



HTCondor's Grid Universe

Jaime Frey (jfrey@cs.wisc.edu)

**Center for High Throughput Computing
Department of Computer Sciences
University of Wisconsin-Madison**

HTCondor-G(rid)

- › HTCondor for the grid
 - Same job management capabilities as a local HTCondor pool
 - Use other scheduling systems' resources

Job Management Interface

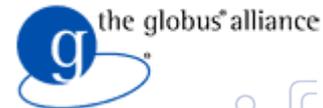
- › Local, persistent job queue
- › Job policy expressions
- › Job activity logs
- › Workflows (DAGMan)
- › Fault tolerance

“Grid” Universe

- › All handled in your submit file
- › Supports a number of “back end” types:
 - Globus GRAM
 - CREAM
 - NorduGrid ARC
 - HTCondor
 - PBS
 - LSF
 - SLURM
 - UNICORE
 - EC2



Condor
High Throughput Computing



UNICORE



“Grid” is a Misnomer

- › Grid universe used any time job is submitted to a different job management/queuing system
 - Grid services
 - GRAM, CREAM, ARC
 - Different HTCondor schedd
 - Local batch system
 - PBS, LSF, SLURM, SGE
 - Cloud service
 - EC2, Google Compute Engine

Gridmanager Daemon

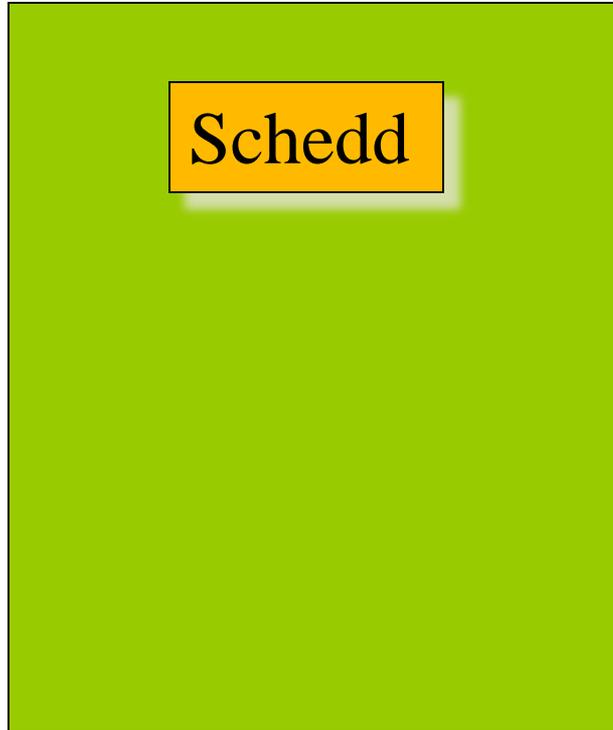
- › Runs under the schedd
- › Similar to the shadow
- › Handles all management of grid jobs
- › Single instance manages all grid jobs for a user

Grid ASCII Helper Protocol (GAHP)

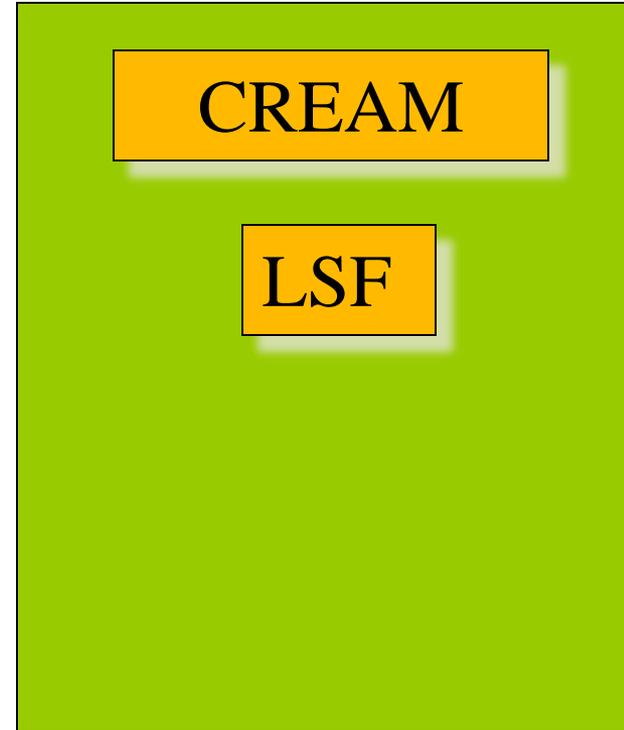
- › Runs under gridmanager
- › Encapsulates grid client libraries in separate process
- › Simple ASCII protocol
- › Easy to use client libraries when they can't be linked directly with gridmanager

