

Integration of the Goettingen HPC cluster Emmy to the WLCG Tier-2 centre GoeGrid and performance tests

Saidev Polisetty

22.10.2024

CHEP 2024, Kraków

GEFÖRDERT VOM



Bundesministerium
für Bildung
und Forschung

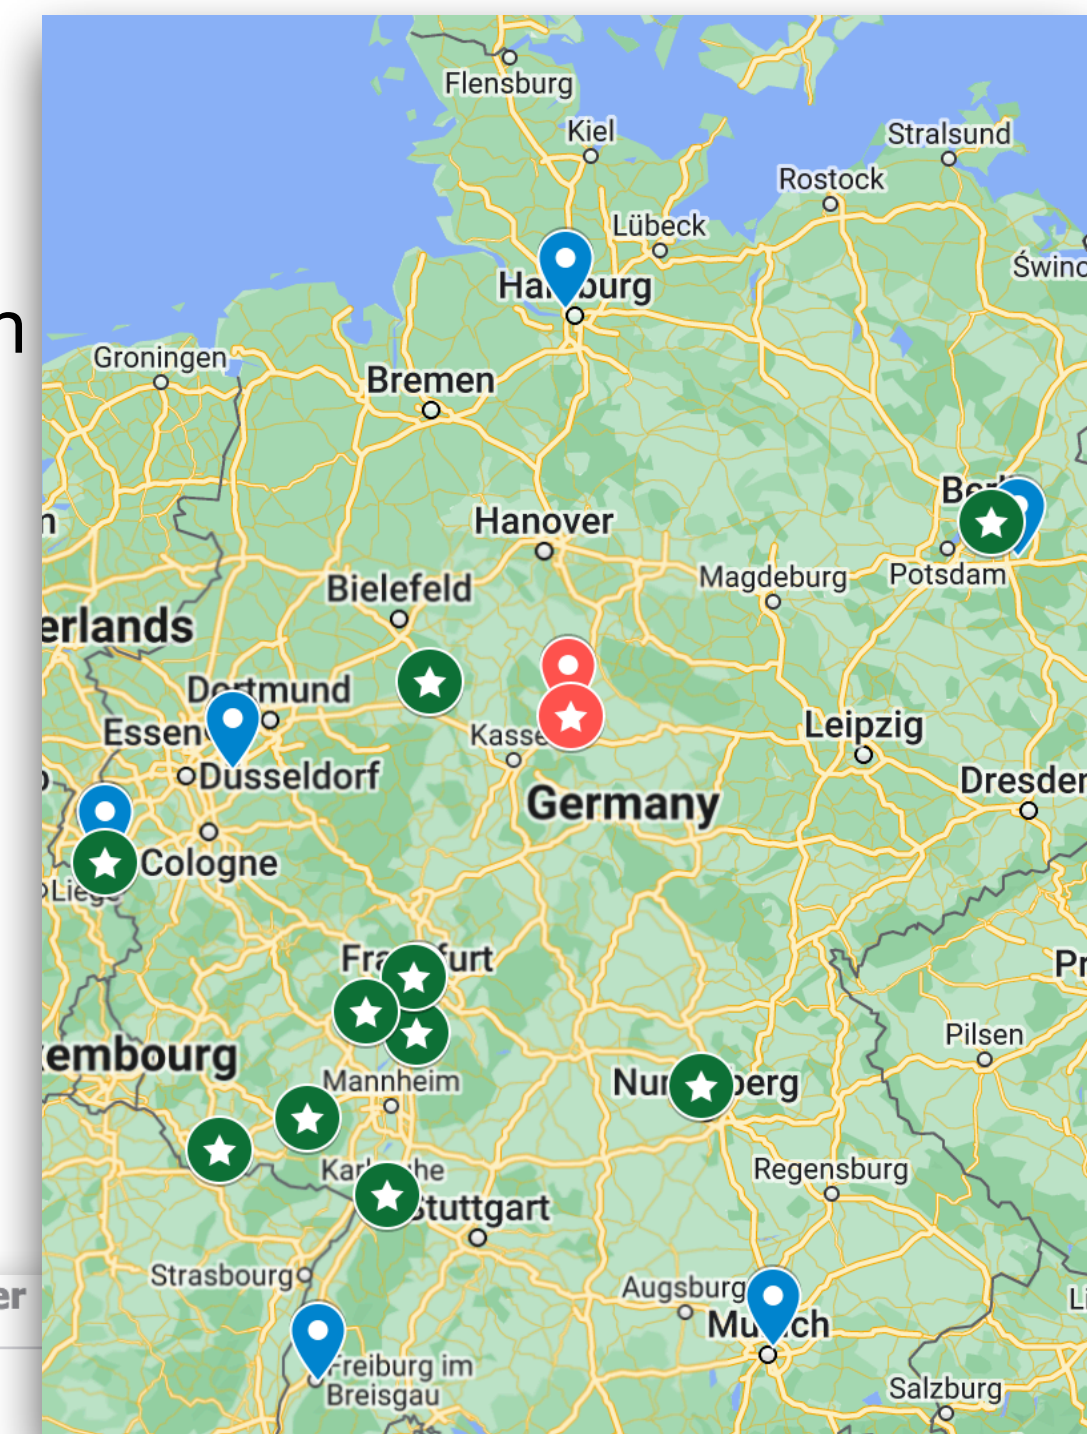


Motivation

- ATLAS needed ~300k cores and ~300 PB storage in 2022.
- New methods of computing and storage are needed as int. lumi. is expected to increase to ~3000 fb⁻¹ for the HL-LHC Run-4.
- WLCG-Tier 2 sites presently run at DESY, MPI and universities within Germany.
- University compute power will move to National High Performance (NHR) sites, mass storage will move to Helmholtz Centres, DESY and KIT.

- Transition will be finished by early 2029.
- NHR, an association of large, university-based, multidisciplinary HPC centres.

- Tier - 2 sites Germany 
 - DESY Hamburg and Zeuthen
 - Uni Freiburg
 - GoeGrid - Uni Goettingen
 - GSI
 - MPI
 - RWTH Aachen
 - Uni Wuppertal



University	High Performance Center
RWTH Aachen	IT Center
Technische Universität Darmstadt	Hochschulrechenzentrum (HRZ)
Universität Nürnberg-Erlangen	Regionales Rechenzentrum Erlangen
Universität Göttingen	Gesellschaft für wissenschaftliche Datenverarbeitung mbH Göttingen (GWDG)
Karlsruher Institut für Technologie	Steinbuch Centre for Computing (SCC)
Technische Universität Dresden	Zentrum für Informationsdienste und Hochleistungsrechnen
Berlin University Alliance	Zuse-Institut Berlin
Universität Paderborn	Paderborn Center für Parallel Computing
Universität Frankfurt am Main	NHR Süd-West
Rheinland-Pfälzische Technische Universität Kaiserslautern-Landau (RPTU)	
Universität Mainz	
Universität des Saarlandes	

The FIDIUM project

- **FIDIUM - Federated digital infrastructures for research in universe and matter**
 - I. Development of tools for the integration of heterogeneous resources.
 - Using external computing sites as additional resources for High Energy Physics studies.
 - II. Data Lakes, Distributed Data, Caching.
 - III. Customisation, testing and optimisation on production and analysis environments.
 - The reliability, scalability and maintainability of the developed solutions are tested under realistic conditions.
- Project funded by the federal government.

SPONSORED BY THE



Göttingen campus - Network managed by GWDG

GoeGrid (HTC)

- Since 2008
- Dedicated resources
- WLCG Tier-2 and 3 for ATLAS
- 17,000 cores
- 3.7 PB disk grid storage (dCache)

NHR : Emmy (HPC)

- Since 2018
- General state or federal resources
- National High Performance (NHR) computing centres
- 140,000 cores
- 8400 GPU cores



- GoeGrid is a Tier-2 and Tier-3 High Throughput Computing (HTC) resource.
- Göttingen model: joint operation of grid centre with colleagues from medicine, humanities, computer science, theoretical physics.
- Emmy is a High Performance Computing (HPC) cluster which is also located in Göttingen.

Göttingen campus - Network managed by GWDG

GoeGrid (HTC)

- Since 2008
- Dedicated resources
- WLCG Tier-2 and 3 for ATLAS
- 17,000 cores
- 3.7 PB disk grid storage (dCache)


NHR : Emmy (HPC)

- Since 2018
- General state or federal resources
- National High Performance (NHR) computing centres
- 140,000 cores
- 8400 GPU cores



- For the integration, resources from both the clusters are utilised.
- Initially, storage is provided by GoeGrid (dCache).
- Jobs are run on the Worker Nodes of the EMMY cluster.

