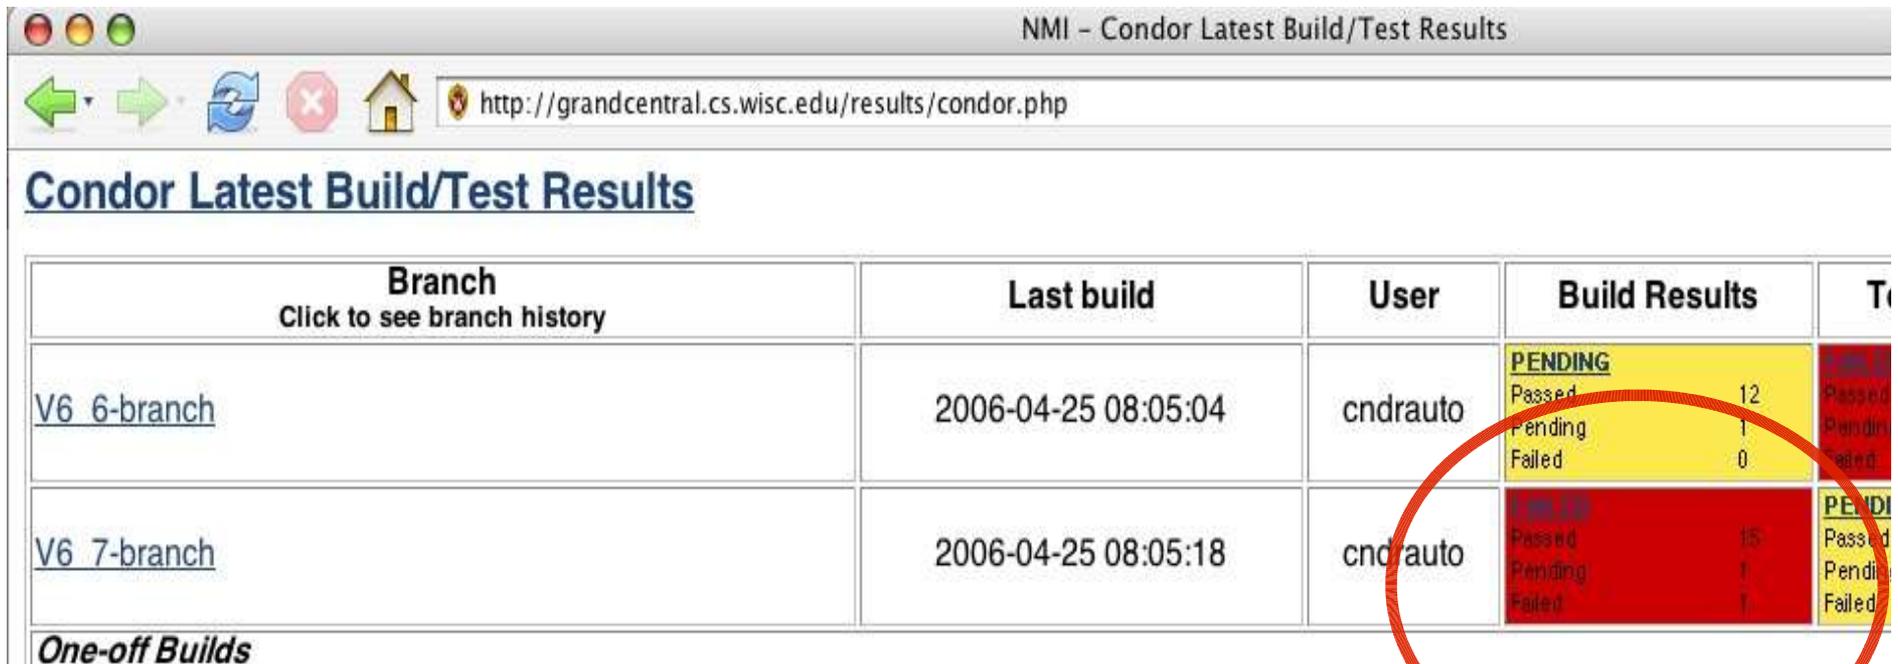


CondorWeek 2006 Summary

Alexander Withers
Technology Meeting
Brookhaven National Laboratory
May 15, 2006

What About 6.8?

