

HTCondor batch @farm.particle.cz

Petr Vokáč

FZU Farm users meeting 2018
20th February

Motivation

- Really ancient batch used @ FZU farm
 - Torque 2.4.16 (2011) + Maui 3.2.6p21 (2009)
 - Scalability issues (even for queued jobs)
 - Slow development (“new” features) & bugfixes
 - Scheduler (FIFO + job priority → no real fairshare without maui/moab)
- Upgrade necessary
 - Torque + ? / PBSPro + ? / HTCondor / SLURM / SGE / LSF
 - HTCondor used by bigger sites
 - replacing other batch system
 - better suited for HTC workload common in HEP
 - scalability, resource usage, advanced features (cgroups, ns, vm, ...)
 - active development with inputs from HEP community
 - very flexible configuration
- FZU farm – HTCondor 8.6.x (upgrade to 8.7.x soon)

Simple HTCondor job

hello.sh

```
#!/bin/sh  
echo "Hello $HOSTNAME"
```

Simple HTCondor job

- Shell

```
$ sh hello.sh  
Hello ui.farm.particle.cz
```

hello.sh

```
#!/bin/sh  
echo "Hello $HOSTNAME"
```

Simple HTCondor job

- Torque / PBS

```
$ qsub hello.sh
42.torque.farm.particle.cz
$ qstat
42 hello.sh user Q queue
...
42 hello.sh user R queue
...
42 hello.sh user C queue
$ cat hello.sh.[oe]42
Hello wn.farm.particle.cz
```

hello.sh

```
#!/bin/sh
echo "Hello $HOSTNAME"
```

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit  
1 job(s) submitted to \  
cluster 42.0
```

hello.sh

```
#!/bin/sh  
echo "Hello $HOSTNAME"
```

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit  
1 job(s) submitted to \  
cluster 42.0
```

hello.sh

```
#!/bin/sh  
echo "Hello $HOSTNAME"
```



?

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit  
1 job(s) submitted to \  
cluster 42.0
```

hello.sh

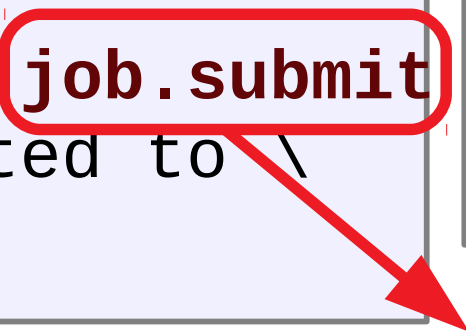
```
#!/bin/sh  
echo "Hello $HOSTNAME"
```

```
#!/bin/sh  
#PBS -q short  
#PBS -l mem=2GB  
#PBS -l nodes=1:ppn=4  
echo "Hello $HOSTNAME"
```


Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit  
1 job(s) submitted to \  
cluster 42.0
```



hello.sh

```
#!/bin/sh  
echo "Hello $HOSTNAME"
```

job.submit

```
executable = hello.sh  
queue
```

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit  
1 job(s) submitted to \  
cluster 42.0
```

JobId = ClusterId.ProcId

hello.sh

```
#!/bin/sh  
echo "Hello $HOSTNAME"
```

job.submit

```
executable = hello.sh  
queue
```

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit
1 job(s) submitted to \
cluster 42.0
$ condor_q
```

hello.sh

```
#!/bin/sh
echo "Hello $HOSTNAME"
```

job.submit

```
executable = hello.sh
queue
```

- waiting for empty job slot

```
[user@ui ~]$ condor_q
-- Schedd: condor.farm.particle.cz : <147.231.25.153:9618> @ 02/20/18
OWNER BATCH_NAME          SUBMITTED   DONE    RUN    IDLE  TOTAL  JOB_IDS
user  CMD: hello.sh      2/20 16:13    _     _     1     1     42.0

1 jobs; 0 completed, 0 removed, 1 idle, 0 running, 0 held, 0 suspended
```

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit
1 job(s) submitted to \
cluster 42.0
$ condor_q
```

hello.sh

```
#!/bin/sh
echo "Hello $HOSTNAME"
```

job.submit

```
executable = hello.sh
queue
```

- job is running on worker node

```
[user@ui ~]$ condor_q
-- Schedd: condor.farm.particle.cz : <147.231.25.153:9618> @ 02/20/18
OWNER BATCH_NAME          SUBMITTED   DONE    RUN    IDLE  TOTAL  JOB_IDS
user  CMD: hello.sh        2/20 16:13    _     1     _     1     42.0

1 jobs; 0 completed, 0 removed, 0 idle, 1 running, 0 held, 0 suspended
```

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit
1 job(s) submitted to \
cluster 42.0
$ condor_q
```

hello.sh

```
#!/bin/sh
echo "Hello $HOSTNAME"
```

job.submit

```
executable = hello.sh
queue
```

- job finished its execution

```
[user@ui ~]$ condor_q
-- Schedd: condor.farm.particle.cz : <147.231.25.153:9618> @ 02/20/18
OWNER BATCH_NAME          SUBMITTED   DONE    RUN    IDLE  TOTAL  JOB_IDS
user  CMD: hello.sh        2/20 16:13    1      _     _     1     42.0

1 jobs; 1 completed, 0 removed, 0 idle, 0 running, 0 held, 0 suspended
```

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit
1 job(s) submitted to \
cluster 42.0
$ condor_q
```

- ...

```
$ ls
hello.sh
job.submit
```

hello.sh

```
#!/bin/sh
echo "Hello $HOSTNAME"
```

job.submit

```
executable = hello.sh
queue
```



?

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit
1 job(s) submitted to \
cluster 42.0
$ condor_q
```

- ...

```
$ ls
hello.sh
hello.stdout
hello.stderr
job.submit
$ cat hello.std{out,err}
Hello wn.farm.particle.cz
```



hello.sh

```
#!/bin/sh
echo "Hello $HOSTNAME"
```

job.submit

```
executable = hello.sh
output = hello.stdout
error = hello.stderr
queue
```

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit
1 job(s) submitted to \
cluster 42.0
$ condor_q
```

- ...

```
$ ls
hello.sh
hello.stdout
hello.stderr
job.submit
$ cat hello.std{out,err}
Hello wn.farm.particle.cz
```

hello.sh

```
#!/bin/sh
echo "Hello $HOSTNAME"
```

job.submit

```
executable = hello.sh
output = hello.stdout
error = hello.stderr
queue
```

Done?

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit
1 job(s) submitted to \
cluster 42.0
$ condor_q
```

- ...

```
$ ls
hello.sh
hello.stdout
hello.stderr
job.submit
$ cat hello.std{out,err}
Hello wn.farm.particle.cz
```

hello.sh

```
#!/bin/sh
echo "Hello $HOSTNAME"
```

job.submit

```
executable = hello.sh
output = hello.stdout
error = hello.stderr
+MaxRuntime = 1200
request_cpus = 1
request_memory = 2GB
queue
```

only 8 core jobs supported now @farm
(restriction will be removed in future)

Simple HTCondor job

- HTCondor

```
$ condor_submit job.submit
1 job(s) submitted to \
cluster 42.0
$ condor_q
```

- ...

```
$ ls
hello.sh
hello.stdout
hello.stderr
job.submit
$ cat hello.std{out,err}
Hello wn.farm.particle.cz
```

hello.sh

```
#!/bin/sh
echo "Hello $HOSTNAME"
```

job.submit

```
executable = hello.sh
output = hello.stdout
error = hello.stderr
+MaxRuntime = 1200
request_cpus = 1
request_memory = 2GB
queue
```

only 8 core jobs supported now @farm
(restriction will be removed in future)

Job management & batch status

Torque / PBS	HTCondor
qsub	condor_submit
qdel	condor_rm
qstat	condor_q
qselect	condor_q
qhold	condor_hold
qrerun	condor_release
qalter	condor_qedit
pbsnodes	condor_status
tracejob	condor_history
ssh	ssh / condor_ssh_to_job
while qselect loop	condor_wait
?	condor_<TAB>

