WURM, 28-10-2024 13:00 CET WURMng Marjolein's office

Present: Paul, Wybren, Bob, Aard, Mark, Marjolein

Plenary announcements:
Marjolein will be away next week, attending the last ORP-PILOT consortium meeting (project ends Feb 2025).
There will be an MT meeting tomorrow: can bring issues up before then for inclusion in discussion.

Paul: for open day upgraded map.jive.eu to https+Py3, need some new LEGO© blocks for the (physical) map. Archive migration is done; now working towards user accounts i.s.o. "jops". Q: which machines to connect w/ 2x25 Gbps [AOB]. The new EEE now boots; some people have expressed interest in testing on that system (e.g. Steven vdV (ASTRON: RADIOBLOCKS WP4). Building bill-of-materials necessary to make all FlexBuffs equal (which is easier for production and e-VLBI). We have support for the 4x SN2100 switches again - three years now, which is about the time scale on which we'll have to replace them. Ongoing configuration and roll-out of binlog for all MySQL instances. Upcoming talk schedule: 12 Nov (TEASER), 14 Nov (Astrofest) and 22 Jan (Astrolunch).

Bob: At Gaia time domain workshop: a patch for the "user" module to support OAuth2 was implemented, but that info not available at the time of creating a form: created an issue+patch to fix this, which was accepted. Investigating Observatory Control System: model seems optical specific, it is Django model-view-controller, and the interface to edit is the admin interface; OCS has Queue Scheduling, interesting? [Yes, as AOB]. A PySCHED issue, seemingly related to Apple M2 silicon, was reported. [MarkK: no, not M2: will elaborate later].

Wybren: JupyterHub upgraded, then roll out Zabbix and then find out Py version on JupHub too new for Ansible, so start upgrading sfxc nodes (resulting in some broken tests, fixed by (Aard|Mark)K; finding that sfxc-k3 has more use cases than expected. Upgrading four FlexBuffs: three OK, fourth (fb0, one of JIVE's own) fails to netboot and option ROMs not seen by BIOS; found bit of metal on the motherboard; temporary solution: PXE-boot from USB drive. cl0 taken out of rack, looted the memory; cl1 now runs Zabbix, ntop-ng, InfluxDB, using 2 SSDs from previous EEE.

Aard: JupHub base image upgraded, find that docker not compatible w/ node.js, requiring an update of the O/S. ASTRON organising fortnightly local RADIOBLOCKS WP4 meeting (besides fortnightly telecon), JIVE invited; after one of those useful 1:1 conversation w/ JohnR on possibility of merging filter and delay component; outcome: unlikely, delay output into device memory, then run filter; performance on own desktop: 20 stations x 32 MHz in realtime, w/ time 50/50 spent between delay and Aart's Own Correlator => TensorCore correlator should be improvement; WP4 wants timing results from the start, so need to look into reworking to single CUDA stream or else performance tooling cannot properly measure. A

VEX problem was reported last week: wrong data stream written in the sfxc config file - BenitoM's script, fix underway, now under version control.)

Mark: sfxc-node upgrade fallout: on Deb12 no libgfortran.so.4; solution: upgrade "home" to new O/S but that can't be done just yet, so temp fix: install newer compiler that links to libgfortran.so.5; this results in not bit-equal output of the correlator. Investigation reveals differences appear exactly in the noise/noisy stuff. Reported issue in importfitsidi fixed and PR to casacore accepted. NeilS reviewed ANTAB to FITS code/algorithm and wants to include them for appending to MS. Multiple phase centres in the GPU: easy to do after TC correlator for small number of phase centers, small FoV. Preparing IVOA RiG meeting in two weeks. Building ParselTongue on Ubuntu 24.04 #FAIL: Py3.12 has no distutils, but setuputils has distutils compatible interface; the issue was not Apple M2 CPU specific.

AOBs:

1) [PaulB: EEE2 can be connected w/ 2x25 Gbps - want to change rack layout, so any other machine(s) foreseen to benefit from 25 Gbps?]
[Mark: if NICs support RDMA would be interesting for tests w/ RADIOBLOCKS cluster]
Machines eligible would be: archive, JupyterHub, out.sfxc, although each time the copy of the raw data is only a one-time overhead in the workflow. Probably check current transfer speed, might be that 25 Gbps is overkill.

2) [MarkK: JupHub uses VO protocol to find data but only public data ends up in the VO catalogue, to protect source positions from the proprietary data. Idea: records in VO catalog are per _source_ and protection is set per source. So even for private data could expose the public source(s), allowing users to find their (proprietary) data - it's in the same FITS file(s); download is password protected anyway and the download tool - Py urllib based - should support passing HttpBasicAuth credentials, i.e. Bob=uncle.]
[all: great idea!]
Need to check how protection of individual sources is administrated: could be db or file system structure, and would need an experiment from an in-house PI to test with.

3) [BobE: Observatory Control System has queue scheduling - interesting?]
After discussing: yes, it's probably not working solving constraints for multiple observatories at the same time, but a) that's interesting and b) even for the optical facilities might be useful if users really want contemporary observations.