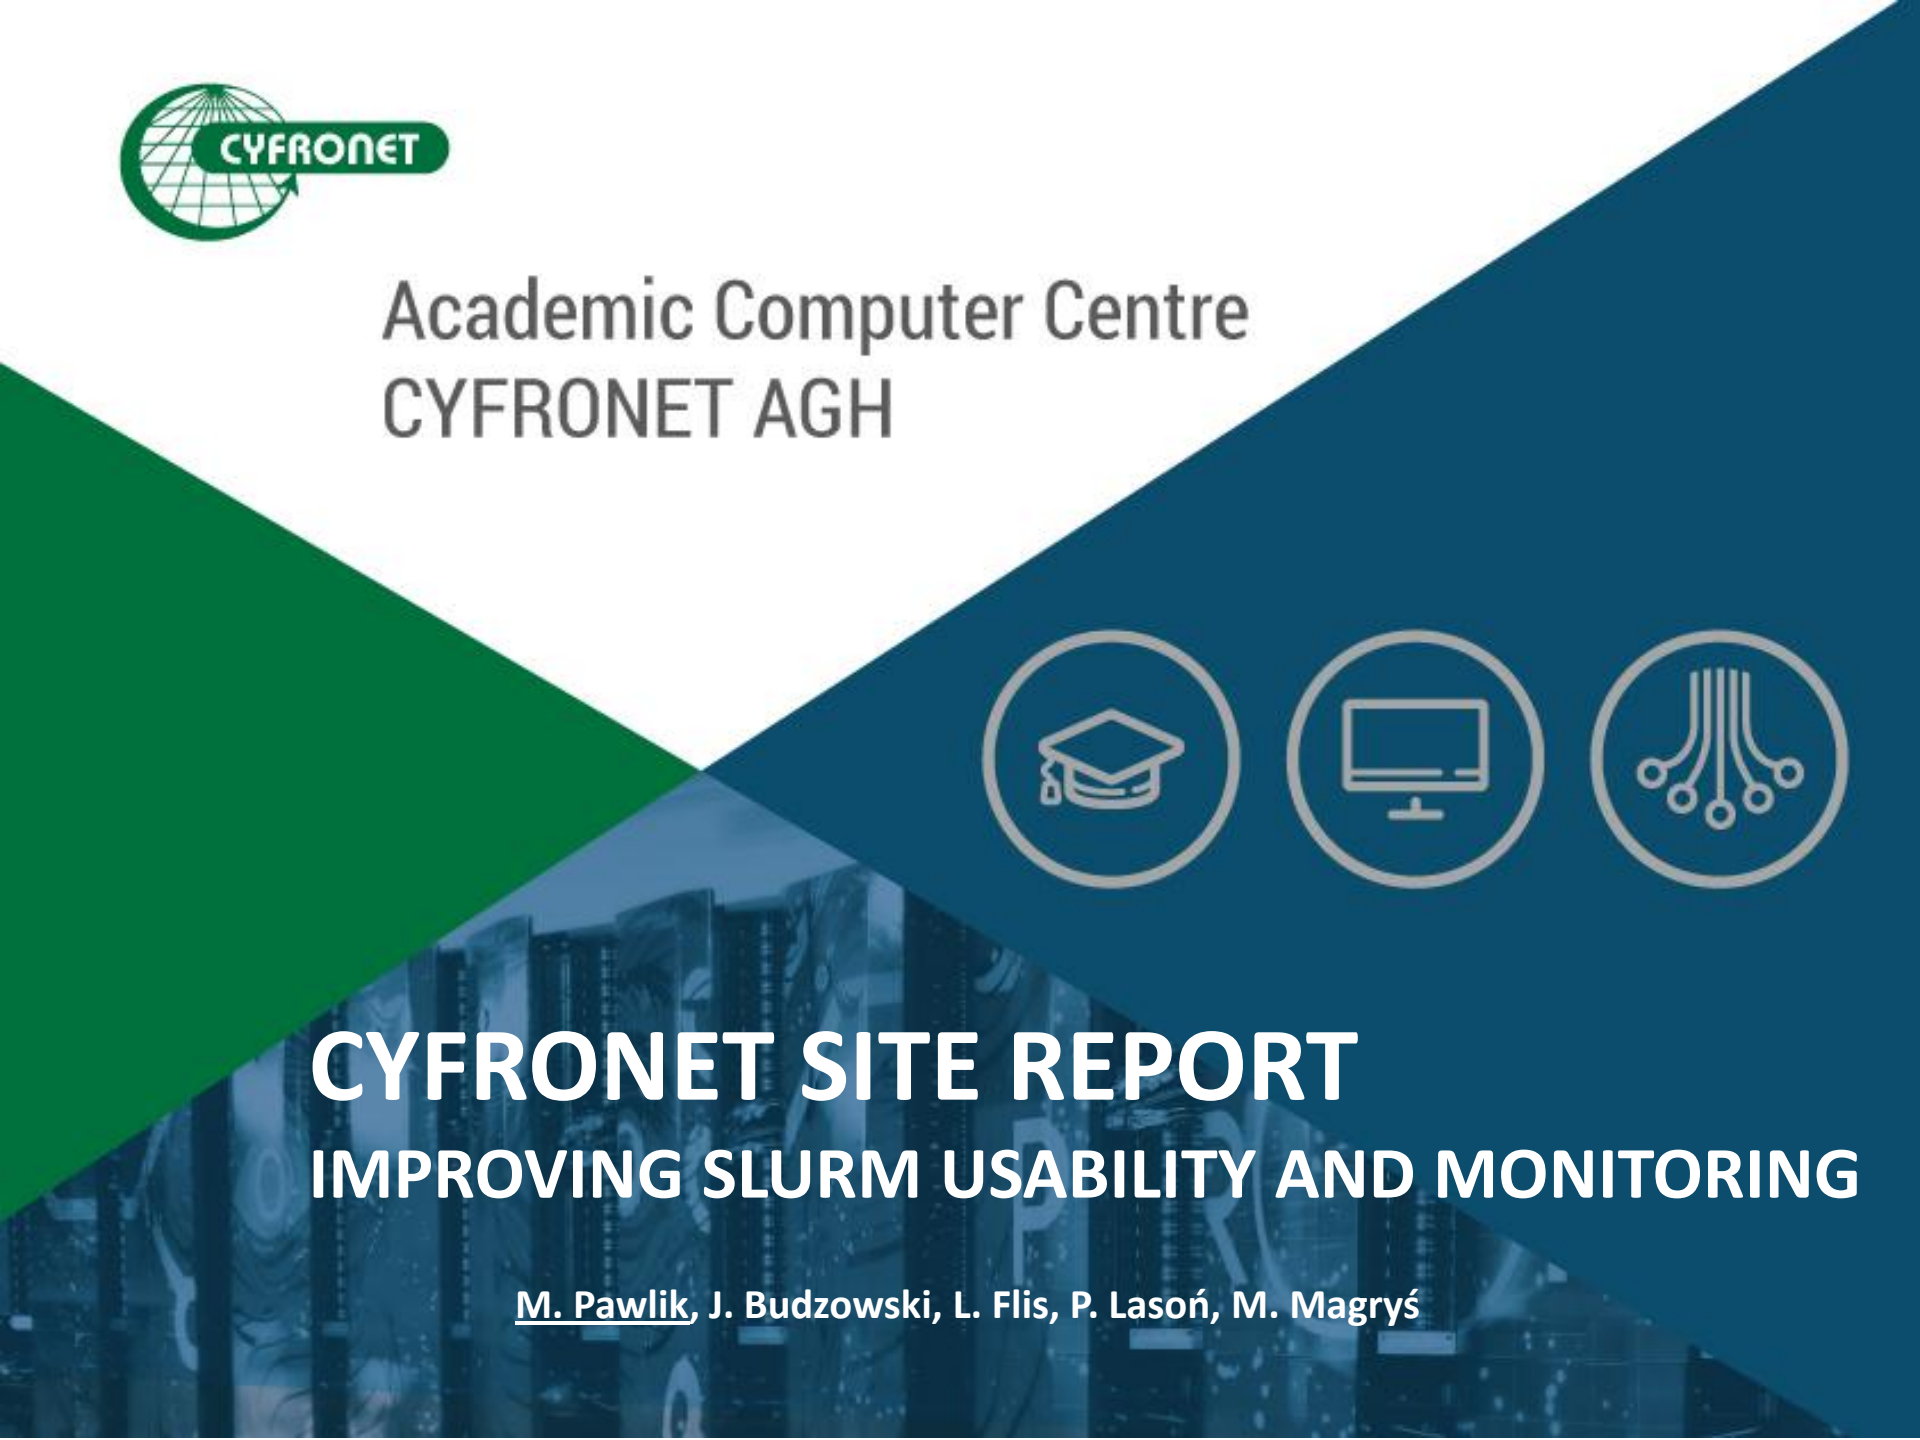

CYFRONET SITE REPORT

IMPROVING SLURM USABILITY AND MONITORING

M. Pawlik, J. Budzowski, L. Flis, P. Lasoń, M. Magryś

Presentation plan

- Cyfronet introduction
- System description
- SLURM modifications
- Job information scripts
- Monitoring


CYFRONET

- established in **1973**
- part of **AGH University of Science and Technology** in Krakow, Poland
- provides **free** computing resources for scientific institutions
- center of competence in **HPC** and **Grid** Computing
- member of **PIONIER National Research and Education Network** and operator of **Krakow Metropolitan Area Network** for research and education
- participants of large EU projects:
- member of international collaborations:

- Polish national IT infrastructure supporting e-Science
  - based upon resources of **most powerful academic resource centers**
  - compatible and interoperable with European Grid
  - offering **grid and cloud computing paradigms**
