



WLCG data lake prototype for HL-LHC

Gavin McCance, Ian Bird, Jaroslava Schovancova, Maria Girone, Simone Campana, Xavier Espinal Currul, Ivan Kadochnikov

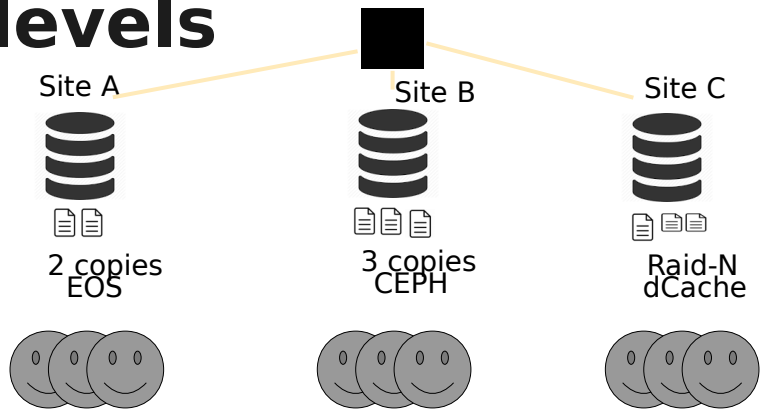
Motivation

HL-LHC demands > expected storage growth
We need a new paradigm to reduce cost



Economy of scale

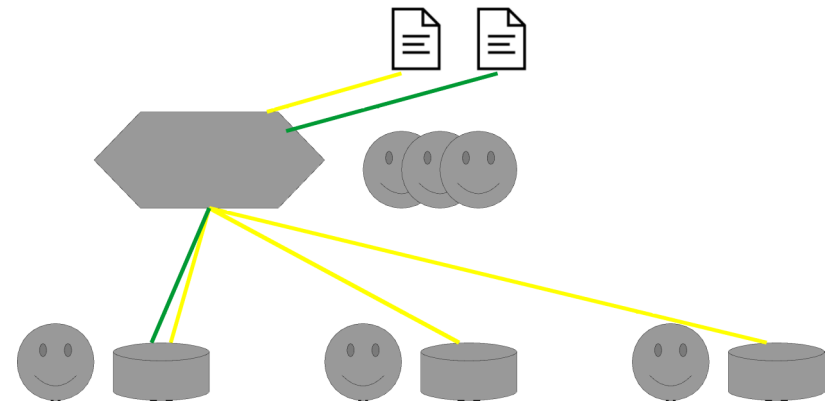
Replication on both levels



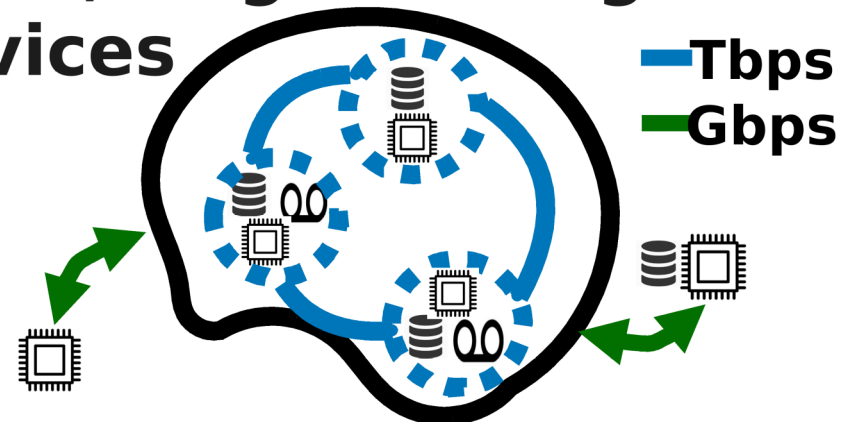
