On-the-fly Capacity Planning in Support of High Throughput (HT) Workloads

Miron Livny
John P. Morgridge Professor of Computer Science
Wisconsin Institutes for Discovery
University of Wisconsin-Madison

The observations I will share with you today are based on broad (and long) experience with real-life High Throughput Computing (HTC) Workloads

I hope to motivate you to work with us on what we do not have while leveraging what is available and works

Perspectives on Grid Computing (2010)

Uwe Schwiegelshohn Rosa M. Badia Marian Bubak Marco Danelutto
Schahram Dustdar Fabrizio Gagliardi Alfred Geiger Ladislav Hluchy
Dieter Kranzlmüller Erwin Laure Thierry Priol Alexander Reinefeld
Michael Resch Andreas Reuter Otto Rienhoff Thomas Rüter Peter Sloot Domenico Talia Klaus Ullmann Ramin Yahyapour Gabriele von Voigt

We should not waste our time in redefining terms or key technologies: clusters, Grids, Clouds... What is in a name? Ian Foster recently quoted Miron Livny saying: "I was doing Cloud computing way before people called it Grid computing", referring to the ground breaking Condor technology. It is the Grid scientific paradigm that counts!

MORGRIDGE INSTITUTE FOR RESEARCH

CENTER FOR HIGH THROUGHPUT COMPUTING

WISCONSIN INSTITUTE FOR DISCOVERY

In **1996** I introduced the distinction between High **Performance** Computing (**HPC**) and High **Throughput** Computing (**HTC**) in a seminar at the NASA Goddard Flight Center in and a month later at the European Laboratory for Particle Physics (**CERN**). In June of 1997 HPCWire published an interview on High Throughput Computing.

```
HIGH THROUGHPUT COMPUTING: AN INTERVIEW WITH MIRON LIVNY                06.27.97
by Alan Beck, editor in chief                                            HPCwire
=============================================================================

  This month, NCSA's (National Center for Supercomputing Applications)
Advanced Computing Group (ACG) will begin testing Condor, a software system
developed at the University of Wisconsin that promises to expand computing
capabilities through efficient capture of cycles on idle machines. The
software, operating within an HTC (High Throughput Computing) rather than a
traditional HPC (High Performance Computing) paradigm, organizes machines
```

Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science and Engineering in 2017-2020

DETAILS

117 pages | 6 x 9 | PAPERBACK
ISBN 978-0-309-38961-7 | DOI: 10.17226/21886

AUTHORS

Committee on Future Directions for NSF Advanced Computing Infrastructure to Support U.S. Science in 2017-2020; Computer Science and Telecommunications Board; Division on Engineering and Physical Sciences; National Academies of Sciences, Engineering, and Medicine

"… many fields today rely on **high-throughput computing** for **discovery**."

"Many fields **increasingly** rely **on high-throughput computing**"

"Recommendation 2.2. NSF should (a) … and (b) broaden the accessibility and utility of these large-scale platforms by allocating **high-throughput** as well as high-performance workflows to them."

High Throughput Computing requires automation as it is a 24-7-365 activity that involves large numbers of jobs

FLOPY ≠ *(60*60*24*7*52)*FLOPS*

100K Hours*1 Job ≠ 1 H*100K J

MORGRIDGE INSTITUTE FOR RESEARCH

CENTER FOR HIGH THROUGHPUT COMPUTING

WISCONSIN INSTITUTE FOR DISCOVERY

Serving the UW-Madison campus

- The Center for High Throughput Computing (**CHTC**) has been serving the campus for more than 10 years
- Delivered more than **320M** core hours in the past 12 month to researchers with HTC workload from more than 50 departments
- ~10% of these hours where provided by the Open Science Grid (**OSG**)

MORGRIDGE INSTITUTE FOR RESEARCH

CENTER FOR HIGH THROUGHPUT COMPUTING

WISCONSIN INSTITUTE FOR DISCOVERY

Operating at a National Scale

- The Open Science Grid (**OSG**) is providing a fabric of distributed HTC services across more than 50 sites in the US
- Supporting a broad variety of Virtual Organizations (VOs) and resource providers
- Delivered more than **1.3B** core hours in the past 12 month
- %65 of these hours were in support of the US-LHC community

MORGRIDGE INSTITUTE FOR RESEARCH

CENTER FOR HIGH THROUGHPUT COMPUTING

WISCONSIN INSTITUTE FOR DISCOVERY

Support International Organizations

- Deployment of HTC technologies at Large Hadron Collider (**LHC**) Tire-0 (**CERN**), Tire-1 and Tier-2 sites
- The *Compact* Muon Solenoid (**CMS**) LHC experiment global computing
- The Laser Interferometer Gravitational-Wave Observatory (**LIGO**) global computing

MORGRIDGE INSTITUTE FOR RESEARCH
HT CENTER FOR HIGH THROUGHPUT COMPUTING
WISCONSIN INSTITUTE FOR DISCOVERY

Courtesy SXS.

OSG helps LIGO confirm Einstein's theory

Einstein predicted gravitational waves over 100 years ago. Open Science Grid (OSG) resources are helping the NSF-funded Laser Interferometer Gravitational-Wave Observatory (LIGO) prove he was right.

> ***Thus far, LIGO has consumed almost four million hours on OSG — 628,602 hours were on Comet and 430,960 on Stampede resources. OSG's Brian Bockelman of the** University of Nebraska-Lincoln **and Edgar Fajardo from the SDSC used** HTCondor **software to help LIGO implement their Pegasus workflow on 16 clusters at universities and national labs across the US.***

MORGRIDGE INSTITUTE FOR RESEARCH

CENTER FOR HIGH THROUGHPUT COMPUTING

WISCONSIN INSTITUTE FOR DISCOVERY

“When a workflow might consist of 600,000 jobs, we don’t want to rerun them if we make a mistake. So we use DAGMan (Directed Acyclic Graph Manager, a meta-scheduler for HTCondor) and Pegasus workflow manager to optimize changes,” added Couvares. “The combination of Pegasus, Condor, and OSG work great together.” Keeping track of what has run and how the workflow progresses, Pegasus translates the abstract layer of what needs to be done into actual jobs for Condor, which then puts them out on OSG.

Using Directed Acyclic Graphs (DAGs) to support declarative automation of interdependent tasks

