Data Challenge 24

NeIC NT1 Manager
Mattias Wadenstein
<maswan@ndgf.org>

2023-11-14
NDGF All Hands fall 2023
Linköping, Sweden

neic
NORDIC E-INFRASTRUCTURE COLLABORATION

Overview

- Transfer plans
- Technology tests
- Other tests
- Conclusions

SPEAKER | Mattias Wadenstein <maswan@ndgf.org>

Plans

- 25% of the realistic HL-LHC data flow for T0→T1 export and some T1→T2 workflows

- February 12th to February 24th (2 weeks)

- T0→T1: 16Gbit/s ATLAS and 4Gbit/s ALICE
  - Most traffic for both expertiments in in other workflows

- ATLAS will inject some extra traffic on top of regular workflow
  - Maybe they'll be noticeable in our monitoring in the case of "flexible" scenario which adds 2x as much

- ALICE will most likely stick with production export of 2023 HI raw data to T1 tape during this period

Tech: Token Auth

- The disk transfers should be done with tokens for auth
- Make sure dCache is upgraded and configured
- Do we need hostcerts on all pools?

Tech: Site network monitoring

- Publish json with max and current usage of external links to/from the site
    - We got a ticket on this some months back
- Needs implementation from NORDUNet and Slovenia
    - Status updates?
- WLCG pushes fairly hard for this to be in place for DC24

Tech: Flow labeling

- Networking people wants flows marked with experiment and activity ("ATLAS, reconstruction")
- UDP fireflies and IPv6 packet header marking
- Support from dCache as of 9.2
- Needs flowd (CERN-developed package) on the pool nodes running as root to do EPF magic
- If we also collect flowd data we could see:
  - (src,dst) pairs, size, speed, etc
  - TCP stats: retransmits, TCP algorithm, etc

Tech: Archive metadata to tape

- Propagated via FTS to dCache
- In order to write data as it is likely to be read
- Needs dCache propagation from door to DB to pool
- ENDIT plugin needs support
- ENDIT daemon needs support and advanced logic
- For future, not DC24 timeline

```
File: `data23_13p6TeV.00452799.physics_Main.daq.RAW._lb0777._SFO-19._0001.data`

```json
archive_metadata = {
    "scheduling_hints": {
        "archive_priority": "100"              # highest priority
    },
    "collocation_hints": {
        "0": "data23_13p6TeV",                                 # project
        "1": "RAW",                                            # datatype
        "2": "00452799",                                       # runnumber
        "3": "data23_13p6TeV.00452799.physics_Main.daq.RAW",   # dataset
    },
    "optional_hints": {
        "activity": "T0 Tape",          # Tier-0/DAQ
        "3": {                          # dataset level
            "length": "19123",                 # total number of files at specified level
            "bytes": "80020799318456"          # total size of files at specified level
        }
    }
}
```
```

Expectations

- The data flows are not expected to be challenging for us
  - Might bottleneck CERN links in some of the stretch goals

- The tech challenges are challenging
  - dCache upgrades and configuration changes
  - Pool requirements updates?
    - IGTF hostcert
    - Flowd
  - Site border network monitoring?
  - When do we want to do interesting tape stuff?

Questions?

neic