

Hybrid WURM, 27-03-2023 13:00 WURM (Library meeting room and JIVE Zoomroom#1)

Present Bob, Aard, Wybren, Des, Mark, Marjolein

Plenary announcements/discussion:

- Marjolein: removed several WURM appointments on account of 2nd *day holidays and work trips by YT

Aard: Attended NVIDIA GTC; take-outs: Hopper GPUs combi of 72-core AMD w/ GPU + unified mem = highly interesting but likely *very* expensive; lots of C++ work done on parallelisation (e.g. multi-node threading C++26) + offloading onto GPU; "tf32" type (18-bit float) performance not so good, about 50% of real 16-bit float, on par with real 32-bit float so what's the point?; 8-bit float avail: 4 + 4 bit mantissa + exp (horrible precision, wider dyn range); availability of cunumeric = numpy drop-in that offloads to GPU. WSRT asked to be added to evn-monitor, added; report that DNS returns internal addr but cannot connect (Wybren talk to Julian); Q: can see flexbuff free space? [Some discussion, benefit unclear, see discussion at TOG; [Bob: if subtract what we already have becomes useful; i.e. easy to add but no real benefit until invest in system w/ more knowledge]. Uploaded JupHub image based on casa v6.5.3 to Zenodo. GV020 data rescue: most disks ok, Mc not ok if read w/ jive5ab, mk5read still ok; Gb pack: completely dead PATA pack, now dd'ing disk-at-a-time over USB.

Wybren: microphone of hybrid setup is *bad*. Zabbix fun: added ASTRON switches to see if we can monitor our own interfaces, editing scripts and making templates: find that using regex not available in all fields and not documented; upgraded, SNMP probs not fixed, but is not faster/more efficient; it's a multi-step process. Auto-adding items to dashboard using ansible: no way, maybe use Zabbix aggregation functio? sfxc-g1: broken disk, installed new and dd'ed but cannot activate remotely, need to take it out in-person. Archive: addition of 10G card: big #FAIL (not booting, booting once, hang/crash during finding disks, after boot find broken disk, use hot spare, 2nd disk breaks ... anyway: now back at original situation, probably have 3wk copying ahead ... Six disks sent to ServerDirect: package was left here, new pick-up date arranged, gone now, but not received yet by vendor? SuperMicro: provided app to read out IPMI power: works; tool can read out other values of interest. Todo: upgrade shared Zabbix to 6.4, add sfxc blades, rename hvvm VM. Q: nginx proxy on services or own VM? [Mark: own VM]

Bob: Tr tested BHTOM automatically generated schedule: seems to work. numpy f2c wrapper bug: added patch that does not break other tests to bug report for discussion: makes some lines EUNREACH, but is unclear what those actually do, maybe for int types, checked and found that code may also be broken for e.g. INT*4. Received pySCHED email: fail to install on recent Mac due to PyQt5 issue(s). Helped Paul w/ webpages on archive2 (replaced mysql() w/ mysql() and home-grown my.cnf). Q: how to do dev env on archive2? [Good question, this not right time+place for discussion].

Des: all casa tickets merged 6.5.5: wide band ffit (marked experimental), combining polarisations in ffit, per-scan interpolation. Not ready: callib bug, VLBA data; failed to build pkg for former, it depends on build environment and Vile; *may* be able to push these through as they're small; weight-based ffit: BenitoM may have another legit use case. PolConvert: can read data and write info for next step but then crashes (mem mgmt). Auth system/keycloak as IdP: MMost client can use SAML, OpenID connect but cost \$moneys?!

Mark: FITS crawler did finish, csv file loaded into V0 service: now have 25 000 records iso 20 000. CASA VLBI meeting: v6.5.5 tickets in; this week create branch, expected release end of Apr, in time for workshop. In-person registrants now finding they have no budget or some are from Russia. NVIDIA GTC: some talks so technical had to bail, need to refresh basics then relisten; large variety of talks; e.g. NVIDIA nsight supports multi-node profiling even if not on GPU. Started to experiment with more efficient read+unpack of (VDIF) data for RADIOBLOCKS using intrinsics - asm instructions as function calls, e.g. use PEXT insn to extract masked bits out of 64-bit number; on AMD cpu actually slower than on CPU, on Intel CPUs fast, can now do ~8 Gbps; wrote versions for 2, 4, 8, 16 ch extraction, some need testing (8ch vsn gives same output as current code), want to build into sfxc but test cluster CPUs do not have this insn. For RADIOBLOCKS/next gen correlator will say: provide VDIF or bust.