Job states



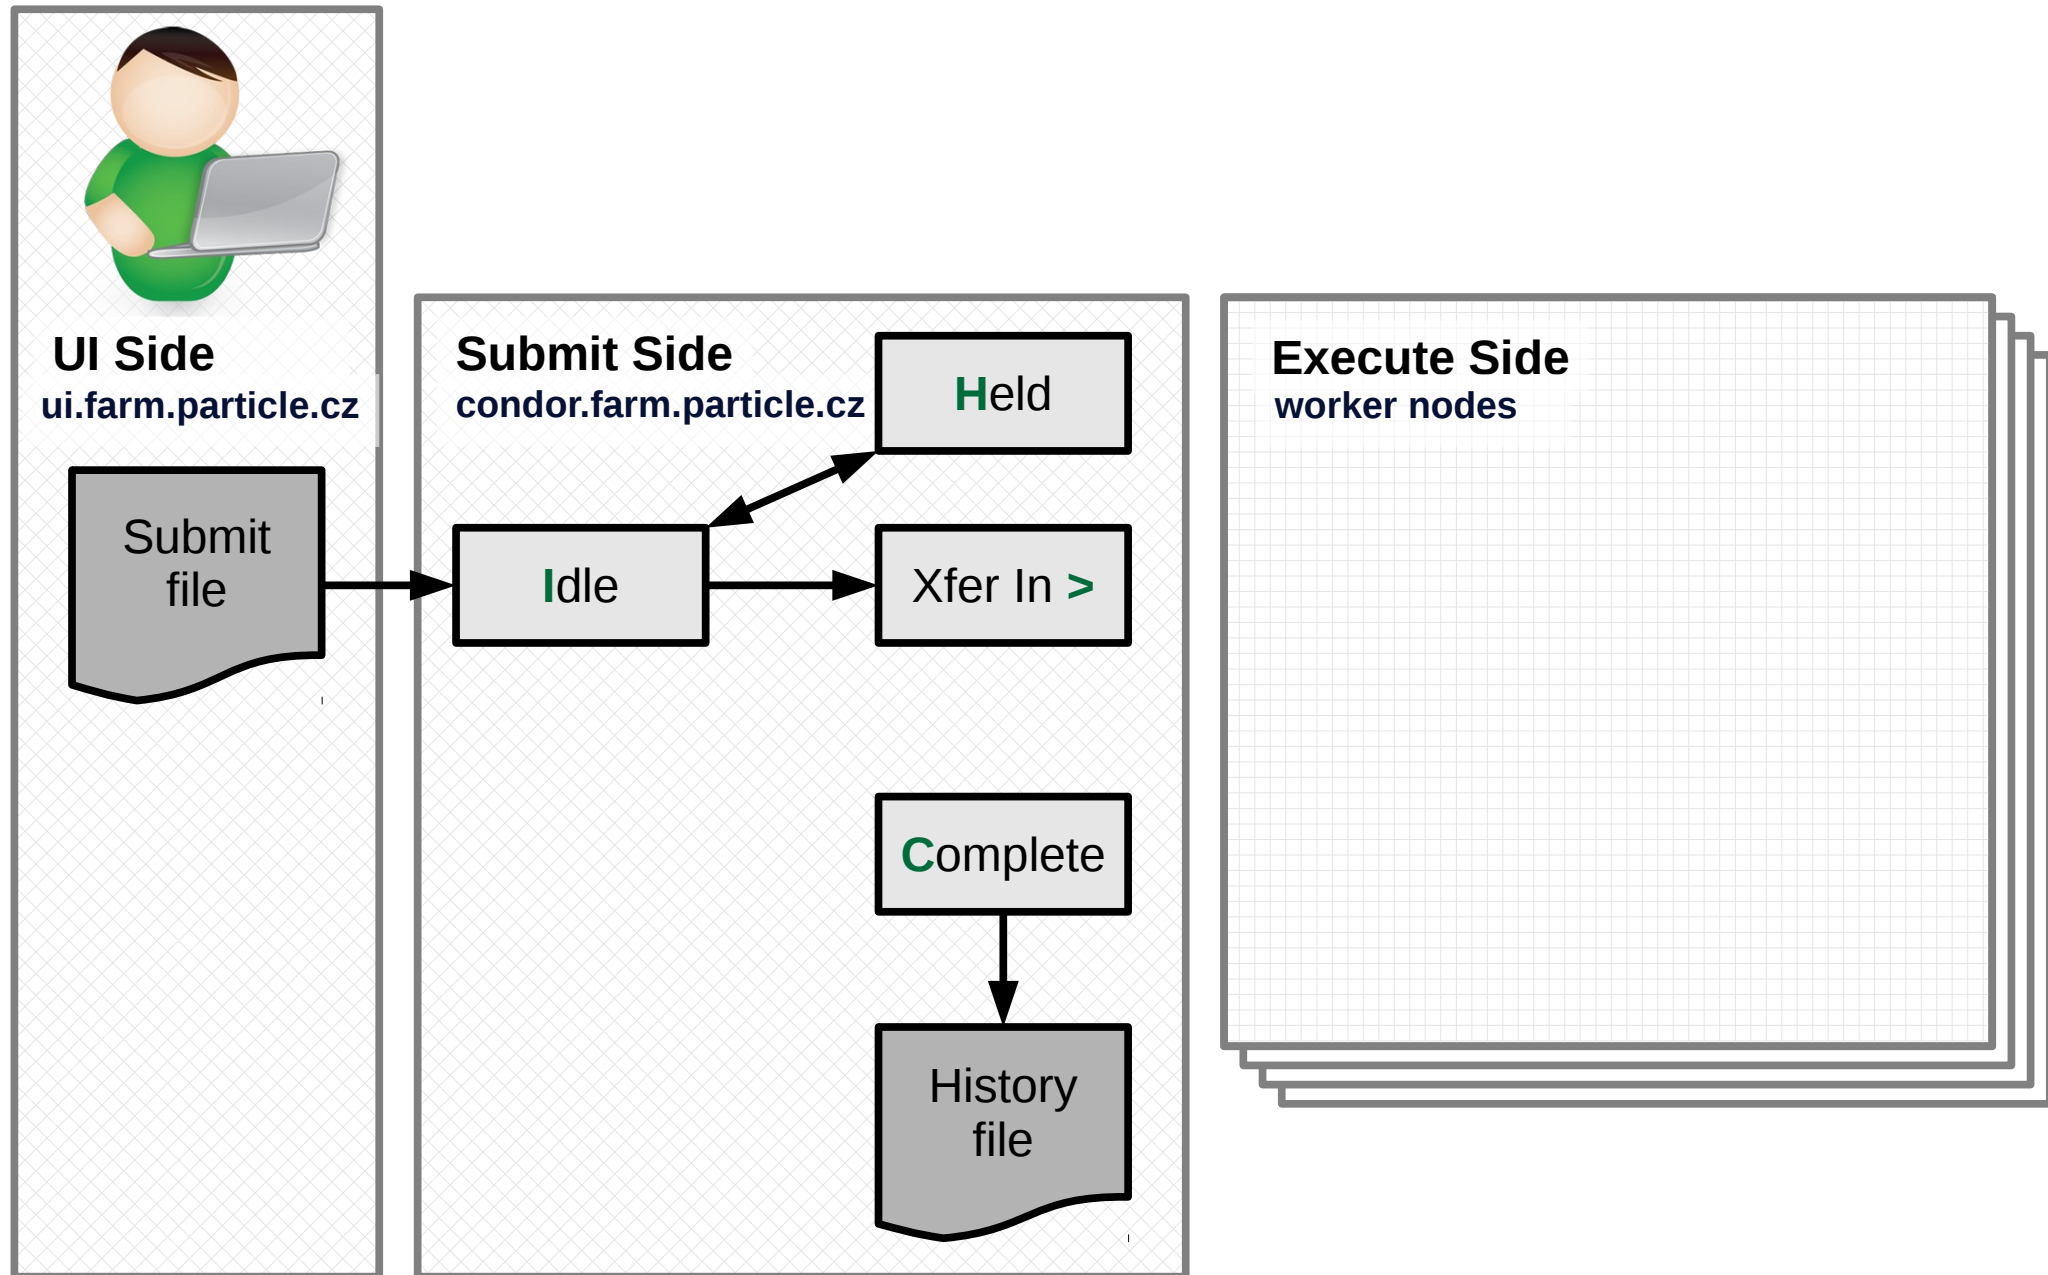
UI Side
`ui.farm.particle.cz`

Submit
file

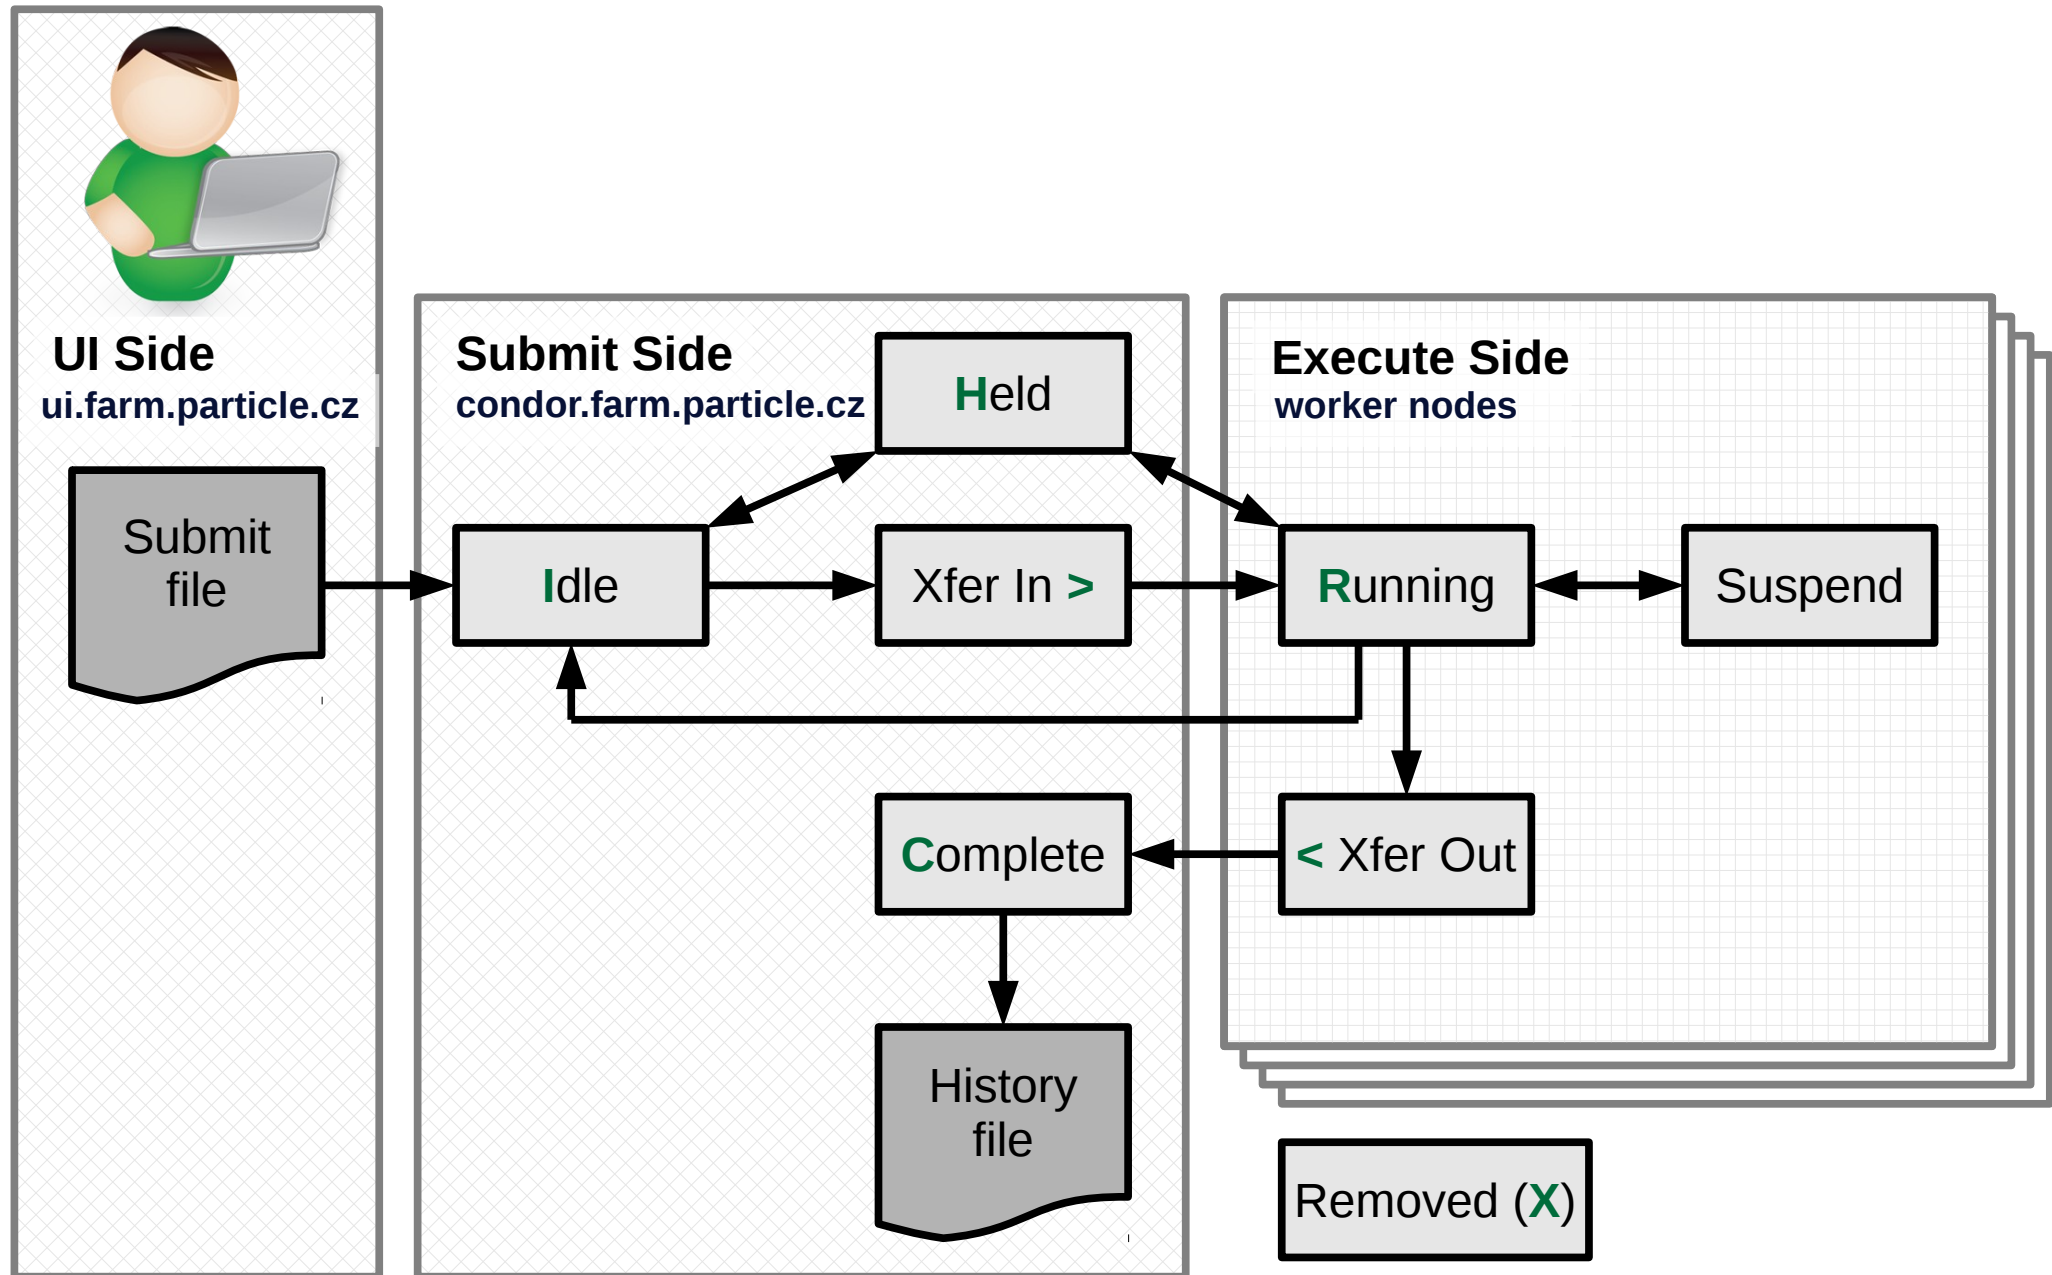
Submit Side
`condor.farm.particle.cz`

Execute Side
worker nodes

Job states




Job states



