Data Recovery Presentation

Øyvind Nyland
Operation Manager

ibas Kroll Ontrack

Varity of jobs Job’s

Columbia Space Shuttle Disaster

- All components had fallen off
- Every piece of plastic on the hard drive melted
- All electronic chips inside had burned and loosened
- Drive was cleaned
- Platters transferred into replacement drive
- Specialists at Ibas Kroll Ontrack recovered 99% of the drive's data

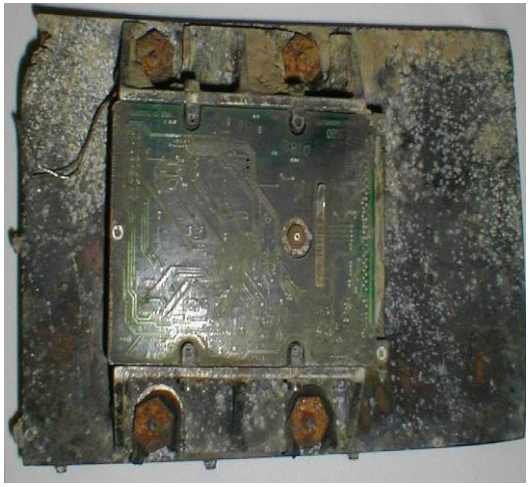

Røyksopp

Ibas recovered 99% of all lyrics and music on the new record from the world famous electronica band Røyksopp.

ROCKNESSALVAGE.COM

On the 19 th of January 2004 the Antigua & Barbuda flagged cargo vessel MV ‘Rocknes’ capsized in a strait south of Bergen, Norway. At the time, the ship was loaded with stones and pebbles that were to be delivered in Emden, Germany. In this tragic accident, 18 people lost their lives.

Ibas was engaged to recover crucial information from the boat’s storing and navigation systems

Private Customers

- Typical
  - External USB
  - NO Backup
  - Thousands of Pictures
  - Damages
    - Deleted
    - Hard drive chrassed
    - Hard drive fallen to floor

Why Data Recovery ?

Financial Impact of Data Loss

Cost of Downtime

Enterprise – Average Cost of Downtime

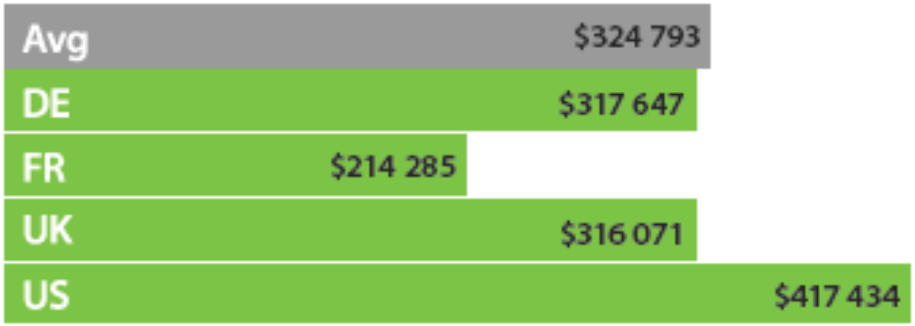

Cost-per-hour of critical servers being down (USD)

Percentage of server backups tested when testing for recoverability (%)

Costs of Downtime and Data Loss

- Impact - Why should you care?
  - IT downtime costs North American companies **$26.5 billion** per year

- Companies with an outage lasting for more than 10 days
  - Will never fully recover financially
  - 50% out of business within 5 years

Data Loss Situations

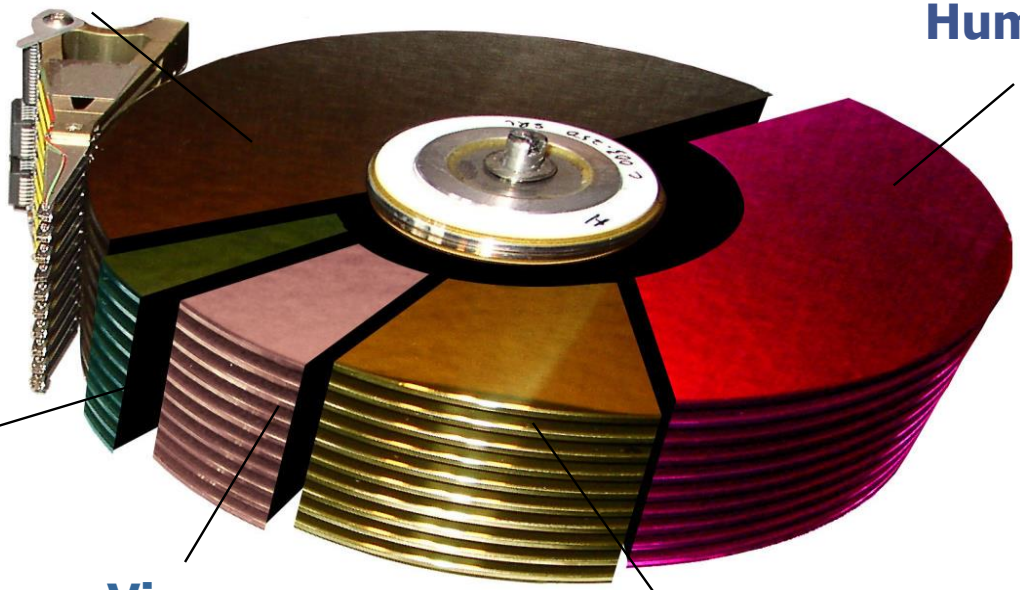

Four Phases of Data Recovery

Repair the drive

Image the drive

Logical recovery of files, partition, MBR and MFT

Repair damaged files that were retrieved

First Phase - Physical Analysis

Physical Analysis

- FIFO queue within each service level
- Register analysis started
- Check all information
  - Consistency/errors
  - Hints on problem
- Physical analysis procedures
- Main goal: Secure as much raw data as possible
- Transfer to Logical Analysis

Anatomy of a Hard Drive

(1)

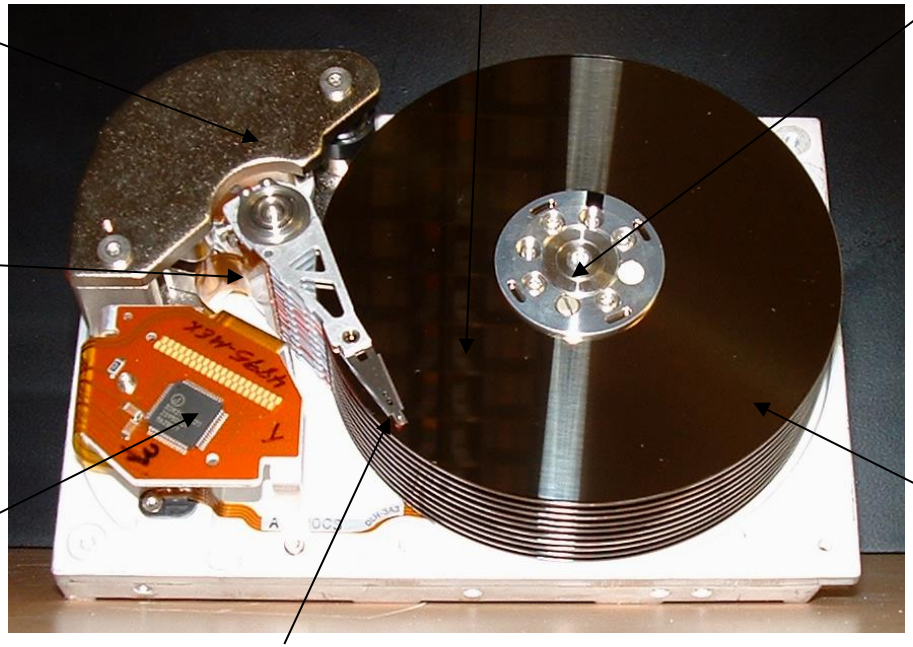

Physical tolerances

- 3 TB Hard Disk
- Arial density 444 Gb/Sqin (1)
- Track density 270.000 TPI
- Linear density 1.383.000 BPI
- Flight distance 5 to 10 nm (2)
- Self-calibration procedure (3) (4)
- Environmental characteristics (5)

Physical tolerances

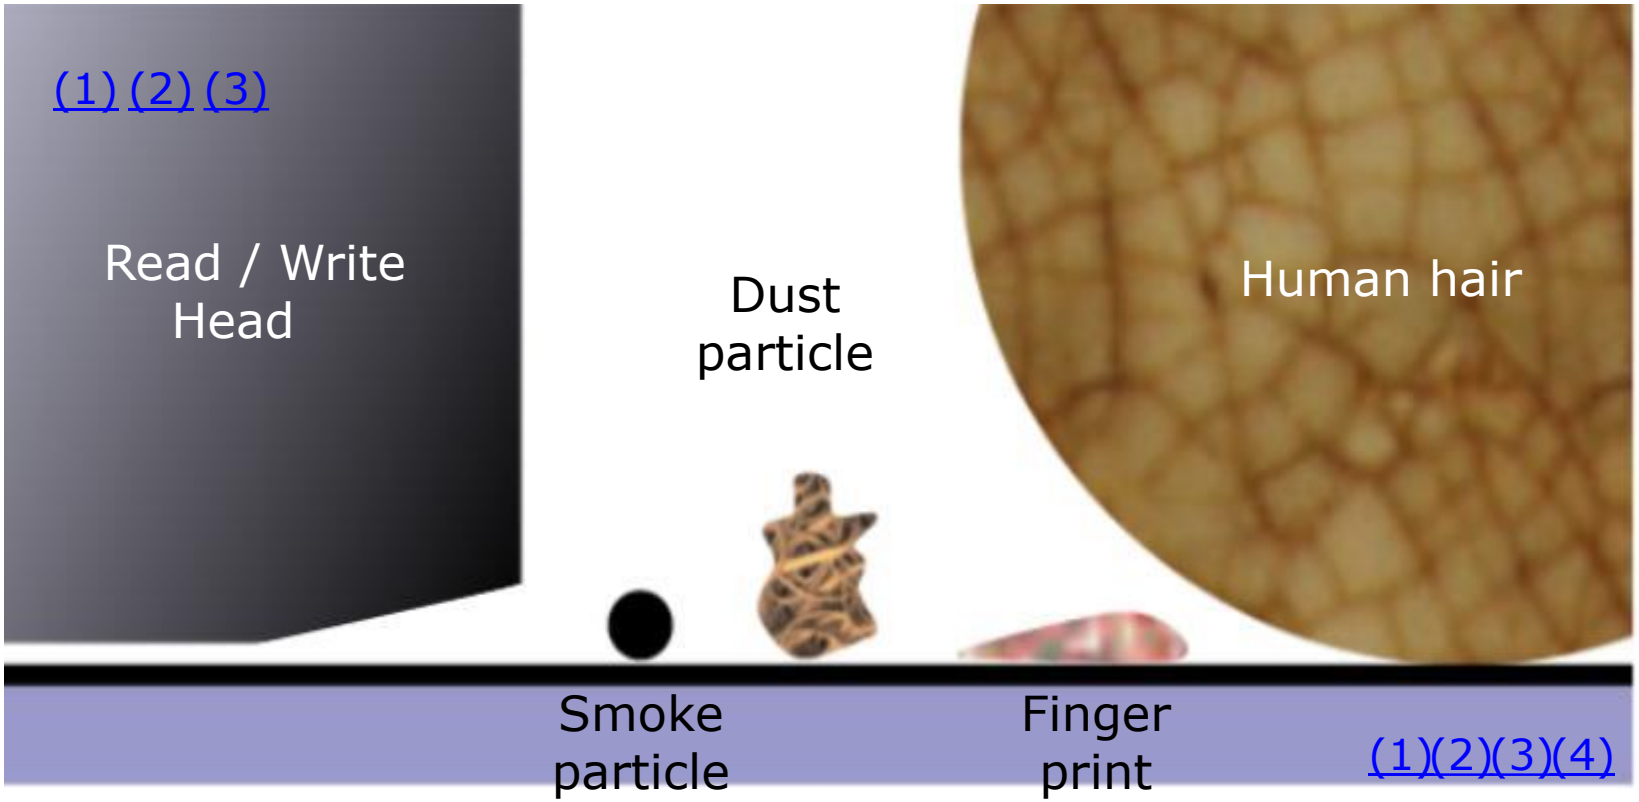

Physical damage, statistics

- Head Crash 45% (1) (2) (3) (4)
- Shock 8% (1) (2)
- Mechanics 17%
- Electronics 16% (1)
- Water/Fire 3% (5) (6) (7) (8)

80 to 85% of all jobs we receive have some kind of physical damage

Second Phase – Secure Raw Data

Second Phase – Secure Raw Data

- Based on the analysis
  - Recover as much raw data as possible
  - Transfer raw data to Ibas datacenter

ibas Kroll Ontrack

Third Phase – Logical Recovery

Logical Analysis

- FIFO queue within each service level
- Register analysis started
- Check all information
  - History
  - Hints on problem
- Logical analysis procedures
- Main goal: Check of file system, and determine possibilities for a data recovery

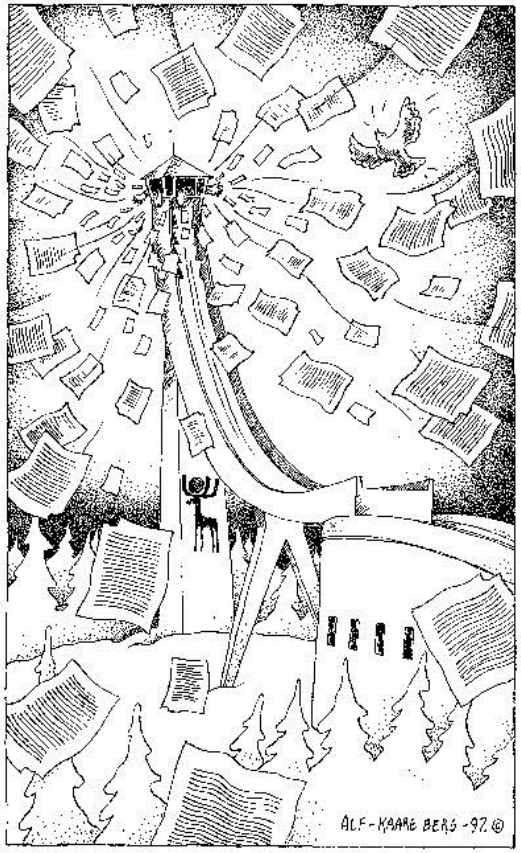

Logical damage, statistics

- Virus 4%
- Overwritten 8%
- Deleted 7%
- Corruption 23%
- Missing Raw data 44%

RAID 5 - NTFS Volume

	Stripe 1	Stripe 2	Stripe 3	Stripe 4	Stripe 5	Stripe 6	Stripe 7	Stripe 8	Stripe 9	Stripe 10
HDD 1	M1	D2	D4	P4	D10			P8		
HDD 2	M2	D3	P3	D7	D11		P7			
HDD 3	D1	P2	D5	D8		P6				P10
HDD 4	P1	M3	D6	D9	P5				P9	

VMware Recovery Overview – RAID Failure

- Common RAID failures:
  - Multiple drive failure
  - Drives forced online
  - Incorrect RAID rebuild
  - RAID re-configured