MORGRIDGE
INSTITUTE FOR RESEARCH

HT
CENTER FOR
HIGH THROUGHPUT
COMPUTING

WISCONSIN
INSTITUTE FOR DISCOVERY

Example of a LIGO Inspiral DAG

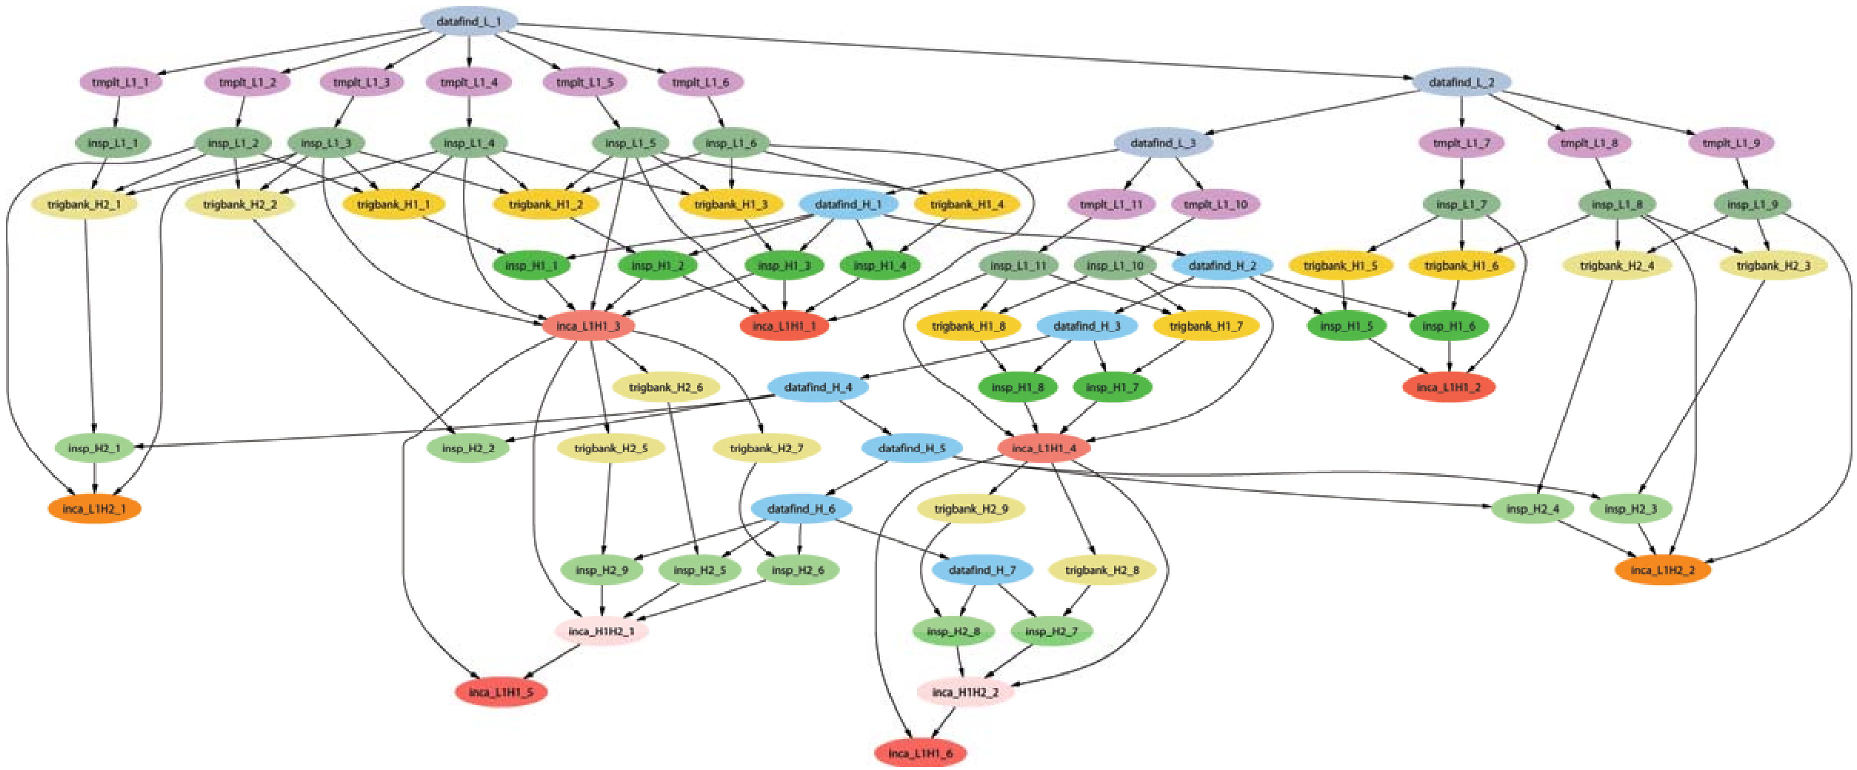

Traditional (low frequency) Capacity Planning

MORGRIDGE
INSTITUTE FOR RESEARCH

CENTER FOR
HIGH THROUGHPUT
COMPUTING

WISCONSIN
INSTITUTE FOR DISCOVERY

Turning $s into computing power

- Collect workload characteristics and customer (performance) metrics
- Understand the cost-performance profile of the hardware and software options
- **Acquire** (select, purchase, install) the resources and place them under the control of a **batch service**
- Live with your decision for (5-8) years

A job submitted to a batch service consists of an Acquisition Request (**AquR**) and a Job Description (**JobD**).

The Provision Manager (**Pman**) of the service provisions the resources and then runs the job on theses resources via the Job Launcher (**JaL**)

Most batch services manage a static collection of resources

HTCondor uses a two phase matchmaking process to dynamically **acquire** resources and then **provision** them to queued jobs. Jobs are **launched** via a task delegation protocol.

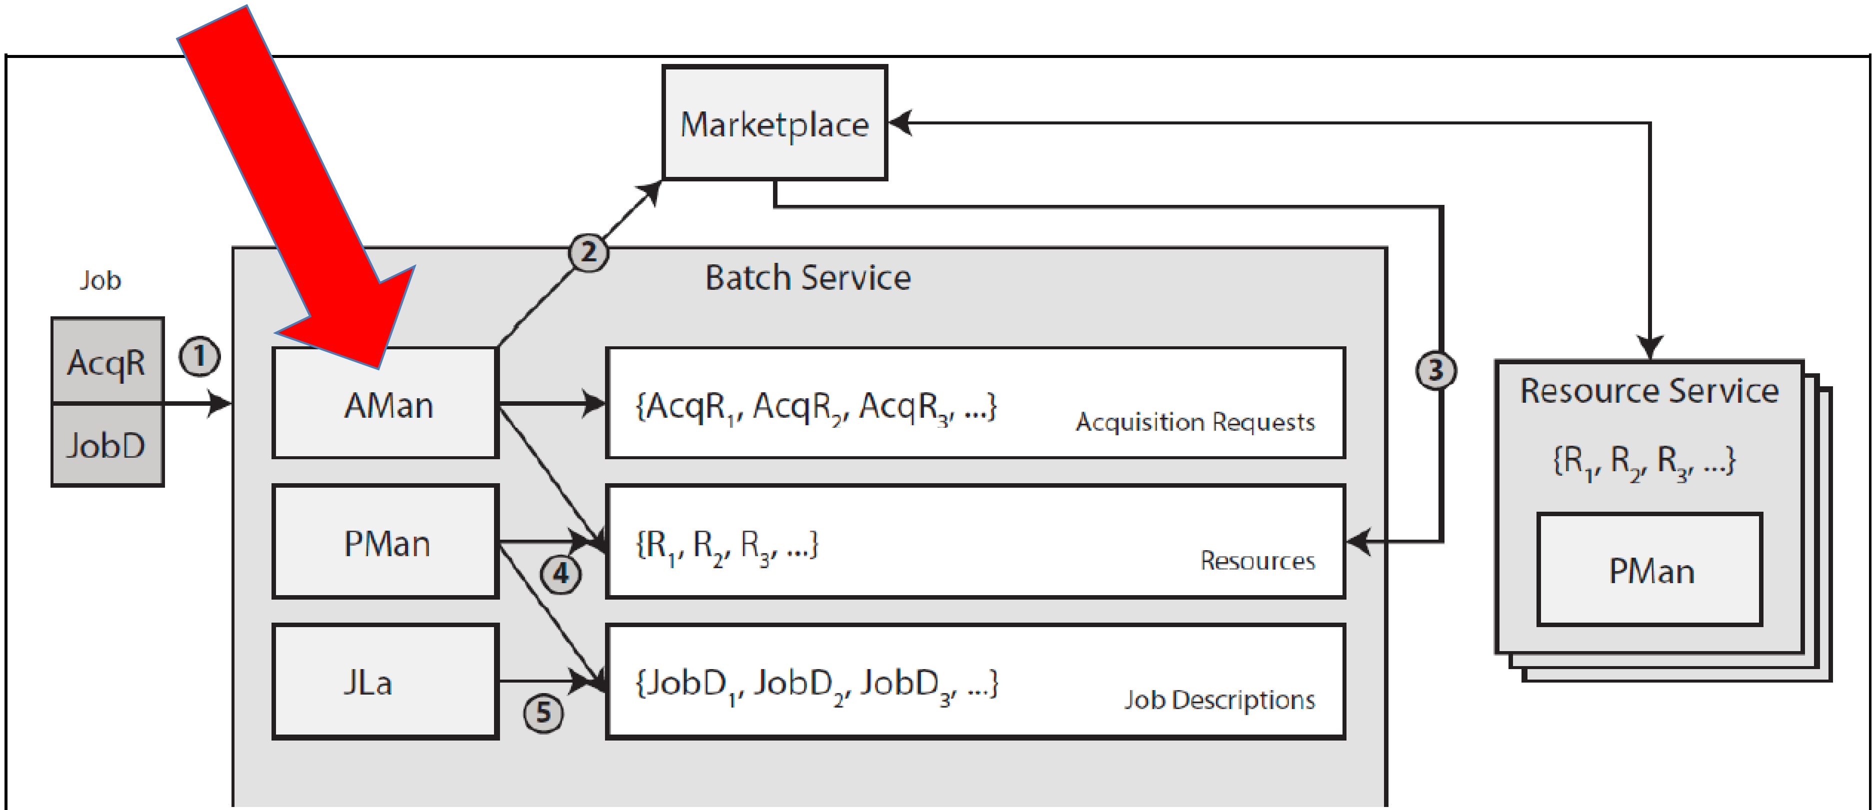

Figure 1: The HTC Framework: (1) A user submits job, composed of an AcqR and a JobD. (2) The AMan requests a resource compatible with an AcqR. (3) The MarketPlace offers a compatible Resource to the batch service. (4) The Pman selects a job description and Resource to send to (5) the JLa, which runs the job.

