

DSS


Data management services

Alberto Pace

Data and Storage Services

EBI-CERN workshop
from Monday, 2 September 2013

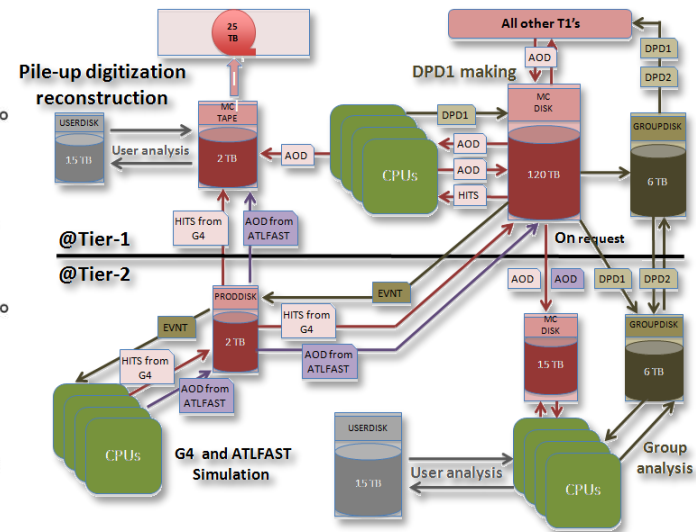


- Scientific research in recent years has exploded the computing requirements
 - Computing has been the strategy to reduce the cost of traditional research
 - Computing has opened new horizons of research not only in High Energy Physics
- Data management is one of the three pillars of scientific computing



- Some numbers
 - 80'000 disks, 125'000 Terabytes, 1500 servers
 - More than 50'000 TB of user data on disks
 - 123 Tape drives, 52'000 tape cartridges, 105'000 TB of capacity, 95'000 TB of physics data
 - Data inflow: ~ 6.5 GB/s, sustained
 - Status pages:
 - <http://eos.cern.ch/>
 - <http://castor.web.cern.ch/>
 - <http://sls.cern.ch/sls/dcbynum.php>

- Data Management solves the following problems
 - Data reliability
 - Access control
 - Data distribution
 - Data archives, history, long term preservation
- In general:
 - Empower the implementation of a workflow for data processing

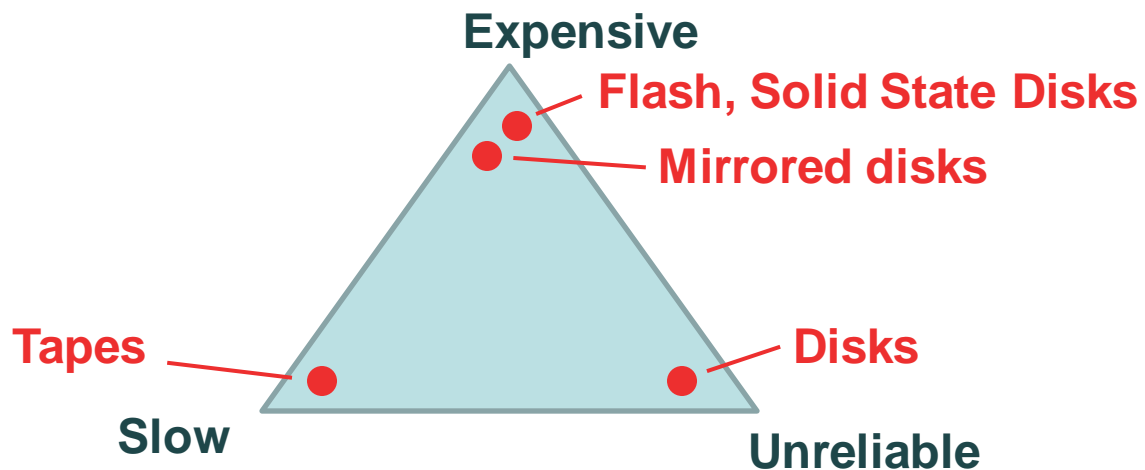
- Examples from LHC experiment data models




◆ Two building blocks to empower data processing

- ◆ Data pools with different quality of services
- ◆ Tools for data transfer between pools



- Different quality of services
 - Three parameters: (Performance, Reliability, Cost)
 - You can have two but not three




- Many ways to split (performance, reliability, cost)



- Performance has many sub-parameters
- Cost has many sub-parameters
- Reliability has many sub-parameters



- Key requirements: Simple, Scalable, Consistent, Reliable, Available, Manageable, Flexible, Performing, Cheap, Secure.
- Aiming for “à la carte” services (storage pools) with on-demand “quality of service”
- And where is scalability ?



