

OSG Campus Grids

Dr. Sebastien Goasguen, Clemson University

Outline

- **A Few examples**
- **Clemson's Pool in details**
 - Windows
 - Backfill
 - OSG
- **Other pools and CI-TEAM**
- **OSG:CElite**
- **Condor and clouds**

**~14,000 CPUs available
US-CMS Tier-2
TeraGrid site
Regional VO**

Campus Condor pool backfills idle nodes in PBS clusters - provided 5.5 million CPU-hours in 2006, all from idle nodes in clusters

Use on TeraGrid: 2.4 million hours in 2006 spent Building a database of hypothetical zeolite structures
million hours allocated to TG

November 1, 2006

Indiana receives funding for regional science computer grid

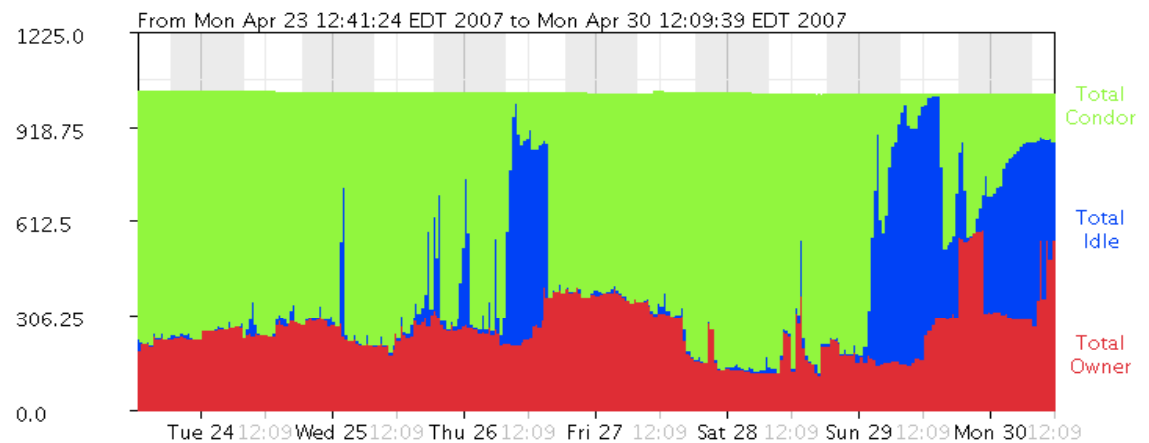
WEST LAFAYETTE, Ind. — The economy of northwest Indiana will be energized thanks to congressional authorization of \$5 million to support a computer grid for the area, officials at Purdue and Notre Dame universities announced Wednesday (Nov. 1).



Gerry McCartney

The U.S. Department of Energy's National Nuclear Security Administration will award the funding for the Northwest Indiana Computational Grid, bringing federal investment in this project to \$6.5 million.

Purdue University Condor Pool Machine Statistics for Week



Grid Laboratory of Wisconsin (GLOW)

- Users submit jobs to their own private or department scheduler as members of a group (e.g. “CMS” or “MedPhysics”)
 - Jobs are dynamically matched to available machines
 - Jobs run preferentially at the “home” site, but may run anywhere when machines are available
 - Computers at each site give highest priority to jobs from same group (via machine RANK)
- Crosses multiple administrative domains
 - No common uid-space across campus
 - No cross-campus NFS for file access

Grid Laboratory of Wisconsin (GLOW)

Housing the Machines

- Condominium Style
 - centralized computing center
 - space, power, cooling, management
 - standardized packages
- Neighborhood Association Style
 - each group hosts its own machines
 - each contributes to administrative effort
 - base standards (e.g. Linux & Condor) to make easy sharing of resources
- GLOW and Clemson have elements of both

Clemson's pool

- **Clemson's Pool**

- Originally mostly Windows, +100 locations on campus.
- Now 6,000 linux slots as well
- Working on 11,500 slots setup, ~120 TFlops
- Maintained by Central IT
- CS dpt tests new configs
- Other dpt adopt the Central IT images
- BOINC Backfill to maximize utilization.
- Connected to OSG via an OSG CE.

