CCTools
New Capabilities

2012-2013

CCTools Suite Overview

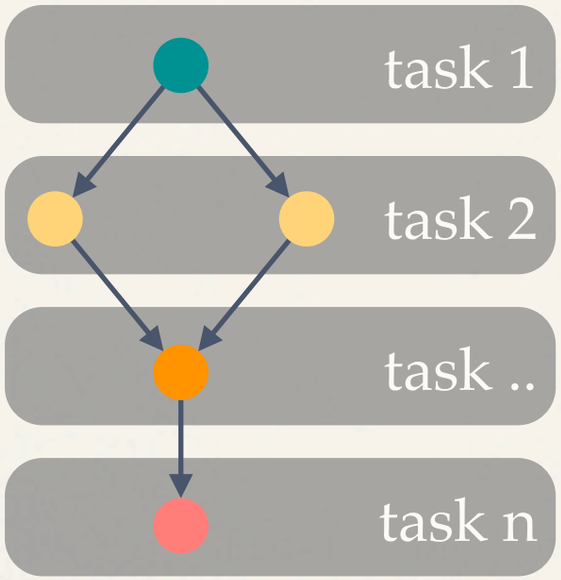

Scientist's Laptop

Demo Running Locally

Local Filesystem

Local to High Throughput

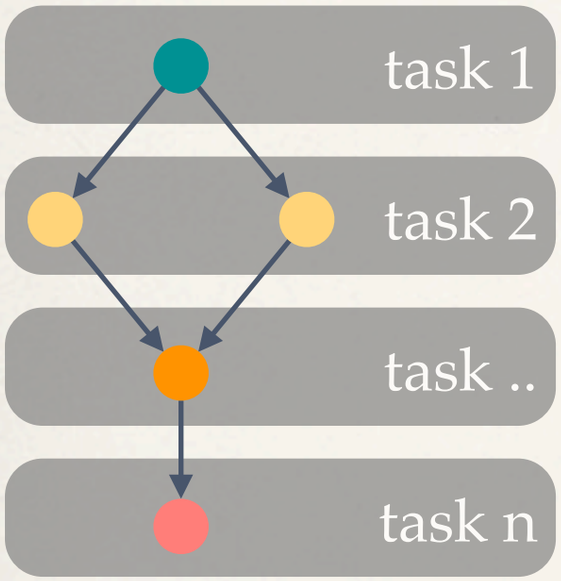

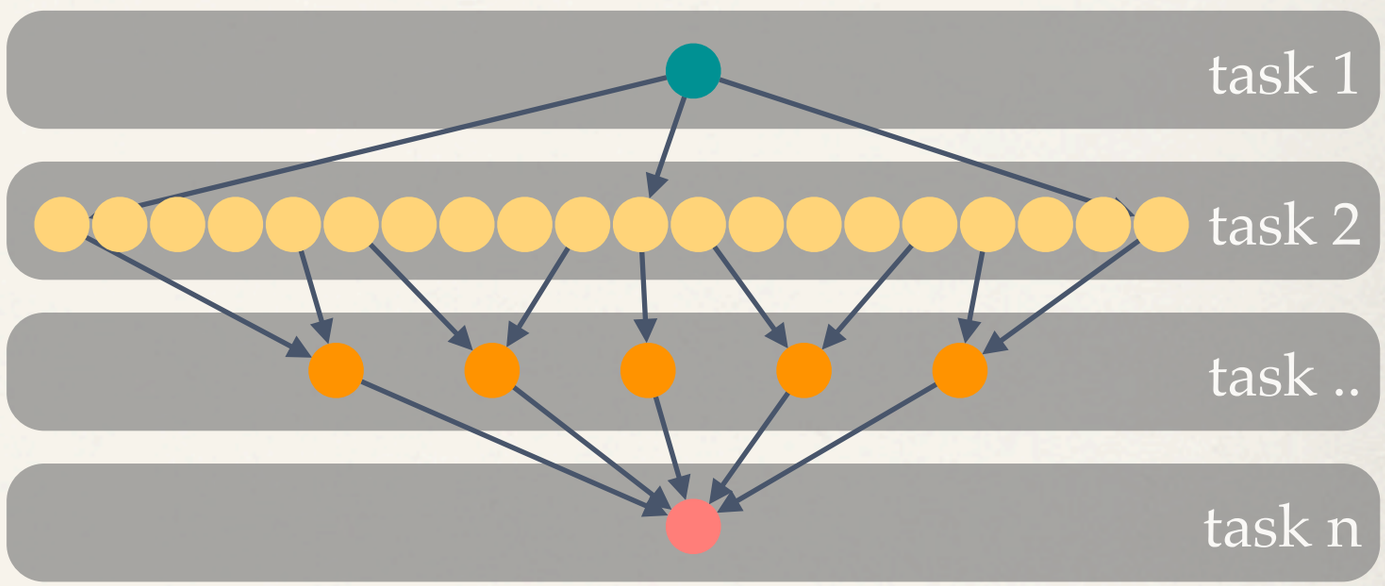

Demo Running Locally

Local Filesystem

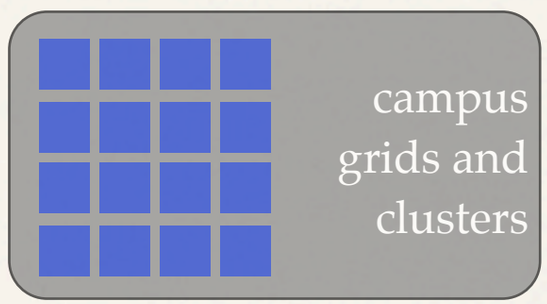

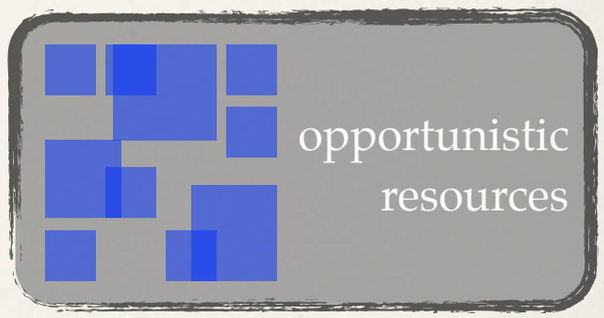

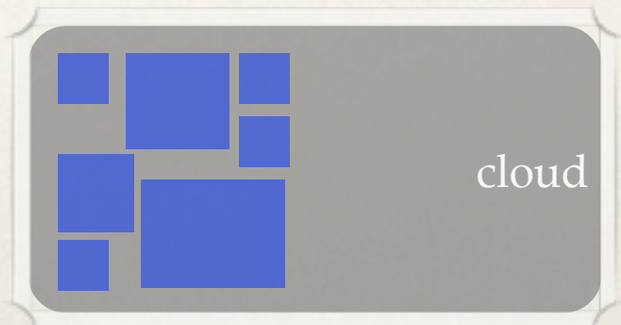

Our Philosophy

- Harness all the resources that are available: desktops, clusters, clouds, and grids.
- Make it easy to scale up from one desktop to national scale infrastructure.
- Provide familiar interfaces that make it easy to connect existing apps together.
- Allow portability across operating systems, storage systems, middleware…
- Make simple things easy, and complex things possible.
- No special privileges required.

CCTools Suite

Makeflow

portable workflow manager

run this task when this other task is completed

make syntax

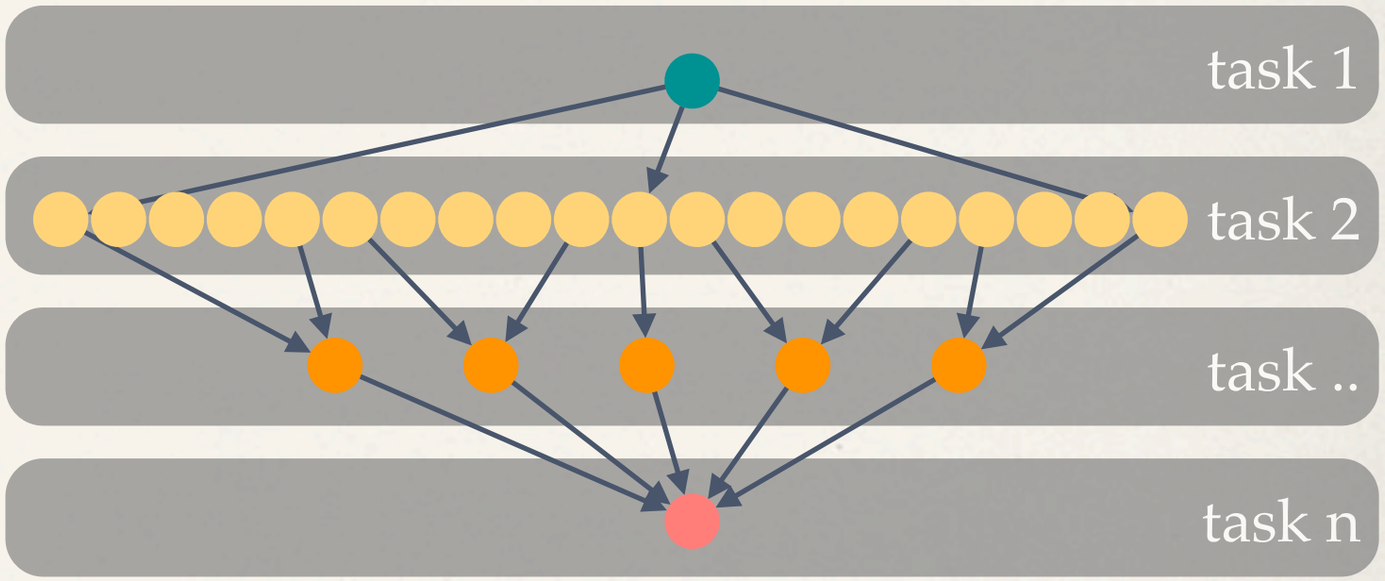

CCTools Suite

Makeflow

portable workflow manager

run this task when this other task is completed

make syntax

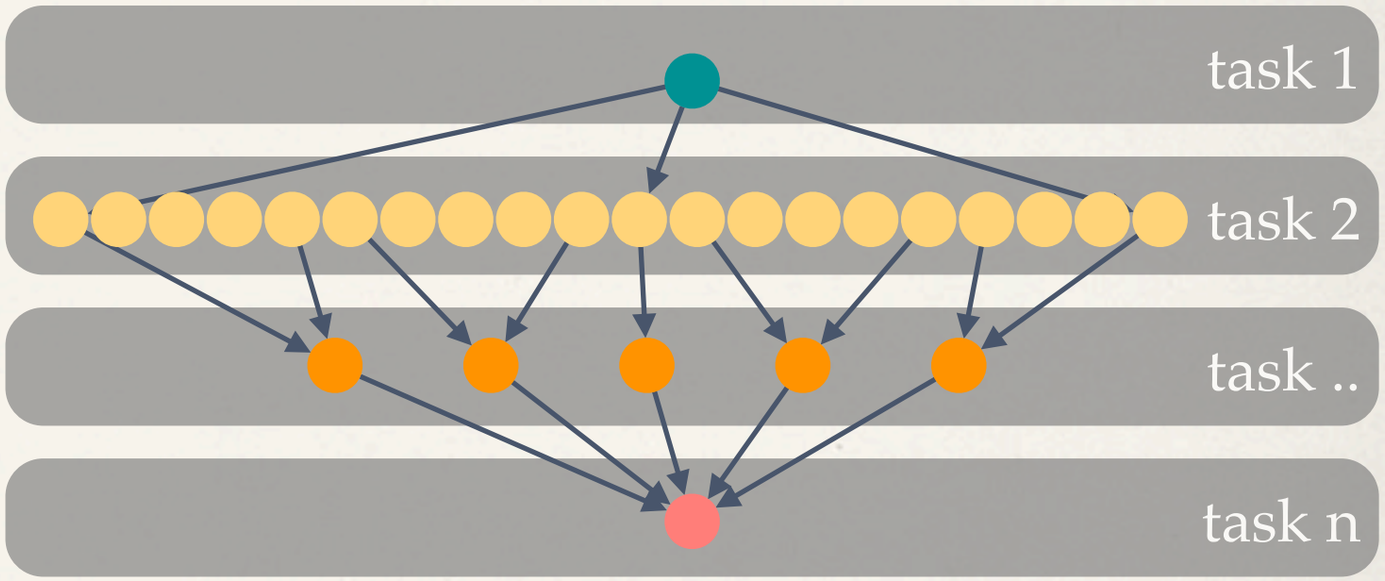

local | work queue | condor | sge | ...

CCTools Suite

AllPairs

specialized execution engine

apply the operation to all combinations

cross products

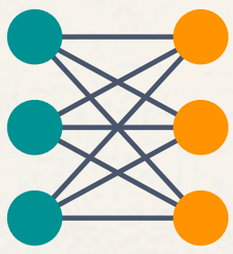

Wavefront

specialized execution engine

run this operation following a wave pattern

dynamic programming

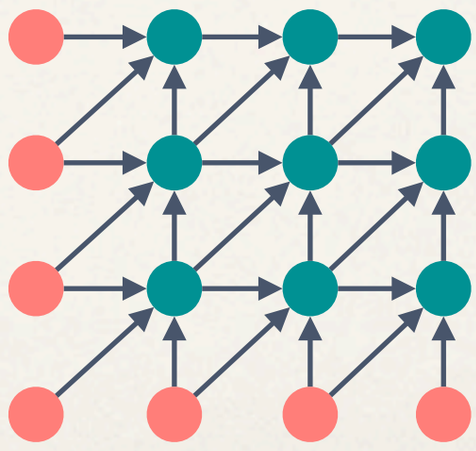

CCTools Suite

Work Queue

lightweight distributed
execution system

run this task in some available resource

master-worker architecture

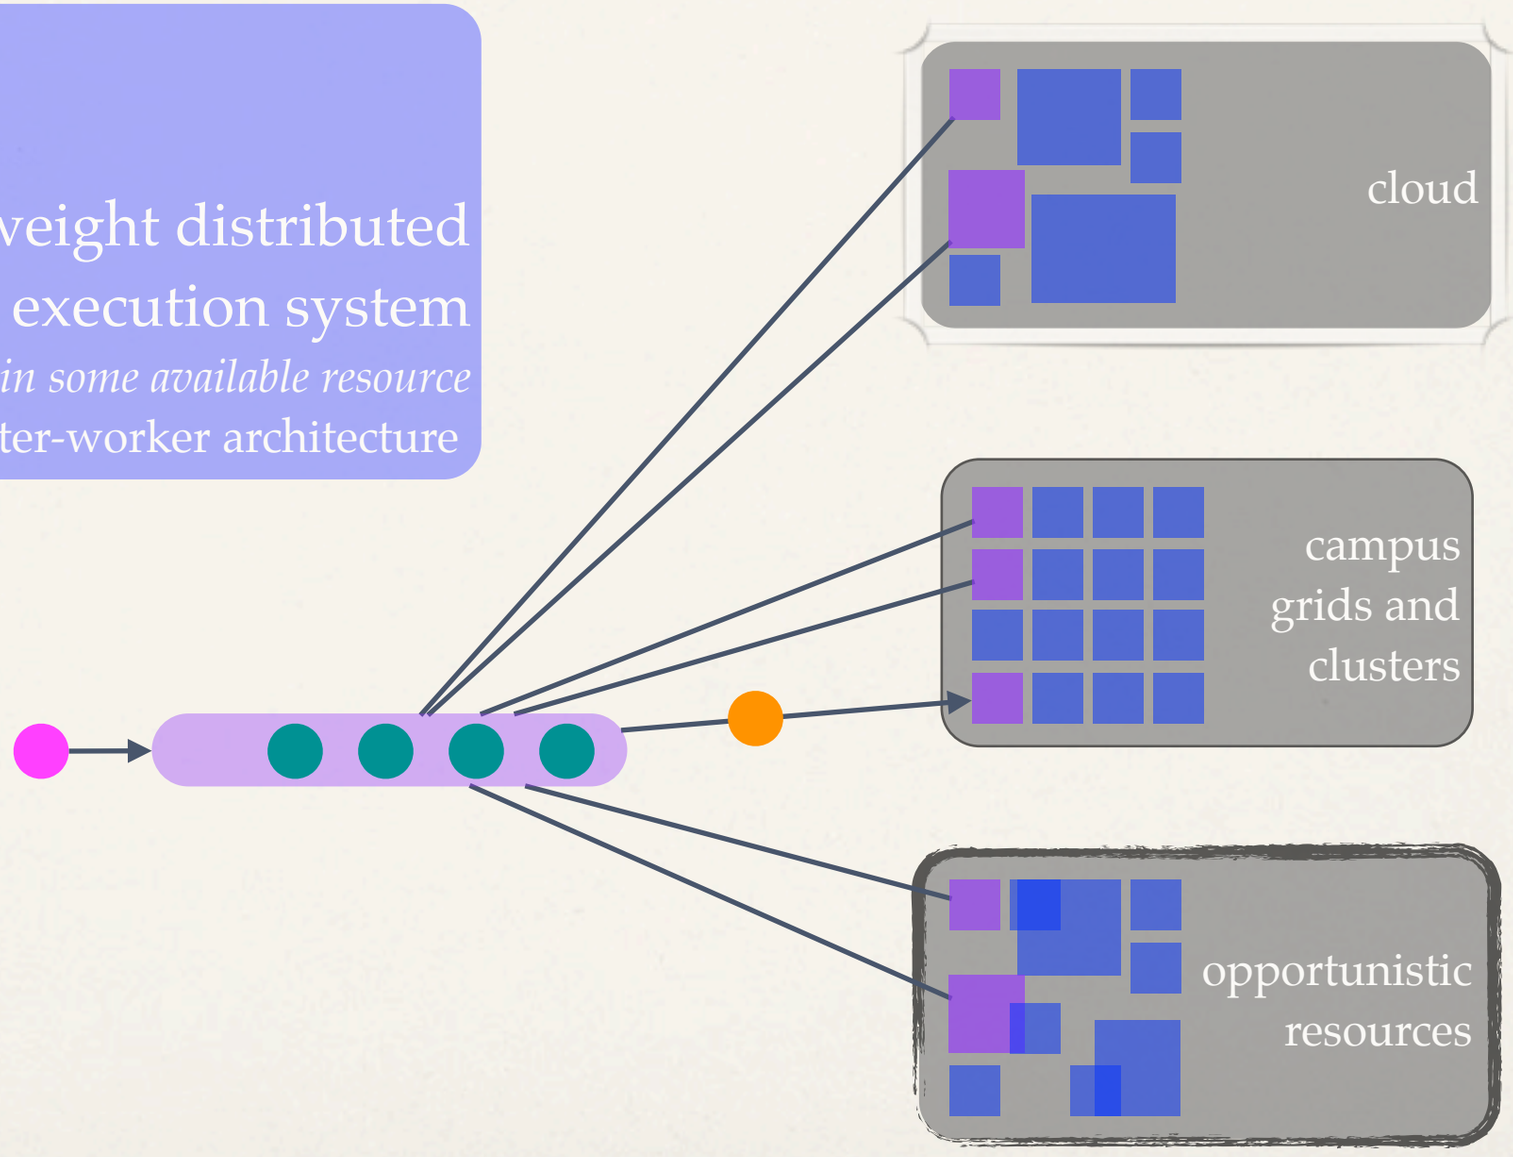

CCTools Suite

Work Queue

lightweight distributed
execution system

run this task in some available resource

master-worker architecture

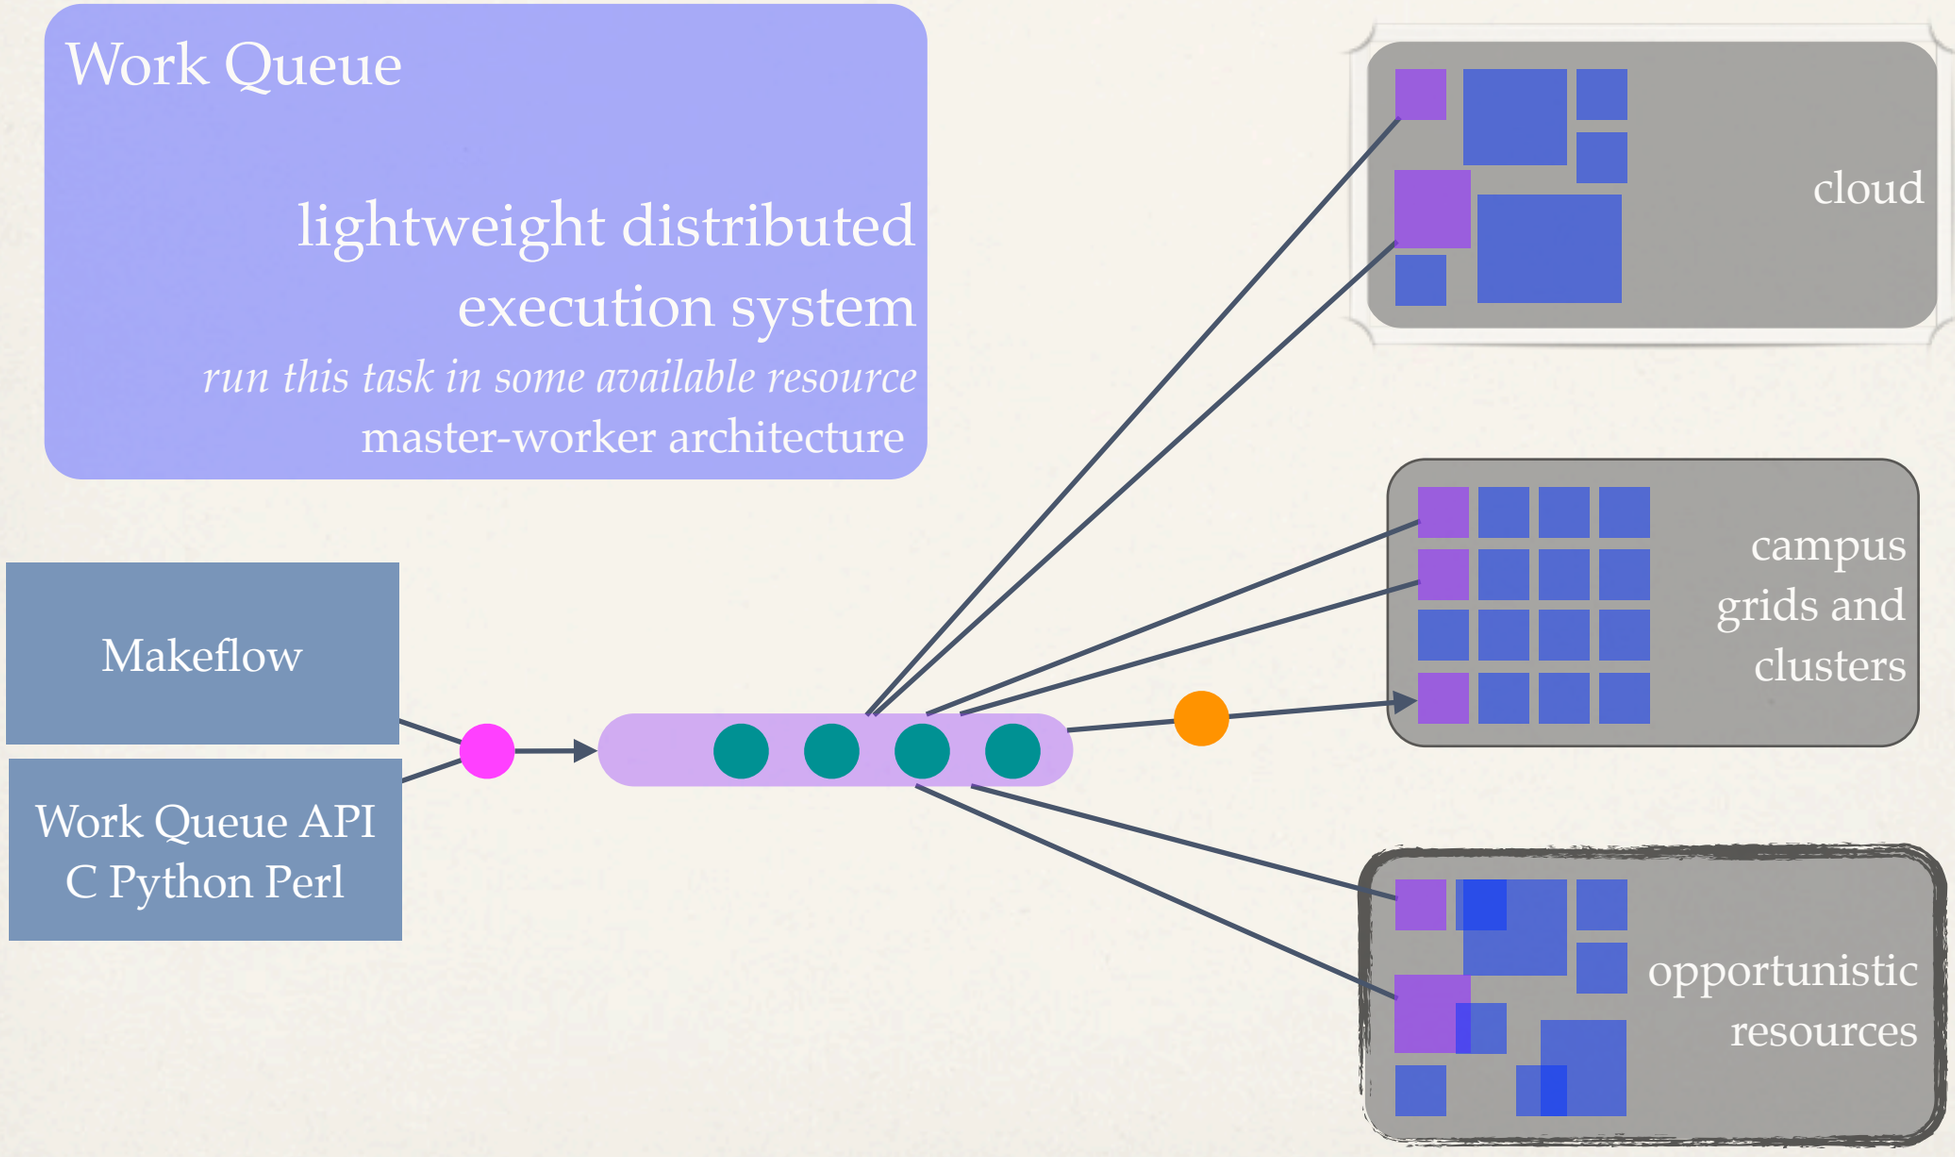

CCTools Suite

Chirp

user-level distributed filesystem

file sharing for scientific workflows

mount volumes without root access

at some_sever.xyz

```
chirp_server -r /home/user/my_files
```

Parrot

personal user-level virtual filesystem

access remote files as if they were local

system call interposition agent

at the local machine

```
parrot_run ls /chirp/some_server.xyz/*.txt
parrot_run cat /http/www.nd.edu/index.html
parrot_run ls /anonftp/ftp.gnu.org/pub
parrot_run cd /cvmfs/cms.cern.ch
```

CCTools Suite

- Open source, GNU General Public License 2.
- Compiles in 1-2 minutes, installs in $HOME.
- All tools may be used independently or in conjunction.
- Runs on Linux, Solaris, OSX, Cygwin, FreeBSD, …
- Interoperates with many distributed computing systems.
  - Condor, SGE, Torque, Globus, iRODS, Hadoop…

New Capabilities

2012-2013

Work Queue Before

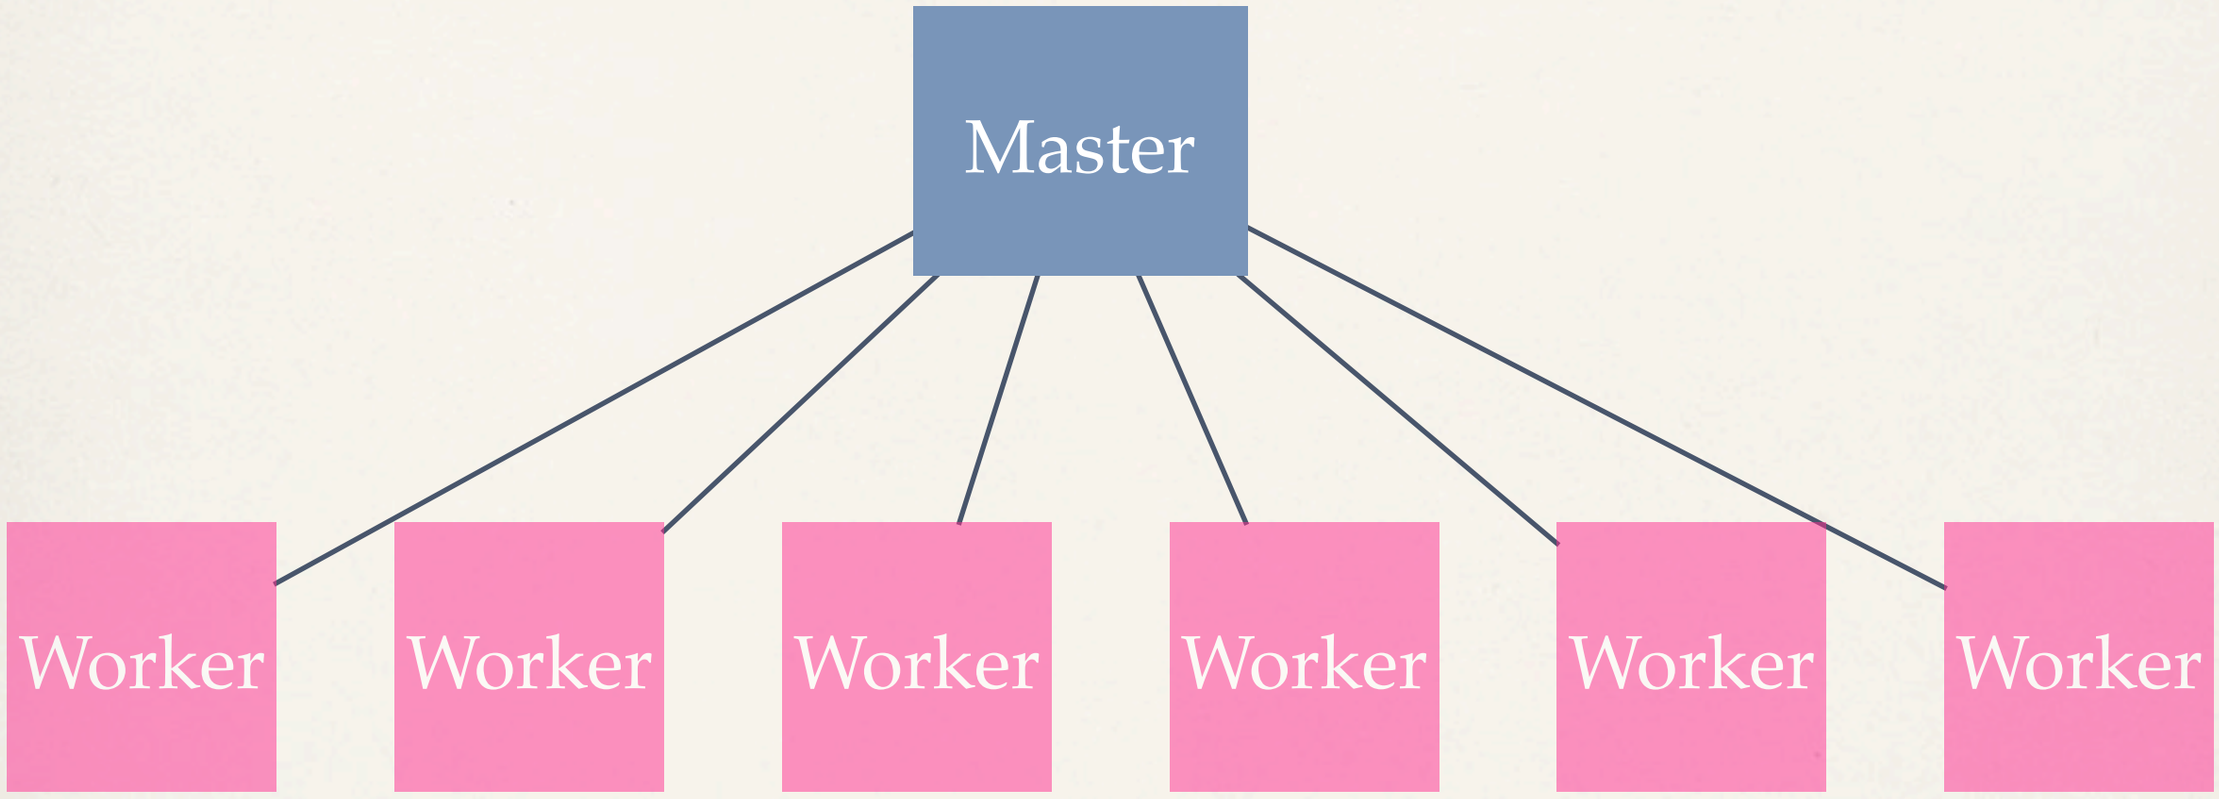

1. Limited scalability: ~1000 (number of open file descriptors)
2. Limited throughput: One task per worker
3. No authentication: Any worker can pull data from master
4. Big file transfers over WAN