- Was told possibly beginning of June 2006
- More features to be added and bug fixes needed



The screenshot shows a web browser window with the title 'NMI - Condor Latest Build/Test Results'. The address bar contains the URL 'http://grandcentral.cs.wisc.edu/results/condor.php'. The main content area is titled 'Condor Latest Build/Test Results' and contains a table with the following data:

Branch Click to see branch history	Last build	User	Build Results	Total
V6 6-branch	2006-04-25 08:05:04	cn drauto	PENDING Passed: 12 Pending: 1 Failed: 0	13
V6 7-branch	2006-04-25 08:05:18	cn drauto	FAILED Passed: 15 Pending: 1 Failed: 1	17

Below the table, the text 'One-off Builds' is visible. A red circle highlights the 'Build Results' column for the V6 6-branch and V6 7-branch rows.

Requests and Issues

- **SYSTEM_PERIODIC_REMOVE**
 - Policy to remove jobs automatically
- Negotiator swamped with too many jobs
 - AutoClustering with **SIGNIFICANT_ATTRIBUTES**
 - Now automatic in 6.7/6.8
- Schedd can now be forced to release its claim
 - Partially solved with **MaxJobRetirementTime**
 - Solved with **CLAIM_LIFETIME**

Requests and Issues, cont.

- Following resulted in “divide by zero” error: $N * (\text{Base Mips} / \text{Mips})$.
 - Can now be fixed with type casting in 6.7/6.8 ClassAds.
- STAR: “*static resource* capabilities a-la-LSF”.
 - Condor developers know of the need to address this.
 - Most likely solved in 6.7/6.8 with Group Quotas.
- Stability for Vanilla universe (more on this later)

Requests and Issues, cont.

- What about COD?
 - Other people want more functionality in COD

New Features and Highlights in 6.8.x

- Virtual Machines (VMware, Zen, etc.)
 - Checkpointing and restart capabilities
 - Several OS's on single machine
 - Finer resource control
- Group Accounting and Quotas
 - More policy nobs
- High availability features for the central manager

New Features and Highlights, cont.

- Condor now supports reestablishing connect between submit and execute node for Grid and Vanilla jobs.
 - Timeout set with JobLeaseTime.
- Condor can support schedd failover.
 - Not sure how this would work with Grid jobs.
 - Jobs should not be evicted with JobLeaseTime but this has not been tested.

New Features and Highlights, cont.

- Schedd crond/HawkEye mechanism like startd
- `BIND_ALL_INTERFACES`
- ClassAd functionality
 - Conditionals
 - String functions: `streat`, `strcmp`, instant root access!
 - More: `StrListContains`, `StrListAppend`, etc.
 - Regular expressions

New Features and Highlights, cont.

- Condor-C
 - Condor's own grid type
 - Provides for central management of all jobs
 - Distributes jobs across N submit machines
 - More efficient way to flock jobs(?)
- Schedd uses non-blocking tcp/ip connects to startd
- Improvements to negotiator and collector for speed

New Features and Highlights, cont.

- Improvements to gridmanager for scalability
- Improvements to I/O reads in schedd
- One condor_dagman for multiple dags
 - Hints that the shadow process is next
- Quill will provide centralized job data and improved response to user queries
- Native 64 bit AMD Linux port

Two Days of Talks...

- Day one: research institutions.
 - Lot of grid related talks.
 - Talks on solving specific problems.
 - Interesting usage of Condor --> obscure.
- Day two: private sector.
 - Talks about how Condor is used for business model X.
 - Low on Condor specific technical details.
 - Desperate need of Condor expertise and features.