How It Works

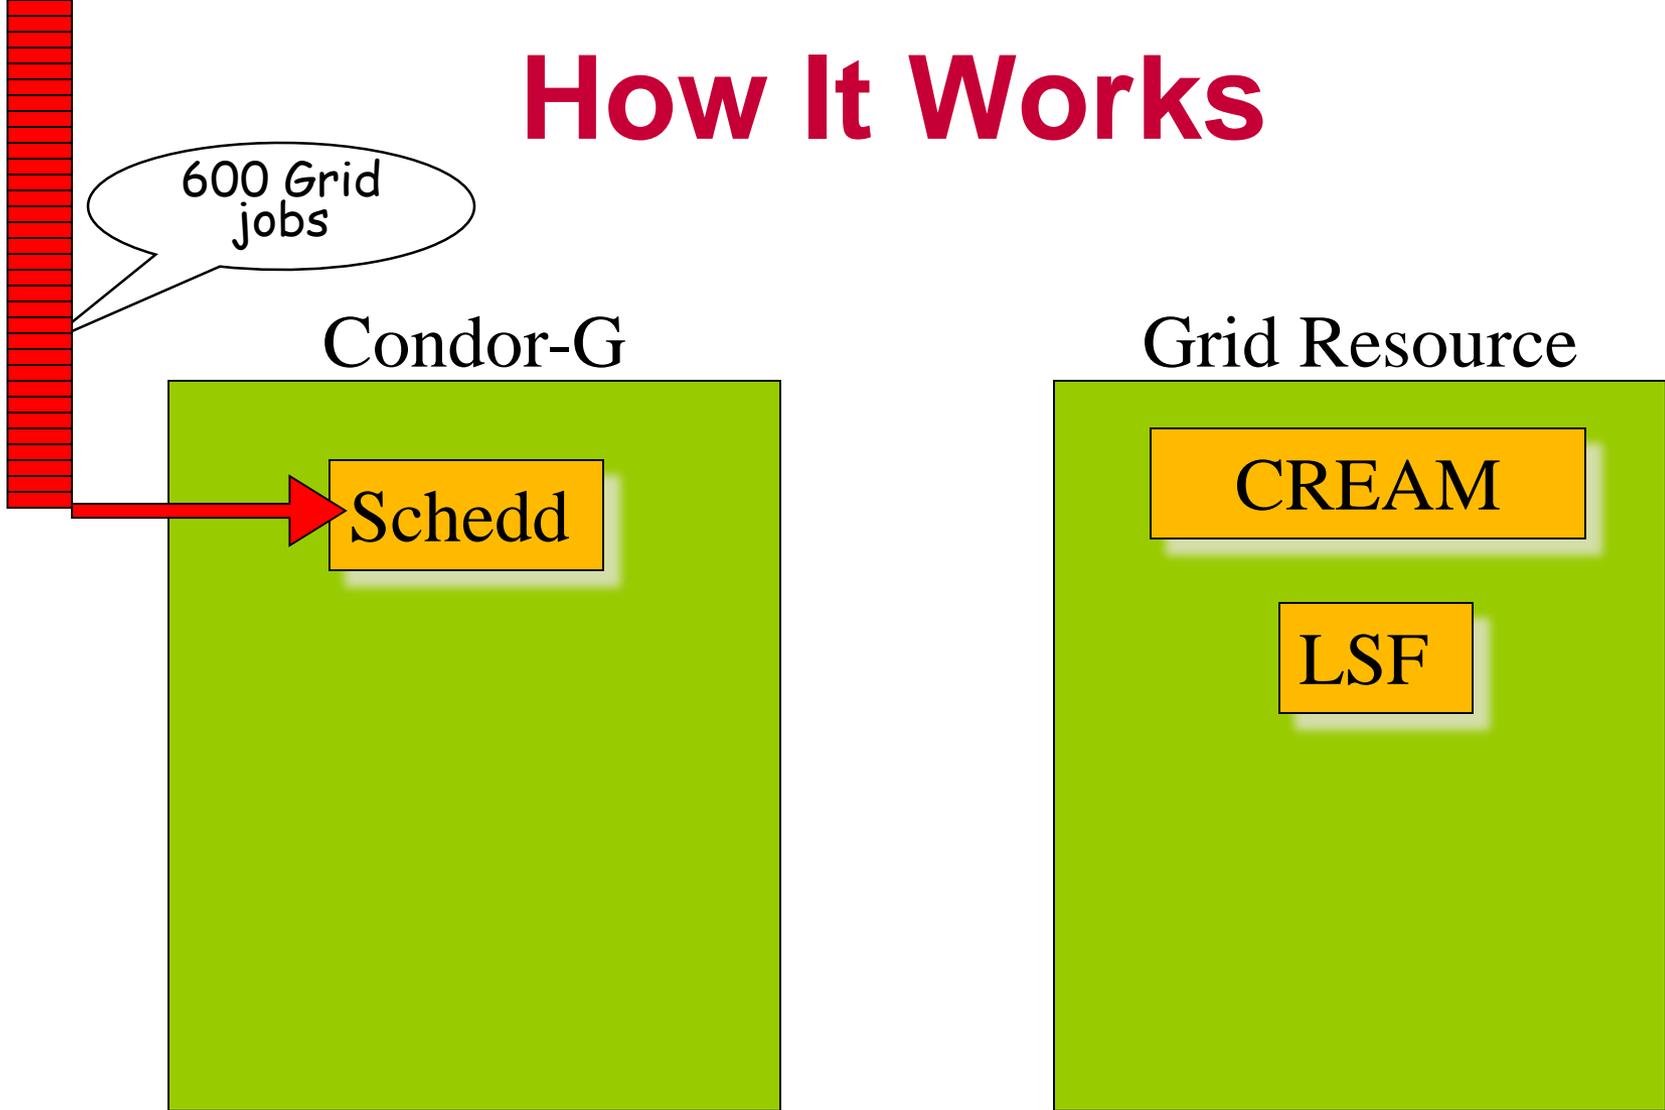
Condor-G



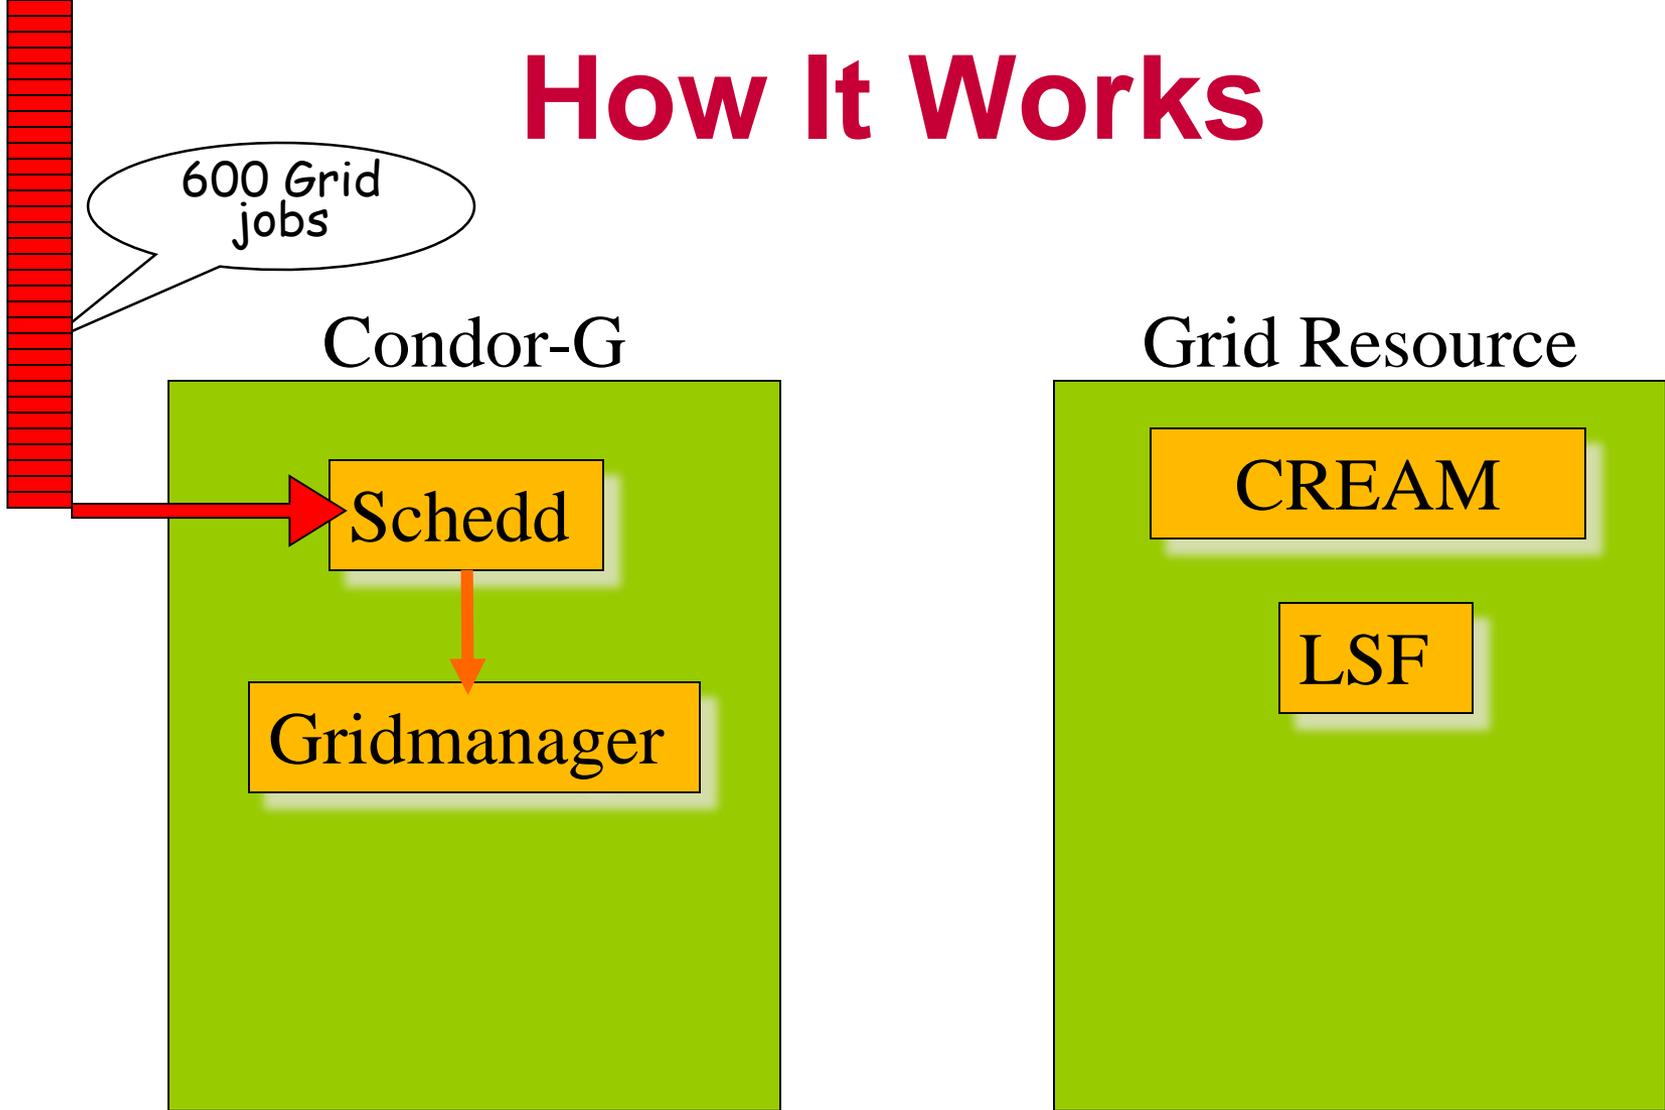
Grid Resource



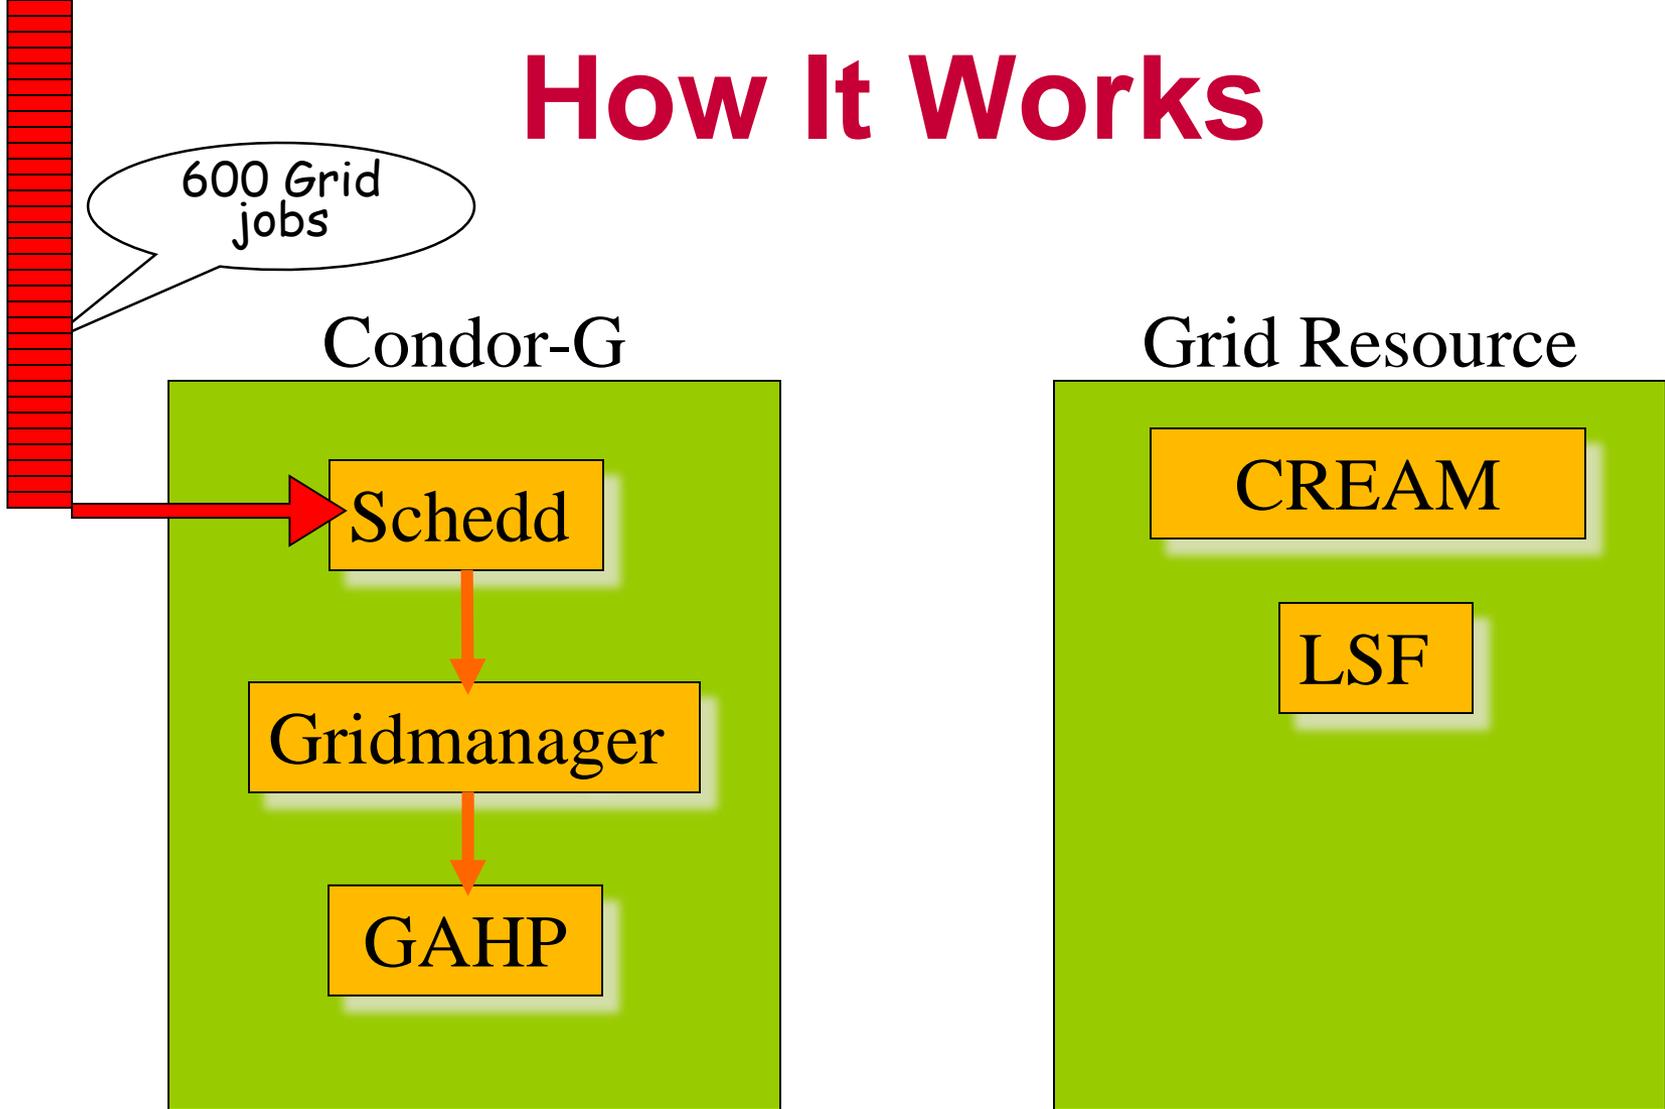
How It Works



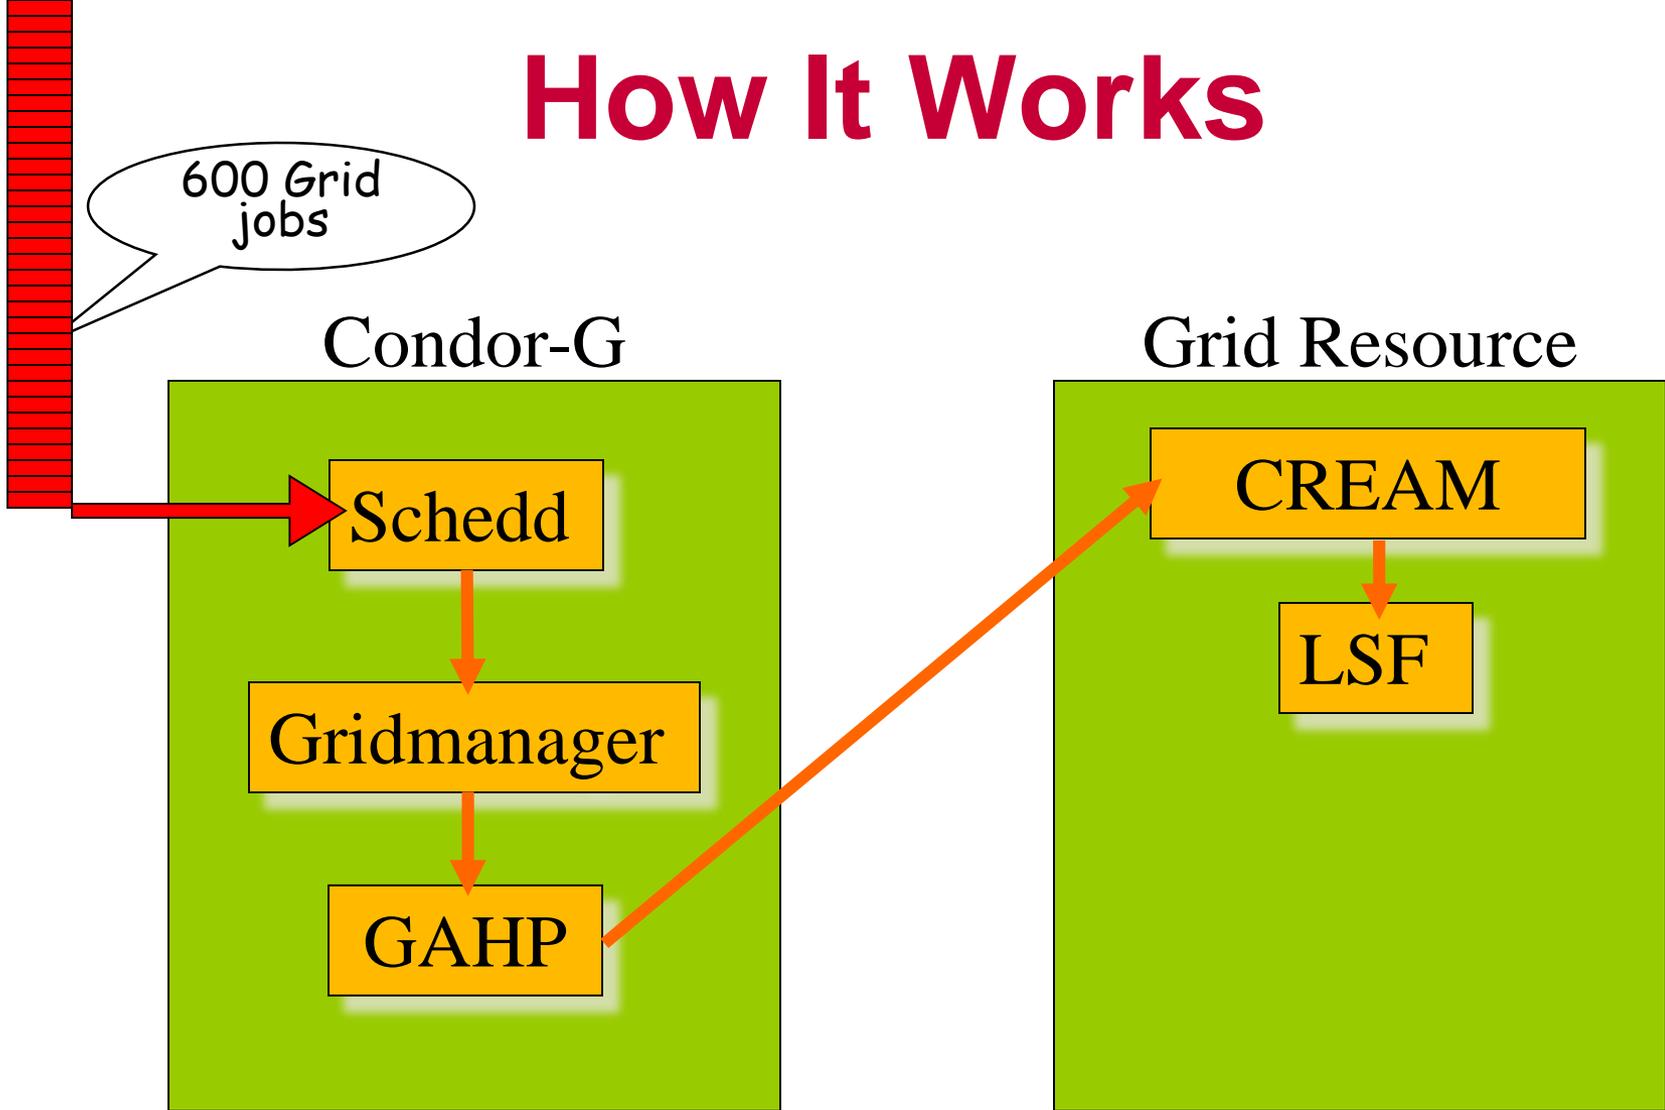
How It Works



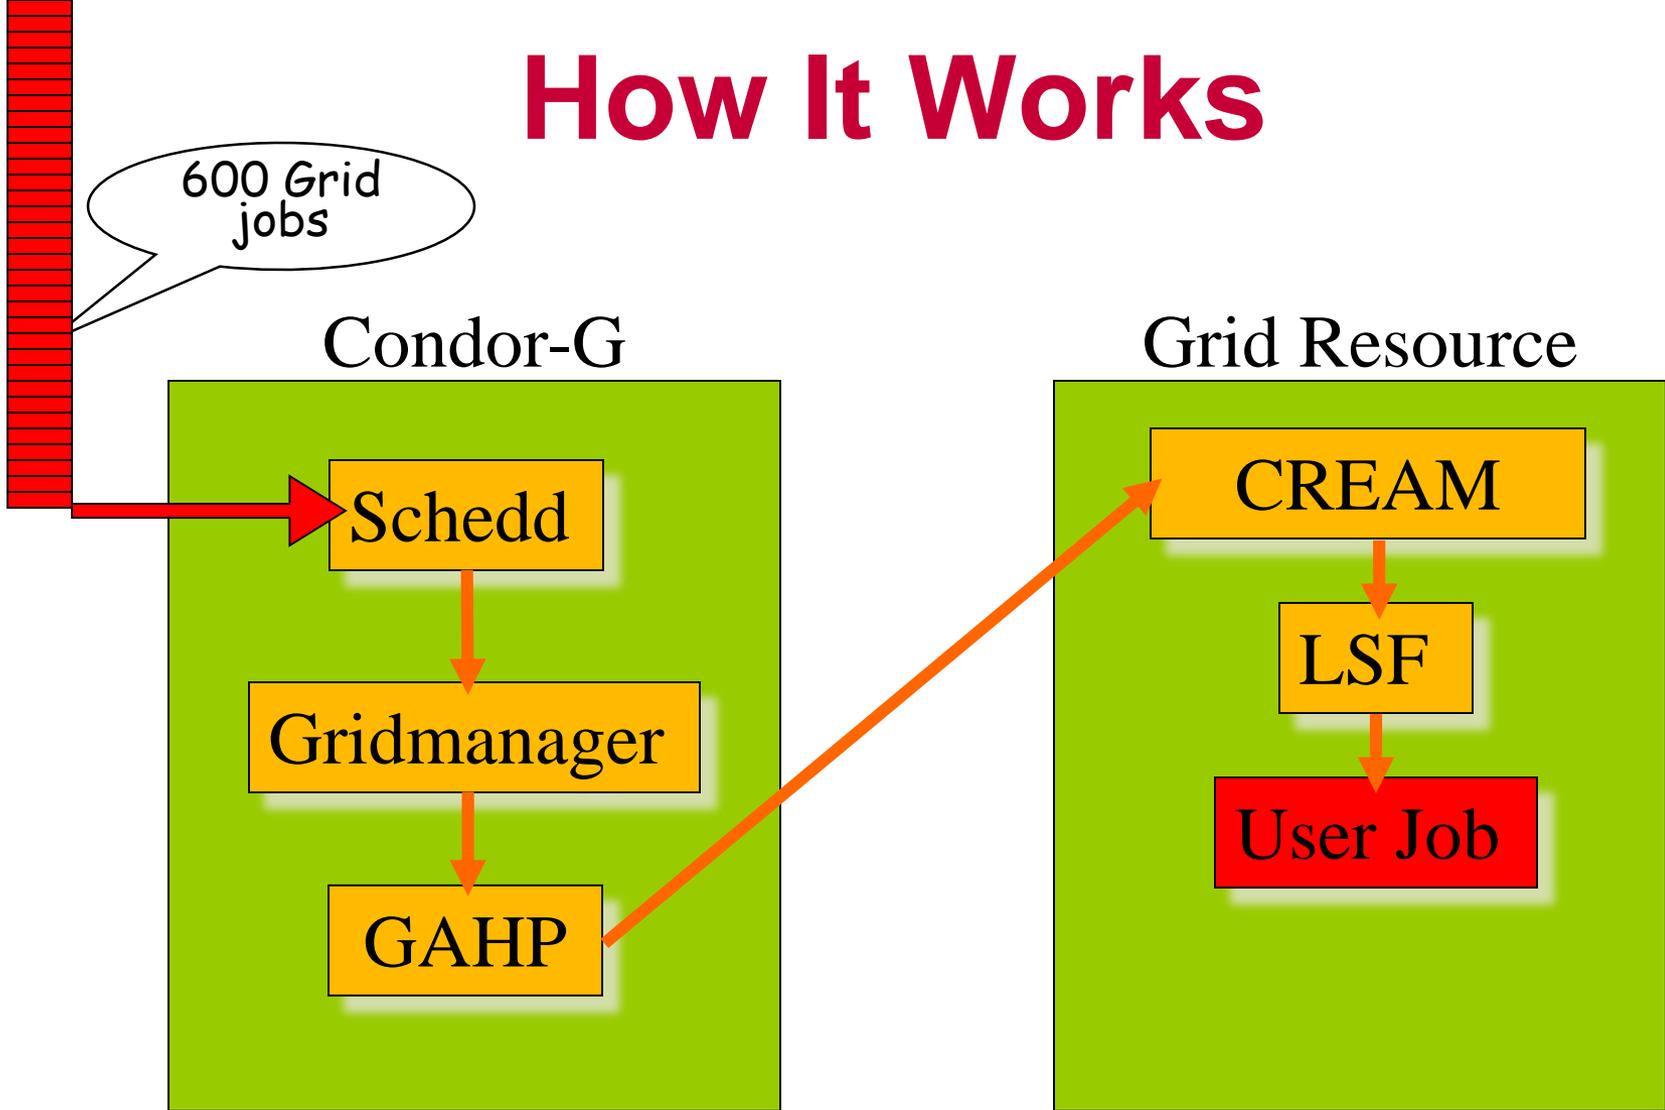
How It Works



How It Works



How It Works



Differences from Normal HTCondor Jobs

- › Vanilla universe
 - Resource acquisition
 - Job sent directly to machine that can start running it immediately
- › Grid universe
 - Job delegation
