

**Federating  
pools – from Clusters to  
Grids**

Greg Thain

# Agenda

Overview of the Condor-CE

basis for GlideinWMS and OSG

- What's a CE
- Why we need one
- Past history
- Possible future work

# First, some terminology

- › CE “Compute Element”
- › SE “Storage Element”
- › WN “Worker Node:

# The Problem

- › You have many pools, all different
  - Different batch systems
  - Different policy
- › Want to use all of them
- › Job sender shouldn't need to know about site differences



# Solution: A common “Door”


- › Door translates common requests into site-specific requests



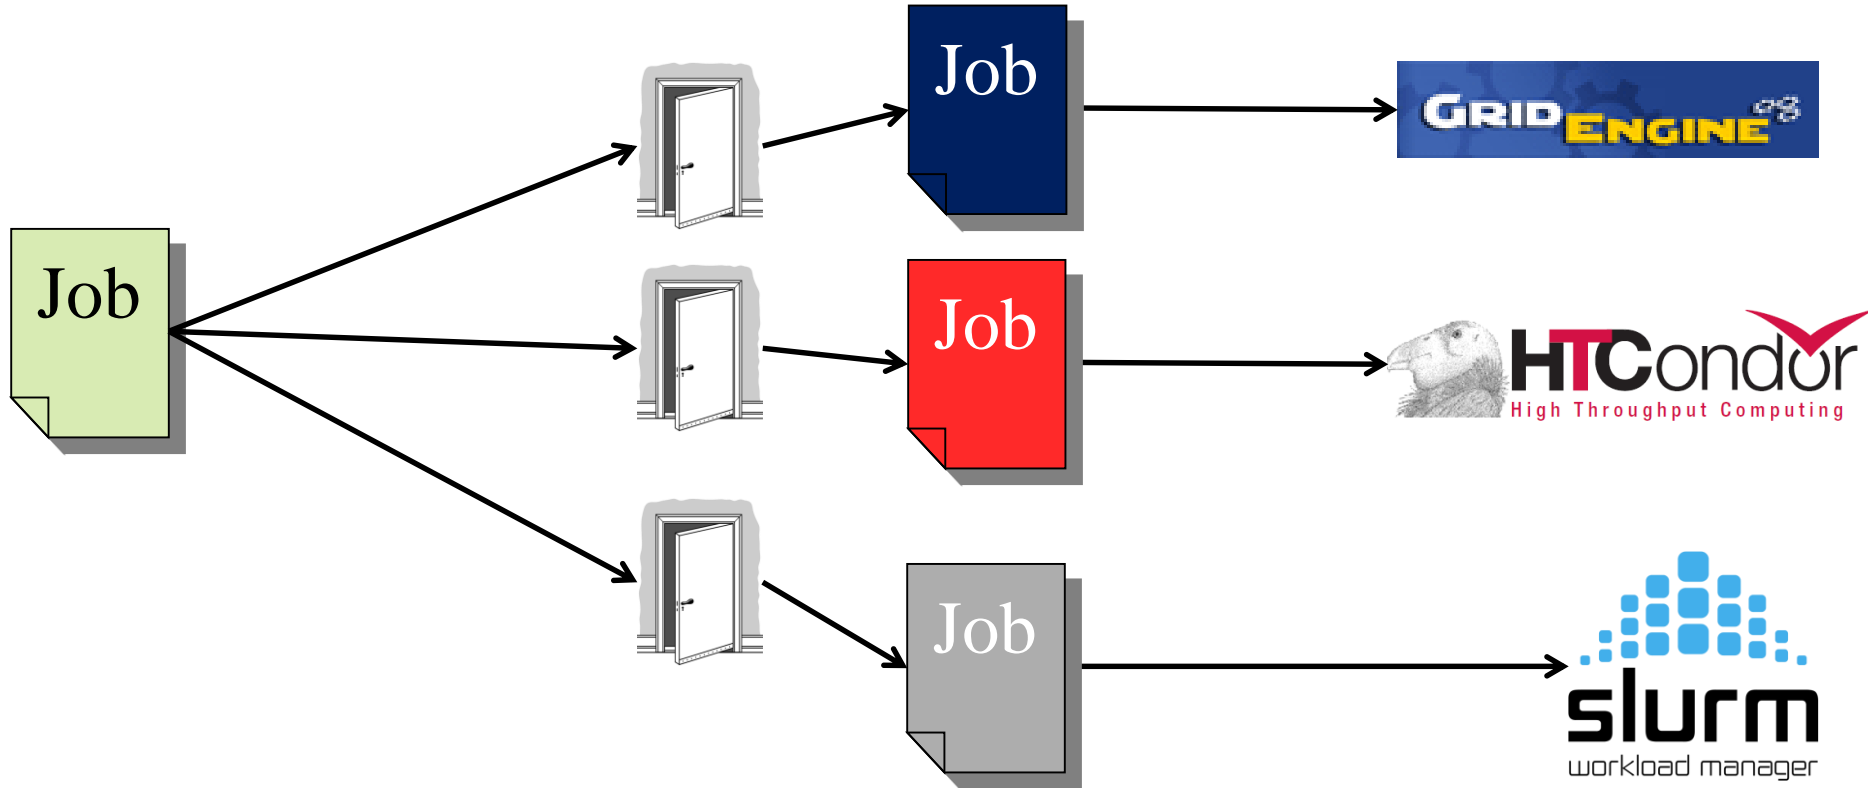
Door applies site-specific policy to common job



