



Collector Tour

EU HTCondor Workshop 2023

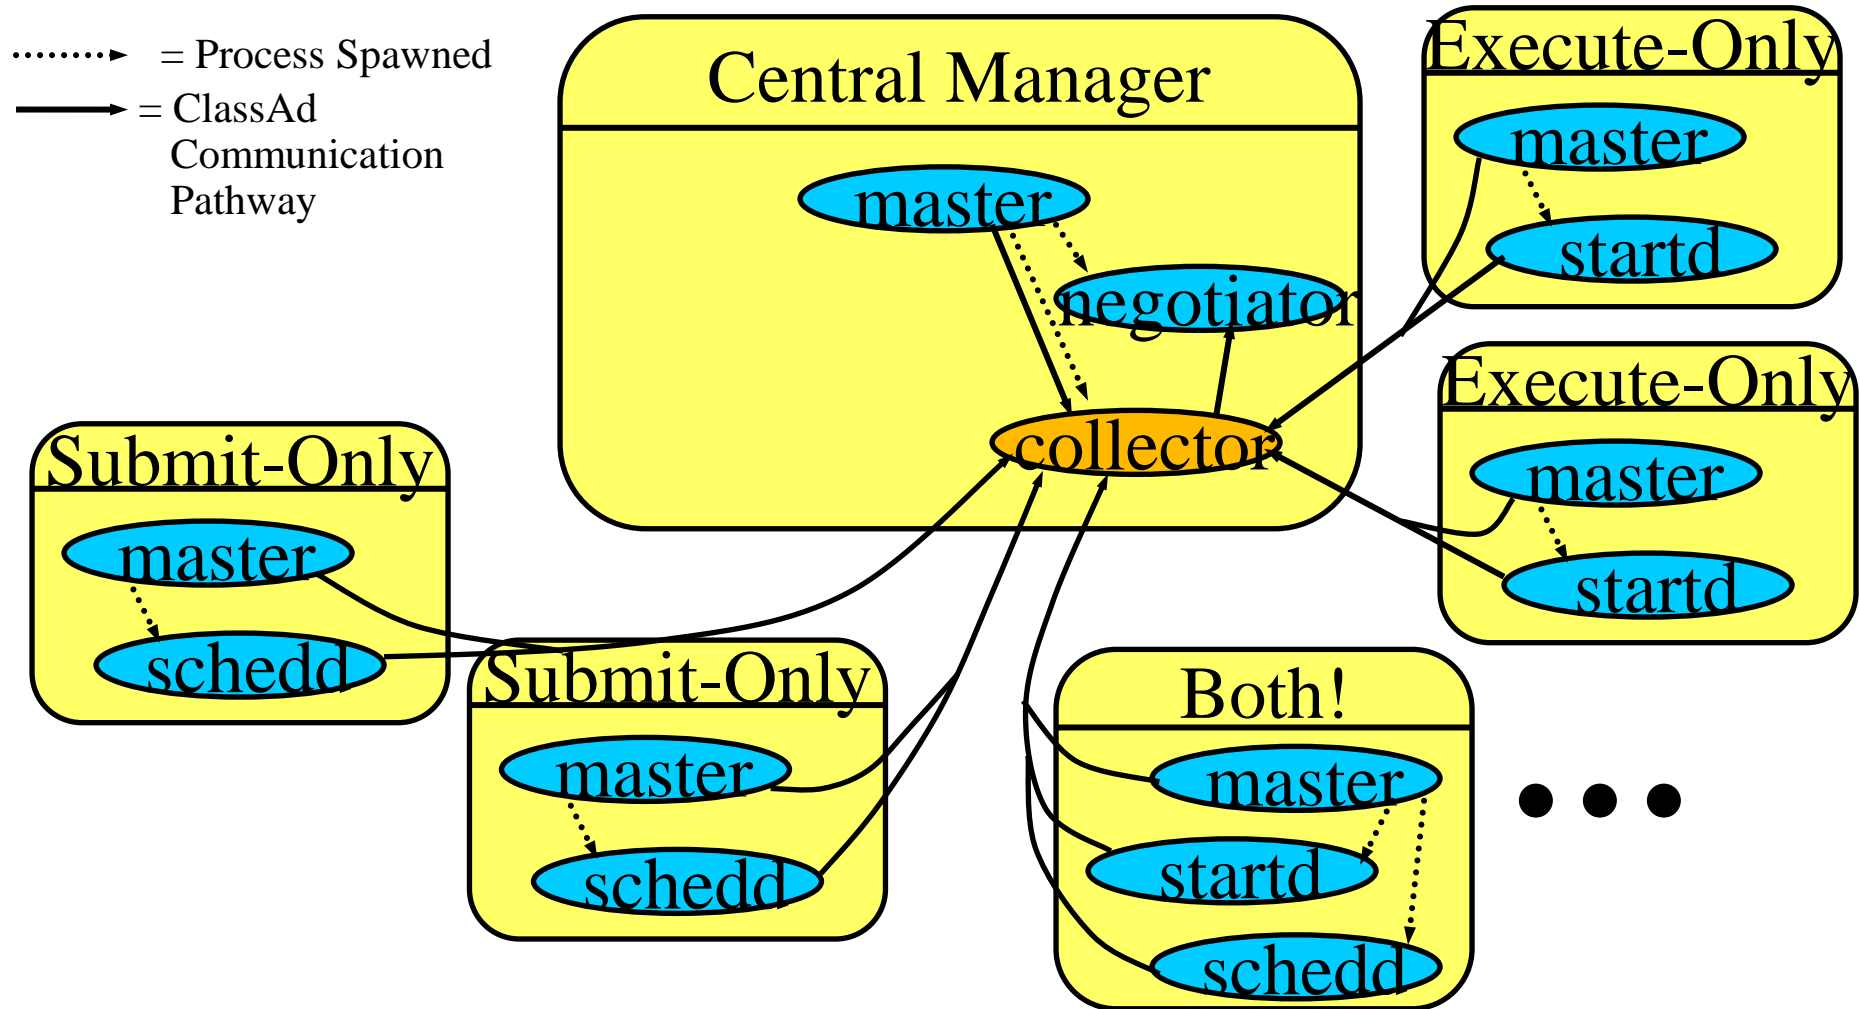
Jaime Frey



Collector Defines an HTCondor Pool

- › Single fixed location in the network
- › Collects information about the pool
 - Receives ClassAds of entities within the pool
- › Decides who can participate in the pool
 - Who can advertise and who can query
- › To join the pool, a daemon must...
 - Know the network location of the collector
 - Gain the trust of the collector

Layout of a General HTCondor Pool



Roles of the Collector

- › Root of all trust
 - Daemons and tools don't need preexisting trust with each other
 - Just trust in the collector
- › Enabler of operations
 - Matchmaking
 - Pool administration
 - Network firewall traversal
- › Monitoring
 - Collector contains data on pool activity and health

Collector is Stateless (mostly)

- › Daemons send regular updates of ads
- › Collector removes ads with no recent update
- › If a daemon, crashes, it is dropped from the pool automatically
- › If the collector restarts, full data of pool is re-established after a short period

Slot Ads

- › Sent by condor_startd
- › Default output of condor_status
- › Multiple ads per EP
 - Coming soon: new startd ad type for whole EP
- › Used for matchmaking
- › Shows available resources
- › Shows details of running job, if any

Submitter Ads

- › Sent by condor_schedd
- › One per owner/group with queued jobs
- › Used by negotiator to know who has jobs that need matching
- › Shows totals of idle, running, held jobs

Schedd Ads

- › Sent by condor_schedd
- › Stats on schedd activity and health
 - RecentDaemonCoreDutyCycle is most useful stat
 - Totals of jobs in various states
 - Totals of file transfers both queued and active