Areas of research in Data Management

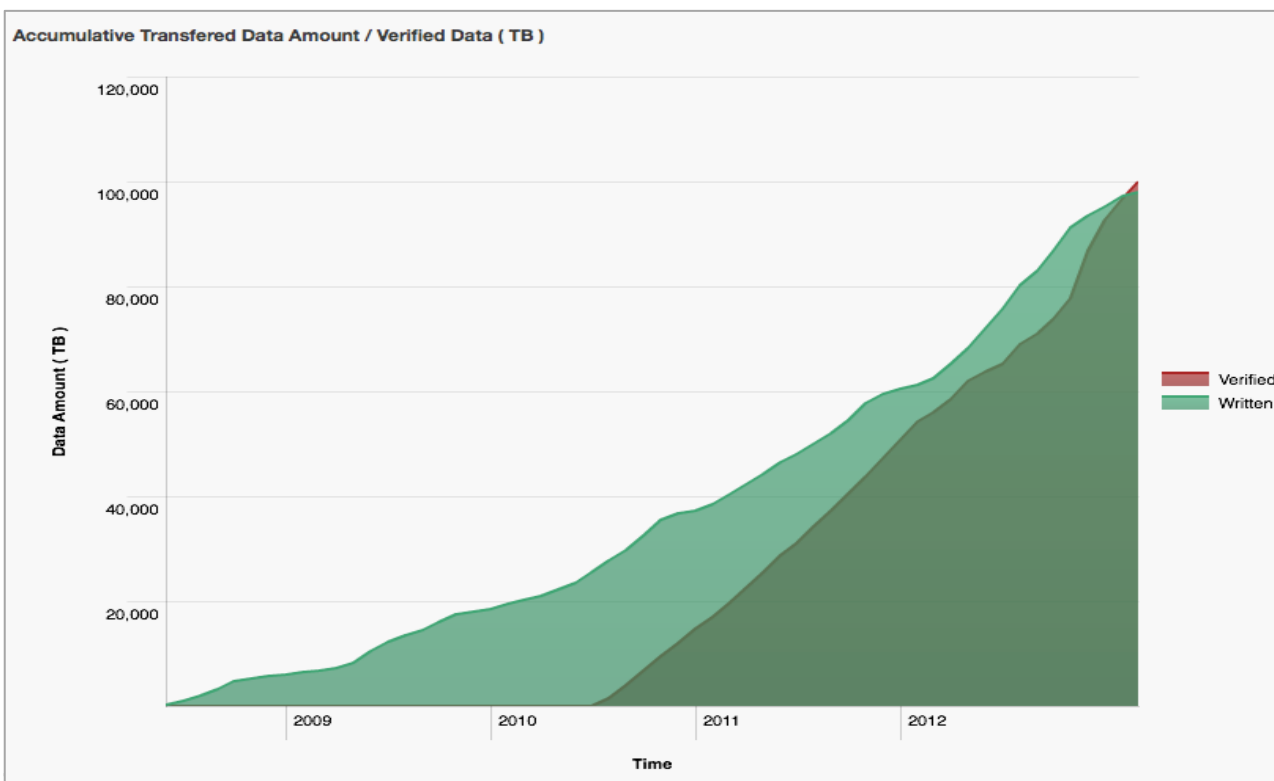
Reliability, Scalability, Security, Manageability

- Reliability is related to the probability to lose data
 - Def: “the probability that a storage device will perform an arbitrarily large number of I/O operations without data loss during a specified period of time”
- Reliability of the “service” starts from the reliability of the underlying hardware / media
 - Example of disk servers with simple disks: reliability of service = reliability of disks
- But data management solutions can increase the reliability of the hardware at the expenses of performance and/or additional hardware / software
 - Example: disk mirroring, Redundant Array of Inexpensive Disks (RAID)