High operation cost

Storage and processors not always co-located

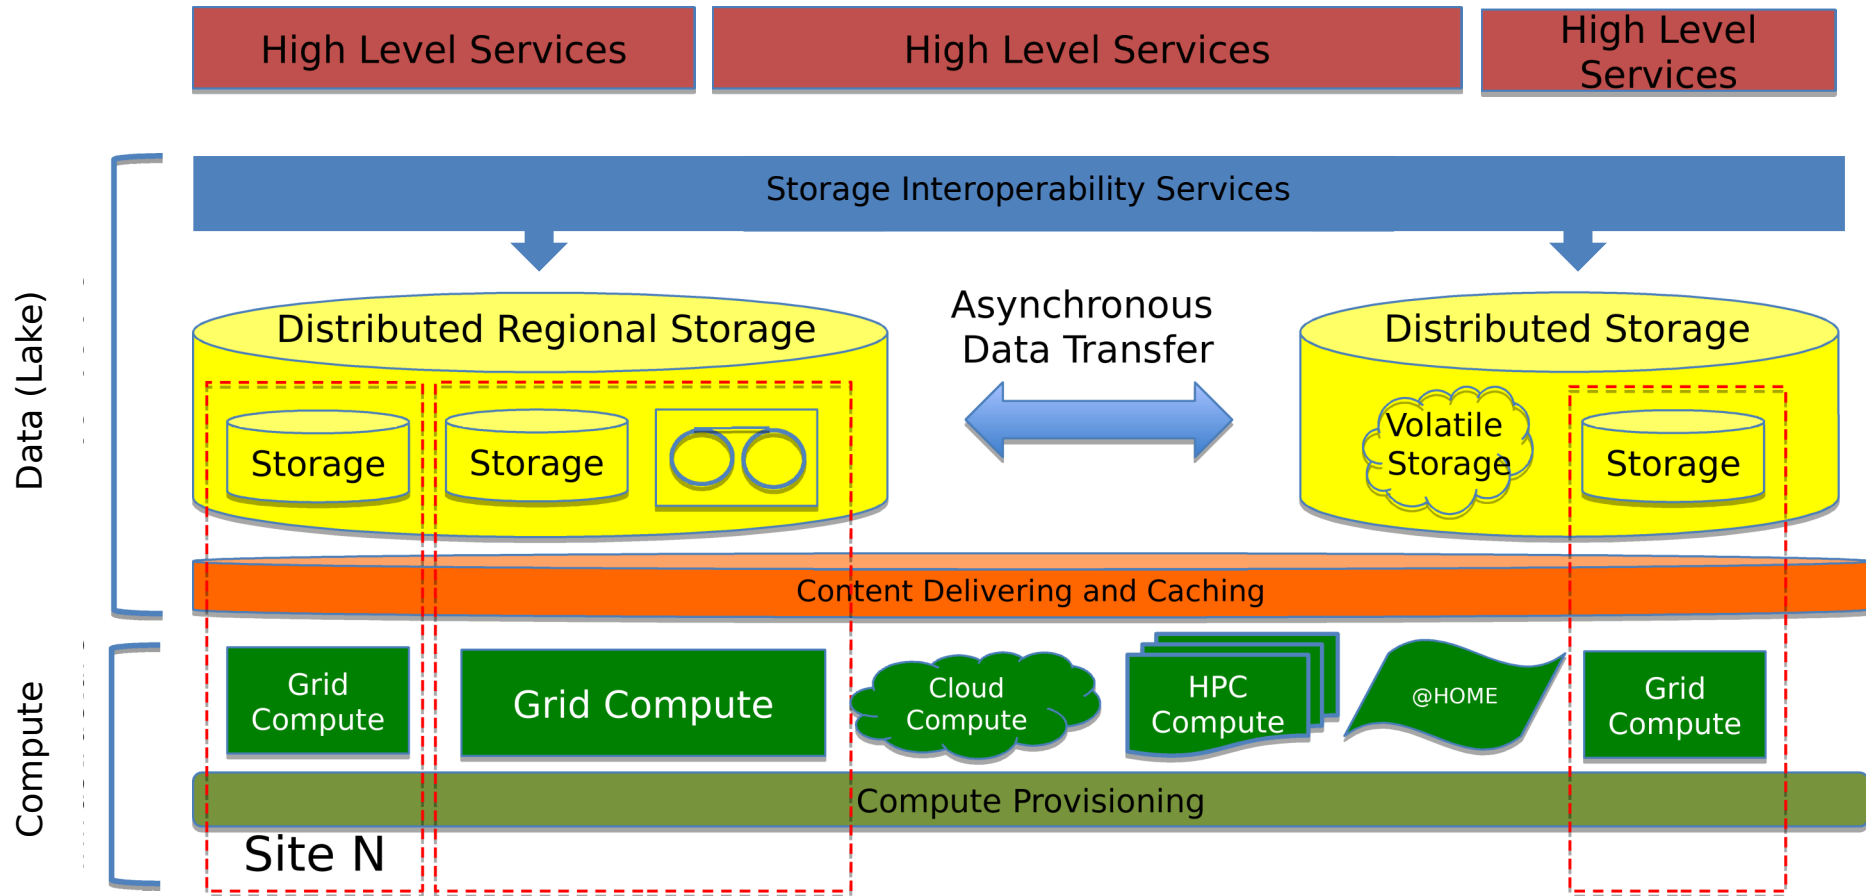
Managed QoS




Fewer, larger storage services



Data and compute infrastructure



The Eulake prototype




Steps done

**Deployment and
comissioning**


Transfer tests

Data replication


Data access tests



eulake I/O (MB/s)




writing into eulake (MB/s)



Quality of service management

