Folding@Home & Rosetta@Home on WLCG

Lukas Heinrich, Alessandra Forti, Mario Lassnig, Thomas Hartmann, Torsten Harenberg

Introduction

(*Folding*|*Rosetta*)@Home: Volunteer Computing Project (like LHC@home)

Science: Molecular Dynamics Simulation to understand Protein Folding

Non Structural Protein 10 (nsp10)

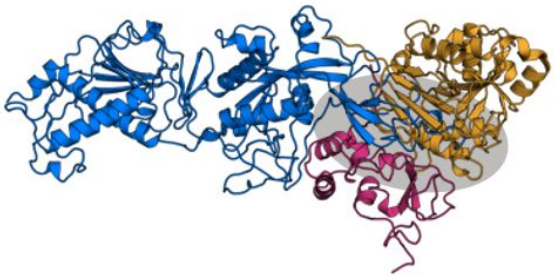

What does Nsp10 do?

Nsp10 helps to activate nsp14 and nsp16, two proteins important for viral replication. Studies from SARS show that disrupting this shared binding interface (shaded) decreases viral replication.

How is F@h helpful here?

We are looking for pockets that open, due to protein motions, where a drug could bind. Drug binding could inhibit the nsp10-nsp14/16 interaction and prevent viral replication.

Non Structural Protein 14 (nsp14)

Non Structural Protein 16 (nsp16)

Bouvet M, and Lugari A, et al. Coronavirus Nsp10, a critical co-factor for activation of multiple replicative enzymes. *J Biol Chem.* 2014;289(37):25783–25796. doi:10.1074/jbc.M114.577353

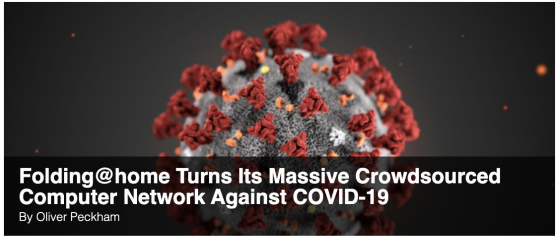

March 16, 2020

For gamers, fighting against a global crisis is usually pure fantasy – but now, it's looking more like a reality. As supercomputers around the world spin up to combat the coronavirus pandemic, the crowdsourced distributed computing platform Folding@home is setting its sights on coronavirus research, spurring a global movement to commit powerful home computers and gaming consoles to the cause.

Rosetta@home

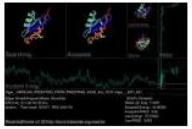

Rosetta@home is a distributed computing project for protein structure prediction on the Berkeley Open Infrastructure for Network Computing platform, run by the Baker laboratory at the University of Washington. Wikipedia

```
prun --containerImage /cvmfs/unpacked.cern.ch/registry.hub.docker.com/lukasheinrich/folding:latest \
--exec 'FAHClient --user=ANALY_MANC_GPU --team=38188 --gpu=true --smp=false' \
--outDS user.lheinric.folding.test.01
```

Proof-of-Principle: running F@H on WLCG resources

Nice side-effect: test new capabilities

- user-defined containers
- cvmfs-based image distribution
- GPU resources on the Grid
- same containers run on CPUs and could be used to run jobs using GridPP DIRAC