Master Ads

- › Sent by condor_master
- › Shows basic machine information, daemon status
- › Locate master to send on/off/restart/reconfig commands

Accounting Ads

- › Sent by condor_negotiator
- › Resource usage and limit info for all users/groups
- › Used to decide how to distribute pool's resources

Grid Ads

- › Sent by condor_gridmanager
- › One per <user, remote scheduler> for grid universe
- › Shows job totals
- › Shows problems communicating with remote scheduler

Collector Ads

- › Sent by condor_collector
 - Sent to itself and potentially other collectors
- › Stats on collector activity and health
 - RecentDaemonCoreDutyCycle is most useful stat
 - Totals of ads received
 - Stats of CCB activity

Generic Ads

› Create new ad types with condor_advertise

```
• % condor_advertise UPDATE_AD_GENERIC - << EOF
  ? MyType="Bar"
  ? Name="Me"
  ? Abc=14
  ? EOF
```

```
Sent 1 of 1 ad to nanya.lan:9619.
```

```
% condor_status -subsystem Bar
```

MyType	TargetType	Name
Bar	None	Me

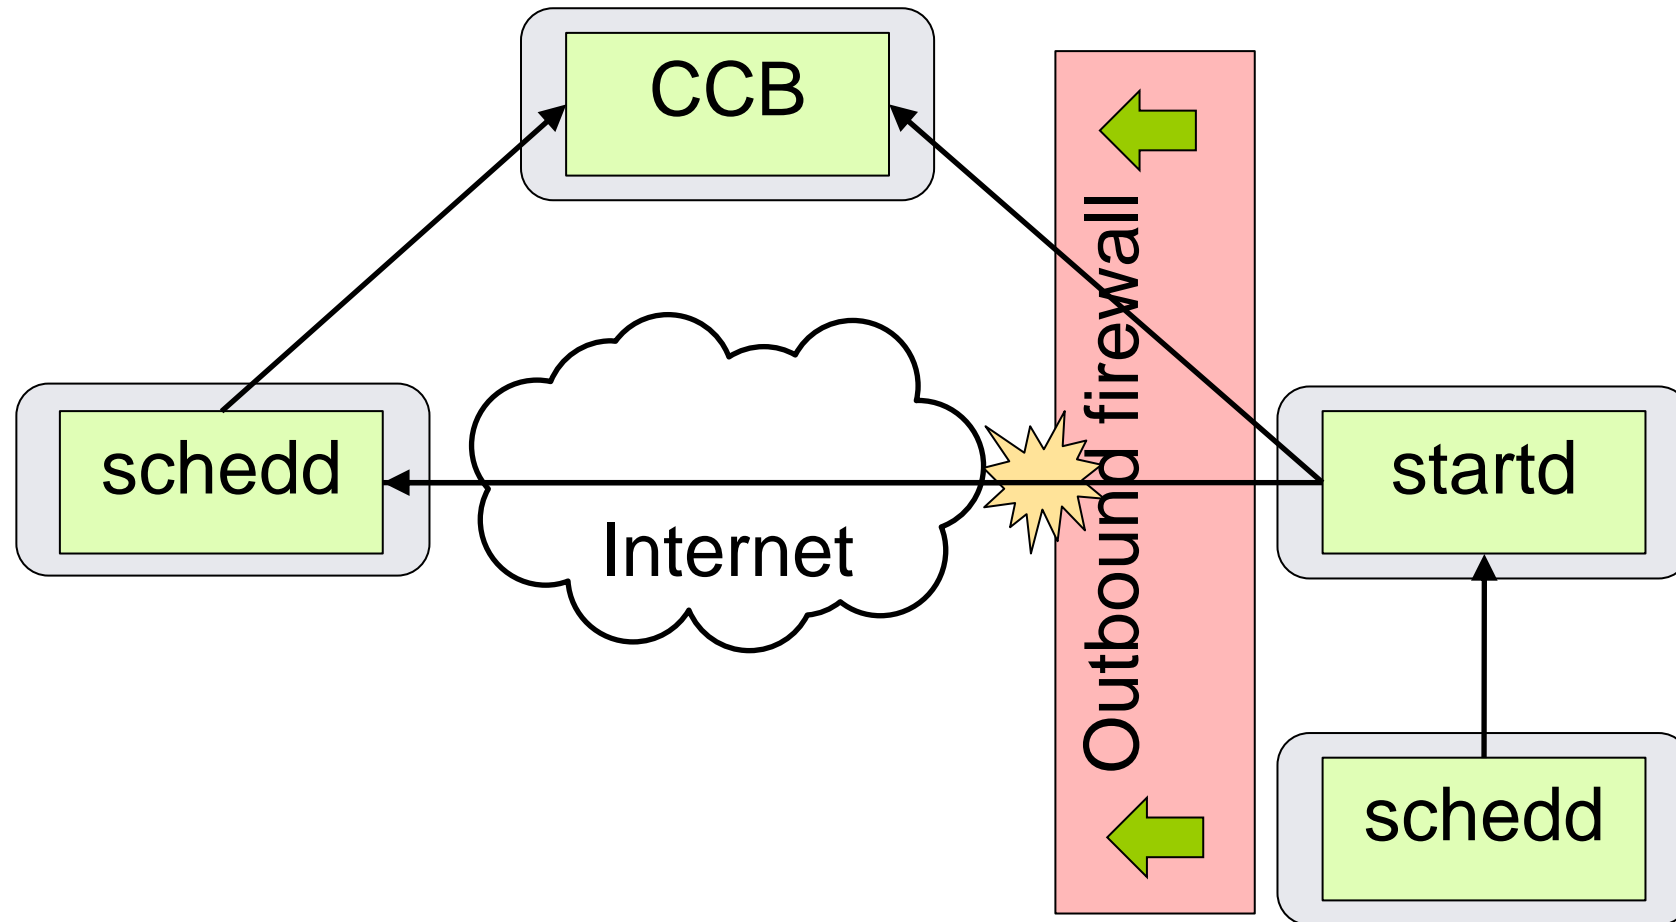
Authentication Levels

- › DAEMON
 - Collector, Generic, Grid
- › ADVERTISE_MASTER (default DAEMON)
 - Master
- › ADVERTISE_SCHEDD (default DAEMON)
 - Schedd, Submitter
- › ADVERTISE_MASTER (default DAEMON)
 - Master
- › NEGOTIATOR
 - Negotiator, Accounting

Condor Connection Broker (CCB)

- › Allows EPs behind a firewall to join a public pool
- › TCP connections to EP are reversed
 - Collector used as a control channel
- › Limitations
 - Collector and AP must be on public internet
 - Firewall must allow outbound connections

CCB: Condor Connection Broker



High Availability

- › Set of primary collectors
- › Advertisers send to all collectors
- › Queriers consult collectors in series
 - Avoid unresponsive collectors in future

Spreading the Load

- › Create secondary collectors to handle some duties
- › Collector dedicated to CCB
- › View collector
 - Primary collector forwards updates to view collector
 - Monitoring queries directed to view collector