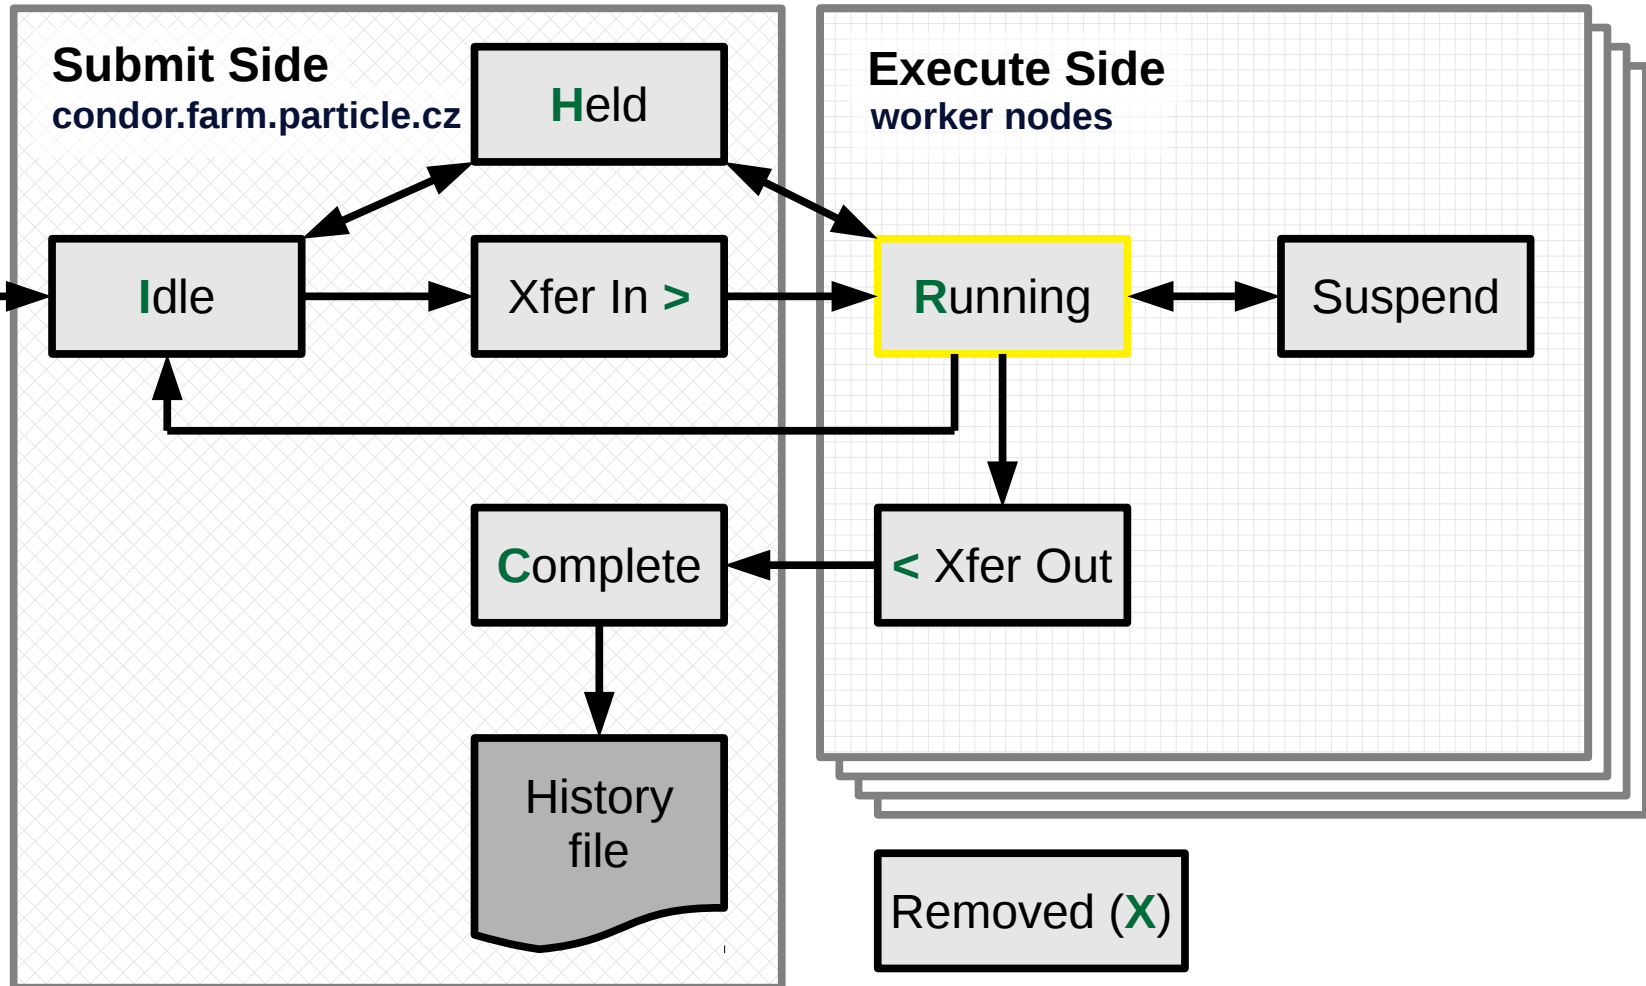
Job states



UI Side
ui.farm.particle.cz

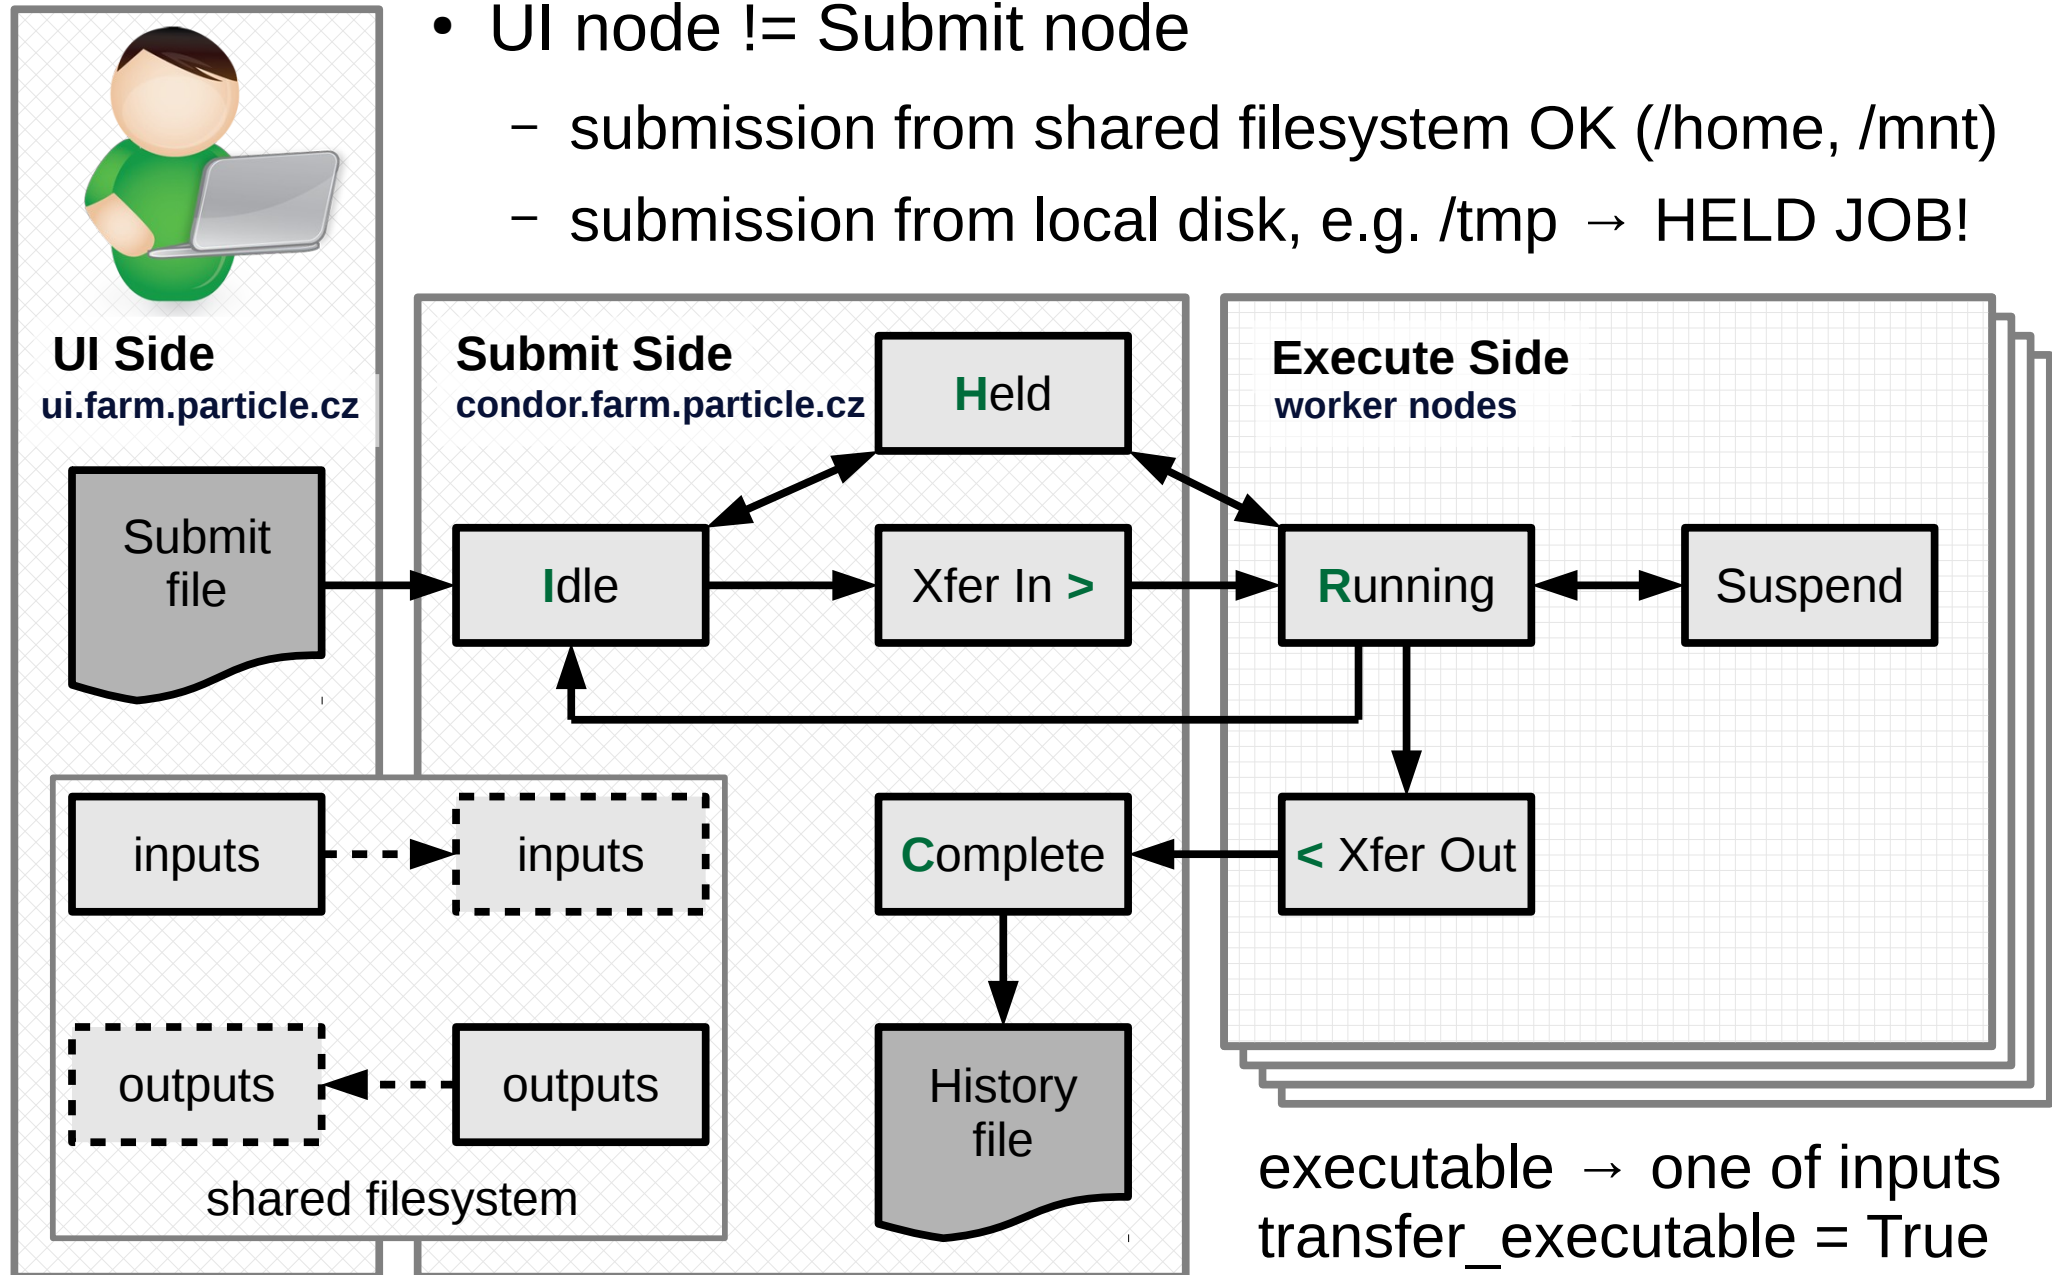


```
[user@ui ~]$ condor_q -nobatch
-- Schedd: condor.farm.particle.cz : <147.231.25.153...
ID OWNER SUBMITTED RUN_TIME ST PRI SIZE CMD
42.0 user 2/20 15:43 0+00:00:04 R 0 0.0 hello.sh
```



Job states

- UI node != Submit node
 - submission from shared filesystem OK (/home, /mnt)
 - submission from local disk, e.g. /tmp → HELD JOB!



