

# Two Extremes: Grid Computing and Globus VS.

## Massively parallel computing for big science

William (Bill) E. Allcock

Director of Operations, Manager, Advanced Integration Group

Argonne National Laboratory, Leadership Computing Facility

With slide assistance from: Raj Kettimuthu (ANL), Miron Livny(UW – Madison), Robert Jacob (ANL), Todd Tannenbaum (UW-Madison), Katherine Riley(ANL)

# Administrivia, disclaimers, etc..

- First, my goal here is to give you what you want. If what I am presenting is not on the mark, please let me know and I will redirect as best I can.
- I am quite capable of talking for the entire time, but this is going to be a lot more fun, and probably useful, if it turns into a conversation.
- We are propeller-heads. We WANT to see bare metal, push all the buttons and twist all the knobs. I have tried to adjust my point of view to what a, perhaps more risk averse commercial point of view might be, but my perspective might be a bit skewed.
- NO ENDORSEMENTS EXPRESS OR IMPLIED
  - I am going to mention vendors at various points in this presentation. That will be for one of two reasons:
    - It is a vendor we use and I am describing our infrastructure
    - It is a vendor I am aware of in a space that seemed to be of interest to you.
  - There are two very good reasons that these are not endorsements
    - Argonne does not permit it and I like my job 😊
    - What works or is good is completely situation dependent and it would be foolish of me to try and endorse one of these products not knowing anything about your situation.



# Argonne National Laboratory (ANL)

- Argonne's mission is to apply a unique mix of world-class science, engineering and user facilities to deliver innovative research and technologies. We create new knowledge that addresses the most important scientific and societal needs of our nation.
- Argonne's strengths include
  - Energy Storage (Chevy Volt, Li Ion technology)
  - Alternative Energy (Solar, Wind, biofuels, fuel and energy efficiency)
  - Nuclear Energy (Safe reactor technology, fuel cycle management)
  - Biological and Environmental systems (genomics, climate)
  - National Security (forensics, decision sciences, sensors, materials, cyber security)
  - Scientific User Facilities
    - Advanced Photon Source (APS)
    - Center for Nanoscale Materials (CNM)
    - Argonne Tandem Linac Accelerator System (ATLAS)
    - Electron Microscopy Center (EMC)
    - Transportation Research and Analysis Computing Center (TRACC)
    - Atmospheric Radiation Measurement Climate Research Facility (ARM)
    - Leadership Computing Facility (LCF - that's my division)
- [http://en.wikipedia.org/wiki/United\\_States\\_Department\\_of\\_Energy\\_National\\_Laboratories](http://en.wikipedia.org/wiki/United_States_Department_of_Energy_National_Laboratories)



# Leadership Computing Facility

- The ALCF's mission is to accelerate major scientific discoveries and engineering breakthroughs for humanity by designing and providing world-leading computing facilities in partnership with the computational science community.

We strive to be the forefront computational center for extending the frontiers of science by solving key problems for the nation that require innovative approaches and the largest-scale systems.

- Currently
  - Intrepid: 557 TFLOP Blue Gene/P, 40K nodes, 160K cores, 80 TB RAM, 8PB of disk, (10) 512 port 10G switches, approximately 300 other servers.
- Deploying this year
  - Mira: 10PFLOP Blue Gene/Q, 48K nodes, 768K cores, 768 TB RAM, 28PB disk, (4) 648 port 40G Infiniband switches, 120+ other servers (file servers are now embedded in the storage).
- Our facility is open to just about anyone (except the “State Sponsors of Terrorism”). This includes commercial companies (Pratt and Whitney, Proctor and Gamble, GE have all run on our machine). There is no charge. It is given as a grant to the best science with the greatest impact. Industry outreach is a focus right now. Talk to me offline if you are interested.





## Two Extremes... Both are big, but moving them would require very different solutions...



- Many small things
- The Sahara –  $8 \times 10^{27}$  grains of sand
- Google – tens of millions of queries a day
- These problems require little or no coordination (communication).
- This is Grid / High Throughput Computing (HTC)



- One huge thing
- Ayers Rock – 1 rock weighing 4M tons
- Quantum Chromodynamics (QCD) – 100s of millions of core hours on one problem.
- These problems are bounded by coordination (communication).
- This is massively parallel computing

# Grid Computing and Globus

# So what is Grid Computing?

- The precise definition was laid out in a paper called “The Anatomy of the Grid” by Ian Foster (ANL), Carl Kesselman (USC/ISI), Steve Tuecke (ANL).
  - flexible, secure, coordinated resource sharing among dynamic collections of individuals, institutions, and resources—what we refer to as virtual organizations.
- What instantiation was chosen?
  - Grid Computing is a protocol / service centric view of the world. The idea was to create a standard set of protocols and services and secondarily, interfaces to them, so that there was a uniform way to access any resource that supported them.
  - This is very similar to the Internet model (IP, TCP, HTTP) and in contrast to things such as Cloud Computing (Infrastructure as a Service, IaaS), Science Portals (Software as a Service, SaaS), etc.. However, they are not mutually exclusive. You could run your Grid services on infrastructure spun up in a cloud (there are companies that do this)
- Where are we today?
  - If you read the paper, most of what is referred to as the Resource Layer was completed and is what is in use today (I will give more details). The next level up, referred to as the collective layer, had many attempts, fits, and starts, but only a couple of things ever really gained traction.
- Anyone considering using Grid Computing should definitely read the Anatomy of the Grid. The details of Globus have changed, but much of it still applies.
  - <http://www.globus.org/alliance/publications/papers/anatomy.pdf>



# Batch Computing

- If everyone in the room is familiar with this, let me know and I will skip this slide, but Grid / Scientific / High Performance Computing in general is based on batch computing and much of what I say won't make sense otherwise.
- Batch vs. Interactive
  - batch<sup>1</sup> (bch)*n.*
    - **4. Computer Science** A set of data or jobs to be processed in a single program run.
    - *tr.v.* **batched, batch·ing, batch·es** To assemble or process as a batch.
  - in·ter·ac·tive (ntr-ktv)*adj.*
    - **1.** Acting or capable of acting on each other.
    - **2. Computer Science** Of or relating to a program that responds to user activity.
- As a metaphor this is a sit-down restaurant vs. a cafeteria
  - In the sit-down restaurant you place your order, wait a while, and it is brought to you.
  - In the cafeteria, it is on demand. No waiting.
- The idea is that you have a valuable resource that is under-provisioned or in contention (compute cluster).
  - To get access to this resource, you execute a command that tells the scheduler how many nodes you want, what executable to run, for how long, etc (you place your order).
    - `qsub -t 20 -n 1024 /bin/sleep 600`
  - The scheduler then plays a big game of Tetris and runs the jobs as optimally as possible according to local policy (your order is prepared and brought to you)



# So, what is Globus?

- It is a project and as with most successful projects, it started over a beer 😊
  - Ian Foster and Carl Kesselman were in San Francisco and discussed how they could merge their complimentary research projects.
- Like any large, long lived project (about 15 years now) it has changed over time.
- It is also a *toolkit* for enabling Grid Computing
- The things that would likely matter to you today are:
  - Security
    - Grid Security Infrastructure (GSI) – X.509 certificates with delegation
  - Job Submission
    - Grid Resource Allocation and Management (GRAM) - Develop a standard job submission interface and then write “shims” to translate to the various schedulers out there (PBS, LSF, LoadLeveler, Maui, Cobalt, etc)
  - Data Management
    - GridFTP - secure, high speed, and reliable movement of large datasets
    - Globus Online (GO) – SasS offering of a highly reliable data transfer service
      - GridFTP was reliable unless the client failed in which case all state was lost. GO addresses that.
  - The Ecosystem that has developed around it
    - Myproxy (NCSA) – now comes with the toolkit; allows more centralized control of X.509 certificates.
    - Workflow management – Swift (UC), Pegasus (USC/ISI), and particularly Condor (UW-Madison)







# Open Science Grid (OSG) HTC at the National Level



Chicago Uniform Meeting, 28 Feb, 2012

# Large Hadron Collider

1800 Physicists, 150 Institutes, 32 Countries



100 PB of data by 20XX; 50,000 CPUs?





# The LIGO Project



- Laser Interferometer Gravitational Wave Observatory
- LIGO instruments in Washington State and Louisiana
- During science runs, produce up to 2 terabytes per day
- Published along with metadata at Caltech (archival site)
- Replicated at up to 10 other LIGO sites
  - LIGO scientists typically move data sets near to computational clusters at their sties
- The LIGO Data Grid



# Earth System Grid Federation (ESGF)

The ESGF consortium's **mission** is to provide climate researchers worldwide with a system of federated science gateways to access data, information, models, analysis tools, and computational capabilities required to evaluate ultrascale data sets. Its goals are to make data more useful to climate researchers by developing collaborative technology that enhances data usability, and provide a universal and secure Web-based data access portal for broad-based multi-model data collections.