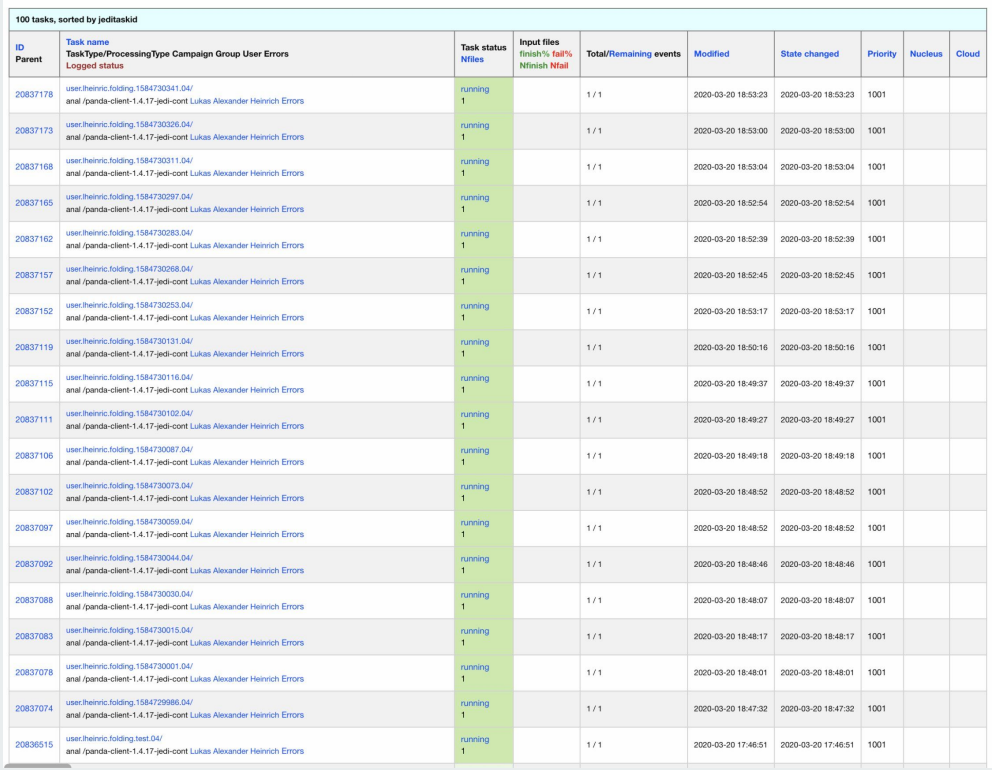

100 tasks, sorted by jeditaskid

ID Parent	Task name TaskType/ProcessingType Campaign Group User Errors Logged status	Task status Nfiles	Input files finish% fail% Nfinish Nfail	Total/Remaining events	Modified	State changed	Priority	Nucleus	Cloud
20837178	user.lheinric.folding.1584730341.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:53:23	2020-03-20 18:53:23	1001		
20837173	user.lheinric.folding.1584730326.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:53:00	2020-03-20 18:53:00	1001		
20837168	user.lheinric.folding.1584730311.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:53:04	2020-03-20 18:53:04	1001		
20837165	user.lheinric.folding.1584730297.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:52:54	2020-03-20 18:52:54	1001		
20837162	user.lheinric.folding.1584730283.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:52:39	2020-03-20 18:52:39	1001		
20837157	user.lheinric.folding.1584730268.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:52:45	2020-03-20 18:52:45	1001		
20837152	user.lheinric.folding.1584730253.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:53:17	2020-03-20 18:53:17	1001		
20837119	user.lheinric.folding.1584730131.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:50:16	2020-03-20 18:50:16	1001		
20837115	user.lheinric.folding.1584730116.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:49:37	2020-03-20 18:49:37	1001		
20837111	user.lheinric.folding.1584730102.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:49:27	2020-03-20 18:49:27	1001		
20837106	user.lheinric.folding.1584730087.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:49:18	2020-03-20 18:49:18	1001		
20837102	user.lheinric.folding.1584730073.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:48:52	2020-03-20 18:48:52	1001		
20837097	user.lheinric.folding.1584730059.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:48:52	2020-03-20 18:48:52	1001		
20837092	user.lheinric.folding.1584730044.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:48:46	2020-03-20 18:48:46	1001		
20837088	user.lheinric.folding.1584730030.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:48:07	2020-03-20 18:48:07	1001		
20837083	user.lheinric.folding.1584730015.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:48:17	2020-03-20 18:48:17	1001		
20837078	user.lheinric.folding.1584730001.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:48:01	2020-03-20 18:48:01	1001		
20837074	user.lheinric.folding.1584729986.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 18:47:32	2020-03-20 18:47:32	1001		
20836515	user.lheinric.folding.test.04/ anal /panda-client-1.4.17-jedi-cont Lukas Alexander Heinrich Errors	running 1		1 / 1	2020-03-20 17:46:51	2020-03-20 17:46:51	1001		

