

Harnessing the Power of the Globus Toolkit

Miron Livny
Computer Sciences Department
University of Wisconsin-Madison
miron@cs.wisc.edu
<http://www.cs.wisc.edu/~miron>



Accomplishments

Opportunities

Challenges



Grid of
Technologies
not just
Resources



Some background ...



The Condor Project (Established '85)

Distributed Computing **research** performed by a team of 30 faculty, full time staff and students who

- face **software engineering** challenges in a UNIX/Linux/NT environment,
- are involved in national and international **collaborations**,
- actively interact with **users**,
- maintain and support a distributed **production** environment,
- and educate and train **students**.

Funding – DoD, DoE, NASA, NIH, NSF, INTEL,
Microsoft and the UW Graduate School

www.cs.wisc.edu/condor

The logo for the Condor project, featuring a large, stylized 'C' followed by the word 'ondor' in a serif font. The 'C' is dark grey with a gold outline, and the 'ondor' is in a gold serif font. A horizontal bar with a gold and grey gradient runs behind the text.

Claims for “benefits” provided by Distributed Processing Systems

P.H. Enslow, *“What is a Distributed Data Processing System?”* Computer, January 1978

- High Availability and Reliability
- High System Performance
- Ease of Modular and Incremental Growth
- Automatic Load and Resource Sharing
- Good Response to Temporary Overloads
- Easy Expansion in Capacity and/or Function



“ ... Since the early days of mankind the primary motivation for the establishment of *communities* has been the idea that by being part of an organized group the capabilities of an individual are improved. The great progress in the area of inter-computer communication led to the development of means by which stand-alone processing sub-systems can be integrated into multi-computer *'communities'*. ... ”

Miron Livny, “ *Study of Load Balancing Algorithms for Decentralized Distributed Processing Systems.*”,
Ph.D thesis, July 1983.

High Throughput Computing

For many experimental scientists, scientific progress and quality of research are strongly linked to computing **throughput**. In other words, they are less concerned about **instantaneous** computing power. Instead, what matters to them is the amount of computing they can harness over a month or a year --- they measure computing power in units of scenarios per **day**, wind patterns per **week**, instructions sets per **month**, or crystal configurations per **year**.



High Throughput Computing
is a
24-7-365
activity

FLOPY \neq (60*60*24*7*52)*FLOPS

Bring Globus technology and services to **end** users as components of an **end-to-end** service.



User/Application


Condor

Globus Toolkit

Condor

Fabric (processing, storage, communication)

Condor Layers



**Condor-G:
Making the
“Customer Agent” of
Condor
Globus aware**



Condor-G

Combine the **inter-domain** resource management protocols of the Globus Toolkit and the **intra-domain** resource management methods of Condor *to allow the user to harness large collections of resources across multiple domains as if they all belong to one personal domain.*

