




# HTCondor and Workflows: Advanced Tutorial

## HTCondor Week 2015

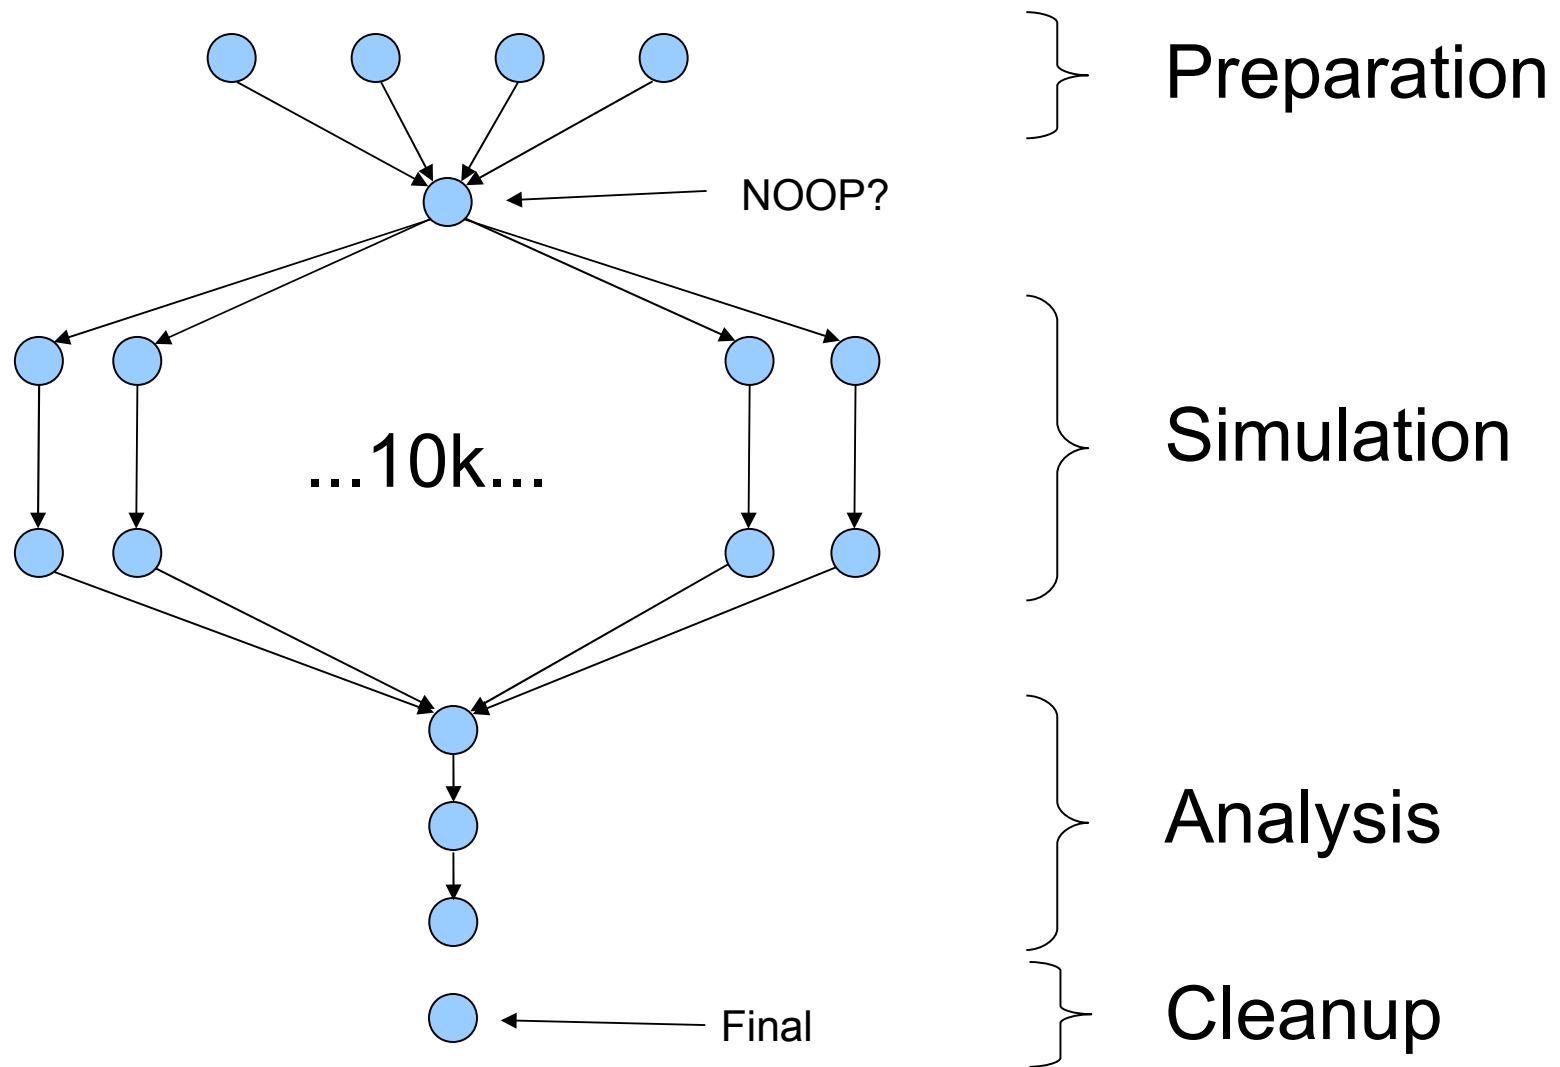
*Kent Wenger*

# Workflows in HTCondor

- This talk: techniques & features
- Please ask questions!




# Example workflow



# How big?

- › We have users running 500k-job workflows in production (user hit 2.2 GB DAG file bug)
- › Depends on resources on submit machine (e.g., memory)
- › “Tricks” can decrease resource requirements (talk to me)

Pretty big!






# Defining a DAG to DAGMan

A **DAG input file** defines a DAG:

```
# file name: diamond.dag
Job A a.submit
Job B b.submit
Job C c.submit
Job D d.submit
Parent A Child B C
Parent B C Child D
```



# Changes since last year



- PRE/POST script retry after delay (DEFER option)
- DAGMan handles submit file “foreach” syntax (TJ will talk about this later)
- Configuration:
  - Maxpre, maxpost default to 20 (was 0)
  - Maxidle defaults to 1000 (was 0)
  - Fixed DAGMan entries in param table

# Changes (cont)



- Node status file:
  - Format is now ClassAds
  - More info (retry number, procs queued and held for each node)
  - Fixed bug: final DAG status not always recorded correctly
  - **ALWAYS-UPDATE** option
  - Now works on Windows

# Changes (cont)



- Log files:

- Added @ (OWNER) and @ (NODE\_NAME) macros in **DAGMAN\_DEFAULT\_NODE\_LOG** (allows use in global config)
- DAGMan ignores per-job logs (avoids FD limit problems)
- **DAGMAN\_SUPPRESS\_JOB\_LOGS** (prevent individual logs from being written)



# Changes (cont)



- `dagman.out` file:
  - Node job hold reason in `dagman.out`
  - DAG\_Status in `dagman.out`
- `-DoRecovery` command-line option for `condor_submit_dag`
- Stricter checking of SPLICE syntax
- No (unused) command socket
- Stork no longer supported

# Changes (cont)



- Other bug fixes:
  - `condor_submit_dag` core dump on non-existent DAG file fixed
  - `condor_dagman` segfault on some DAG file syntax errors fixed
  - Duplicate node job removal by DAGMan and schedd eliminated

# Organization of files and directories



# Files and directories

- By default, all paths in a DAG input file *and the associated submit files* are relative to the current working directory when `condor_submit_dag` is run.
- Modified by `DIR` directive in `JOB` command
- Also by `-usedagdir` on `condor_submit_dag` command line