Work Queue Hierarchy

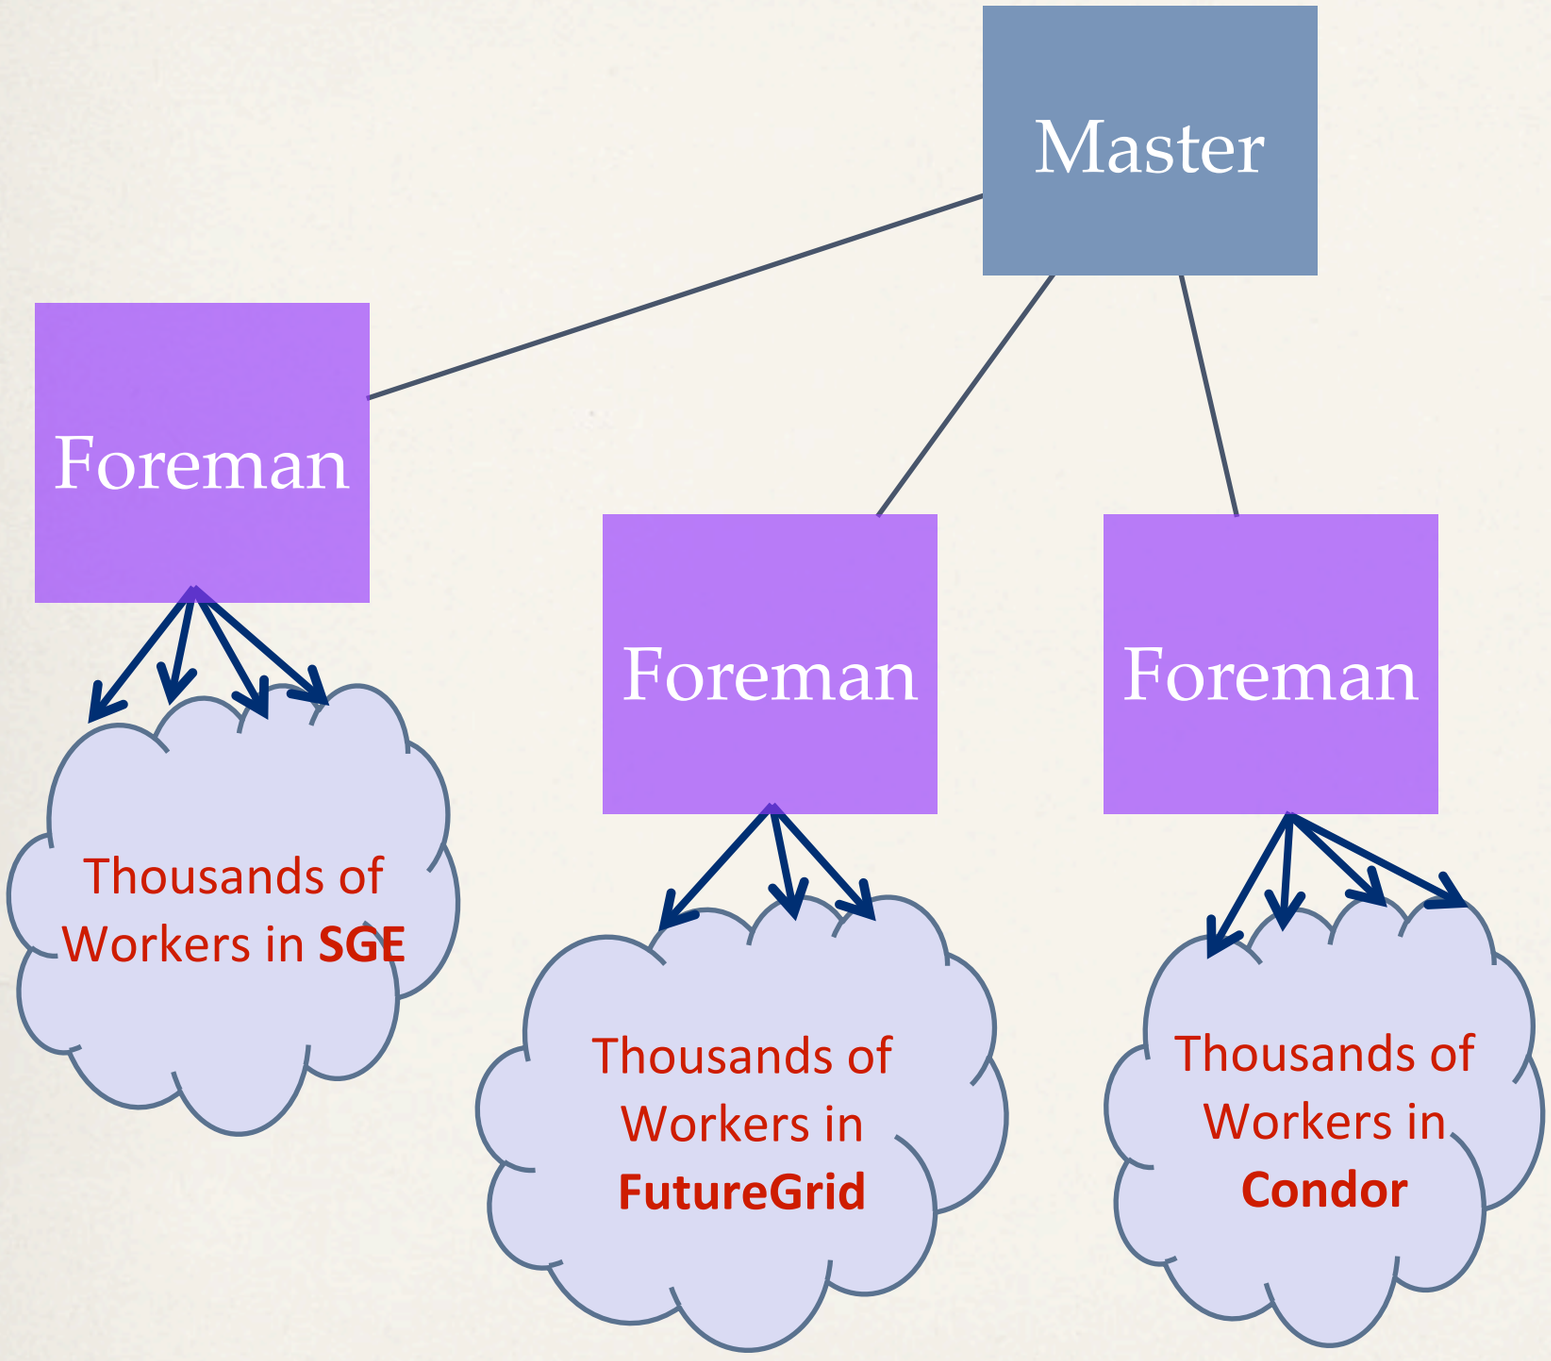

Work Queue Hierarchy

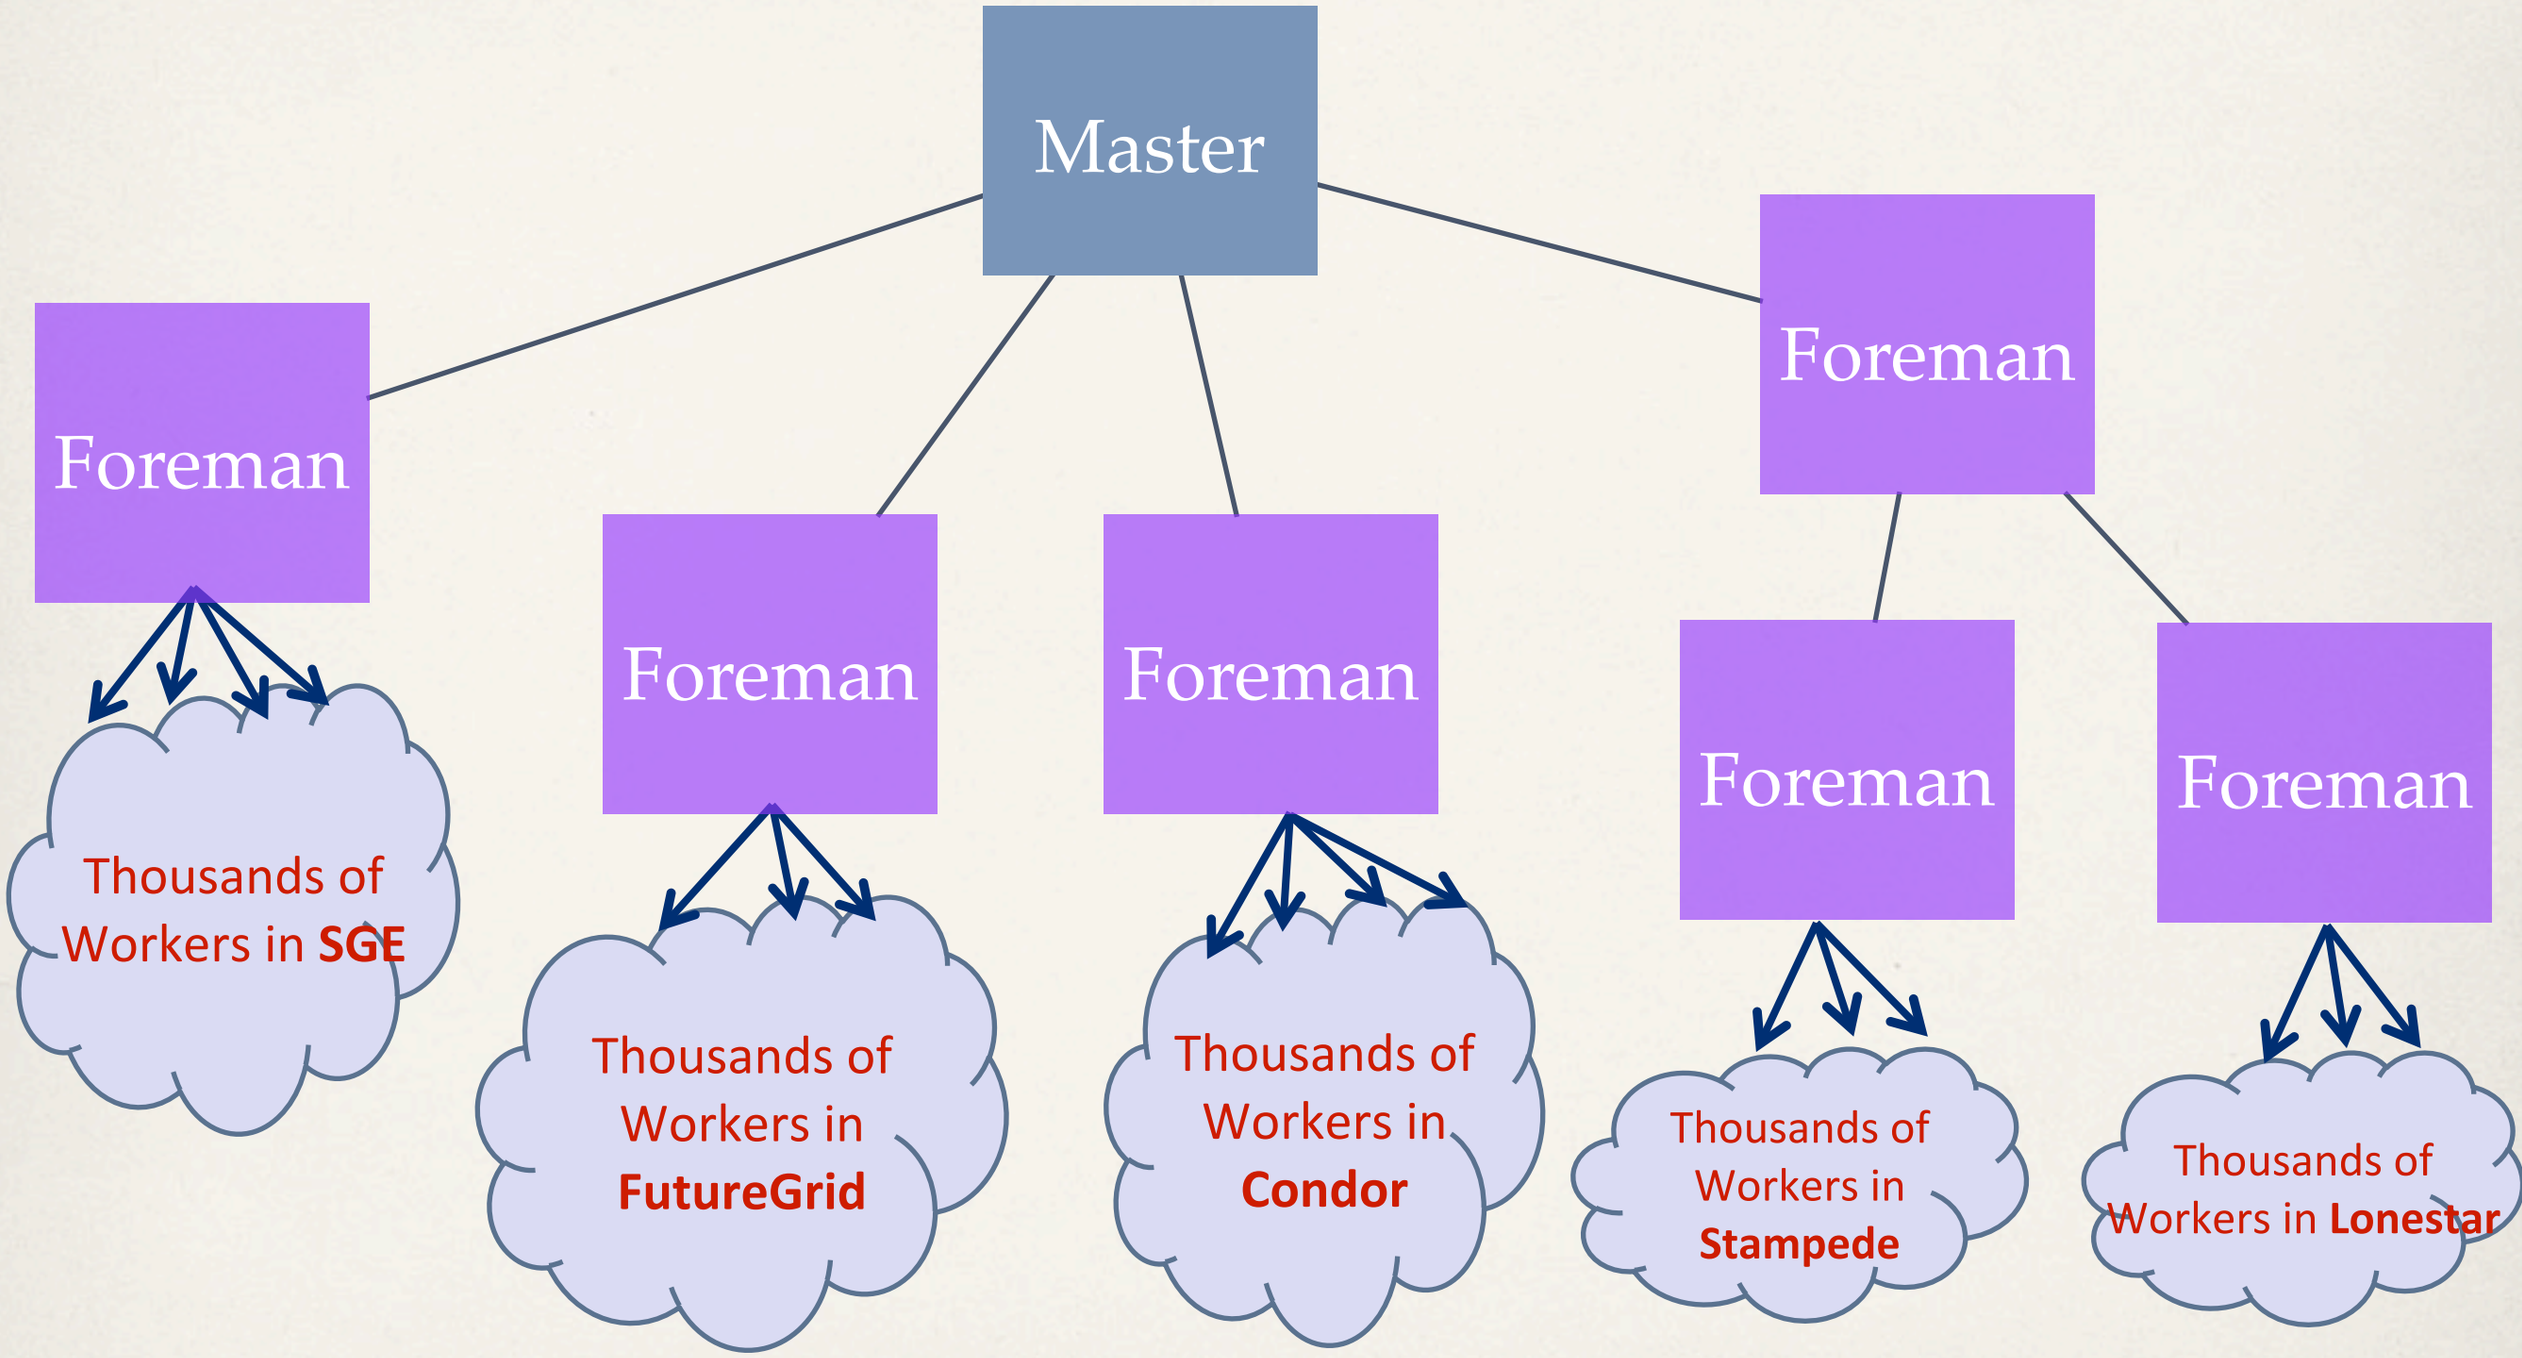

Work Queue Hierarchy

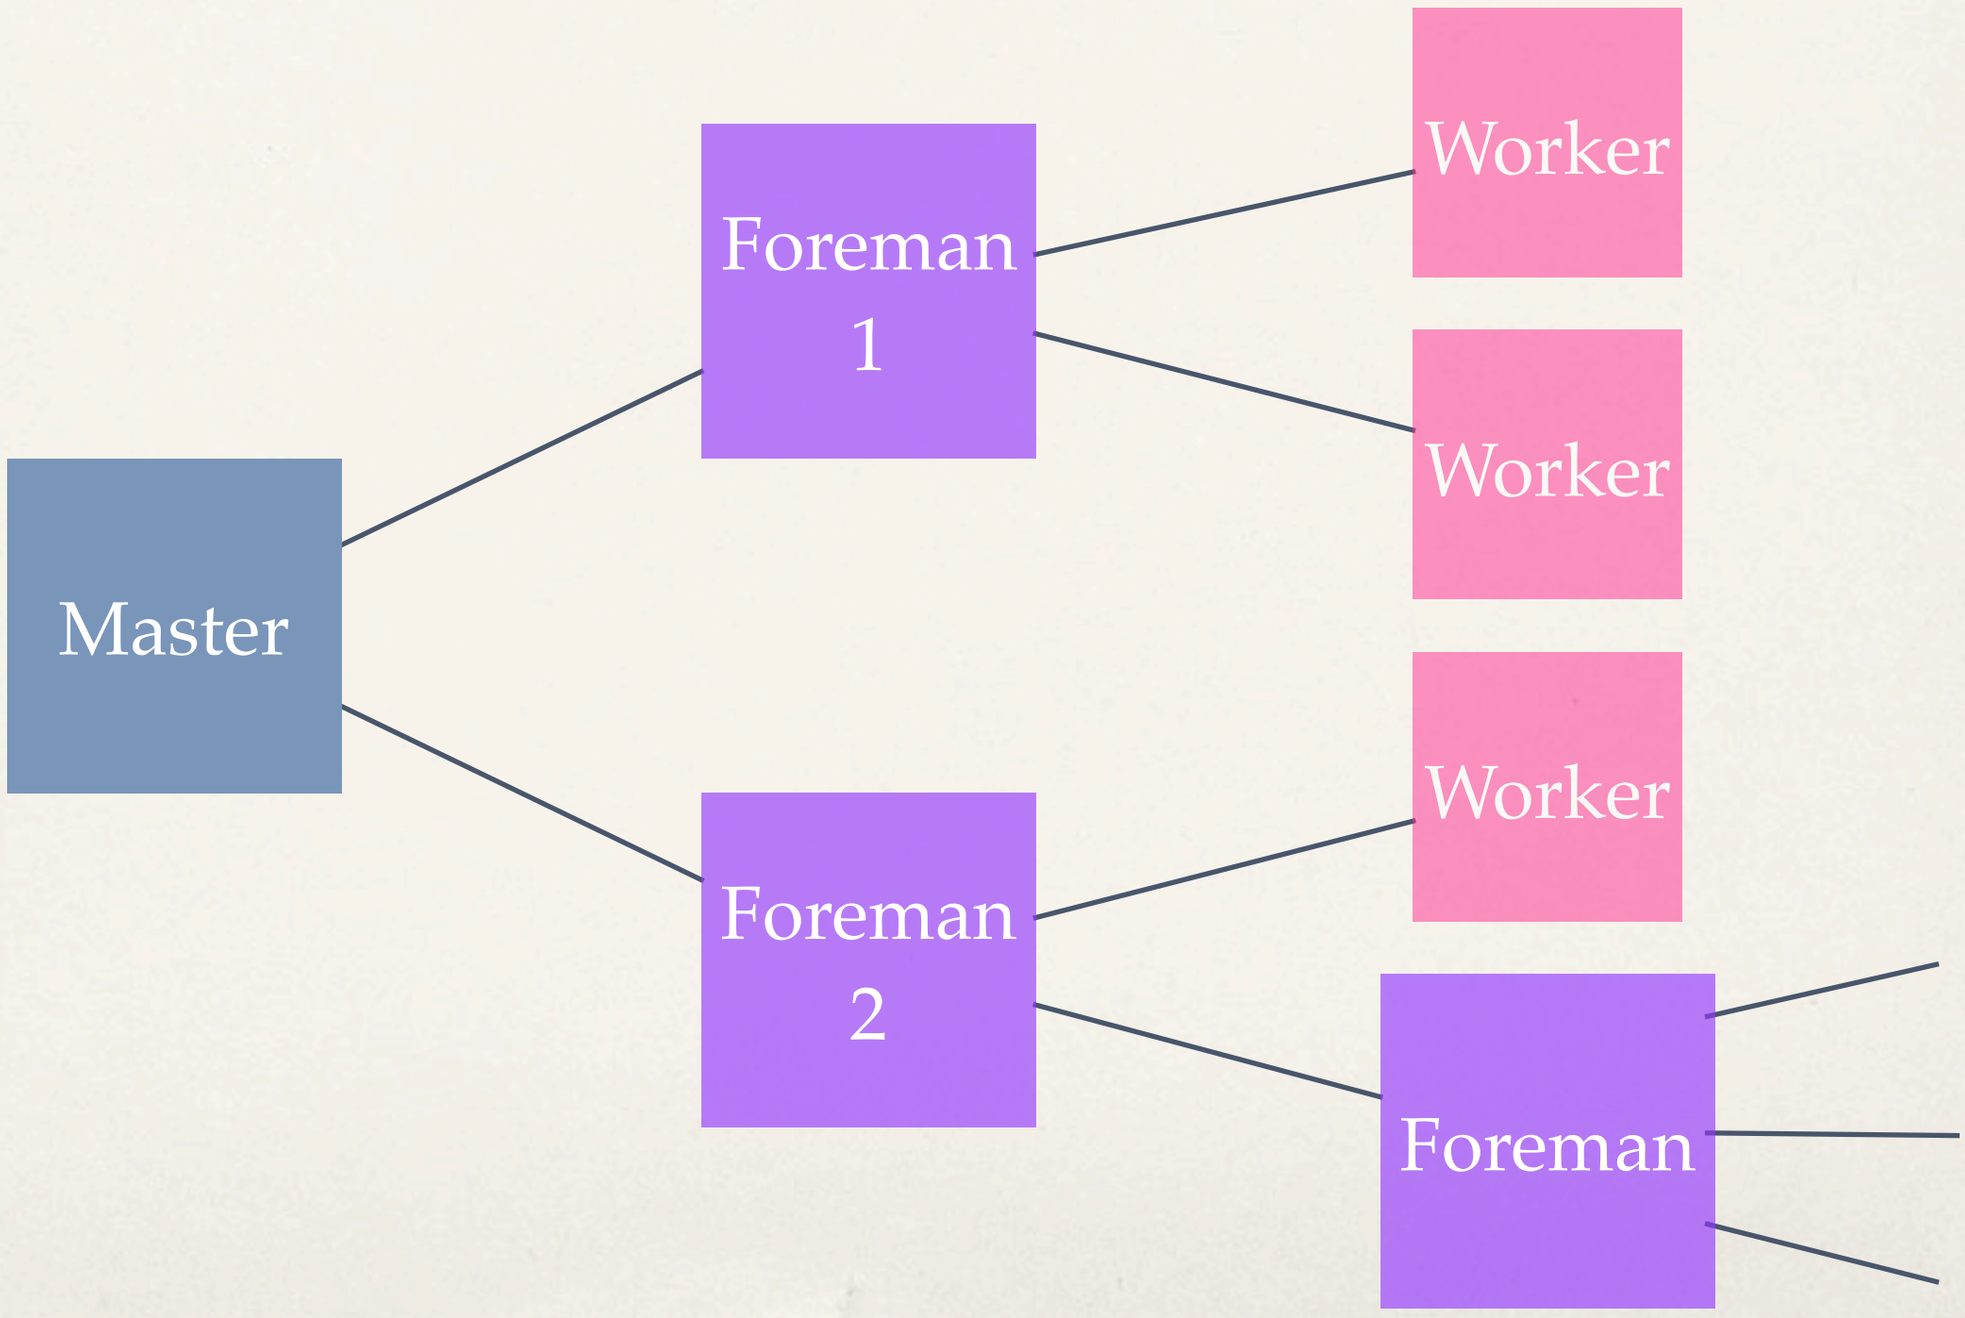

Work Queue Hierarchy

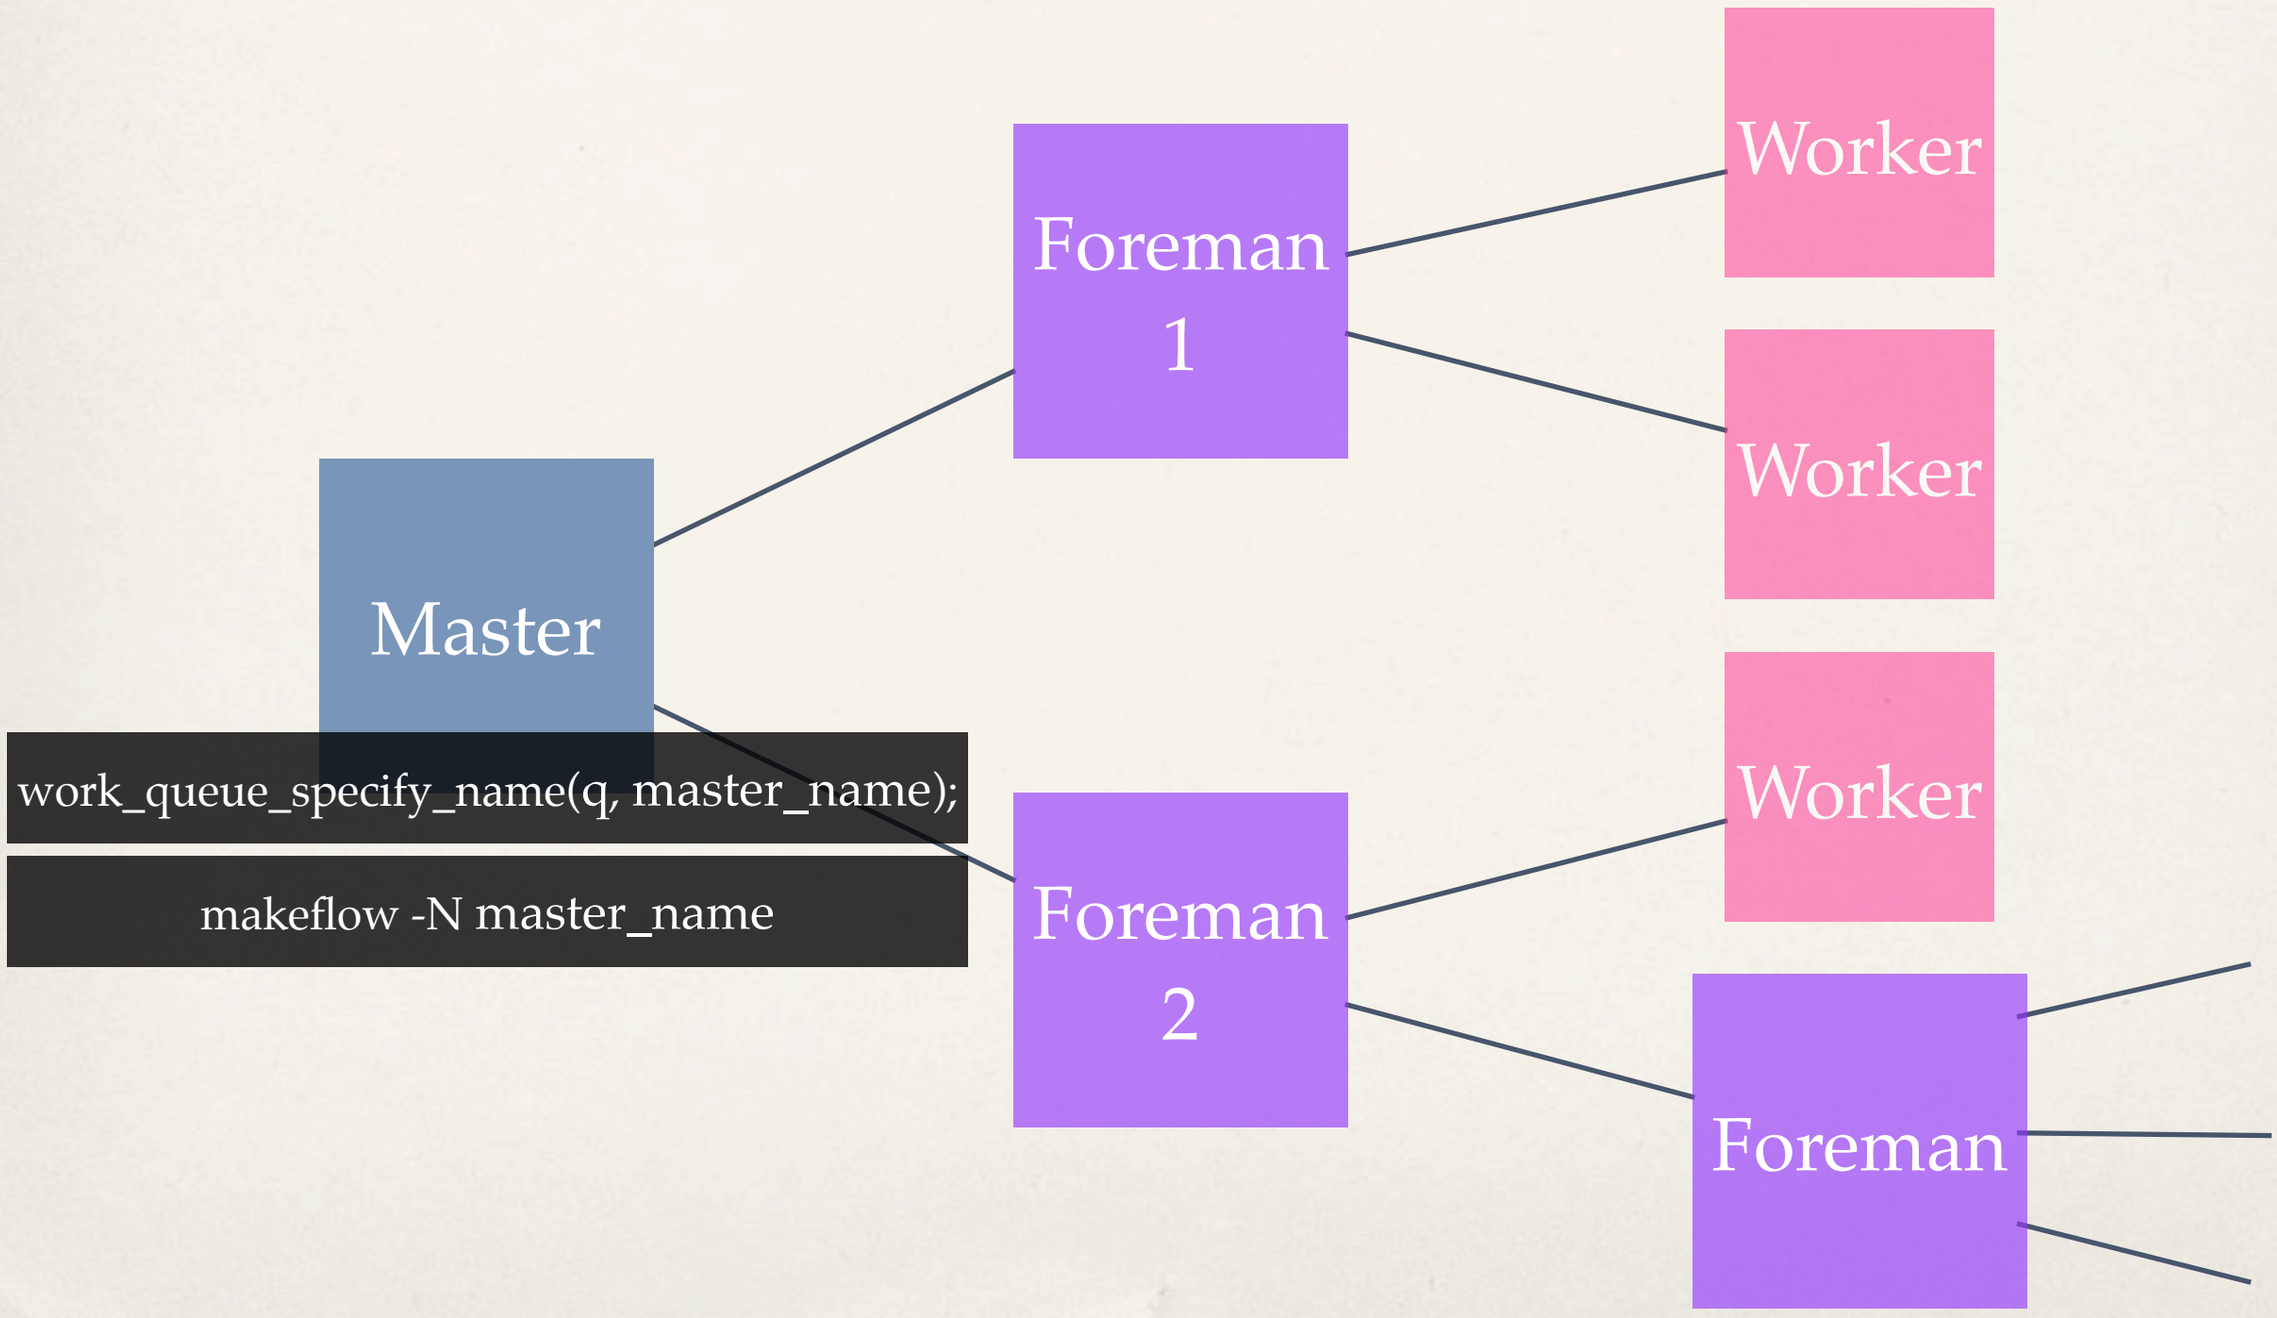

Work Queue Hierarchy

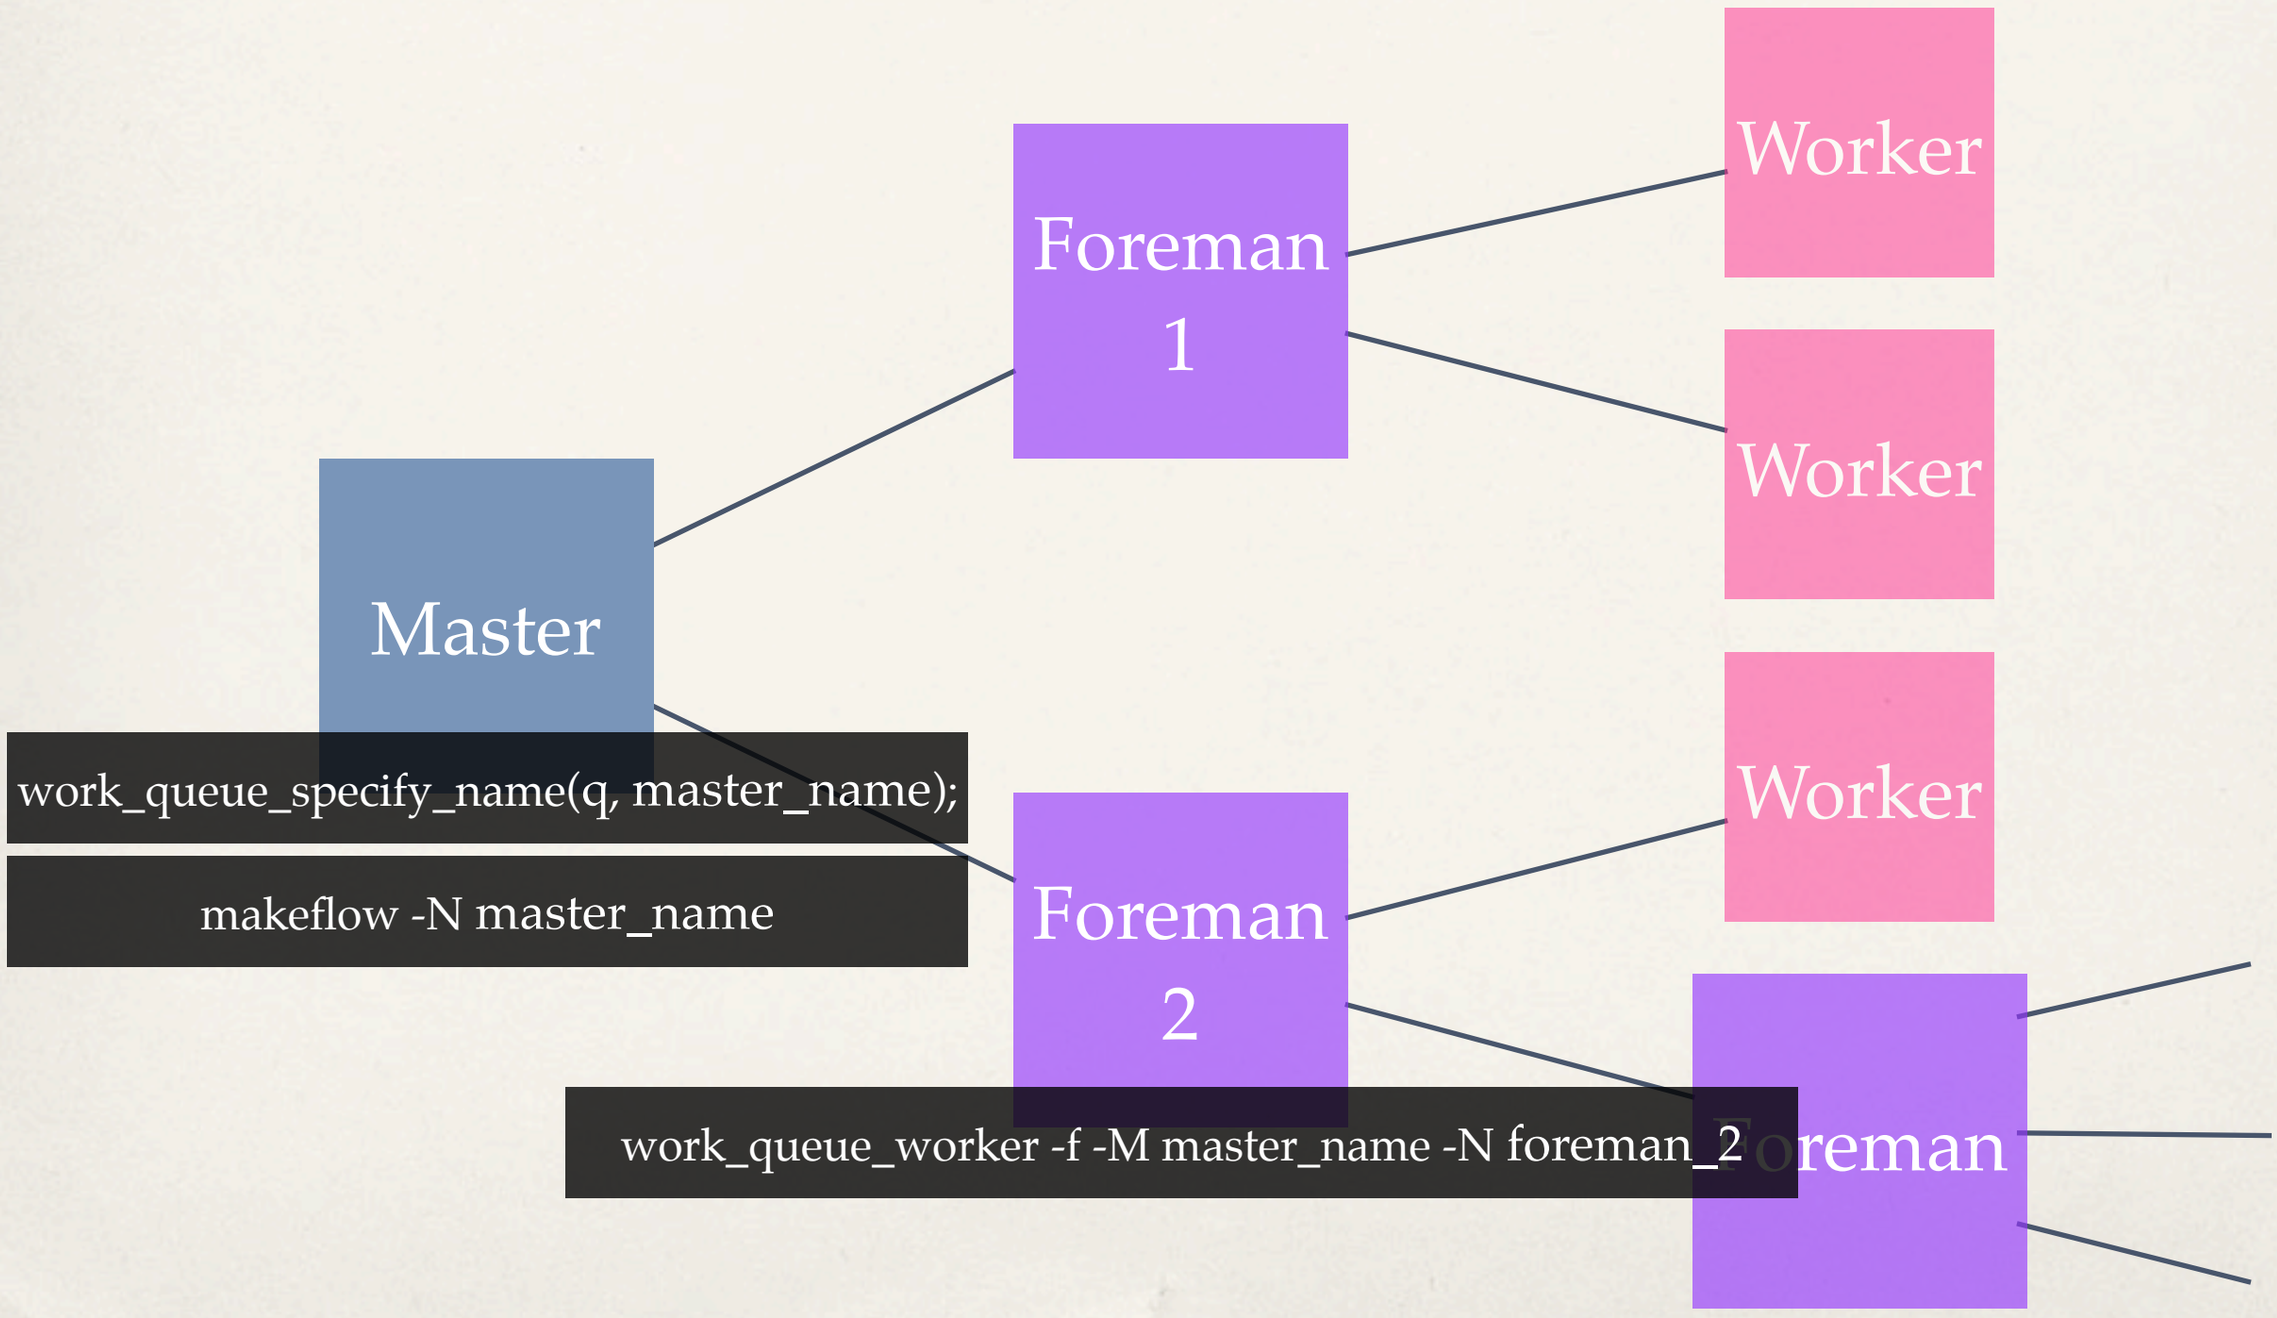

Work Queue Hierarchy

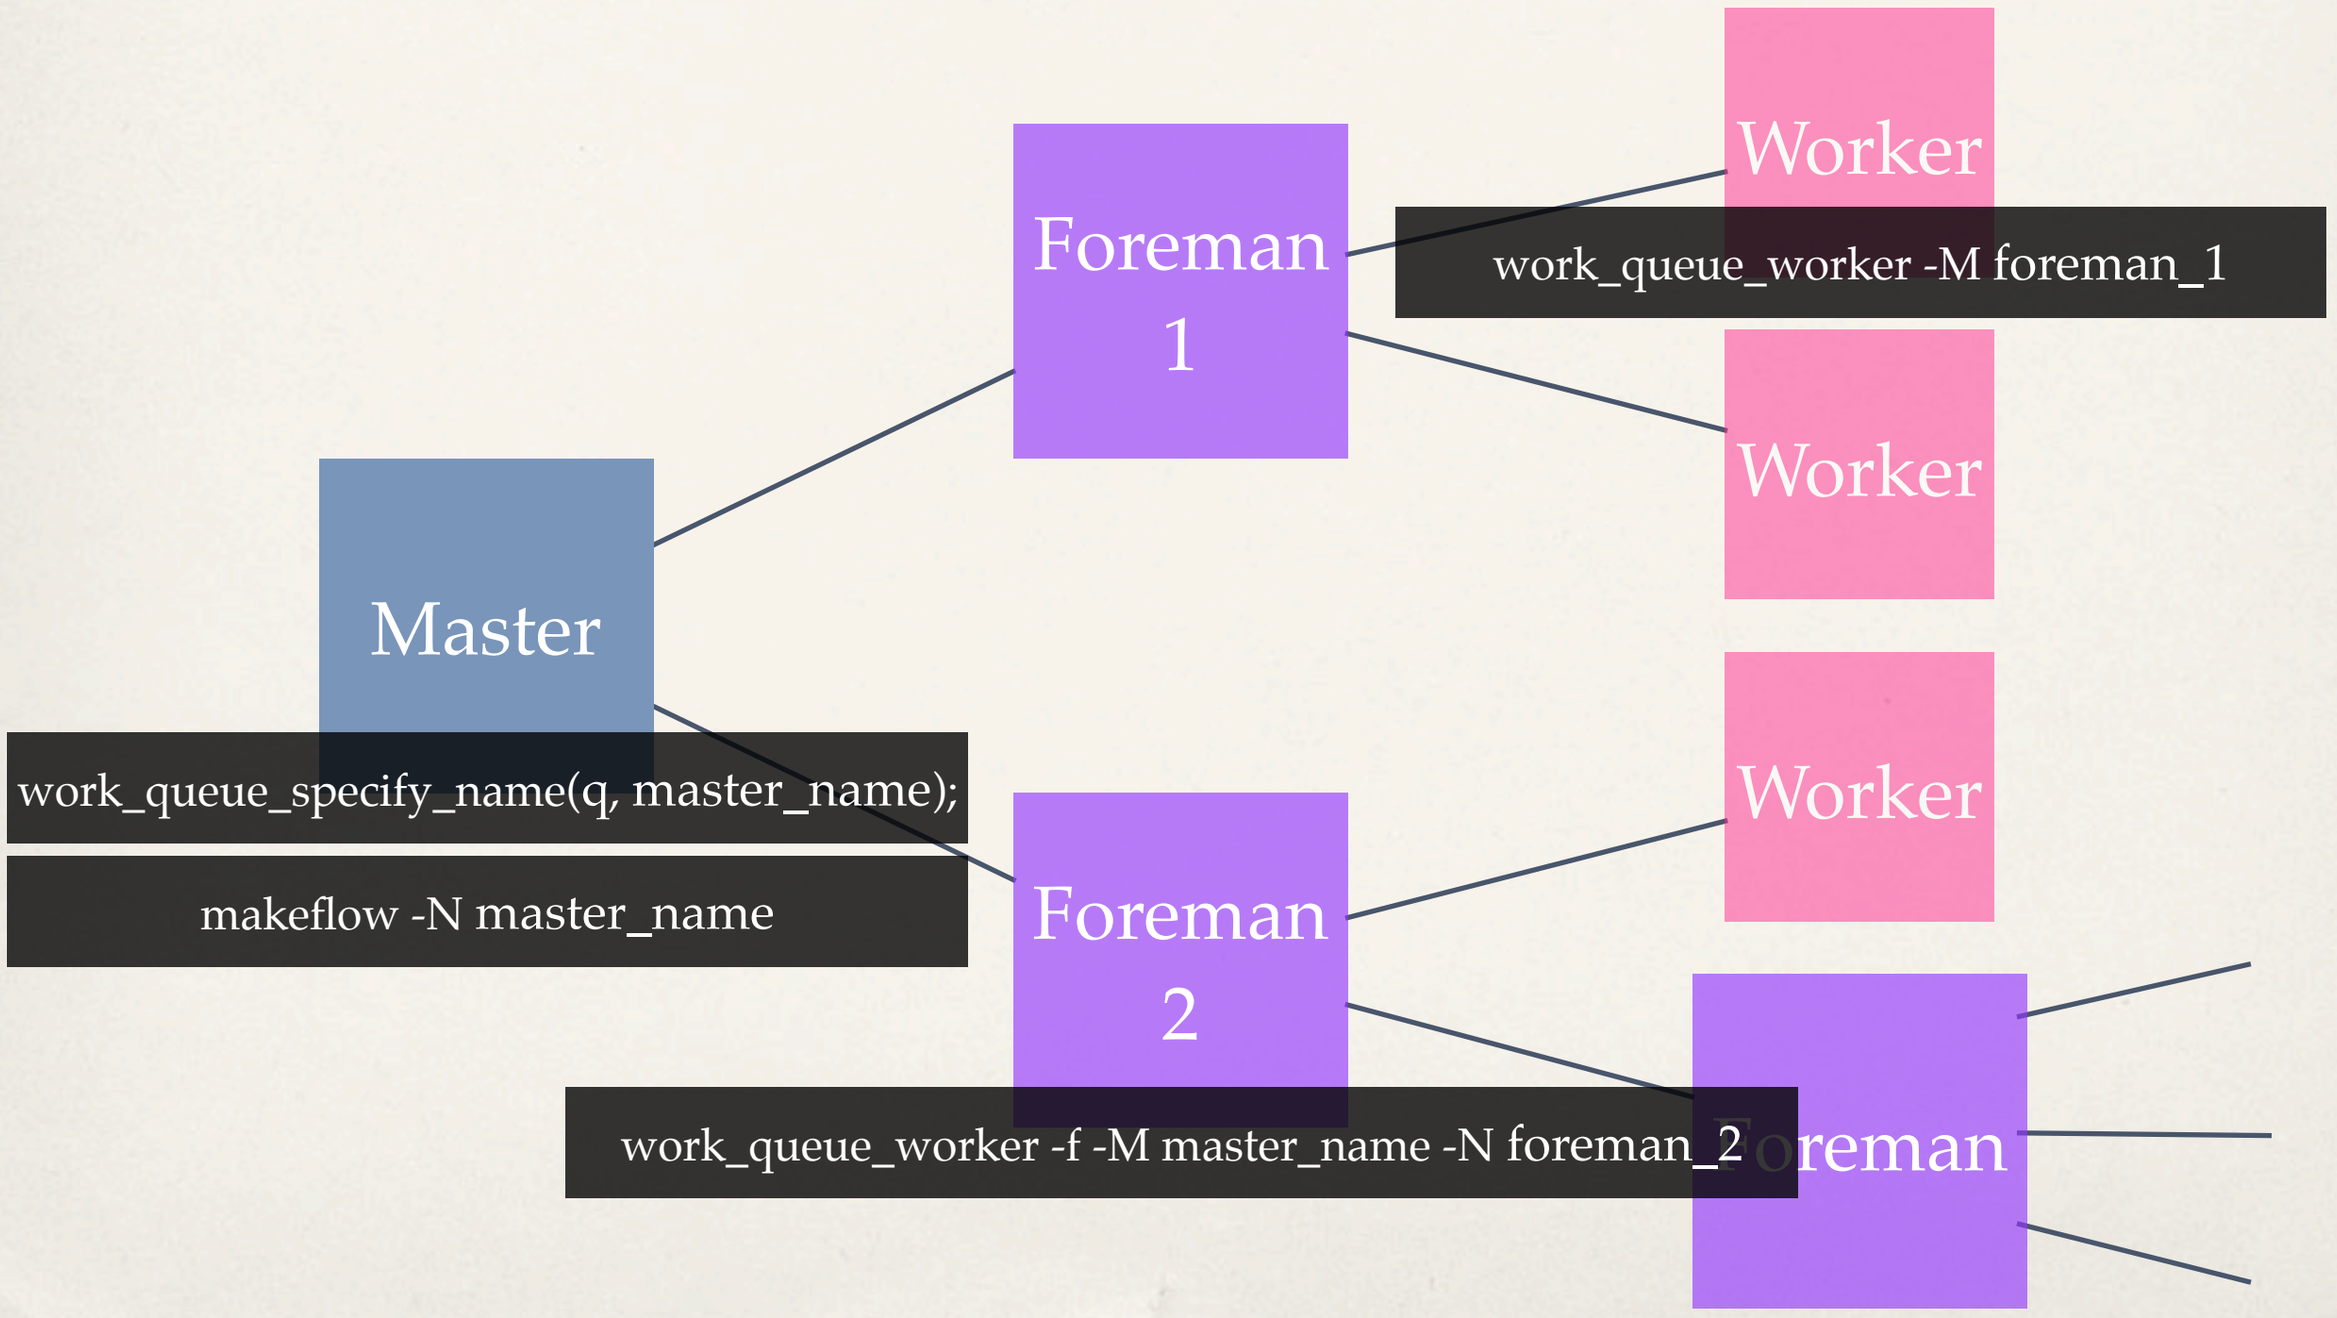

Work Queue Hierarchy

- Foremen cache files

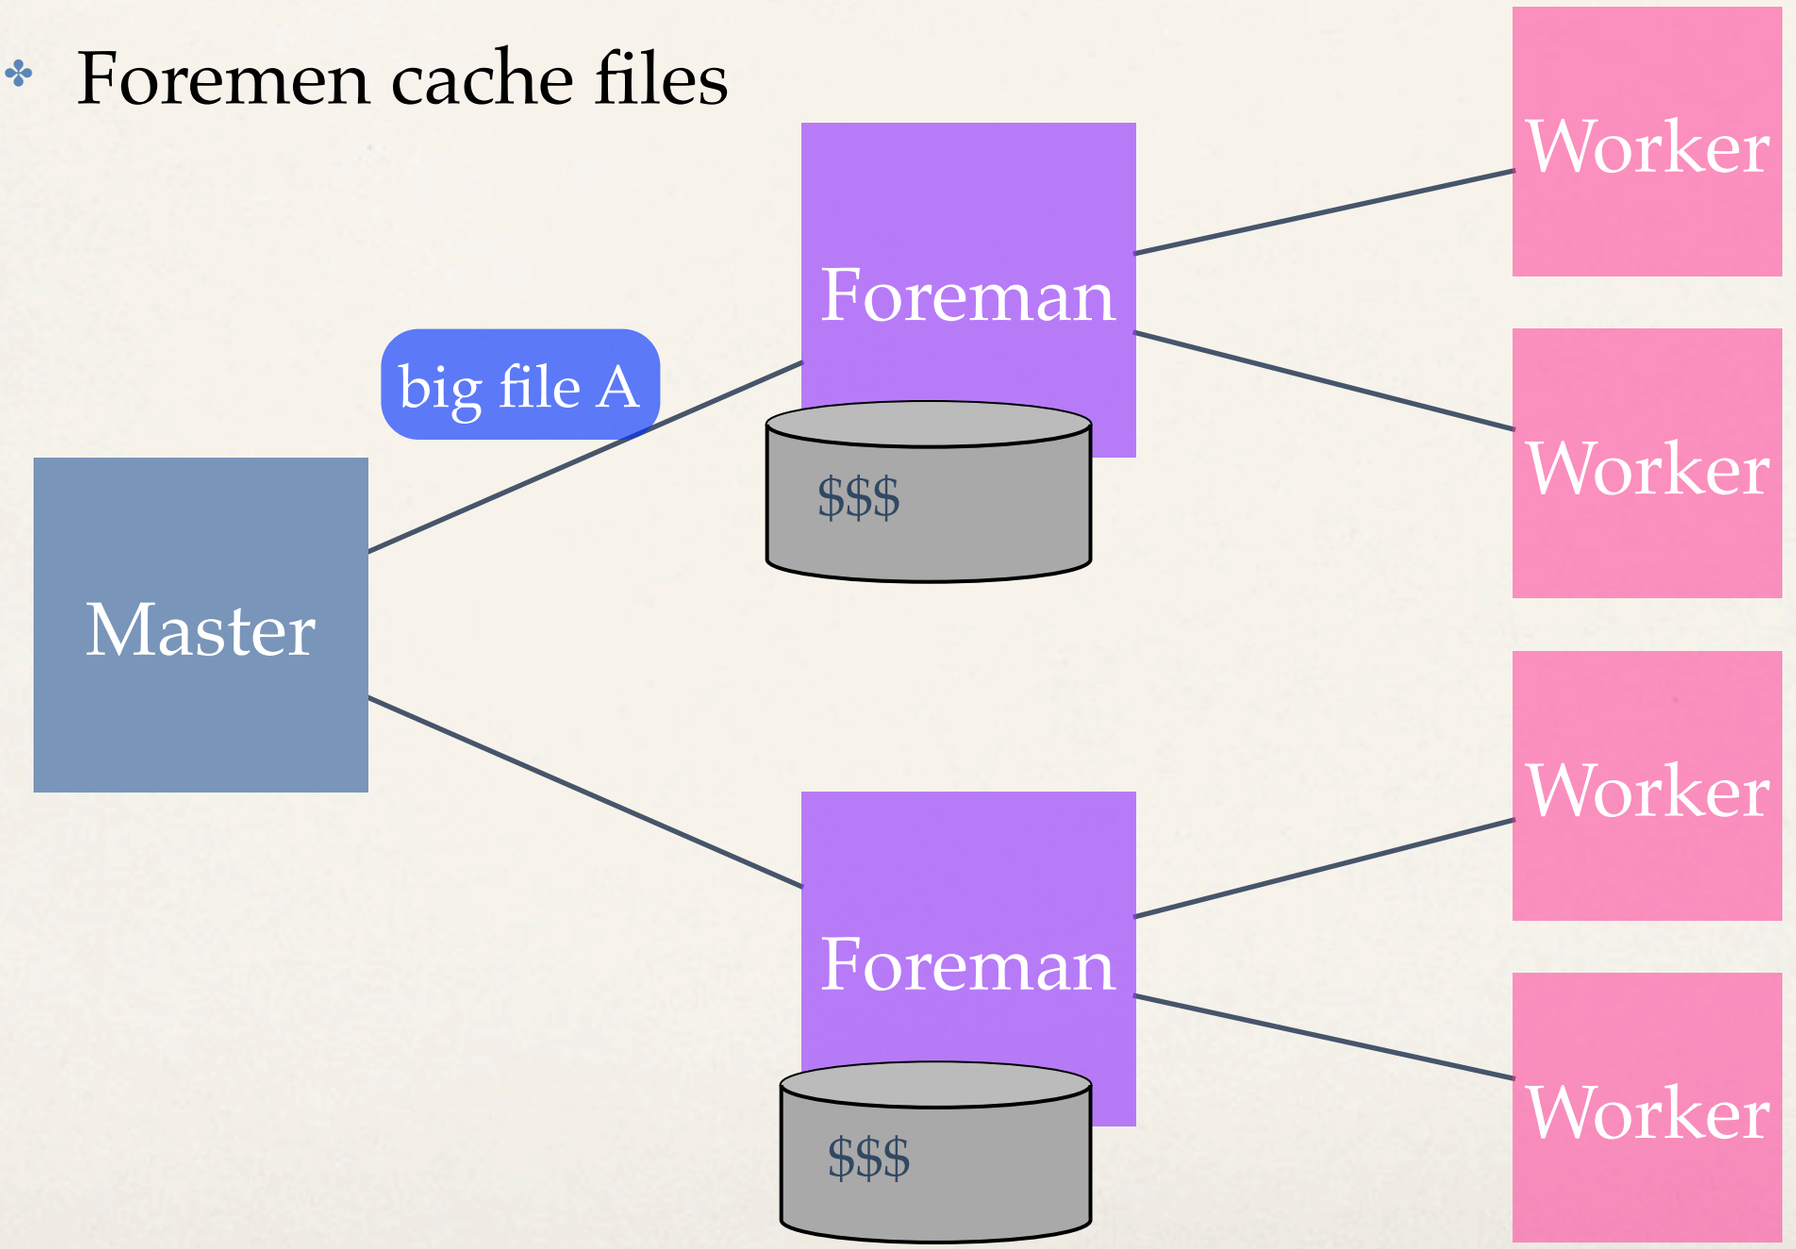

Work Queue Hierarchy

- Foremen cache files

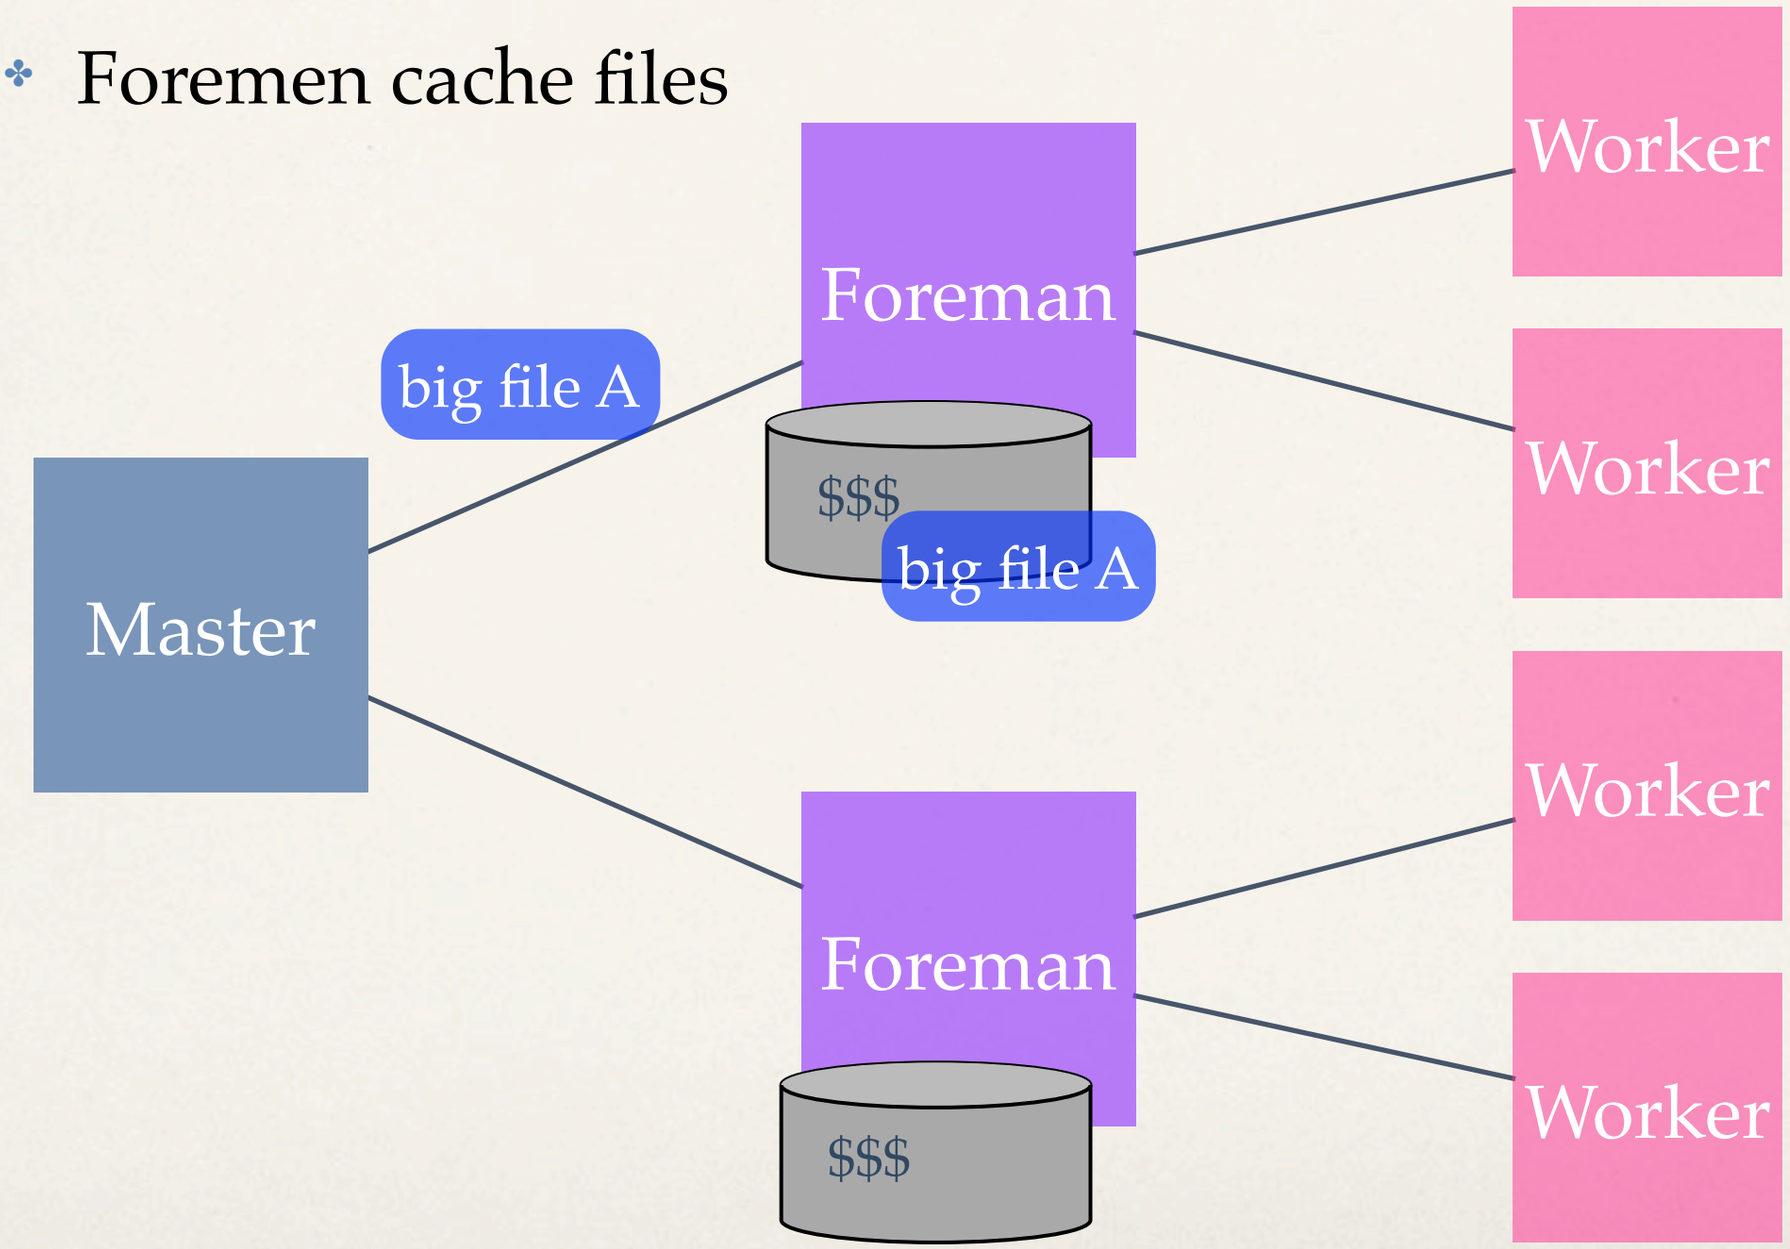

Work Queue Hierarchy

- Foremen cache files

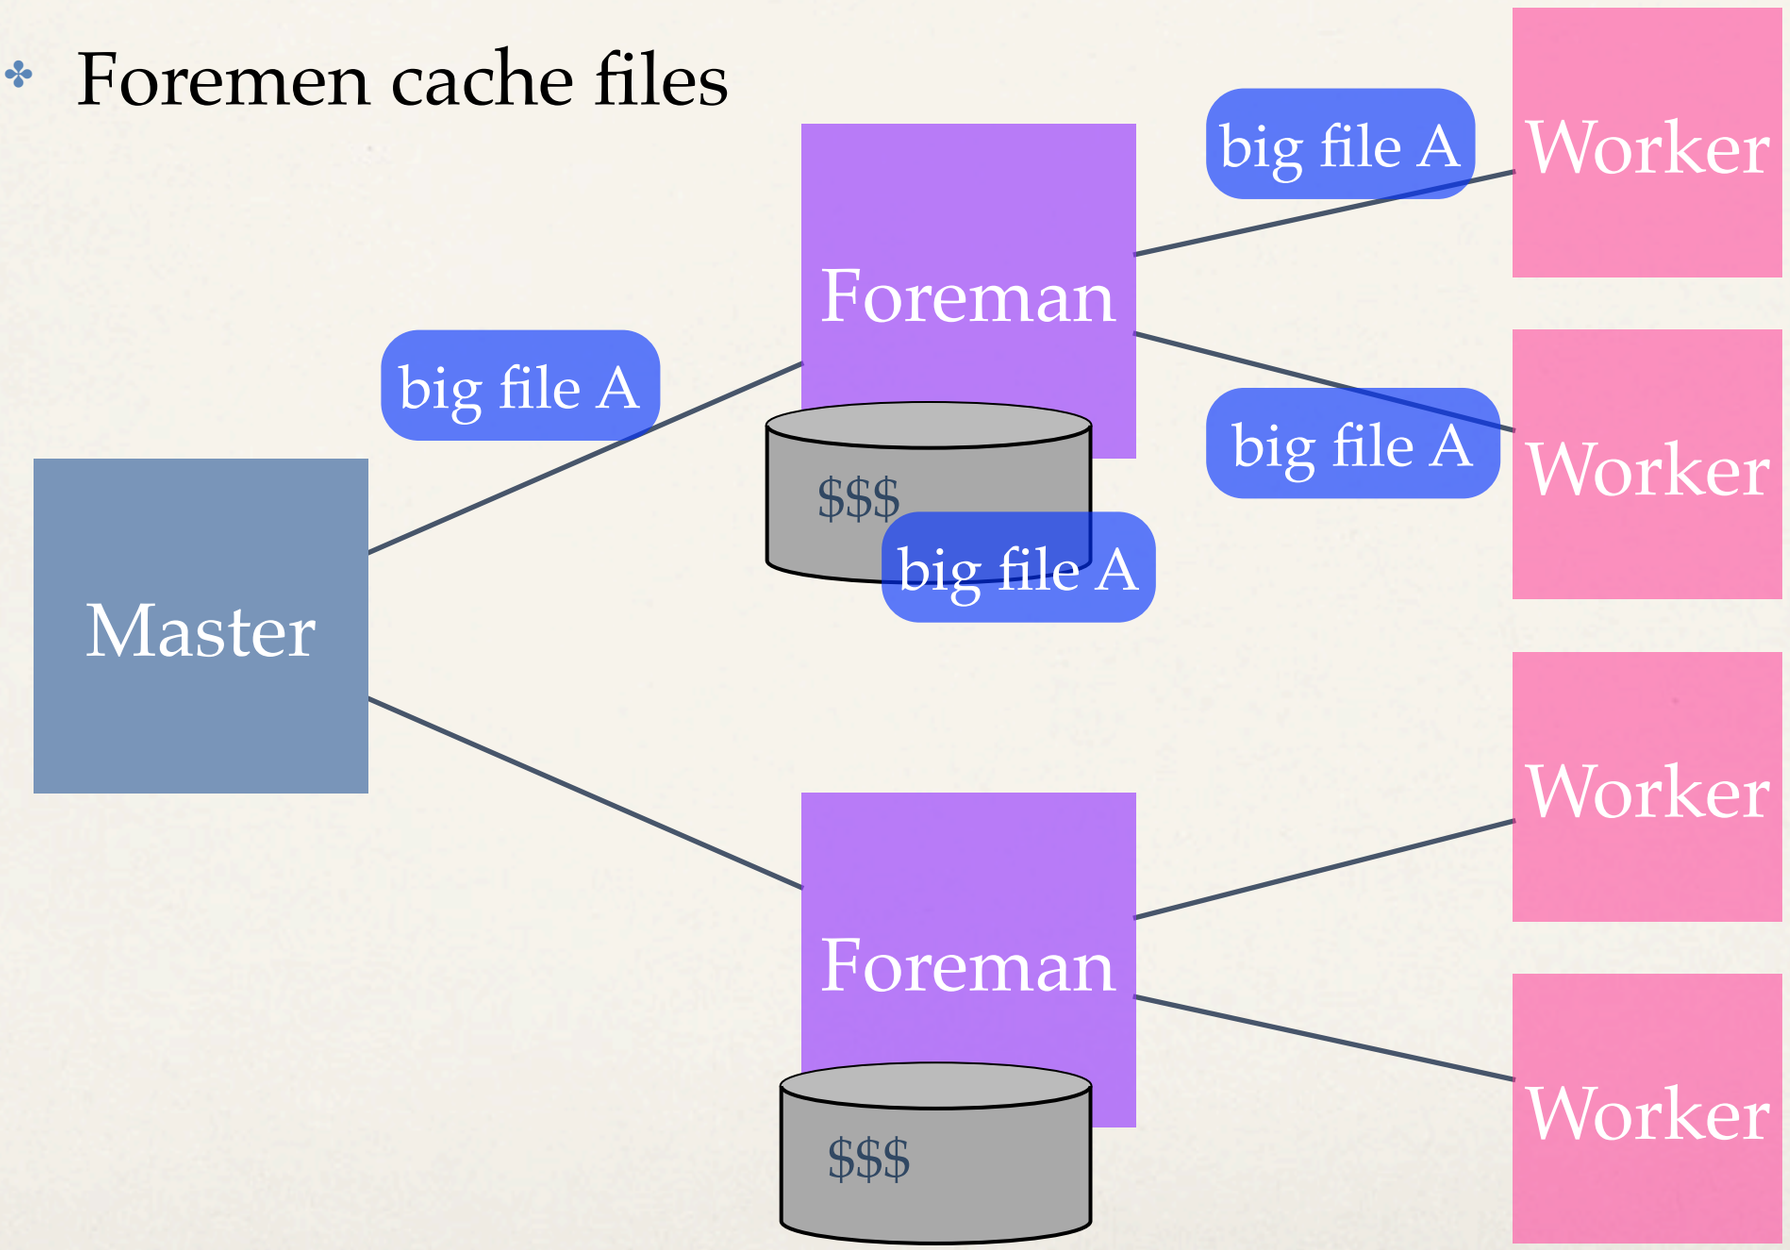

Work Queue Hierarchy

- Foremen cache files

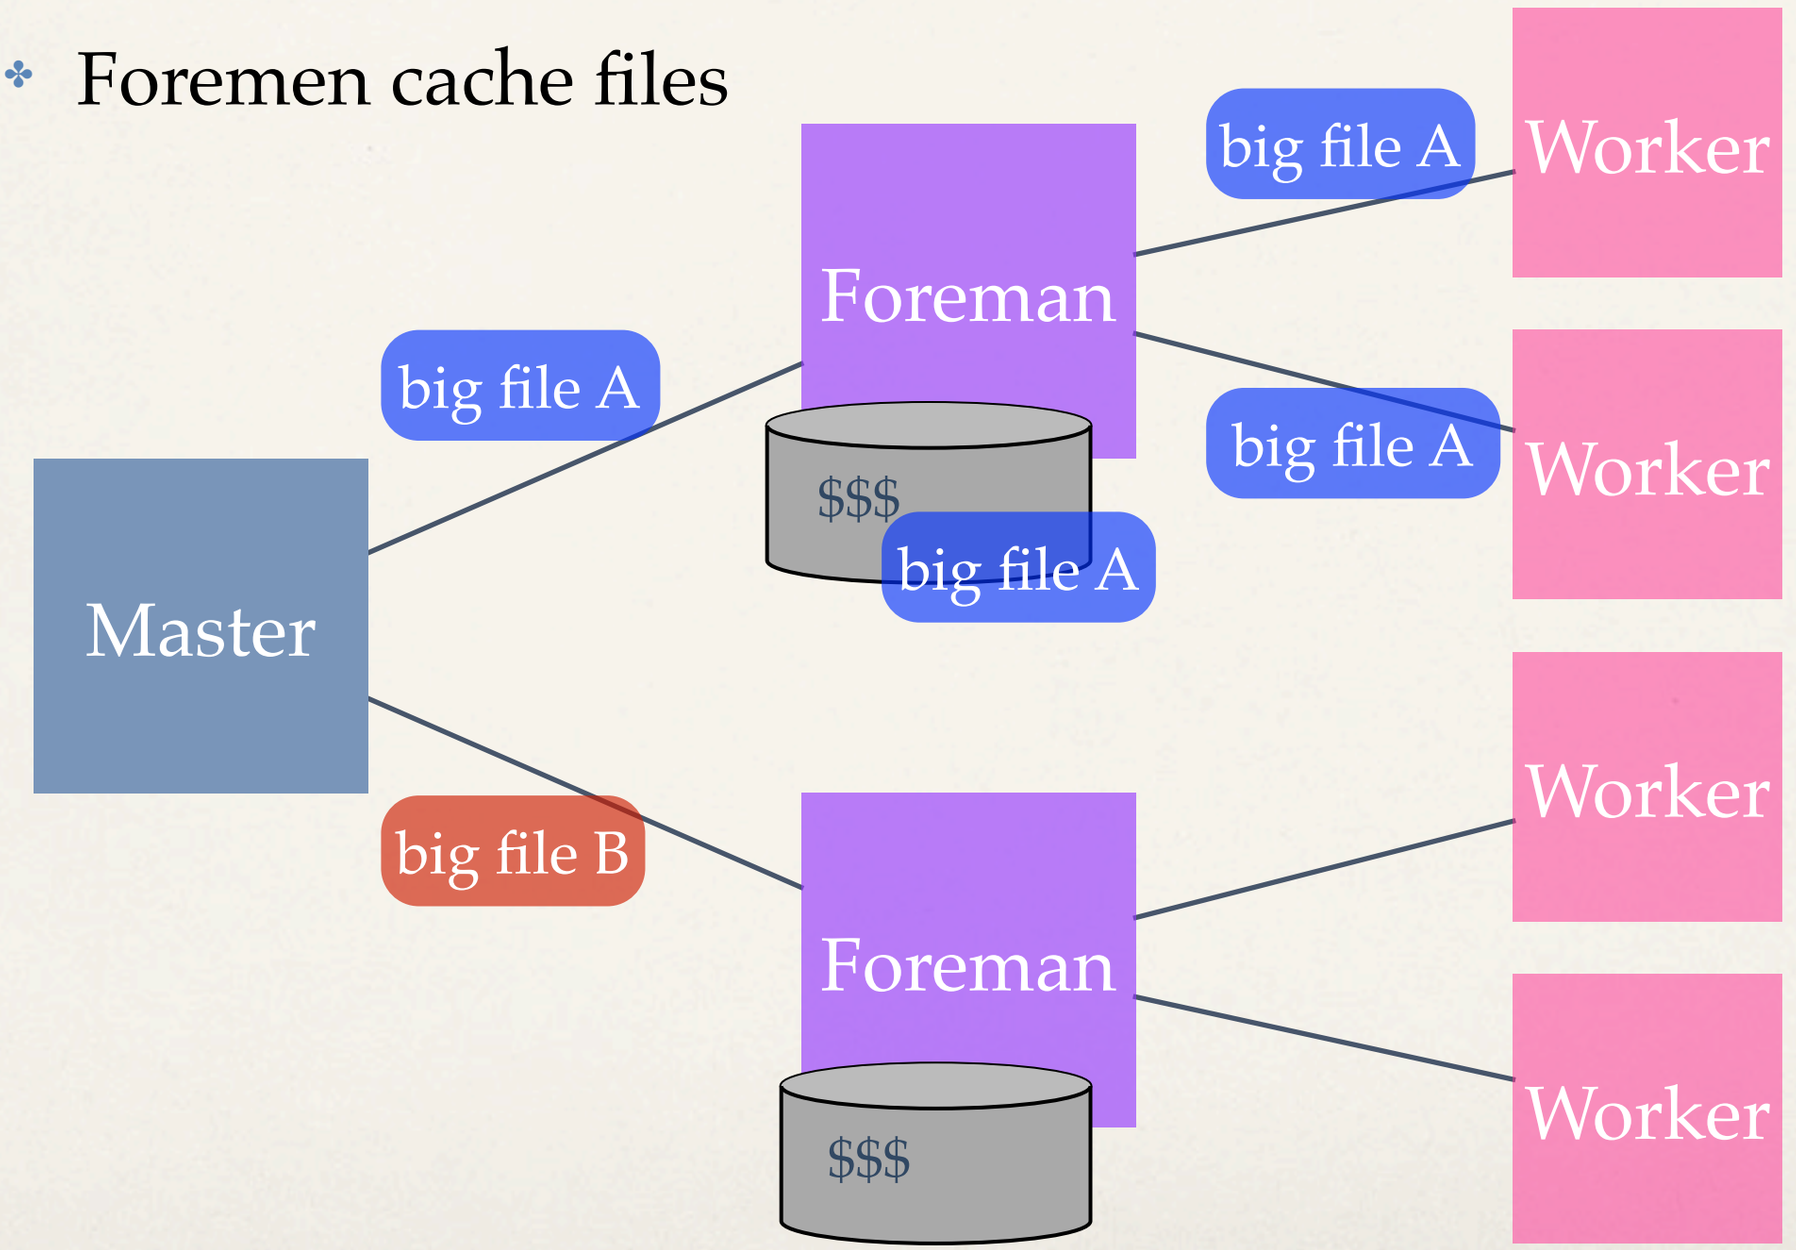

Work Queue Hierarchy

- Foremen cache files

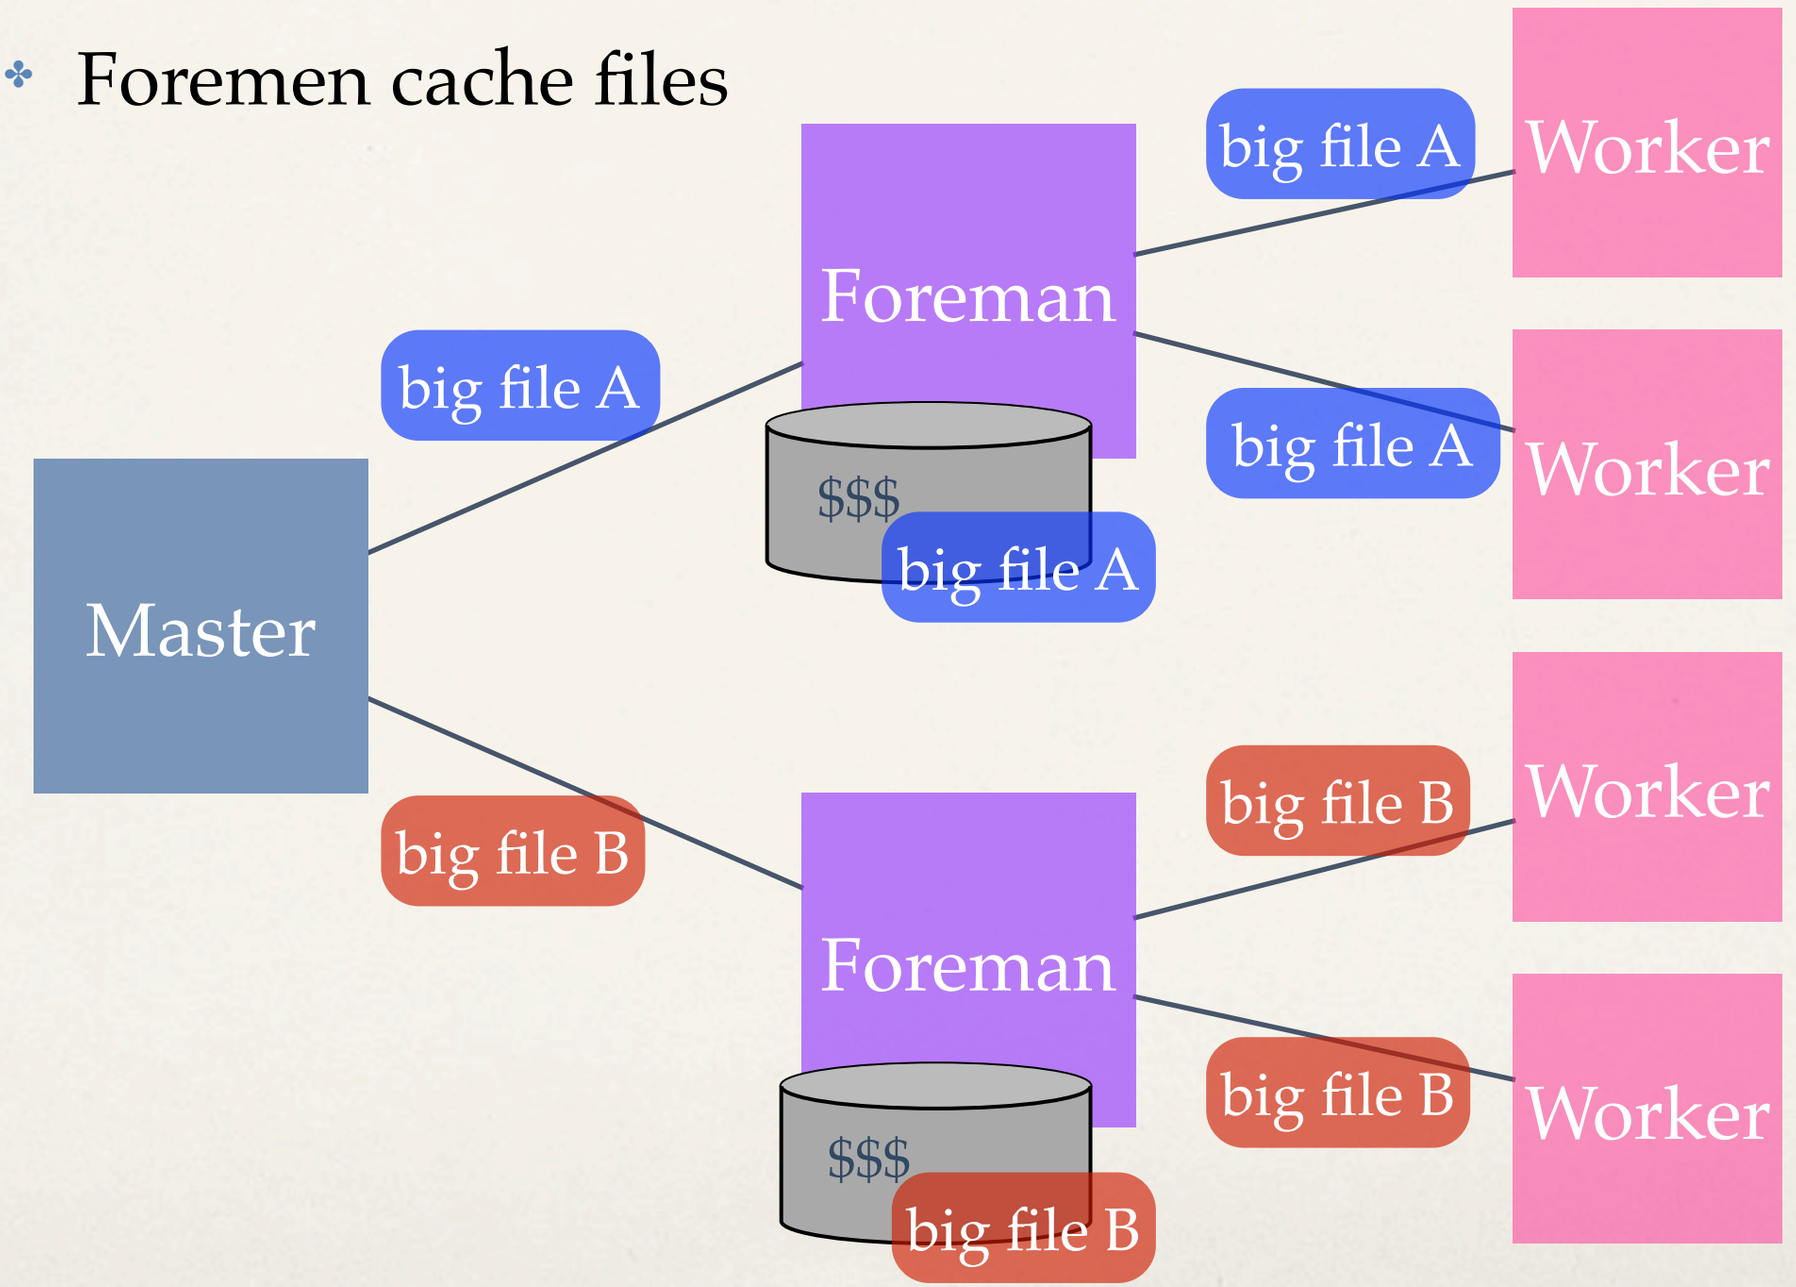

Work Queue Hierarchy

- Foreman running on head nodes

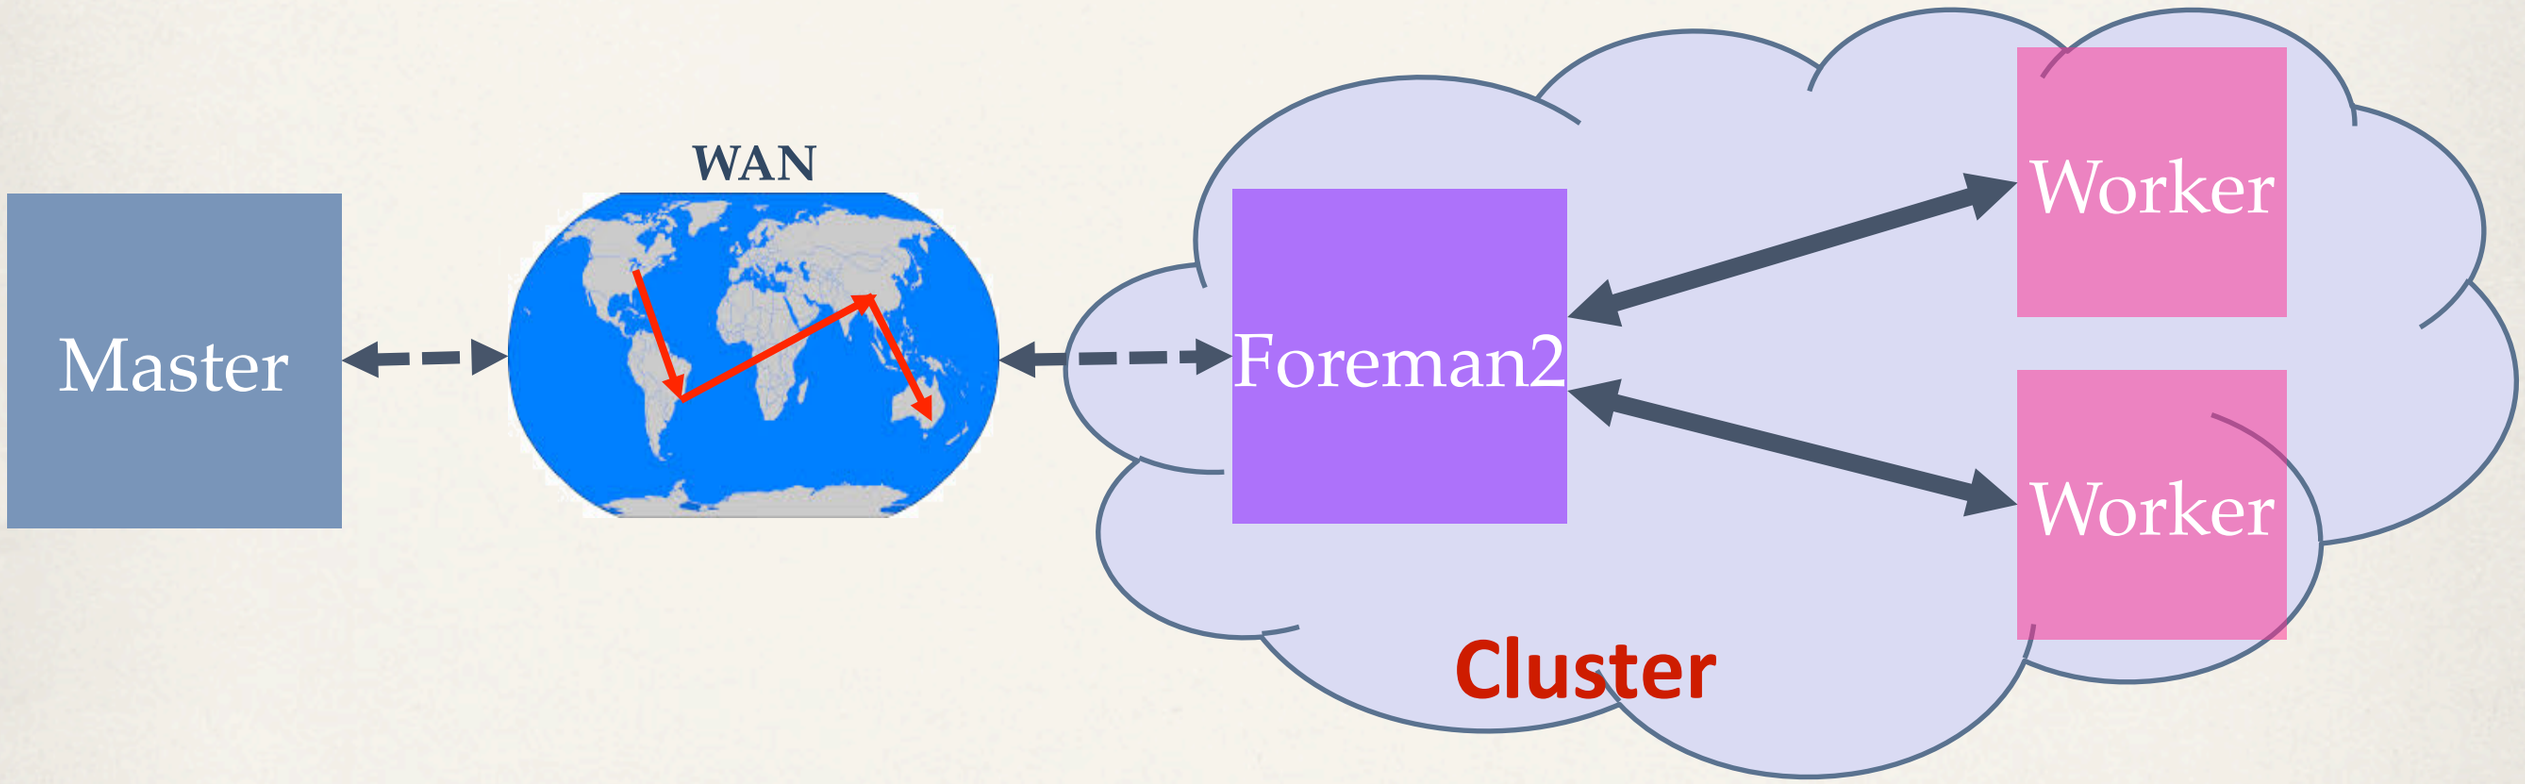

Work Queue Status

PROJECT	HOST	PORT	WAITING	BUSY	COMPLETE	WORKERS
wq_chief	cclws16.cse.nd.edu	1024	1	0	0	0
>wq-foreman-a	cclws03.cse.nd.edu	1024	0	0	0	0
->wq-foreman-a1	cclws02.cse.nd.edu	1024	0	0	0	0
