



HTCondor Offline: Running on Isolated HPC Systems

Run Everywhere

- › Use all available resources
 - Submit locally
 - Run globally

Foreign Languages


- › Using remote resources
- › Easy with HTCondor and friendly admins
 - Flocking
- › Harder otherwise
 - Lost jobs
 - Unknown queue times
 - Different allocation expectations
 - Different level of service

Paper Over the Differences

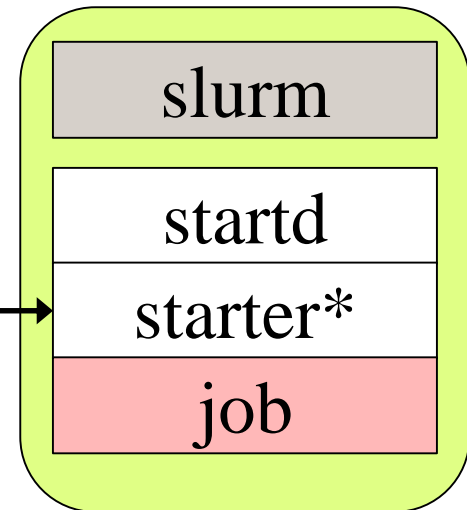
- › Don't send user jobs directly
- › Make everything act like HTCondor
- › Glideins
 - Run HTCondor startd as a job/image/container
 - Send jobs to it
- › Turn remote resources into temporary members of your own HTCondor pool
- › See also Annex

Glidein Example

Home



Away



* One shadow, starter per running job

Isolated HPC Systems

- › Limited remote access
 - Job submission
 - Data transfer
 - Limited or no general network communication
- › High-bandwidth shared filesystem
 - Both external and internal access
 - We can move jobs in addition to their data