 - Job sent to alternate scheduling system
 - Job may sit idle while resources available elsewhere

Differences from Normal HTCondor Jobs

- › No matchmaking
 - Specify destination with GridResource
 - No Requirements, Rank
 - Resource requests often ignored
- › Run-time features unavailable
 - condor_ssh_to_job
 - condor_tail
 - condor_chirp

Differences from Normal HTCondor Jobs

- › Information about job execution often lacking
 - Job exit code
 - Runtime
 - Resource usage
- › Usually must name output files

Grid Job Attributes

- › GridResource = “<grid type> <server> ...”
- › Specify service type and server location
- › Examples
 - `condor submit.foo.edu cm.foo.edu`
 - `gt5 portal.foo.edu/jobmanager-pbs`
 - `cream creamce.foo.edu/cream-pbs-glow`
 - `nordugrid arc.foo.edu`
 - `batch pbs`

Grid Job Attributes

- › GridJobId
 - Job ID from remote server
- › GridJobStatus
 - Job status from remote server
- › LastRemoteStatusUpdate
 - Time job status last checked
- › GridResourceUnavailableTime
 - Time when remote server went down

Network Connectivity

- › Outbound connections only for most job types
- › GRAM requires incoming connections
 - Need 2 open ports per <user, X509 DN> pair

Authentication

- › GAHP acts as the user with the remote service
- › Destination service is local
 - UID-based authentication
 - E.g. PBS, LSF, SLURM, Condor
- › Destination service is remote
 - X.509 proxy (possibly with VOMS attributes)
 - Automatically forward refreshed proxy
 - E.g. Condor, GRAM, CREAM, ARC

HELD Status

- › Jobs will be held when HTCondor needs help with an error
- › On release, HTCondor will retry
