ATLAS ANALYSIS PERFORMANCE ON THE GRID

monitoring and improving

Fourth meeting

A FEW STRANGE EFFECTS

- Observed curiosities:
  - 0 stage-out time when in direct access mode
  - Larger than 1 efficiencies for non xrootd sites
  - Disappearance of AGLT2 24 core machines in direct access mode
  - Low number of processed jobs at MWT2
- Explanation:
  - Someone re-committed to svn repository code with the line turning ON AsyncPrefetch.
  - AP is still off in case of local hard disk access
  - AP often but not always craches the job with Out Of Memory message.
    - Depending on sites (on some storage technologies it just works)
    - Depending on how much events has been read
    - Depending on how much memory there is in the machine
  - This message is not picked up by Panda which finds that job finished ok.
  - But it is detected in the stage out procedure which than aborts.

CURRENT SITUATION

- AP has been removed
- HC tests are run in a direct access mode since 10th July 10:00
- Will have to re-run copy-to-scratch mode for a week
- Problem with too slow (usually 120 seconds) stage-out remains
  - It's not just timeout or something very simple:

```
17 Jul 04:09:21|LocalSiteMov| Executing command: source /afs/atlas.umich.edu/OSGWN/setup.sh;
lsm-put -t ATLASUSERDISK --size 1180 --checksum adler32:41ab5ddd --guid 405639bc-37f8-4638-b6ce-290bd3a8f7f1
/tmp/Panda_Pilot_2409773_1342497052/PandaJob_1547573356_1342497055/user.gangarbt.4548401.EXT0._02656.info.txt
srm://head01.aglt2.org:8443/srm/managerv2?SFN=/pnfs/aglt2.org/atlasuserdisk/user/gangarbt/hc20008010/user.gangarbt.hc20008010.ANALY_AGLT2.188/user.gang
arbt.4548401.EXT0._02656.info.txt
17 Jul 04:10:19|pilot.py    | --- Main pilot monitoring loop (job id 1547573356, state:stageout, iteration 20)
17 Jul 04:10:28|LocalSiteMov| Elapsed time: 67

17 Jul 04:10:29|SiteMover.py| lfc-mkdir -m 0775 -p /grid/atlas/users/pathena/user/gangarbt/hc20008010/user.gangarbt.hc20008010.ANALY_AGLT2.188
17 Jul 04:11:19|pilot.py    | Payload stdout (athena_stdout.txt) within allowed size limit (2147483648 B): 125834 B

17 Jul 04:11:19|pilot.py    | Completed output file size verification
17 Jul 04:11:20|pilot.py    | --- Main pilot monitoring loop (job id 1547573356, state:stageout, iteration 21)
17 Jul 04:11:29|SiteMover.py| _ec: 0
17 Jul 04:11:29|SiteMover.py| telapsed: 60.0140919685

17 Jul 04:11:31|runJob.py   | Payload cleanup has finished
17 Jul 04:11:31|runJob.py   | runJob (payload wrapper) has finished
17 Jul 04:12:20|pilot.py    | Production job is done
17 Jul 04:12:20|pilot.py    | Clean up the ended job: [2409904, <Job.Job instance at 0x511b128>, 2409595]
```

CURRENT STATUS

- Some sites don't run any jobs if direct access is asked for (RHUL, Glasgow, ECDF). Need to see why.
- No results from ANALY_HU_ATLAS_Tier2 test dataset could not be found – still not clear if DS is really missing. Why jobs are starting at all?
- As most sites were running very efficiently (a lot of jobs was finished) I'll take this week as a Base line for future comparisons.