MORGRIDGE INSTITUTE FOR RESEARCH

CENTER FOR HIGH THROUGHPUT COMPUTING

WISCONSIN INSTITUTE FOR DISCOVERY

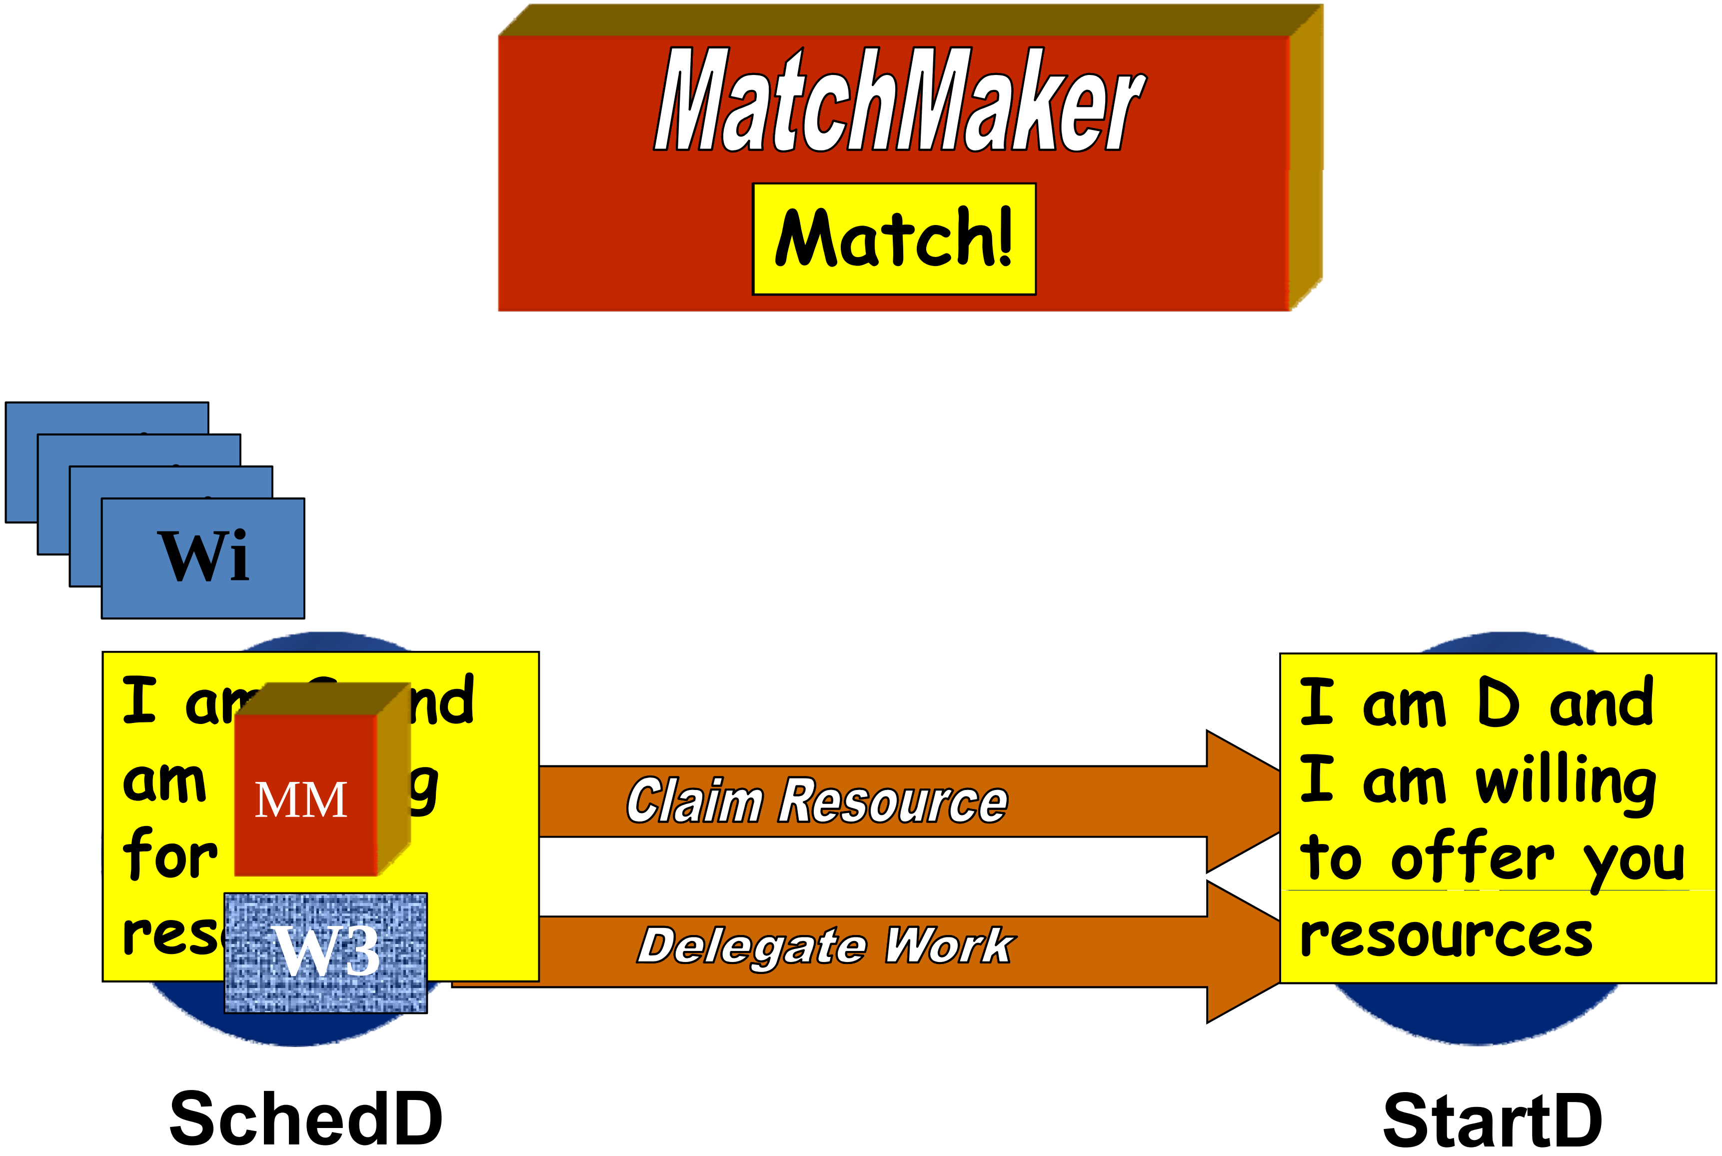
MatchMaker
Match!
Wi
I am S and am looking for resources
MM
W3
Claim Resource
Delegate Work
I am D and I am willing to offer you resources
SchedD
StartD

What Did We Learn From Serving a Quarter of a Million Batch Jobs on a Cluster of Privately Owned Workstations

1992

Miron Livny

Computer Sciences Department
University of Wisconsin — Madison
Madison, Wisconsin
{miron@cs.wisc.edu}

User Prospective

- Maximize the capacity of resources accessible via a single interface
- Minimize overhead of accessing remote capacity
- Preserve local computation environment


desktop

Researchers and VOs want to “Submit Locally and run Globally”. This principal has been underpinning HTCondor and OSG since inseption

MORGRIDGE INSTITUTE FOR RESEARCH
HT CENTER FOR HIGH THROUGHPUT COMPUTING
WISCONSIN INSTITUTE FOR DISCOVERY

“Run Globally” implies “Acquire Globally”

Here is what the OSG offers today with the support of HTCondor technologies

HTCondor 101

- Jobs are submitted to the HTCondor **SchedD**
- The **SchedD** can **Flock** to additional **Matchmakers**
- The **SchedD** can delegate a job for execution to a **HTCondor StartD**
- The **SchedD** can delegate a job for execution to a another **Batch system**.
- The **SchedD** can delegate a job for execution to a **Grid Compute Element (CE)**
- The **SchedD** can delegate a job for execution to a **Commercial Cloud**