2 copies, 1 copy, striped: 3+2 chunks

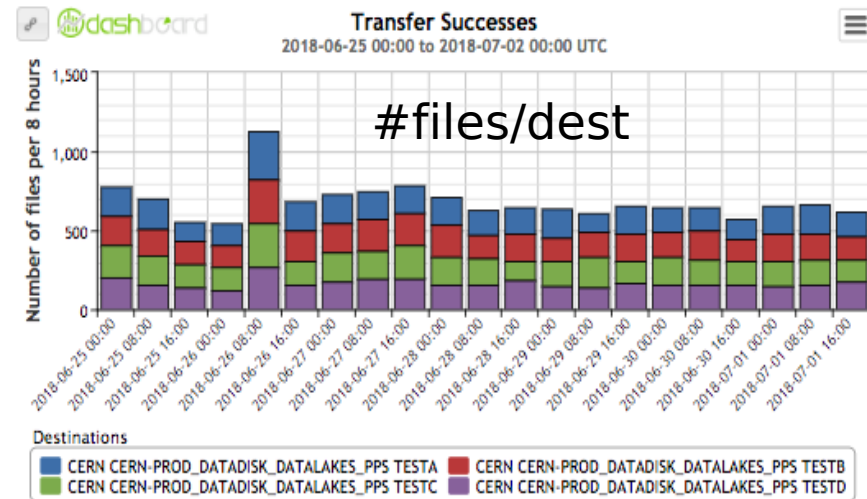
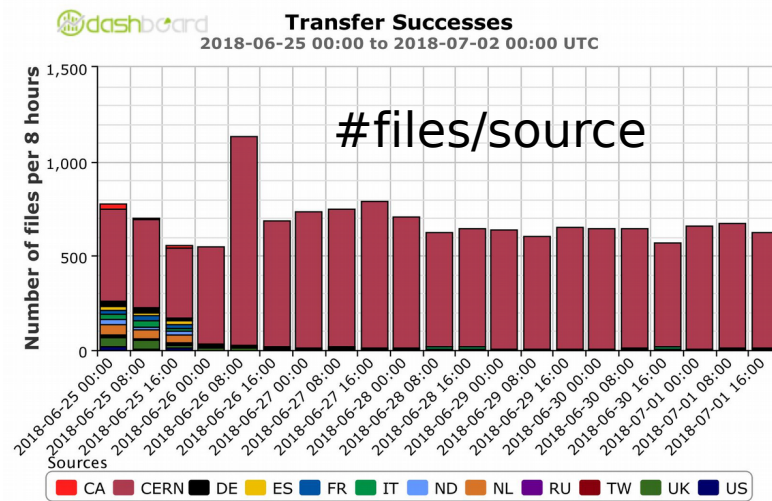


Dataset: 100 files of 1GB - Single client writing (VM)

Integration with data management

Eulake is registered as a storage element in Rucio (ATLAS data management)

4 sample datasets transferred for processing tests



HammerCloud tests

Run HammerCloud framework tests on eulake data. Initial focus on ATLAS.