  - coordinated by Cyfronet
- Benefits for users
  - unified infrastructure from 5 separate compute centres
  - unified access to software, compute and storage resources
  - non-trivial **quality of service**
- Challenges
  - unified monitoring, computing grants, accounting, security
  - create environment of **cooperation rather than competition**
- Federation – the key to success

Zeus (older system)

- 374 TFlops
- Several times on top500
- Repurpose:
  - Torque/Moab -> SLURM
  - Cloud services

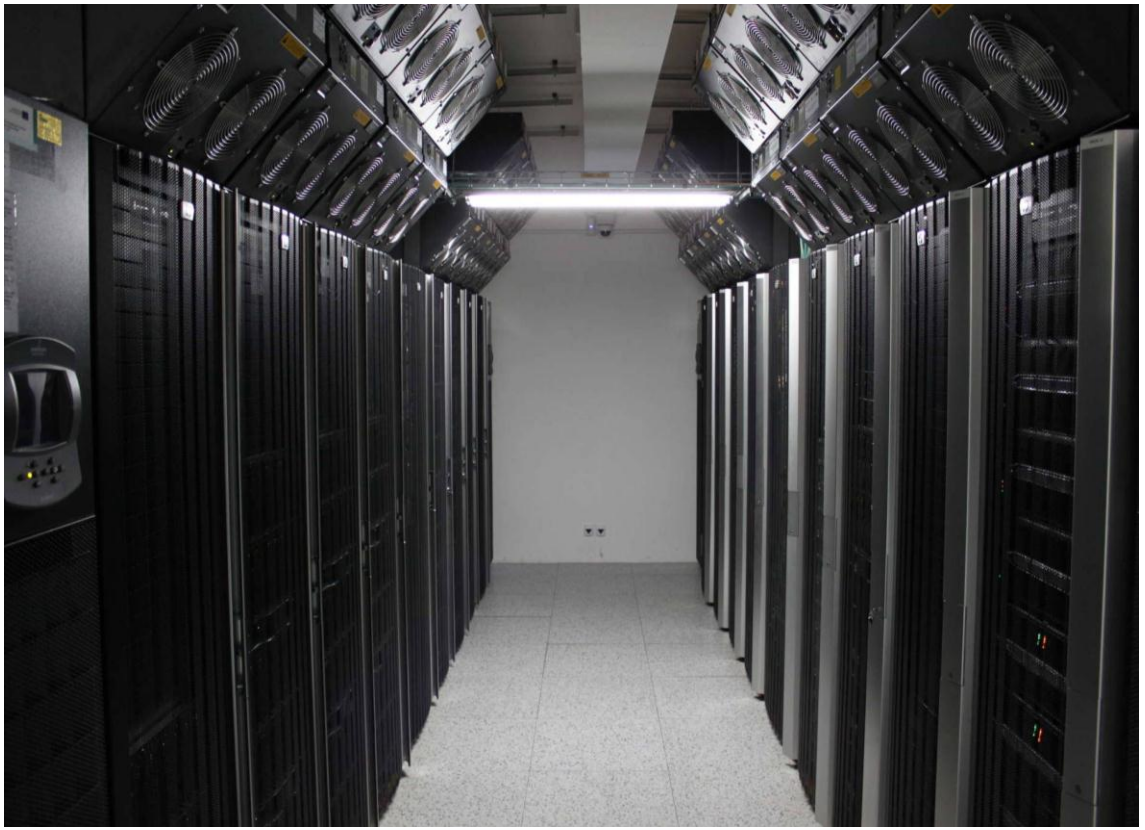

2.4 PFLOPS, #38 TOP500, #72 GREEN 500

Prometheus Cluster

- Installed in Q4 2015
- Centos 7 + SLURM 17.02
- HP Apollo 8000
  - 20 racks (4 CDU, 16 compute)
- 2232 nodes, 53568 CPU cores (Haswell), 279 TB RAM
  - 2160 regular nodes (2 CPUs, 128 GB RAM)
  - 72 nodes with GPGPUs (2x NVIDIA Tesla K40 XL)
  - 4 islands
- 2.4 PFLOPS total performance (Rpeak)
  - 2140 TFLOPS in CPUs
  - 256 TFLOPS in GPUs
- <850 kW power (including cooling)