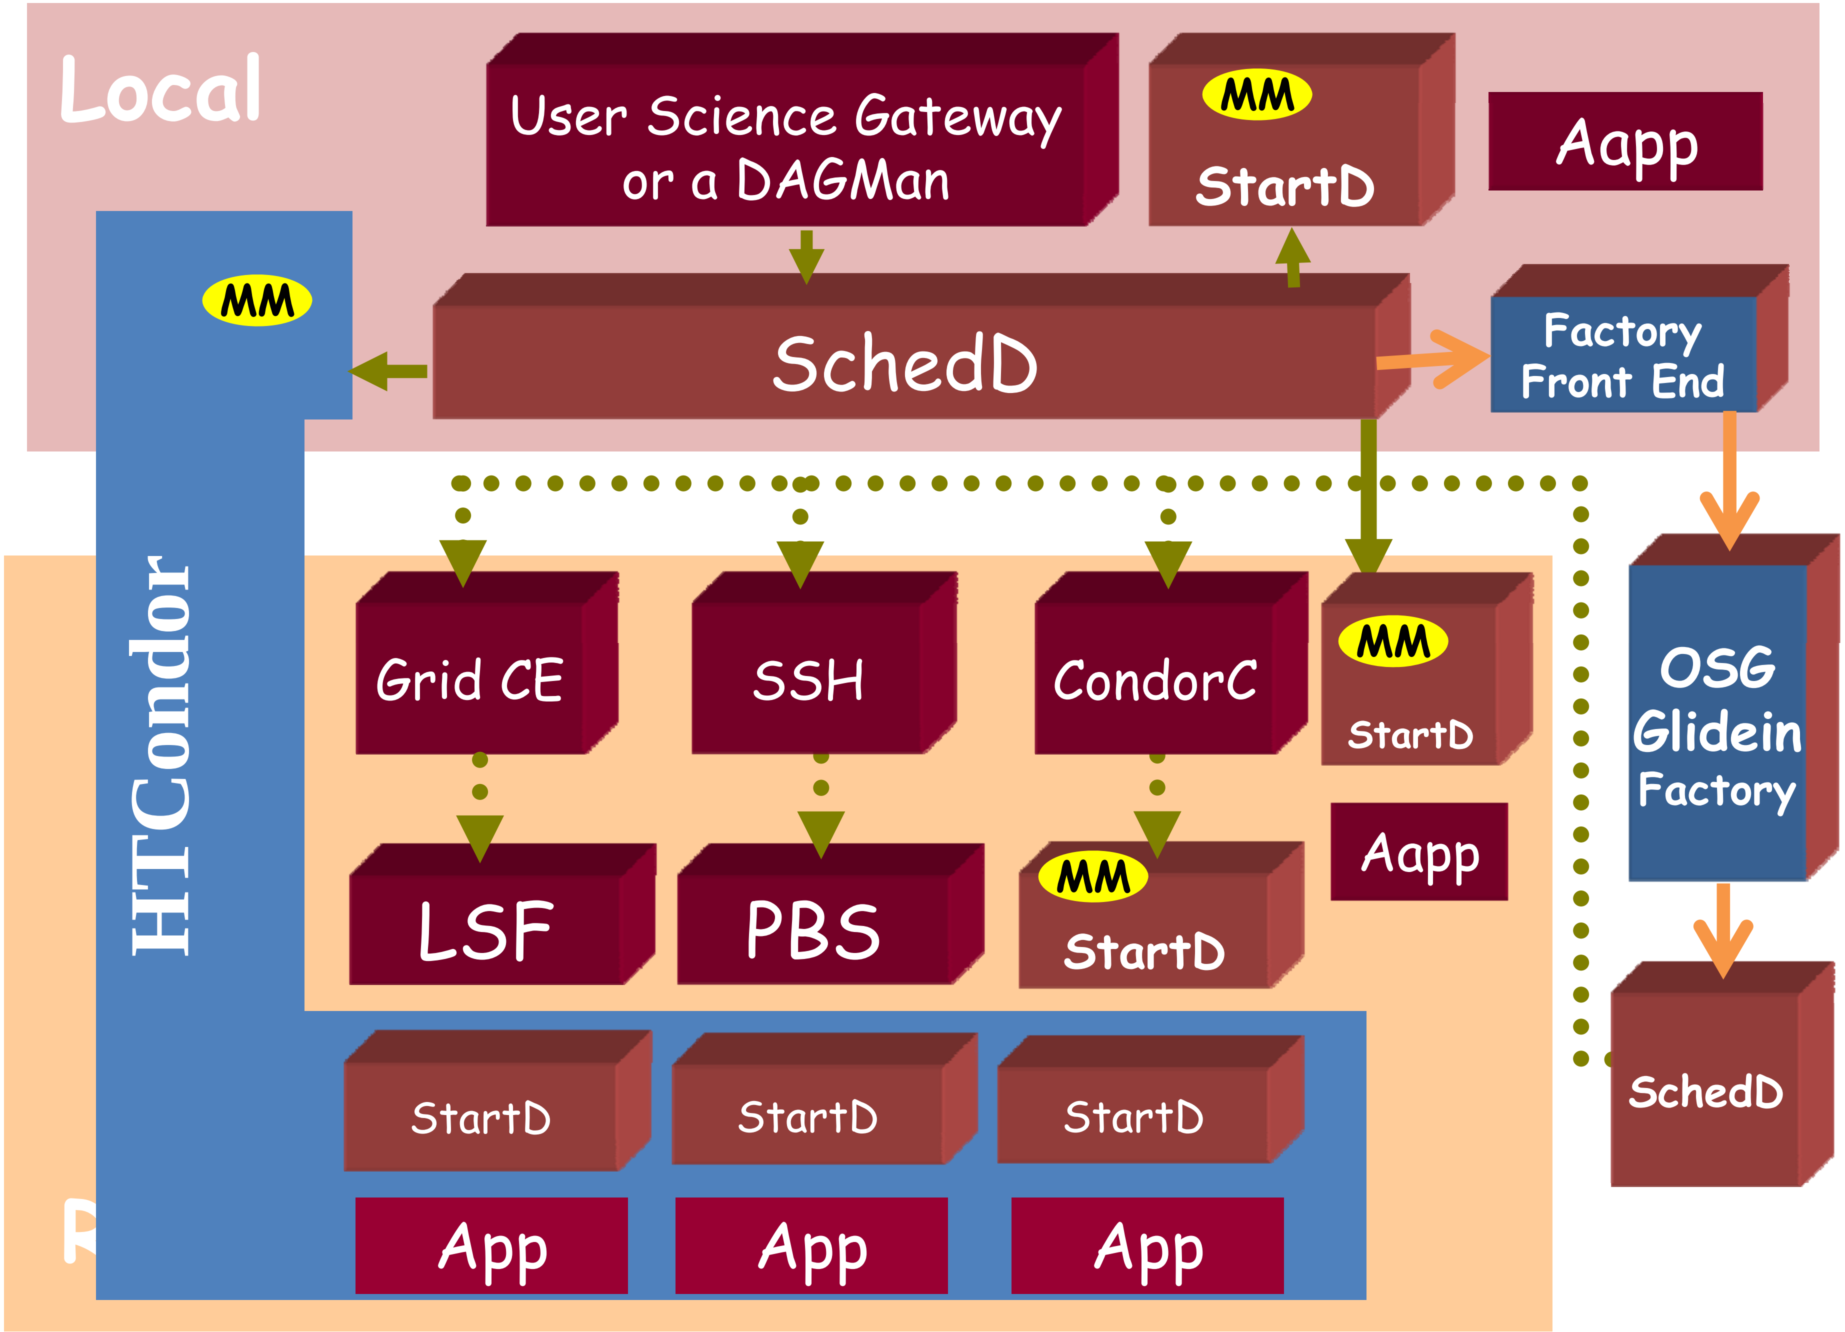
Local
User Science Gateway or a DAGMan
MM
StartD
Aapp
MM
SchedD
Factory Front End
HTCondor
Grid CE
SSH
CondorC
MM
StartD
OSG Glidein Factory
LSF
PBS
MM
StartD
Aapp
SchedD
StartD
StartD
StartD
App
App
App

Researcher or VOs may have …

- **Resources** they own and therefore fully control
- An **allocation** of resources on shared campus/national computing facility
- **“Fair Share” privileges** on shared campus/national computing facilities
- **Opportunistic Resources** provided by collaborators
- **Funding** to purchase resources from a commercial cloud provider

Commercial clouds offer to individuals with money …

- **Unbounded** on demand capacity for (almost) as long as needed
- A **variety** of cost/performance option for processing and storage resources
- **Dynamic cost** structures that track demand and supply
- **Diverse** (and competing) suppliers of computing resources and associated services