Cluster layout: GoeGrid



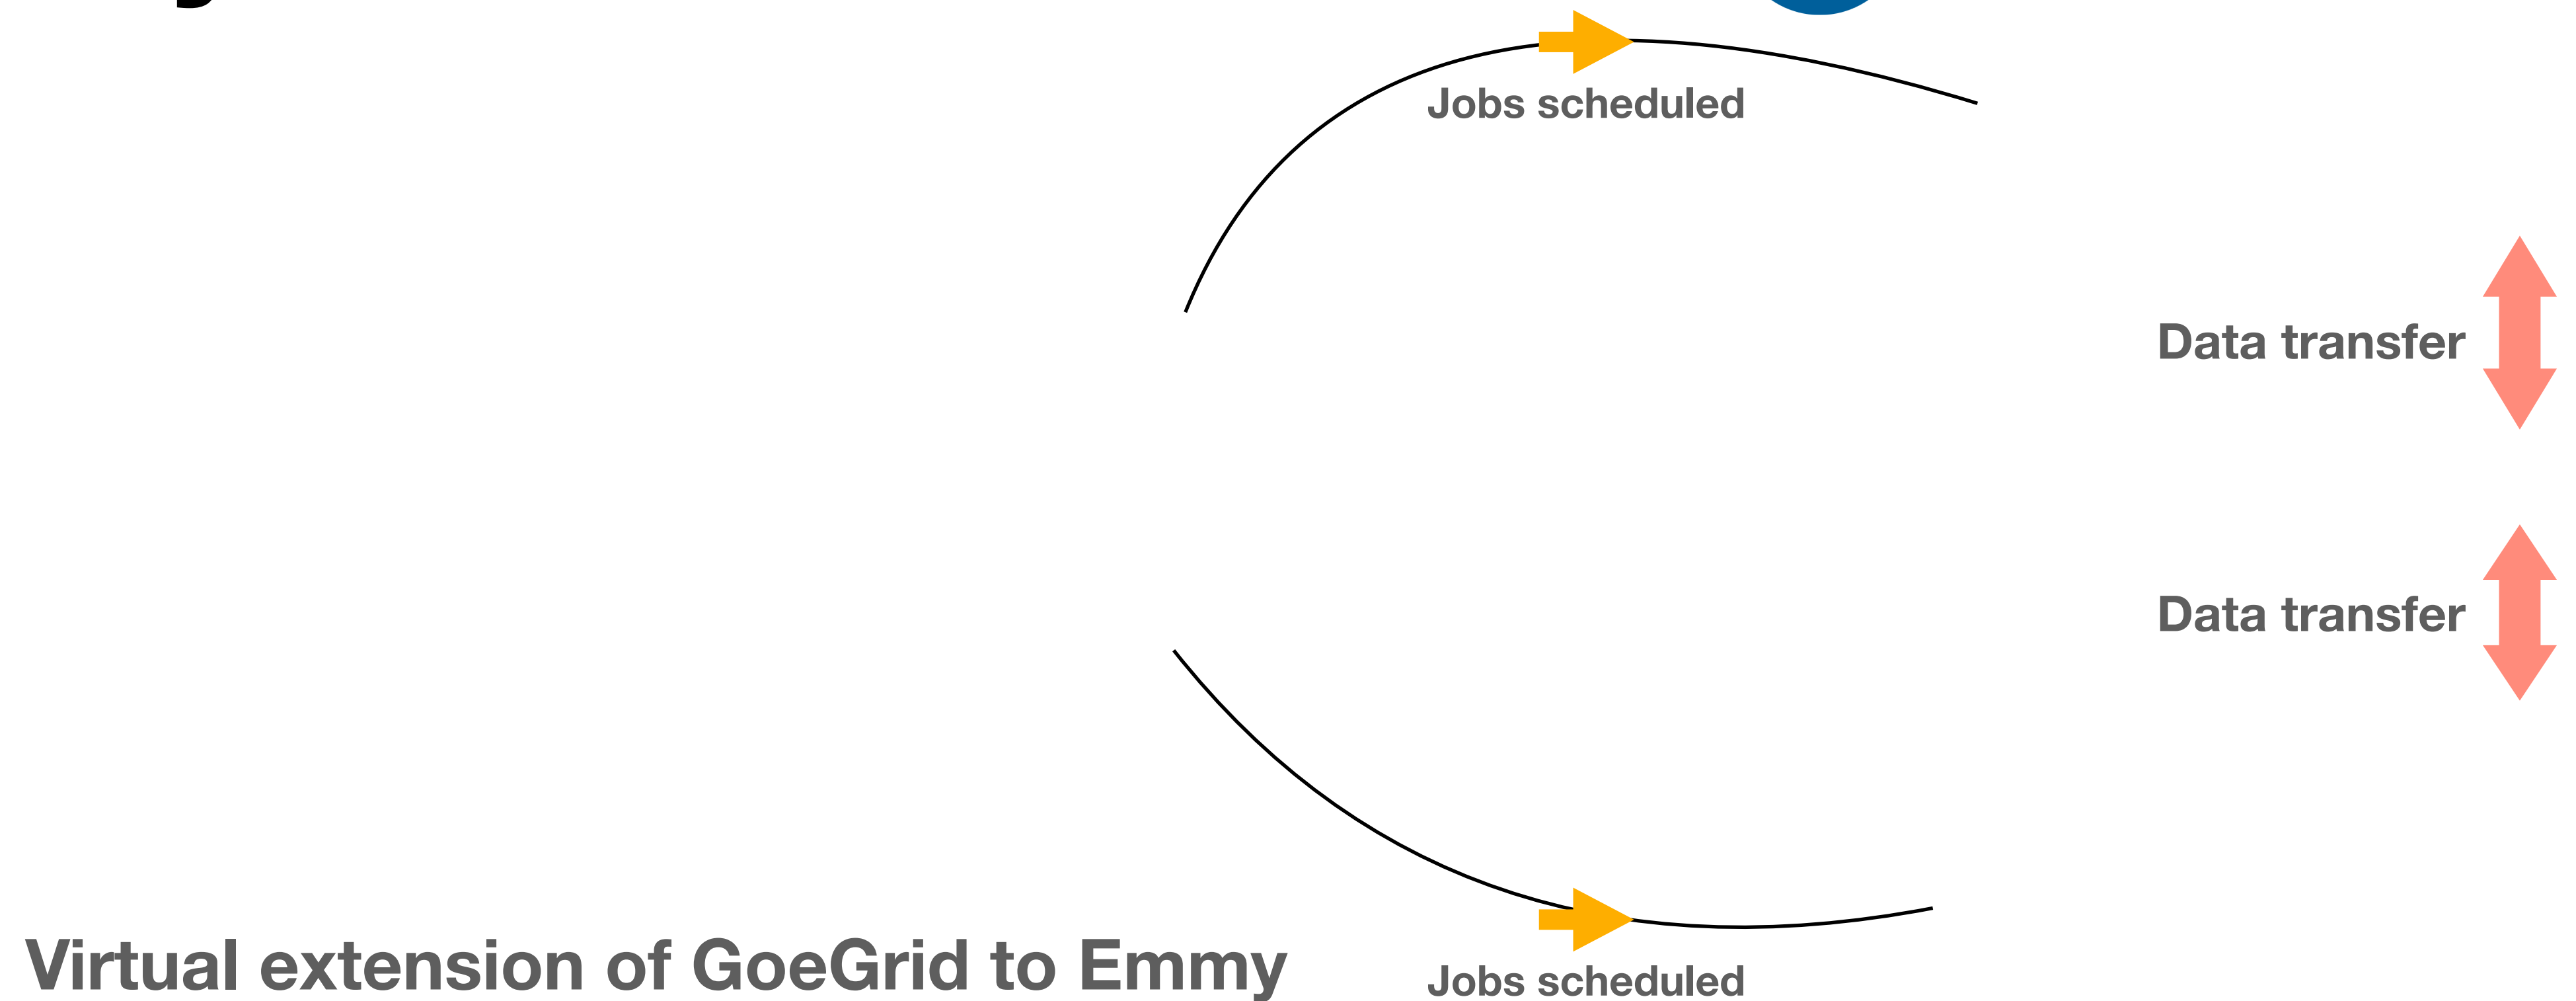
Workflow of jobs on GoeGrid

- ATLAS submits jobs to the PanDA queue.
- HTCondor receives jobs from the Panda Queues and schedules them to WNs.
- Data is read/written to the local GoeGrid data storage.