- › The reason for the hold will be saved in the job ad and user log

Hold Reason

> condor_q -held

```
161.0    jfrey          2/13 13:58 CREAM_Delegate Error:  
Received NULL fault;
```

> cat job.log

```
012 (161.000.000) 02/13 13:58:38 Job was held.
```

```
    CREAM_Delegate Error: Received NULL fault; the error  
is due to another cause..
```

> condor_q -format '%s\n' HoldReason

```
CREAM_Delegate Error: Received NULL fault; the error is due  
to another cause..
```

Common Errors

› Authentication

- Hold reason may be misleading
- User may not be authorized by CE
- Condor-G may not have access to all Certificate Authority files
- User's proxy may have expired

Common Errors

- › CE no longer knows about job
 - CE admin may forcibly remove job files
 - Condor-G is obsessive about not leaving orphaned jobs
 - May need to take extra steps to convince Condor-G that remote job is gone

Thank You!

Throttles and Timeouts

- › Limits that prevent Condor-G or CEs from being overwhelmed by large numbers of jobs
- › Defaults are fairly conservative

Throttles and Timeouts

- > **GRIDMANAGER_MAX_SUBMITTED_JOBS_PER_RESOURCE = 1000**
 - You can increase to 10,000 or more
- > **GRIDMANAGER_JOB_PROBE_INTERVAL**
 - Default is 60 seconds
 - Can decrease, but not recommended

Throttles and Timeouts

- > **GRIDMANAGER_MAX_PENDING_REQUESTS = 50**
 - Number of commands sent to a GAHP in parallel
 - Can increase to a couple hundred
- > **GRIDMANAGER_GAHP_CALL_TIMEOUT = 300**
 - Time after which a GAHP command is considered failed
 - May need to lengthen if pending requests is increased