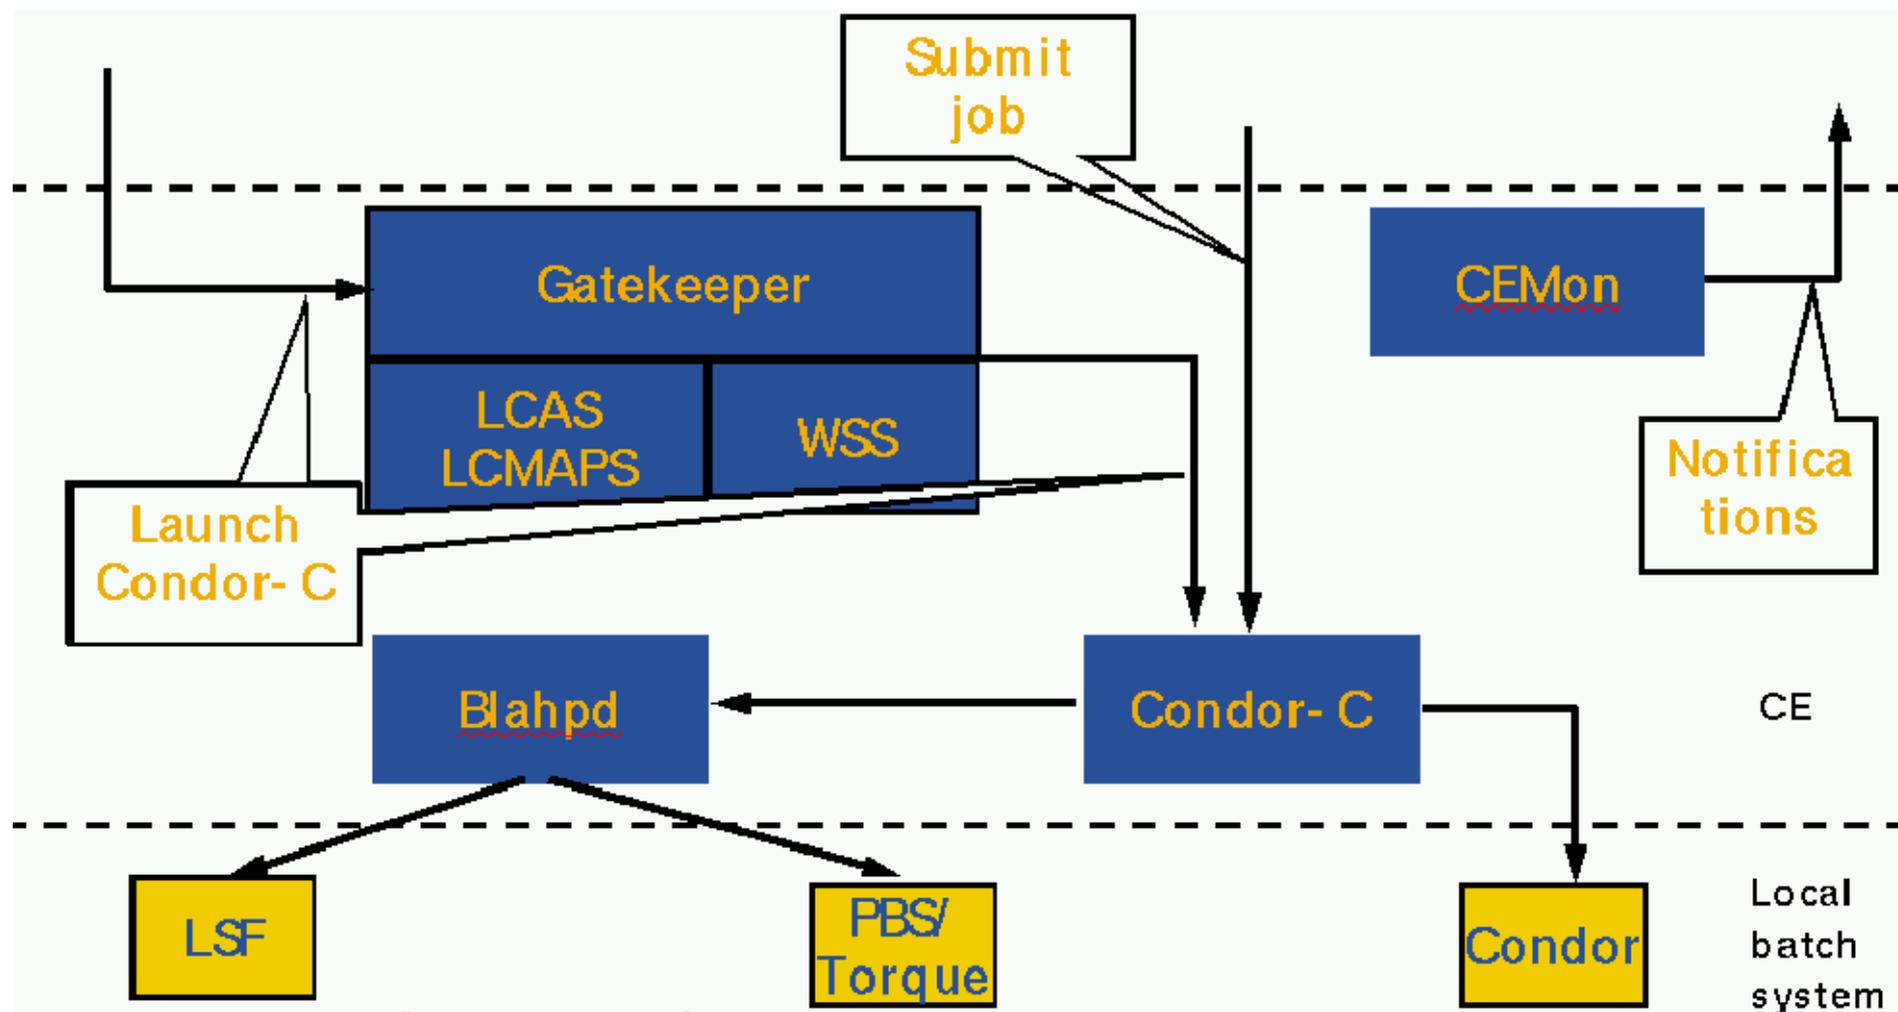
Condor Research

- Inventive uses of grid functionality to solve interesting problems
- Tactical storage
 - Allows a node to act as a file server/client without needing special privileges
 - Semantic: access via meaning not location
 - Easy on existing applications: no kernel modules, relinking, etc.

Condor Research, cont.

- Schedd on the side: routing jobs
- Vulnerability assessment being performed
- gLite
 - Using Condor-C with BLAHP for submitting jobs
 - Investigating scheme to switch user IDs on worker node
 - Investigating pilot job model using Glide-In