Three Approaches to Moving Jobs via the FileSystem

- › Condor-F
- › Split-Starter
- › Lumberjack

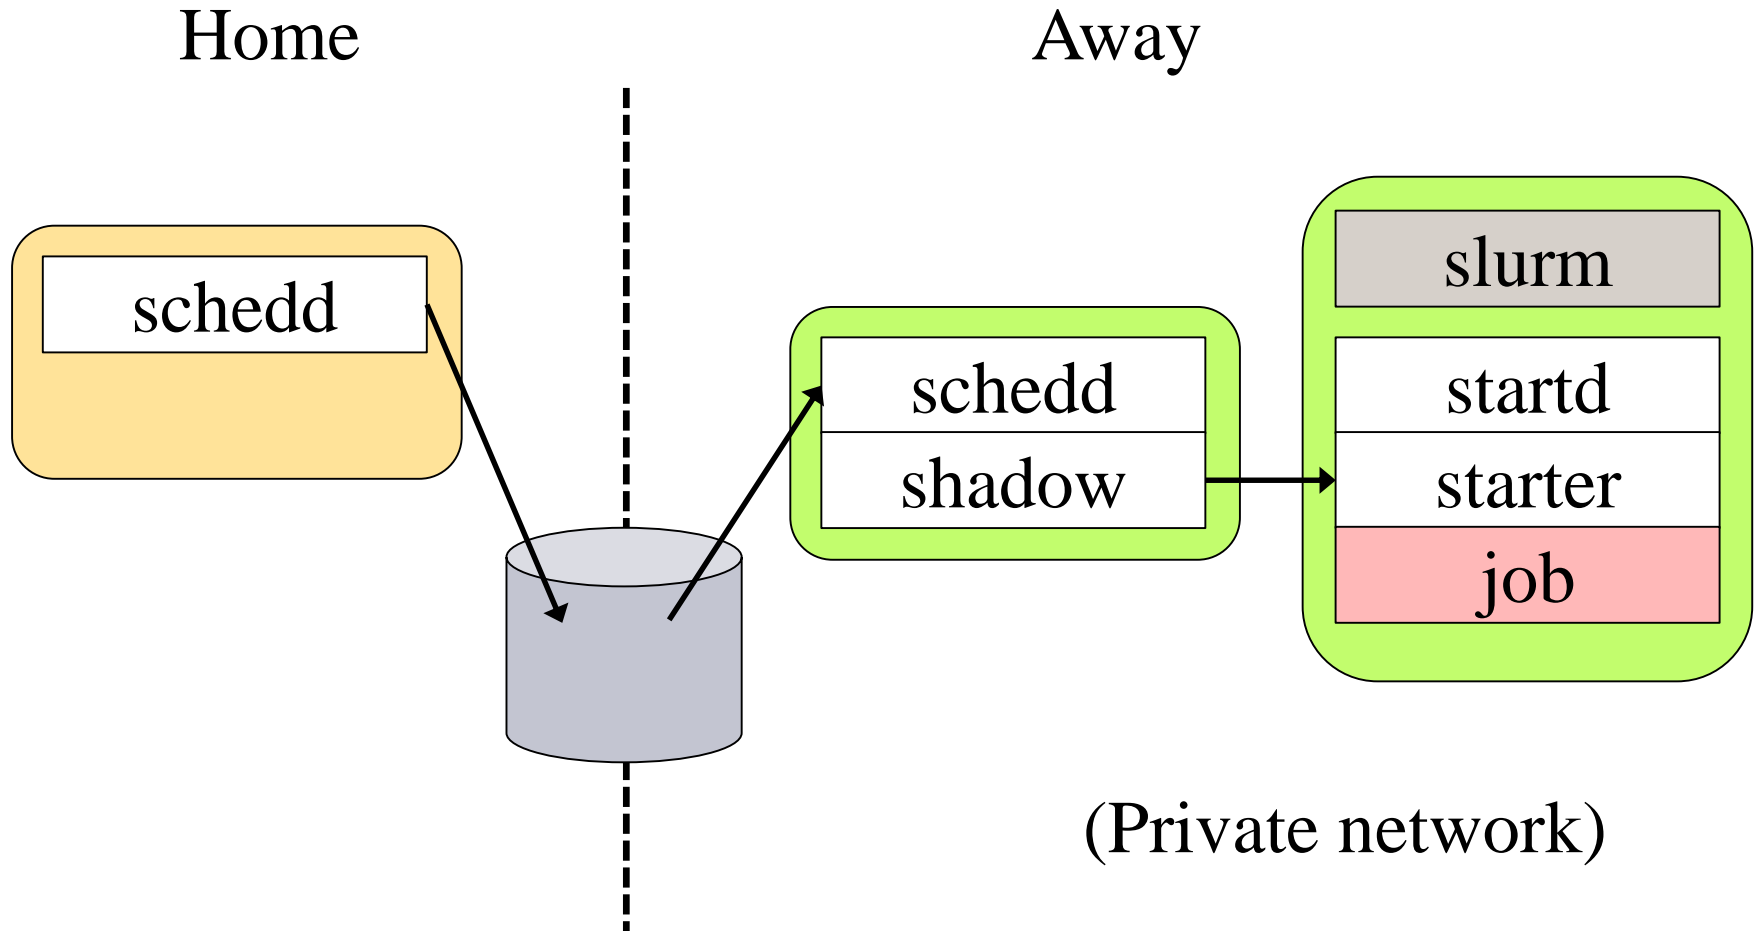
Condor-F(filesystem)

- › Similar to Condor-C
- › Communicate via files
- › Support any file sharing mechanism
 - NFS, Gluster, GPFS
 - Box.com, Google Drive, Dropbox
 - Gridftp, xrootd, rsync

