

Towards Near-Data Processing in Deep and Cold Storage Hierarchies

Marcus Paradies, German Aerospace Center (DLR)

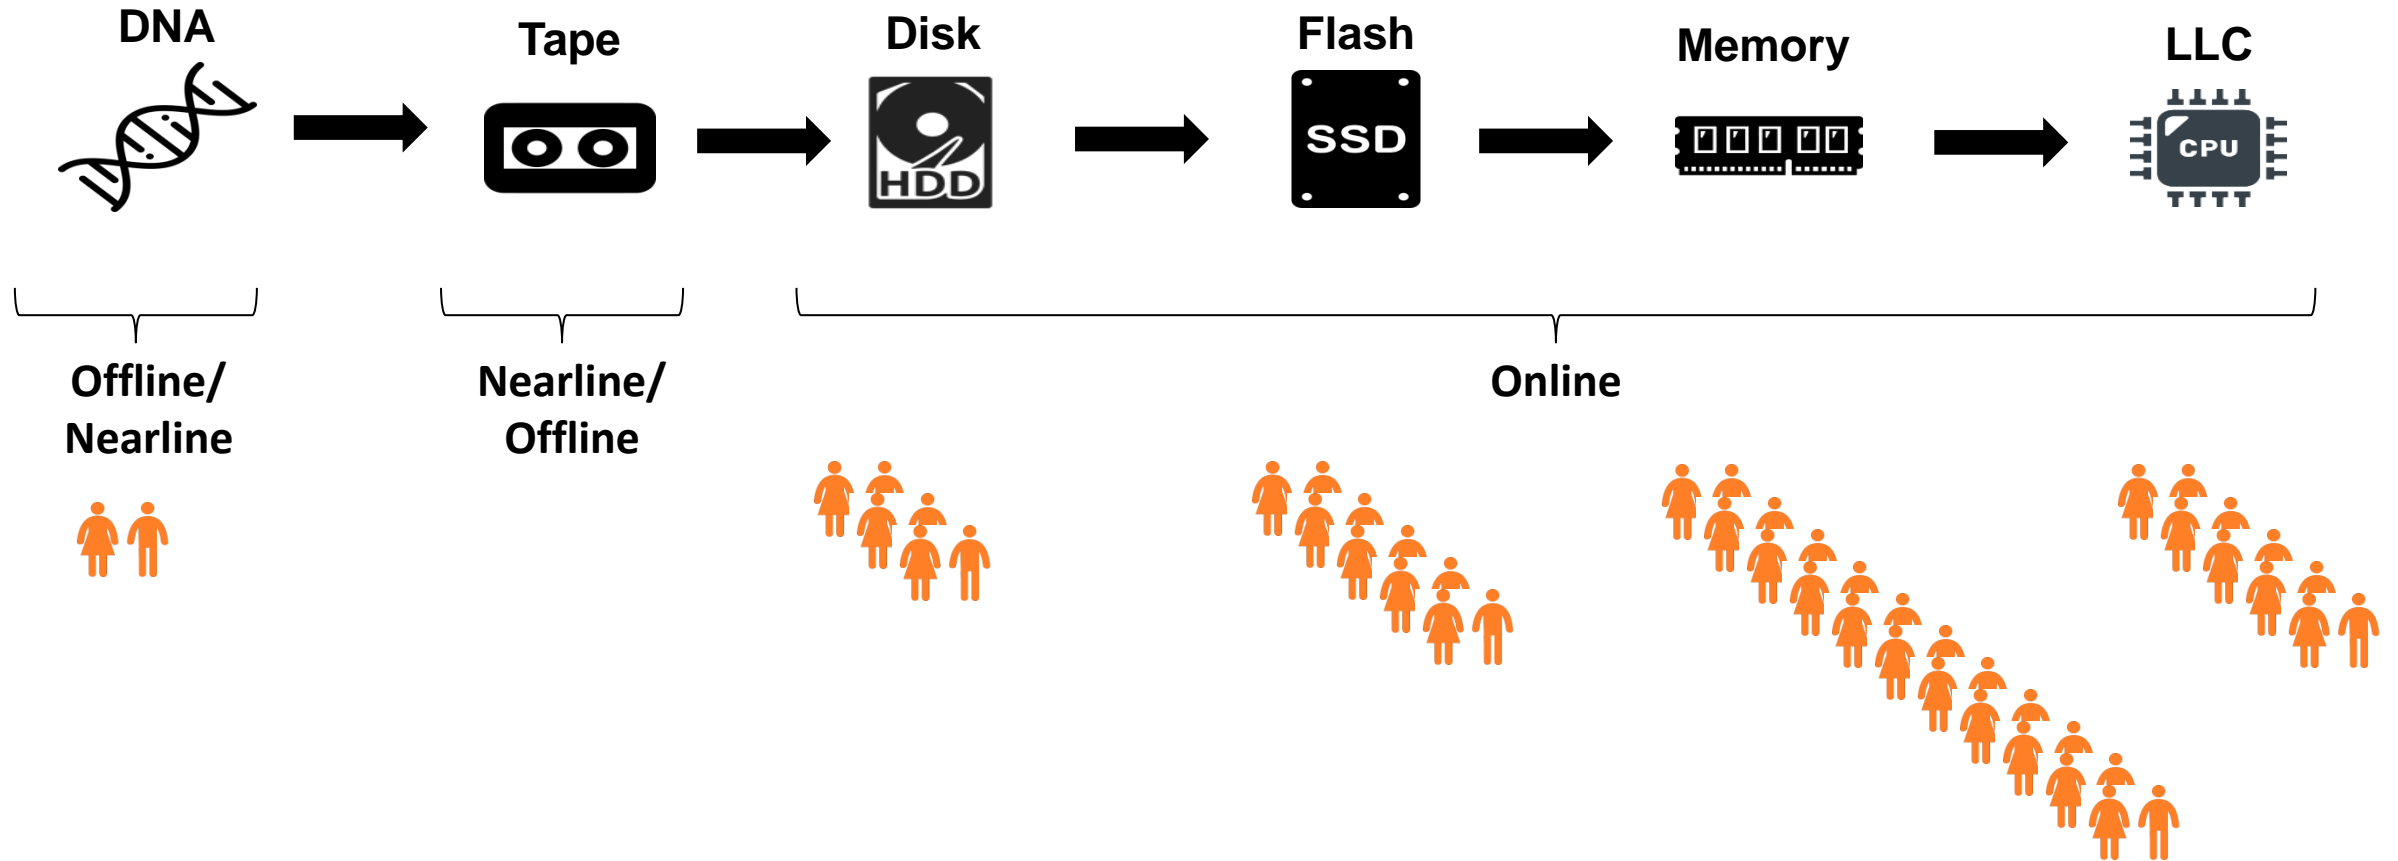
04/03/2019, XLDB



Knowledge for Tomorrow



The Storage Hierarchy and Its (Current) Coverage in DB Research



A young man with curly hair and a shocked expression is the central focus of the image. He is wearing a white t-shirt with the word "BILLBOARD" printed on it. He is surrounded by a dense crowd of people at what appears to be a party or concert. The lighting is dim and warm, typical of an indoor event. The text "NEARLINE STORAGE?" is overlaid at the top in large, white, bold letters with a black outline. At the bottom, the text "I THOUGHT TAPES ARE DEAD" is also overlaid in the same style. A small watermark "memegenerator.n" is visible in the bottom right corner.

NEARLINE STORAGE?