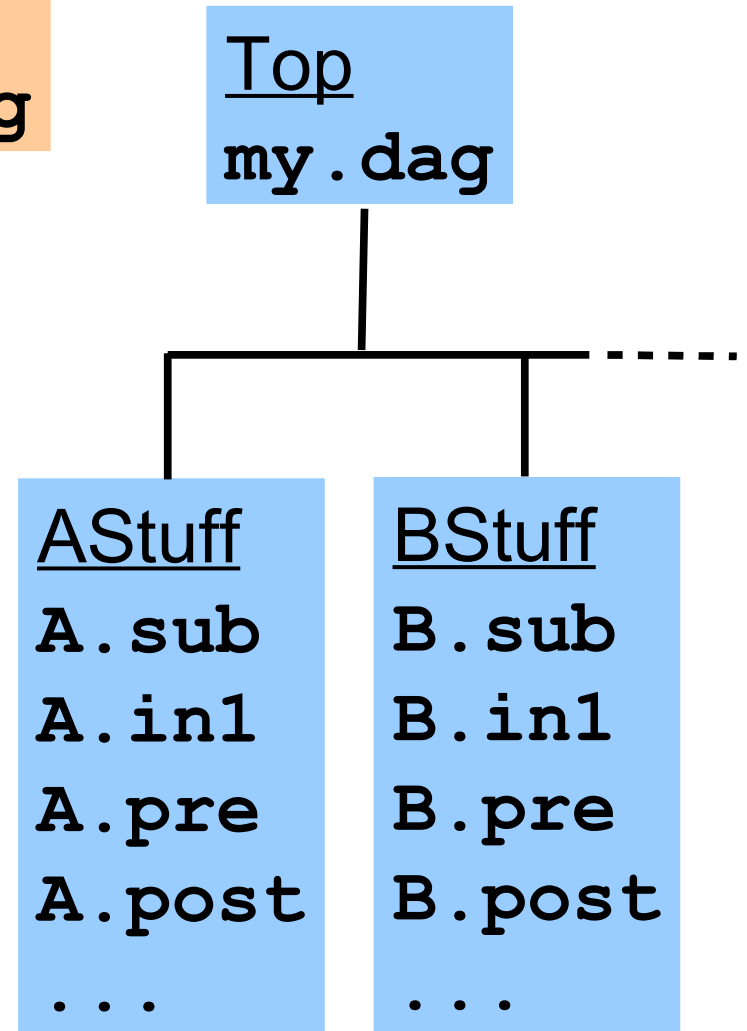


# Nodes in subdirectories

```
# in Top:  
condor_submit_dag my.dag
```

```
# my.dag:  
Job A A.sub Dir AStuff  
Script Pre A A.pre  
Job B B.sub Dir BStuff  
...
```


```
# A.sub:  
...  
input = A.in1  
...
```



# DAGs in subdirectories


```
# in Top:  
condor_submit_dag  
-usedagdir  
Dag1/my.dag  
Dag2/my.dag ...
```

```
# Dag1/my.dag:  
Job A A.sub  
Job B B.sub  
...
```



# Multiple independent jobs

- Why use DAGMan?:
  - Throttling
  - Retry of failed jobs
  - Rescue DAG
  - PRE/POST scripts
  - Submit file re-use (VARS)
- DAG file has JOB commands but not PARENT...CHILD





# Rescue DAGs





# Rescue DAGs (cont)



# Rescue DAGs (cont)

- › Save the state of a partially-completed DAG
- › Created when a **node fails** (after maximal progress) or the **condor\_dagman job is removed** with **condor\_rm** or when **DAG is halted** and all queued node jobs finish
  - DAGMan makes as much progress as possible in the face of failed nodes
- › DAGMan immediately exits after writing a rescue DAG file
- › Automatically run when you re-run the original DAG (**unless -force is passed to condor\_submit\_dag**)

# Rescue DAGs (cont)

- › The Rescue DAG file, by default, is only a partial DAG file (complete rescue DAGs will go away).
- › A partial Rescue DAG file contains only information about which nodes are done, and the number of retries remaining for nodes with retries.
- › Does not contain information such as the actual DAG structure and the specification of the submit file for each node job.
- › Partial Rescue DAGs are automatically parsed in combination with the original DAG file, which contains information such as the DAG structure.

# Rescue DAGs (cont)

- › If you change something in the original DAG file, such as changing the submit file for a node job, that change will take effect when running a partial rescue DAG.

# Rescue DAG naming

- › *DagFile.rescue001*, *DagFile.rescue002*, etc.
- › Up to 100 by default (last is overwritten once you hit the limit)
- › Newest (highest number) is run automatically when you re-submit the original *DagFile*
- › `condor_submit_dag -dorescuefrom number` to run specific rescue DAG
  - Newer rescue DAGs are renamed

# DAGMan configuration




Photo courtesy American Locomotive Co.

# DAGMan configuration (cont)

- › A few dozen DAGMan-specific configuration macros (see the manual...)
- › From lowest to highest precedence
  - HTCondor configuration files
  - User's environment variables:
    - `_CONDOR_macroname`
  - DAG-specific configuration file (preferable)
  - `condor_submit_dag` command line (recorded in `dagman.out` file)



# Per-DAG configuration

- › In DAG input file:

**CONFIG** *ConfigFileName*

or command line:

**condor\_submit\_dag -config**  
*ConfigFileName ...*

- › Generally prefer **CONFIG** in DAG file over **condor\_submit\_dag -config** or individual arguments
- › Specifying more than one configuration file is an error.