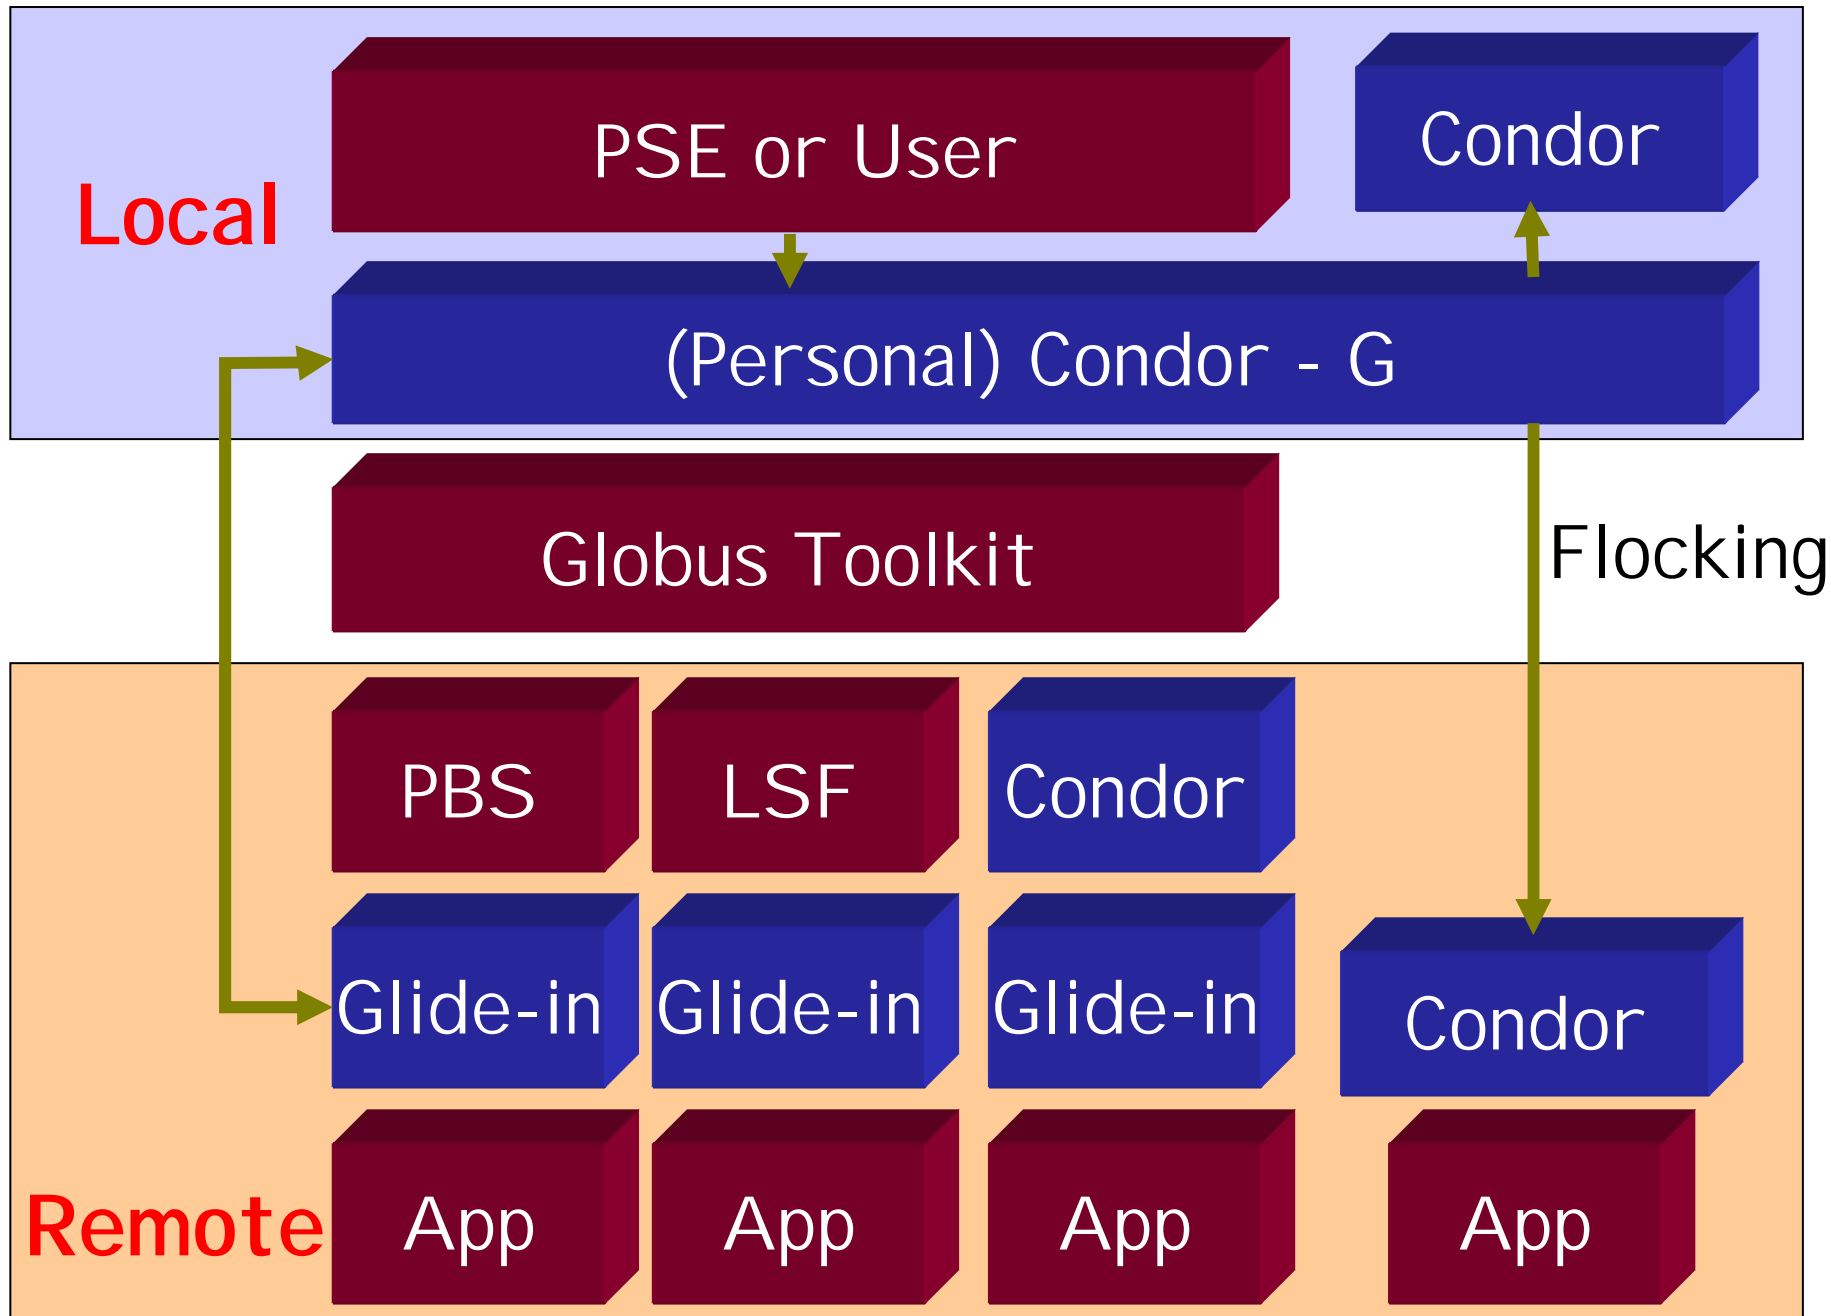
Globus Toolkit Services

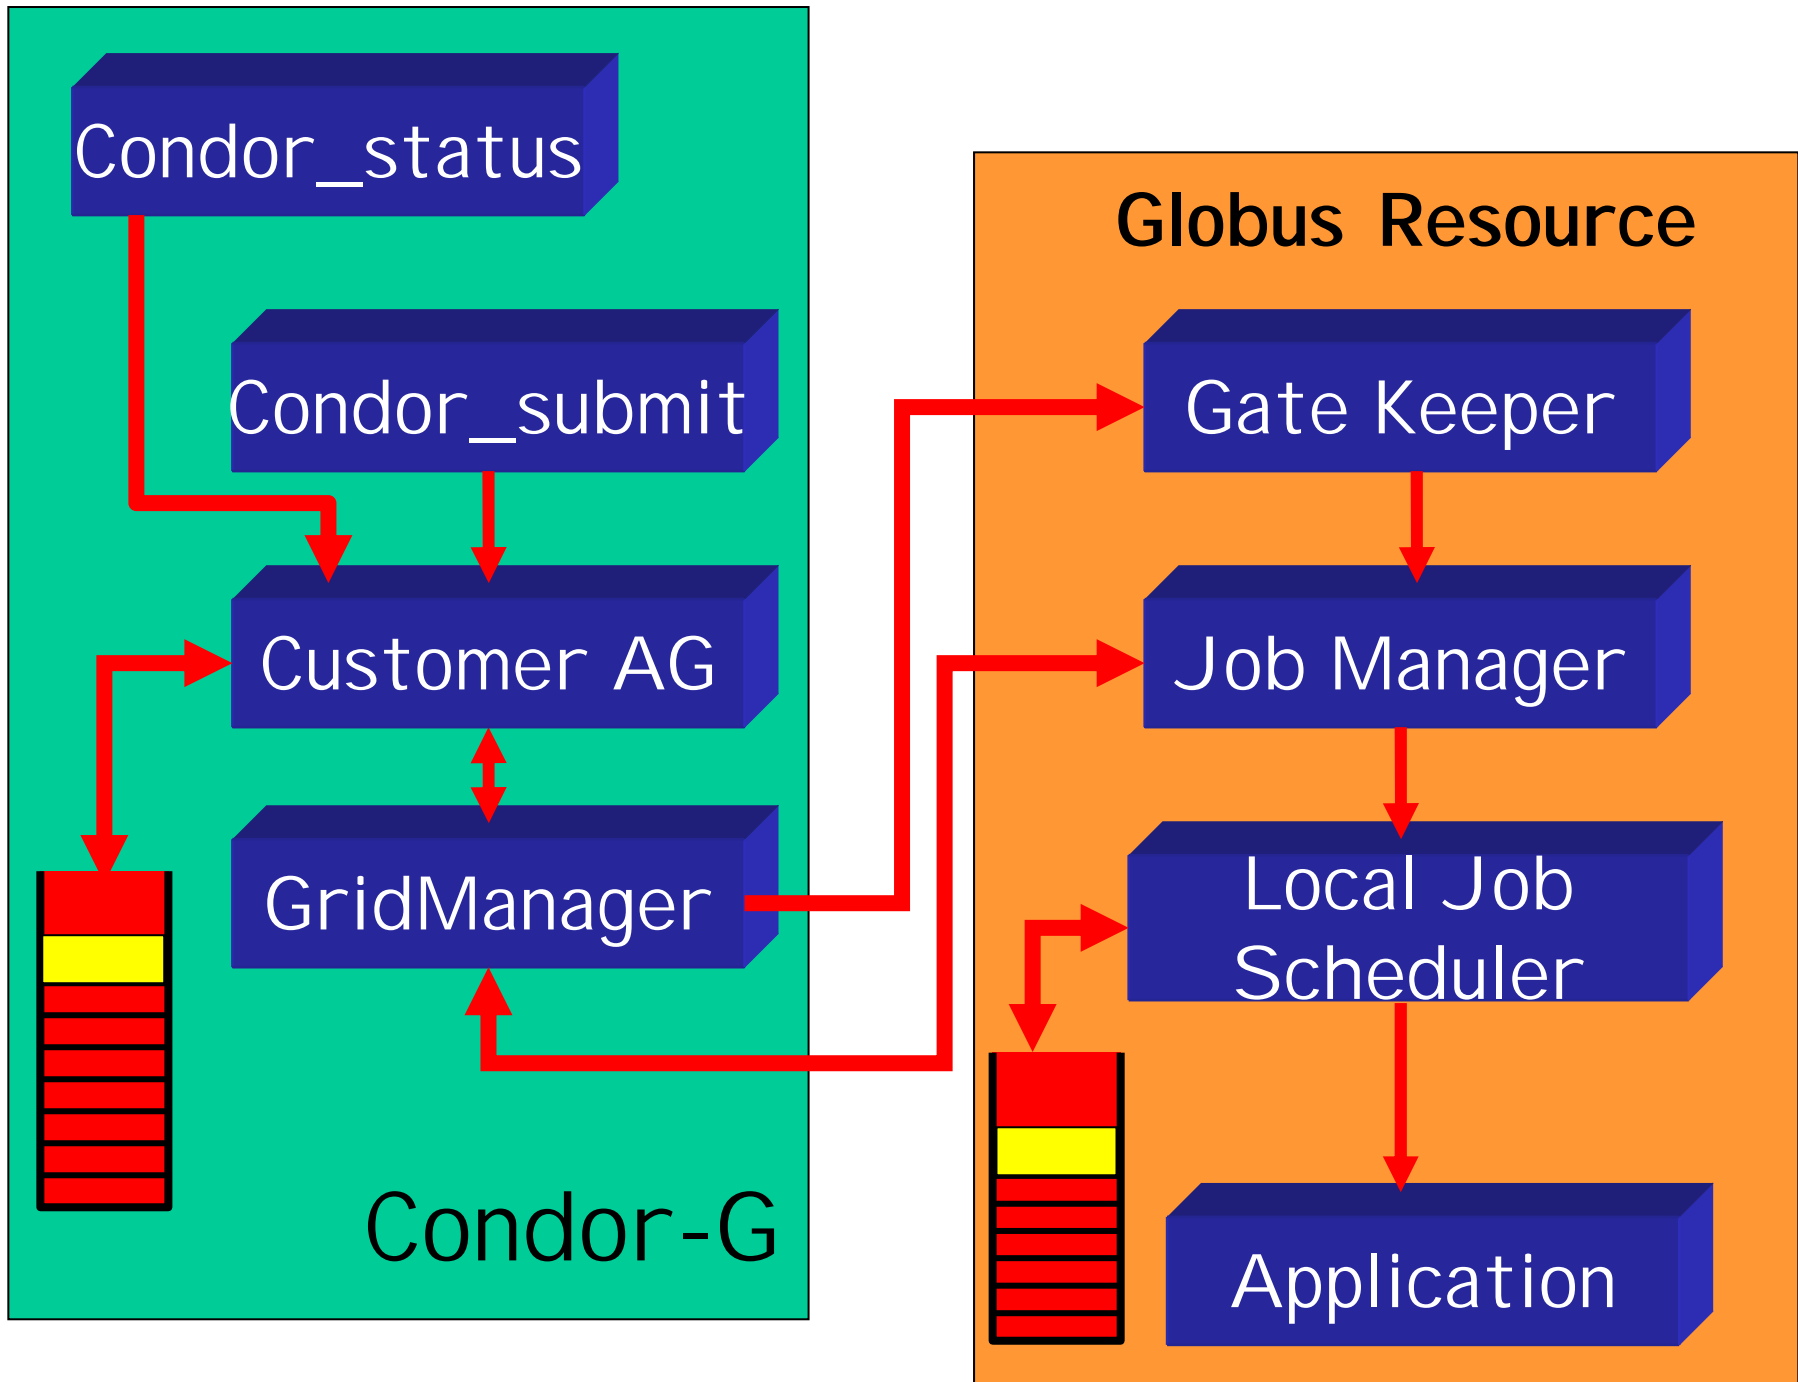
- › GSI – Grid Security Infrastructure
- › GRAM – Grid Resource Allocation and Management protocol
- › GASS – Global Access to Secondary Storage

Job Manager⁺⁺ (GRAM 1.5)


