

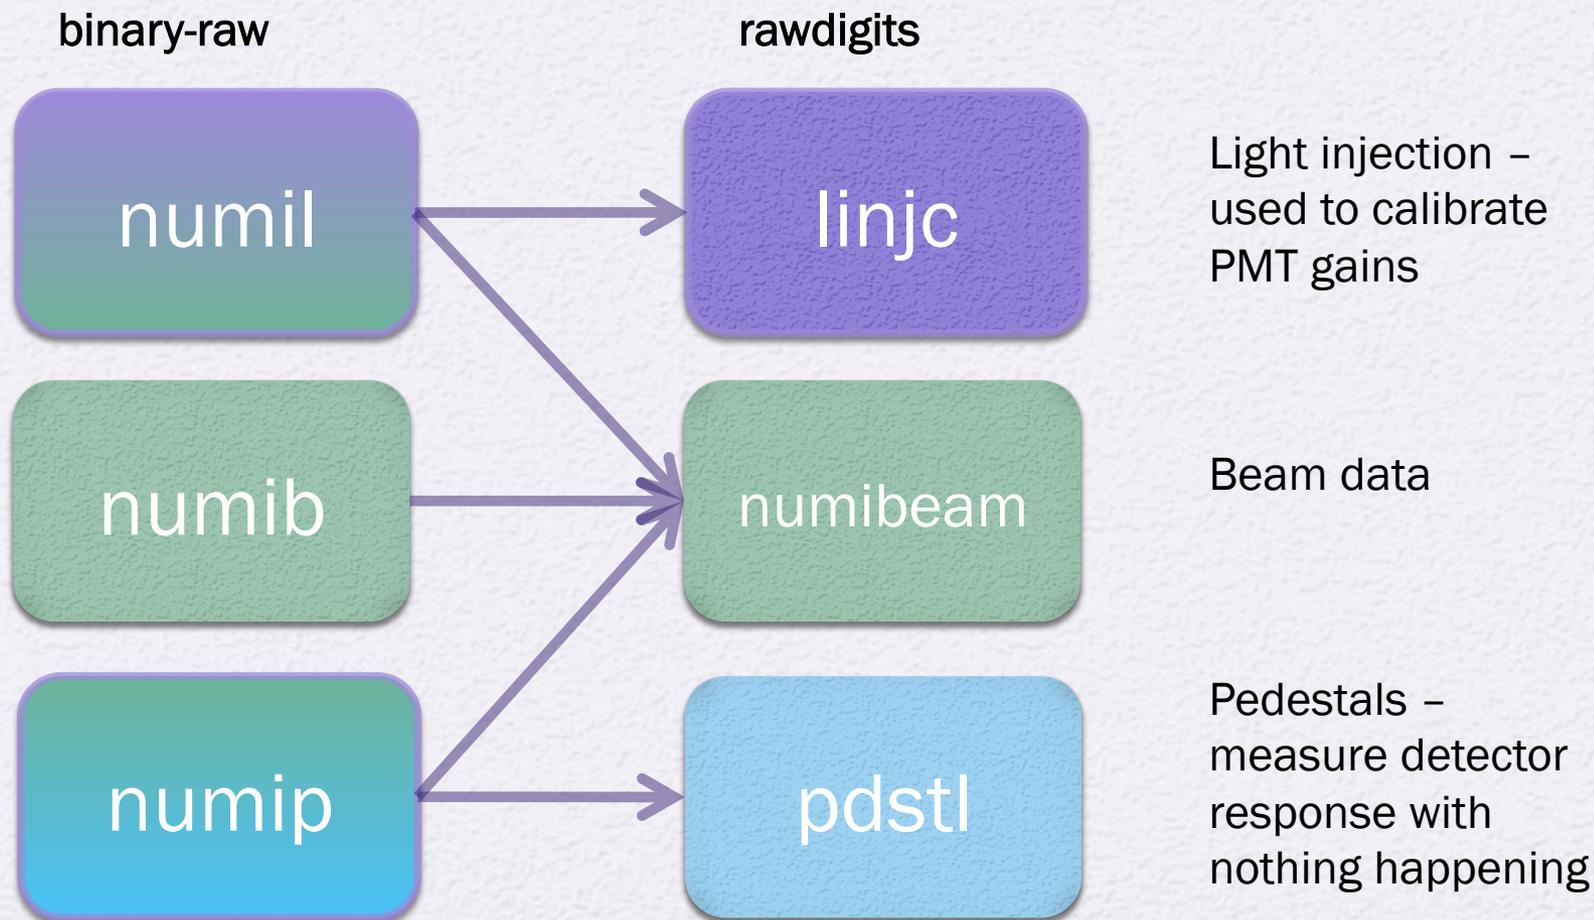


# Keep-up processing