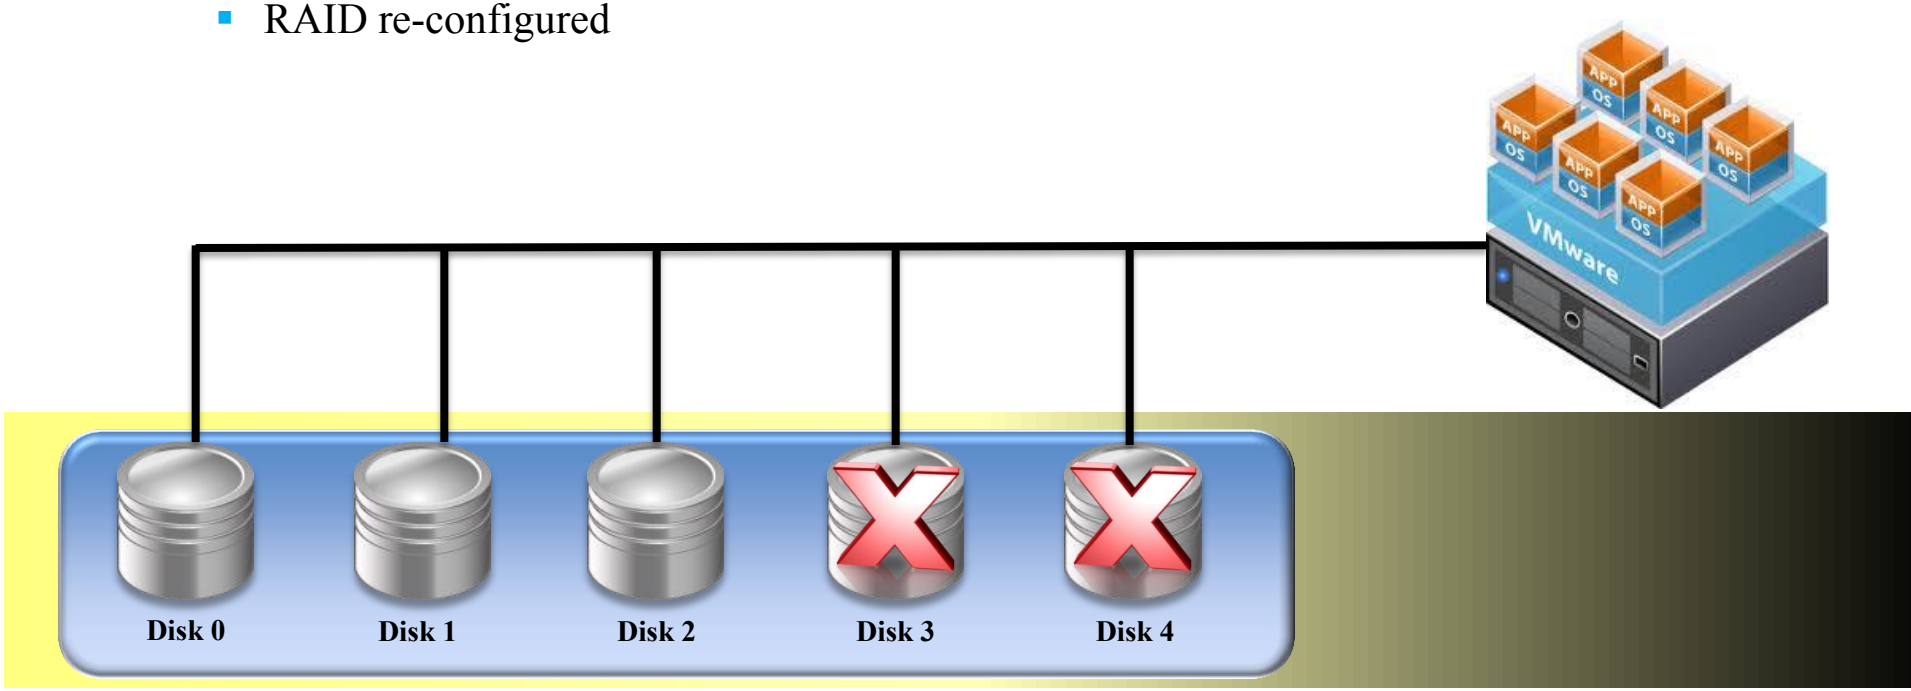

VMware Recovery Overview – RAID Failure

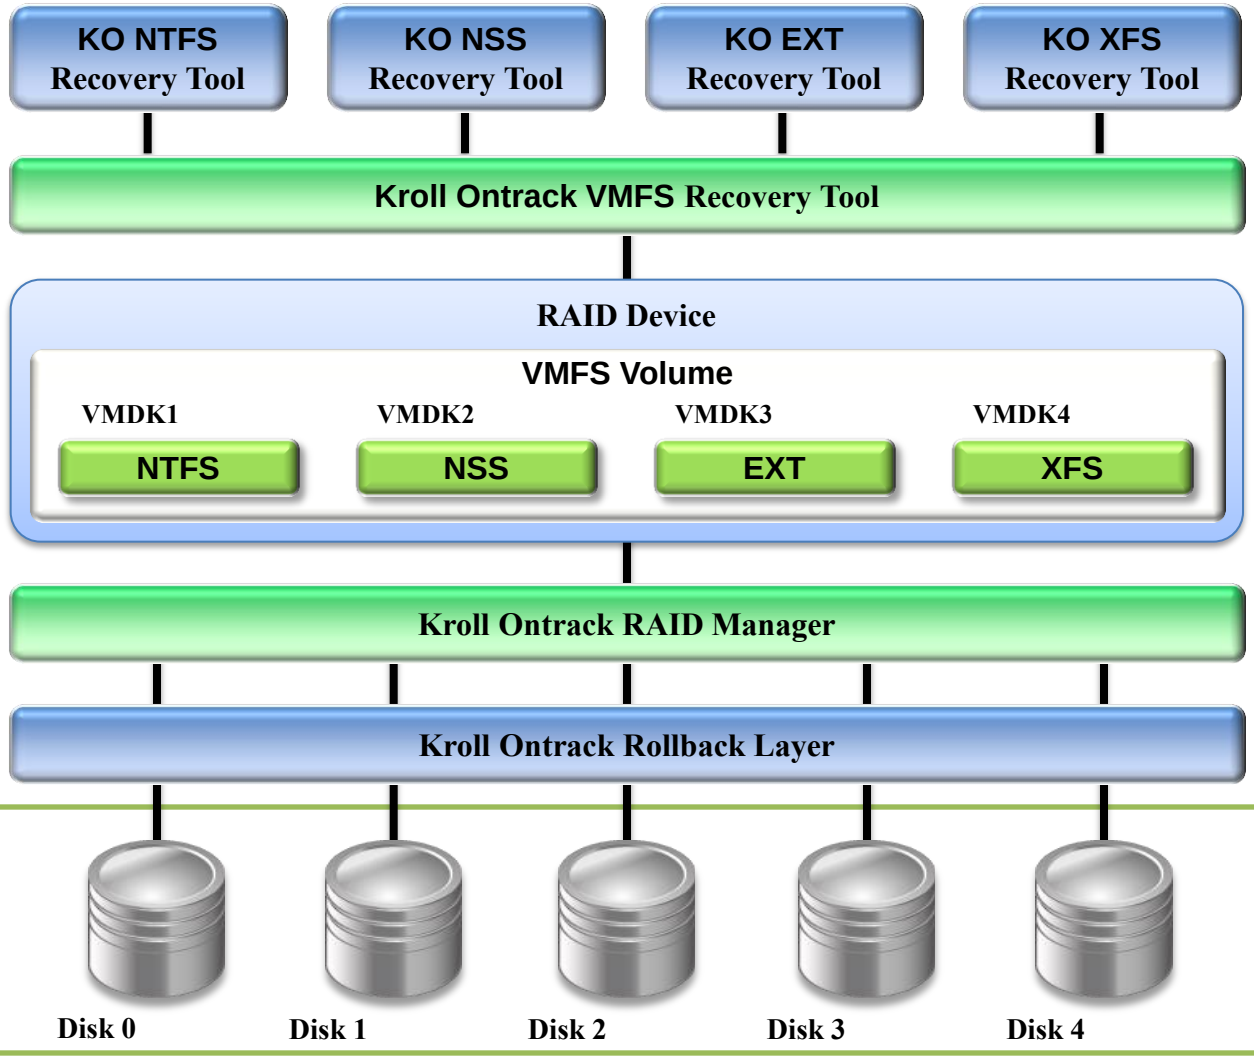

- Specialized file system recovery tools. All repair applications are developed by Kroll Ontrack
- Works like a VMware snapshot to catch any changes or repairs needed
- Locally attached drives, separated from RAID controller

ibas Kroll Ontrack

Traditional Disk Array Approach

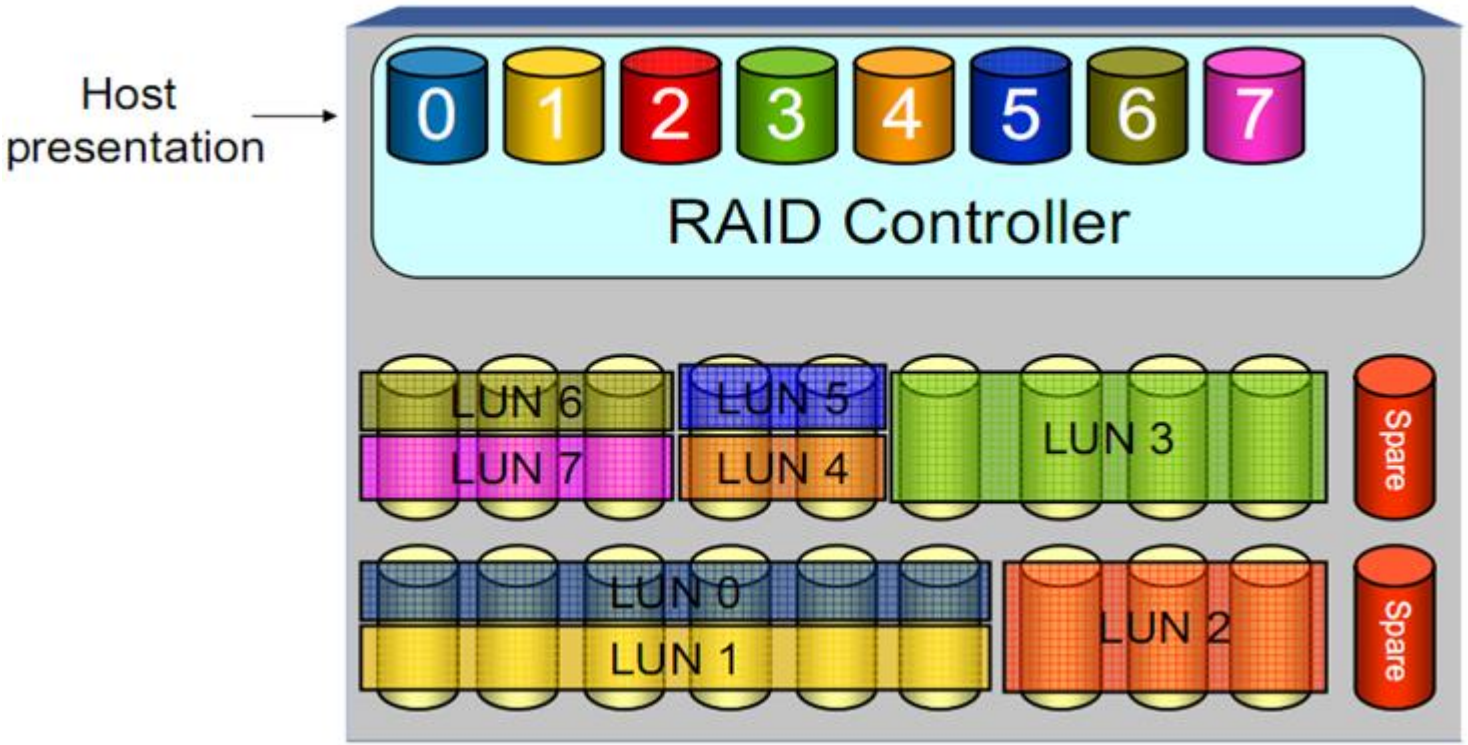

Recovery of Software Defined Storage

- Many storage platforms now use custom data mapping instead of normal Raids
  - NetApp – WAFL
  - HP EVA, 3Par, and Lefthand
  - Dell EqualLogic and Compellent
  - EMC VMAX, VNX and Isilon

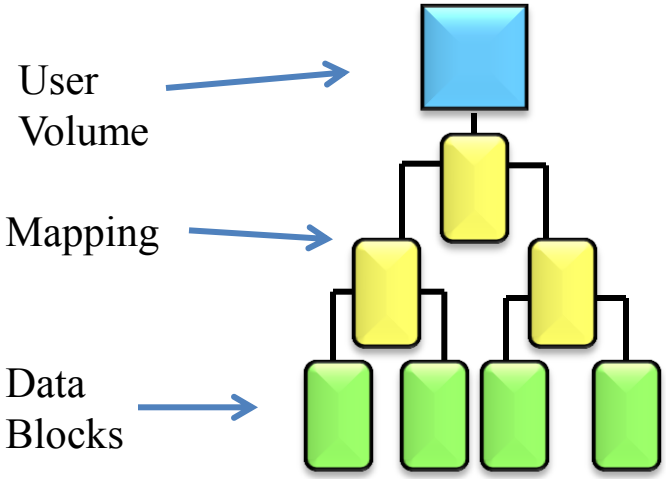

- In some recovery scenarios, we must use special tools to assemble the data
- We have over 100 developers who create custom recovery tools for new systems like these

Virtual Array Approach

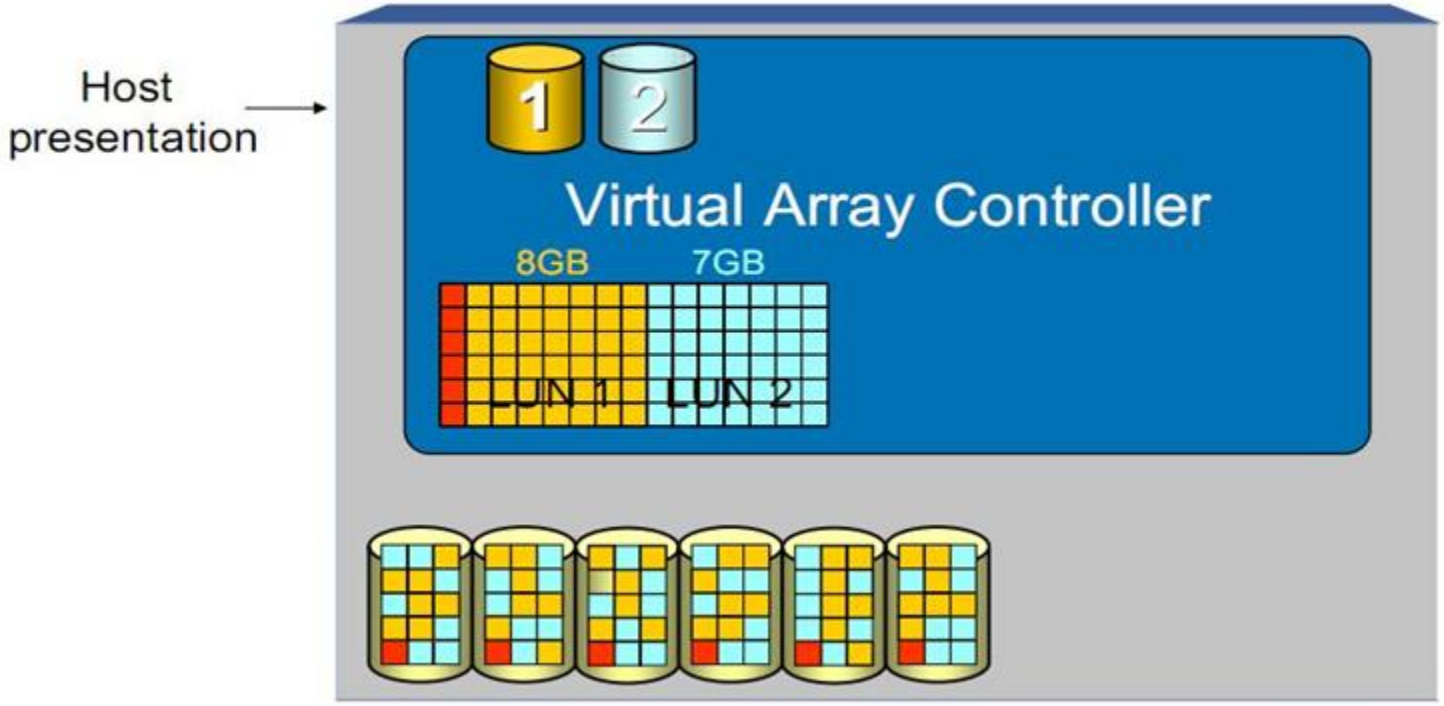

Data Loss Scenarios
Deleted VM Recovery

Critical VMFS Metadata Files

- All VMFS metadata exists in first 1GB of Datastore
- FDC - File Descriptor Cluster
  - Contains Inodes for all files
- PBC - Pointer Block Cluster
  - Contains pointers to all large file
- SBC – Sub Block Cluster
  - Contains data for all small files

