

DPM MIGRATION: BERNE EXPERIENCE

Gianfranco Sciacca

AEC - Laboratory for High Energy Physics, University of Bern, Switzerland

GDB - 15 September 2022

u^b

^b
UNIVERSITÄT
BERN

AEC
ALBERT EINSTEIN CENTER
FOR FUNDAMENTAL PHYSICS

LABORATORIUM FÜR HOCHENERGIEPHYSIK
LHEP
UNIVERSITÄT BERN

STATUS AT FEBRUARY GDB

<https://indico.cern.ch/event/1096027/contributions/4706375/attachments/2388029/4081517/Swiss-DPM-migration%20plans.pdf>

▶ Three options were identified for the Bern DPM migration

2.3 PB for ATLAS Tier-2/3 ... and a bit more

- A. In place migration to dCache
- B. Consolidation of resources at the national level (dCache)
- C. Consolidation of resources at the international level (dCache)

* *all solutions come with pros and cons*

▶ Decision taken at end of February

u^b

DECISION ON DPM MIGRATION 1/2

▶ ATLAS Tier-2 storage

C. Pledged capacity to be integrated with the NDGF-T1 dCache

- * 1.7 PB on 3 servers
- * ATLASDATADISK and ATLASCRATCHDISK
- * Tier-1 storage provides better value to researchers vs. Tier-2 storage
- * Decrease the fragmentation of data endpoints (data lake model)
- * Uncomplicated integration process

DECISION ON DPM MIGRATION 2/2

▶ ATLAS Tier-3 storage and neutrino experiments

A. In place migration to dCache

- * 0.6 PB on 7 servers
- * Wish to retain local expertise on SE, learn how to run dCache
- * Lower QoS, so less stringent requirements on hardware

▶ **ATLAS Tier-2 storage**

- * **PoC in production since early 2022 with 80 TB integrated with NDGF-T1_ATLASDATADISK**
 - * setup and operation uncomplicated
 - * communication and support over a dedicated channel for chat with NDGF ops
- * **Core storage migration started in March 2022**
- * **Core storage commissioning completed in early September 2022 (target was end 2022)**

▶ ATLAS Tier-2 storage

* UniBE checklist

- * ATLAS data on the DPM drained by ATLAS central DDM team
- * Internal re-factorisation of data areas on the disk servers (*all servers were shared among other users*)
- * CPU, memory, internal network upgrades to get as close as possible to the recommended specs
- * Fresh installation of servers with minimal CentOS 7
- * Apply some settings to allow the remote deployment of the dCache stack (*wiki checklist*)
- * Apply TCP and OS tuning, as recommended by NDGF for performance (*wiki checklist*)
- * Integration with the NDGF central monitoring (*wiki checklist*)
- * Firewall / Net ACLs settings (*wiki checklist*)
- * **Handed over to NDGF for deployment of dCache**
- * (... summer break ...)