- › Two-phase commit job submission protocol
- › Attach to existing job
- › Credential refresh (GRAM 1.6)

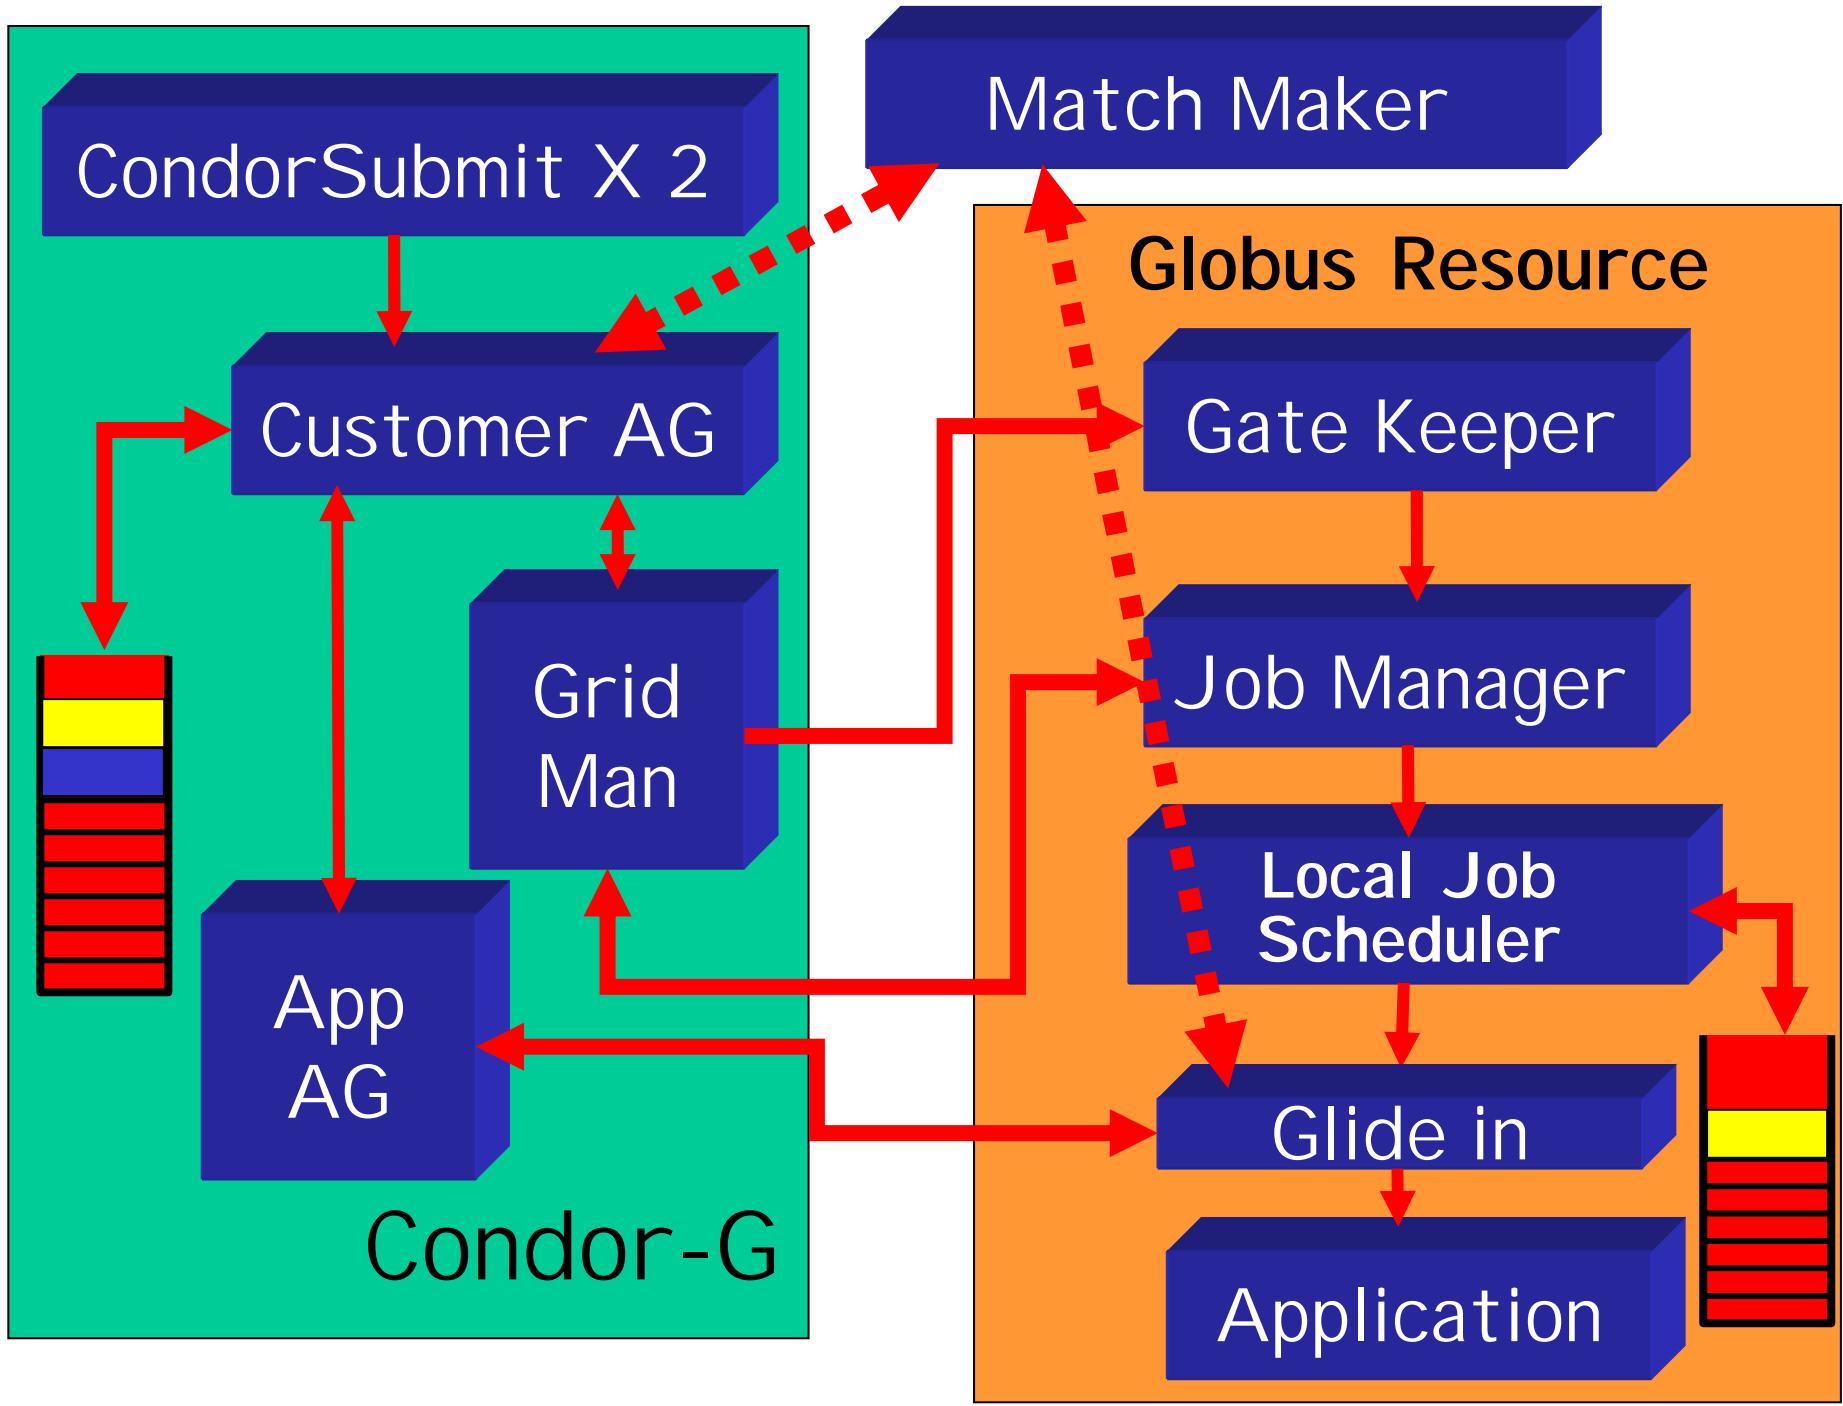
New version transferred to the Globus Team and is part of 2.0.