Job Management Primitives

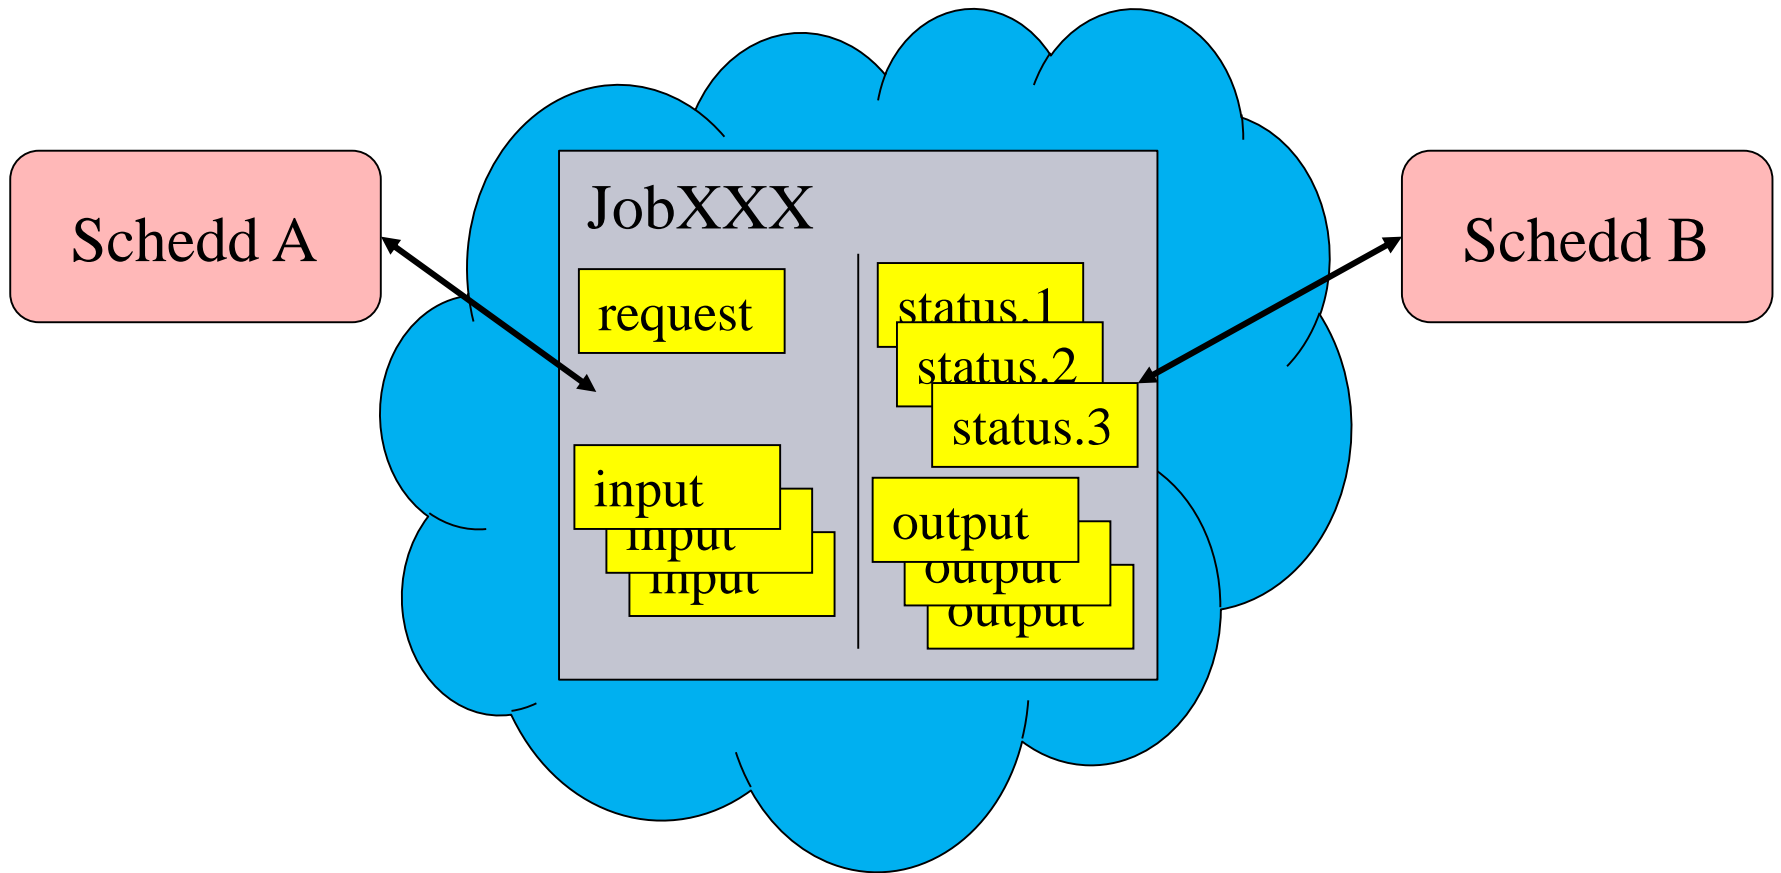
- › Submit
 - Write job description and input sandbox
- › Status (optional)
 - Write status description
- › Completion
 - Write final status description and output sandbox
- › Cleanup or Removal
 - Delete job description and input sandbox