▶ ATLAS Tier-2 storage

* NDGF dCache administrator checklist




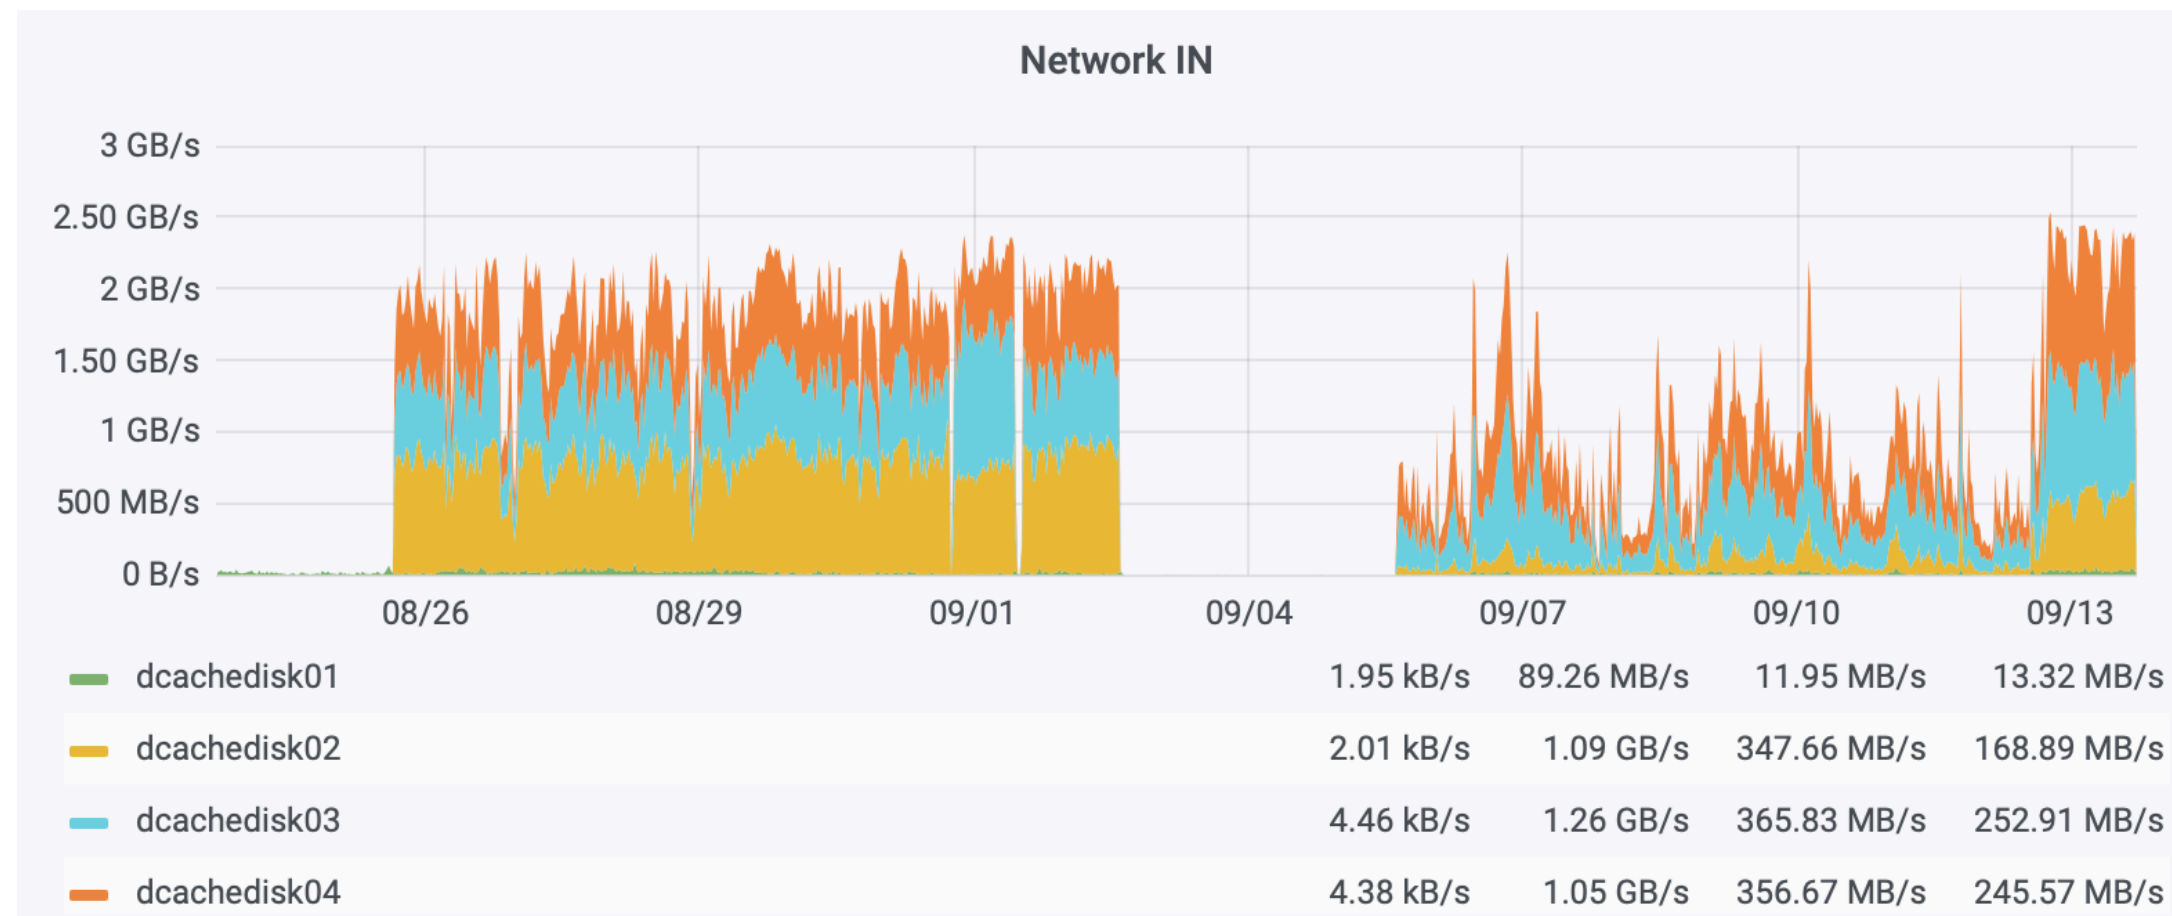

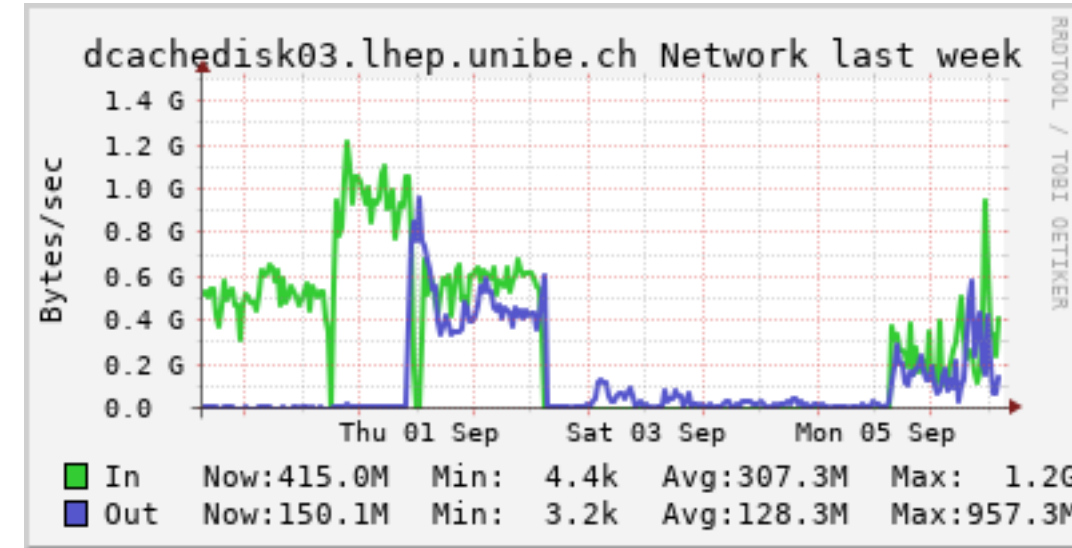
- * dCache stack deployed via Ansible and an unprivileged user account
- * Workname is "tarpool" (*tar file distribution as opposed to a deb/rpm package to install locally*)
 - * Install and upgrade Java and dCache and apply the required configuration of dCache
 - * Unprivileged access via ssh keys
- * Ensure communication with the remote pools can be established, add to the monitoring
- * Load-test pools before taking them to production
 - * Fill them up with production data from nearby/fast pools, then checksum all data on disk
 - * See how fast you can read data out (migration cache to a different set of pools)
 - * Migrate in some new data, see if writes starves reads, or vice versa.
 - * Add to a read-only pool group to get real client reads from the pool

