SUT TREE

ALICE

A JOURNEY OF DISCOVERY

# Asian sites in ALICE – storage role in the context of Run 3

**4-th Asia Tier Center Forum**
**Bangkok, Thailand**

21 November 2018

Latchezar Betev

# ALICE today

- Last data taking period of RUN2: Pb-Pb
- Expected to collect 9PB of data – 3x more than any other Pb-Pb period prior
- Possible due to a partial upgrade of the readout
- Will be followed up by data processing for

~4-5 months

# Data taking projection

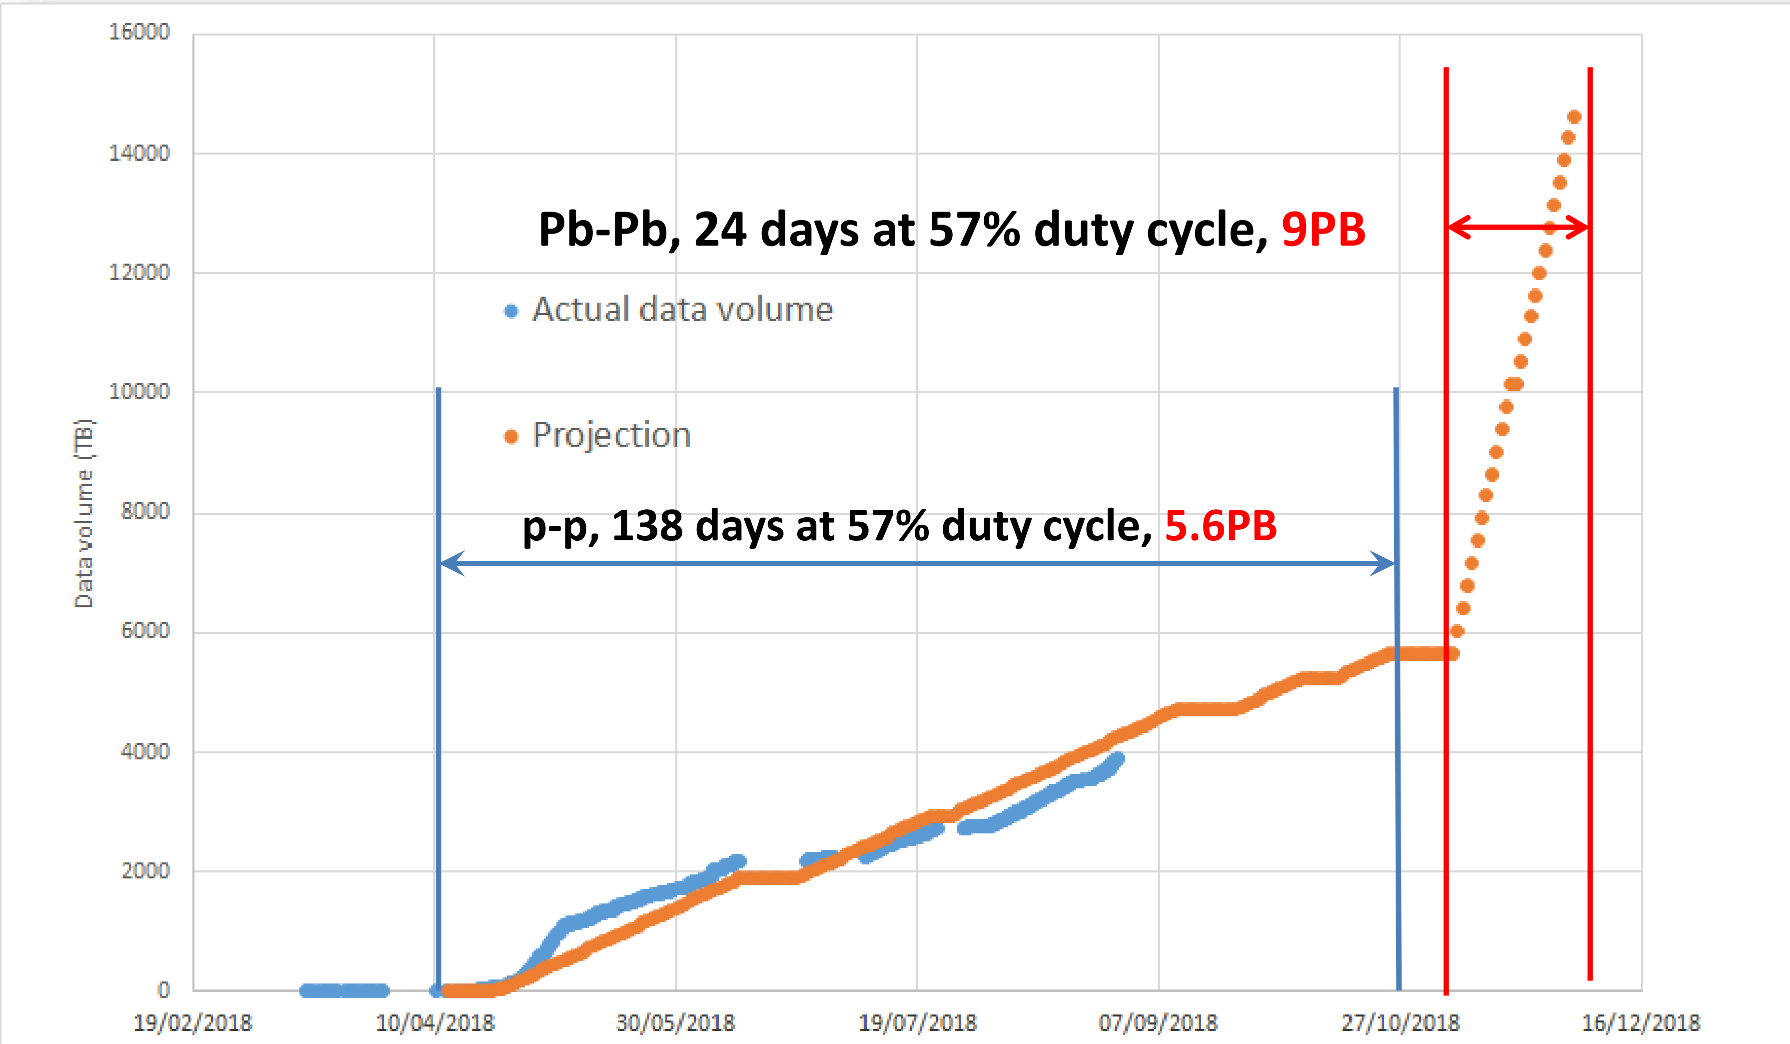

# What follows?

- Intensive upgrade period – detector, readout and software
- Run3 – ALICE will run with upgraded everything and new physics programme
  - Focused on charm physics
  - Substantially higher (x50) data rates from detectors
  - Compression facility O2 to bring the data back into a manageable volume for offline storing/processing
  - A 20% growth in computing resources + the O2 facility will be enough to process the Run3+ data

