



Open Science Grid

# Grid Computing Introduction (based on a talk by Brian)

Brian Bockelman, [bbockelm@cse.unl.edu](mailto:bbockelm@cse.unl.edu)

OSG Staff

University of Nebraska-Lincoln

# Part I: Introduction to Grids



# Outline

---

- Grid Computing, definitions and implementation.
- How security and handling works on the OSG.
- The OSG approach to grid computing.



# Grid Computing

- Per usual, wikipedia offers a decent starting point:
  - **Grid computing** is the combination of computer resources from multiple administrative domains for a common goal.
- Grid computing is used to perform computations which may not be feasible otherwise. Reasons may be:
  - Practical (one site can't hold all the computers)
  - Opportunistic (an organization wants to take advantage of more computing resources)
  - Political (multiple big sites working together).

# Grid Computing

---

- Important aspects of the definition:
  - “Combination of computing resources”:
    - Implies each resource can function separately.
    - Overall, extra layer of difficulty to handle compared to using a single resource (even if this is hidden from the end user!).
  - “Multiple administrative domains”:
    - There must be some level of trust between the user and sites. These trust relationships can be very complex!

# Original Idea

---

- The original idea behind grid computing was to make computing power as easy to access as the electrical grid.
  - You could take your job and plug it in to “the grid”.
  - Everyone can use the same interface.
  - This metaphor also implied grids would be as easy to use as the power system...
    - (... which might have been a pipe dream)

# What makes it unique?

---

- Food for thought:
  - What's the difference between grid computing and cloud computing?
  - What's the difference between grid computing and the capabilities Condor provides?

# Grids in the US

---

- The two largest grids in the US are the Teragrid and the OSG.
  - Both are formed by taking traditional computing sites and allowing users to access resources in a somewhat uniform manner.
  - Resources include:
    - Unique supercomputers: IBM Blue Genes
    - Linux clusters: Loosely-coupled Intel/Linux
    - Data archives: Long-term tape storage
    - Large-scale data systems: Distributed or clustered file systems, providing hundreds of TB to multiple petabytes.



# Teragrid, in a nutshell

---

- The Teragrid is formed by a small number (less than 10) of computing centers.
  - These are some of the largest computing centers in the world.
  - Often multiple, unique resources per center.
    - Clearly favors a few incredibly powerful resources. By invite only.
  - Compute time is allocated by committee.
  - Access via both grid protocols and ssh logins.

# The Open Science Grid

---

- The OSG is a grid formed by 80 sites across the US and the world.
  - Most sites are small-to-medium Linux clusters, with a few large clusters.
    - Primary stakeholders are the LHC and LIGO.
  - Focus is on data-intensive, high-throughput processing (not “traditional supercomputing”).
  - Compute time is allocated by individual site policy. Strong emphasis on decentralization.
  - No SSH logins allowed to remote sites.