->wq-foreman-a2	cclws17.cse.nd.edu	1028	0	0	0	0
>wq-foreman-b	cclws15.cse.nd.edu	1025	0	0	0	0
arab_cnvrt	crcfe01.crc.nd.edu	9551	0	88	902	213
forcebalance	fire-4-0-ext.slac.stan	9793	1418	18	43187	18
alldat_T_7	fugu.biostat.wisc.edu	9140	10	10	1	10
dihedral	leeping.Stanford.EDU	7329	524	0	77343	856
quad_cnvrt	newcell.crc.nd.edu	9552	481	153	1608	153
cclosdci8	workspace.crc.nd.edu	9907	0	0	0	0
cclosdci	workspace.crc.nd.edu	9908	0	0	354	6

Work Queue Status

PROJECT	HOST	PORT	WAITING	BUSY	COMPLETE	WORKERS
wq_chief	cclws16.cse.nd.edu	1024	1	0	0	0
>wq-foreman-a	cclws03.cse.nd.edu	1024	0	0	0	0
->wq-foreman-a1	cclws02.cse.nd.edu	1024	0	0	0	0
->wq-foreman-a2	cclws17.cse.nd.edu	1028	0	0	0	0
>wq-foreman-b	cclws15.cse.nd.edu	1025	0	0	0	0
arab_cnvrt	crcfe01.crc.nd.edu	9551	0	88	902	213
forcebalance	fire-4-0-ext.slac.stan	9793	1418	18	43187	18
alldat_T_7	fugu.biostat.wisc.edu	9140	10	10	1	10
dihedral	leeping.Stanford.EDU	7329	524	0	77343	856
quad_cnvrt	newcell.crc.nd.edu	9552	481	153	1608	153
cclosdci8	workspace.crc.nd.edu	9907	0	0	0	0
cclosdci	workspace.crc.nd.edu	9908	0	0	354	6