Job states



- UI node != Submit node
 - submission from shared filesystem OK (/home, /mnt)
 - submission from local disk, e.g. /tmp → HELD JOB!

```
[user@ui /tmp]$ condor_submit job.submit
1 job(s) submitted to cluster 42.0
```

```
[user@ui /tmp]$ condor_q
-- Schedd: condor.farm.particle.cz : <147.231.25.153:9618> @ 02/20/18
OWNER BATCH_NAME          SUBMITTED DONE RUN  IDLE HOLD TOTAL JOB_IDS
User  CMD: /tmp/hello.sh    2/20 15:43  _   _   _   1   1   42.0
1 jobs; 0 completed, 0 removed, 0 idle, 0 running, 1 held, 0 suspended
```

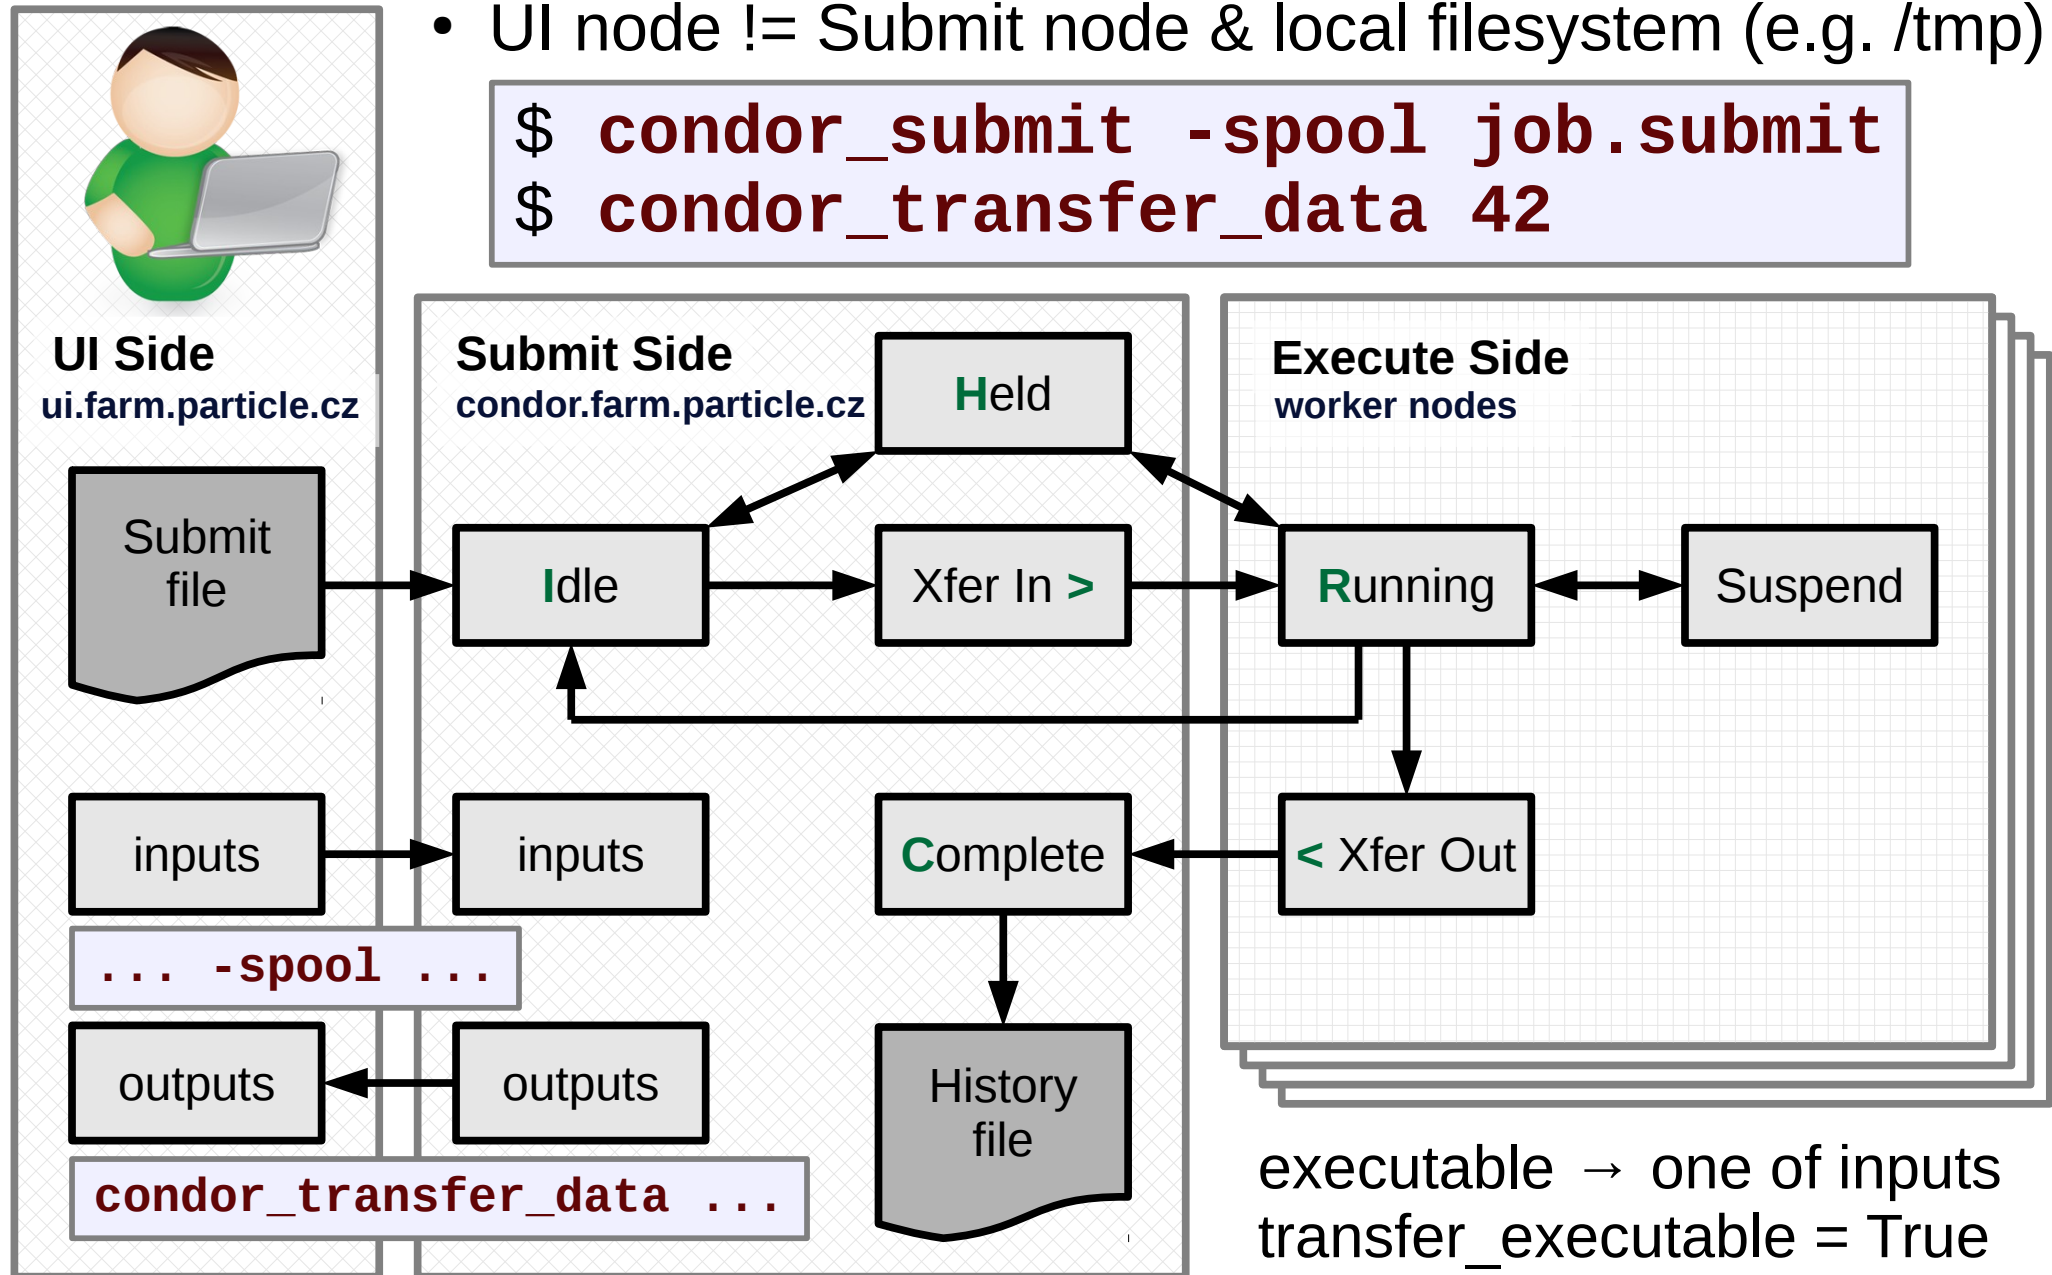
```
[user@ui /tmp]$ condor_q -hold
-- Schedd: condor.farm.particle.cz : <147.231.25.153:9618> @ 02/20/18
ID      OWNER      HELD_SINCE  HOLD_REASON
42.0    user       2/20 15:43    Error from slot1@aplex31.farm.particle.cz
```

```
[user@ui /tmp]$ condor_q 42.0 -af HoldReason
Error from slot1@aplex31.farm.particle.cz: SHADOW at 147.231.25.153 \
failed to send file(s) to <172.16.12.31:42337>: error reading from \
/tmp/hello.sh: (errno 2) No such file or directory; STARTER failed \
to receive file(s) from <147.231.25.153:9618>
```

Job states

- UI node != Submit node & local filesystem (e.g. /tmp)

```
$ condor_submit -spool job.submit
$ condor_transfer_data 42
```



Submit file

- Submission file parsed & translated by `condor_submit`
 - standard parameters
 - ``man condor_submit``
 - custom options prefixed by `+`
- Job classAd
 - key → value
 - full job description
 - include job state
 - queue, start, end, ... timestamp
 - resource usage
 - worker node info
 - transformations (add/del/mod)
 - accounting groups
 - JobFlavour → MaxRuntime

```
executable = job.sh
arguments = arg1 arg2 arg3

input = job.stdin
output = job.stdout
error = job.stderr
log = job.log

request_cpus = 1
request_memory = 2GB

# required @ FZU (CERN)
+MaxRuntime = 24*60*60
#+JobFlavour = "longlunch"

queue
```