- Infiniband 4x FDR (56Gb/s)
- Diskless nodes
  - Improves realiability
- Lustre FS as main storage:
  - Scratch: 5 PB @ 120 GB/s
  - Archive: 5 PB @ 60 GB/s
- NFS for:
  - $HOME dirs
  - software

Prometheus Island

Secondary loop

- Academic workload
  - Lots of small/medium jobs
  - Few big jobs
- 330 projects
- 750 users
- Main fields:
  - Chemistry
  - Biochemistry (farmaceuticals)
  - Astrophysics

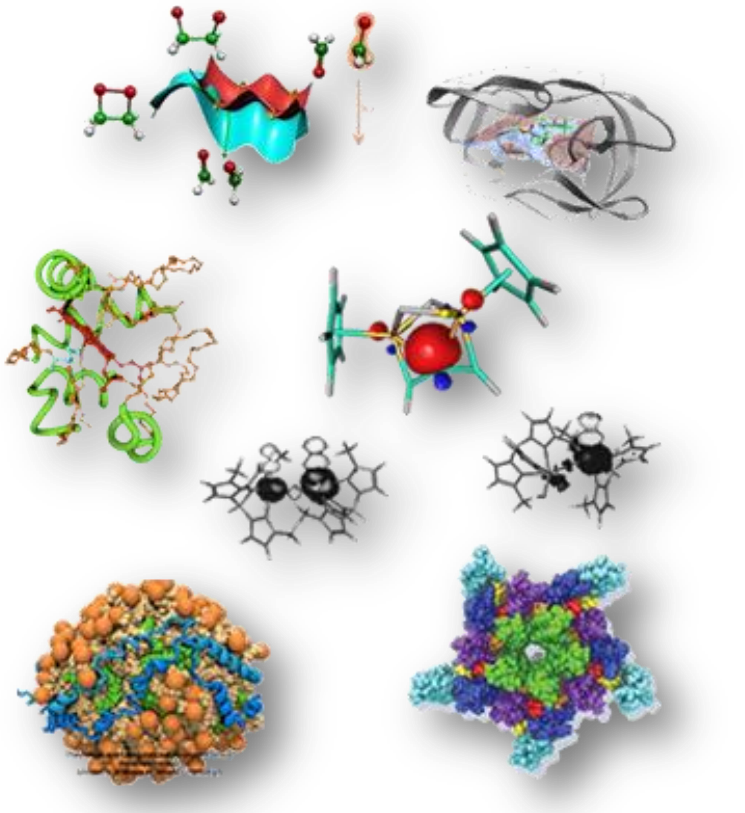

- Really happy with it, openness, community
- Power saving
  - Full shutdown/bootup instead of suspend/resume
  - Don’t power down „downed” nodes
  - Patched some race conditions in slurmctld (deadlock during config read, fix coming in 17.11)
- Proper handling of longer account names (>20 chars)
- Kmem patch – cgroups accounted for kmem (task/cgroup)

- Integration with PL-Grid:
  - SLAs import (sacctmgr)
    - SLA translates to limits/FS/priority
  - Accounting data reports (sacct)

- Deeper FS tree
- No static resource allocations

- Account names have to be unique:
  - Use „domain like” names:
    grid.lhc (FS:10)
      grid.lhc.atlas (FS:5)
        grid.lhc.atlas.prd
        grid.lhc.atlas.sgm
      grid.lhc.alice (FS:5)
- Long account names are a challenge:
  - Display in command line tools

- Squeue/sstat/scontrol/sacct *(arguments)*
- User centric scripts (wrappers), important information at a glance:
  - Pro-jobs - display information about running jobs
  - Pro-jobs-history - display information about past jobs
- Support for:
  - Basic filtering
  - Sorting
  - etc…

Pro-jobs


```
[prometheus][plgszaleniec@login01 ~]$ pro-jobs
ID          Partition       Name     State   Nodes   Cores   Decl. mem   Max. node mem.   Mem. % usage   Eff.     Walltime
--          ---------       ----     -----   -----   -----   ---------   --------------   ------------   ----     --------
7293498        plgrid   Mo_pterin   RUNNING       1      24     48.0GiB    2.1GiB(p0855)           4.4%   99.4%   2-08:11:45
7296097        plgrid   Mo_pterin   RUNNING       1      24     48.0GiB    2.1GiB(p1654)           4.4%   99.4%     12:41:01
================================================================================================================================
Statistical data of jobs which has not already ended is not always correct.
To get more accurate statistics you need to wait till job's completion and then use 'pro-jobs-history' command.
================================================================================================================================

[prometheus][plgszaleniec@login01 ~]$
```

