



SSH To Job: How it Works (and Sometimes Doesn't)

EU HTCondor Workshop 2022

Jaime Frey



Debugging Can Be Hard

- › Remote debugging is harder
 - Or at least time-consuming
- › Foreign environment
- › Long round-trip time



Solution: SSH To Job

- › A terminal session where your job is running
- › In the same environment as the job
 - Same environment variables, etc
 - In the same container instance, if any

```
[jfrey@submit-1 test]$ condor_ssh_to_job 1138.0
Welcome to slot1_24@e2565.chtc.wisc.edu!
Your condor job is running with pid(s) 799083.
[jfrey@e2565 ~]$
```

Make It Simple to Use

- › Want to use ssh programs
- › Don't want to require admin to configure ssh/sshd
 - Different user account on Execution Point
 - Firewalls
 - Host and user credentials
- › HTCondor can make a network connection with strong authentication
 - Used to run the job

SSH To Job Workflow

- › Ask schedd for capability to EP
- › Connect to starter on EP and present capability
- › Generate and exchange ssh credentials with starter
 - Ssh_to_job writes private user key and known_hosts file
 - Starter writes server key and authorized_users file
- › Launch ssh and sshd and hand them the existing network connection
 - They communicate and establish ssh session

Complications

- › Sshd can inherit socket from parent
- › Ssh client insists on initiating new network connection
- › ProxyCommand option
 - Tells ssh to run arbitrary command that establishes connection
- › We provide proxy command
 - Connects to ssh_to_job process
 - Receives FD of socket connection to starter/sshd
 - Proxies data

More Complications

- › Sshd reads /etc/passwd for executing user account info
 - User account may be absent
 - User's login shell may be /bin/nologin
- › We wrote a small shared library
 - Defines alternate getpwnam() that returns a valid struct passwd with usable login shell
- › We use LD_PRELOAD to intercept sshd's getpwnam() calls

It Works Most of the Time

- › Some failures of `ssh_to_job` will require admin attention

Sshd Is Missing

- › Common when running job in a container
- › Add sshd to container image

AP Firewall and CCB

- › When EP is using CCB, network connection goes from EP to AP
- › User processes can't use HTCondor's shared port by default
 - Ssh_to_job will bind to dynamic TCP port, which firewall may block
- › Create directory `/var/lock/condor/daemon_sock`
 - Enable write permissions for HTCondor users

Clever Firewalls

- › Firewall with Deep Packet Inspection
 - May be configured to block ssh on all ports
- › Talk to your network admin

Job Wrapper Script

- › Wrapper script places job in a container
 - Ssh session can end up in separate, identical container
- › See earlier talk about avoiding wrapper scripts
- › Have HTCondor launch the job container
 - EP will launch the sshd in the same container instance

Thank You!



This work is supported by NSF under Cooperative Agreement OAC-2030508 as part of the PATH Project. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the NSF.