**Glide-in: Expending your
Condor pool
“on the fly”
and executing your jobs
on the remote resources in
a “friendly”
environment.**



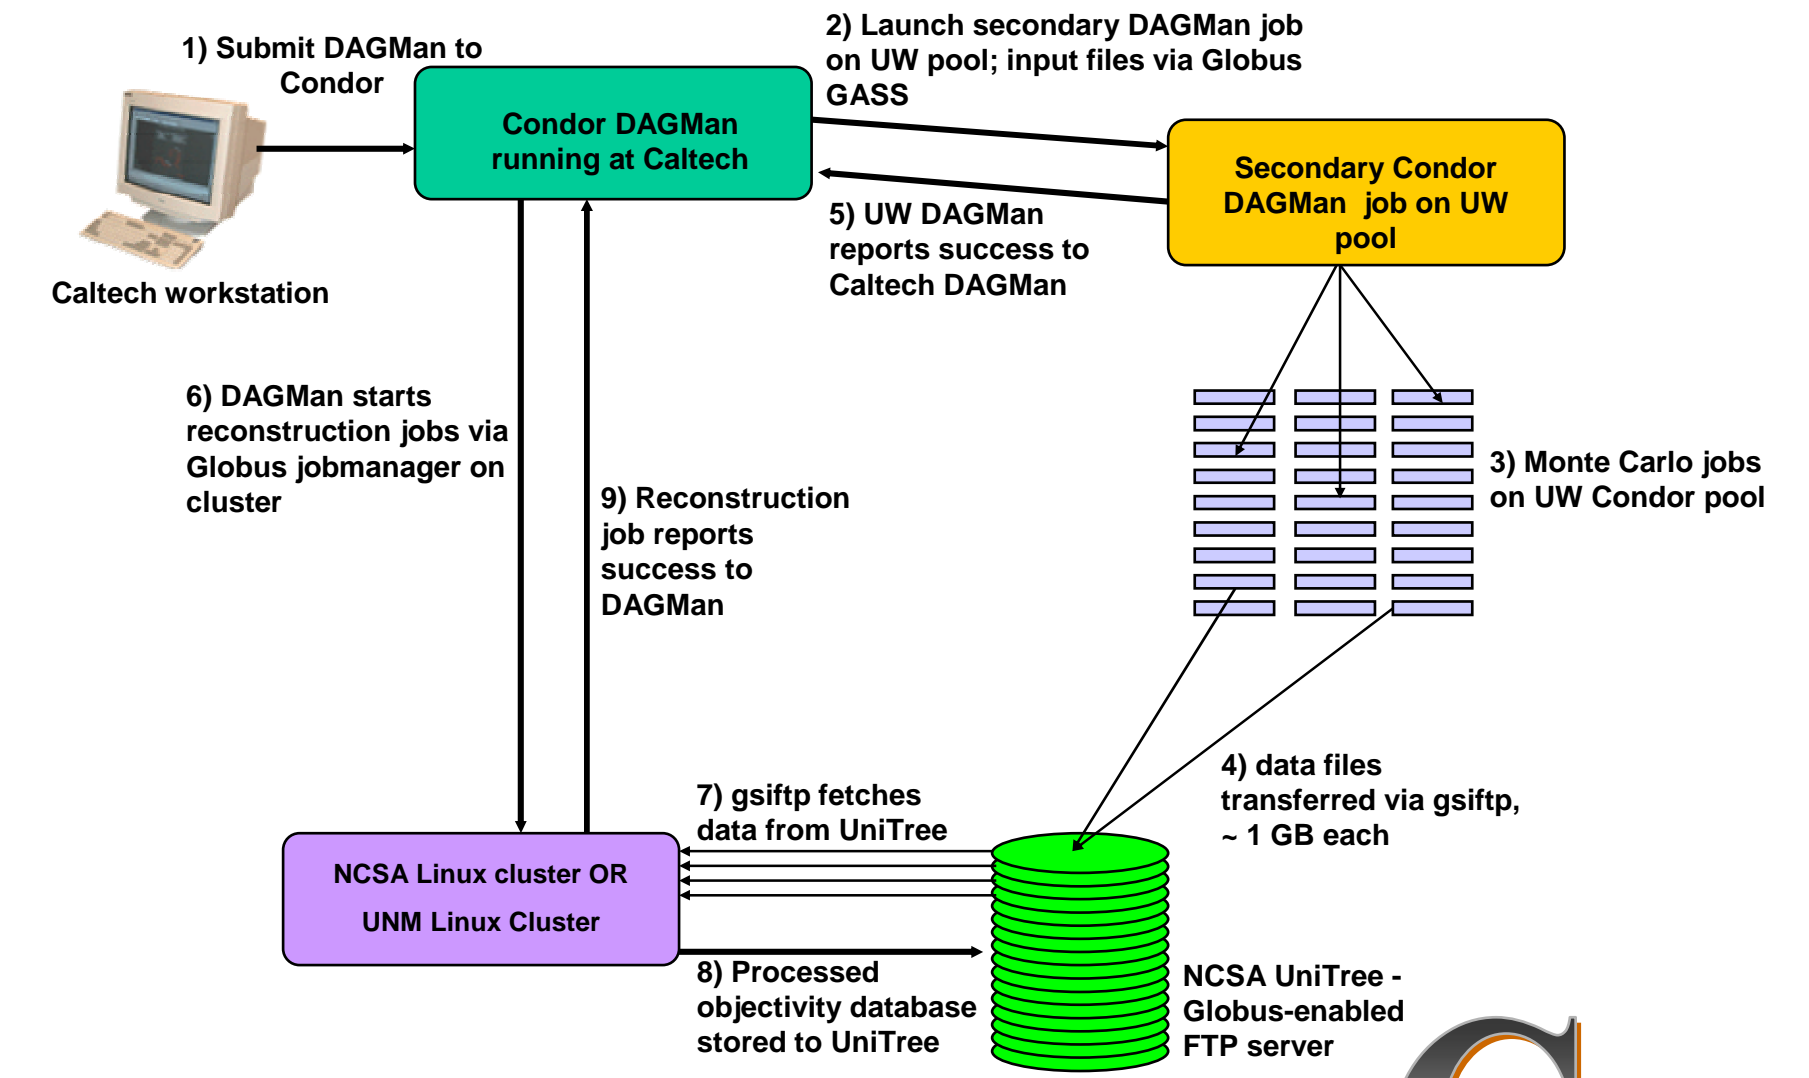


The CMS Production Story

A collaboration* between:

- Physicists & Computer Scientists
 - Vladimir Litvin ([Caltech CMS](#))
 - Scott Koranda, Bruce Loftis, John Towns ([NCSA](#))
 - Miron Livny, Peter Couvares, Todd Tannenbaum, Jamie Frey ([UW-Madison Condor](#))
- Software
 - [Condor](#), [Globus](#), [CMS](#)

* members of the GriPhyN project




**X509 Authentication
for all connections
established by
Condor**




Condor Security Infrastructure