Door added to sites



# Solution: A common “Door”



# Door Principles

- › Need powerful transformations
- › Need to hide site dependencies
- › Want language to be **declarative**
  - **Terse, Analyzable, Usable**
- › Need all work to be **Manageable**

# Focus on non-functional requirements

- › Doors must be
  - Scalable
  - Manageable
  - Accountable
  - Secureable
  - Extensible



# Fundamental Principle

- › “Submit locally, Run Globally”
  - Same job description, local or grid
  - Global grid is not a separate system, it **extends** the existing system

# The 2000s solution to Doors



Globus GRAM

In 2017, Globus announced dropping support for GRAM  
Required new implementation

Allow us to learn from lessons of GRAM

# GRAM lessons learned

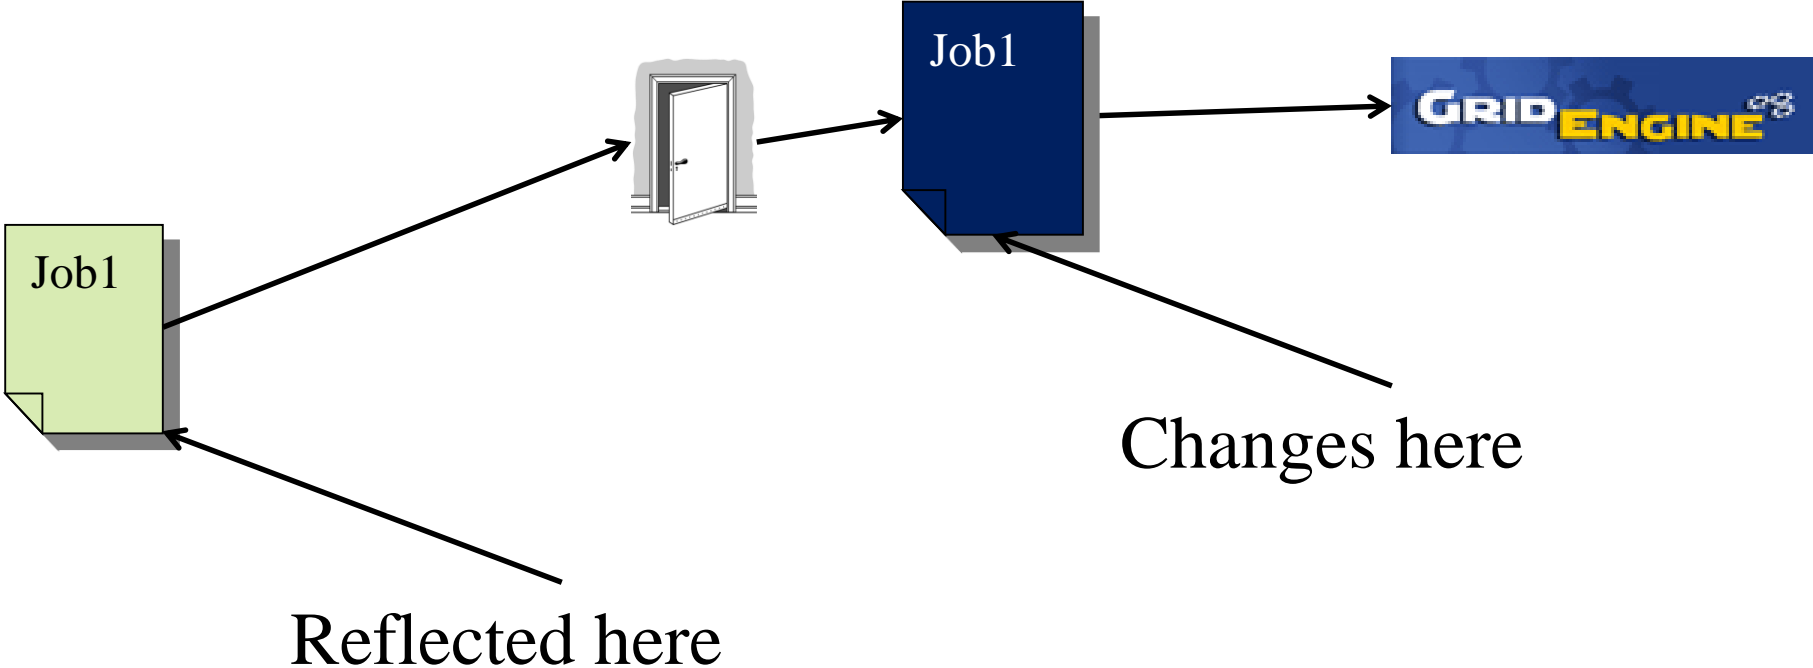
- › GRAM imperative, not declarative
  - Adaptation provided by perl scripts
    - Modified by every local site
    - Upgrade nightmares
- › Scalability problems
- › Accounting and manageability problems