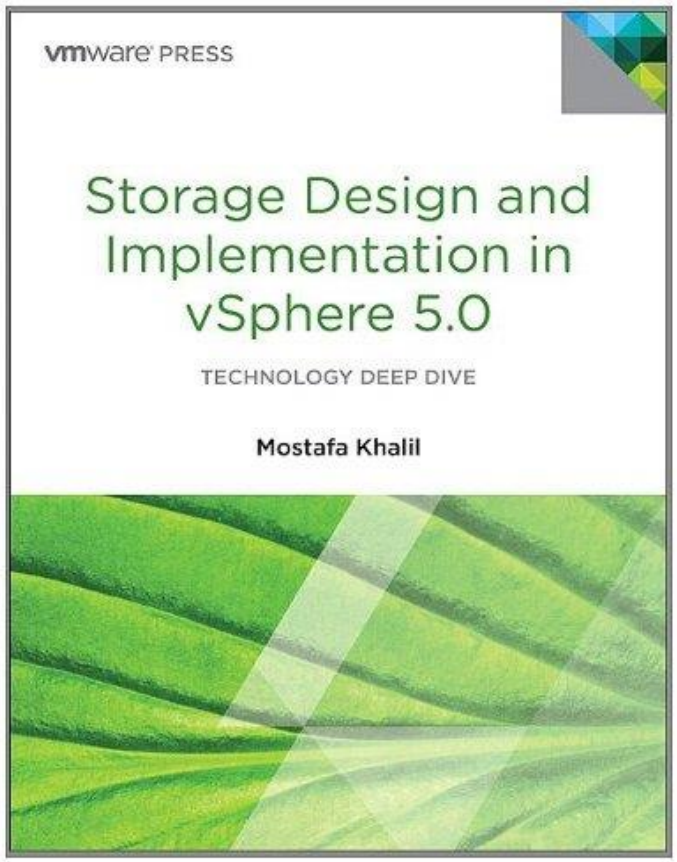

Deleted Virtual Machines

- When a file on a VMFS volume is deleted, the inode is zeroed, but the PBC pointers and VMDK data still exist until overwritten.

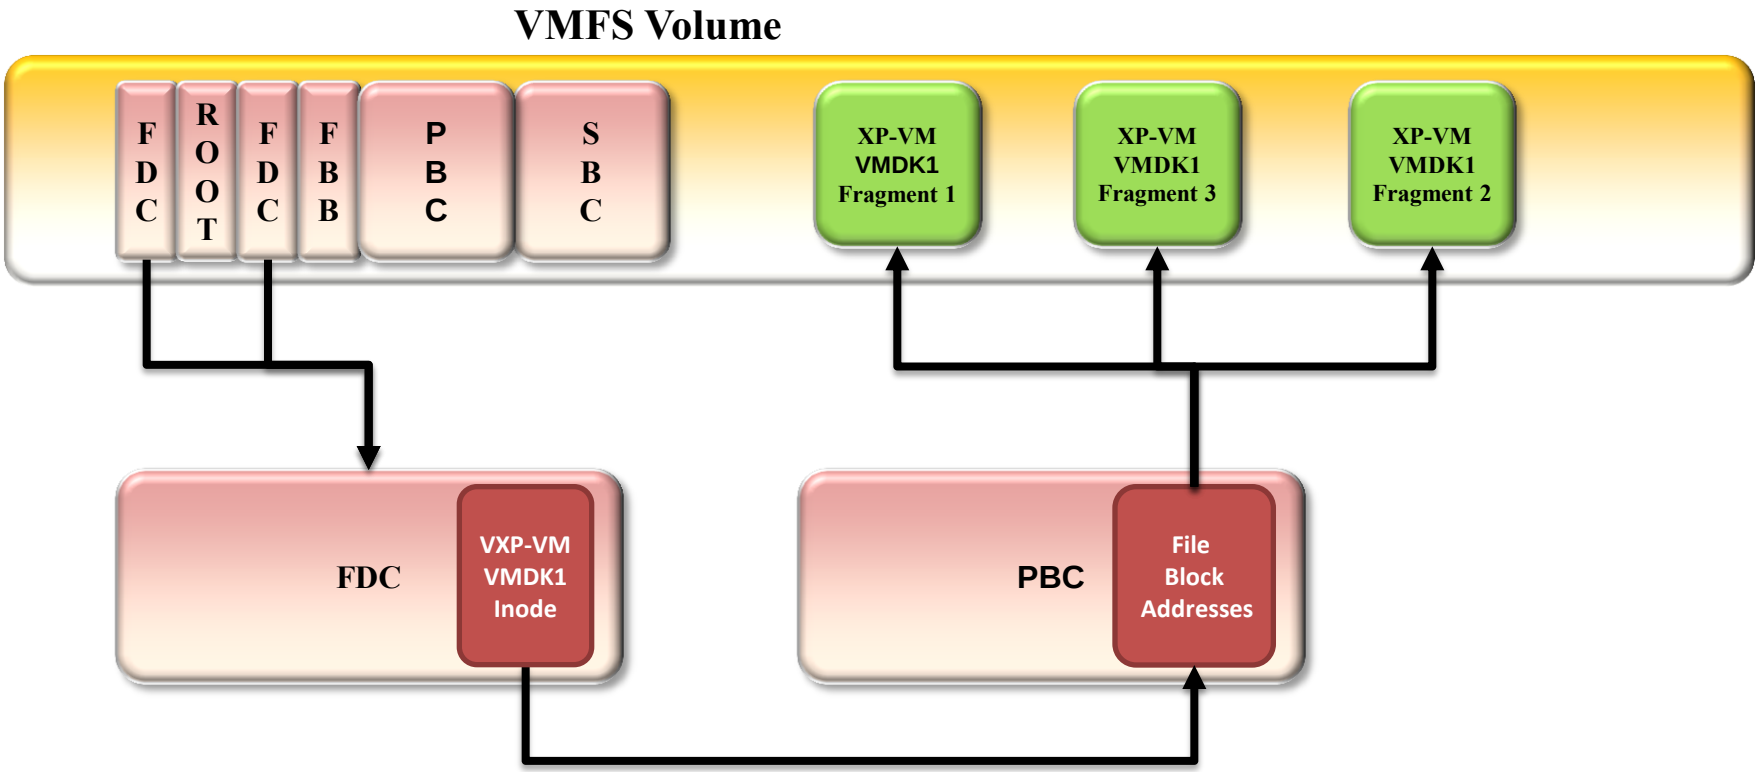

Deleted VMDK file recovery

Kroll Ontrack engineers can use structures from the guest FS as well as st[illegible] use in re[illegible]

These blocks can be allocated in a contiguous area of the datastore, or fragmented in thousands of pieces scattered across the device.

[illegible] VMDK file data in areas no longer referenced by the file system.

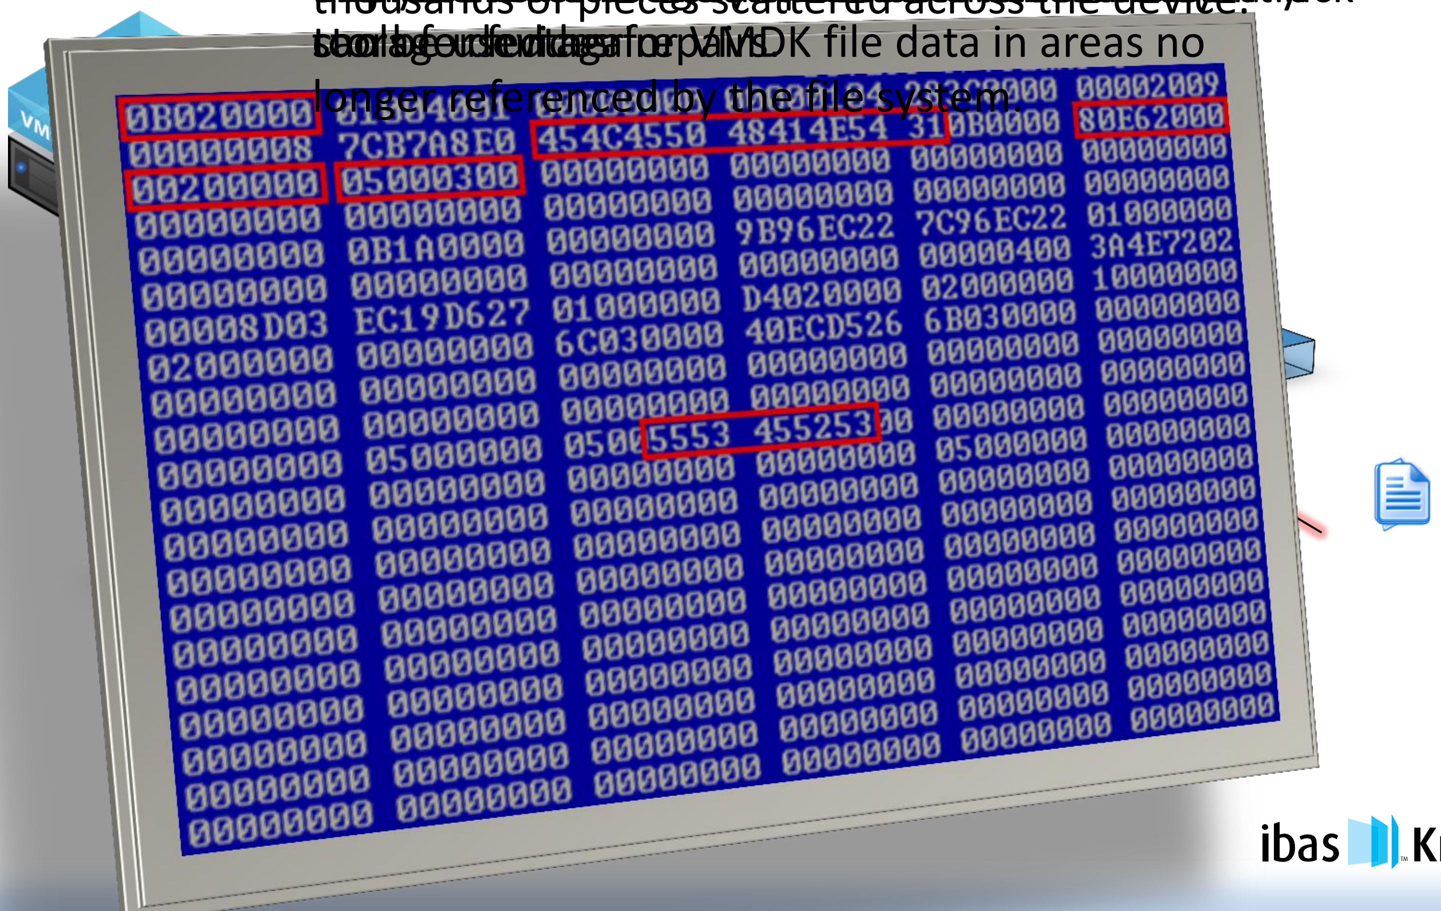

Remote Data Recovery Connection

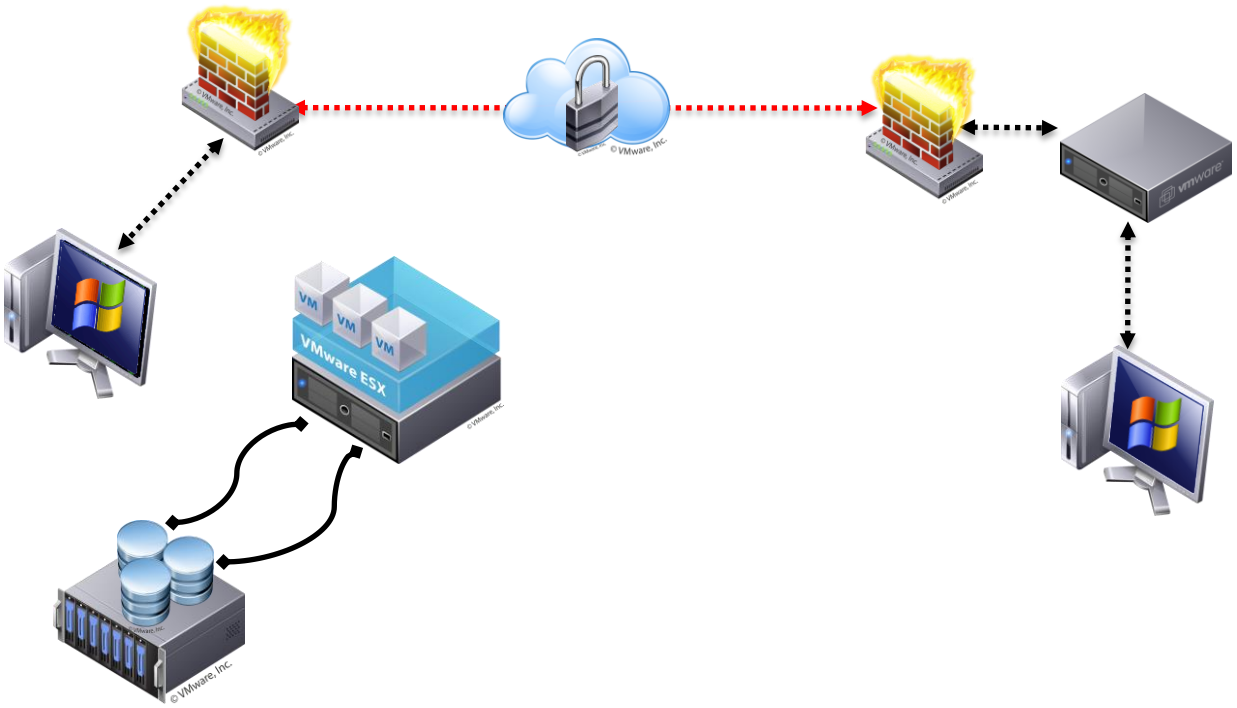

- 128-bit RSA encryption
- Proprietary network protocol wrapped in firewall-friendly HTTP packets
- Only screenshots and keystrokes are sent across the connection
- The tools are run on the client's side

Fourth Phase – File repair

- Files can be
  - Incomplete
  - Missing raw data
  - Corrupted
  - Un mountable
- Databases
  - Extract tables/rows to new DB
- Virtual Machines
  - Repair structures
  - Extract content

Final Phase – Data Return

Final Phase – Data Return

- FTP download
  - 500GB – 30MB/s line
  - Delivery time
    - 1 day 13 Hours 56 min
- 99% of all jobs
  - Delivered on a NTFS formatted USB 3 Hard Drive
- 30 Days Backup of raw data in Ibas Datacenter Included

ibas Kroll Ontrack

Data Recovery Tips

- How does thick and thin provisioning effect the data recovery possibilities
- ESX 5 provides three options to provision a VMDK file on VMFS.
- Flat Disk (Thick Provision Lazy Zeroed)
  - This is the default which creates a normal flat disk. It does not write anything to data blocks until the VM writes to them.
- Thick Provision (Eagar Zero)
  - This option pre-zeros all blocks in file. It is required to use the ESX Fault-Tolerance feature. It is not good for recovery jobs since it zeros all blocks when the VMDK is created. If an Eager Zeroed VMDK file is created after the deletion of a VMDK file, it can overwrite data that belonged to the deleted file.
- Thin Provision
  - This option creates a flat disk that does not allocate any blocks on the VMFS volume until they are needed. This usually causes lots of fragmentation, and a deleted recovery of a thin disk is more complex than on a tick disk.