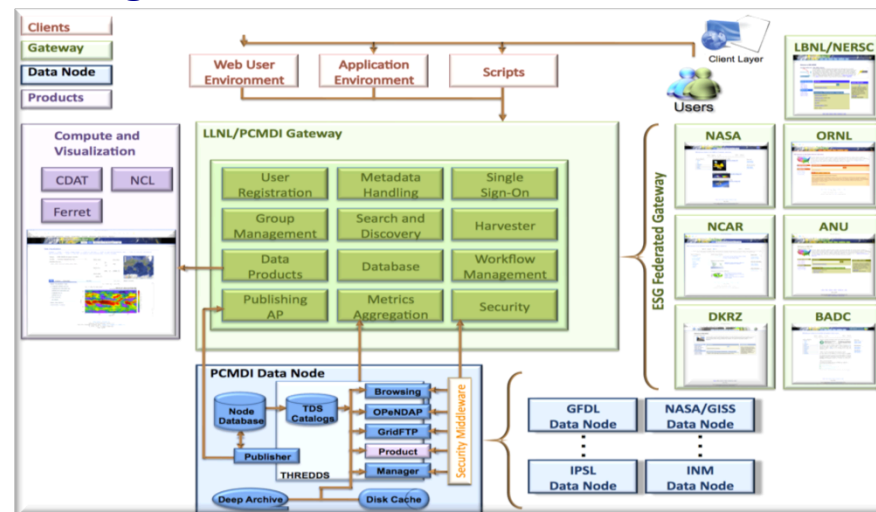

## Objectives

- Meet the specific needs of national and international climate projects for distributed databases, data access, and data movement.
- Provide a wide range of climate data-analysis tools and diagnostic methods to international climate centers and U.S. government agencies.
- Integrate in a collaborative problem-solving environment highly publicized climate data sets using distributed storage management, remote high-performance units, high-bandwidth wide-area networks, and user desktop platforms.

## Impact

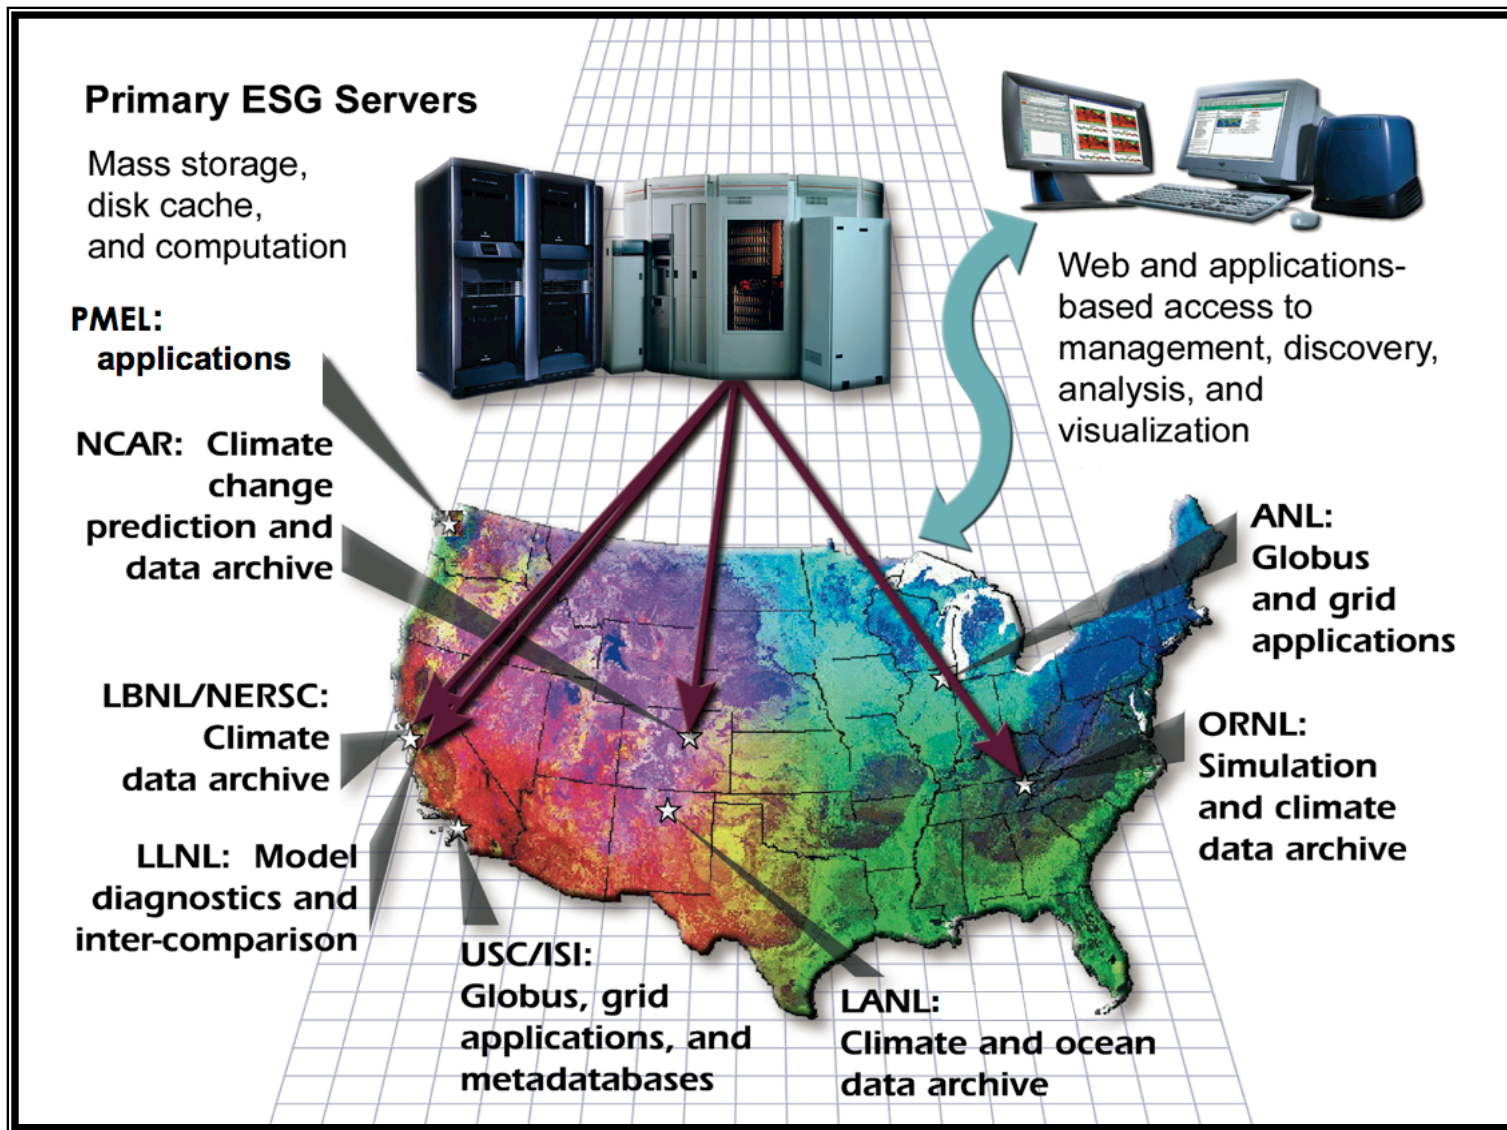
- Massive data archives (PB moving to XB)
- Multiple data centers worldwide
  - Existing IT infrastructure and separate security domains
- Heterogeneous data sources (models, observations, reanalysis)
- Multiple physical realms (atmosphere, ocean, land, sea ice)
- Multiple data, metadata formats and conventions
- Multiple scales (global, regional and local)
- Cyber security
- Multiple audiences (scientists, policy makers, students, educators)

## Deployment of Climate Data Management and Access


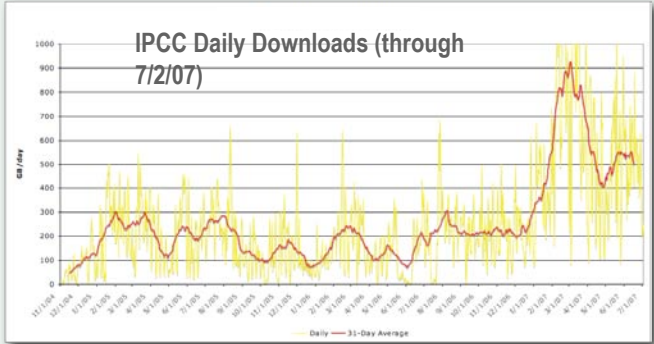


Contact Dean N. Williams ([williams13@llnl.gov](mailto:williams13@llnl.gov)) for more information

