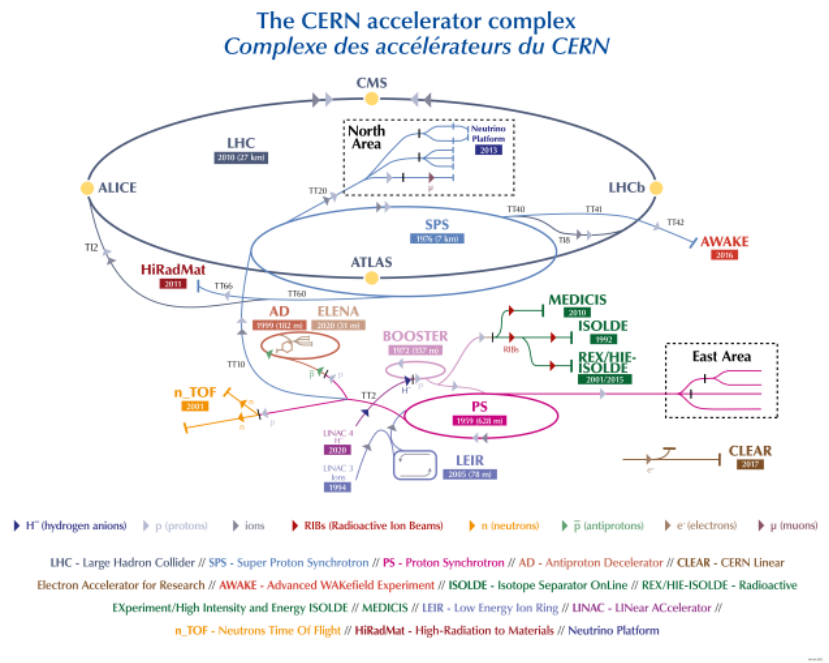


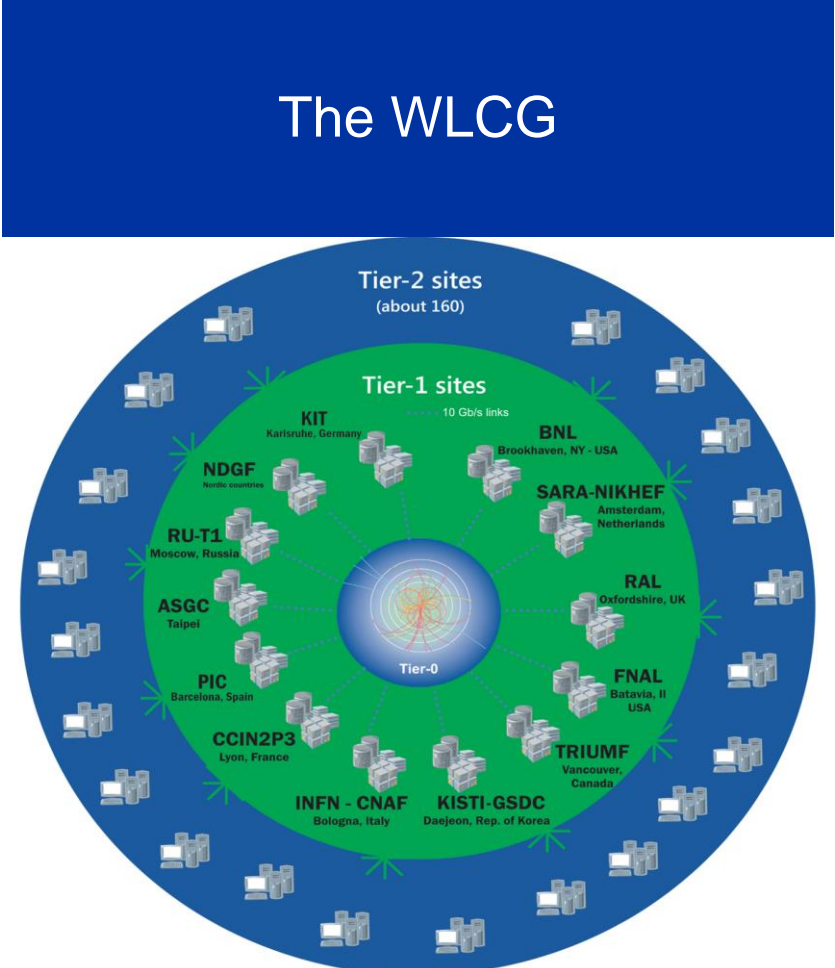
HTCondor @ CERN

Update & Challenges

CERN Batch System



Local Production



User analysis



Batch stats

- 265k cores (3.4MHS06)
- ~2 Condor pools
- 5k worker nodes (EPs?)
- 20 Local schedds / 22 Grid CEs
- Infra (+/-) 9.0.17 / workers -> 9.0.17
- 300-350 unique daily submitters
- **Compute capacity stable, number of workers quartered (v2p)**