Offline Ads

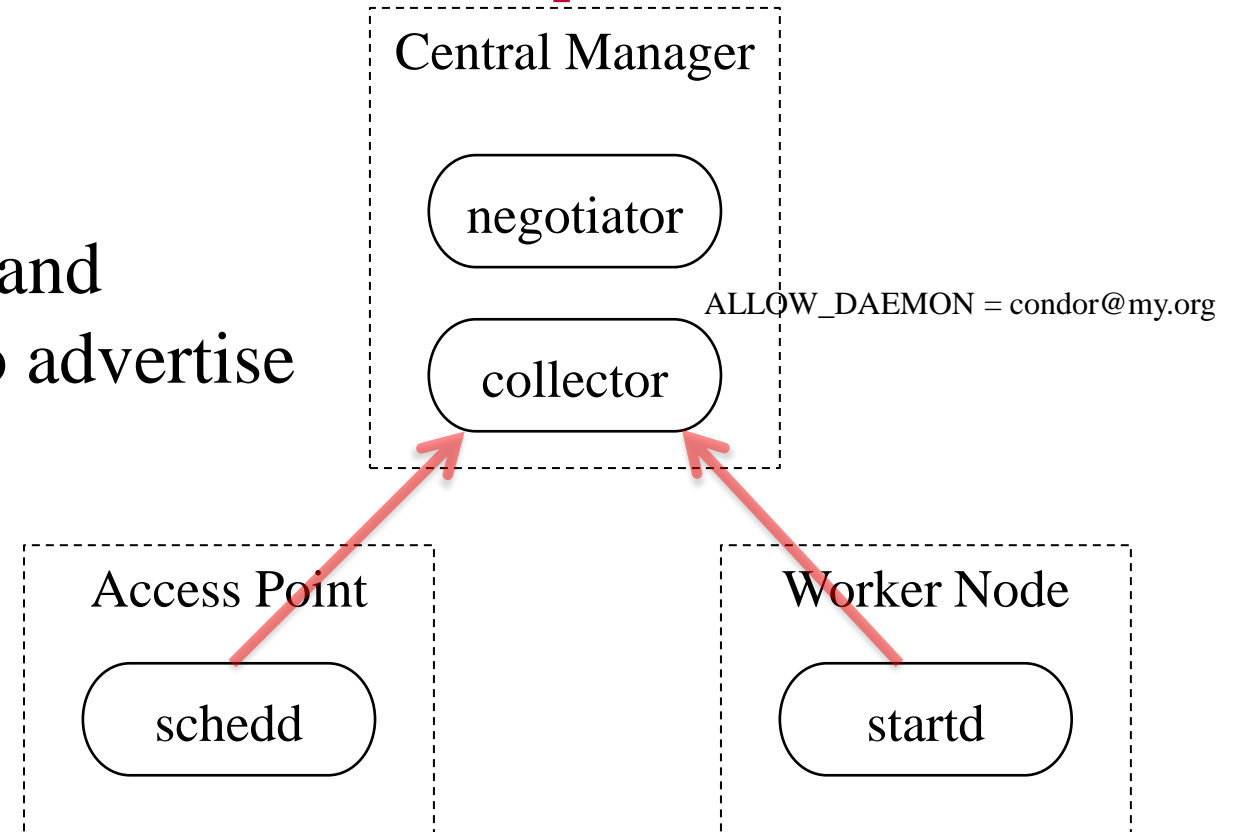
- › When an EP hibernates, it marks its ads as Offline
- › Offline ads don't expire
- › Offline ads are saved to disk and restored on a collector restart
- › Considered for matchmaking

Absent Ads

- › When an EP stops updating, collector can mark its slot ads as Absent
- › Absent ads only expire after a month
- › Absent ads are saved to disk and restored on a collector restart
- › Hidden from queries unless -absent is requested