# O2 in a nutshell

- Custom-build cluster (CPU/GPU/FPGA) for efficient data compression

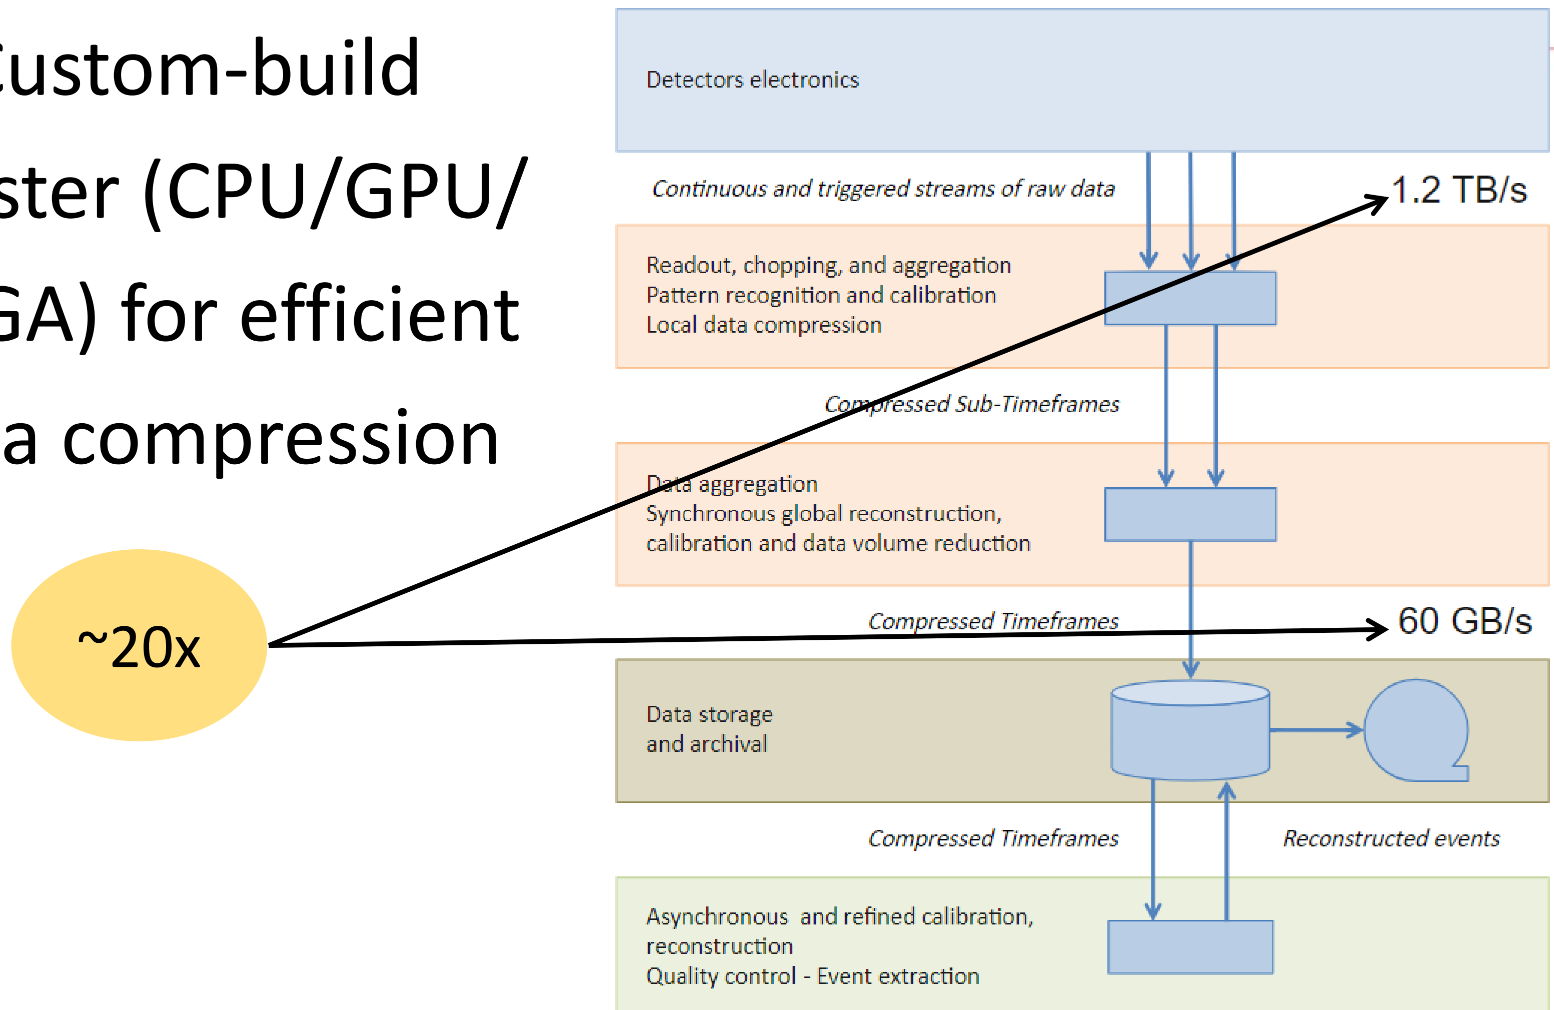

# What changes in the computing model

- High reliance on the synchronous data compression – without it, the numbers will not match
- The role of the Grid computing remains largely familiar with some changes to
  - Additional analysis facility as a separate entity
  - More load on storage and redefined function (details to follow)