- Tapes have a bad reputation in some use cases
 - Slow in random access mode
 - high latency in mounting process and when seeking data (F-FWD, REW)
 - Inefficient for small files (in some cases)
 - Comparable cost per (peta)byte as hard disks
- Tapes have also some advantages
 - Fast in sequential access mode
 - > 2x faster than disk
 - physical read after write verification (makes > 4x faster than disks)
 - Several orders of magnitude more reliable than disks
 - Few hundreds GB loss per year on 80 PB tape repository
 - Few hundreds TB loss per year on 50 PB disk repository
 - No power required to preserve the data
 - Less physical volume required per (peta)byte
 - Inefficiency for small files issue resolved by recent developments
 - Nobody can delete hundreds of PB in minutes
- Bottom line: if not used for random access, tapes have a clear role in the architecture



- Turnaround time: ~2.6 years @ ~1.26GB/s for 100 PB
 - 10-12 drives (~10%) @ 90% efficiency
- Measured data loss: ~ 65GB lost over 69 tapes




- “Disk pool” reliability is higher than “raw disk” reliability because we increase it within the service (replication, error correction,)
- Tape reliability still $\sim O(1)$ higher than disk with error correction !
 - Note: single tape copy vs. redundant copies on disk



- Plain (reliability of the service = reliability of the hardware)




- Plain (reliability of the service = reliability of the hardware)
- Replication
 - Reliable, maximum performance, but heavy storage overhead
 - Example: 3 copies, 200% overhead

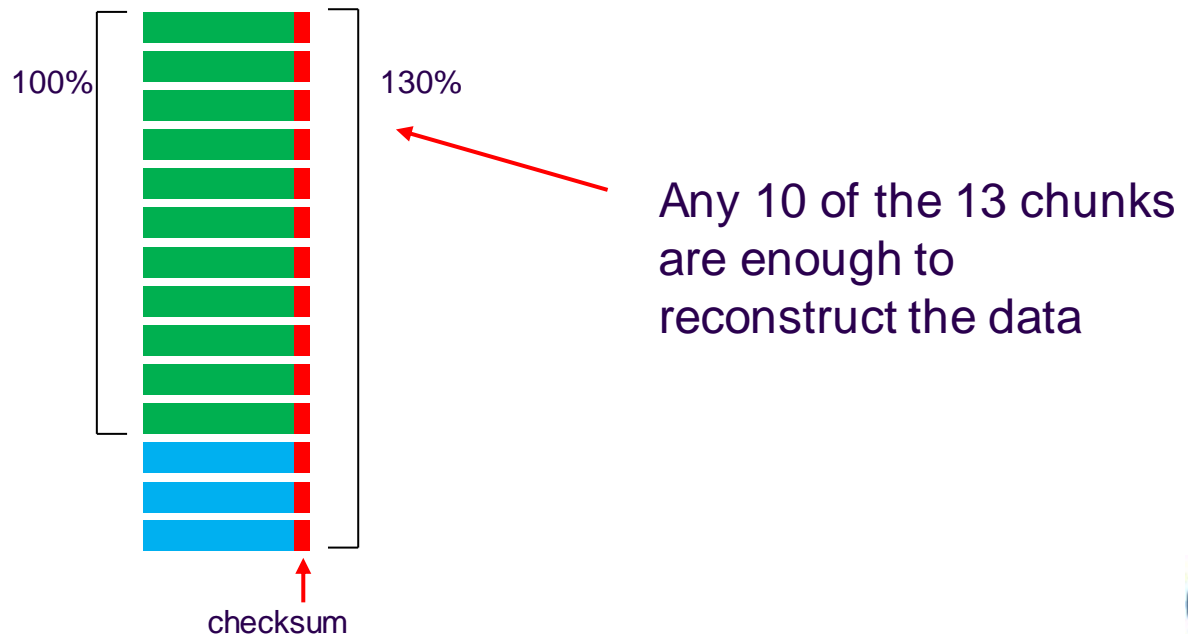


Any of the 3 copies is enough to reconstruct the data

- Double parity / Diagonal parity
 - Example 4+2, can lose any 2, remaining 4 are enough to reconstruct, only 50 % storage overhead





- Plain (reliability of the service = reliability of the hardware)
- Replication
 - Reliable, maximum performance, but heavy storage overhead
 - Example: 3 copies, 200% overhead
- Reed-Solomon, double, triple parity, NetRaid5, NetRaid6
 - Maximum reliability, minimum storage overhead
 - Example 10+3, can lose any 3, remaining 10 are enough to reconstruct, only 30 % storage overhead