Cluster layout: EMMY extension

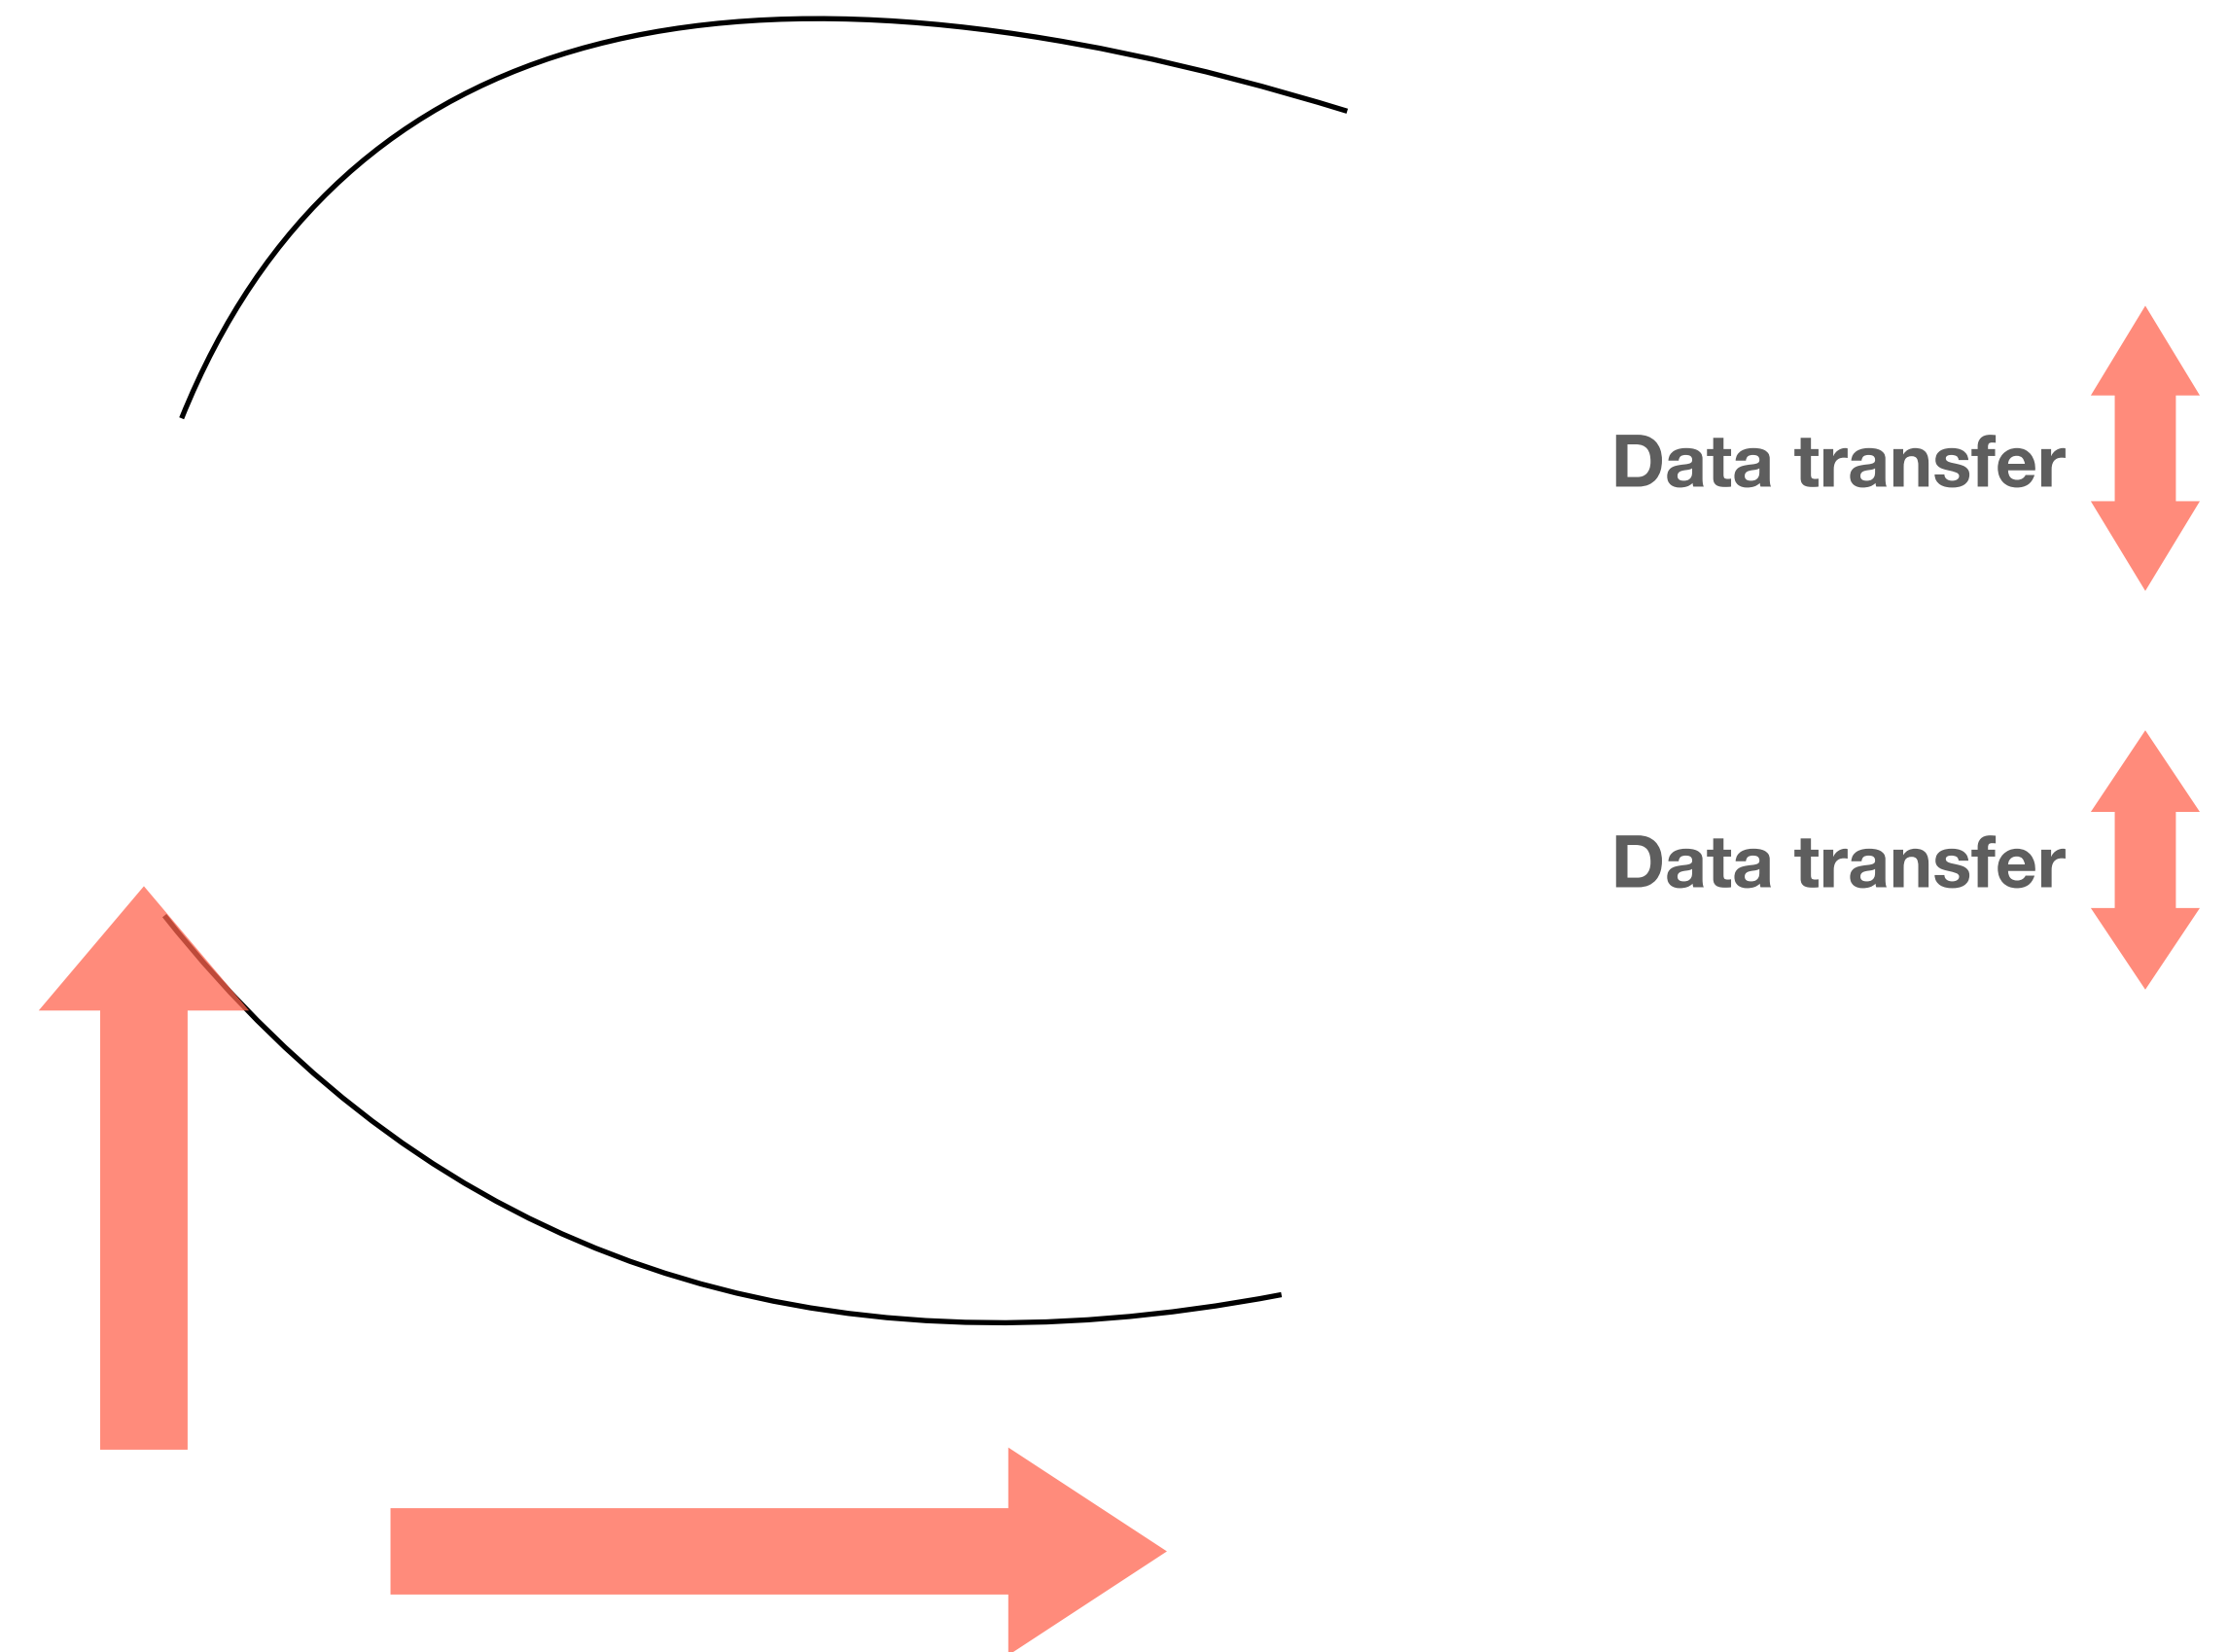


GEORG-AUGUST-UNIVERSITÄT
GÖTTINGEN IN PUBLICA COMMODA
SEIT 1737



- Singularity/Apptainer container containing HTCondor is launched on the HPC.
 - Slurm is modified to have whole node scheduling, allowing partitionable slots.
- Unprivileged CVMFS client (cvmfsexec) is installed in container due to lack of FUSE permissions.
 - But allows us to have control over cvmfs.

