

Data Challenge 2024

--

European Site Feedback

Overall Perception

- Useful and necessary exercise
- Feedback ranging from
 - “stress test greatly impacting sites and overloading storage endpoints”
 - to
 - “no issues expected – only minimal issues encountered”
- Material
 - Slides provided by sites
 - DRAFT WLCG/DOMA Data Challenge 2024 - Final Report
 - WLCG/DOMA DC24 Retrospective presentations

DC24 at DESY

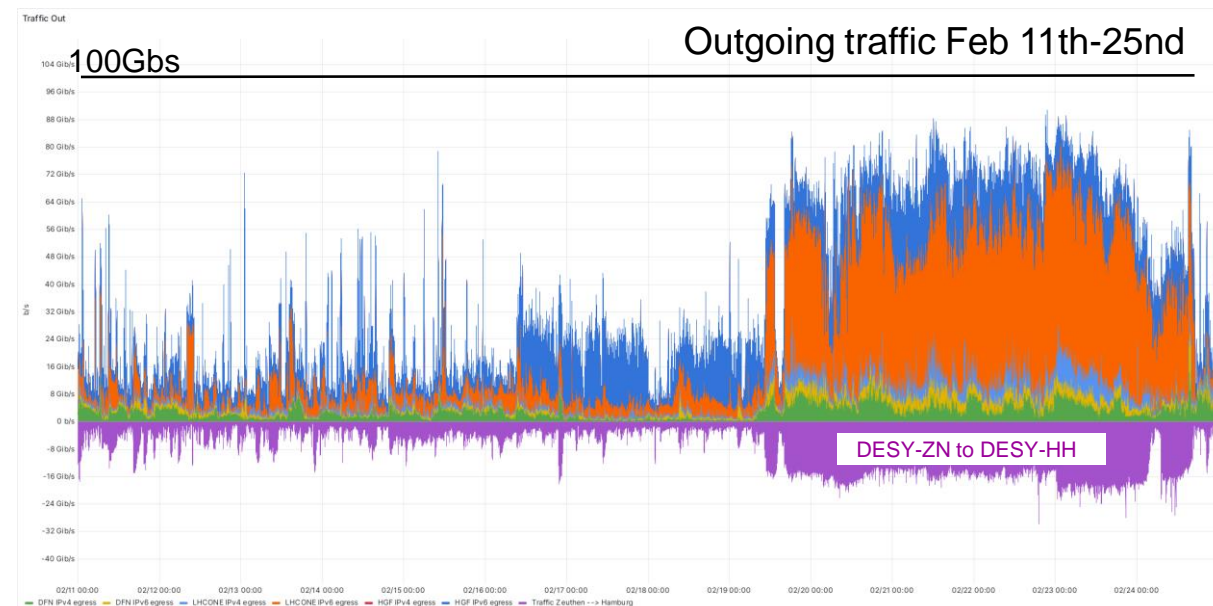

TLDR: No Problems, but WAN well utilized

DESY has two sites serving as WLCG Tier-2 and Belle-II RAW data center; SDMC for CTA and (national) Analysis Facility for various experiments & communities

- Site **DESY-HH**: ATLAS, Belle-II, CMS, LHCb, ILC, ILDG, local experiments, photon science
- Site **DESY-ZN**: ATLAS, IceCube, CTA, ILDG, astro-particle experiments


DC24

- No issues for storage systems
- WAN link saturated – needs upgrade for future DCs and HL-LHC running
- LHC1 traffic to DESY-ZN is routed through DESY-HH
 - Presently not captured by WLCG WAN monitoring
 - DESY-ZN and DESY-HH separated on layer-3 wrt WLCG monitoring, both share the same layer-2 link



INFN-T1


- DC24 very useful to identify bottlenecks – very much in favour of running such exercises + preparatory tests and repetitions of the tests
- However, DC is a stress test greatly impacting sites and overloading storage endpoints
 - SAM tests should be aware of ongoing DC.
 - Significant production load during the challenge – sometimes heavy impact
- Wish for FTS to regulate injection based on success rate to avoid overloading storage endpoints
- Differences between FTS and site monitoring observed



- Planning to align the StoRM WebDAV instances dedicated to CMS (observer higher failure rates on servers with less cores)
- Re-think LHCb hardware configuration so to accommodate their workflow
- Recently upgraded StoRM WebDAV in all the endpoints to improve efficiency + planning to introduce performance markers in StoRM WebDAV

