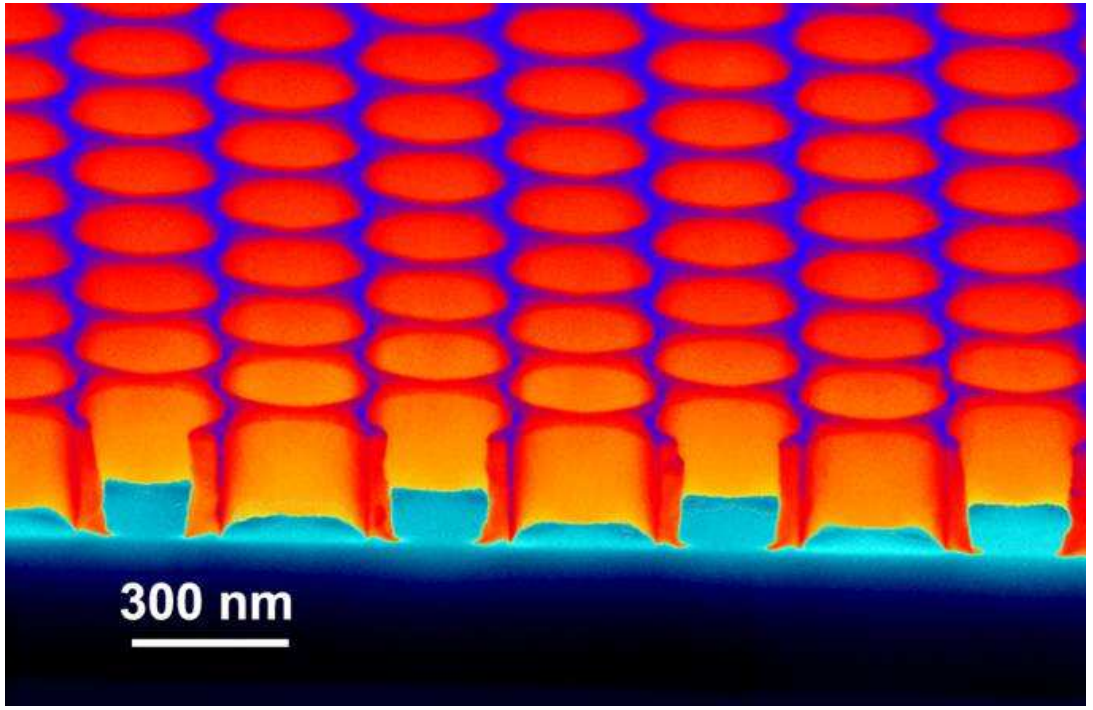

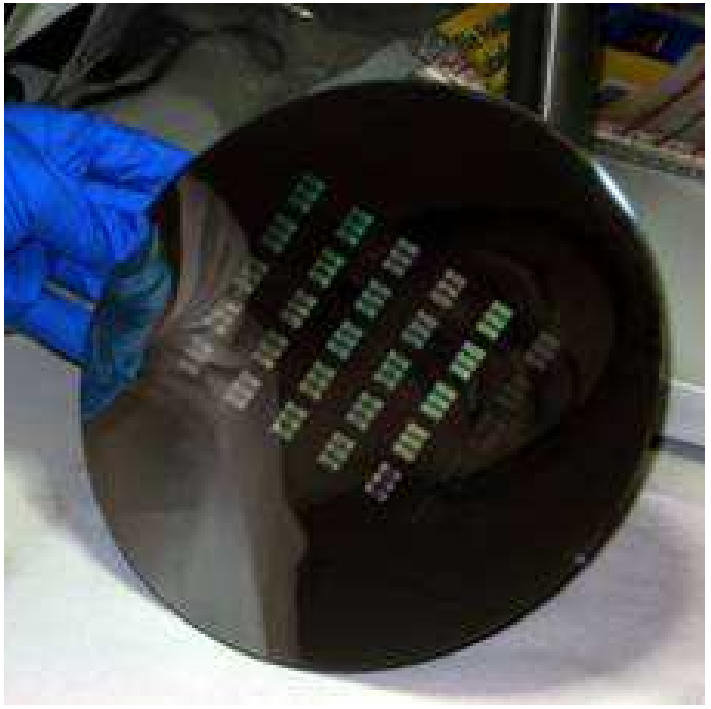

Predictive Summary

Future Computers will contain…

- "Slow light" storage *(maybe)*
- 3D nano structures
- Spintronics *(for sure)*