Pro-jobs-history


```
[prometheus][plgszaleniec@login01 ~]$ pro-jobs-history
ID                    Name      Partition   Nodes  Cores  GPUs  Decl. mem  Mem. % usage  Eff.   CPU. used     GPUtime [h]  Wall. Used  Wall. Req.  End Time
--                    ----      ---------   -----  -----  ----  ---------  ------------  ----   ---------     -----------  ----------  ----------  --------
7272137               Mo_pterin plgrid-short    1     24     0   48.0GiB          3.3%  99.3%  04:56:00             --     00:12:20    01:00:00  2017-09-19 15:25:00
7272764               Mo_pterin plgrid-short    1     24     0   48.0GiB          3.3%  99.6%  15:12:24             --     00:38:01    01:00:00  2017-09-19 17:26:11
7273740               Mo_pterin plgrid-short    1     24     0   48.0GiB          0.0%  22.3%  00:00:24             --     00:00:01    01:00:00  2017-09-19 20:02:22
7274536               Mo_pterin plgrid-short    1     24     0   48.0GiB          4.5%  98.2%  1-00:06:24           --     01:00:16    01:00:00  2017-09-19 23:09:27
7274963               Mo_pterin       plgrid    1     24     0   48.0GiB          4.6%  98.3%  1-12:33:12           --     01:31:23 1-00:00:00  2017-09-20 01:21:34
7278899     BSS.Amber.14.gpu     plgrid-gpu    1     24     2  120.0GiB          0.0%   7.3%  00:10:00           0.01     00:00:25 3-00:00:00  2017-09-20 16:01:18
7278983     BSS.Amber.14.gpu     plgrid-gpu    1     24     2  120.0GiB          0.0%   7.3%  00:10:00           0.01     00:00:25 3-00:00:00  2017-09-20 16:43:40
7279350     BSS.Amber.14.gpu     plgrid-gpu    1     24     2  120.0GiB          0.0%   7.3%  00:12:48           0.02     00:00:32 3-00:00:00  2017-09-20 20:02:50
7278884               Mo_pterin       plgrid    1     24     0   48.0GiB         94.0%  99.4%  11-17:54:24          --     11:44:46 2-00:00:00  2017-09-21 03:27:19
7276154               Mo_pterin       plgrid    1     24     0   48.0GiB          4.9%  97.9%  25-15:11:12          --  1-01:37:58 2-00:00:00  2017-09-21 09:44:26
7282602               Mo_pterin       plgrid    1     24     0   48.0GiB          0.0%  27.6%  00:01:12             --     00:00:03 2-00:00:00  2017-09-21 09:59:05
7286001               Mo_pterin       plgrid    1     24     0   48.0GiB          4.8%  98.0%  24-05:12:00          --  1-00:13:00 2-00:00:00  2017-09-22 12:52:09
7291593            I_Mo_pterin       plgrid    1     24     0   48.0GiB         50.7%  99.4%  8-11:57:12           --     08:29:53 3-00:00:00  2017-09-22 22:23:14
7290756               Mo_pterin       plgrid    1     24     0   48.0GiB          4.4%  99.4%  48-00:00:48          --  2-00:00:02 2-00:00:00  2017-09-23 21:20:08
7296017               Mo_pterin       plgrid    1     24     0   48.0GiB          0.0%   3.2%  00:00:48             --     00:00:02 3-00:00:00  2017-09-24 16:14:27
[prometheus][plgszaleniec@login01 ~]$
```

- Set of services gathers data from SLURM and feeds it to Graphite/Redis monitoring system