BASE LINE – CPU TIME

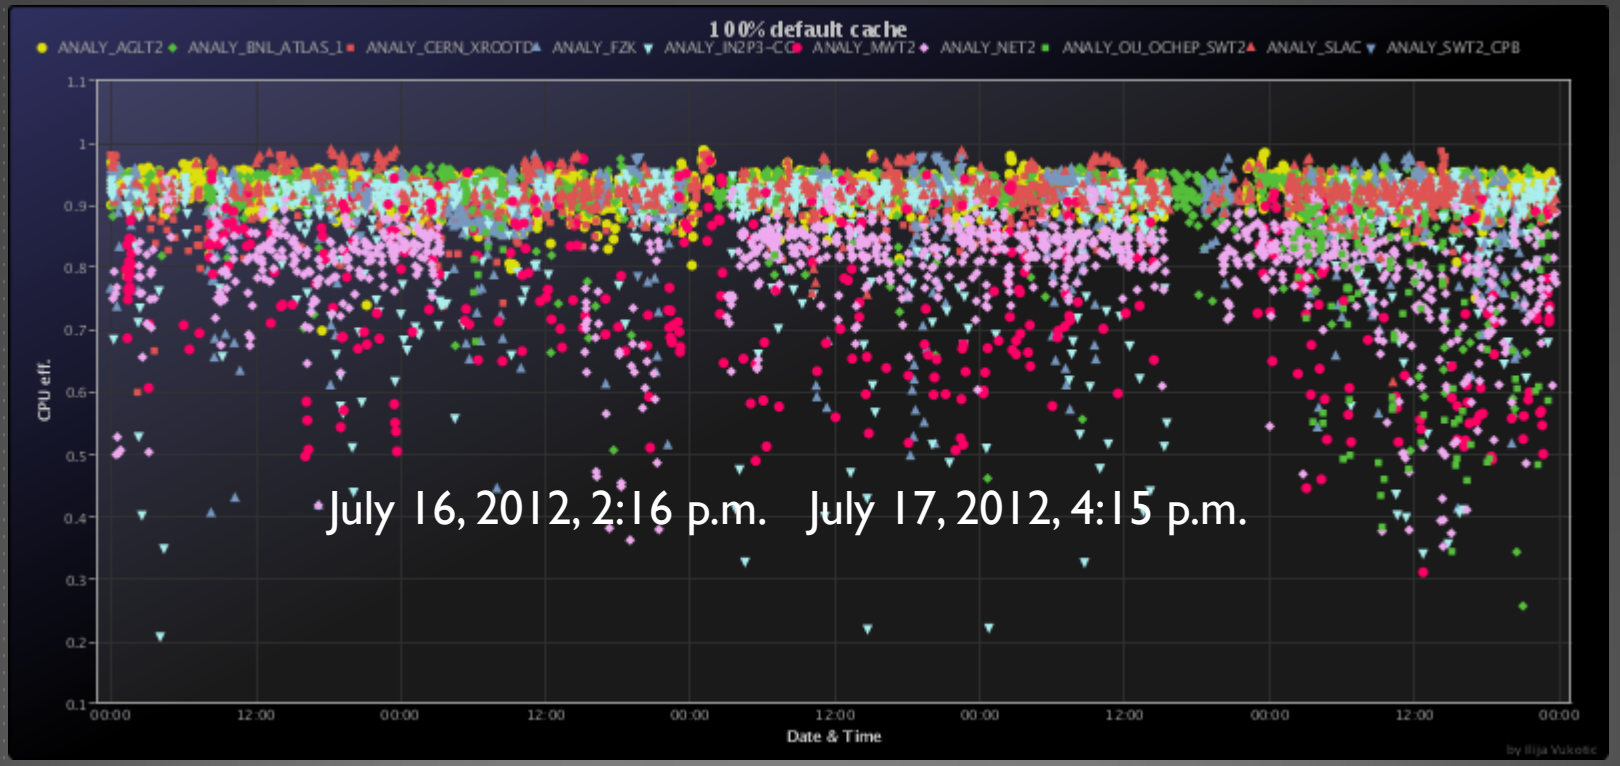

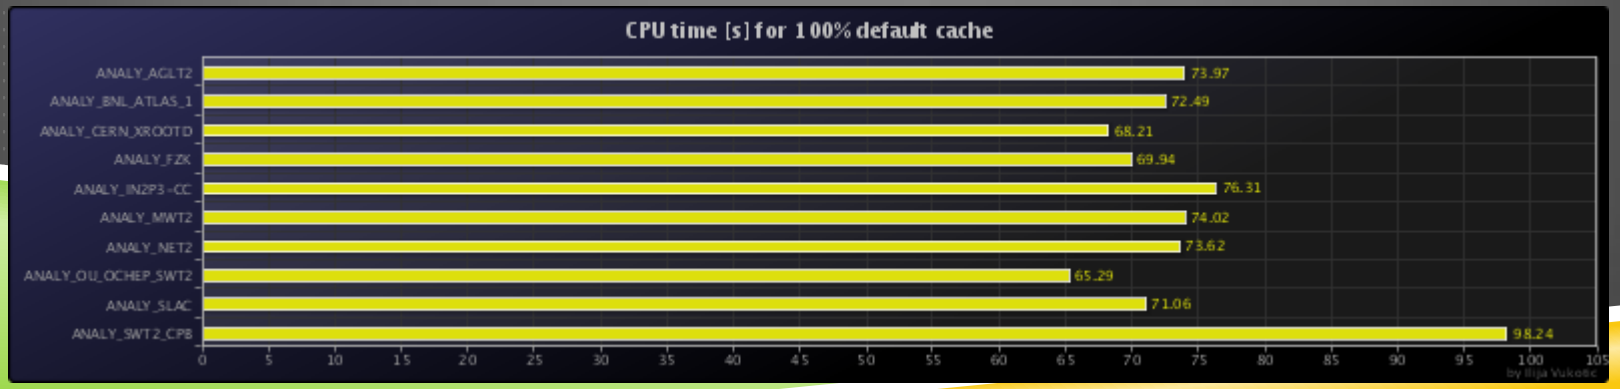

BASE LINE – WALL TIME

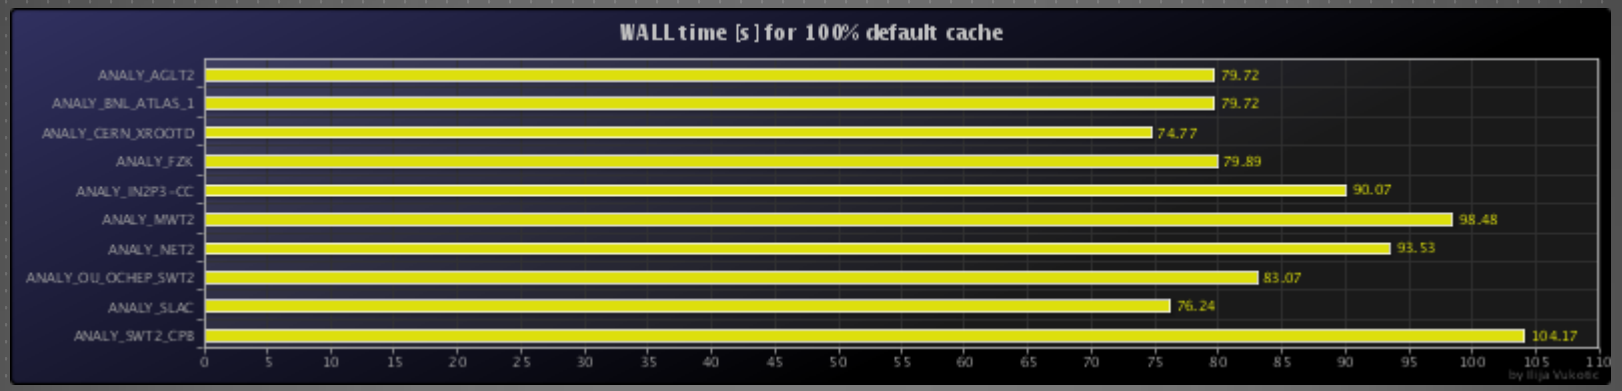