* Add to production poolgroup


IMPLEMENTATION: FROM TESTING TO PRODUCTION



IMPLEMENTATION: FROM TESTING TO PRODUCTION



PRODUCTION SNAPSHOT



Snapshot ▾

🕒 Last 24 hours

Endpoint	Files	Bytes	Storage used	Storage total ▾
NDGF-T1_DATADISK	91536990	20.58 PB	21.04 PB	22.20 PB
BNL-OSG2_DATADISK	50577284	21.78 PB	21.79 PB	22.10 PB
RAL-LCG2-ECHO_DATADISK	31323173	16.59 PB	-1.00 B	16.93 PB
IN2P3-CC_DATADISK	32843152	15.22 PB	15.60 PB	16.03 PB
FZK-LCG2_DATADISK	31137124	11.75 PB	12.13 PB	12.44 PB
TRIUMF-LCG2_DATADISK	28838504	9.95 PB	10.60 PB	11.00 PB
INFN-T1_DATADISK	20710356	8.28 PB	8.28 PB	8.70 PB
TAIWAN-LCG2_DATADISK	13182498	6.03 PB	6.05 PB	7.48 PB
SARA-MATRIX_DATADISK	17231565	3.70 PB	3.79 PB	4.97 PB
RRC-KI-T1_DATADISK	15779404	4.19 PB	4.21 PB	4.56 PB
PIC_DATADISK	19137711	3.55 PB	3.57 PB	3.87 PB
NIKHEF_DATADISK	6966214	2.63 PB	2.63 PB	2.94 PB

OUTSTANDING WORK

▶ ATLAS Tier-2 storage

* Outstanding work: accounting

* Implementation of the SSR for federated storage accounting

- * allows tracking of remote site storage contribution in the monthly WLCG/CRIC reports
- * discussed and agreed on at the WLCG ops coordination meeting in July 2022
- * add an additional share per site to the SSR, e.g. : **xxx_admin_UNIBE-LHEP**
- * the WSSA (WLCG storage space accounting) application handles these shares accordingly to account the corresponding space to the correct site
- * the experiment operation is instead only concerned about the whole NDGF storage area

OUTSTANDING WORK

▶ **ATLAS Tier-2 storage**

* **Outstanding work: network**

* **Upgrade and possibly connect Uni Bern to LHCONE**

* current connectivity for the Bern Science DMZ is 2x10Gb/s

* saturation following the full storage migration

* upgrading to 40 Gb/s now (*100Gb/s at beginning of 2023*)

* negotiating with SWITCH + UniBE NOC connection to LHCONE (or an equivalent solution)

* discussions in progress ...

TIMELINE AND IMPLEMENTATION

▶ **ATLAS Tier-3 storage and neutrino experiments**

- * **ATLAS_LOCALGROUPDISK, MicroBooNE, DUNE storage areas**
 - * DPM to dCache "in place" migration targeted for Q2 2023
 - * Hope to get it done by end of 2022

CONCLUSIONS

- ▶ **The ATLAS Tier-2 DPM storage at UniBE has been integrated with the NDGF-T1 dCache**
 - * 2nd Tier-2 site storage integrated with NDGF
 - * Smooth integration procedure
 - * In production since about a week
 - * Accounting and network upgrades being implemented

- ▶ **The ATLAS Tier-3 and neutrino storage to be migrated in place from DPM to dCache**
 - * Not anticipating particular complications
 - * Lower QoS means lower pressure