# Earth System Grid

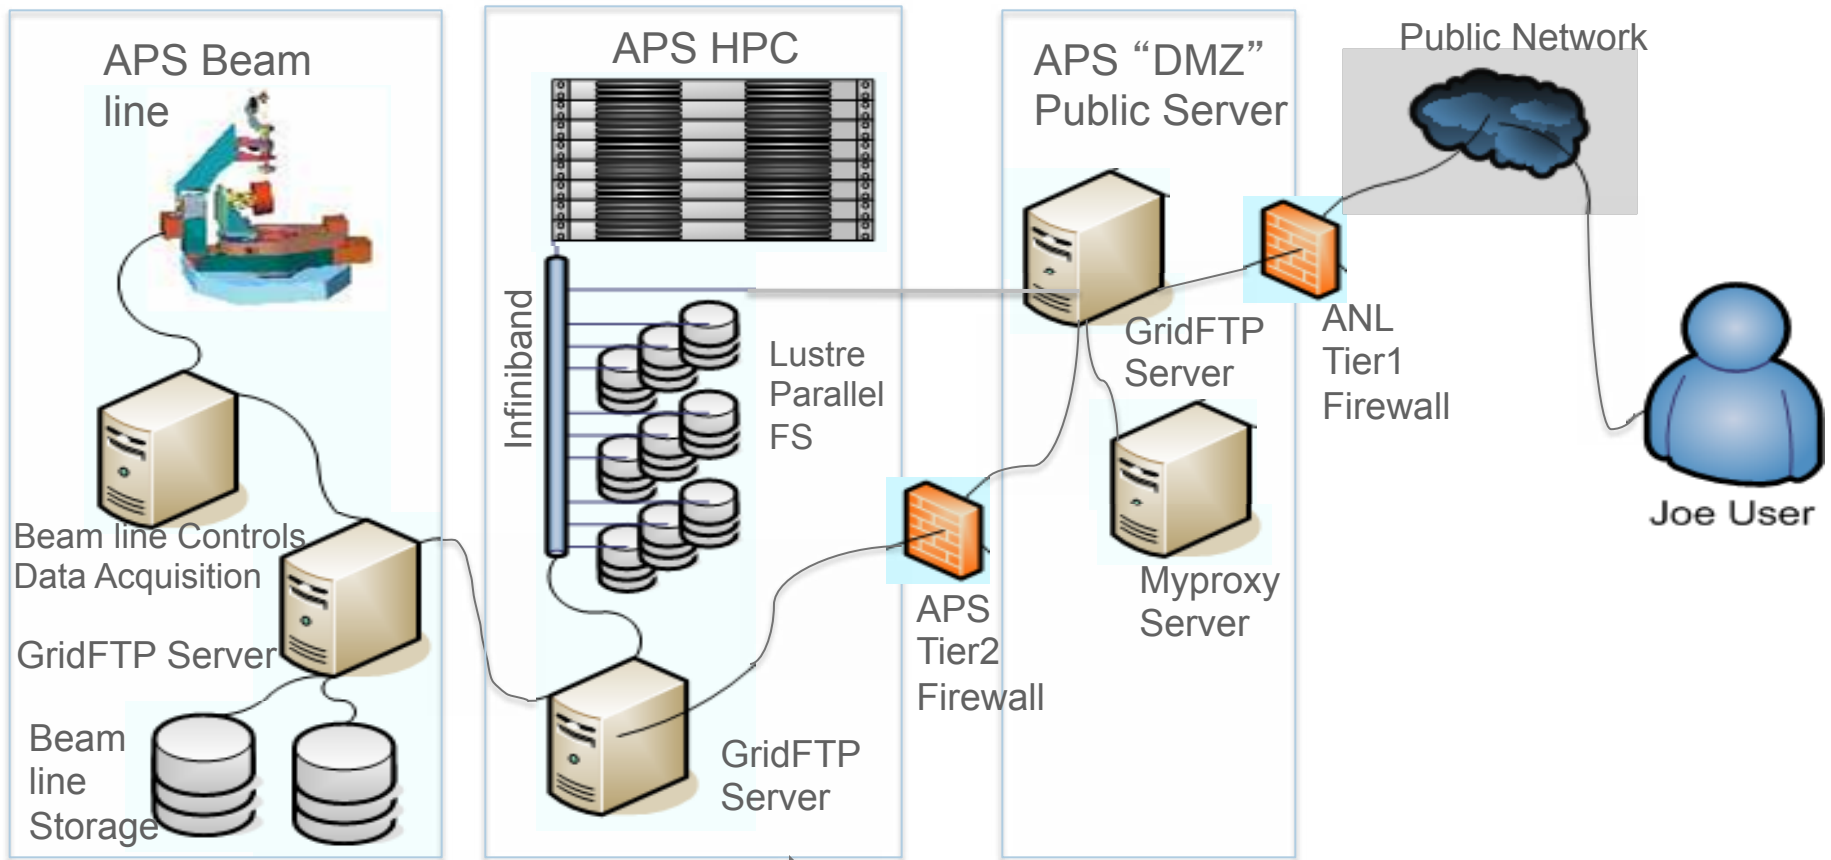


# ESG facts and figures

Main ESG Portal	IPCC AR4 ESG Portal
<p>146 TB of data at four locations</p> <ul style="list-style-type: none"> <li>● 1,059 datasets</li> <li>● 958,072 files</li> <li>● Includes the past 6 years of joint DOE/NSF climate modeling experiments</li> </ul>	<p>35 TB of data at one location</p> <ul style="list-style-type: none"> <li>● 77,400 files</li> <li>● Generated by a modeling campaign coordinated by the Intergovernmental Panel on Climate Change</li> <li>● Model data from 13 countries</li> </ul>
4,910 registered users	1,245 registered analysis projects
<p>Downloads to date</p> <ul style="list-style-type: none"> <li>● 30 TB</li> <li>● 106,572 files</li> </ul>	<p>Downloads to date</p> <ul style="list-style-type: none"> <li>● 245 TB</li> <li>● 914,400 files</li> <li>● 500 GB/day (average)</li> </ul>
 <p>Worldwide ESG user base</p>	 <p><b>&gt; 300 scientific papers published to date based on analysis of IPCC AR4 data</b></p>



# Advanced Photon Source



Beam line controls pushes data to HPC

User gets credentials from Myproxy server

User pulls data from public GridFTP server

Chicago Uniform Meeting, 28 Feb, 2012






# Dark Energy Survey



Telescope



Extremely red-sensitive  
500 megapixel camera

Data acquisition system fast enough  
to take images in 17s



# Dark Energy Survey

- The Dark Energy Survey (DES) will have roughly 525 observing nights over a five-year survey period (2011-2016)
- Each of these nights will produce approximately 300 GB of raw data and 2.5 TB total data products after processing
- Within each of its recent Data Challenges, the Data Energy Survey Data Management (DESDM) team has processed 10 such nights of simulated DES data, producing 25-30 TBs data

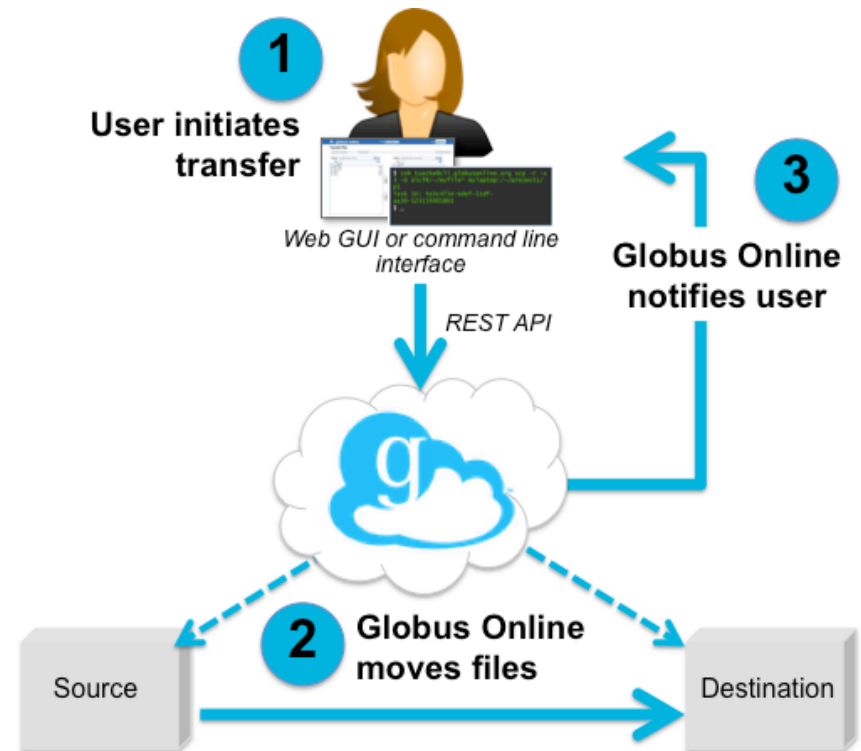


# Let's get down to brass tacks

- From a business perspective, what would I suggest you take a hard look at:
- If large data movement is a problem for you, then look at GridFTP and Globus Online
  - GridFTP is mature; It has been around for 10 years and the sites that report to us average about 15M operations per day world wide. It is available in several of the Linux distributions.
    - Security is either via GSI (X.509 proxies) or ssh
  - Globus Online; It is a relative new comer, but it solves a real problem, provides a security solution that most businesses can work with, it runs on Amazon so the infrastructure is pretty reliable and so far, the software has been pretty solid.
    - Still uses X.509, but authentication to the local myproxy server to get the proxy is via whatever local security you have in place (PAM).
- If High Throughput Computing / Cycle scavenging is of interest then look at Condor
  - Not part of Globus, but a partner of Globus. Can interface to sites with GRAM running, but GRAM is not required
  - Very mature, available in at least RedHat, has some successful commercial users.