Work Queue Multi-slot Workers

Work Queue Multi-slot Workers

BEFORE

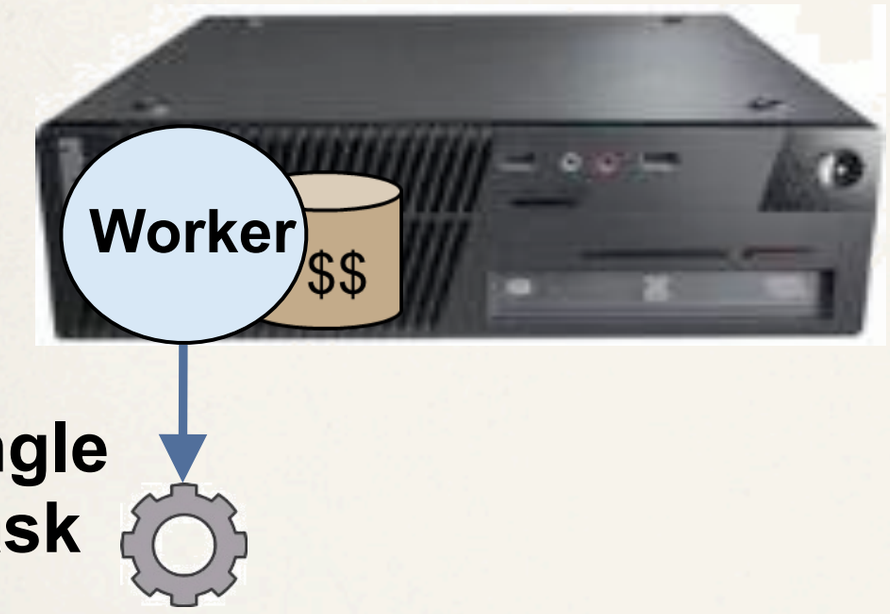

Work Queue Multi-slot Workers

BEFORE

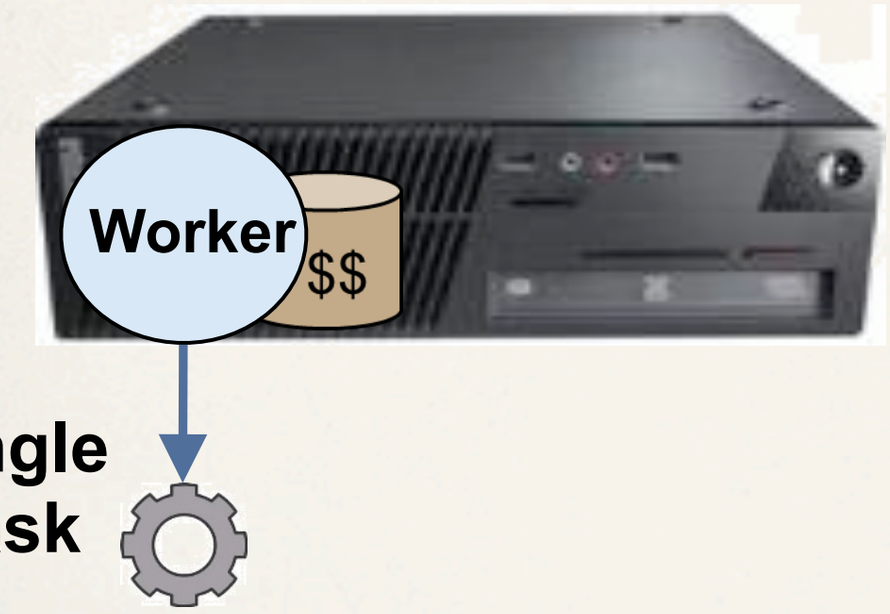

NOW

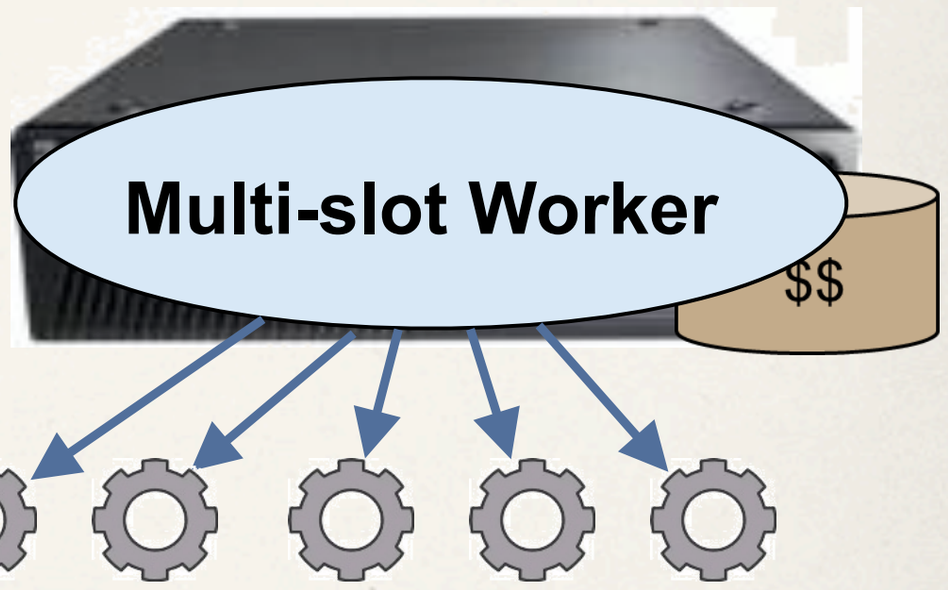

Work Queue Multi-slot Workers

BEFORE

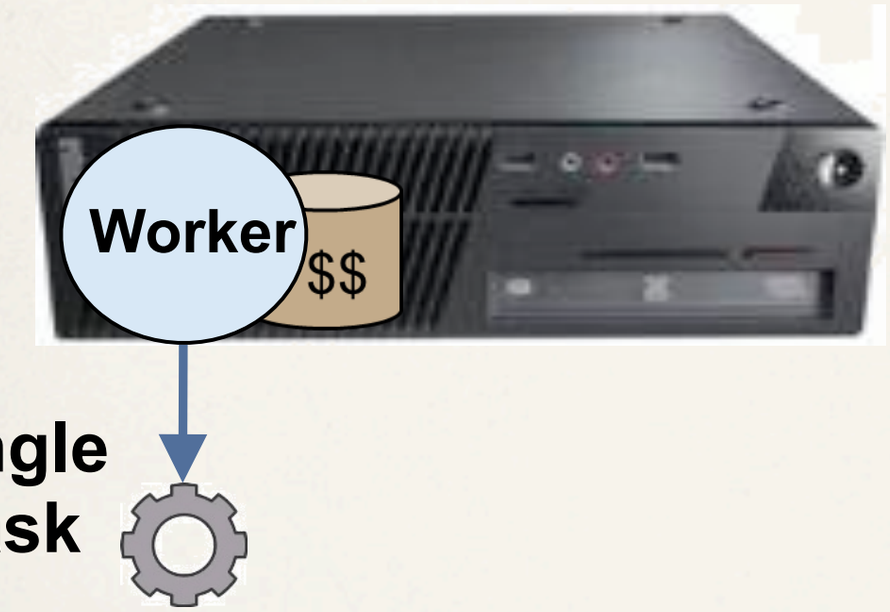

NOW

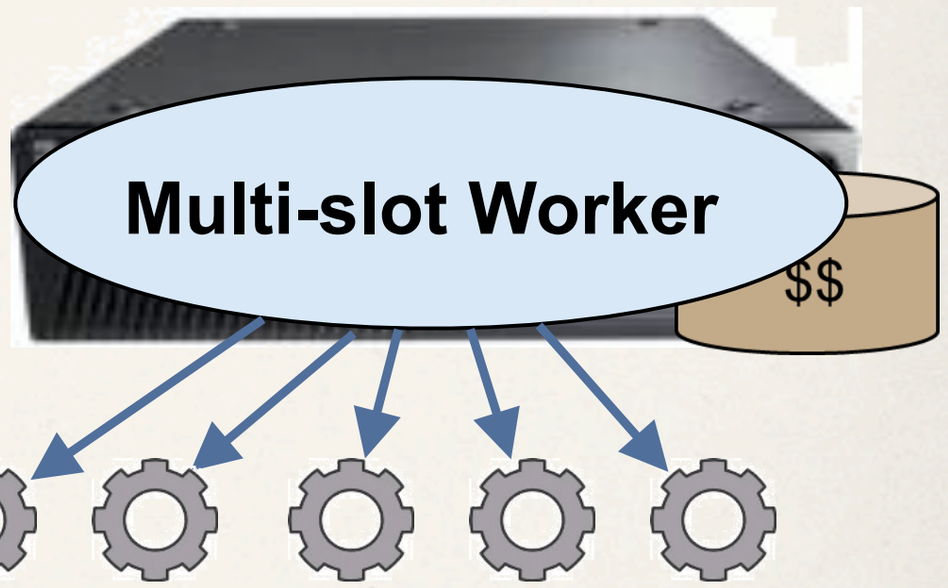

1. Worker can have multiple slots for running tasks depending on available resources
2. Simultaneously runs multiple tasks if they fit within available resources
3. Tasks must be specified with their resource requirements

Work Queue Mutli-slot Workers

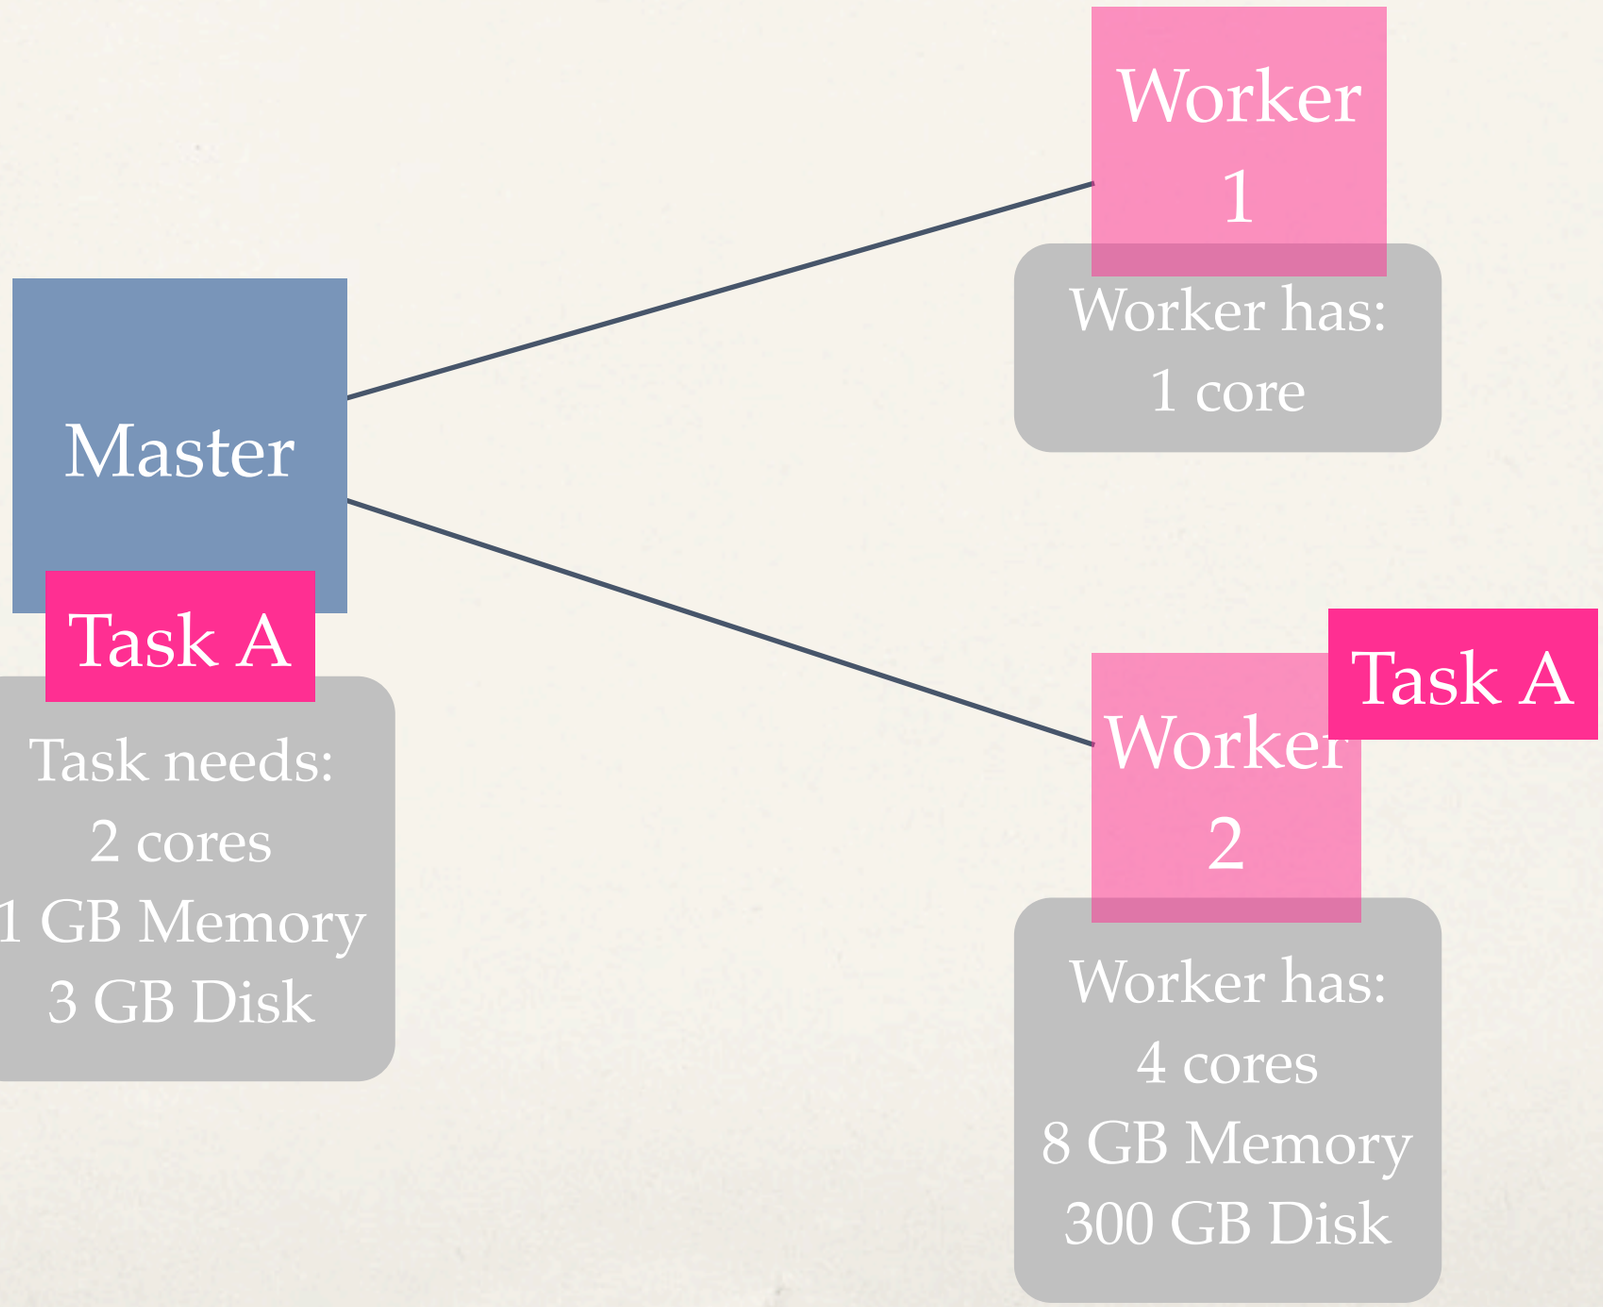

Work Queue Mutli-slot Workers

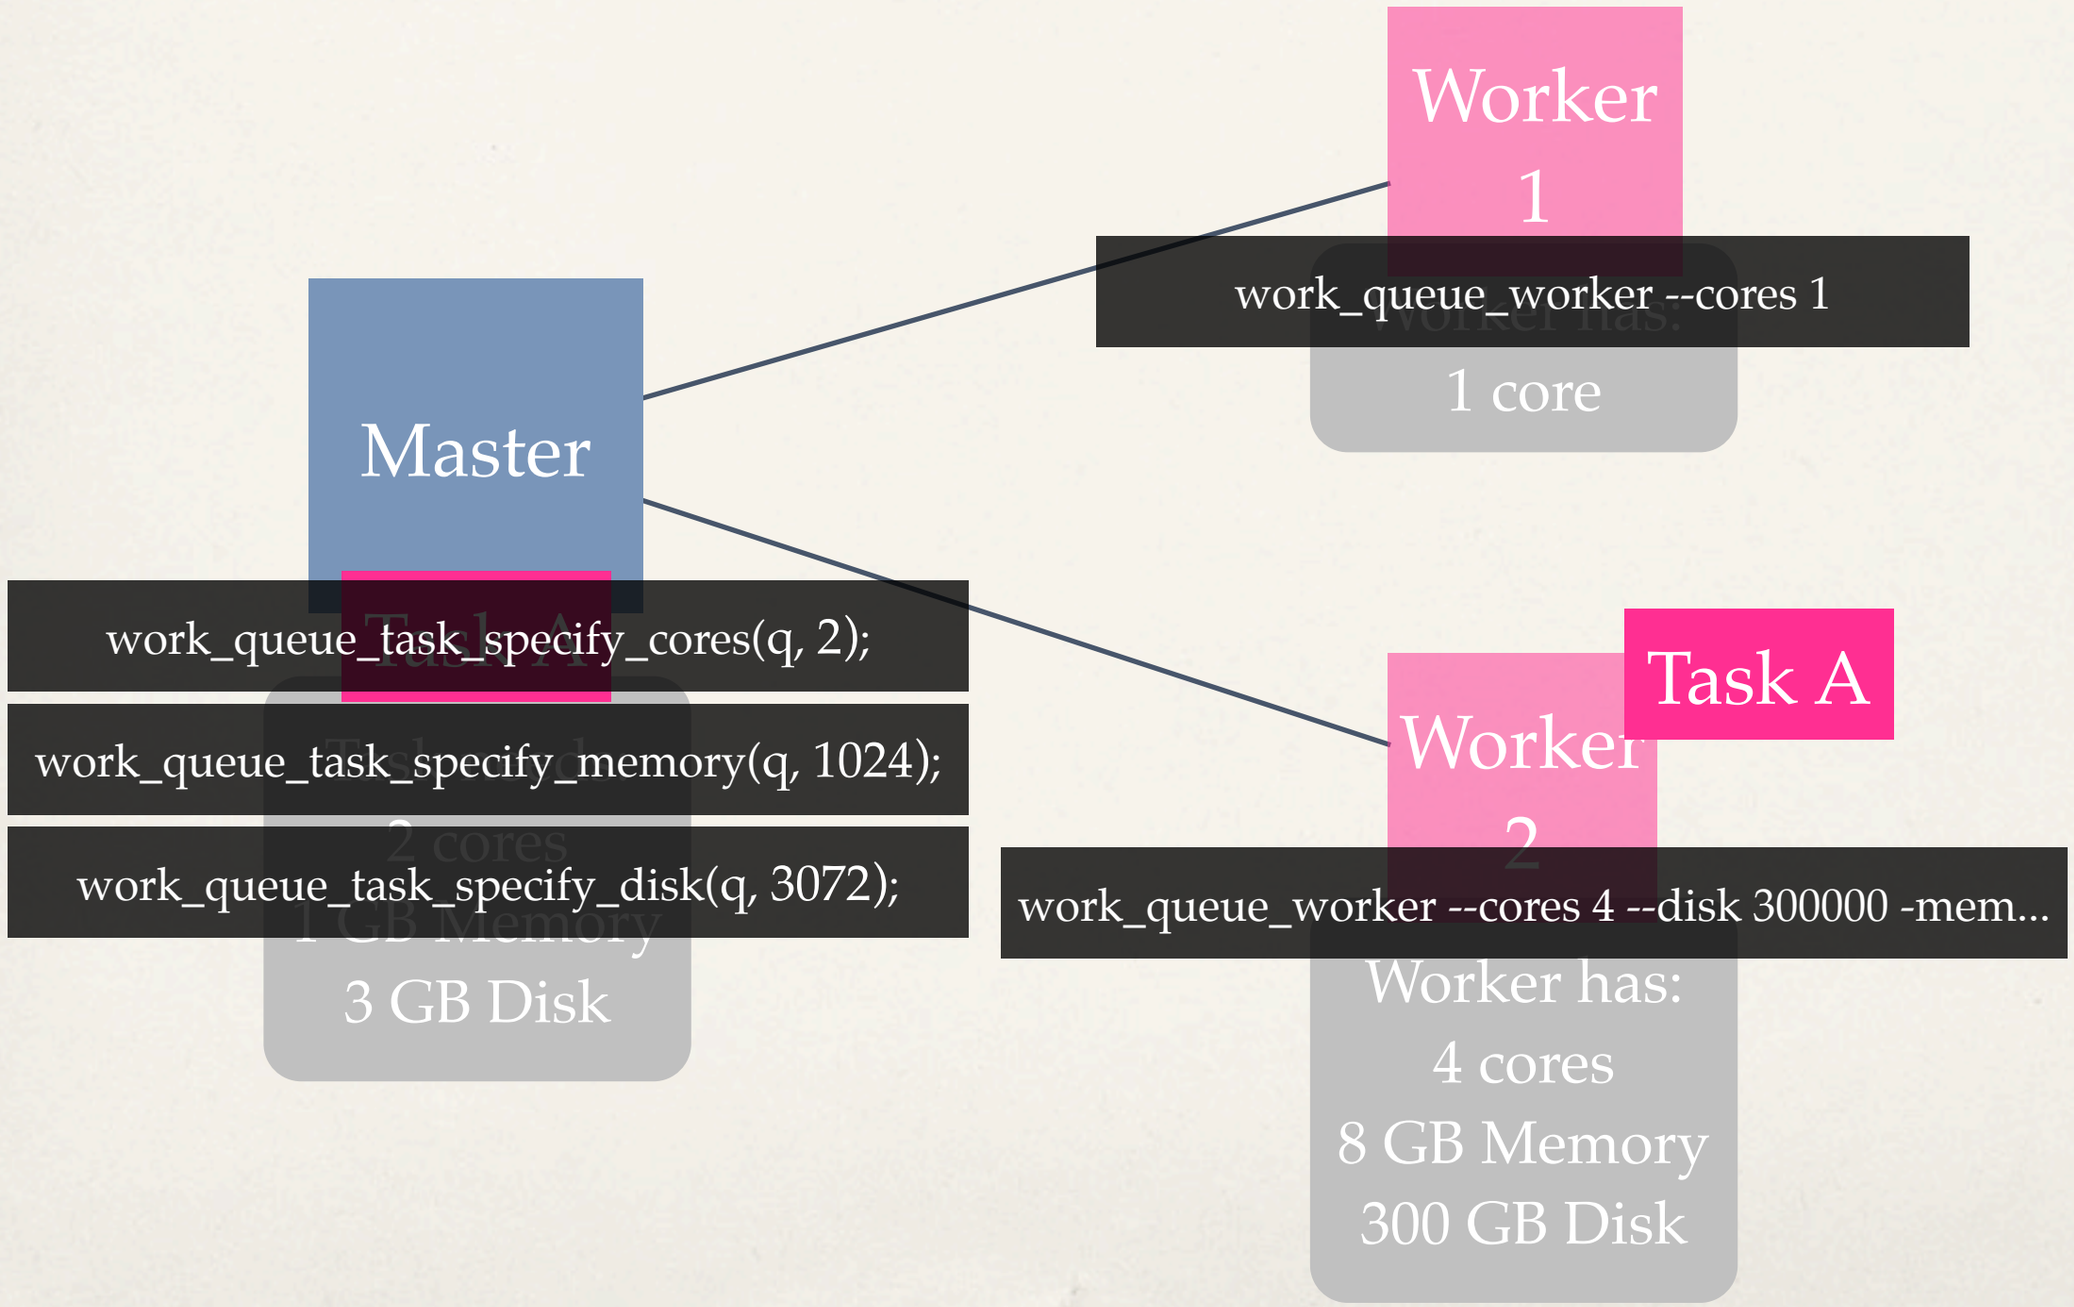

Work Queue Multi-slot Workers

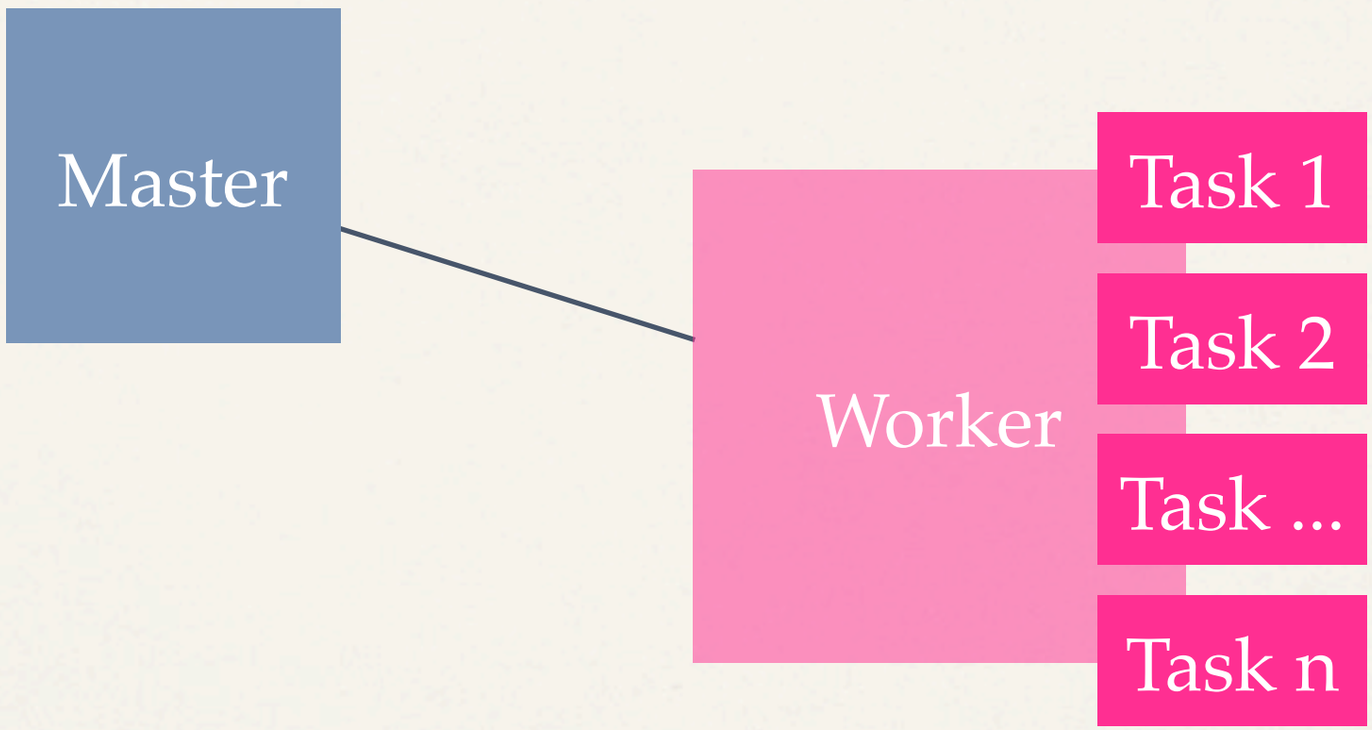

Work Queue Multi-slot Workers

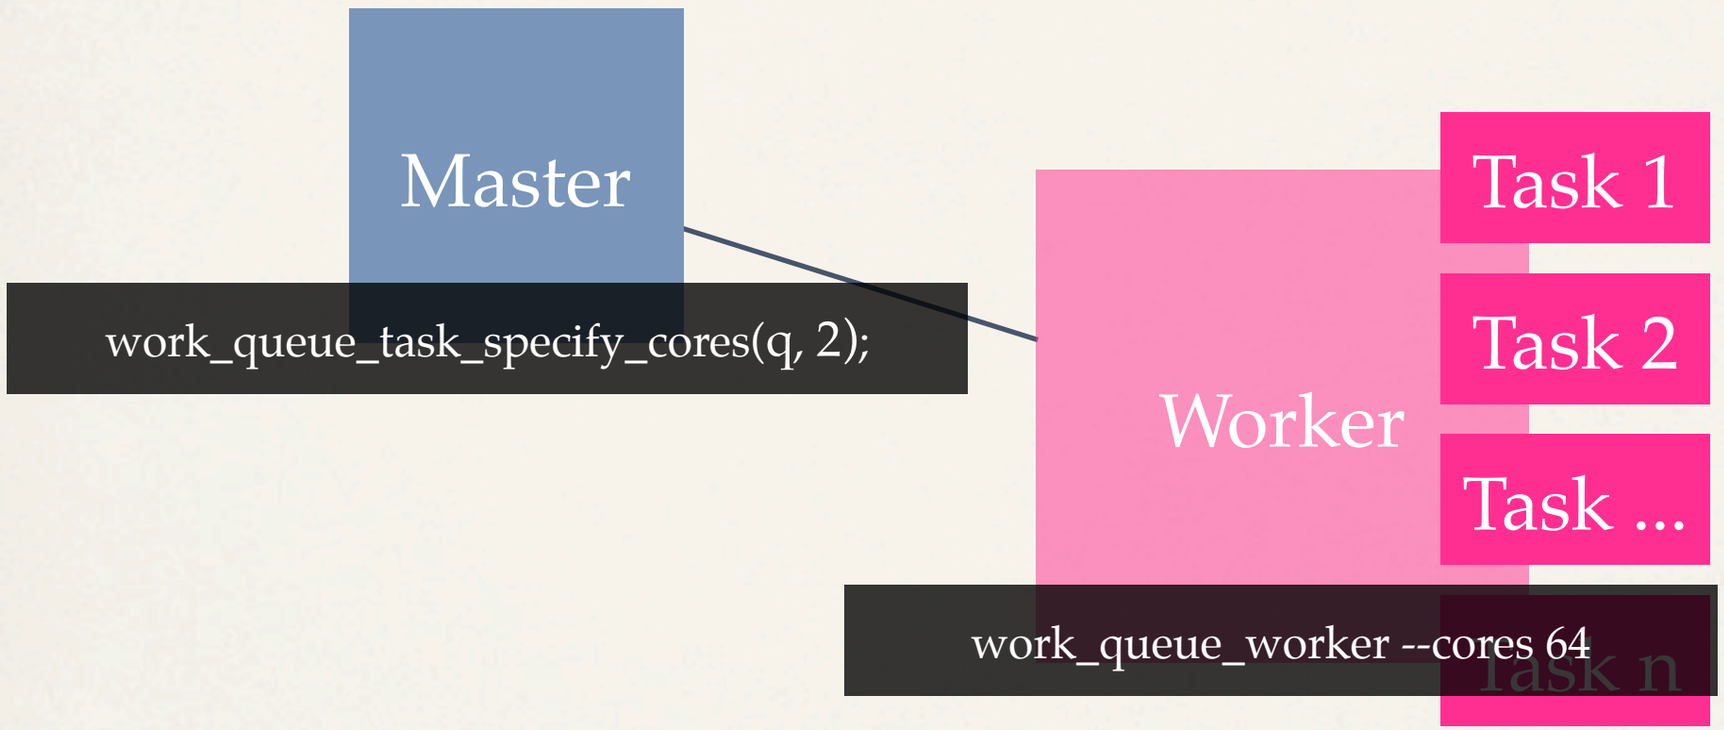

Work Queue Status

work_queue_status -R

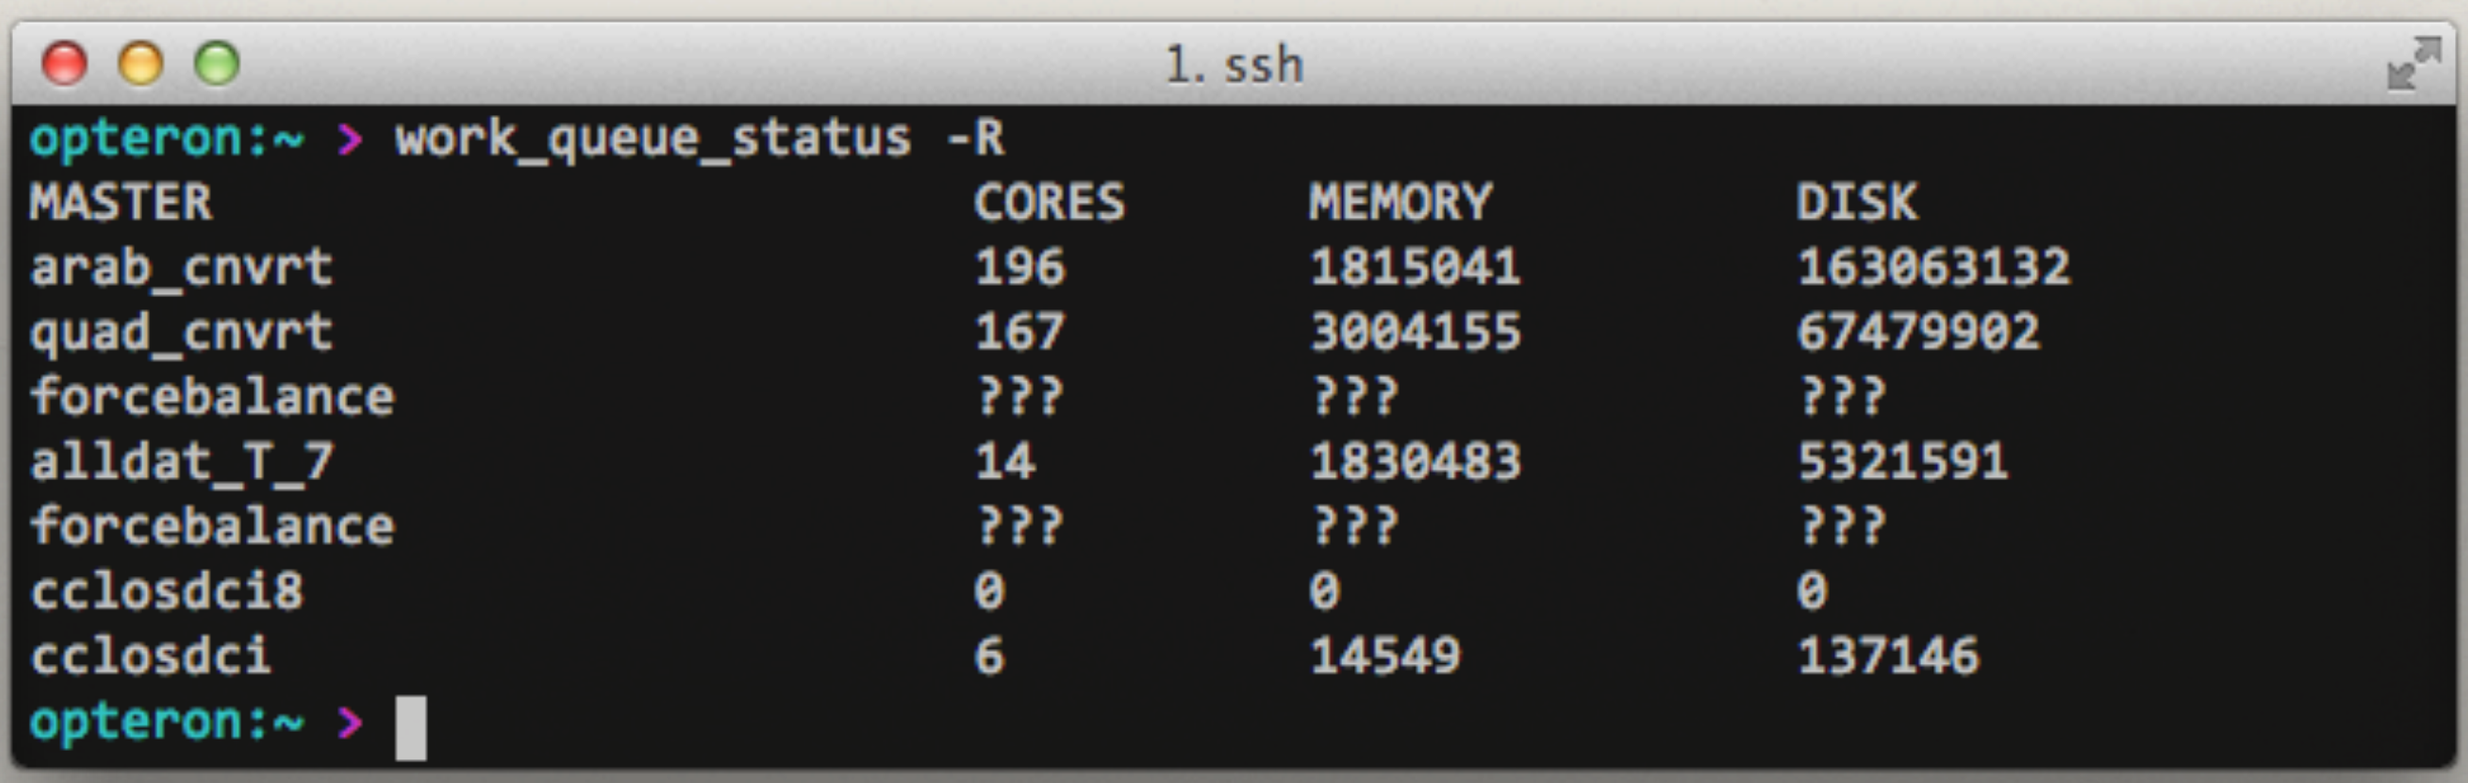

WQ Using Project Names

WQ Using Project Names

“myproject”