Total Owner Claimed Unclaimed Matched Preempting Backfill

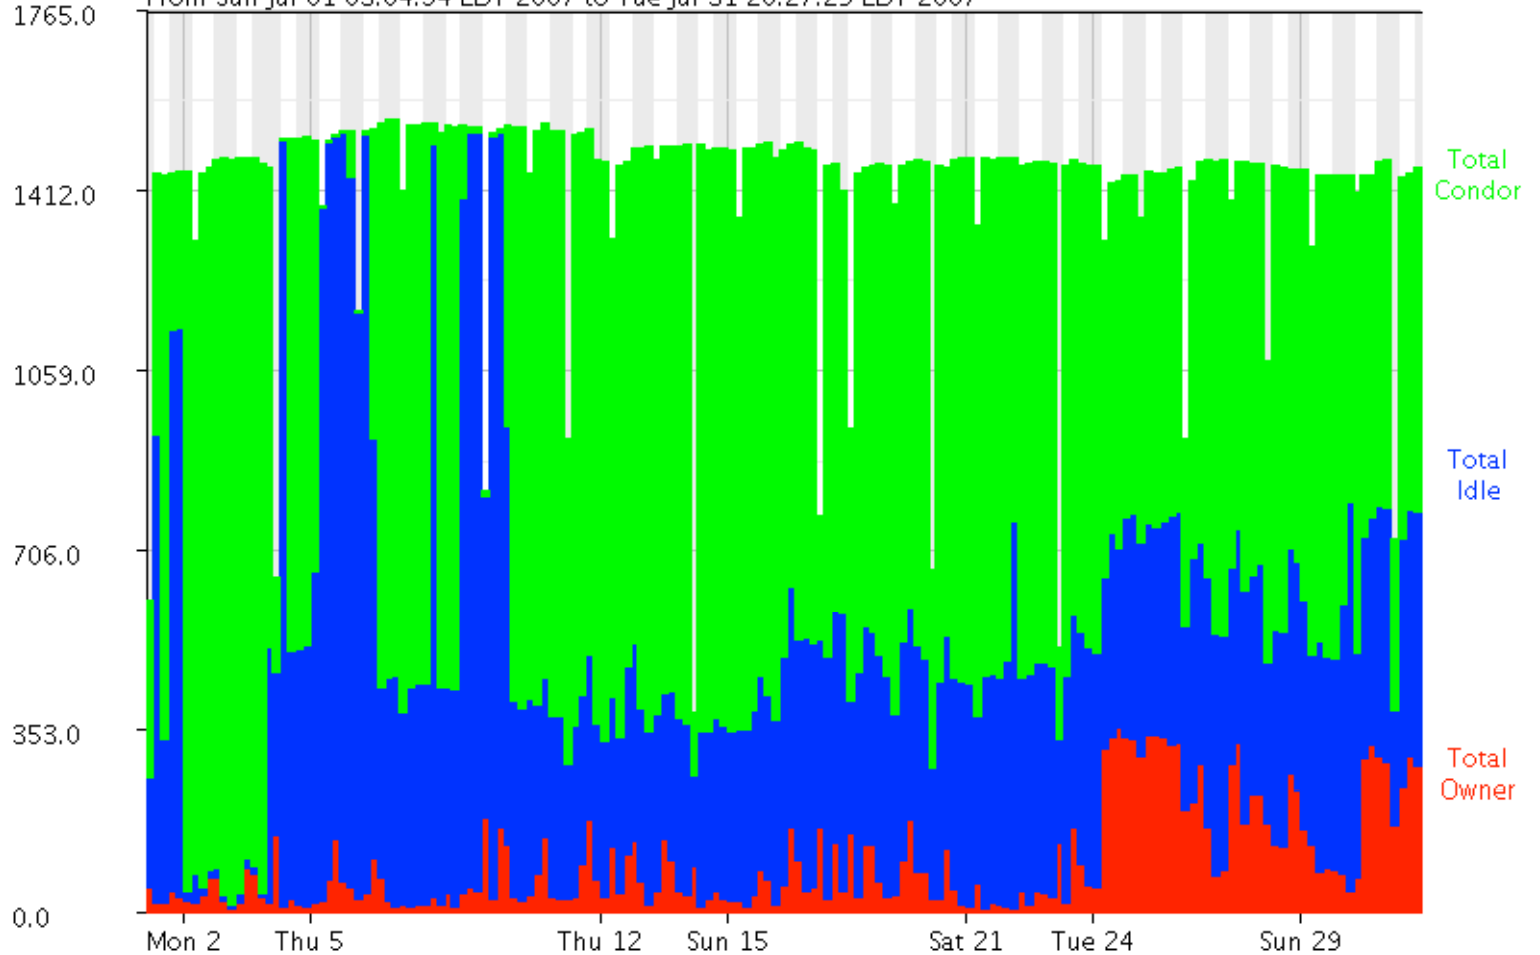
INTEL/LINUX	4	0	0	4	0	0	0
INTEL/WINNT51	895	448	3	229	0	0	215
INTEL/WINNT60	1246	49	0	2	0	0	1195
SUN4u/SOLARIS5.10	17	3	0	14	0	0	0
X86_64/LINUX	26	2	3	21	0	0	0
Total	2188	502	6	270	0	0	1410

Clemson's pool history




Clemson University Condor Pool Machine Statistics for Jul

From Sun Jul 01 03:04:54 EDT 2007 to Tue Jul 31 20:27:29 EDT 2007



Clemson University



Configure...