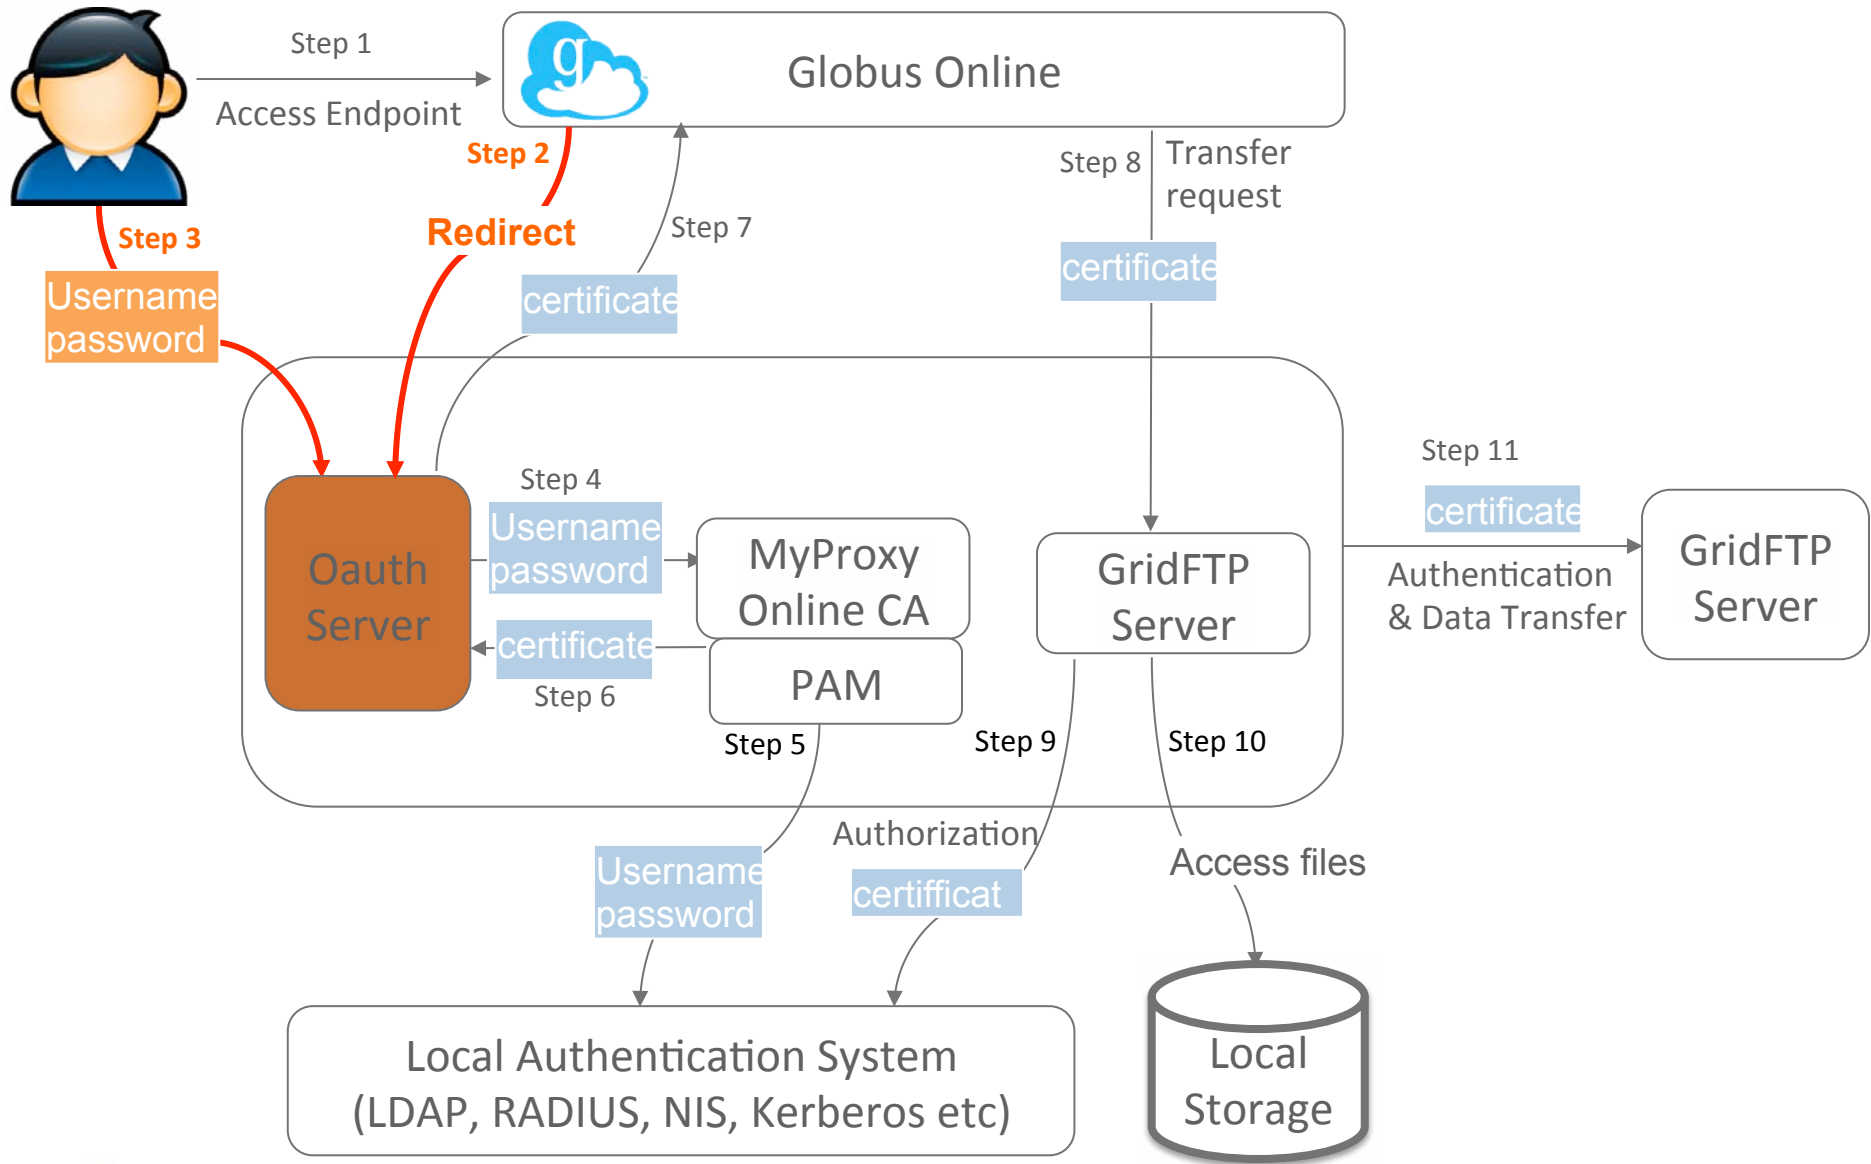
# What is Globus Online?

- Reliable file transfer.
  - Easy “fire-and-forget” transfers
  - Automatic fault recovery
  - High performance
  - Across multiple security domains
- No IT required.
  - Software as a Service (SaaS)
    - No client software installation
    - New features automatically available
  - Consolidated support & troubleshooting
  - Works with existing GridFTP servers
  - Globus Connect solves “last mile problem”
- A recommended service of XSEDE, NERSC, ESnet, UK NGS, UCLA, U. Michigan, ...