Submit file

```
[user@ui ~] condor_q -long 42.0
AcctGroup = "group_fermilab.user"
AcctGroupUser = "vokac"
Args = "arg1 arg2 arg3"
ClusterId = 42
Cmd = "/home/user/job.sh"
EnteredCurrentStatus = 1518885268
Environment = ""
Err = "job.stderr"
GlobalJobId = "condor.farm.particle.cz#42.0#1518..."
In = "job.stdin"
Iwd = "/home/user"
JobPrio = 0
JobStatus = 1
JobUniverse = 5
MaxRuntime = 24 * 60 * 60
MyType = "Job"
Out = "job.stdout"
Owner = "user"
Procid = 0
QDate = 1518885268
RequestCpus = 1
RequestDisk = DiskUsage
RequestMemory = 2048
Requirements = ( TARGET.Arch == "X86_64" ) && ...
User = "user@farm.particle.cz"
UserLog = "/home/user/job.log"
```

```
executable = job.sh
arguments = arg1 arg2 arg3

input = job.stdin
output = job.stdout
error = job.stderr
log = job.log

request_cpus = 1
request_memory = 2GB

# required @ FZU (CERN)
+MaxRuntime = 24*60*60
#+JobFlavour = "longlunch"

queue
```

Condor Job classAd

Submit file

```
[user@ui ~] condor_q -long
...
[user@ui ~] condor_q -run
...
[user@ui ~] condor_q -hold
...
[user@ui ~] condor_q -cputime
...
[user@ui ~] condor_q -global
...
[user@ui ~] condor_q -all
...
[user@ui ~] condor_q -af ClusterId ProclD Owner
42 0 user
[user@ui ~] condor_q -af:r ClusterId Requirements
42 (TARGET.Arch == "X86_64") \
  && (TARGET.OpSys == "LINUX") \
  && (TARGET.Memory >= RequestMemory) \
  ...
[user@ui ~] condor_q -format "%d" ClusterId \
-format ".%d" ProclD \
-format "%s\n" Owner \
-format "%v" Requirements \
-format "%r" Requirements
42.0 user undefined (TARGET.ARCH == "X86_64")...
```

```
executable = job.sh
arguments = arg1 arg2 arg3

input = job.stdin
output = job.stdout
error = job.stderr
log = job.log

request_cpus = 1
request_memory = 2GB

# required @ FZU (CERN)
+MaxRuntime = 24*60*60
#+JobFlavour = "longlunch"

queue
```

Submit file

- executable
 - script
 - binary executable
- arguments
 - passed to executed process
- transfer_executable
 - by default copied to WN
- MaxRuntime
 - walltime
 - max 14 days

```
executable = job.sh
arguments = arg1 arg2 arg3
```

```
#input = job.stdin
output = job.stdout
error = job.stderr
log = job.log
```

```
+MaxRuntime = 15*60
```

```
queue
```

- executed command:

\$PWD/job.sh arg1 arg2

Submit file

- executable
 - script
 - binary executable
- arguments
 - passed to executed process
- transfer_executable
 - by default copied to WN
- MaxRuntime
 - walltime
 - max 14 days
- executed command:
/bin/sleep 60

```
executable = /bin/sleep
arguments = 60

#input = job.stdin
output = job.stdout
error = job.stderr
log = job.log

transfer_executable = False

+MaxRuntime = 15*60

queue
```

Submit file

- executable
 - script
 - binary executable
- arguments
 - passed to executed process
- transfer_executable
 - by default copied to WN
- MaxRuntime
 - walltime
 - max 14 days

- executed command:

/bin/sleep 900

```
executable = /bin/sleep
arguments = $$([MaxRuntime])

#input = job.i$(ClusterId)
output = job.o$(ClusterId)
error = job.e$(ClusterId)
log = job.l$(ClusterId)

transfer_executable = False

+MaxRuntime = 15*60

queue
```


Submit file

- job process stdin/stdout/stderr
 - by default discarded
 - user defined output names
- log file
 - condor info for one ClusterId
 - job(s) state and failure reason
 - resource usage for finished
 - condor_wait
 - excessive condor_q can overload batch server
 - especially long outputs
 - only monitor changes in job condor log file
 - only for simple “dependencies” (DAG much more robust)

```
executable = job.sh
arguments = arg1 arg2 arg3


```

Submit file

```
000 (42.000.000) 02/20 13:51:50 Job submitted from host: <147.231.25.153:9618?addr
...
001 (42.000.000) 02/20 13:52:09 Job executing on host: <172.16.6.35:9618?addrs=172
...
006 (42.000.000) 02/20 13:52:18 Image size of job updated: 750
1 - MemoryUsage of job (MB)
476 - ResidentSetSize of job (KB)
...
006 (42.000.000) 02/20 13:53:00 Image size of job updated: 832
1 - MemoryUsage of job (MB)
476 - ResidentSetSize of job (KB)
...
005 (42.000.000) 02/20 13:53:01 Job terminated.
(1) Normal termination (return value 0)
  Usr 0 00:00:50, Sys 0 00:00:00 - Run Remote Usage
  Usr 0 00:00:00, Sys 0 00:00:00 - Run Local Usage
  Usr 0 00:00:50, Sys 0 00:00:00 - Total Remote Usage
  Usr 0 00:00:00, Sys 0 00:00:00 - Total Local Usage
3960 - Run Bytes Sent By Job
565324 - Run Bytes Received By Job
3960 - Total Bytes Sent By Job
565324 - Total Bytes Received By Job
Partitionable Resources : Usage Request Allocated
  Cpus : 1 1
  Disk (KB) : 750 750 2880075
  Memory (MB) : 1 1 1
...
```

Submit file

- environment variables
 - no variables passed by default
 - including HOME variable!
 - e.g. ``cd`` rely on \$HOME
 - adds condor specific env. vars.
 - `_CONDOR_*`
 - `_CONDOR_SCRATCH_DIR`
 - job and machine classAd file
 - `_CONDOR_JOB_AD`
 - `_CONDOR_MACHINE_AD`
 - `_X509_USER_PROXY`
- Namespaces
 - `mount – /tmp, /var/tmp`
 - `pid`

```
executable = job.sh
arguments = arg1 arg2 arg3

getenv = True
#environment = "KEY1=VAL1"
#environment = "KEY2=VAL2"
#environment = \
#           "HOME=$ENV(HOME)"

+MaxRuntime = 15*60

queue
```

Submit file

- priority [-20, 20]
 - user job priority (not global)
- job state changes notifications
- accounting groups
 - per experiment / exp. group
 - group_alice.user
 - group_atlas.user
 - ...
 - selected automatically
 - use unix group info
 - can be overwritten
 - special groups
 - group_user
 - group_hiprio
 - no access by default

```
executable = job.sh
arguments = arg1 arg2 arg3

priority = 0

notification = Complete
notify_user = user@fzu.cz

accounting_group = \
group_exp.user

+MaxRuntime = 15*60

queue
```

Submit file

- stagein / stageout job data
 - use condor built-in transfers
 - Initiated from/to Schedd
condor.farm.particle.cz
 - spool & transfer from UI
 - executable stagein by default
 - condor transfers limit

```
executable = job.sh
arguments = arg1 arg2 arg3


```

Submit file

- Submit multiple jobs
 - N same jobs
 - Cluster.ProcId for output files
 - different job directories
 - for each input file

```
executable = job.sh
arguments = arg1 arg2 arg3

queue 10
```

Submit file

- Submit multiple jobs
 - N same jobs
 - Cluster.ProcId for output files
 - different job directories
 - for each input file

```
executable = job.sh
arguments = arg1 arg2 arg3

initialdir = job1
queue
initialdir = job2
queue

# or

initialdir = job$(ProcId)
queue 10
```

Submit file

- Submit multiple jobs
 - N same jobs
 - Cluster.ProcId for output files
 - different job directories
 - for each input file

```
executable = job.sh
arguments = $(infile)

#input = job.i$(ClusterId)
output = job.o$(ClusterId).$(ProcId)
error = job.e$(ClusterId).$(ProcId)
log = job.l$(ClusterId)

+MaxRuntime = 15*60

queue infile in (a b c)
queue infile from list.txt
queue infile *.dat
```


Submit file

- Requirements
 - constrain machines suitable for our jobs
 - standard options
 - request_*
 - user defined
 - expression matching for job / machine classad

```
executable = job.sh
arguments = arg1 arg2 arg3

request_cpus = 2
request_memory = 2GB
request_disk = 10GB
requirements = has_sse4_1 && HasVM

queue
```

```
[user@ui ~]$ condor_q -af:r Requirements
(TARGET.Arch == "X86_64") && (TARGET.OpSys == "LINUX") \
&& (TARGET.Disk >= RequestDisk) && (TARGET.Memory >= RequestMemory) \
&& ((TARGET.HasFileTransfer) || (TARGET.FileSystemDomain == MY.FileSystemDomain))
```

Submit file

```
[user@ui ~]$ condor_status -compact
Machine Platform Slots Cpus Gpus TotalGb FreCpu FreeGb
aplex01.farm.particle.cz x64/SL6 31 32 76.17 1 64.30
aplex02.farm.particle.cz x64/SL6 27 32 76.17 1 18.80
...
rubus23.farm.particle.cz x64/SL6 25 32 76.17 0 11.42
wntestcc7.farm.particle.cz x64/CentOS7 0 2 5.86 2 5.86
```

```

Machines Owner Claimed Unclaimed Matched Preempting Drain
x64/CentOS7 1 1 0 0 0 0 0
x64/SL6 4443 11 4389 1 0 0 42
Total 4444 12 4389 1 0 0 42
```

```
[user@ui ~]$ condor_status -startd -long
```


```
...
[user@ui ~]$ condor_status -startd \
  -constraints '(Machine == "db1.farm.particle.cz")||(Machine ==
  "db2.farm.particle.cz")'
```

```
...
[user@ui ~]$ condor_status -startd \
  -constraints 'stringListMember(Machine,
  "db1.farm.particle.cz,db2.farm.particle.cz")'
```

```
...
[user@ui ~]$ condor_status -startd \
  -constraints 'regexp("^db[0-9]+\.\farm\.particle\.cz$", Machine, "I")'
```