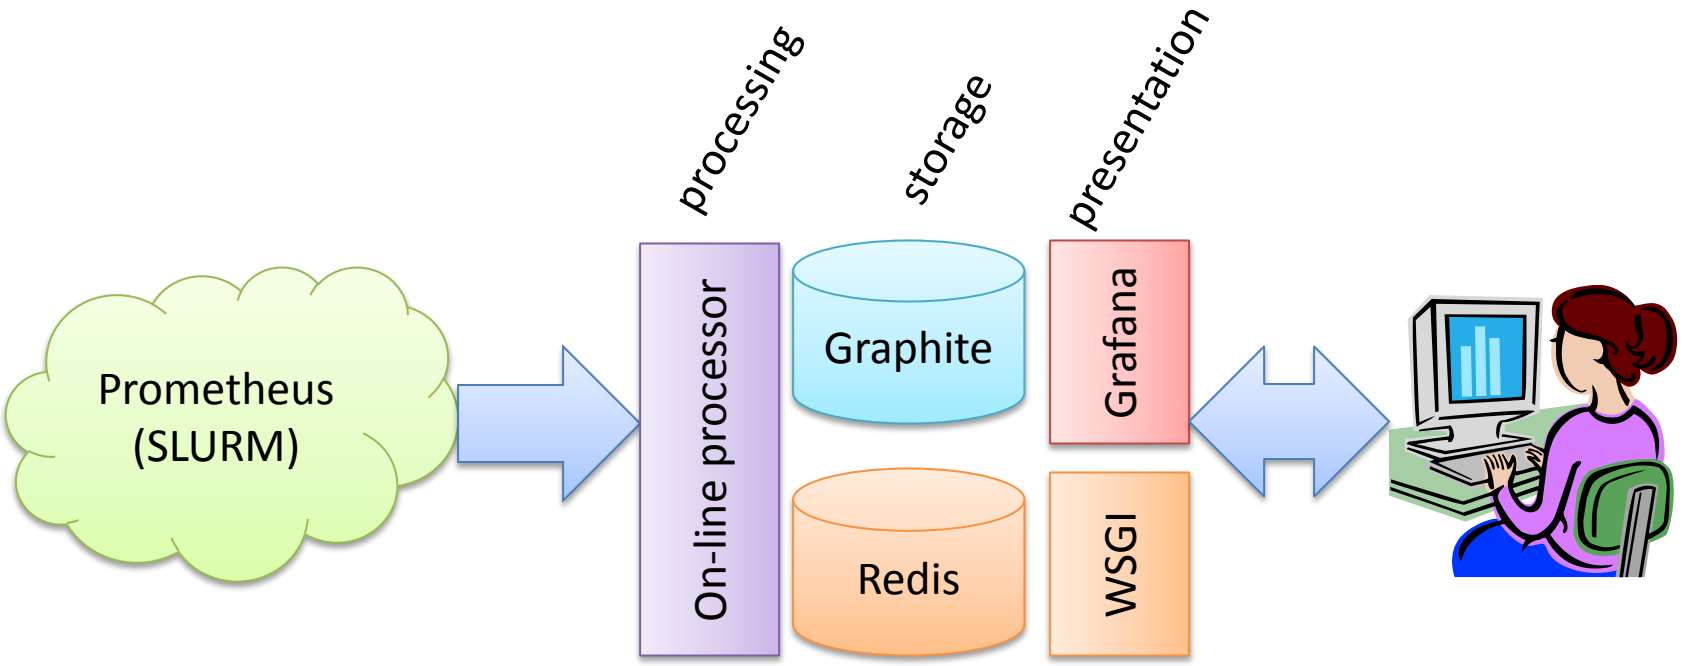

Monitoring integration and dashboards

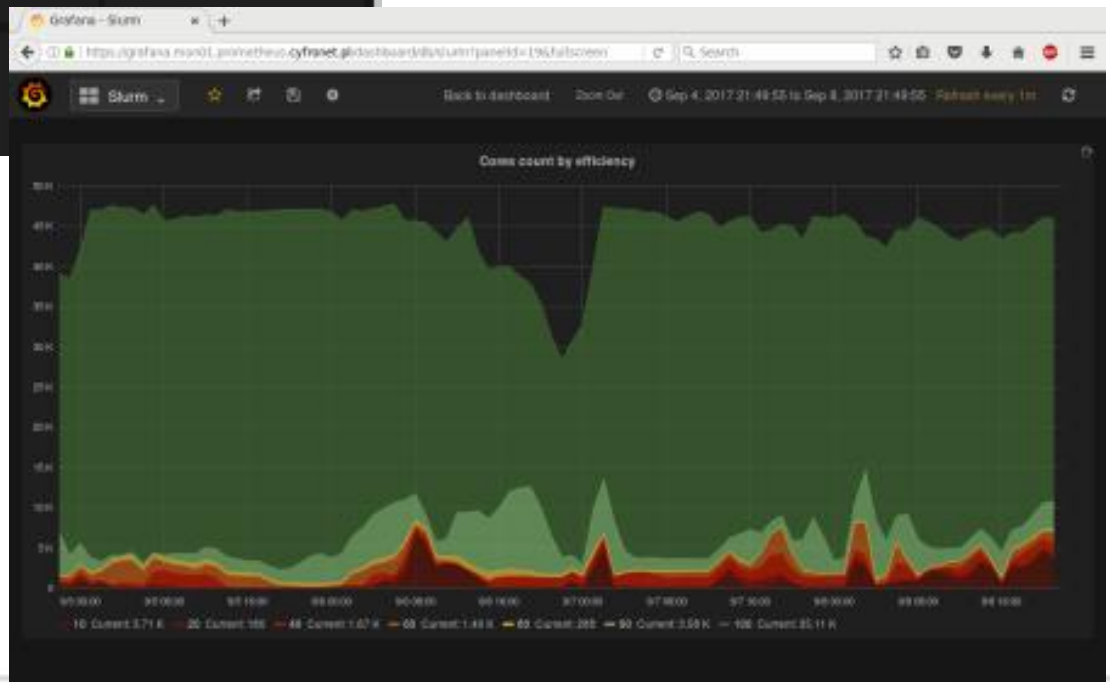

Node centric dashboards

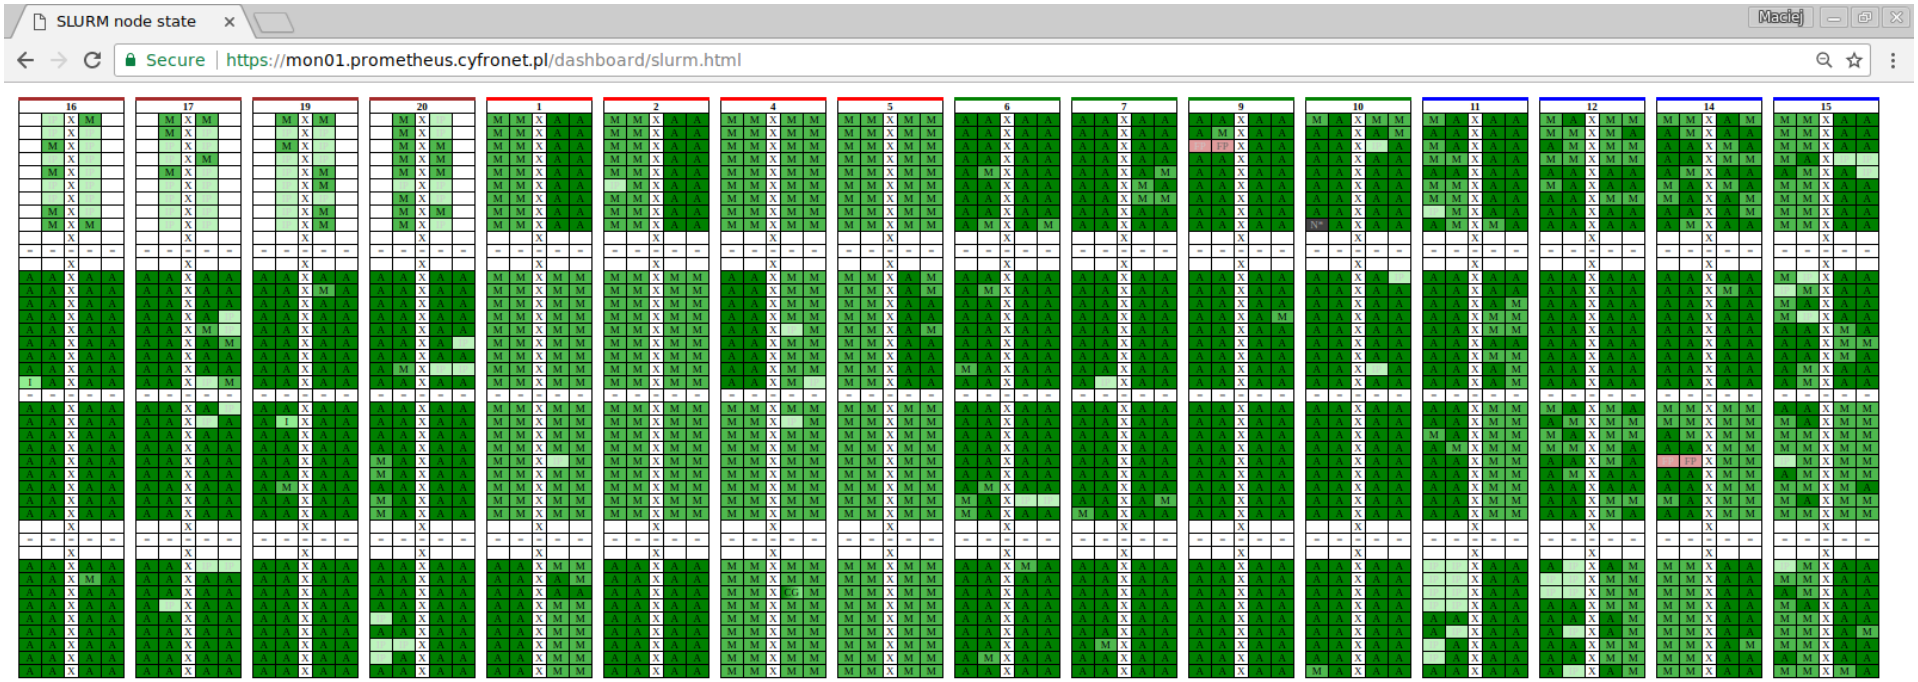

- Node:
  - State
    - Reason
  - Powered down
  - Responding

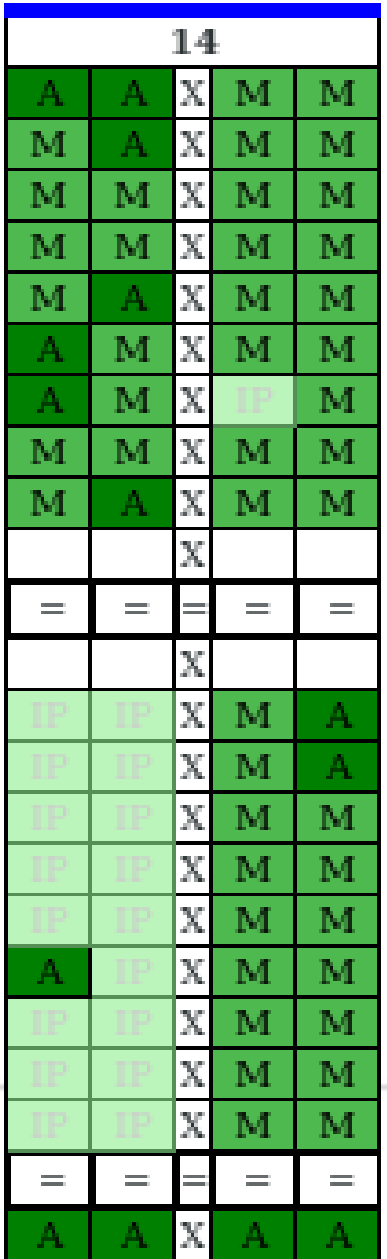

Node centric dashboards

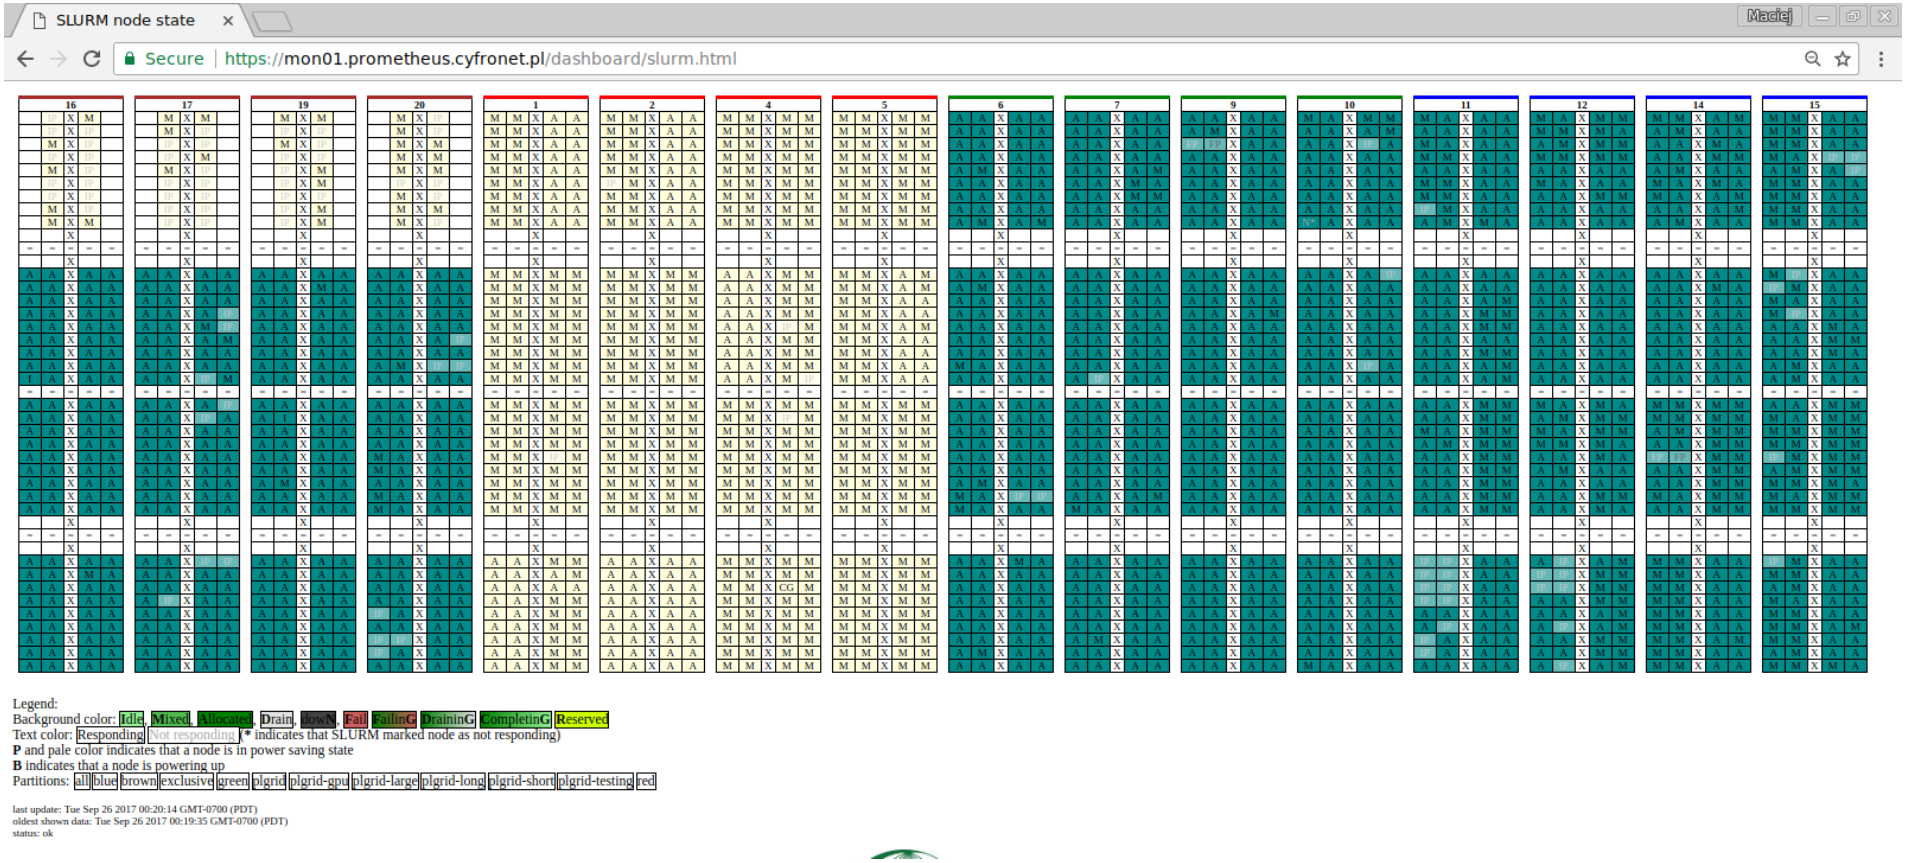

Node centric dashboards

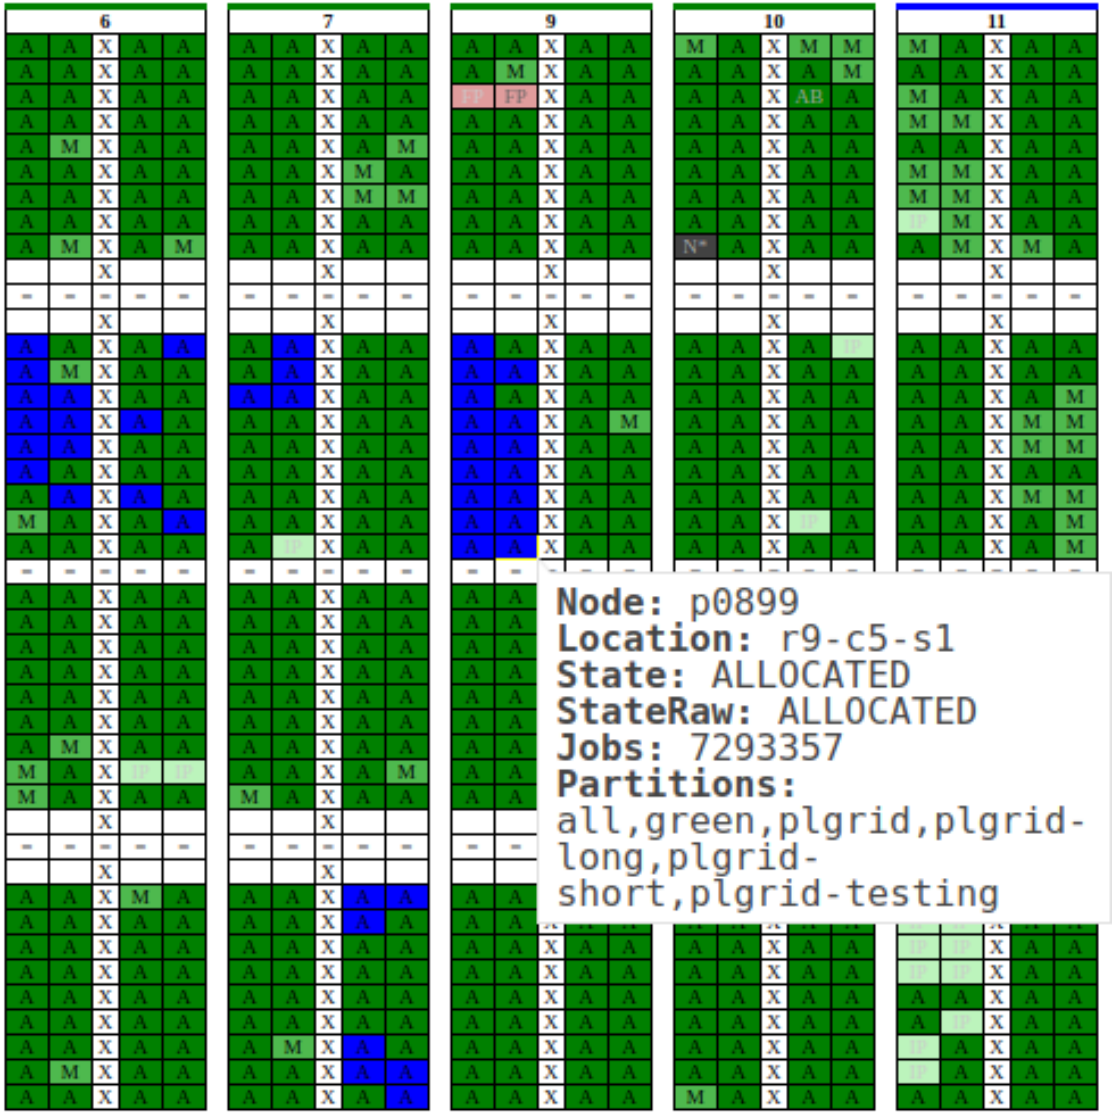

Job centric dashboards

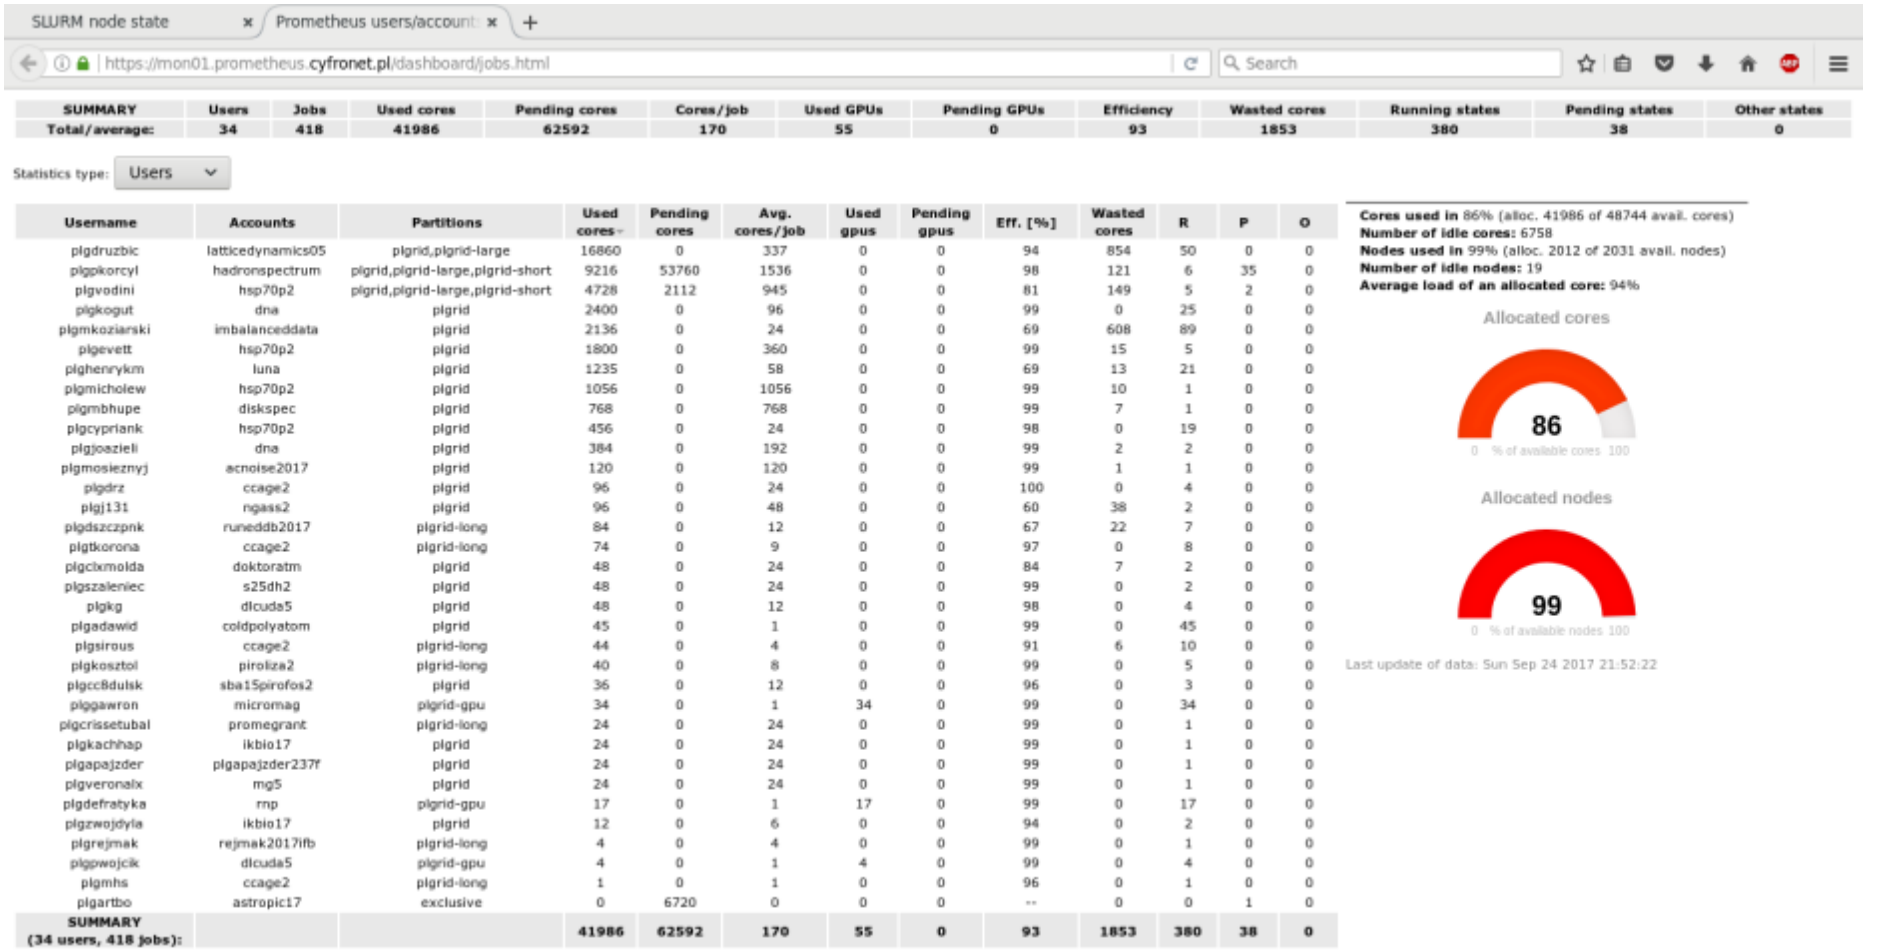

SUMMARY	Users	Jobs	Used cores	Pending cores	Cores/job	Used GPUs	Pending GPUs	Efficiency	Wasted cores	Running states	Pending states	Other states
Total/average:	34	418	41986	62592	170	55	0	93	1853	380	38	0

Statistics type: Users

Username	Accounts	Partitions	Used cores	Pending cores	Avg. cores/job	Used gpus	Pending gpus	Eff. [%]	Wasted cores	R	P	O