Graph Fints: The Y-axis is number of machines, press

Arch

Total

X86_64/LINUX

INTEL/WINNT51

Clemson's pool BOINC backfill



- **Put Clemson in World Community Grid, LHC@home and Einstein@home.**
- **Reached #1 on WCG in the world, contributing ~4 years per day when no local jobs are running**

```
# Turn on backfill functionality, and use BOINC
ENABLE BACKFILL = TRUE
BACKFILL_SYSTEM = BOINC
```

```
BOINC Executable = C:\PROGRA~1\BOINC\boinc.exe
BOINC_Universe = vanilla
```

```
BOINC Arguments = --dir $(BOINC HOME) --attach project
http://www.worldcommunitygrid.org/
cbf9dNOTAREALKEYGETYOUROWN035b4b2
```



Clemson's pool BOINC backfill



- Reached #1 on WCG in the world, contributing ~4 years per day when no local jobs are running = Lots of pink

Statistics

By Members



Member Name:

Total Run Time (y:d:h:m:s):

[clemsontiger](#)

3:196:10:29:02

[marist_college](#)

3:122:18:16:54

[RCTCGrid](#)

3:042:15:06:14

OSG VO through BOINC



LHC @home

New participants

- [Join LHC@home](#)
- [About LHC@home](#)

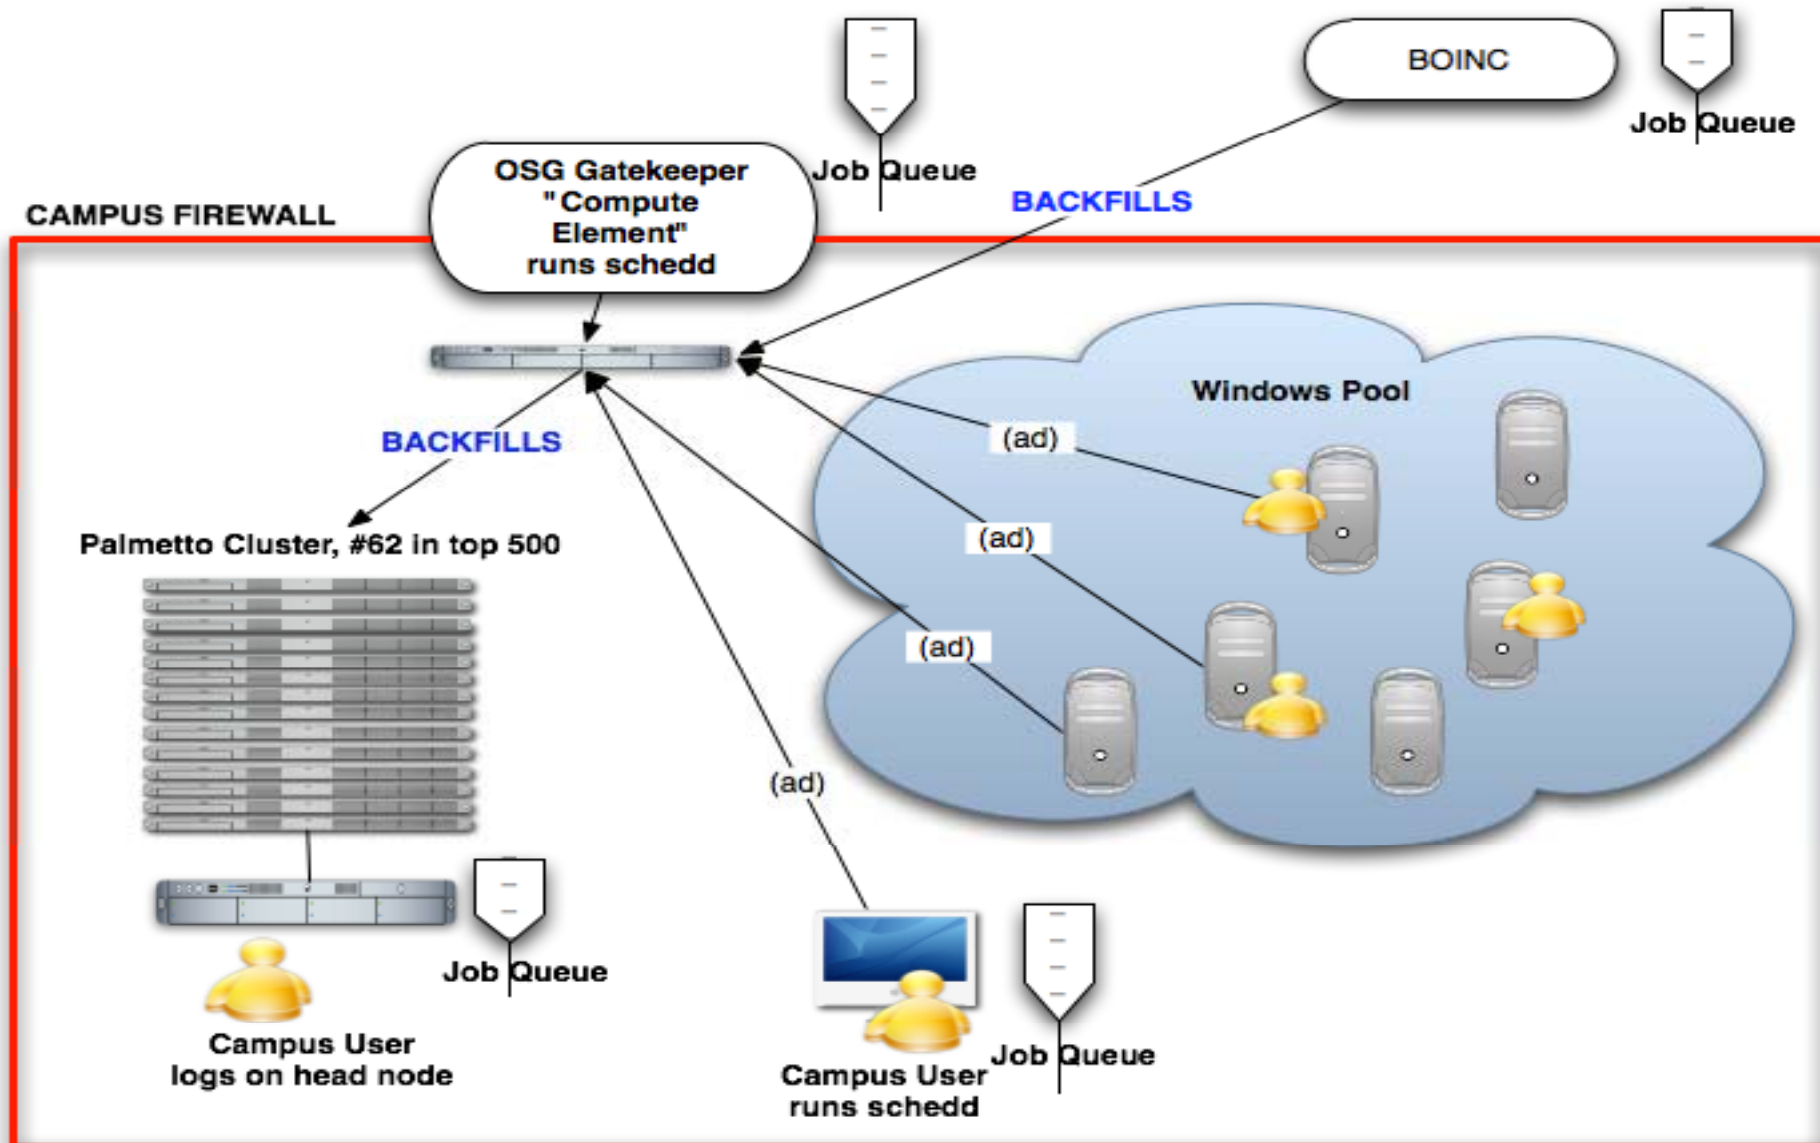
Rank	Name	Recent average credit	Total credit	Country	Participant since
1	neu-innova	9,531.74	259,340	Cyprus	19 Oct 2008 14:50:00 UTC
2	clemsonTiger	7,350.43	238,703	United States	6 Oct 2008 14:55:13 UTC