Advanced features

- DAG
- Python bindings
 - for `condor_q`, `condor_stat`, `condor_history`
 - `import htcondor, classad`
- job universe
 - default – vanilla
 - VM (KVM, Xen, VMWare)
 - Containers (docker, singularity)
 - not yet supported
 - CentOS7 WN required
 - MPI
 - commercial cloud provider



Advanced features

- DAG
- Python bindings
 - for condor_q, condor_stat, condor_history
 - import htcondor, classad
- job universe
 - default – vanilla
 - VM (KVM, Xen, VMWare)
 - Containers (docker, singularity)
 - not yet supported
 - CentOS7 WN required
 - MPI
 - commercial cloud provider

```
universe = docker
executable = job.sh
arguments = arg1 arg2 arg3

output = job.o$(ClusterId)
error = job.e$(ClusterId)
log = job.l$(ClusterId)

#+OS_IMAGE = SLC5
+OS_IMAGE = SLC6
#+OS_IMAGE = CC7
#+OS_IMAGE = Ubuntu

+MaxRuntime = 15*60

queue
```

Advanced features

- DAG
- Python bindings
 - for condor_q, condor_stat, condor_history
 - import htcondor, classad
- job universe
 - default – vanilla
 - VM (KVM, Xen, VMWare)
 - Containers (docker, singularity)
 - not yet supported
 - CentOS7 WN required
 - MPI
 - commercial cloud provider

```
universe = vanilla
executable = job.sh
arguments = arg1 arg2 arg3
```

```
output = job.o$(ClusterId)
error = job.e$(ClusterId)
log = job.l$(ClusterId)
```

```
Requirements = \
                HasSingularity
```

```
+SingularityBind = /mnt,...
+SingularityImage = \
    /cvmfs/atlas.cern.ch/repo\
    /containers/fs/singularity
```

```
+MaxRuntime = 15*60
```

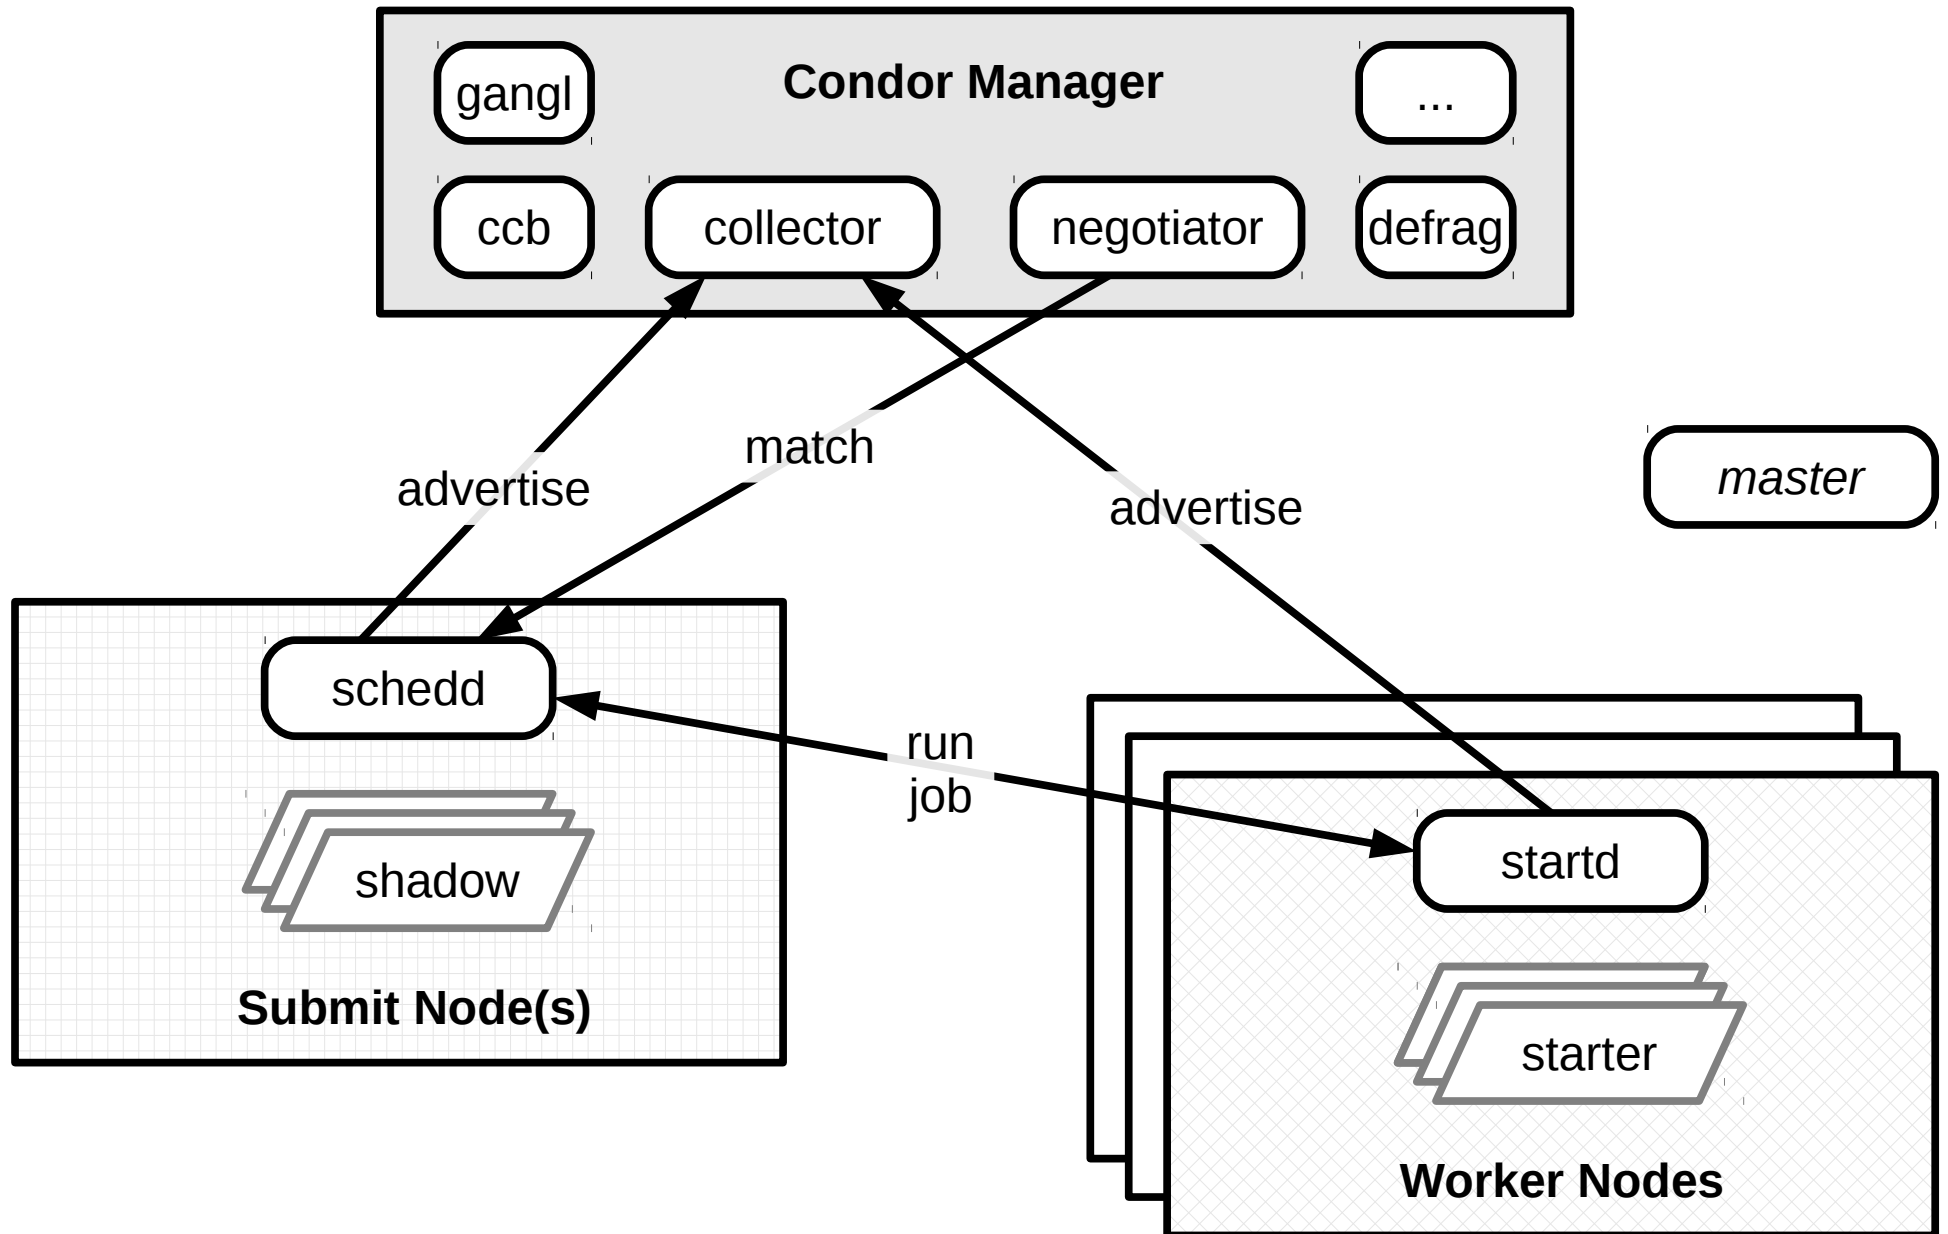
```
queue
```

References & documentation

- FZU Farm HTCondor twiki
 - <https://www.farm.particle.cz/twiki/bin/view/VS/VsGoliasCondorUsersEn>
- CERN HTCondor Quick Start Guide for Ixplus
 - <http://batchdocs.web.cern.ch/batchdocs/local/quick.html>
- Past HTCondor Week Meetings → HTCondor User Tutorial
 - https://indico.cern.ch/event/611296/contributions/2604376/attachments/1471164/2276521/TannenbaumT_UserTutorial.pdf
- User documentation & examples
 - <http://vivaldi.ll.iac.es/sieinvens/siepedia/pmwiki.php?n=HOWTOs.CondorUsefulCommands>
- Official HTCondor documentation (not really end user guide)
 - <https://research.cs.wisc.edu/htcondor/manual/current/>
- FZU Farm monitoring
 - <http://monitor.farm.particle.cz>