<https://twiki.grid.iu.edu/bin/view/ReleaseDocumentation/WhatIsOSG>



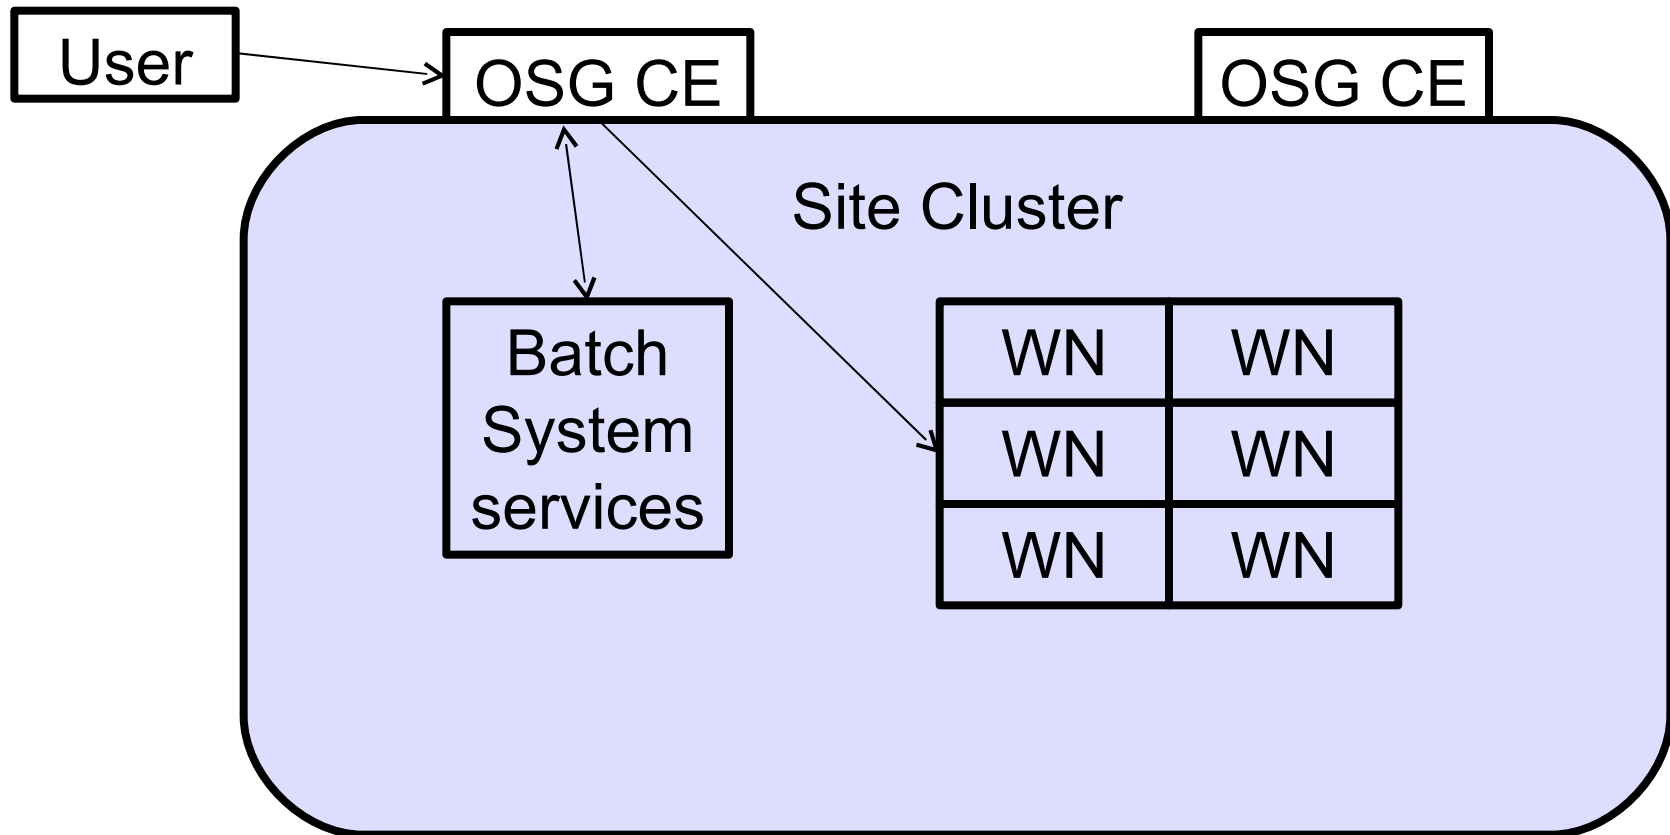
# Map of OSG sites



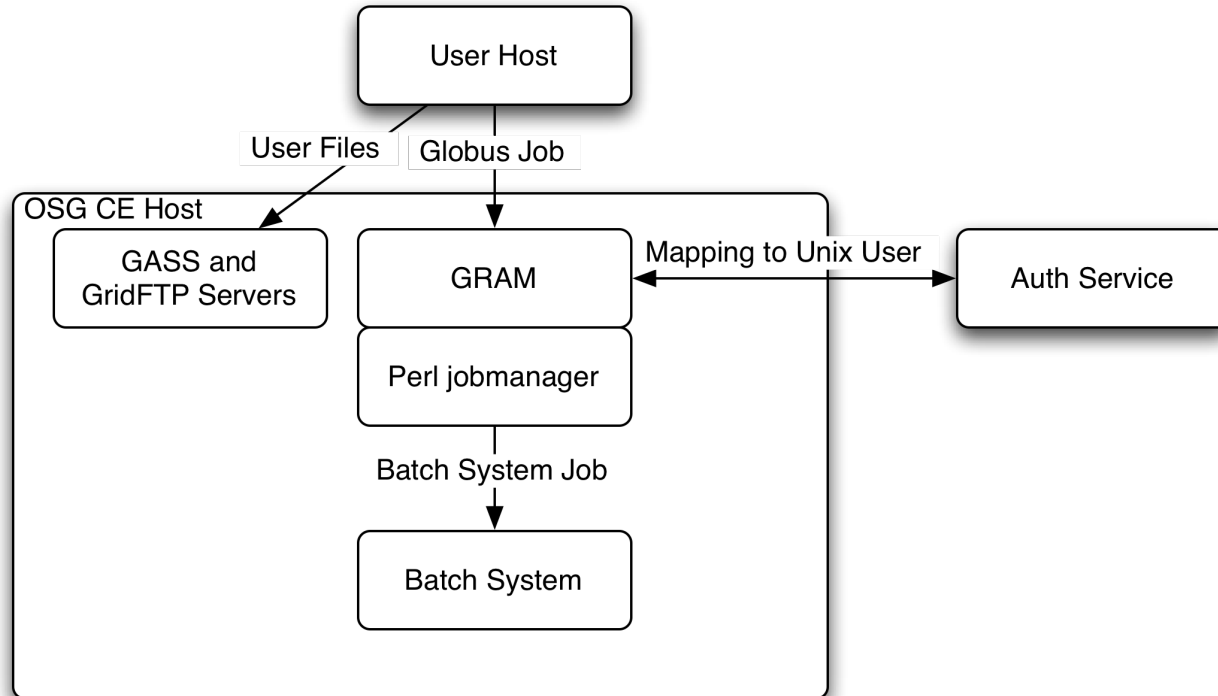
To give you a feel for the distribution of the OSG sites...