Allows test real workflows and data access patterns.

Four test scenarios. Read from:

1. Local access (no eulake)
2. eulake, data@CERN, WN@CERN
3. eulake, data NOT@CERN, WN@CERN
4. eulake, 4+2 stripes, WN@CERN

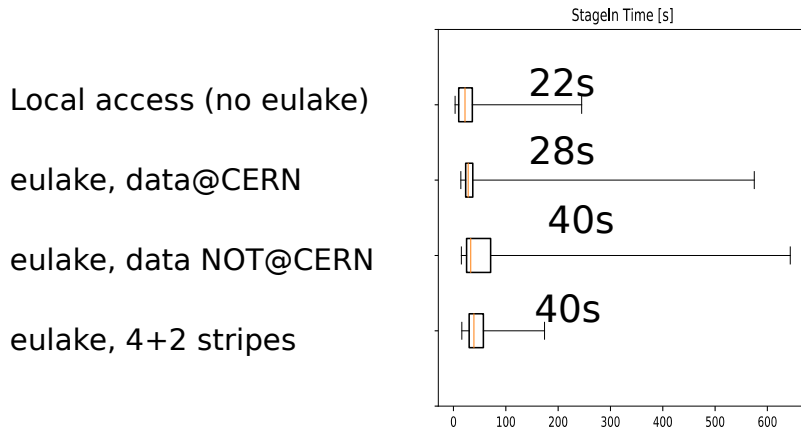
Data is copied from storage to WN

Running Tests backed by the WLCG Data Lake												
State	Id	Host	Template	Start (Europe/Zurich)	End (Europe/Zurich)	Sites	subm jobs	run jobs	comp jobs	fail jobs	fail %	tot jobs
running	20116971	hammercloud-ai-12	1006: benchmark derivation AthDerivation/21.2.8.0 1k events - WLCG Data Lakes - local data clone.977 EULAKE folder CERN	01/Jul, 10:24	02/Jul, 11:52	CERN-PROD_DATA LAKES, CERN-PROD_DATA LAKES_MCORE, CERN-PROD_DATA LAKES_TESTA, 3 more...	1	6	121	8	6	137
running	20116974	hammercloud-ai-12	1007: benchmark digi+reco derivation Athena/21.0.53 5 events - WLCG Data Lakes - local data clone.987 EULAKE folder CERN	01/Jul, 12:58	02/Jul, 13:50	CERN-PROD_DATA LAKES, CERN-PROD_DATA LAKES_MCORE, CERN-PROD_DATA LAKES_TESTA, 3 more...	0	6	71	42	35	119
running	20116984	hammercloud-ai-12	1012: A.F.T. AtlasDerivation 20.7.6.4 clone.808 clone.845 EULAKE folder CERN	01/Jul, 19:02	02/Jul, 16:39	ANALY_CERN-PROD_DATA LAKES, ANALY_CERN-PROD_DATA LAKES_TESTA, ANALY_CERN-PROD_DATA LAKES_TESTB, 2 more...	5	0	0	0	0	5
running	20116988	hammercloud-ai-12	1005: P.F.T. mc16 Sim_tf 21.0.16 - WLCG Data Lakes - local data clone.989 EULAKE folder CERN	02/Jul, 2:58	03/Jul, 5:01	CERN-PROD_DATA LAKES, CERN-PROD_DATA LAKES_MCORE, CERN-PROD_DATA LAKES_TESTA, 3 more...	0	6	63	6	8	77