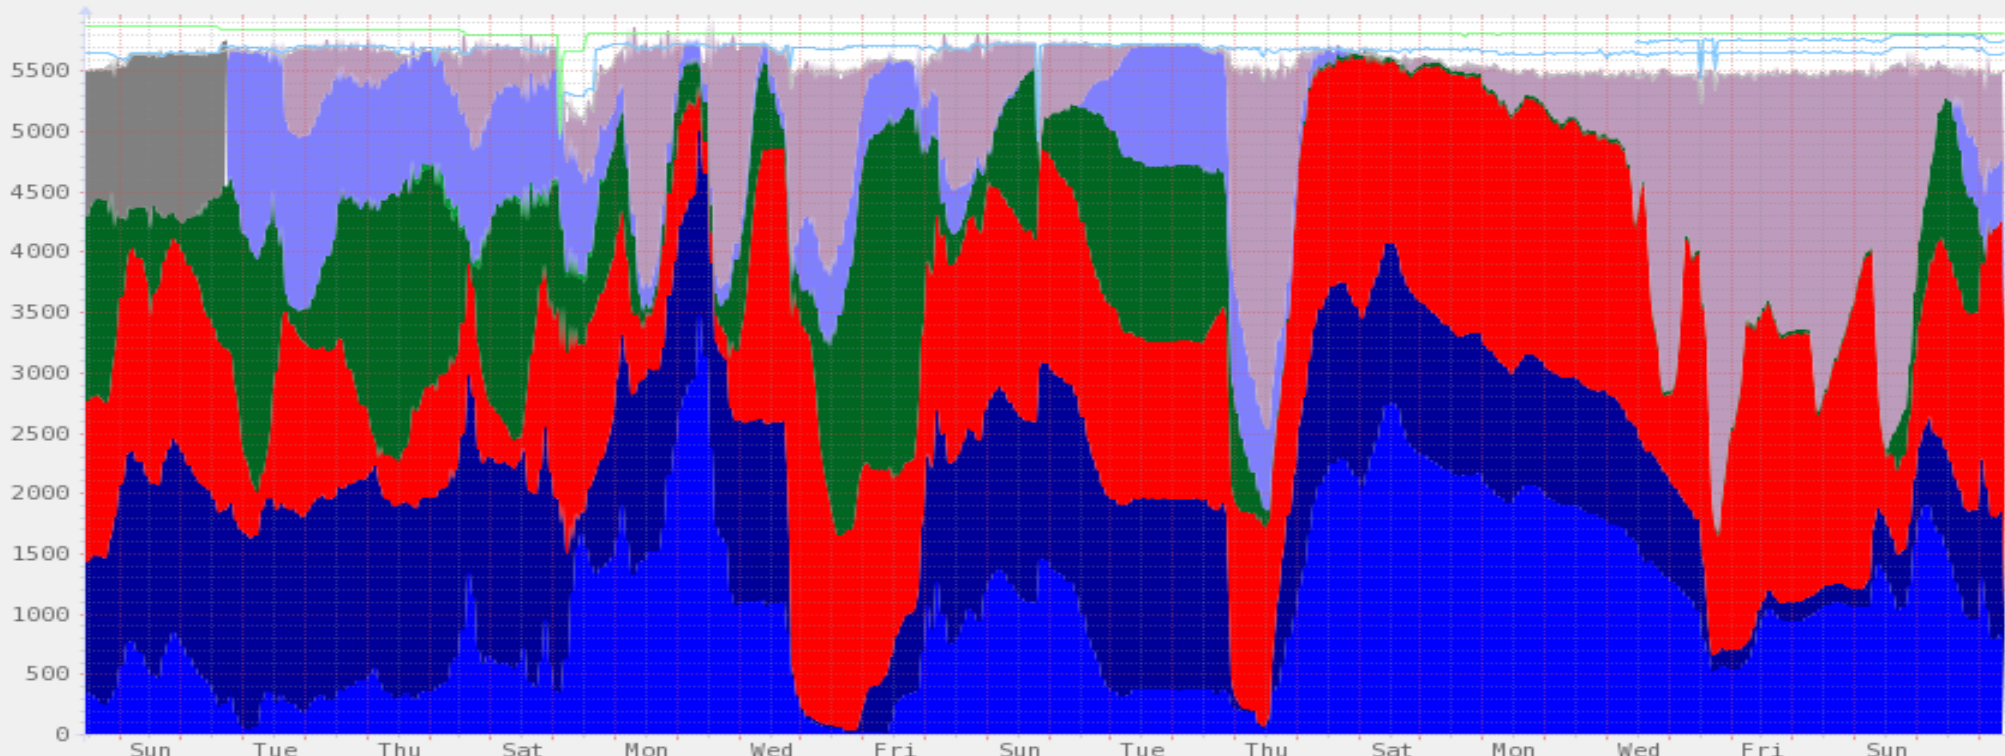
BACKUP

HTCondor Pool



FZU HTCondor batch usage

Used condor cores per project - from Sat Jan 20 10:00:00 2018 to Tue Feb 20 10:00:00 2018



	Cur:	Min:	Avg:	Max:
ATLAS VO	827	0	1029.31	3580
ATLAS VO (multicore)	1015	0	1133.91	1680
ALICE VO	2404	0	1526.66	3209
Auger VO	0	0	0.00	0
CTA VO	0	0	0.00	0
NOvA VO	0	0	704.21	2945
Dune VO	0	0	4.85	136
ops	0	0	0.01	3
Unknown VO	0	0	0.00	0
ATLAS User	498	0	399.32	1743
ALICE User	0	0	0.00	0
Auger User	745	0	789.45	3850
Fermilab User	0	0	0.00	1
Priority User	0	0	0.00	0
User	0	0	0.00	0
Unknown	0	0	0.00	0
Local (deprecated)	0	0	91.69	1526
Total jobs	4353	0	4390.07	5434
Used cores	5499	0	5591.51	5741
Drained cores	150	0	85.34	355
Offline cores	104	104	104.00	104
Total cores	5818	1032	5820.58	5874

Munin 2.0.33
Last update: Tue Feb 20 10:00:59 2018

List of FZU worker nodes

Golias - torque

iberis

iberis01

magic

magic01 magic02 magic03 magic04 magic05 magic06 magic07 magic08 magic09 magic10
magic11 magic12 magic13 magic14 magic15 magic16 magic17 magic18 magic19 magic20
magic21 magic22 magic23 magic24 magic25 magic26 magic27 magic28 magic29 magic30
magic31 magic32 magic33 magic34 magic35 magic36

malva

malva01 malva02 malva03 malva04 malva05 malva06 malva07 malva08 malva09 malva10
malva11 malva12

saltix

saltix01 saltix02 saltix03 saltix04 saltix05 saltix06 saltix07 saltix08 saltix09 saltix10
saltix11 saltix12 saltix13 saltix14 saltix15 saltix16 saltix17 saltix18

Golias - condor

aplex

aplex01 aplex02 aplex03 aplex04 aplex05 aplex06 aplex07 aplex08 aplex09 aplex10
aplex11 aplex12 aplex13 aplex14 aplex15 aplex16 aplex17 aplex18 aplex19 aplex20
aplex21 aplex22 aplex23 aplex24 aplex25 aplex26 aplex27 aplex28 aplex29 aplex30
aplex31 aplex32

db

db2

ib

ib01 ib02 ib03 ib04 ib05 ib06 ib07 ib08 ib09 ib10
ib11 ib12 ib13 ib14 ib15 ib16 ib17 ib18 ib19 ib20
ib21 ib22 ib23 ib24 ib25 ib26

ibis

ibis01 ibis02 ibis03 ibis04 ibis05 ibis06 ibis07 ibis08 ibis09 ibis10
ibis11 ibis12 ibis13 ibis14 ibis15 ibis16 ibis17 ibis18 ibis19 ibis20
ibis21 ibis22 ibis23 ibis24 ibis25 ibis26 ibis27 ibis28 ibis29 ibis30
ibis31 ibis32 ibis33 ibis34 ibis35 ibis36 ibis37 ibis38 ibis39 ibis40
ibis41 ibis42 ibis43 ibis44 ibis45 ibis46 ibis47 ibis48 ibis49 ibis50
ibis51 ibis52 ibis53 ibis54 ibis55 ibis56 ibis57 ibis58 ibis59 ibis60
ibis61 ibis62 ibis63 ibis64 ibis65

mff

mff39 mff40 mff41 mff42 mff71 mff72 mff73 mff74

mikan

mikan01 mikan02 mikan03 mikan04 mikan05 mikan06 mikan07 mikan08 mikan09 mikan10
mikan11 mikan12 mikan13 mikan14 mikan15 mikan16 mikan17 mikan18 mikan19 mikan20
mikan21 mikan22 mikan23 mikan24 mikan25 mikan26 mikan27 mikan28 mikan29 mikan30
mikan31 mikan32 mikan33 mikan34 mikan35 mikan36 mikan37 mikan38 mikan39 mikan40
mikan41 mikan42 mikan43 mikan44 mikan45 mikan46 mikan47 mikan48 mikan49 mikan50
mikan51

rubus

rubus01 rubus02 rubus03 rubus04 rubus05 rubus06 rubus07 rubus08 rubus09 rubus10
rubus11 rubus12 rubus13 rubus14 rubus15 rubus16 rubus17 rubus18 rubus19 rubus20
rubus21 rubus22 rubus23