# Per-DAG configuration (cont)

- › Configuration entries not related to DAGMan are ignored
- › Syntax like any other HTCondor config file

```
# file name: bar.dag  
CONFIG bar.config
```

```
# file name: bar.config  
DAGMAN_ALWAYS_RUN_POST = False  
DAGMAN_MAX_SUBMIT_ATTEMPTS = 2
```

# Configuration: workflow log file

- DAGMan now uses a single log file for all node jobs (must be unique per DAG instance!)
- Put workflow log file on local disk
- **DAGMAN\_DEFAULT\_NODE\_LOG**
- Options for global config:
  - **@ (DAG\_DIR) /@ (DAG\_FILE) .nodes.log**
  - **/localdisk/@ (DAG\_FILE) .nodes.log**
  - **/localdisk/@ (OWNER) .nodes.log**

# Script deferral



- Re-try failed script (not entire node) *after a specified deferral time*
- Deferred scripts don't count against maxpre/maxpost
- In the DAG file:
  - **SCRIPT [DEFER *status time*] PRE| POST *node script\_path args...***
  - If script exits with *status*, re-try after *time* (or more) seconds

# Pre skip



# DAG node with scripts:

## PRE\_SKIP

- › Allows PRE script to immediately declare node successful (job and POST script are not run)
- › In the DAG input file:

```
JOB A A.cmd  
SCRIPT PRE A A.pre  
PRE_SKIP A non-zero_integer
```
- › If the PRE script of A exits with the indicated value, the node succeeds immediately, and the node job and POST script are skipped.
- › If the PRE script fails with a different value, the node job is skipped, and the POST script runs (as if PRE\_SKIP were not specified).

# DAG node with scripts: PRE\_SKIP (cont)

- › When the POST script runs, the **\$PRE\_SCRIPT\_RETURN** variable contains the return value from the PRE script. (See manual for specific cases)



# No-op nodes



# No-op nodes (cont)

- › Appending the keyword **NOOP** causes a job to not be run, without affecting the DAG structure.
- › The PRE and POST scripts of NOOP nodes will be run. If this is not desired, comment them out.
- › Can be used to test DAG structure

# No-op nodes (ex)


- › Here is an example:

```
# file name: diamond.dag
Job A a.submit NOOP
Job B b.submit NOOP
Job C c.submit NOOP
Job D d.submit NOOP
Parent A Child B C
Parent B C Child D
```

- › Submitting this to DAGMan will cause DAGMan to exercise the DAG, without actually running node jobs.

# No-op nodes (ex 2)

Simplify dag structure



# Node retries



- › For possibly transient errors
- › Before a node is marked as failed.

- Retry N times. In the DAG file:

**Retry C 4**


(to retry node C four times before calling the node failed)

- Retry N times, unless a node returns specific exit code. In the DAG file:


**Retry C 4 UNLESS-EXIT 2**

# Node retries, continued

- › Node is retried as a whole



# Node variables





# Node variables (cont)

- › To re-use submit files for multiple nodes
- › In DAG input file:  
***VARs JobName varname="value"***  
***[ varname="value" . . . ]***
- › In submit description file:  
***\$(varname)***
- › ***varname*** can only contain alphanumeric characters and underscore
- › ***varname*** cannot begin with “queue”
- › ***varname*** is not case-sensitive
- › ***varname*** beginning with “+” defines classad attribute (e.g., ***+State = "Wisconsin"***)

# Node variables (cont)

- › *Value* cannot contain single quotes; double quotes must be escaped
- › The variable **\$ (JOB)** contains the DAG node name
- › **\$ (RETRY)** contains retry count
- › Any number of VARS values per node
- › DAGMan warns if a VAR name is defined more than once for a node

# Node variables (ex)

```
# foo.dag
```

```
Job B10 B.sub
```

```
Vars B10 infile="B_in.10"
```

```
Vars B10 +myattr="4321"
```

```
# B.sub
```


```
input = $(infile)
```

```
arguments = $$([myattr])
```

# Nested DAGs



# Nested DAGs (cont)



# Nested DAGs (cont)

- › Multiple DAG files in a single workflow (runs the sub-DAG as a job within the top-level DAG)
- › In the DAG input file:  
**SUBDAG EXTERNAL *JobName DagFileName***
- › Any number of levels
- › Sub-DAG nodes are like any other (can have PRE/POST scripts, retries, DIR, etc.)
- › Each sub-DAG has its own DAGMan
  - Separate throttles for each sub-DAG
  - Separate rescue DAGs (run automatically)