# The upgrade challenge

- Order of magnitude more storage and files to manage
- Higher client-server (tasks) interaction rates
  - Estimated to be ~200 kHz (from about 20kHz today)
  - Affects the load on existing services – not expected to scale
- O2 facility addition to the system will bring new resources
  - Integration is a challenging problem
- Preserve the data model as much as possible
  - Storage provisioning and management is one of the toughest, most expensive and most time-consuming items on the Grid

# The file catalogue

- All files on the Grid are annotated in the catalogue
  - Logical File Names (LFN) namespace with ACLs

-rwxr-xr-x alidaq alidaq 264403565 Sep 09 22:10 /alice/data/2016/LHC16n/000261088/raw/16000261088034.205.root

  - Pointer to file location and GUIDs for SE filenames

root://alice-tape-se.gridka.de:1094//10/33903/76cebd12-76a0-11e6-9717-0d38a10abeef

root://voalice10.cern.ch//castor/cern.ch/.../16000261088034.205.root

- 6 billion LFNs, 3.5 billion physical files (archives of multiple logically combined LFNs)
- Running on MySQL (master-slave) DB with cache
  – Somewhat outdated technology by today's standards
- For 150k running jobs – 20KHz read / 1KHz change/delete

# The file catalogue (2)

- Will continue (together with the common task queue) to be the backbone of the ALICE distributed computing
- The basic functionality will remain unchanged
- The backend will be changed
  - To be able to cope with increased load

# The new backend

- Apple-ALICE collaborative work on Cassandra
- Answers all general and specific requirements:
  - Horizontal scaling, no single point of failure
  - High query rate, high availability
  - Consistency, easy setup
  - Drawback – SQL to NoSQL operations

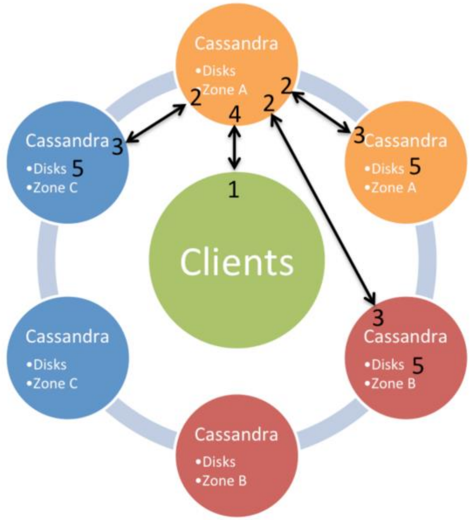

# Setup and benchmarking

- 6-node test ring, replication factor 3, ~cheap hardware
- Any node (+1) can go away without performance degradation

**=> Success!**

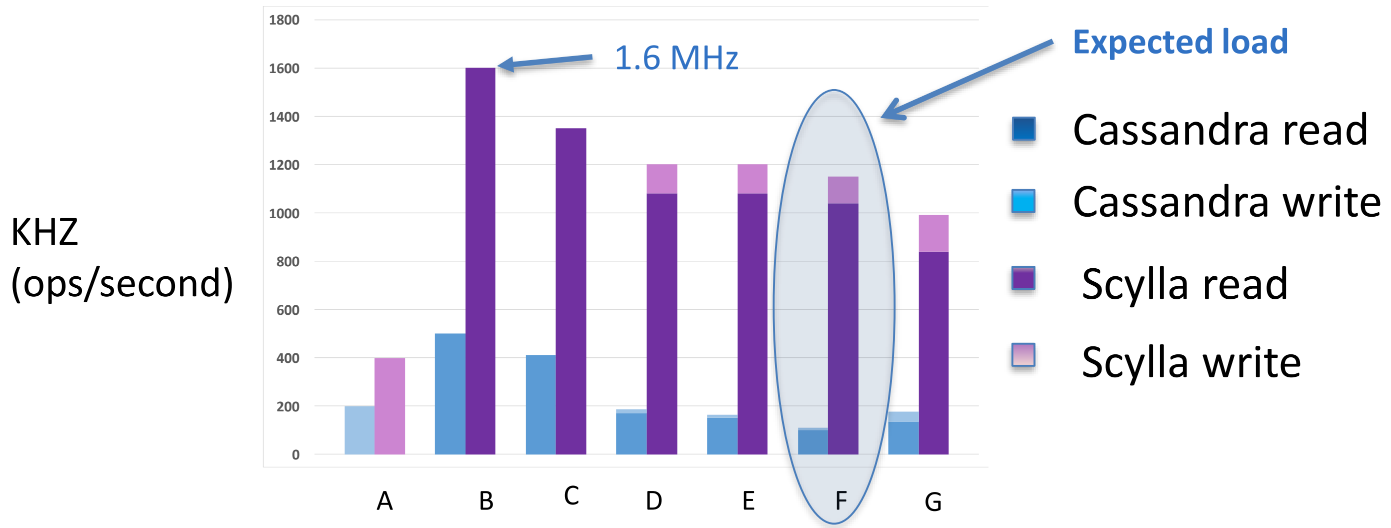


| | Benchmark (default cassandra-stress) | Cassandra | Scylla | Diff |
|---|---|---|---|---|
| A | Insert only | 18 %util, 2% iowait | 100 %util | x2 |
| B | Read-only Gauss(5B,2.5B,10K) | Disk idle, 50% cpu | 100 %cpu | x3 |
| C | Read-only Gauss(5B,2.5B,1M) | 11 %util, 40% cpu | 11 %util, 100 %cpu | x3.28 |
| D | Mixed (10r,1w) Gauss(5B,2.5B,100K) | 2 %util, 45% cpu | 10 %util, 100 %cpu | x5.8 |
| E | Mixed (10r,1w) Gauss(5B,2.5B,1M) | 5 %util, 50% cpu | 16 %util, 100 %cpu | x6.4 |
| **F** | **Mixed (10r,1w) Gauss(5B,2.5B,10M)** | **9% util, 45% cpu** | **40 %util, 100 %cpu** | **x6.1** |
| G | Mixed 2K thrd. read, 200 write, G(5B,2.5B,100K) | 8 %util, 40% cpu, no iowait | 26 %util, 100 %cpu | x5.62 |