Here is what we can do today with the **Condor-Annex** Utility

What $1.5K can do for you

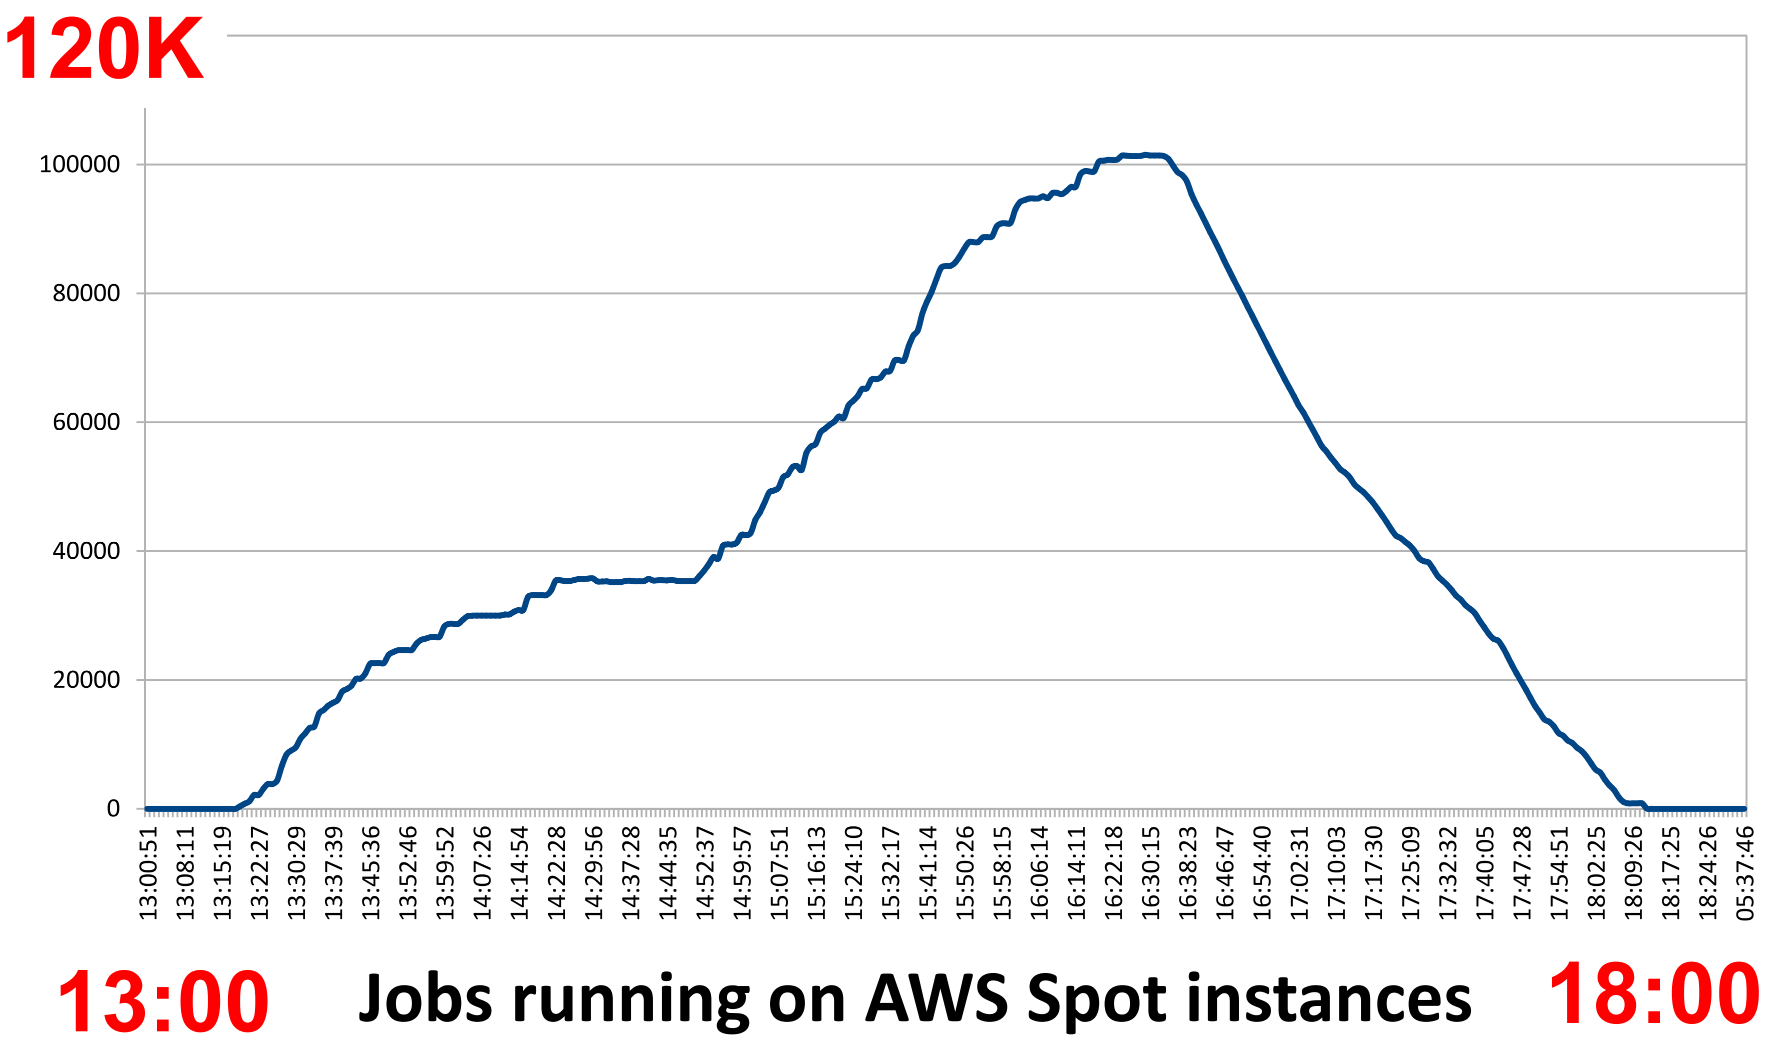

Annex means Addition

- An annex is "a building joined to main building, providing additional space or accommodations."
  - extra cores
  - GPUs or larger main memories
  - specialized policies
- To `condor_annex` is "to append or add as an extra or subordinate pool."

Annex Lifecycle

1. User requests resources (number, duration).
2. Then `condor_annex` starts instances.
3. Instances join pool.
4. Instances stop spending your money:
   – if they become idle, or
   – after the duration.

MORGRIDGE
INSTITUTE FOR RESEARCH
HT CENTER FOR HIGH THROUGHPUT COMPUTING
WISCONSIN
INSTITUTE FOR DISCOVERY

Here is what we can do today with the **GlideIn Workflow Management System (WMS)**