Rank	Name	Recent average credit	Total credit	Country	Participant since
1	ATLAS AEI Hannover	1,620,293.07	87,238,600	Germany	4 Sep 2008 13:31:20 UTC
2	AEI eScience group, for the German Grid (D-Grid) and the Open Science Grid (OSG)	1,560,489.91	684,343,248	Germany	1 Feb 2007 17:05:25 UTC
3	Armin Burkhardt speaking for MPI/FKF	334,166.28	60,726,407	Germany	21 Feb 2005 16:22:20 UTC
4	 Nemo	333,212.23	26,069,038	United States	30 Jul 2008 21:00:41 UTC
5	clemsonTiger	235,696.84	12,984,027	United States	6 Oct 2008 14:56:26 UTC

- Einstein@home, LIGO VO
- LHC@home, very little jobs to grab
- Could we count BOINC work for OSG VO led project into OSG accounting. A.k.a count jobs not coming through the CE.

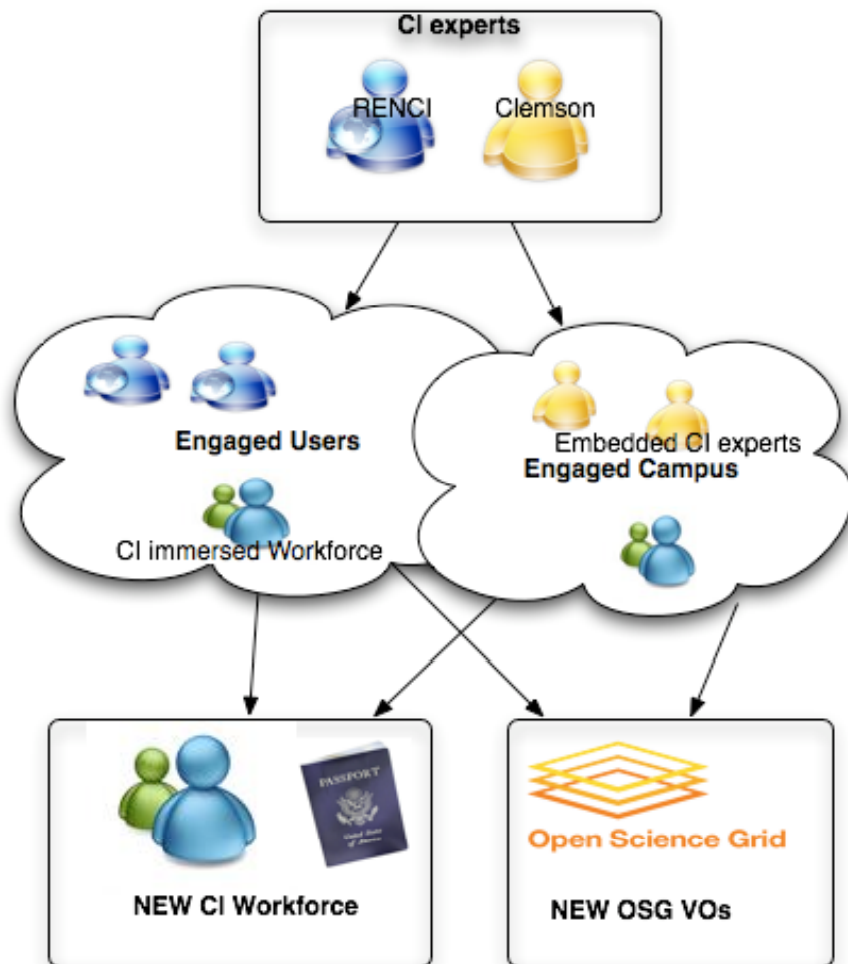
Clemson's pool on OSG

- Multi-tier job queues to fill the pool
- Local users, then OSG, then BOINC



Other Pools and CI-TEAM

- **CI-TEAM** is a NSF award to outreach to campuses, help them build their cyberinfrastructure and make use of it as well as the national OSG infrastructure. *“Embedded Immersive Engagement for Cyberinfrastructure, EIE-4CI”*




- Provide help to build cyberinfrastructure on campus
- Provide help to make your application run on “the Grid”
- Train experts
- <http://www.eie4ci.org>

Other Pools and CI-TEAM

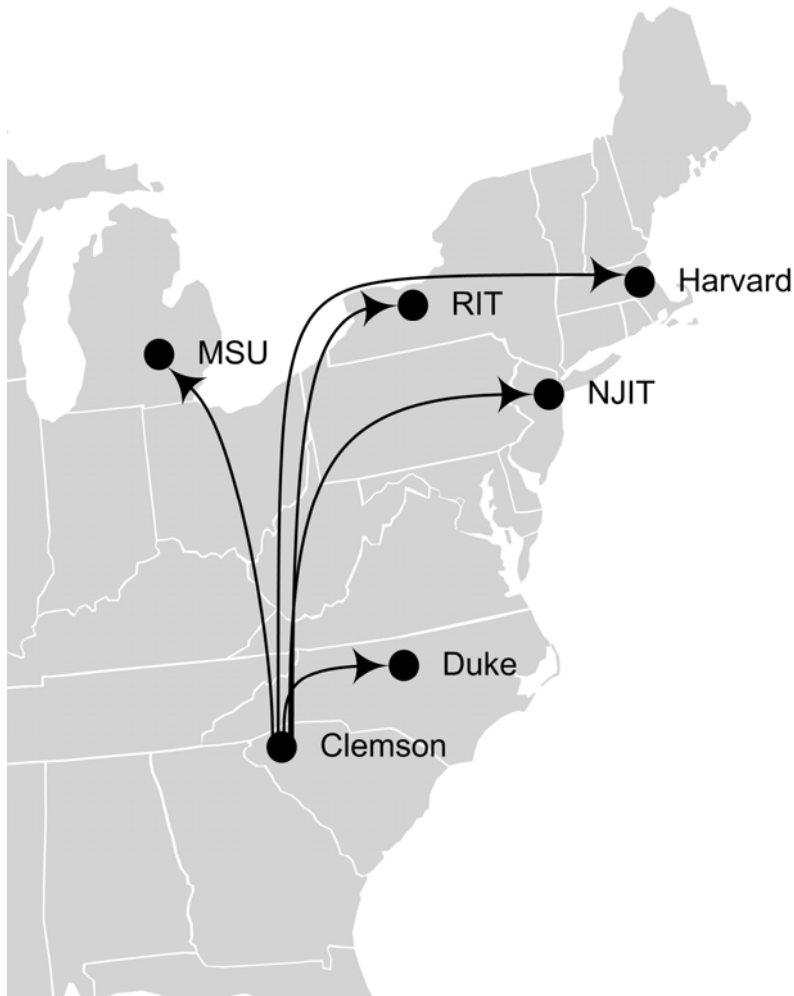
- **Other Large Campus Pools**
 - Purdue –14,000 slots (Led by US-CMS Tier-2).
 - GLOW in Wisconsin (Also US-CMS leadership).
 - FermiGrid (Multiple Experiments as stakeholders).
- RIT and Albany have created +1,000 pools after CI-days in Albany in December 2007