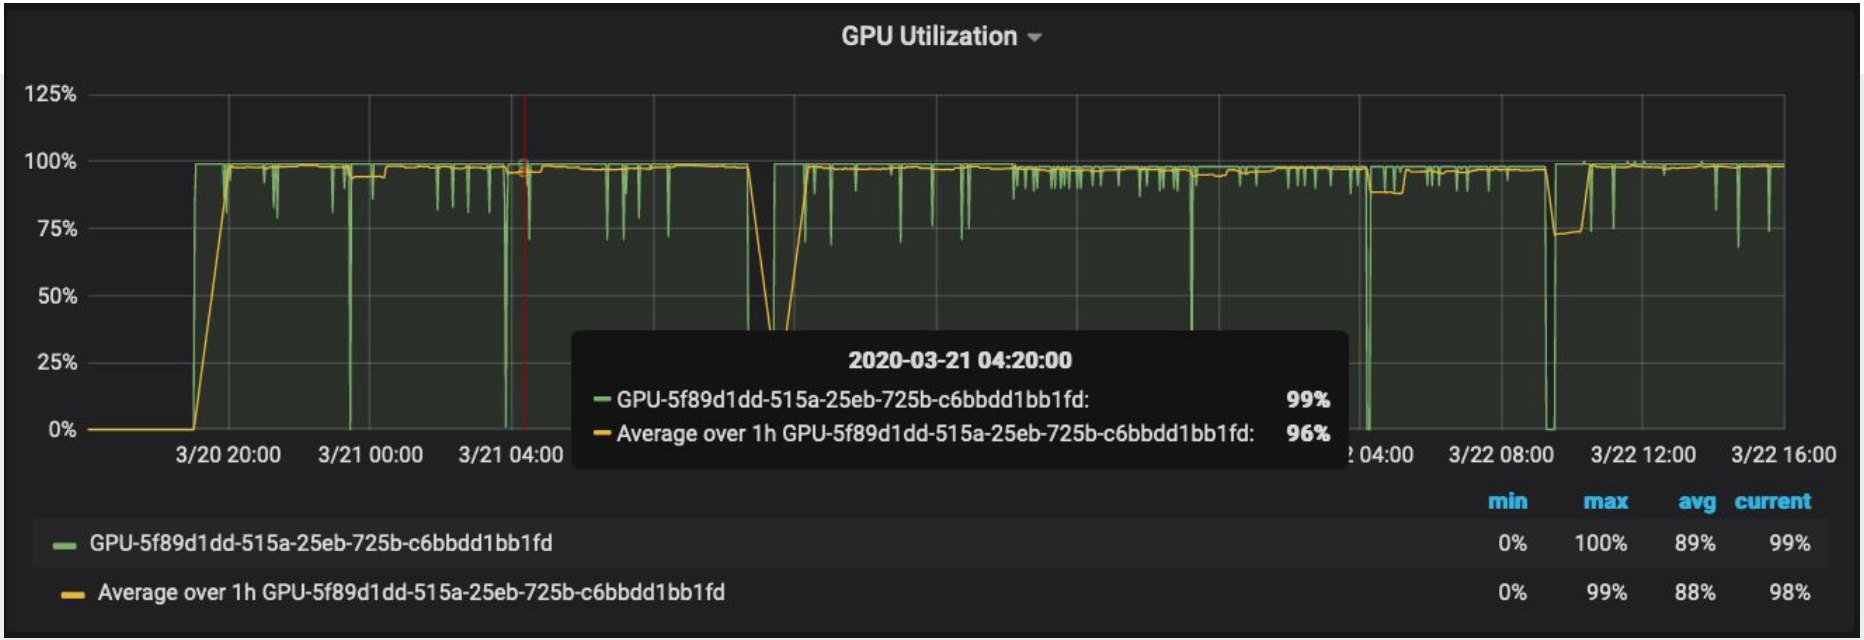
GPU Utilization
125%
100%
75%
50%
25%
0%
3/20 20:00
3/21 00:00
3/21 04:00
3/22 08:00
3/22 12:00
3/22 16:00
2020-03-21 04:20:00
GPU-5f89d1dd-515a-25eb-725b-c6bbdd1bb1fd: 99%
Average over 1h GPU-5f89d1dd-515a-25eb-725b-c6bbdd1bb1fd: 96%
min max avg current
GPU-5f89d1dd-515a-25eb-725b-c6bbdd1bb1fd 0% 100% 89% 99%
Average over 1h GPU-5f89d1dd-515a-25eb-725b-c6bbdd1bb1fd 0% 99% 88% 98%