Upcoming activity

- **EL9 for next platform**
- **ARM (some, maybe)**
- **Move pool auth from GSI (probably to kerberos)**
- **CEs to token submission only**
- **HTCSS 10**



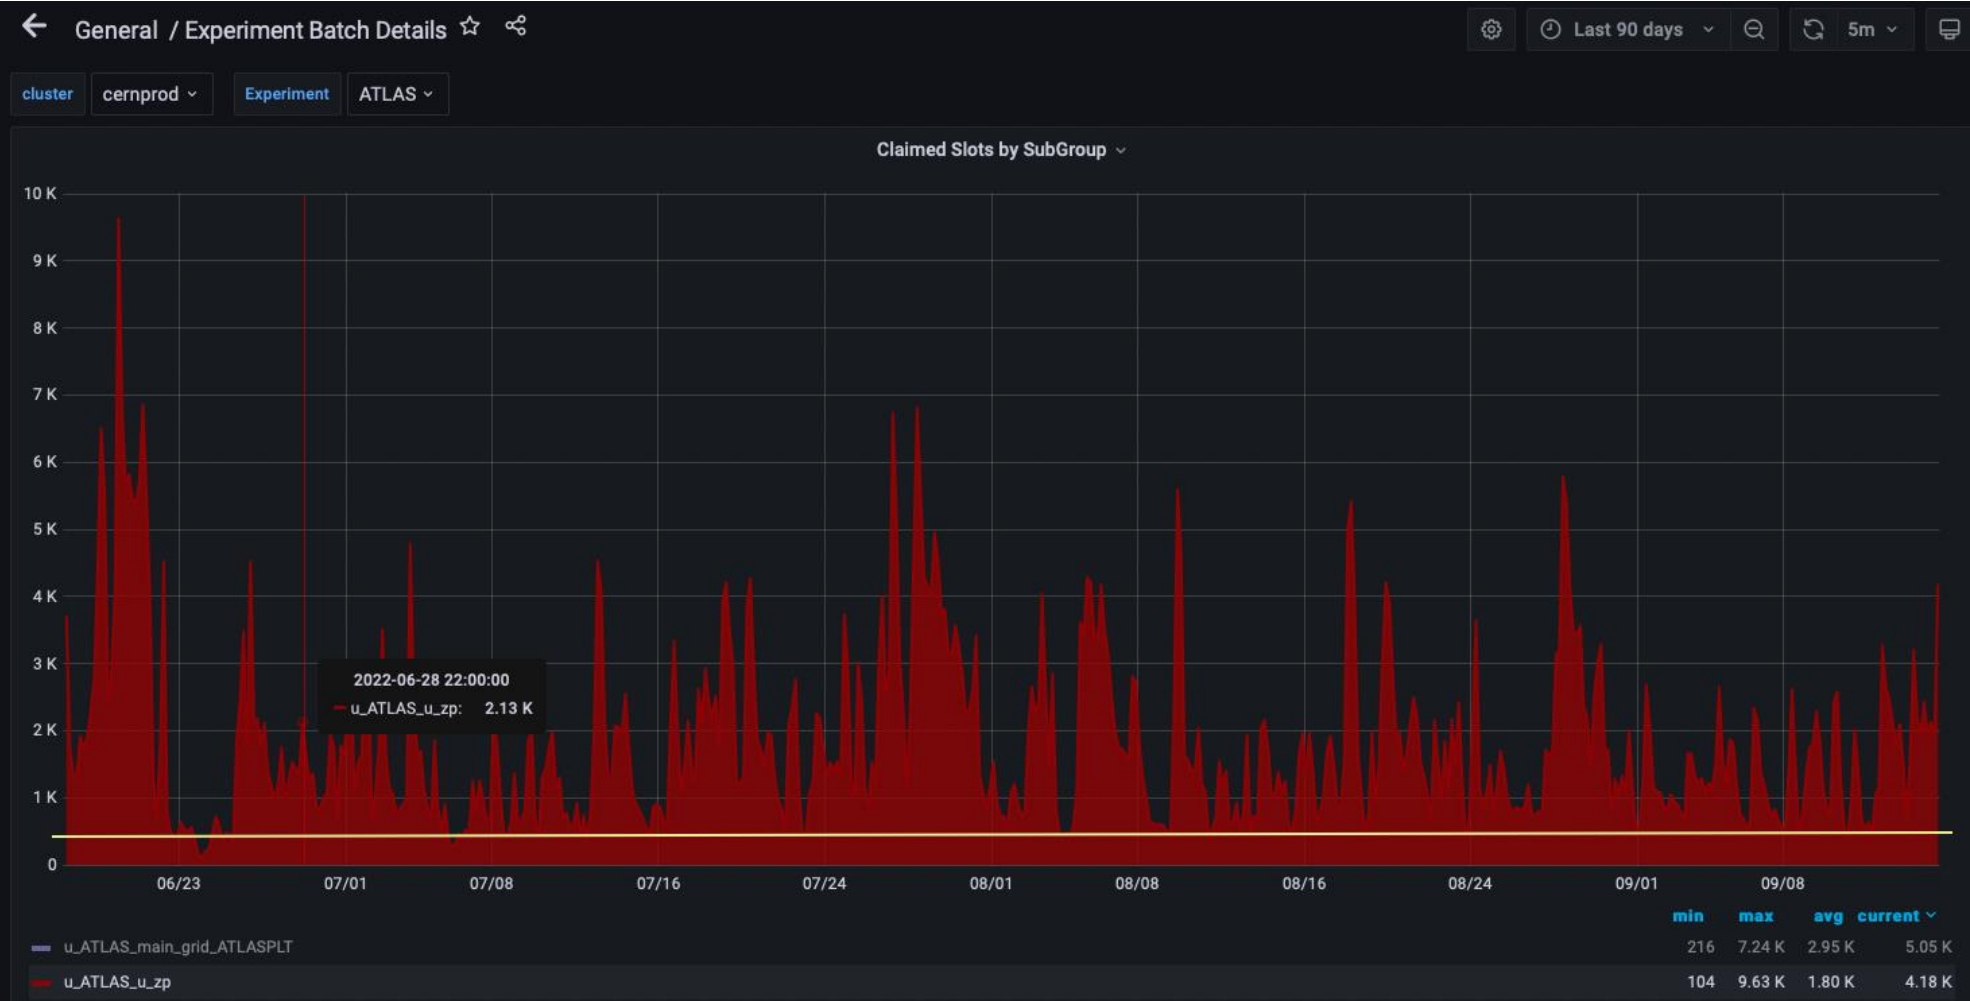
Local v Grid



*“Why do my jobs take so long to schedule?
Last week I got 1000s of cores, this week
none. How can we use the system more
efficiently?”*

User Story...

Inefficient?