Running Tests backed by the standard storages, copy-to-scratch												
State	Id	Host	Template	Start (Europe/Zurich)	End (Europe/Zurich)	Sites	subm jobs	run jobs	comp jobs	fail jobs	fail %	tot jobs
running	20116979	hammercloud-ai-12	977: benchmark derivation AthDerivation/21.2.8.0 1k events - WLCG Data Lakes - local data	01/Jul, 13:42	02/Jul, 14:21	CERN-PROD-preprod, NIKHEF-ELPROD, SARA-MATRIX, 6 more...	0	4	96	0	0	100
running	20116980	hammercloud-ai-12	987: benchmark digi+reco derivation Athena/21.0.53 5 events - WLCG Data Lakes - local data	01/Jul, 13:54	02/Jul, 12:20	CERN-PROD-preprod, NIKHEF-ELPROD, SARA-MATRIX, 6 more...	0	3	96	0	0	100
running	20116986	hammercloud-ai-73	845: AFT AtlasDerivation 20.7.6.4 clone.808	01/Jul, 21:42	02/Jul, 19:38	ANALY_AGLT2_SL6, ANALY_AGLT2_TEST_SL6-condor, ANALY_AARNES, 142 more...	180	237	6603	267	4	7339
running	20116991	hammercloud-ai-12	989: P.F.T. mc16 Sim_tf 21.0.16 - WLCG Data Lakes - local data	02/Jul, 8:04	03/Jul, 6:16	CERN-PROD-preprod, NIKHEF-ELPROD, SARA-MATRIX, 6 more...	0	2	8	0	0	10

Test job performance

Low I/O intensity workflow: ~40MB input (1 file), 2 events, ~5 mins/event




Local access (no eulake)

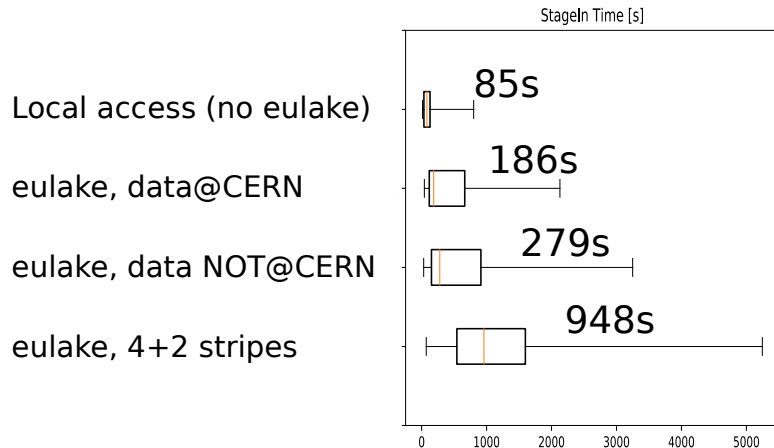
eulake, data@CERN

eulake, data NOT@CERN

eulake, 4+2 stripes



High I/O intensity workflow: ~6GB input (1 file), 1000 events, ~2 seconds/event

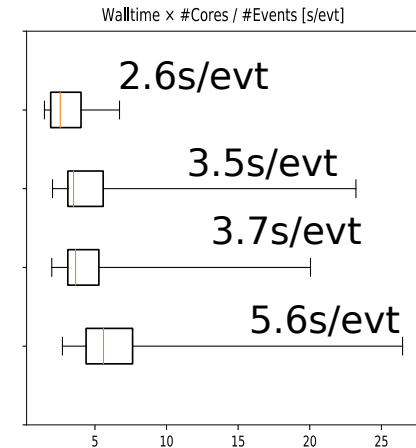


Local access (no eulake)

eulake, data@CERN

eulake, data NOT@CERN

eulake, 4+2 stripes



Conclusions

Distributed storage instance based on EOS technology

Small in storage volume, large in the geographical sense

Deployment varieties: Native EOS, EOS on Docker containers, volume export (CEPH)

Integrated with the ATLAS distributed computing services and HammerCloud

Next step is integration with CMS

Prototype in place and performance metrics taken

Only first results, not enough for conclusions

The prototype allows to test many ideas in preparation for HL-LHC

Acknowledgements

This work has been done by a team at CERN in collaboration with many participant WLCG sites, with the help of the ATLAS Distributed Computing and particularly the Rucio team, and the EOS experts at CERN

This report is based on the presentation at CHEP by Simone Campana