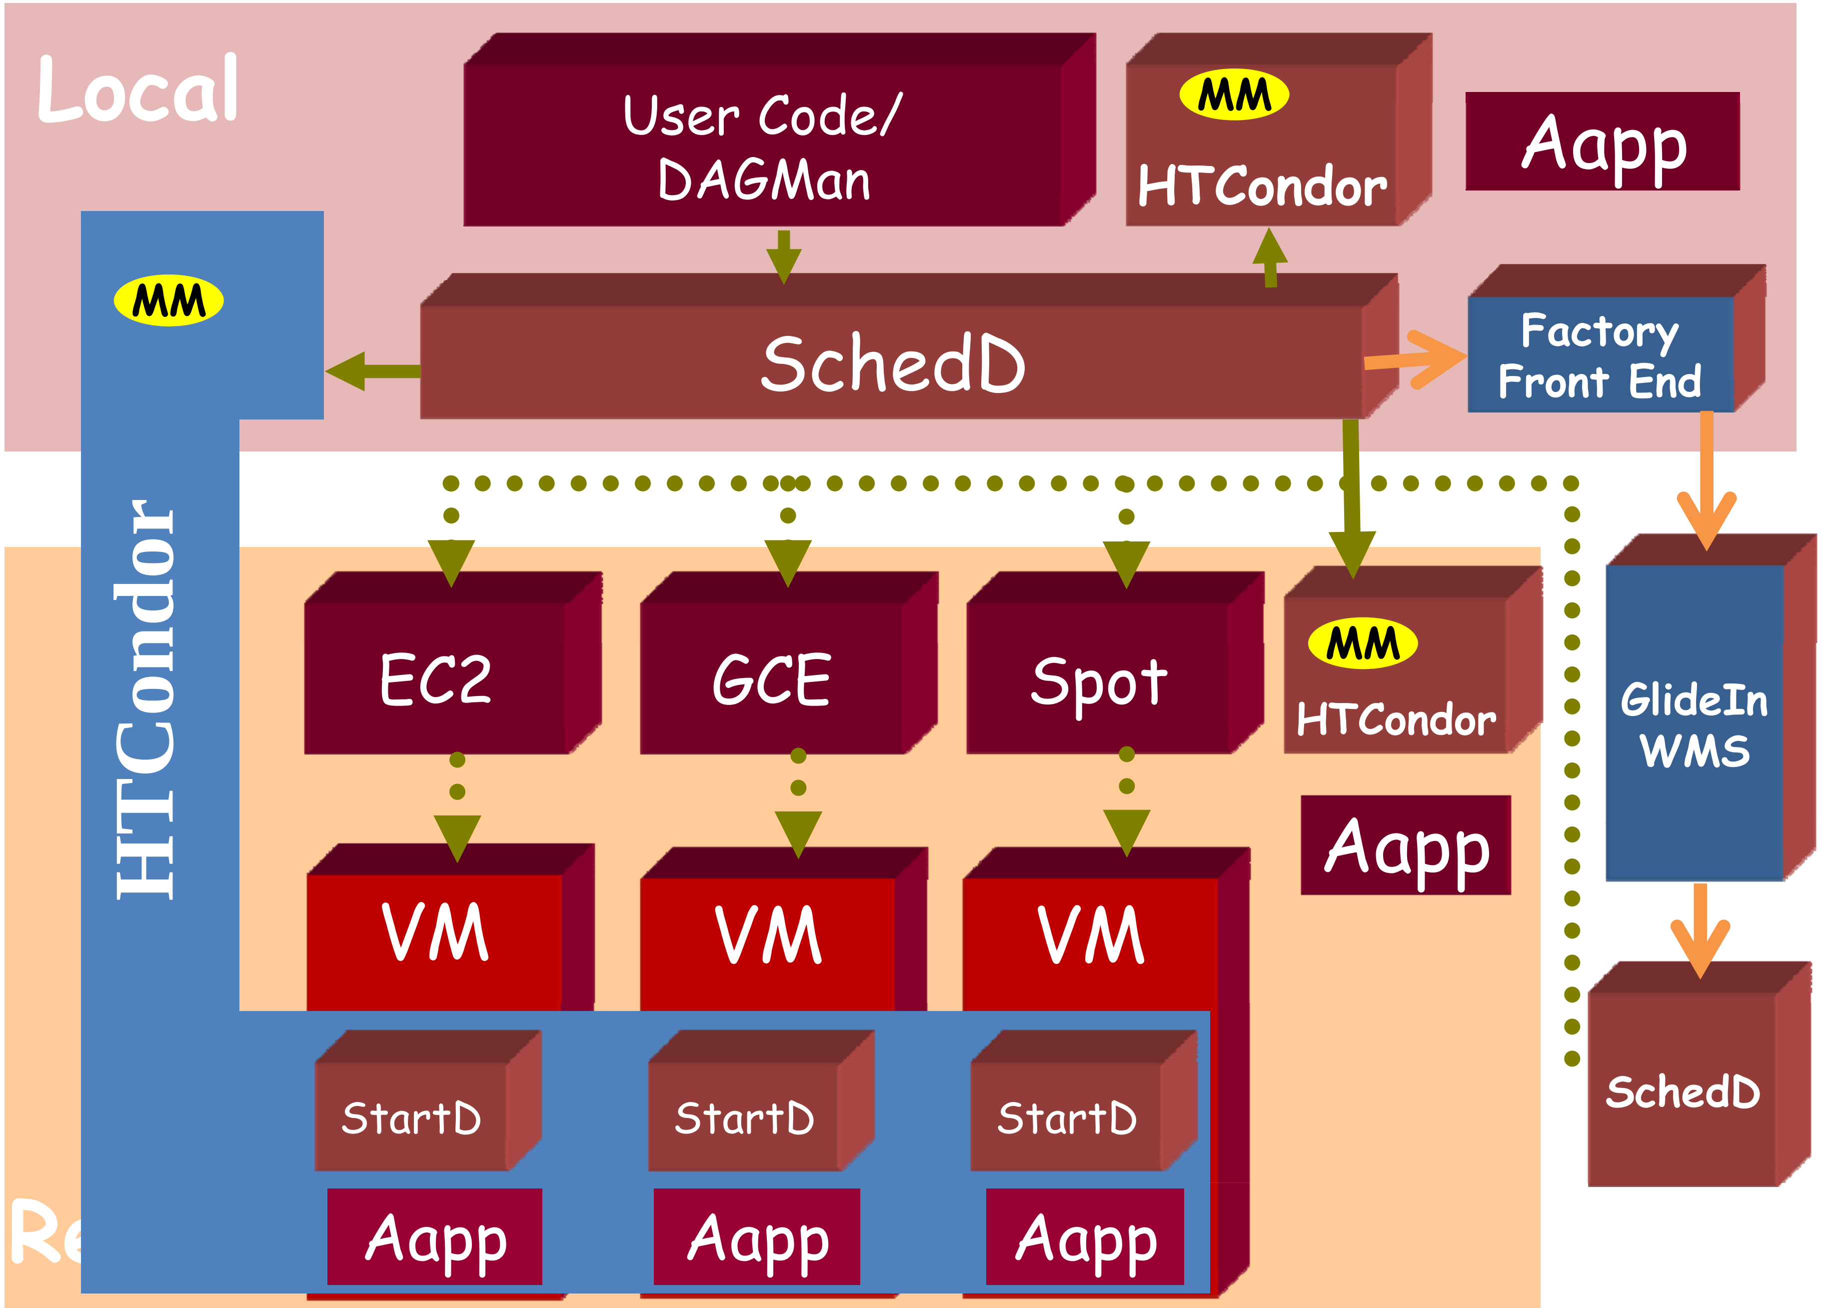

Local
User Code/
DAGMan
MM
HTCondor
Aapp
MM
SchedD
Factory
Front End
HTCondor
EC2
GCE
Spot
MM
HTCondor
GlideIn
WMS
Aapp
VM
VM
VM
StartD
StartD
StartD
SchedD
Aapp
Aapp
Aapp

CMS @ Google – preliminary numbers

- 6.35 M wallhours used; 5.42 M wallhours for completed jobs.
  - 730172 simulation jobs submitted; only 47 did not complete through the CMS and HEPCloud fault-tolerant infrastructures
  - Most wasted hours during ramp-up as we found and eliminated issues; goodput was at 94% during the last 3 days.
- Used ~$100k worth of credits on Google Cloud during Supercomputing 2016
  - $71k virtual machine costs
  - $8.6k network egress
  - $8.5k disk attached to VMs
  - $3.5k cloud storage for input data
- 205 M physics events generated, yielding 81.8 TB of data