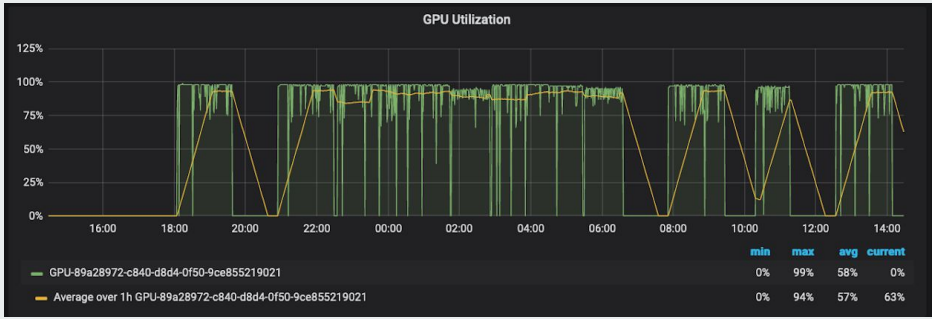
GPU Utilization
125%
100%
75%
50%
25%
0%
16:00
18:00
20:00
22:00
00:00
02:00
04:00
06:00
08:00
10:00
12:00
14:00
min max avg current
GPU-89a28972-c840-d8d4-0f50-9ce855219021 0% 99% 58% 0%
Average over 1h GPU-89a28972-c840-d8d4-0f50-9ce855219021 0% 94% 57% 63%

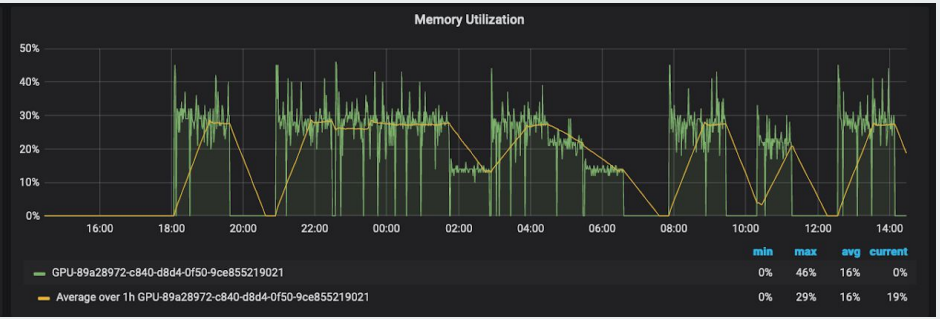
Memory Utilization
50%
40%
30%
20%
10%
0%
16:00
18:00
20:00
22:00
00:00
02:00
04:00
06:00
08:00
10:00
12:00
14:00
min max avg current
GPU-89a28972-c840-d8d4-0f50-9ce855219021 0% 46% 16% 0%
Average over 1h GPU-89a28972-c840-d8d4-0f50-9ce855219021 0% 29% 16% 19%