# HTCondor to the rescue

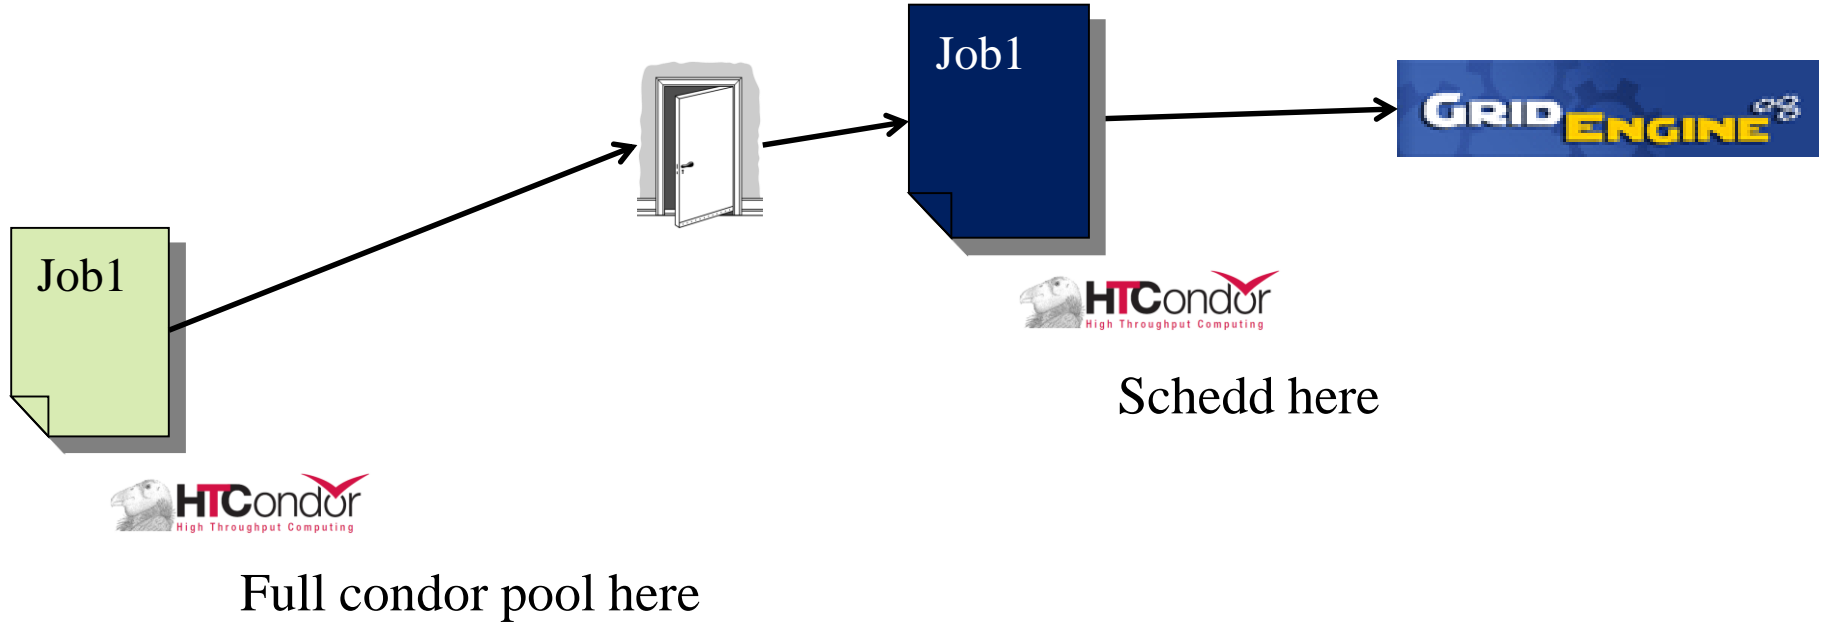
- › With HTCondor tools described yesterday...
  - Condor-C
  - Job Router
- › We built Condor-CE with existing parts