Doubling CMS compute capacity

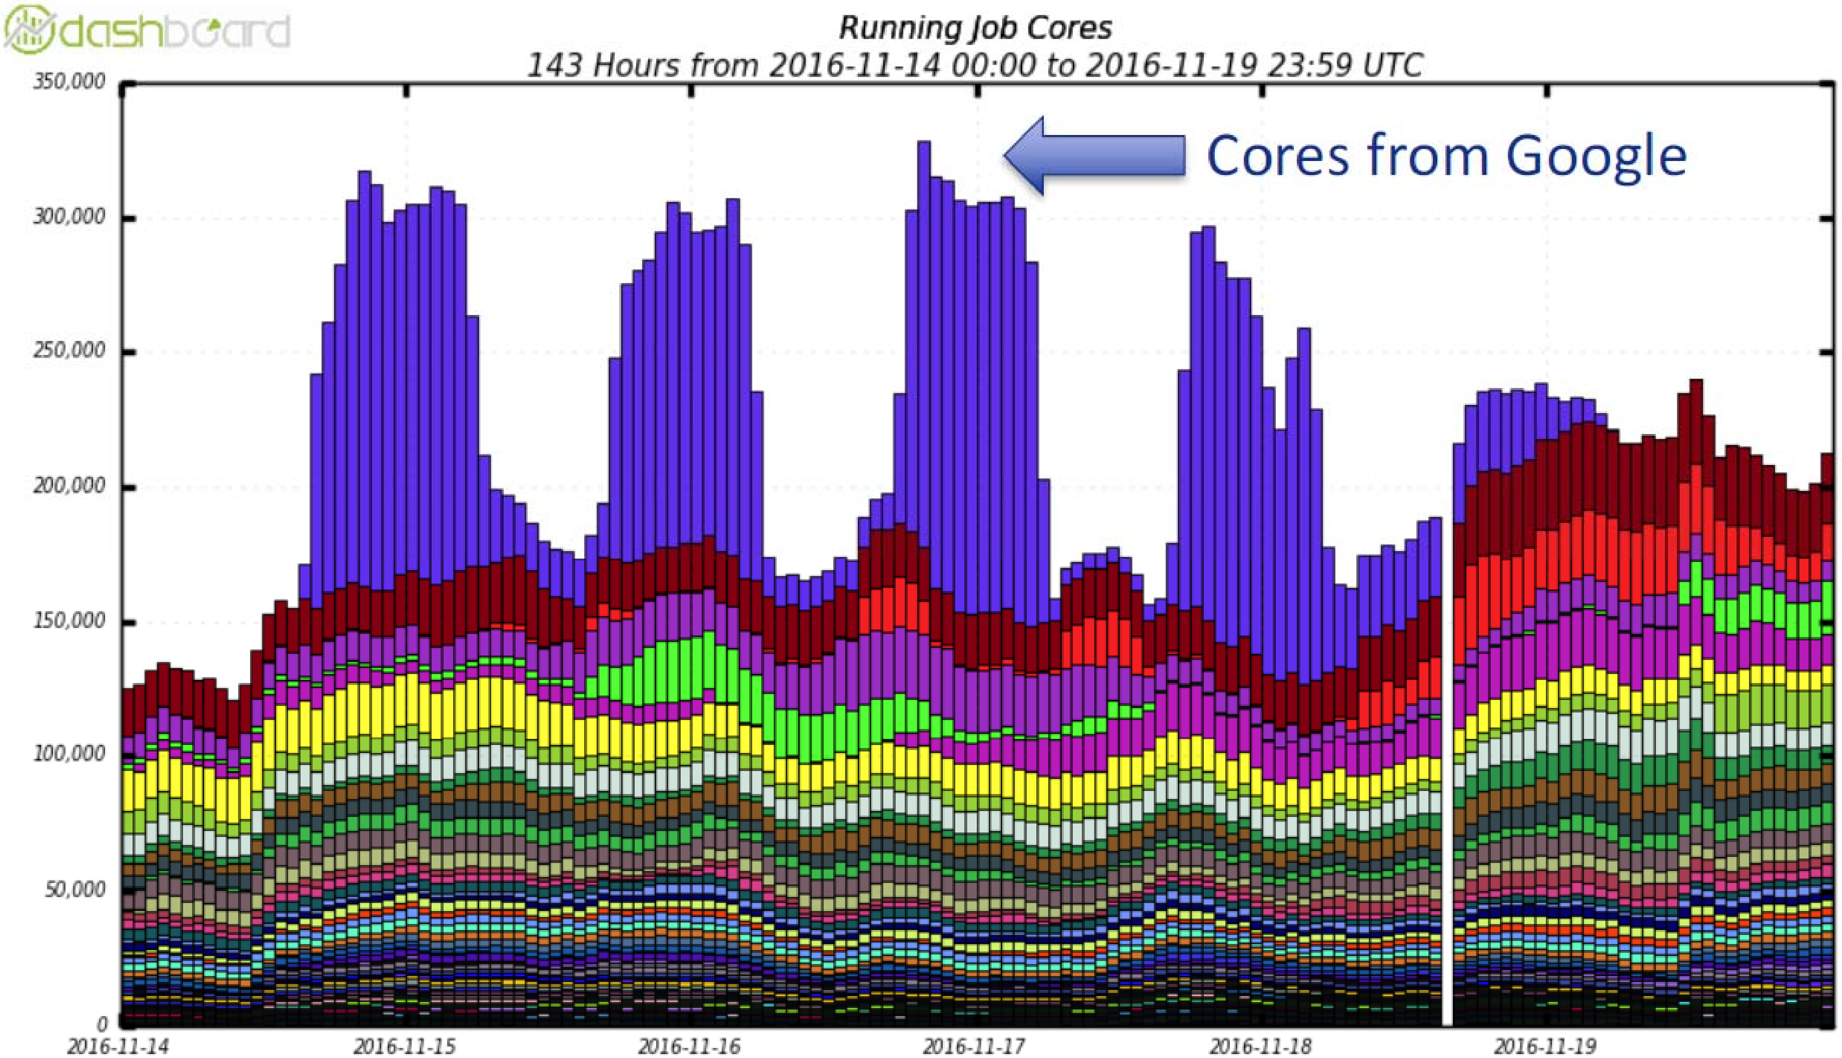

Cost < $100K

HEPCloud is an R&D project led by the Fermi computing division

MORGRIDGE INSTITUTE FOR RESEARCH

CENTER FOR HIGH THROUGHPUT COMPUTING

WISCONSIN INSTITUTE FOR DISCOVERY

HEPCloud – glideinWMS and HTCondor

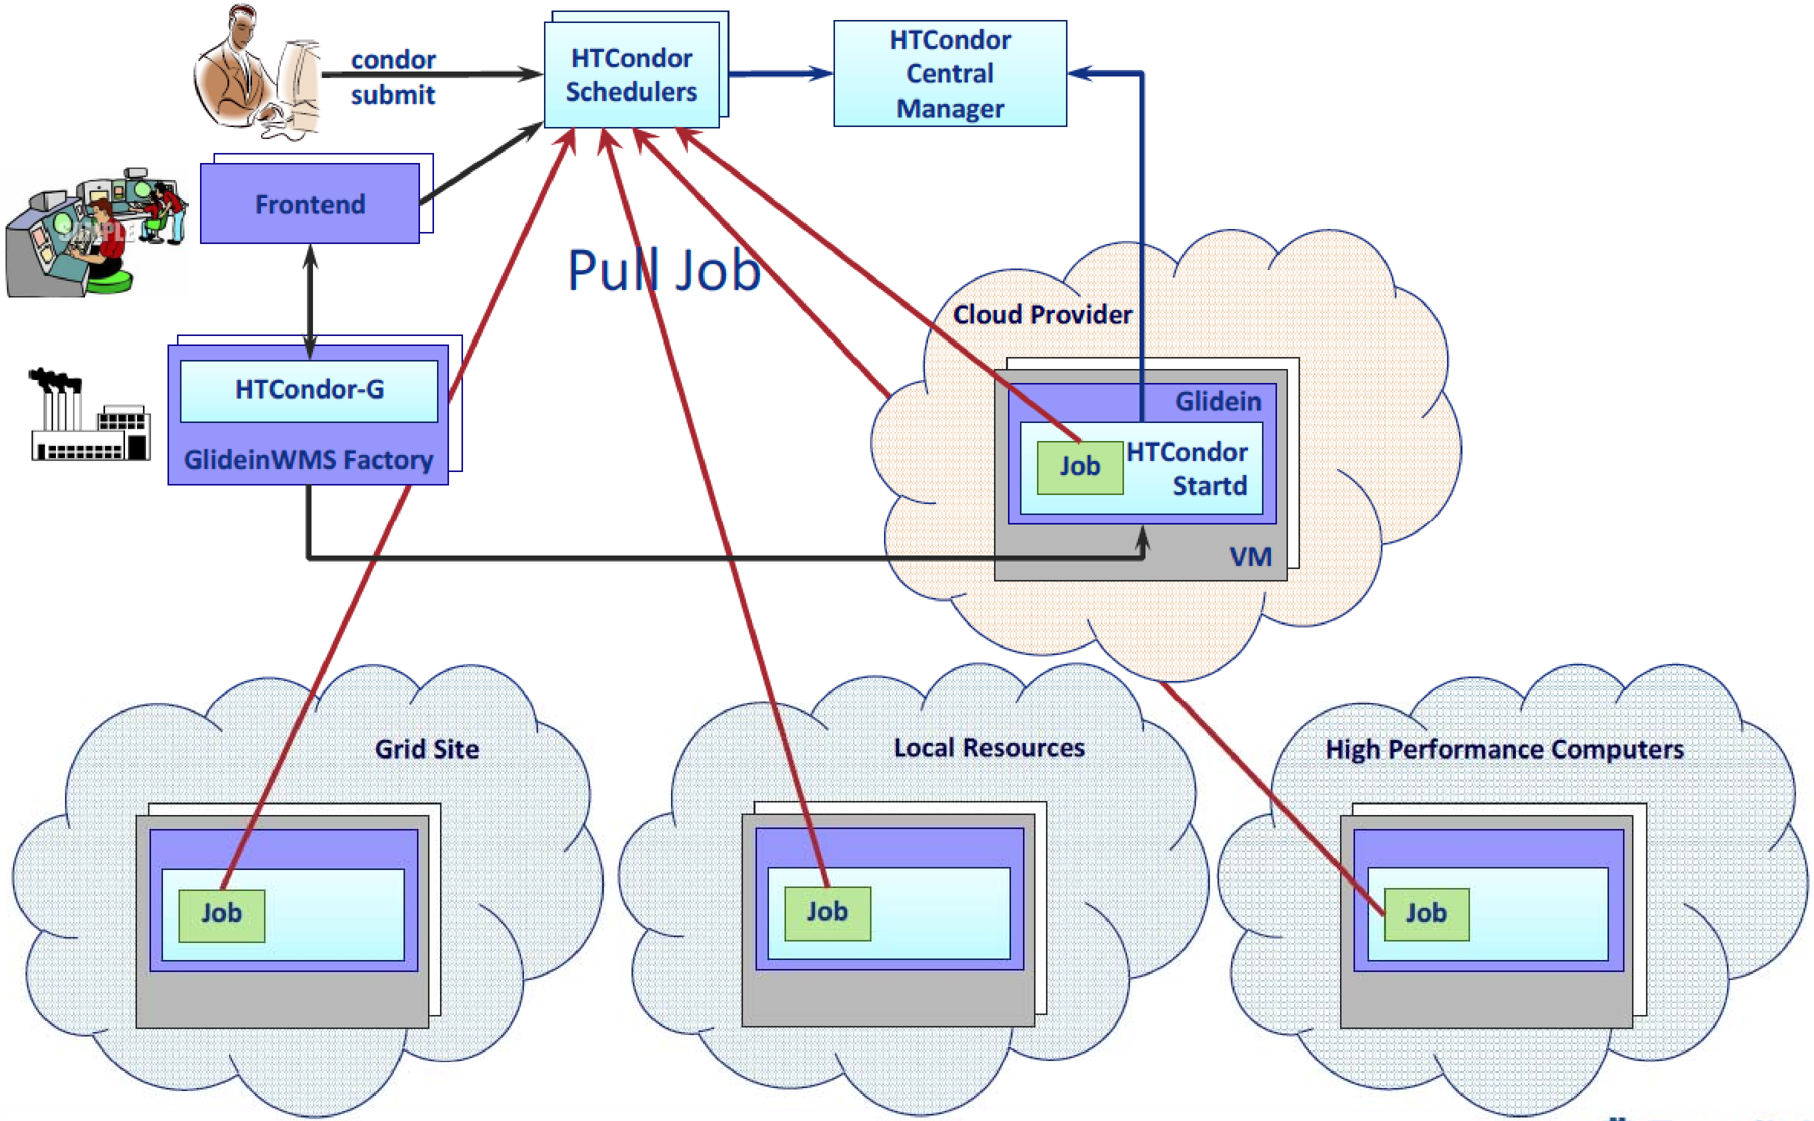

HEPCloud Architecture

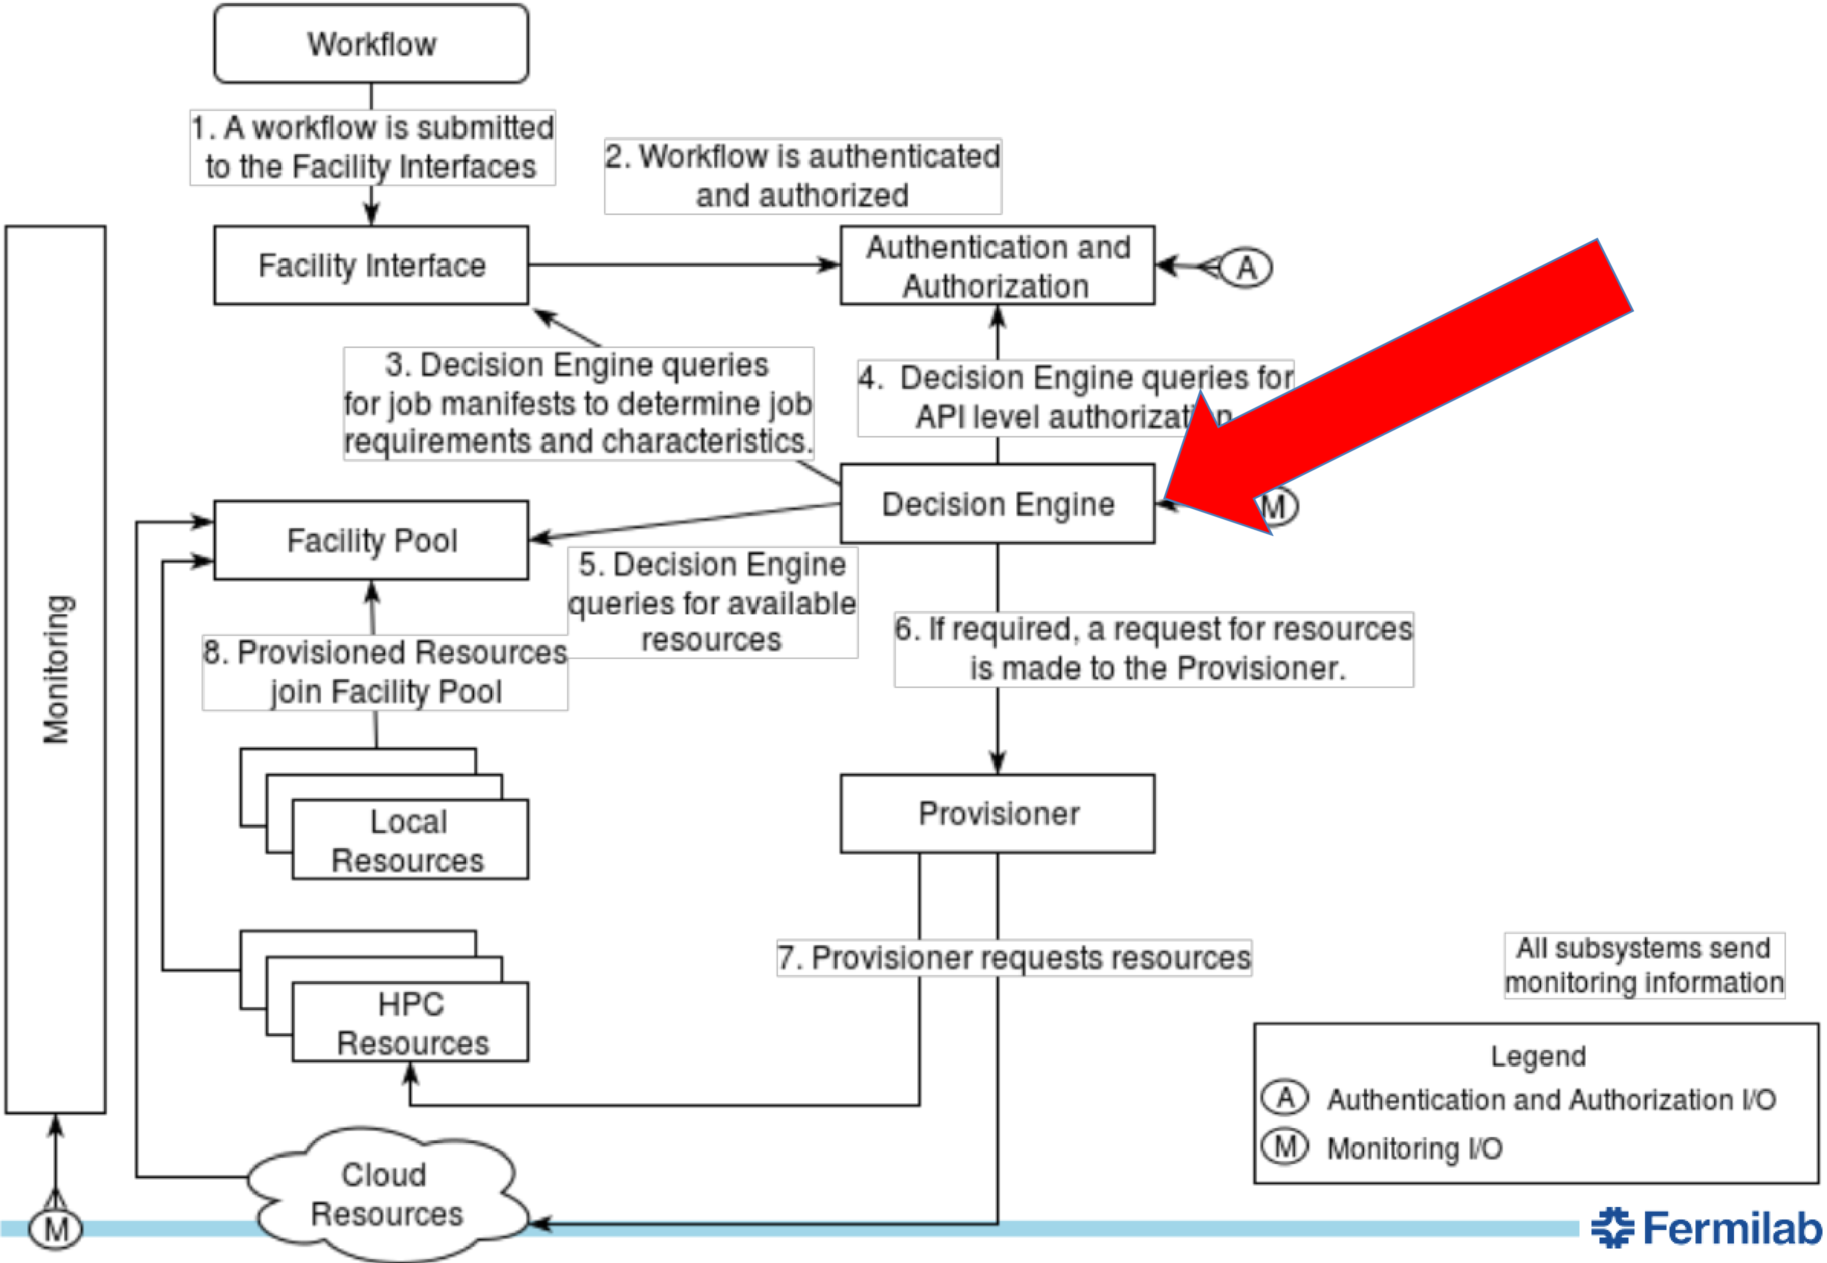

Decision Engine will have to implement **on-the-fly capacity planning** to control acquisition and release of resources

Many Challenges Ahead

- Language to define policies
- Software that manages (real) money
- Validation and verification of policies and software tools
- Risk management
- Accounting and Auditing
- Integration with DAGMan (workflows)

Global Scientific Computing via a Flock of Condors

CERN 92

Miron Livny

Computer Sciences Department
University of Wisconsin — Madison
Madison, Wisconsin
{miron@cs.wisc.edu}

MISSION

Give scientists effective and efficient access to large amounts of cheap (if possible free) CPU cycles and main memory storage

THE CHALLENGE

How to turn existing privetly owned clusters of *workstations*, *farms*, *multiprocessors*, and *supercomputers* into an efficient and effective Global Computing Environment?

In other words, how to minimize wait while idle?

APPROACH

Use wide-area networks to transfer batch jobs between Condor systems

- Boundaries of each Condor system will be determined by physical or administrative considerations

TWO EFFORTS

☐ ***UW CAMPUS***

Condor systems at Engineering, Statistics, and Computer Sciences

☐ ***INTERNATIONAL***

We have started a collaboration between CERN-SMC-NIKHEF-Univ. of Amsterdam, and University of Wisconsin-Madison