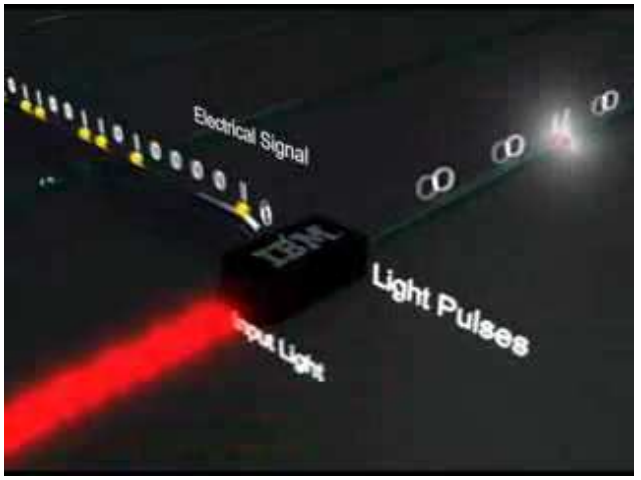

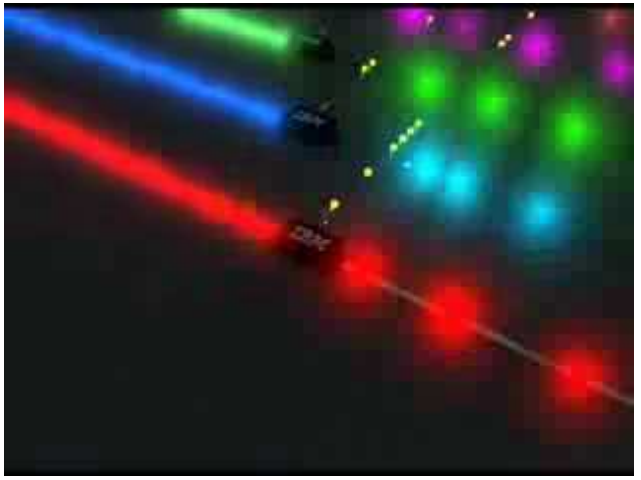

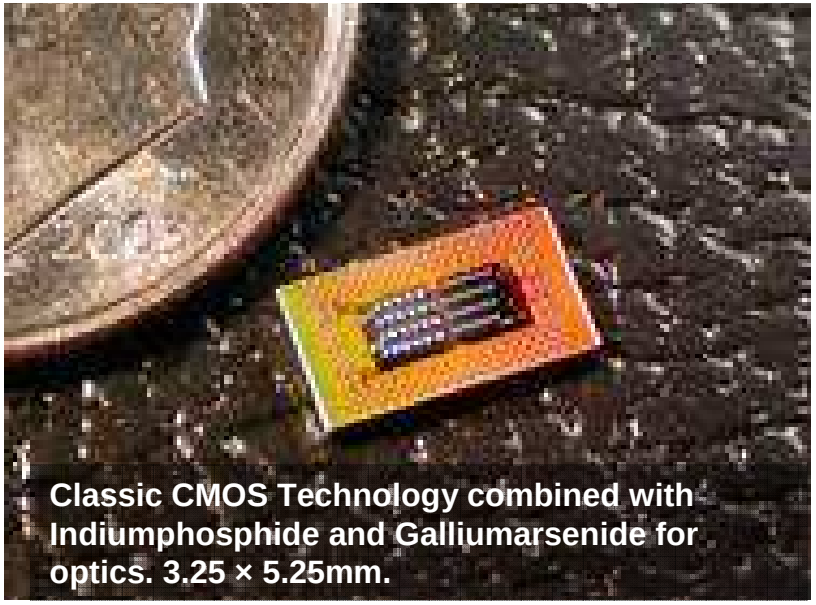

Slow Light : Beware of the Consequences…

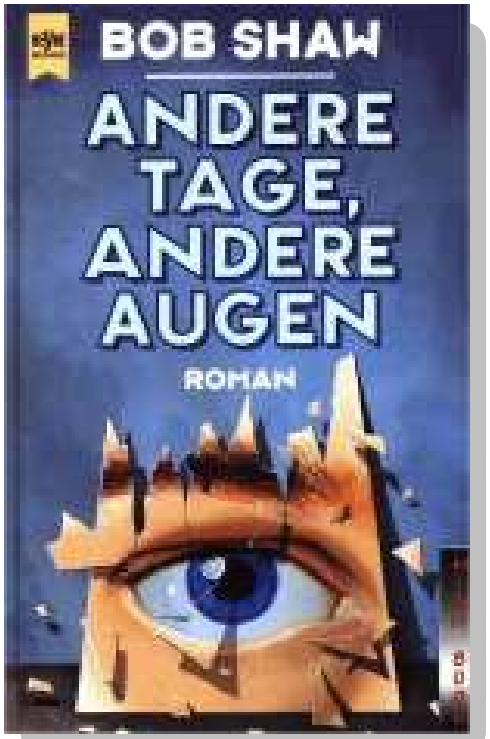

"Slow Glas"

axel.koester@de.ibm.com

Disclaimer