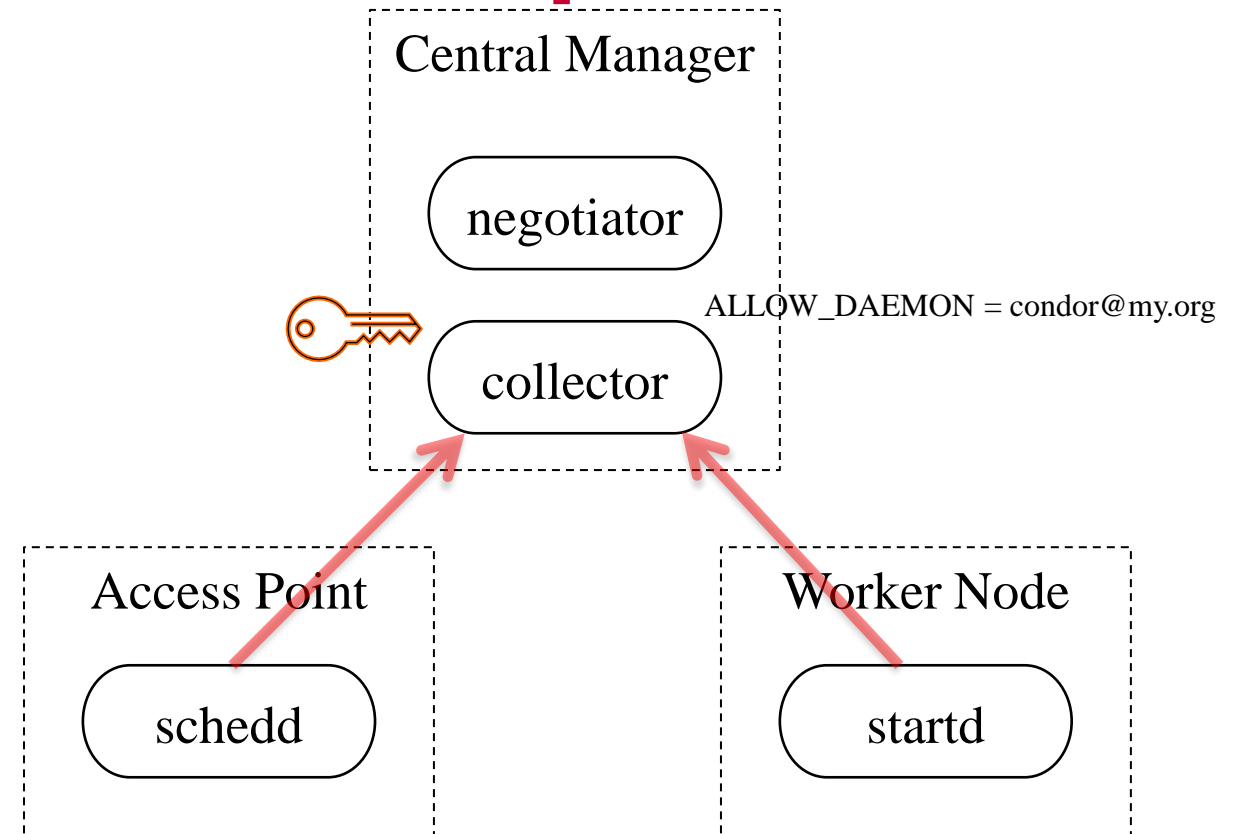
IDTOKENS to secure a pool

The key to a secure pool is to only allow trusted (authenticated and authorized) startds and schedds to advertise into the collector.