# Why nested DAGs?

- › DAG re-use
- › Scalability
- › Re-try multiple nodes as a unit
- › Short-circuit parts of the workflow (ABORT-DAG-ON in sub-DAG)
- › Dynamic workflow generation (sub-DAGs can be created “on the fly”)



# Splices



# Splices (cont)

- › Multiple DAG files in a single workflow (directly includes splice DAG's nodes within the top-level DAG)
- › In the DAG input file:  
***SPLICE JobName DagFileName***
- › Splices can be nested (and combined with sub-DAGs)

# Why splices?

- DAG re-use
- Advantages of splices over sub-DAGs:
  - Reduced overhead (single DAGMan instance)
  - Simplicity (e.g., single rescue DAG)
  - Throttles apply across entire workflow
- Limitations of splices:
  - Splices cannot have PRE and POST scripts (for now)
  - Splices cannot have retries
  - Splice DAGs must exist at submit time

# Throttling



# Throttling (cont)

- › Limit load on submit machine and pool
  - **Maxjobs** limits jobs in queue
  - **Maxidle** submit jobs until idle limit is hit
    - Can get more idle jobs if jobs are evicted
  - **Maxpre** limits PRE scripts
  - **Maxpost** limits POST scripts
- › All limits are *per DAGMan*, not global for the pool or submit machine (sub-DAGs count separately)
- › Limits can be specified as arguments to **condor\_submit\_dag** or in configuration

# Throttling (cont)

- › Example with per-DAG config file

```
# file name: foo.dag
```

```
CONFIG foo.config
```

```
# file name: foo.config
```

```
DAGMAN_MAX_JOBS_SUBMITTED = 100
```

```
DAGMAN_MAX_JOBS_IDLE = 5
```

```
DAGMAN_MAX_PRE_SCRIPTS = 3
```

```
DAGMAN_MAX_POST_SCRIPTS = 15
```




# Node categories





# Node categories (cont)



# Node category throttles

- › Useful with different types of jobs that cause different loads
- › In the DAG input file:  
**CATEGORY *JobName* *CategoryName***  
**MAXJOBS *CategoryName* *MaxJobsValue***
- › Applies the ***MaxJobsValue*** setting to only jobs assigned to the given category
- › Global throttles still apply

# Cross-splice node categories

- › Prefix category name with “+”

**MaxJobs +init 2**

**Category A +init**

- › See the Splice section in the manual for details

# Node priorities



# Node priorities (cont)

- › In the DAG input file:  
**PRIORITY *JobName* *PriorityValue***
- › Determines order of submission of ready nodes
- › DAG node priorities are copied to job priorities (including sub-DAGs)
- › Does *not* violate or change DAG semantics
- › Higher numerical value equals “better” priority

# Node priorities (cont)

- › Better priority nodes are not *guaranteed* to run first!
- › Effective node prio =  $\max(\text{explicit node prio, parents' effective prios, DAG prio})$
- › For sub-DAGs, pretend that the sub-DAG is spliced in
- › Overrides priority in node job submit file
- › *Not* relative to other users

# DAG abort



- In DAG input file:

***ABORT-DAG-ON JobName AbortExitValue  
[RETURN DagReturnValue]***

- If node value is *AbortExitValue*, the entire DAG is aborted *immediately*, implying that queued node jobs are removed, and a rescue DAG is created.
- Can be used for conditionally skipping nodes (especially with sub-DAGs)



# FINAL nodes



# FINAL nodes (cont)

- › FINAL node *always* runs at end of DAG (even on failure)
- › Use **FINAL** in place of **JOB** in DAG file:
  - **FINAL NodeName SubmitFile**
- › At most one FINAL node per DAG
- › FINAL nodes cannot have parents or children (but can have PRE/POST scripts)

# FINAL nodes (cont)

- › Success or failure of the FINAL node determines the success of the entire DAG
- › PRE and POST scripts of FINAL (and other) nodes can use **\$DAG\_STATUS** and **\$FAILED\_COUNT** to determine the state of the workflow
- › **\$(DAG\_STATUS)** and **\$(FAILED\_COUNT)** available as VARS

# Advanced workflow monitoring



# Status in DAGMan's ClassAd

```
> condor_q -1 59 | grep DAG_  
DAG_Status = 0  
DAG_InRecovery = 0  
DAG_NodesUnready = 1  
DAG_NodesReady = 4  
DAG_NodesPrerun = 2  
DAG_NodesQueued = 1  
DAG_NodesPostrun = 1  
DAG_NodesDone = 3  
DAG_NodesFailed = 0  
DAG_NodesTotal = 12
```