- Provide a flexible, powerful and maintainable security infrastructure
- Enable:
 - Authentication – identifying user, resources and software modules
 - Authorization – providing access control
 - Encryption – providing secure channel




Condor Security Infrastructure

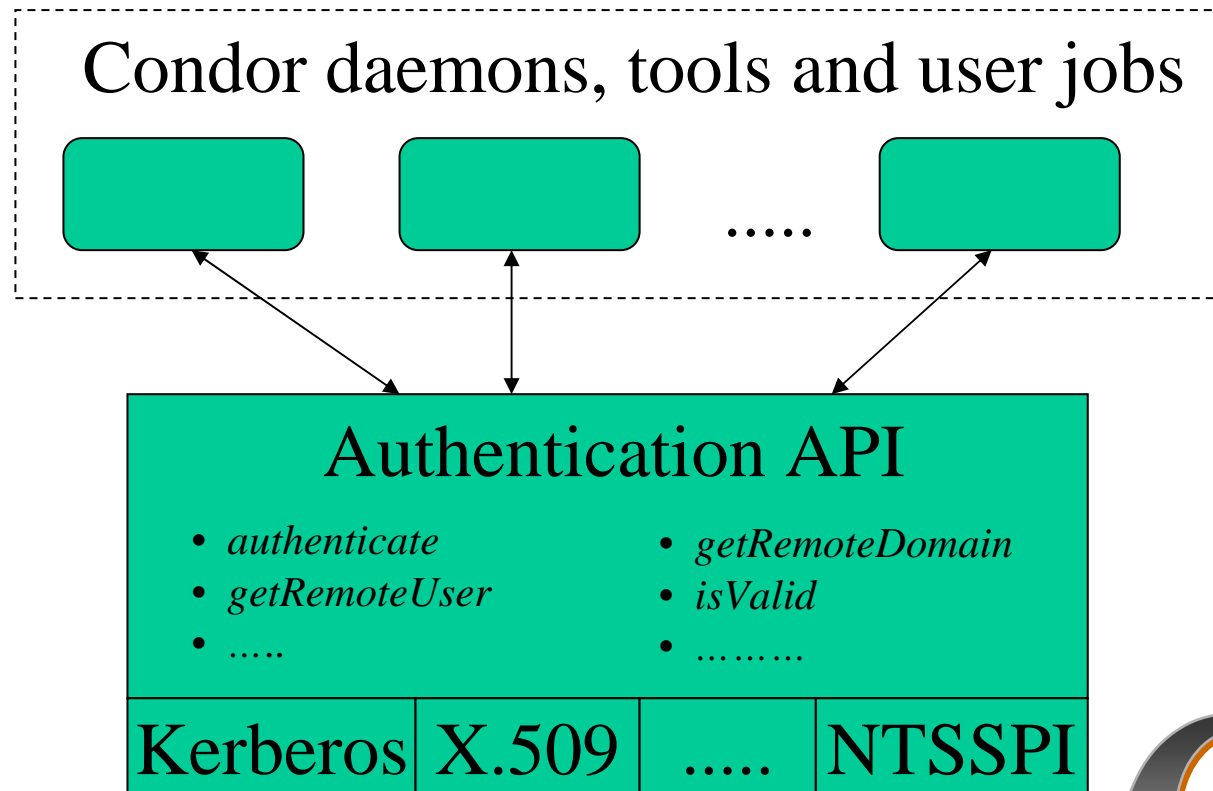


Authentication in Condor

- Flexible
 - Condor's authentication infrastructure supports multiple authentication protocols
 - X.509, Kerberos, NTSSPI , etc.
- Simple and extensible
 - A common set of API to provide consistency and hide complexity of the protocols
- Negotiable
 - Handshake process allows two parties to negotiate the authentication protocols to use