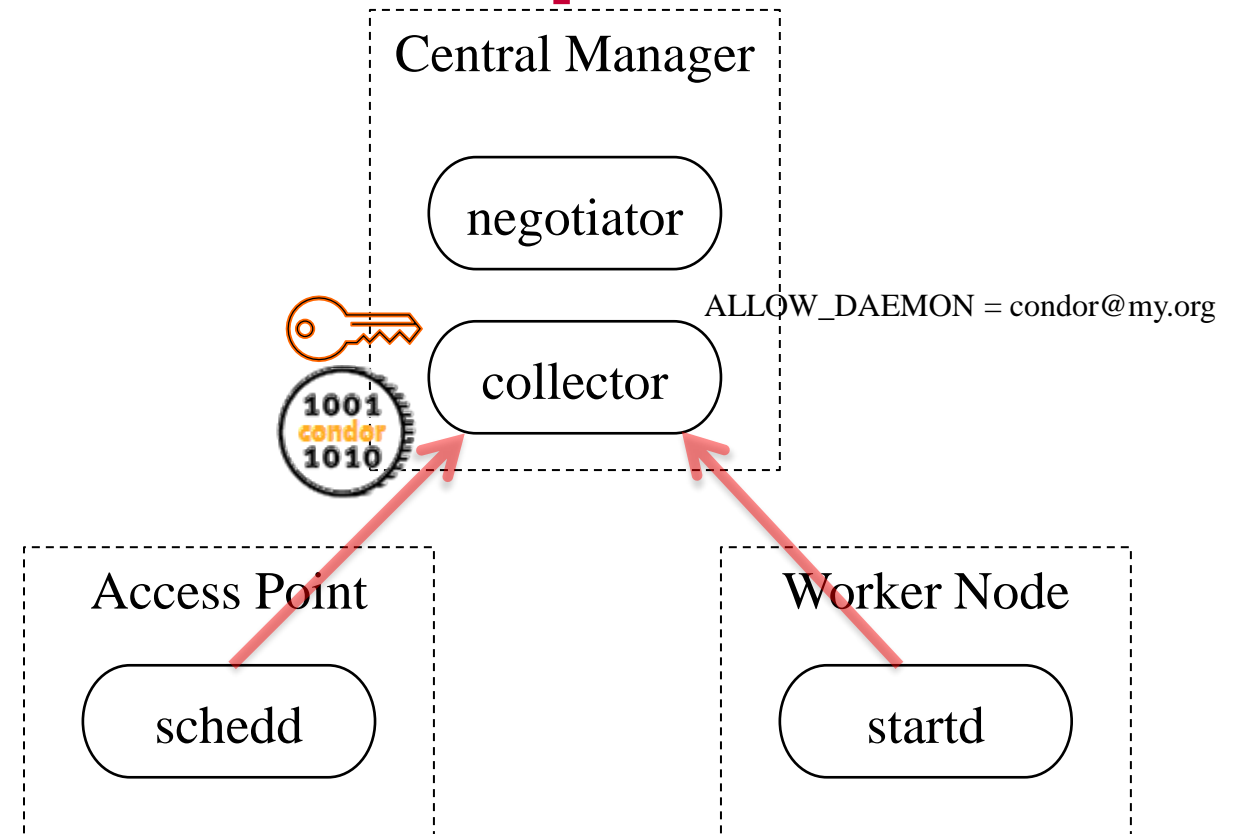
IDTOKENS to secure a pool

- › 1. Create a signing key file on the central manager
 - A random key will be created by default at collector startup, or explicitly create with tool:
`condor_store_cred add -c`
 - Signing key is stored by default in file
`/etc/condor/passwords.d/POOL`