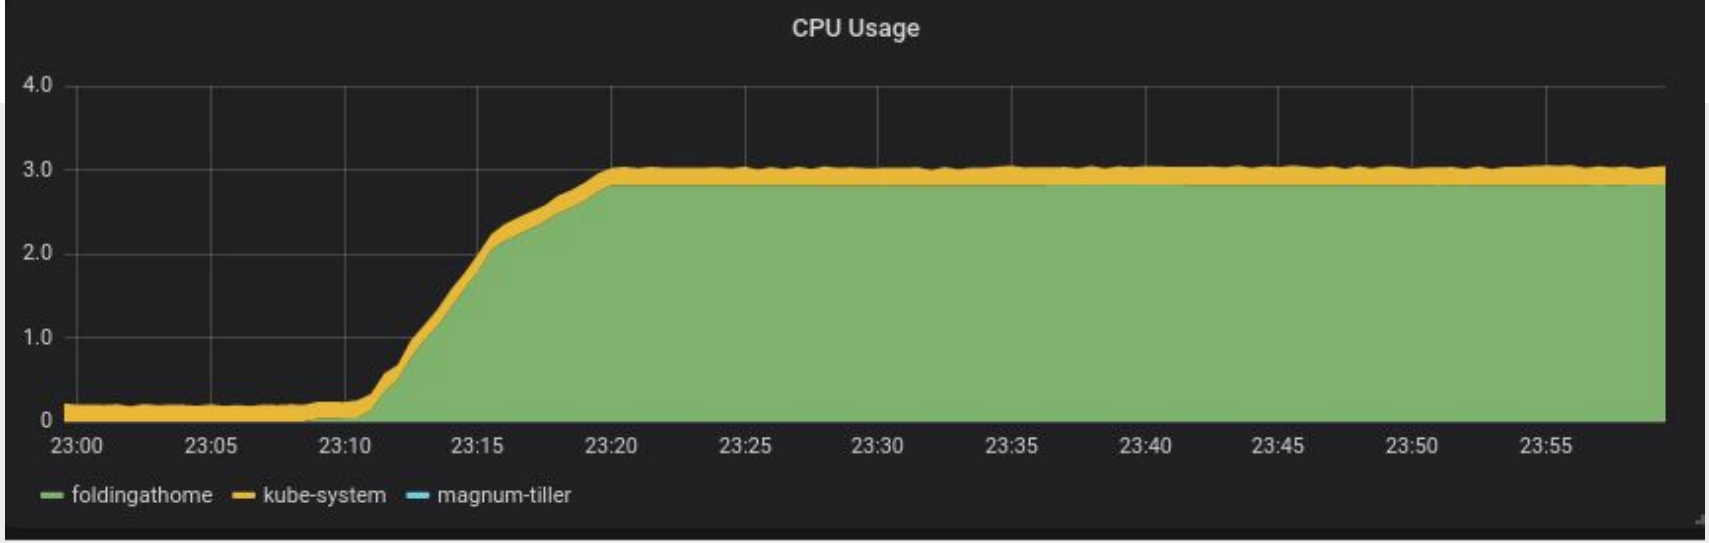
CPU Usage
4.0
3.0
2.0
1.0
0
23:00
23:05
23:10
23:15
23:20
23:25
23:30
23:35
23:40
23:45
23:50
23:55
foldingathome
kube-system
magnum-tiller