Cluster layout: Integrated



- COBaLD/TARDIS manages the provision of virtual nodes depending on resource needs and availability.
- High bandwidth connection established between Emmy and GoeGrid (4x100 Gbit/s).



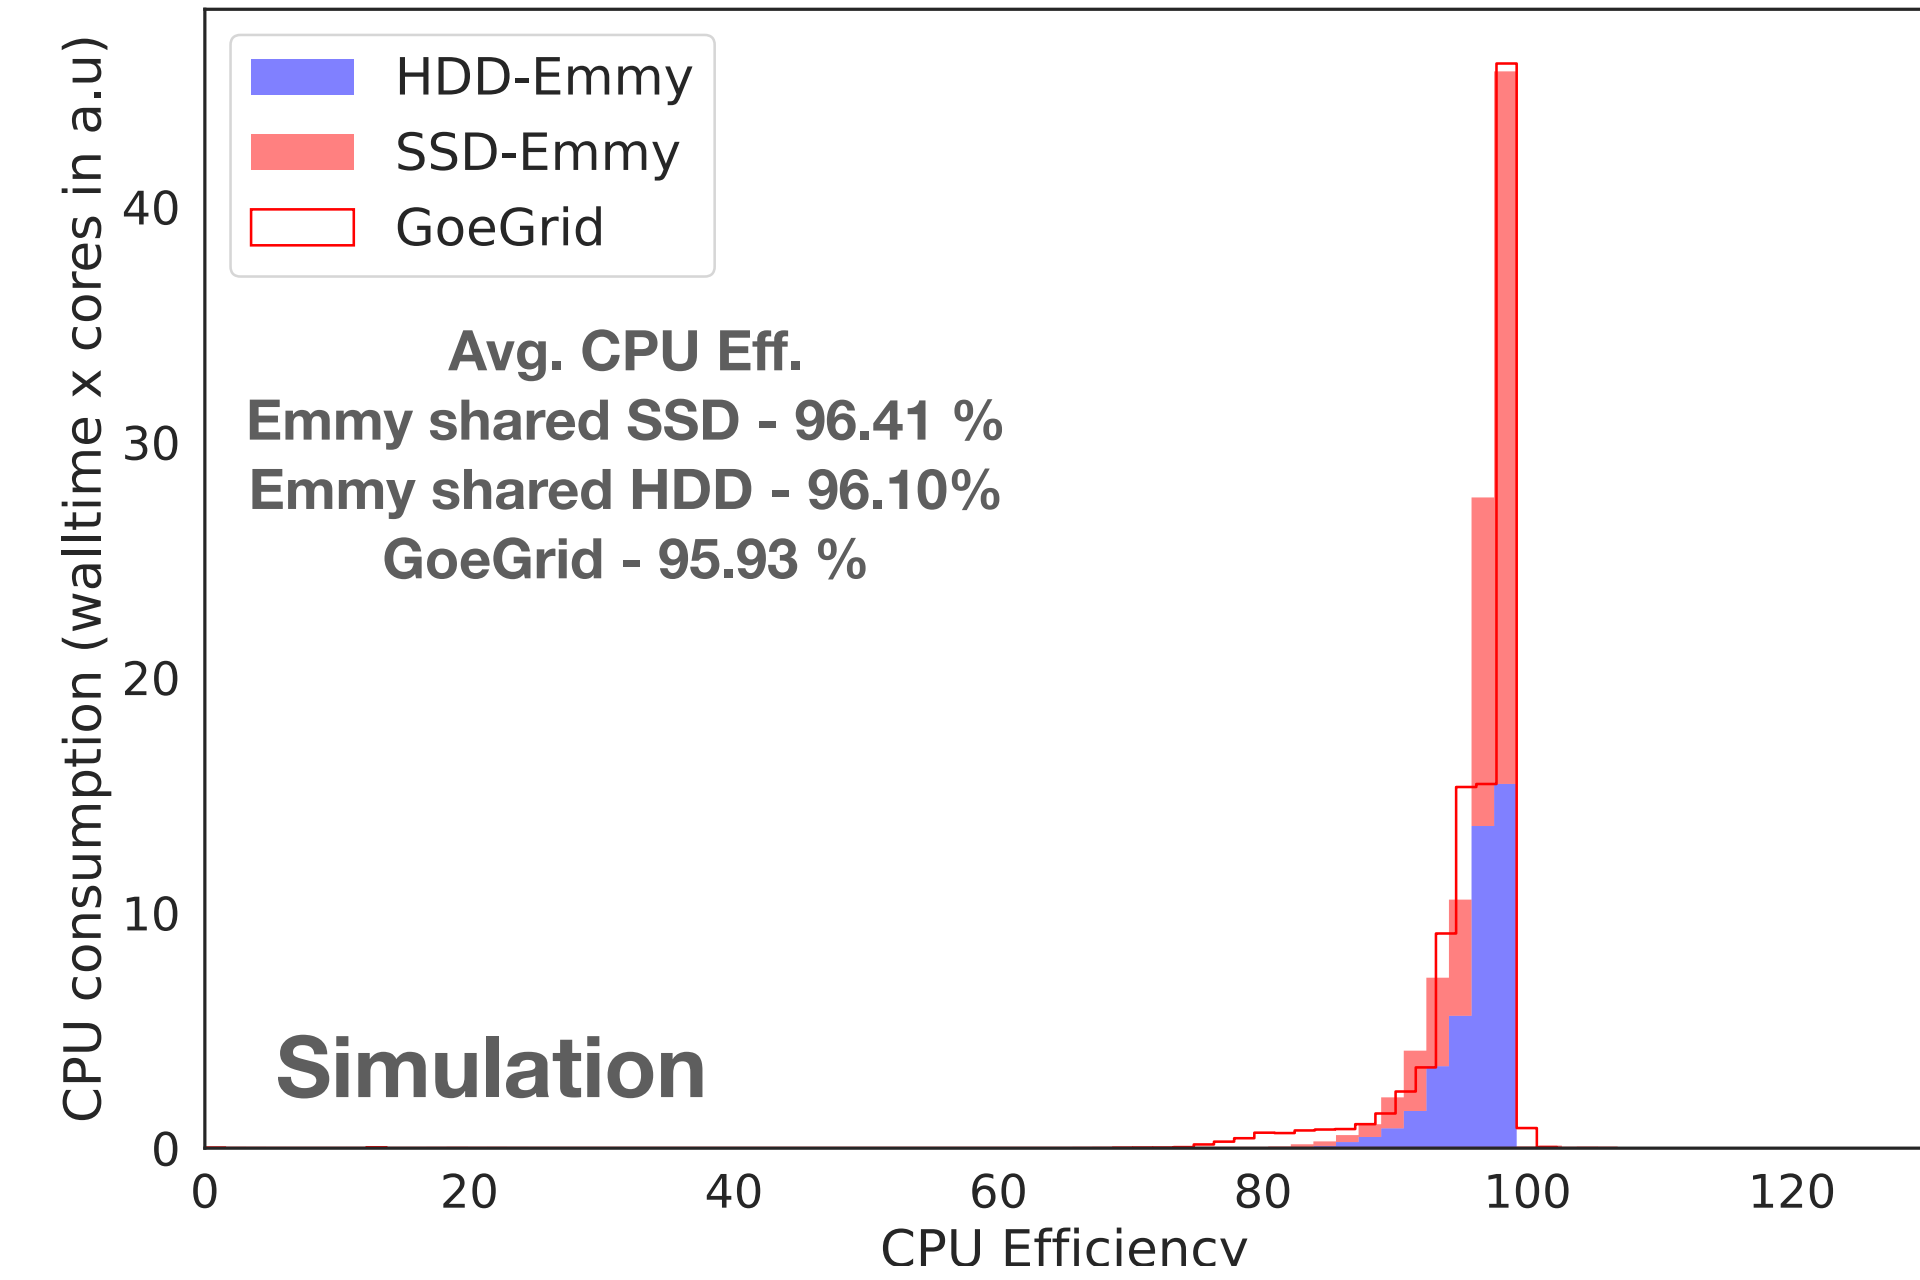
This virtually extends the GoeGrid batch system into Emmy using containers turning HPC nodes into virtual nodes with own job scheduling.

Tests on EMMY

- Tests compare jobs on “EMMY” to the Tier-2 cluster of GoeGrid.
 - Shows comparison between containerised setup and a regular Tier-2 WLCG site.
 - Effects of storage types on CPU performance.
 - Sustainability of the cluster on heavy loads.
 - Provides an insight on finding ways to optimise resource consumption for max. efficiency.
- Jobs contain different ATLAS job types and user analysis jobs.
- Lifetimes of jobs are up to a max. 4 days and the drones run for 7 days.

Benchmark

- HEPscore23 (HS23) benchmark
 - GoeGrid: ~2380 HS23 on 256 cores with HT.
 - Emmy:
 - Phase 2: Shared SSD nodes: ~1900 HS23 on 192 cores with HT.
 - Phase 3: Local SSD nodes: ~2900 HS23 on 192 cores with HT.