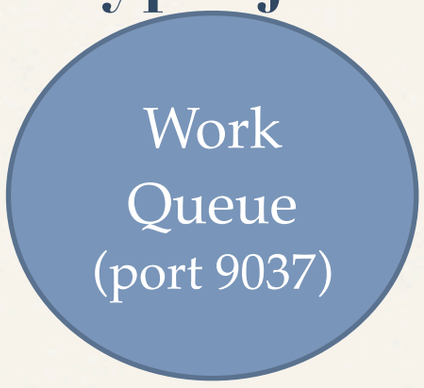

WQ Using Project Names

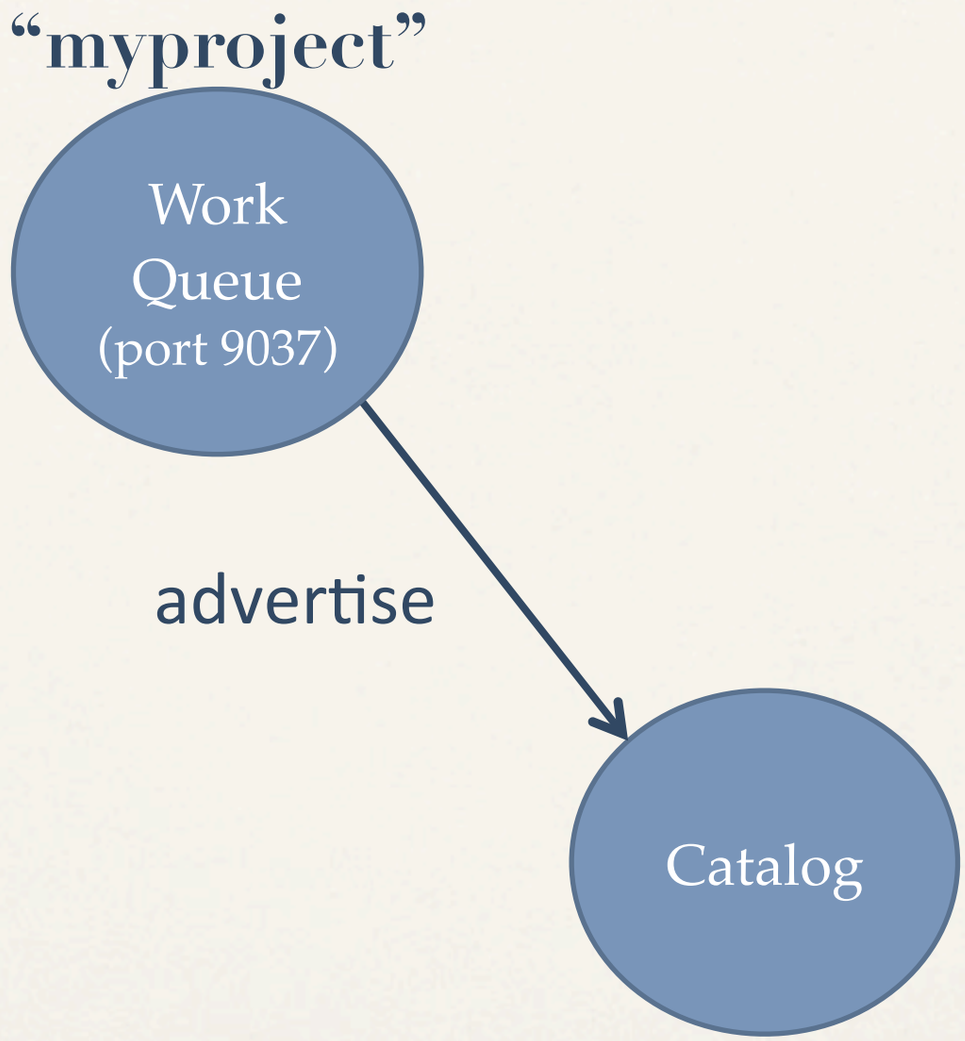

WQ Using Project Names

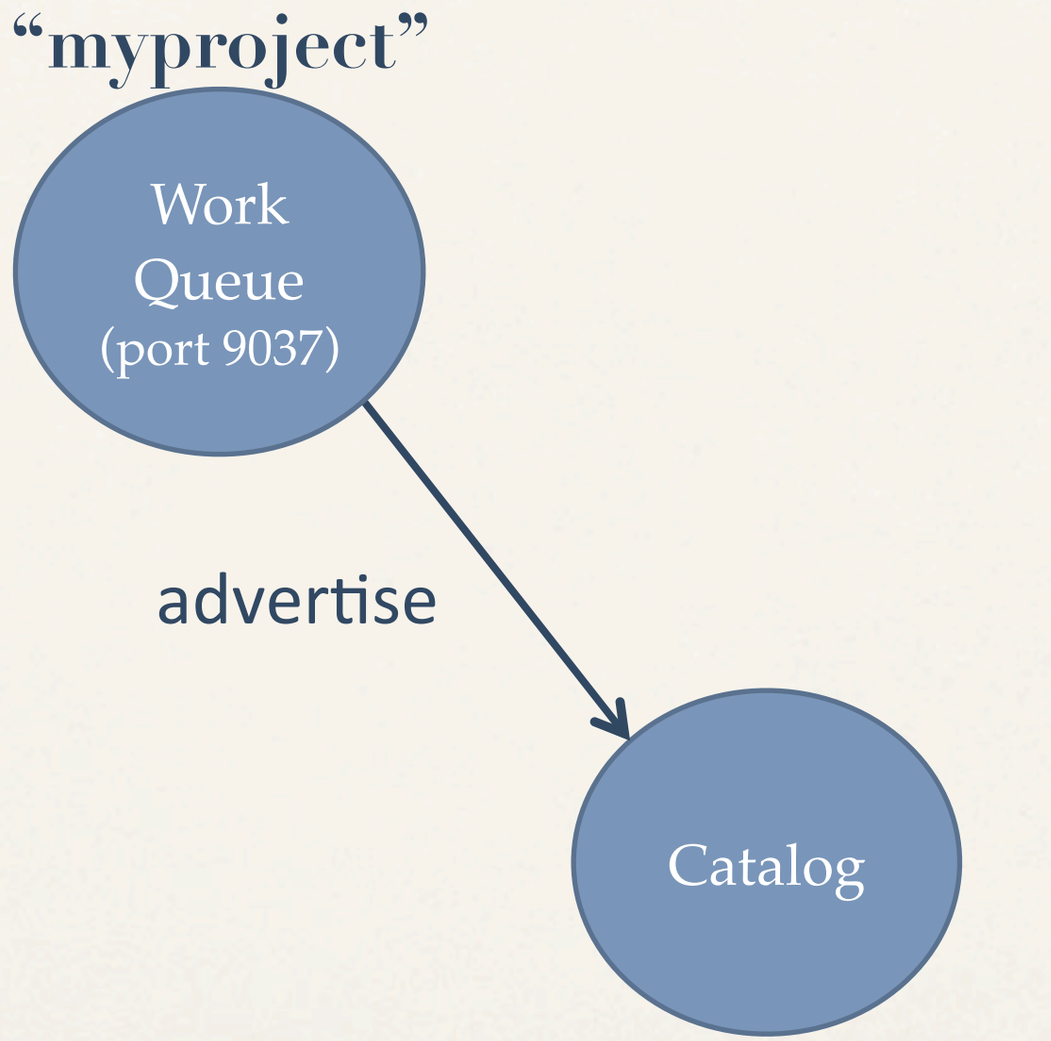

WQ Using Project Names

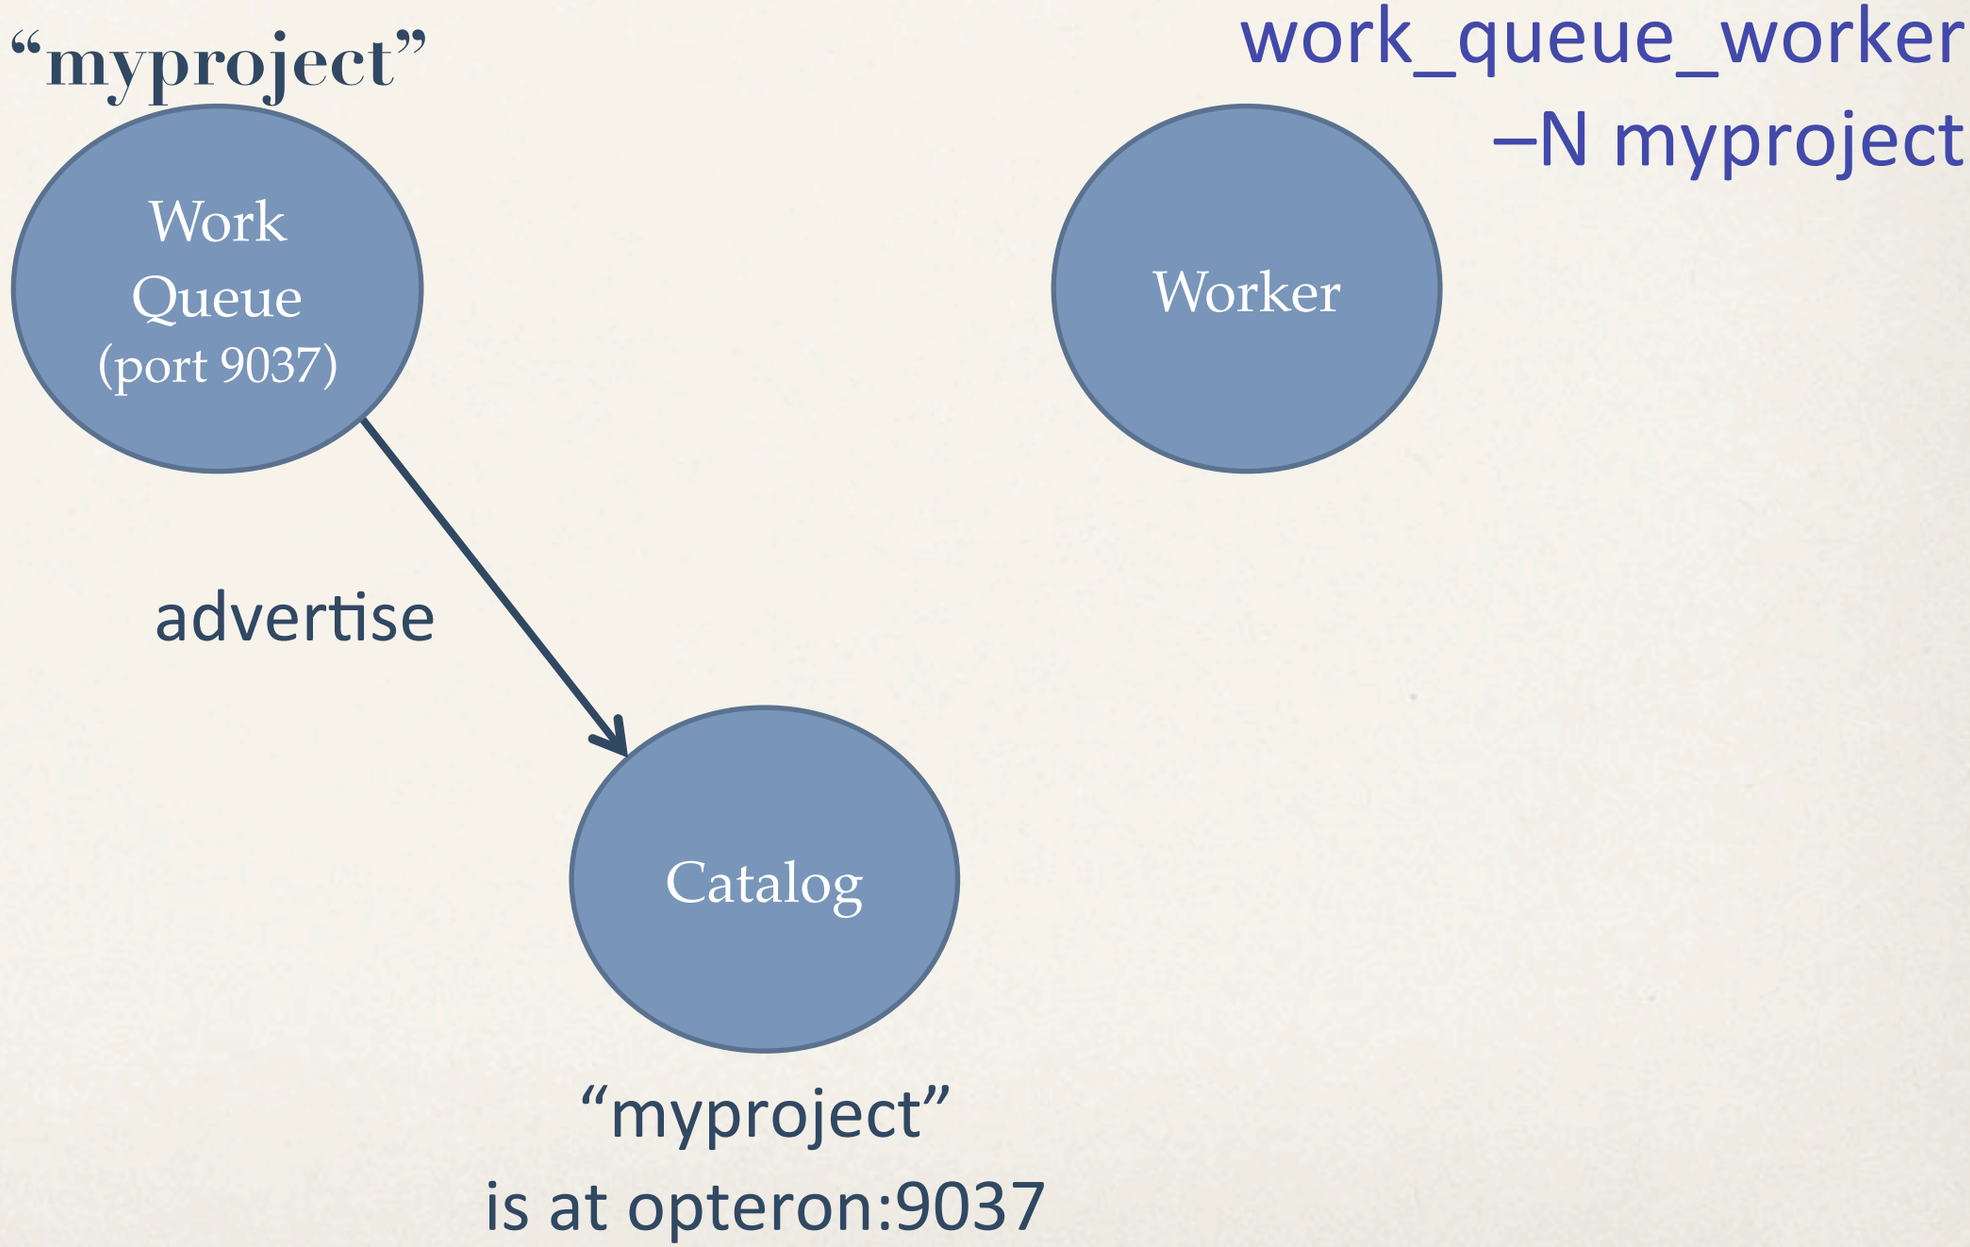

WQ Using Project Names

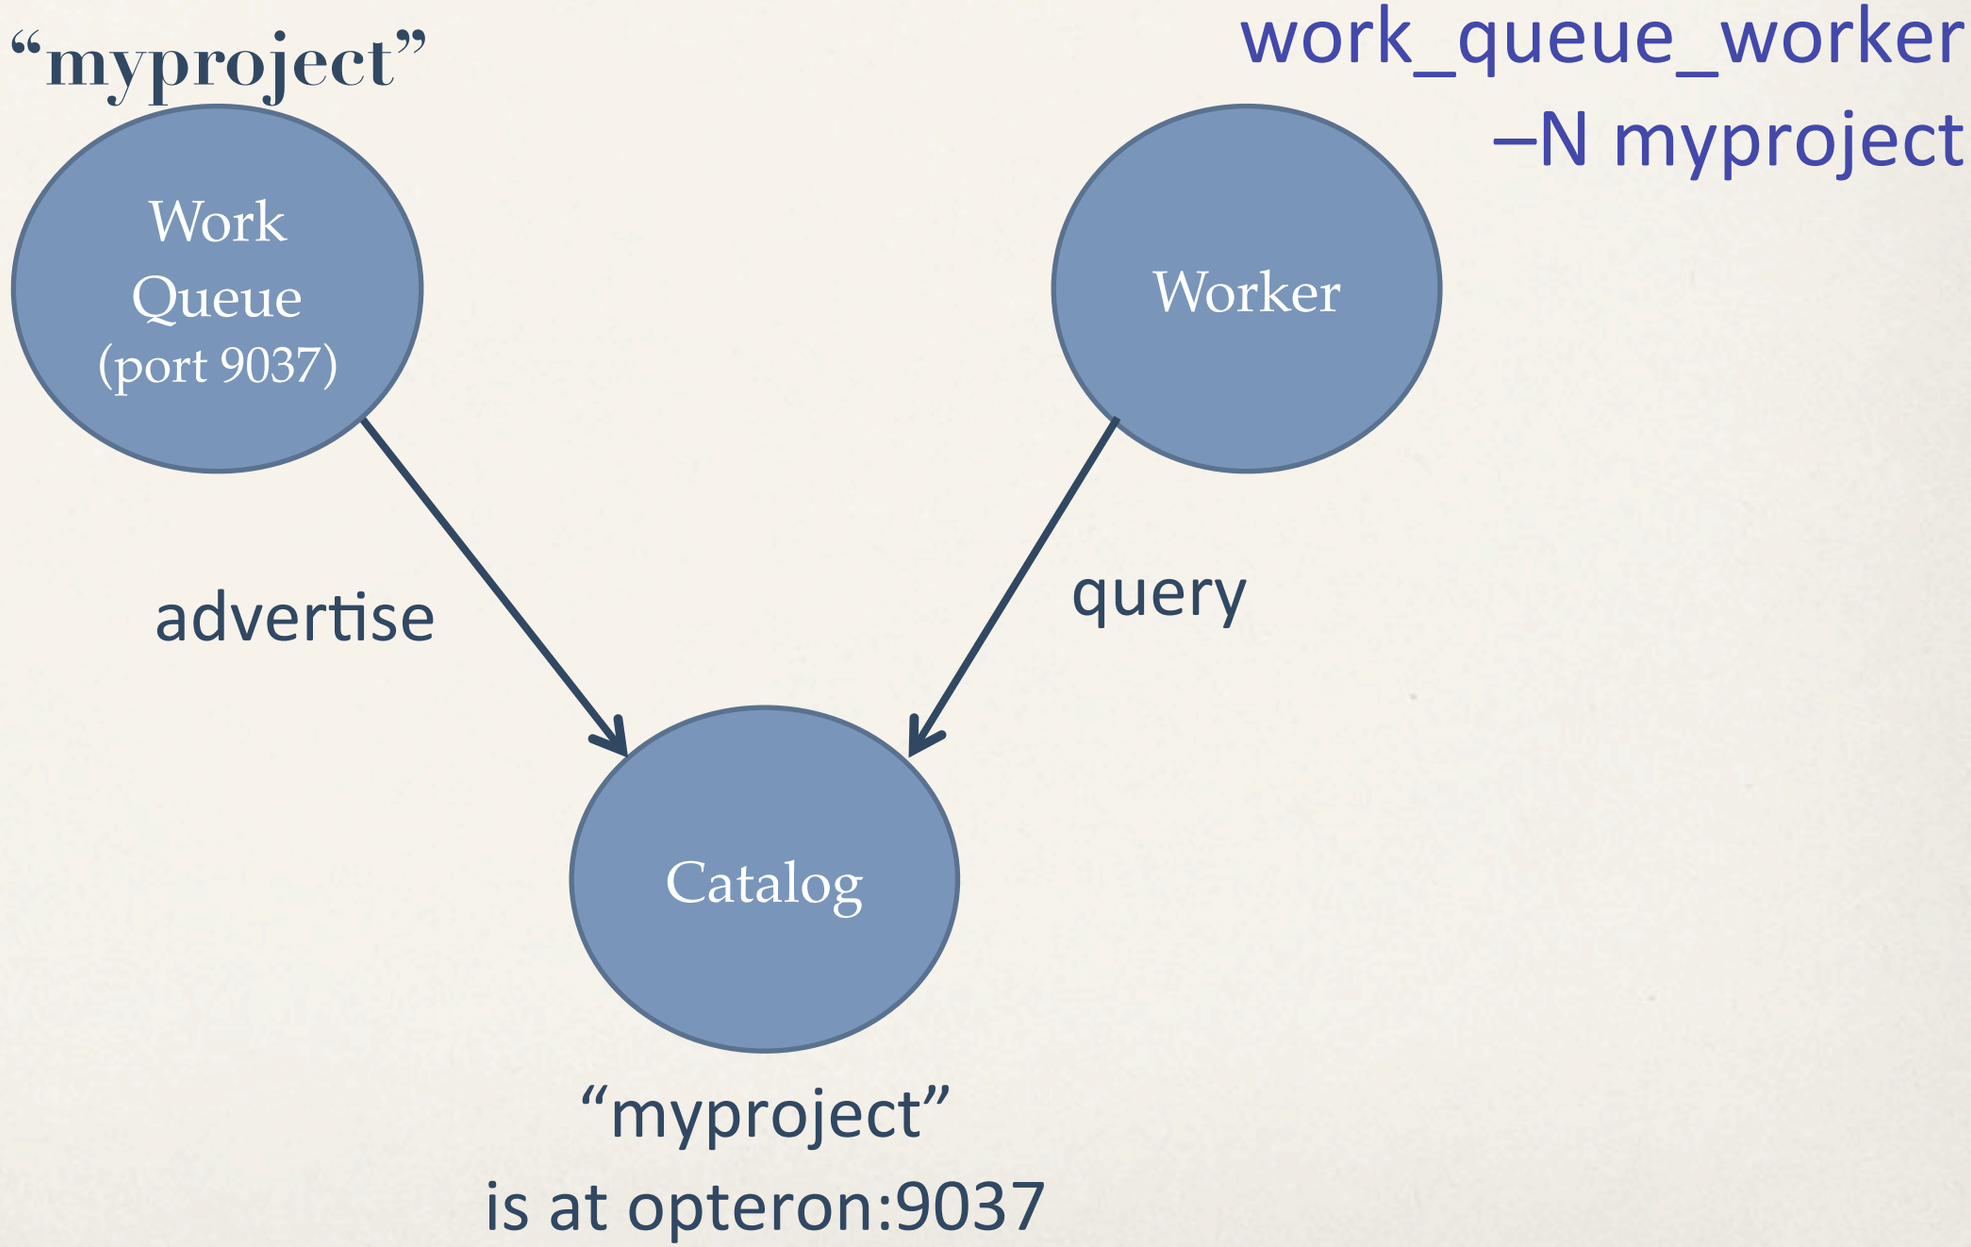

WQ Using Project Names

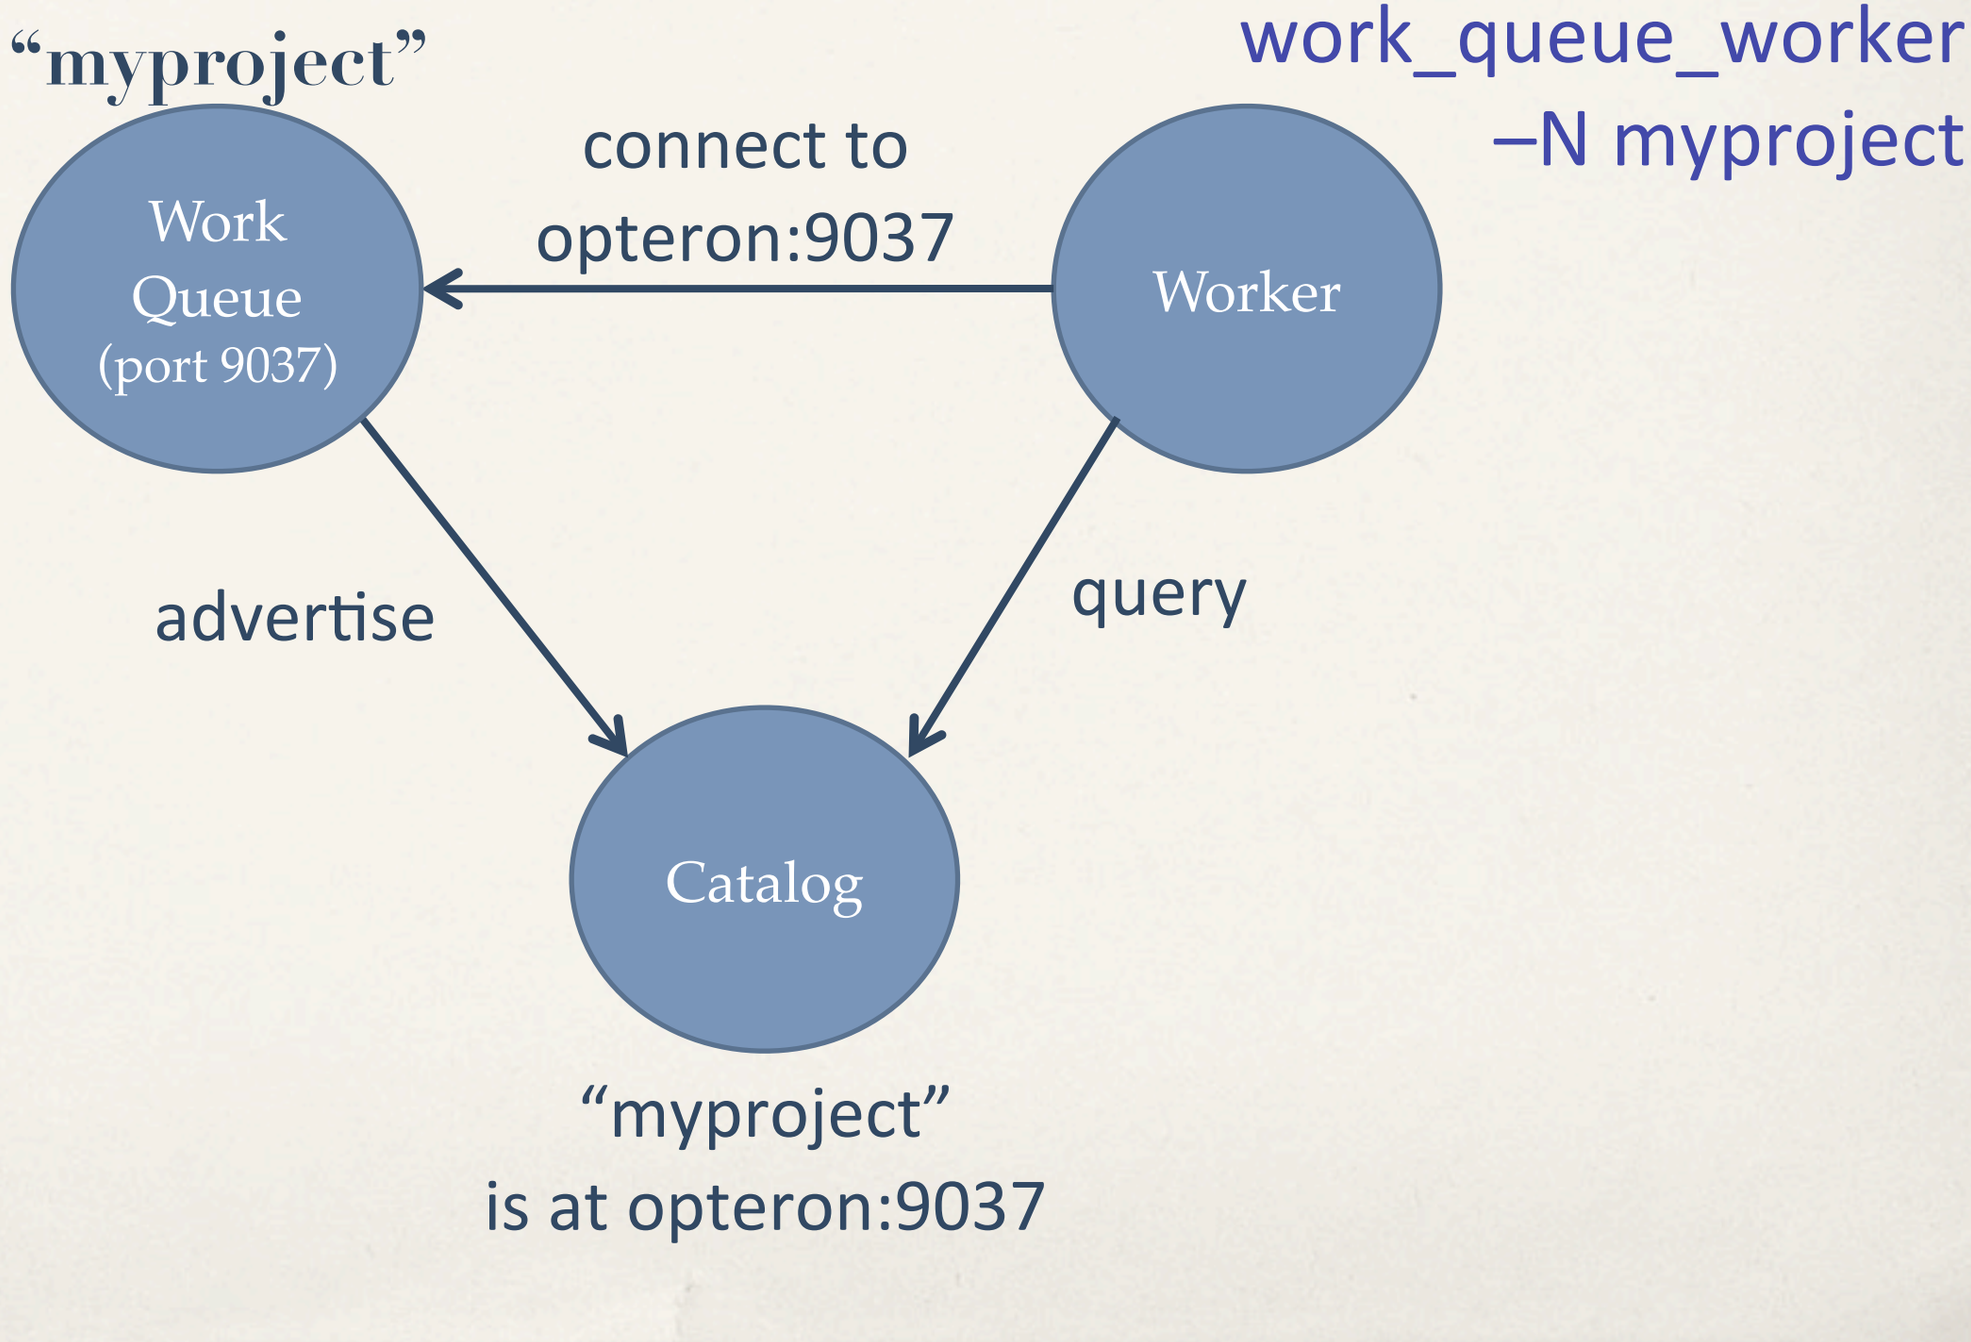

WQ Using Project Names

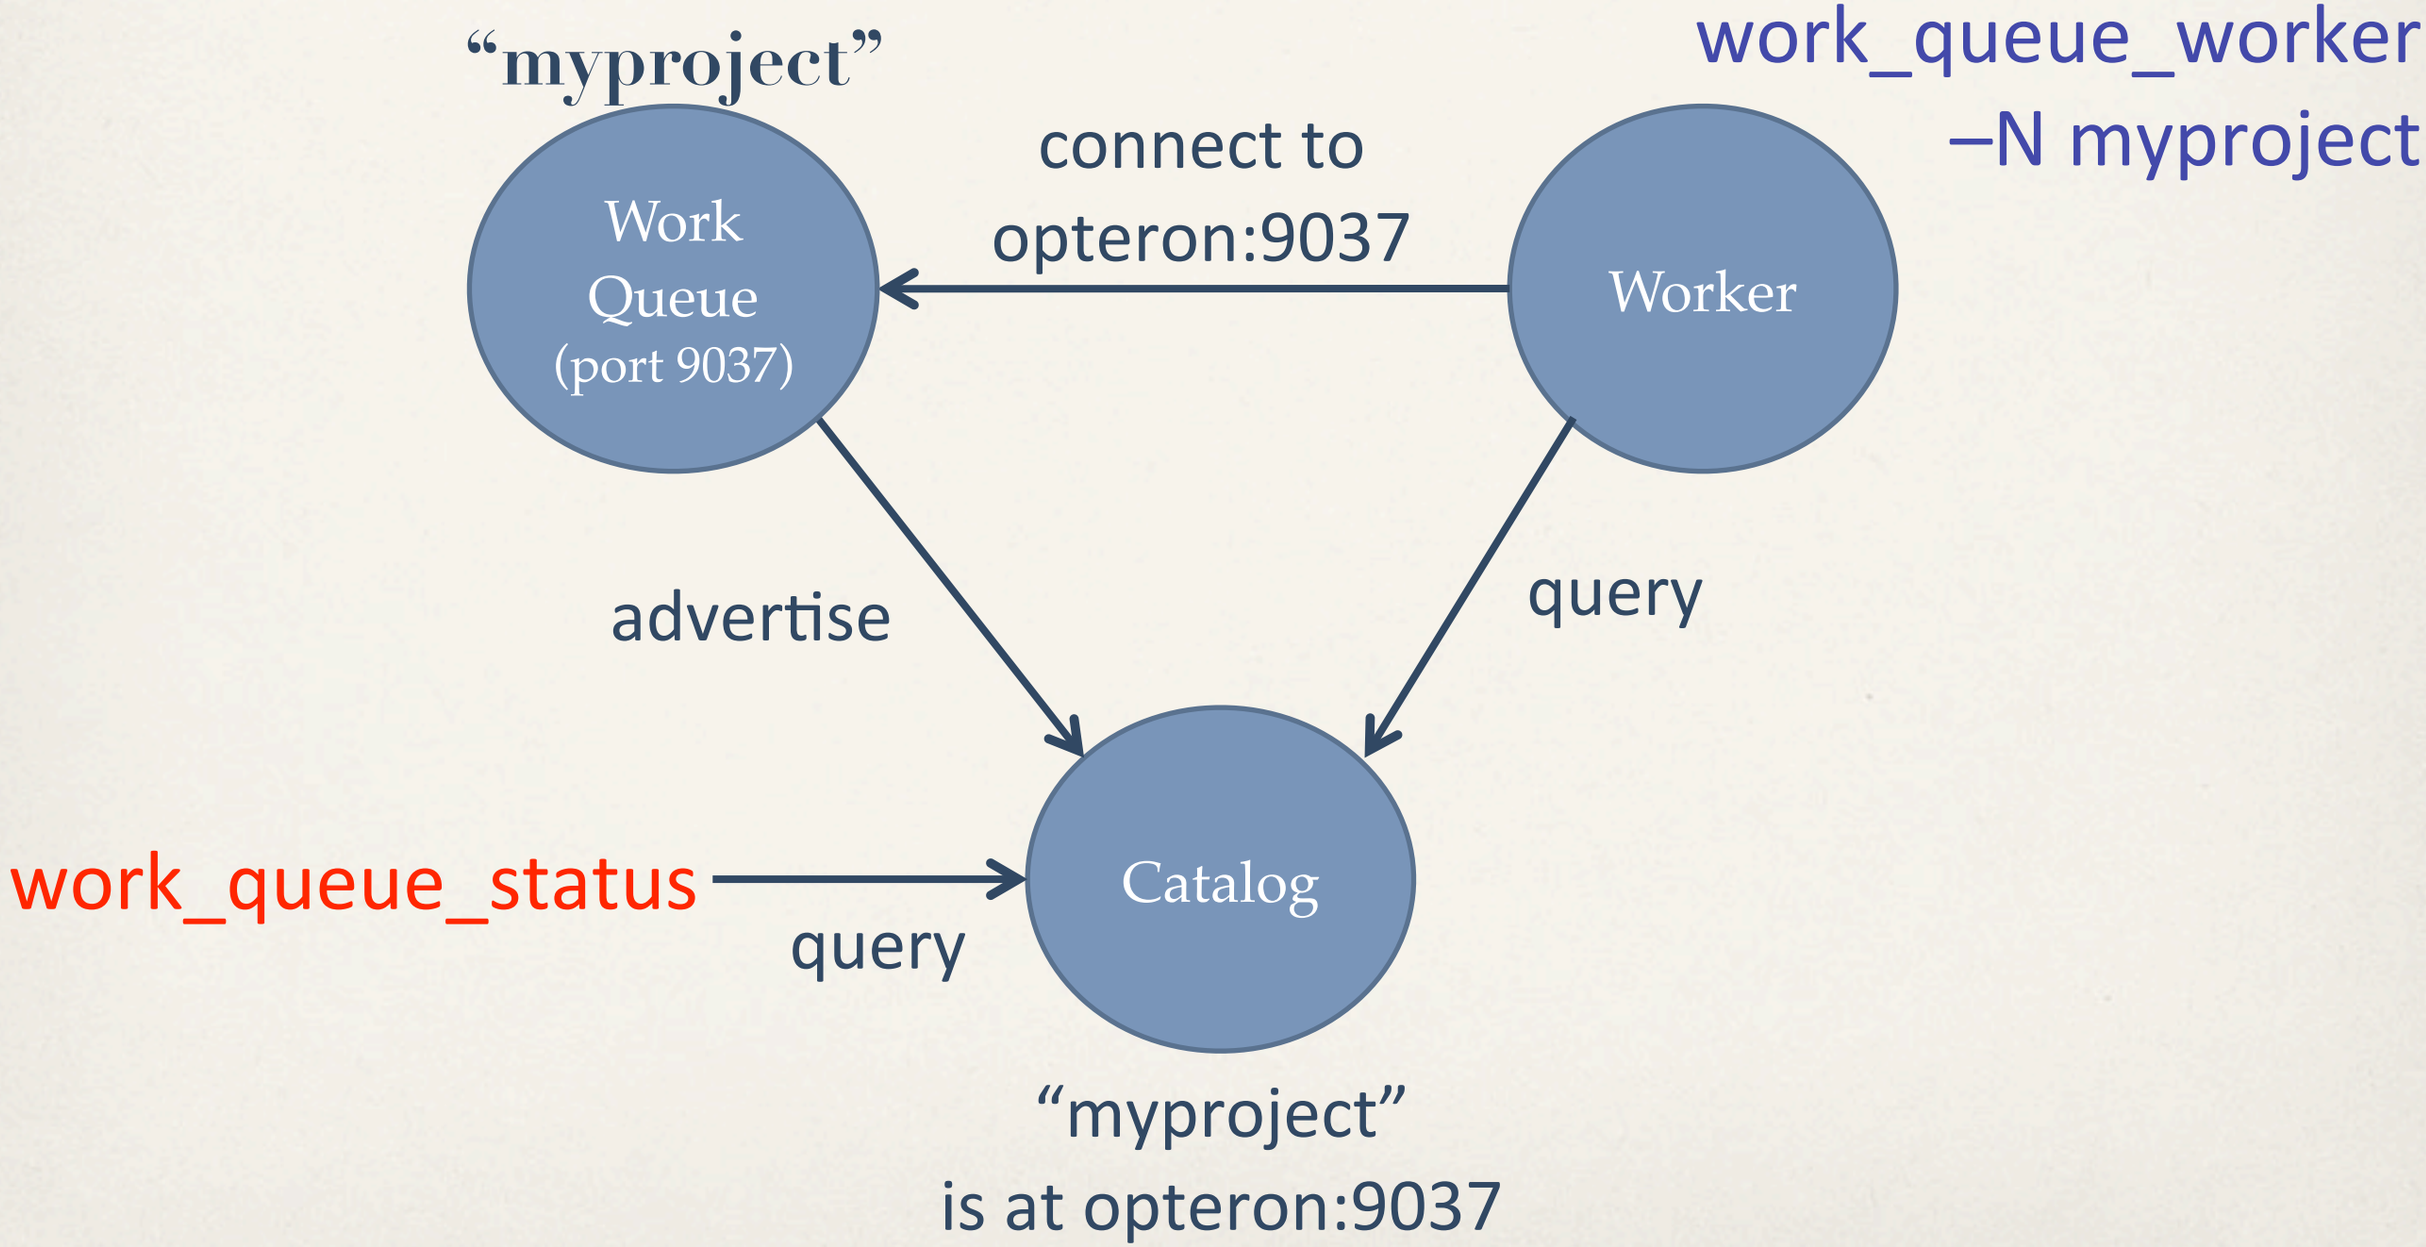

Find master location in work_queue_status

wizard.cse.nd.edu - PuTTY

```
% ./work_queue_status
PROJECT              NAME                        PORT  WAITING  BUSY COMPLETE WORKERS
awe-fip35            fahnd04.crc.nd.edu          1024      719  1882  1206967    1882
hfeng-gromacs-10ps   lclsstor01.crc.nd.edu       1024     4980     0  1280240     111
hfeng2-ala5          lclsstor01.crc.nd.edu       1025     2404   140  1234514     140
forcebalance         leeping.Stanford.EDU        5817     1082    26      822      26
forcebalance         leeping.Stanford.EDU        9230        0     3      147       3
fg-tutorial          login1.futuregrid.tacc      1024        3     0        0       0
%
```