BASE LINE – CPU EFFICIENCY

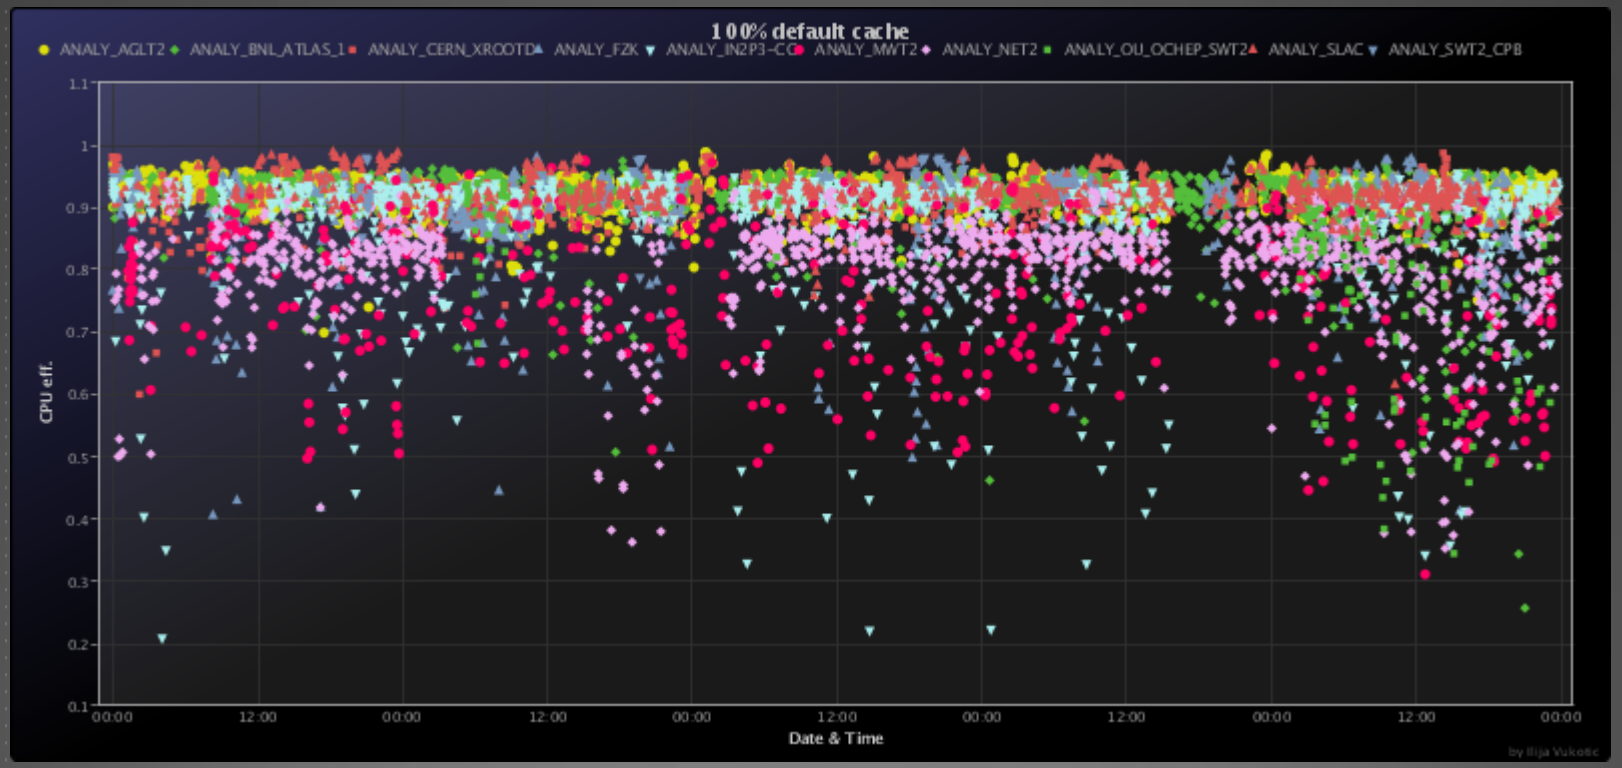

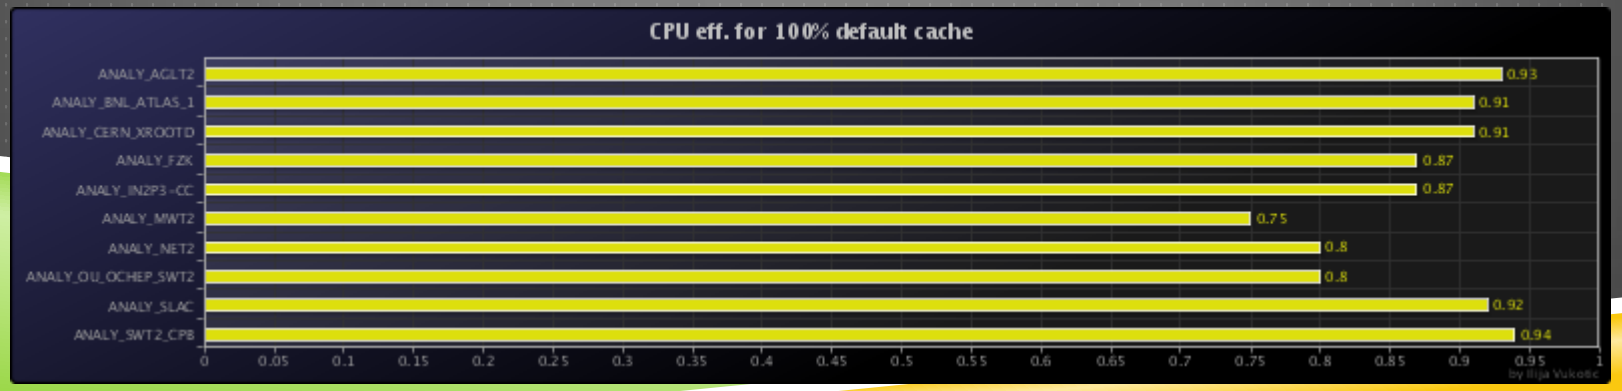

BASE LINE – STAGE IN TIME

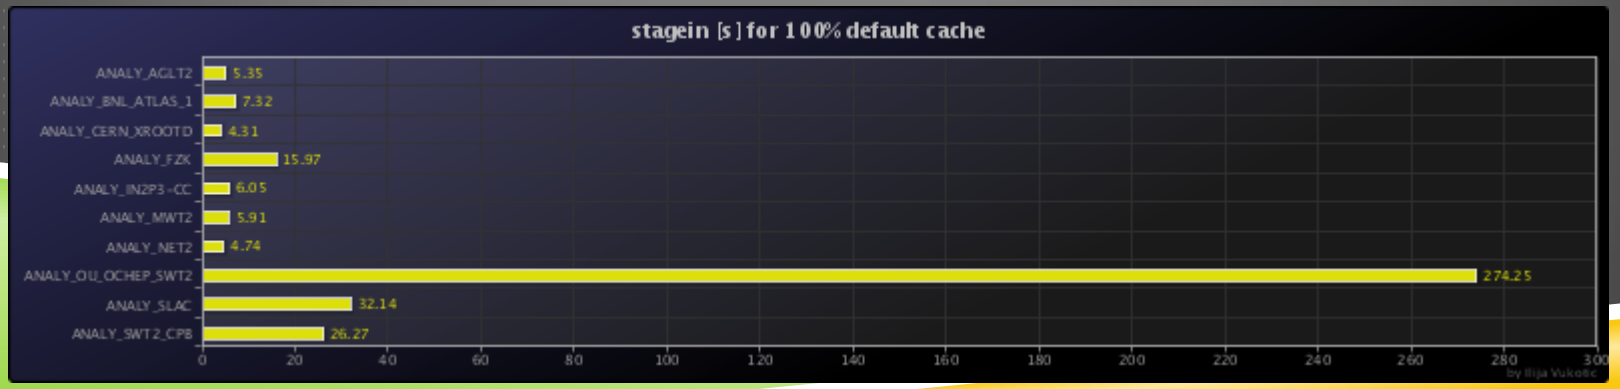