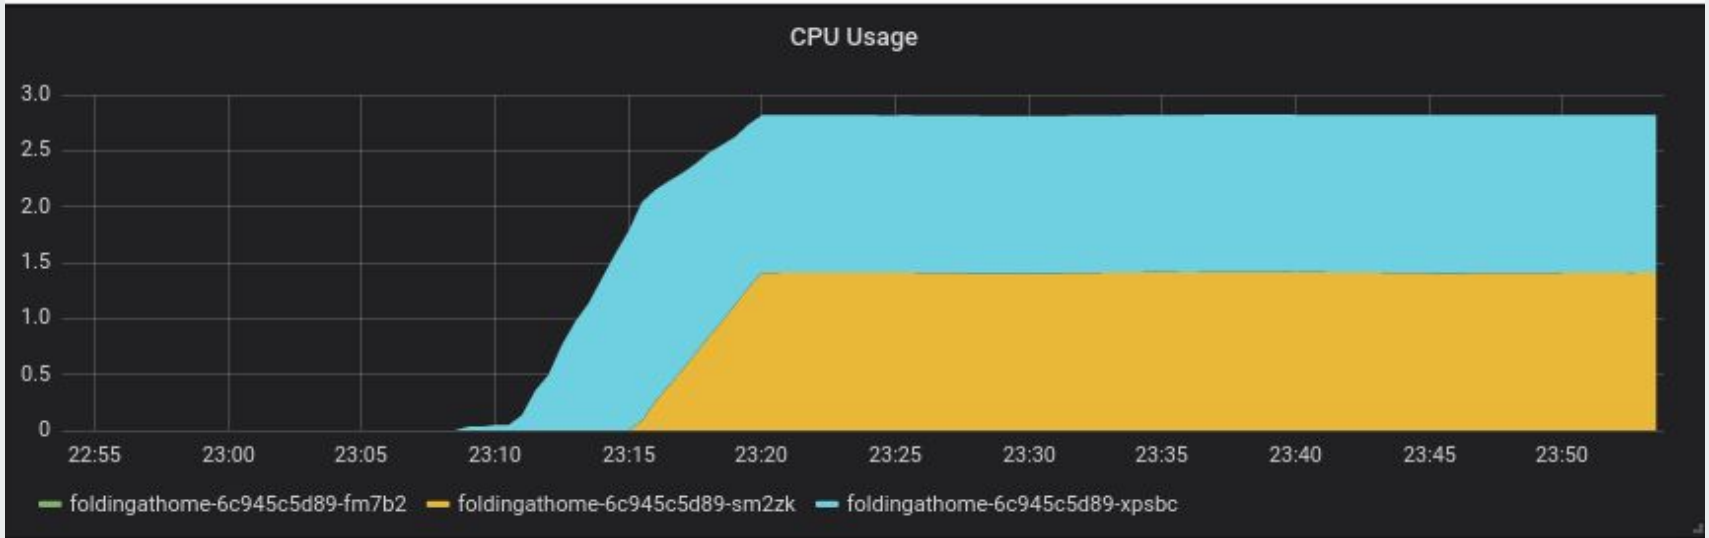
CPU Usage
3.0
2.5
2.0
1.5
1.0
0.5
0
22:55
23:00
23:05
23:10
23:15
23:20
23:25
23:30
23:35
23:40
23:45
23:50
foldingathome-6c945c5d89-fm7b2
foldingathome-6c945c5d89-sm2zk
foldingathome-6c945c5d89-xpsbc

Call with F@H team

There seem to be three workloads:

- Running WU servers that distribute work units (needs large storage)
  - Needed to keep the large # of volunteer clients occupied
  - Identified as bottleneck
- Running non-volunteer workloads (“adaptive sampling”)
  - Current code uses classic SLURM submissions
  - Waiting for code to be released
- Running clients
  - Given exaFLOP capability of current volunteers perhaps lowest priority

Team Monthly | Team | Donor | OS Stats

Team Statistics for March, 2020

Previous Month | Next Month

Name: is exactly | CERN

Team #: | Search

Rank	Name	Credit	WUs	Team	Change
	Name is 'CERN'				
689	CERN	23,780,396	7,757	38188	

Team Monthly | Team | Donor | OS Stats

Team: CERN

Date of last work unit	2020-03-28 01:25:59
Active CPUs within 50 days	2,270
Team Id	38188
Grand Score	23,780,401
Work Unit Count	7,757
Team Ranking	2041 of 244560
Homepage	http://public.web.cern.ch/public/