› Sub-DAGs count as one node in parent



# Node status file

- › Shows a snapshot of workflow state
  - Overwritten as the workflow runs
  - Updated atomically
- › In the DAG input file:  
**NODE\_STATUS\_FILE**     *statusFileName*  
*[minimumUpdateTime]*   **[ALWAYS-UPDATE]**
- › Not enabled by default
- › Separate files for sub-DAGs
- › As of 8.1.6, *in ClassAd format* (a set of ClassAds)

# Node status file contents

```
[  
  Type = "DagStatus";  
  DagFiles = {  
    "job_dagman_node_status.dag"  
  };  
  Timestamp = 1397683160; /* "Wed Apr 16 16:19:20  
    2014" */  
  DagStatus = 3; /* "STATUS_SUBMITTED ()" */  
  NodesTotal = 12;  
  NodesDone = 0;  
  NodesPre = 0;  
  NodesQueued = 1;  
  ...
```



# Node status file contents (cont)

```
[  
  Type = "NodeStatus";  
  Node = "C";  
  NodeStatus = 6; /* "STATUS_ERROR" */  
  StatusDetails = "Job proc (1980.0.0)  
failed with status 5";  
  RetryCount = 2;  
  JobProcsQueued = 0;  
  JobProcsHeld = 0;  
]
```

# Jobstate.log file


- › Shows workflow history
- › Meant to be machine-readable (for Pegasus)
- › Basically a subset of the `dagman.out` file
- › In the DAG input file:  
**`JOBSTATE_LOG`** *JobstateLogFileName*
- › Not enabled by default
- › Separate files for sub-DAGs

# Jobstate.log contents

```
1302884424 INTERNAL *** DAGMAN_STARTED 48.0
***
1302884436 NodeA PRE_SCRIPT_STARTED - local
- 1
1302884436 NodeA PRE_SCRIPT_SUCCESS - local
- 1
1302884438 NodeA SUBMIT 49.0 local - 1
1302884438 NodeA SUBMIT 49.1 local - 1
1302884438 NodeA EXECUTE 49.0 local - 1
1302884438 NodeA EXECUTE 49.1 local - 1
...
```

# DAGMan metrics

- *Anonymous* workflow metrics (for Pegasus)
- Metrics file (JSON format) generated at end of run (***dagfile.metrics***)
- Reported by default (can be disabled)
- **Dagman.out** tells whether metrics were reported



# DAGMan metrics example

```
{  
  "client": "condor_dagman",  
  "version": "8.3.5",  
  ...  
  "start_time": 1396448008.138,  
  "end_time": 1396448047.596,  
  "duration": 39.458,  
  "exitcode": 0,  
  ...  
  "total_jobs": 3,  
  "total_jobs_run": 3,  
  "total_job_time": 0.000,  
  "dag_status": 0  
}
```

# More information

- › There's much more detail, as well as examples, in the DAGMan section of the online HTCondor manual.
- › DAGMan:  
<http://research.cs.wisc.edu/htcondor/dagman/dagman.html>
- › For more questions:  
[htcondor-admin@cs.wisc.edu](mailto:htcondor-admin@cs.wisc.edu), [htcondor-users@cs.wisc.edu](mailto:htcondor-users@cs.wisc.edu)

# Extra slides




# DAGMAN\_HOLD\_CLAIM\_TIME

- › An optimization introduced in HTCondor version 7.7.5 as a configuration option
- › If a DAGMan job has child nodes, it will instruct the HTCondor schedd to hold the machine claim for the integer number of seconds that is the value of this option, which defaults to 20.
- › Next job starts w/o negotiation cycle, using existing claim on startd

# Dot file

- › Shows a snapshot of workflow state
- › Updated atomically
- › For input to the dot visualization tool
- › In the DAG input file:  
`DOT DotFile [UPDATE] [DONT-OVERWRITE]`
- › To create an image  
`dot -Tps DotFile -o PostScriptFile`

# Dot file example



DAGMan Job status at Mon Apr 18 16:57:33 2011

# Node priorities (upcoming changes)

- Priority change to DAGMan job “trickles down” to nodes
- Different “inheritance” policy:
  - Effective node prio = explicit node prio + DAG prio?
  - Effective node prio = average(explicit node prio, parents' effective prios, DAG prio)?