# Synchronous processing data path

- Simplified schema during data taking periods

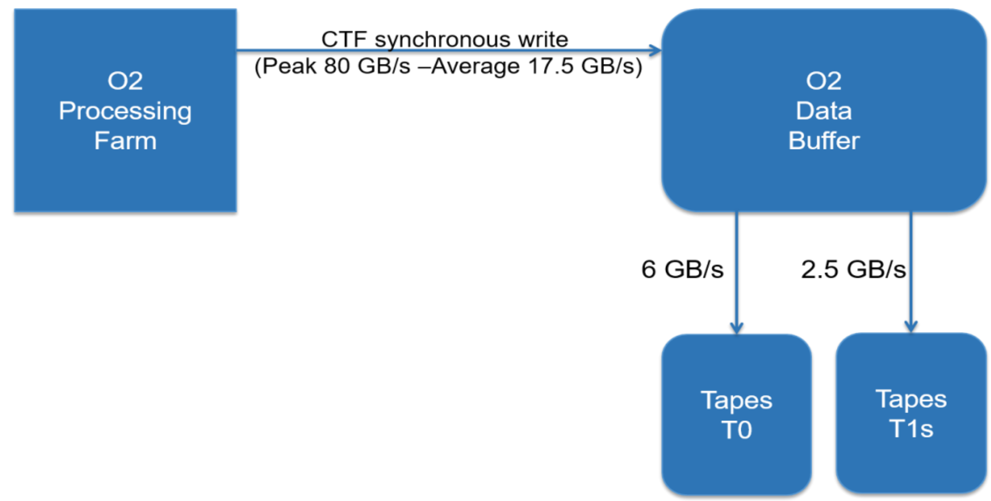


- Defining element for the O2 data buffer
  - Data rates from O2
  - Number of events => buffer size
  - Data rates to T0 tapes and T1s (export)

# Asynchronous processing

- During beam off periods and after data taking at O2 facility
- All the time at T0/T1s
- Asynchronous processing is similar to synchronous, but with better detector calibration/software
- Output from asynchronous processing is AODs