# ... with MyProxy OAuth (new feature)



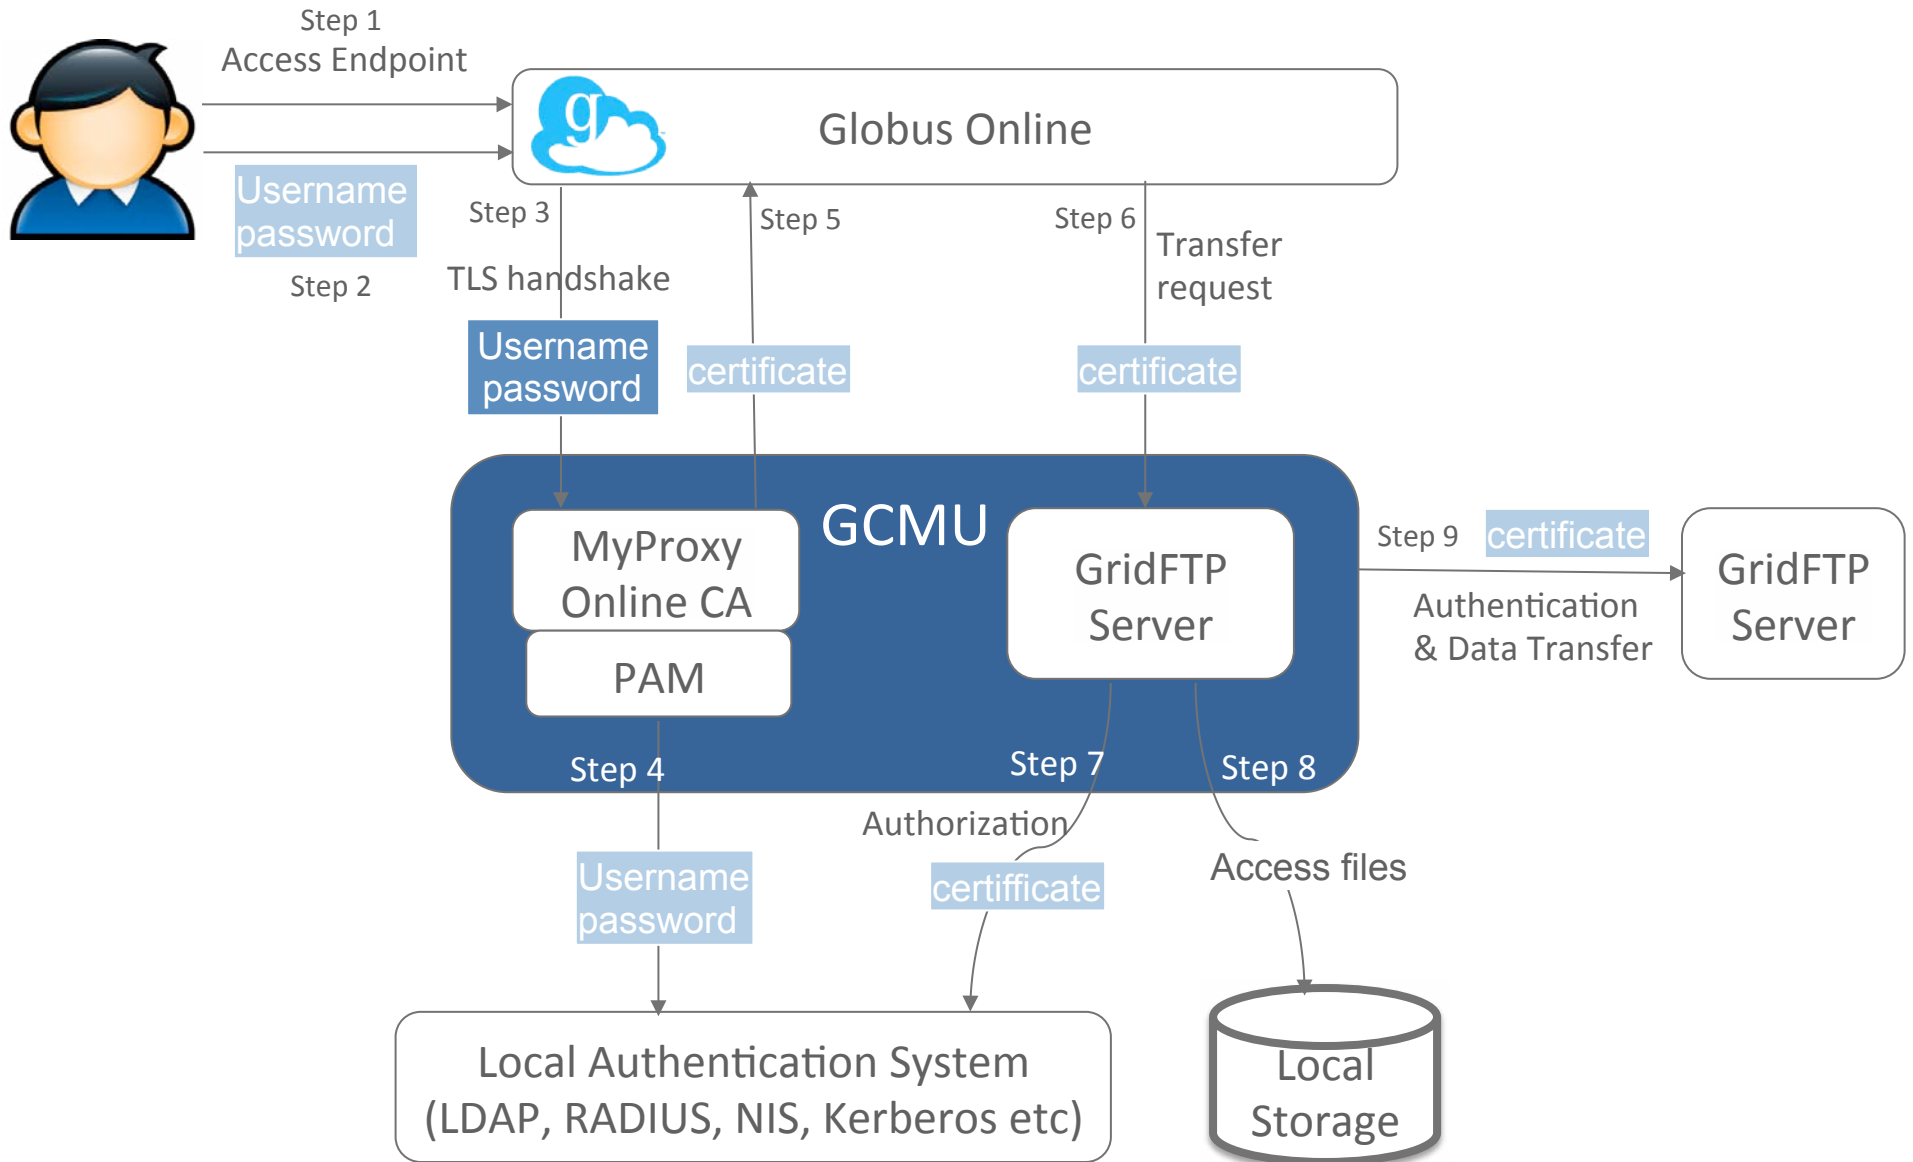
# Globus Online in Action



Astrophysics simulation data generated in Tennessee, moved to Illinois for visualization (Enzo, UCSD; Futures Lab, Argonne)



# Globus Connect Multi User

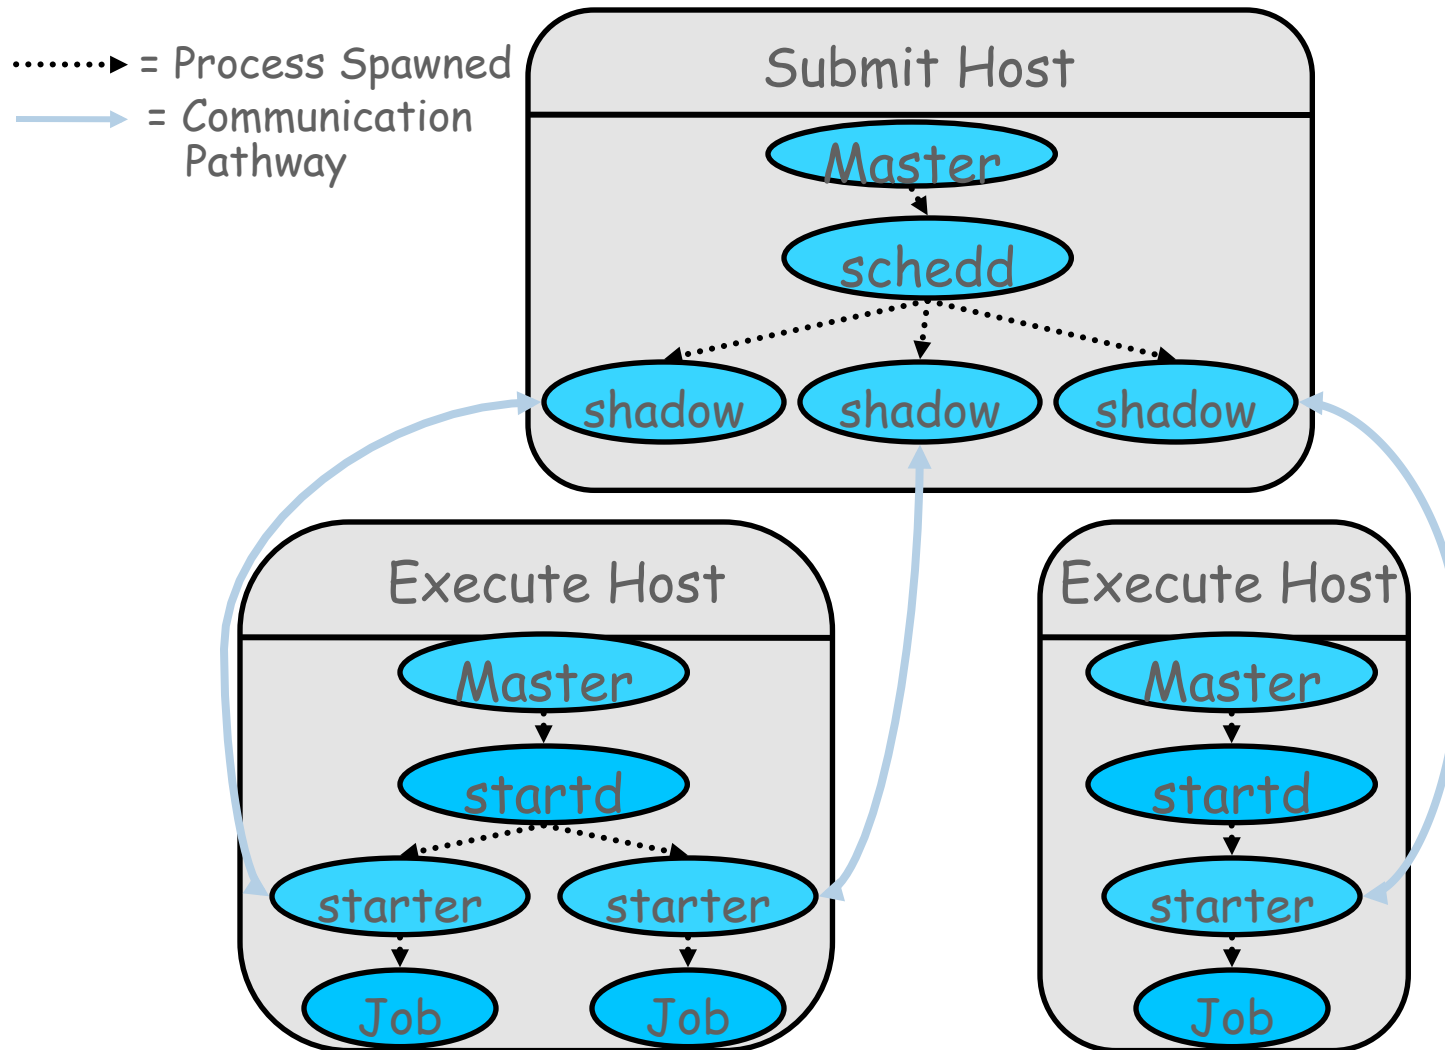


# High Throughput Computing

- BOINC (<http://boinc.berkeley.edu/> )
  - Grew out of the SETI at Home project
  - Infrastructure to set up any “at home” project
  - Has teams and competitions which encourages participation
  - Need to run each job multiple times to have confidence in the results
- Intel (Internal Development Project)
  - Every non-mobile machine in the company was available
  - I have heard claims of high 80’s for utilization
  - Followed the diurnal cycle
    - Singapore ran chip design verification codes on secretaries machine in California at night.
- Condor (<http://research.cs.wisc.edu/condor/> )
  - Designed for large ( $10^5$ ) tasks and large ( $10^4$ ) resources
  - Highly resilient
  - Matchmaking / Class Advertisements
    - Resources advertise their capabilities, jobs advertise their needs, condor arranges a marriage made in silicon 😊
    - No one else that I am aware of has a good system for this.
  - Glide-in’s
    - The job you start on the cluster is Condor startd, which registers itself and then becomes a resource.
    - <http://www.uscms.org/SoftwareComputing/Grid/WMS/glideinWMS/doc.prd/index.html>