Effective Quota

User Challenges

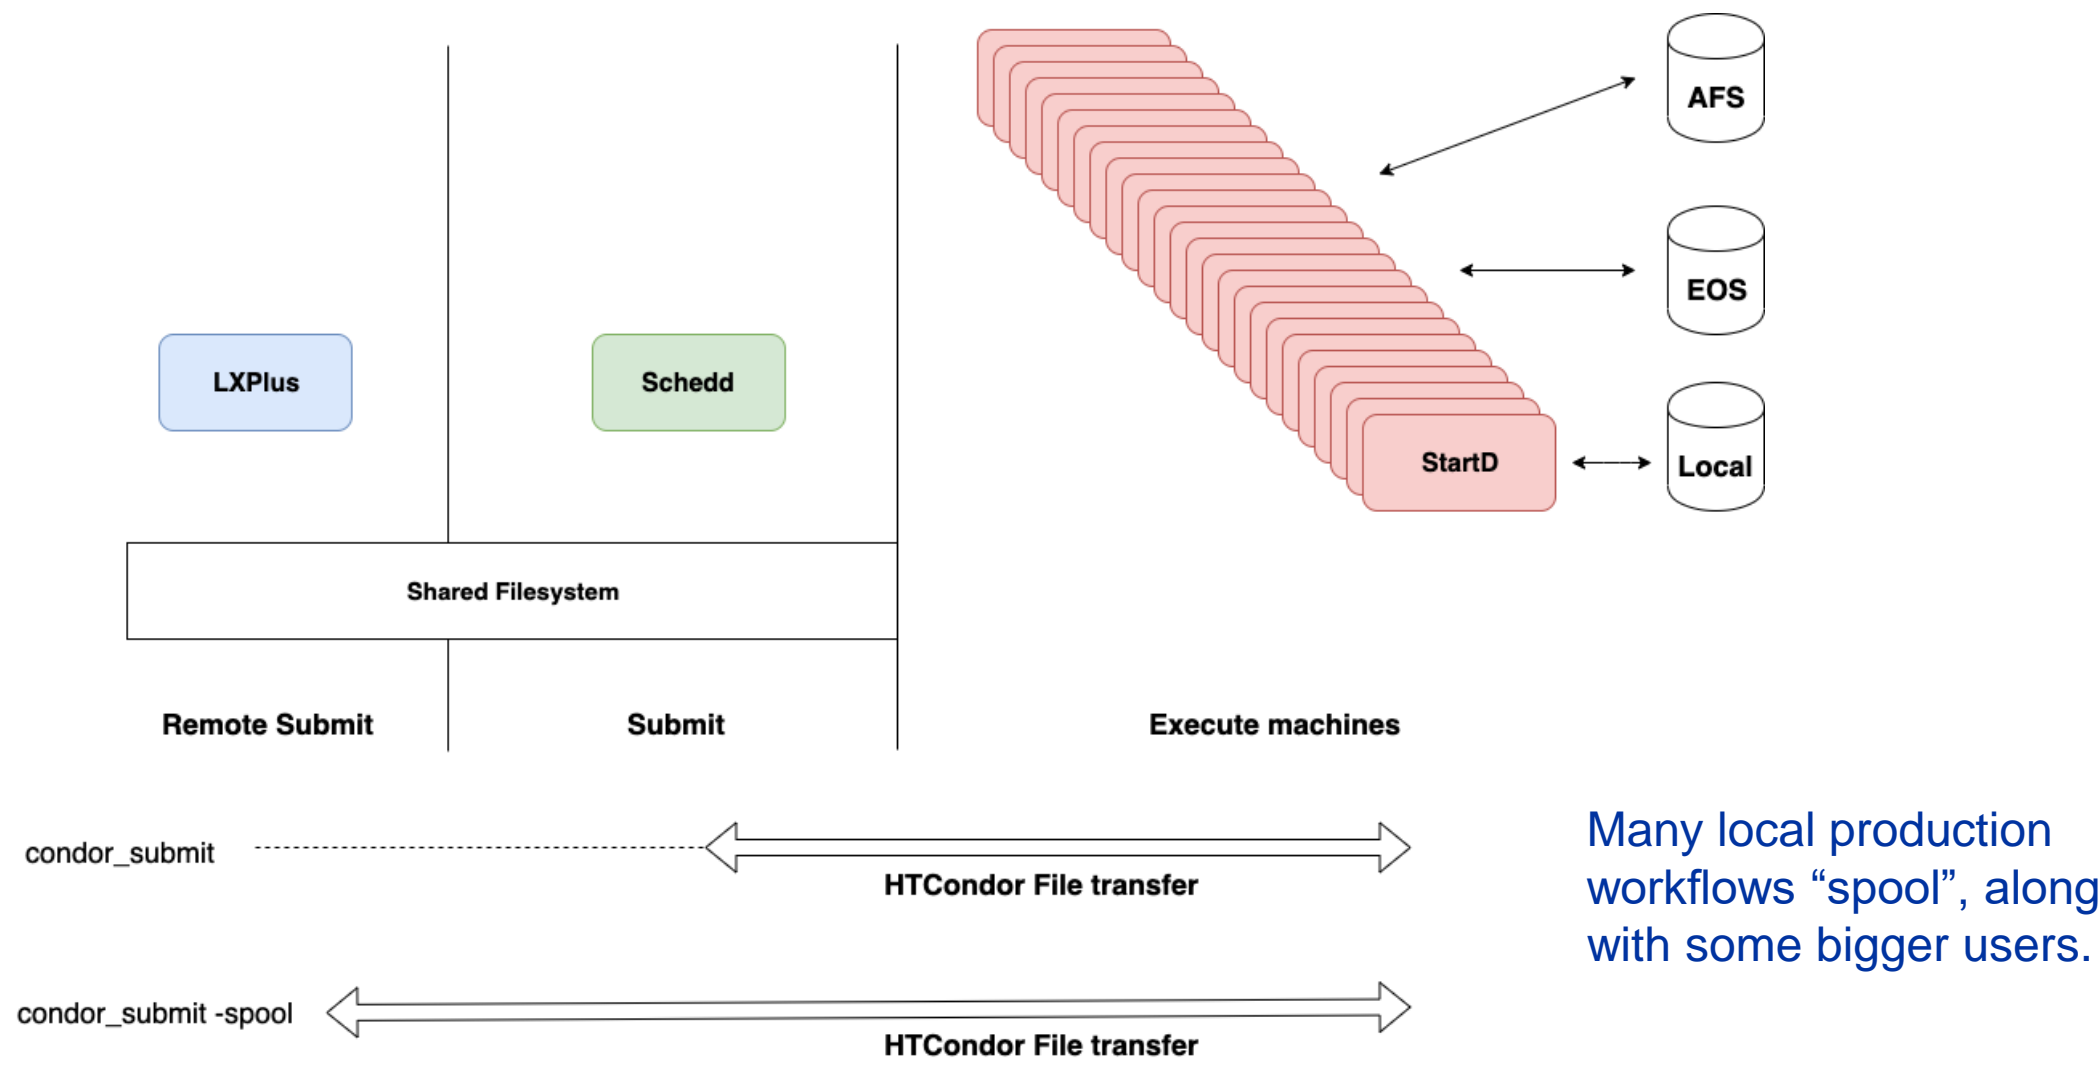
Question does illustrate some of the challenges in running our batch system

- Most of capacity is intended for “production”
- User submitted jobs with low nominal quota can acquire significant bursts of surplus
- Interactivity / responsiveness expectations different with “users” vs “production”
- User jobs more likely to have issues with sudden bursts of scale
- More support requirement for users v “production” or grid.
- User analysis is 10-20% of jobs, but 80% of support overhead*