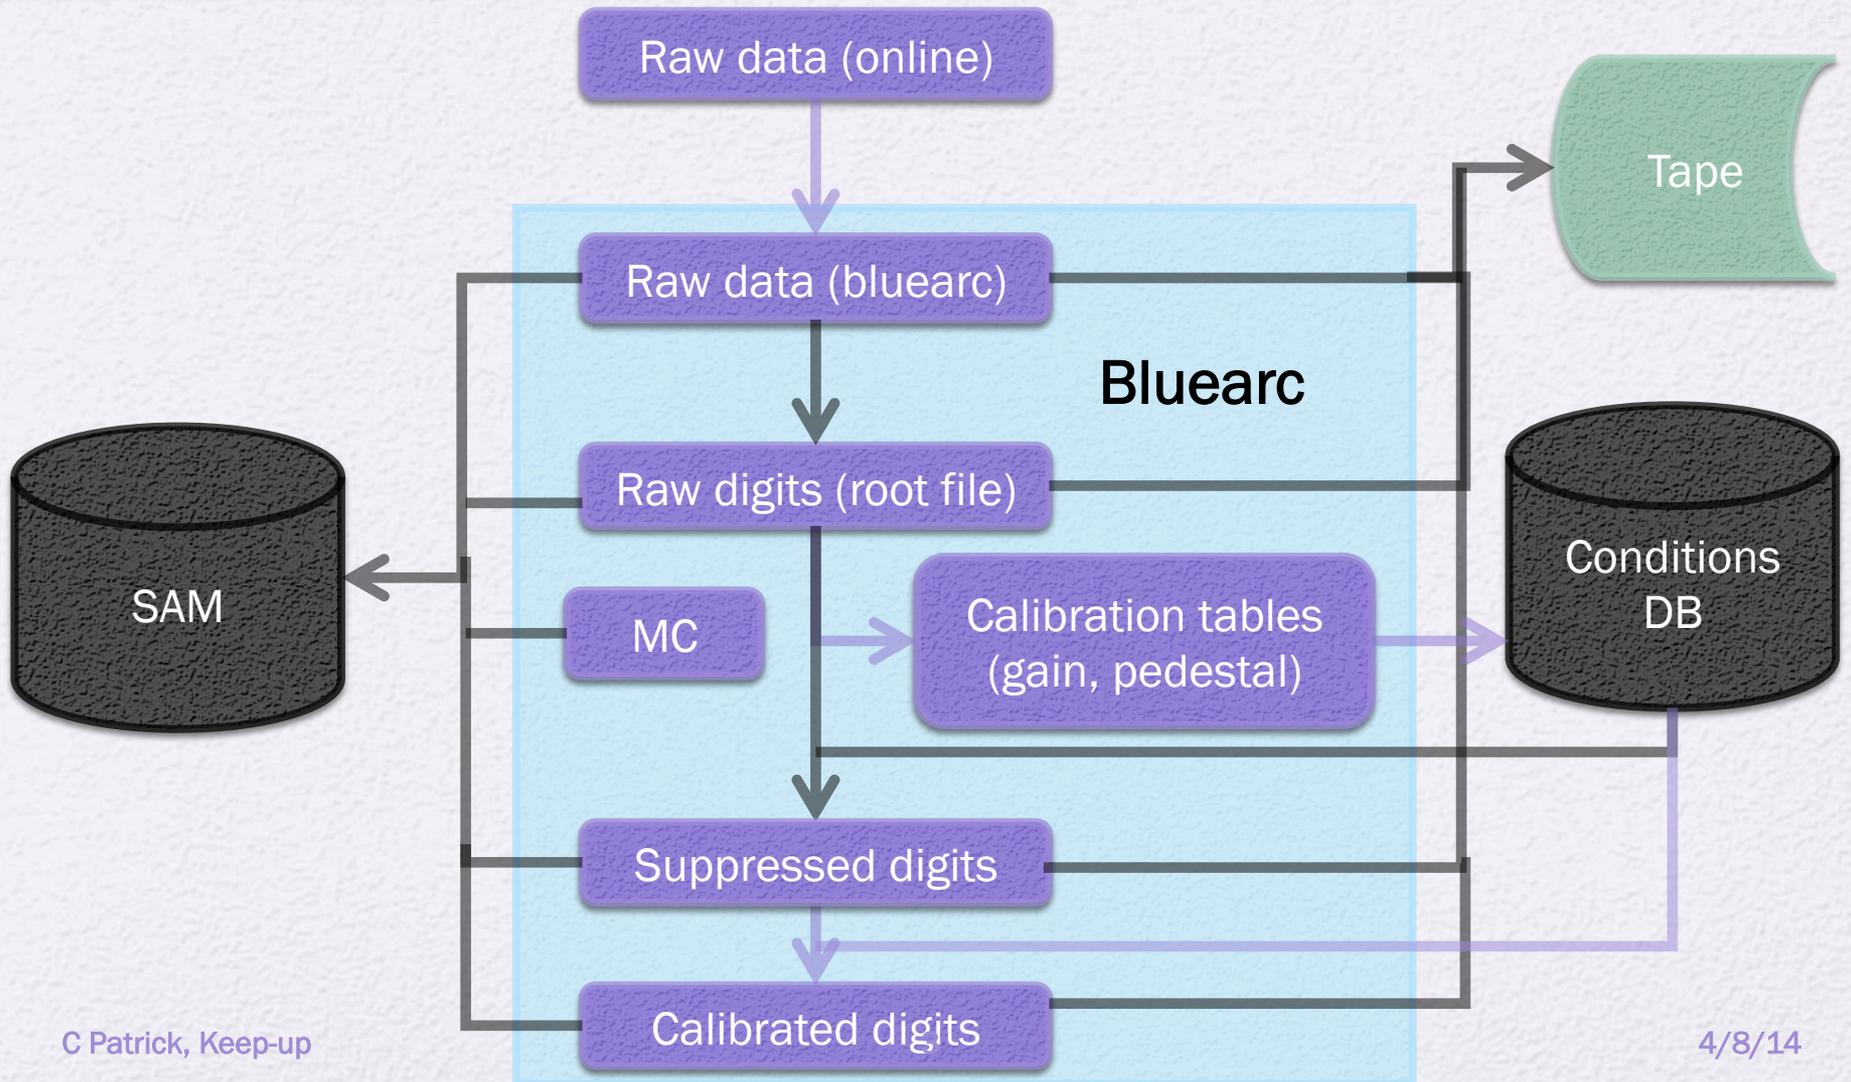
---

Cheryl Patrick  
MINERvA data