University at Albany - Research IT Group Condor Pool Machine Statistics for Aug



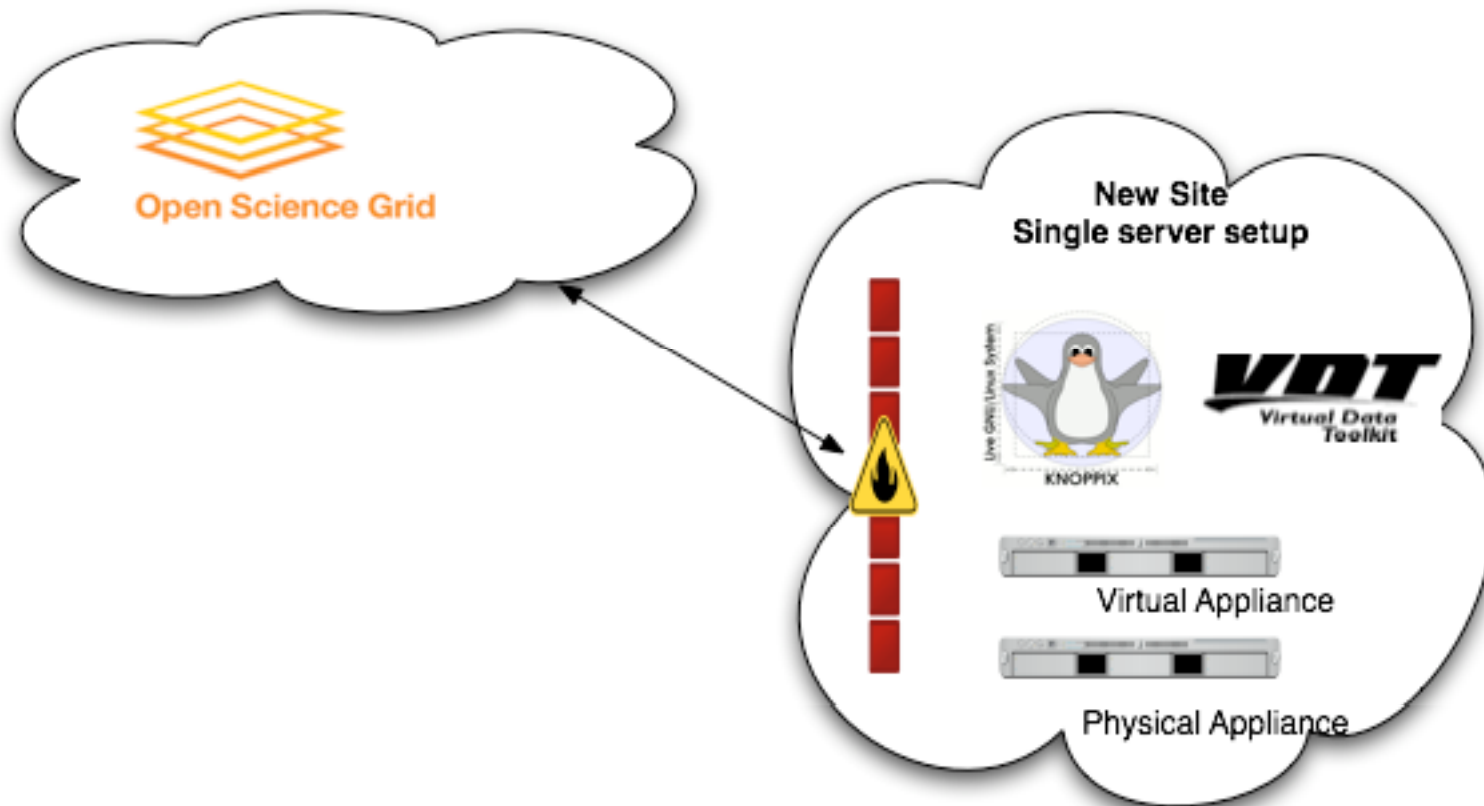
Campus sites “levels”