Condor-F Example



File-Based Job Submission

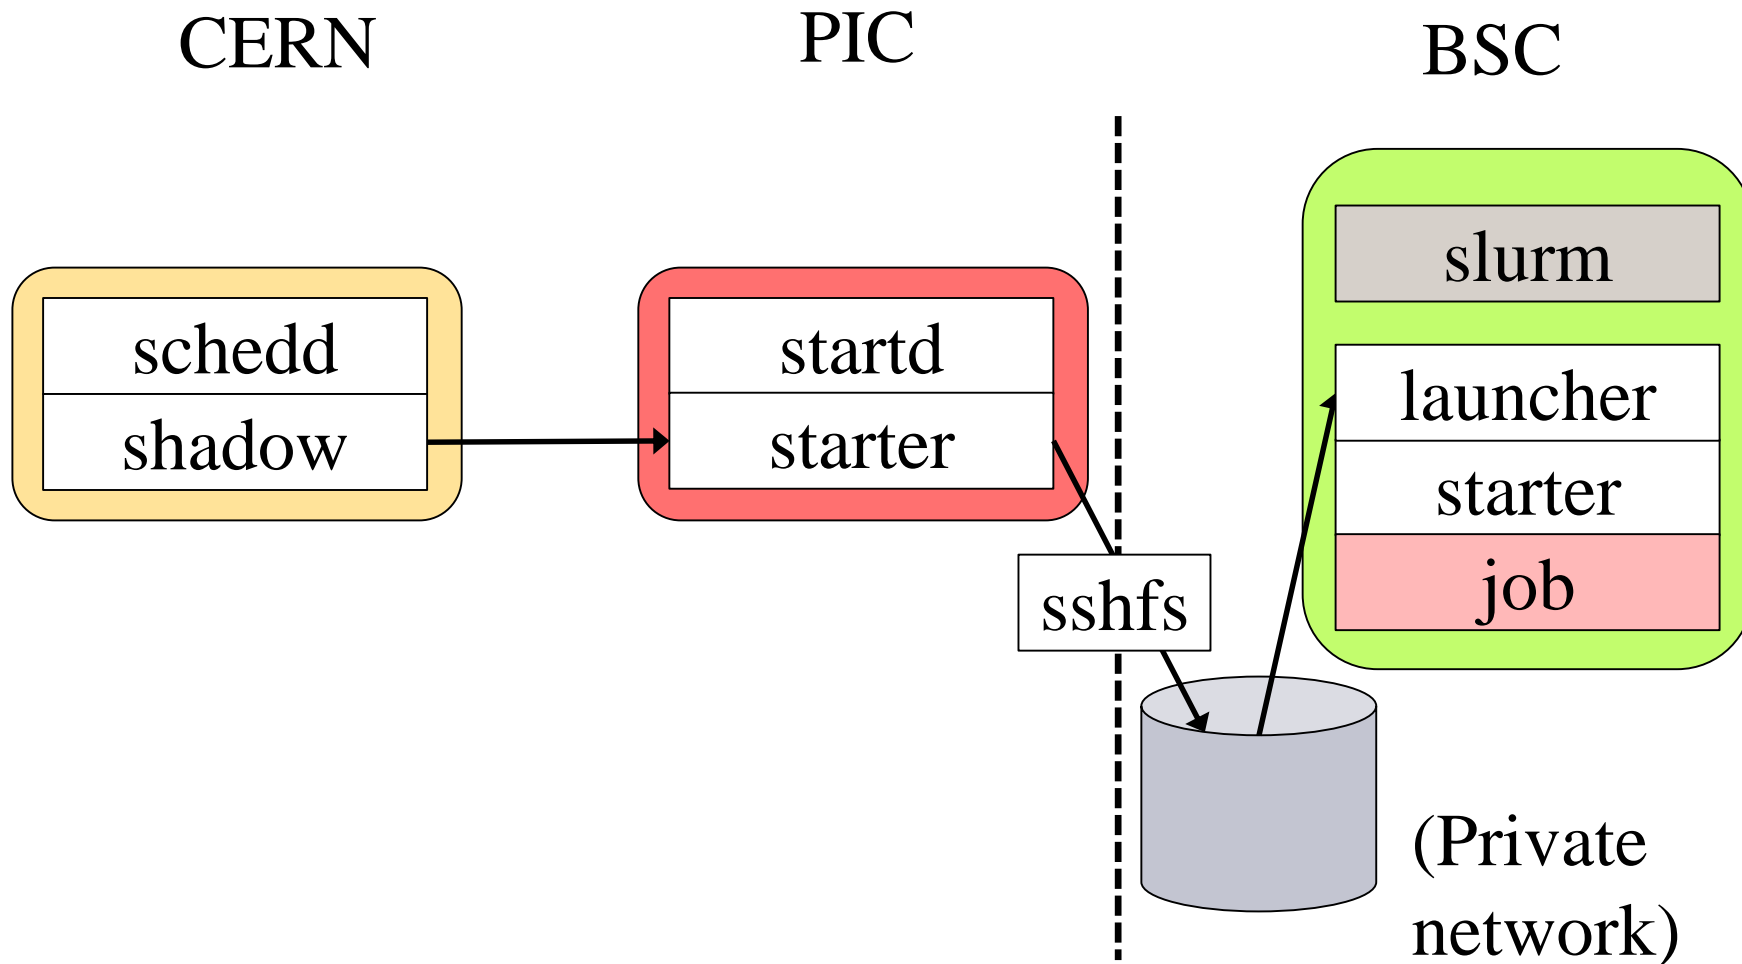
Box Cloud Storage



Split-Starter

- › Proxy startd on public network
 - Accepts jobs
 - Starter writes job + data to filesystem
- › On execute node
 - Job + data copied to local disk
 - Standalone starter runs job
 - Results written back to filesystem