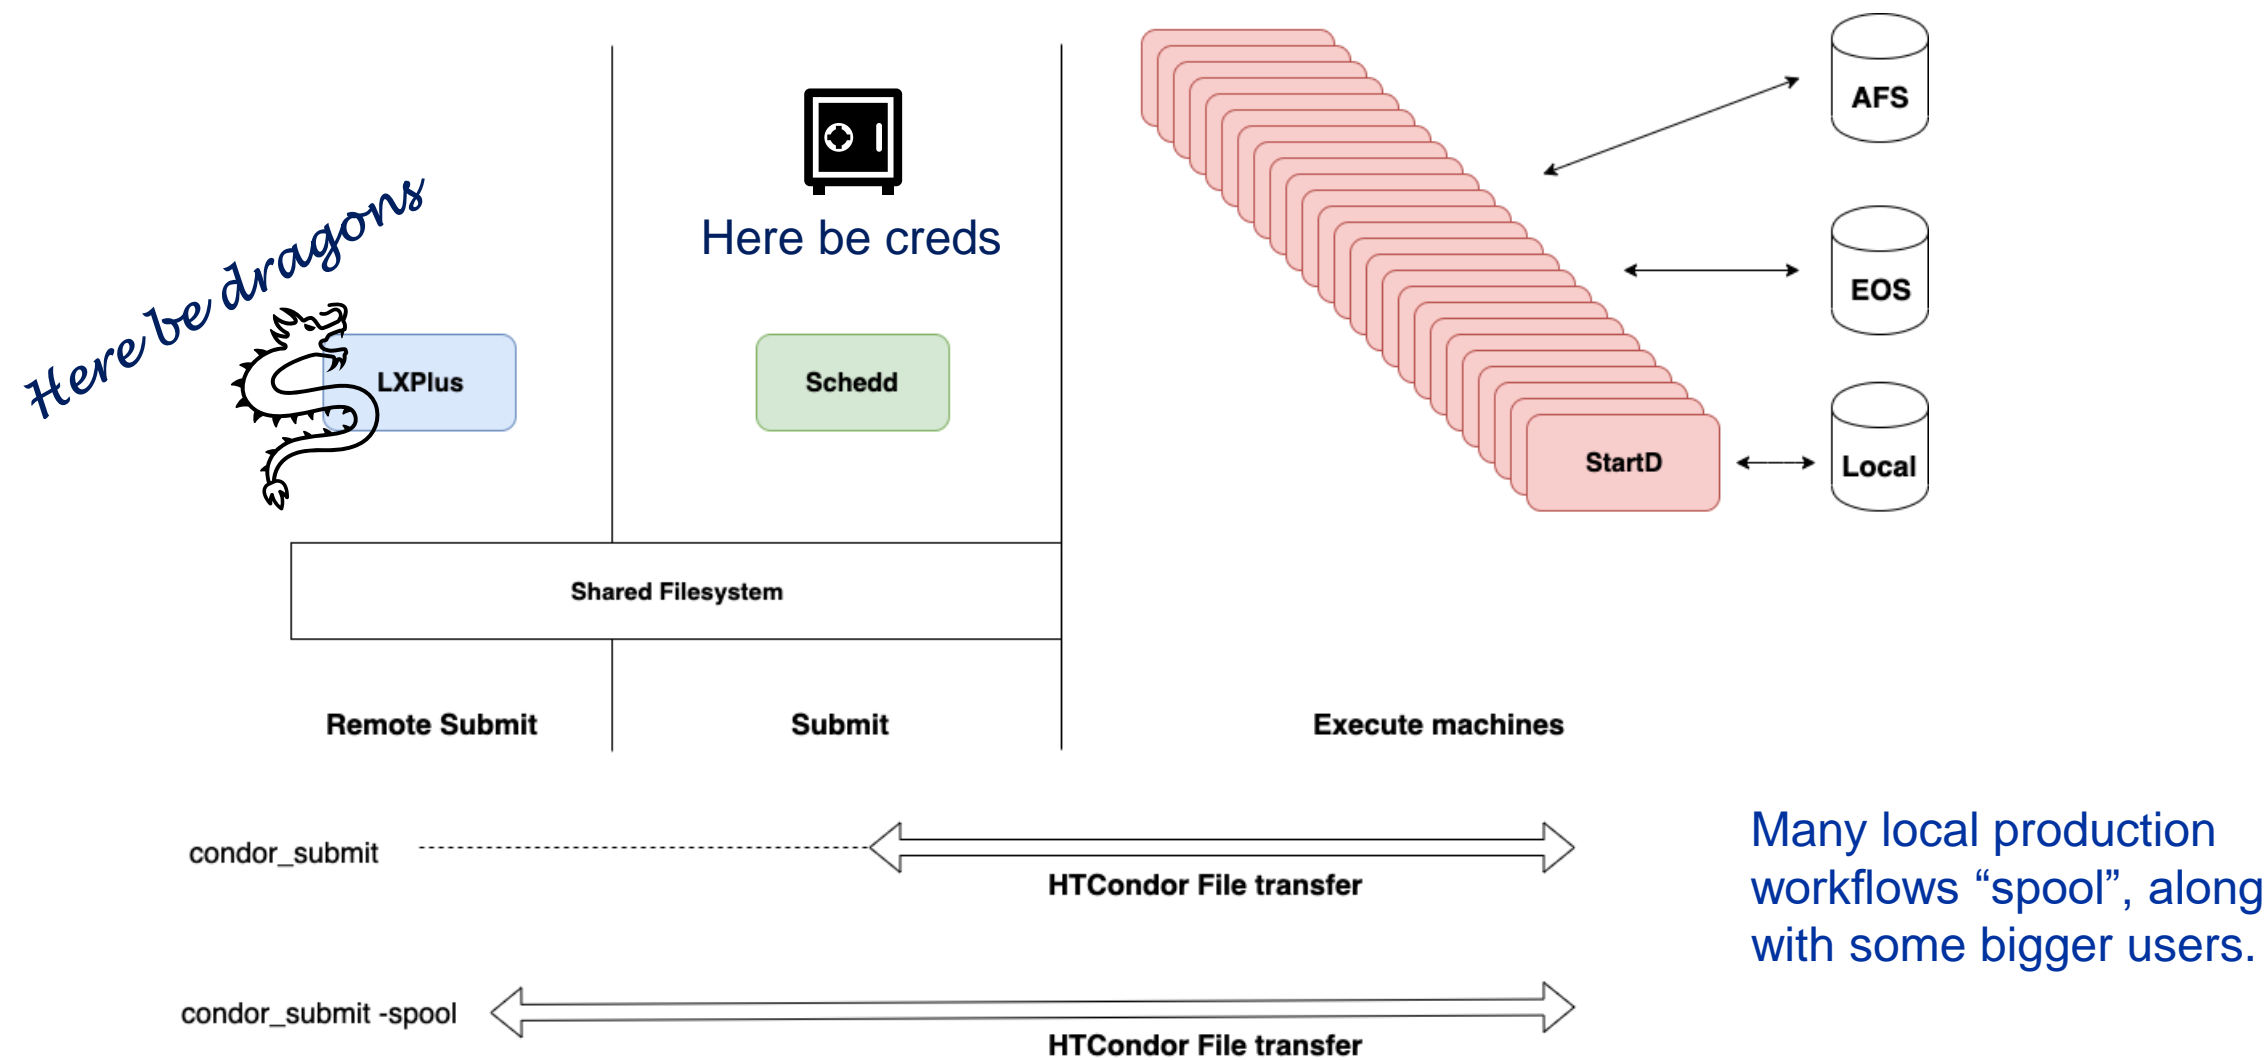
*Yes of course I made up this stat



Remote submission with shared filesystem



Remote submission with shared filesystem



Basic user workflow

- `ssh lxplus.cern.ch`
- `vim supervisors_file.sub`
- `condor_submit supervisors_file.sub`
- `[... wait ...]`
- Results appears as if by magic in `pwd!`



To make that “I type a command here , the data comes back here” interface work, we use shared filesystems

- HTCondor devs hate shared filesystems
- Batch service managers aren't massively fond of shared filesystems
- People who run the shared filesystems aren't overjoyed about batch systems using their filesystems either
- It's really really easy for users to understand*

* until it scales a bit more than they'd anticipated



File transfer plugins...

- **xrdcp plugin for EOS**
- **+/- replicate the simple submission interface**

```
executable = script.sh
log = xfer.$(ClusterId).log
error = yf.$(ClusterId).$(ProcId).err
output = yf.$(ClusterId).$(ProcId).out
output_destination = root://eosuser.cern.ch//eos/user/b/bejones/condor/xfer/$(ClusterId)/
transfer_input_files = root://eosuser.cern.ch//eos/user/b/bejones/condor/file.txt
MY.XRDCP_CREATE_DIR = True
queue
```

Some plugin hacks

- HTCondor makes it very easy to work around road bumps
- Plugin originally transferred contents of sandbox, but Out/Err named `_condor_std*`
- Now plugin just inspects `.job.ad` to name files correctly
- When using `–spool condor_transfer_data` has issues unless we fiddle with the attribute
- With remote submission UserLog still a problem

```
JOB_TRANSFORM_OutputDest @=end
    NAME OutputDest
    REQUIREMENTS (jobUniverse =?= 5 &&
!isUndefined(OutputDestination))
    COPY Out SubmittedOut
    COPY Err SubmittedErr
    COPY OutputDestination
SubmittedOutputDestination
    SET OutputDestination ifThenElse(JobStatus ==
4, undefined, SubmittedOutputDestination)
    SET Out ifThenElse(JobStatus == 4,
"/dev/null", SubmittedOut)
    SET Err ifThenElse(JobStatus == 4,
"/dev/null", SubmittedErr)
@end
```

Less traditional entry points...