plgdruzbic	latticedynamics05	plgrid,plgrid-large	16860	0	337	0	0	94	854	50	0	0
plgpkorcyl	hadronspectrum	plgrid,plgrid-large,plgrid-short	9216	53760	1536	0	0	98	121	6	35	0
plgvodini	hsp70p2	plgrid,plgrid-large,plgrid-short	4728	2112	945	0	0	81	149	5	2	0
plgkogut	dna	plgrid	2400	0	96	0	0	99	0	25	0	0
plgmkoziarski	imbalanceddata	plgrid	2136	0	24	0	0	69	608	89	0	0
plgevett	hsp70p2	plgrid	1800	0	360	0	0	99	15	5	0	0
plghenrykm	luna	plgrid	1235	0	58	0	0	69	13	21	0	0
plgmicholew	hsp70p2	plgrid	1056	0	1056	0	0	99	10	1	0	0
plgmbhupe	diskspec	plgrid	768	0	768	0	0	99	7	1	0	0
plgcypriank	hsp70p2	plgrid	456	0	24	0	0	98	0	19	0	0
plgjoazieli	dna	plgrid	384	0	192	0	0	99	2	2	0	0
plgmosieznyj	acnoise2017	plgrid	120	0	120	0	0	99	1	1	0	0
plgdrz	ccage2	plgrid	96	0	24	0	0	100	0	4	0	0
plgj131	ngass2	plgrid	96	0	48	0	0	60	38	2	0	0