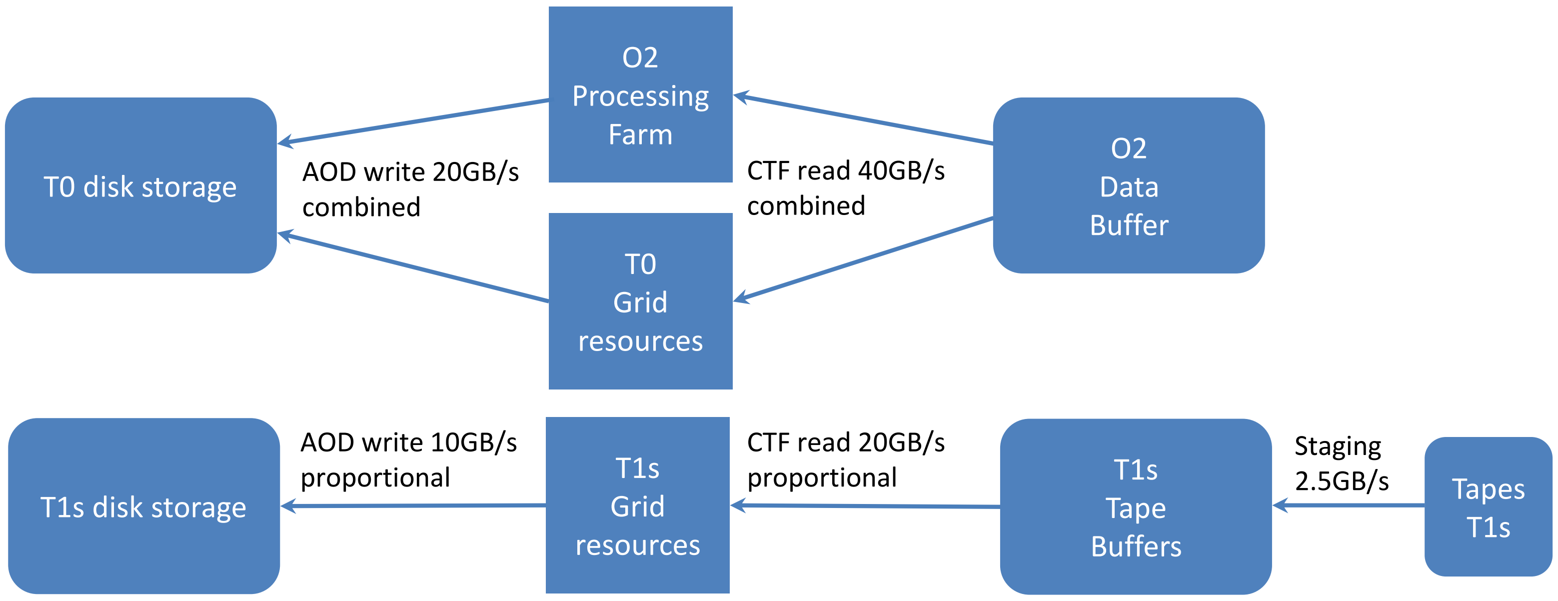

# Data rates

- Final (projected) data rates in Run 3

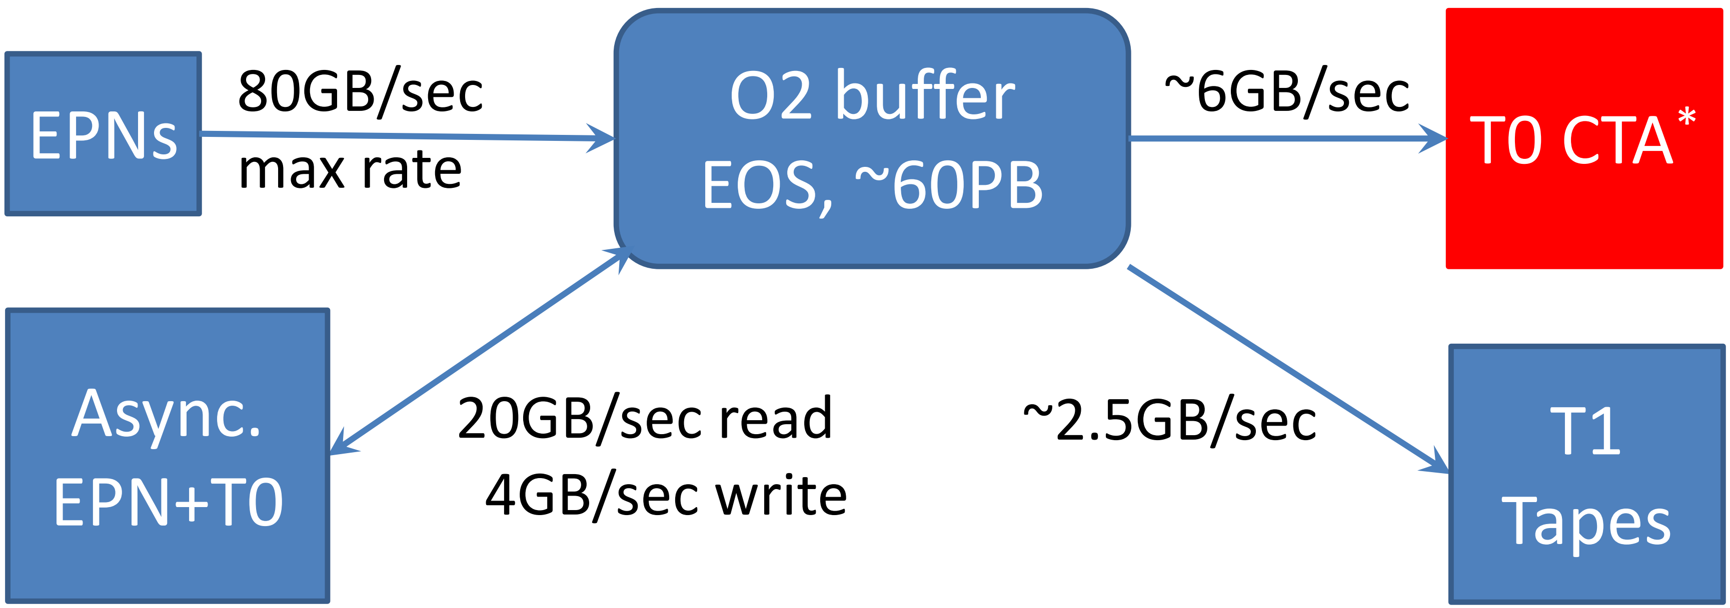


* CTA – CERN Tape Archive

# Data handling on global level

- Replication across sites

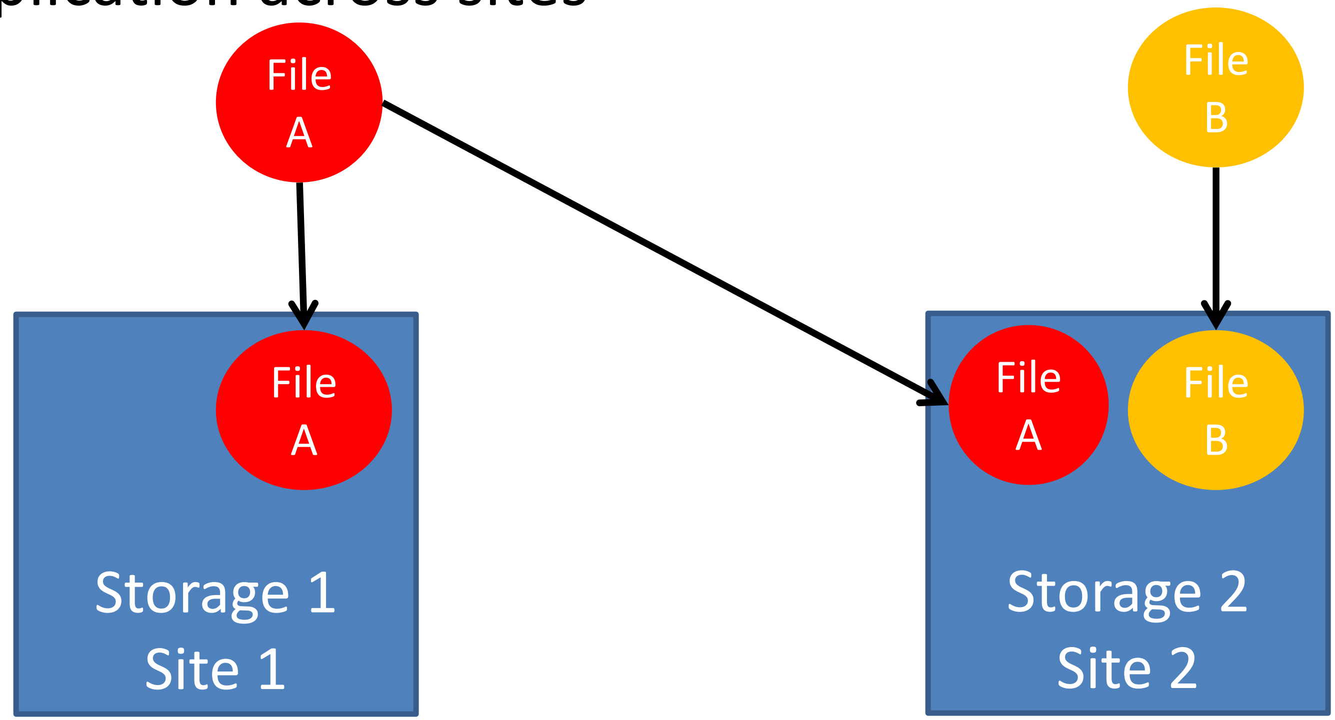


- Two (or more) copies – safe, but costly
- One copy of bulk of the data is now a norm
  - With the exception of RAW (2x), smaller analysis objects, frequently accessed and user files

# Data handling on site level

- Most common – fileservers with some RAID level
  - Usually large (hundred TB) filesystems per server
  - Classical technology, supposedly safe
  - All redundancy within same filesystem