Data Recovery Tips

- Storage Vmotion delete VM’s
  - If we are going to recover a deleted VMDK file, many customers need to move their working VMs over to another storage before we start the analysis
  - Storage Vmotion allows a virtual machine’s VMDK files to be moved from one VMFS volume to another while the machine is still running.

StorrageVmotion – “move during operation”

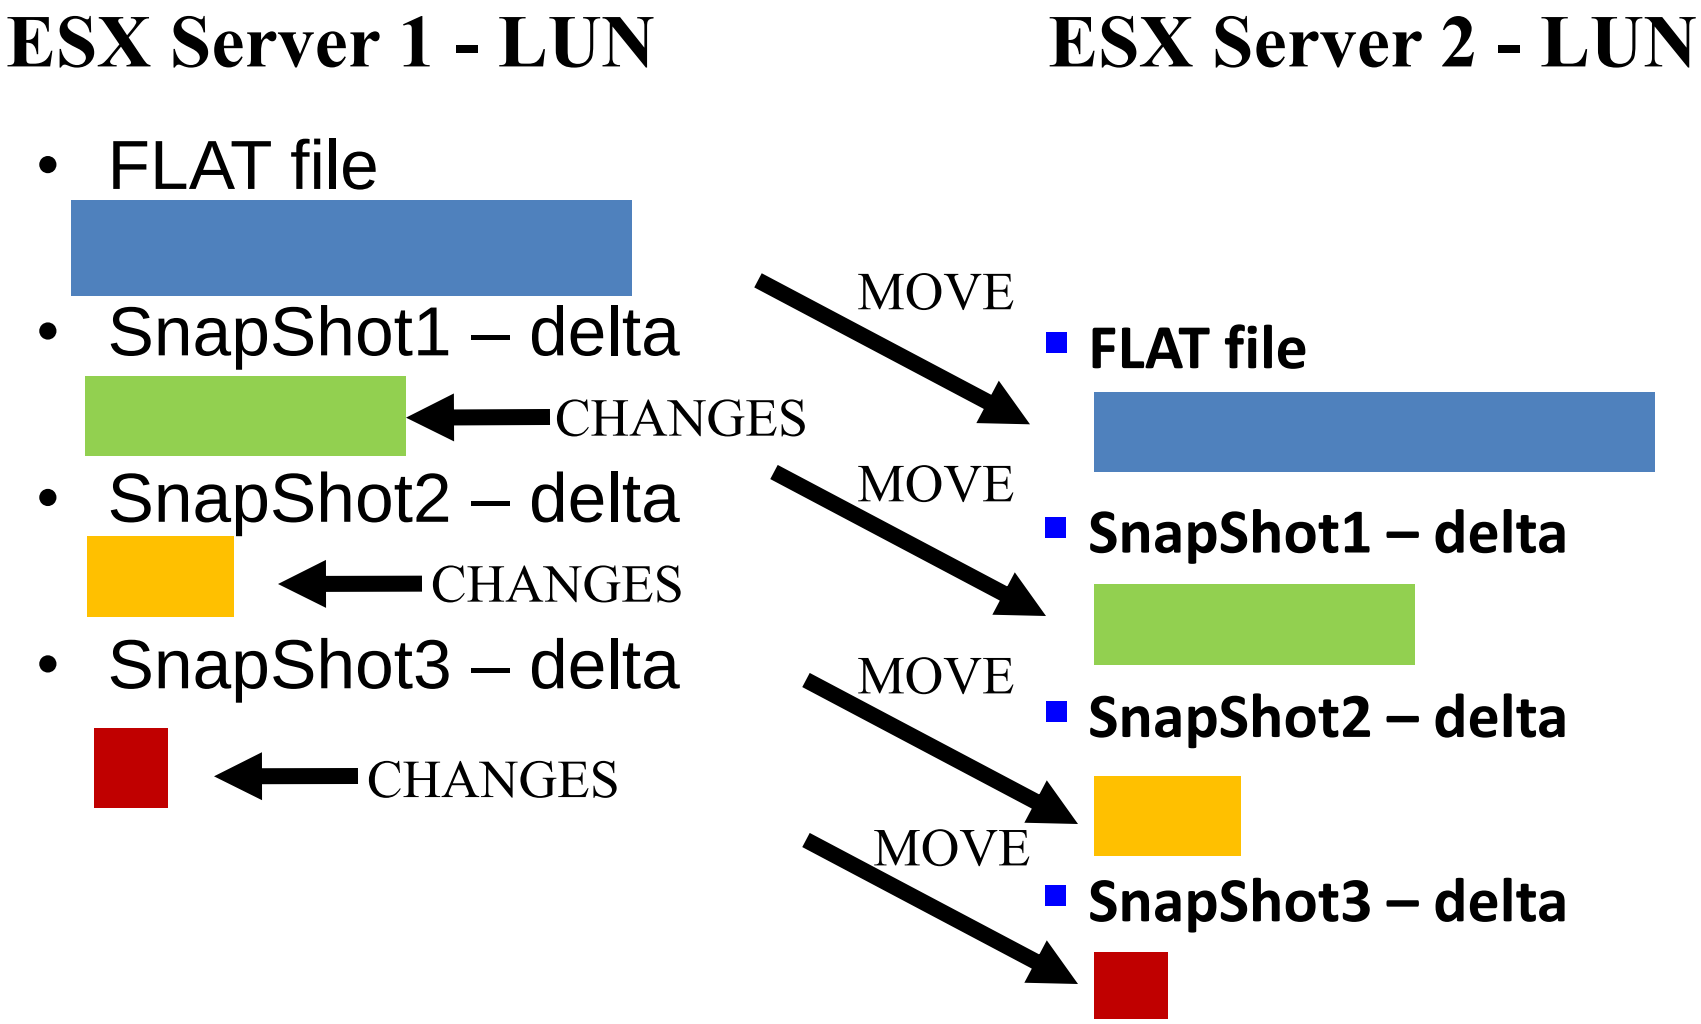

Data Recovery Tips

Storage Vmotion - Example.

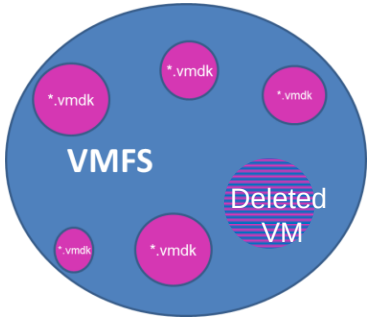

One VM is deleted, 5 VMs are OK and in use. VMFS don't keep track of the pieces of deleted VMs. So it will be fragmented like this

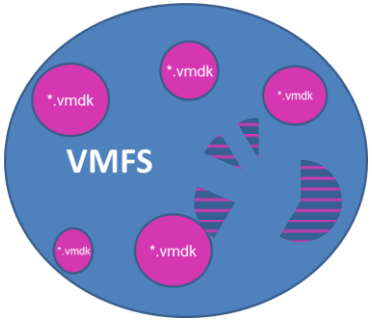

Storage Vmotion deletes the VMs from the source datastore in the same process as they are moved

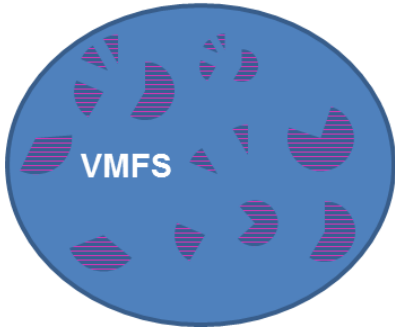

This recovery will be much more complex than if the remaining VMs had not been deleted.

Data Recovery Tips

- Storage Vmotion delete VM’s
  - We suggest cloning the VMs instead. Cloning will keep the VMDK files on the original VMFS volume, while the VMs can continue to run on the copies on the destination VMFS volume.
  - We like to have the original VMFS volume kept as “original” as possible.

Data Recovery Tips

- Cloning a VMFS volume
  - If it is not possible to clone the existing VMs over to another VMFS volume, and the original VMFS volume must be kept online, it is possible to clone the entire VMFS volume.
  - This can be done with the command “dd” in ESX server.
  - If the VMFS volume is presented to a Windows machine, it can be cloned with for example FTK Imager Lite.
  - We can work with both raw clones (disk to disk) and file clones (disk to file). However, if the VMFS volume is cloned to a file, the file must be stored on an NTFS file system that is mounted on a Windows machine for a Remote Data Recovery to be possible.

Case study

HP EVA – Multiple Drive Fail

- Hardware Details
  - HP EVA 6000 SAN
  - 80 drives, 40TB
  - 2 disk groups, 18 vDisks

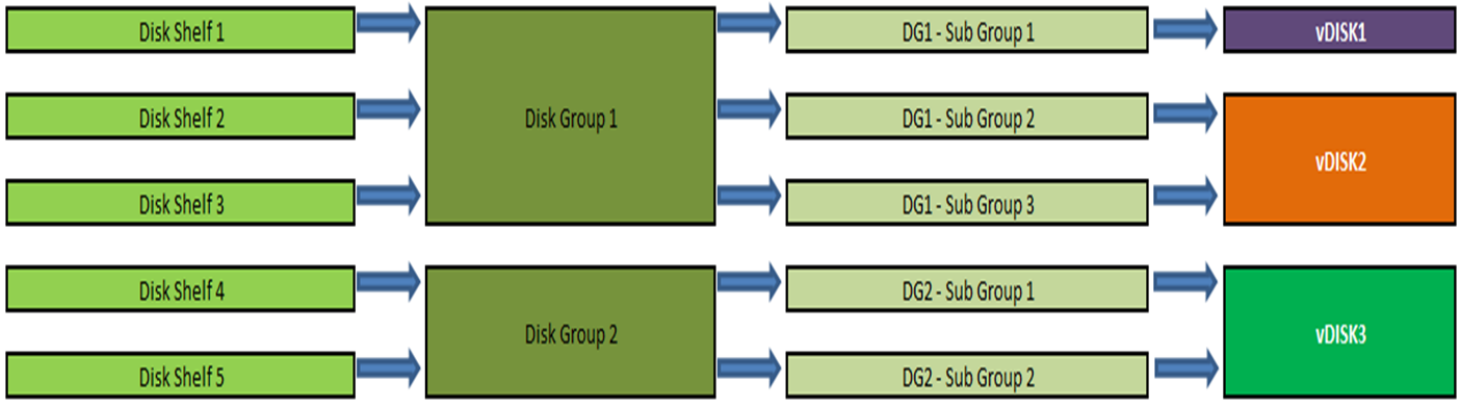

HP EVA – Multiple Drive Fail

- Data Loss Event
  - Server room flooded at the Gothenburg Rescue Service
- Impact
  - Technology worth millions of Swedish Kroner damaged
  - Rebuild of one SQL database estimated at 20,000 man hours
  - Non-functional backups

HP EVA – Multiple Drive Fail

- Ibas Kroll Ontrack Recovery
  - All but 24 of the drives were recovered in clean room
  - Critical volumes rebuilt even with missing drives
  - Over 4TB of data recovered
  - 100% of critical data recovered
    86% of total data recovered

Thank you for your attention

Øyvind Nyland
Operation Manager Ibas

ibas Kroll Ontrack

[Return]

Return

[Return]

[Return]

Return

Temperature tolerances

700°C

- **Cobalt/Chrome,**
- **Cobalt/Nickel**

300°C

- **Iron oxide**
- **Soldering**

150°C

- **Plastic**

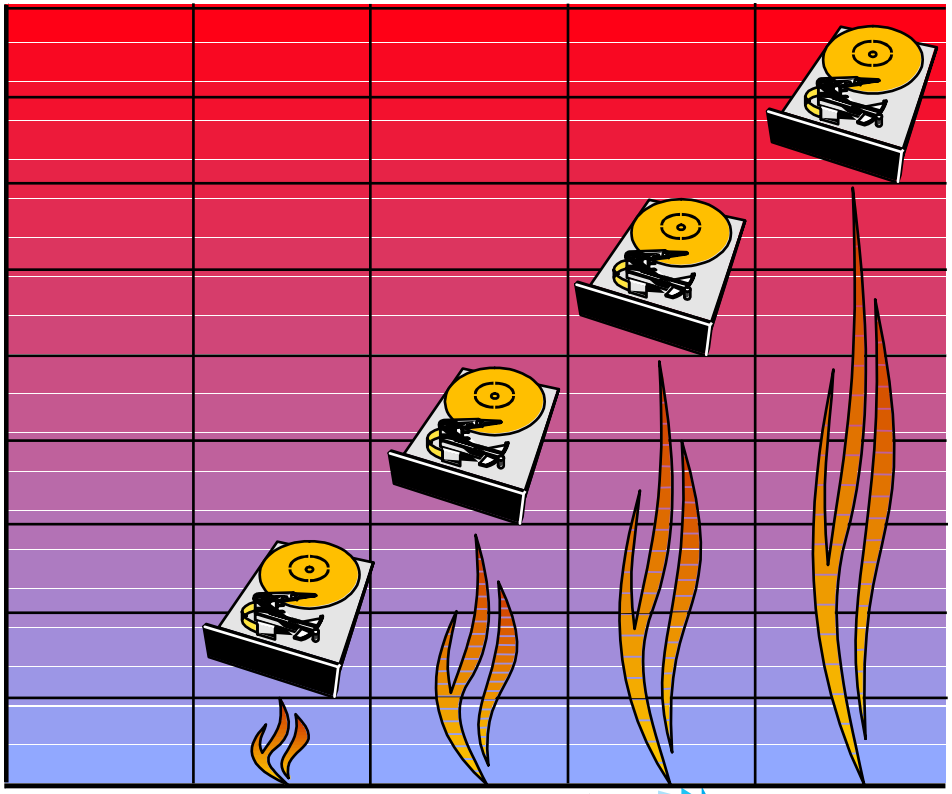

Water & Fire Damages

1. Handling, Logistics and Routines
2. Incident scene, drawings of premises etc
3. Documentation and labelling
4. Decide priority list
5. Time factor
6. Handling, packaging and shipment
7. Ibas personnel on site