- Plain (reliability of the service = reliability of the hardware)
- Replication
 - Reliable, maximum performance, but heavy storage overhead
 - Example: 3 copies, 200% overhead
- Reed-Solomon, double, triple parity, NetRaid5, NetRaid6
 - Maximum reliability, minimum storage overhead
 - Example 10+3, can lose any 3, remaining 10 are enough to reconstruct, only 30 % storage overhead
- Low Density Parity Check (LDPC) / Fountain Codes / Raptor Codes
 - Excellent performance, more storage overhead
 - Example: 8+6, can lose any 3, remaining 11 are enough to reconstruct, 75 % storage overhead (See next slide)

0 .. 7: original data

8 .. 13: data xor-ed following the arrows in the graph



Any 11 of the 14 chunks are enough to reconstruct the data using only XOR operations (very fast)

138% (min size required to reconstruct)

You are allowed to lose any 3 chunks (21 %)

- Plain (reliability of the service = reliability of the hardware)
- Replication
 - Reliable, maximum performance, but heavy storage overhead
 - Example: 3 copies, 200% overhead
- Reed-Solomon, double, triple parity, NetRaid5, NetRaid6
 - Maximum reliability, minimum storage overhead
 - Example 4+2, can lose any 2, remaining 4 are enough to reconstruct, 50 % storage overhead
 - Example 10+3, can lose any 3, remaining 10 are enough to reconstruct, only 30 % storage overhead
- Low Density Parity Check (LDPC) / Fountain Codes
 - Excellent performance, more storage overhead
 - Example: 8+6, can lose any 3, remaining 11 are enough to reconstruct, 75 % storage overhead
- In addition to
 - File checksums (available today)
 - Block-level checksums (available today)

- Reliability relates to the probability to loose data
 - However, you can have service interruptions and temporary unavailability of data without data loss
- Consequences
 - High data reliability does not imply high service availability
 - When hardware fails, the data on it becomes unavailable. It is a service failure
 - When experiencing service failure, an intervention must happen as fast as possible.
 - Requires a piquet with engineers on call
 - Draining and restore operations take lot of time, affecting also availability

DSS

Data & Storage Services


CERN IT
Department




Reserve slides on Availability




- We have “sets” of T independent storage
 - This example has $T=6$
- The storage pool is configured to replicate files R times, with $R < T$
 - This example: $R=3$ every file is written 3 times on 3 independent storage out of the 6 available
 - When a client read a file, any copy can be used
 - Load can be spread across the multiple servers to ensure high throughput (better than mirrored disks, and much better than Raid 5 or Raid 6)




- The loss of a storage component is detected. The storage component is disabled automatically
- File Read requests can continue if $R > 1$ (at least 1 replica), at reduced throughput
 - The example has $R=3$
- File Creation / Write requests can continue
 - New files will be written to the remaining $T - 1 = 6 - 1 = 5$ storage components
- File Delete request can continue
- File Write / Update requests can continue
 - Either by just modifying the remaining replicas or by creating on the fly the missing replica on another storage component
- Service operation continues despite hardware failure. (independent storage)





- The disabled faulty storage is not used anymore
- There is “Spare Storage” that can be used to replace faulty storage
 - manually or automatically
- The lost replicas are regenerated from the existing replicas
 - Manually or automatically



- To drain a server, just power it off
- Will be seen as faulty and disabled (it will not be used anymore)
- The available “Spare Storage” will be used to replace faulty storage
 - manually or automatically
- The lost replicas are regenerated from the existing replicas
 - Manually or automatically





Internet Services

- Ensure that there is enough spare storage to cope with:
 - Hardware failures
 - Planned replacement
- No need to intervene timely when Hw fails
 - **Asynchronous** interventions only
 - No engineers **on call or piquet** needed
- Create and configure data pools
- Arbitrary level of services: Monitor reliability, availability and performance and adapt the various parameters (storage size, T, R, ...) to optimize the operational experience and meet the service level agreement
- Identify what responses can be automated and what needs to be handled manually



- Beware of the approach “all data in one pool”
 - Offers only one, common, quality of service
- It is important above a certain volume of data, to be able to optimize the cost to the real requirements, and go beyond “generic” solutions.