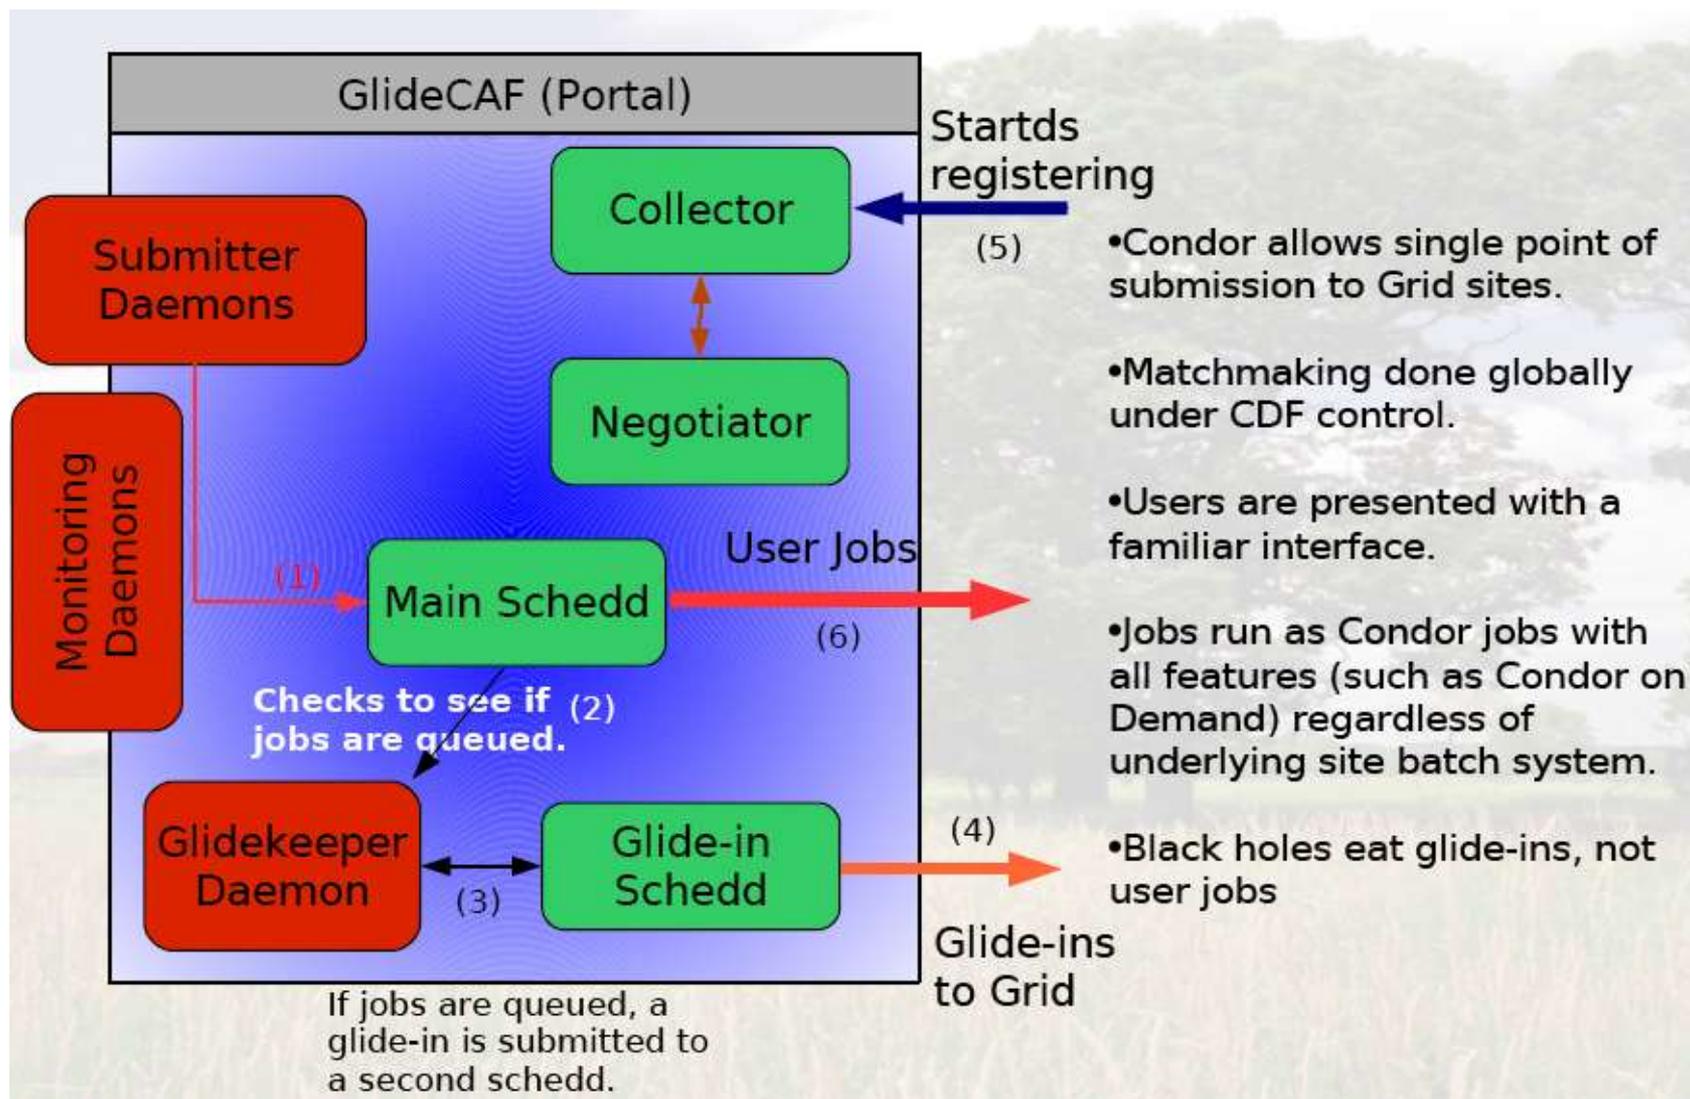
gLite schematic



Condor Research, cont.

- CDF
 - Using Condor-G, GCB, and Glide-In.
 - Use existing resources as temporary private Condor resources.
 - Glide-In validates node before accepting job.
 - Use TCP (for collector) and GCB to get around NAT and firewall issues.

CDF schematic



CondorDB

- Quill ready for production.
 - Centralized location for current queue state and history.
 - Decrease user load on the schedd.
- Quill++.
 - Add machine states and other info into the mix.
 - Better DB layout.
- CondorDB J2.
 - Everything in the database: config, state, etc.

Condor Enterprise

- Adding features to Condor
 - Contribute through third party
 - Contributions would be open and available to all
- Thousands of employees --> thousands of idle cycles
- Uses:
 - Yahoo!: web page indexing, contributions
 - Finance: quick computing once numbers are available
 - Hardware: problems with 4×10^{43} search space

Birds of a Feather Sessions

- Managing large Condor pools.
 - Developers understand that current design will not scale as clusters become larger and larger.
 - Problems reside in communication between negotiator and schedd and negotiator and startd.
 - Removing blocking I/O calls.
 - Removing superfluous tcp/ip connections.
 - Rethinking schedd/negotiator protocol.
 - Suggest dividing load between more than one central manager and using flocking or Condor-C.

Birds of a Feather Sessions, cont.