Authentication API



Credential Management

> Problems

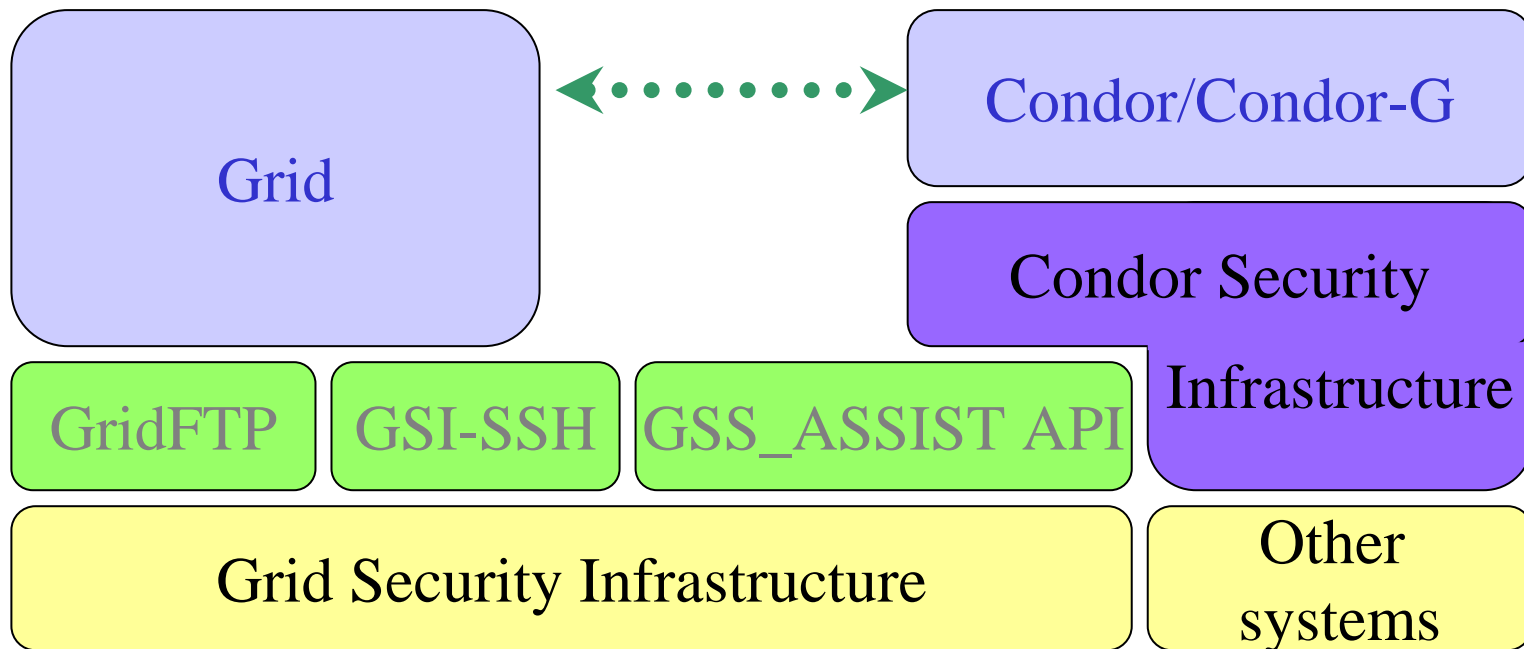
- Need to deal with credential expiration
 - Very important for unattended jobs
- Need to manage all necessary credentials
- Delegation – jobs move around

> Solutions?

- Periodically check credential and automatically notify user when the credential is about to expire
- Delegation support – only works with certain protocols such as X.509 and Kerberos



The Big Picture



An Example ...