IN2P3-CC


- 200Gbit/s each for LHCOPN, LHCONE
- ALICE, ATLAS, CMS, LHCb and Belle II
- No issues identified in network infrastructure and on storage services.
- ALICE, CMS and Belle II have reported no issues with the IN2P3 site, and targets were achieved.
- LHCb
 - Usage of Tape Rest API was successfully tested
 - A buffer overflow occurred during the LHCb staging. Increasing its size solved the issue.
- ATLAS
 - Due to an overload of storage connections, the ATLAS hammer cloud mechanism put the site on test which stopped ATLAS production on site. To recover it, ATLAS had to reduce the number of connections on FTS which in turn reduced the throughput rates.
 - In the table below, flexible (T0 ->Tier1) and flexible (T0+T1->T1) phases highlight bad results for all Tier1s but notably for IN2P3 in comparison with other T1



| T1 Site | Minimal (T0→T1) | | | Flexible (T0→T1) | | | Flexible (T0+T1→T1) | | |
|-------------|-----------------|---------|-----|------------------|---------|-----|---------------------|---------|-----|
| | model | reality | [%] | model | reality | [%] | model | reality | [%] |
| BNL-ATLAS | 60.0 | 25.9 | 43 | 68.4 | 21.2 | 31 | 82.1 | 57.1 | 70 |
| FZK-LCG2 | 32.0 | 34.1 | 107 | 39.0 | 13.2 | 34 | 59.4 | 43.2 | 73 |
| IN2P3-CC | 38.0 | 36.4 | 96 | 44.2 | 1.4 | 3 | 59.1 | 21.4 | 36 |
| INFN-T1 | 23.0 | 22.0 | 96 | 28.3 | 8.9 | 31 | 39.4 | 47.6 | 121 |
| NDGF-T1 | 45.0 | 0.7 | 5 | 24.4 | 0.0 | 0 | 52.2 | 0.0 | 0 |
| SARA-MATRIX | 15.0 | 17.9 | 119 | 19.3 | 32.8 | 170 | 36.2 | 84.6 | 234 |
| pic | 11.0 | 13.8 | 126 | 13.3 | 4.2 | 32 | 18.1 | 35.7 | 198 |
| RAL-LCG2 | 38.0 | 12.5 | 33 | 44.4 | 29.7 | 67 | 56.9 | 48.4 | 85 |
| TRIUMF-LCG2 | 25.0 | 26.0 | 104 | 29.3 | 12.5 | 43 | 38.6 | 54.0 | 140 |
| Σ (no NDGF) | 242.0 | 188.6 | 78 | 286.3 | 123.9 | 43 | 389.8 | 392.0 | 101 |

Data Challenge 2024 @ GridKa

- Additional temporary 100Gbit/s link to CERN for DC24
 - 300Gbit/s to LHCOPN + 200Gbit/s LHCONE
 - Max LHCOPN throughput observed **before** start of DC24
- ALICE, ATLAS, Belle II, CMS, LHCb
- No bottlenecks expected for GridKa → no bottlenecks observed!
 - One small dCache tuning issue for ATLAS&CMS





NDGF

- Minor dCache Bug & HAProxy config broke 3rd party copies under load during DC period
- Some network bottlenecks observed between CERN and NDGF
 - Around expect HL-LHC data rates or above
- Later DC test went smoothly
 - No problems for either networking nor storage at NDGF


PIC

- ATLAS (T1+T2), CMS, LHCb and DUNE
- No issues were detected for network infrastructure, storage services
- NOTED successfully active at PIC
- Impressive DUNE data streaming from PIC to European sites
- Requested NREN to separate LHCOPN and LHCONE paths



RAL

- DC24 was very successful
- By end of DC24, throughput was 2 – 3 times higher than ever achieved before
- LHCOPN link cut under the sea, fixed by February 16
 - Failover to LHCONE needed tuning
- Tape services during DC24 meet expectations. Post DC24 tape tests for ATLAS helped to expose areas for improvement.
- Hit 25Gbit/s limitations of gateway hardware
 - 100Gbit/s capable hardware has been purchased
- Lots of experience gained, trying to balance the load across the gateways



SARA

- Minor dCache (tuning) related issues for LHCb
- ATLAS: 21 GB/s for 5 days in a row
 - Saturated internal 200Gbit/s uplink
 - Affected LHCb transfers to RAL
- Post-DC tests with ATLAS and Nikhef
 - **temporary 1200 Gbit/s** connection with Nokia equipment
 - transfers from the Atlas EOS to the SARA dCache up to 45 GB/s, while doing normal production at the same time.
 - **1h 530 Gbit/s, with a peak of 661 Gbit/s**

Summary

- No major issues observed at sites
- DC24 was very useful to flush out bottlenecks → many lessons learned
- Real life intervenes → opportunity to test/tune failover mechanisms
- European T1s achieved targeted rates (with very few exceptions)