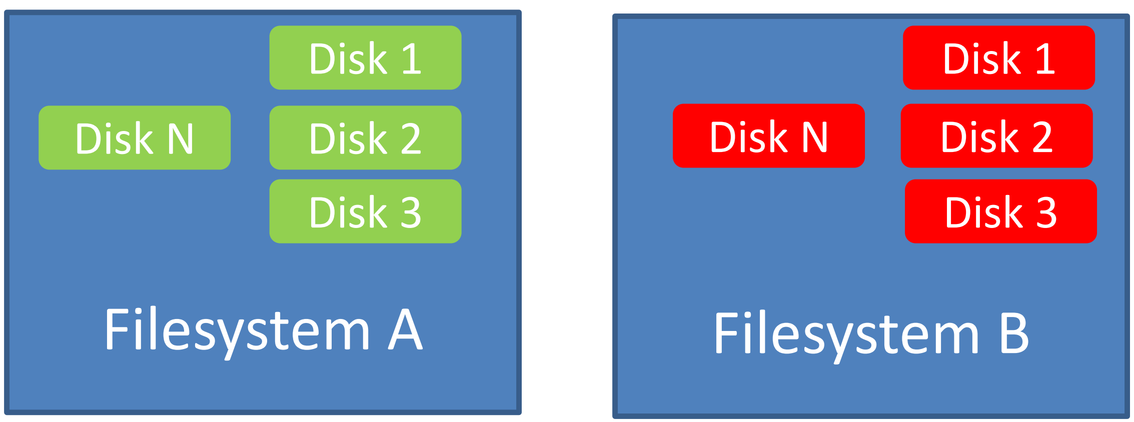

# Data loss recovery on global level

- Re-replication (or re-generation) of lost data

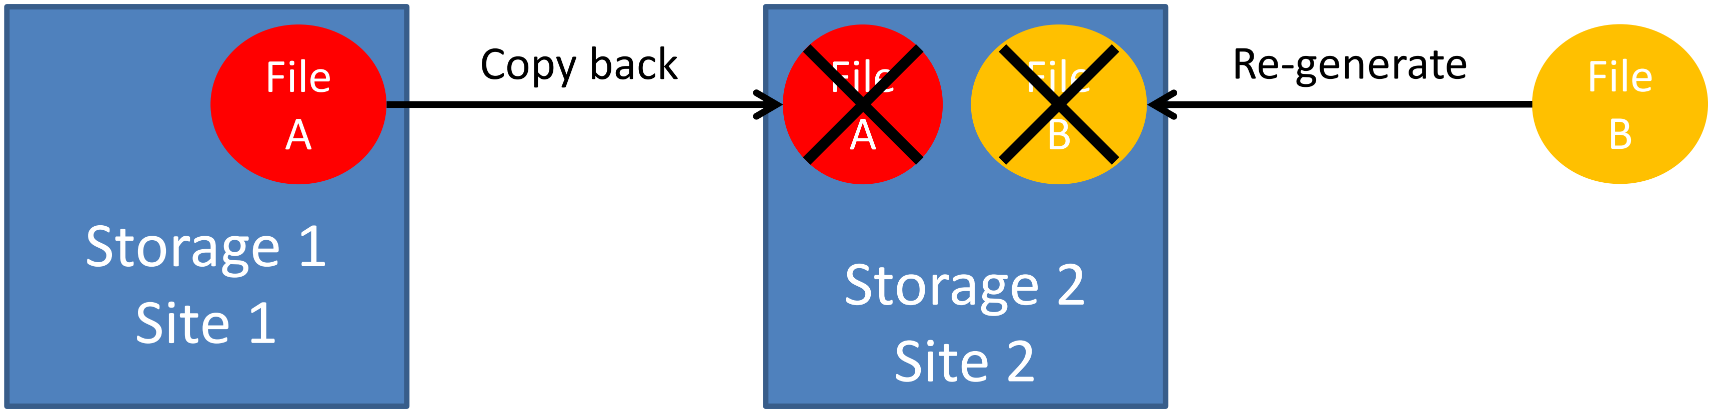


- Both methods external to site and used in case of major or locally unrecoverable data loss

# Data loss recovery on site level

- RAID recovery

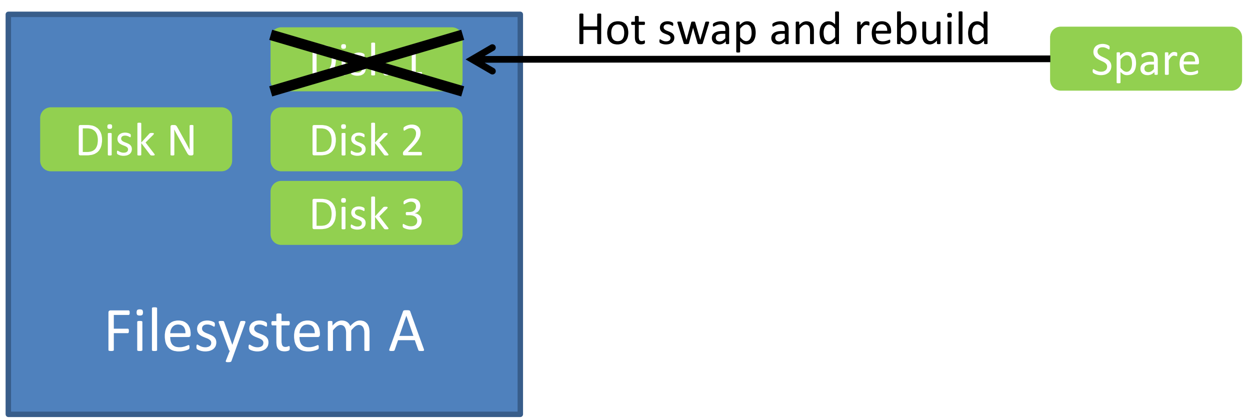


- Long (days!) process with large-volume disks
- Assumes broken disks are less than RAID level – may be worst with power cuts
- If RAID rebuild fails – full filesystem loss
  - Recovery mechanism on previous slide