- Different level of efforts, different commitments, different results. How much to do?
 - Duke, ATLAS Tier-3. One day of work, not registered on OSG
 - *Harvard, SBGrid VO. Weeks/Months of work, registered, VO members highly engaged*
 - *RIT, NYSgrid VO, regional VO. Windows based Condor pool, BOINC backfill.*
 - *SURAGRID, interop partnership, different CA policies.*
 - ***Trend towards Regional “Grids” (NWICG, NYSGRID, NJEDGE SURAGRID, LONI...) leverage OSG framework to access more resources and share there own resources.***



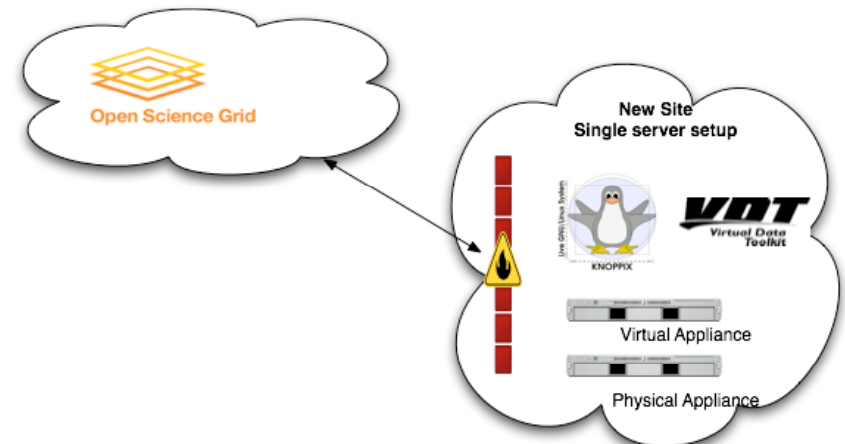
OSG:CElite

- **Low-level entry to OSG CE (and SE in the future). What is the minimum required set of software to setup a OSG CE ?**
- **Physical appliance, Virtual appliance, Live CD, new VDT cache or new P2P network with separate security.**



OSG:CElite

- **Physical appliance:** Prep machine, configure software, ship machine, receiving site just turns it on.
- **Virtual appliance:** Same as physical but no shipping, no buying of machines
- **Live CD:** size of the image ?
- **VDT cache:** pacman `-get OSG:CElite`
- **Problems:** Drop in valid certificates for hosts, registration of the resource. Use a different CA to issue these certs ?
- **P2P network of Tier-3s,** create a “VPN” and create an isolated testbed for sys admin ...more of an academic exercise..



What software ?

```
● # vdt-control -list
```

● Service	Type	Desired State
● -----	+-----	+-----
● fetch-crl	cron	enable
● vdt-rotate-logs	cron	enable
● gris	init	do not enable
● globus-gatekeeper	inetd	enable
● gsiftp	inetd	enable
● mysql	init	enable
● globus-ws	init	enable
● edg-mkgridmap	cron	do not enable
● gums host cron	cron	do not enable
● MLD	init	do not enable
● vdt-update-certs	cron	do not enable
● condor-devel	init	enable
● apache	init	enable
● osg rsv	init	do not enable
● tomcat-5	init	enable
● syslog-ng	init	enable
● gratia-condor	cron	enable