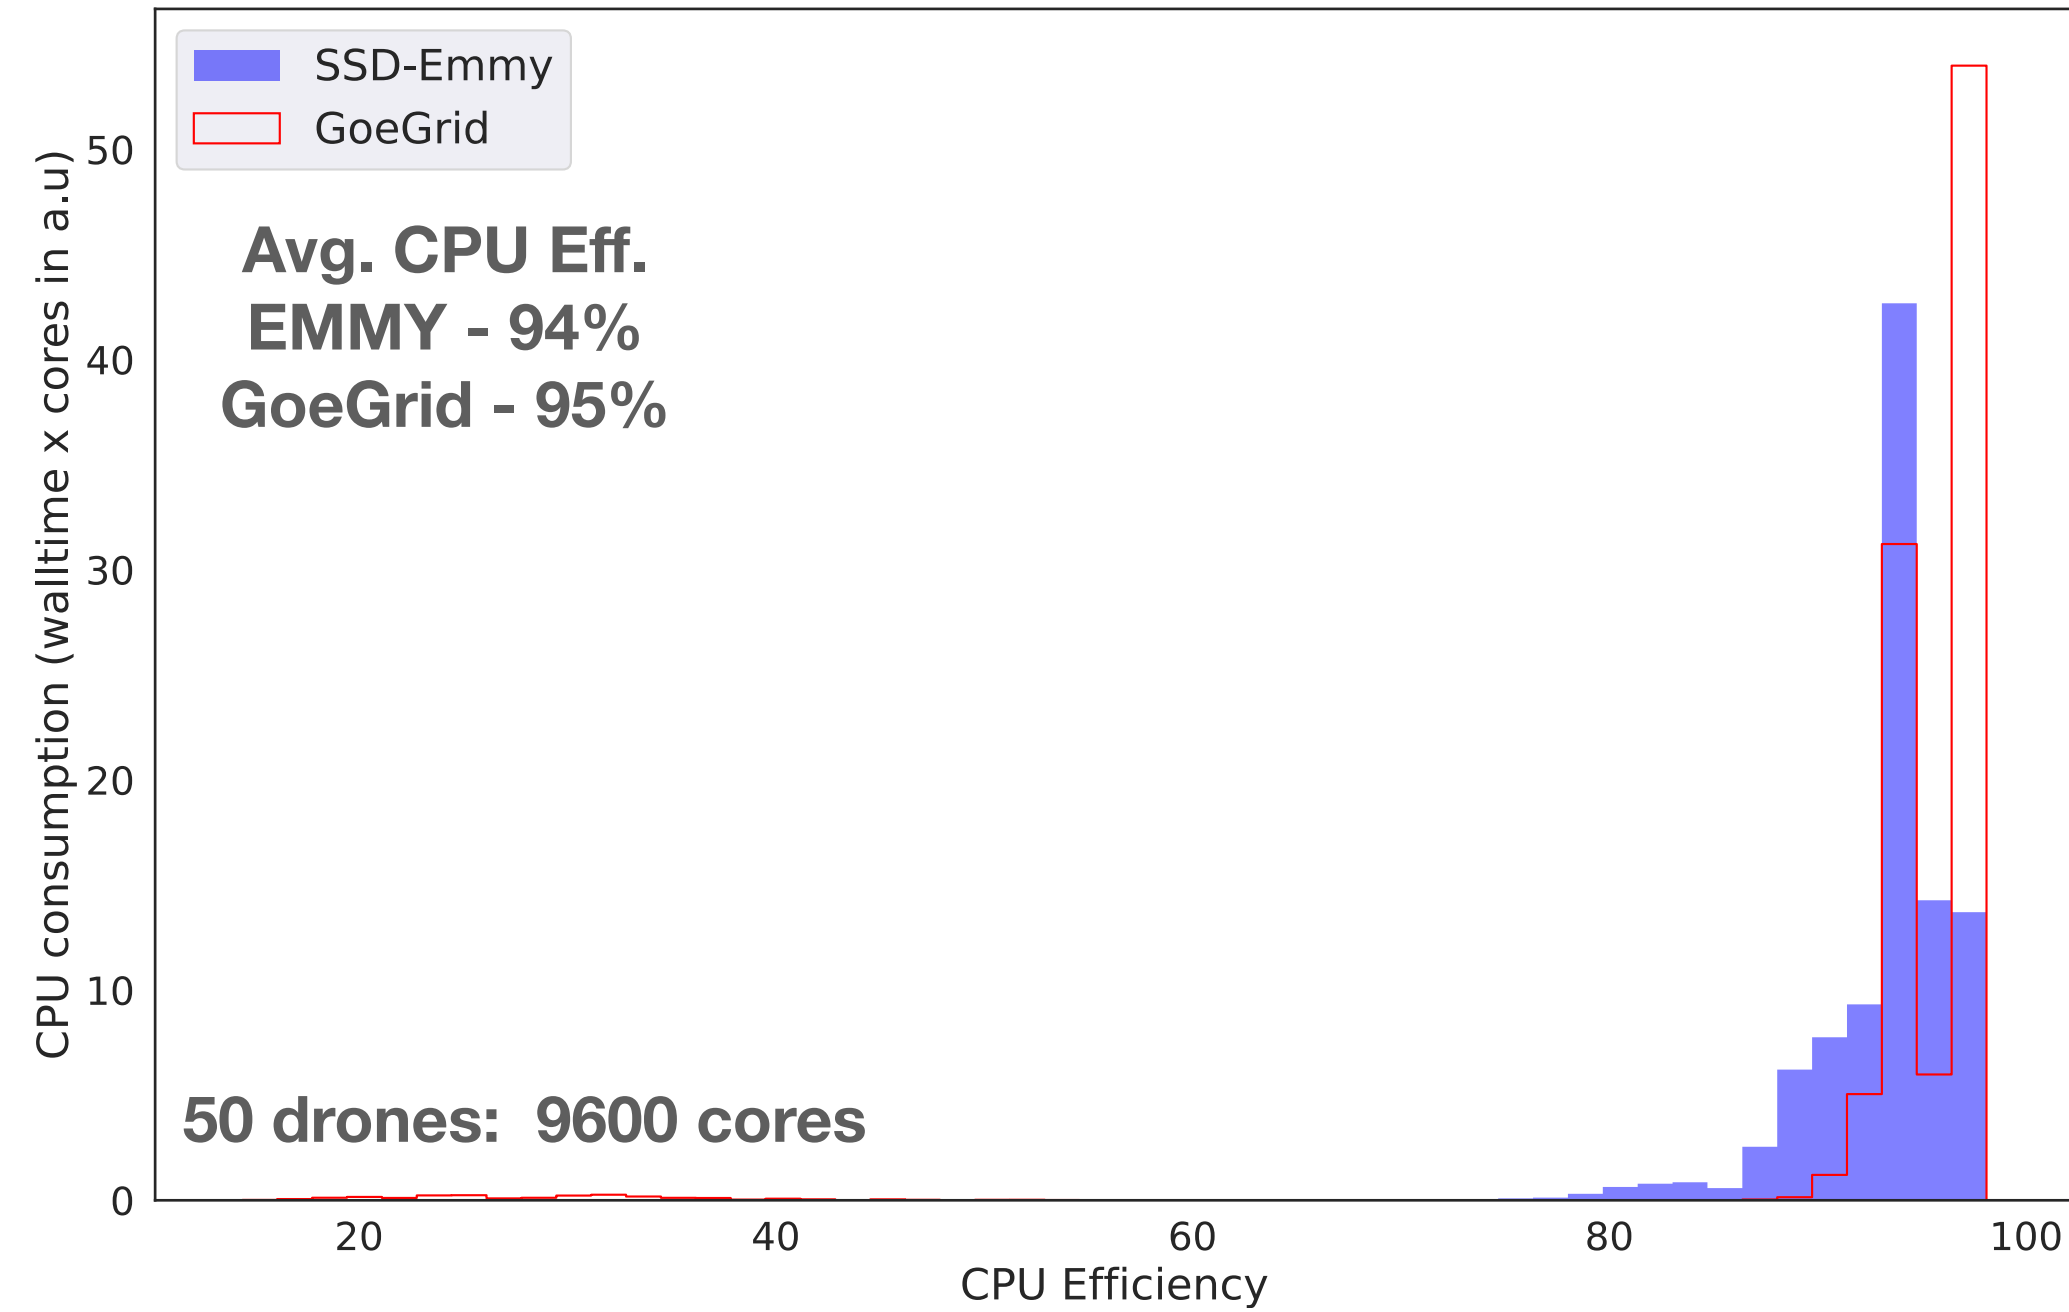


- Normalised CPU consumption of jobs for phase 2 Emmy is shown.
- >100% is due to background proc. of evgen jobs (monitoring) and higher parallelisation than booked cores.
- Higher IO intensive jobs have an effect on the CPU efficiencies (Pile, recon and analysis jobs from users).

Remarks




- The data is taken under realistic conditions.
- GoeGrid has local disks.
- EMMY has a large fraction of shared disks (lustre filesystem).
 - Shared SSDs and shared HDDs.
 - Lustre filesystem has an effect on shared SSD/HDD performance based on no. of users.
 - Data Intensive jobs are affected.
- Under ideal conditions (with only ATLAS jobs) shared SSD perform better than HDD.
- Under these shared nodes, performance of EMMY is comparable to GoeGrid.

Scaling tests



- CPU efficiencies of diff. no. of drones for **simulation** jobs are shown.
- Efficiency remains constant with minor fluctuations.
 - Most CPU resources are used by jobs at higher efficiencies.


Scaling tests



- CPU efficiencies of diff. no. of drones for **all** jobs are shown.
- Efficiency remains constant with minor fluctuations on scaling of the drones.


Data caching

- ATLAS typically runs job near the data.
- Rucio manages transfer, deletion of the data.
- Storage will be moved to remote sites.
 - Wide Area Network, pre-caching, ...
- **Pre-caching**, where small cache instances are used near sites to run jobs (testing: EMMY).
 - Performance to be compared to GoeGrid.
- Tests are done with 50TB disk space.
 - Constant flow of jobs is observed.
 - Data is transferred from remote mass storage.
- Testing with different sizes and no. of drones.