# OSG Compute Resources

- The OSG CE is layered on top of a traditional batch system.



# Inside the OSG CE



- The current core of the OSG CE is the Globus Toolkit. Important parts are:
  - Globus GRAM: Translates jobs in RSL format to batch system jobs; allows generic batch system commands to be translated to the site.
  - GASS server: Used to stage small files in and out of the
  - GridFTP server: Used to move large files in and out of the host.

# Anatomy of a Grid Job

---

- All job creations will have the following steps:
  - User creates job and determines which OSG CE it will be sent to.
  - User submits job description and files to Globus on the OSG CE.
  - Globus converts the job to a batch system job, and submits that.
  - Job starts running on the worker node.
- Job finishing goes up in the reverse direction.
- For the user to know the job status, there are **3** different systems (user, Globus, batch system) that must be in sync.

# OSG CE, In Summary

---

- The core of the CE is the Globus middleware (using the GT2 protocol).
  - Globus allows an abstract interface to your system to be exported to the world.
  - Provides the means to allow users to submit to multiple administrative domains.
- The CE interacts with your cluster's batch system; grid jobs are converted to and run as batch system jobs.
  - The component doing the conversion is called the "Job Manager". Condor sites use "jobmanager-condor", PBS sites use "jobmanager-pbs", etc.

# And everything else

---

- There are quite a few more OSG components:
  - Monitoring
  - Accounting
  - Information services
  - Storage and transfer
- Which I will not be covering in this talk.
- You'll learn them throughout the week (assuming you don't skip class).



# Review

---

- Grid computing allows one to utilize multiple computing resources from multiple administrative domains.
  - This is more complex than traditional batch systems.
- OSG is one implementation of a grid; its technology is based upon the Globus Toolkit and Condor.
  - It has almost 100 computing resources (clusters) and 50 storage resources.



Open Science Grid

# Part II: Trust Relationships

# Trust Relationships

---

- What kind of trust relationships do we encounter in the airport?
  - Passports
  - Tickets
  - “Secured area” inside terminal

# Trust Relationships in the Grid

---

- In the grid, we usually think of the users/organizations as the consumers and the sites as the producers.
  - How does the site know you are who you say you are?
  - How does the site know you are allowed to submit jobs?

# Compare to the Cloud

---

- What kind of trust relationship is there in the Amazon EC2 cloud?
  - You trust a SSL connection with the Amazon SSL certificate has Amazon on the other side.
  - Amazon allows you to use the compute resources if your credit card number is valid.
  - You trust Amazon gives you a certain amount of computing for your money.
  - What else?
- Note the trust is 2 way!

# The one you forgot about!

---

- How do you trust the site?
  - After all, this is your data! How do you know they aren't going to steal your Ph.D. thesis? Your new novel protein?
    - (How is this different from any case of using computing you don't own?)
    - The answer is that you trust the organization running the resource.
    - Often, you trust the OSG to only allow reputable organizations to join.
  - In this case, the trust relationship is based on society, not technology.
    - Keep this in mind – the societal aspects are equally as important as the technology sometimes.
    - Read “Reflections on Trusting Trust”!

# Authn and Authz

---

- In order to establish trust relationships on the grid, two things need to happen:
  - Authentication (authn): The process of establishing an identity for your job.
  - Authorization (authz): Determining that your job is allowed to run at the site.
- Think: What authentication and authorization need to happen at an airport?

# X509 and GSI Security

- In the OSG, authentication happens using a *grid certificate*. This is simply a personal SSL certificate.
  - The grid certificate is signed by a trusted authority – the Certificate Authority – and vouches for your identity.
  - When you need to temporarily delegate your rights to elsewhere – like to a remote job – you can use your certificate to form a *proxy certificate*.
    - This authenticates a grid job as belonging to you. Any grid job with your proxy you get the blame for.

Your grid certificate identifies who you are!



# Authorization on the OSG

---

- It would be very hard for sites to authorize each user independently (think: CMS has around 2000).
- Instead, each site authorizes the organizations they want to partner with to run at their site.
  - And the organization securely informs the site who is in their organization.
  - Because these organizations don't always deal with a physical entity (like a single campus or lab), they're referred to as a “*virtual organization*” or a VO.

# Take Home Message

---

- You are identified by your certificate. This is your *authentication*.
- You must join a VO to use the OSG.
- A site makes *authorization* decisions based upon a VO.
- With this model, we tend to minimize the number of communications between the site, user, and organization.
- The OSG implements authorization and authentication on top of x509 certificates and PKI.