Return

Return

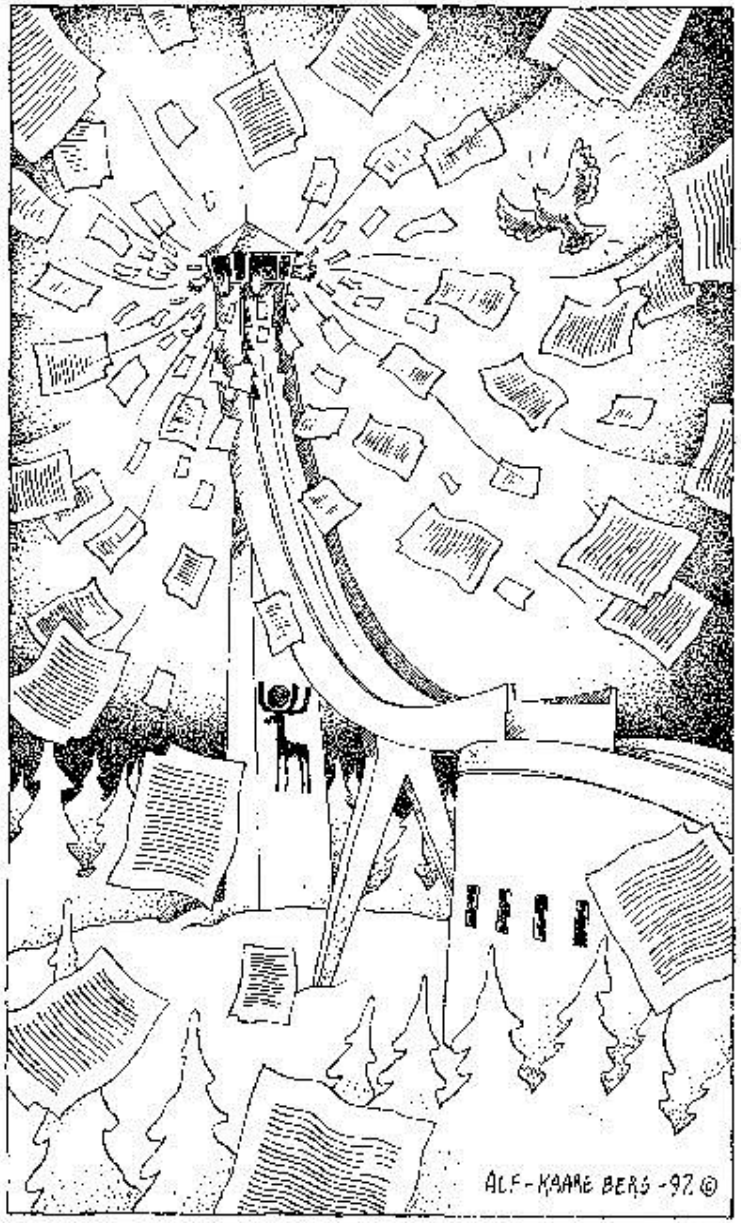

Return

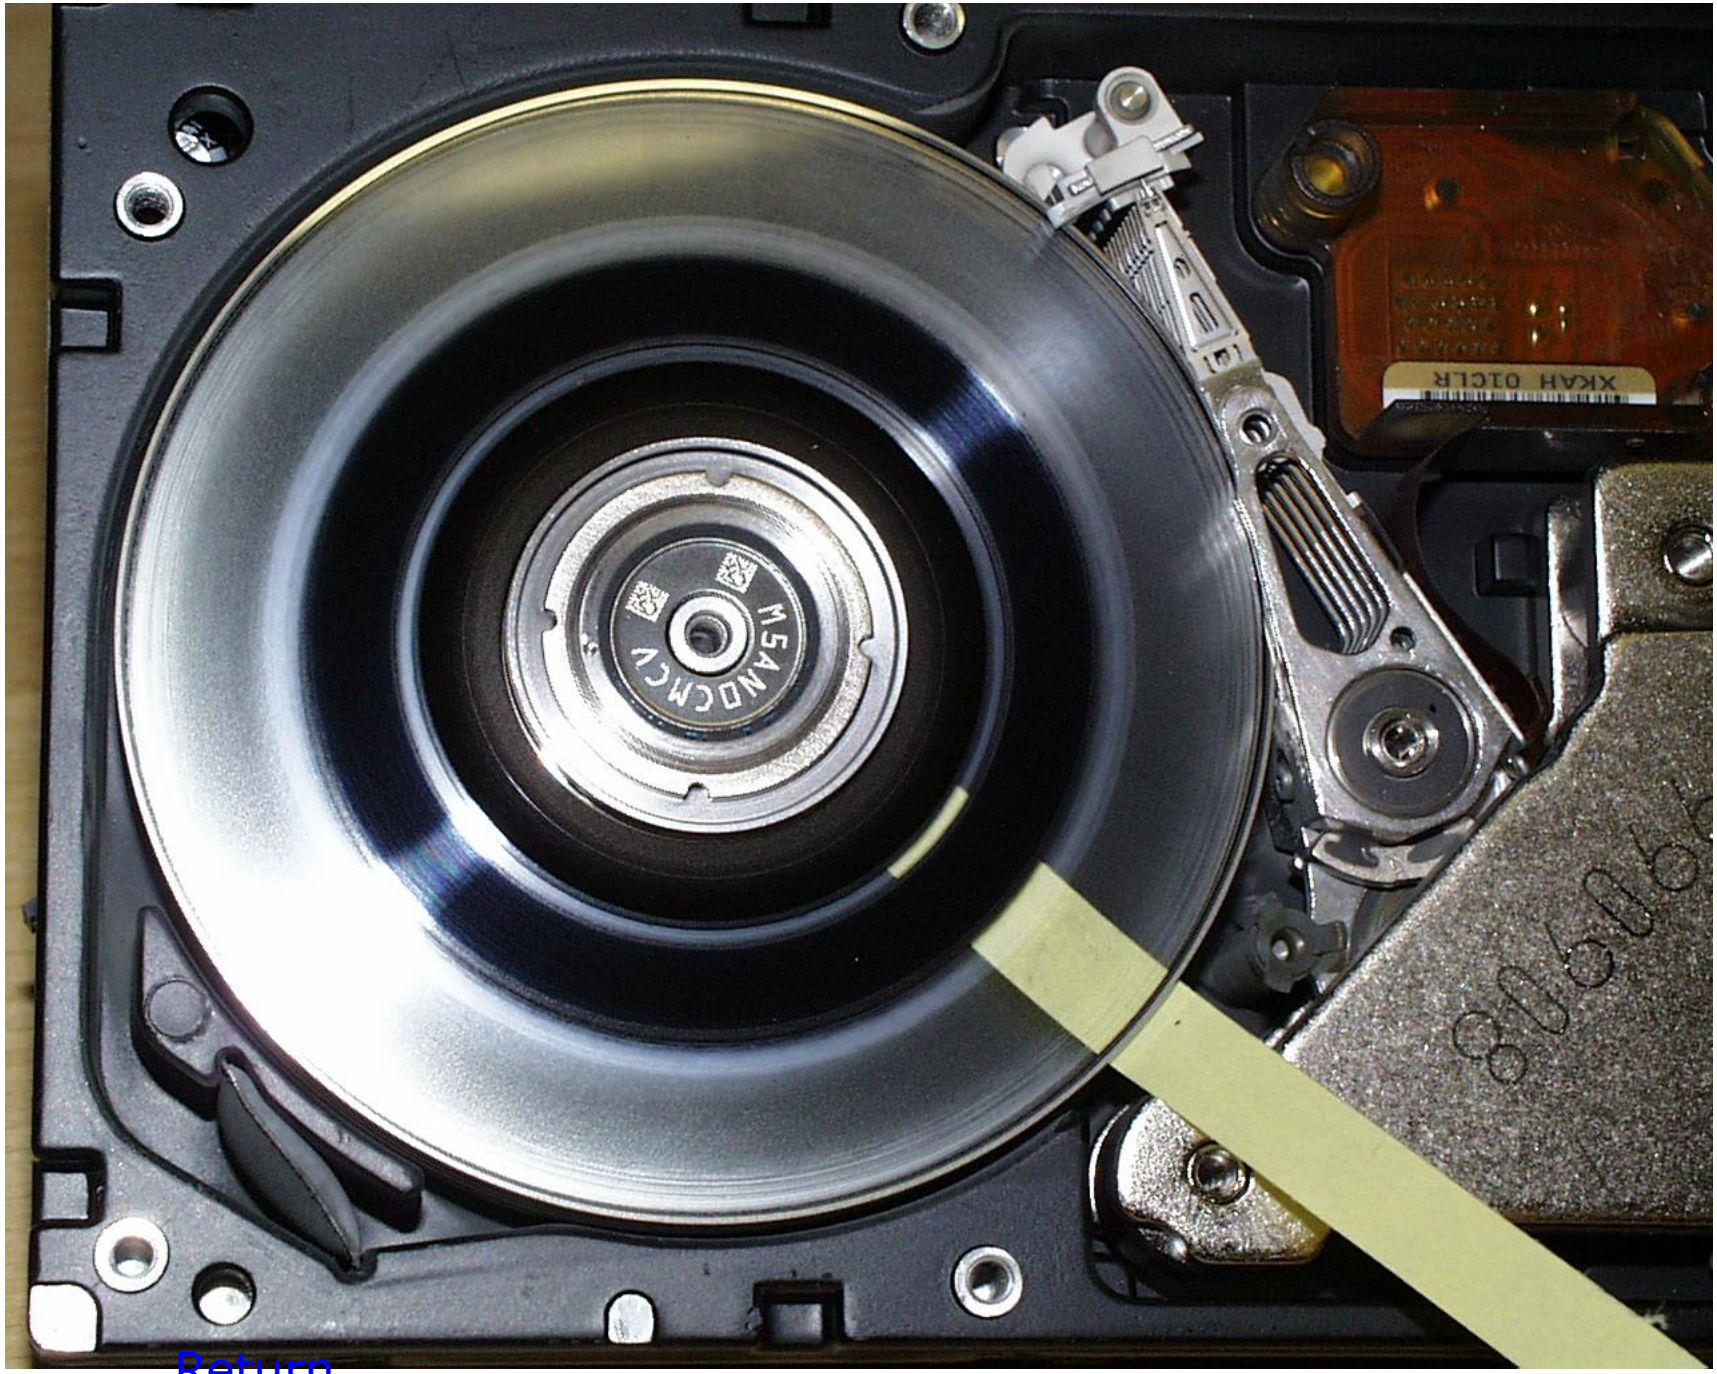

[Return]

[Return]

Return

Physical tolerances

- 270000 Track pr. inch → 10639 tracks pr. mm.
- 1383000 Bit pr. inch → 54448 Bit per mm.
- One A4 page filled with “m” formatted with Calibri size 11
  - This gives 2295 characters – One letter is 8 bit
  - This gives a total of 18360 bit pr. A4 side
- If you cut the media in pieces of 5 x 5 mm.
  - One data track from this piece will contain 14,8 A4 pages
  - The whole piece will contain 787286 A4 pages
- One A4 sheet is approx. 0.1 mm

- If we stack these sheets on top of each other we get a pile that is approx. 80 meters high

- Return to Physical tolerances

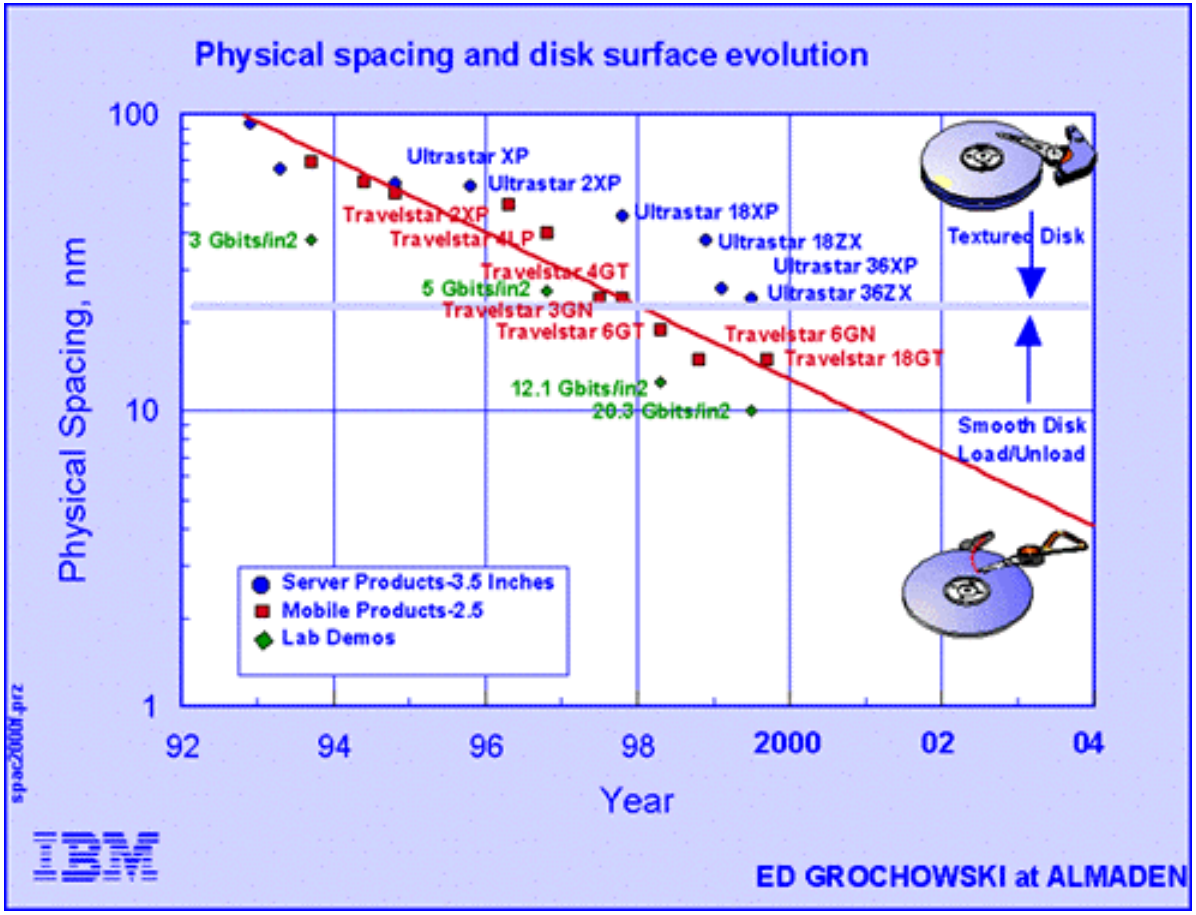
Physical spacing and disk surface evolution
Physical Spacing, nm
100
10
1
92
94
96
98
2000
02
04
Year
Ultrastar XP
Ultrastar 2XP
Travelstar 2XP
Travelstar 4LP
Ultrastar 18XP
Ultrastar 18ZX
Ultrastar 36XP
Ultrastar 36ZX
3 Gbits/in2
5 Gbits/in2
Travelstar 4GT
Travelstar 3GN
Travelstar 6GT
Travelstar 6GN
Travelstar 18GT
12.1 Gbits/in2
20.3 Gbits/in2
Textured Disk
Smooth Disk
Load/Unload
Server Products-3.5 Inches
Mobile Products-2.5
Lab Demos
spac2000f.prz
IBM
ED GROCHOWSKI at ALMADEN