I THOUGHT TAPES ARE DEAD


*Active Data Archives for
Scientific Data*

STORAGETEK



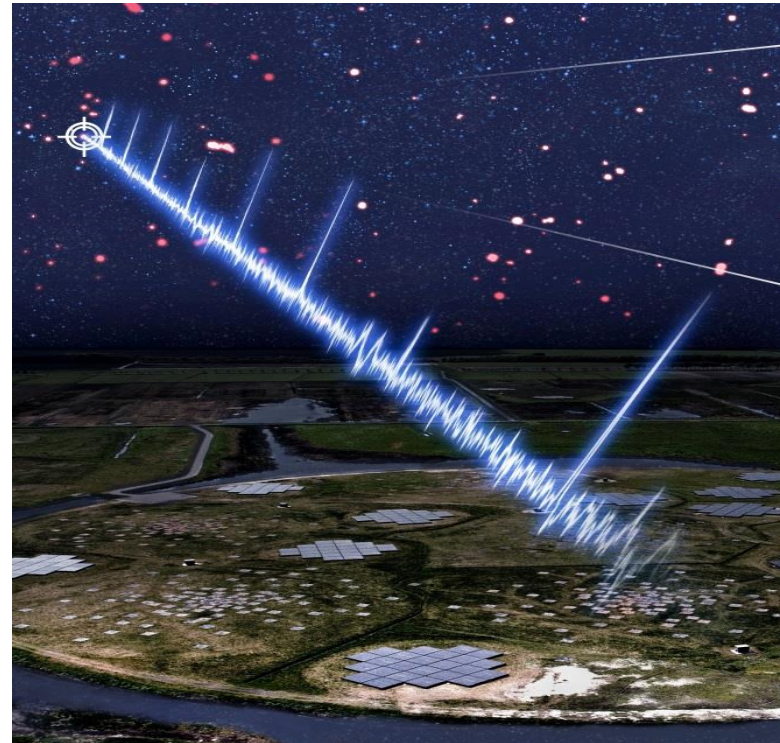
Scientific Application Domains

Earth Observation



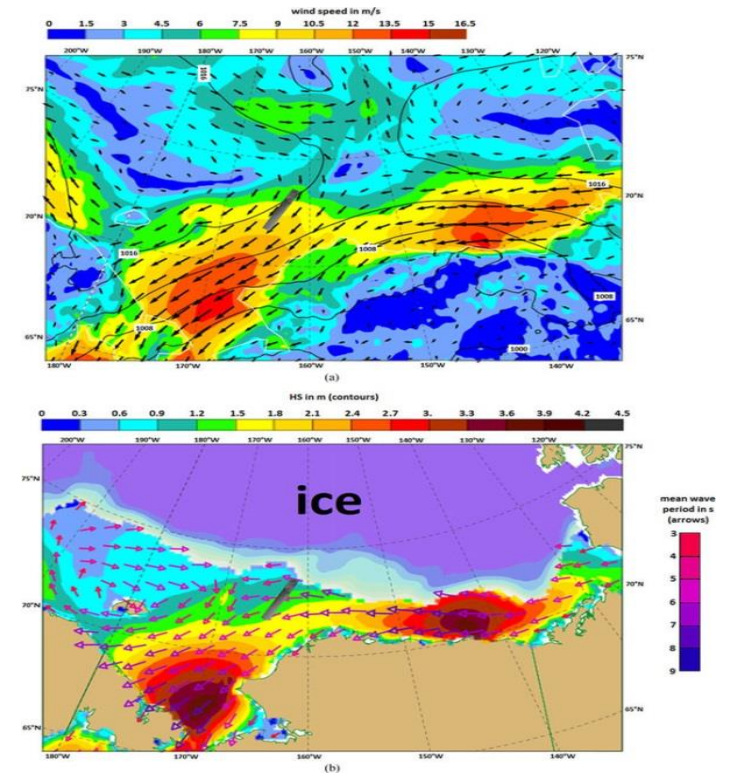
Archive: 14 PB
Disk Cache: 175 TB

Radio Astronomy



Archive: 50 PB
Disk Cache: 750 TB

Weather Forecasting




Archive: 100 PB
Disk Cache: 1.34 PB

Data Movement as Major Performance Bottleneck

Active data archives (and their catalogs) are like Amazon, but just for data.



- ✗ No SLAs on access latency, usually between minutes and hours
- ✗ Tail latency can be multiple days
- ✗ Historic data analysis can easily request 100s of TB



CryoDrill---Near-Data Processing for Cold Storage

- ❄ Focus on nearline storage (archival disks, tape)
- ❄ Consider all NDP opportunities (in-network, in-storage)
- ❄ Push data reduction ops down the storage hierarchy