Summary and Outlook

- First phase of Integration of EMMY to GoeGrid is achieved.
 - Automation implemented using COBaLD/TARDIS.
 - Production pilot phase is running currently, transition phase till 2029.
- Large scale test on Emmy (~38,400 v. cores) has been successful.
- **Performance of the clusters is quite comparable with each other and latest nodes on the HPC outperforming GoeGrid.**
 - **No drawback observed with containerised setup.**
- First results from local data caching are promising.
- Further tests on ATLAS job load scaling, storage scaling and network load will be conducted.



- With support from Inga Lakomic, Arnulf Quadt, Daniel Schindler, Sebastian Wozniewski and the network, administration team of Emmy and the partners of FIDIUM project.

Backup slides

Remarks

- The data is taken under realistic conditions.
 - GoeGrid has local disks where as NHR (Emmy) has large fraction of shared disks (lustre filesystem).
 - There are always third party users using the Emmy resources affecting the performance of ATLAS jobs.
 - This also effects SSD/HDD performance.
- Under the ideal conditions (where only ATLAS jobs run), SSD perform significantly better than HDD.

CPU Efficiency (normalised) of ATLAS jobs on EMMY NHR HPC site

**Regular run
2 months**

**Exclusive run
2 days**

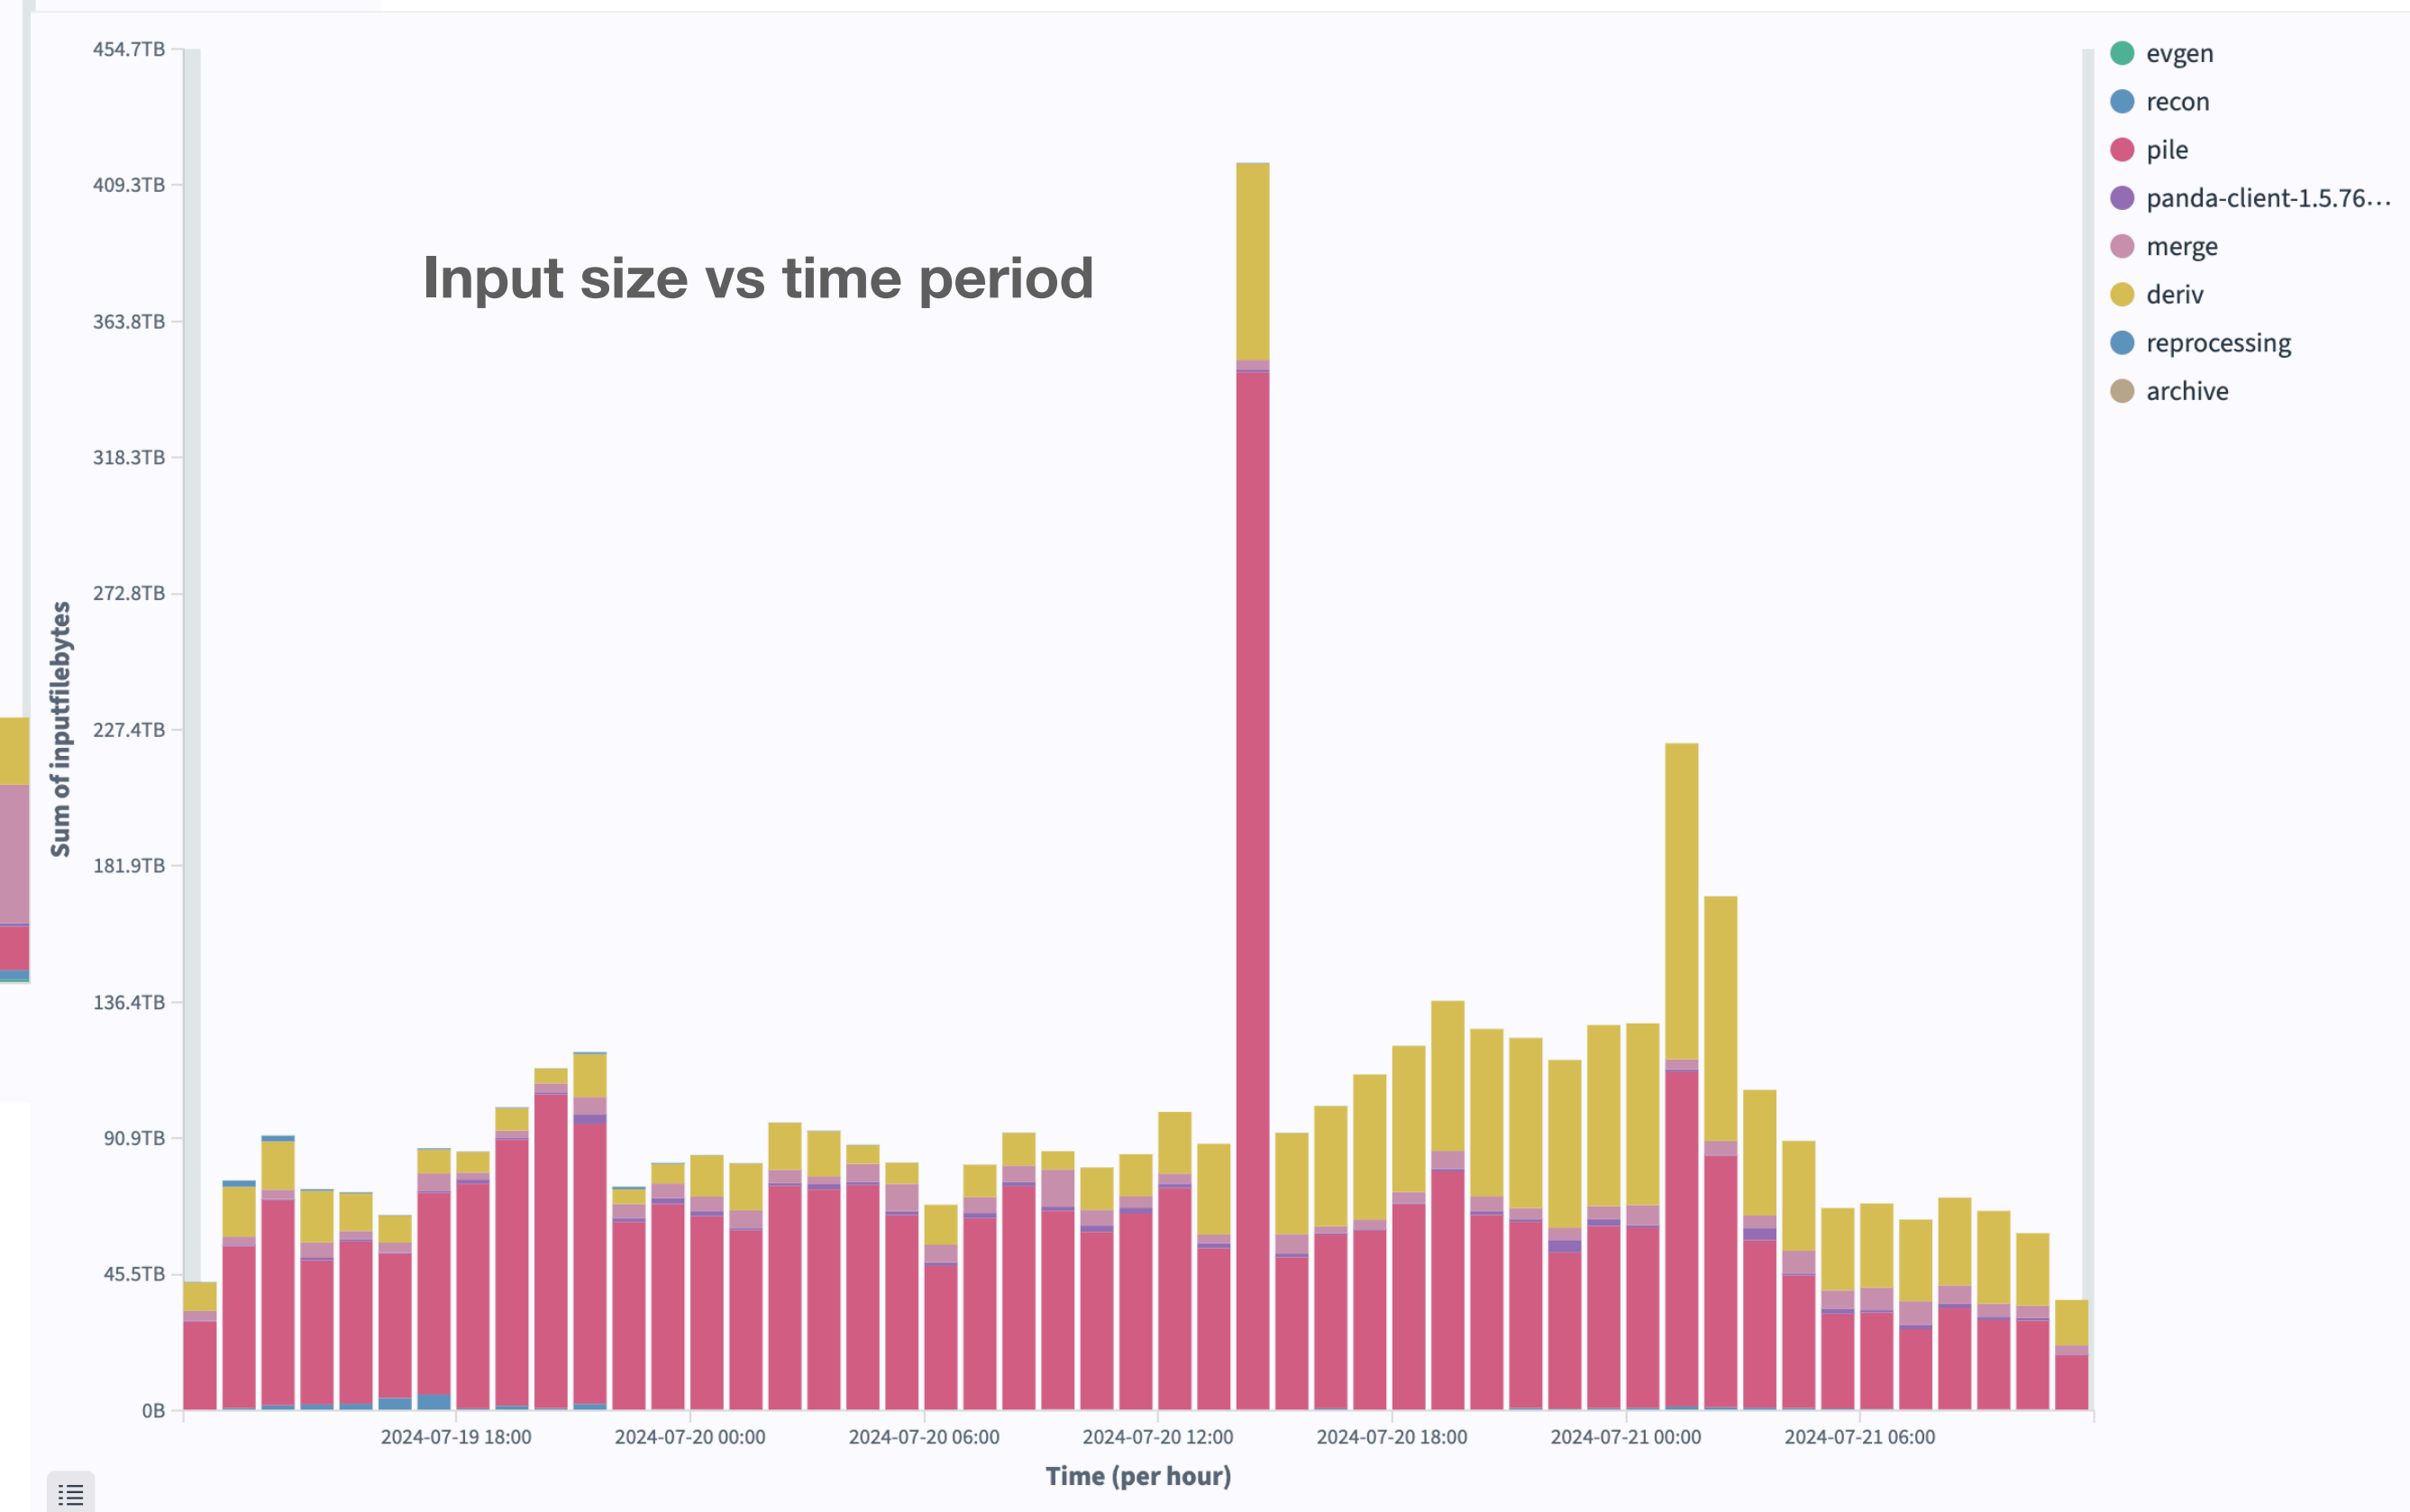
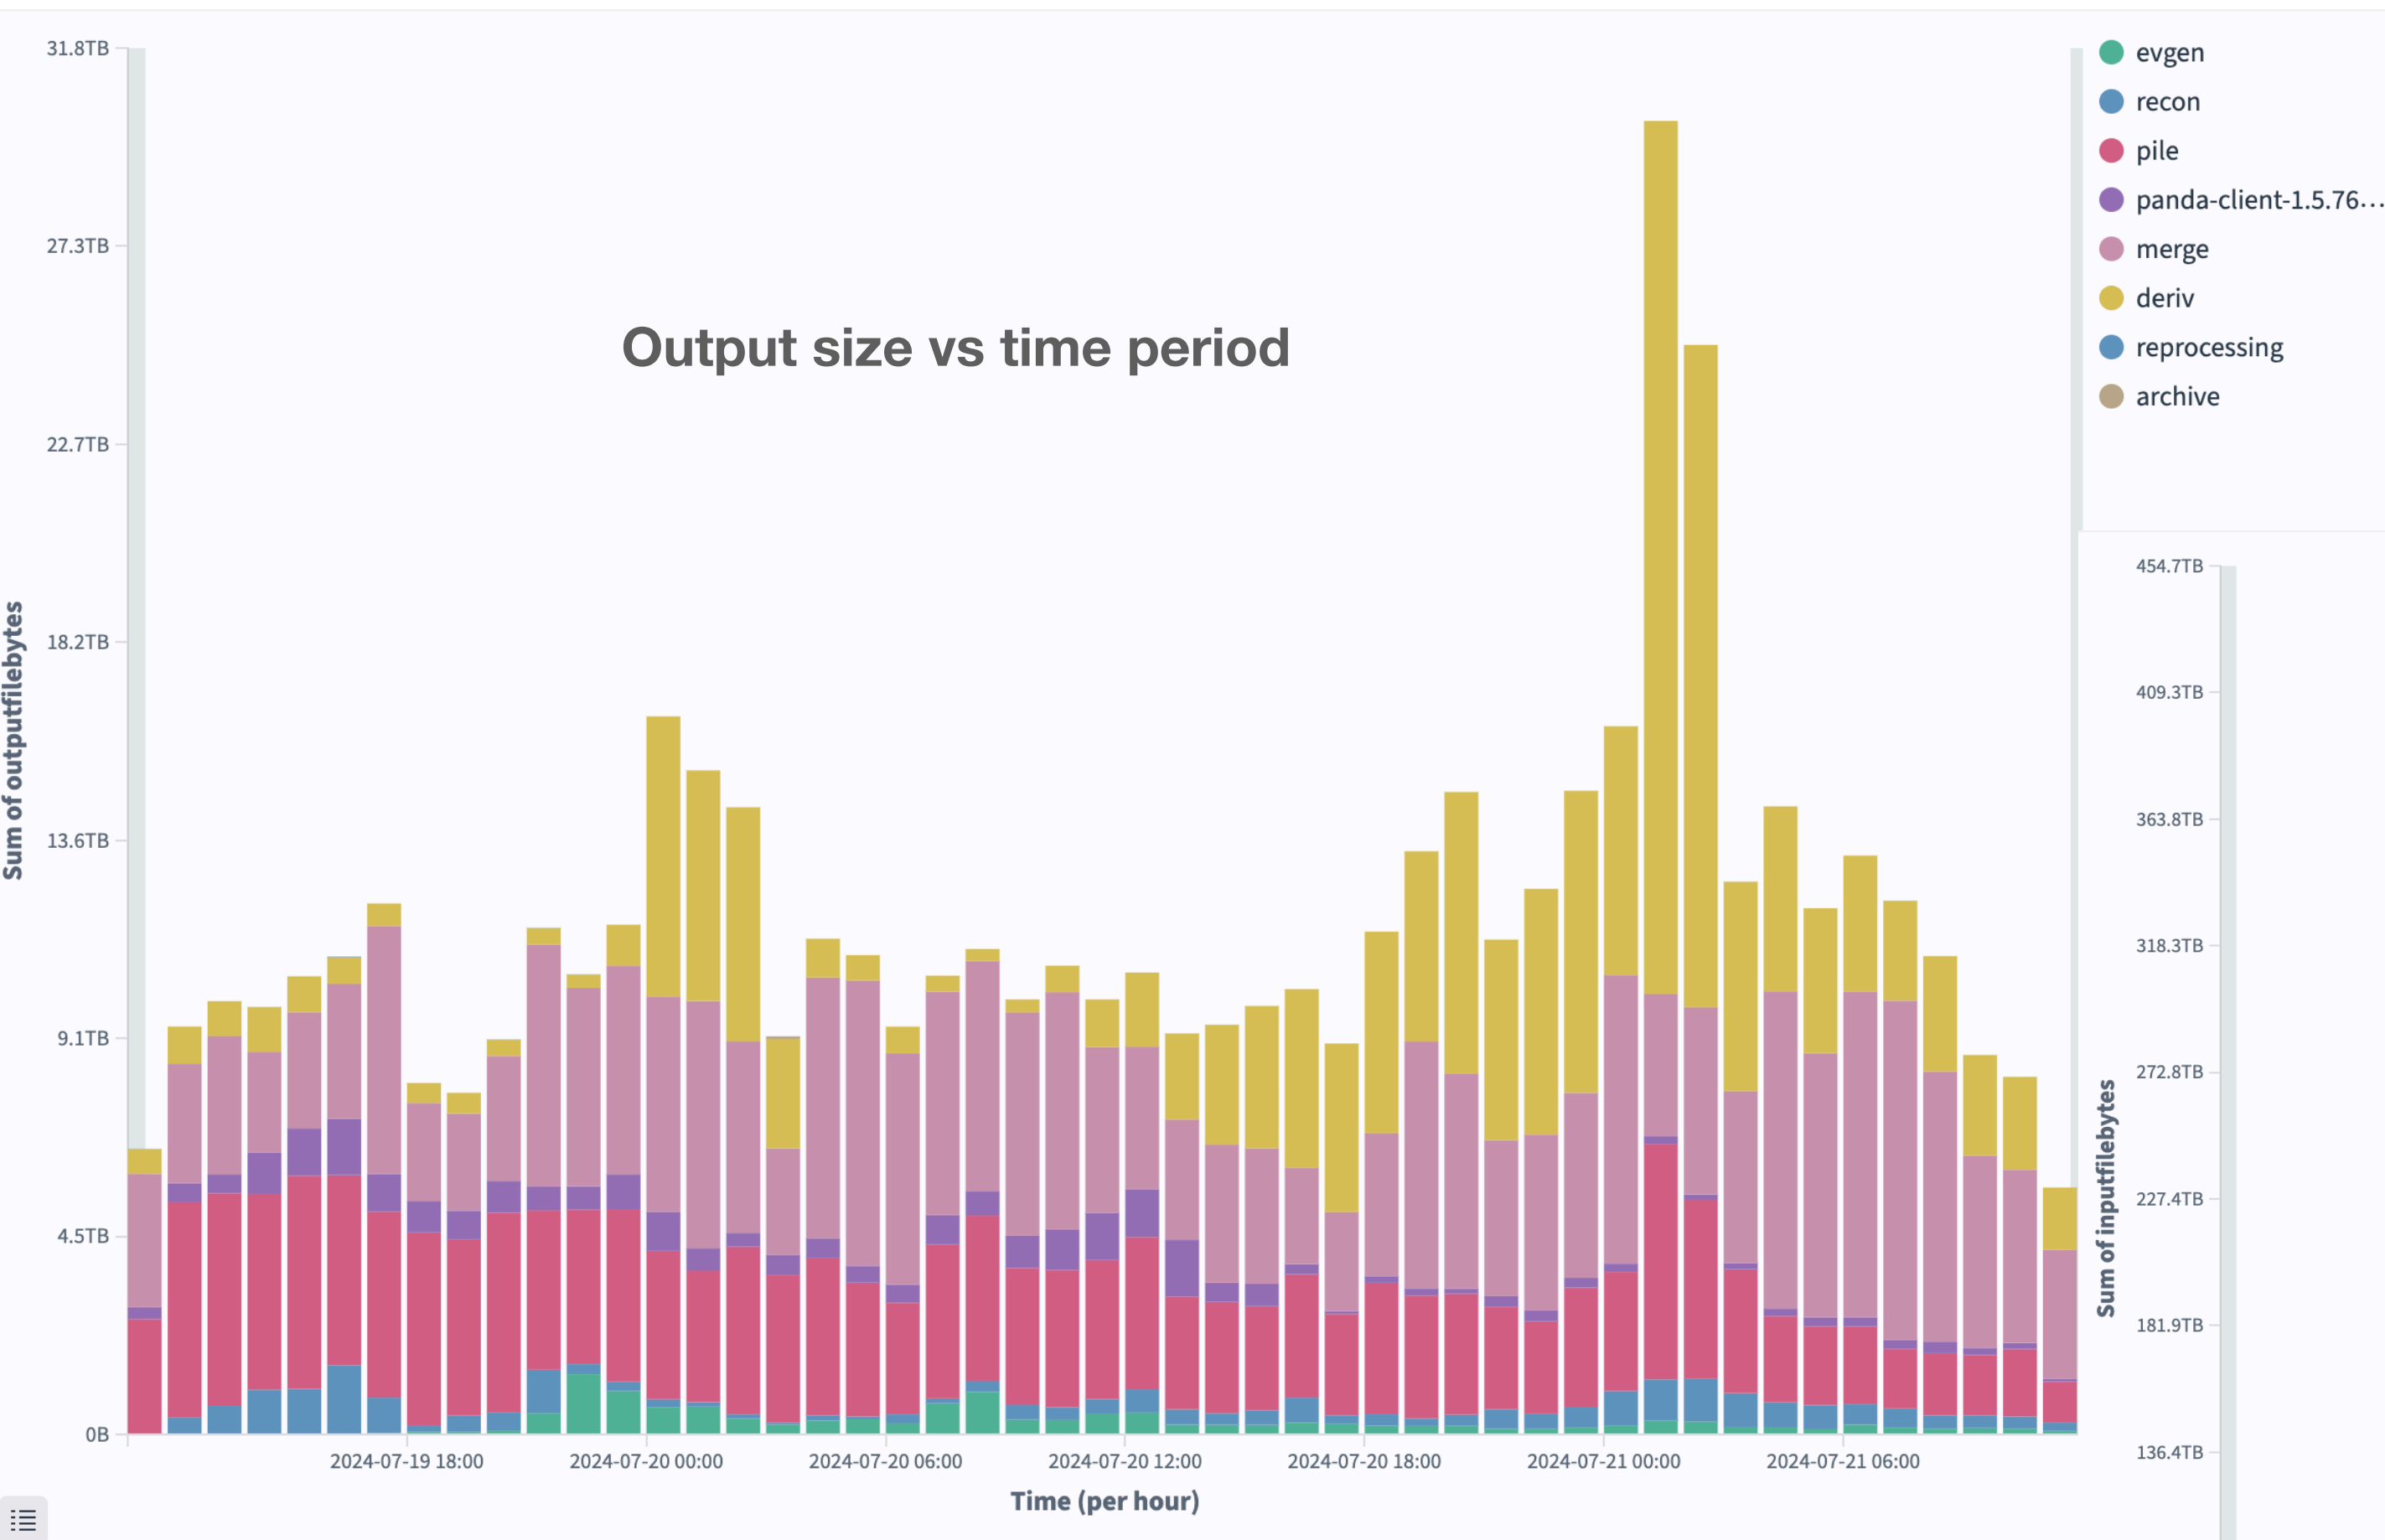
Nodes	Simul	Pile	Evgen	All
HDD/SSD	0.99	0.91	0.97	0.94
HDD/SSD	0.97	0.87	1.01	0.94