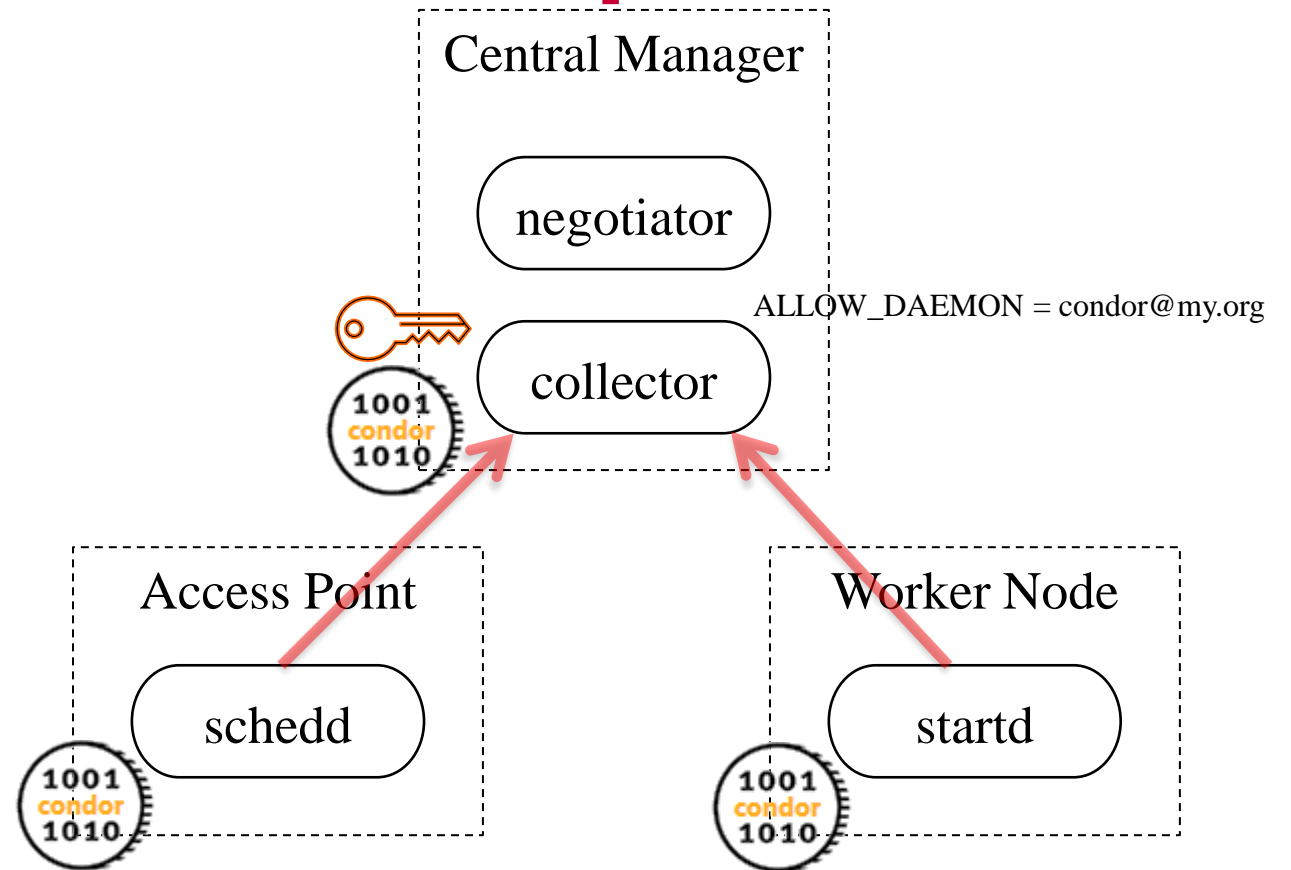
IDTOKENS to secure a pool

- › 1. Create a signing key file on the central manager
- › 2. Create an IDTOKEN file with identity "condor" on central manager
 - Use tool
`condor_token_create -identity condor@cm.myorg`



IDTOKENS to secure a pool

- › 1. Create a signing key file on the central manager
- › 2. Create an IDTOKEN file with identity "condor" on central manager
 - Use tool
`condor_token_create -identity condor@cm.myorg`
- › 3. Copy this IDTOKEN file to each trusted server you want to join your pool
 - Place file into directory
`/etc/condor/tokens.d`



Daemon Capabilities

- › Daemons include a capability when advertising to the collector
 - Similar to a token, but in-memory and daemon-specific
 - Anyone possessing the capability can issue commands to the daemon
- › Collector will give capability to trusted queriers
- › Allows daemons to communicate with no preexisting trust
 - Trust is vested in the collector

Daemon Capabilities

- › During matchmaking
 - Negotiator receives capabilities for schedds and startds
 - Negotiator sends startd capability to schedd for each match
- › Admin tools
 - Tool receives capabilities for daemons it will act on

Thank You!

PATH PARTNERSHIP to ADVANCE
THROUGHPUT
COMPUTING

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.