# One Logical job, in many places

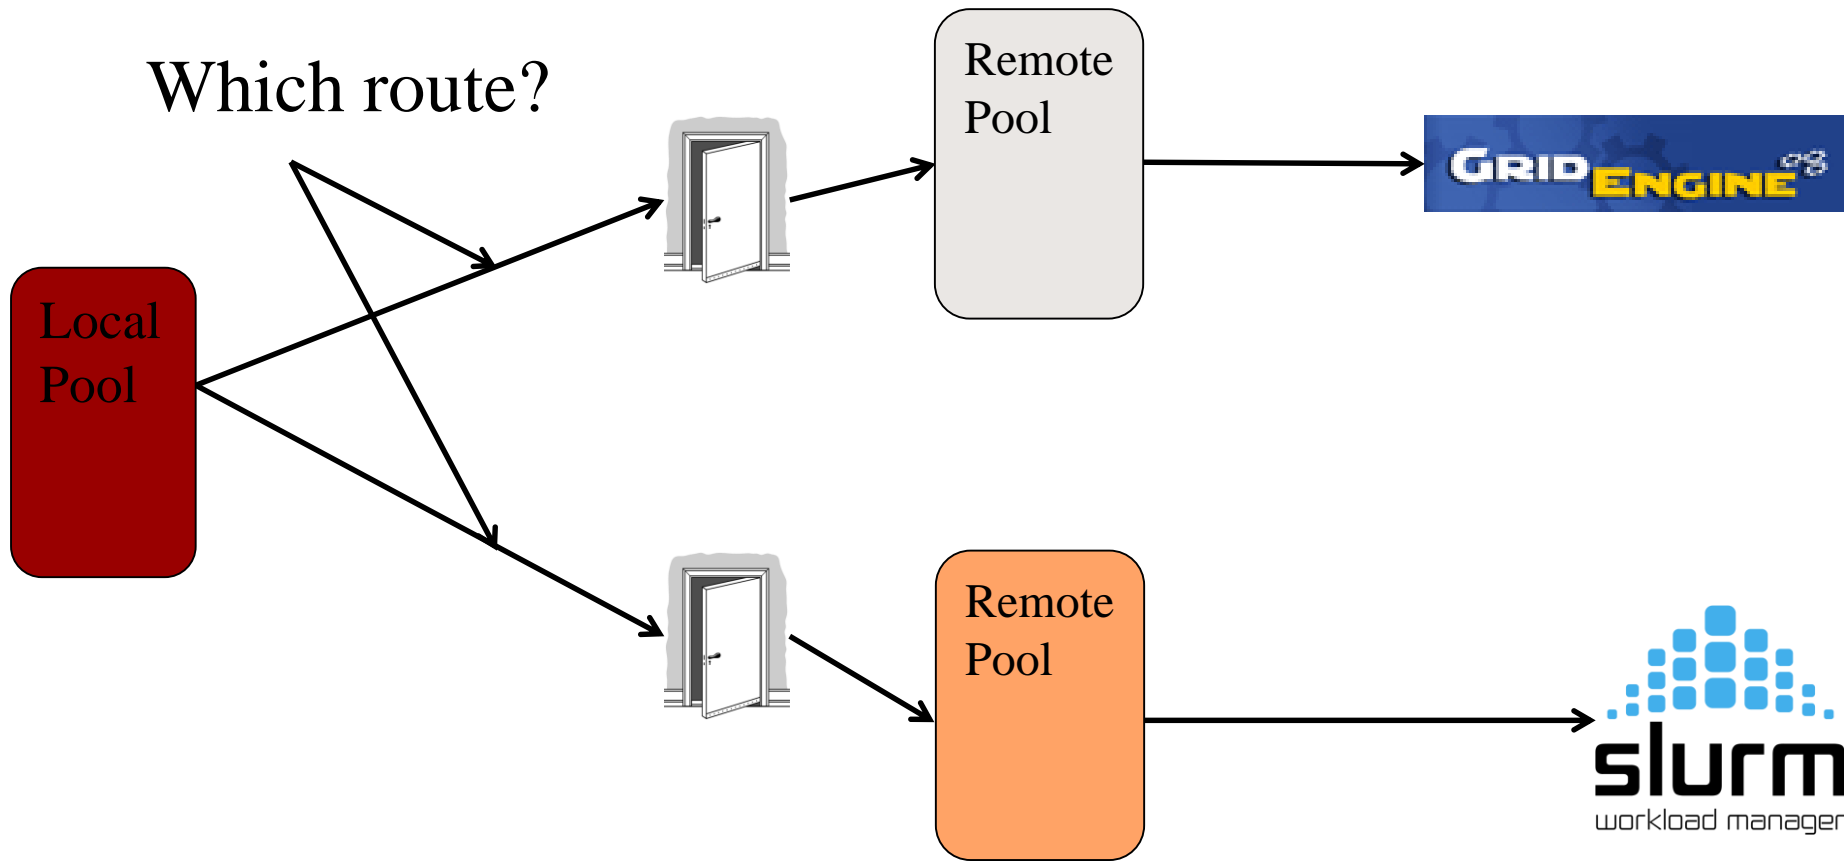


# Condor power on both sides




# All is not well: source routing

Which route?