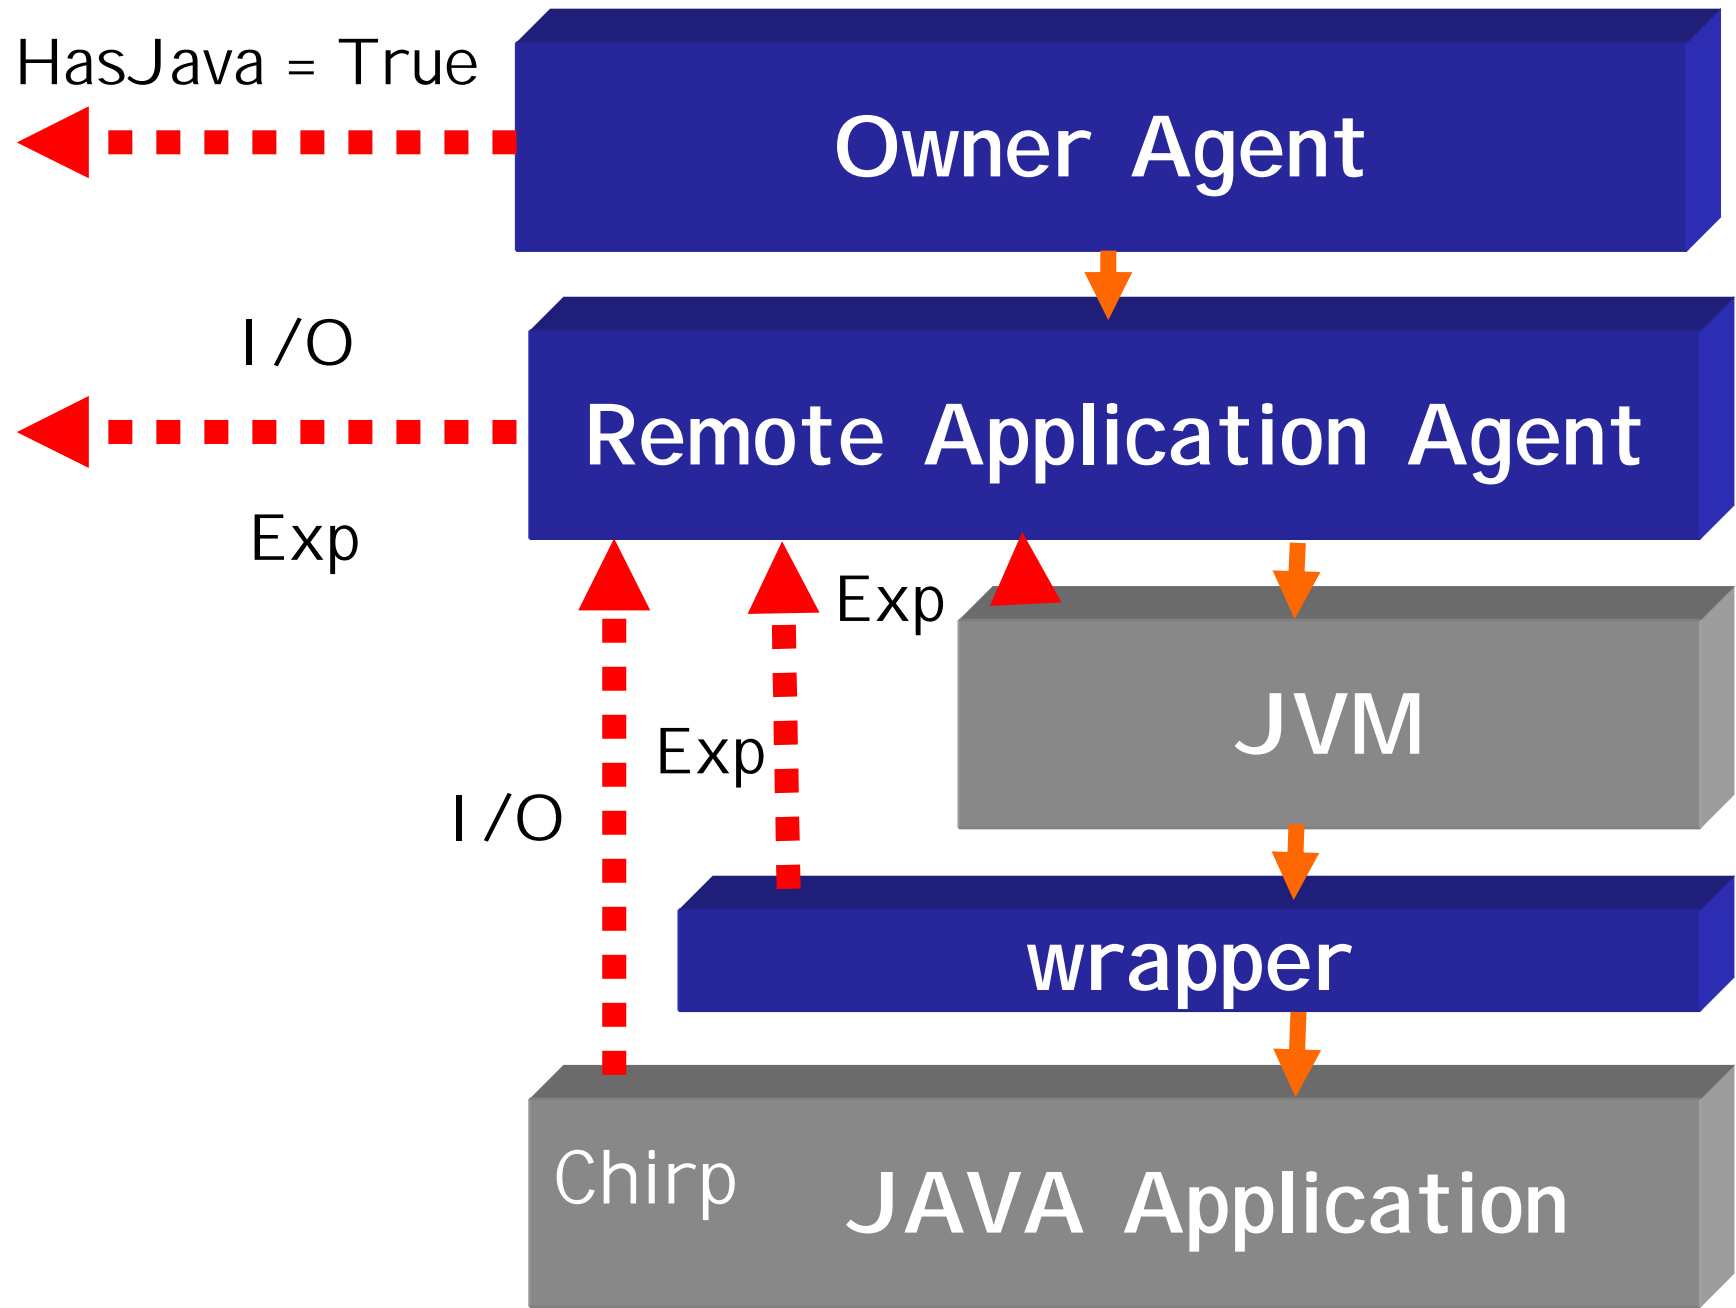


JAVA Universe

Make Condor "JAVA Aware"

- Route jobs to "JAVA capable" resources (HasJava = TRUE)
- Intercept, identify and communicate JVM errors and application exceptions back to submission point
- Support secure I/O to remote storage via a lightweight fine-grained I/O protocol (Chirp) that resembles the UNIX interface





GridFTP - Interface to Remote I/O services and Storage Appliances



GridFTP Libraries

- Used to implement server capabilities for NeST – a Grid enabled storage appliance - Provides space management and file transfer services.
- Re-implemented client services for Remote I/O support using ByPass technology and supporting Kangaroo distributed I/O services.



NeST

Develop a portable self-contained storage “appliance”

- Lot management
- User Management
- File Management
- File transfer support
- POSIX support



The Data Management Challenge ...



Data Placement (DaP) Jobs

Define, manage and schedule DaPs like any other "CPU" job.

- Includes space/lot management
- Integrate with Kangaroo technology to support "lazy" operations
- Take advantage of "disk routers" to improve throughput of "bulk" data transfers
- Logging and error recovery



Visit us at
www.cs.wisc.edu/Condor
and/or join us in
Madison for the
Paradyn/Condor meeting
(March 4-6 2002)



**Grid
Computing
is a Grid of
Technologies**