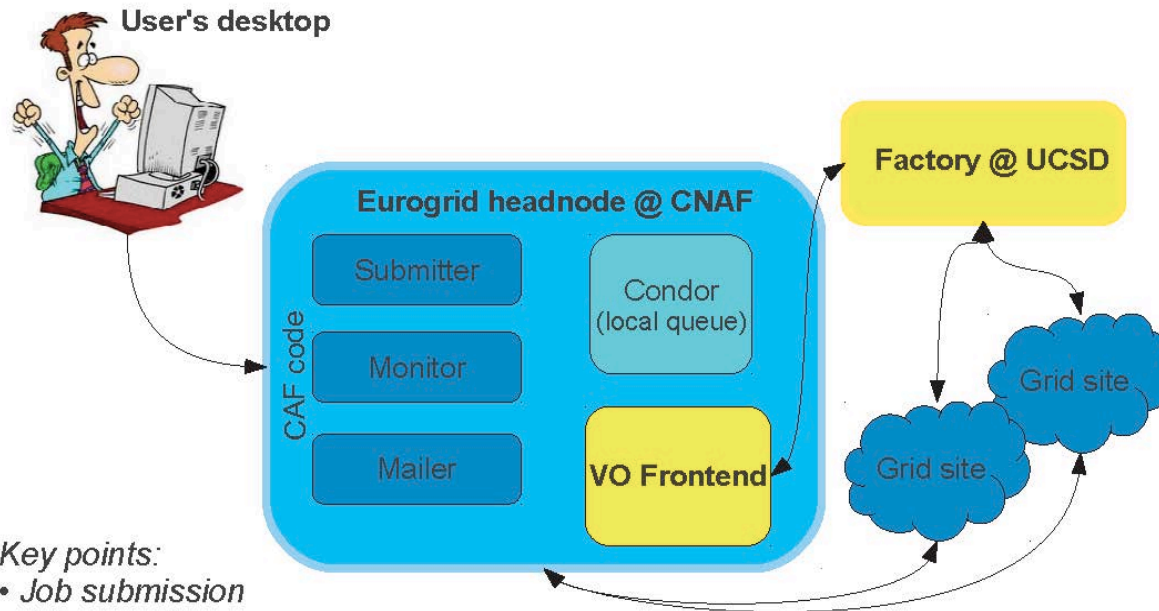


# Condor Pool Layout: running a job



# Dynamically deployed on grid resources

## Eurogrid at a glance



### Key points:

- Job submission
- Job execution
- Authentication
- Code distribution and remote DB access
- Monitoring
- Output retrieval



# A Cool Condor Example

- A pharmaceuticals company needed to run 10,000 jobs.
  - It was estimated to have taken weeks, or possibly even months on their existing infrastructure.
  - HTC / Grid is often a good match for Big Pharma due to their need for parameter sweeps.
- For a proof of concept, they teamed with
  - Cycle Computing (tries to hide the complexity)
  - Condor (to handle scheduling, resiliency)
  - Amazon (to get the actual infrastructure)
- 12:35 – 10,000 Jobs submitted and requests for batches of cores are initiated
- 12:45 – 2,000 cores acquired
- 1:18 – 10,000 cores acquired
- 9:15 – Cluster shut down



# DreamWorks Animation Utilizes the Red Hat Portfolio to Build a Private Cloud for Financial and Creative Applications

**Customer:** DreamWorks Animation SKG

**Industry:** Entertainment

**Geography:** Glendale, California, and Redwood City, California

**Business Challenge:** Enable rapid and cost-efficient scalability of computing and resources for HPC (high-performance compute) workloads and optimal use of a multi-datacenter footprint. Enhance the reliability and availability of mission-critical services.

**Solution:**

**Deploy Red Hat Enterprise MRG as the grid infrastructure of DreamWorks Animation's 30,000+ core render farm**

Deploy Red Hat Enterprise Virtualization as the platform to run mission-critical services on Red Hat Enterprise Linux





# So, what's not to like?

- steep learning curve; significant investment in human effort.
  - The things that have been primarily successful, like GO and Condor, have tried to address this, but it is still non-trivial
- Support
  - This is often a concern for commercial users.
  - The Globus Project is a best effort support model. You can file bugs, or send mail to the list, but there is no SLA in regards to what the response time / quality will be.
  - There is / was some commercial support available
    - Univa – based in Lisle, started by some of the original Globus people. I would describe them as (among other things) trying to be the RedHat of the Grid world (better packaging, commercial support, etc..)
    - IBM used to provide this and ultimately, if you pay them money they will sell you services, but I don't know how good they would be in that arena any more.
    - I am not positive, but presumably you can get support for Condor through RedHat since it is the heart of their cloud bundle.
- QoS
  - If you are not diligent, a resource can get overwhelmed by Grid workloads.
- Debugging
  - With services and jobs running literally all over the planet, figuring out what actually went wrong can be tough.
- Co-Scheduling – if you need multiple resources at multiple places your problem gets exponentially more difficult.



## Questions?

*Knowing the other and knowing oneself,  
In one hundred battles no danger.  
Not knowing the other and knowing oneself,  
One victory for one loss.  
Not knowing the other and not knowing oneself,  
In every battle certain defeat.  
- Sun Tzu, The Art of War*

# Massively Parallel Computing and Argonne's Storage System

# It's all about communication...

- A true statement in so many contexts
  - Business (you probably know this better than I do)
  - Marriage (my wife would probably say you certainly know this better than I do 😊 )
  - And today's topic, parallel computing
- What makes a supercomputer super?
  - The interconnect, i.e., the communications network for processor to processor communication.
    - In fact, the Blue Gene/P runs 850MHz PPC450 chips, but the network is fast enough and low enough latency that you can efficiently throw lots of them at the problem.
  - Remember that the difference between Grid and Parallel computing is independent tasks vs. coordinated tasks
  - Coordination requires communication and if you are spending time communicating (Think I/O wait) then you aren't computing
    - While in theory you can overlap communication and computation, it is often the nature of the beast that you have to compute something, send that result to your neighbors, wait for them to respond, and then based on that response, compute some more. You have to wait.
    - Latency is as important, if not more important than bandwidth
  - The biggest machines use a custom interconnect (Blue Gene, Cray SeaStar or Gemini, Dawning). If they don't use a custom interconnect, Infiniband is the COTS interconnect of choice.



# Blue Gene/P at ALCF



Chicago Uniform Meeting, 28 Feb, 2012




# ALCF Resources – Current

## Cloud

### *Magellan*

504 nodes  
4032 cores  
Retasked Oct, 2011




Summer 2010

USER TEAMS  
(via ESnet, UltraScienceNet,  
Internet2 now, ANI in 2011)

## Production


### *Intrepid*

40 racks  
160K cores  
556 TF



8 PB

DDN 9900s



640 TB

DDN 9550s




8 PB

Spectra Logic T950

### *Eureka (Viz)*

800 cores  
50 NVIDIA S4 GPUs  
100 TF




## Test & Development

### *Surveyor*

1 rack  
4K cores  
13.9TF



128 TB DDN 9550



Gadzoos (Viz)  
32 cores  
2 NVIDIA S4 GPUs





## Mira



**BGQ Compute**  
10 PF  
768 TB RAM  
768 K cores



**BGQ IO**  
384 PCI slots  
6 TB RAM  
6 K cores



IB Federated Switch



### Data Storage

240 GB/s write BW (aggregate)  
28 PB capacity

### 2014 Storage Expansion

~478 GB/s write BW (aggregate for both phases)  
~74 PB aggregate capacity both phases




Login Nodes



GridFTP Servers  
12 servers

### 2014 Tape Expansion

100 PB



**Viz**  
3 TB RAM  
1536 cores  
48 K GPU cores

## Planned 2012 Delivery

## Test & Development



**BGQ Compute**  
208 TF  
16 TB RAM  
16 K cores



**BGQ IO**  
8 PCI slots  
128 GB RAM  
128 cores



IB

**Data Storage**  
2.4 PB  
20 GB/s write BW (aggregate)



Login Nodes



Ethernet

**WAN**  
100 Gb/s

- BG/Q Links (2x2GB/s)
- IB QDR (4GB/s)
- Ethernet (10Gb/s)

# BlueGene/P Overview

## Intrepid System


40 Racks



557 TF/s  
80 TB RAM


**Rack** Cabled 8x8x16

32 Node Cards  
1024 chips, 4096 procs



14 TF/s  
2 TB

**Node Card**  
(32 chips 4x4x2)  
32 compute, 0-2 IO cards



435 GF/s  
64 GB

**Compute Card**  
1 chip, 20  
DRAMs



**Chip**  
4 processors



850 MHz  
8 MB EDRAM

13.6 GF/s  
2.0 GB DDR  
Supports 4-way SMP



**Front End Node / Service Node**  
System p Servers  
Linux SLES10


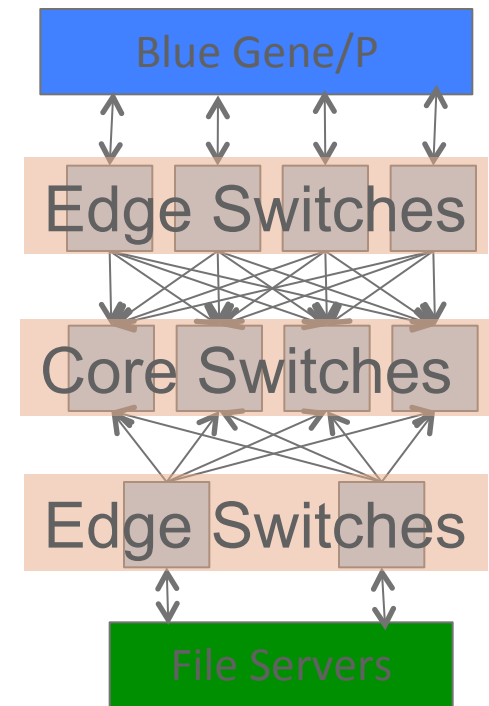
**HPC SW:**  
Compilers  
GPFS  
ESSL  
Cobalt





# Fabric Implementation

- Used Myrinet 10g for the network
  - Over 1000 ports in use
  - Ethernet from BG/P to Edge, MX elsewhere
  - Uses CLOS / Fat tree topology with wormhole switching
  - Pros
    - 2z chip integrated standard Ethernet
    - “dumb” switch, much cheaper (like 80%)
    - Very high density (512 ports in 21U)
    - low latency (not critical for us)
  - Cons
    - Very high density 😊; hard to work with
    - Black box; management tools weak; debugging difficult
    - Needed close coordination with vendor to achieve (aggressive) performance goals
    - Typical new product issues



# Disk Storage

- Key Drivers
  - Performance is paramount
  - Large files, streaming access is the norm
- Use of SATA for cost / capacity, but still needed performance
- Corruption issues a significant concern.
- Configured a storage “cell”
  - (2) storage arrays with a file server / switch rack in between
  - Developed models that could ripple changes through the design (impact on switches, ports, power, space...)