- Increasing interest around “analysis facilities” and metaschedulers
- Most interest around Dask, but users often pick / develop other projects
- With Dask we found that we (also!) needed to wrap/subclass it to avoid some assumptions + add some policy
- May help break some of the dependencies on filesystems, but makes some interactivity questions more difficult
- Dask (+ coffea) used from CLI via plus, but trying to improve with notebooks

SWAN: the interface

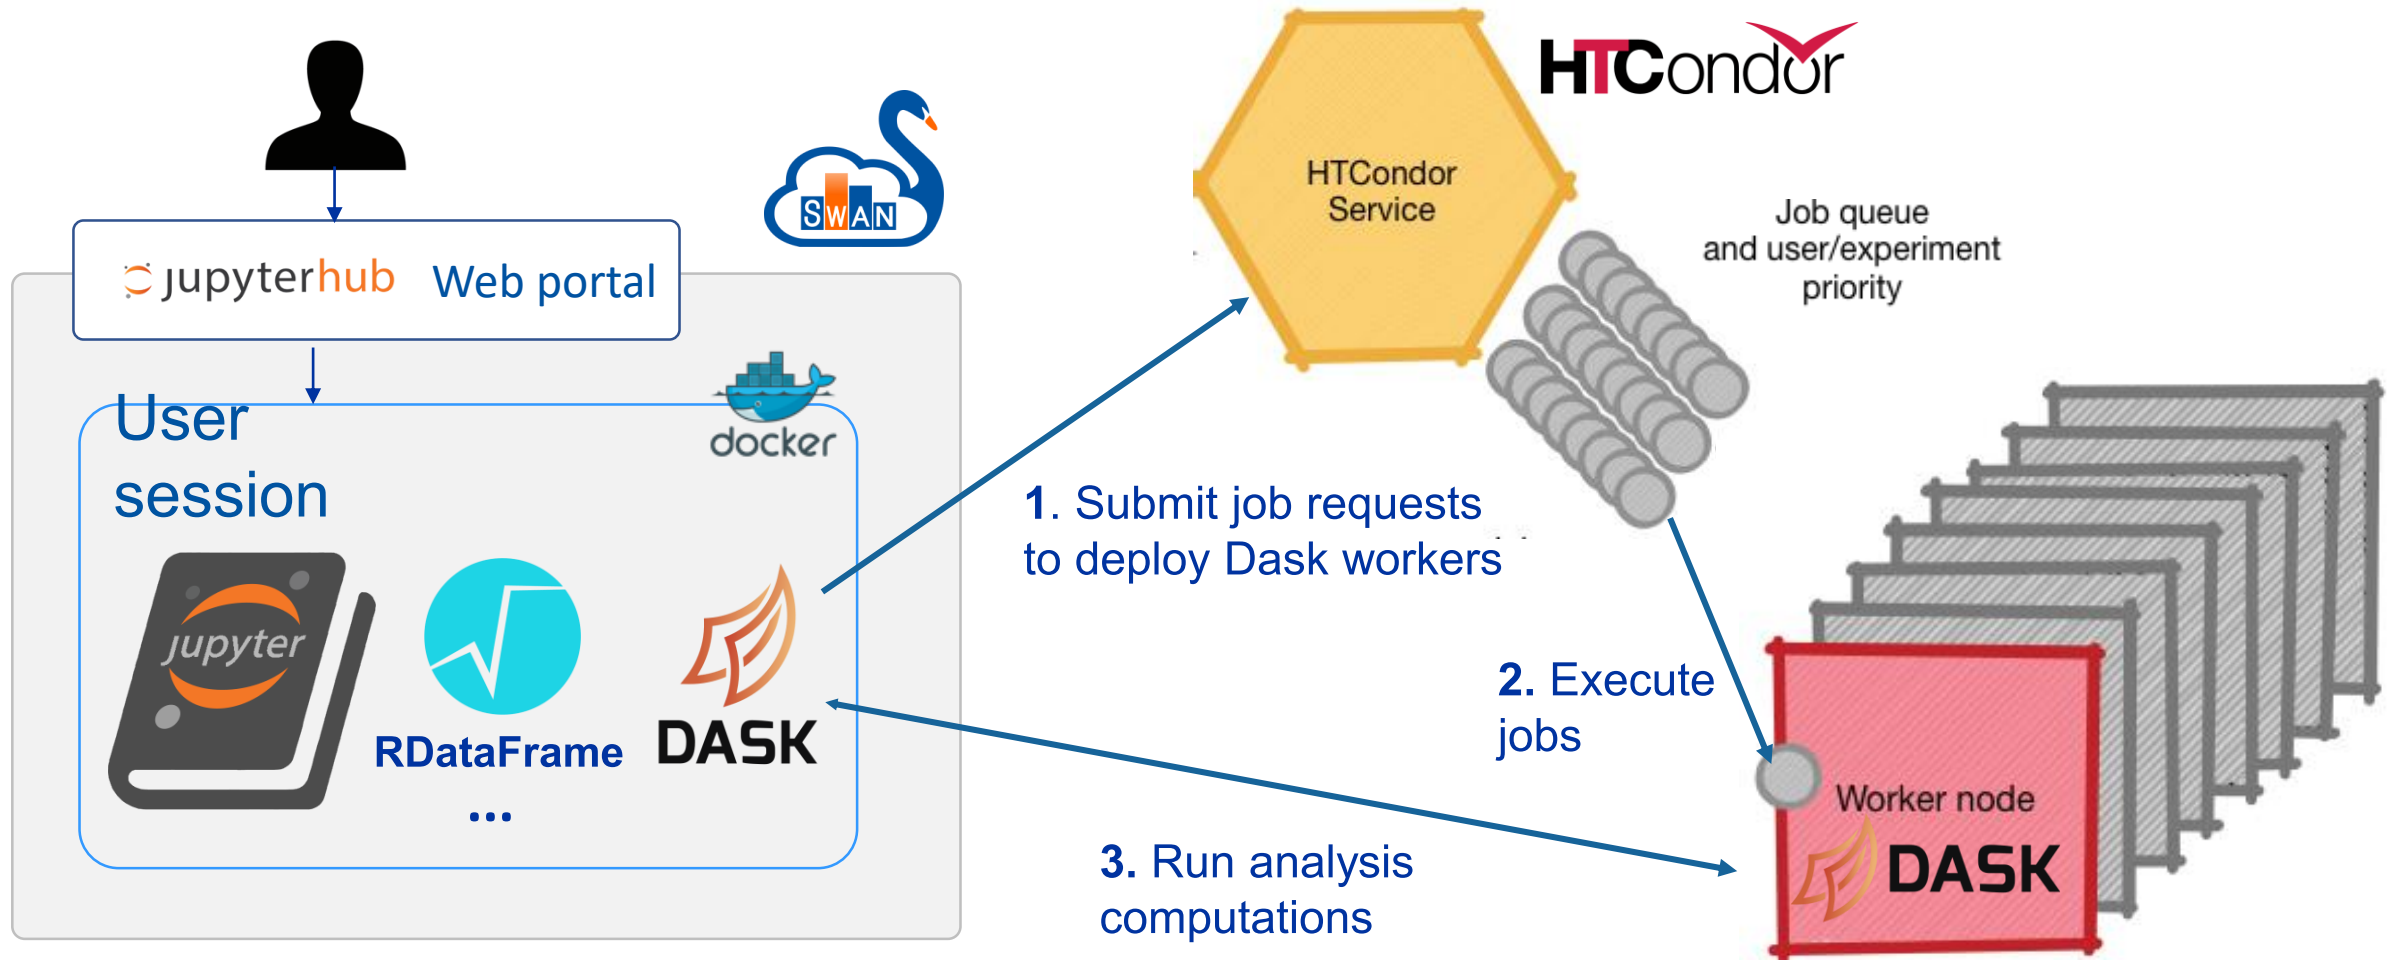
- › **SWAN: Service for Web-based Analysis**
- › **CERN's Jupyter notebook service**
 - Created in 2016
 - Managed jointly by EP and IT
 - Used by 200-250 people daily
- › **Jupyter interface + federation of CERN services → added value!**
 - Software (CVMFS)
 - Storage (EOS, CERNBox)
 - Computing resources (GPU, Spark, HTCondor)
- › **Platform for physics analysis: supports both *single-node* and *distributed* analysis**