BASE LINE – STAGE IN TIME ZOOMED

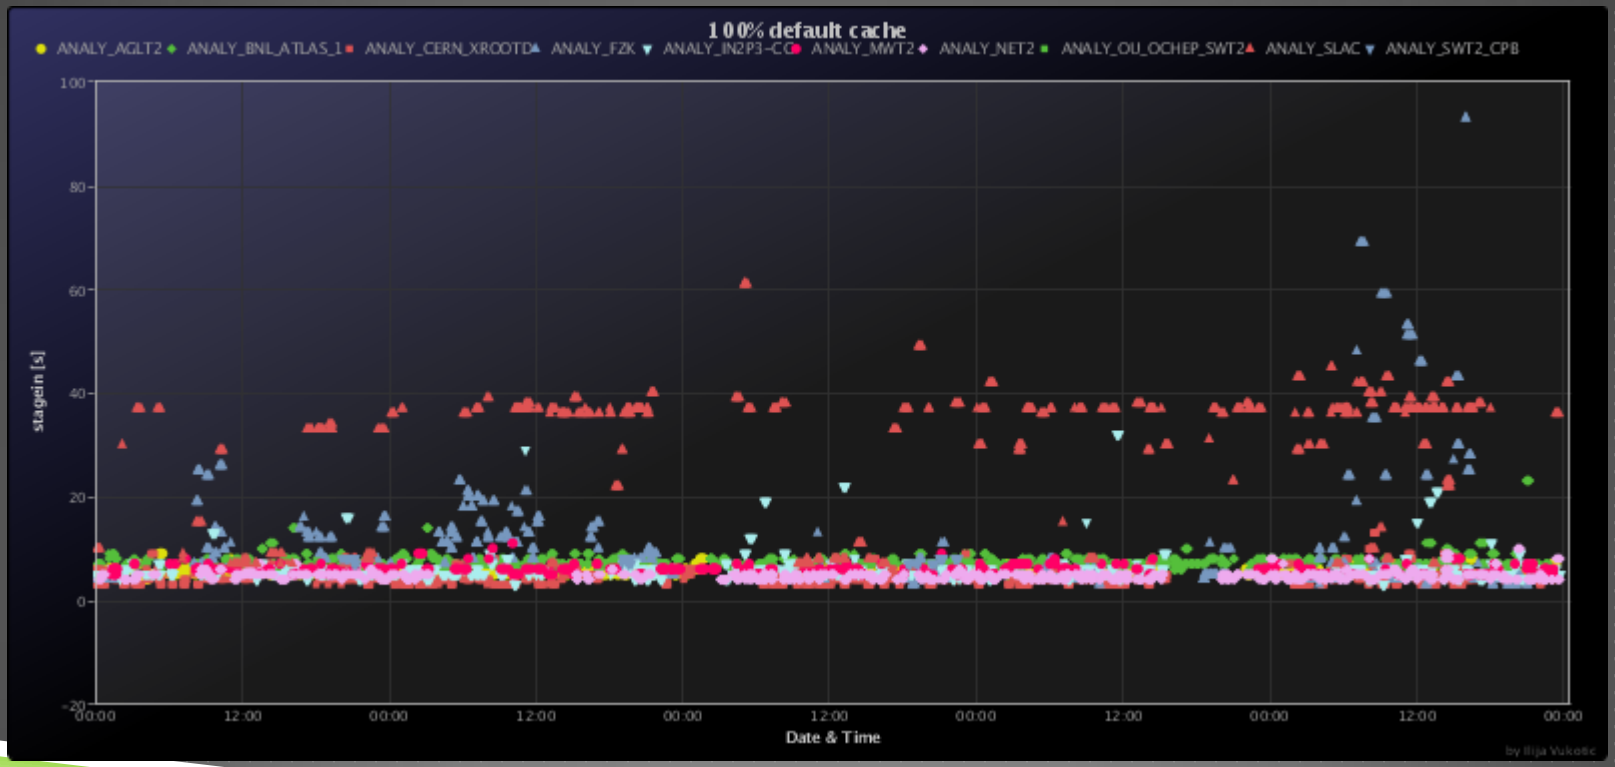

BASE LINE – STAGE OUT TIME

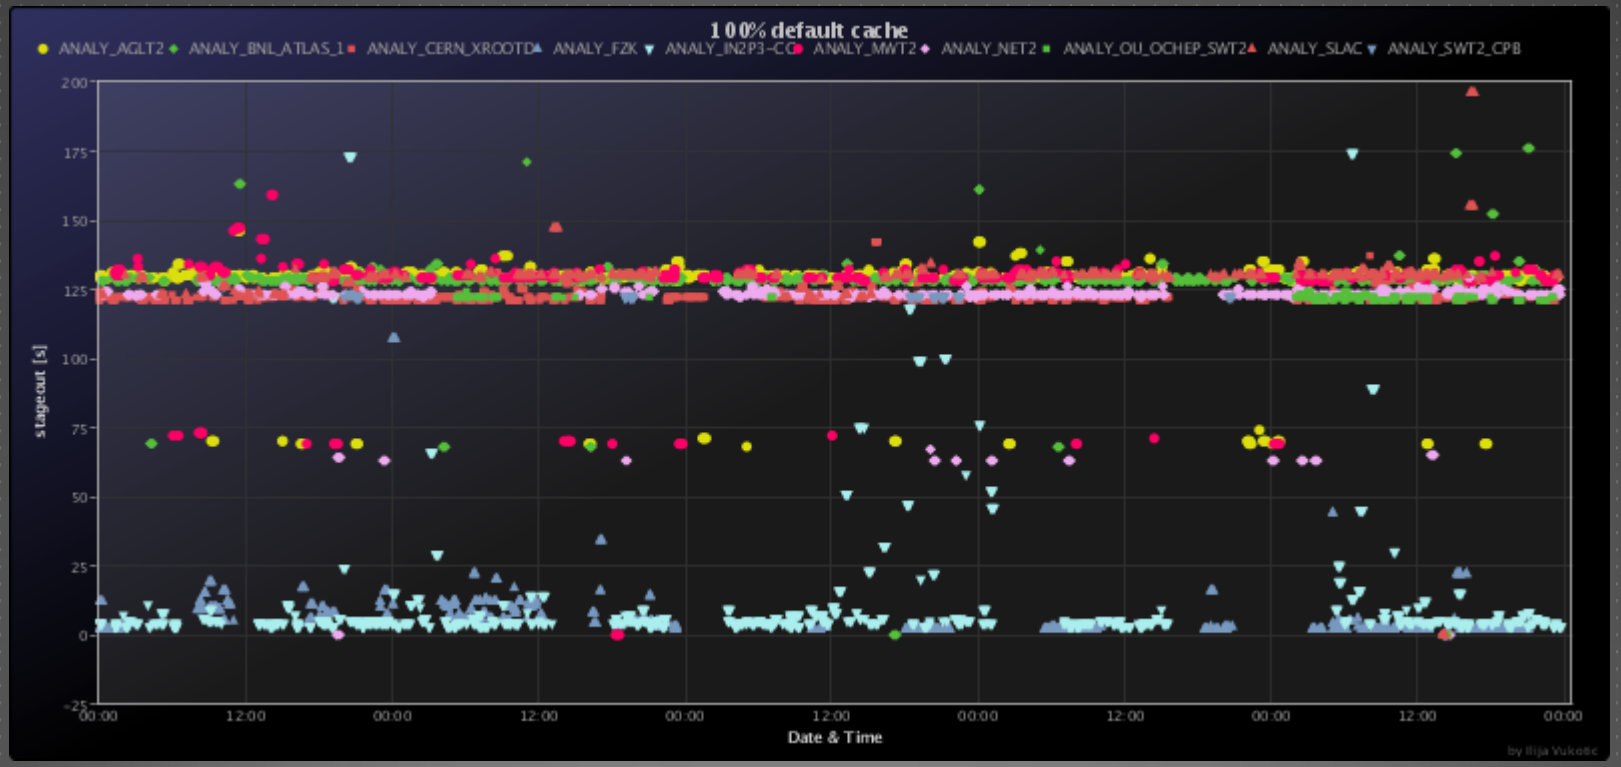

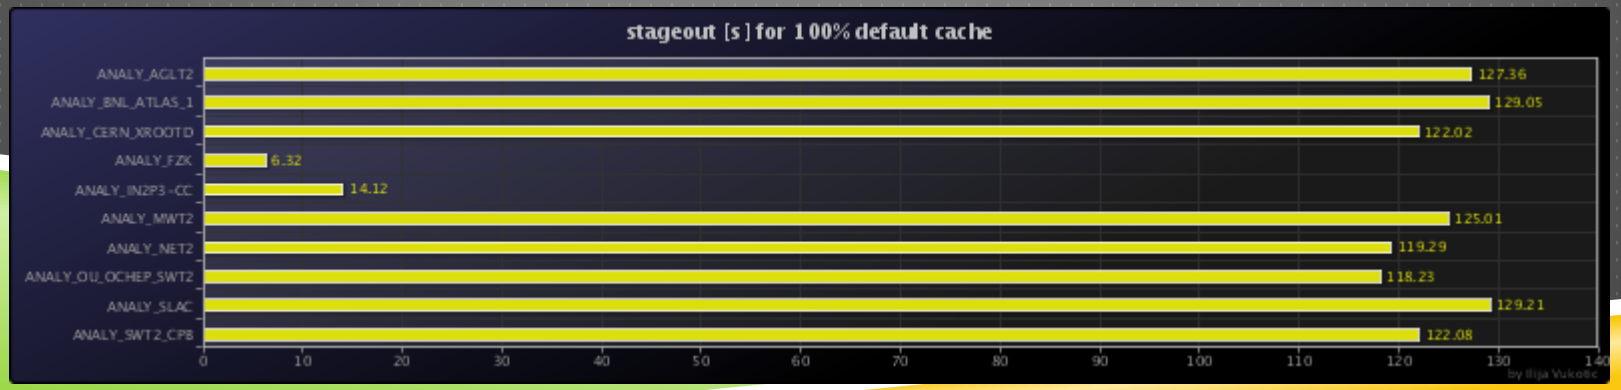