# Data loss recovery on site level (2)

- The recovery of RAID-ed filesystem is heavy and can potentially fail
  - Resulting in major data loss (100s TBs)
  - Source of embarrassment and long recovery procedure, out of site control
- RAIDs are typically expensive… and there is a better solution
  - Coincidentally, it simplifies the setup and operation

# Data handling on site level

- (not) New system – RAIN
  - Data is stripped/replicated across physical nodes
  - RAID-x – like
  - Most efficient is RAIN-1, other levels require substantial network traffic between servers
  - RAIN-1 configuration is almost as safe as out of site copy

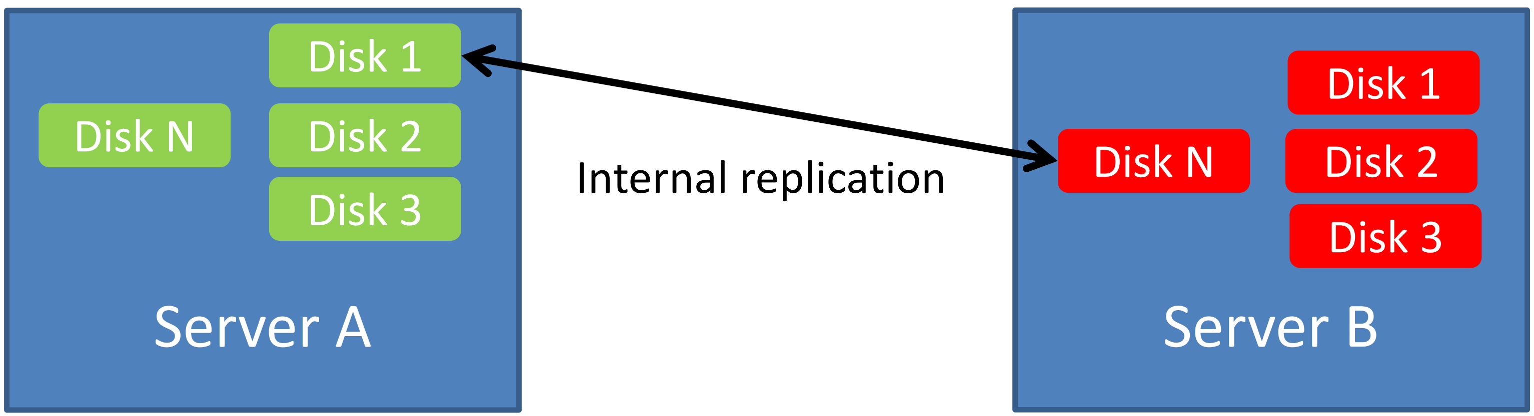

# Summary on data handling

- Sites provide disk storage capacity, usually secured with RAID (think ‘I don’t want to lose data’)
  - In  a typical (RAID 6) configuration, the capacity overhead is 15%, the cost of RAID controller is on top
  - Provides some security, but in case of failure there is a massive data loss
- Experiments mostly disregard site redundancy (think ‘Data loss is inevitable’)
  - Precious data is replicated
  - Data is spread around, single site storage failure results only in ‘statistics’ loss
  - Data loss is unpleasant, but tolerated and there are recovery mechanisms in place

# Additional considerations - ALICE

- To stay within ‘flat budget’ on storage ALICE will
  - Have a single copy of RAW data on custodial storage
  - Give up one derived data format (ESD) and remain with a single data format (AOD) for all subsequent data analysis
- Even with the above - at the limit of capacity
  - Provided the growth will not be reduced by unforeseen circumstances