# Tape is dead... Long Live Tape

- The reports of my death are greatly exaggerated... at least for our application
- Space
  - Both of these pictures show approximately 8PB of high density storage
- Power
  - (2) L6-30P plugs < 1KW in normal operation
  - (128) L6-30P plugs approx 250KW in normal operations



# Scientists are Pack Rats...

- Key Drivers
  - Large, unique, relatively uncompressible data
  - Data loss is undesirable, but not catastrophic
  - Restore time is not critical
  - Once archived, it is rarely accessed
  - We do our own maintenance
- Three primary tasks
  - User initiated archive: Uses HPSS. Moves are user initiated primarily via the HSI command, but some users use HTAR. We intended to make HTAR the primary client, but it has file size and path length limits that we hit regularly. This is the bulk of the volume. Scientists are pack rats. They never throw anything away. So when they are "done" with the data, they put it on tape in case they decide they want to re-analyze it 20 years from now.
  - Home backups: First line of defense is GPFS snapshots. Tape uses HPSS. These are initiated by us as root. Client story is the same as above, mostly HSI, we wish we could use HTAR. This is the bulk of the file count (\*lots\* of small files). We estimate it would take us 6 weeks to do a complete restore of a catastrophic failure of our home file system :-).
  - Critical system backups: Uses Amanda. These are for our internal use only. Nightly incremental backups run 15 hours. Monthly full is about 4TB and runs about 4 days. We plan to write dual copies, one to each library. We had intended to be doing so already, but networking between the data centers prevented it. We should have that resolved this year. For now, we "offsite" the tapes in our offices.



# How does tape fit in my world?


- Tape is currently not mission critical to us. If we lose access to a library, we record it, but it causes no significant issues. On the other hand, scheduled availability for HPSS was 99.7%. Note that our facility takes a maintenance outage every other Monday.
- Our data use has been no where near what we projected. I believe it is because we don't run a "scratch file system" like other facilities. We let our users keep there data on disk MUCH longer than other facilities. One particular user has had nearly 0.5PB on disk for 2 years.
- Though not initially planned, we took advantage of the library virtualization features. We host two other divisions on our libraries. Saved the purchase of small libraries and the associated administration and maintenance costs.
- Changes coming:
  - We hope to use GHI, which uses HPSS as a HSM for GPFS.
    - This will mean tape usage will go up dramatically.
    - This will also mean that tape access will become much more critical.
  - server backup (Amanda) load will drop. New servers will mount drives from our NetApp and we will use snapshots and reduce the backup interval.
- My problems / concerns
  - 230 million small files. The metadata kills HPSS. The performance on small files is poor (estimated 6 weeks to restore home)
  - Data Integrity: Is the data still good; replication vs. RAIT, etc..



# Tape Block Diagram


## Critical System Backups

Local Server Disk < 1TB  
Typical Backup




## "Scratch"

8 PB / 60+ GB/s  
Archive Only



## Home

640 TB 8+ GB/s  
90% Archive / 10% Backup



x2 for Backup

Tape Server

- Quad core, 4GB RAM
- (1) 10Gbs Myricom
- (2) 4 Gbs FC HBA

(2) LTO4

x10 for Archive

Tape Server


- Quad core, 4GB RAM
- (1) 10Gbs Myricom
- (2) 4 Gbs FC HBA

(2) LTO4


x8 for Both

Disk Server

- Quad core, 4GB RAM
- (1) 10Gbs Myricom
- (1) IB 4x DDR HCA



Spectralogic  
8 frame T950  
10,000 slots  
(24) LTO4  
tape drives



DDN 9900  
(480) 1TB drives






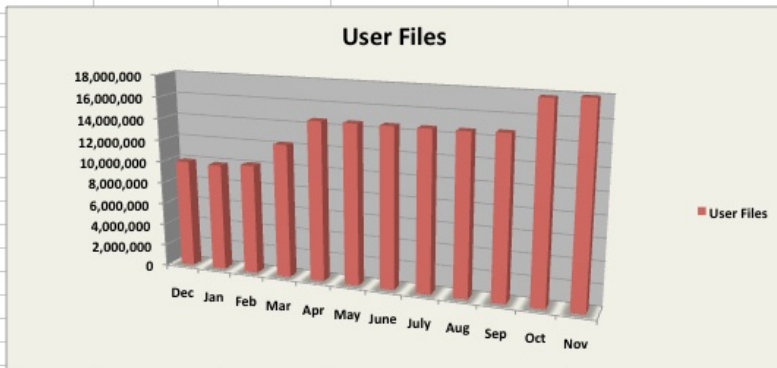
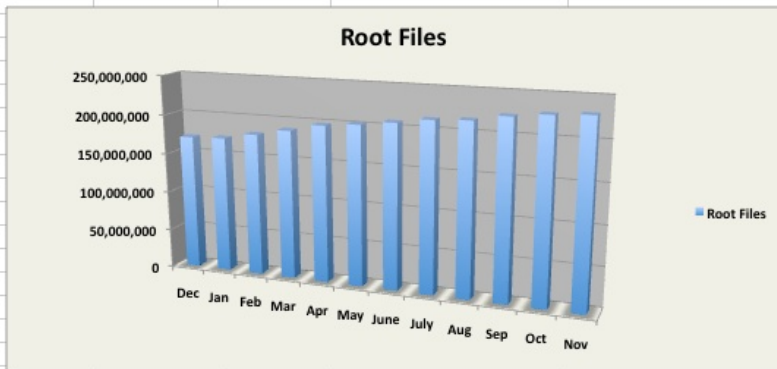
# Tape Software

- Using HPSS ([www.hpss-collaboration.org](http://www.hpss-collaboration.org))
  - large files are a key driver for us
  - Implements tape striping (think RAID 0)
    - Allows reasonable access times to TB sized files
    - Unexpected benefit when paired with Spectralogic: Single tape pick can get all the tapes
      - Actually need a feature added to realize this because HPSS puts the tape back in nearest open slot not where it came from
  - Single name space
    - Hides multiple libraries
    - Uses DB2; Can scale to billions of filenames
  - Weak on the client side.
- Using Amanda for backup
  - Our backup needs are small, but...
  - There are significant limitations to Amanda and we are investigating commercial solutions (TSM, NetBackup, etc)



# A recent snapshot of our HPSS system

HPSS Utilization (over 12 months)				
	Root Files	User Files	Root Archive Size (TB)	User Archive Size (TB)
Dec	171,405,060	9,957,080	260	1223
Jan	172,912,371	9,869,243	284	1276
Feb	180,660,909	10,113,198	534	1389
Mar	188,787,711	12,264,404	685	1413
Apr	198,195,367	14,637,639	723	1440
May	202,293,988	14,640,479	725	1458
June	207,641,399	14,644,211	728	1470
July	213,826,722	14,650,240	737	1492
Aug	216,118,347	14,653,008	742	1509
Sep	223,617,989	14,741,888	754	1681
Oct	228,259,146	17,784,081	763	1988
Nov	230,798,607	17,976,574	766	2097



Chicago Uniform Meeting, 28 Feb, 2012



# Impact of Multiple Cores

- HPC requires IO nodes
- Over the next 5 years you will see 16, 32, 64, maybe even more cores
- Two possibilities
  - Make the storage dumber
    - Bring the storage closer and move functionality out of HW and into SW
    - Software is the key
      - Reliability, manageability, Metadata
    - Google File System is a variant of this
  - Make the storage smarter
    - Put the IO node(s) in the storage
    - Ensure good match between IO node and storage
    - Needs to be flexible on amount of compute
      - Ratio of IO to storage varies widely



# Data Integrity

- RAIT: Nothing new, but some interesting possibilities...
  - Stripe over  $n$  tapes, but need only  $m$  to restore
  - Write  $n+1$  at first, asynchronously add extra redundancy
  - If drives are busy, load  $n$  or  $n+1$  and compute the rest
    - Note the multi-core makes this even more attractive
- Cleversafe
  - DISCLAIMER: I am on the technical advisory board and I have stock in this company.
  - Arbitrary level of parity protection: Split the data across 57 (potentially geographically distributed) devices and as long as I have 39 I can rebuild it
    - The numbers 57 and 39 are just examples. Can trade storage space for reliability
    - Undetected data integrity faults can mathematically be in the hundreds of millions of years.
    - Their focus is reliable and cheap, but they are not fast... sounds kind of like tape...
      - Still random access, high reliability, but requires power.
      - No file system, object based storage system
    - Possible replacement for tape?
      - Particularly if the restore time from tape is not acceptable (photo sites)
      - Shutterfly uses this for their photo storage



# Laser-Plasma Interactions for the National Ignition Facility

Denise Hinkel, LLNL

INCITE  
50M

## Science and Approach

- Using code pF3D simulating both radiation and hydrodynamics
- A large-scale simulation of laser-plasma interaction in the National Ignition Facility (NIF) looking at non-uniform cross-beam energy transfer
- Double quad-beam simulations have shown an increase of reflectivity out of the target
- Expanding to more beams is critical to improving target design


PHYSICS OF PLASMAS 18,  
056312 (2011)

## Key Impact

- Impacted design of target for future National Ignition Facility experiments.
- Fusion energy promises to bring sustainable, green energy production to the world

## ALCF Contributions

- Debugging a few critical issues
- Coordinated large campaigns



Single-quad beam propagation

Double-quad beam propagation



# Lattice QCD

## Paul Mackenzie, Fermilab

### Science and Approach

- Provide crucial high precision lattice QCD calculations needed for new or in-progress experiments and analyze results from experiments
- Used in my high energy and nuclear physics

### Key Impact


- Lattice QCD has delivered essential results to experimental programs
- RHIC - Helped to firmly constrain heavy-ion collision models for the first time. First and second most cited papers in LQCD since 2006
  - Fermilab – Calculations of USQCD members combined with experimental results have allowed many of the fundamental parameters of the Standard Model to be determined more accurately than ever before.