BASE LINE – SETUP TIME

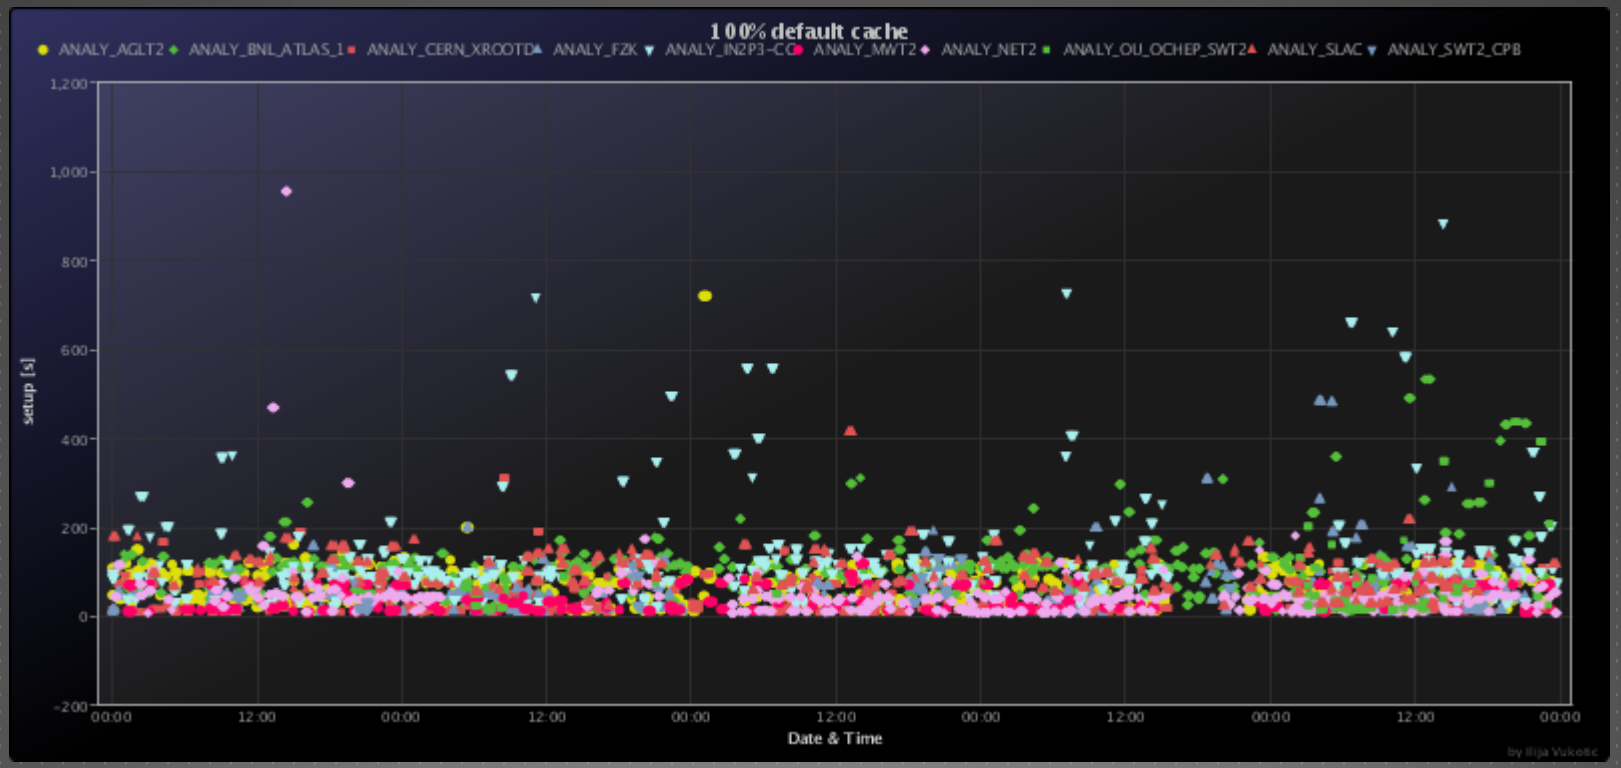

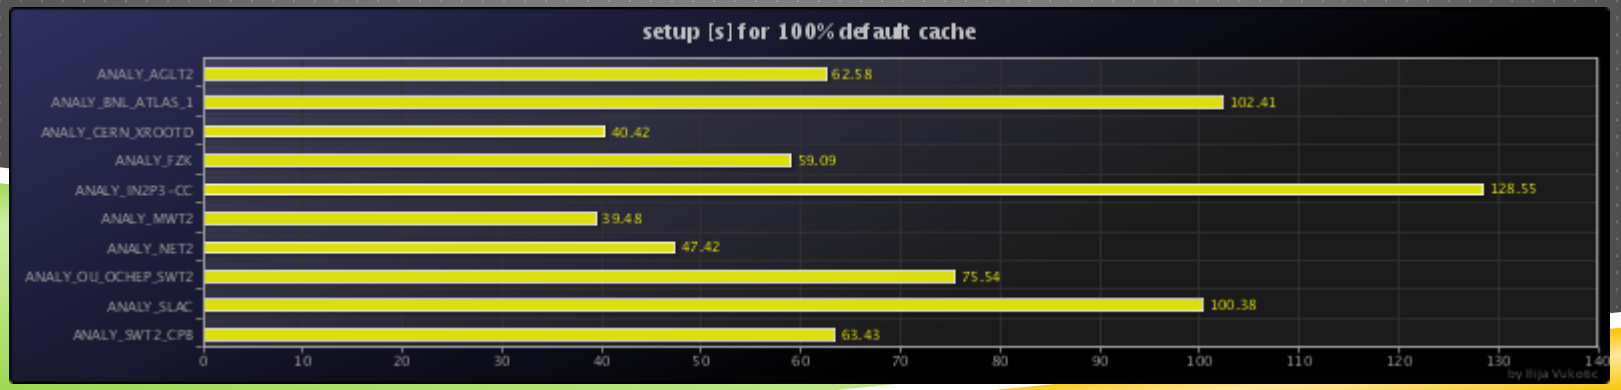

TASK LIST

- HU test data set is missing. previously had bad performance
- BNL - issue was a long stage-in time. have to recheck with next week's copy-to-scratch
- MWT2 - bad cpu eff. very fluctuating.
- NET2 - bad cpu eff.
- OU_OCHEP_SWT2 - long stage-in. SLAC could do better.
- repeat week of copy2scratch
- having plots of currently running and currently queued jobs. Need to contact panda people to get access to real-time info.
- simple and easy to understand weekly plots for sys admins to look at are mostly there but still not finished.
- automatic e-mails - nothing done yet