SWAN's building blocks



SWAN + HTCondor for interactive analysis



SWAN / Dask integration

- **One issue early on resolved: needed Dask to set a contact_address for the workers to call back to a scheduler running on a k8s cluster**
 - Again: HTCondor transforms gave us flexibility to work around the problem till we had an upstream fix
- **CVMFS helps us to ensure that the Dask scheduler and the workers are running the same software**
 - Simplicity / cache layer
 - Dask / python env in /cvmfs/sft.cern.ch/lcg/... with setup for PYTHONPATH, LD_LIBRARY_PATH etc
 - One of the only times we've found a use for `getenv = True`
- **Still issues around interactivity. SWAN / Jupyter implies more interactivity, where many resources available to users are opportunistic**

Interactivity / low latency

- **Current strategy is to reserve some resources for shorter jobs**
 - ie we use MaxRuntime (as now implemented as `allowed_{job,execute}_duration`)
 - Some machines will only accept jobs $< \$time$
 - Ordinary users encouraged to submit shorter jobs
 - We have sometimes used separate negotiator for very small amounts of resources
- **Other ideas**
 - Dual Startds to allow for a small amount of slots that will take more interactive jobs?
 - Buffer partition with separate netgotiator with low ceiling per user?
 - ?
 - ...
 - No, nobody wants to have jobs preempted

...just one thing about the schedds

- **Random wish list for users & schedds**
 - condor_now but for different schedd
 - condor_now but AccountingGroup based super-users for now-job / vacate-job
 - condor_move_job_to_other_schedd
 - I didn't workshop the name
 - I/O on behalf of user to those nasty shared filesystems can still break schedd (condor_sos is very good though)



Questions?