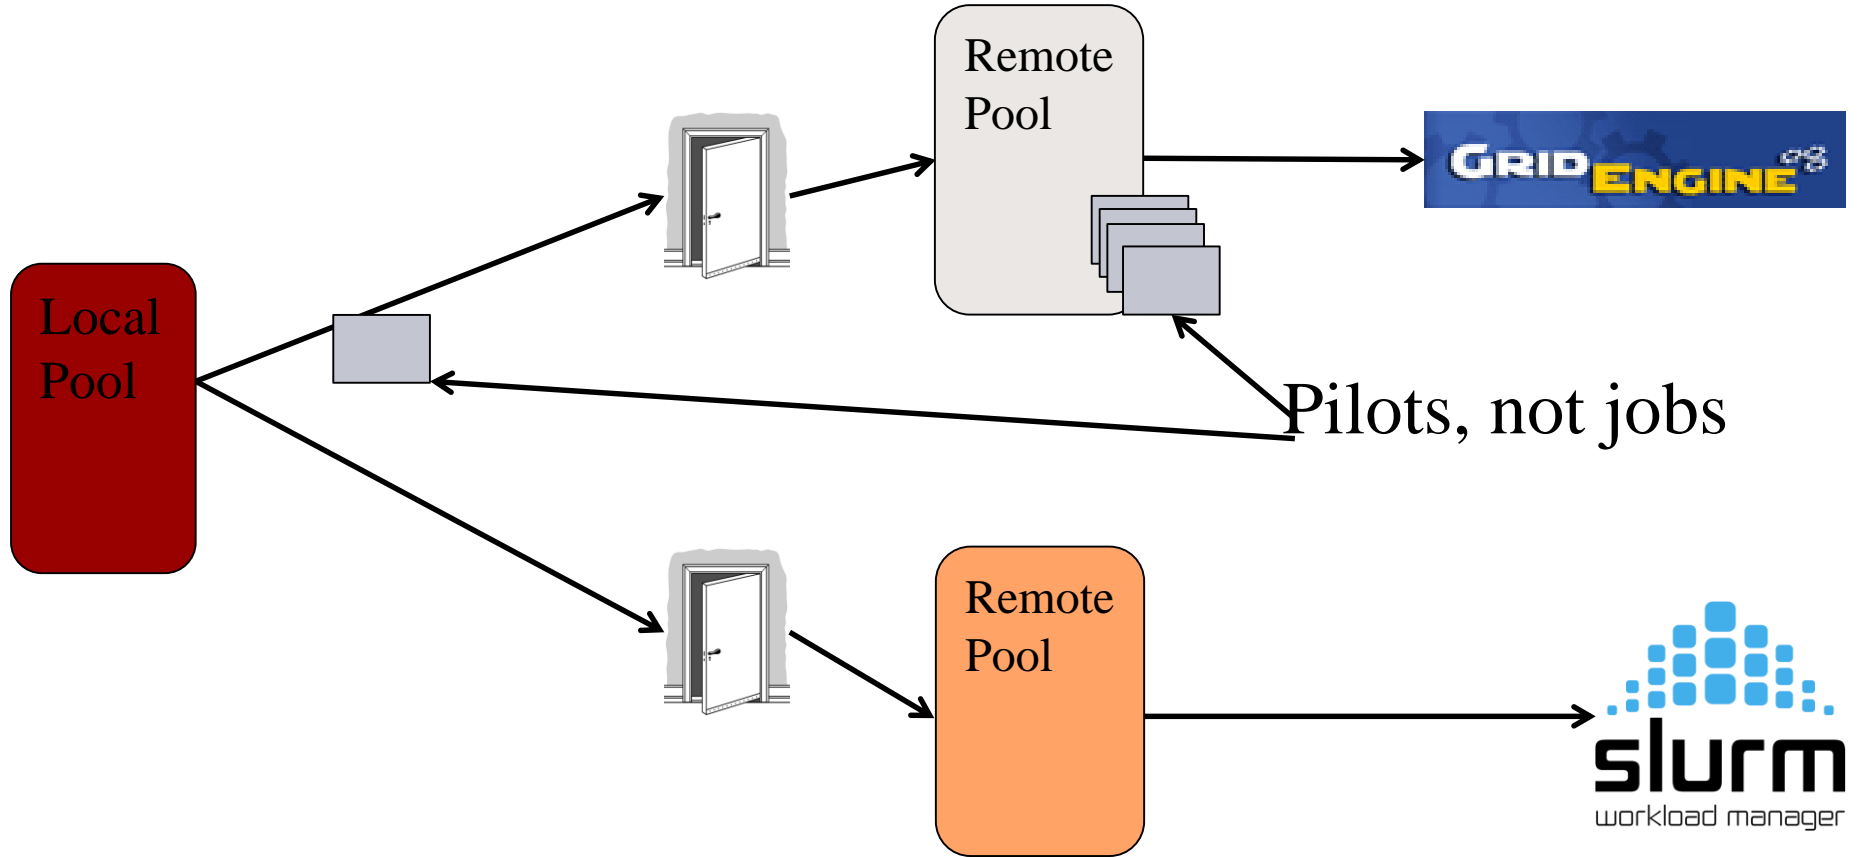







# Wait while idle...



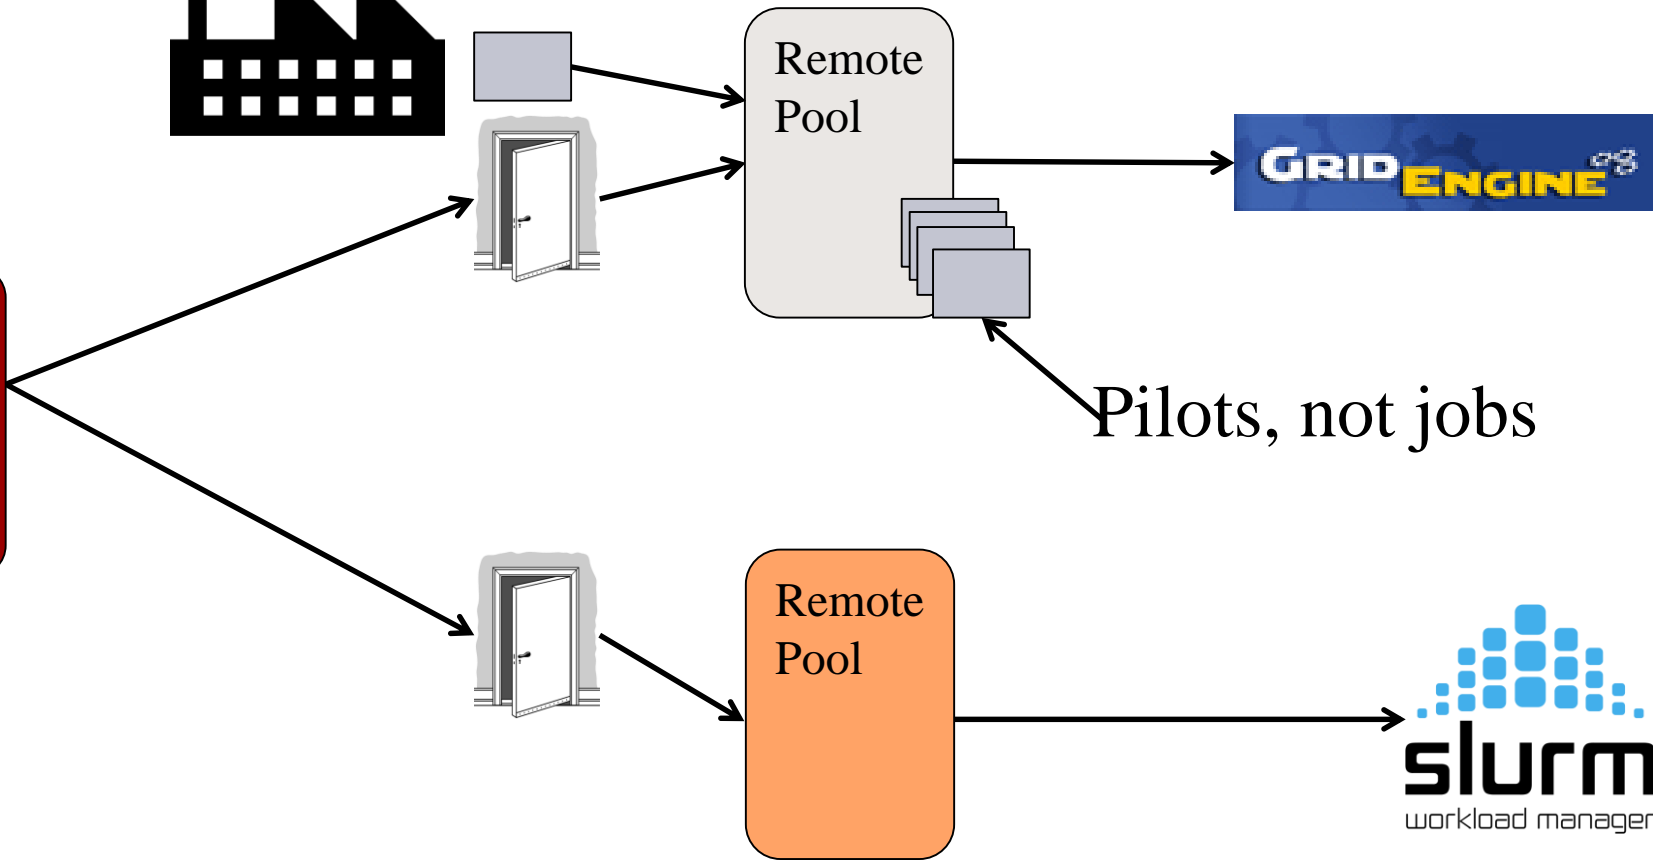



# Solution: Late binding



# Factory submission



Pilots, not jobs



# kubernetes

- › Kubernetes
  - Could be just another grid type
- › allows interesting possibilities
- › Runs containers
- › Runs as root
- › Provisions its own network



# Conclusion

- › Condor-CE allows us to build grids
- › Built on years of experience in months