Work Queue Authentication

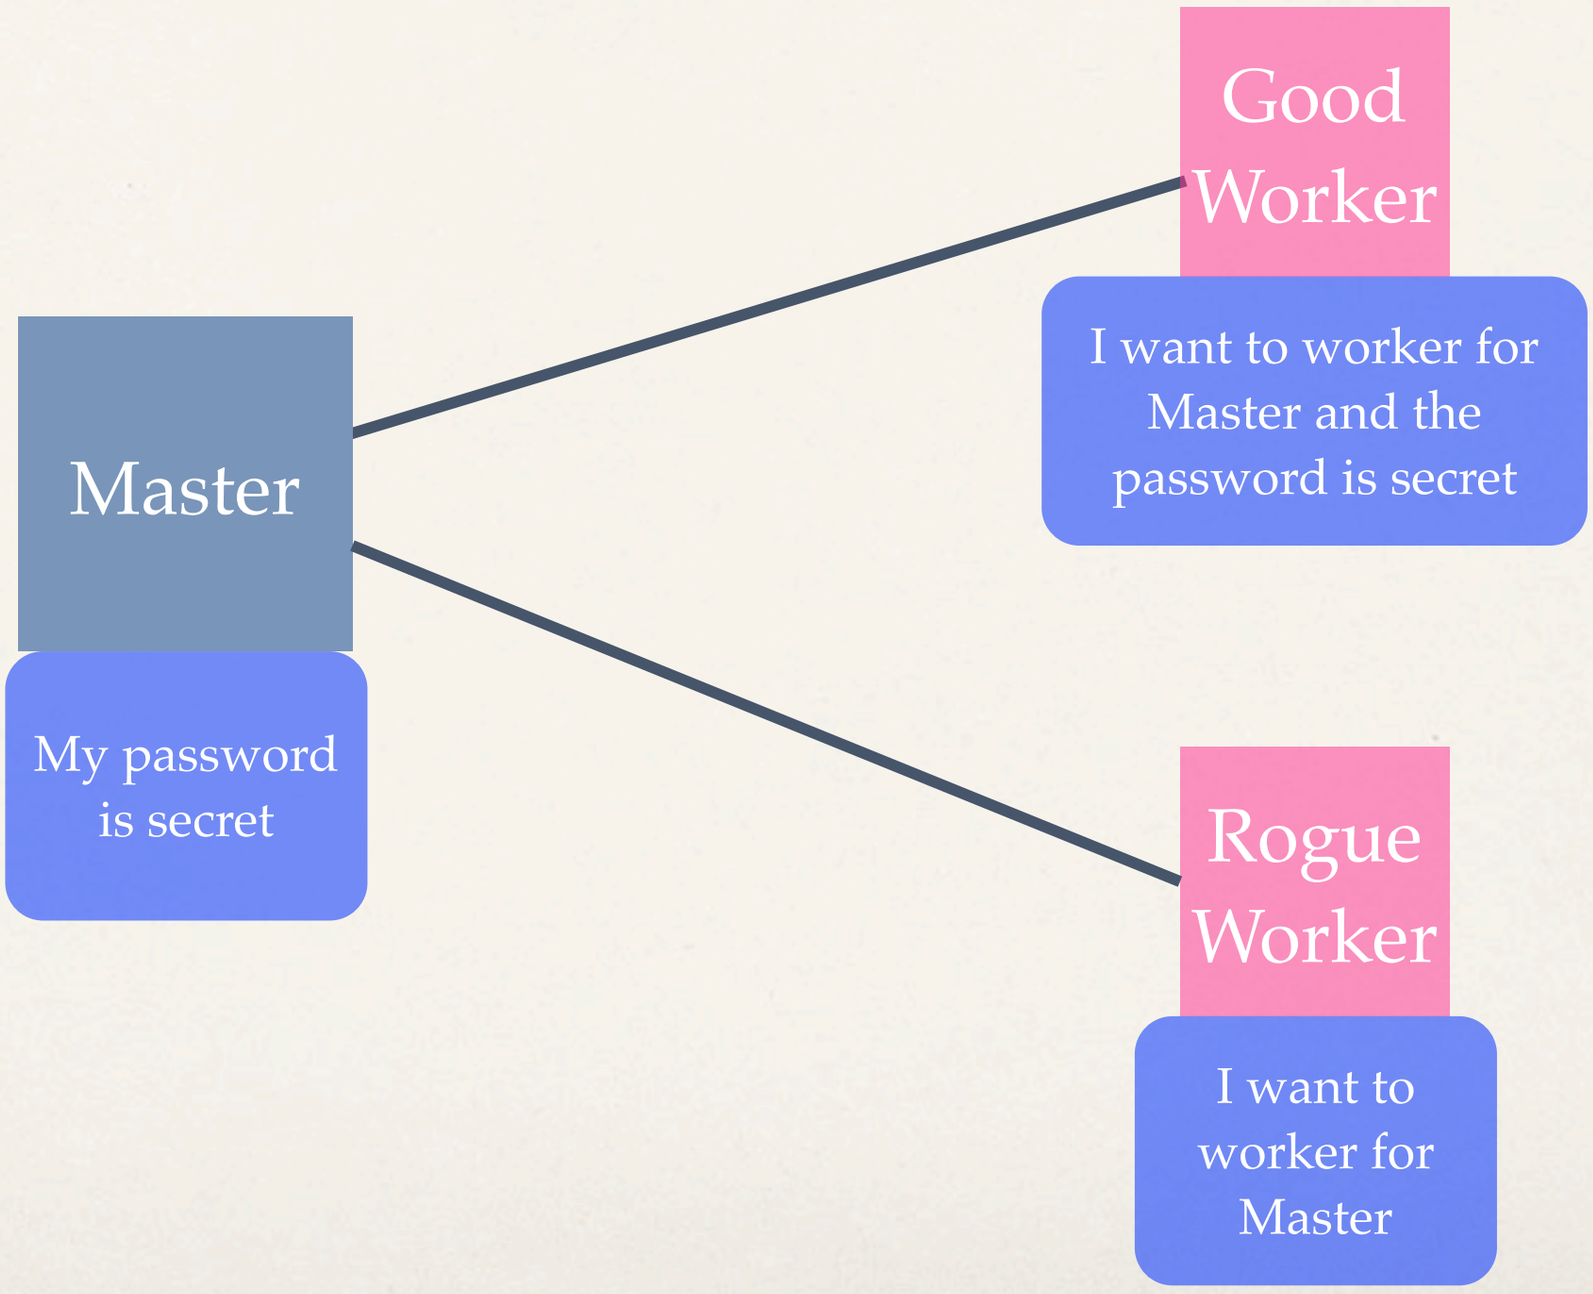

Work Queue Authentication

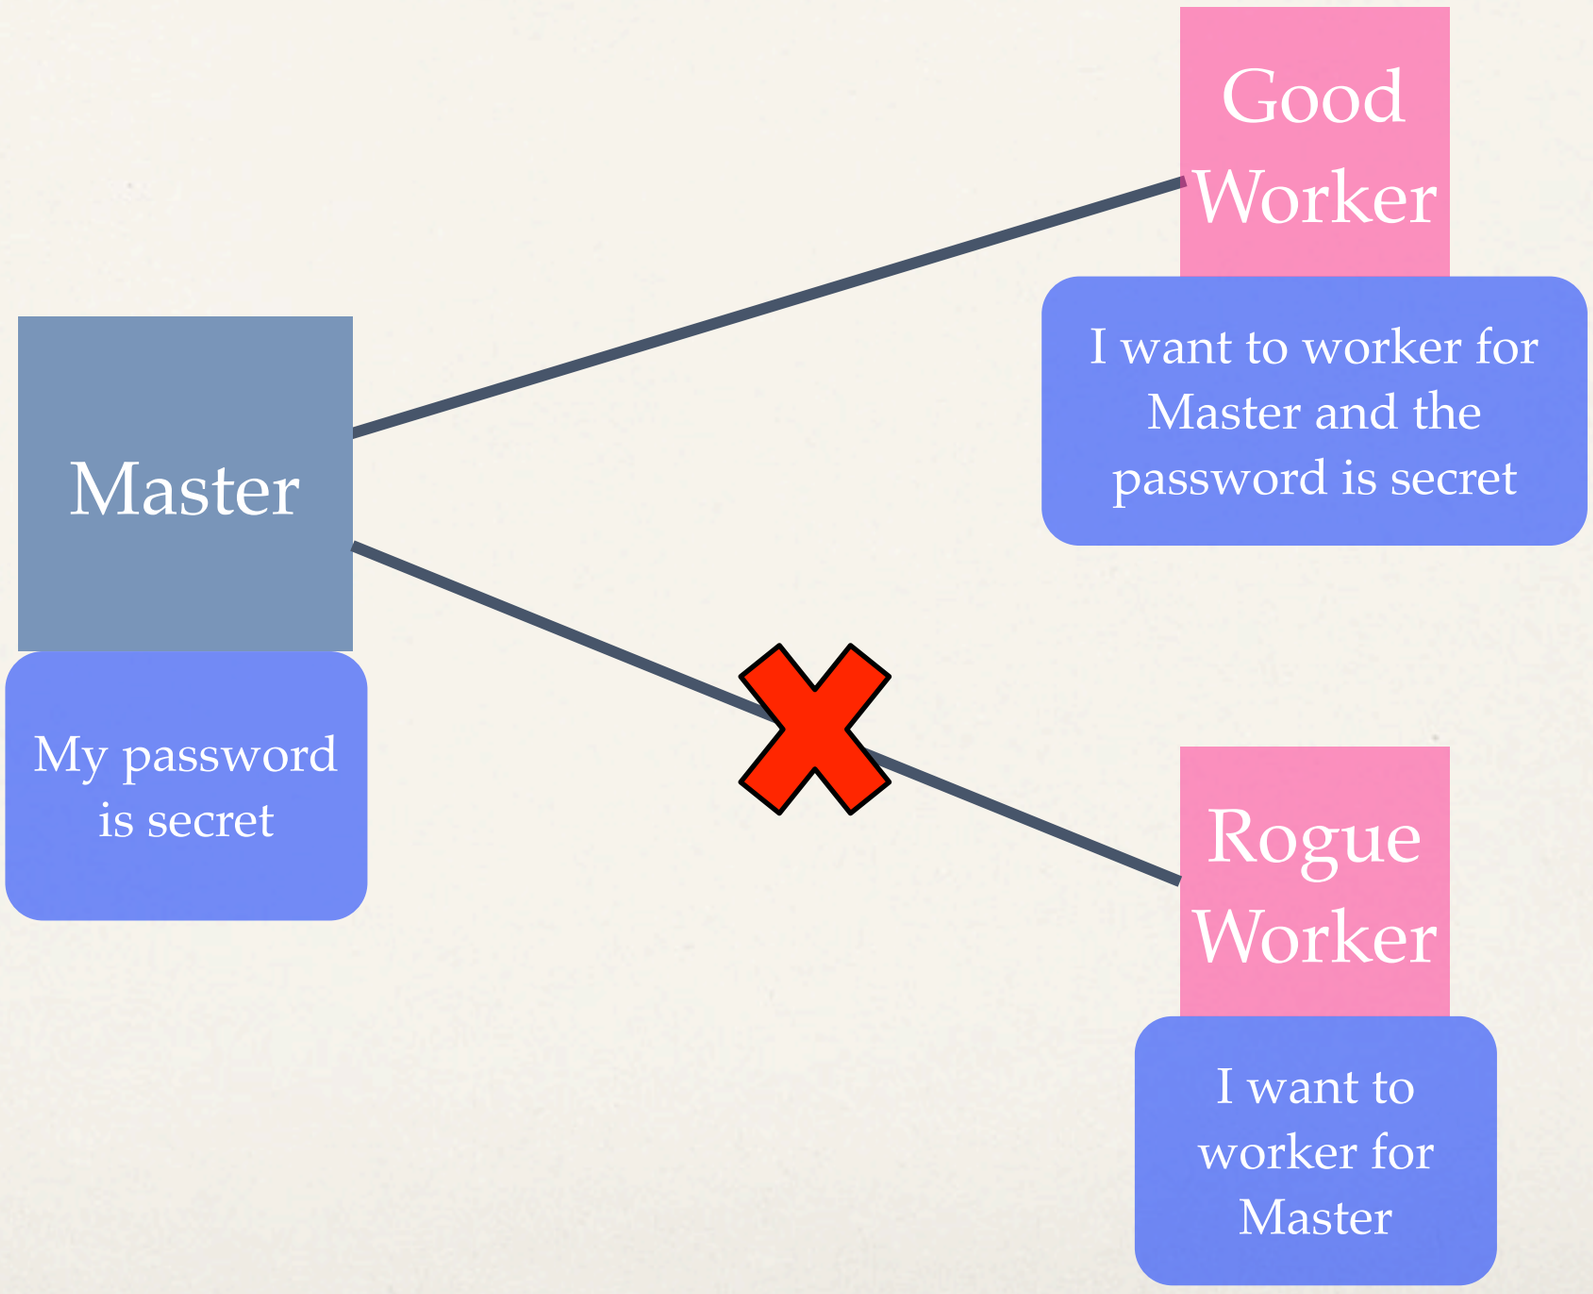

Work Queue Authentication

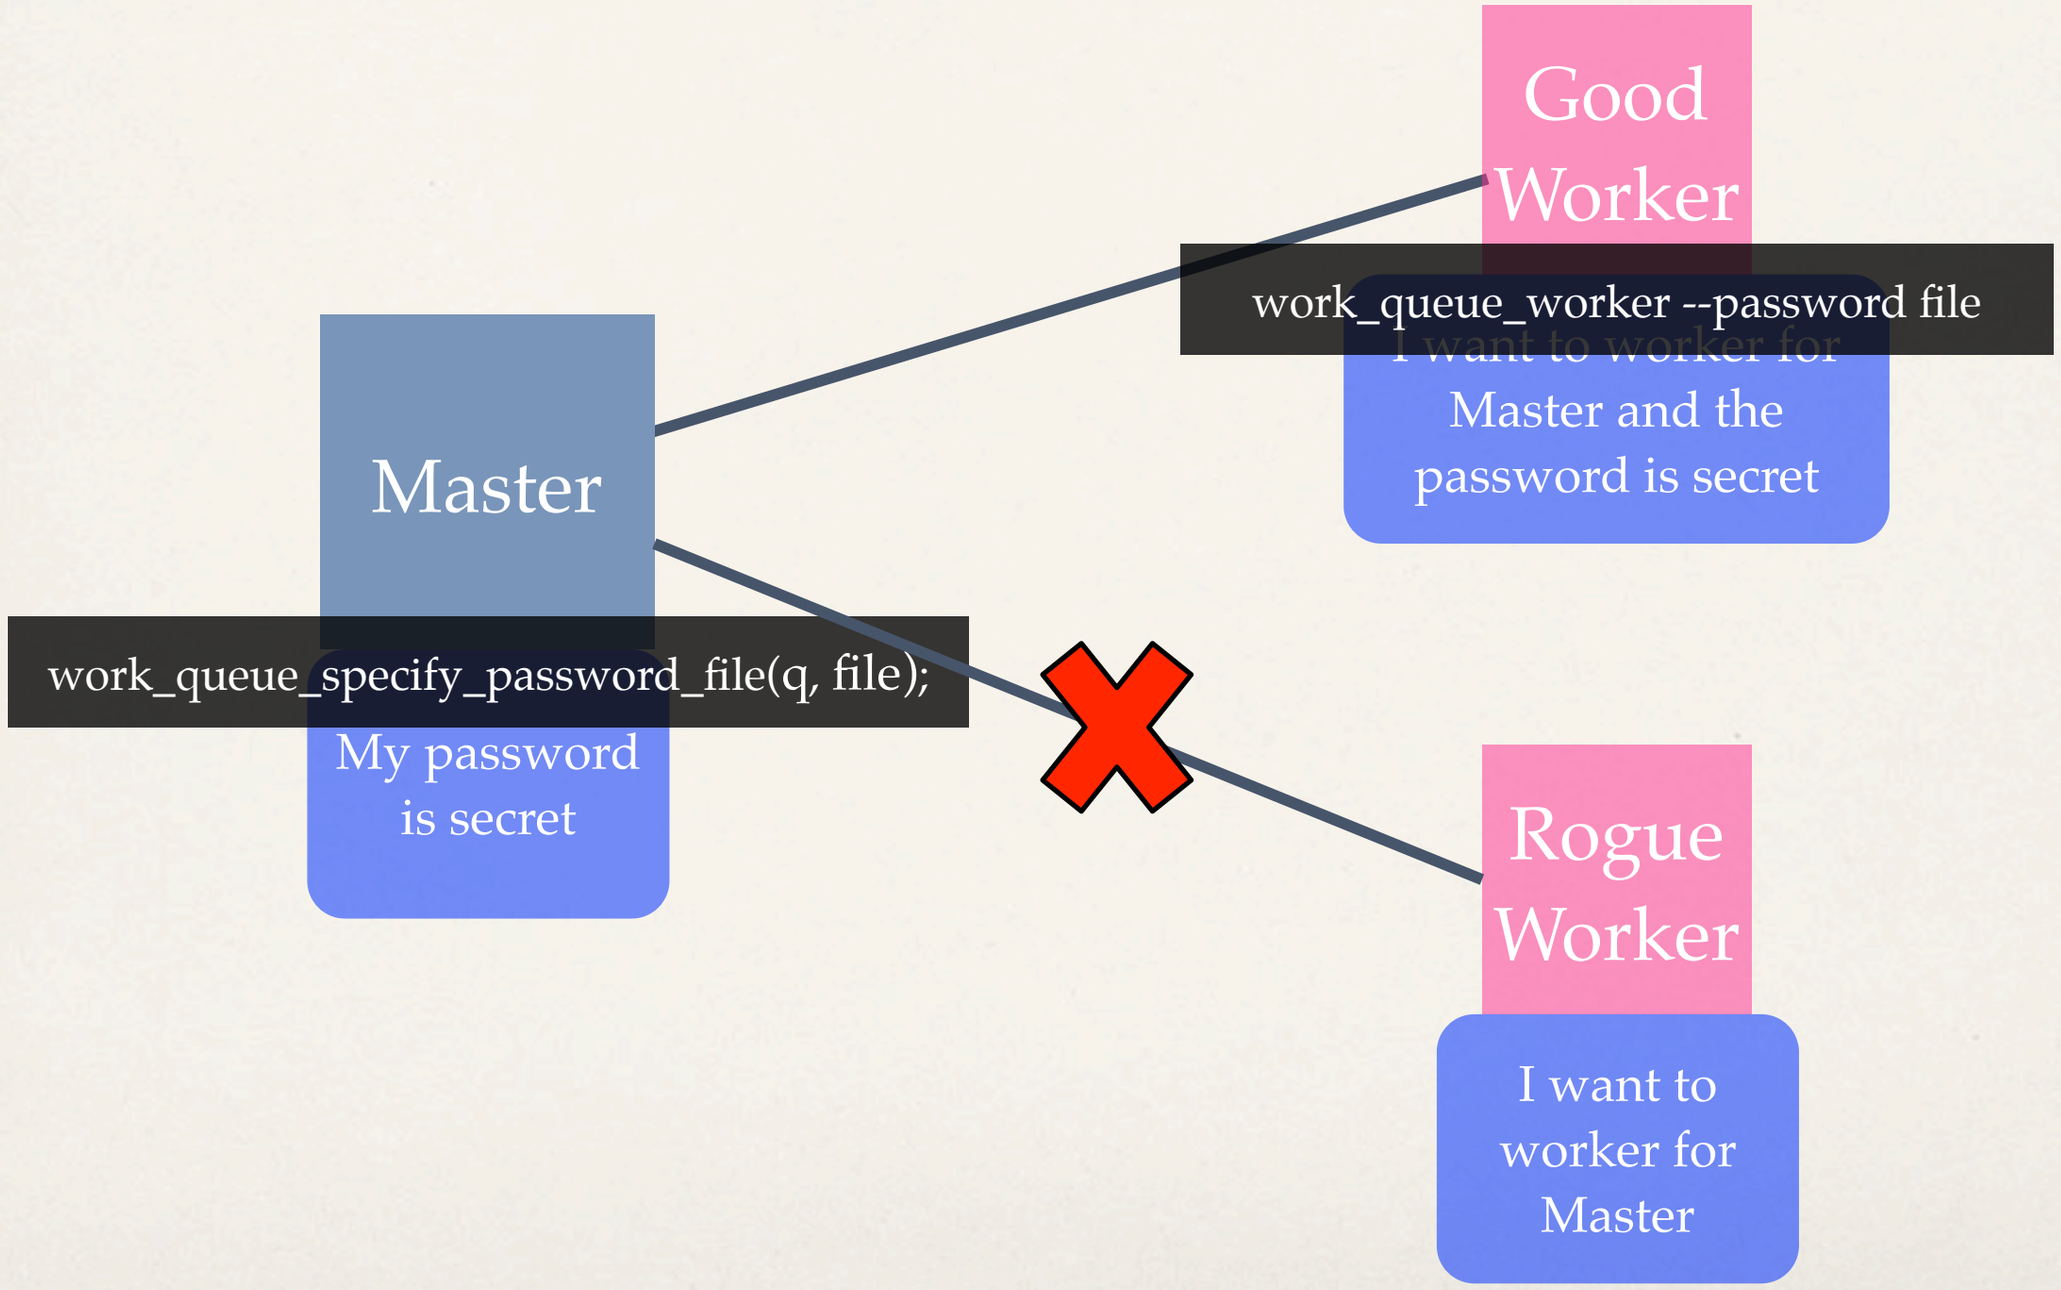

Work Queue Partial Files

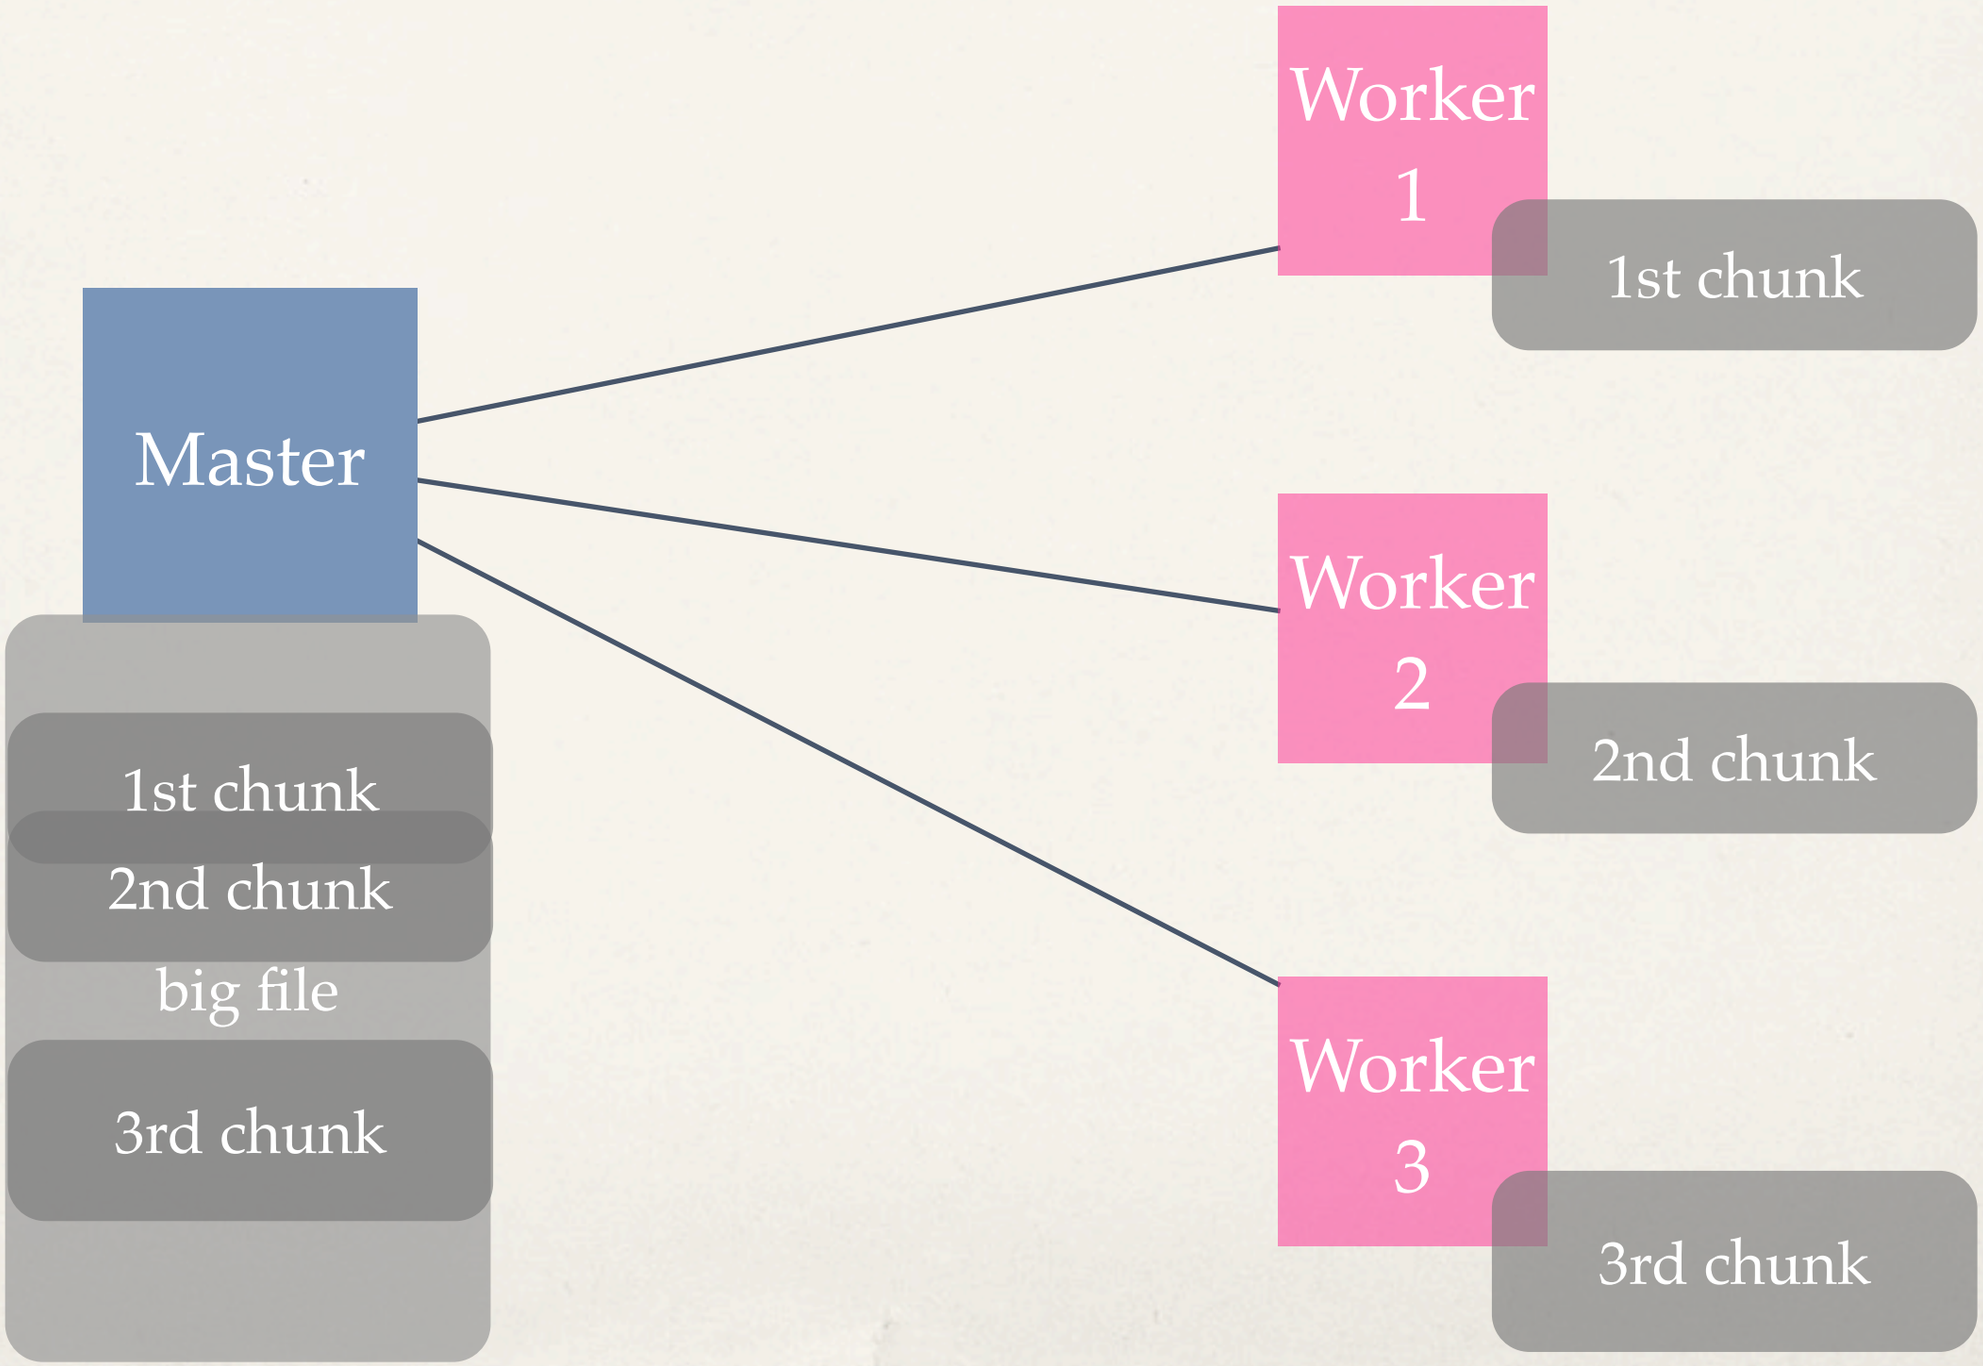

Work Queue Partial Files

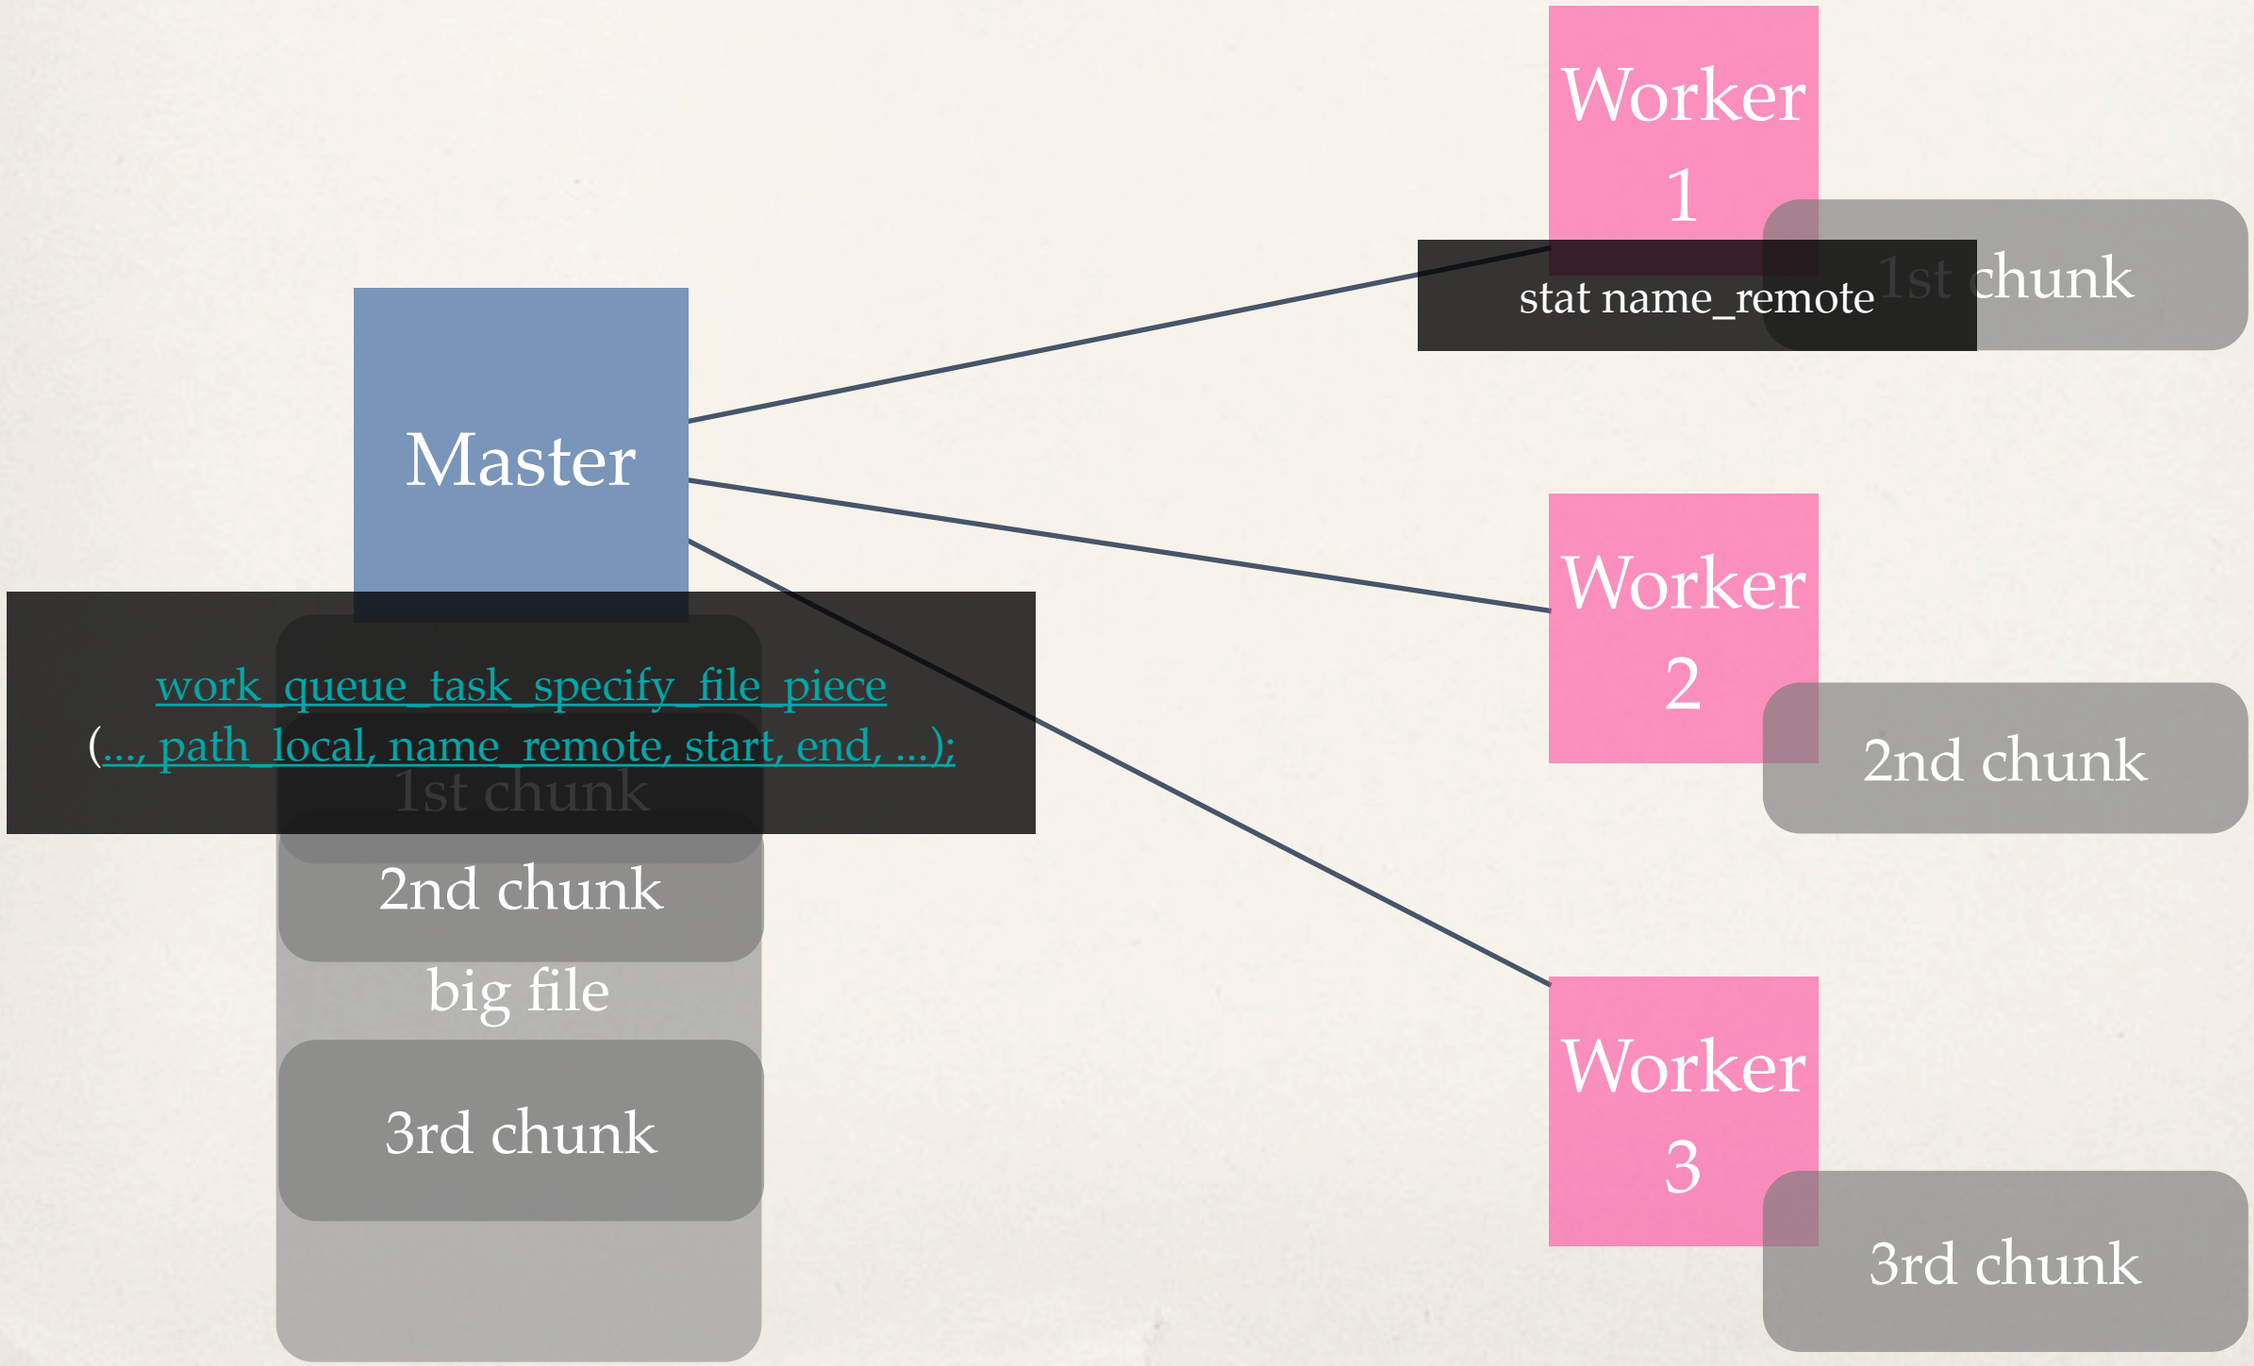

Resource Monitor

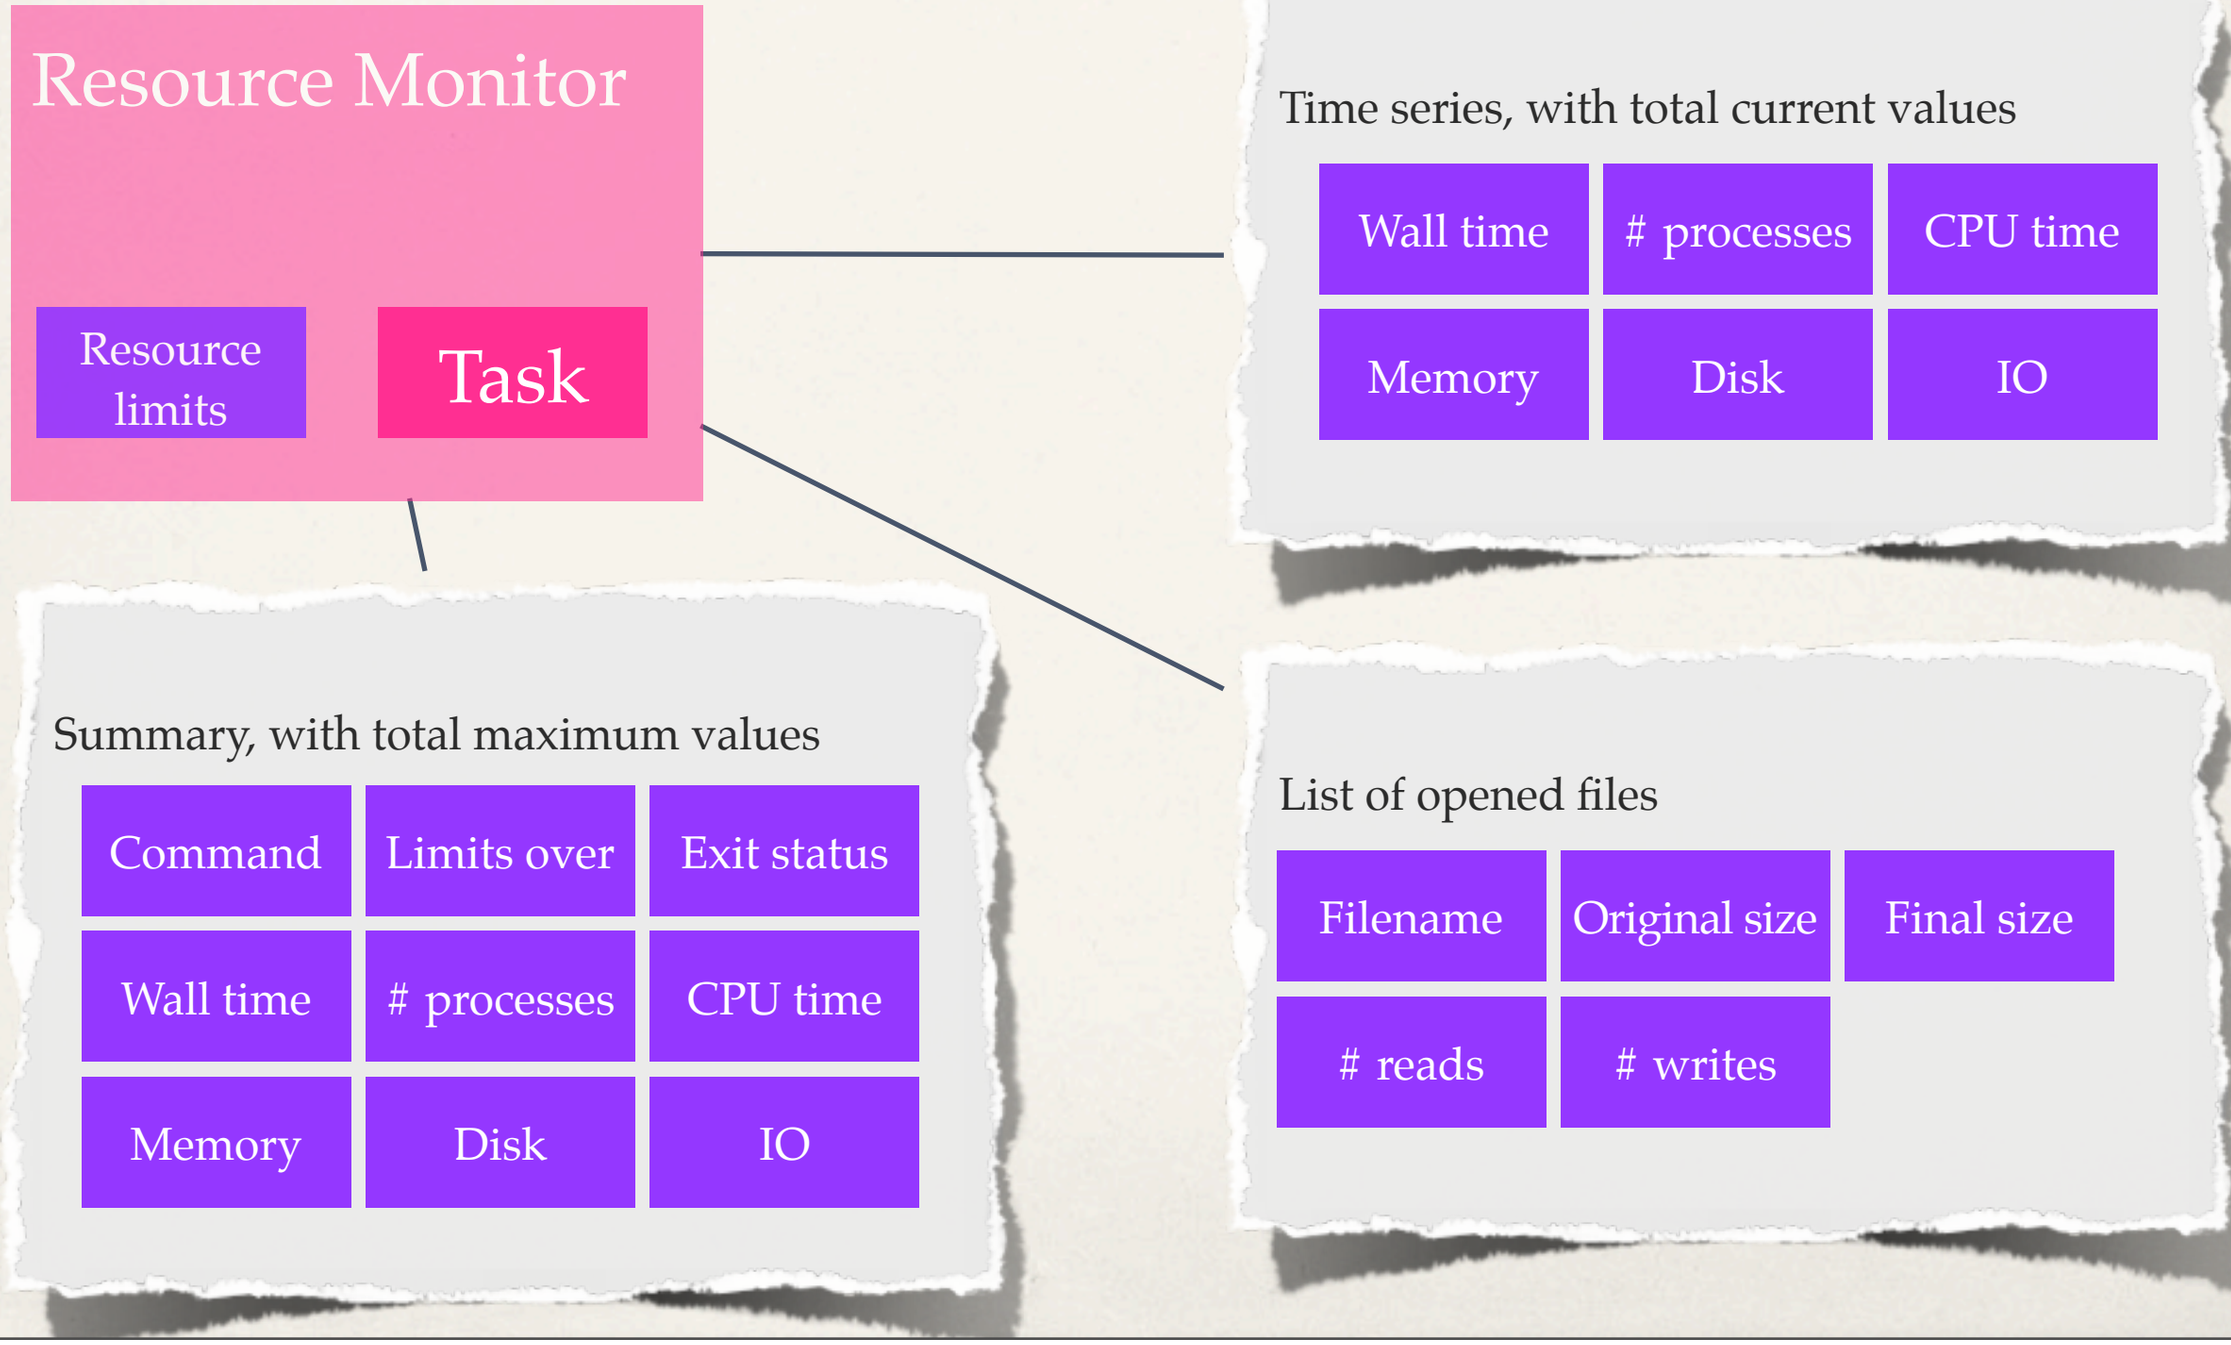

Resource Monitor

Time series, with total current values

Wall time	# processes	CPU time
Memory	Disk	IO

Summary, with total maximum values

Command	Limits over	Exit status
Wall time	# processes	CPU time
Memory	Disk	IO

List of opened files

Filename	Original size	Final size
# reads	# writes	

Resource Monitor and Makeflow

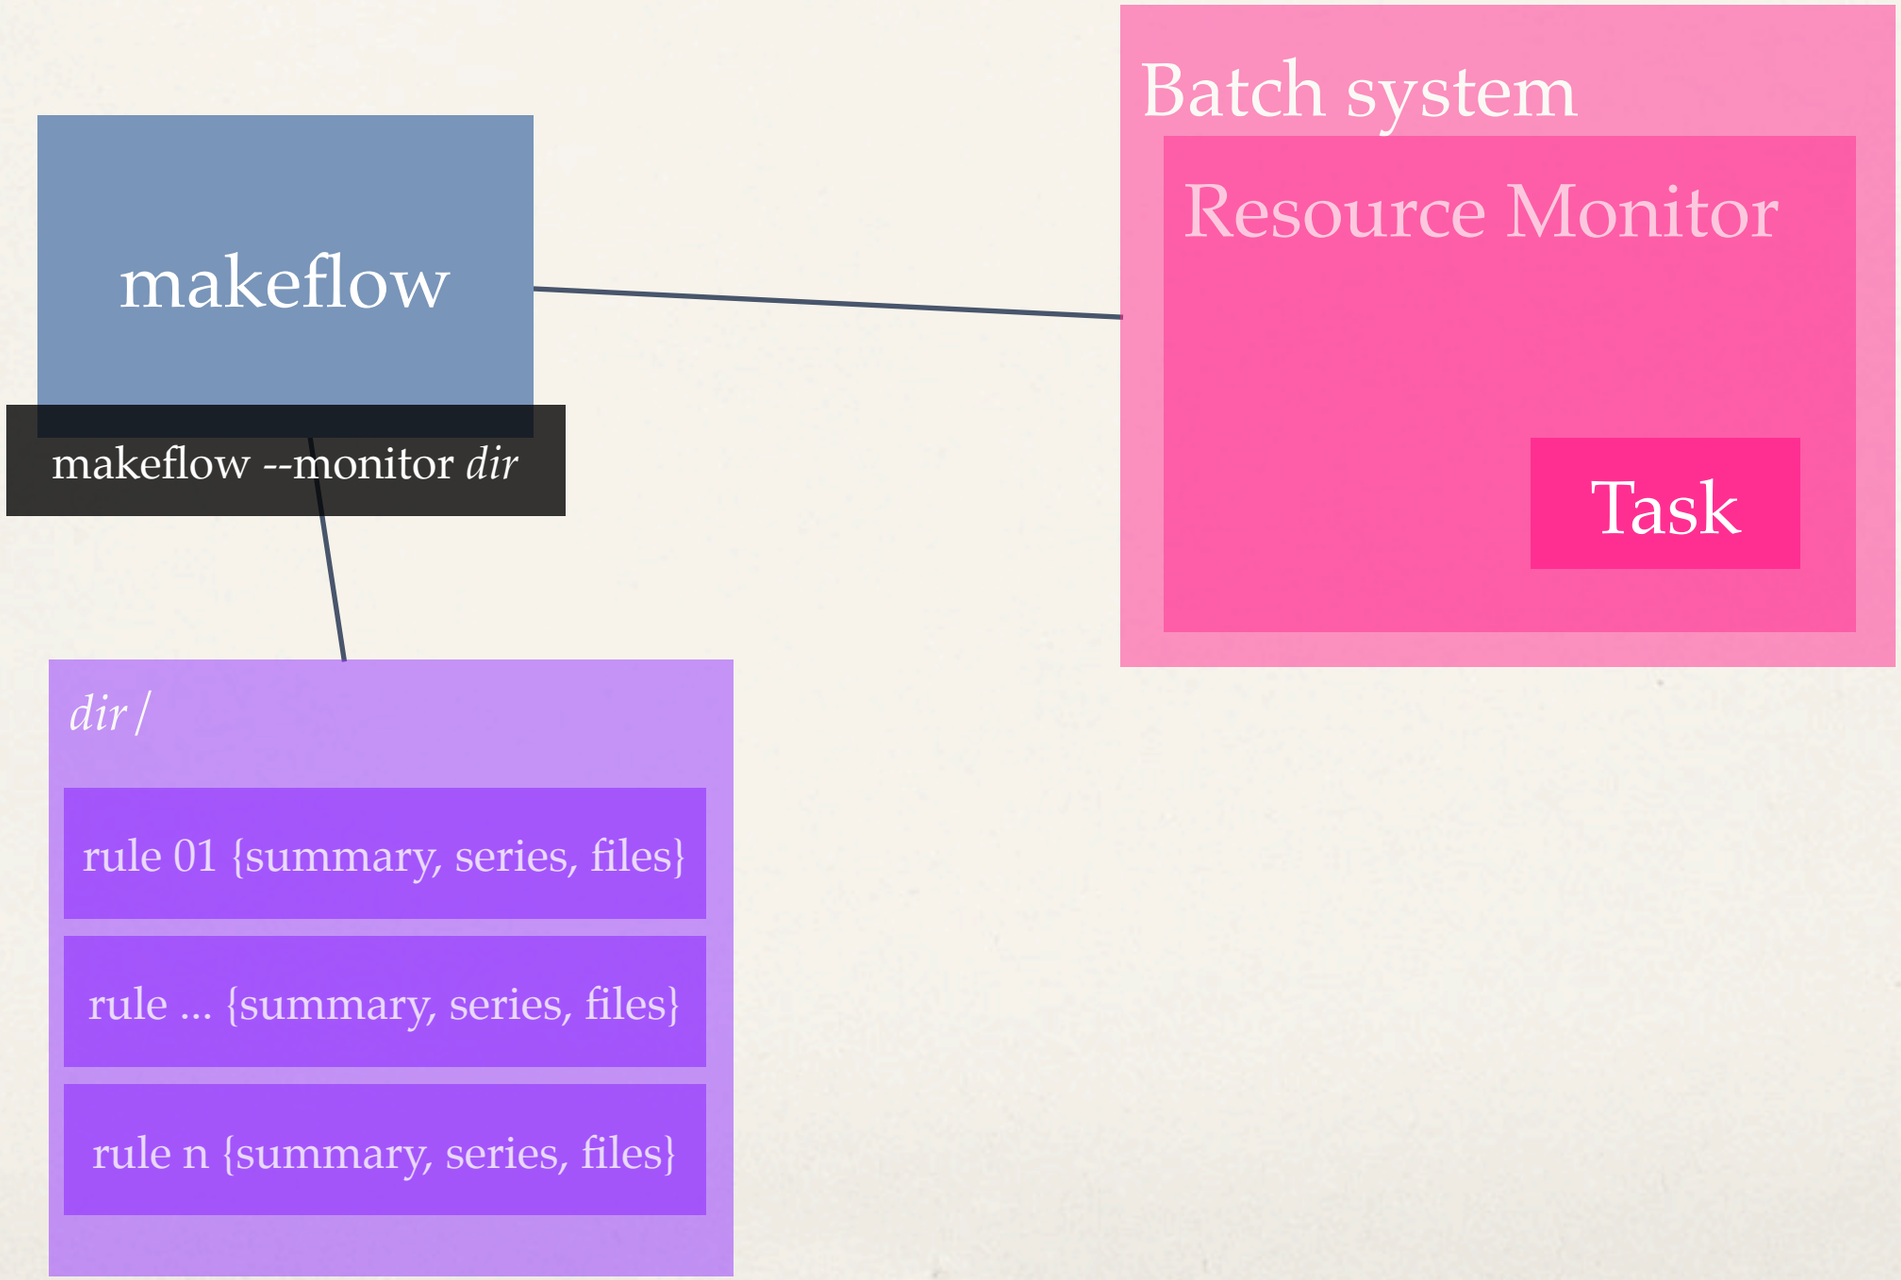

Resource Monitor and WQ

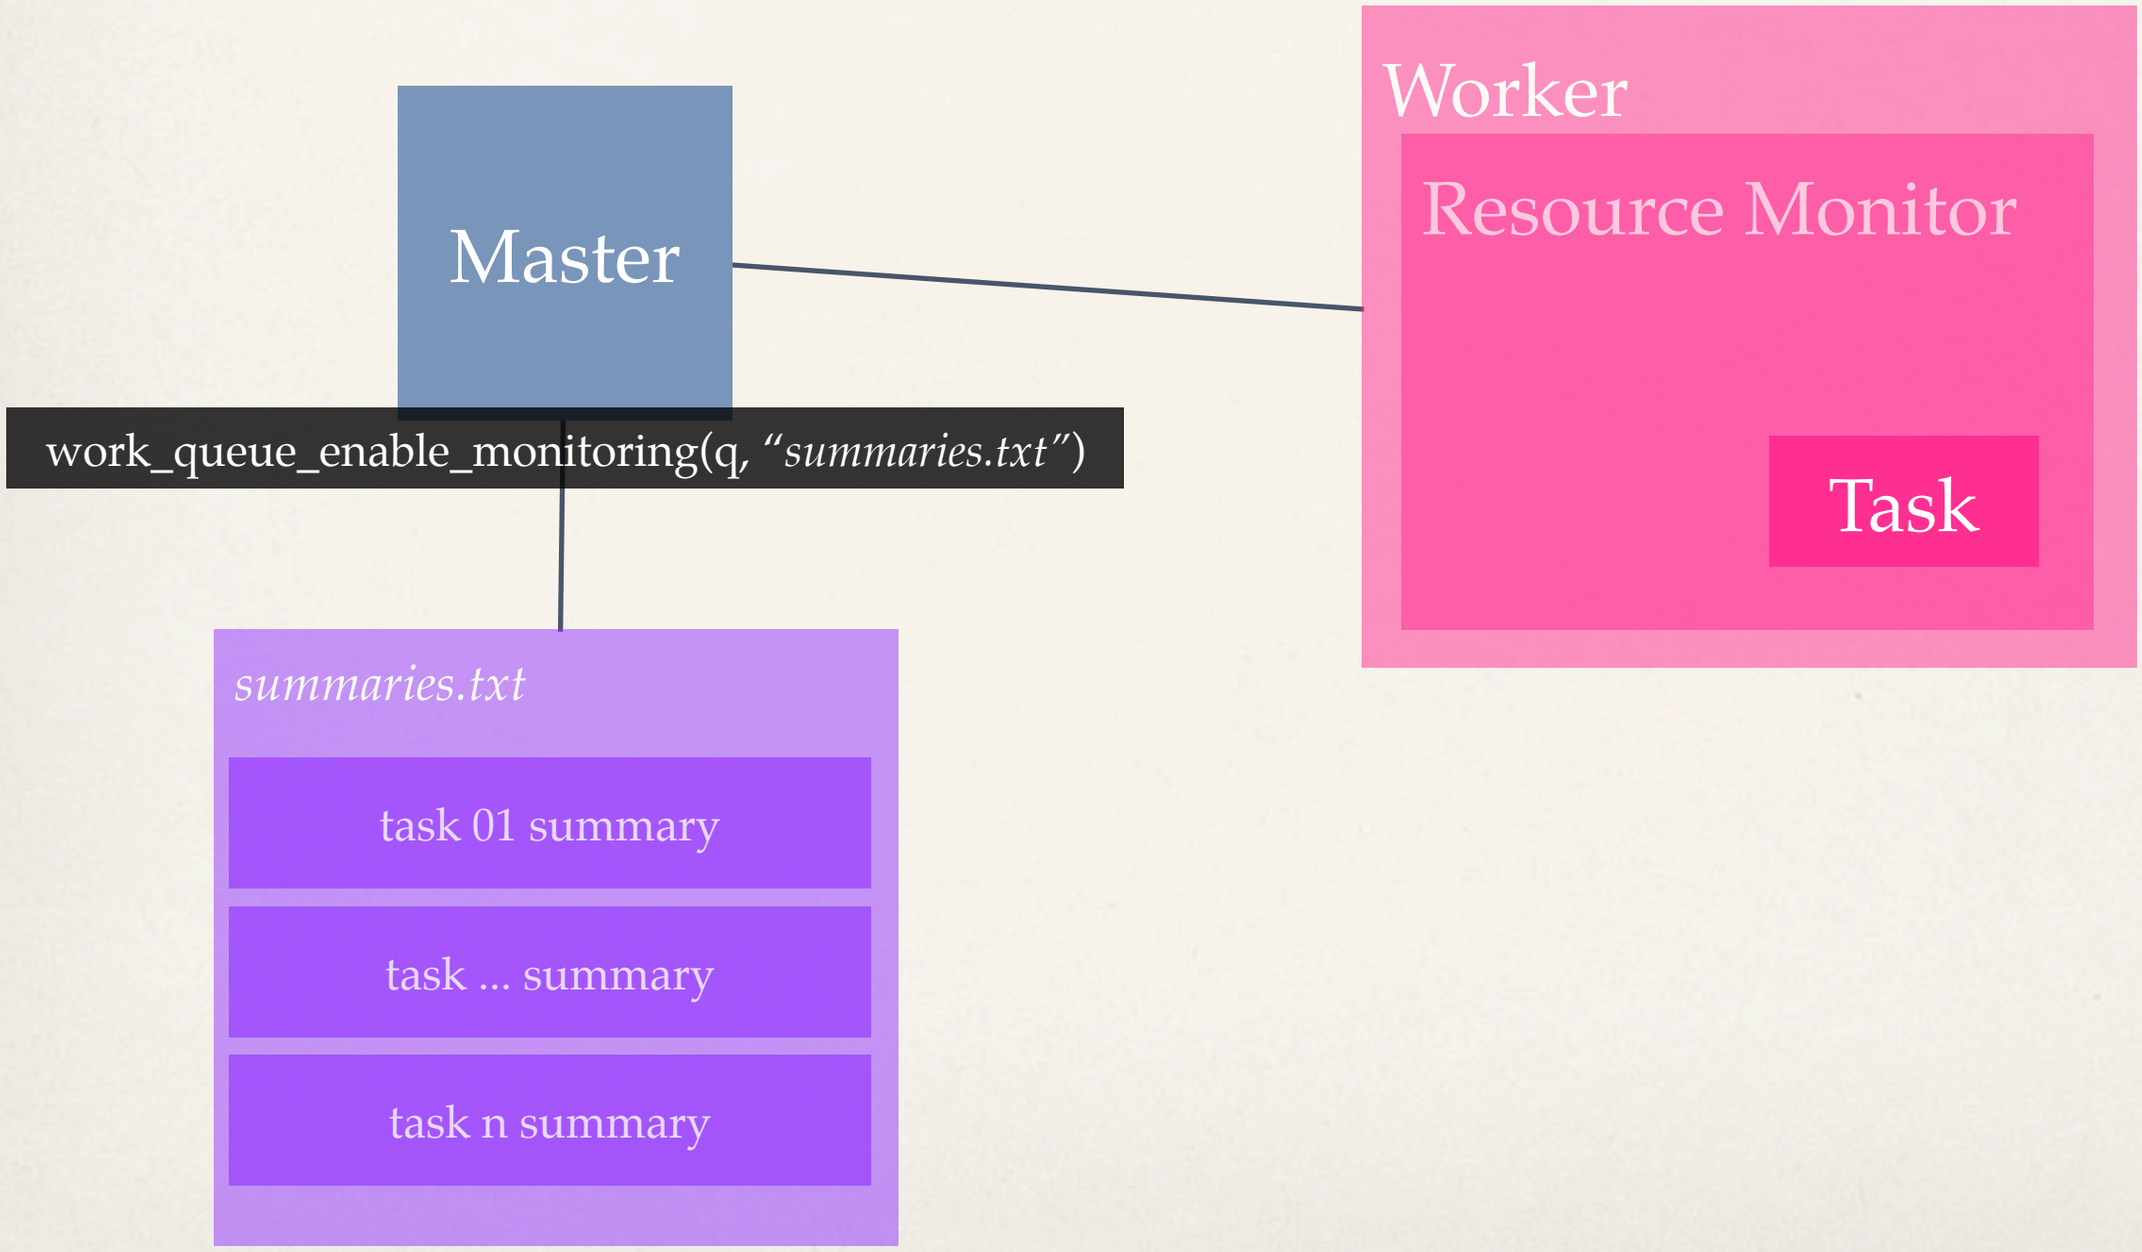

Resource Monitor Visualizer

dir /

task 01 {summary, series}

task ... {summary, series}

task n {summary, series}

visualizer

Resource Monitor Visualizer

dir /

task 01 {summary, series}

task ... {summary, series}

task n {summary, series}

resource_monitor_visualizer *logs* *output* *title*

visualizer

Makeflow Task Categories

Makeflow file

CATEGORY=*preprocess*
DISK=100
MEMORY=10
CORES=2
rule: ...

CATEGORY=*analysis*
...

Resource monitor as watchdog

Work Queue task requirements

Condor ClassAd requirements

Makeflow Remote Renaming

Makeflow file

output_local->output_remote: input_local->input_remote cmd_local-> cmd_remote
    cmd_remote --in input_remote --out output_remote