Remarks

- Efficiency remains constant with minor fluctuations parallel to scaling of drones.
 - Maximum resources are utilised by jobs at higher efficiencies.
- There is an effect on the CPU efficiency based on the job mix running on the cluster.
 - Higher IO intensive in the job mix has an effect on the CPU efficiencies of other job types.
 - Typically high IO jobs are of the type - pile, recon and analysis jobs from users.

	Simul			User jobs			Pile			Recon		
DAY	2	4	5	2	4	5	2	4	5	2	4	5
CPU Eff.	93 %	80 %	95 %	85 %	83 %	81 %	88 %	88 %	86 %	-	27 %	25 %
Avg. IO int.	~16k	~14k	~7.2k	~5M	~5.4M	~5.7M	~80k	~160k	~163k	-	~50k	40k
Total jobs	5k	11k	1.7k	~500	~400	43	3	300	1k	-	2k	2.6k

Job IO sizes



Computing clusters in Göttingen

- GoeGrid:
 - Has range of CPU clusters with both Intel Xenon and AMD EPYC.
 - Hyper threading is available and turned on as a default on all the computing nodes.
 - Turbo boost option available and turned on for Intel alone.

- Emmy:
 - Used compute node on Emmy is Intel Platinum 9242 @ 2.3 GHz with 96 physical cores.
 - Additional new phase is added with Intel Sapphire Rapids 8468 CPU (also used).
 - Hyper threading is available and turned on on as a default on all computing nodes (with choice to turn it off).
 - Turbo boost option available and turned off as default.

Benchmarking

- For Benchmarking, HEPscore23 software is used.
- All workloads include HEP applications from ATLAS, CMS, LHCb, ALICE and BELLE2.

- On EMMY:
 - With Hyper threading enabled: ~ 1900 HS23 on 192 cores, per thread: 9.89 HS23 (on all workloads) on Intel Platinum 9242 CPU.
 - With Hyper threading enabled: ~ 2900 HS23 on 192 cores, per thread: 15.1 HS23 (on all workloads) on Intel Sapphire rapids 8468 CPU.
 - The numbers on Emmy are extrapolated with an uncertainty of 2%.

- On GoeGrid:
 - With Hyper threading enabled: ~ 2380 HS23 on 256 cores, per thread: 9.29 HS23 (on all workloads).

Benchmarking

- All workloads include HEP applications from ATLAS, CMS, LHCb, ALICE and BELLE2.
- ATLAS
 - Generation of MC ttbar events.
 - Reconstruction of run 2 real data input.
- CMS
 - Generation and simulation of the ttbar events.
 - Reconstruction of ttbar events.
- LHCb
 - Generation and simulation of inclusive B events.
- ALICE
 - Digitisation and core reconstruction for ALICE Run 3 PbPb time frames.
- BELLE2
 - This workload executes event generation, detector and trigger simulation, and track reconstruction.

Benchmarking

- For Benchmarking, HEPsScore23 software is used.
- On EMMY:
 - Intel platinum 9242
 - With Hyper threading turned off: ~ 1500 HS23 on 96 physical cores.
 - With Hyper threading enabled: ~ 1900 HS23 on 192 cores.
 - With Hyper threading and Turbo boost: ~ 2260 HS23 on 192 cores.
 - Intel Sapphire rapids 8468
 - With Hyper threading: ~2900 HS23 on 192 cores.
 - The numbers on Emmy are extrapolated with an uncertainty of 2%.
- On GoeGrid:
 - With Hyper threading enabled: ~ 2380 HS23 on 256 cores.

Requirements for the integration

- HTCondor and COBaID/TARDIS are launched on the HPC cluster (Emmy) using containers.
- User namespaces are required to have stacked containers.
 - User namespaces is a linux feature that allows mapping users in the container to different users in the host.
 - It is useful for providing root access inside of a container.
- FUSE (Filesystem in Userspace) is required to install CVMFS as a user.
 - Unprivileged CVMFS client (CVMFS exec) is installed for the ATLAS jobs as the FUSE is unavailable on Emmy.