plgdszczpnk	runeddb2017	plgrid-long	84	0	12	0	0	67	22	7	0	0
plgtkorona	ccage2	plgrid-long	74	0	9	0	0	97	0	8	0	0
plgclxmolda	doktoratm	plgrid	48	0	24	0	0	84	7	2	0	0
plgszaleniec	s25dh2	plgrid	48	0	24	0	0	99	0	2	0	0
plgkg	dlcuda5	plgrid	48	0	12	0	0	98	0	4	0	0
plgadawid	coldpolyatom	plgrid	45	0	1	0	0	99	0	45	0	0
plgsirous	ccage2	plgrid-long	44	0	4	0	0	91	6	10	0	0
plgkosztol	piroliza2	plgrid-long	40	0	8	0	0	99	0	5	0	0
plgcc8dulsk	sba15pirofos2	plgrid	36	0	12	0	0	96	0	3	0	0
plggawron	micromag	plgrid-gpu	34	0	1	34	0	99	0	34	0	0
plgcrissetubal	promegrant	plgrid-long	24	0	24	0	0	99	0	1	0	0
plgkachhap	ikbio17	plgrid	24	0	24	0	0	99	0	1	0	0
plgapajzder	plgapajzder237f	plgrid	24	0	24	0	0	99	0	1	0	0
plgveronalx	mg5	plgrid	24	0	24	0	0	99	0	1	0	0
plgdefratyka	rnp	plgrid-gpu	17	0	1	17	0	99	0	17	0	0
plgzwojdyla	ikbio17	plgrid	12	0	6	0	0	94	0	2	0	0
plgrejmak	rejmak2017ifb	plgrid-long	4	0	4	0	0	99	0	1	0	0
plgpwojcik	dlcuda5	plgrid-gpu	4	0	1	4	0	99	0	4	0	0
plgmhs	ccage2	plgrid-long	1	0	1	0	0	96	0	1	0	0
plgartbo	astropic17	exclusive	0	6720	0	0	0	--	0	0	1	0
SUMMARY (34 users, 418 jobs):			41986	62592	170	55	0	93	1853	380	38	0

Cores used in 86% (alloc. 41986 of 48744 avail. cores)
Number of idle cores: 6758
Nodes used in 99% (alloc. 2012 of 2031 avail. nodes)
Number of idle nodes: 19
Average load of an allocated core: 94%

Allocated cores

Allocated nodes

Last update of data: Sun Sep 24 2017 21:52:22

- All of the scripts are (going to be) open sourced
  - Toolkit rather than a complete solution
- Even more openness
  - SLURM community could benefit from sharing software/knowledge
    - Knowledge – already happening on mailing list
    - Software – not yet?
- Questions?