Local filesystem

output_local
input_local
cmd_local

Remote filesystem

output_remote
input_remote
cmd_remote

Makeflow Linker - Coming Soon

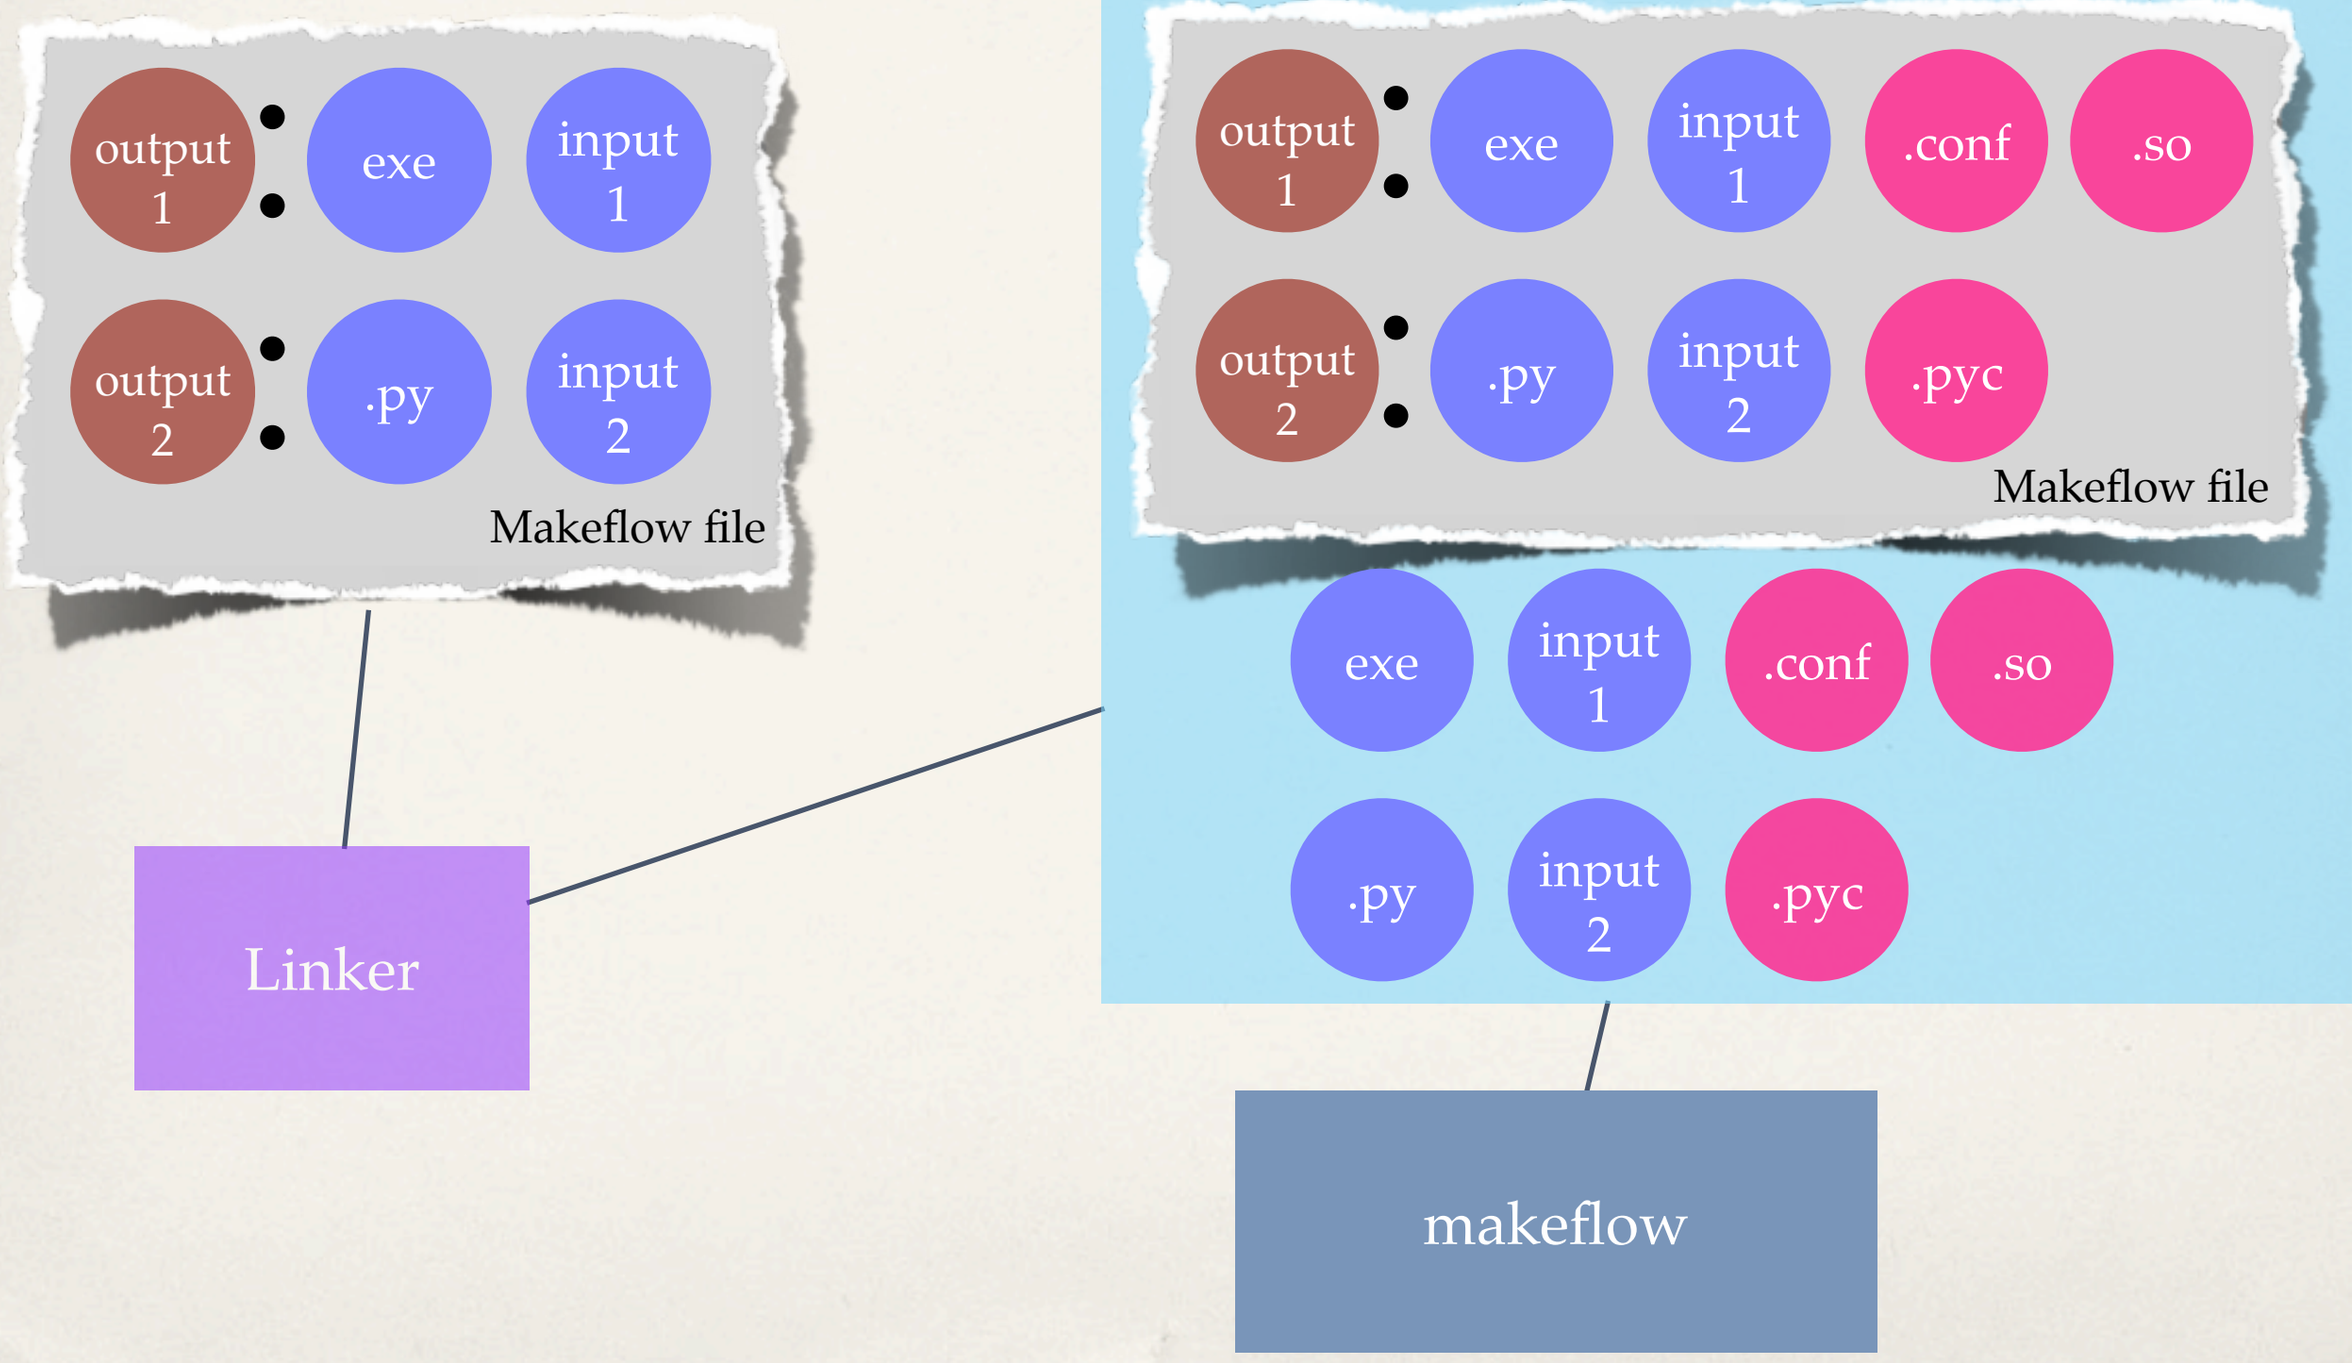

Chirp: Unix Authentication

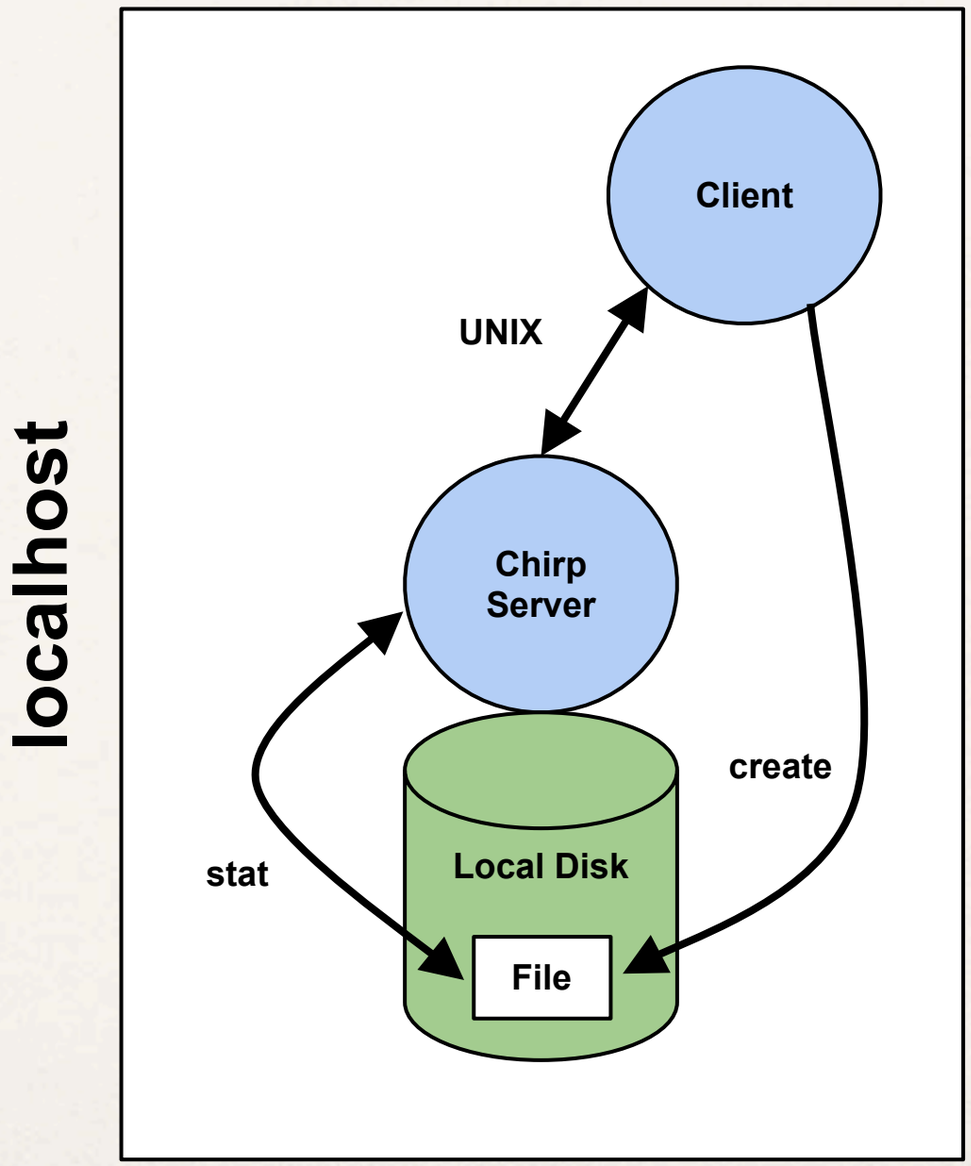

unix:pdonnel3

Chirp: Hostname/Address Authentication

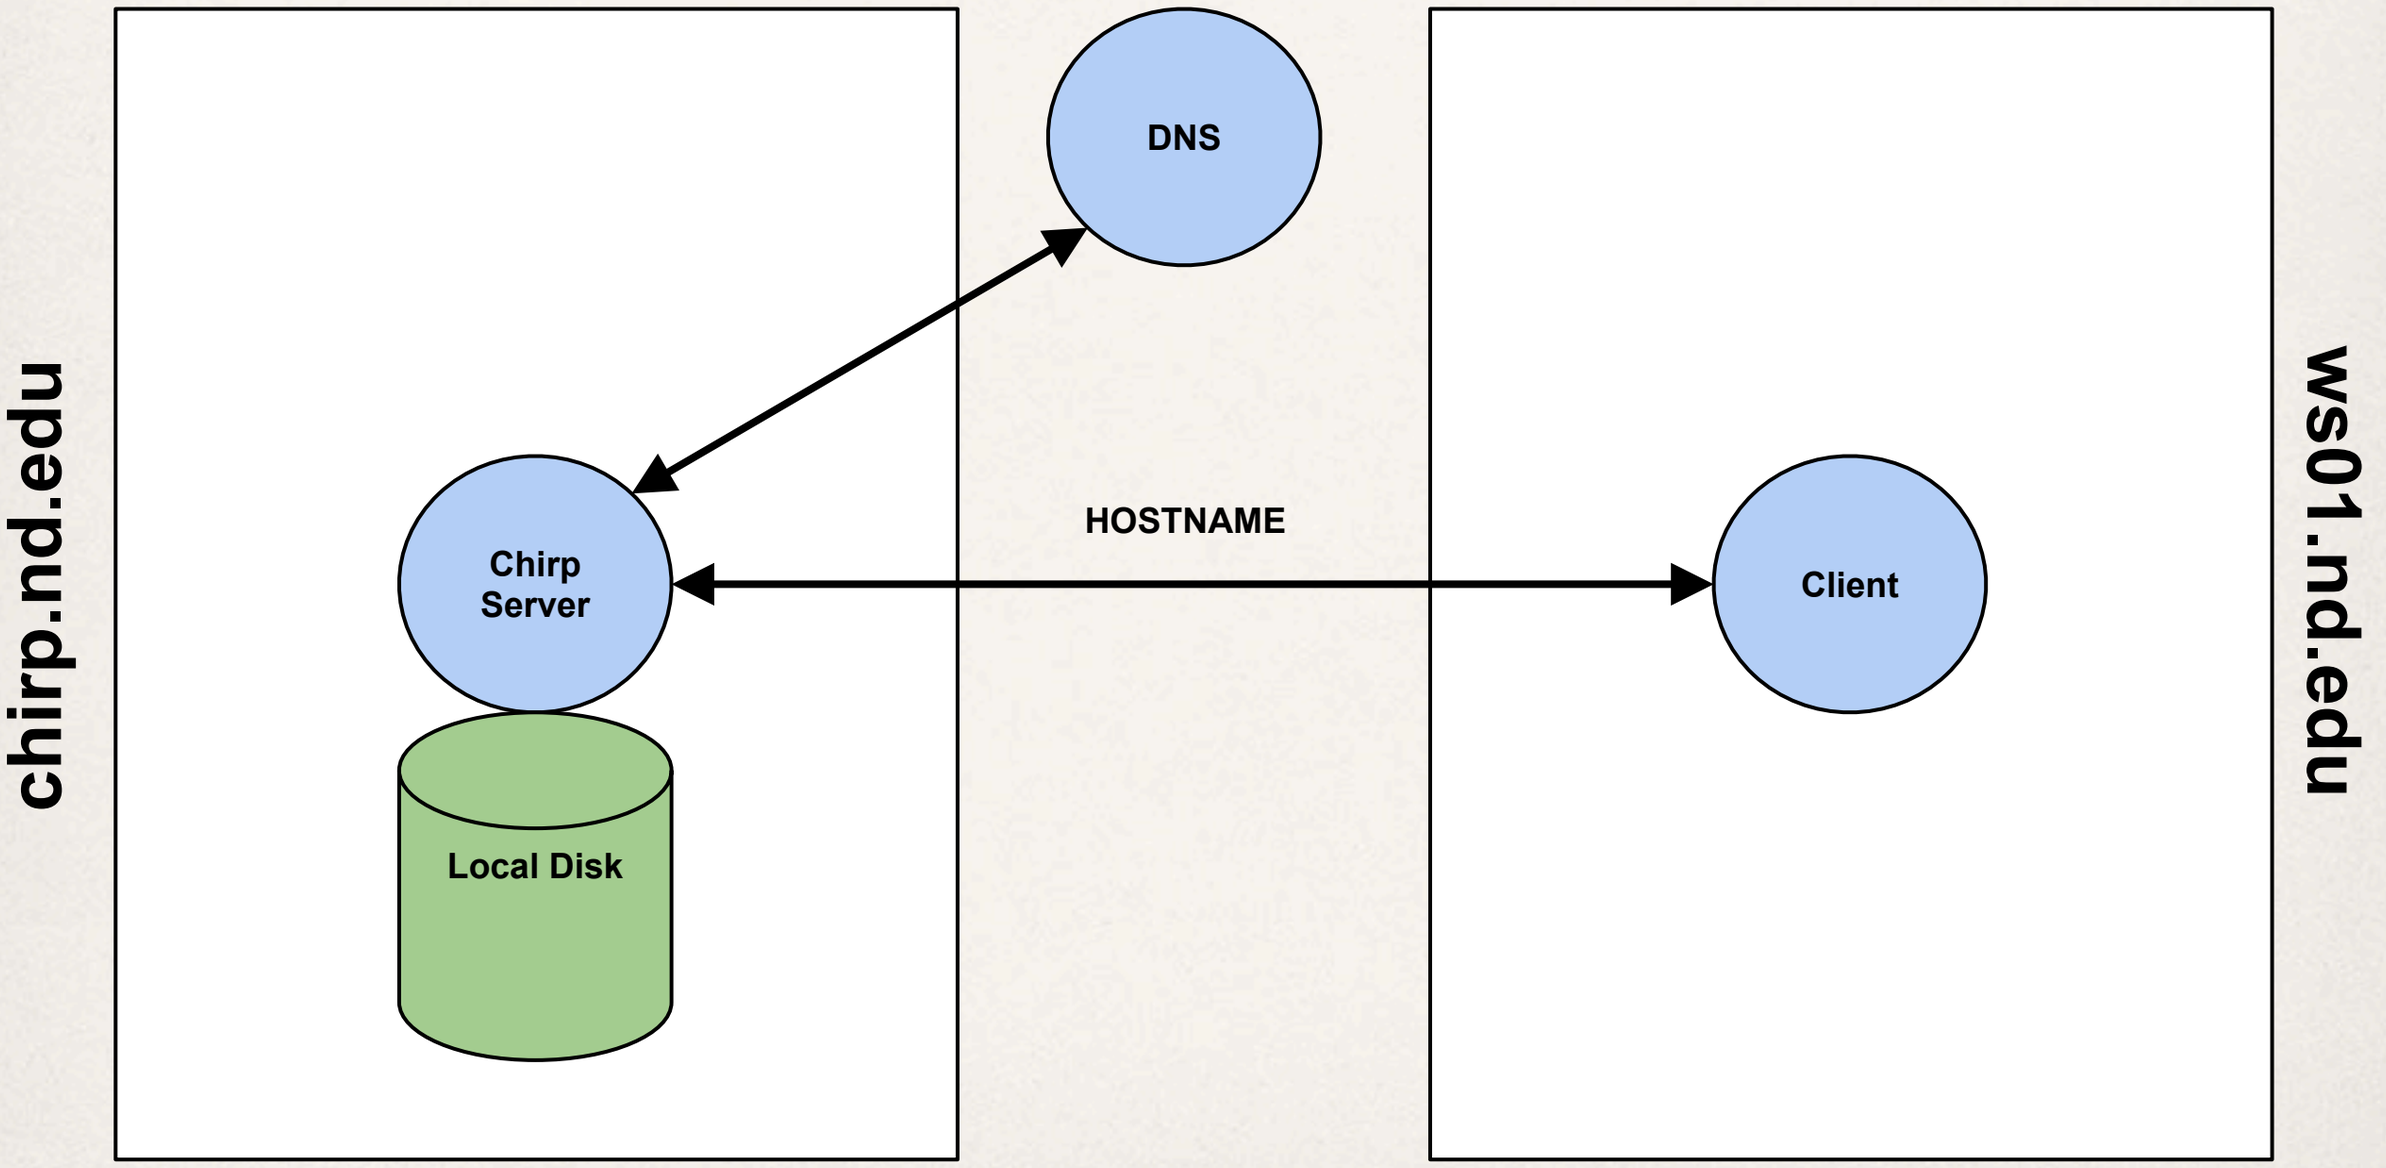

hostname:ws01.nd.edu

Chirp: Kerberos Authentication

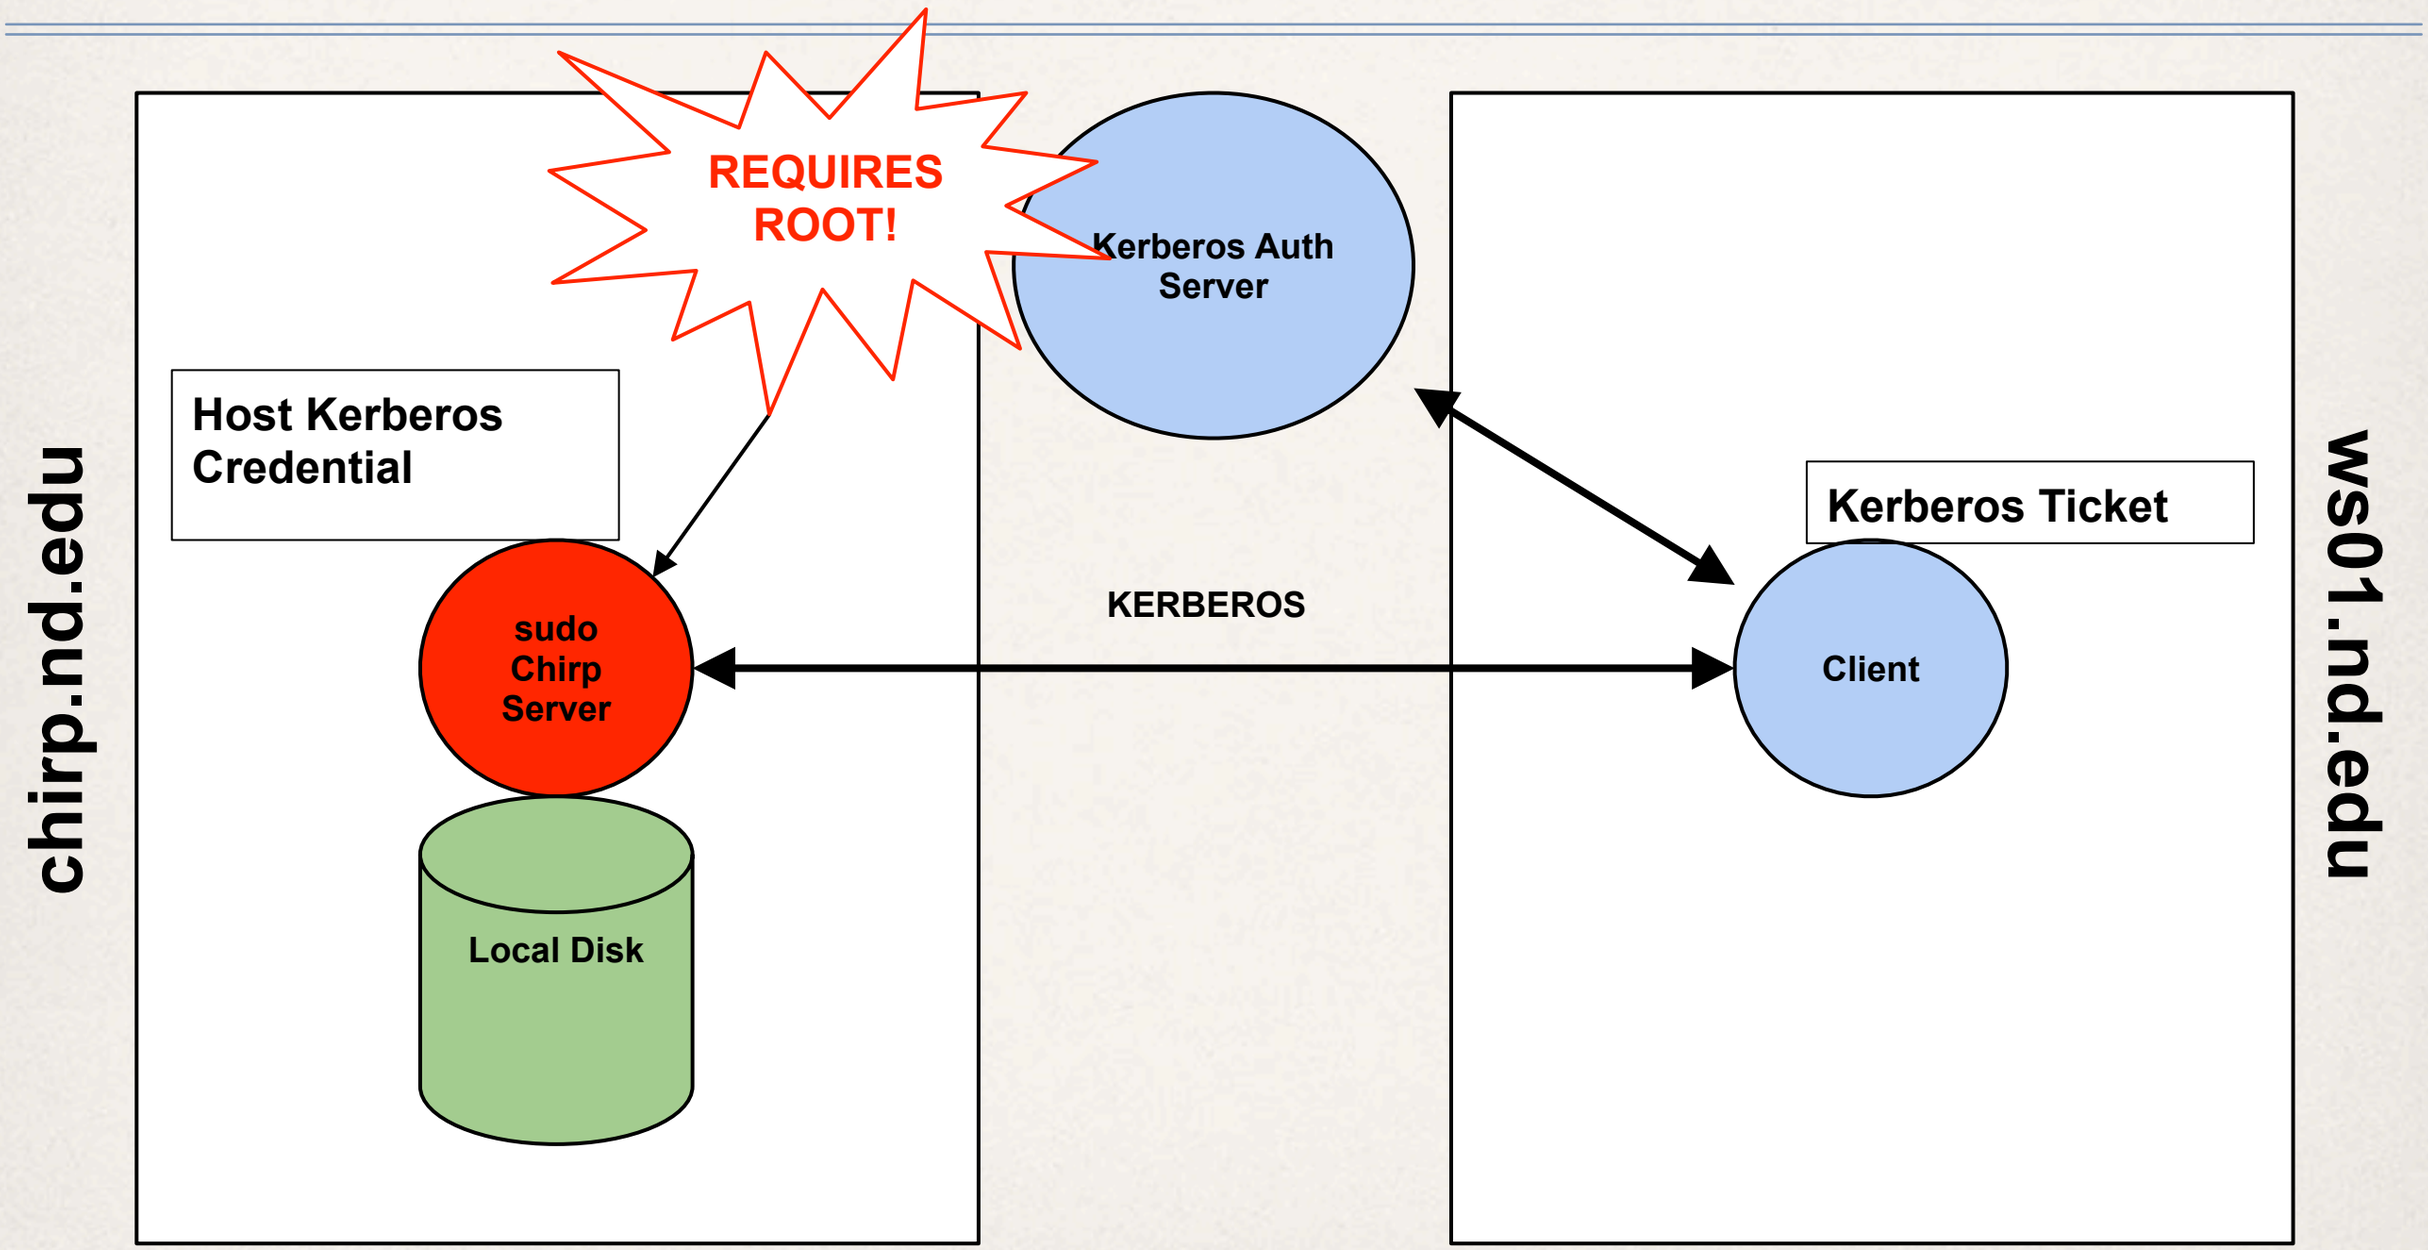

kerberos:PDONNEL3@ND.EDU

Chirp: Globus Authentication

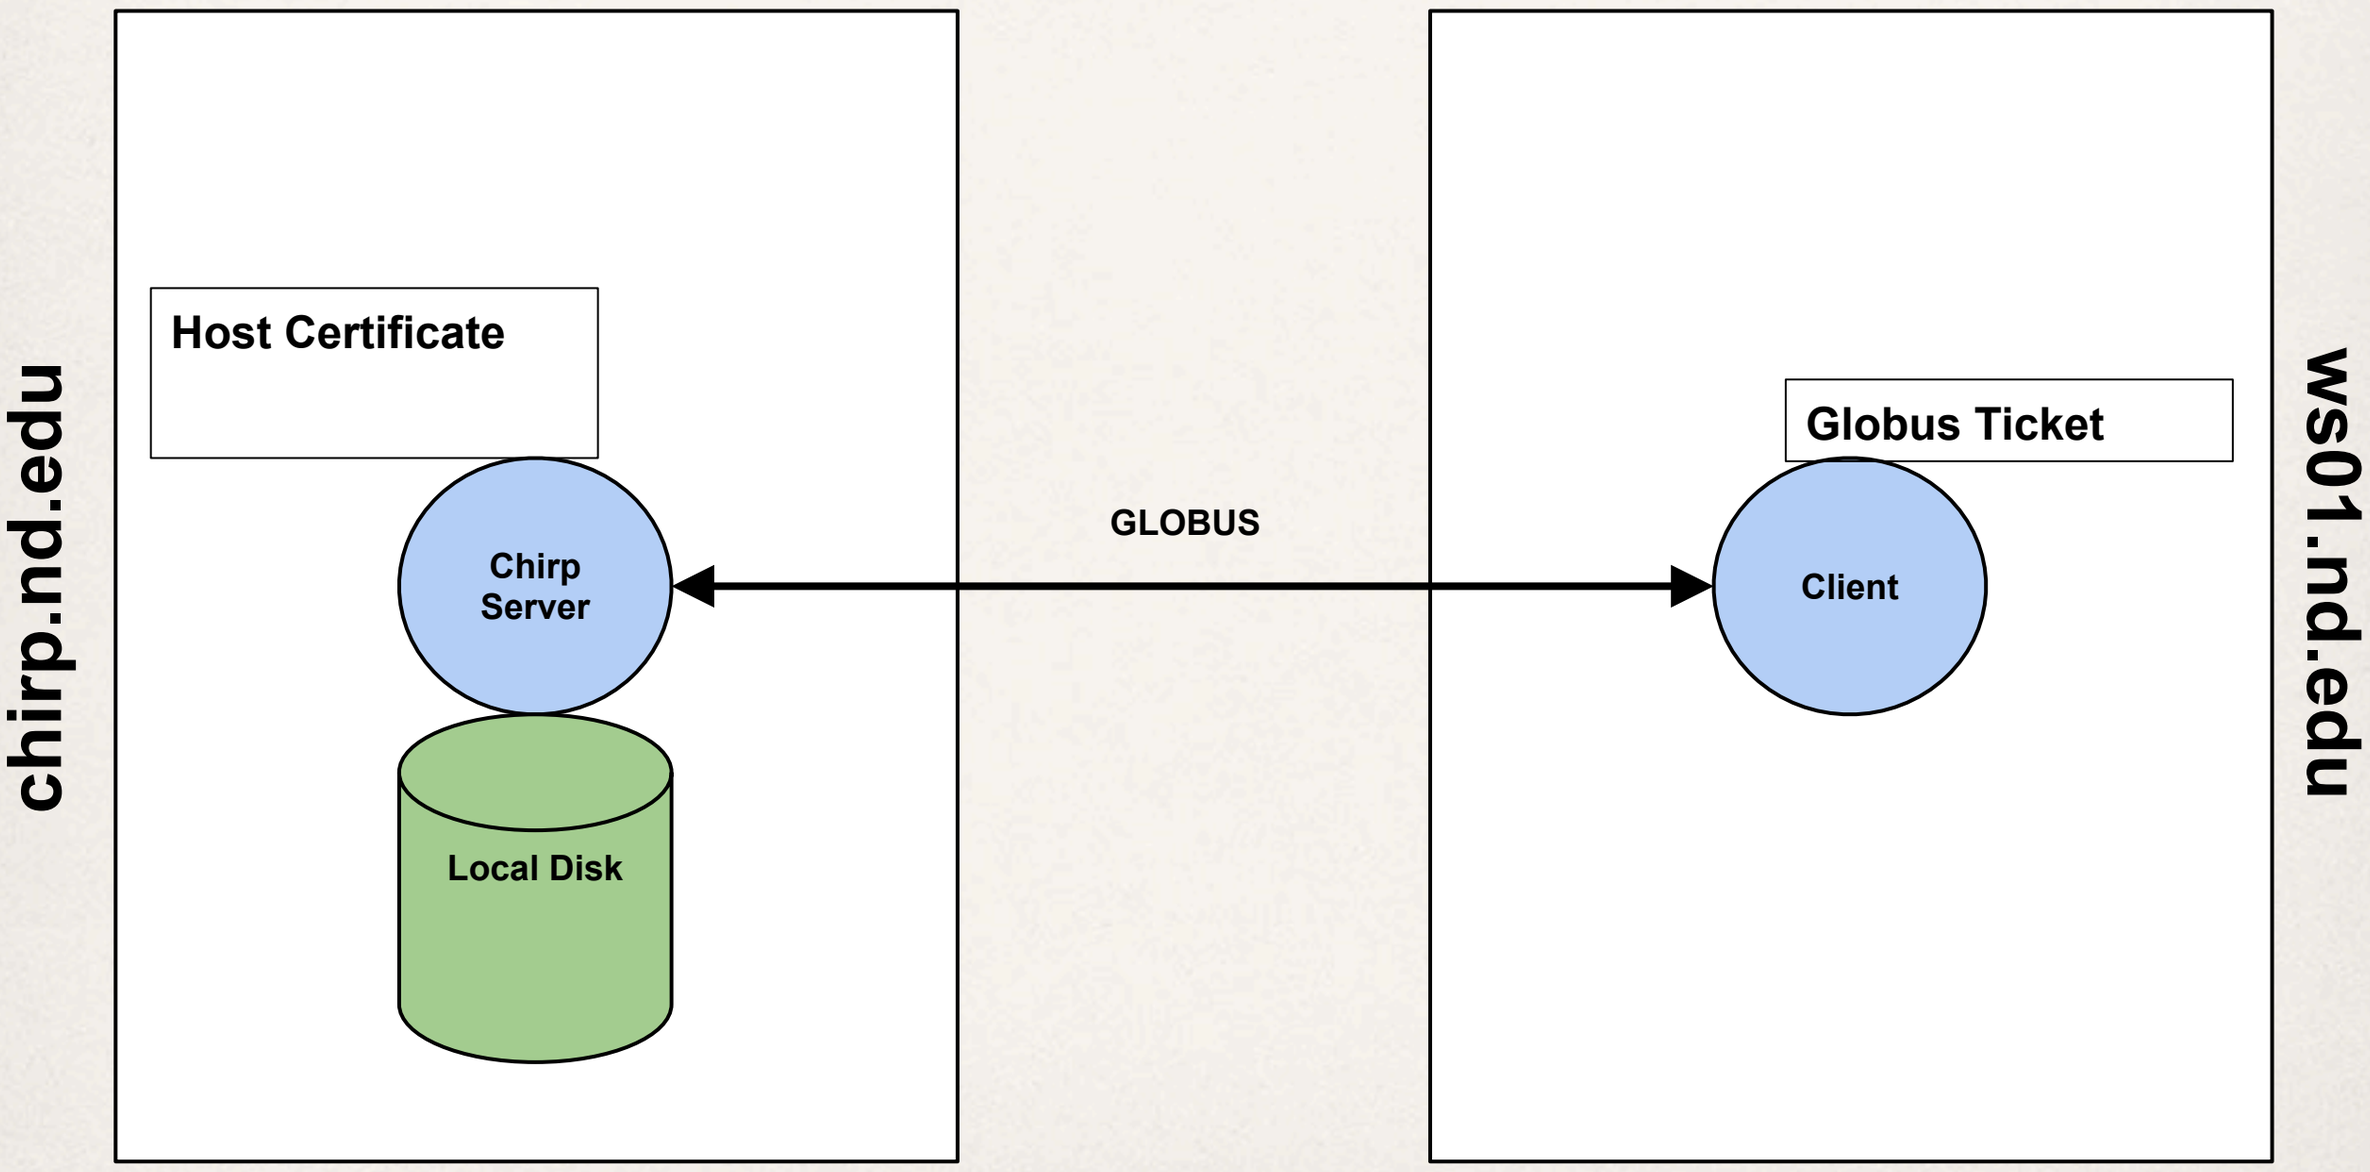

globus:/O=Cooperative_Computing_Lab/CN=Patrick_Donnelly

Chirp: Challenges

Problems to solve:

- Low value (low risk) credential to ship with jobs.
- Disposable & Time Limited.
- Works across multiple infrastructures.

Chirp: Ticket Authentication

```
chirp chirp.nd.edu ticket_create -output my.ticket
                                 -bits 512
                                 -duration 86400
                                 / rl
                                 /data rwl
```

1. Create Ticket

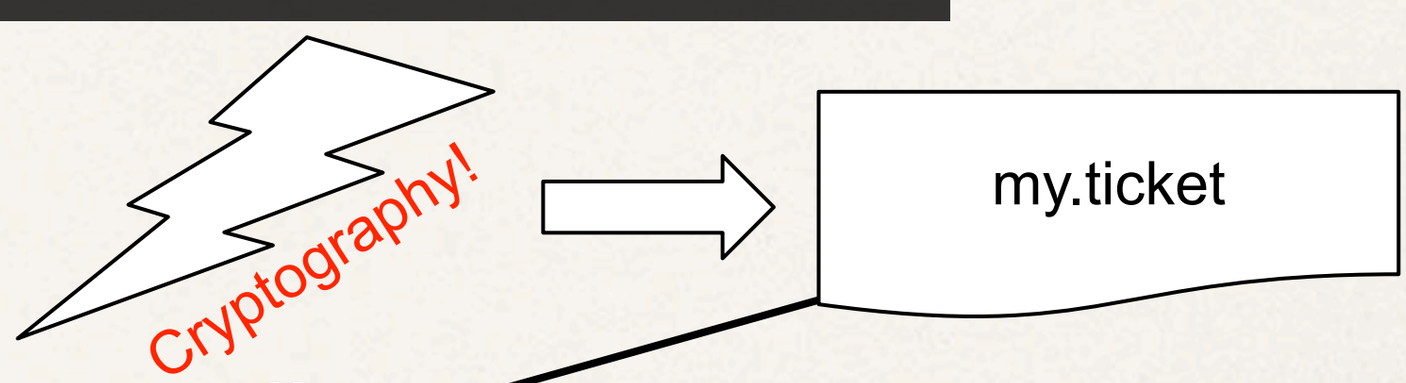

2. Register Ticket

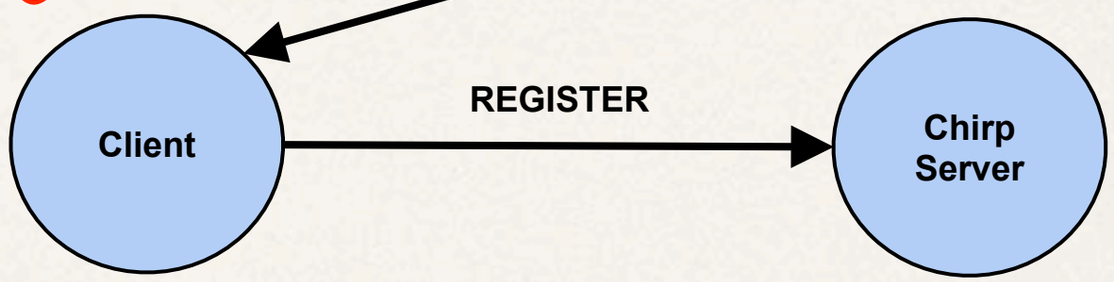