- Security in Condor.
 - Goal of reaching privilege separation.
 - Kerberos is fine but would like to move towards GSI and SSL for all authentication, authorization, and communications.
 - Security bugs.
 - Moving towards greater disclosure of security vulnerabilities.
 - Patching current versions instead of releasing new ones.

Birds of a Feather Sessions, cont.

- Condor and Grids.
 - Better Grid manager, monitor, etc.
 - New format for submit files, i.e.:
 - `grid_resource = gt4 foo.edu PBS`
`globus_xml = <queue>long</queue>`
 - Matchmaking on the Grid.
 - MyProxy support.

Matchmaking on the

- Machine Ad **Grid**

```
MyType = "Machine"  
TargetType = "Job"  
Name = "foo.edu"  
Machine = "foo.edu"  
ResourceName = "gt4 foo.edu PBS"  
UpdateSequenceNumber = 4  
Requirements = TARGET.JobUniverse == 9 && \  
    CurMatches < 10  
CurMatches = 0  
Rank = 0.0  
CurrentRank = 0.0  
WantAdReevaluate = True
```

- Job Ad

```
grid_resource = $$ (ResourceName)  
requirements = TARGET.ResourceName \  
    != Undefined  
rank = random()
```

MyProxy Support

- Submit file

```
MyProxyHost = foo.edu:12345
MyProxyServerDN = /DC=org/DC=doegrids...
MyProxyCredentialName = proxy_file
MyProxyRefreshThreshold = 240 #mins
MyProxyNewProxyLifetime = 12 #hrs
MyProxyPassword = password
```

- Or give password on command line

```
condor_submit -p password submit.desc
```

More Information

- <http://www.cs.wisc.edu/condor/CondorWeek2006/>