management meeting  
April 8-9, 2014

# Raw data files for different run types



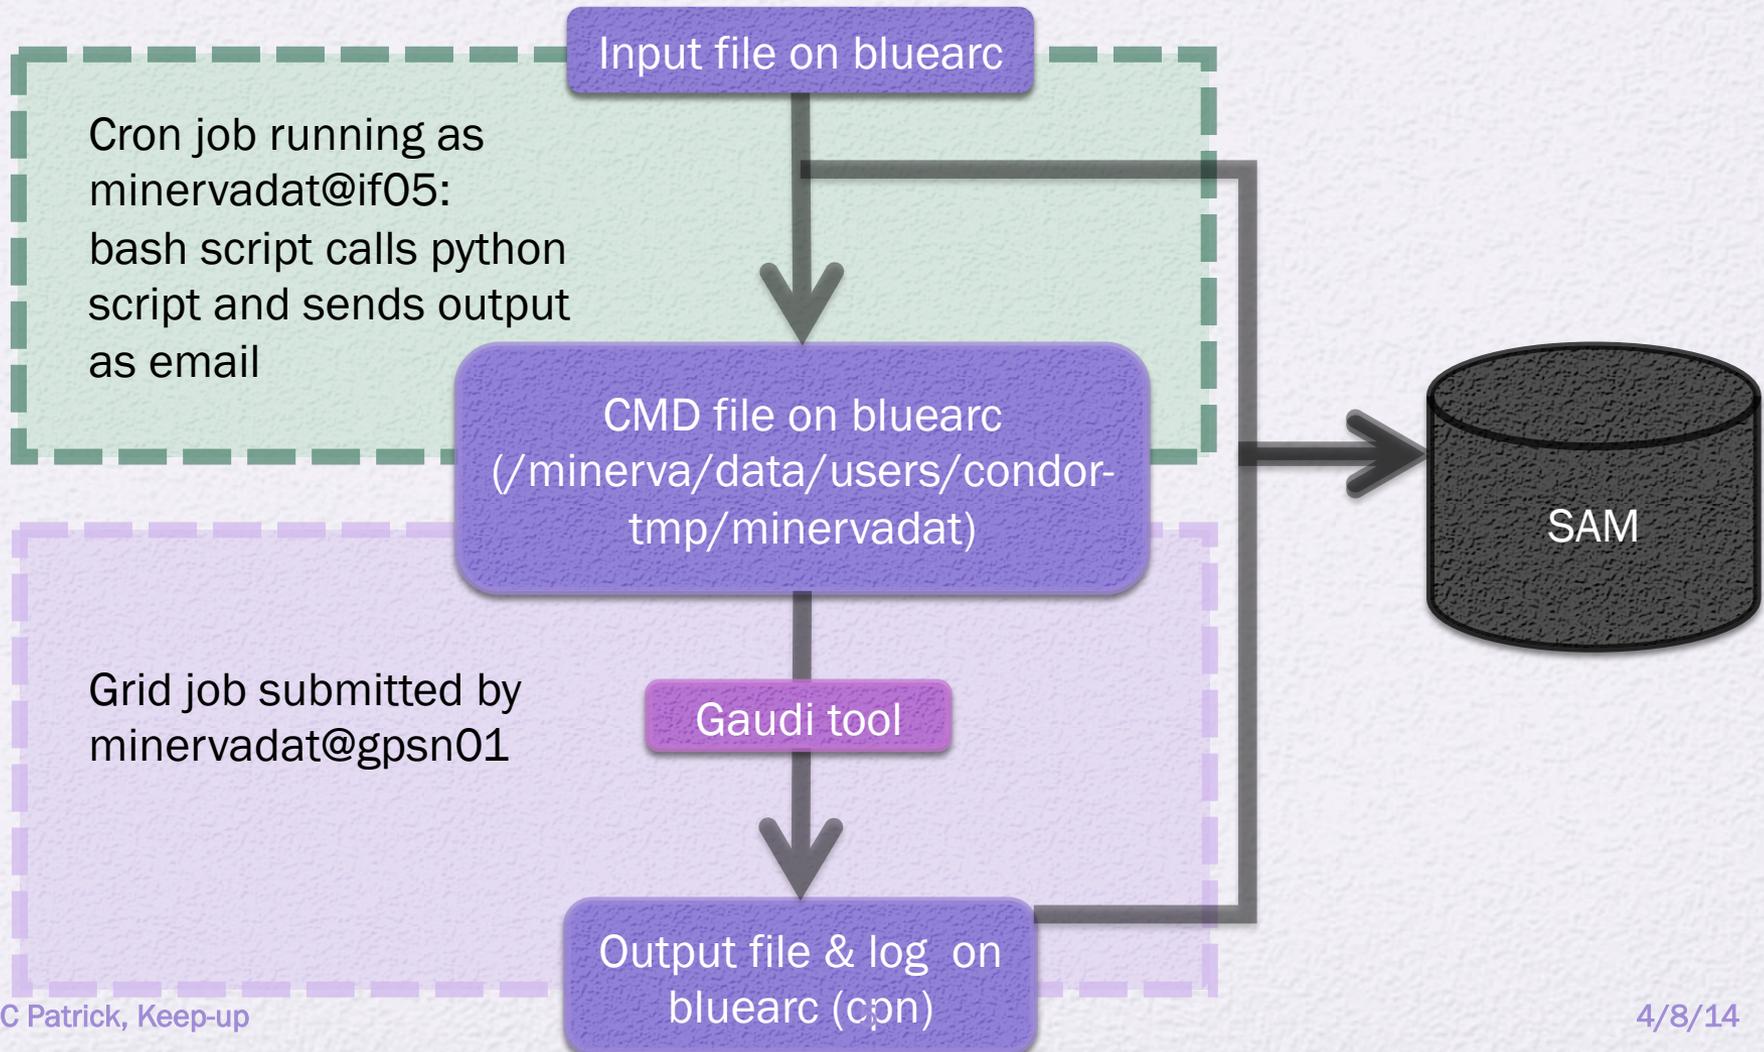
# What's keep-up?