### ALCF Contributions

- Discovered a new, much more efficient, way to calculate the HISQ “fermion force”.
- Reduces FLOPS 10x
  - Improves parallel efficiency
  - Represents 10-50% of total time

- Similar improvements were made for the gauge force routines.
- ~10% of runtime and is now 2-3x faster in the MILC code

Changes are captured in the LQCD SciDAC libraries.





# Deflagration-to-Detonation Transition in Reactive Gases

Alexei Khoklov, University of Chicago

## Science and Approach


- First principle, reactive flow direct numerical simulations
- 7 micron resolution shock tube in CO<sub>2</sub> with heat conduction and isothermal walls has excellent agreement with experiment
- More complicated physics for next suite of simulations

## Key Impact

- Design safe systems for future hydrogen energy systems by gaining insight into the physical mechanisms of the burning and detonation of hydrogen-oxygen mixtures

## ALCF Contributions

- Revealed a strong load imbalance and presented a recommended solution. Collaborating on implementation.
- Visualization team helped visualize data in VISIT
- To prepare for initial INCITE award, ALCF: Implemented OpenMP allowing use of all cores on the node giving 3x speedup and rewrote I/O.



Temperature in a three-dimensional Navier-Stokes first-principles direct numerical simulation of a Mach=3 reflected shock bifurcation in a hydrogen-oxygen mixture in a square channel.

# High Reynolds Number Turbulent Boundary Layers

Robert Moser, University of Texas

## Science and Approach

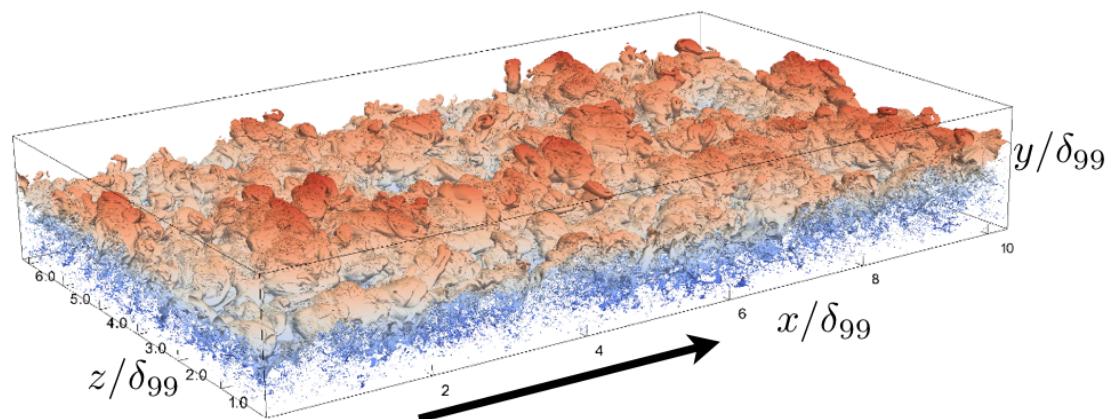
- Highest ever Reynolds number simulations for spatially evolving incompressible turbulent boundary layers
- Direct numerical simulation (DNS) of high Reynolds number turbulent flows in a boundary layer
- Gathering the required turbulence statistics in analysis requires significant computation

## Key Impact

- Boundary layers are central to the energy losses in transportation
- Understanding these flows can improve the design and efficiency of vehicles. These designs are impeded by the lack of accurate models of the turbulent phenomena

## ALCF Contributions

- Initial port and performance optimizations for Intrepid which drove further optimizations by Moser's team.
- Testing threading performance for Mira




This shows the turbulent/non-turbulent interface at the boundary layer; red are turbulent zones and blue are non-turbulent zones. Rotational flow is a key difference between boundary layer flow and internal flows (in channels and pipes). Understanding the rotational flow will reveal more information about the structure of the flow shown here.

# Reducing Toxic Gas through Metal Catalysis

Jeff Greeley, Argonne National Laboratory

INCITE  
16M



## Science and Approach


- Determining the size of nanoparticles to use for catalysts for optimal results is difficult with physical experiments
- Platinum nanoparticles converge much quicker than gold particles (147 atoms vs 561 atoms)
- $O(N^3)$ , real space, grid-based Density Functional Theory (DFT) code, GPAW, for our nanocatalytic modeling efforts.

## Key Impact

- Electrocatalysis for fuel cells
- Metal catalysis of carbon monoxide to carbon dioxide
- Enable design of improved catalytic systems (like catalytic converters) for a range of industrial uses

## ALCF Contributions

- Reduced memory footprint and added layer of parallelization
- Scaled from 512 cores to 131,072 cores
- Parallel I/O implementation running 40% of peak at 32K cores facilitating easier use of Intrepid
- Identified very difficult ScaLAPACK memory bug




The charge density difference of a carbon monoxide molecule – one oxygen atom and one carbon atom -- adsorbed (adhered to the surface) on a gold nanoparticle of 309 atoms



# Multiscale Blood Flow Simulations

George E. Karniadakis, Brown University

INCITE  
50M




## Science and Approach

- Study **blood flow** in a patient-specific cerebral vasculature with a **brain aneurysm** and analyze the interaction of blood cells with the arterial walls that lead to **thrombus formation** and eventual aneurysm rupture.
- Mapping exactly how **red blood cells** move through the brain and study SS and CM diseases.
- Using **computers, scientists and cardiologists** will conduct virtual experiments to study cerebral blood flow and improve **diagnosis and treatment**.

## Key Impact

- **Cerebral aneurysms** occur in up to **5%** of the general population, leading to strokes in over **40,000** Americans each year.
- **Sickle cell anemia** (a chronic inflammatory disease) affects **72,000** individuals in USA.

*Biconcave and sickle shape of the red blood cells (RBC)*




- **Cerebral malaria** is a red blood cell disease, with significant mortality and mostly affects children.

## ALCF Contributions

- Performance tuning and optimization:
  - improved messaging rate
  - improved instruction scheduling
  - reduced memory accesses
  - Improved data alignment
- Scaled up to **131,072 cores (BG/P)**.
- Gordon Bell finalist SC 2011.



Award to recognize outstanding achievement in HPC applications



*Simulation of RBCs in a vessel*

52



# Materials Design from First Principles

Larry A. Curtiss, Argonne National Laboratory

ALCC  
30M

## Science and Approach

Reduce the significant chlorinated waste from producing propylene

- Subnanometer Ag clusters may provide pathways to overcome this waste problem.
- Structure of key molecule optimized and study completed on part of the process


Study the transformation of biofuels with efficient catalysts

- Structure of  $ZrO_2$  nanobowls in dry and hydroxylated alumina have been optimized
- In support of the Institute for Atomic-efficient Chemical Transformation (IACT), an Energy Frontiers Research Center (EFRC)

- J. Phys. Chem. C, 2011, 115 (44), pp 21785–21790

## Key Impact


- Provide the fundamental understanding and predictions needed to design new materials for catalysis and energy storage



Thirty-three atom silver cluster that is being studied as a new catalyst for propylene epoxidation.

## ALCF Contributions

- Significant time was spent mentoring 4 post-docs on using GPAW and how to use Intrepid
- Building off of significant effort on the performance and scalability of GPAW




Nanobowl in  $ZrO_2$  surface (1nm diam.) that is being studied for biomass conversion to fuels



# Faster Design of Better Jet Engines


## Science

- Save cost and time by designing engines through simulation rather than building models
- Technologies from simulations now being applied to next generation high-efficiency low-emission engines
- A key enabler for the depth of understanding needed to meet emissions goals



## Challenges

- I/O algorithm redesign speeds up simulations by 3x





# Insight into Parkinson's Disease


Igor Tsingelny  
University of California, San Diego

## Science


- Parkinson's Disease affects 5 million in US and Europe
- Increased aggregation of  $\alpha$ -synuclein protein is thought to lead to harmful pore-like structures in human membranes
- Simulations show  $\alpha$ -synuclein complexes, and  $\beta$ -synuclein prevents creation of propagating  $\alpha$ -synuclein complexes

## Methods and Challenges


- Using NAMD and MAPAS on Blue Gene at ALCF and SDSC
- $\alpha$ -syn proteins are unstructured



$\alpha$ -synuclein coformer at membrane (left), completed pentamer (above), evolution of  $\alpha$ -synuclein over 5ns (right)



Chicago Uniform Meeting, 28 Feb, 2012



## Questions?

*Knowing the other and knowing oneself,  
In one hundred battles no danger.  
Not knowing the other and knowing oneself,  
One victory for one loss.  
Not knowing the other and not knowing oneself,  
In every battle certain defeat.  
- Sun Tzu, The Art of War*