Condor and Clouds

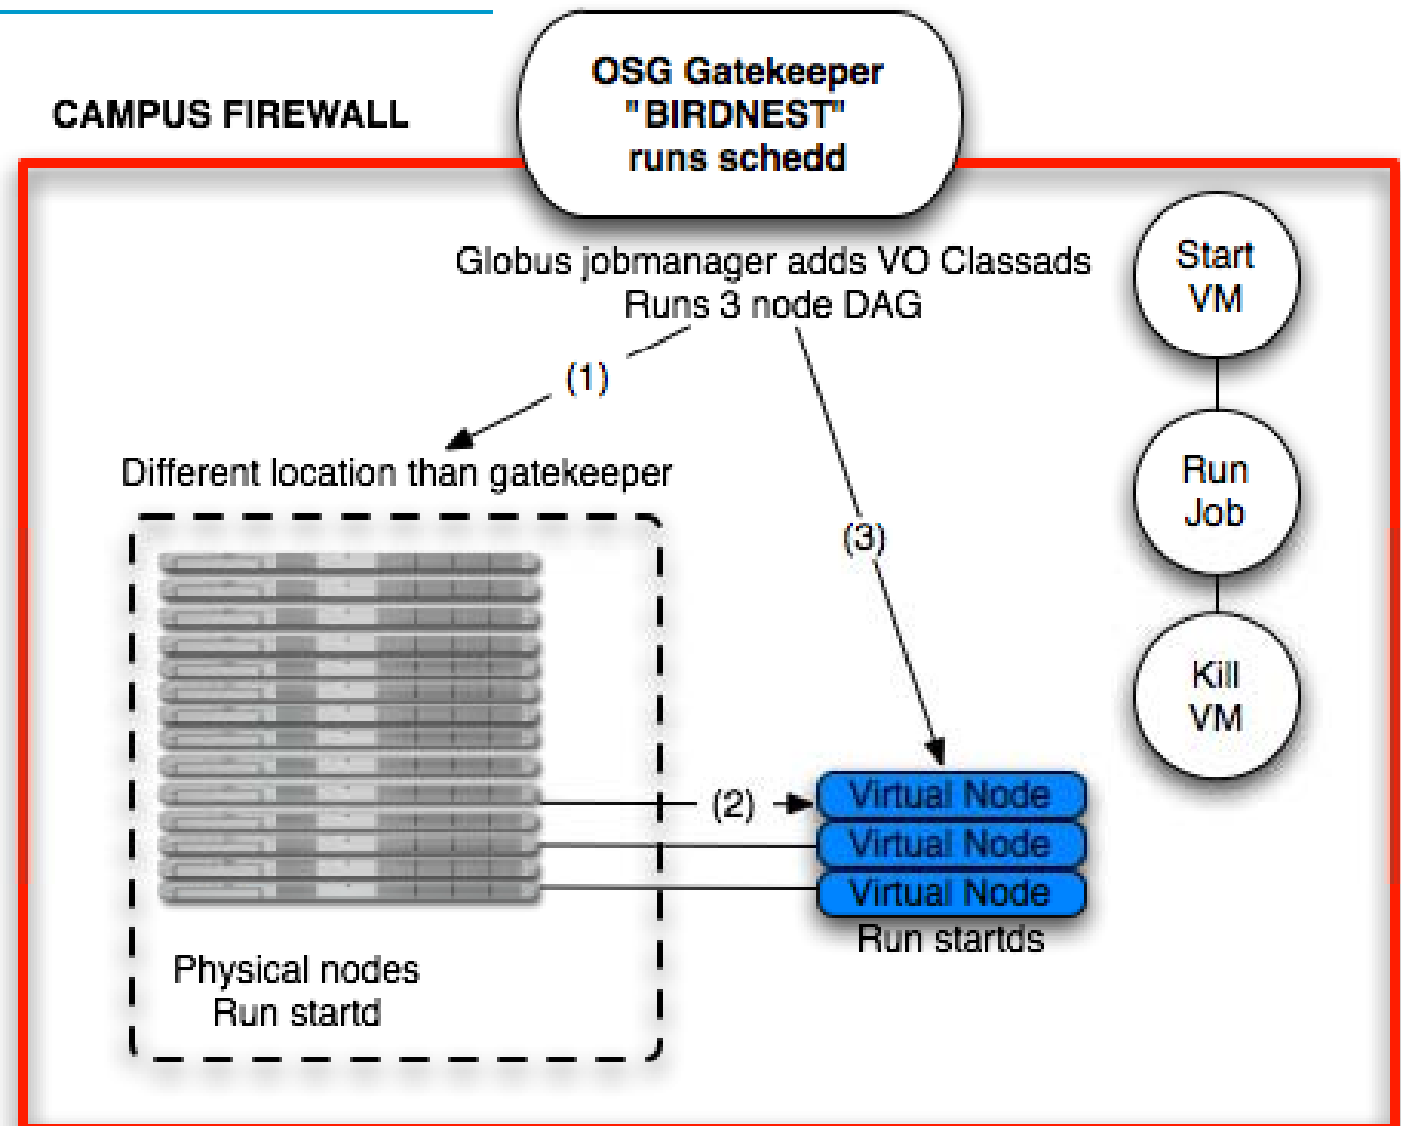
- **For us clouds are clusters of workstations/servers that are dedicated to a particular Virtual Organization.**
- **Their software environments can be tailored to the particular needs of a VO.**
- **They can be provisioned dynamically.**
- **Condor can help us build clouds:**
 - **Ease to target specific machines for specific VOs with classads**
 - **Ease of having adding nodes to clouds by sending ads to collectors.**
 - **Ease to integrate with existing grid computing environments, OSG for instance.**
- **Implementation:**
 - **Use virtual machine (VM) to provide different running environment for each VO. Each VM advertized with different classads**
 - **Run Condor within the VMs**
 - **Start and Stop VMs depending on job load**

Condor and Clouds



- VM as a job
- Job “glides” in VM
- VM destroyed
- “VPN” for all VMs
- Different OS/sw for each VO
- Use EC2...
- Use VM universe...


- Under test as we speak



- Use IPOP (<http://www.grid-appliance.org/>) to build WAN “VPN” that traverse NATs. Ability to isolate clouds in different address space.

Acknowledgements

lots of folks at clemson...Dru, Matt, Nell, John-Mark, Ben...
lots of condor folks: Miron, Todd, Alain, Jaime, Dan, Ben,
Greg....



questions?

sebgoa@clemson.edu

<http://cirg.cs.clemson.edu>

yum repo for condor...