# Keep-up data files on bluearc

File type	SAM data tier	Number per day	Typical size	Notes
Raw data (numip, numil,numib)	Binary-raw	~ 80 (Current config 1 numip: 16 numil:28 numib)	6-700MB (numip)	Numip runs typically 1500 gates, others 700
			4-500MB (numil, numib)	
Raw digits (pdstl, linjc, numib)	Rawdigits	~ 100-150 (1 per numib, 2 per numil or numip)	240 MB (pstl)	Number of files depends on run configuration
			220-270MB (linjc)	
			420MB (numibeam)	
Pedestal tables	Not in SAM	~2	2MB (x3 files)	
Suppressed digits	supdigits	~80 (numip, numil,)	22MB	Generated monthly, 200ish per day

# A common keep-up step



# Proxies

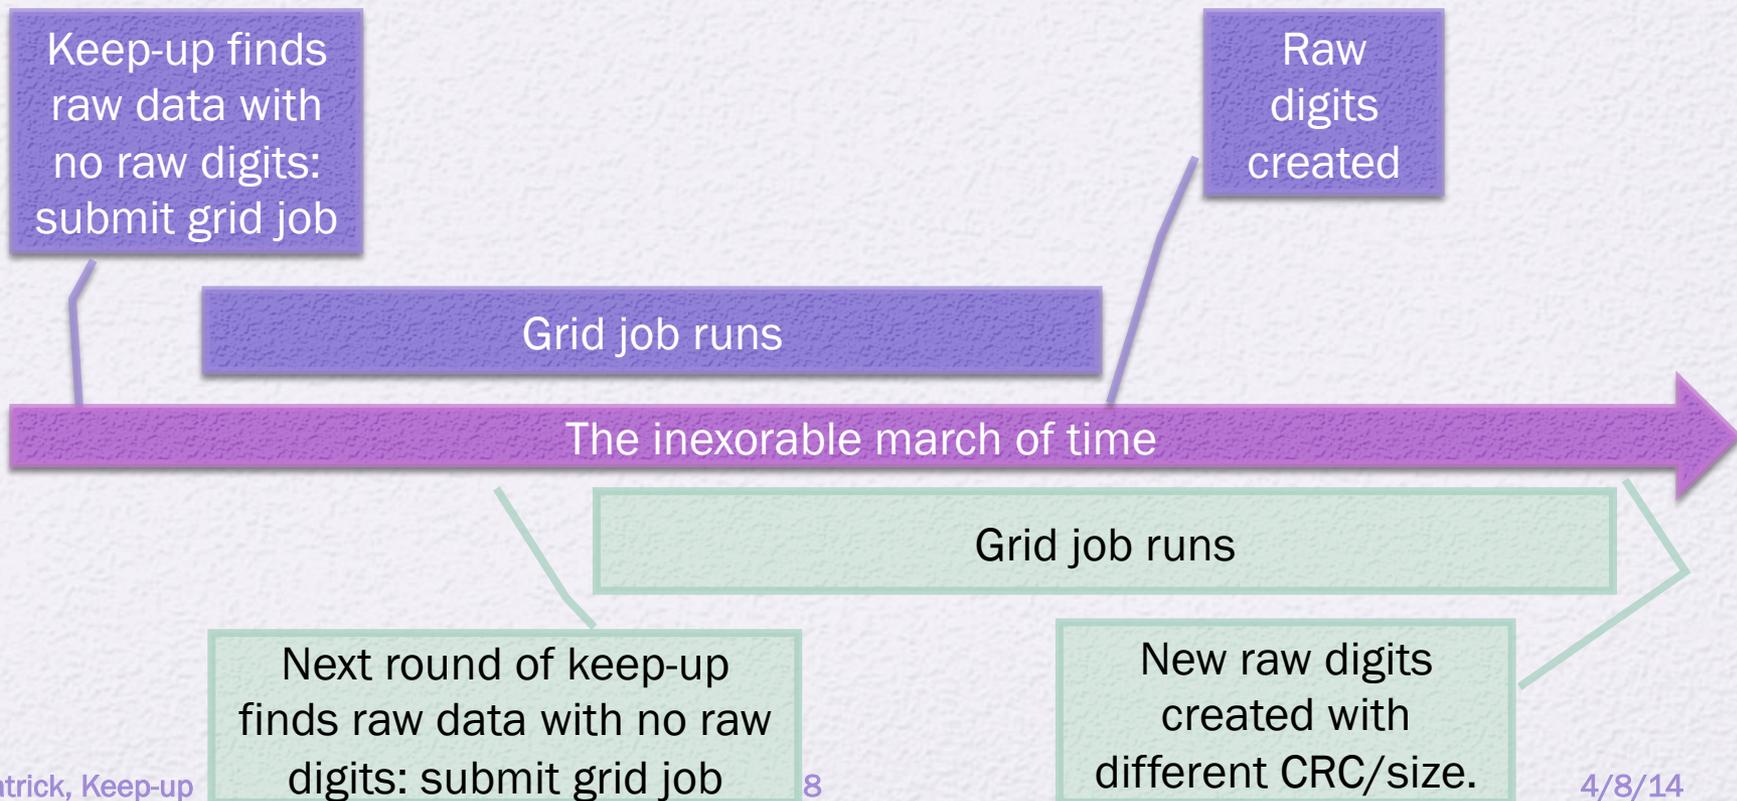
- We write to the SAM catalog (declare files, add application families, add locations etc.) from
  - Shared minervadat account, on if05, from a cron job
  - Shared minervadat account, from grid jobs (submitted on gpsn01)
- As minervadat is a shared account, we use Cheryl's (cpatrick) account to create proxies
  - Created every 2 hours by cron job on cpatrick@gpsn01
  - Copied to scratch areas for minervacal, minervadat, minervapro (data and production proxies)
  - Data proxy also copied to minervadat@if05 scratch area
  - \$X509\_USER\_PROXY set by our setup scripts, and passed to grid jobs