3. Set ACL Limits

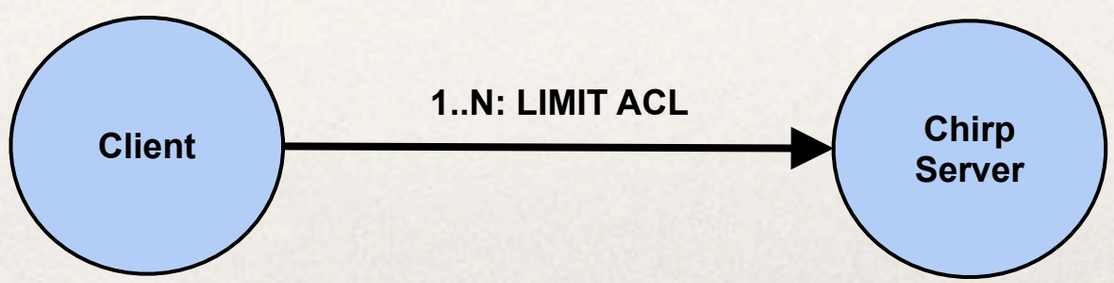

Since CCTools 3.4.0

Chirp: Using a Ticket

```
chirp -a ticket -i my.ticket chirp.nd.edu ls /data
```

or

```
parrot_run -a ticket -i my.ticket /bin/ls /chirp/chirp.nd.edu/data
```

Since CCTools 3.4.0

Chirp/Parrot: File Extended Attributes

```
parrot_run getfattr -n user.instrument sensor.dat
```

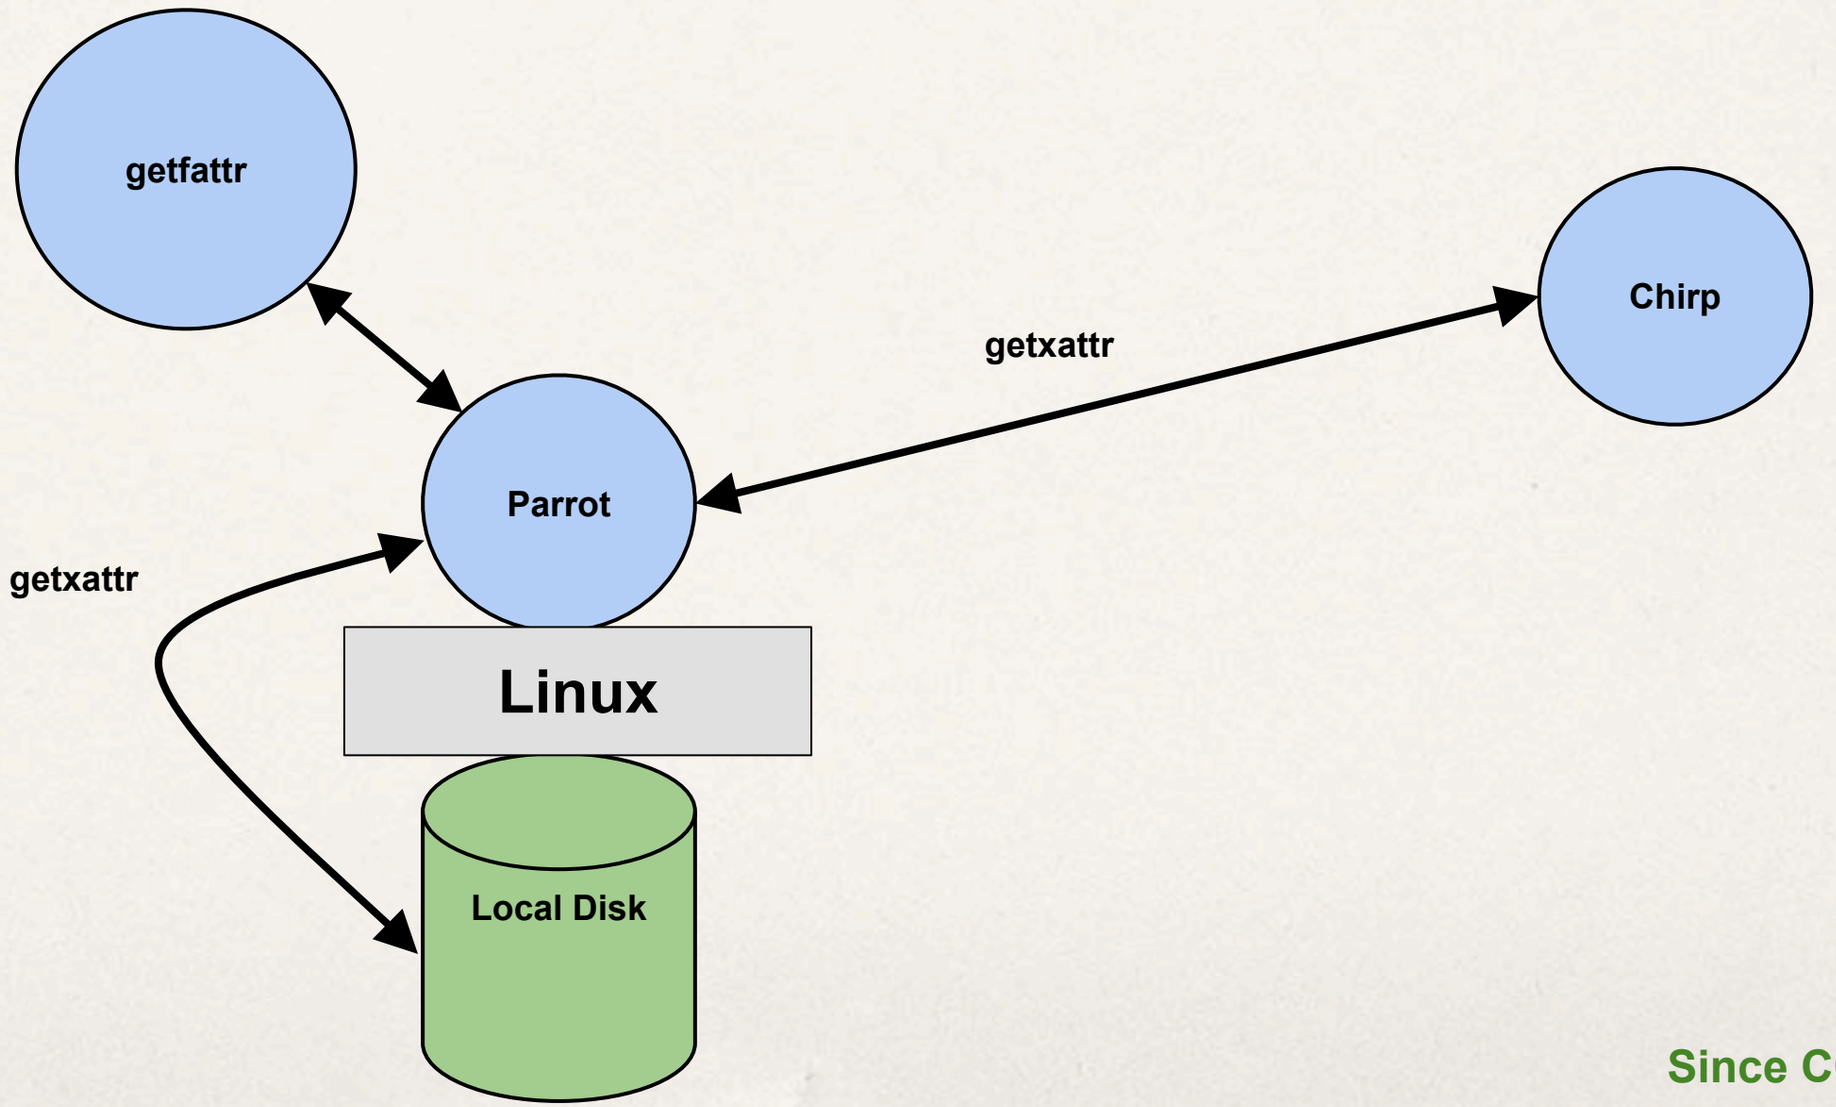

Since CCTools 3.6.0

Parrot and CVMFS

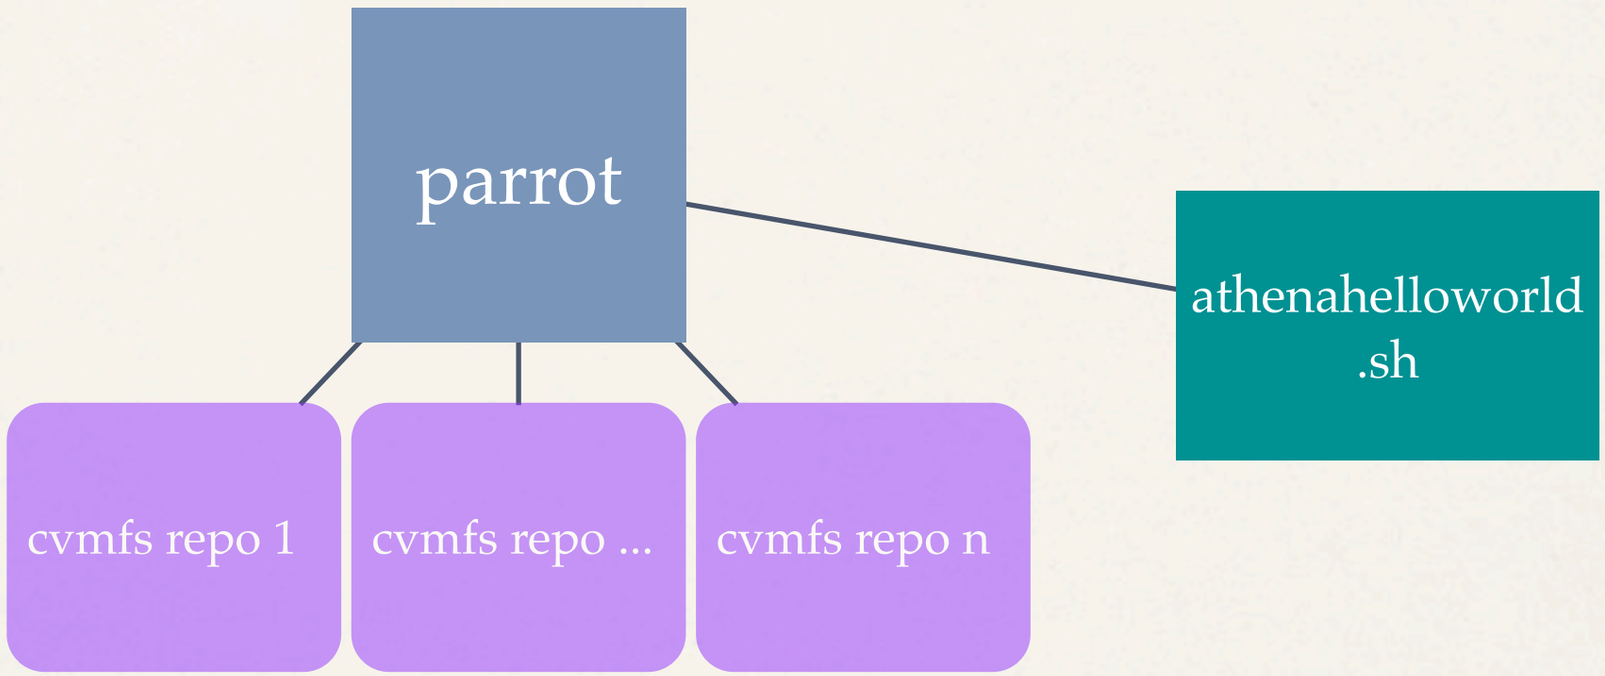

Parrot and CVMFS

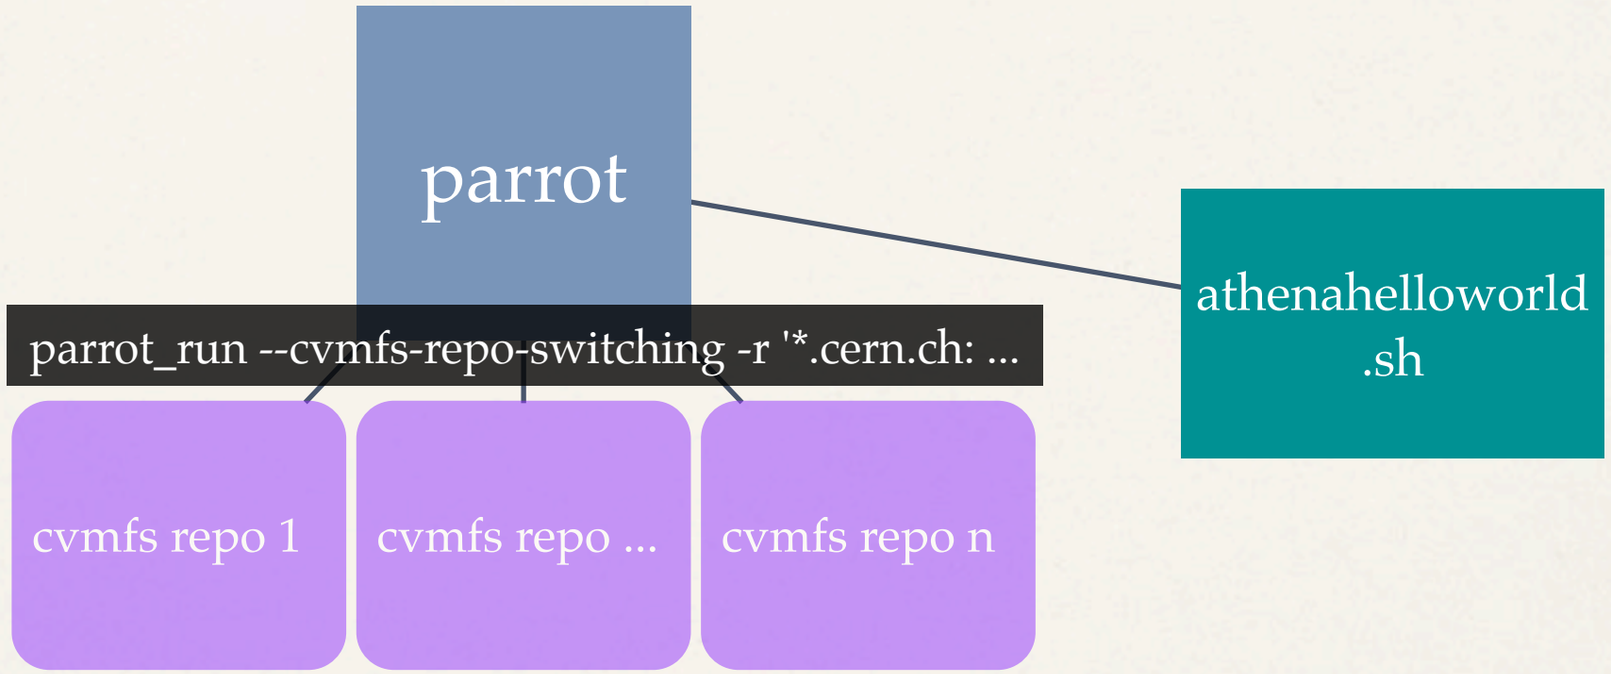

Parrot Search System Call

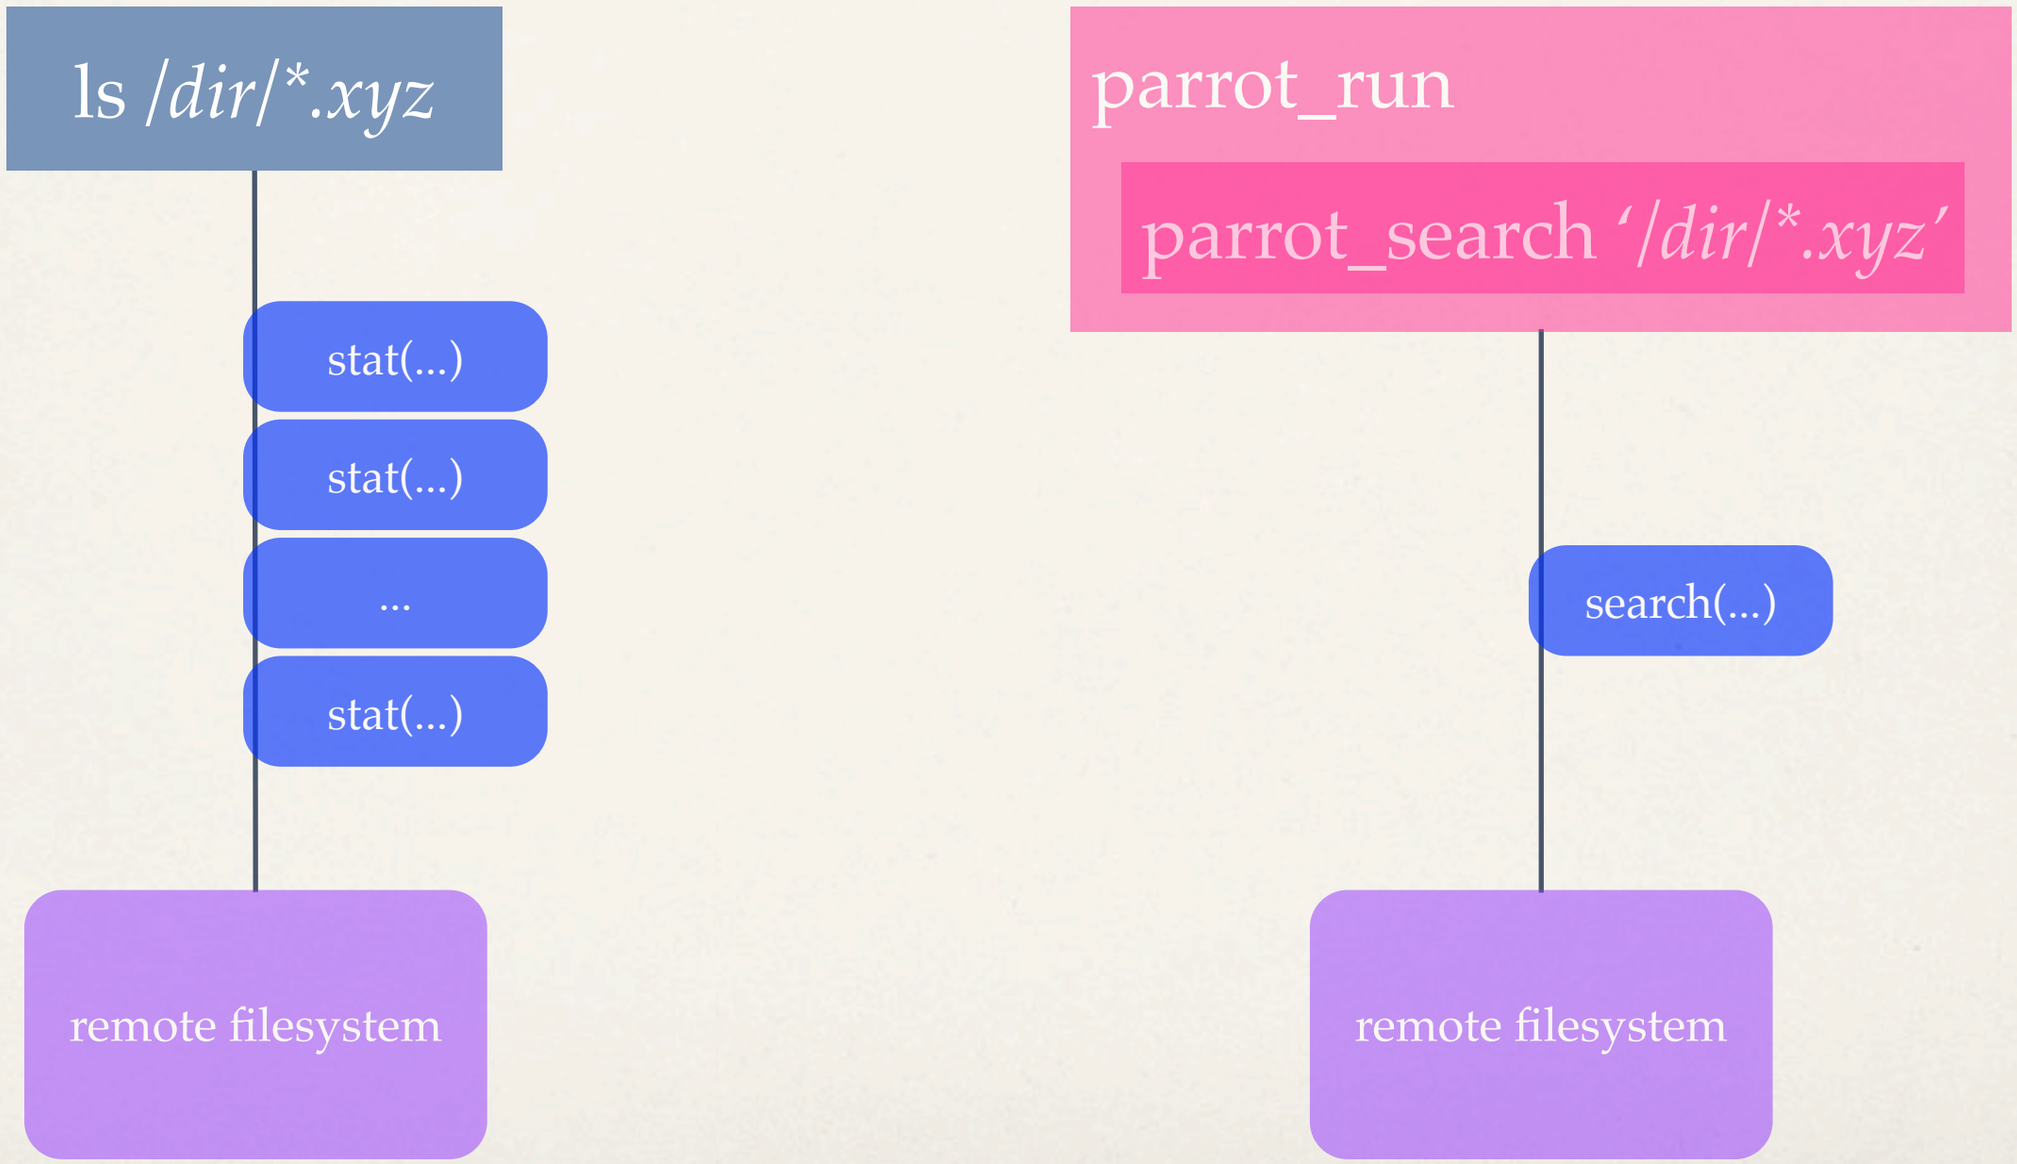

Development in Github

https://github.com/cooperative-computing-lab/cctools

git clone https://github.com:cooperative-computing-lab/cctools.git

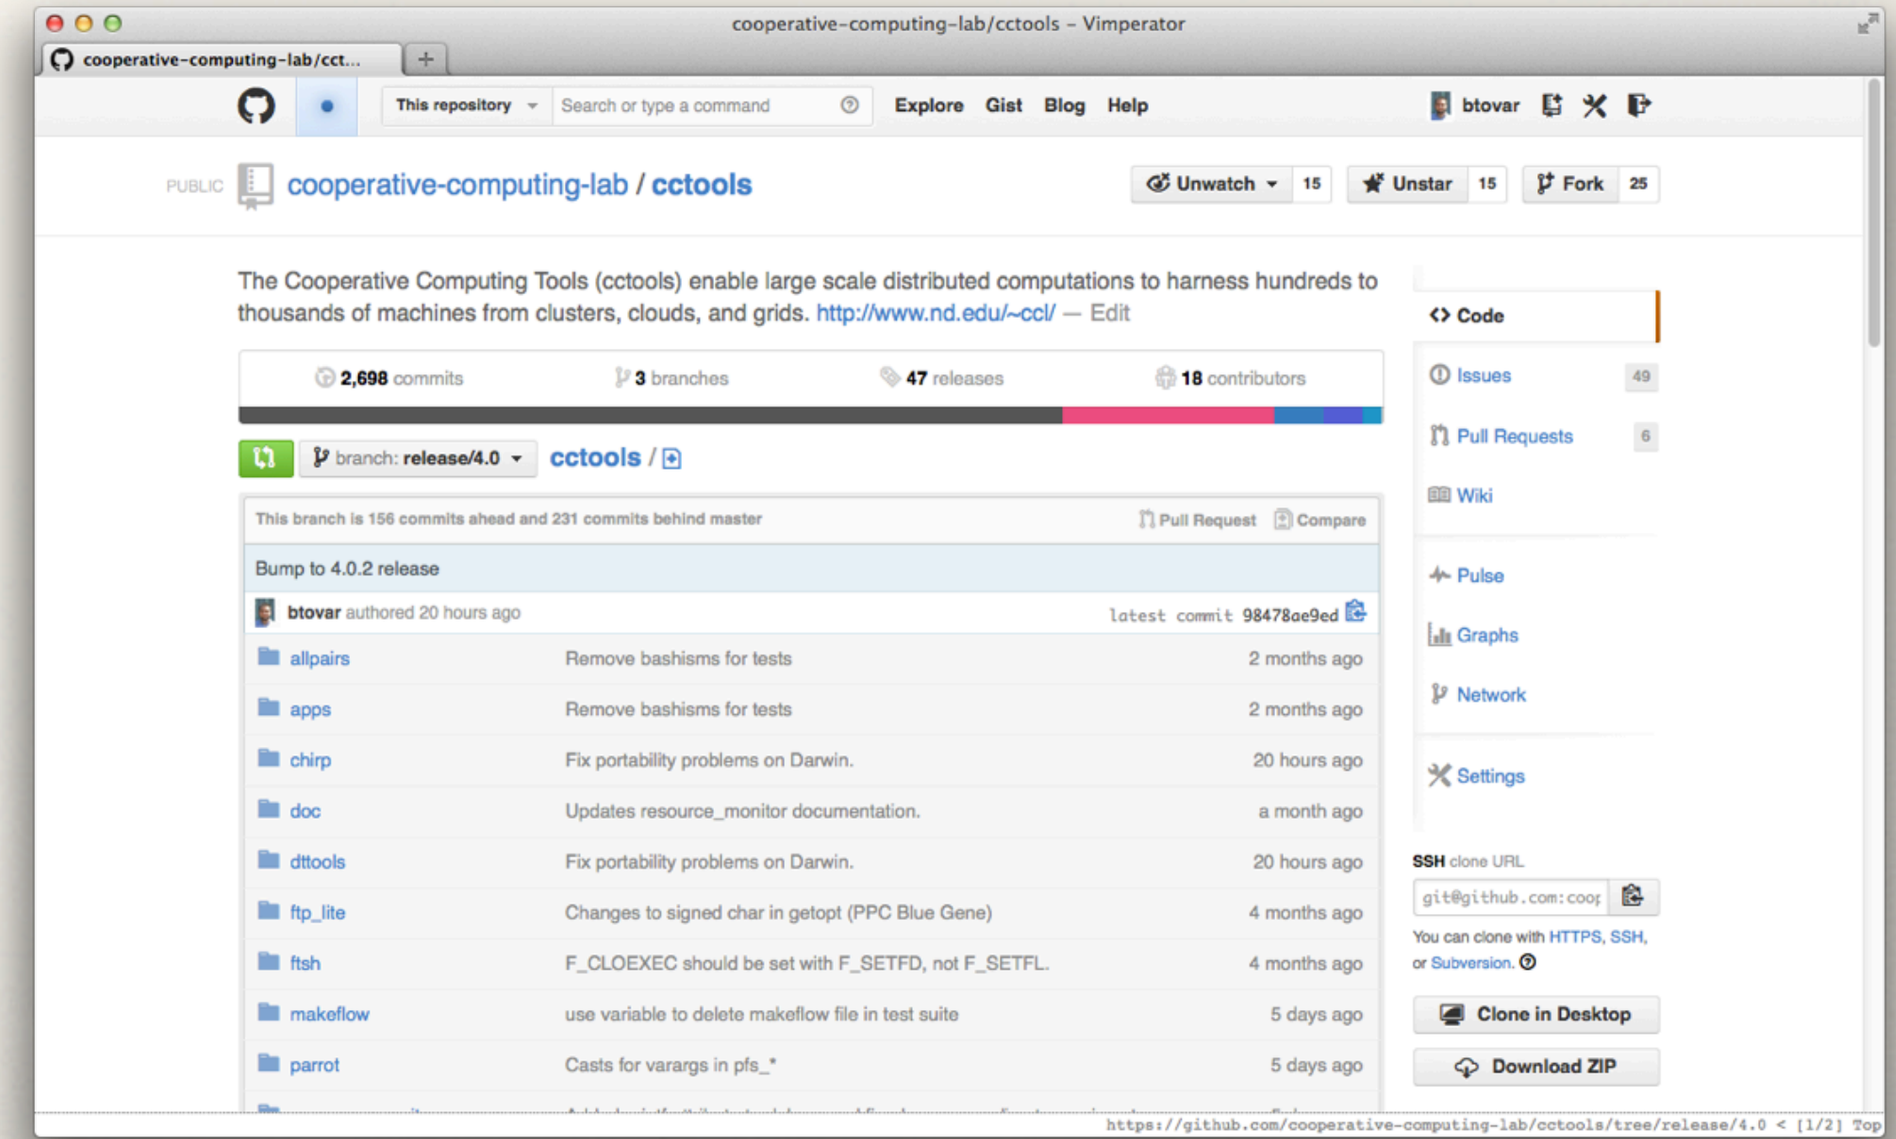

Thanks to All of the Contributors!

Michal Albrecht

DeVonte Applewhite

Neil Best

Brian Bockelman

Dan Bradley

Peter Bui

Iheanyi Ekechukwu

Patrick Donelly

Brian Du Sell

Brenden Kokoszka

Kyle Mulholland

Francesco Prelz

Dinesh Rajan

Casey Robinson

Peter Sempolinski

Douglas Thain

Andrew Thrasher

Benjamin Tovar

Li Yu