[Return]

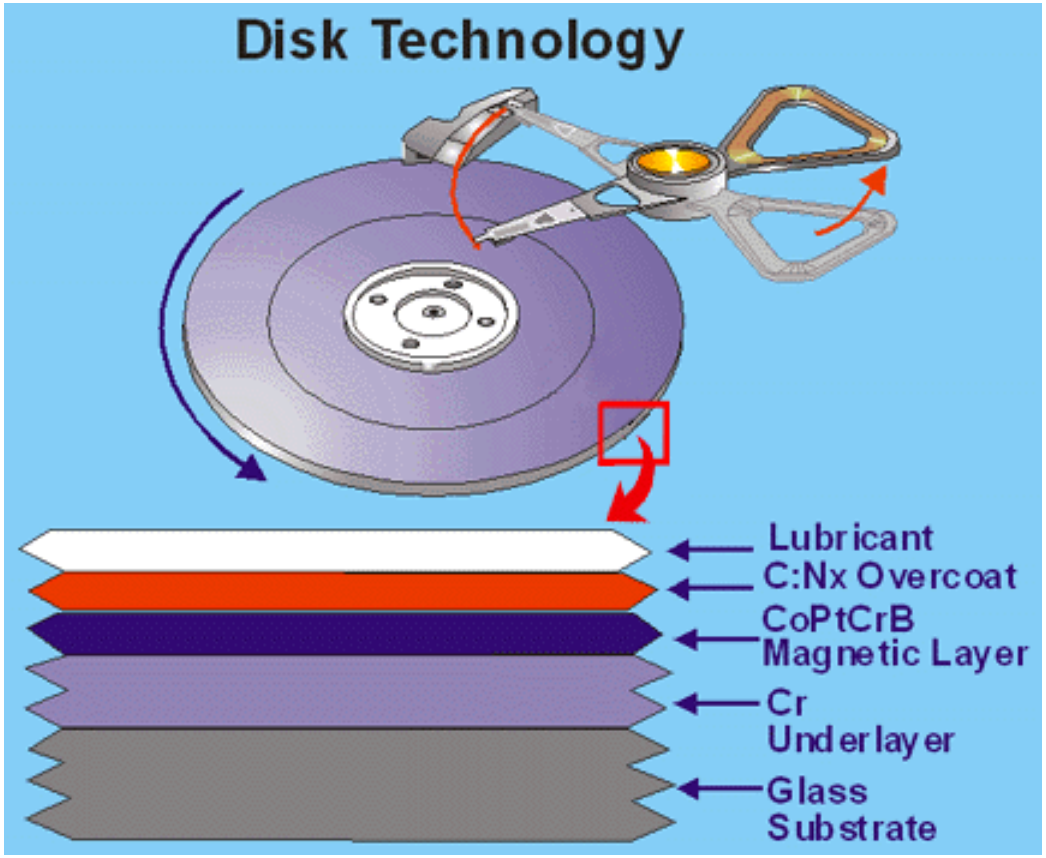
Disk Technology
Lubricant
C:Nx Overcoat
CoPtCrB
Magnetic Layer
Cr
Underlayer
Glass
Substrate

Return

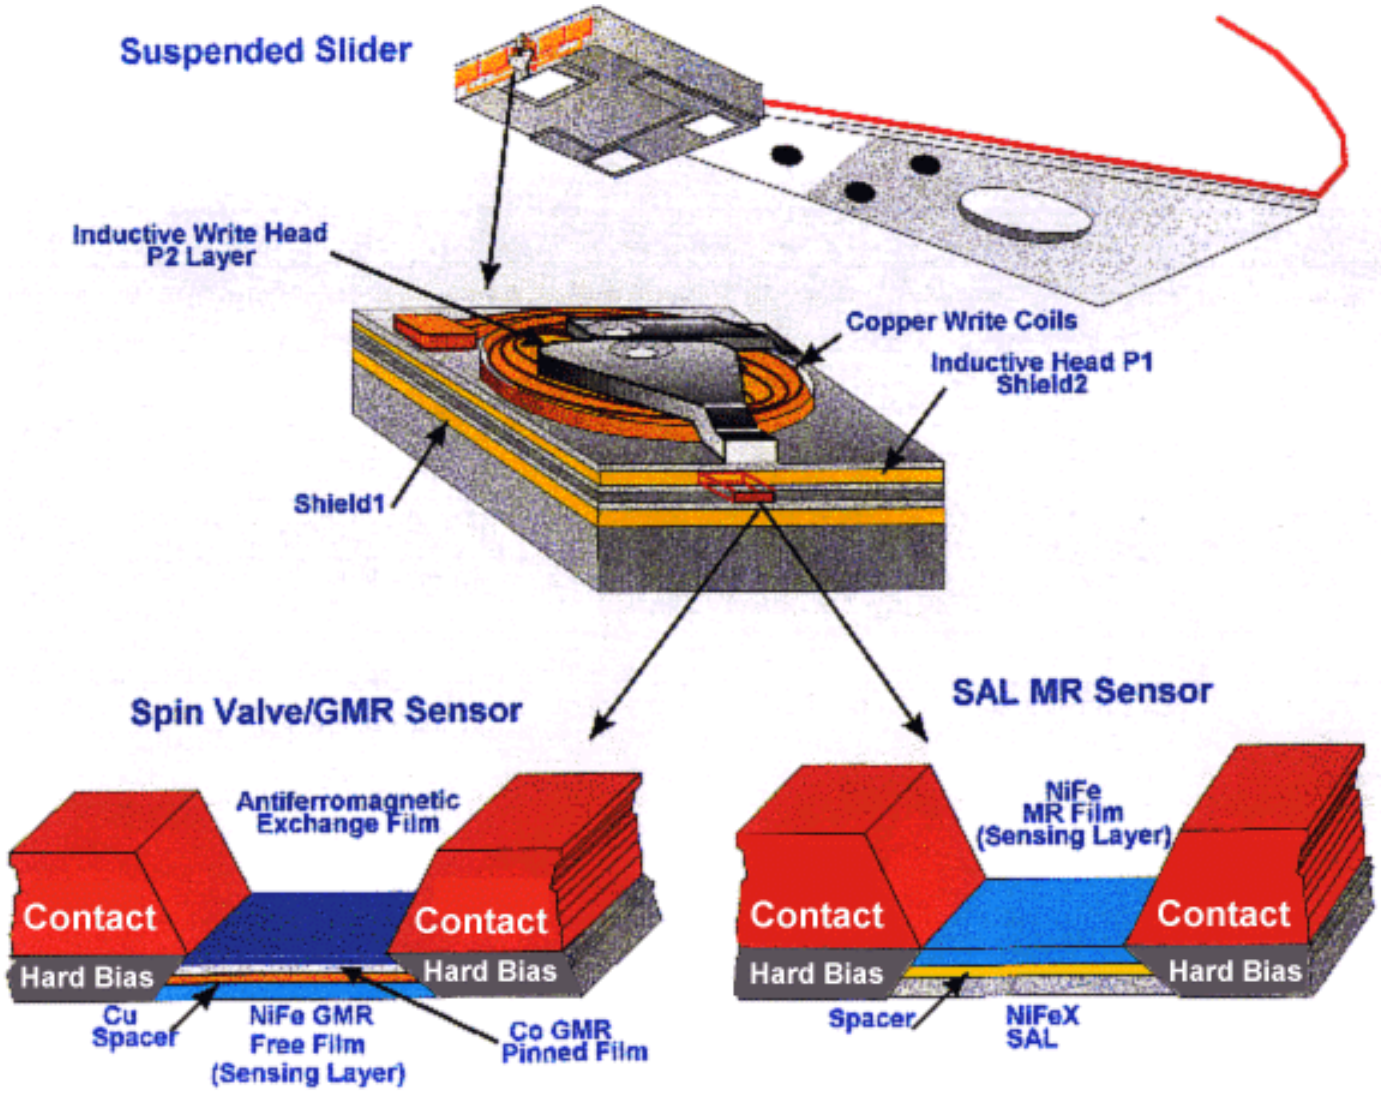

Figure 2. MR and GMR head structures.

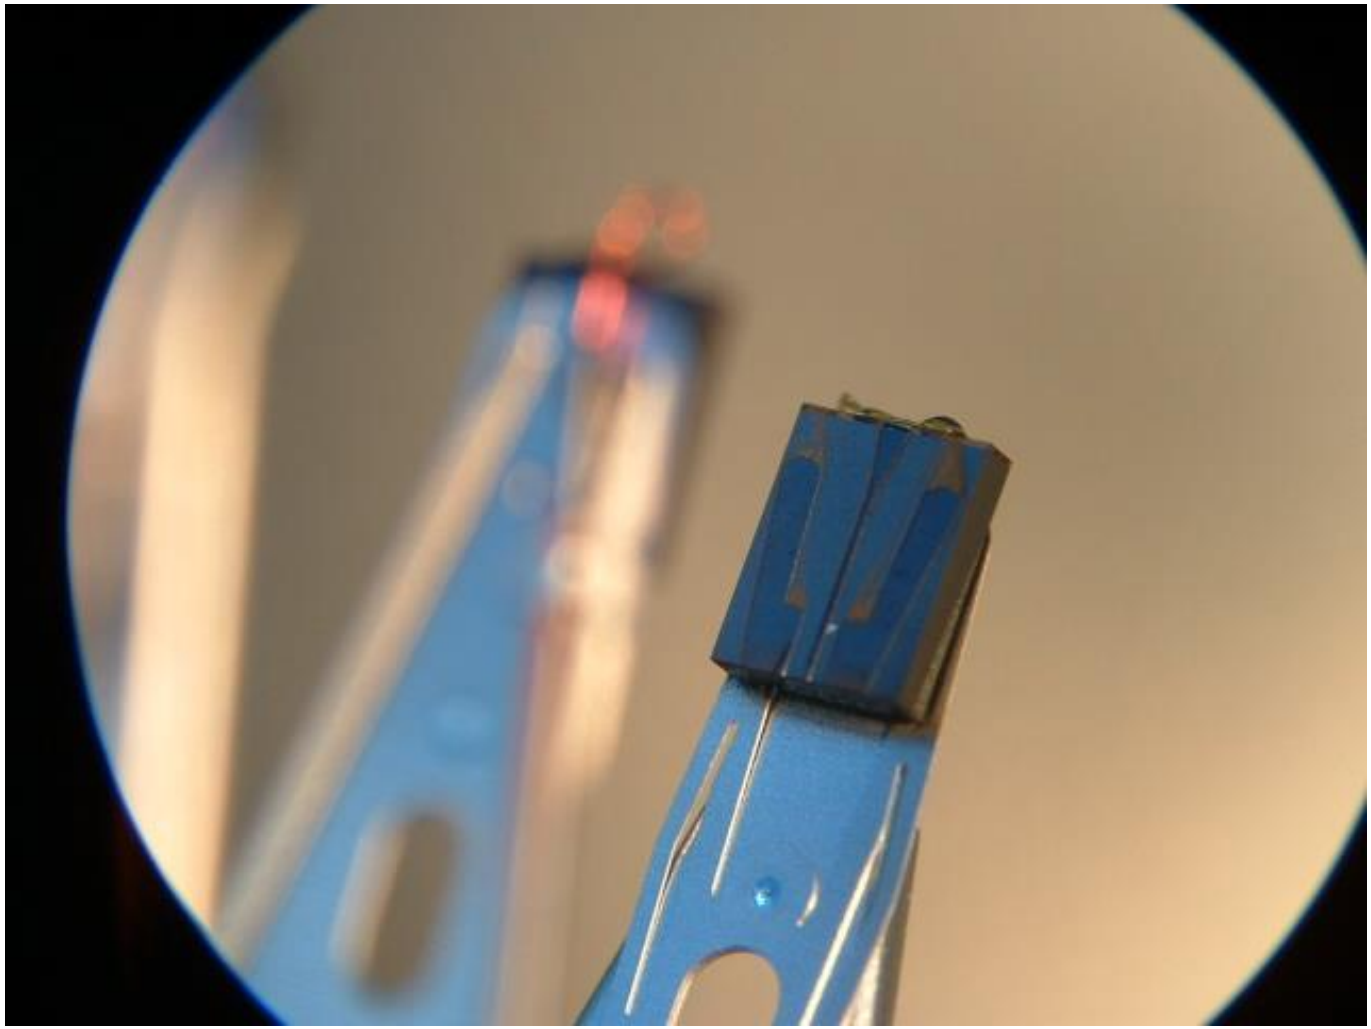

Return

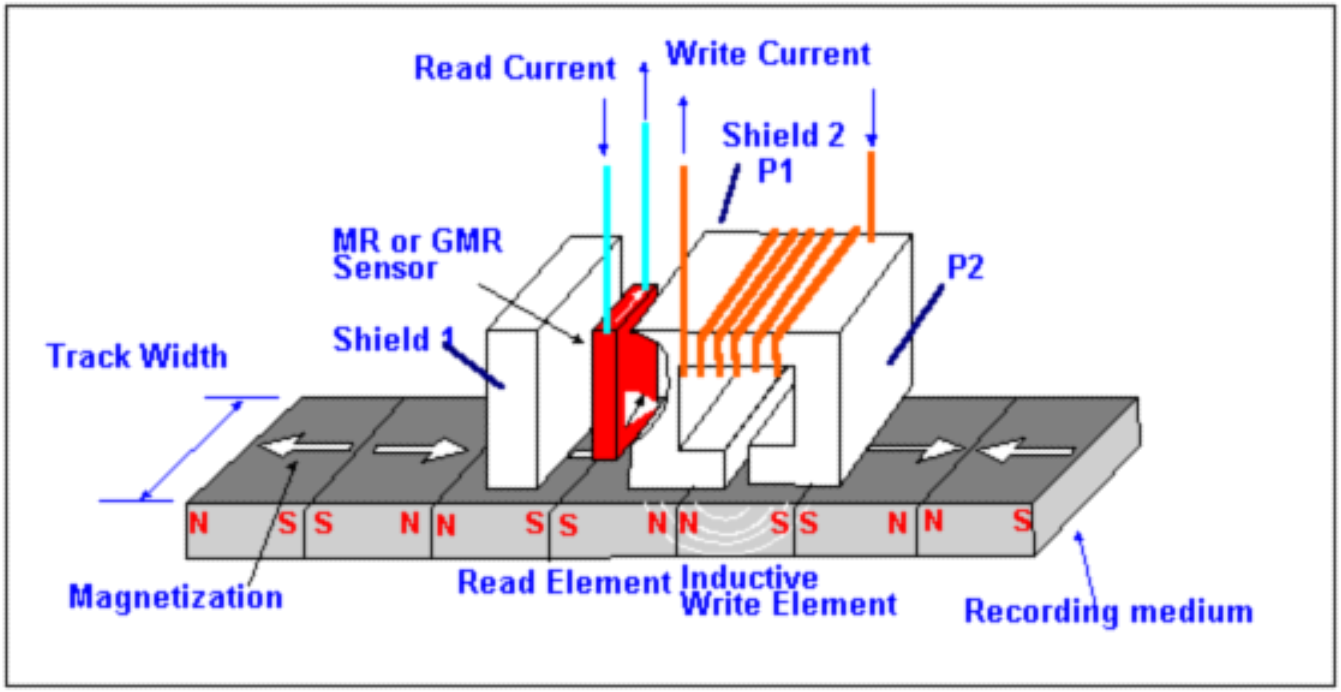

Figure 3. Magnetic recording process.