# Backups (dailyLink.py)

- List files in data tier with no PNFS location
- Check: on bluearc, CRC and size match SAM
  - Size difference small (5kB) – update SAM
  - Size difference larger – warn and skip
- Make soft link to FTS upload directory (max 5000 per directory)
- File backed up, SAM location updated, link deleted from upload area

# Duplicate files and locks

**Problem:** some raw digits files' timestamps/CRCs did not match SAM, and here's why



# Solution: lock files

- Create a 0-byte file <filename>.lock in minervadat's home area when creating cmd file for a given input file
- If a lock file exists for a file with a timestamp less than 1 week ago, don't attempt to reprocess it
- This has (almost) fixed the duplicate file problem

# Lock files: drawbacks

- Problems with the grid can cause jobs to hang and eventually be removed from the queue
  - Locking prevents the job from being submitted for a whole week
  - Can we distinguish between a grid job in progress (don't retry) and a grid job that failed (should retry)?
- Even with the lock files, we saw duplicate raw digits after a grid stall in March.
  - Original files were created and written to FTS
  - Bluearc version was then overwritten
  - Should we prevent the second write in this case?

# Timeline/latency

Numbers are taken from condor log files and represent typical times for a single job on a grid node (elapsed time can be much longer depending on queue)

Process	Usr time	Sys time
Raw data to raw digits	7-8 hours	15-20 minutes
Raw digits to pedestals	10-20 minutes	Seconds
Raw digits to supdigits	5 minutes	Seconds

# Outstanding issues

- We don't have a sophisticated method for dealing with hung grid jobs
  - Would like to run high\_prio but don't know how
  - How to tell if they are not running
  - Locking doesn't discriminate so they don't re-run for a whole week
  - And despite all this, after some grid trouble we STILL got duplicate raw digits and I don't know why
- We should upgrade to the latest samweb\_client which has better retry mechanism
- Proxy method is working but hacky

# Other keep-up processes

- Declaring Monte Carlo and reconstructed data to SAM
- Generating metadata for MC, reconstructed data and flux files by parsing file names and paths
  - Complicated and prone to failure due to inconsistency in naming
  - Very slow as it necessitates finding files NOT in SAM on very big disks
  - Doesn't really work at all for flux files

# Check scripts

- How full is disk (is there a faster way to do that than a du? It takes DAYS)
- Quick-and-dirty check scripts
  - Any raw data declared with start time yesterday?
  - Any raw digits with start time yesterday?
  - Any raw data from yesterday that isn't declared?