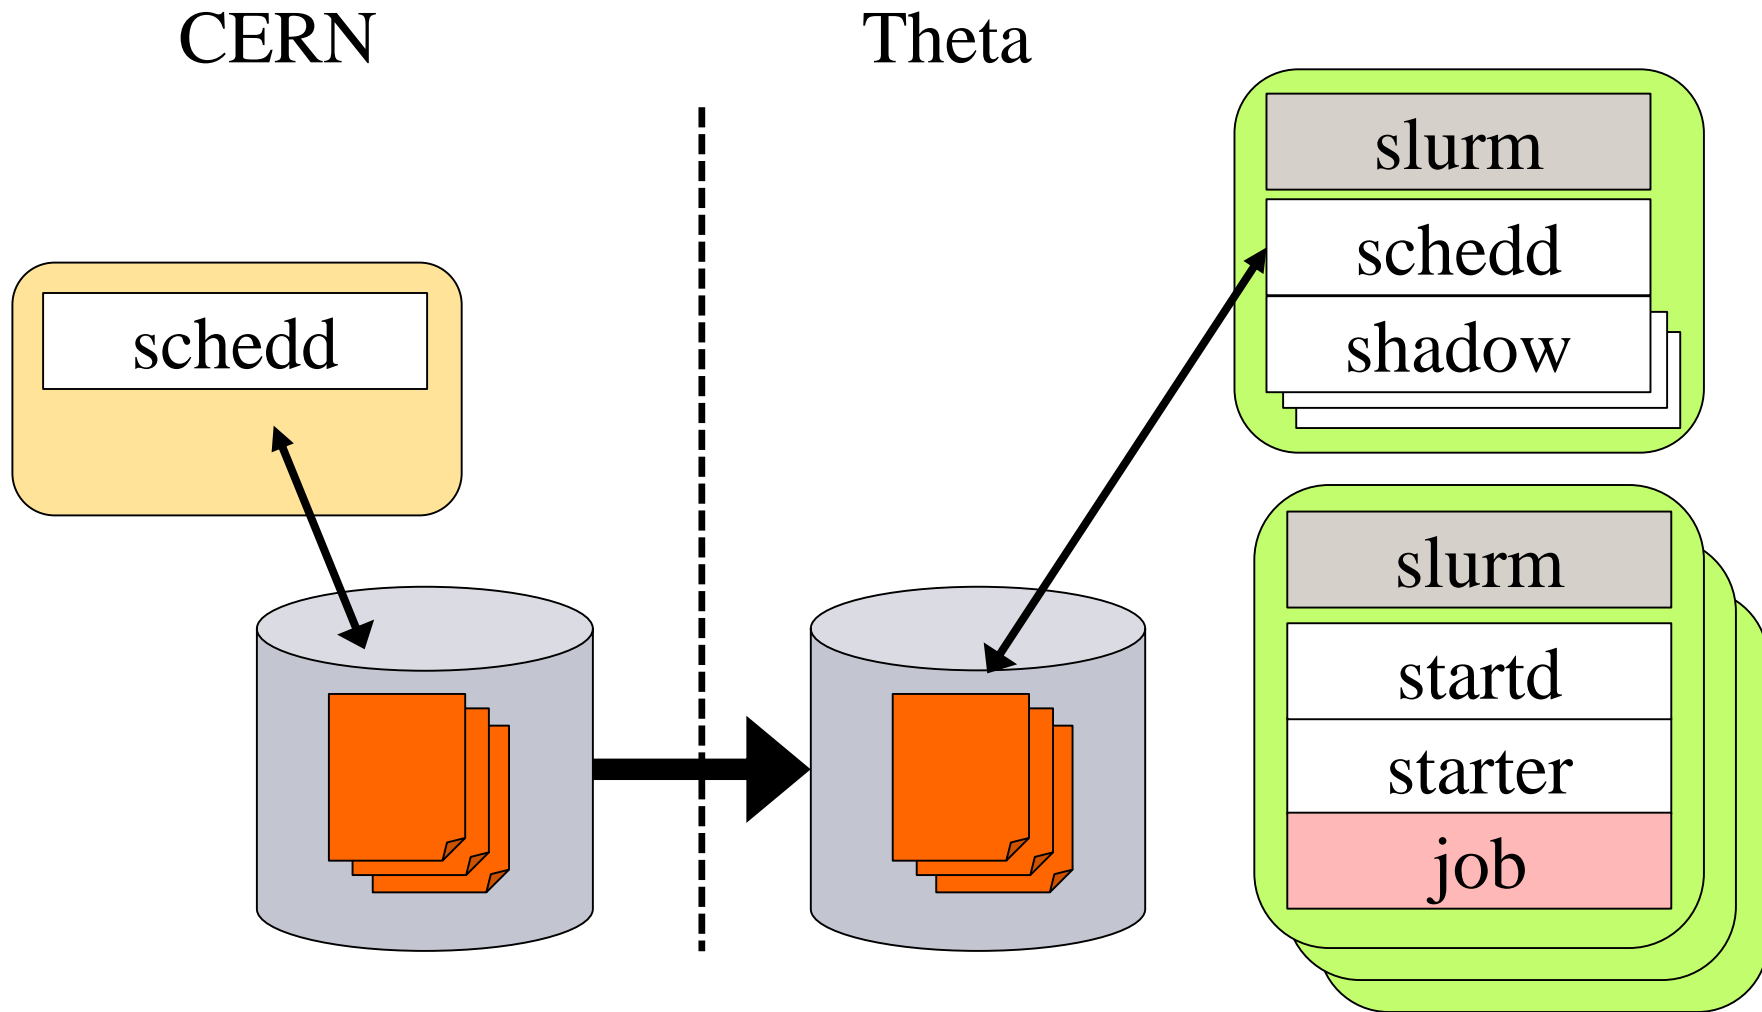
Split-Starter Example



Lumberjack

- › “It moves the logs”
- › Pause jobs in home schedd
- › Transfer job queue and input files
 - Create export job queue file
- › Run an HTCondor pool as HPC job
- › Transfer job queue and output files
 - Merge into home schedd’s queue/spool
- › Unpause jobs in home schedd

Lumberjack Example



Comparison

- › Condor-F
 - Like Condor-C
 - Uses login node (bad)
- › Split-Starter
 - Invisible outside of starter
 - Keeps late binding
- › Lumberjack
 - Just two bulk data transfers

Questions?