Return

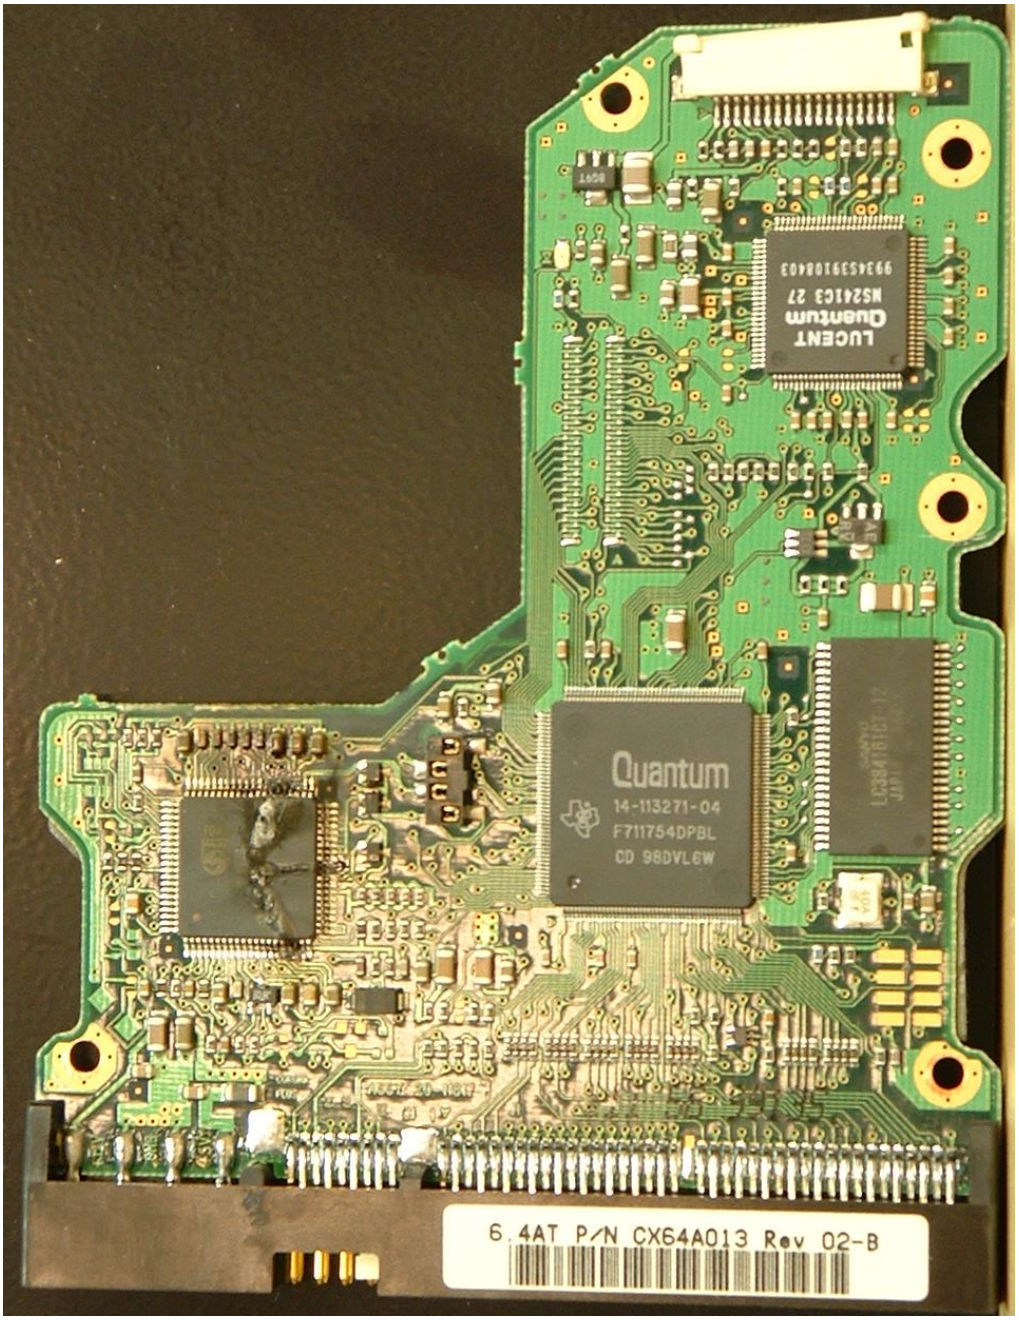
LUCENT
Quantum
MS2413 27
Quantum
14-113271-04
F7111754DPBL
CD 98DVL6W
6.4AT P/N CX64A013 Rev 02-B

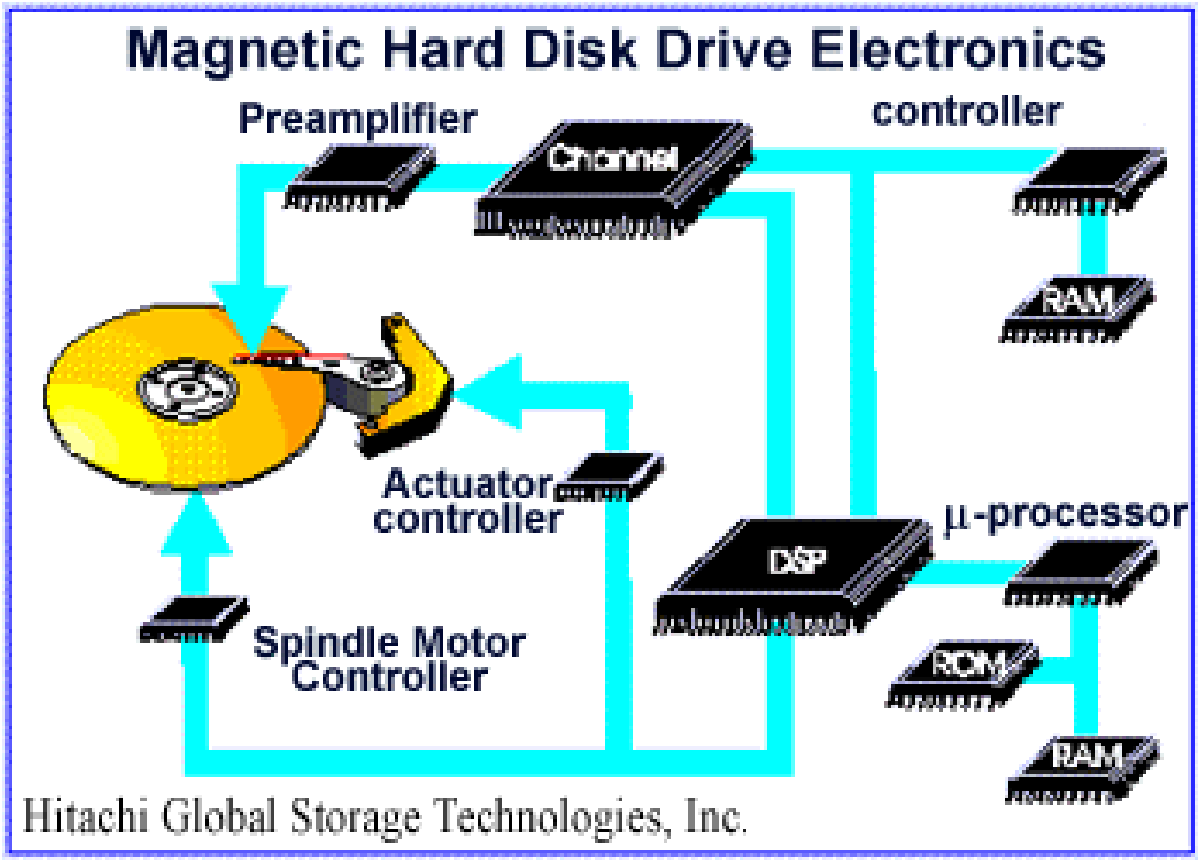
Magnetic Hard Disk Drive Electronics
Preamplifier
Channel
controller
RAM
Actuator
controller
µ-processor
DSP
Spindle Motor
Controller
ROM
RAM
Hitachi Global Storage Technologies, Inc.

Return

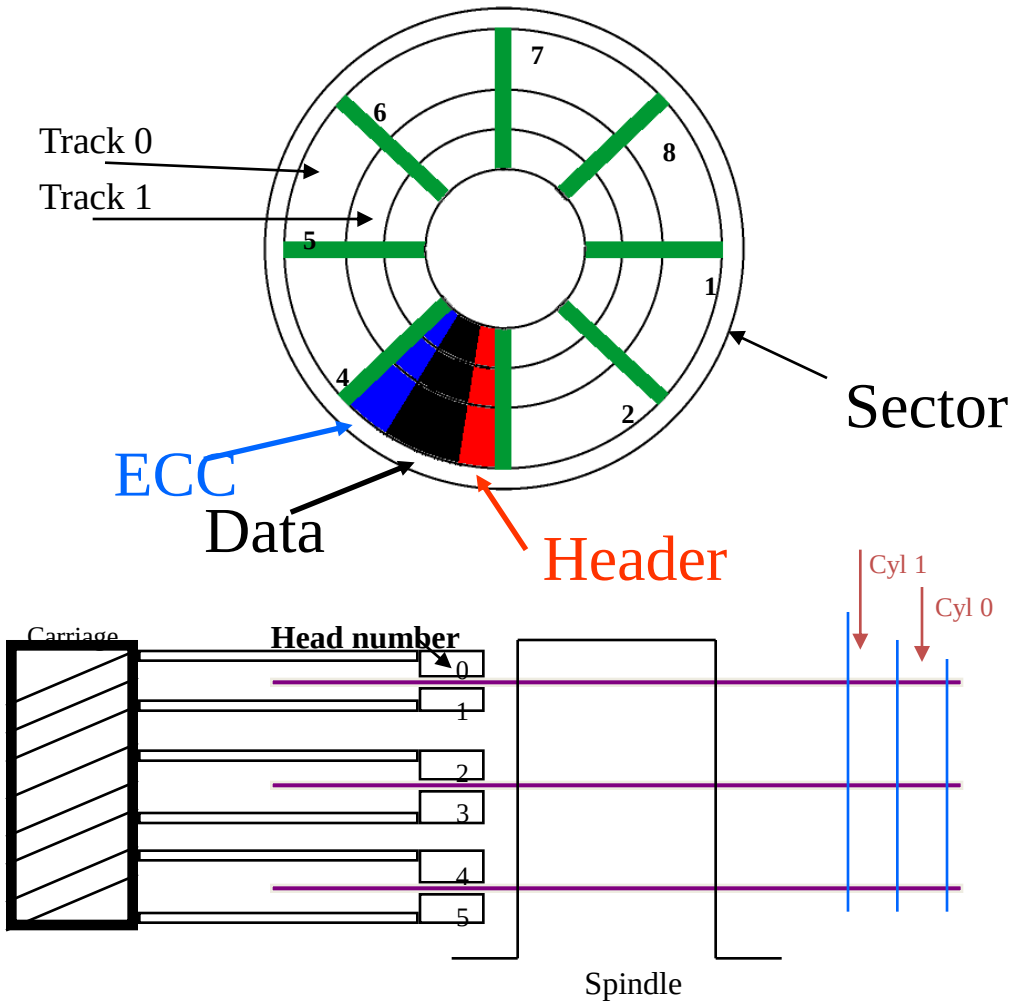
Track 0
Track 1
7
6
8
5
1
4
2
Sector
ECC
Data
Header
Cyl 1
Cyl 0
Carriage
Head number
0
1
2
3
4
5
Spindle

Return

Longitudinal vs. perpendicular

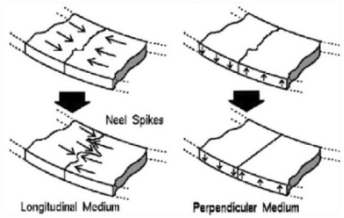

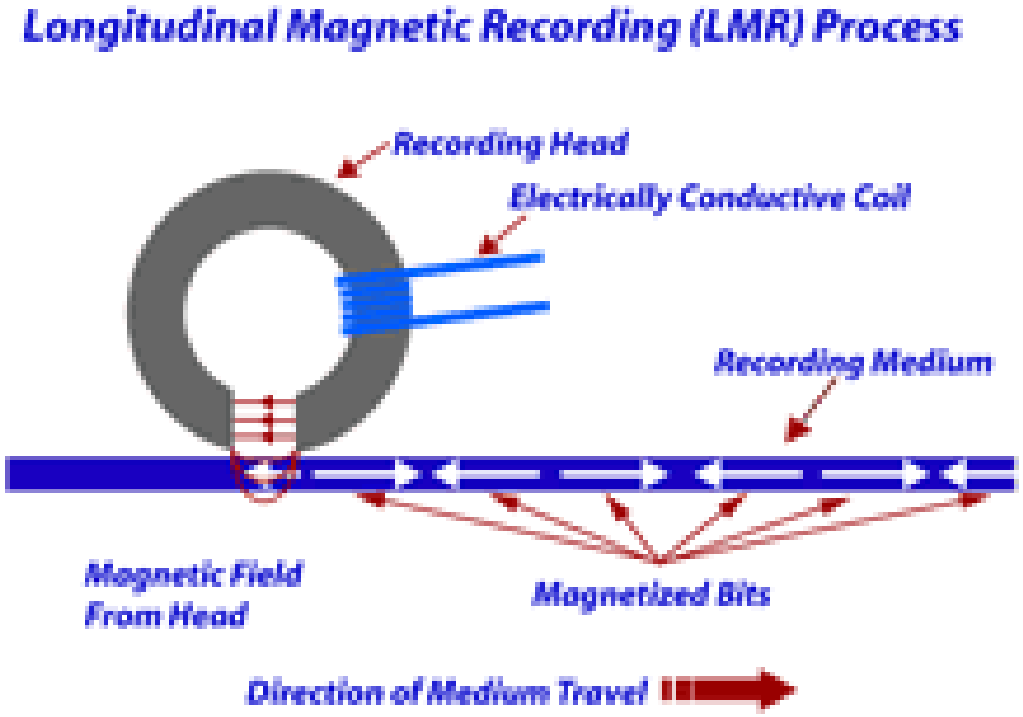

LMR places bits lying flat in the plane of the disk

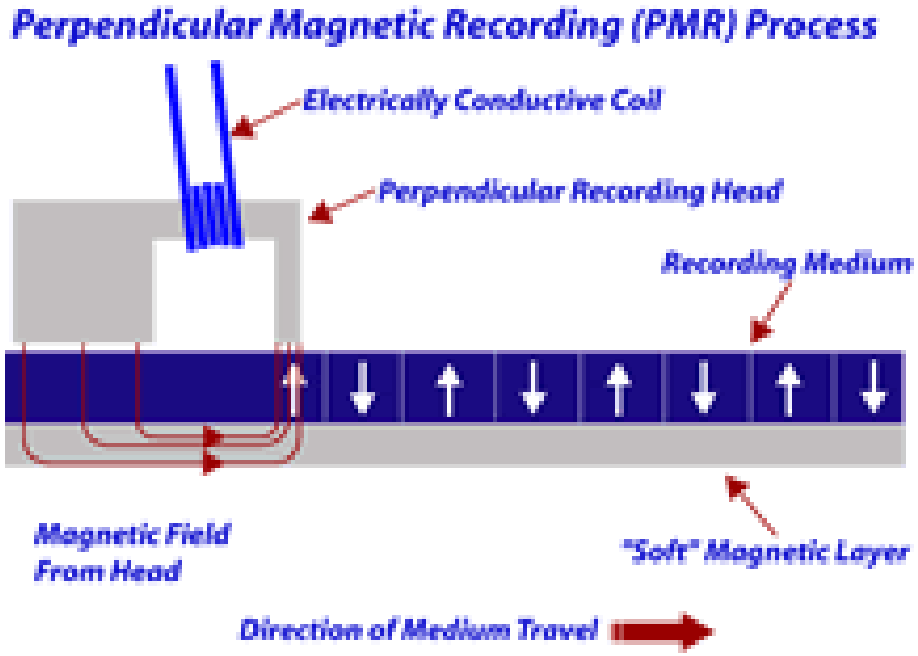

PMR configuration uses new heads and disks that record bits perpendicular to the plane of the disk