Team members

Rank	Name	Credit	WUs
7	Anonymous	7,906,608	1,198
32,272	TheLaboratoire	6,362,987	150
41,377	CERN_Cloud	4,222,683	5,208
60,094	Cloverfield	2,237,882	42
57,839	Shaba-kun	748,008	10
142,686	ANALY_MANC_UCORE	455,910	635
150,009	Corne_Lukken	407,266	31
178,916	ANALY_MANC_GPU_TEST	301,525	37
195,211	Jarek	238,659	26
187,185	lucagardi	189,781	28
231,752	CERN	159,047	277
266,488	P.Mogg	102,784	9
N/A	szak1	95,983	21
N/A	Plam	92,609	9
N/A	ANALY_MWT2_GPU	54,016	6
N/A	gogidecker	38,281	9
177,787	b2ag	25,145	4
N/A	Jackalo_Moss	19,350	2
N/A	ANALY_INFN-T1_GPU	18,810	2
N/A	OneArmedRobbery	17,090	2
N/A	Dommy	15,810	3
N/A	WizeDom	15,802	11
N/A	Clovefield	15,216	2
N/A	Pouet	10,108	3
N/A	ANALY_QMUL_GPU_TEST	9,553	3
N/A	Nantho	7,894	2
N/A	alphacc	3,581	15
8,331	Jimmy	2,110	3
44,286	ff	2,043	3
47,360	YanNick	2,000	2
N/A	ogsv	1,000	1
N/A	kikileplus	420	1
N/A	Tquast	220	1
N/A	ANALY_FZK	220	1

GridPP experience

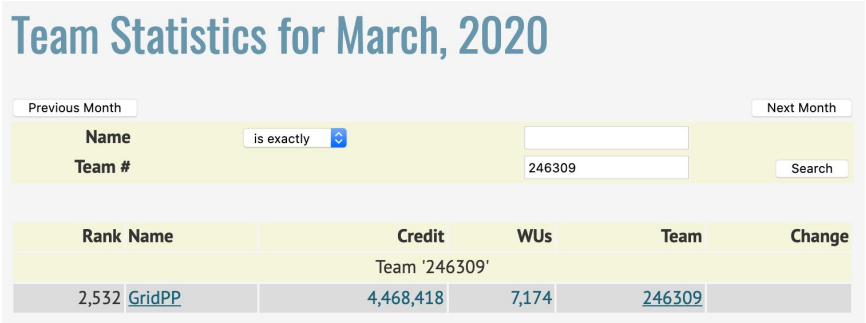

- UK community DIRAC
  - Same containerized FAH client used on ATLAS GPU queues
  - Advantage all sites run without having to do anything

```
Arguments = "--containerImage /cvmfs/unpacked.cern.ch/registry.hub.docker.com/lukasheinrich/folding:latest \
-p 'FAHClient --user=GridPP.UKI-LT2-IC-HEP --team=246309 --gpu=false --cpus=8 --smp=true --max-units=3 --exit-when-done";
OutputData = "LFN:/gridpp/user/a/alessandra.forti/fah_test/fah_gridpp_output_%j.tar"
```

- Glasgow dedicated 512 cores off the grid
  - Same client installed with an rpm
  - Machines weren't in production yet
- **Containerized payloads** key for quick deployment on production resources

Rosetta@Home configuration

- Standard BOINC client (*like ATLAS@Home*) running protein science project payloads
- Understanding BOINC particularities and dynamically generating configuration files
  - **1-core/8-core** Client limits tailored for main WLCG node core configuration
  - **Desktop** Reasonable backfill on user desktops and login nodes
  - **Custom** Allowing for own CPU, memory & disk settings
- Many things hard-coded - some polish needed

Rosetta@Home workloads

- Running jobs since last week using Singularity container
  - **Singularity container:** /cvmfs/grid.desy.de/container/boinc.d
  - **Documentation:** https://confluence.desy.de/display/grid/BOINC+user+computing
- Running jobs since last week
  - Reassigned a handful out-of-warranty batch nodes @ DESY
  - Login nodes @ UWuppertal
  - A few user desktops @ DESY
  - O(100) MCORE jobs in the batch system @ UWuppertal
- Ongoing
  - Investigating backfilling of the Maxwell HPC @ DESY
  - Draining further out-of-warranty batch nodes for reassignment