Analysis launches on the OSG

- Try to reuse as much infrastructure as possible

SW + data under /cvmfs

OSG

2017-01
REST files

OSG stash

(UConn hosted)

std. container
+ modified
launch.py

JLab farm

Tape

standard merging scripts

scosg20

Start
here

HTCondor submit
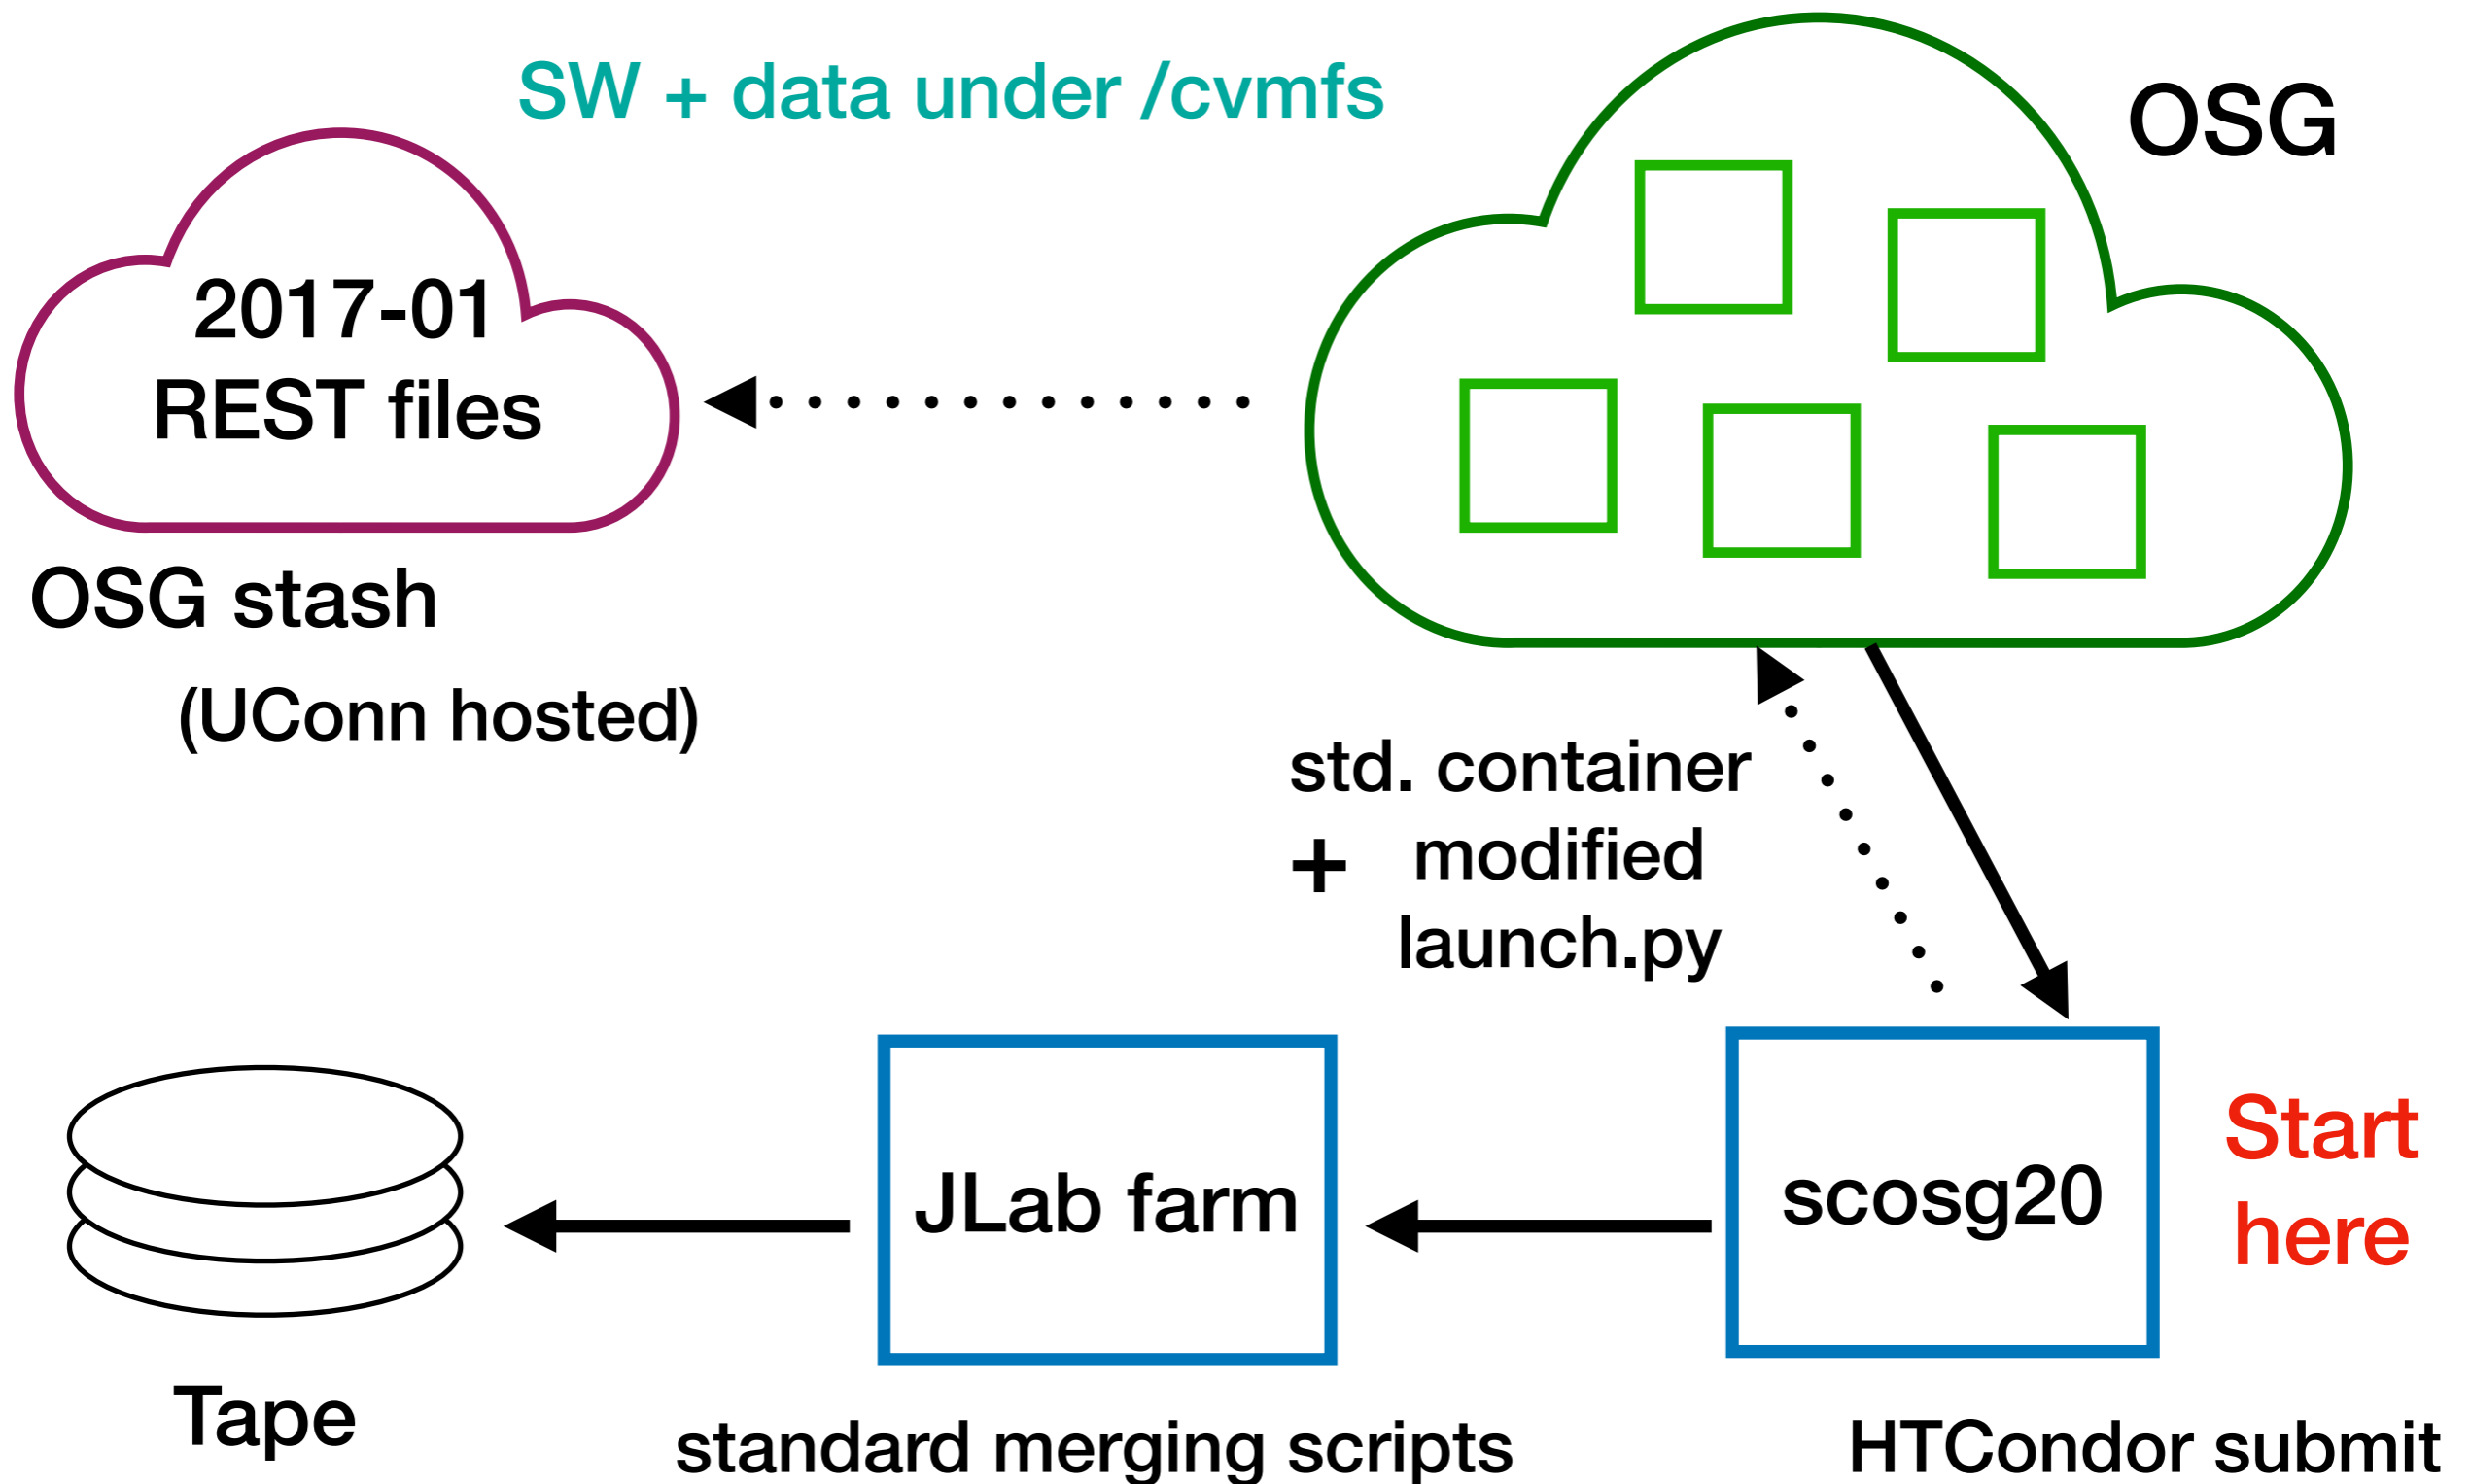
Analysis launches on the OSG

- Started with request from Igal

  - $\gamma p \rightarrow 4\gamma p,\ 4\gamma(p),\ \gamma\gamma(p),\ \gamma\gamma\pi^+\pi^-(p)$

Analysis launches on the OSG

- Started with request from Igal

  - $\gamma p \rightarrow 4\gamma p$, $4\gamma(p)$, $\gamma\gamma(p)$, $\gamma\gamma\pi^+\pi^-(p)$

- Ran 30% of the jobs in 2 days

  - Started filling up disk on scosg20

Analysis launches on the OSG

- Ran 30% of the jobs in 2 days
  - Started filling up disk on scosg20
  - 4 TB / user quota established
  - Requested resources: 1 core + 4 GB of RAM
    - ~5% of jobs exceed this limit
- Wrote data transfer script, copied files to volatile
- Resubmitted jobs, adding better checks to deal with problem nodes that suck up jobs
- Next steps:
  - Looking into reusing parts of MCWrapper for job accounting
  - Change merging scripts to push results to OSG Stash?