Return

Performance Bottlenecks

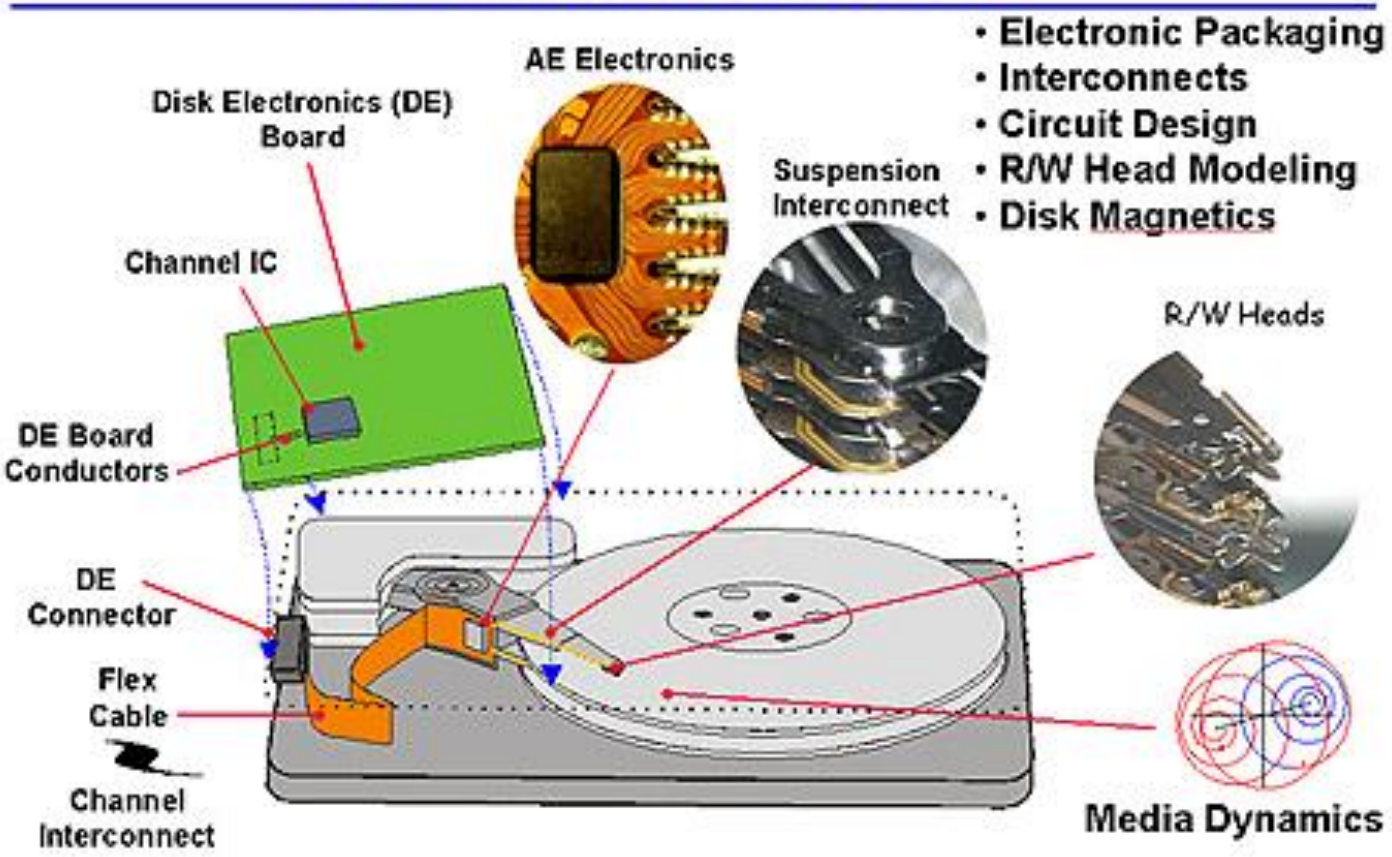

[Return]

MFM Images

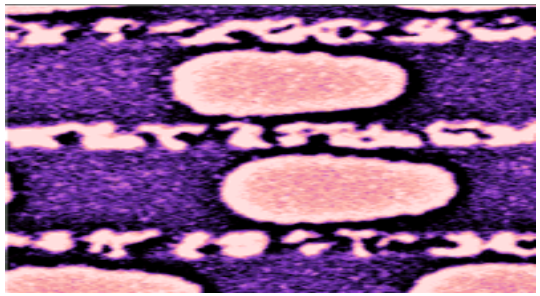

MO bits

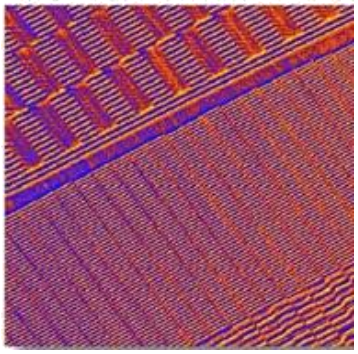

Servo pattern

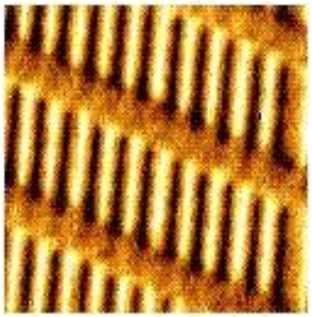

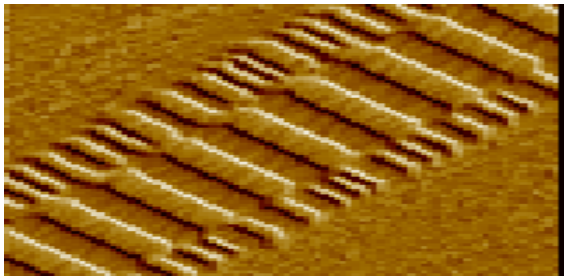

Partly Overwritten Track

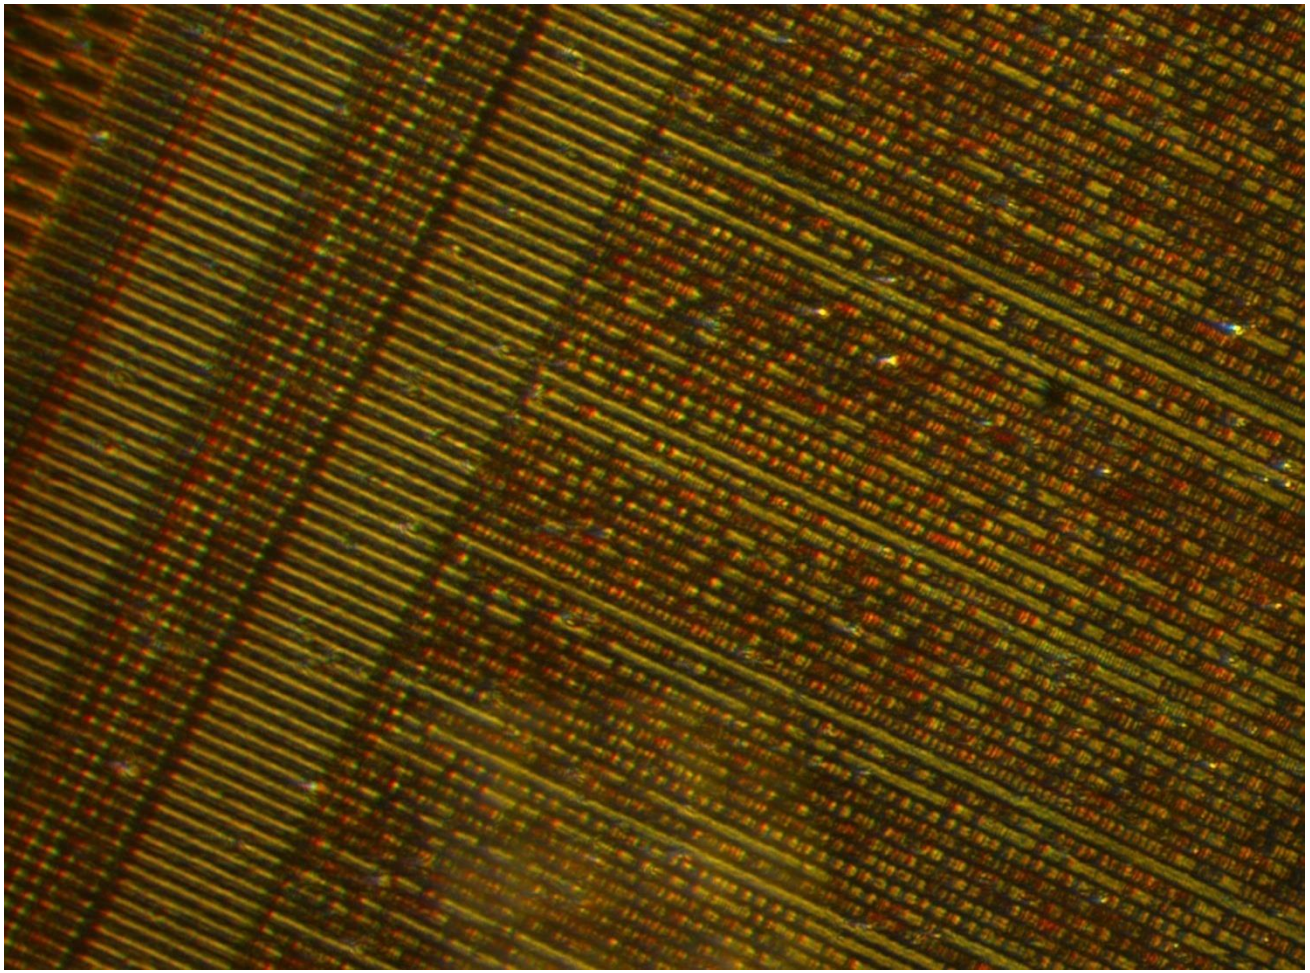

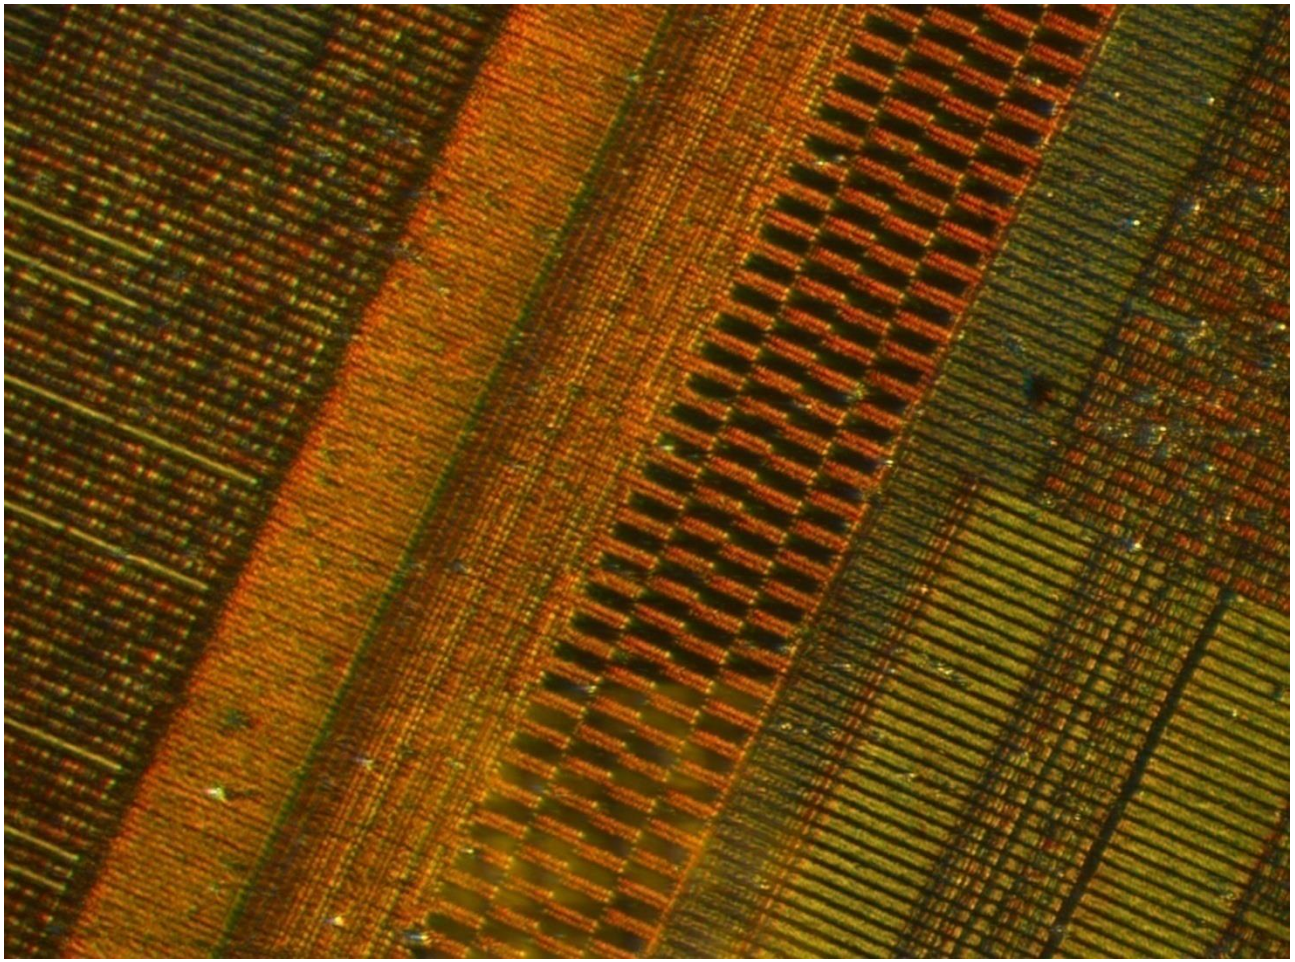

Return

Operating

• Ambient temperature	5° to 55° C
• Relative humidity (non-condensing)	5% to 90%
• Max. wet bulb (non-condensing)	29° C
• Shock (half sine wave, 2ms)	15G (11 ms)
• Vibration (random (RMS))	1.0 G, all axis

Non-operating

• Ambient temperature	-40° to 70° C
• Relative humidity (non-condensing)	5% to 90%
• Max. wet bulb (non-condensing)	29° C
• Shock (half sine wave, 2ms)	250G (2 ms) / 75G (11ms)
• Vibration (random (RMS))	5.0 G, all axis

Return

Return

Return