# Minimize impact of data loss and tune performance

- Site storage structure

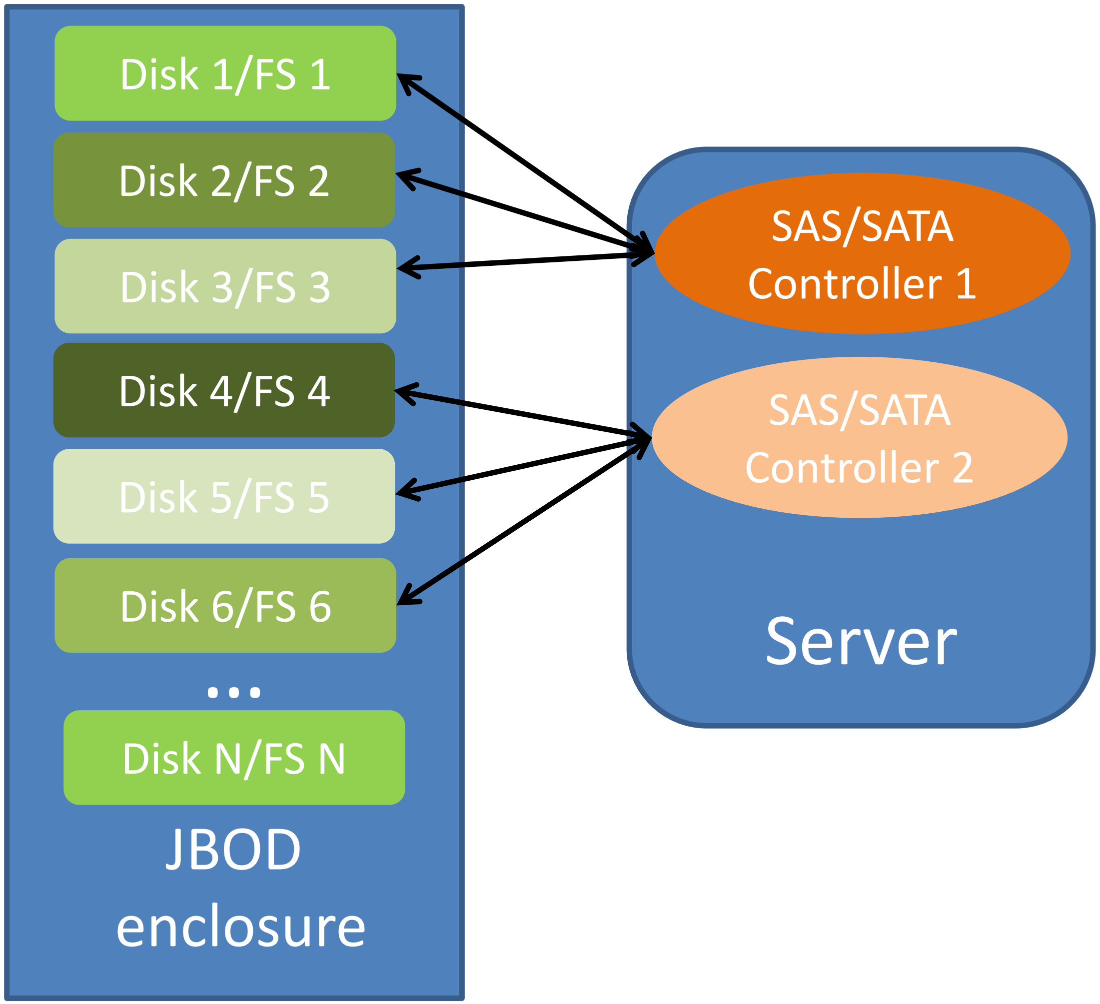


- Each disk is a filesystem
- Max loss of data is one disk
- Number of disk combined under a controller defines the throughput

- Basic configuration – no data redundancy within the storage element

- Software configuration – can define 2 or more copies within the same server for data redundancy

# Technical solutions

- Setup of EOS storage at CERN – RAIN-1 for the bulk data
  - Even higher redundancy – directory based
- Setup of EOS storage at Kolkata – RAIN-6 for the bulk data
  - Works, but is not trivial to update
- Setup of EOS storage at ORNL – no redundancy for the bulk data
  - Will serve as a test for the future operation and policy
  - Cheapest storage – RAW=usefull

# Technical solutions (2)

- The replication policy will remain the same
  - Most important aspect is the wide distribution of any bulk data – avoid concentration in the same storage
  - Extremely important to limit impact of unavailable (or lost) storage capacity
- Foreseen bulk data with 2 copies will be the AODs
  - Assures continuous availability and quick recovery in case of loss
  - Any other data will 'suffer' statistical reduction in case of storage unavailability, potentially can be regenerated

# Operational considerations

- Reduced local redundancy is taken into account by experiment replication policy
- Local data loss will be limited to single disk capacity
  - Site cannot recuperate data (understood and OK)
  - Should be reported to the experiment quickly to assure proper mitigation (at the discretion of the experiment)
  - Should simplify the local storage operation – replace faulty disk and continue

# Operational considerations (2)

- The sites can determine storage performance, as needed by the function
  - Tuned for max throughput per server/network infrastructure – vital for analysis
  - … or high capacity with low throughput – custodial storage/tape replacement (?)

# Storage distribution across Asian sites

- Displayed current capacities (disk only)
- Expected to grow by 20% per year + other sites to join

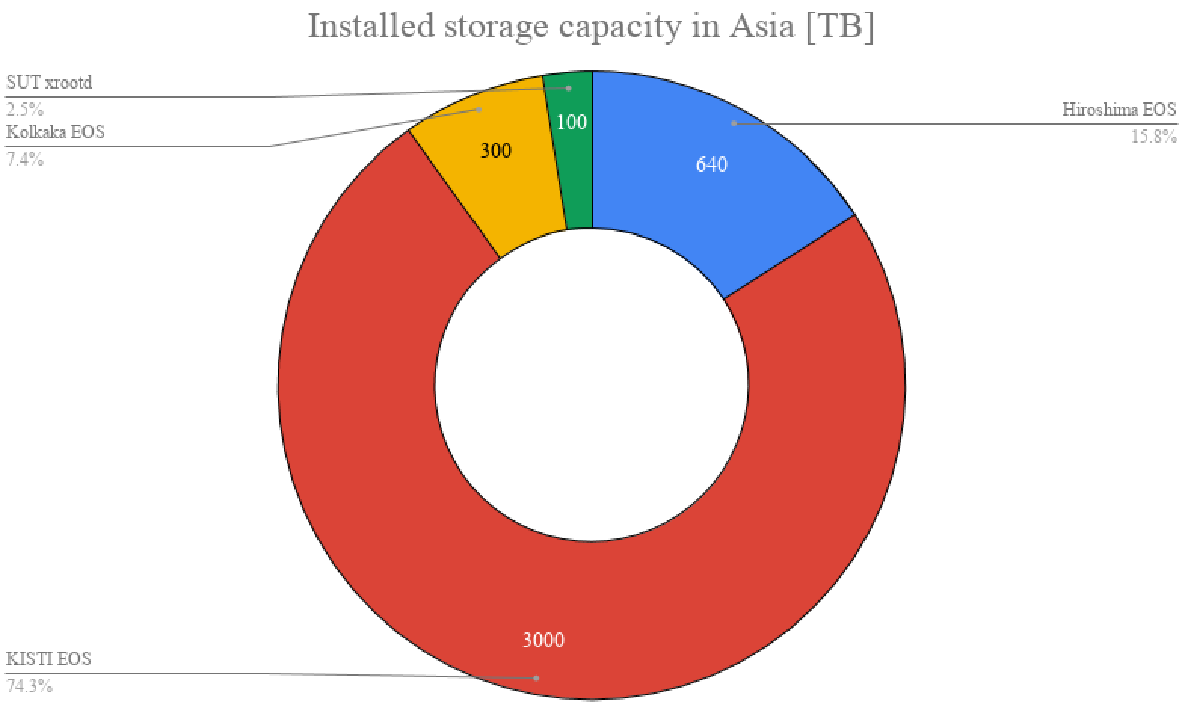

# Future cooperation

- Substantial storage capacity, expected to grow
- Possible areas of cooperation:
  - Selection of common hardware, synchronised purchase (economy of scale)
  - Installation procedure and sharing of experience, perhaps common experts pool
- Single entry point between storage elements with similar/same setup

# Summary

- Substantial storage capacity, expected to grow
- Possible areas of cooperation:
  - Selection of common hardware, synchronised purchase (economy of scale)
  - Installation procedure and sharing of experience, perhaps common experts pool
- Single entry point between storage elements with similar/same setup

**Many thanks to the SUT team for organizing the workshop**

**As with the previous installments – it provides an opportunity for a focused discussion of a relevant topic**

**Looking forward to your questions and discussion**